



Hitachi Universal Storage Platform V

Hitachi Universal Replicator User's Guide

FASTFIND LINKS

[Document Organization](#)

[Product Version](#)

[Getting Help](#)

[Contents](#)



Contents

Preface.....	xiii
Intended Audience	xiv
Product Version	xiv
Document Revision Level	xiv
Changes in this Revision.....	xiv
Document Organization.....	xv
Referenced Documents	xv
Document Conventions	xvii
Convention for Storage Capacity Values.....	xviii
Getting Help.....	xviii
Comments	xix
Overview of Universal Replicator.....	1-1
Hitachi Universal Replicator (UR).....	1-2
Features.....	1-3
Benefits.....	1-3
Business Solutions.....	1-4
About Universal Replicator Operations.....	2-1
Functionality Overview.....	2-2
Journal Obtain.....	2-4
Journal Copy	2-4
Journal Restore	2-4
UR Components	2-5
USP V Storage Systems	2-7
Universal Replicator Software	2-7
Logical Disk Controllers	2-7
Main and Remote Control Units (Primary and Secondary Storage Systems) ..	2-8
Consistency Groups	2-9
Remote Copy Connections.....	2-10

Initiator Ports and RCU Target Ports	2-11
Journal Groups	2-11
Data Volume Pairs	2-12
Journal Volumes	2-13
The Number of Journal Volumes	2-13
Specifications of Journal Volumes	2-13
Restrictions on Journal Volume	2-14
Journal Volume Areas	2-15
Host Failover Software	2-15
UR Remote Copy Operations	2-17
Initial Copy Operation	2-17
Update Copy Operation	2-19
Journal Group Operations	2-19
Journal Group Operations	2-19
Read and Write I/O Operations During UR Operations	2-20
Secondary Data Volume Write Option	2-21
Difference Management	2-21
Journal Processing	2-22
Journal Processing at the Primary Storage System	2-22
Types of Journal	2-23
Journal Processing at the Secondary Storage System	2-23
Storing Journal at the Secondary Storage System	2-23
Selecting and Restoring Journal at the Secondary Storage System	2-23
UR Delta Resync Operation	2-26
Journal Obtain in TrueCopy Synchronous Secondary Site	2-26
Switching the Master Journal Group of UR	2-28
Pair Status	2-32
Split Types (PSUS)	2-35
Suspend Types	2-37
Suspension Condition	2-39

Preparing for Universal Replicator Operations 3-1

Requirements and Restrictions	3-2
System Requirements	3-2
One-to-One Volume Copy Operations	3-3
Logical Unit (LU) Types	3-4
The Maximum Number of Pairs	3-6
Journal Groups	3-8
Accessing UR Primary Data Volume and Secondary Data Volume	3-9
Cache and Nonvolatile Storage (NVS)	3-9
Host Failover Software	3-9
Duplicate Volumes	3-9
Installing the Hardware	3-10
Setting up Remote Copy Connections	3-11

Configuring the Primary and Secondary Storage Systems for UR Operations	3-14
Enabling the UR Option	3-15
Interoperability with Other Products and Functions	3-15
LUSE	3-17
Virtual LUN	3-18
Cache Residency Manager	3-18
LUN Manager	3-18
Server Priority Manager	3-18
ShadowImage	3-18
TrueCopy Synchronous (3DC Cascading Configuration)	3-25
Basic Behavior	3-26
Hardware Configuration	3-27
Setup Procedure	3-27
TrueCopy Synchronous (3DC Multi-target Configuration)	3-28
Hardware Configuration	3-30
Setup Procedure	3-31
Requirements for Creating UR Pair for Delta Resync Operation	3-32
Requirements for Performing Delta Resync Operation	3-32
Changing to 3DC Multi-target Configuration after Recovering from Primary Site Failures	3-34
Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC Cascading Configuration)	3-34
Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC multi-target configuration)	3-35
Transferring Business Tasks from TrueCopy Secondary Site to Primary Site (When Delta Resync Operation is Performed in 3DC Multi-target Configuration)	3-36
Recovering from Failures in the Primary Site and the TrueCopy Synchronous Secondary Site	3-39
Transferring Business Tasks from the UR Secondary to the Primary Site	3-39
Data Retention Utility	3-41
Copying Access Attribute	3-42
License Capacity of Data Retention Utility	3-45
Performance Monitor	3-45
Planning of Journal Volumes	3-47
Data Transfer Speed Required for Journal Volumes	3-47
RAID Group Configuration and Journal Group Configuration	3-48
Arranging Journal Volumes	3-49
Journal Volume Capacity	3-49
Journal Volumes in Restore Journal Groups	3-51
Factors that Determine Data Transfer Speed between the Primary and Secondary Storage Systems	3-51
Bandwidth of Data Transfer Paths	3-52
Journal Transfer Speed of the DKC	3-52
Performing Remote Copy Between TagmaStore USP/NSC and USP V	3-53

System Option Mode	3-53
Logical Storage System (LDKC) Connected to TagmaStore USP/NSC	3-53
Volume that Can Be Used to Create Pairs.....	3-54
Connection with TagmaStore USP/NSC for 3DC Remote Copy Configuration .	3-55

Using the Universal Replicator GUI 4-1

Journal Operation Window	4-2
Pair Operation Window	4-7
DKC Operation Window.....	4-14
Displaying Information about Remote Storage Systems.....	4-16
Displaying Information about Logical Paths	4-17
Displaying Information about Ports on the Local Storage System	4-18
Usage Monitor Window	4-20
History Window.....	4-22
Optional Operation Window.....	4-28

Configuring Storage Systems and Logical Paths 5-1

Reviewing System Configuration Tasks.....	5-2
Configuring Port Attributes	5-4
Configuring Storage System Options	5-7
Establishing Relationship between Primary and Secondary Storage Systems (Add DKC).....	5-9
Changing Options for Logical Paths and Storage Systems	5-13
Adding Logical Paths.....	5-15
Viewing the Status of Logical Paths.....	5-17
Deleting Logical Paths.....	5-20
Managing Power for Storage Systems and Network Relay Devices	5-21
When the Power Stops Unexpectedly.....	5-21
When the Power is Removed from the Primary Storage System.....	5-21
When the Power is Removed from the Secondary Storage System.....	5-21
When the Power is Removed from Network Relay Devices	5-22
Turning Power Off Intentionally	5-22
Power Off the Primary Storage System.....	5-22
Power Off the Secondary Storage System.....	5-23
Power Off the Primary and Secondary Storage Systems at Same Time...	5-23
Power Off Network Relay Devices	5-24
Removing Relationship Between Primary and Secondary Storage Systems.....	5-24

Using Journal Groups..... 6-1

Reviewing Administrator Tasks for Managing Journals	6-2
Registering Journal Volumes in a Journal Group	6-3
Deleting Journal Volumes from a Journal Group	6-8
Displaying Detailed Information about a Journal Group.....	6-10

Changing Options for a Journal Group.....	6-15
Deleting a Journal Group	6-20
Splitting a Mirror (Pairsplit-r: Suspending a copy operation).....	6-21
Restoring a Mirror (Pairresync: Resuming a copy operation)	6-23
Deleting Data Volumes from a Mirror (Pairsplit-S: Ending a copy operation)	6-25
Performing UR Pair Operations	7-1
Reviewing the Pair Operation Window.....	7-2
Filtering Information in the List in the Pair Operation Window	7-3
Saving Pair Status Information into a Text File	7-5
Creating a Pair of Data Volumes (Paircreate)	7-7
Displaying Detailed Information about a Pair of Data Volumes	7-13
Changing Options for a Data Volume Pair.....	7-17
Splitting a Pair of Data Volumes (Pairsplit-r)	7-19
Restoring a Pair of Data Volumes (Pairresync)	7-22
Releasing a Pair of Data Volumes (Pairsplit-S)	7-25
Creating a Point-in-Time Copy of Data Volumes.....	7-28
Recovering a Pinned Track.....	7-28
Recovering a Pinned Track on a UR Data Volume	7-28
Recovering a Pinned Track on a UR Master Data Volume	7-29
Usage Monitor Operations	8-1
Reviewing the Usage Monitor Window	8-2
Starting and Stopping Usage Monitoring.....	8-3
Displaying the Usage Monitor Graph	8-4
Saving Monitoring Data in Text Files	8-7
Viewing Pair Operation History	8-7
Saving Operation History to a Text File	8-8
Performing Disaster Recovery Operations	9-1
Preparing for Disaster Recovery.....	9-2
File and Database Recovery Procedures.....	9-2
Switching Operations to the Secondary Site	9-3
Transferring Operations Back to the Primary Site.....	9-3
Resuming Normal Operations at the Primary Site	9-4
Recovering from a Disaster in a 3DC Cascading Configuration.....	9-5
Recovering from Primary Site Disaster in 3DC Multi-Target Configuration.....	9-7
Recovering from Primary Site Failures (When Delta Resync Operation is Performed)	9-8
Troubleshooting	1
Checking Universal Replicator Error Codes.....	7

Checking Service Information Messages (SIMs).....	7
Calling the Hitachi Data Systems Support Center.....	9
Acronyms and Abbreviations	1
Index	1



Preface

This document describes and provides instructions for using the Hitachi Universal Replicator software to configure and perform Universal Replicator operations on the Hitachi Universal Storage Platform V (USP V) storage system.

Please read this document carefully to understand how to use this product, and maintain a copy for reference purposes.

This preface includes the following information:

- [Intended Audience](#)
- [Product Version](#)
- [Document Revision Level](#)
- [Changes in this Revision](#)
- [Document Organization](#)
- [Referenced Documents](#)
- [Document Conventions](#)
- [Convention for Storage Capacity Values](#)
- [Getting Help](#)
- [Comments](#)

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

Intended Audience

This document is intended for system administrators, Hitachi Data Systems representatives, and Authorized Service Providers who are involved in installing, configuring, and operating the Hitachi Universal Storage Platform V storage system.

This document assumes the following:

- The user has a background in data processing and understands RAID storage systems and their basic functions.
- The user is familiar with the Hitachi Universal Storage Platform V storage system and has read the *Universal Storage Platform V User and Reference Guide*.
- The user is familiar with the Storage Navigator software for the Universal Storage Platform V and has read the *Storage Navigator User's Guide*.
- The user is familiar with the operating system and web browser software on the system hosting the Storage Navigator software.

Product Version

This document revision applies to Universal Storage Platform V microcode 60-01-0x and higher.

Document Revision Level

Revision	Date	Description
MK-96RD624-P	February 2007	Preliminary Release
MK-96RD624-00	April 2007	Initial Release, supersedes and replaces MK-96RD624-P
MK-96RD624-01	May 2007	Revision 1, supersedes and replaces MK-96RD624-00

Changes in this Revision

Not applicable to this release.

Document Organization

The following table provides an overview of the contents and organization of this document. Click the [chapter title](#) in the left column to go to that chapter. The first page of each chapter provides links to the sections in that chapter.

Chapter	Description
Chapter 1 Overview of Universal Replicator	This chapter describes the Universal Replicator (UR) feature which enables you to create and maintain duplicate copies of all user data stored on a distant Hitachi Universal Storage Platform V (USP V) storage system for data duplication, backup, and disaster recovery purposes.
Chapter 2 About Universal Replicator Operations	This chapter UR provides a RAID storage-based hardware solution for disaster recovery which enables fast and accurate system recovery. Once UR operations are established, duplicate copies of data are automatically maintained for backup and disaster recovery purposes
Chapter 3 Preparing for Universal Replicator Operations	This chapter describes the prerequisites and preparations to make to begin UR operations.
Chapter 4 Using the Universal Replicator	This chapter describes how to use the UR graphical user interface:
Chapter 5 Configuring Storage Systems and Logical Paths	This chapter introduces you must consider when introducing Universal Replicator (henceforth abbreviated as UR) into your system and configuring your system for remote copy operations; for example, deciding which data should be copied from which storage system to which storage system.
Chapter 6 Using Journal Groups	This chapter introduces how to set up UR in your system and configure your system for remote copy operations. The primary site administrator and the secondary site administrator must configure journal groups, as well as ports and logical paths.
Chapter 7 Performing UR Pair Operations	This chapter discusses using the Pair Operation window of UR.
Chapter 8 Usage Monitor Operations	This chapter describes using the Usage Monitor window to follow the collection and display of performance monitoring statistics.
Chapter 9 Performing Disaster Recovery Operations	This chapter describes how to determine the best approach for disaster recovery. The user should anticipate and plan for all types of failures and disasters.
Troubleshooting	This chapter provides general troubleshooting information for Universal Replicator.
Acronyms and Abbreviations	Defines the acronyms and abbreviations used in this document.
Index	Lists the topics in this document in alphabetical order.

Referenced Documents

Hitachi Universal Storage Platform V:

- *Command Control Interface (CCI) User and Reference Guide*, MK-90RD011
- *Data Retention Utility User's Guide*, MK-94RD210
- *LUN Expansion User's Guide*, MK-96RD616
- *Performance Manager User's Guide*, MK-94RD218
- *Shadow Image User's Guide*, MK-96RD618

- *Shadow Image for IBM z/OS User's Guide*, MK-96RD619
- *USP V Storage Navigator User's Guide*, MK-96RD621
- *True Copy User's Guide*, MK-96RD622
- *True Copy for IBM z/OS User's Guide*, MK-94RD623
- *Universal Replicator for IBM z/OS User's Guide*, MK-94RD625
- *Universal Volume Manager User's Guide*, MK-94RD626
- *Virtual Partition Manager User's Guide*, MK-94RD629








Document Conventions

The terms “Universal Storage Platform V” and “USP V” refer to all models of the Hitachi Universal Storage Platform V, unless otherwise noted.

This document uses the following typographic conventions:

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

This document uses the following icons to draw attention to information:

Icon	Meaning	Description
	Note	Calls attention to important and/or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (e.g., disruptive operations).
	WARNING	Warns the user of severe conditions and/or consequences (e.g., destructive operations).
	DANGER	Dangers provide information about how to avoid physical injury to yourself and others.
	ELECTRIC SHOCK HAZARD!	Warns the user of electric shock hazard. Failure to take appropriate precautions (e.g., do not touch) could result in serious injury.
	ESD Sensitive	Warns the user that the hardware is sensitive to electrostatic discharge (ESD). Failure to take appropriate precautions (e.g., grounded wriststrap) could result in damage to the hardware.

Convention for Storage Capacity Values

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

- 1 KB = 1,000 bytes
- 1 MB = 1,000² bytes
- 1 GB = 1,000³ bytes
- 1 TB = 1,000⁴ bytes
- 1 PB = 1,000⁵ bytes

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

- 1 KB = 1,024 bytes
- 1 MB = 1,024² bytes
- 1 GB = 1,024³ bytes
- 1 TB = 1,024⁴ bytes
- 1 PB = 1,024⁵ bytes
- 1 block = 512 bytes

Getting Help

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any error messages displayed on the host system(s).
- The exact content of any error messages displayed by Storage Navigator.
- The Storage Navigator configuration information (use the FD Dump Tool).
- The service information messages (SIMs), including reference codes and severity levels, displayed by Storage Navigator.

The Hitachi Data Systems customer support staff is available 24 hours/day, seven days a week. If you need technical support, please call:

- United States: (800) 446-0744
- Outside the United States: (858) 547-4526

Comments

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

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

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

Overview of Universal Replicator

This chapter provides an overview of the Hitachi Universal Replicator software and describes its features and benefits. This chapter covers the following key topics:

- [Hitachi Universal Replicator](#)
- [Features](#)
- [Benefits](#)

Hitachi Universal Replicator (UR)

The Hitachi Universal Replicator software intelligently replicates data among storage environments controlled through the Hitachi Universal Storage Platform V, satisfying the most demanding disaster recovery and uptime requirements. Since its introduction on the Hitachi TagmaStore® Universal Storage Platform and Network Storage Controller, the Universal Replicator software has set a new standard for data protection by redefining the way asynchronous replication is performed.

Reliable data storage and recovery systems are essential in today's market climate where downtime can be very costly. Businesses must manage increasing amounts of data across a variety of storage systems and operating environments in various locations, while optimizing usage of storage hardware resources and minimizing the management burden.

To address these needs, Hitachi Universal Replicator software provides the enterprise-class performance associated with storage system-based replication while delivering resilient business continuity. Through the Hitachi RAID storage systems, Universal Replicator provides a powerful data management and recovery solution that replicates data to a variety of storage platforms at one or multiple remote sites. Data is replicated asynchronously over any distance without the need for redundant servers or replication appliances, thus significantly reducing resource consumption.

The Hitachi Universal Replicator software helps organizations to:

- Lower the cache and resource consumption on production/primary storage systems
- Improve bandwidth utilization
- Simplify bandwidth planning
- Mitigate the impact of network failures
- Gain more flexibility in trading off between Recovery Point Objective and cost
- Implement advanced multi-data center support more easily
- Move data among levels of tiered storage systems more easily
- Fully leverage the Universal Storage Platform V and optimize the storage infrastructure

Features

Hitachi Universal Replicator provides the following key features:

- Heterogeneous Storage System Support
 - Used with the Universal Storage Platform or Network Storage Controller, Universal Replicator software enables storage management and disaster recovery in heterogeneous systems, providing maximum flexibility and support of enterprise-class environments.
 - Universal Replicator software supports any storage connected to a Universal Storage Platform or Network Storage Controller, permitting data to be copied from any supported device to any other supported device, regardless of operating system or protocol differences. This ensures maximum flexibility for data distribution as well as increased storage utilization and failover options.
- More Efficient Replication
 - Universal Replicator software uses asynchronous replication driven by the remote site to minimize impact on primary production systems and takes advantage of journaling rather than cache files to mitigate the high resource usage of other asynchronous approaches.
 - Storage usage on the Universal Storage Platform or Network Storage Controller can be minimal, just enough for the journals.
 - Limited use of cache leaves cache for production application usage, further restoring primary site storage to its intended role as a transaction processing resource, not a replication engine.
 - Advanced three data center capabilities provide a choice of cascade or multi-target configurations (teams with TrueCopy Synchronous software for advanced configurations).
 - Consistency groups can span multiple storage systems for large enterprise-class applications requiring unmatched scalability and data integrity.

Note: Please check with your Hitachi Data Systems representative for detailed feature availability information.

Benefits

The business benefits of Hitachi Universal Replicator include:

- Ensure Business Continuity
 - Simplifies implementation to meet the most demanding disaster recovery and uptime requirements, regardless of the type of supported storage platform hosting the business-critical data
 - Supports availability of up-to-date copies of data in dispersed locations by leveraging Hitachi TrueCopy® Synchronous software

- Maintains integrity of a replicated copy without impacting processing, even when replication network outages occur or optimal bandwidth is not available
- Works with Universal Storage Platform V replication technology to greatly enhance administrative productivity and response to and proactive aversion of crises
- Optimize Resource Usage
 - Leverages advanced technology to maintain data integrity and optimize the storage/IT infrastructure for protection of data from any application across a variety of hardware and software platforms
 - Optimizes storage resources for more efficient data protection over any distance
 - Significantly reduces cache utilization and increases bandwidth utilization by leveraging performance-optimized disk-based journals
 - Reduces overhead and application impact at production site by placing more of the workload on remote site
 - Centralizes operations for management resources and provides secure management of data-related operational risk
- Improve Operational Efficiency and Resiliency
 - Simplifies consolidation/aggregation and mapping of data value to the cost of storage
 - Supports planned site outages
 - Keeps logging changes in the event of network problems between sites
 - Reduces costs—requires only one product to provide asynchronous copy across all attached storage systems
- Synergy with Hitachi Business Continuity Framework
 - Builds on the data integrity heritage of Hitachi open-systems and mainframe remote replication software
 - Provides unified, simplified management via Hitachi HiCommand® Device Manager and Hitachi Business Continuity Manager software for IBM® z/OS®
 - Integrates tightly with other Hitachi software products supporting business continuity, further expanding capabilities

Business Solutions

Hitachi Data Systems and its Hitachi TrueNorth™ Channel Partners provide cost-effective storage products and solutions that leverage world-renowned Hitachi global R&D resources to deliver performance, availability, and scalability—supporting business-critical applications and strengthening competitive advantage.

- Complementary solutions for Universal Replicator software include:

- Hitachi HiCommand® Replication Monitor software
- Hitachi TrueCopy® Synchronous software, which duplicates data between like Hitachi storage systems
- Hitachi ShadowImage™ Heterogeneous In-System Replication software for non-disruptive, high-speed data replication within any Hitachi storage system

About Universal Replicator Operations

This chapter UR provides a RAID storage-based hardware solution for disaster recovery which enables fast and accurate system recovery. Once UR operations are established, duplicate copies of data are automatically maintained for backup and disaster recovery purposes:

- Functionality Overview
- UR Components
- UR Remote Copy Operations
- Journal Processing
- UR Delta Resync Operation
- Pair Status

Hitachi Universal Replicator represents a unique and outstanding disaster recovery solution for large amounts of data that span multiple volumes. The UR 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. The user-defined UR journal groups (volume groups) at the secondary site can be recovered with full update sequence consistency but behind the primary site due to asynchronous copy operations. This functionality also provides protection for write-dependent applications in the event of a disaster.

Functionality Overview

UR enables you to create duplicate volumes by copying data from the primary data volumes in the primary storage system to the secondary data volumes in the secondary storage system at the remote location. To perform this function, the journal obtain function at the primary site, the journal copy function between the primary and secondary sites, and the journal restore function at the secondary site are performed sequentially with the primary and secondary data volumes and the journal volumes. Write sequence consistency for the primary data volume at the primary site is also maintained for the secondary data volume at the secondary site by the write sequence number to be assigned for the journal data with the journal obtaining function, enabling you to configure the duplicate system which has data integrity. UR reduces the occurrence of pair suspensions due to the restrictions of data transfer from the primary site to the secondary site by storing the write data from the host in the master and restore journal volumes, providing a highly-reliability duplication system.

Figure 2-1 shows an overview of UR operations.

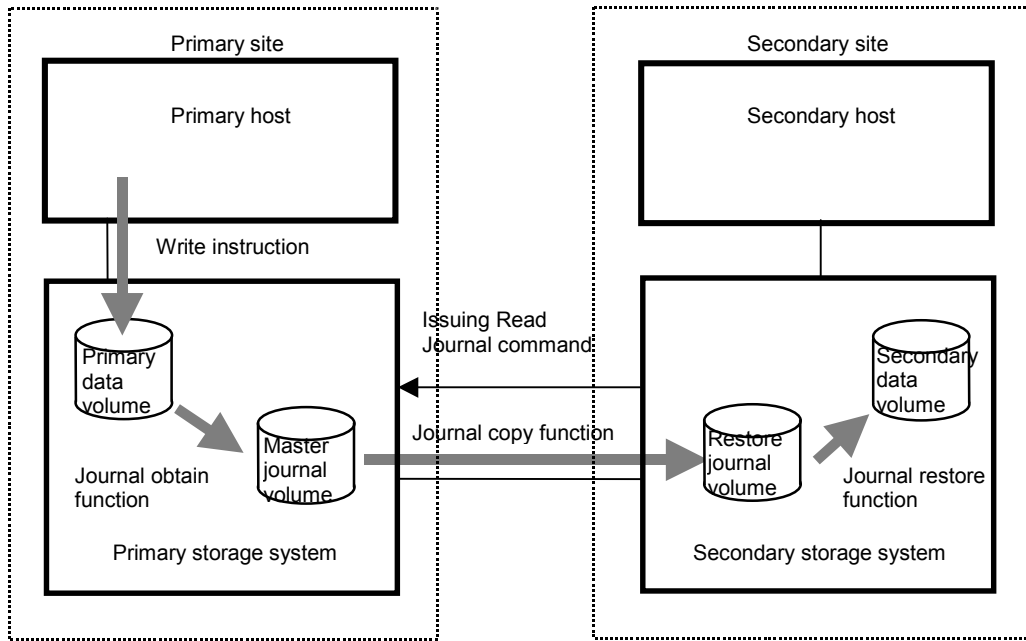


Figure 2-1 UR Components for Fiber-Channel Connection

Journal Obtain

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. And then, 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 secondary 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.

Journal Copy

Journal copy is the function to copy the data in the master journal volume at the primary site to the restore journal volume at the secondary site. The secondary storage system issues the read journal command to the primary storage system to request to transfer the data that is stored in the master journal volume according to the pair create or pair resync operation instruction from the primary site. The primary storage system transfers the data in the journal volume to the secondary site according to the read journal command if it has the journal data that should be sent. If the primary storage system does not have the journal data, the information is sent. The secondary storage system stores the journal volume data that is sent from the primary site in the restore journal volume at the secondary site. The read journal commands are issued repeatedly and regularly from the secondary site to the primary site until the journal operation is stopped. After the data are restored, the journal sequence numbers are informed from the secondary site to the primary site when the read journal command is issued. According to this information, the journal data at the primary site are discarded.

Journal Restore

Journal restore is the function to reflect the stored data in the restore journal volume to the secondary data volume at the secondary site. The data in the restore journal volume are restored to the secondary data volume according to the write sequence number. This will ensure the write sequence consistency between the primary and secondary data volumes. After the journal data are restored to the secondary data volume, the journal data are discarded at the secondary site.

UR Components

UR operations involve the *USP V* storage systems at the primary and secondary sites, the physical communications paths between these storage systems, and the *USP V* UR remote console software. UR copies the original online data at the primary site to the offline backup volumes at the secondary site via the dedicated Fiber-channel remote copy connections using a journal volume. You can operate the UR software with the user-friendly GUI environment using the *USP V* UR remote console software. **Note:** Host failover software is required for effective disaster recovery with UR. Figure 2-2 shows the UR components and their functions:

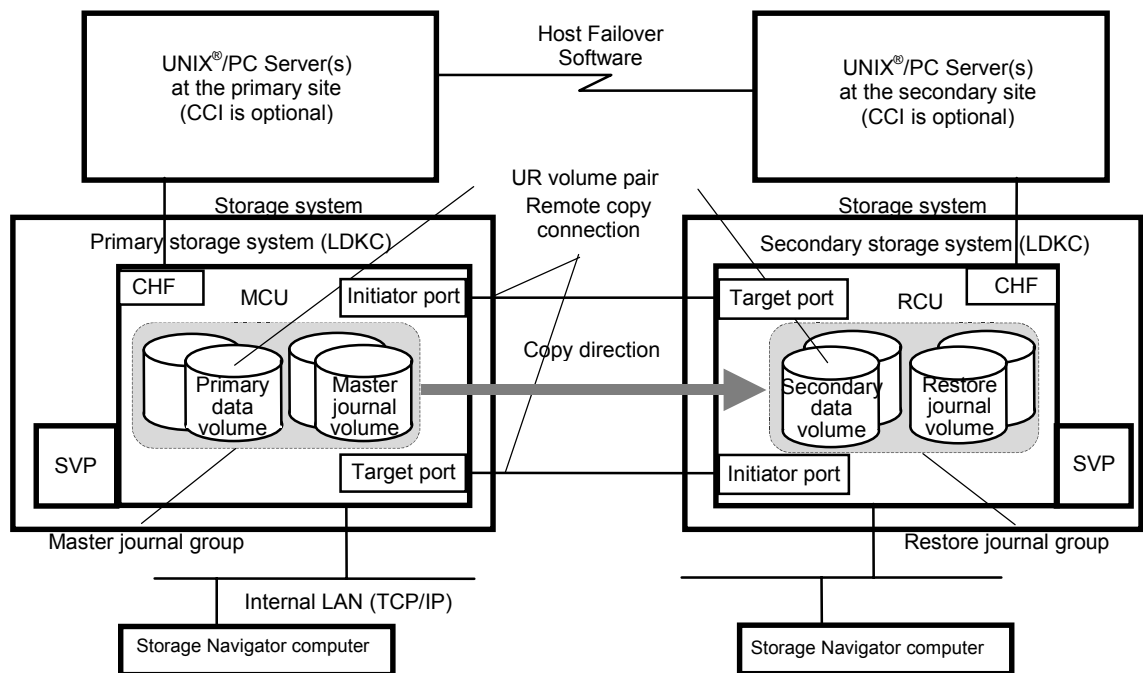


Figure 2-2 UR components

Table 2-3 shows the plural secondary storage systems connection configuration of UR. By connecting one primary storage system with more than one secondary storage systems, you can create a volume pair that has a one-to-one relationship for each journal group.

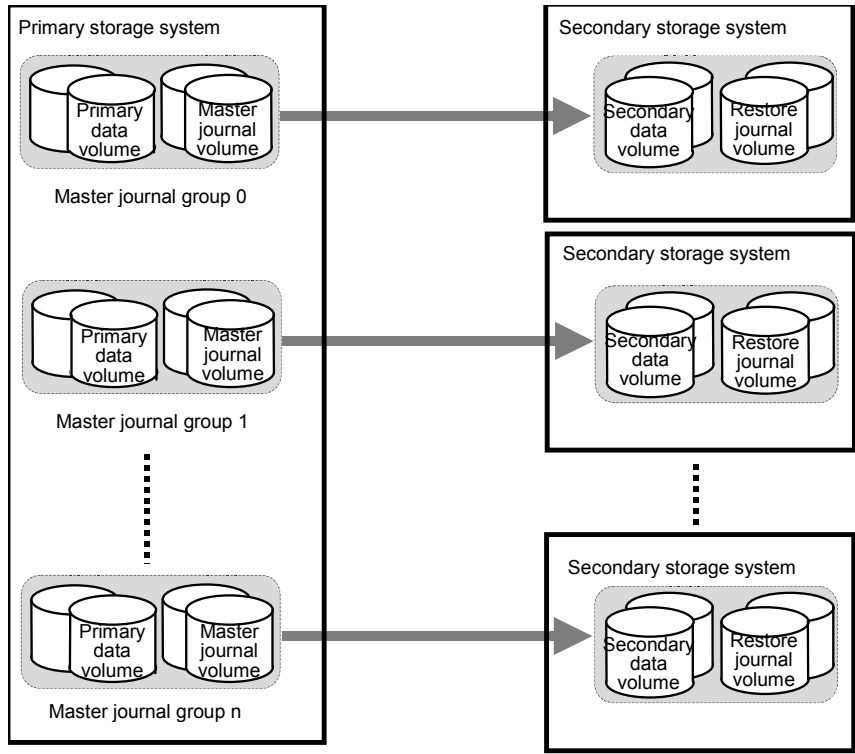


Figure 2-3 Connection Configuration of Plural Secondary Storage Systems

USP V Storage Systems

UR operations involve the USP V storage systems at the primary and secondary sites. The primary storage system consists of the main control unit (primary storage system) and SVP. The secondary storage system consists of the remote control unit (secondary storage system) and SVP.

To provide greater flexibility and to enable the USP V to be tailored to unique customer operating requirements, operational parameters, or optional modes, are available for the USP V storage system. At installation, the USP V modes are set to their default values, so make sure to discuss these settings with your Hitachi team. The USP V modes can only be changed by your Hitachi representative.

Universal Replicator Software

USP V Storage Navigator Java™ application program product includes UR for the USP V storage system. The 9900V Storage Navigator software communicates with the SVP of each USP V storage system via defined TCP/IP connections. For further information on USP V Storage Navigator operations, please refer to the *Storage Navigator User's Guide*, or contact your Hitachi account team.

The Storage Navigator PC at the primary site must be attached to the primary storage system. You should also attach a Storage Navigator PC at the secondary site to all secondary storage systems. Having a Storage Navigator PC at the secondary site enables you to change the UR parameter of the secondary storage system and access the UR secondary data volume (e.g. for the maintenance of media). If you need to perform UR operations in the reverse direction from the remote site to the primary site (e.g., disaster recovery), the *USP V* UR software simplifies and expedites this process.

Note: If the *USP V* Storage Navigator remote console PC is not installed, please contact your Hitachi account team for information on UR configuration services.

Logical Disk Controllers

CUs (control units) in the USP V storage system are divided into groups of 255 CUs. Each group of CUs is controlled by a virtual storage system that logically exists in USP V (logical storage system). This virtual storage subsystem is called a "logical DKC" or an "LDKC (Logical disk controller)". There are 2 LDKCs in the USP V storage system and number "00" or "01" is assigned to each LDKC.

Each LDKC controls 255 CUs. However the number of CUs that can be used for USP V program products is up to 255. Therefore, the maximum number of volumes that can be used for USP V program products is 130,560 (65,280 volumes for an LDKC).

Main and Remote Control Units (Primary and Secondary Storage Systems)

The main control unit (primary storage system) and remote control unit (secondary storage system) control UR operations:

- The primary storage system is the control unit in the primary storage system which controls the primary data volume of the UR pairs and master journal volume. The Storage Navigator remote console PC must be LAN-attached to the primary storage system. The primary storage system communicates with the secondary storage system via the dedicated remote copy connections. The primary storage system controls the host I/O operations to the UR primary data volume and the obtaining journal operation of the master journal volume as well as the UR initial copy and update copy operations between the primary data volumes and the secondary data volumes.
- The secondary storage system is the control unit in the secondary storage system which controls the secondary data volume of the UR pairs and restore journal volume. The secondary storage system controls copying of journals and restoring of journals to secondary data volumes. The secondary storage system assists in managing the UR pair status and configuration (e.g., rejects write I/Os to the UR secondary data volumes). The secondary storage system issues the read journal command to the primary storage system and executes copying of journals. The secondary Storage Navigator PC should be connected to the secondary storage systems at the secondary site on a separate LAN. The secondary storage systems should also be attached to a host system to allow sense information to be reported in case of a problem with a secondary data volume or secondary storage system and to provide disaster recovery capabilities.

The USP V can function simultaneously as a primary storage system for one or more primary data volumes and as an secondary storage system for one or more secondary data volumes, provided the remote copy connections and Fiber-channel interface ports are properly configured. The UR software allows you to specify the secondary storage system from the connected primary storage system. UR operations can be performed on all LDEVs except for the USP V command device. For further information on the USP V command device, please refer to the *Hitachi Command Control Interface User and Reference Guide*.

Note: When you configure an UR journal group pair, you have to specify the serial numbers of primary storage systems and secondary storage systems. You have to specify the different serial numbers of primary storage system and secondary storage system for the same UR journal group pair. If you have to specify the same serial number, please contact your Hitachi account team.

Consistency Groups

The concept of UR consistency group number equals to TrueCopy Asynchronous consistency group numbers. Therefore, you can create a UR pair with the consistency group numbers from CCI as is the case with TrueCopy Asynchronous. In UR operations, the consistency group numbers that are specified from CCI are assigned to the journal group numbers.

The specifications of consistency group number are different from the specifications of journal group number. One storage system has total 128 consistency groups for the primary and secondary data volumes (No. 0 to 127 are assigned).

Note: The consistency group number of the primary and secondary data volumes to be paired must be same.

Note: When UR and UR for z/OS® coexist in the same USP V storage system, the consistency group number must be assigned to the different number of the TrueCopy Asynchronous consistency group number..

Figure 2-4 shows the corresponding number of consistency group.

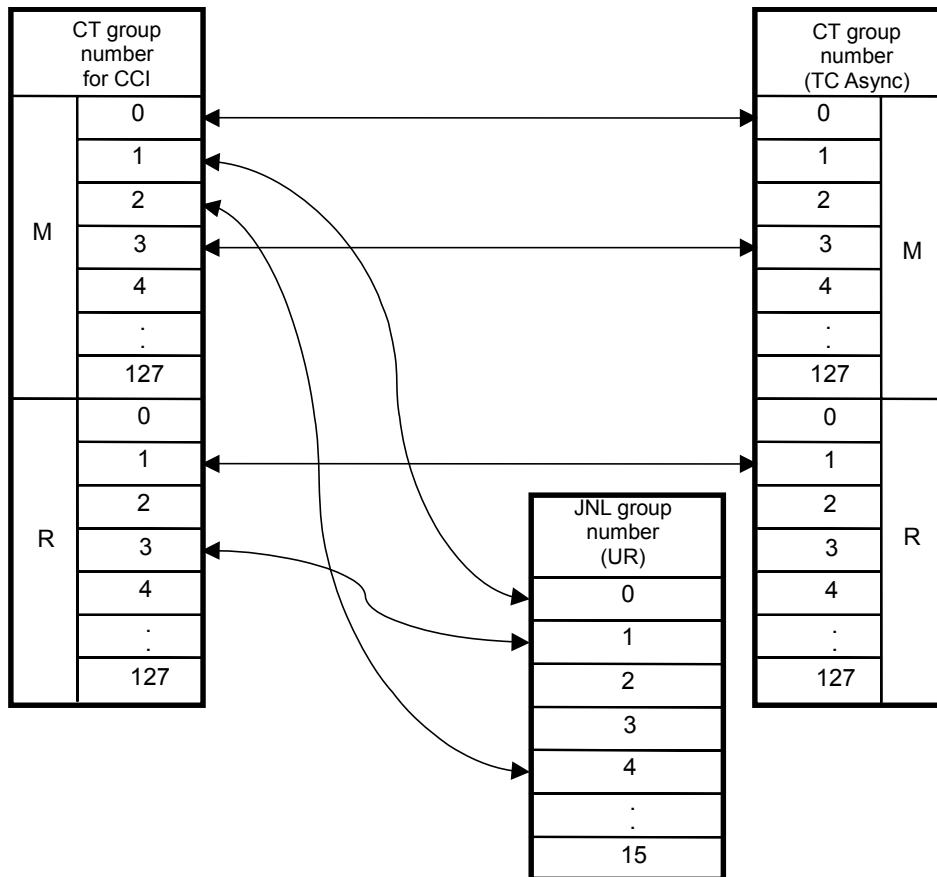


Figure 2-4 Corresponding Consistency Group Number and Journal Group Number

Remote Copy Connections

The remote copy connections are the physical paths used by the primary storage systems to communicate with the secondary storage systems. Remote copy connections enable communication between the primary and secondary storage systems. The primary storage systems and secondary storage systems are connected via Fiber-channel interface cables. You must establish paths from the primary to the secondary storage system, and also from the secondary to the primary storage system. Up to eight paths can be established in both of these directions.

When Fiber-channel interface (optical 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. If the distance between the primary and secondary sites is greater than 1.5 km, the optical single mode long wave interface connections are required. When Fiber-channel interface (single-mode long wave) 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.

See section Setting up Remote Copy Connections for further information on installing and configuring the FC remote copy connections.

The UR remote copy configuration between primary storage system and secondary storage system has the following requirement:

UR supports 1-to-1 remote copy connection in one journal group pair. In one journal group pair, one primary storage system can be connected to only one secondary storage system. This configuration ensures the backup data consistency of two or more volumes (e.g., large databases) within the same storage system.

Note: Hitachi strongly recommends that you establish at least two independent remote copy connections from the primary storage system to the secondary storage system and vice versa to provide hardware redundancy for this critical communications path.

Initiator Ports and RCU Target Ports

The initiator port and the RCU target port are required at both the primary storage system and secondary storage system. The initiator port at the primary storage system is connected to the RCU target port at the secondary storage system via the Fiber channel interface. The initiator port at the secondary storage system is connected to the RCU target port at the primary storage system. The initiator port at the secondary storage system issues a "read journal" command to the primary storage system, and then the RCU target port at the primary storage system sends journal data to the secondary storage system in response to the "read journal" command.

Any Fiber-channel interface port of the USP V can be configured as an initiator port. The initiator ports cannot communicate with the host processor channels. The host channel paths must be connected to the Fiber-channel interface port other than the initiator port.

Note: Two or more initiator ports must be configured before you can add the secondary storage systems and create the UR volume pairs.

The Fiber-channel interface ports that are assigned for the RCU target ports can be connected to the host channel paths via the Fiber-channel switch.

See section Configuring Port Attributes for the information on configuring host interface port.

Journal Groups

Journal group consists of two or more data volumes and journal volumes. To enable UR to be tailored to unique business content of user is available by using this journal group. The journal volume in the primary storage system is referred to as the master journal volume. The journal volume in the secondary storage system is referred to as the restore journal volume. The data volume in the master journal group is also called the primary data volume. The journal volume in the master journal group is called the master journal volume. The data volume in the restore journal group is similarly called the secondary data volume. The journal volume in the restore journal group is called the restore journal volume.

The data update sequence from the host is managed per the journal group. The data update sequence consistency between the master and restore journal groups to be paired is maintained and ensured. The master and restore journal groups are managed according to the journal group number. The journal numbers of master and restore journal groups that are paired can be different. One data volume and one journal volume can belong to only one journal group.

Caution: Data volumes and journal volumes that belong to different LDKC cannot coexist in a journal group.

For detailed information about the specification of journal groups, see Table 3-7.

Data Volume Pairs

UR performs remote copy operations for data volume pairs created by the user. Each UR pair consists of one primary data volume and one secondary data volume which can be located in different storage systems. The UR primary data volumes are the primary volumes (LUs) which contain the original data, and the UR secondary data volumes are the secondary volumes (LUs) which contain the backup or duplicate data. During normal UR operations, the primary data volume remains available to all hosts at all times for read and write I/O operations. During normal UR operations, the secondary storage system rejects all host-requested write I/Os for the secondary data volume. The secondary data volume write enable option allows write access to a secondary data volume while the pair is split and uses the secondary data volume and primary data volume track maps to resynchronize the pair (see section Secondary Data Volume Write Option).

UR also supports LUN Expansion (LUSE) Virtual LVI/LUN (VLL) and Cache Residency Manager features, so that UR meets a variety of user needs and facilitates data copying and data migration. This ensures that all user data can be backed up or duplicated. See section Logical Unit (LU) Types for further information on LU requirements and support.

Journal Volumes

When UR is used, updates to primary data volumes can be stored in other volumes, which are called *journal volumes*. The updates (which are sometimes referred to as *update data*) that will be stored in journal volumes are called *journal 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 secondary storage system fails temporarily, remote copy operations can continue after 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 secondary storage system, remote copy operations between the primary storage system and the secondary storage system can continue. Because journal volumes can contain a lot more *update data* than the cache memory can contain, 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 secondary storage system.

The Number of Journal Volumes

One journal group can contain up to 64 journal volumes. Each of the journal volumes can have different volume sizes and different RAID configurations. Journal data will be stored sequentially and separately into each journal volume in the same journal group.

Specifications of Journal Volumes

Types of logical units (LUs):

The following DKU emulation types are allowed for journal volumes:

Table 2-1 Emulation Types for Journal Volumes

Emulation Category	Supported Emulation Types
DKU (drive)	OPEN-V Note: You can only use OPEN-V volumes for journal volumes.

Volumes and their capacity:

You can use VLL volumes for journal volumes.

Journal volumes in the same journal group can be of different capacity. A master journal volume and the corresponding restore journal volume can be of different capacity.

A journal volume consists of two areas: one area is used for storing journal data, and the other area is used for storing metadata for remote copy.

RAID configuration:

Journal volumes support all RAID configurations that are supported by USP V. Journal volumes also support all physical volumes that are supported by USP V.

Support for program products:

The volumes that are set Cache Residency Manager can be used for journal volumes.

Caution: Volumes containing a VMA (volume management area) cannot be used as journal volumes. For detailed information about a VMA, please refer to the *Data Retention Utility User's Guide* or for the latest information on the availability of the delta resync function, please contact your Hitachi Data Systems account representative.

Restrictions on Journal Volume

Registering journal volumes

Caution: You must register journal volumes in a journal group before you create a data volume pair for the first time in the journal group.

You can add journal volumes under any of the following conditions:

- When the journal group does not contain data volumes (i.e., before you create a data volume pair for the first time in the journal group, or after all data volume pairs are released)
- When all data volume pairs in the journal group are suspended.
- When processing for changing the status of a data volume pair (for example, release or suspension of a data volume pair) is not in progress

Note: If a path is defined from a host to a volume, you cannot register the volume as a journal volume.

You can use Storage Navigator computers to register journal volumes.

If you add a journal volume when a remote copy operation is in progress (i.e., when at least one data volume pair exists for data copying), the metadata area of the journal volume (see section Journal Volume Areas) 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. To check the journal data usage rate, use the Usage Monitor window (see Usage Monitor Window).

Deleting journal volumes:

You can delete journal volumes under any of the following conditions:

- When the journal group does not contain data volumes (i.e., before you create a data volume pair for the first time in the journal group, or after all data volume pairs are released)
- When all data volume pairs in the journal group are suspended.

You can use Storage Navigator computers to delete journal volumes.

Access from hosts to journal volumes:

If a path is defined from a host to a volume, you cannot register the volume as a journal volume.

You cannot define paths from hosts to journal volumes. This means that hosts cannot read from and write to journal volumes.

Journal Volume Areas

The journal volume consists of the metadata area and the journal data area. The ratio of metadata area to journal data area is common in the journal volumes within the journal group.

In the metadata area, the metadata that manages the journal data is stored. For further information on the metadata area, see Table 2-2. The journal data that the metadata manages is stored in the journal data area.

Note: If the metadata or the journal data cannot be stored for a given length of time because the metadata or journal data areas have become full with the metadata or the journal data that had not been discarded, the pair is suspended according to a failure. Users can use a Storage Navigator computer to specify this timeout period (Data overflow watch) as a journal group option. This timeout period must be within the range of 0 to 600 seconds. For details on journal group options, see section Changing Options for a Journal Group.

Host Failover Software

Host failover software, which transfers information between host servers at the primary and remote sites, is a critical component of any disaster recovery effort. Host failover is configured using the desired host failover software product for the platform (e.g., Microsoft Cluster Server), depending on your installation requirements and standards. The *USP V* Storage Navigator software products do not provide any host failover functions.

When UR is used as a data migration tool, host failover is recommended but not required. When UR is used as a disaster recovery tool, host failover is required to ensure effective disaster recovery operations. When a UR pair is suspended due to an error condition, the primary storage system generates sense information which should be transferred to the remote site via the host failover software for effective disaster detection and recovery.

UR Remote Copy Operations

Figure 2-5 illustrates the two types of UR remote copy operations: initial copy and update copy.

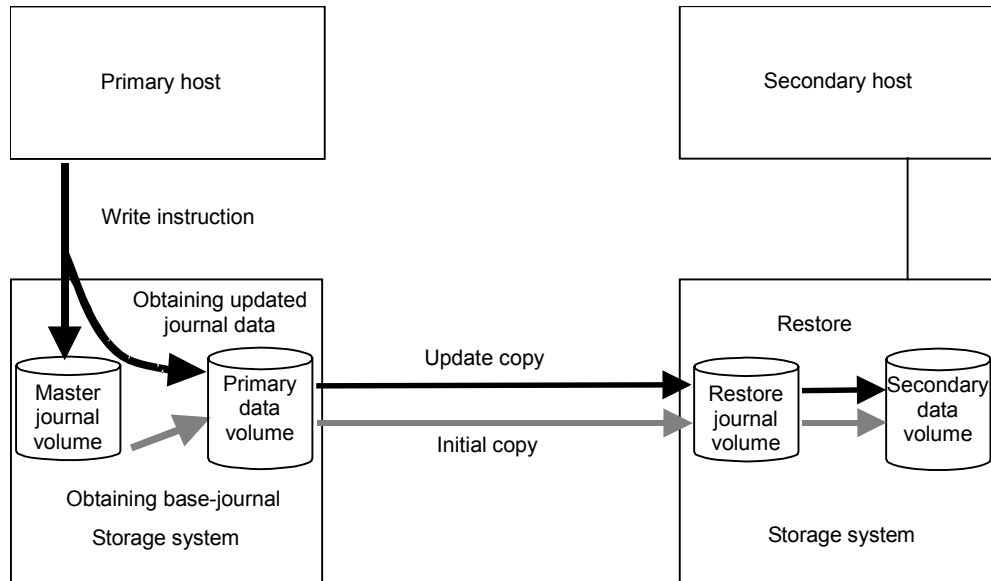


Figure 2-5 Remote copy operations

This section describes the following topics that are related to remote copy operations with UR.

Initial Copy Operation

Initial copy operations synchronize data in the primary data volume and data in the secondary data volume. Initial copy operations are performed independently from host I/Os. Initial copy operations 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, and then restores the base-journal to the secondary data volume.

If the journal-obtain operation starts at the primary data volume, the primary storage system obtains 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 will be copied from the primary storage system to the secondary storage system after the secondary storage system issues a read-journal command. After a base-journal is copied to the secondary storage system, the base-journal will be stored in a restore journal volume in a restore journal group where the secondary data volume belongs. After that, the data in the restore journal volume will be restored to the secondary data volume, so that the data in the secondary data volume synchronizes with the data in the primary data volume.

The base-journal data is stored in the entire data volume or the area for the difference. The area for the difference is used when the difference resynchronization operation is performed. The journal data for the entire data volume is created when the data volume pair is created. The difference journal data is obtained when the pair status of the data volume changes from the **Suspending** status to the **Pair resync** status. Merging the difference bitmaps that are recorded on both primary and secondary data volumes enables you to obtain the journal data for only difference. When a data volume pair is suspended, the status of data that is updated from the host to the primary and secondary data volumes is recorded to the difference bitmap.

The base-journal data of primary storage system is stored to the secondary storage system journal volume according to the read command from the secondary storage system. After that, the base-journal data is restored from the journal volume to the secondary data volume. The initial copy operation will finish when all base-journals are restored.

Note: If you manipulate volumes (not journal groups) to create or resynchronize two or more data volume pairs within the same journal group, the base journal of one of the pairs will be stored in the restore journal volume, and then the base journal of another pair will be stored in the restore journal volume. Therefore, the operation for restoring the latter base journal will be delayed.

Note: You can specify **None** as the copy mode for initial copy operations. If the **None** mode is selected, initial copy operations will not be performed. The **None** mode must be used at your responsibility only when you are sure that data in the primary data volume is completely the same as data in the secondary data volumes.

Update Copy Operation

When a host performs a write I/O operation to a primary data volume of a data volume pair, an update copy operation will be performed. During an update copy operation, the update data that is written to the primary data volume is obtained as an *update journal*. The update journal will be copied to the secondary storage system, and then restored to the secondary data volume.

The primary storage system obtains update data that the host writes to the primary data volume as *update journals*. Update journals will be stored in journal volumes in the journal group that the primary data volume belongs to. When the secondary storage system issues "read journal" commands, update journals will be copied from the primary storage system to the secondary storage system asynchronously with completion of write I/Os by the host. Update journals that are copied to the secondary storage system will be stored in journal volumes in the journal group that the secondary data volume belongs to. The secondary storage system will restore the update journals to the secondary data volumes in the order write I/Os are made, so that the secondary data volumes will be updated just like the primary data volumes are updated.

Journal Group Operations

UR 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 secondary 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 volumes.

Journal Group Operations

UR provides the following group-based operations to simplify and expedite disaster or failure recovery procedures:

- Group operations at the primary storage system
 - Split all pairs in a journal group. See section Splitting a Pair of Data Volumes (Pairsplit-r) for a description of the Suspend Range-Group pairsplit option.
 - Resume all suspended pairs in a journal group. See section Restoring a Mirror (Pairresync: Resuming a copy operation) for a description of the UR Resume Range-Group pairresync option.

- Release all pairs in a journal group. See section Deleting Data Volumes from a Mirror (Pairsplit-S: Ending a copy operation) for a description of the Delete Range-Group pairsplit -S option.
- Group operations at the secondary storage system
 - Split (pairsplit-r) all pairs in a journal group. See section Splitting a Mirror (Pairsplit-r: Suspending a copy operation) for a description of the Suspend Range-Group pairsplit-r option.
 - Release (pairsplit-S) all pairs in a journal group regardless of their consistency status. See section Deleting Data Volumes from a Mirror (Pairsplit-S: Ending a copy operation) for a description of the Delete Range-Group pairsplit-S option.

Read and Write I/O Operations During UR Operations

When a primary storage system receives a read I/O for an UR primary data volume, the primary storage system performs the read from the primary data volume. If the read fails, the redundancy provided by RAID-1 or RAID-5 technology recovers the failure. The primary storage system does not read the UR secondary data volume for recovery.

When a primary storage system receives a write I/O for the primary data volume with PAIR status, the primary storage system performs the update copy operation, as well as writing to the primary data volume.

The primary storage system completes the primary data volume write operations independently of the update copy operations at the secondary data volume. The secondary storage system updates the data in the secondary data volume according to the write sequence number of journal data. This will maintain the data consistency between the primary and secondary data volumes. If the primary data volume write operation fails, the primary storage system reports a unit check and does not create the journal data for this operation. If the update copy operation fails, the secondary storage system suspends either the affected pair or all UR pairs in the journal group, depending on the type of failure. When the suspended UR pair or journal group is resumed (pairresync), the primary storage system and secondary storage system negotiate the resynchronization of the pair(s). See section Suspend Types for further information on UR suspend types.

During normal UR operations, the secondary storage system does not allow UR secondary data volumes to be online (mounted), and therefore hosts cannot read from and write to secondary data volumes. The UR secondary data volume write enable option allows write access to a secondary data volume while the pair is split (see section Secondary Data Volume Write Option). The secondary data volume write option can only be enabled when you split the pair from the primary storage system.

Note: To reduce the overhead associated with these remote copy activities and maximize data transfer, the *USP V* storage system utilizes a special write command which is allowed only for UR initial and update copy operations. This command transfers the control parameters and the FBA-format data for consecutive updated records in a track using a single write operation. The special UR write command eliminates the overhead required for performing FBA-to-CKD and CKD-to-FBA conversions.

Secondary Data Volume Write Option

For additional flexibility, UR provides a secondary data volume write option (S-Vol. Write) which enables write I/O to the secondary data volume of a split UR pair. The secondary data volume write option can be selected by the user during the pairsplit-r operation and applies only to the selected pair(s). The secondary data volume write option can be accessed only when you are connected to the primary storage system. When you resync a split UR pair which has the secondary data volume write option enabled, the secondary storage system sends the secondary data volume track bitmap to the primary storage system, and the primary storage system merges the primary data volume and secondary data volume bitmaps to determine which tracks are out-of-sync. This ensures proper resynchronization of the pair.

Difference Management

The differential data (updated by write I/Os during split or suspension) between the primary data volume and the secondary data volume is stored in each track 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 storage system. For details on the maximum possible number of pairs, see section The Maximum Number of Pairs.

Journal Processing

The UR journal data contains the primary data volume updates and the metadata information (associated control information), which enables the secondary storage system to maintain update consistency of the UR secondary data volumes. UR journal processing includes:

- Creating and storing journals at the primary storage system
- Copying journals to the secondary storage system
- Storing journals at the secondary storage system
- Selecting and restoring journals at the secondary storage system
- Types of journals

Journal Processing at the Primary Storage System

When a primary storage system performs an update (host-requested write I/O) on a UR primary data volume, the primary storage system creates a journal data to be transferred to secondary storage system. The journal data will be stored into the cache at first, and then into the journal volume.

Metadata information will be attached to journal data (see Table 2-2). When base-journal is obtained, only metadata information is created and stored in UR cache or the journal volume.

Table 2-2 Metadata Information

Type	Description
Journal type	Type of journal (e.g., base-journal or update journal)
LDEV No. (data)	The number of primary data volume that stores the original data
Original data storing position	The primary data volume slot number, and the start and end of sub-block number (data length)
LDEV No. (journal)	The volume number of master journal volume that stores the journal data
Journal data storing position	The slot number of master journal volume, and the start sub-block number
Journal sequence number	The sequence number that is assigned when the journal is obtained
Timestamp	The time when the journal data is obtained

The journal sequence number indicates the primary data volume write sequence that the primary storage system has created for each journal group. The journal data is transferred to the secondary storage system asynchronously with the host I/O. The secondary storage system updates the secondary data volume in the same order as the primary data volume according to the sequence number information in the journal.

Note: UR processing continues uninterrupted if the SVP reboots or even if the SVP fails.

Types of Journal

In addition to the journal data for updating, the primary storage system sends control information to the secondary storage system. This control information indicates when volume pair status changes and when a primary storage system power-off sequence is initiated, and also maintain sequence numbers in periods of low host activities.

Journal Processing at the Secondary Storage System

When a primary storage system receives a read journal command from a secondary storage system, the primary storage system sends the journal data to the secondary storage system. The secondary storage system's initiator ports act as host processor channels and issue special I/O operations, called remote I/Os (RIOs), to the primary storage system. The RIO transfers the journal data in FBA format using a single channel command. The primary storage system can send several journal data using a single RIO, even if their sequence numbers are not contiguous. Therefore, the journal data are usually sent to the secondary storage system in a different order than the journal data were created at the primary storage system. The secondary storage system ensures that the journal data are applied to the secondary data volume in the correct sequence. This method of remote I/O provides the most efficient use of primary storage system-to-secondary storage system link resources.

Note: You must make sure that your channel extenders are capable of supporting remote I/O. For further details, please contact your Hitachi account team.

Storing Journal at the Secondary Storage System

A secondary storage system receives the journal data that is transferred from a primary storage system according to the read journal command. The journal data will be stored into the cache at first, and then into the journal volume.

Note: The primary storage system does not remove the target journal data from its master journal volume until it receives the sequence numbers of restored journal which is give to the read journal command from the secondary storage system. This is true even if the primary storage system and secondary storage system are connected via a channel extender product.

Selecting and Restoring Journal at the Secondary Storage System

The secondary storage system selects journal data to be promoted to formal data (or "restored") as follows:

1. The secondary storage system gives the number as the management information to distinguish the journal data arrival to the sequence number that is assigned to the journal data from the primary storage system. If the number is 1, the journal data arrived at the secondary storage system. If the number is 0, the journal data has not arrived yet. The secondary storage system determines whether the journal data should be settled or not according to this number. If the journal data has not arrived yet, the secondary storage system waits for the journal data.
2. When the top of queue in the journal group indicates the journal data arrival, the secondary storage system selects the journal data which has the lowest sequence number, and then settles this journal data.
3. The secondary storage system repeats steps (1) and (2) to select and settle the journal data.

Figure 2-6 illustrates the journal data selection and settling at the secondary storage system. This diagram shows that journal data S1 arrives at the secondary storage system because the management information indicates 1. The secondary storage system selects journal data S1 to be settled, because S1 is the lowest sequence number. When S1 is removed from the queue of sequence numbers, journal data S2 becomes the top entry, but it has not arrived yet. The management information of journal data S2 is 0. The secondary storage system waits journal data S2. When journal data S2 arrives, the secondary storage system selects S2 as the next journal data to be settled. The journal data selected by the secondary storage system is marked as "host-dirty" and treated as formal data.

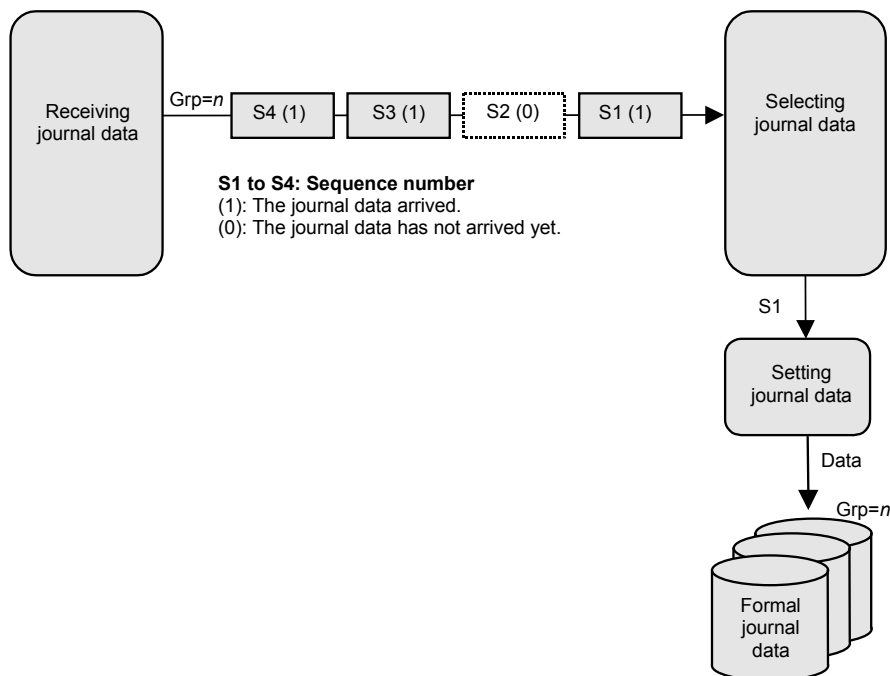


Figure 2-6 Selecting and Settling Journal at the Secondary Storage System

The secondary storage system settles and restores the journal data to the secondary data volume as follows:

- Journal data stored in the cache

The journal data is copied to the corresponding cached track and promoted to formal data.

- Journal data stored in the restore journal volume

The journal data is read from the restore journal volume to cache. The journal data that is read to cache is copied to the existing cache track and promoted to formal data. After that, the space for the restore journal volume is released.

UR Delta Resync Operation

When you are using UR and TrueCopy Synchronous in a 3DC multi-target configuration, UR provides delta resync operation as one of the solutions for failures in primary site. In a 3DC multi-target configuration, there are one primary site and two secondary sites; TrueCopy Synchronous and UR secondary sites. For detailed information about 3DC multi-target configuration, see section TrueCopy Synchronous (3DC Cascading Configuration).

If a failure occurs on the primary site in 3DC multi-target configuration, you need to use CCI **horctakeover** command to use the TrueCopy Synchronous secondary site as the primary site. If you perform a delta resync operation after the TrueCopy Synchronous secondary site becomes a primary site. The UR pair will be restored quickly by the delta resync operation, you will not need to wait for a long time before you can use the UR data volumes again.

Delta resync operation consists of the two processings; one is the processing for the preparation before the failure occurs, the other is the processing for the recovery after the failure occurs.

- Processing for the preparation before the failure occurs (see the next section)
- Processing for the recovery after the failure occurs (see section Switching the Master Journal Group of UR)

Note: For the latest information on the availability of the delta resync function, please contact your Hitachi Data Systems account.

Journal Obtain in TrueCopy Synchronous Secondary Site

To perform delta resync operation when a failure occurs, you also need to obtain the journal data in the TrueCopy Synchronous secondary site of 3DC multi-target configuration. Specify the TrueCopy Synchronous S-VOL in TrueCopy Synchronous secondary site as the primary data volume, and specify the data volume in UR secondary site as the secondary data volume, in order to create a UR pair for the delta resync operation.

When you create a UR pair for delta resync operation, the differential data of data volumes in TrueCopy Synchronous primary site and secondary site will be stored in the journal volumes in TrueCopy Synchronous secondary site as journal data. The following figure shows an example.

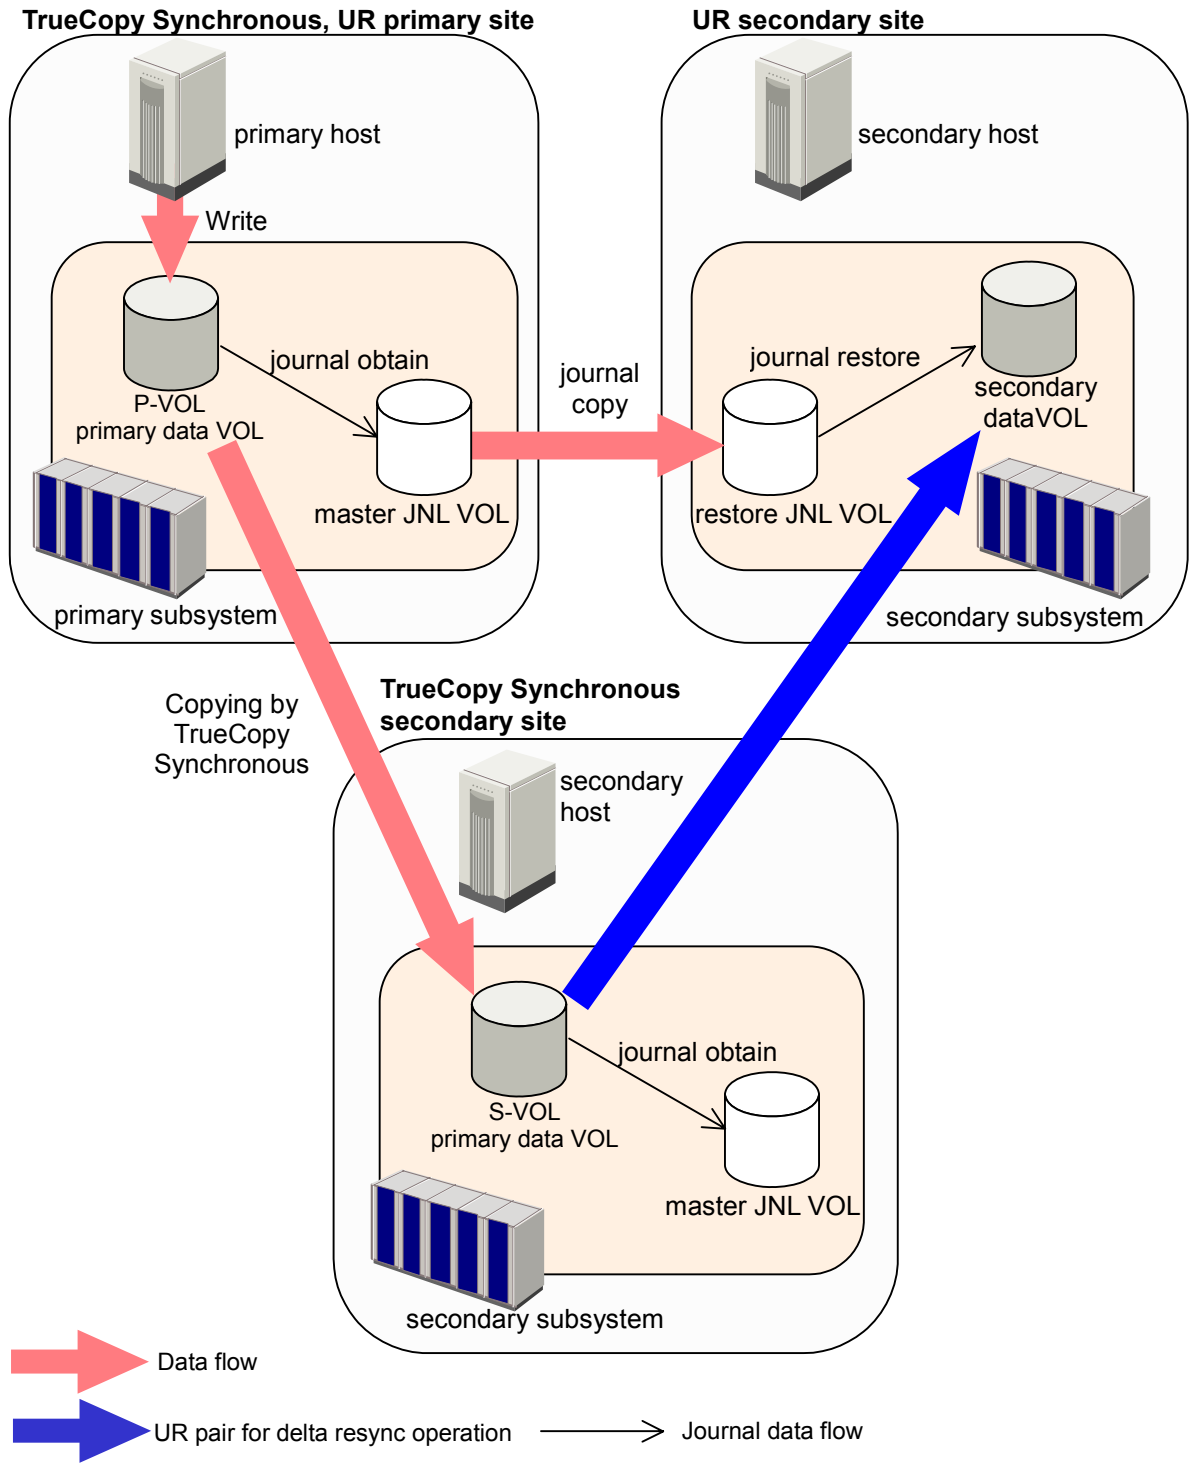


Figure 2-7 Delta Resync Setting 3DC Multi-target Configuration - Before Failure Occurs

As shown in Figure 2-7, a UR pair created with the delta resync option is defined as a pair but no copy operation is performed (HOLD status). Actual copy operation will not be performed until when the failure occurs and delta resync operation is performed. Note that there are several requirements to create a UR pair for delta resync operation, such as you need to specify the unused mirror ID. For detailed information about the requirements of creating a UR pair for delta resync operation, see section Requirements for Creating UR Pair for Delta Resync Operation.

For the information about the delta resync operation that will be performed when a failure occurs in the configuration shown in Figure 2-7, see the next section. Note that the UR pair need to be in HOLD status if you want to perform the delta resync operation when the failure occurs. However, the UR pair status may be changed to HLDE for example when the cache memory or shared memory error occurs in TrueCopy Synchronous secondary site, or when no journal cannot be obtained in TrueCopy Synchronous secondary site because of the failure in the master journal volume or occurrence of the pinned track. If the status of the UR pair for delta resync operation changes to HLDE, follow the steps in section Restoring a Pair of Data Volumes (Pairresync) and change the pair status to HOLD again.

Switching the Master Journal Group of UR

When a failure occurs on the primary site in Figure 2-7 (3DC multi-target configuration), the UR pair for delta resync operation can use the journal group in TrueCopy Synchronous secondary site as the master journal group. To switch the master journal group, first change TrueCopy Synchronous secondary site to the primary site by CCI's **horctakeover** command, then perform the delta resync operation on the primary data volume of the UR pair in HOLD status. The following figure shows an example.

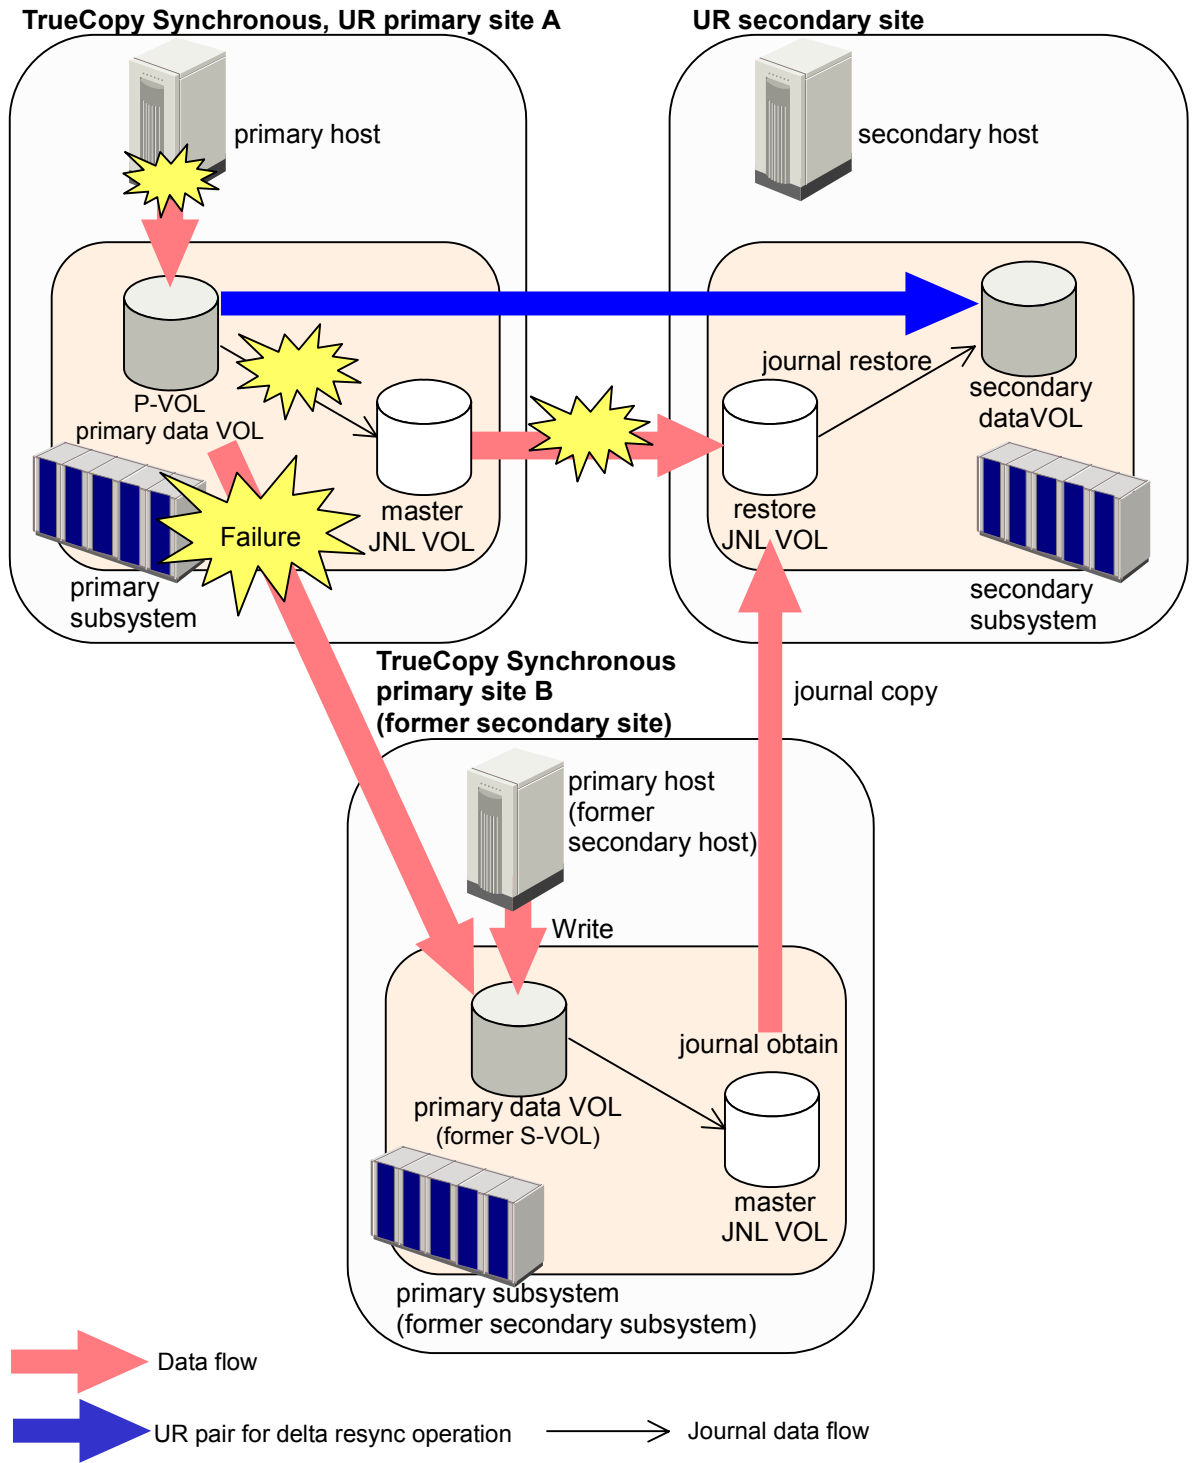


Figure 2-8 Delta Resync Setting 3DC Multi-target Configuration - After Failure Occurs

Note: For the latest information on the availability of the delta resync function, please contact your Hitachi Data Systems account.

In Figure 2-8, because a failure occurs in the primary site A, the **horctakeover** command is used to change the former TrueCopy Synchronous secondary site to the primary site B. If you perform delta resync operation in this situation, the UR pair for delta resync operation in Figure 2-7 will be synchronized and usable.

When you perform delta resync operation, first the journal data in the primary site B are copied to the UR secondary site by journal copy. In this journal copy, only the journal data which is not yet restored to the secondary data volume in the UR secondary site are copied in chronological order. When the journal copy completes, journal restore takes place in the UR secondary site.

In delta resync operation, the status of the UR pair will not change to COPY but PAIR. This is because the delta resync operation copies only the really necessary data by journal copying. Compared to the usual way which copies the whole data in the primary data volume, delta resync operation requires shorter time for the recovery of the UR pair after a failure occurs.

Note that when the total capacity of stored journal data exceeds 80% of the journal volume of TrueCopy Synchronous secondary site, old journal data will be deleted automatically. Therefore, if the total capacity of the journal data which is not restored to the UR secondary data volume exceeds 80% of the journal volume, the secondary data volume will not be able to be restored completely by copying the journal data to the restore journal volume in the UR secondary site. In that case, according to the setting of the journal group option, whole data in the primary data volume will be copied to the secondary data volume, 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 which is not restored to the UR secondary site will not exceed 80% of the journal volume. Though, for example if the UR 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 UR secondary data volume. In such case, note that 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 updated the P-VOL

When you restore the UR pair between the primary site and the UR secondary site, then updated 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

As shown in Figure 2-8, after delta resync operation is performed properly and the primary site A is recovered from the failure, then the UR pair between the primary site A and the UR secondary site will be the pair for delta resync operation and become prepared for the failure in the primary site B. For information about the requirements to perform delta resync operation properly, see section Requirements for Performing Delta Resync Operation.

Pair Status

UR displays the pair status for each data volume of specified CU Image (CUI) in the connected *USP V* storage system. Table 2-3 lists and describes the UR data volume pair status descriptions. The primary storage system maintains the status of the primary data volume and can change the pair status of the primary data volume and secondary data volume. The secondary storage system maintains the status of the secondary data volume and can change the pair status of the primary data volume and secondary data volume but not the primary data volume. The primary storage system will detect when the secondary storage system changes the secondary data volume status (if the path status is normal) and will change the primary data volume status accordingly. You can display the detailed pair status information at the Storage Navigator remote console PC (UR Detailed Information window) or at the UNIX[®]/PC server host (CCI Pairedisplay command).

A volume which is not assigned to a UR data volume pair has the status **SMPL**. When a UR data volume pair is started, the primary storage system changes the status of the primary data volume and secondary data volume to **COPY**. When the initial copy operation is complete, the primary storage system changes the status of both data volumes to **PAIR**. When a pair is suspended due to an error condition, the primary storage system changes the primary data volume and secondary data volume status to **PSUE** (if the path status is normal). When a UR pair is split by the user (pairsplit-r), the primary storage system or secondary storage system changes the status of the primary data volume and secondary data volume to **PSUS** (if the path status is normal). When a pair is split from the secondary storage system, the secondary storage system changes the secondary data volume status to **PSUS**, and the primary storage system detects the pair deletion (if path status is normal) and changes the primary data volume status to **PSUS**. When a pair is released from the primary storage system, the primary storage system changes the status of both data volumes to **SMPL**. When a pair is released from the secondary storage system, the secondary storage system changes the secondary data volume status to **SMPL**, and the primary storage system detects the pair release (if path status is normal) and changes the primary data volume status to **PSUS**.

When a UR data volume pair is split or suspended, the primary storage system generates a service information message (SIM) to notify the host(s). If SNMP is installed and operational for *USP V*, this SIM results in an SNMP trap which indicates the reason for suspension.

UR Pair Status

The UR Suspending and Deleting (releasing) transitional states occur when a request to change UR pair status has been accepted, but the change to the requested status (PSUS, PSUE, or SMPL) is not yet complete. These states are not reported to the host. In the case of Suspending, both the user and the primary storage system can request the status change. In the case of Deleting (releasing), only the user can request the status change. If the user requested the status change, the final status is reported at the end of the transition. If an error caused the status to change to PSUE, the suspended status is reported at the beginning of the transition.

After a storage system receives a request for splitting or releasing a pair in **Flush** mode, the status of the pair will remain **Suspending** or **Deleting** until the journal in the master journal group is restored into the restore journal group and the pair is completely split or released. To calculate the time during which the pair remains **Suspending** or **Deleting**, use the following equation:

$$C \times (u \div 100) \times 1,024 \div V \text{ (The unit is seconds)}$$

where:

C is the total capacity of the master journal volume. The unit is megabytes.

u is the usage rate of data in the master journal volume. The unit is percent.

V is the data transfer speed between the primary and the secondary storage system. The unit is MB/s (megabytes per second).

To find the usage rate of a journal volume, use the monitoring feature (see Usage Monitor Window).

This UR pair status describes:

UR split type (PSUS) (see section Split Types (PSUS)),

UR suspend type (see section Suspend Types),

UR suspend condition (see section Suspension Condition).

Table 2-3 UR Data Volume Pair Status

Pair Status	Description	Primary Data Volume Access	Secondary Data Volume Access
SMPL (simplex)	This volume is not currently assigned to a UR data volume pair. This volume does not belong in the journal group. When this volume is added to a UR 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	Read Only
PAIR	This data volume pair is synchronized. Updates to the primary data volume are duplicated on the secondary data volume.	Read / Write	Read Only

<p>PSUS (pair suspended-split) (see Table 2-4 for split types)</p>	<p>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 secondary storage system (pairsplit-S). For UR pairs, the primary storage system and secondary 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 secondary storage system keep track of the primary data volume and secondary data volume tracks which are updated.</p> <p>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 secondary storage system, the secondary 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.</p> <p>When you release a pair from the secondary storage system, the secondary 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.</p>	<p>Read / Write</p>	<p>Read Only; Read and write, if write option is enabled ("Disable" is set for the default setting).</p>
<p>PSUE (pair suspended-error) (see Table 2-6 for suspend types)</p>	<p>This data volume pair is not synchronized, because the primary storage system or secondary storage system has suspended the pair due to an error condition. For UR pairs the primary storage system and secondary 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.</p> <p>For UR pairs, if the primary storage system detects a UR suspension condition (see section Suspension Condition), the primary storage system changes the secondary data volume status to PSUE.</p> <p>For UR pairs, if the primary storage system detects a UR suspension condition, the secondary 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.</p>	<p>Read/write; Read only if fenced.</p>	<p>Read Only</p>
<p>PFUL</p>	<p>UR 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).</p> <p>Note: The PFUL status is displayed by CCI (Command Control Interface). Storage Navigator displays this status as PAIR.</p>	<p>Read / Write</p>	<p>Read Only</p>

PFUS	If the time period exceeds the specified monitoring time period, the pair status changes from PFUL to PFUS, and the pair is suspended. Note: 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 ("Disable" is set for the default setting).
SSWS	The data can be written into the secondary data volume that is reassigned from the primary data volume during processing of resynchronization (Takeover). Note: The SSWS status is displayed by CCI (Command Control Interface). Storage Navigator displays this status as PSUS or PSUE.	Read Only	Read / Write
Suspending	This pair is not synchronized. 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	Read Only
Deleting (releasing)	This pair is not synchronized. 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 (releasing). When the pairsplit-S operation is complete, the status changes to SMPL.	Read / Write	Read Only
HOLD	The pair is prepared for delta resync operation. When the status of primary data volume is HOLD, the write data for the TrueCopy Synchronous S-VOL is stored in the master journal volume. Only the delta resync operation, releasing operation, or changing pair option operation are allowed on the pairs in HOLD status.	Read / Write*1	Read / Write*2
HLDE	An error occurred on the pair in HOLD status. When the status of primary data volume is HLDE, the write data for the TrueCopy Synchronous S-VOL will not be stored in the master journal volume. Only recovering pair status to standby (HOLD) operation, releasing operation, or changing pair option operation are allowed on the pairs in HLDE status.	Read / Write*1	-

*1: Whether you can or cannot access the primary data volume depends on the status of the TrueCopy Synchronous pair which shares the volume.

*2: The secondary data volume of UR pair in HOLD status uses two mirrors. Whether you can or cannot access the secondary data volume depends on the status of the mirror whose status is not HOLD.

Split Types (PSUS)

Table 2-4 lists and describes the UR split types. A user can split a UR pair at any time after the initial copy operation is complete. The user must split a UR pair in order to perform media maintenance on the primary data volume or to access the secondary data volume (secondary data volume write enable).

When an UR pair is split by the user, the primary storage system and secondary storage system ensure synchronization by either completing or discarding any pending update copy operations according to the user-specified pending update option (purge/flush). The primary storage system and secondary storage system keep track of any journal data that were discarded during the pairsplit operation. After the pair status changes to PSUS, the primary storage system stops performing journal obtaining operations for the pair, continues accepting write I/Os for the split UR primary data volume, and keeps track of the primary data volume tracks which are updated while the pair status is split. If you enabled the secondary data volume write option when you split the pair, the secondary storage system keeps track of the secondary data volume tracks which are updated while the pair is split. When the pair is resumed (pairresync), the secondary storage system sends the secondary data volume track bitmap to the primary storage system, and the primary storage system merges the primary data volume and secondary data volume bitmaps to determine which tracks are out-of-sync.

A split (or suspended) UR secondary data volume has an additional status called the consistency status. The consistency status is displayed only at the secondary storage system and indicates the update sequence consistency of the secondary data volume with respect to the other secondary data volume in the same group. Table 2-5 lists and describes the consistency status descriptions for split or suspended UR secondary data volumes (the consistency status is displayed only at the secondary storage system).

Table 2-4 UR Split Types (PSUS)

Type	Applies to	Description
PSUS, P-VOL by Operator	Primary data volume	The secondary data volume split type is PSUS-by MCU .
PSUS, S-VOL by Operator	Primary data volume, Secondary data volume	The user split the pair from the primary storage system or secondary storage system using the secondary data volume option.
PSUS, by MCU	Secondary data volume	The secondary storage system received a request from the primary storage system to split the pair. The primary data volume split type is PSUS-P-VOL by Operator or PSUS-S-VOL by Operator .
PSUS, by RCU	Primary data volume, Secondary data volume	The primary storage system detected an error condition at the secondary storage system which caused the primary storage system to suspend the UR data volume pair. The secondary data volume suspend type is PSUE-S-VOL Failure .
PSUS, Pairsplit-S to RCU	Primary data volume	The primary storage system detected that the secondary data volume status changed to SMPL because the user released the pair (pairsplit-S) from the secondary storage system. The pair cannot be resumed because the secondary data volume does not have the PSUS/PSUE status.
PSUS, JNL Cache Overflow	Primary data volume, Secondary data volume	The data volume pair was suspended because it was highly likely that journal data will overflow. Note: CCI (Command Control Interface) displays this split type as PFUS .

Table 2-5 Consistency Status for Split/Suspended UR Secondary Data Volume

Consistency Status (displayed at secondary storage system)	Description
Volume	<p>This UR data volume pair was split or suspended alone. Update sequence consistency between this secondary data volume and other secondary data volume in this journal group is not ensured. This secondary data volume cannot be used for disaster recovery at the secondary site. This status is indicated when:</p> <p>This pair was split by the user using the Suspend Range-Volume (LU) pairsplit option.</p> <p>This pair was suspended due to a failure that did not affect the entire journal group.</p>
Group	<p>This UR data volume pair was split or suspended along with the other pairs in its journal group. Update sequence consistency between this secondary data volume and other secondary data volume in this journal group is ensured. This secondary data volume can be used for disaster recovery at the secondary system (after releasing the UR data volume pair from the secondary storage system). This status is indicated when:</p> <p>This pair was split by the user using the Suspend Range-Group pairsplit option.</p> <p>All pairs in this journal group were suspended due to a failure that affected the entire journal group (not just one pair) (e.g., primary storage system-secondary storage system communication failure).</p> <p>One pair in the group was suspended due to a failure that did not affect the entire group.</p>

Suspend Types

Table 2-6 lists and describes the UR suspend types, which indicate the reason for the suspension. UR pairs can be suspended by either primary storage system or secondary storage system. For further information on UR suspension conditions, see section Suspension Condition.

The primary storage system suspends a UR pair when it detects any of the following suspension conditions:

When the primary storage system detects that the user has released the data volume pair from the secondary storage system,

When the primary storage system detects an error condition related to the secondary storage system, secondary data volume, or a UR update copy operation,

When the primary storage system is unable to communicate with the secondary storage system

When the primary storage system detects a UR suspension condition.

When a UR pair is suspended, the primary storage system stops performing journal-obtaining operations for the pair. The primary storage system continues accepting write I/Os for the suspended primary data volume and keeps track of the primary data volume cylinders/tracks which are updated while the pair is suspended. The primary storage system and secondary storage system also keep track of any journal data that were discarded during the pair suspension. When the suspended UR pair is resumed (pairresync), the secondary 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/tracks are out-of-sync. The primary storage system merges the primary data volume and secondary data volume bitmaps to identify the out-of-sync cylinders/tracks. This method ensures that all cylinders/tracks which contain journal data that were discarded at the secondary storage system are resynchronized at this time.

A split (or suspended) UR secondary data volume has an additional status called the consistency status. The consistency status is displayed only at the secondary storage system and indicates the update sequence consistency of the secondary data volume with respect to the other secondary data volume in the same group. Table 2-5 lists and describes the consistency status descriptions for split or suspended UR secondary data volumes (the consistency status is displayed only at the secondary storage system).

Table 2-6 Suspend Types (PSUE)

Type	Applies to	Description
PSUE, S-VOL Failure	Primary data volume, Secondary data volume	The primary storage system detected an error during secondary storage system communication or an error during update copy. In this case, the secondary data volume suspend type is usually PSUE-S-VOL Failure .
PSUE, MCU IMPL	Primary data volume, Secondary data volume	The primary storage system could not find valid control information in its nonvolatile memory during IMPL. This condition occurs only if the primary storage system is without power for more than 48 hours (i.e., power failure and fully discharged backup batteries).
PSUE, Initial Copy Failed	Primary data volume, Secondary data volume	The UR pair was suspended before the initial copy operation was complete. The data on the secondary data volume is not identical to the data on the primary data volume.
PSUE, MCU P/S OFF	Secondary data volume	The primary storage system is powered off.

Suspension Condition

UR operations involve additional suspension conditions related to the asynchronous journal operations. Both the primary storage system and secondary storage system can detect UR suspension conditions and suspend UR pairs. The primary storage system stores the differential bitmap per UR primary data volume in the shared memory. The secondary storage system stores the differential bitmap per UR secondary data volume in the shared memory. When a UR pair is suspended, the tracks which contain the following journal are marked in the differential bitmap as modified (to be copied during the pairresync operation):

- The journal data that were created by the primary storage system but not yet sent to the secondary storage system.

After marking these primary data volume tracks as modified, the primary storage system discards these journal data.

- The journal data that were sent to the secondary storage system but not acknowledged by the secondary storage system.

After marking these primary data volume tracks as modified, the primary storage system discards these journal data. This ensures that journal data lost during transmission to the secondary storage system are identified and marked.

- The journal data that reached the secondary storage system but have not yet been restored.

After marking these primary data volume tracks as modified, the primary storage system discards these journal data. Therefore, the secondary storage system can check and mark journal data that have not been restored.

The primary storage system does not discard journal data from the master journal volume until the primary storage system receives sequence numbers of restored journal volumes; the sequence numbers are included in read journal commands from the secondary storage system.

- The primary data volume records updated by host-requested write I/Os after the pair was suspended. The UR suspension conditions and indicates described by Table 2-7 which CU detects the condition and which pairs are suspended. See Troubleshooting for information on UR suspension conditions.

Table 2-7 UR Suspension Condition

Suspension Condition	Detected by:	UR Pairs to be Suspended
The secondary storage system could not copy the journal data successfully due to a hardware failure or logic error.	RCU	All UR secondary data volumes in the journal groups, or the affected secondary data volume.
The secondary storage system detected a logical error while selecting the journal data to be restored.	RCU	All UR secondary data volume in the journal group, or only the affected secondary data volume.
The secondary storage system could not restore the journal data due to a hardware failure, track condition, or logical error.	RCU	All UR secondary data volumes in the journal groups, or the affected secondary data volume.

Preparing for Universal Replicator Operations

This chapter describes the prerequisites and preparations to make to begin UR operations:

- Requirements and Restrictions
- Installing the Hardware
- Configuring the Primary and Secondary Storage Systems for UR Operations
- Enabling the UR Option
- Interoperability with Other Products and Functions
- TrueCopy Synchronous (3DC Multi-target Configuration)
- Planning of Journal Volumes
- Factors that Determine Data Transfer Speed between the Primary and Secondary Storage System
- Performing Remote Copy Between TagmaStore USP/NSC and USP V

Requirements and Restrictions

UR has the following requirements and restrictions:

- One-to-one volume copy operations
- Logical unit (LU) type
- Journal group
- Accessing UR primary data volumes and secondary data volumes
- Cache and NVS
- Host failover software
- Duplicate volume

System Requirements

UR operations involve the USP V primary storage systems and secondary storage systems containing the primary and secondary data volumes, the remote copy connections between the primary storage systems and secondary storage systems, the host(s) at the primary and remote sites, and the licensed UR remote console software. The UR system requirements are:

- Primary storage system: USP V storage system with UR installed.
- Secondary storage system: USP V storage system with UR installed.

Note: UR can coexist with UR for z/OS® in the same *USP V* storage system.

- Remote copy connections – Fiber channel (see section Setting up Remote Copy Connections):
 - Multimode or single-mode optical Fiber cables are required at both the primary storage system and secondary storage system.
 - For distance up to 0.5 km, multimode optical shortwave fiber cables are required between the primary storage system and secondary storage system.
 - For distances from 0.5 km to 1.5 km (1,640 to 4,920 feet), multimode shortwave Fiber-channel interface cables with up to two switches are required.
 - For distance up to 10 km, single optical long wave fiber cables are required between the primary storage system and secondary storage system.
 - For distances from 10 km to 30 km (6.2 to 18.6 miles), single-mode long wave Fiber-channel interface cables with up to two switches are required.
 - For distances greater than 30 km (18.6 miles), approved third-party channel extender products and telecommunications lines are required. Long-distance UR solutions are provided based on user requirements and workload characteristics.

- Host platforms:
USP V supports UR operations for a variety of UNIX[®] based and PC-server platforms, including HP-UX[®] 11.0, HP-UX[®] 11i, Sun[®] Solaris[®] 9, Windows[®] 2000, Windows[®] 2003, and IBM[®] AIX[®] 5.1 platforms.
Please contact your Hitachi account team for the latest information on platform support for UR.
- A computer that runs Storage Navigator(Storage Navigator computer):
The USP V Storage Navigator remote console software is required for USP V UR operations. The UR remote console software is a component of the USP V Hitachi Storage Navigator product. The UR license key(s) are required to enable the UR option(s) on the USP V storage system (see section Enabling the UR Option). Separate license keys are required for each USP V storage system. For further information on USP V Storage Navigator operations, please refer to the *Storage Navigator User's Guide*, or contact your Hitachi account team.
Note: Administrator or UR write access to the USP V Storage Navigator software is required to perform UR operations. Users without Administrator or UR write access can only view UR information.
- About the license of Universal Replicator:
If you want to use Universal Replicator, you must install not only a license for Universal Replicator but also a license for TrueCopy.
- Connection with TagmaStore USP / NSC
UR can execute remote copy operations by connecting USP V with TagmaStore USP/NSC. Specifically, the following configurations are supported.
- System configuration for remote copy operation using UR from USP V to TagmaStore USP.
- System configuration for remote copy operation using UR from TagmaStore USP/NSC to USP V.
Note: For detailed information about the connection with TagmaStore USP/NSC, see section 3.7.10.

One-to-One Volume Copy Operations

UR requires a one-to-one relationship between the volumes of the volume pairs. A volume (LU) can only be assigned to one UR pair at a time. However, when creating a UR pair for delta resync operation, you can specify the secondary data volume of a UR pair that is not for delta resync operation as the secondary data volume of the UR pair for delta resync operation. In that case, you need to create a mirror for the delta-resync pair and the non-delta-resync pair. For detailed information about delta resync operation, see section UR Delta Resync Operation and TrueCopy Synchronous (3DC Multi-target Configuration).

Note: UR does not support operations in which one primary data volume is copied to more than one secondary data volume, or more than one primary data volume is copied to one secondary data volume.

Because UR operates on volumes rather than on files, multivolume files require special attention. For complete duplication and recovery of a multivolume file (e.g., a large database file which spans several volumes), make sure that all volumes of the file are copied to UR secondary data volume, and use UR to ensure update sequence consistency across this group of secondary data volume.

Logical Unit (LU) Types

The following DKC and DKU emulation types can be used for the UR software.

Table 3-1 Supported Emulation Type

Emulation	Support type
DKC	3990-6E, 2105, or 2107
DKU (Drive)	Open-V Note: Only Open-V can be used for journal volumes.

UR displays emulation types of primary and secondary data volumes.

Table 3-2 lists the volumes and the volume capacity that can be used for the UR data volume and journal volume.

Note: The capacity of journal volume is not included in the accounting capacity.

Table 3-2 Supported Data Volume and Journal Volume

Type	Support specifications	
	Data Volume	Journal Volume
VLL volume	Available	
The volume that is set Cache Residency Manager	Available	
Maximum volume capacity	The maximum capacity of volume for each emulation type.	
Minimum volume capacity	The minimum capacity for a VLL volume.	The minimum capacity for a VLL volume. Note: A journal groups consists of areas containing journal data and an area containing metadata for remote copy.

Notes:

- Specify a port in SLPR 0 when you use LUN Manager to define an LU path to a data volume. It is possible but not recommended that you specify a port in another SLPR when you define an LU path. For detailed information about SLPR 0, please refer to the *Virtual Partition Manager User's Guide*.
- Volumes containing a VMA (volume management area) cannot be used as journal volumes. For detailed information about a VMA, please refer to the *Data Retention Utility User's Guide*.

The table below explains emulation types and the capacity of volumes that can form pairs. For details on the maximum number of pairs, see the next section.

Table 3-3 Specifications of Volumes that can Form Pairs

Item	Support specifications
Emulation type	Same type of pair can be created for OPEN-V.
Volume capacity	Same capacity

UR supports the LUN Expansion (LUSE) feature of the USP V storage system, which allows you to configure a LUSE volume by using 2 to 36 sequential LDEVs. If two LUSE volumes are assigned to a UR pair, the capacity and configuration of the secondary data volume must be the same as the primary data volume. For example, when the primary data volume is a LUSE volume in which 1-GB, 2-GB, 3-GB volumes are combined in this order, the secondary data volume must be a LUSE volume in which 1-GB, 2-GB, 3-GB volumes are combined in this order. In addition, RAID1, RAID5, and RAID6 can coexist in the LUSE volume.

UR also supports the Virtual LVI/LUN feature of the USP V storage system, which allows you to configure custom-size LUs which are smaller than standard-size LUs. When custom-size LUs are assigned to a UR pair, the secondary data volume must have the same capacity as the primary data volume. The emulation types and capacities of the master and restore journal volumes that can be used for a UR software are shown next..

Table 3-4 Journal Volume Specifications

Item	Support specifications
Emulation type	Same. Note: Only Open-V can be used for journal volumes.
Volume capacity	Same or different amount of capacity is available.

Table 3-5 shows the RAID level combination of data volume and journal volume in the journal group that can be used for UR.

Table 3-5 RAID Level Configuration of UR

Item	Support specifications
RAID configuration of data volume and journal volume	RAID1, RAID5, and RAID6 can be coexisted.

RAID1, RAID5, and RAID6 can be coexisted in the same journal group.

The Maximum Number of Pairs

Note: The number of pairs that can be created in a storage system is limited.

Use the number of cylinders and bitmap areas to calculate the maximum number of pairs that can be created in a storage system.

The number of cylinders:

The number of pairs of a primary data volume and a secondary data volume are limited. The number of pairs depends on the volume capacity (i.e., the number of cylinders in the volumes that form pairs). If VLL is used, the number of pairs depends on the number of cylinders specified by VLL. Limit on the number of pairs is applied to both the primary storage system and the secondary storage system. Use the following formulae to calculate the number of cylinders for different emulation types:

For OPEN-V:

$$\text{The number of cylinders} = (\uparrow ((\uparrow (\text{Max. LBA}/512) \uparrow) / 15) \uparrow) + 1$$

- The $\uparrow \dots \uparrow$ symbols around a value indicate that the value should be rounded up to the nearest integer.

- LBA : Logical Block Address

The number of the required bitmap areas:

The number of bitmap areas to be used by all data volumes that form pairs can be calculated out of the number of cylinders. The number is called the required number of bitmap areas. Use the following formula to calculate the required number of bitmap areas for a data volume. The $\uparrow \dots \uparrow$ symbols around a value indicate that the value should be rounded up to the nearest integer.

$$\text{The required number of bitmap areas} = \left(\left(\text{number of cylinders} \times 15 \right) \div 122,752 \right) \uparrow$$

" number of cylinders × 15" indicates the number of slots

122,752 is the number of slots that a bitmap area can manage

Note: If the required number of bitmap areas calculated, exceeds the total number of bitmap areas in the storage system, the number of pairs that can be created will be limited.

Note: If you use LUSE volumes, you must add plus 1 to the value calculated by the above formula.

The maximum number of pairs that can be created:

The maximum possible number of pairs that can be created is defined by the number of bitmap areas of the storage system and the required number of bitmap areas required to create pairs.

The number of bitmap areas of the storage system is defined by the capacity of shared memory. The relationship between the area number of shared memory and the number of bitmap areas in the storage system is described in Table 3-6.

Table 3-6 The Relationship between Additional Shared Memory and Total Number of Bitmap Areas of Storage System

Additional Shared Memory for UR	Total Number of Bitmap Areas of Storage System
No additional shared memory for UR	0
Additional shared memory for UR is installed	7,424
Extension 1	16,384
Extension 2	32,768
Extension 3	44,256
Extension 4	65,536

Use the following formulae to calculate the maximum possible number of pairs that can be created, on the basis of the number of bitmap areas described in Table 3-6 and the required number of bitmap areas you calculated:

$$\text{The maximum number of pairs} = \left(\downarrow \text{Number of bitmap areas} \div \text{required number of bitmap areas} \downarrow \right)$$

The ↓...↓ symbols around a value indicate that the value should be rounded down to the nearest integer.

Note: If the calculated maximum number of pairs exceeds 32,768, the actual maximum number of pairs is limited to 32,768.

Note: The capacity of OPEN-V is variable. To get the maximum number of pairs of OPEN-V, calculate by using the formula.

Caution: The bitmap areas that are used for Universal Replicator are also used for TrueCopy. If you use both TrueCopy and Universal Replicator, use the total number of both pairs.

Journal Groups

The UR journal groups have the following requirements:

Each UR pair must be assigned to only one journal group. The next table shows the specifications of journal group.

Table 3-7 Journal Group Specifications

Item	Support specifications
Number of journal groups	Up to 256 journal groups (No. 0 - 255) per one disk subsystem Note: The recommended number of journal groups is up to 16.
Number of data volumes in journal group	Up to 4,096
Number of journal volumes in journal group	Up to 64
Number of Mirror IDs	Up to 4 (ID No.: 0 to 3) Note: If TrueCopy Sync. uses No. 0, No. 1 to 3 are available for UR.

The same number of journal volumes is not required in the master journal group and the restore journal group that are paired.

Mirror ID is required for the configuration that will be supported in the future within the 3-data center (3DC). In the future, we might support that two or more restore journal group will be able to set for one master journal group. Each pair relationship in a journal group is called "Mirror". Mirror ID identifies two or more mirrors that one journal group has. The same Mirror ID is applied for the pair of data volume. See section TrueCopy Synchronous (3DC Cascading Configuration) for 3DC configurations.

Table 3-8 shows the specifications of relationship between the data volumes, between the journal volumes, and between the data volume and journal volume in the journal group.

Table 3-8 Volume Specifications in Journal Group

Item	Support specifications
Emulation type	Same or different types of emulation type is available. Note: 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. Data volumes must also belong to the same CLPR. Note: A primary journal group and the corresponding secondary journal group need not belong to the same CLPR.

Note: When UR and UR for z/OS[®] coexist in the same *USP V* storage system, each journal group must contain either UR pairs or UR for z/OS[®] pairs (not both).

Accessing UR Primary Data Volume and Secondary Data Volume

To ensure maximum data integrity during normal UR operations, the secondary storage system rejects all write operations issued by a host to a UR secondary data volume. If you need write operation to an UR secondary data volume, you must set the secondary data volume write option (see section Secondary Data Volume Write Option). When you resume (Pairresync) the split pair, the secondary storage system will send the secondary data volume track bitmap to the primary storage system to ensure proper resynchronization of the pair.

Cache and Nonvolatile Storage (NVS)

Cache and nonvolatile storage (NVS) must be operable for both the primary storage system and secondary storage system of a UR data volume pair. If not, the UR paircreate operation will fail. The remote storage system cache should be configured to adequately support the UR remote copy workloads as well as any local workload activity.

Host Failover Software

The USP V Storage Navigator software product does not provide any host failover functions for disaster recovery. Host failover software is a critical component of any disaster recovery effort. When a primary storage system fails to maintain synchronization of a UR pair, the primary storage system generates sense information. It is very important that this information be transferred to the remote site via the host failover software for effective disaster recovery.

Duplicate Volumes

Since the contents of the primary data volume and secondary data volume of a UR pair are identical, the secondary data volume can be considered a duplicate of the primary data volume. Since the host operating system does not allow duplicate volumes, the host system administrator must take precautions to prevent system problems related to duplicate volumes. You must define the UR secondary data volume so they do not auto-mount or come online to the same host at the same time as the primary data volume (see **WARNING** below).

UR does not allow the secondary data volume to be online (except while the pair is split). If the secondary data volume is online, the UR paircreate operation will fail.

WARNING: If the UR secondary data volume are physically attached to the same host server(s) as the UR primary data volumes, the following problem can occur:

When a UR pair is released, the old secondary data volume is usually offline. If the host is then restarted, the system administrator may be offered both volumes and asked which volume should be left offline. This can be confusing and is prone to error.

If the UR secondary data volumes and primary data volumes are connected to the same host(s), Hitachi strongly recommends that the secondary data volumes are defined to remain offline to avoid this problem.

Installing the Hardware

Initial installation of the UR hardware is performed by the user and the Hitachi representative. To install the hardware required for UR operations:

1. **User:** Identify the locations of the UR primary and secondary data volumes (primary data volumes and secondary data volumes), so that the UR hardware can be installed and configured properly.
2. **Hitachi Representative:** Make sure that the primary storage system(s) and secondary storage system(s) are configured for USP V Storage Navigator operations (e.g., SVP connected to LAN). Refer to the *Storage Navigator User's Guide* for information and instructions on setting up Storage Navigator operations.
3. **Hitachi Representative:** Make sure that the primary storage systems and secondary storage systems are properly configured for UR operations (e.g., cache, NVS) (see section Cache and Nonvolatile Storage (NVS)). Make sure that the desired system option modes are enabled (see Table 2-2). Make sure that adequate cache is installed and available for UR operations. You must also consider the amount of Cache Residency Manager data to be stored in cache when determining the required amount of cache.
4. **Hitachi Representative:** Make sure the primary storage systems are configured to report sense information to the host(s). The secondary storage systems should also be attached to a host server to enable reporting of sense information in case of a problem with an secondary data volume or secondary storage system. If the remote site is unattended, the secondary storage systems should be attached to a host server at the primary site, so that the system administrator can monitor the operational condition of the secondary storage systems.

5. **Hitachi Representative:** If power sequence control cables are used, set the power select switch for the cluster to LOCAL to prevent the primary storage system from being powered off by the host. Also make sure the secondary storage system will not be powered off during UR operations. See Managing Power for Storage Systems and Network Relay Devices for further information on powering off/on the primary storage systems and secondary storage systems.
6. **Hitachi Representative:** Install the UR remote copy connections between the primary storage system(s) and secondary storage system(s). This hardware (optical Fiber cables, switches, etc.) is supplied by the user.
See the next section for remote copy configurations. Distribute the paths between different storage clusters and switches to provide maximum flexibility and availability. The logical paths between the primary storage system and secondary storage system must be separate from the logical paths between the host and secondary storage system.

Setting up Remote Copy Connections

Figure 3-1 shows the remote copy connection configurations for UR operations. The primary storage system and secondary storage system of each UR pair must be connected via optical fiber cables. If you use multimode shortwave optical fiber cables, Fiber cables up to 1.5 km in length and up to two switches are required for distances greater than 0.5 km. If you use single-mode longwave optical fiber cables, Fiber cables up to 30 km in length and up to two switches are required for distances greater than 10 km. UR operations can be performed at distances of up to 30 km (18.6 miles) using standard single-mode longwave support. For further distance, the channel extender connections are required. UR operations can be performed at distances of up to 30 km (18.6 miles) using standard single-mode longwave support. For distances greater than 43 km (26.7 miles), approved channel extender products and telecommunications lines are required.

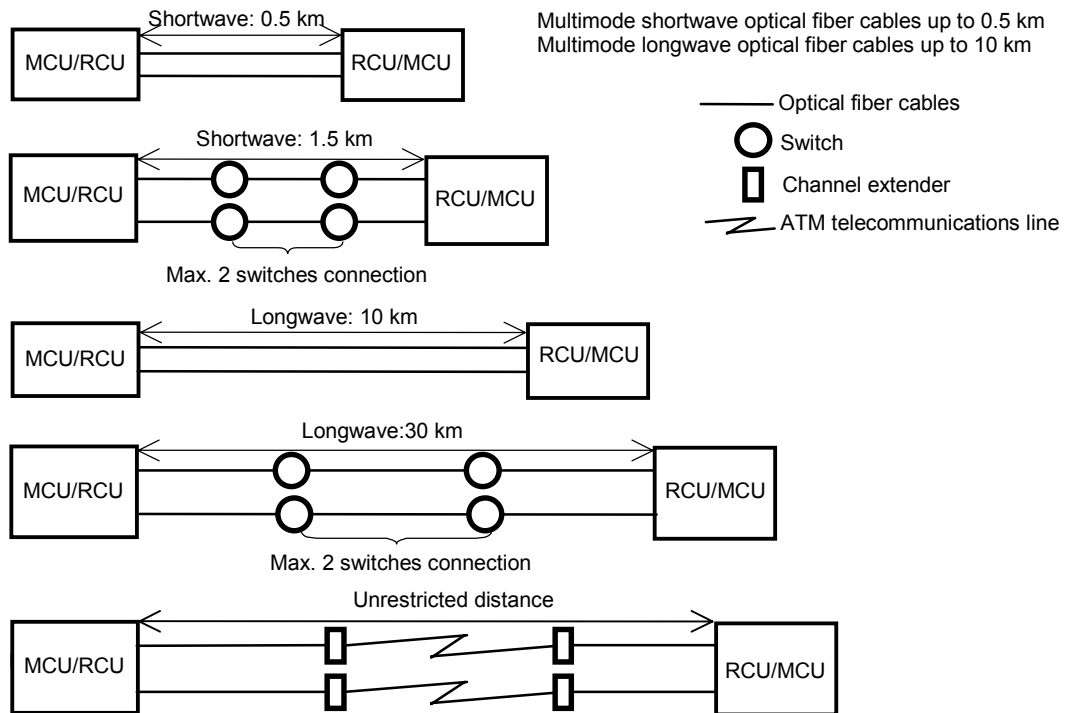
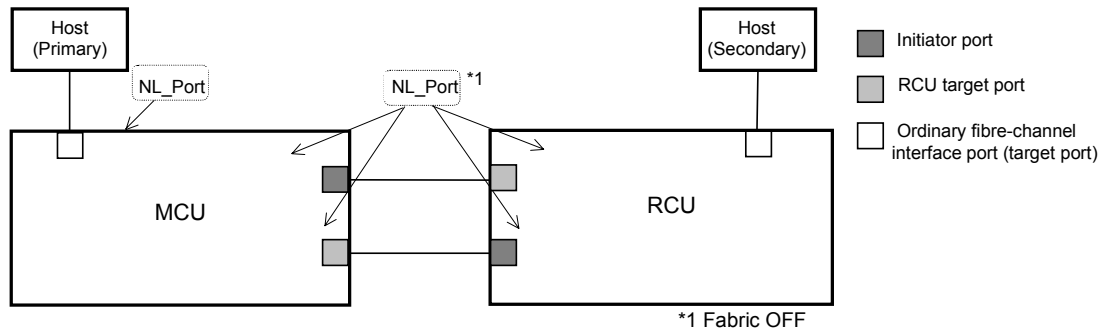


Figure 3-1 UR Remote Copy Connection Configuration

The remote copy connection between primary storage system and secondary storage system provides three different configurations:

- Direct connection (see Figure 3-2),
- Switch connection (see Figure 3-3),
- Extender connection (see Figure 3-4).



* To set ports, use LUN Manager and set port topology to: Fabric off, FC-AL.

Figure 3-2 Direct Remote Copy Connections

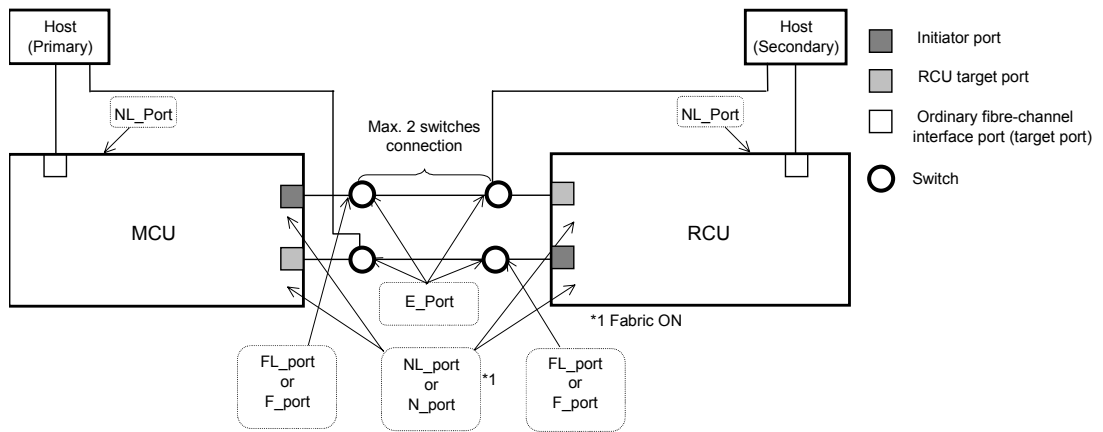


Figure 3-3 Switch Remote Copy Connection




Figure 3-4 Extender Remote Copy Connection

Caution: When a MCU and RCU are connected via switches with channel extender, and multiple remote copy paths are assembled, the capacity of data to be transmitted may concentrate on particular switches, depending on the configuration and the settings of switch routing.

Configuring the Primary and Secondary Storage Systems for UR Operations

After you have installed the UR hardware and software, you can configure the primary storage systems and secondary storage systems for UR operations.

To configure the primary storage systems and secondary storage systems for UR operations:

1. Identify the volumes that will become the UR data volumes and journal volumes in primary storage system and secondary storage system. You need to know the storage system S/N and port ID, so that you can configure the primary storage systems and secondary storage systems correctly for your desired pairs. When you create the pairs, you will need to know the port, group ID(GID), and LUN of each volume.
2. Stop Performance Monitor to avoid overloaded traffic of TCP/IP, if this program product is running. On a primary storage system and a secondary storage system, you will stop Performance Monitor.
3. Connect to the *USP V* storage system that you want to use. For further information on *USP V* Storage Navigator installation and operations, please refer to the *Storage Navigator User's Guide*, or contact your **Hitachi** account team. Operate at a primary storage system and a secondary storage system.

Note: You must log in with Administrator or UR write access, and you must be in modify mode. If you are in the view mode, you can refer to the UR pair status, but you cannot modify the settings.

4. When you click **Go, Universal Replicator** and then **UR Monitor** on the menu bar of the Storage Navigator main window, Universal Replicator starts and the Universal Replicator window is active.
5. Select the **Port** radio button on the DKC Operation window, set the initiator and RCU target ports (see section Configuring Port Attributes). Operate at a primary storage system and a secondary storage system.
6. Select the DKC radio button on the DKC Operation window to display the Add DKC window. Set the S/N of the storage system which will be connected, controller ID, path group ID, the initiator port in this storage system, the RCU target port in the storage system which will be connected (see section Establishing Relationship between Primary and Secondary Storage Systems (Add DKC)). Operate at a primary storage system and a secondary storage system.
7. Select the Option button on Add DKC window, and then set the desired options of local storage system. Operate at a primary storage system and a secondary storage system.
8. Specify the journal group number from the Journal Operation window. Register the journal volumes from the Edit JNL Volumes window (see section Registering Journal Volumes in a Journal Group). Operate at a primary storage system and a secondary storage system.

9. Register the primary and secondary data volumes and the journal groups that are paired from the Paircreate window in Pair Operation (see section Creating a Pair of Data Volumes (Paircreate)).
Note: UR pairs can only be registered by the primary storage system.
10. Exit the UR remote console software, and disconnect from the storage system. Operate at a primary storage system and a secondary storage system.
11. Restart Performance Monitor, when necessary. On a primary storage system and a secondary storage system, you will start Performance Monitor.
12. Repeat steps (3)–(9) for each storage system which will function as a UR primary storage system. After you have configured the primary storage systems, added the secondary storage systems, and configured the UR options and journal groups, you are ready to begin UR volume pair operations.

Enabling the UR Option

To operate the UR software, PC for the USP V Storage Navigator is required. For further information on USP V Storage Navigator operations, please refer to the *Storage Navigator User's Guide*, or contact your **Hitachi** account team.

Interoperability with Other Products and Functions

Some types of volumes used by non-UR functions can also be used as UR data volumes and/or journal volumes. Table 3-9 explains whether non-UR volumes are also usable as UR volumes.

Table 3-9 Whether Non-UR Volumes Can Be Used as UR Volumes

Functions and Volumes	Can the Volumes be Used as Primary Data Volumes?	Can the Volumes be Used as Secondary Data Volumes?	Can the Volumes be Used as Journal Volumes?
LUN Expansion (LUSE)			
LUSE volume	Yes.	Yes.	No.
ShadowImage			
P-VOL in PSUS status	Yes.	Yes.	No.
P-VOL in COPY(RS-R) status	No.	No.	No.
P-VOL that is also used as a TC P-VOL or TC S-VOL	Yes.	Yes.	No.
P-VOL (none of the above)	Yes.	Yes.	No.
S-VOL in PSUS status	Yes.	No.	No.
S-VOL (none of the above)	No.	No.	No.
Reserved volume	No.	No.	No.

Functions and Volumes	Can the Volumes be Used as Primary Data Volumes?	Can the Volumes be Used as Secondary Data Volumes?	Can the Volumes be Used as Journal Volumes?
Copy-on-Write Snapshot			
P-VOL in COPY(RS-R) status	No.	No.	No.
P-VOL that is also used as a TC P-VOL or TC S-VOL	No.	No.	No.
P-VOL (none of the above)	Yes. *1	No.	No.
V-VOL	No.	No.	No.
Pool-VOL	No.	No.	No.
Volume Migration			
Source volume (when volume migration is in progress)	Yes. Note: If the source volume is used as a primary data volume, the volume migration will stop.	Yes. Note: If the source volume is used as a secondary data volume, the volume migration will stop.	No.
Source volume (after volume migration is finished)	Yes.	Yes.	No.
Reserved volume to which no path is defined	No.	No.	No.
TrueCopy			
P-VOL in COPY status	No.	No.	No.
P-VOL in PAIR status	Yes.	No.	No.
P-VOL in PSUS status	Yes.	No.*21	No.
P-VOL in PSUE status	Yes.	No.*21	No.
S-VOL in COPY status	No.	No.	No.
S-VOL in PAIR status	Yes.	No.	No.
S-VOL in PSUS status	Yes.	No.	No.
S-VOL in SSWS status	Yes.	No.*21	No.
S-VOL in PSUE status	Yes.	No.	No.
TrueCopy Asynchronous			
TrueCopy Asynchronous volume	No.	No.	No.
Data Retention Utility			
Volume with Read/Write attribute	Yes.	Yes.	Yes.
Volume with Protect attribute	Yes.	Yes.	No.
Volume with Read Only attribute	Yes.	Yes.	No.
Volume that is disabled for use as an S-VOL	Yes.	No. However, the volume can be used as a secondary data volume only when	No.

Functions and Volumes	Can the Volumes be Used as Primary Data Volumes?	Can the Volumes be Used as Secondary Data Volumes?	Can the Volumes be Used as Journal Volumes?
		you restore a UR pair.	
Volume including VMA	Yes.	Yes.	No.
Cross-OS File Exchange			
Volume usable by both mainframe and open systems	No.	No.	No.
Database Validator			
Database Validator volume	No.	No.	No.
Cache Residency Manager			
The volume that is set Cache Residency Manager	Yes.	Yes.	Yes.
Serverless Backup Enabler			
Serverless Backup Enabler volume	Yes.	Yes.	No.
Server Priority Manager			
Server Priority Manager volume	Yes.	Yes.	No.
Virtual LUN			
Virtual LUN volume	Yes.	Yes.	Yes.
LUN Manager			
Volume to which a path is defined	Yes.	Yes.	No.
Volume to which no path is defined	No.	No.	No.
Volume to which LUN security is applied	Yes.	Yes.	No.

Note*1: You cannot use the volume as a data volume of a UR pair for delta resync operation.

Note*1: The volume can be used as a secondary data volume only when you restore a UR pair or perform a CCI **horctakeover** operation. However, even in that case, you cannot use the volume as a secondary data volume of a UR pair for delta resync operation.

LUSE

Two LUSE volumes can be assigned to a UR pair. Both of LUSE volumes that are assigned to primary and secondary data volumes must consist of the same number of LDEVs and must have 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), please refer to *LUN Expansion User's Guide*.

Virtual LUN

You can perform Virtual LUN operations on primary and secondary data volumes in UR pairs. If you need to perform Virtual LUN operations on a UR primary data volume or secondary data volume, you must delete the pair first to return the volume to SMPL status.

When creating a UR pair consisting of two Virtual LUN volumes, make sure that the primary data volume and the secondary data volumes have the same capacity.

Cache Residency Manager

You can perform Cache Residency Manager operations on UR primary data volumes and secondary data volumes.

LUN Manager

LUN Manager operations do not affect UR operations. Volumes which are under secure ports and/or which are assigned to host groups or iSCSI targets can also be assigned to UR pairs. Volumes which are assigned to UR pairs can be protected by LUN Manager.

Note: UR secondary data volumes cannot be accessed by any UNIX[®] or PC server host except when the pair is split.

Server Priority Manager

The initiator ports in the primary storage system do not support the Server Priority Manager option.

ShadowImage

UR and ShadowImage can be used together in the same storage system and on the same volumes to provide multiple copies of data at the primary and/or remote sites.

The host pair status reporting for UR volumes described by Table 3-10, ShadowImage volumes, and UR/ShadowImage shared volumes. Table 3-11 shows the currency of the data on a shared UR/ShadowImage volume based on UR and ShadowImage pair status.

- For shared UR/ShadowImage volumes, the UR pair status is reported to the host if you query the UR primary data volume or secondary data volume. To obtain the ShadowImage pair status, query the P-VOL of the ShadowImage pair.
- ShadowImage supports multiple S-VOLs for each P-VOL. If you issue a pair status query to a ShadowImage P-VOL (e.g., pairdisplay), the status for only one ShadowImage pair is reported (the pair with the S-VOL with the lowest LDEV ID). To obtain the pair status for the ShadowImage pair(s) with the other S-VOL(s), you must direct the host query to the specific S-VOL using the S-VOL's LDEV ID in the host command. The ShadowImage remote console software displays the port, GID, LUN, LDEV ID and ShadowImage pair status of all S-VOLs associated with a P-VOL.

Table 3-10 Host Pair Status Reporting for UR/ShadowImage Shared Volumes

Number of UR Pairs	Number of ShadowImage S-VOLs	Pair Status Reported by USP V
0	0	Simplex
0	1	ShadowImage pair status
0	2 or more	ShadowImage pair status for the pair whose S-VOL has the lowest LDEV ID
1	0	UR pair status
1	1	UR pair status
1	2 or more	UR pair status

Table 3-11 Data Currency of a Shared UR/ShadowImage Volume

UR pair status	ShadowImage Pair Status					
	COPY (PD)	PAIR	COPY (SP)	PSUS	COPY (RS/RS-R)	PSUE
COPY	Not current	Not current	Not current	CURRENT	Not current	Not current
PAIR	Not current	Not current	Not current	CURRENT	Not current	Not current
PSUS/PSUE	Not current	CURRENT	CURRENT	CURRENT	CURRENT	Not current

Figure 3-5 through Figure 3-8 show the various UR/ShadowImage configurations which share volumes.

- UR/ShadowImage configurations which share the UR primary data volume and ShadowImage P-VOL

Figure 3-5 shows an example of an UR primary data volume which is also functioning as a ShadowImage P-VOL. This configuration allows you to use ShadowImage for on-site data backup in case of an UR failure, and to use UR to provide remote backup of the ShadowImage P-VOL in case of a ShadowImage failure.

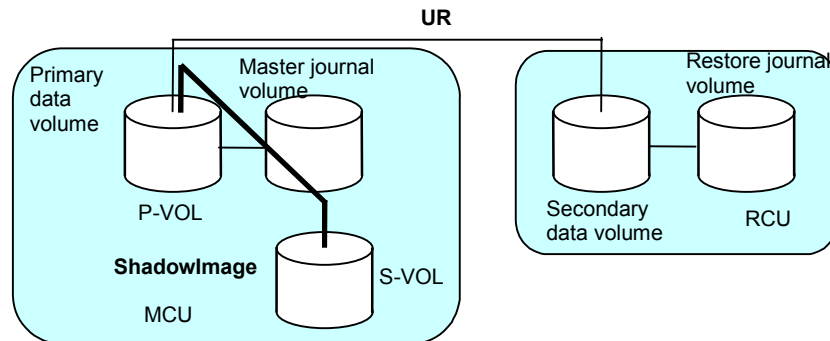


Figure 3-5 Shared UR Primary Data Volume and ShadowImage P-VOL

- **UR/ShadowImage configurations which share the UR secondary data volume and ShadowImage P-VOL**

Figure 3-6 shows an example of an UR secondary data volume which is also functioning as a ShadowImage P-VOL. This configuration allows you to use ShadowImage to provide multiple backup copies of a single UR primary data volume.




Figure 3-6 Shared UR Secondary Data Volume and ShadowImage P-VOL

Caution: When you share a UR secondary data volume with a ShadowImage P-VOL as shown in Figure 3-6, the write operation to the UR primary data volume takes time. Especially, when the ShadowImage pair is in the PSUS(SP) status, the write operation to the UR primary data volume may takes extra time according to the time for copying process of the ShadowImage pair.

In addition, note that if the journal volume size is small, the UR pair may be suspended by failure because of the shortage of the capacity of its journal volume.

- UR/ShadowImage configuration which share the UR primary data volume and ShadowImage P-VOL, and UR secondary data volume and ShadowImage P-VOL

Figure 3-7 combines the configurations shown in Figure 3-5 and Figure 3-6. Within a single UR pair, the primary data volume and secondary data volume are both functioning as ShadowImage P-VOLs, providing multiple copies at the primary and remote sites.




Figure 3-7 Shared UR primary data volume and ShadowImage P-VOL, and UR secondary data volume and ShadowImage P-VOL

- **UR/ShadowImage configuration where a ShadowImage S-VOL in PSUS status is used as a UR primary data volume**

In the following example, the ShadowImage S-VOL in PSUS status is also functioning as a UR primary data volume. This configuration allows UR to make a remote backup copy of the ShadowImage S-VOL.

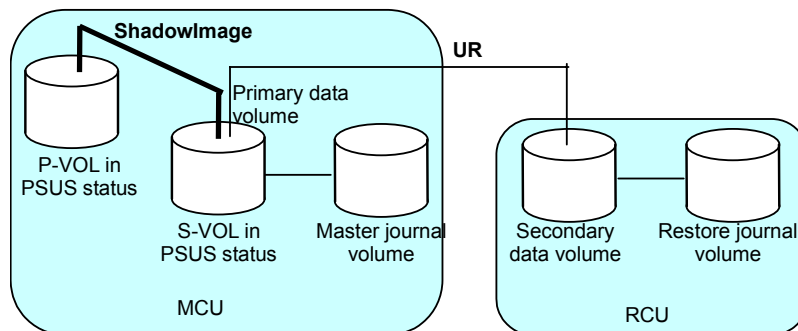


Figure 3-8 ShadowImage S-VOL in PSUS Status Functioning as UR Primary Data Volume

If a failure occurs and the ShadowImage P-VOL is damaged in Figure 3-8, take the following steps to copy data from the UR secondary data volume to the ShadowImage P-VOL so that data can be restored to the ShadowImage P-VOL:

1. Execute the CCI **pairsplit -S** command on the ShadowImage pair to release the pair (see Figure 3-9).

2. Execute the CCI **horctakeover** command on the UR pair to reverse the copy direction and re-establish the pair (see Figure 3-10).
3. Execute the CCI **horctakeover** command on the UR pair to change the copy direction to the original direction and re-establish the pair (see Figure 3-11).
4. Execute the CCI **pairsplit -r** command on the UR pair to suspend the pair (see Figure 3-12).
5. Execute the CCI **paircreate** command on the ShadowImage pair to perform copying in the reverse direction (see Figure 3-13).
6. Execute the CCI **pairsplit -S** command on the ShadowImage pair to release the pair (see Figure 3-14).
7. Execute the CCI **paircreate** command on the ShadowImage pair to perform copying in the original direction (see Figure 3-15).
8. Execute the CCI **pairsplit -r** command on the ShadowImage pair to put the pair in PSUS status (see Figure 3-16).
9. Execute the CCI **pairresync** command on the UR pair to resynchronize the pair (see Figure 3-17).




Figure 3-9 Restoring a ShadowImage P-VOL - Step 1




Figure 3-10 Restoring a ShadowImage P-VOL - Step 2




Figure 3-11 Restoring a ShadowImage P-VOL - Step 3




Figure 3-12 Restoring a ShadowImage P-VOL - Step 4




Figure 3-13 Restoring a ShadowImage P-VOL - Step 5




Figure 3-14 Restoring a ShadowImage P-VOL - Step 6

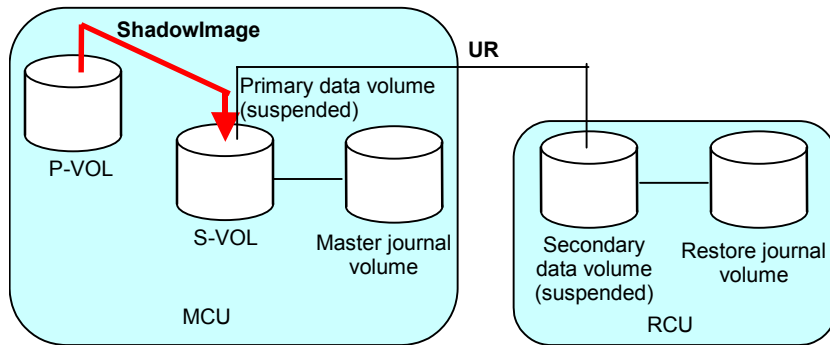


Figure 3-15 Restoring a ShadowImage P-VOL - Step 7




Figure 3-16 Restoring a ShadowImage P-VOL - Step 8

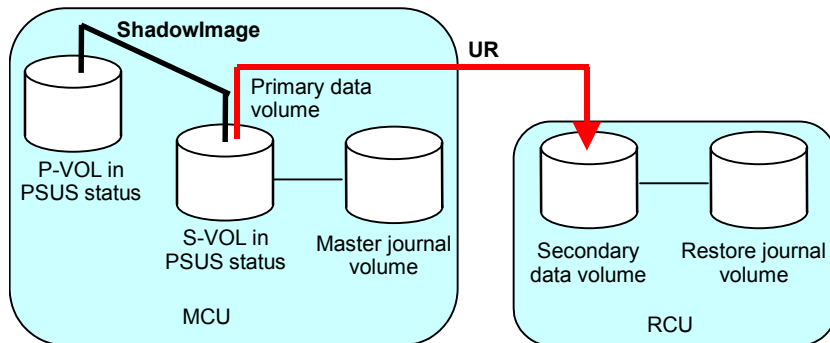


Figure 3-17 Restoring a ShadowImage P-VOL - Step 9

TrueCopy Synchronous (3DC Cascading Configuration)

The USP V storage system provides the function to combine UR and TrueCopy Synchronous. This combination is intended to ensure that the response time against host I/Os is comparable, regardless of whether the distance between the primary and the secondary sites are short or long. This combination is also intended to ensure that the secondary site stores data that has been stored in the primary site even when a failure occurs in the primary site. These intentions will be fulfilled if remote copy operations are performed using cascading connections and a three data center (3DC) configuration; in a 3DC configuration, an intermediate site is located between the primary and secondary sites.

Figure 3-18 shows an example of a 3DC cascading configuration in which UR and TrueCopy Synchronous are combined. In this configuration, the volume in the primary site is a TrueCopy Synchronous primary volume (P-VOL). The corresponding TrueCopy secondary volume (S-VOL) is the volume in the intermediate site, which is within a short distance from the primary site. This TrueCopy Synchronous secondary volume is also used as a UR primary data volume. The corresponding UR secondary data volume is the volume in the secondary site, which is within a long distance from the intermediate site. If a host issues an I/O to the TrueCopy Synchronous primary volume in the primary site, data will be transferred via TrueCopy Synchronous and UR to the secondary site and will be restored to the secondary data volume.




Figure 3-18 3DC Cascading Configuration (Combining UR with TrueCopy Synchronous)

Data that is copied by TrueCopy Synchronous is stored in the UR primary data volume. If a failure occurs at the TrueCopy Synchronous P-VOL, business tasks can continue by using data in the UR secondary data volume containing data in the TrueCopy Synchronous P-VOL. Data in the TrueCopy Synchronous P-VOL is restored via UR to the secondary data volume in the secondary site, with update sequence integrity ensured. If a failure occurs at both TrueCopy Synchronous P-VOL and S-VOL due to a disaster or some other reason, disaster recovery can be performed using the UR secondary data volume.

Note: A TrueCopy Synchronous P-VOL must be created in a USP V storage system. The fence level of the TrueCopy Synchronous P-VOL must be **Data**.

Basic Behavior

This section explains the basic behavior of a 3DC cascading configuration under the following conditions:

- The status of the TrueCopy Synchronous pair is **PAIR**. The status of the UR pair is also **PAIR**.
- A host in the primary site issues an I/O request to the TrueCopy Synchronous P-VOL

As illustrated in Figure 3-19, if a host in the primary site issues an I/O request, data will be stored into the TrueCopy Synchronous P-VOL. The data will also be stored into the TrueCopy Synchronous S-VOL (i.e., the UR primary data volume). After that, the fact that data has been stored will be reported to the primary site, and then the fact that I/O has been finished will be reported to the host. Also, data in the UR primary data volume (i.e., TrueCopy Synchronous S-VOL) will be stored as journal data into the master journal volume in the intermediate site. Sequence numbers, which indicates the order that data will be written, will be added to these journal data. Also, journal data in the master journal volume in the intermediate site will be transferred, independently from and asynchronously with the behavior mentioned above, to the restore journal volume in the secondary site. Next, the journal data in the restore journal volume will be restored to the UR secondary data volume in the secondary site; the journal data will be restored in the order of sequence numbers that are added to journal data. These sequence numbers ensure that updates to UR secondary data volumes in the secondary site will be made in the same order that updates to TrueCopy Synchronous P-VOL in the primary site are made.

As explained above, when the host in the primary site issues an I/O request to the TrueCopy Synchronous P-VOL, the data will be restored asynchronously to the UR secondary data volume in the secondary site in a remote location.

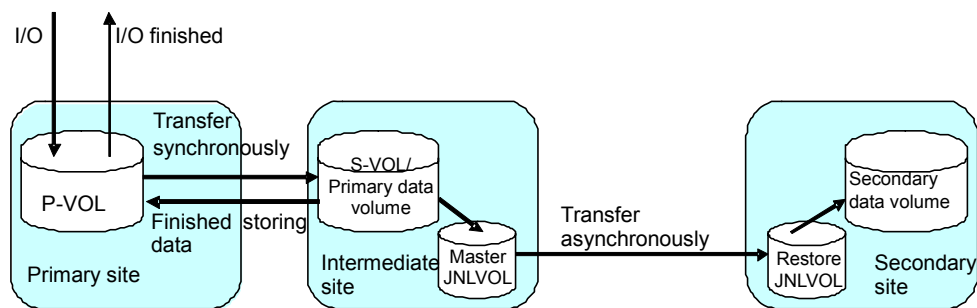


Figure 3-19 Basic Behavior in 3DC Cascading Configuration

If a TrueCopy Synchronous pair is suspended in a 3DC cascading configuration, you will be able to resynchronize the pair by using differential copy. If a UR pair is suspended, you will be able to resynchronize the pair by using differential copy.

Note: To perform disaster recovery from a failure in the primary site in a 3DC cascading configuration, you must ensure that the fence level of the TrueCopy Synchronous P-VOL is **Data**. If this fence level is specified, you can perform disaster recovery by using the UR secondary data volume that contains the data stored in the TrueCopy Synchronous P-VOL.

Hardware Configuration

This section explains hardware configuration for a 3DC cascading configuration using TrueCopy Synchronous and UR. In a 3DC cascading configuration, three USP V storage systems are required. It is recommended that Command Control Interface (CCI) is installed on hosts in the primary site, the intermediate site, and the secondary site. Storage Navigator computers are required for these sites. The following explains hardware configuration for these sites.

- Primary site
 - The USP V storage system must have TrueCopy Synchronous installed.
 - The USP V storage system must have TrueCopy Synchronous P-VOL.
- Intermediate site
 - The USP V storage system must have TrueCopy Synchronous installed.
 - The USP V storage system must have UR installed.
 - The USP V storage system must have a volume used as both a TrueCopy Synchronous S-VOL and a UR primary data volume.
 - The USP V storage system must have a UR master journal volume.
- Remote site
 - The USP V storage system must have UR installed.
 - The USP V storage system must have a UR secondary data volume.
 - The USP V storage system must have a UR restore journal volume.

Setup Procedure

This section explains the procedure for setting up a 3DC cascading configuration using TrueCopy Synchronous and UR.

To set up a 3DC cascading configuration:

1. Install TrueCopy Synchronous and UR. Configure ports and journal groups.
2. Issue a request for creating a TrueCopy Synchronous pair to the primary storage system where TrueCopy Synchronous is installed.
3. Wait until the status of the TrueCopy Synchronous pair becomes PAIR.
4. Issue a request for creating a UR pair to the UR primary storage system.

5. When creating a UR pair, be sure to select a value from 1 to 3 as the mirror ID.
6. Wait until the status of the UR pair becomes **PAIR**.

Note: You must ensure that the TrueCopy Synchronous pair status changes to **PAIR** before the UR pair status changes to **PAIR**. If you have made a request for changing the UR pair status to **PAIR** before the TrueCopy Synchronous pair status changes to **PAIR**, the request will be rejected.

For detailed information about installing and using TrueCopy Synchronous, please refer to *TrueCopy for IBM z/OS User's Guide*. For detailed information about UR journal group configuration, see section Requirements for Creating UR Pair for Delta Resync Operation. For detailed information about port configuration, see section Configuring Port Attributes.

TrueCopy Synchronous (3DC Multi-target Configuration)

The USP V storage system provides the function to combine UR and TrueCopy Synchronous. In a system where UR and TrueCopy Synchronous are combined, there are one primary site and two secondary sites. One of the secondary sites is for synchronous remote copy, and is usually located at a short distance from the primary site. The other secondary site is for asynchronous remote copy, and is usually located at a long distance from the primary site. Such a system configuration is called *three data center (3DC) multi-target configuration*.

3DC multi-target configuration provides solution for disaster recovery and business continuity. Even when a disaster or failure occurs at two sites at the same time, you can restart your business tasks by using data in the remaining secondary site.

The figure below illustrates a 3DC multi-target configuration where UR and TrueCopy Synchronous are combined. In this configuration, there are three sites as follows:

- Primary site for both TrueCopy Synchronous and UR
- Secondary site for TrueCopy Synchronous
- Secondary site for UR

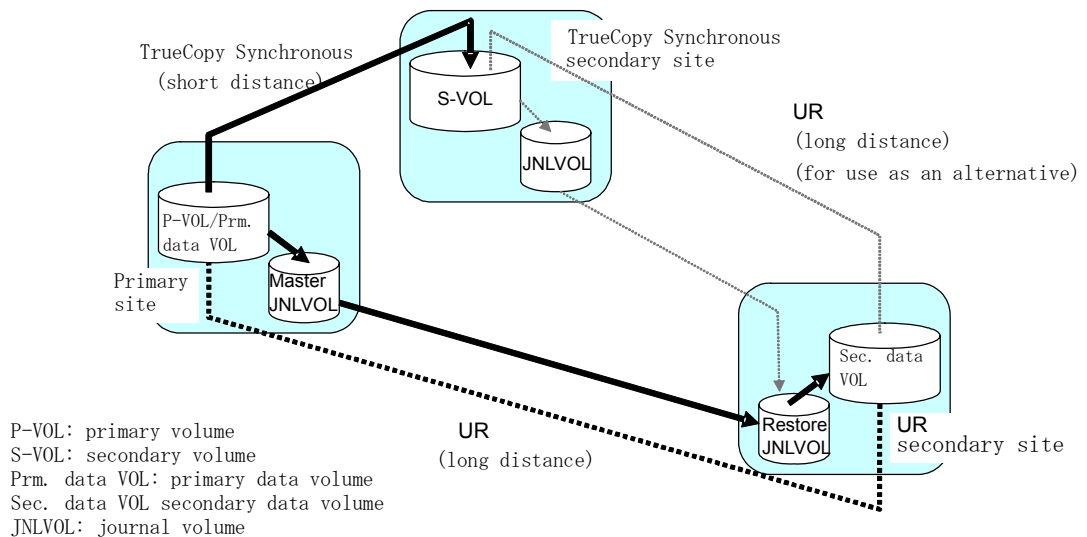


Figure 3-20 3DC Multi-target Configuration (Combining UR with TrueCopy Synchronous)

In this configuration, the primary volume in the primary site is paired with the secondary volume for TrueCopy Synchronous. This primary volume is also paired with the secondary volume for UR. If a host in the primary site writes data to the primary volume, the data will be written to the TrueCopy Synchronous secondary volume in real time, synchronously with writing from the host. The data will also be written to the UR secondary data volume, asynchronously with writing from the host. Update sequence consistency will be maintained with the UR secondary data volume.

Since writing from the UR primary volume to the UR secondary volume is performed asynchronously with writing from the host to the primary volume, this configuration can prevent delay in response to the host even if the UR primary site and the UR secondary site are far from each other.

In addition, you can create the UR pair for delta resync operation in 3DC multi-target configuration, by specifying TrueCopy Synchronous S-VOL in the TrueCopy Synchronous secondary site as the UR primary data volume, and by specifying the volume in the UR secondary site as the UR secondary data volume. If you create the UR pair for delta resync operation, you can resynchronize the UR pair by using journal data in the TrueCopy Synchronous secondary site when a failure occurs in the primary site.

For detailed information about the hardware configuration required for 3DC multi-target configuration, see the next section. For detailed information about the requirements to create a UR pair for delta resync operation, see section Requirements for Creating UR Pair for Delta Resync Operation. As for the requirements to perform delta resync operation, see section Requirements for Performing Delta Resync Operation.

If a failure occurs in a TrueCopy Synchronous primary volume, which is also used as a UR primary data volume, you can resume business tasks by using the TrueCopy Synchronous secondary volume. After you resume business tasks, you need to remove the failure from the primary volume and transfer business tasks back to the primary site. For details, see sections Recovering from Primary Site Disaster in 3DC Multi-Target Configuration through Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC Cascading Configuration).

If a failure occurs in both the primary volume and the TrueCopy Synchronous secondary volume, you can resume your business tasks by using the UR secondary data volume. After you resume business tasks, you need to remove the failure from the primary volume and the TrueCopy Synchronous secondary volume, and transfer business tasks back to the primary site. For details, see sections Recovering from Failures in the Primary Site and the TrueCopy Synchronous Secondary Site through Transferring Business Tasks from the UR Secondary to the Primary Site.

Hardware Configuration

A computer system in 3DC multi-target configuration requires the following three sites:

- Primary site for both TrueCopy Synchronous and UR
- Secondary site for TrueCopy Synchronous
- Secondary site for UR

Each of these sites requires one USP V storage system and Storage Navigator computers. It is recommended that Command Control Interface (CCI) should be installed in host computers in each of these sites.

The following explains hardware configuration for USP V storage systems in these sites.

- **USP V storage system in the primary site:**
 - The USP V storage system must have TrueCopy Synchronous installed.
 - The USP V storage system must have UR installed.
 - The USP V storage system must have a volume used as both a TrueCopy Synchronous primary volume and a UR primary data volume
 - The USP V storage system must have a UR master journal volume.
- **USP V storage system in the TrueCopy Synchronous secondary site:**
 - The USP V storage system must have TrueCopy Synchronous installed.
 - The USP V storage system must have UR installed.
 - The USP V storage system must have a TrueCopy Synchronous secondary volume.

Note: This volume might be used as a UR primary data volume when disaster recovery is performed at the primary site. For details, see section Recovering from Primary Site Disaster in 3DC Multi-Target Configuration.

- The USP V storage system must have a UR master journal volume.
Note: This volume might be necessary for creating a UR pair when disaster recovery is performed at the primary site or when you create a UR pair for delta resync operation.

USP V storage system in the UR secondary site:

- The USP V storage system must have UR installed.
- The USP V storage system must have a UR secondary data volume.
- The USP V storage system must have a UR restore journal volume.

Setup Procedure

This section explains the procedure for setting up a 3DC multi-target configuration using TrueCopy Synchronous and UR.

To set up a 3DC multi-target configuration:

1. Install TrueCopy Synchronous and UR. Configure ports and journal groups. Journal groups must be configured for storage systems in the three sites.
2. In the primary site, create a TrueCopy Synchronous pair.
3. Wait until the status of the TrueCopy Synchronous pair becomes PAIR.
4. In the primary site, create a UR pair.
When creating a UR pair, be sure to select a value from 1 to 3 as the mirror ID.
5. Wait until the status of the UR pair becomes PAIR.
If you want to create a UR pair for delta resync operation, go to the next step.
6. Specify the TrueCopy Synchronous S-VOL in TrueCopy Synchronous secondary site as the primary data volume, and specify the secondary data volume in UR secondary site as the secondary data volume to create the UR pair for delta resync operation.
Select any mirror ID from 1 to 3 except the one used in step 4.

Note: If you release the TrueCopy Synchronous pair that is created in step 2, the UR pair created in step 6 will be released as well. In addition, if you release the UR pair created in step 4, the secondary data volume of the UR pair created in step 6 will be deleted. In this case, you can only perform deleting operation on the remaining primary data volume.

For detailed information about installing and using TrueCopy Synchronous, please refer to *TrueCopy User's Guide*.

For detailed information about UR journal group configuration, see section Journal Group Operations. For detailed information about port configuration, see section Configuring Port Attributes.

Requirements for Creating UR Pair for Delta Resync Operation

To create a UR pair for delta resync operation, the followings are required.

- Create the pair in 3DC multi-target configuration
- Use TrueCopy Synchronous S-VOL in PAIR status as the primary data volume
- Use UR data volume in PAIR status as the secondary data volume
- Use the mirror ID from 1 to 3 that is not used by the secondary data volume
- Fulfill the conditions to combine UR with other functions in Table 3-9
- The system option mode 506 must be set to ON at all site

In addition to the abovementioned requirements, all UR pairs in the journal group must satisfy the following requirements when you create more than one UR pair for delta resync operation.

- Use the same mirror ID for all pairs
- Use the same restore journal group for all pairs

Note: For the latest information on the availability of the delta resync function, please contact your Hitachi Data Systems account.

Requirements for Performing Delta Resync Operation

To perform delta resync operation properly, you need to specify a UR pair that belongs to the journal group whose UR pairs meet the following requirements. Note that if any UR pair in the journal group does not meet the requirements, an error occurs even if the specified UR pair meets the requirements.

- Status of the primary data volume is HOLD
- Using the volume of TrueCopy Synchronous pair as the primary data volume

If you use TrueCopy Synchronous P-VOL, the pair status should be PAIR.
If you use TrueCopy Synchronous S-VOL, the pair status should be SSWS.

Note: Check the pair status after the CCI horctakeover command is executed.

- Secondary data volume uses two mirror IDs; the status of the secondary volume which uses one of the two mirror IDs is HOLD, and the status of the secondary volume which uses another mirror ID is PAIR, PSUS, or PSUE

- All differential data between the primary and secondary data volume is stored in the master journal group

Note: Right after you have created the UR pair for delta resync operation, or if the failure occurred at the primary site right after the recovery of TrueCopy Synchronous or UR pair, only a part of the differential data between the primary and secondary data volume may be stored in the master journal group.

When either of the abovementioned requirements is not fulfilled, usually an error occurs and delta resync operation will fail. Especially in the following cases, delta resync operation will not be performed since the necessary journal data does not exist:

- When the primary data volume in the UR pair for delta resync operation is updated after creating the UR pair, but UR primary data volumes that are also used as TrueCopy Synchronous primary data volumes are not updated in the primary site
- When the volumes in the primary site are updated after splitting TrueCopy Synchronous pair
- When the secondary data volume is updated after splitting UR pair
- As a result of the update of the primary data volume after splitting UR pair, when the capacity of the journal data exceeds 70% of the journal volume in TrueCopy Synchronous secondary site
- When the primary data volume in the UR pair for delta resync operation is updated and then the capacity of the journal data exceeds 70% of the journal volume in the primary site of UR pair for delta resync operation
- When no volumes (including volumes after failover or failback) in the primary site are updated after creating the UR pair for delta resync operation
- When delta resync operation is performed within about one minute after a volume in the primary site is updated after the UR pair for delta resync operation is created

Note: By specifying an option, you can copy the whole data in the primary volume to the secondary volume and avoid the error when only a part of the differential data between primary and secondary data volume is stored in the journal volume. For detailed information about the option, see section Changing Options for a Journal Group.

Changing to 3DC Multi-target Configuration after Recovering from Primary Site Failures

If you follow the instructions in section Recovering from Primary Site Disaster in 3DC Multi-Target Configuration and then the system does not change into 3DC cascading configuration, you can change the system into 3DC multi-target configuration after removing failures from the primary site and other locations. In the resulting 3DC multi-target configuration, the former TrueCopy secondary volume will be the primary volume.

To change the system into 3DC multi-target configuration:

1. Use CCI to release the former UR pair.
2. If pair resynchronization has not been performed to reverse the primary/secondary relationship of a TrueCopy Synchronous pair (i.e., if the pair is suspended), use CCI to resynchronize the pair, so that the primary/secondary relationship will be reversed.

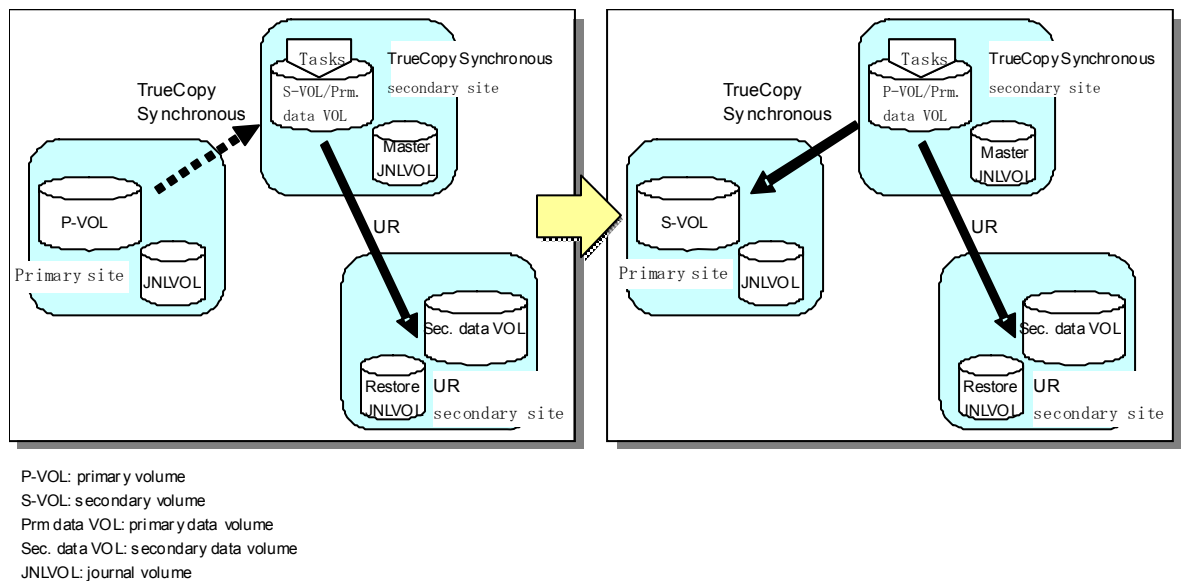


Figure 3-21 Changing into 3DC Multi-target Configuration After Recovering from Primary Site Failures

Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC Cascading Configuration)

If you follow the instructions in section Recovering from Primary Site Disaster in 3DC Multi-Target Configuration and then the system is changed to 3DC cascading configuration, you can transfer your business tasks back to the primary site after removing failures from the primary site and other locations.

To transfer your business tasks back to the primary site, follow the procedure below. Command Control Interface (CCI) is used in this procedure:

1. Stop business tasks at the TrueCopy Synchronous secondary site.
2. Use CCI to execute the horctakeover command on the TrueCopy Synchronous pair.

The system configuration returns to 3DC multi-target configuration.

3. Use the primary volume in the primary site to resume your business tasks.

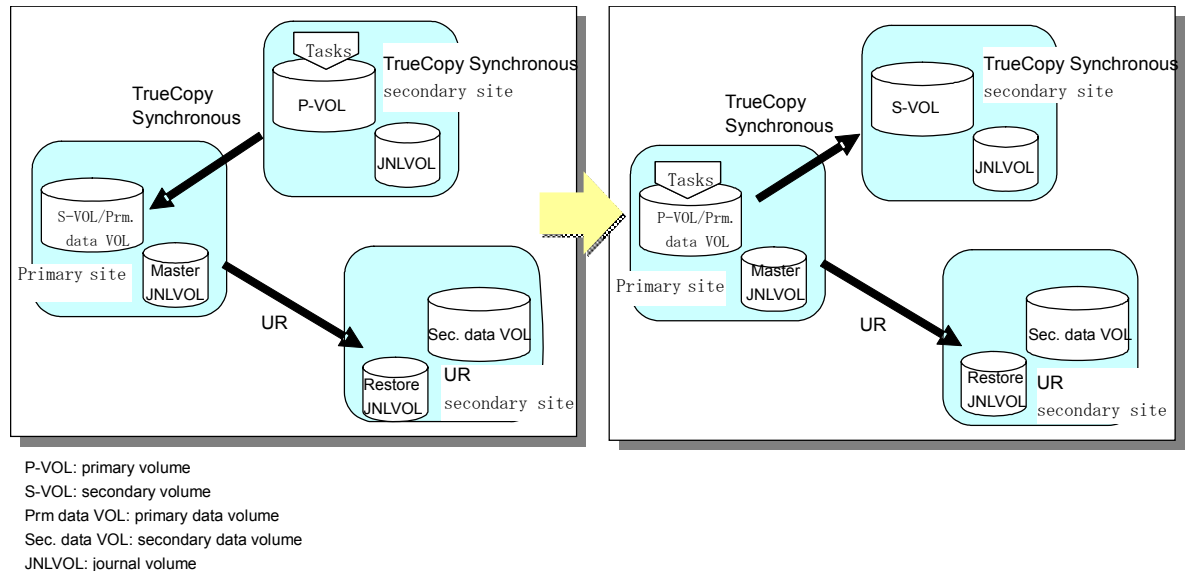


Figure 3-22 Transferring Business Tasks from the TrueCopy Secondary Site to the Primary Site (in 3DC Cascading Configuration)

Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC multi-target configuration)

If you remove failures from the primary site and other locations and then the system is changed to 3DC multi-target configuration, you can transfer your business tasks back to the primary site.

To transfer your business tasks back to the primary site, follow the procedure below. Command Control Interface (CCI) is used in this procedure:

1. Use CCI to release the current UR pair, which extends between the TrueCopy secondary site and the UR secondary site.
2. Stop business tasks at the TrueCopy Synchronous secondary site.
3. Use CCI to execute the **horctakeover** command on the TrueCopy Synchronous pair.
4. Use the primary volume at the primary site to resume your business tasks.

5. Use CCI to create a UR pair. Ensure that the UR primary data volume is a volume in the primary site, and that the UR secondary data volume is a volume in the UR secondary site.

The system configuration returns to the original 3DC multi-target configuration.

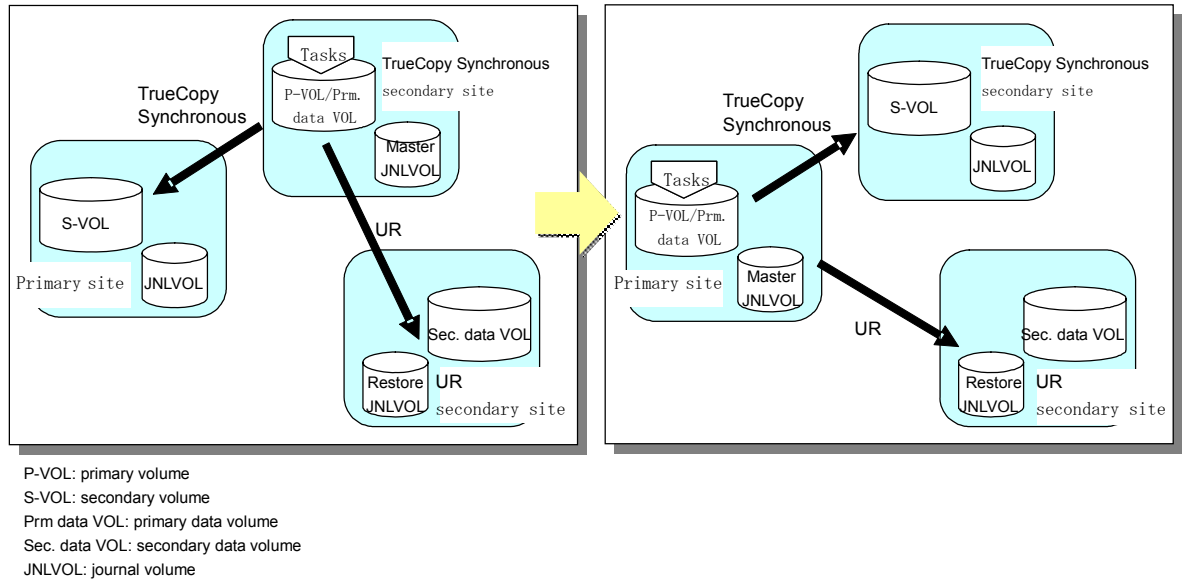


Figure 3-23 Transferring Business Tasks from the TrueCopy Secondary Site to the Primary Site (in 3DC Multi-target Configuration)

Transferring Business Tasks from TrueCopy Secondary Site to Primary Site (When Delta Resync Operation is Performed in 3DC Multi-target Configuration)

If you remove failures from the primary site and other locations and then the system is changed to 3DC multi-target configuration (see section Recovering from Primary Site Failures (When Delta Resync Operation is Performed)), you can transfer your business tasks back to the primary site.

To transfer your business tasks back to the primary site, follow the procedure below. Command Control Interface (CCI) is used in this procedure:

1. If the TrueCopy Synchronous pair is suspended, use CCI to resynchronize the pair and reverse the primary and secondary volume.
2. Stop business tasks at the TrueCopy Synchronous secondary site.
3. Use CCI to execute the horctakeover command on the TrueCopy Synchronous pair.
4. Use the primary volume at the primary site to resume your business tasks.

- Use CCI to resynchronize the volumes in TrueCopy Synchronous secondary site. When you perform delta resync operation on the volume in TrueCopy Synchronous secondary site, the status of the data volumes of UR pair changes as explained in the following table.

Table 3.12 Changes of UR Pair Status by Delta Resync Operation (When Business Tasks are Transferred from TrueCopy Secondary Site to the Primary Site)

UR Pair	Pair Status before Delta Resync Operation		Pair Status after Delta Resync Operation	
	Primary Data Volume	Secondary Data Volume	Primary Data Volume	Secondary Data Volume
UR pair between TrueCopy Synchronous primary site and UR secondary site	HOLD	HOLD	PAIR or COPY	PAIR or COPY
UR pair between TrueCopy Synchronous secondary site and UR secondary site	PAIR, PSUS, or PSUE	PAIR, PSUS, or PSUE	HOLD	HOLD

Caution on transferring business tasks back to the primary site when using delta resync:

After you remove failures from the primary site and other locations and then the system is changed to 3DC multi-target configuration (see section Recovering from Primary Site Failures (When Delta Resync Operation is Performed)), the status of a pair between the primary site and the UR secondary site may be abnormal. If the status of the pair remains abnormal, you will be unable to transfer your business tasks back to the primary site. To solve this problem, please see Table 3.13, check the pair status and then perform a problem-solving operation before transferring business tasks back to the primary site.

Table 3.13 Pair Status and Operation after Recovery of the Primary Site

No.	If the Status of the Pair in the Primary Site is	and the Status of the Pair in the UR Secondary Site is	Please Perform the Following Operation before Transferring Business Tasks Back to the Primary Site
1	COPY	HOLD	Make sure that the status of the pair in the primary site is PSUE or PSUS , and then perform the operation explained in the row No. 3 of this table.
2	PAIR	HOLD	Perform the operation explained in the row No. 1 of this table.
3	PSUE or PSUS	HOLD	Step 1: Release the UR pair from the primary site. Step 2: Make sure that all the pairs belonging to the journal group in the primary site are released. Step 3: Create a UR delta resync pair that extends from the primary site to the UR secondary site. For information about requirements for creating the pair, see section Requirements for Creating UR Pair for Delta Resync Operation.
4	HOLD	HOLD	The pair is in normal condition. You can transfer business tasks as described earlier in this section.

5	HLDE	HOLD	Change the status of the HLDE pair back to HOLD . For information about how to do this, see section Restoring a Pair of Data Volumes (Pairresync).
6	SMPL	HOLD	Step 1: Release the pair in HOLD status from the UR secondary site. Step 2: Create a UR delta resync pair that extends from the primary site to the UR secondary site. For information about requirements for creating the pair, see section Requirements for Creating UR Pair for Delta Resync Operation.

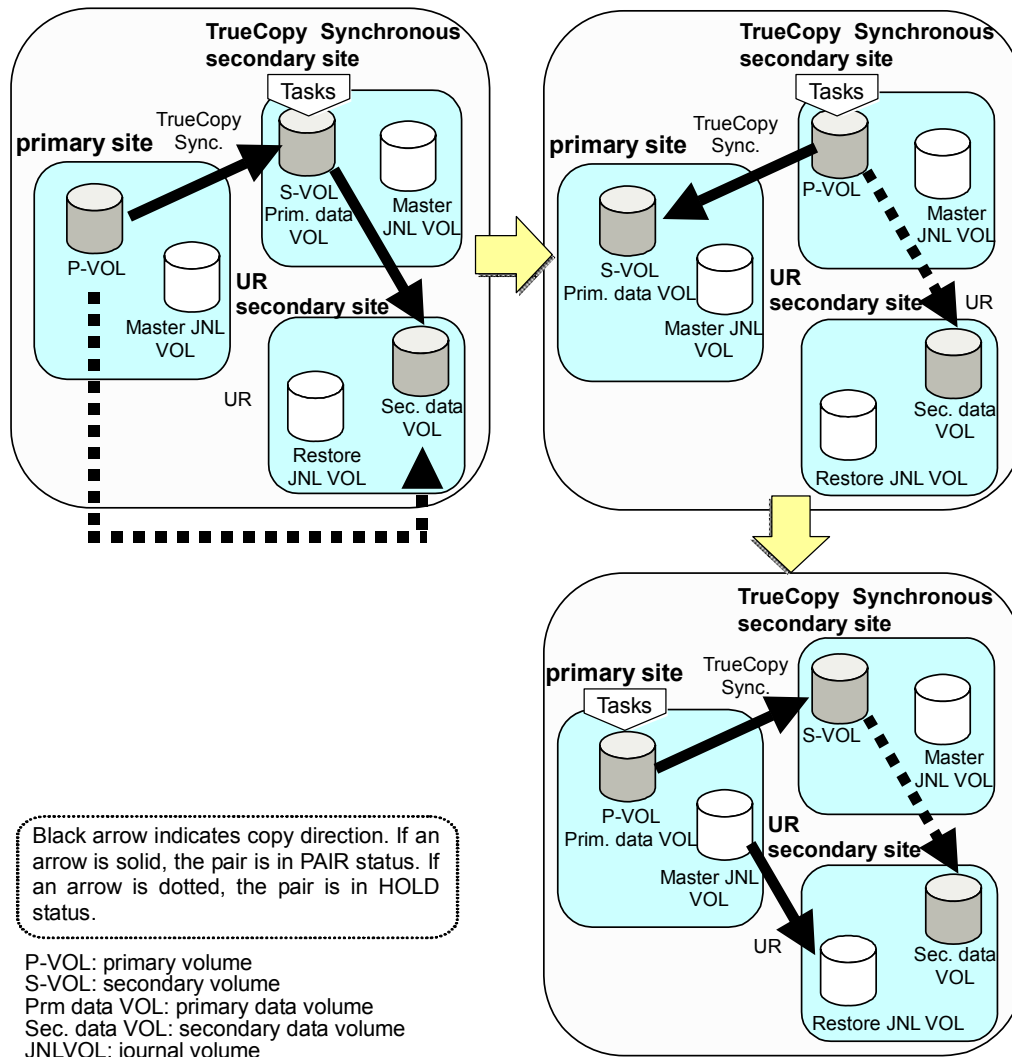


Figure 3-24 Transferring Business Tasks from the TrueCopy Secondary Site to the Primary Site (When Delta Resync Operation is Performed in 3DC Multi-target Configuration)

Note: For the latest information on the availability of the delta resync function, please contact your Hitachi Data Systems account.

Recovering from Failures in the Primary Site and the TrueCopy

Synchronous Secondary Site

If a disaster or failure occurs in both the primary site and the TrueCopy Synchronous secondary site in a 3DC multi-target configuration, you can resume your business tasks by using the secondary volume in the UR secondary site. After you transfer business tasks to the UR secondary site, you need to remove failures from the primary site and the TrueCopy Synchronous secondary site, so that you will be able to transfer business tasks back to the primary site.

To transfer your business tasks to the UR secondary site, follow the procedure below. Command Control Interface (CCI) is used in this procedure:

1. Use CCI to execute the **horctakeover** command on the UR pair.
2. Use the volume in the UR secondary site to resume your business tasks.

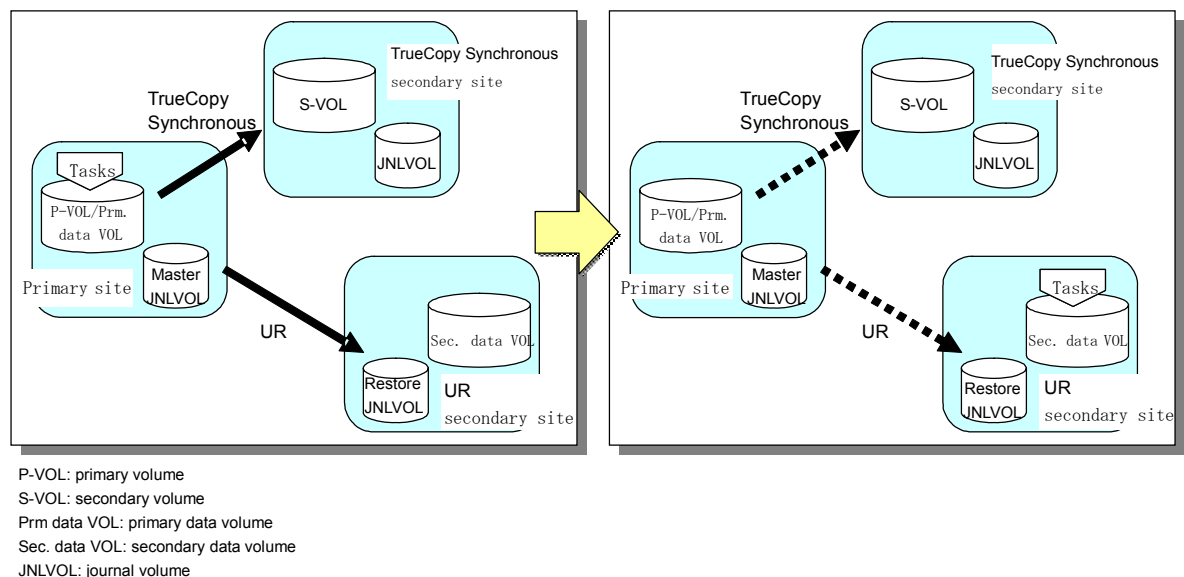


Figure 3-25 Recovering from Failures in the Primary Site and the TrueCopy Synchronous Secondary Site

Transferring Business Tasks from the UR Secondary to the Primary Site

If you follow the instructions in the previous section and then remove failures from the primary site and the TrueCopy Synchronous secondary site, you can transfer your business tasks back to the primary site.

To transfer your business tasks back to the primary site, follow the procedure below. Command Control Interface (CCI) is used in this procedure:

1. Use CCI to release the TrueCopy Synchronous pair.

2. If the UR pair has not been resynchronized to reverse the primary volume and the secondary volume and is suspended, use CCI to resynchronize the pair to reverse the primary and secondary volumes.

The secondary data volume in the former UR secondary site changes to the primary data volume. The primary data volume in the former primary site changes to the secondary data volume (see the upper right picture in Figure 3-26).

3. Stop business tasks at the UR secondary site.
4. Use CCI to execute the **horctakeover** command on the UR pair.
The UR primary and secondary volumes are reversed.
5. Use the primary volume in the primary site to resume business tasks.
6. Use CCI to resynchronize the TrueCopy pair.

The system configuration changes to the original 3DC multi-target configuration (see the lower-right picture in Figure 3-26).

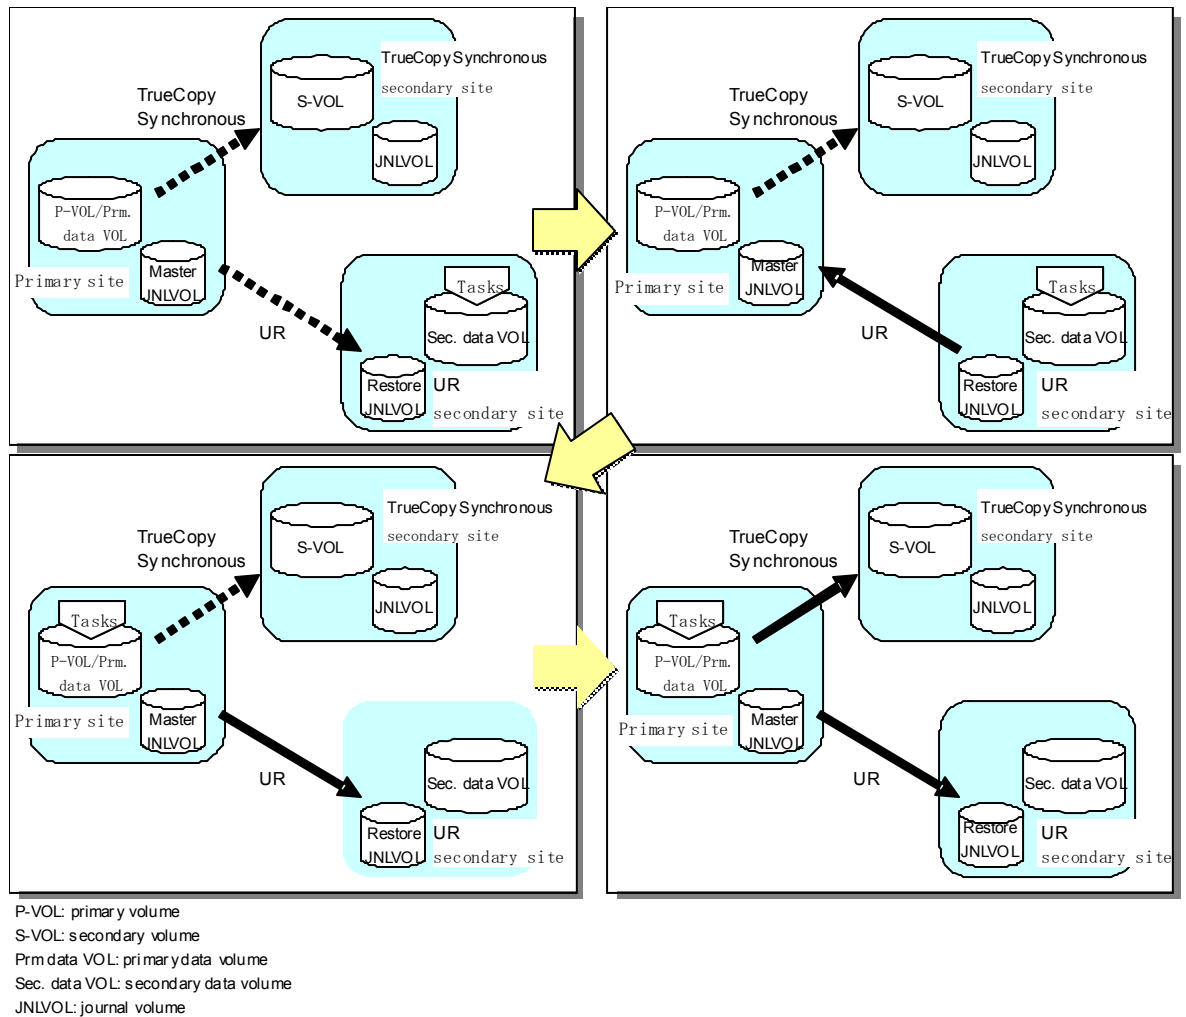


Figure 3-26 Transferring Business Tasks from the UR Secondary Site to the Primary Site

Data Retention Utility

You can create a Universal Replicator pair by using volumes to which access attribute is assigned by Data Retention Utility. However, you cannot specify a volume with "S-VOL Disable" attribute as the secondary data volume of the pair.

Table 3-12 shows pair status and availability of Data Retention Utility operation. Table 3-13 shows whether access attribute can be change when pair status of Universal Replicator primary data volume is PAIR or COPY.

Table 3-12 Pair Status and Availability of Data Retention Utility Operation

Volume	Pair Status	Data Retention Utility Operations	
		Changing Access Attribute	Referencing Access Attribute
Primary data volume	SMPL	OK	OK
	COPY	*1	OK
	PAIR	*1	OK
	PSUS	OK	OK
	PSUE	OK	OK
Secondary data volume	SMPL	OK	OK
	COPY	No	OK
	PAIR	No	OK
	PSUS	No*2	OK
	PSUE	No	OK

*1: See Table 3-13.

*2: If the Write option is set to the secondary data volume, the access attribute can be changed.

Table 3-13 Whether Access Attribute Can be Changed When Pair Status of Universal Replicator Primary Data Volume is PAIR or COPY

Access Attribute before Change	Resulting Access Attribute			
	Read/Write	Read Only	Protect	S-VOL Disable
Read/Write	-	No	No	OK
Read Only	OK	-	OK	OK
Protect	OK	OK	-	OK
S-VOL Disable	OK	No	No	-

Copying Access Attribute

If you create a Universal Replicator pair by using volumes to which the VMA is set with Data Retention Utility, the access attribute recorded in the VMA can be copied from of the primary data volume to the secondary data volume. The VMA, which is necessary for copying the access attribute, can be set to OPEN-V volumes only.

For further information on Data Retention Utility, please refer to the *Data Retention Utility User's Guide* or for the latest information on the availability of the VMA function, please contact your Hitachi Data Systems account team. Table 3-14 shows the VMA setting on the volume and the availability of Paircreate operation.

Table 3-14 VMA Setting and Availability of Paircreate Operation

Primary data volume	Secondary data volume
---------------------	-----------------------

	Volume with VMA	Volume without VMA
Volume with VMA	OK	NO
Volume without VMA	NO	OK

The following figure shows the process of the copying of access attribute when you create a pair using the volumes with VMA.

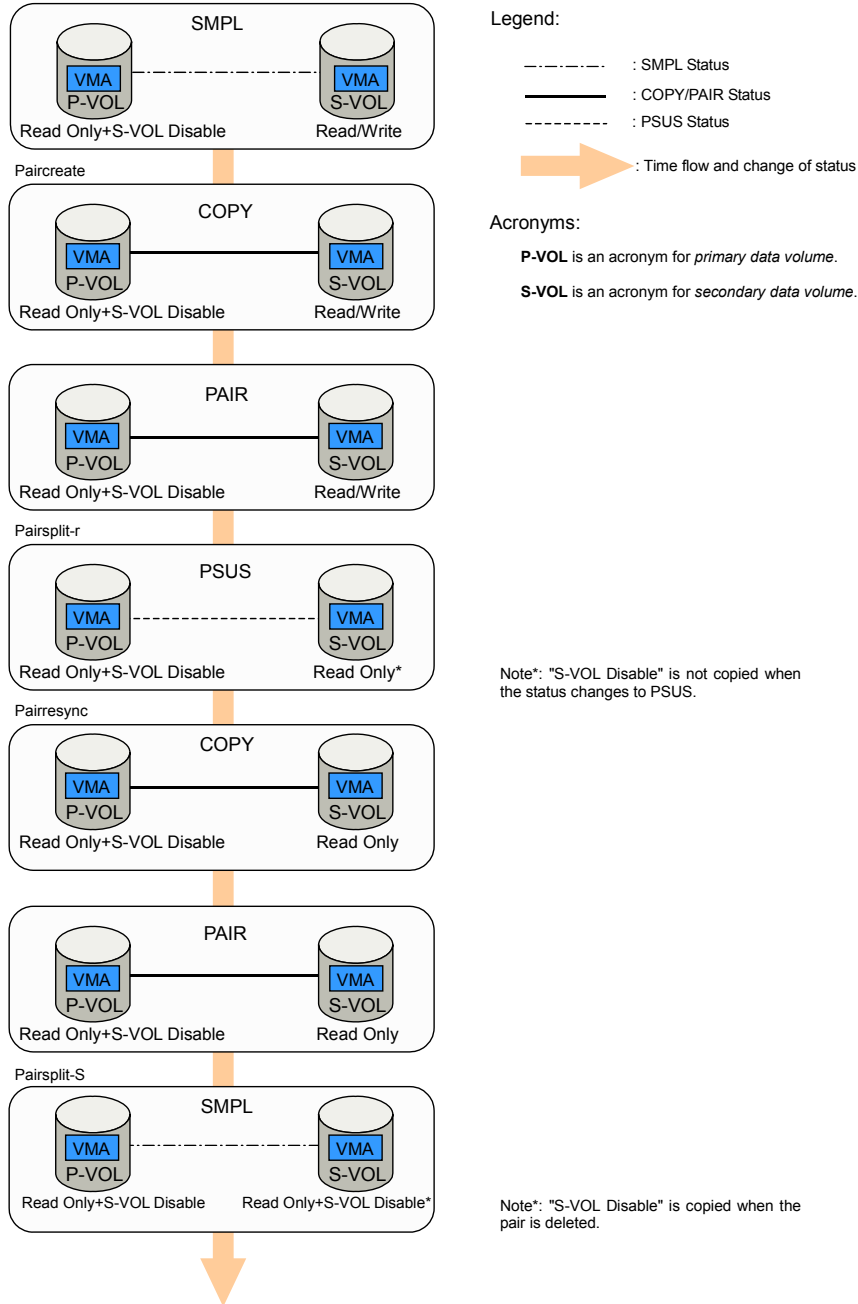


Figure 3.27 Copying Access Attribute

As you can see in Figure 3-27, if both primary and secondary data volumes have VMA, the VMA of the primary data volume is copied to the secondary data volume. In Figure 3-27, access attribute of the secondary data volume is Read/Write when the pair status is SMPL. After the VMA of the primary data volume is copied to the secondary data volume, access attribute of the secondary data volume becomes Read Only, which is same as the access attribute of the primary data volume. Note that when the pair status is PAIR, copied information of access attribute is not effective, so that access attribute of the secondary data volume is still Read/Write. To make the copied information of access attribute effective, you need to suspend or release the pair with the Pairsplit-r or Pairsplit-S operation and make the pair status PSUS or SMPL. However, even if the pair is suspended, S-VOL Disable does not become effective. To make S-VOL Disable effective, you need to release the pair. For details about pair status and the operations related to access attribute, see Table 3-15.

Note: If the pair is suspended or released during the COPY status, the access attribute of the primary data volume is not copied to the secondary data volume. To copy the access attribute of the primary data volume to the secondary data volume, make sure that the pair status is PAIR, and then suspend or release the pair.

Note: When the primary data volume is in the COPY or PAIR status, there are two access attributes which can be changed and cannot be changed according to the existing access attribute. For more information about the variable access attribute, see Table 3-16.

Table 3-15 Pair Status and Availability of Data Retention Utility Operation

Volume	Pair Status	Data Retention Utility Operations	
		Changing Access Attribute	Referencing Access Attribute
Primary data volume	SMPL	OK	OK
	COPY	*1	OK
	PAIR	*1	OK
	PSUS	OK	OK
	PSUE	OK	OK
Secondary data volume	SMPL	OK	OK
	COPY	No	OK
	PAIR	No	OK
	PSUS	No*2	OK
	PSUE	No	OK

*1: See Table 3.17.

*2: If the Write option is set to the secondary data volume, the access attribute can be changed.

Table 3-16 Variable Access Attribute (Pair Status of Universal Replicator Primary Data Volume is PAIR or COPY)

Access Attribute before Change	Variable Access Attribute			
	Read/Write	Read Only	Protect	S-VOL Disable
Read/Write	-	No	No	OK
Read Only	OK	-	OK	OK
Protect	OK	OK	-	OK
S-VOL Disable	OK	No	No	-

License Capacity of Data Retention Utility

To create a Universal Replicator pair using the volumes with VMA, if the primary data volume has the access attribute for which the license of Data Retention Utility is necessary, you need the license capacity of Data Retention Utility for secondary data volumes in the storage system on the secondary site. This is because when the copying of access attributes is completed, primary and secondary data volumes will have the same access attributes. For details about copying access attribute, see Figure 3-27.

Note: If the license capacity of Data Retention Utility is insufficient in the storage system on the secondary site, you may not be able to create Universal Replicator pairs. For details about the license capacity of Data Retention Utility, please refer to the *Data Retention Utility User's Guide*.

Performance Monitor

The Performance Monitor software product provides detailed information on the I/O activity and hardware performance of the USP V storage systems. Performance Monitor can be used to monitor the USP V storage systems which will be (or already are) performing UR operations. The storage system usage and performance data collected and displayed by Performance Monitor enables you to:

- Identify the optimum timings to perform UR data duplexing operations (e.g., during periods of light system I/O activity),
- Determine the best locations for the UR S-VOLs (e.g., in parity groups with less frequently accessed volumes to avoid bottlenecks of backend activity),
- Monitor storage system performance during UR operations and during testing activities.

Performance Monitor data collection does not affect storage system operations in any way. However, Performance Monitor data collection can cause a significant amount of traffic on the USP V-internal LAN, especially when Performance Monitor is collecting lots of LDEV data. To prevent timeouts from occurring while you are performing UR (or other) operations on the Storage Navigator PC, you may want to decrease the Performance Monitor data collection activities before using the USP V Storage Navigator software. This reduces USP V-internal LAN traffic, so that inquiries and commands issued by the Storage Navigator PC can be processed more quickly.

For further information on Performance Monitor, please refer to the *Performance Manager User's Guide*.

To perform UR (or other) Storage Navigator operations while Performance Monitor is collecting data for one or more storage systems on the same USP V-internal LAN:

1. If Performance Monitor is collecting lots of LDEV data, consider disabling Performance Monitor LDEV data collection for one or more storage systems before using the Storage Navigator software. Please refer to the Performance Manager User's Guide for instructions on disabling LDEV data collection.
2. If Performance Monitor is collecting data for more than three storage systems on the USP V-internal LAN, consider disconnecting Performance Monitor from one or more storage systems before using the Storage Navigator software. Please refer to the Performance Manager User's Guide for instructions on disconnecting Performance Monitor.
3. After you have disabled LDEV data collection and/or disconnected Performance Monitor, you can connect to the storage system using USP V Storage Navigator software, launch the UR software, and perform TrueCopy operations (e.g., create pairs, view pair status).
4. When you are finished performing UR operations, exit the UR software, and disconnect USP V Storage Navigator software from the storage system.
5. After you have disconnected the USP V Storage Navigator PC from the storage system, you can re-enable Performance Monitor data collection.

Planning of Journal Volumes

Data transfer speed for journal groups is affected by specifications of journal volumes that the journal groups use. Therefore, you need to think about specifications of journal volumes to be used by journal volumes, in order to achieve the data transfer speed that you want.

This section explains what you need to consider when deciding on journal volume specifications.

The following factors affect the speed for writing to and reading from journal volumes, particularly when a temporary communication path failure occurs between the primary storage system and the secondary storage system, or when the amount of data transferred from hosts to the primary storage system is increasing:

- RAID configuration of the RAID groups that will contain journal volumes
- Types of physical volumes in the RAID groups that will contain journal volumes
- Frequency of access to non-journal volumes in the RAID groups that will contain journal volumes
- Data transfer speed required for the non-journal volumes mentioned above
- Disk usage rate for RAID groups

The following factor affects the time during which data transfer with hosts can continue without being influenced by a temporary communication path failure between the primary storage system and the secondary storage system, or without being influenced by an increase in the data to be transferred from hosts to the primary storage system.

- Capacity of journal volumes

Data Transfer Speed Required for Journal Volumes

Figure 3-27 illustrates how the data transfer speed (i.e., the amount of data to be transferred per unit of time) will change as time elapses, citing different types of data transfer speed with UR. Data transfer speed between hosts and the primary storage system goes through two phases; in one phase the data transfer speed remains almost unchanged, and in another phase the data transfer speed increases temporarily.




Figure 3-27 Data Transfer Speed with UR (Data Transfer Speed of Journal Volumes)

As illustrated in Figure 3-27, the data transfer speed (i.e., the speed for reading and writing) of journal volumes in the master journal group must exceed the amount of temporary increase in data to be transferred. If the data transfer speed of journal volumes is below the amount of temporary increase in data to be transferred, the journal data for the temporary increase in data to be transferred will not be stored into journal volumes timely.

In Figure 3-27, data transfer speed between the primary storage system and the secondary storage system indicates transfer speed of journal data between the primary storage system and the secondary storage system. For details, see section Journal Volumes in Restore Journal Groups.

RAID Group Configuration and Journal Group Configuration

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

The data transfer speed of a journal volume depends on the data transfer speed of the RAID group that the journal volume belongs. One RAID group can consist of one or more volumes, including journal volumes. Therefore, if frequent accesses are made to non-journal volumes in a RAID group, relatively fewer accesses can be made to journal volumes in the same RAID group, and this can cause drop in the data transfer speed of journal volumes. To avoid drop in the data transfer speed of journal volumes, you will need to consider, for example, relocating the journal volumes and the frequently accessed non-journal volumes (i.e., placing the journal groups in one RAID group and placing the frequently accessed non-journal volumes in another RAID group).

Arranging Journal Volumes

The following indicates the data transfer speed of journal volumes mentioned earlier in section Data Transfer Speed Required for Journal Volumes.

If a host attempts to write data of "*aa*" kilobytes and the number of write I/Os per second is "*bb*", the data transfer speed required for journal volumes is calculated as follows:

$$\text{Data transfer speed of journal volumes (MB/sec)} > aa \times bb \div 1,024$$

The data transfer speed of journal volumes must exceed the data transfer speed from hosts to the primary storage system. You must consider the data transfer speed (i.e., the speed for writing and reading) required for journal volumes, determine physical volume configuration and RAID configuration of RAID groups, and create journal volumes in the RAID groups.

Journal Volume Capacity

In Figure 3-28, the size of the shaded area indicates the amount of journal data to be stored in journal volumes as a result of temporary increase in data transferred.

If a temporary communication path failure occurs between the primary storage system and the secondary storage system, journal transfers between the primary storage system and the secondary storage system will stop temporarily. If hosts transfer data to the primary storage system while journal transfers between the primary storage system and the secondary storage system are stopped temporarily, all the data transferred from hosts to the primary storage system will be stored as journal data in journal volumes.




Figure 3-28 Data Transfer Speed with UR (Influence on Journal Volume Capacity)

The following factors determine the required journal volume capacity:

- the period of time during which data transfer can continue between hosts and the primary storage system when a temporary increase in transferred data occurs or when a communication path failure occurs between the primary and secondary storage system.
- the data transfer speed for the period of time that is mentioned above

To determine the journal volume capacity, use the following formula. In this formula, V_{H-M} is the data transfer speed between hosts and the primary storage system. V_{M-R} is the data transfer speed between the primary storage system and the secondary storage system. t is the length of time during which data transfer can continue.

$$\text{Journal volume capacity} > V_{H-M} - V_{M-R} \times t$$

If you want to calculate journal volume capacity that is required when a communication path failure occurs between the primary storage system and the secondary storage system, please assign 0 (zero) to V_{M-R} .

The total capacity of journal volumes in each journal group must exceed the value illustrated above.

Journal Volumes in Restore Journal Groups

When a failure occurs at a host, one of the failure recovery measures is to reverse the primary data volume and the secondary data volume (i.e., change the copy direction). If you want to reverse the primary data volume and the secondary data volume, usually you must ensure that the data transfer speed is the same before and after you reverse these data volumes. Therefore, when you do planning on master journal volumes and restore journal volumes, you need to apply the same scheme to both master and restore journal volumes.

If you do not want to reverse the primary data volume and the secondary data volume, you will be able to cope with a temporary increase in data transfers and a communication path failure between the primary storage system and the secondary storage system, if the master journal volume satisfies conditions mentioned earlier. Therefore, the data transfer speed and the volume capacity required for restore journal volumes are smaller than those required for master journal volumes. Note that the data transfer speed for journal volumes must exceed the data transfer speed in "normal status".

Factors that Determine Data Transfer Speed between the Primary and Secondary Storage Systems

The speed of data transfer between the primary storage system and the secondary storage system depends on the following factors:

- The bandwidth of data transfer paths (see the next sections)
- The journal transfer speed of the DKC (see the next section)

The data transfer speed between the primary storage system and the secondary storage system must exceed the data transfer speed in "normal status" mentioned in Figure 3-29 below. If a temporary increase in transferred data occurs and the data transfer speed between the primary storage system and the secondary storage system is exceeded, the excess data will be stored temporarily in journal volumes.




Figure 3-29 Data Transfer Speed with UR (Data Transfer Speed between the Primary Storage System and the Secondary Storage System)

Bandwidth of Data Transfer Paths

The bandwidth (i.e., the data transfer speed) of data transfer paths between the primary storage system and the secondary storage system depends on the user's environment. The bandwidth must exceed the data transfer speed in "normal status" mentioned earlier.

Journal Transfer Speed of the DKC

The journal transfer speed of the DKC depends on the number of paths used for transferring journals. The number of paths for transferring journals must be so enough that the journal transfer speed exceeds the data transfer speed in the "normal status" mentioned earlier. The maximum number of paths for each pair of the primary DKC and the secondary DKC is 8.

Performing Remote Copy Between TagmaStore USP/NSC and USP V

UR can be used to perform remote copy operations between USP V and TagmaStore USP/NSC. Data can be copied from USP V to TagmaStore USP/NSC, or vice versa.

This section explains the following:

- System option mode (see the next section)
- Logical Storage system (LDKC) that can be connected to TagmaStore USP/NSC (see section Logical Storage System (LDKC) Connected to TagmaStore USP/NSC)
- Volume pair that can create pairs (see section Volume that Can Be Used to Create Pairs)
- Connection with TagmaStore USP/NSC in 3DC remote copy configuration (see section Connection with TagmaStore USP/NSC for 3DC Remote Copy Configuration)

System Option Mode

When you execute remote copy by connecting USP V with TagmaStore USP/NSC, system option mode XXXXX needs to be set to OFF. For detailed information about the system option mode XXXXX, see Table 3-17.

Table 3-17 System Option Mode XXX

Mode	Description
XXXXX	<p>ON: The format of the serial number of the logical storage system (LDKC) 00 of USP V is five-digit alphanumeric characters (serial number of the LDKC). The serial number reported to or displayed in the host to SVP is the is five-digit alphanumeric value.</p> <p>OFF: The format of the serial number of the logical storage system (LDKC) 00 of USP V is five-digit alphanumeric characters (serial number of the storage system). The serial number reported to or displayed in the host to SVP is the is five-digit alphanumeric value. The serial number for LDKC00 and the serial number for the storage system is the same.</p> <p>Note: The default value for mode XXXXX is OFF.</p>

Logical Storage System (LDKC) Connected to TagmaStore USP/NSC

To execute remote copy from USP V to TagmaStore USP/NSC or vice versa, configure a logical path between LDKC of USP V and TagmaStore USP/NSC. In this case, the logical path can be configured only between LDKC00 of USP V and TagmaStore USP/NSC. Therefore, you cannot configure a logical path between LDKC01 of USP V and TagmaStore USP/NSC. (see Figure 3-30)

More than one TagmaStore USP/NSC can be connected to LDKC00 of USP V.

Note: There are only two LDKCs for the current version, which are LDKC00 and LDKC01. LDKC01 and TagmaStore USP/NSC cannot be connected.

Note: The steps to configure a logical path between USP V and TagmaStore USP/NSC are the same as the steps to set logical paths between volumes of USP V. For detailed information about the steps to configure logical paths, see *Configuring Storage Systems and Logical Paths*.




Figure 3-30 Logical Path Between LDKC00 of USP V and TagmaStore USP/NSC

Volume that Can Be Used to Create Pairs

When you create a UR volume pair consisting of a USP V volume and a TagmaStore USP/NSC volume, the USP V volume must be a volume of LDKC00. The CU:LDEV number of LDKC00 is between 00:00 to 3F:FF.

Note: 16,384 volumes of LDKC00 (The CU:LDEV number is between 00:00 to 3F:FF) can be used for volume pairs in the current version.

Note: The steps to use data volume pairs in a system configuration where USP V and TagmaStore USP/NSC are connected are the same as the steps to use data volume pairs of USP V. For detailed information about the steps to use data volume pairs, see Data Volume Pairs.

Connection with TagmaStore USP/NSC for 3DC Remote Copy Configuration

USP V can execute remote copy in 3DC cascading configuration or 3DC multi-target configuration by combining UR with TC Synchronous. USP V and TagmaStore USP/NSC can be connected in 3DC cascading configuration or 3DC multi-target configuration. When you connect USP V and TagmaStore USP/NSC, there are no restrictions for the combination of the primary and secondary sites.

Note: For detailed information about the 3DC cascading configuration, see section TrueCopy Synchronous (3DC Cascading Configuration). For detailed information about the 3DC multi-target configuration, see section TrueCopy Synchronous (3DC Multi-target Configuration).

Using the Universal Replicator GUI

This chapter describes how to use the UR graphical user interface:

- Journal Operation Window
- Pair Operation Window
- DKC Operation Window
- Usage Monitor Window
- History Window
- Optional Operation Window

Journal Operation Window

To configure journal groups, use the Journal Operation window of UR.

To display the Journal Operation window, do either of the following:

- If Universal Replicator has not been started:
 - a. Use your Web browser to display the storage device list. In the storage device list, select the storage system to log in. Enter a user name and the password, and then click **OK**.

The Storage Navigator main window is displayed.

Note: For detailed instructions on this step, please refer to *Storage Navigator User's Guide*.

- b. Click **Go, Universal Replicator** and then **Journal Operation** on the menu bar of the Storage Navigator main window.

UR starts and the Journal Operation window is displayed.

- **If Universal Replicator has already been started:**

1. Select the Journal Operation tab.

The Journal Operation window is displayed.

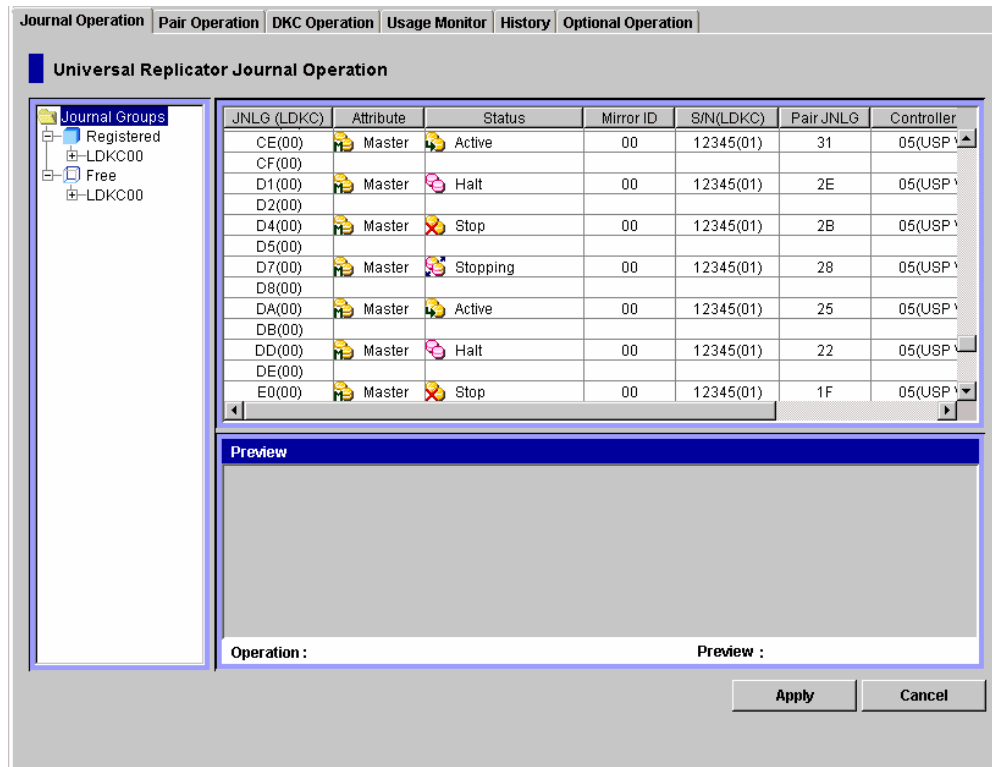


Figure 4-1 Journal Operation Window

The Journal Operation window displays the following:




- **Tree:** Lists journal groups in the local storage system, which is a storage system you are logged in. **Note:** The tree displays journal groups used with UR, and does not display journal groups used with Universal Replicator for z/OS®.

- **Journal Groups:** This item is located at the top of the tree. When you select this item, the upper-right list displays journal groups in the local storage system.
- **Registered:** When you select this item, the upper-right list displays journal groups in which journal volumes are registered.

When you double-click this item, LDKCs are displayed in the tree. When you select an LDKC in the tree, the upper-right list displays information about the LDKC.

When you double-click the LDKC, the tree displays journal groups in which journal volumes are registered. When you select a journal group in the tree, information about the journal group appears in the upper-right list. If you have selected the journal group which uses multiple mirror IDs, the list displays information about the data volume which belongs to the journal group whose mirror ID is not in **Hold** or **Hold(Failure)** status.

The journal group icons are:

-  A journal group in initial status.
Journal volumes are registered in this journal group, but no data volumes (primary data volumes nor secondary data volumes) are registered in this journal group.
-  A master journal group.
Journal volumes and primary data volumes are registered in this journal group.
-  A restore journal group.
Journal volumes and secondary data volumes are registered in this journal group.

- **Free:** When you select this item, the upper-right list displays journal groups in which no journal volumes are registered.

When you double-click this item, LDKCs are displayed in the tree. When you double-click the LDKC, the tree displays journal groups in which no journal volumes are registered.




For instructions on how to register journal volumes in journal groups, see section Registering Journal Volumes in a Journal Group.

- **Upper-right list:** If a master journal group or a restore journal group is selected in the tree, the upper-right list displays a list of *mirrors*. A mirror is a combination of a master journal group and a restore journal group.










If another journal group is selected in the tree, the upper-list displays information about the selected journal group.

One row in this list represents one mirror (or one journal group).

- **JNLG (LDKC):** Indicates the number of a journal group in the local storage system. The LDKC number is enclosed in the parentheses following the journal group number.
- **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 no data volumes (primary data volumes nor secondary data volumes) are registered in this journal group.
 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.

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 no data volumes (primary data volumes nor secondary data volumes) are registered in this journal group. 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: <ul style="list-style-type: none"> 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. <p>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 window. For details on how to restore them, see section Restoring a Pair of Data Volumes (Pairresync).</p>
 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: <ul style="list-style-type: none"> 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.
 Hold	A UR pair for delta resync operation is created.
 Hold(Failure)	An error occurred with the UR pair for delta resync operation.
Blank	No data volumes are registered in this journal group.

- **Mirror ID:** Indicates a mirror ID.

This column is blank if the attribute of the journal group is neither Master nor Restore.

- **S/N (LDKC):** Indicates the serial number of the remote storage system. The LDKC number is enclosed in the parentheses following the serial number.
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 (model name that indicates the model) 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 USP V storage system is **5**.
- **Path Gr. ID:** Displays *path group IDs*, which are used to identify path groups (i.e., groups of logical paths). Up to eight logical paths can be registered in one path group.
This column is blank if the attribute of the journal group is neither Master nor Restore.
- **CTG:** Indicates the number of a consistency group to which the mirror belongs.
This column is blank if the journal group does not have any volume pair.

- **Preview:** Displays detailed information about operations that will occur when you click **Apply**.
- **List:** Displays changes that have been made in the window. When you change settings in the window, the changes will be displayed in the **Preview** list before the changes are applied to storage systems. If you are sure that information in the **Preview** is correct, click the **Apply** button to apply the settings that you have made.

Note: When the **Preview** list displays changes that have been made, you can only perform the same type of operation that you have been doing and cannot perform most of other operations (note that you can only view detailed information about journal groups even when the **Preview** list displays changes). For example, if you are trying to delete journal groups and the **Preview** list is displaying the journal groups that you want to delete, you are unable to perform any other operations; for example, you are unable to split mirrors and restore mirrors.

- **Operation:** Indicates the operation that will occur when you click **Apply**.

Edit JNL Volumes	Register or delete journal volumes
Change JNL Option	Change journal group options
Delete JNL Groups	Delete journal group
Pairsplit -r	Split a mirror
Pairresync	Restore a mirror

Pairsplit -S	Delete a mirror
Blank	Nothing will occur when you click Apply

- **Preview:** The number to the left of the slash (/) indicates the number of items displayed in the **Preview** list.

The number to the right of the slash indicates the maximum number of items that can be displayed in the **Preview** list.

For example, if **Preview 010/256** is displayed, you can add another 246 items to the **Preview** list.

- **Apply:** Applies settings in the Preview list to the storage systems.
- **Cancel:** Cancels settings in the **Preview** list.

Note: Information on the Journal Operation window will be updated when you do one of the following:

2. Select another tab and then reselect the **Journal Operation** tab.
3. Click **File, Refresh** on the menu bar of the Storage Navigator main window.
4. Select the **Apply** button.
5. Select modify mode when you are in view mode.

Pair Operation Window

To manipulate pairs of primary and secondary data volumes, use the Pair Operation window.

To display the Pair Operation window, do either of the following:

If Universal Replicator has not been started:

1. Use your Web browser to display the storage device list. In the storage device list, select the storage system to log in. Enter a user name and the password, and then click **OK**.

The Storage Navigator main window is displayed.

Note: For detailed instructions on this step, please refer to *Storage Navigator User's Guide*.

2. Click **Go, Universal Replicator** and then **Pair Operation** on the menu bar of the Storage Navigator main window.

UR starts and the Pair Operation window is displayed.

If Universal Replicator has already been started:

1. Select the **Pair Operation** tab.

The Pair Operation window is displayed.

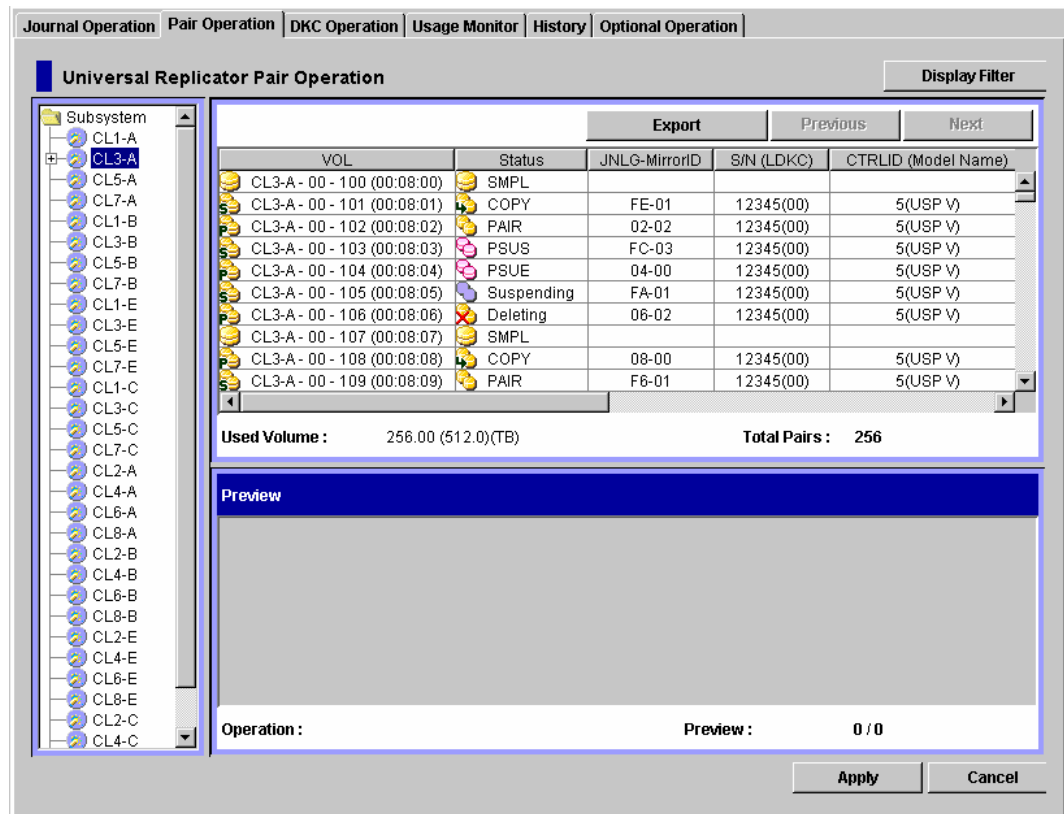


Figure 4-2 Pair Operation Window

Using the Universal Replicator GUI

The Pair Operation window displays the following:

- **Tree:** Lists ports in the local storage system (🔗). Below each port appear host groups or iSCSI targets (🔗). If you select a port, or a host group, or an iSCSI target, the list displays volumes for the port, or the host group, or the iSCSI target. You can select only one port, or one host group, or one iSCSI target at one time and cannot select two or more simultaneously.
- **Upper-right buttons:** The following buttons appear at the upper-right corner of the window:
 - **Display Filter:** Displays the Display Filter window (see Figure 7-1). This window allows you to narrow information in the list, so that the list displays only the volumes that satisfy certain conditions. For detailed information, see section Filtering Information in the List in the Pair Operation Window.
 - **Export:** Saves information about volume pairs in a text file. For detailed information, see section Saving Pair Status Information into a Text File.
 - **Previous:** The list can display up to 1,024 volume pairs simultaneously. If the number of volume pairs exceeds 1,024, you can use the **Previous** and **Next** buttons to display the remaining volume pairs in the list. The **Previous** button displays the previous 1,024 volume pairs in the list.
 - **Next:** The list can display up to 1,024 volume pairs simultaneously. If the number of volume pairs exceeds 1,024, you can use the **Previous** and **Next** buttons to display the remaining volume pairs in the list. The **Next** button displays the next 1,024 volume pairs in the list.
- **List:** Displays volumes in the local storage system. One row represents one volume.

By default, information in the list is arranged in order of port numbers. For volumes with more than one path, each path is displayed in a separate row.

The maximum number of rows in the list is 1,024. If the number of volumes exceeds the number of rows, you can use the **Previous** and **Next** buttons to view information about volumes that do not appear in the list currently.

VOL	Status	JNLG-MirrorID	S/N (LDKC)	CTRLID (Model Name)
CL3-A-00-100 (00:08:00)	SMPL			
CL3-A-00-101 (00:08:01)	COPY	FE-01	12345(00)	5(USP V)
CL3-A-00-102 (00:08:02)	PAIR	02-02	12345(00)	5(USP V)
CL3-A-00-103 (00:08:03)	PSUS	FC-03	12345(00)	5(USP V)
CL3-A-00-104 (00:08:04)	PSUE	04-00	12345(00)	5(USP V)
CL3-A-00-105 (00:08:05)	Suspending	FA-01	12345(00)	5(USP V)
CL3-A-00-106 (00:08:06)	Deleting	06-02	12345(00)	5(USP V)
CL3-A-00-107 (00:08:07)	SMPL			
CL3-A-00-108 (00:08:08)	COPY	08-00	12345(00)	5(USP V)
CL3-A-00-109 (00:08:09)	PAIR	F6-01	12345(00)	5(USP V)

Paired VOL	Pair JNLG	CTG	Err Lv.	Sync.	CLPR	Pair Copy Time
CLA-R-EE-7EE (00:EE:EE)	01	01	LU		00:CLPR0	
CL1-E-DD-7DD (00:DD:DD)	FD	02	Group	100 %	00:CLPR0	001:00:00
CLA-R-EE-7EE (00:EE:EE)	03	03	LU	028 %	00:CLPR0	
CL1-E-DD-7DD (00:DD:DD)	FB	04	Group	098 %	00:CLPR0	
CLA-R-EE-7EE (00:EE:EE)	05	05	LU		00:CLPR0	
CL1-E-DD-7DD (00:DD:DD)	F9	06	Group	000 %	00:CLPR0	
					00:CLPR0	
CL1-E-DD-7DD (00:DD:DD)	F7	08	Group	042 %	00:CLPR0	
CLA-R-EE-7EE (00:EE:EE)	09	09	LU		00:CLPR0	

Figure 4-3 The List in the Pair Operation Window

Note: Use the vertical and horizontal scroll bar if the information that you want to view is hidden and invisible.

– **VOL:** Indicates volumes in the local storage system. The icons are:

- This volume not paired with any other volume.
This volume is neither a primary data volume nor a secondary data volume.
- This volume is a primary data volume.
- This volume is a secondary data volume.

To the right of the icon appears information about an LU path, which is a path from hosts to the volume.

CL3-A - 00 - 002 (00:07:02)


① ② ③ ④


- ①Port that receives read and write requests from hosts.
In this example, read and write requests will be sent via the CL3-E port to the logical volume.
- ②Group number of the host group. A host group is a group of hosts that issue read requests and write requests.
- ③LUN (logical unit number).
- ④Logical volume ID. Indicate the number of "LDKC : CU : LDEV".


Figure 4-4 Information in The VOL Column in the Pair Operation Window


Note: If a 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*.


– **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.

 **HOLD** A UR pair for delta resync operation is created.

 **HLDE** An error occurred with the UR pair for delta resync operation.

– **JNLG-MirrorID:** Indicates a journal group number and a mirror ID. The number on the left of the hyphen (-) is a journal group number. The number on the right of the hyphen is a mirror ID.

This column is blank if the volume in the local storage system is neither a primary data volume nor a secondary data volume.

– **S/N (LDKC):** displays the serial number and the LDKC number of the remote storage system.

This column is blank if the volume of local storage system is neither a primary data volume nor a secondary data volume..

Note: The column of **S/N (LDKC)** can be blank while the pair is in transition to the Simplex status. To display the latest information in this column, refresh the screen.

– **CTRLID (Model Name):** Indicates the serial number and the controller ID of the remote storage system. The model name is enclosed by parentheses.

Note: The controller ID for a USP V storage system is **5**.

Note: The column of **CTRL ID** can be blank while the pair is in transition to the Simplex status. To display the latest information in this column, refresh the screen.

– **Paired VOL:** Indicates a data volume in the remote storage system.

This column indicates a primary data volume if the remote storage system is a primary storage system.

This column indicates a secondary data volume if the remote storage system is a secondary data volume.

The following is an example of an entry in this column:

CL1-E - 02 - 70D (01:FE:F2)

① ② ③ ④

① Port

② Group number of the host group

③ LUN (logical unit number).

④ Logical volume ID. Indicates the number of "LDKC : CU : LDEV".

Note: If a 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*.

Note: If a volume is a volume of TagmaStore USP/NSC, "00" is displayed as for the LDKC number

This column is blank if the volume in the local storage system is neither a primary data volume nor a secondary data volume.

- **Pair JNLG:** Indicates the journal group number for the remote storage system.

This column is blank if the volume in the local storage system is neither a primary data volume nor a secondary data volume.

- **CTG:** Indicates the consistency group number.

This column is blank if the volume in the local storage system is neither a primary data volume nor a secondary data volume.

- **Err Lv.:** Indicates the range of pair split on error.

Group If an error occurs with this pair, all the pairs in the mirror where this pair belongs will be split.

LU If an error occurs with this pair, only this pair will be split.

- **Sync.:** If the volume in the local storage system is a primary data volume, this column displays progress of an initial copy operation.

If the volume in the local storage system is a secondary data volume, this column displays information in the following ways:

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

Caution: If a failure in the initial copy operation causes the volume pair to be split, this column displays nothing. If a failure occurs in the initial copy operation, the Detailed Information window (Figure 7-4) displays the phrase "Initial copy failed".

In the following cases, this column will be blank.

- When the volume in the local storage system is neither a primary data volume nor a secondary data volume.
- When the status of the volume pair is **HOLD** or **HLDE**.
- **CLPR:** Indicates the number and the name of the CLPR where the data volume belongs.
- **Pair Copy Time:** The time taken for the copy operation (from the start of the operation to the end).

Note: The time that is shown in **Pair Copy Time** differs from the time that is shown in **Copy Time** on the **History** window.

To create a pair:

- a. MCU receives a request to create a pair.
- b. MCU receives a request to start the paircreate operation.
- c. The paircreate operation is started according to the conditions of initial copy priority and maximum initial copy activities.

Pair Copy Time on the Detailed Information dialog box shows the amount of time that has been needed from step 3 to the completion of the paircreate operation (i.e., the progress of the paircreate operation reaches 100%).

Copy Time on the **History** window shows the amount of time that has been needed from step 1 to the completion of the paircreate operation.

- **Used Volume:** Indicates the size of used volumes, and also indicates the licensed capacity. For example, if **12.34 (15.0) (TB)** is displayed, the licensed capacity for UR is 15.0 terabytes, and 12.34 terabytes of volumes are used.
Note: If the licensed capacity is unlimited, the sized of used volume is not displayed.
- **Total Pairs:** Indicates the current total number of data volume pairs.
- **Preview:** Displays detailed information about operations that will occur when you click **Apply**.
 - **List:** Displays changes that have been made in the window. When you change settings in the window, the changes will be displayed in the **Preview** list before the changes are applied to storage systems. If you are sure that information in the **Preview** is correct, select the **Apply** button to apply the settings that you have made.

Note: When the **Preview** list displays changes that have been made, you can only perform the same type of operation that you have been doing and cannot perform most of other operations (note that you can only view detailed information about pairs even when the **Preview** list displays changes). For example, if you are trying to release pairs and the **Preview** list is displaying the pairs that you want to release, you are unable to perform any other operations; for example, you are unable to split pairs and restore pairs.

- **Operation:** Indicates the operation that will occur when you select **Apply**.

Paircreate	Create pairs
Pairsplit-r	Split pairs
Pairresync	Restore pairs
Pairsplit-S	Release pairs
Change Pair Option	Change pair option(s)
Blank	Nothing will occur when you click Apply

- **Preview:** The number to the left of the slash (/) indicates the number of items displayed in the **Preview** list.

The number to the right of the slash indicates the maximum number of items that can be displayed in the **Preview** list.

For example, if **Preview 010/1096** is displayed, you can add another 4,086 items to the **Preview** list.

- **Apply:** Applies settings in the Preview list to the storage systems.
- **Cancel:** Cancels settings in the **Preview** list.

Note: Information on the Pair Operation window will be updated when you do one of the following:

2. Select another tab and then reselect the Pair Operation tab.
3. Click File, Refresh on the menu bar of the Storage Navigator main window.
4. Click the Apply button.
5. Select modify mode when you are in view mode.
6. Update the Display Filter window.
7. Click the **Previous** button or the **Next** button.

DKC Operation Window

To configure storage systems and logical paths, use the DKC Operation window.

To display the DKC Operation window, do either of the following:

- If Universal Replicator has not been started:
 - a. Use your Web browser to display the storage device list. In the storage device list, select the storage system to log in. Enter a user name and the password, and then click **OK**.
The Storage Navigator main window is displayed.
Note: For detailed instructions on this step, please refer to *Storage Navigator User's Guide*.
 - b. Click **Go, Universal Replicator** and then **DKC Operation** on the menu bar of the Storage Navigator main window.
UR starts and the DKC Operation window is displayed.

- **If Universal Replicator has already been started:**

1. Select the **DKC Operation** tab.

The DKC Operation window is displayed.

Note: DKC is an acronym for disk controller, which controls an entire storage system. UR windows use the word "DKC" to indicate a storage system. LDKC is an acronym for logical disk controller, and it also may be called logical DKC. LDKC is a controller that controls a logical storage system that exists in USP V. In the Universal Replicator window, the term "LDKC" indicates a logical storage system.

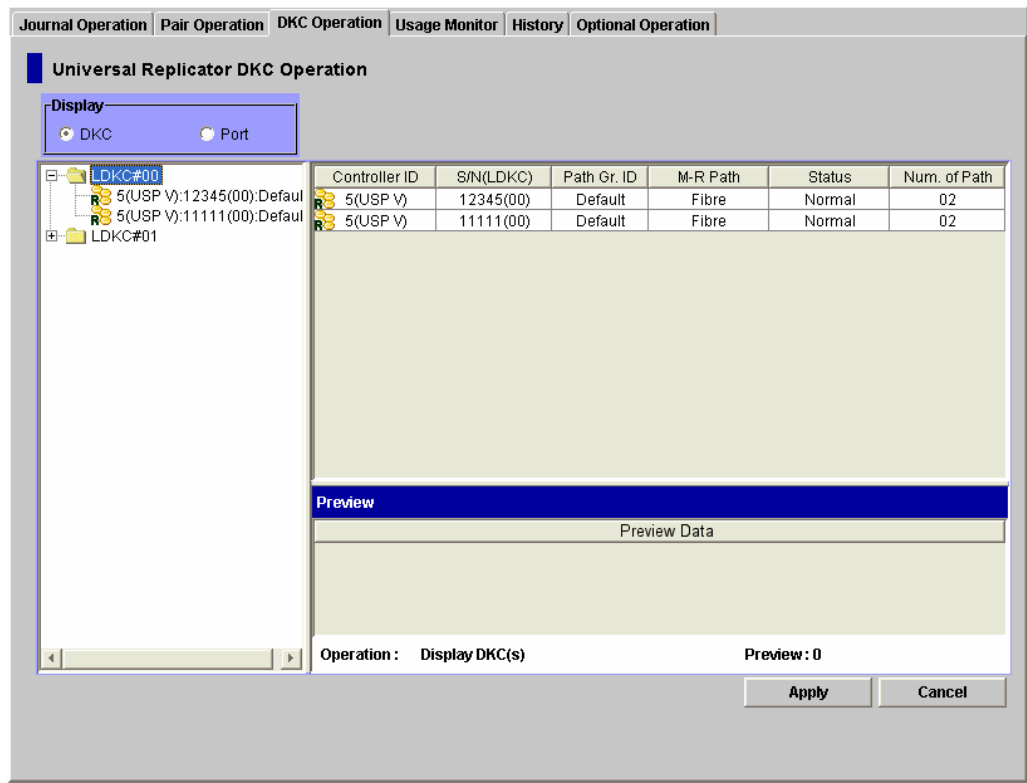


Figure 4-5 DKC Operation Window

The DKC Operation window displays the following:

- **Display:** Changes information in the DKC Operation window.
 - If you select **DKC**, the DKC Operation window displays information about the remote storage systems and the logical paths.
 - If you select **Port**, the DKC Operation window displays information about ports on the local storage system.
- Lists either of the following:
 - remote storage systems (refer to the next section)
 - channel adapters on the local storage system and port attributes (refer to section Displaying Information about Ports on the Local Storage System)
- **List:** Displays one of the following:
 - detailed information about remote storage systems (refer to the next sections)
 - detailed information about logical paths (refer to the next sections)
 - detailed information about ports on the local storage system (refer to the next sections)

- **Preview:** Displays changes that have been made in the window. When you change settings in the window, the changes will be displayed in the **Preview** list before the changes are applied to storage systems. If you are sure that information in the **Preview** is correct, select the **Apply** button to apply the settings that you have made.
 - **Operation:** Indicates the operation in progress in the DKC Operation window.
 - **Apply:** Applies settings in the **Preview** box to the storage systems.
 - **Cancel:** Cancels the settings displayed in the **Preview** box.
- Note:** Information on the DKC Operation window will be updated when you do one of the following:
2. Select another tab and then reselect the DKC Operation tab.
 3. Click File, Refresh on the menu bar of the Storage Navigator main window.
 4. Select the Apply button.
 5. Select modify mode when you are in view mode.
 6. Close the **DKC Status** window.

Displaying Information about Remote Storage Systems

To display information about remote storage systems:

1. Display the DKC Operation window.
2. In the **Display** box, select **DKC**.

The tree and the list display remote storage systems.




Figure 4-6 Remote Storage System Information in the DKC Operation Window

- **Tree:** Lists the remote storage systems of each LDKC in the local storage system. The following information appears to the right of the icon of the remote storage system:
 controller ID of the remote storage system (The model name of the remote storage system)
 - serial number of the remote storage system
 - path group ID

Note: The **LDKC#01** cannot be used in this version.

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



All the logical paths are in normal status.



A failure occurred at some of the logical paths.

List: Displays information about remote storage systems:

- **Controller ID** displays the controller ID and the model name of the storage system of a remote storage system. The icon indicates the status of logical paths between the local storage system and the remote storage system:



All the logical paths are in normal status.



A failure occurred at some of the logical paths.

Note: The controller ID for a USP V storage system is **5**. The model name of USP V is **USP V**.

- **S/N** displays the five-digit serial number and the LDKC number of the remote storage system. "00" is displayed as the LDKC number when the remote storage system is TagmaStore USP/NSC.
- **Path Gr. ID** displays the path group ID. All the path group IDs are displayed as **Default**.
- **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 local storage system and the remote storage system. This column always displays **Fibre**.
- **Status** indicates whether logical paths fail.

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.

Displaying Information about Logical Paths

Logical paths are communication paths that connect the local storage system ports and the remote storage system ports.

To display information about logical paths in the DKC Operation window:

1. In the **Display** box, select **DKC**.
2. Select a remote storage system from the tree.

The list displays logical paths. Each row in the list indicates one logical path (Figure 4-7).




Figure 4-7 Logical Paths Information in the DKC Operation Window

The list displays the following:

- **Path Gr. ID:** Displays the path group ID.

The icon indicates the status of the path:



The logical path is in normal status



A failure occurred at the logical path.

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

- **Status:** indicates whether the logical path is in normal status.

Normal The logical path is in normal status. No failure occurs at the logical path

Failed A failure occurs at the logical path.

- **Port:** Indicates a port number of the local storage system.
- **Pair-Port:** indicates a port number of the remote storage system.

Displaying Information about Ports on the Local Storage System

To display information about ports on the local storage system in the DKC Operation window:

1. In the **Display** box, select **Port**.

The tree displays a list of channel adapters and port attributes (Figure 4-7).

2. Do either of the following in the tree:
 - Select **Subsystem**. The list displays all the ports on the local storage system.
 - Select a channel adapter. The list displays ports on the channel adapter.
 - Select a port attribute. The list displays ports that have the selected port attribute.

Usage Monitor Window

The **Usage Monitor** window (see Figure 4-9) displays the remote copy I/O statistics for all volumes to be monitored on the connected storage system, and, displays the status of remote copy usage monitoring. To display the Usage Monitor window, do either of the following:

- If Universal Replicator has not been started:
 - a. Use your Web browser to display the storage device list. In the storage device list, select the storage system to log in. Enter a user name and the password, and then click **OK**.
The Storage Navigator main window is displayed.
Note: For detailed instructions on this step, please refer to *Storage Navigator User's Guide*.
 - b. Click **Go, Universal Replicator** and then **Usage Monitor** on the menu bar of the Storage Navigator main window.

UR starts and the Usage Monitor window is displayed.

- **If Universal Replicator has already been started:**

3. Select the **Usage Monitor** tab.

The Usage Monitor window is displayed.



Figure 4-9 Usage Monitor Window

- **Monitoring Switch:** displays **Enable** when monitoring is on; displays **Disable** when monitoring is off.
Note: When monitoring is stopped, the usage monitor graph is closed.
- **Gathering Interval:** displays the data collection interval

- **Update:** displays the most recent data sample time of the data on the graph.
- **Graph:** displays the remote I/O statistic information and the status of remote copy monitor.

History Window

The **History** window (see Figure 4-10) display history of pf operations for data volume pairs. For example, the window displays the date and time when data volume pairs are created or released.

To display the History window, do either of the following:

- If Universal Replicator has not been started:
 - a. Use your Web browser to display the storage device list. In the storage device list, select the storage system to log in. Enter a user name and the password, and then click **OK**.
The Storage Navigator main window is displayed.
Note: For detailed instructions on this step, please refer to *USP V Storage Navigator User's Guide*.
 - b. Click **Go, Universal Replicator** and then **History** on the menu bar of the Storage Navigator main window.
UR starts and the History window is displayed.
Caution: Here, the History window may not display latest operation history. To view the latest operation history, go to the next step.
 - c. Click **File, Refresh** on the menu bar of the Storage Navigator main window.
The operation history is updated.
- **If Universal Replicator has already been started:**
 - d. Click the **History** tab to display the History window.
Caution: Here, the History window may not display latest operation history. To view the latest operation history, go to the next step.
 - e. Click **File, Refresh** on the menu bar of the Storage Navigator main window.
The operation history is updated.

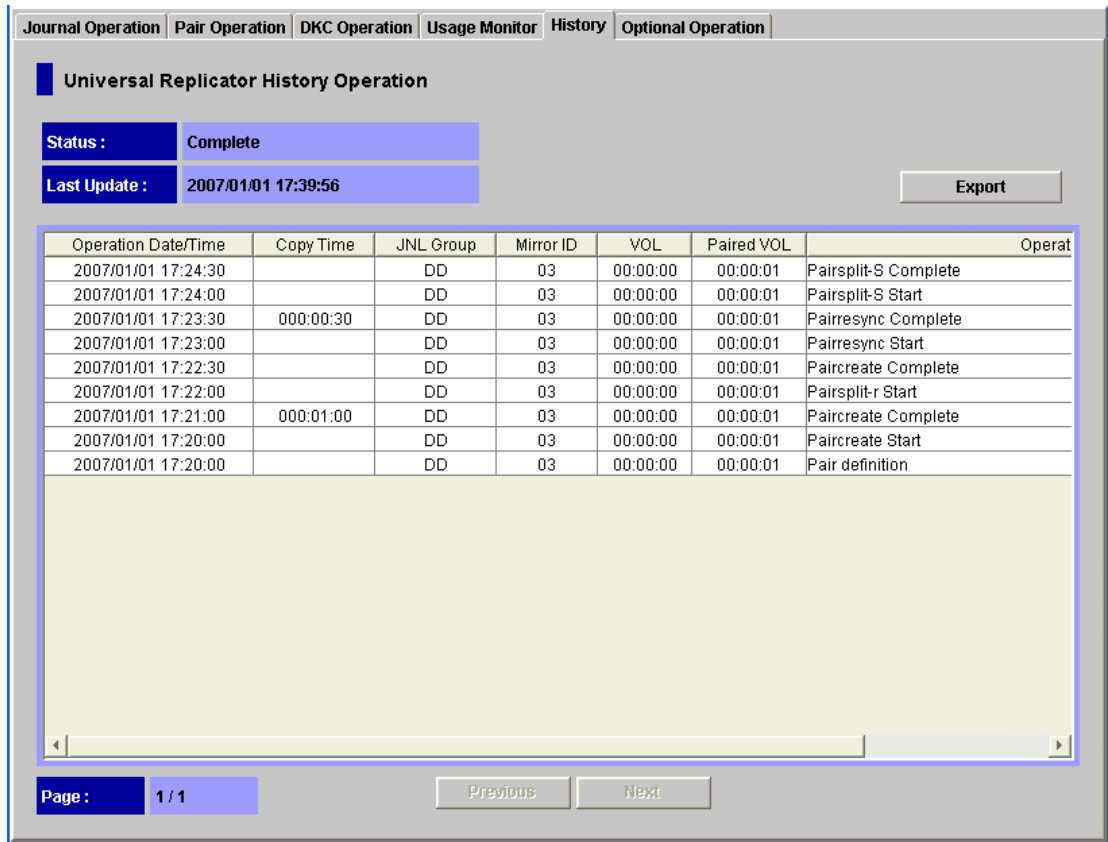


Figure 4-10 History Window

The **History** window presents:

- **Status** displays the current status of operation history:
 - **No history file exists:** Operation history does not exist.
 - **Reading a history file failed:** An attempt to read operation history failed.
 - **Updating ... n (%):** Updating of operation history is now in progress. "n (%)" indicates the progress (in %) of the updating process of operation history.
 - **Complete:** Updating of operation history has been completed.

Note: When the updating process is in progress, the checking process automatically continues until the updating process finishes. The updating process is checked at ten-second intervals.

- **Last Update:** Displays the date and time when operation history was last updated.
- **Export:** Saves operation history in a CSV file. For detailed information, see section Saving Operation History to a Text File.

Note: You cannot save the history file while operation history is being updated. Please save operation history in a text file after operation history is updated.

- **The list** displays history of operations on data volume pairs in the connected storage system.

The list consists of the following columns:

- **Operation Date/Time:** The date and time when the operation was completed.
- **Operation:** The operations are:
 - Pair definition:** A data volume pair was defined.
 - Paircreate Start:** Creation of the data volume pair was started.
 - Paircreate Complete:** Creation of the data volume pair was finished.
 - Pairresync Start:** Restoring of the data volume pair was started.
 - Pairresync Complete:** Restoring of the data volume pair was finished.
 - Pairsplit-r Start:** Splitting (Suspending) of the data volume pair was started.
 - Pairsplit-r Complete:** Splitting (Suspending) of the data volume pair was finished.
 - Pairsplit-r(Failure):** The data volume pair was split (suspended) because of a failure.
 - Pairsplit-S Start:** Release of the data volume pair was started.
 - Pairsplit-S Complete:** Release of the data volume pair was finished.
 - Status Change by MCU(SMPL to COPY):** The status of the data volume pair was changed from SMPL to COPY because of an operation from the primary storage system.
 - Status Change by MCU(SMPL to PAIR):** The status of the data volume pair was changed from SMPL to PAIR because of an operation from the primary storage system.
 - Status Change by MCU(COPY to PAIR):** The status of the data volume pair was changed from COPY to PAIR because of an operation from the primary storage system.
 - Status Change by MCU(COPY to PSUS/PSUE):** The status of the data volume pair was changed from COPY to PSUS or PSUE because of an operation from the primary storage system.
 - Status Change by MCU(PAIR to PSUS/PSUE):** The status of the data volume pair was changed from PAIR to PSUS or PSUE because of an operation from the primary storage system.
 - Status Change by MCU(PAIR to SMPL):** The status of the data volume pair was changed from PAIR to SMPL because of an operation from the primary storage system.
 - Status Change by MCU(COPY to SMPL):** The status of the data volume pair was changed from COPY to SMPL because of an operation from the primary storage system.

Status Change by MCU(PSUS/PSUE to SMPL): The status of the data volume pair was changed from PSUS or PSUE to SMPL because of an operation from the primary storage system.

Status Change by MCU(PSUS/PSUE to COPY): The status of the data volume pair was changed from PSUS or PSUE to COPY because of an operation from the primary storage system.

Status Change by RCU(Pairsplit-r Start): The status of the data volume pair was changed because an operation for splitting a pair started at the secondary storage system.

Status Change by RCU(Pairsplit-r Complete): The status of the data volume pair was changed because an operation for splitting a pair finished at the secondary storage system.

Status Change by RCU(PSUS/PSUE to SMPL; Pairsplit-S Start): An operation for releasing a pair has been started at the secondary storage system. The status of the data volume pair will change from PSUS or PSUE to SMPL.

Status Change by RCU(COPY to SMPL; Pairsplit-S Start): An operation for releasing a pair has been started at the secondary storage system. The status of the data volume pair will change from COPY to SMPL.

Status Change by RCU(PAIR to SMPL; Pairsplit-S Start): An operation for releasing a pair has been started at the secondary storage system. The status of the data volume pair will change from PAIR to SMPL.

Status Change by RCU(Pairsplit-S Complete): The status of the data volume pair was changed because an operation for releasing a pair finished at the secondary storage system.

Ready for Delta resync: A data volume pair became ready for delta resync.

Ready for Delta resync(Failure): The failure occurred with the data volume pair which was ready for delta resync.

Status Change for Delta resync: The status of the primary data volume was changed to HOLD because of a delta resync operation.

Status Change by MCU(SMPL to HOLD): The status of the data volume pair was changed from SMPL to HOLD because of an operation from the primary storage system.

Status Change by MCU(HOLD to PAIR): The status of the data volume pair was changed from HOLD to PAIR because of an operation from the primary storage system.

Status Change by MCU(HOLD to COPY): The status of the data volume pair was changed from HOLD to COPY because of an operation from the primary storage system.

Status Change by MCU(HOLD to SMPL): The status of the data volume pair was changed from HOLD to SMPL because of an operation from the primary storage system.

Status Change by RCU(HOLD to SMPL, Pairsplit-S Start): An operation for releasing a pair has been started at the secondary storage system. The status of the data volume pair will change from HOLD to SMPL.

Status Change to HOLD: The status of the secondary data volume was changed to HOLD because of a delta resync operation.

Unknown: The storage system could not identify the type of the operation.

- **JNL Group:** The journal group number.
- **Mirror ID:** The mirror ID.
- **VOL:** The volume manipulated during the operation. This volume is located in the local storage system. This number indicates LDKC:CU:LDEV (the LDKC number, the CU number, and the LDEV number).
- **Paired VOL:** The volume paired with the manipulated volume. This volume is located in the remote storage system. This number indicates LDKC:CU:LDEV (the LDKC number, the CU number, and the LDEV number).
- **Copy Time:** The time **taken** for the operation (from the start of the operation to the end). Displayed only for **Paircreate Complete** and **Pairresync Complete** operations.

Notes:

- The history information (i.e., rows in the list) might not be displayed in chronological descending order. To sort the information in descending (or ascending) order, click a header of the list.
- The History window can display history information of up to 524,288 operations. However, the window can only display up to 16,384 operations at one time. To display the remaining operations, click the **Previous** or **Next** button.
- The History window does not display history information older than seven days. Also, if the number of operations exceeds 524,288, the oldest operations are deleted in chronological order to keep the number at 524,288, even if the operations occurred within the last week. The history file always contains the most recent operations up to a maximum of 524,288 operations. The history information for an entire week may not always be available.
- If a failure occurs with two or more data volume pairs at the same time, only up to two rows showing "**Pairsplit-r(Failure)**" or "**Ready for Delta resync(Failure)**" will be displayed.
- If a data volume pair consists of LUSE volumes, the list only displays the top LDEV numbers of the LUSE volumes.
- The copy time might not be displayed in the Copy Time column, even though **Paircreate Complete** or **Pairresync Complete** is displayed in the Operation column. In such case, you can confirm the copy time at the volume list in the Pair Operation window.

- **Page:** displays the number of current page and total number of pages. The display format of **Page** is "the number of current page / total number of pages". If there is no history file, "-/-" is displayed.
- **Previous:** The list displays up to a maximum of 16,384 operations at a time. If the number of operations exceeds 16,384, the **Previous** and **Next** buttons allows you to display the remaining operations. The **Previous** button allows you to display the previous 16,384 operations.
- **Next:** The list displays up to a maximum of 16,384 operations at a time. If the number of operations exceeds 16,384, the **Previous** and **Next** buttons allows you to display the remaining operations. The **Next** button allows you to display the next 16,384 operations.

Optional Operation Window

The **Optional Operation** window (see Figure 4-11) allows you to set options.

To display the **Optional Operation** window, do either of the following:

- **If Universal Replicator has not been started:**
 - a. Use your Web browser to display the storage device list. In the storage device list, select the storage system to log in. Enter a user name and the password, and then click **OK**.

The Storage Navigator main window is displayed.

Note: For detailed instructions on this step, please refer to *Storage Navigator User's Guide*.

- b. Click **Go, Universal Replicator** and then **Optional Operation** on the menu bar of the Storage Navigator main window.

UR starts and the Optional Operation window is displayed.

- **If Universal Replicator has already been started:**
 1. Select the Optional Operation tab to display the Optional Operation window.

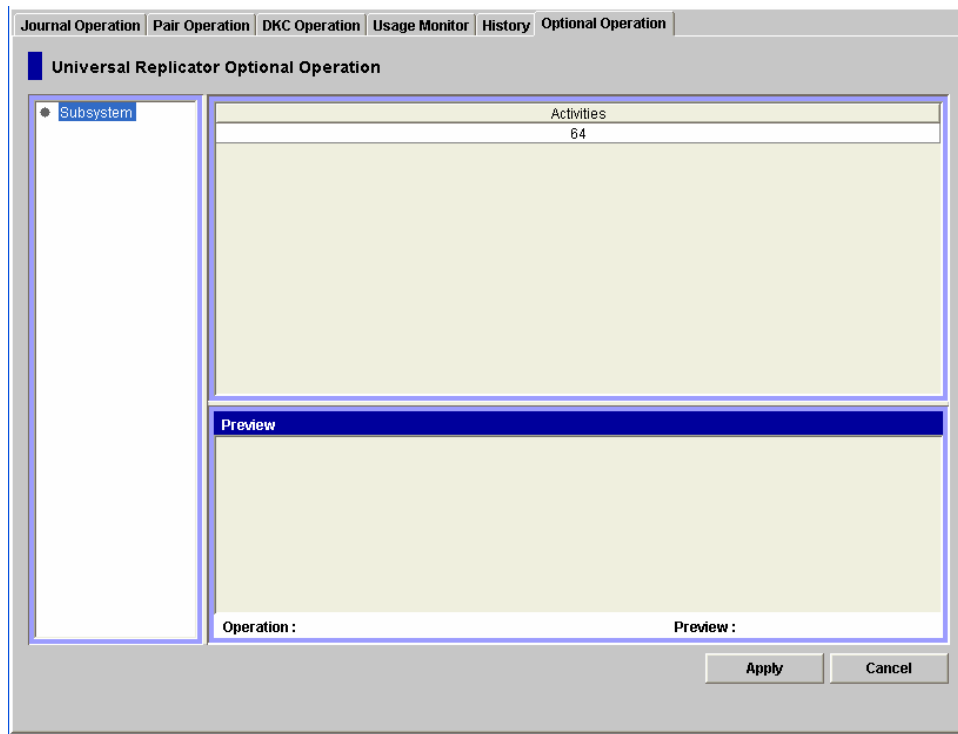


Figure 4-11 Optional Operation Window

- **Tree:** allows you to select **Subsystem**.
- **Activities List:** displays storage system option settings.

- **Preview List:** When you change settings in the Optional Operation window, the **Preview** list displays the changes. Here, the changes have not been applied to the storage system. If you select the **Apply** button, the changes will be applied to the storage system.
- **Operation:** Indicates the operation that will occur when you select **Apply**.

Change System Option	Change storage system options
Blank	Nothing will occur when you click Apply

- **Preview:** The number to the left of the slash (/) indicates the number of items (i.e., rows) displayed in the **Preview** list.
The number to the right of the slash indicates the maximum number of items (i.e., rows) that can be displayed in the **Preview** list.
- The **Apply** button: applies settings in the **Preview** list to the storage system.
- The **Cancel** button: cancels the settings in the **Preview** list.

Note: Information on the Optional Operation window will be updated when you do one of the following:

2. Select another tab and then reselect the **Optional Operation** tab.
3. Click **File, Refresh** on the menu bar of the Storage Navigator main window.
4. Select the **Apply** button.
5. Select modify mode when you are in view mode.

Configuring Storage Systems and Logical Paths

This chapter introduces you must consider when introducing Universal Replicator (henceforth abbreviated as UR) into your system and configuring your system for remote copy operations; for example, deciding which data should be copied from which storage system to which storage system:

- Reviewing System Configuration Tasks
- Configuring Port Attributes
- Configuring Storage System Options
- Establishing Relationship between Primary and Secondary Storage Systems (Add DKC)
- Changing Options for Logical Paths and Storage Systems
- Adding Logical Paths
- Viewing the Status of Logical Paths
- Deleting Logical Paths Managing Power for Storage Systems and Network Relay Devices
- Removing Relationship Between Primary and Secondary Storage Systems

Reviewing System Configuration Tasks

The storage system containing data to be copied to another storage system is referred to as the *primary storage system*. The storage system to which data are copied from another storage system is referred to as the *secondary storage system*. The place where the primary storage system exists is referred to as the *primary site*. The place where the secondary storage system exists is referred to as the *secondary site*.

When performing system configurations for remote copy operations, the primary site administrator and the secondary site administrator must do the following:

- The primary site administrator must configure some of the ports on the primary storage system to connect hosts. The primary site administrator must also configure some other ports on the primary storage system to communicate with remote storage systems.
- The secondary site administrator must configure some ports on the secondary storage system to communicate with the primary storage system.
- After finishing port configurations, the primary and secondary site administrators must configure logical paths between the communication ports on the primary site and the communication ports on the secondary site, so that the two storage systems can communicate with each other.

When performing system configurations, the administrators must do the following:

- Display the DKC Operation window
- Configure port attributes
- Associate the primary and the secondary storage systems and configure logical paths between the two storage systems
- When performing system configurations, the administrator should do the following when necessary:
 - Change options for logical paths and storage systems
 - Add logical paths
 - Delete logical paths
 - Displaying status of logical paths
 - Delete the relationship between the primary and the secondary storage systems

Note: Throughout this chapter, the primary and the secondary storage systems are sometimes referred to as *local storage systems* or *remote storage systems*. If you are logged into the primary storage system and are using UR, the primary storage system is a *local* storage system and the secondary storage system is a *remote* storage system. If you are logged into the secondary storage system and are using UR, the secondary storage system is a *local* storage system and the primary storage system is a *remote* storage system.

Note: When performing system configuration, the administrator also needs to configure journal groups. For detailed information about configuring journal groups, see Using Journal Groups.

Configuring Port Attributes

The administrator at both the primary site and the secondary site must configure ports on the primary storage system and the secondary storage system in the following ways:

- The administrator at the primary site must configure target ports, initiator ports, and RCU target ports on the primary storage system.
- The administrator at the secondary site must configure initiator ports and RCU target ports on the secondary storage system.

Ports on USP V storage system have the *target* attribute by default. If you are configuring USP V storage systems, you must determine ports that should be used as initiator ports or RCU target ports, and then you must change the attribute of the ports from *target* to *initiator* or *RCU target*. The port attributes are explained below:

- **Target:** A target port is a Fiber Channel port used for connecting a storage system and a host. When the host issues a write request, the request will be sent to a volume via a target port on the storage system. Target ports must be configured on primary storage systems for UR operations.
- **Initiator:** An initiator port is a Fiber Channel port that sends commands to a remote storage system. Initiator ports must be configured on both primary and remote storage systems for UR operations.
- **RCU target:** An initiator port is a Fiber Channel port that receives commands from a remote storage system. RCU target ports must be configured on both primary and remote storage systems for UR operations.
- **External:** An external port is a port that is required for Universal Volume Manager copy operations. This port will not be used for Universal Replicator copy operations.

If necessary, you can change this port to an target port, an initiator port, or an RCU target port.

- **Initiator/external:** This setting shows the combination of ports that have an initiator port and an external port attributes. Only port that has an initiator port attribute within these ports can be used for Universal Replicator. This attribute cannot be changed by the Universal Replicator software. For details on the port to which an initiator/external mix mode is set, please refer to the *Universal Volume Manager User's Guide*.

Note: If initiator ports and RCU target ports are *not* configured on both primary and remote storage systems, the following problems will occur and remote copy operations will not be performed:

- The administrator at the primary site will be unable to create any pair of a primary data volume and a secondary data volume.
- The secondary storage system will be unable to read data from journal volumes in the primary storage system.

WARNING: Before changing a Fiber Channel port to an initiator port, disconnect the port from the host, release all affected data volume pairs, delete all logical paths from the initiator port to the remote storage system, and then remove all channel paths to the port.

WARNING: Before changing a Fiber Channel port from initiator to target or RCU target, release all affected data volume pairs, delete all logical paths from the initiator port to the remote storage system, and then disconnect the local storage system and the remote storage system.

Caution: Limit the number of hosts connected to a target port to 128 or fewer to avoid mistaken disconnection. If more than 128 hosts are connected to a target port, some hosts may be disconnected after changing the type from target to RCU target.

Note: The port assigned to SLPR other than SLPR0 can be set to the Target port attribute only.

To configure attributes of Fiber Channel ports, follow the procedure below. Both the primary storage system administrator and the secondary storage system administrator must perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.

For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.

2. Ensure that the DKC Operation window is displayed.
3. In **Display**, select **Port**.

The tree displays channel adapters in the local storage system and port attributes.

4. Do either of the following:

Select a channel adapter from the tree.

Select a port attribute (i.e., target, RCU target or initiator) from the tree.

5. Select and right-click the port that you want to configure.
6. From the pop-up menu, select the desired port type (i.e., initiator, RCU target, or target).

The rightmost column of the list displays "Modified" to indicate that you are modifying the attribute of the port.

The **Preview** list displays the changes that you have made (Note that these changes are not applied to the storage system yet).

7. See the **Preview** 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 upper-right 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 **Preview** list and then select **Cancel**.

WARNING: Read and follow the important warnings and caution above before applying the port attribute changes.

8. Select **Apply** to apply port attribute changes to the storage system.

Configuring Storage System Options

The Optional Operations window allows you to change storage system option settings

To change the option settings of the storage system:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Optional Operation window is displayed (see section Optional Operation Window).
3. Select Subsystem in the tree.
The list displays storage system options (see Figure 5-1).
4. Do either of the following:
 - Right-click Subsystem in the tree, and then select Change System Option from the pop-up menu.
 - Right-click information in the list to display the pop-up menu, and select **Change System Option** from the pop-up menu.
5. In the System Option window (Figure 5-2), change storage system options.
6. Select the Set button.
7. See the **Preview** list to verify the changes that you have made.
 - If you want to modify a change, right-click the change and then select Modify. A window appears and allows you to modify the change.
 - If you want to cancel a change, right-click the change and then select **Cancel**.
8. Select **Apply** to apply the changes.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.




Figure 5-1 Storage System Option Settings in the Optional Operation Window

- **Activities:** The number of volumes that can be copied concurrently during an initial copy operation.




Figure 5-2 System Option Window

- **Maximum Initial Copy Activities:** Specifies the number volumes that can be copied concurrently during an initial copy operation. A value within the range of 1 to 128 can be specified. The default setting is 64 volumes.

Maximum Initial Copy Activities can impact the performance of the primary storage system, depending on the amount of I/O activity and the number of pairs being registered at the same time. If **Maximum Initial Copy Activities** is 64 volumes, it allows you to limit the impact of initial copy activities on storage system performance. For example, if you set the maximum initial copy activities to 64 volumes and then add 65 pairs at the same time, the primary storage system starts the first 64 pairs and will not start the 65th pair until one of the pairs is synchronized. When you change the maximum initial copy activities setting, the new setting applies to pairs created after the setting was changed, not to existing pairs.

- **Set:** Closes the System Option window and then adds the settings to the **Preview** list.
- **Cancel:** Cancels the settings.

Establishing Relationship between Primary and Secondary Storage Systems (Add DKC)

The Add DKC window (see Figure 5.3) allows you to associate the primary storage system and the secondary storage system, and also allows you to establish logical paths between the two storage systems. *Logical paths* are used to connect the primary and secondary storage systems. Both the primary storage system and the secondary storage system use logical paths to communicate with each other.

One primary storage system can be associated with up to 64 secondary storage systems. Up to eight logical paths can be configured between one primary storage system and one secondary storage system. Therefore, one primary storage system can have up to 512 logical paths to secondary storage systems.

To associate the primary and secondary storage systems and then configure logical paths between the two storage systems, both the primary storage system administrator and the secondary storage system administrator must perform this operation:

1. Make sure the remote copy connections and ports are properly configured. Also, get the serial number of the remote storage system.

Note: To get the serial number of the remote storage system, ask the administrator of the remote storage system. The serial number will be displayed in the Storage Navigator main window when the Storage Navigator user logs into the storage system.

2. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
3. Ensure that the DKC Operation window is displayed (see section DKC Operation Window).
4. In Display, select DKC.
5. Select LDKC from the tree.
6. Right-click the upper-right list, and then select **DKC Operation** and **Add DKC**

The Add DKC window appears (refer to Figure 5.3).

7. In the Add DKC window, enter the following:
 - the serial number of the remote storage system
 - the LDKC number of the remote storage system
"00" is displayed when the remote storage system is TagmaStore USP/NSC.
 - the controller ID of the remote storage system (model name that indicates the model)
Note: The controller ID for a USP V storage system is **5(USP V)**.
 - the path group ID

The path group ID is within the range of 1-FF (hexadecimal). Up to 64 path group IDs can be registered per storage system.

Note: In the current version, the path group IDs always take the default values and cannot be changed.

8. In the Add DKC window, select a primary storage system port and a remote storage system port to configure a logical path.
9. Select Option.
10. In the DKC Option window, configure logical path options and storage system options.
For detailed information about the options, see the next section and read explanations below Figure 5-4.
11. Select **Set** to close the DKC Option window.
12. See the **Preview** list to check the settings that you have made.
 - If you want to modify a setting, select and right-click the setting from the **Preview** list and then select **Modify**. A window appears and allows you to modify the setting.
 - If you want to cancel a setting, select and right-click the setting in the **Preview** list and then select **Cancel**.
13. Select **Apply** to apply the settings that you have made.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Add DKC

S/N:

LDKC:

Controller ID:

Path Gr. ID: Default

M-R Path	
Port	Pair-Port
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>

Figure 5.3 Add DKC Window

The Add DKC window displays the following:

- **S/N:** allows you to enter the five-digit serial number of the remote storage system.
- **LDKC:** allows you to enter the LDKC number of the remote storage system. Enter "00" when the remote storage system is TagmaStore USP/NSC.
Note: The **LDKC#01** cannot be used in this version.
- **Controller ID:** allows you to enter the controller ID (model name that indicates the model) of the remote storage system.
Note: The controller ID for a USP V storage system is **5(USP V)**.
- **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.

Path group IDs are within the range of 0-FF (hexadecimal). If you select the **Default** check box, the default path group ID will be set.

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 local storage system to RCU target ports on the remote storage system.
- **Port:** displays a list of initiator ports on the local storage system. Select an initiator port from this drop-down list.
- **Pair-Port:** displays a list of all ports on the remote storage system. Select an RCU target port on the remote storage system from this drop-down list.
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 window.
- **Cancel:** cancels the settings you made on the Add DKC window and then closes the window.

Changing Options for Logical Paths and Storage Systems

The DKC Option window (see Figure 5-4) allows you to set the logical path options and storage system options. These options will apply to the local storage system and the corresponding remote storage system.

Note: DKC is an acronym for disk controller, which controls an entire storage system. UR windows use the word "DKC" to indicate a storage system.

To change logical path options and storage system options, follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the DKC Operation window is displayed (see section DKC Operation Window).
3. In Display, select DKC.
4. Select LDKC from the tree.
The list displays a list of remote storage systems.
5. From the list, select and right-click the desired remote storage system,
6. From the pop-up menu, select **DKC Operation** and **Change DKC Option**.
The DKC Option window appears (Figure 5-4).
7. On the DKC Option window, change the options as desired. For detailed information about the options, refer to explanations below Figure 5-4.
8. Select Set to close the DKC Option window.
9. See the Preview list to check the settings that you have made.
 - If you want to modify a setting, select and right-click the setting from the Preview list and then select Modify. A window appears and allows you to modify the setting.
 - If you want to cancel a setting, select and right-click the setting in the **Preview** list and then select **Cancel**.
10. Select **Apply** to apply the settings that you have made.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

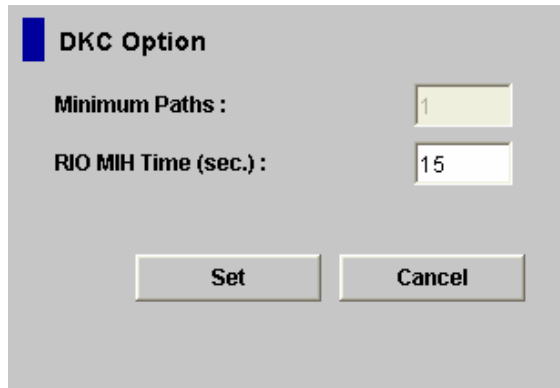


Figure 5-4 DKC Option Window

The DKC Option window displays the following:

- The **Minimum Paths** option specifies the minimum number of paths required for each remote storage system connected to the local storage system (default value is 1).
Note: You cannot change the **Minimum Paths** for the current version.
- The **RIO MIH Time** setting specifies the RIO MIH timer value, which is the wait time until data transfer from the local 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.

Adding Logical Paths

The Add Path window (see Figure 5-5) allows you to add one or more logical paths between the local and remote storage systems.

One primary storage system can be associated with up to 64 secondary storage systems. Up to eight logical paths can be configured between one primary storage system and one secondary storage system. When you add a logical path, UR will automatically start using the new logical path to perform UR copy activities.

To add one or more new logical paths between the local and remote storage systems, follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Make sure the remote copy connections are properly configured.
2. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
3. Ensure that the DKC Operation window is displayed (see section DKC Operation Window).
4. In Display, select DKC.
5. Do either of the following:
 - In the tree, select a remote storage system.
 - In the list, select and right-click a remote storage system and then select **Edit Path(s)** from the pop-up menu. The list displays information about logical paths.
6. Right-click the list and then select **Add Path** from the pop-up menu.
The Add Path window appears (Figure 5-5).
7. Use the Add Path window to configure new logical paths.
Up to eight paths can be configured.
8. Select Set to close the Add Path window.
9. See the **Preview** list to check the settings that you have made.
 - If you want to modify a setting, select and right-click the setting from the **Preview** list and then select **Modify**. A window appears and allows you to modify the setting.
 - If you want to cancel a setting, select and right-click the setting in the **Preview** list and then select **Cancel**.
10. Select **Apply** to apply the settings that you have made.
Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Figure 5-5 Add Path Window

The Add Path window displays the following:

- **Port:** allows you to select an initiator port on the local storage system.
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.
- **Pair-Port:** allows you to select an RCU target port on the remote storage system.
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.
- **Set:** closes the Add Path window and adds the settings to the **Preview** list.
- **Cancel:** cancels the settings.

Viewing the Status of Logical Paths

The DKC Status window (see Figure 5-6) displays the status of logical path(s).

To view the status of logical path(s), follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the DKC Operation window is displayed (see section DKC Operation Window).
2. In Display, select DKC.
3. Execute a following operation between two ways.
 - Click **LDKC** from the tree. In the upper-right list, click and right-click a remote storage system.
 - Click the remote storage system from the tree. In the upper-right list, click and right-click a logical path.
4. Select **DKC Status** from the pop-up menu.

The DKC Status window appears (Figure 5-6).

- If a remote storage system is selected in the list, the DKC status window displays all the logical paths between the local and the remote storage systems and shows the status for each of the paths.
 - If a logical path is selected in the list, the DKC status window only displays the status of the selected logical path.
5. Check the status of the logical path(s), and then select **Close** to close the DKC status window.

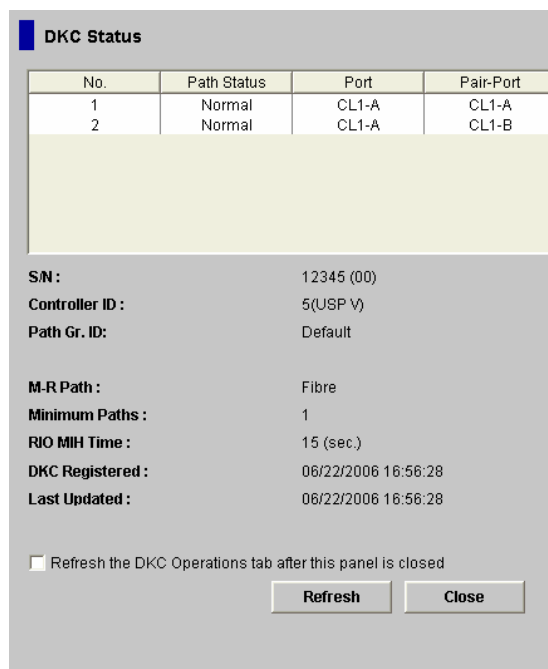


Figure 5-6 DKC Status Window

The DKC Status window displays the following:

- **List:**
 - **No.:** displays serial numbers used for rows in the list
 - **Path Status:** indicates status of a logical path (For details, see Table 5-1)
 - **Port:** indicates a port on the local storage system
 - **Pair-Port:** indicates a port on the remote storage system.
 - **S/N:** indicates the serial number and LDKC number of the remote storage system.
- **Controller ID:** displays the controller ID and the model name of the remote storage system.
- **Path Gr. ID:** indicates a path group ID.
- **M-R Path:** indicates the type of channel interface between the local and the remote storage systems. This column displays Fiber.
- **Minimum Paths:** Indicates the minimum possible number of paths between the local and the remote storage systems.
- **RIO MIH Time:** indicates the RIO MIH timer value, which is the wait time until data transfer from the local storage system to the remote storage system is complete.

Note: RIO MIH is an acronym for remote I/O missing interrupt handler.
- **DKC Registered:** Indicates the date and time when the local 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 Operation tab after this window is closed:** If you select this check box, information in the DKC Operation window will be refreshed after you close the DKC Status window.
- **Refresh:** refreshes the information in the DKC Status window.
- **Close:** closes the DKC Status window.

Table 5-1 Logical Path Status

Status	Remarks
Normal	This path has been successfully established and can be used for UR 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 local and the remote storage system. The probable causes are: <ul style="list-style-type: none"> • No cable is connected to the local storage system. • No cable is connected to the remote storage system. • No cable is connected to the network device that comes between the local and the remote storage system.

Communication Time Out	This status indicates one of the following: <ul style="list-style-type: none"> • A timeout error has occurred between the primary and the secondary storage systems. • A logic error is detected between the primary and the secondary storage systems.
Resource Shortage (Port)	The local storage system rejected the function for configuring logical path connection. All logical path resources in the local 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 window.
Invalid Port	The port is not an initiator port.
Pair-Port Number Mismatch	This status indicates one of the following: <ul style="list-style-type: none"> • The specified port number is incorrect. • The port in the remote storage system is physically disconnected from the local storage system.
Pair-Port Type Mismatch	The port on the remote storage system is not an RCU target port.
Communication Failed	A communication timeout error has occurred on the path between the primary and the secondary storage systems.

Deleting Logical Paths

Before deleting logical path(s), make sure that the remaining number of logical paths will be greater than the minimum number of paths setting (see section Changing Options for Logical Paths and Storage Systems). The delete path operation will fail if the number of remaining paths is equal to or less than the minimum number of paths.

To delete logical path(s), follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the DKC Operation window is displayed (see section DKC Operation Window).
3. In **Display**, select **DKC**.
Do either of the following:
 - In the tree, select a remote storage system.
 - In the list, select and right-click a remote storage system and then select **Edit Path(s)** from the pop-up menu. The list displays information about logical paths.
4. In the list, select the logical path(s) that you want to delete.
5. Right-click the list and then select **Delete Path** from the pop-up menu.
A confirmation message appears.
6. Select OK to close the message.
7. In the DKC Operation window, locate the upper-right list and the **Preview** list and then check the settings that you have made.
 - The right-most column of the upper-right list displays the word "Delete" to indicate logical path(s) to be deleted.
 - The **Preview** list displays logical path(s) to be deleted.
 - If you want to cancel deleting logical path(s), select and right-click the logical path(s) in the **Preview** list and then select **Cancel**.
8. Select **Apply** to delete logical path(s).
Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Managing Power for Storage Systems and Network Relay Devices

This appendix explains power management for storage systems and network relay devices during remote copy operations. In particular, this appendix discusses the following:

- What happens when power is removed from storage systems or network relay devices due to some unexpected reason (see the next section)
- What should be noted when you power off storage systems or network relay devices intentionally (see section Turning Power Off Intentionally)

Here, the term "network relay devices" refers to hardware used for connecting the primary and secondary storage systems, such as ESCON directors, channel extenders, and switches.

When the Power Stops Unexpectedly

This section explains what happens when power is removed from storage systems or network relay devices due to some unexpected reason.

When the Power is Removed from the Primary Storage System

If power is removed from the primary storage system during remote copy operations, the primary storage system assumes that a failure occurs, and then splits all data volume pairs. If the primary storage system splits data volume pairs, the secondary storage system also assumes that a failure occurs, and then splits all data volume pairs.

If power is removed from the primary or secondary storage system and its backup batteries are fully discharged while data volume pairs are split, differential data (i.e., update data) will not be retained. In this unlikely case, the primary or secondary storage system assumes all the suspended data volumes are updated. If you restore the data volume pairs at the primary site, all the primary data volumes will be copied to the secondary data volumes.

When the Power is Removed from the Secondary Storage System

If power is removed from the secondary storage system during remote copy operations, the secondary storage system assumes that a failure occurs, and then splits all data volume pairs. If the secondary storage system splits data volume pairs, the primary storage system also assumes that a failure occurs, and then splits all data volume pairs.

If power is removed from the primary or secondary storage system and its backup batteries are fully discharged while data volume pairs are split, differential data (i.e., update data) will not be retained. In this unlikely case, the primary or secondary storage system assumes all the suspended data volumes are updated. If you restore the data volume pairs at the primary site, all the primary data volumes will be copied to the secondary data volumes.

When the Power is Removed from Network Relay Devices

If power is removed from a network relay device during remote copy operations, the primary and secondary storage systems assume that a failure occurs, and then split all data volume pairs.

Turning Power Off Intentionally

This section explains what should be noted when you power off storage systems or network relay devices intentionally.

Power Off the Primary Storage System

If you want to power off the primary storage system during remote copy operations, you must ensure, before powering off the storage system, that all data volume pairs or mirrors (i.e., pairs of journal groups) are split, and that the status of all the data volume pairs is **PSUS**. **Caution:** If you want to power off the primary storage system, please contact your Hitachi account team and ask them to power off the storage system.

If data volume pairs are split and the primary storage system is powered off as mentioned above, you can resume remote copy operations by taking the following steps:

1. Power on the primary storage system.
2. If the primary storage system is ready to resume remote copy operations, please restore the data volume pairs that have been split, and then confirm that the status of the data volume pairs is **COPY** or **PAIR**. This operation must be performed at the primary site.

After you take these steps, remote copy operations will be resumed. The primary storage system will be able to respond to *read journal commands* issued by the secondary storage system.

Note: If you want to power off both the primary and secondary storage systems at the same time, please see section Power Off the Primary and Secondary Storage Systems at Same Time.

Note: If you turn the power off without changing the status of all data volume pairs to **PSUS**, and then you turn the power back on, the status of all data volume pairs could be suspended according to a failure.

Power Off the Secondary Storage System

If you want to power off the secondary storage system during remote copy operations, you must ensure, before powering off the storage system, that all data volume pairs or mirrors (i.e., pairs of journal groups) are split, and that the status of all the data volume pairs is **PSUS**. **Caution:** If you want to power off the secondary storage system, please contact your Hitachi account team and ask them to power off the storage system.

If data volume pairs are split and the secondary storage system is powered off as mentioned above, you can resume remote copy operations by taking the following steps:

1. Power on the secondary storage system.
2. If the secondary storage system is ready to resume remote copy operations, please restore the data volume pairs that have been split, and then confirm that the status of the data volume pairs is **COPY** or **PAIR**. This operation must be performed at the primary site.

After you take these steps, remote copy operations will be resumed. The secondary storage system will be able to issue *read journal commands* to the primary storage system.

Note: If you want to power off both the primary and secondary storage systems at the same time, please see the next section.

Note: If you turn the power off without changing the status of all data volume pairs to **PSUS**, and then you turn the power back on, the status of all data volume pairs could be suspended according to a failure.

Power Off the Primary and Secondary Storage Systems at Same Time

If you want to power off the primary and secondary storage systems at the same time, the primary storage system must be powered off before the secondary storage system is powered off. If you want to power on the primary and secondary storage systems again, the secondary storage system must be powered on before the primary storage system is powered on.

The following is the procedure for powering off the primary and secondary storage systems at the same time:

1. Split data volume pairs that will be affected by powering off of the primary and secondary storage systems.

For example, if two primary storage systems and one secondary storage systems are connected to each other and you want to power off one of the primary storage system and the secondary storage system, you must split data volume pairs that are defined between the remaining one primary storage system and the primary storage system, because these data volume pairs will be affected by powering off of the storage systems.

2. After the splitting of the pairs finishes, confirm that the status of all the pairs is changed to PSUS at the primary storage system.

3. Power off the primary storage system. Do not power on the primary storage system yet.
4. Power off the secondary storage system.
5. Power on the secondary storage system.
6. If the secondary storage system is ready to resume remote copy operations, power on the primary storage system.
7. If the primary storage system is ready to resume remote copy operations, restore the data volume pairs that have been split at the primary storage system. Confirm that the status of the data volume pairs is changed to **COPY** or **PAIR**.

Note: After you turned the power on or off both the primary and secondary storage system at the same time, if a status of a data volume pair of primary storage subsystem is PSUS and a status of a data volume pair of secondary storage subsystem is PAIR, you must suspend the data volume pair of secondary storage system by using Storage Navigator. After confirming that the statuses of the data volume pair of primary storage system and secondary storage system are PSUS, restore the data volume pair that has been split at the primary storage system.

Power Off Network Relay Devices

If you want to power off a network relay device during remote copy operations, you must ensure, before powering off the device, that all data volume pairs or mirrors (i.e., pairs of journal groups) are split, and that the status of all the data volume pairs is **PSUS**.

If data volume pairs are split and the network relay device is powered off as mentioned above, you can resume remote copy operations by taking the following steps:

1. Power on the network relay device.
2. If the network relay device is ready for remote copy operations, please restore the data volume pairs that have been split, and then confirm that the status of the data volume pairs is **COPY** or **PAIR**. This operation must be performed at the primary site.

Removing Relationship Between Primary and Secondary Storage Systems

To remove the relationship between the primary and the secondary storage systems and then remove all the logical paths between the two storage systems, both the primary storage system administrator and the secondary storage system administrator must perform this operation:

1. Make sure that all affected UR volume pairs have been released.
2. Ensure that the Storage Navigator main window is in Modify mode.

3. For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
4. Ensure that the DKC Operation window is displayed (see section DKC Operation Window).
5. In **Display**, select **DKC**.
6. Select **LDKC** from the tree.
7. The list displays a list of remote storage systems.
8. Select and right-click remote storage system(s) from the list.
9. Select **DKC Operation** and **Delete DKC** from the pop-up menu.
10. A confirmation message appears.
11. Select **OK** to close the message.
12. See the **Preview** list to check the settings that you have made.
 - The **Preview** list displays the specified remote storage system(s). One row indicates one remote storage system.
 - If you want to cancel removal of the relationship, select and right-click remote storage system in the **Preview** list and then select **Cancel**.
13. Select **Apply** to remove the relationship.

Using Journal Groups

This chapter introduces how to set up UR in your system and configure your system for remote copy operations. The primary site administrator and the secondary site administrator must configure journal groups, as well as ports and logical paths.

- Reviewing Administrator Tasks for Managing Journals
- Registering Journal Volumes in a Journal Group
- Deleting Journal Volumes from a Journal Group
- Displaying Detailed Information about a Journal Group
- Changing Options for Logical Paths and Storage Systems
- Deleting a Journal Group
- Splitting a Mirror (Pairsplit-r: Suspending a copy operation)
- Restoring a Mirror (Pairresync: Resuming a copy operation)
- Deleting Data Volumes from a Mirror (Pairsplit-S: Ending a copy operation)

Reviewing Administrator Tasks for Managing Journals

Administrators must register journal volumes in journal groups. During UR remote copy operations, data in primary data volumes are copied via journal volumes to secondary data volumes.

When configuring journal groups, the administrators must do the following:

- Display the Journal Operation window
- Register journal volumes in journal groups

When configuring journal groups, the administrator should do the following when necessary:

- Delete journal volumes from journal groups
- View detailed information about journal groups
- Change journal group options
- Delete journal groups

When UR is used for remote copy operations, the administrator should do the following when necessary:

- Split mirrors
- Restore mirrors
- Delete data volumes (i.e., the primary and the secondary data volumes) from mirror

Note: Throughout this chapter, the primary and the secondary storage systems are sometimes referred to as *local storage systems* or *remote storage systems*. If you are logged into the primary storage system and are using UR, the primary storage system is a *local* storage system and the secondary storage system is a *remote* storage system. If you are logged into the secondary storage system and are using UR, the secondary storage system is a *local* storage system and the primary storage system is a *remote* storage system.

Registering Journal Volumes in a Journal Group

To be able to perform remote copy operations with UR, the primary site administrator must register journal volumes in journal groups in the primary storage system. Also, the secondary site administrator must register journal volume in journal groups in the secondary storage system. One journal group can contain up to 64 journal volumes.

To register journal volumes, follow the procedure below. Both the primary storage system administrator and the secondary storage system administrator must perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Journal Operation window is displayed (see section Journal Operation Window).
3. In the tree of the Journal Operation window, do either of the following:
 - If you want to register new volumes to a journal group in which journal groups are already registered, select the journal group from below **Registered**.
 - If you want to register volumes to a journal group in which journal groups have not been registered, select the journal group from below **Free**.
4. Do either of the following:
 - In the tree, right-click the selected journal group and then select **Edit JNL Volumes** from the pop-up menu.
 - In the upper-right list, right-click the desired journal group and then select **Edit JNL Volumes** from the pop-up menu.
5. In the **Free Volumes** list of the Edit JNL Volumes window (Figure 6-1), select the volumes that you want to register.

In the **Free Volumes** list, one row represents one volume. If you cannot find the volumes that you want to register, do any of the following:

- Select the **PG** radio button, enter a parity group number in the text boxes to the right, and then select **Show**. The list displays volumes in the specified parity group. Finally, select the volumes that you want to register.

- To register external volumes, select the **PG(Ext.)** radio button, enter a parity group number in the text boxes to the right, and then select **Show**. The list displays volumes in the specified parity group. Finally, select the volumes that you want to register. **Note:** Parity group numbers for external volumes start with the letter "E", but you do not need to enter "E" in the text boxes. For example, if you want to specify the parity group number "E1-2", you only need to enter 1 and 2 into the text boxes.
- Select the **CU** radio button, select a CU number from the drop-down list to the right. The list displays volumes in the specified CU. Finally, select the volumes that you want to register.

Note: In the current version, the **Free Volume** list only displays OPEN-V volumes.

6. Select **Add**.

The **JNL Volumes** list displays the volumes that you want to register. The **Operation** column of the list displays **Add**.

7. Select Set to close the Edit JNL Volumes window.

8. See the **Preview** list to check the settings that you have made.

- If you want to add volumes to register, select and right-click a volume in the **Preview** list and then select **Modify**. A window appears and allows you to add volumes.
- If you want to cancel registering a volume, select and right-click the volume and then select **Cancel**.
- If necessary, you can repeat steps 3 to 8 to add volumes to other journal groups.

9. Select **Apply** to register journal volumes.

After volumes are registered, the journal groups will be displayed below **Registered** in the tree.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

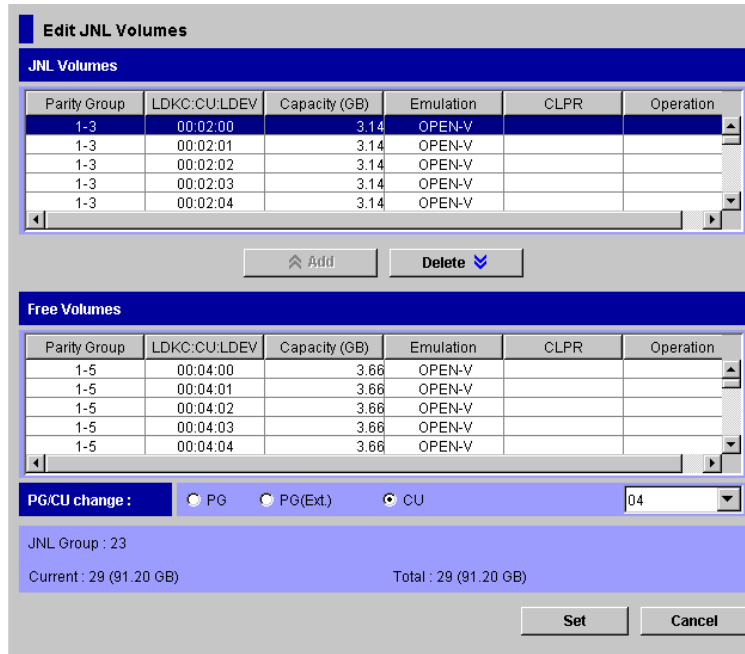


Figure 6-1 Edit JNL Volumes Window

The Edit JNL Volumes window displays the following:

- **JNL Volumes:** Displays information about journal volumes.
 - **Parity Group:** indicates the parity group where a journal volume belongs.

Note: If the letter "E" is displayed at the beginning of a parity group name, the journal volume is an external volume.
 - **LDKC:CU:LDEV:** Indicates the LDKC number, the CU number and the LDEV number of a journal volume.

Note: If a sharp symbol (#) is displayed at the end of a volume, the volume is an external 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.
 - **Operation:** Displays one of the following:
 - Blank This column usually displays a blank.
- **Add** Indicates a volume to be added to a journal group
- **Delete** Indicates a volume to be deleted from a journal group
- **Add:** Use this button when you register volumes in a journal groups. When registering volumes, you select the volumes from **Free Volumes** and then select **Add** to add the volumes to **JNL Volumes**.

- **Delete:** Use this button when you delete volumes from a journal groups. When deleting volumes, you select the volumes from **JNL Volumes** and then select **Delete**.
- **Free Volumes:** Displays information about *free volumes*, which are *not* registered in journal groups.

Note: LUN Expansion (LUSE) volumes are not displayed.

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

Note: If the letter "E" is displayed at the beginning of a parity group name, the volume is an external volume.
- **LDKC:CU:LDEV:** Indicates the LDKC number, the CU number and the LDEV number of a volume.

Note: If a sharp symbol (#) is displayed at the end of a volume, the volume is an external volume.
- **Capacity:** Indicates the capacity of a volume in gigabytes.
- **Emulation:** Indicates the emulation type of a volume.
- **CLPR:** Indicates the number and the name of the CLPR where the volume belongs.
- **Operation:** Displays one of the following:

Blank This column usually displays a blank.

Add Indicates a volume to be added to a journal group

Delete Indicates a journal volume to be deleted from a journal group

- **PG/CU change:** The following radio buttons enables you to switch information in the **Free Volumes** list.
 - **PG:** Use this radio button if you want the **Free Volumes** list to display volumes belonging to a parity group.

If you select this radio button, specify a parity group number in the text boxes to the right, and then select the **Show** button, **Free Volumes** will display volumes in the specified parity group (see Figure 6-2).
 - **PG(Ext.):** Use this radio button if you want the **Free Volumes** list to display external volumes belonging to a parity group.

If you select this radio button, specify a parity group number in the text boxes to the right, and then select the **Show** button, **Free Volumes** will display volumes in the specified parity group.
 - **CU:** Use this radio button if you want the **Free Volumes** list to display volumes belonging to a CU.

If you select this radio button and then select a CU from the drop-down list to the right, **Free Volumes** will display volumes in the selected CU (see Figure 6-3).

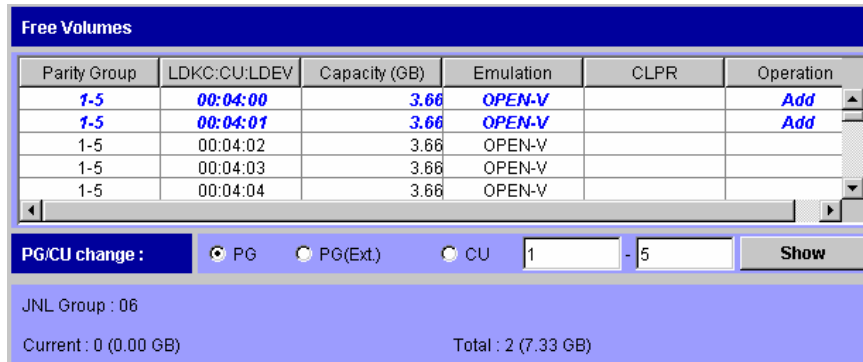


Figure 6-2 The Free Volumes List and the PG Button in the Edit JNL Volumes Window

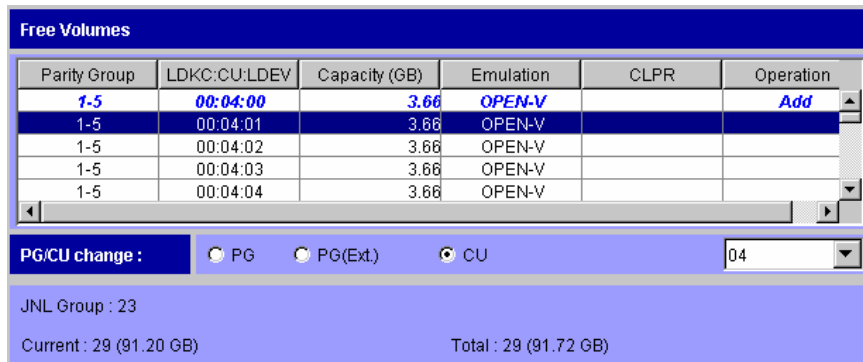


Figure 6-3 The Free Volumes List and the CU Button in the Edit JNL Volumes Window

- **JNL Group:** Indicates the number of a journal group.
- **Current:** Indicates the number and the capacity of journal volumes that currently exist in the journal group.
- **Total:** Estimates the number and the capacity of journal volumes after addition or deletion, and then displays the estimated number and the capacity.
- **Set:** Applies the settings in the window to the Journal Operation window (Figure 4-1).
- **Cancel:** Cancels the settings.

Deleting Journal Volumes from a Journal Group

To delete journal volumes from the primary storage system, log into the primary storage system and then perform the following operation.

To delete journal volumes from the secondary storage system, log into the secondary storage system and then perform the following operation.

1. Ensure that the Storage Navigator main window is in Modify mode.

For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.

2. Ensure that the Journal Operation window is displayed (see section Journal Operation Window).
3. In the tree of the Journal Operation window, select a journal group from below **Registered**.

Note: You cannot select more than one journal group.

4. Do either of the following:
 - In the tree, right-click the selected journal group and then select **JNL Volumes** from the pop-up menu.
 - In the upper-right list, select and right-click the desired journal group and then select **JNL Volumes** from the pop-up menu.

Note: You can delete journal volumes only when the attribute of the journal group is **Initial** or when the status of the journal group is **Stop**.

Caution: As a general rule, you can delete journal volumes only when the attribute of the journal group is **Initial** or when the status of the journal group is **Stop** or **Hold (Failure)**. However, when one journal group uses multiple mirror IDs, you can delete journal volumes only in the case shown in the following table.

Table 6-1 Requirement to Delete Journal Volumes (When One Journal Group Uses Two Mirror IDs)

Status of Journal Group		Can the Journal Volumes be Deleted?
Mirror ID 1	Mirror ID 2	
Active	Hold	No.
Active	Hold(Failure)	No.
Stop	Hold	Yes.
Stop	Hold(Failure)	Yes.

5. In the **JNL Volumes** list of the Edit JNL Volumes window (Figure 6-1), select the volumes that you want to delete.

In the **JNL Volumes** list, one row represents one volume.

6. Select **Delete**.

The volumes that you want to delete will be displayed in blue italics. The **Operation** column of the list displays **Delete**.

7. Select **Set** to close the Edit JNL Volumes window.
8. See the **Preview** list to check the volumes that you want to delete.

If you want to add volumes to delete, select and right-click a volume in the **Preview** list and then select **Modify**. A window appears and allows you to specify volumes to delete.

- If you want to cancel deleting a volume, select and right-click the volume and then select **Cancel**.
- If necessary, you can repeat steps 3 to 8 to specify volumes that should be deleted from other journal groups.

9. Select **Apply** to delete journal volumes.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Displaying Detailed Information about a Journal Group

To display detailed information about journal groups in the primary storage system, log into the primary storage system and then perform the following operation.

To display detailed information about journal groups in the secondary storage system, log into the secondary storage system and then perform the following operation.

1. Ensure that the Journal Operation window is displayed (see section Journal Operation Window).
2. In the tree of the Journal Operation window, select a journal group.

Note: You cannot select more than one journal group.
3. Do either of the following:
 - In the tree, right-click the selected journal group and then select **JNL Groups** and **JNL Status** from the pop-up menu.
 - In the upper-right list, right-click the desired journal group and then select **JNL Groups** and **JNL Status** from the pop-up menu.
4. In the JNL Group Detail window (Figure 6-4), view detailed information about the journal group.
5. After you finish viewing the information, select **Close** to close the JNL Group Detail window.

JNL Group Detail

JNL Group (LDKC) :	C2 (00)		
Attribute :	Master		
JNL Volumes :	62	JNL Capacity (GB) :	194.97
Data Volumes :	256	Data Capacity (GB) :	0.50
Inflow Control :	Yes	Data Overflow Watch (sec) :	40
Copy Pace :	Medium		
Use of Cache :	Use	Speed of Line (Mbps) :	100

JNL Volumes

Parity Group	LDKC:CU:LDEV	Capacity (GB)	Emulation	CLPR
1-3	00:02:00	3.14	OPEN-V	
1-3	00:02:01	3.14	OPEN-V	
1-3	00:02:02	3.14	OPEN-V	
1-3	00:02:03	3.14	OPEN-V	

Mirrors

Mirror ID	Attribute	Status	CTG	S/N (LDKC)	Pair JNLG	CTRLID (Model Na
00	Master	Active	3D	12345 (01)	3D	5(USP V)

Figure 6-4 JNL Group Detail Window

The JNL Group Detail window displays the following:

- **JNL Group(LDKC):** Indicates the number of a journal group and the LDKC number. The LDKC number is enclosed in the parentheses following the number of a journal group.
- **Attribute:** Indicates the attribute of the journal group. When one journal group uses multiple mirror IDs, **Attribute** indicates the attribute of the data volume in the journal group whose mirror ID is not **Hold** or **Hold(Failure)**.

Initial	A journal group in initial status. Journal volumes are registered in this journal group, but no data volumes (primary data volumes nor secondary data volumes) are registered in this journal group.
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.

- **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. When one journal group uses multiple mirror IDs, **Data Volumes** indicates the number of the data volumes in the journal group whose mirror ID is not **Hold** or **Hold(Failure)**.
- **Data Capacity:** Indicates the total capacity of all the data volumes. When one journal group uses multiple mirror IDs, **Data Capacity** indicates the total capacity of the data volumes in the journal group whose mirror ID is not **Hold** or **Hold(Failure)**.
- **Meta/Data Ratio:** Indicates the ratio of metadata area to journal data area. For example, if **16** is displayed, the ratio of metadata area to journal data area is 1:16. If **32** is displayed, the ratio of metadata area to journal data area is 1:32.
- **Extent:** Indicates the number of extents.
- **Inflow Control:** Indicates whether to restrict inflow of update I/Os to the journal volume (in other words, whether to slow delay response to hosts). **Yes** indicates inflow will be restricted. **No** indicates inflow will not be restricted.
- **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**.
- **Copy Pace**: Indicates the pace for an initial copy activity for one volume. One of the following is displayed: **High**, **Medium**, or **Low**.

Notes:

- By default, the pace for an initial copy activity is **Low**.
- **Copy Pace** displays nothing if 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 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.

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. If you set **Use**, this setting only takes effect on the journal volumes of RAID-5 or RAID-6 that are in the journal group. For external volumes, non-RAID-5 journal volumes, and non-RAID-6 journal volumes, **Use** works the same as **Not Use**.

- **Delta resync Failure**: Indicates the processing that would take place when delta resync operation cannot be performed.
 - **Entire**: Whole data in primary data volume will be copied to secondary data volume when delta resync operation cannot be performed.
 - **None**: No processing will take place when delta resync operation cannot be performed. Therefore, the secondary data volume will not be updated.
- **Speed of Line**: Indicates the line speed of data transfer. The unit is Mbps (megabits per second).

One of the following is displayed: **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.

- **JNL Volumes**: Displays a list of registered journal volumes.
 - **Parity Group**: indicates the parity group where a journal volume belongs.
 - **LDKC:CU:LDEV**: Indicates the LDKC number, 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 the attribute of a mirror ID.

Initial A mirror in initial status. No data volumes are registered in the journal groups of local storage system.

Master A mirror to whom primary data volumes of local storage system are registered.

Restore A mirror to whom secondary data volumes of local storage system are registered.

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

Initial A journal group in initial status.
Journal volumes are registered in this journal group, but no data volumes (primary data volumes nor secondary data volumes) are registered in this journal group.
When you create an 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: 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 window. For details on how to restore them, see section Restoring a Pair of Data Volumes (Pairresync).

HaltAccept An operation for splitting the mirror has been started. The status of the mirror 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.

Hold A UR pair for delta resync operation is created.

Hold(Failure) An error occurred with the UR pair for delta resync operation.

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(LDKC):** Indicates the serial number of the remote storage system and the LDKC number.
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 (model name that indicates the model) 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 USP V storage system is **5**.
- **Path Watch Time:** Indicates the time for monitoring blockade of paths to the remote storage system. If the status of the mirror where secondary data volumes in local storage system are registered is **Hold** or **Hold(Failure)**, this column will be blank.
- **Volume:** Indicates the number of data volumes that are registered in the mirror.
- **Capacity:** Indicates the total capacity of data volumes that are registered in the mirror.
- **Previous:** Displays detailed information about the previous journal group.
- **Next:** Displays detailed information about the next journal group.
- **Close:** Closes the JNL Group Detail window.

Changing Options for a Journal Group

To change options for a journal group, follow the procedure below. Both the primary storage system administrator and the secondary storage system administrator must perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Journal Operation window is displayed.
3. Do either of the following:
 - In the tree, right-click a journal group from below **Registered** 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.

Caution: You can select **Change JNL Option** only when one of the following conditions is satisfied:

- The attribute of the journal group is **Initial**.
- The status of the journal group is **Active**.
- The status of the journal group is **Stop**.
- The status of the journal group is **Hold**.
- The status of the journal group is **Hold(Failure)**.

When one journal group uses multiple mirror IDs, whether you can or cannot select **Change JNL Option** depends on the status of the journal group whose mirror is not in **Hold** or **Hold(Failure)** status.

4. In the Change JNL Option window (Figure 6-5), change journal group options and then select **Set**.

Caution: If the status of the journal group is **Active**, or if the status of the master journal group is **Hold**, you can change only the **Line of Speed** option.

Note: If you have selected two or more journal groups in step 3, and want to make some options in the Change JNL Options window unchanged, then you must specify blank for the options. If you specify blank for an option, the option value will remain unchanged. If you specify a value other than blank, the specified value will be set to the selected journal groups.

5. See the **Preview** list in the Journal Operation window to check the settings that you have made.
 - If you want to modify a setting, select and right-click the setting in the **Preview** list and then select **Modify**.
 - If you want to cancel a setting, select and right-click the setting and then select **Cancel**.
6. Select **Apply** to apply the settings.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Figure 6-5 Change JNL Option Window

The Change JNL Option window displays the following:

- **Meta/Data Ratio:** Allows you to specify the ratio of metadata area to journal data area.

For example, if you select **16**, the ratio of metadata area to journal data area will be 1:16. If you select **32**, the ratio of metadata area to journal data area will be 1:32.

Note: In the current version, you cannot change the ratio of metadata area to journal data area.
- **Extent:** Allows you to specify the number of extents.

Note: In the current version, the number of extents is always 32 and cannot be changed.
- **Inflow Control:** Allows you to specify whether to restrict inflow of update I/Os to the journal volume (in other words, whether to slow delay response to 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.
- **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 want to specify **Medium**, please ensure that the amount of update I/Os (i.e., write requests from hosts to primary data volumes) is 10 MB/s or less per one parity group. If it exceeds 10 MB/s, data volume pairs may become split (suspended).
 - **High:** The speed of the initial copy activity is faster than **Low** and **Medium**.

If you want to specify **High**, please ensure that update I/Os (i.e., write requests from hosts to primary data volumes) will not occur. If update I/Os occurs, data volume pairs may become split (suspended).

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

- **Unit of Path Watch Time:** Allows you to specify the unit of path watch time. You can specify one of the following: **minute**, **hour**, or **day**.
- **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 59 minutes, 1 to 23 hours, or 1 to 30 days. You can specify a numeric value in **Path Watch Time**.

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.

Important: If you want a mirror to get split (suspended) immediately after a path gets blocked, please ask the Support Center to set the system option mode 448 to ON and set the system option mode 449 to OFF. For detailed information about the modes 448 and 449, see the next table.

Table 6-2 System Option Modes

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

- **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 from the master journal group to the restore journal group, 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 secondary site. Select **Yes** in the primary site of UR pair for delta resync operation since you need to set path watch time in the secondary site to perform delta resync operation.

- **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. If you set **Use**, this setting only takes effect on the journal volumes of RAID-5 or RAID-6 that are in the journal group. For external volumes, non-RAID-5 journal volumes, and non-RAID-6 journal volumes, **Use** works the same as **Not Use**.

- **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:** Whole data in primary data volume will be copied to secondary 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 secondary data volume will not be updated.

Caution:

- This option cannot be specified in the secondary site.
 - If the journal group is created with the DKCMAIN micro program whose version is older than 50-07-6X-XX/XX, the default is **None**.
- **Set:** Applies the settings in the window to the Journal Operation window (Figure 4-1).
 - **Cancel:** Cancels the settings.

Deleting a Journal Group

To delete a journal group from the primary storage system, log into the primary storage system and then perform the following operation.

To delete a journal group from the secondary storage system, log into the secondary storage system and then perform the following operation.

Note: You cannot delete master journal groups and restore journal groups. You can only delete journal groups whose attribute is **Initial**.

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Journal Operation window is displayed (see section Journal Operation Window).
3. Do either of the following:
 - In the upper-right list, right-click a journal group whose attribute is Initial.
 - In the tree, select and right-click a journal group in initial status (🟡).
4. Select JNL Groups and Delete JNL from the pop-up menu.
5. See the **Preview** list to check the journal groups that you want to delete.
 - If you want to cancel deleting a journal group, select and right-click the journal group and then select Cancel.
 - If necessary, you can repeat steps 3 to 5 to specify other journal groups.
6. Select **Apply** to apply the settings.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Splitting a Mirror (Pairsplit-r: Suspending a copy operation)

A mirror is a combination of a master journal group and a restore journal group. If you split a mirror, the copying of data from the master journal group to the restore journal group will be suspended.

To split a mirror, follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Journal Operation window is displayed.
3. Do either of the following:
 - In the tree, select a master journal group (M) or a restore journal group (R) from below Registered.
 - In the tree, select **Journal Groups** or **Registered** and then select a master journal group (M) or a restore journal group (R) from the list to the right.

The list displays a list of mirrors. One row in the list represents one mirror.

4. Locate mirrors that are in Active status.
5. From the mirrors, select and right-click the mirror(s) that you want to split.
6. From the pop-up menu, select JNL Groups and Pairsplit-r.
7. In the Pairsplit-r window (Figure 6-6), specify parameters and then select Set.
8. See the **Preview** list to check the mirrors that you want to split.
 - If you want to modify parameters of a mirror, select and right-click the mirror and then select Modify.
 - If you want to cancel splitting a mirror, select and right-click the mirror and then select Cancel.
 - If necessary, you can repeat steps 3 to 8 to specify other mirrors.
9. Select **Apply** to split the mirror(s).

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

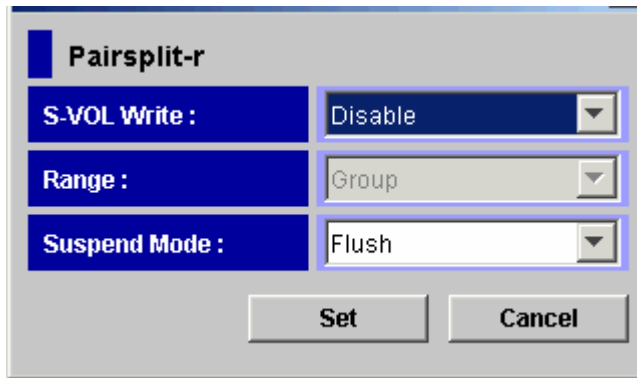


Figure 6-6 Pairsplit-r Window

The Pairsplit-r window displays the following:

- **S-VOL Write:** Allows you to specify whether hosts can issue read I/O and write I/O to the secondary data volume after the mirror is split. The default is **Disable**.

If you select **Enable**, hosts can issue read I/O and write I/O to the secondary data volume after you split the mirror.

If you select **Disable**, hosts cannot issue read I/O and write I/O to the secondary data volume after you split the mirror.

This parameter takes effect only when the selected volume is a master journal group.

- **Range:** Allows you to specify the split range. This parameter is always set to **Group** and cannot be changed.
- **Suspend Mode:** Allows you to specify how to handle update data that are not reflected to the secondary data volume. The default is **Flush**.

- If you select **Flush**, update data will be reflected to the secondary data volume when you split the mirror.

When the secondary storage system receives a request for splitting a mirror, all the journal data (i.e., update data) that the mirror retains will be written to the secondary data volume. After that, the status of the mirror will change from **Suspending** to **PSUS** if the mirror does not retain any journal data for a certain period of time.

- If you select **Purge**, update data will not be reflected to the secondary data volume when you split the mirror. If you restore the mirror, the update data will be reflected to the secondary data volume.

Restoring a Mirror (Pairresync: Resuming a copy operation)

A mirror is a combination of a master journal group and a restore journal group. If you split a mirror, the copying of data from the master journal group to the restore journal group will be suspended. To restart the copying of data, you need to restore the mirror. As for a mirror in **Hold** status, you need to restore the mirror in order to restart the copying of data.

To restore a mirror, follow the procedure below. Only the primary storage system administrator can perform this operation; the secondary storage system administrator cannot perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Journal Operation window is displayed (see section Journal Operation Window).
3. Do either of the following:
 - In the tree, select a master journal group (M) or a restore journal group (R) from below Registered.
 - In the tree, select **Journal Groups** or **Registered** and then select a master journal group (M) from the list to the right.

The list displays a list of mirrors. One row in the list represents one mirror.

4. Locate mirrors that are in Stop, **Hold**, or **Hold(Failure)** status.
5. From the mirrors, select and right-click the mirror(s) that you want to restore.

When you select the mirror in **Stop** status, the copying of data will be restarted after the recovery.

The result of this operation differs according to the status of the selected mirror.

- When you select the mirror in Stop status, the copying of data will be restarted after the recovery.
 - When you select the mirror in Hold status, delta resync operation will be performed after the recovery.
 - When you select the mirror in **Hold(Failure)** status, the status will change to **Hold** after the recovery.
6. From the pop-up menu, select JNL Groups and Pairresync.
 7. See the **Preview** list to check the mirrors that you want to restore.
 - If you want to cancel restoring a mirror, select and right-click the mirror and then select **Cancel**.
 - If necessary, you can repeat steps 3 to 7 to specify other mirrors.

8. Select **Apply** to restore the mirror(s).

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Deleting Data Volumes from a Mirror (Pairsplit-S: Ending a copy operation)

A mirror is a combination of a master journal group and a restore journal group. If you delete data volumes from a mirror, the copying of data from the master journal group to the restore journal group will end.

When one journal group uses two mirror IDs, the result of the operation differs as follows according to the status of the specified mirror.

When you specify a mirror in **Hold** or **Hold(Failure)** status, only the UR pair of the specified mirror will be deleted.

When you specify a mirror that is not in **Hold** or **Hold(Failure)** status, UR pairs of both mirrors (including the mirror which you do not specify) will be deleted.

To delete data volumes from a mirror, follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Journal Operation window is displayed (see section Journal Operation Window).
3. Do either of the following:
 - In the tree, select a master journal group (M) or a restore journal group (R) from below **Registered**.
 - In the tree, select **Journal Groups** or **Registered** and then select a master journal group (M) or a restore journal group (R) from the list to the right.

The list displays a list of mirrors. One row in the list represents one mirror.

4. Locate mirrors that are not in Initial status.
5. From the mirrors, select and right-click the mirror(s) from which you want to delete data volumes.
6. From the pop-up menu, select JNL Groups and Pairsplit-S.
7. See the **Preview** list to check the settings that you have made.
 - If you want to cancel a setting, select and right-click the setting and then select **Cancel**.
 - If necessary, you can repeat steps 3 to 7 to specify other mirrors.
8. Select **Apply** to delete data volumes from the mirror(s).

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

Performing UR Pair Operations

This chapter discusses using the Pair Operation window of UR.

- Reviewing the Pair Operation Window
- Creating a Pair of Data Volumes (Paircreate)
- Displaying Detailed Information about a Pair of Data Volumes
- Changing Options for a Data Volume Pair
- Splitting a Pair of Data Volumes (Pairsplit-r)
- Restoring a Pair of Data Volumes (Pairresync)
- Releasing a Pair of Data Volumes (Pairsplit-S)
- Creating a Point-in-Time Copy of Data Volumes
- Recovering a Pinned Track

Reviewing the Pair Operation Window

To perform remote copy operations with UR, you must create pairs of a *primary data volume* and a *secondary data volume*. A primary data volume is a data volume to be copied. A secondary data volume is a data volume that is the copy destination for a primary data volume. To manipulate pairs of these data volumes, use the Pair Operation window of UR.

Note: This manual sometimes uses the term "volume pair" when referring to a pair of a primary data volume and a secondary data volume.

Filtering Information in the List in the Pair Operation Window

The Display Filter window (Figure 7-1) enables you to filter information in the list in the Pair Operation window. The list will display only the volumes that satisfy certain conditions.

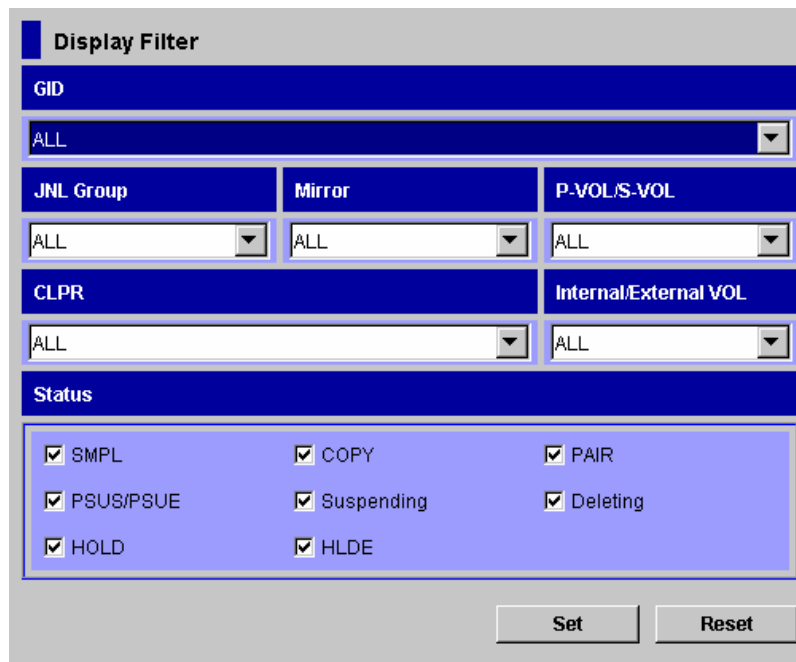
To filter information in the list in the Pair Operation window:

1. In the Pair Operation window (Figure 4-2), select Display Filter.
2. Use the Display Filter window to specify information that you want to display in the list.
3. Select **Set** to close the Display Filter window.

The list in the Pair Operation window displays the specified information only.

Notes:

- To restore the list to the default status (i.e., the status before filtering), select **Reset** in the Display Filter window and then **Set** in the same window.
- If you change settings in the Display Filter window, the changes will remain effective during the UR session until you reset the settings.



The screenshot shows the 'Display Filter' window with the following sections:

- GID:** A dropdown menu set to 'ALL'.
- JNL Group:** A dropdown menu set to 'ALL'.
- Mirror:** A dropdown menu set to 'ALL'.
- P-VOL/S-VOL:** A dropdown menu set to 'ALL'.
- CLPR:** A dropdown menu set to 'ALL'.
- Internal/External VOL:** A dropdown menu set to 'ALL'.
- Status:** A list of status items with checkboxes, all of which are checked:
 - SMPL
 - PSUS/PSUE
 - HOLD
 - COPY
 - Suspending
 - HLDE
 - PAIR
 - Deleting

At the bottom of the window are two buttons: 'Set' and 'Reset'.

Figure 7-1 Display Filter Window

The Display Filter window displays the following:

- **Port:** Allows you to select the port to be displayed. If you select **ALL**, all ports will be displayed in the list.

Note: You can specify the port number with two characters. For instance, you can abbreviate CL1-A to 1A. You can also enter the port number in both lowercase and uppercase characters.

- **GID:** Allows you to select the group number of a host group or an iSCSI target to be displayed.
If you select **ALL**, all host groups and iSCSI targets will be displayed in the list.
- **JNL Group:** Allows you to select the journal group to be displayed.
If you select **ALL**, all journal groups will be displayed in the list.
- **Mirror:** Allows you to select the mirror to be displayed.
If you select **ALL**, all mirrors will be displayed in the list.
- **P-VOL/S-VOL:** Allows you to select the type of volumes to be displayed.
If you select **P-VOL**, primary data volumes will be displayed in the list.
If you select **S-VOL**, secondary data volumes will be displayed in the list.
If you select **ALL**, all primary data volumes and all secondary data volumes will be displayed in the list.
- **CLPR:** Allows you to select the CLPR to be displayed.
If you select **ALL**, all CLPRs will be displayed in the list.
- **Internal/External VOL:** Allows you to specify whether internal volumes or external volumes will be displayed in the list.
If you select **ALL**, all the internal and external volumes will be displayed in the list.
- **Status:** Allows you to specify pair statuses. The volumes of the specified statuses will be displayed in the list.
- **Set:** Applies the settings in the Display Filter window to the list.
- **Reset:** Restores all options in the Display Filter window to the default state. All the drop-down lists will display **ALL**. All the check boxes will be selected.

Saving Pair Status Information into a Text File

The export function of UR enables you to save information about pair status into a text file. The example of the text file is described in Figure 7-2.

Note: You can use the export function only while the client file access is enabled. When the client file access is not enabled, the **Export** button is not displayed on the Pair Operation window. For the further information of the client file access settings, please refer to the *Storage Navigator User's Guide*.

The following is an example of an exported text file.

```

<Pair Status at 30/04/2005 12:34:56>
Port Gr (Name) LUN Vol Status JNG Mr S/N(LDKC) CNL Port Gr LUN Vol JNG Type CTG ErrLv
Sync CLPR CopyTime
CL3-A 00(3A-G00 ) 100 00:08:00 PAIR 01 02 11111(00) 05 CLA-R FF 7EE 00:3F:FF F6 P-VOL 09 Group
100 00:CLPR0 001:00:00
CL3-A 00(3A-G00 ) 101 00:08:01 SMPL
01:CLPR1
CL3-A 00(3A-G00 ) 102 00:08:02 PSUS 10 02 11111(00) 05 CL1-E FF 7EE 00:3F:F6 ED P-VOL 12 LU
80 02:CLPR2

```

(1) (2) (3) (4) (5) (6) (7) (8) (9) (10) (11) (12) (13)
 (14) (15) (16) (17) (18) (19) (20)

- | | |
|--|---|
| <p>(1) Port in the local storage system</p> <p>(2) Group number and name of host group or iSCSI target in the local storage system
(The length of a name is limited to a maximum of 16 characters)</p> <p>(3) LU number in the local storage system</p> <p>(4) LDKC number, CU number and LDEV number in the local storage system</p> <p>(5) Status of the pair</p> <p>(6) Journal group number in the local storage system</p> <p>(7) Mirror ID</p> <p>(8) Serial number of the remote storage system</p> <p>(9) Controller ID and model name of the remote storage system</p> <p>(10) Port number in the remote storage system</p> | <p>(11) Group number of host group or iSCSI target in the remote storage system</p> <p>(12) LU number in the remote storage system</p> <p>(13) LDKC number, CU number and LDEV number in the remote storage system</p> <p>(14) Journal group number in the remote storage system</p> <p>(15) Type of volumes in the local storage system</p> <p>(16) Consistency group number</p> <p>(17) Error level</p> <p>(18) Progress of copy operation. Or, synchronization rate (i.e., concordance rate) between the data volume before the pair was split and the data volume after the pair was split.</p> <p>(19) CLPR number and CLPR name</p> <p>(20) Copy time</p> |
|--|---|

Note: If a 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*.

Note: "00" is indicated as the LDCK number when the volume is a volume of TagmaStore USP/NSC.

Figure 7-2 Example of an Exported Text File

To save information in a text file:

1. If necessary, follow the instruction in section Filtering Information in the List in the Pair Operation Window to filter information in the list.
2. In the Pair Operation window (Figure 4-2), click **Export**.
A message appears and asks whether you want to save volume pair information in a file.
3. Click **OK** to close the message.
A window appears and prompts you to save a file.
4. Enter the name of the file and then click **Save**.
A message appears when UR finishes saving the file.
5. Click **OK** to close the message.
The export processing stops if **Subsystem** or an LDKC is selected in the tree, and you click **Cancel** in the dialog box that is displayed during export processing.




Creating a Pair of Data Volumes (Paircreate)

Using the Paircreate window, the administrator of the primary storage system must create pairs of a primary data volume and a secondary data volume. When creating a pair, the administrator can specify options for initial copy activities.


Before creating one or more pairs, the administrator must make sure that:

- The volumes to be used as secondary data volumes are offline to all hosts.
- Ports have already been configured (see section Configuring Port Attributes)
- The primary storage system is associated with the secondary storage system (see section Establishing Relationship between Primary and Secondary Storage Systems (Add DKC))

After all these points are confirmed, the administrator of the primary storage system can follow the procedure below to create a pair of a primary data volume and a secondary data volume. Note that the administrator of the secondary storage system cannot perform the following operation.

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Pair Operation window is displayed (see section Pair Operation Window).
3. In the tree, select a port , or a host group , or an iSCSI target .
4. In the list, select and right-click the volume that you want to use as a primary data volume.

Notes:

- Volumes with the  icon are already used as primary data volumes.
 - You can select and right-click more than one volume if you want to create more than one pair at one time. Note that you will need to choose all the secondary data volumes from the same secondary storage system.
5. From the pop-up menu, select **Paircreate**.
The Paircreate window appears (Figure 7-3).
 6. In **DKC**, specify the secondary storage system.
 7. In S-VOL, specify the secondary data volume.
 - Use the leftmost drop-down list to specify a port.
 - Use the drop-down list in the middle to specify a GID, which is a group number of a host group or an iSCSI target.
 - Use the rightmost drop-down list to specify an LUN.

Note: If more than one primary data volume is specified in the list, you can specify secondary data volumes in three ways. For detailed information, refer to the explanation about the **Select Other S-VOL(s)** option later in this section.

8. Use the **Mirror** drop-down lists to specify the following items sequentially:
 - Master journal group
 - Mirror ID
 - Restore journal group
9. In **CTG**, specify a consistency group number.
10. Do one of the following to make settings on the initial copy operation:
 - Select **Entire** if you want to start the initial copy operation immediately after the pair is created.
 - Select **None** if you do not want to start the initial copy operation immediately after the pair is created.
 - Select **Delta** if you want to create a UR pair for delta resync operation.

Note: You need to meet the requirements listed in section Requirements for Creating UR Pair for Delta Resync Operation to create a UR pair for delta resync operation.

11. Use **Error Level** to specify the range used for splitting a pair when a failure occurs.
12. Select **Set** to close the Paircreate window.
13. See the **Preview** list in the Pair Operation window 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**.
 - If you want to cancel a setting, select and right-click the setting and then select **Cancel**.
14. Select **Apply** to start pair creation and the initial copy operation.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

15. In the list of the **Pair Operations** window, verify that the status of the pair is correct.

If you selected **Entire** or **None** in step 10, the status of the pair will be **COPY** or **PAIR**. If you selected **Delta**, the status of the pair will be **HOLD**.
16. To monitor the progress of the paircreate operation(s), click **File, Refresh** on the menu bar of the Storage Navigator main window to update the information in the list, or use the Detailed Information window (see Figure 7-4) to monitor the detailed status of each pair.

The screenshot shows the Paircreate window with the following fields and values:

- Port-GID-LUN:** CL3-A - 00 - 062 (00:07:62)
- P-VOL:** CL3-A - 00 - 062 (00:07:62)
- S-VOL:** CL1-A, 00, 000
- Select Other S-VOL(s):** (empty dropdown)
- Mirror:** M-JNL: 02, Mirror ID: 00, R-JNL: 00
- CT Group:** 00
- DKC:** S/N, CTRLID, Path: 10000(01), 5(USP V), Fibre
- Initial Copy:** Entire
- Priority:** 32
- Error Level:** Group
- M-JNL Information:** Current Mirror(s) : 0, Total Mirror(s) : 0

Buttons at the bottom: Set, Cancel

Figure 7-3 Paircreate Window

The Paircreate window displays the following:

- **P-VOL:** Indicates a primary data volume. The numbers indicates the port number, the GID, and the LUN of the primary data volume. The numbers in parentheses are the LDKC number, the CU number and the LDEV number. The GID is a group number for a host group or an iSCSI target.

Notes:

- **P-VOL** displays only one primary data volume even when two or more primary data volumes are selected in the Pair Operation window. **P-VOL** only displays the primary data volume that has the smallest LUN (LU number).
- If a 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*.
- **S-VOL:** Allows you to select a secondary data volume. When specifying a secondary data volume, you must use the three drop-down lists to select a port, a GID, and an LUN of the secondary data volume. The GID is a group number for a host group or an iSCSI target.

Note: You can select only one secondary data volume even when two or more primary data volumes are selected in the Pair Operation window. UR automatically specifies the other secondary data volumes. For detailed information, refer to the explanation about the **Select Other S-VOL(s)** option later in this section.

Note: If you selected more than one primary data volume, select the secondary data volume for the primary data volume being displayed. The secondary 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 secondary data volumes for the two other primary data volumes will be LUN02 and LUN03.

Note: When creating more than one pair at a time, make sure that the S-VOLs to be assigned automatically are available. If an S-VOL is not available, you will have to select the S-VOL individually.

Note: You can specify the port number with two characters. For instance, you can abbreviate CL1-A to 1A. You can also enter the port number in both lowercase and uppercase characters.

- **Select Other S-VOL(s):** Specify how UR automatically will assign secondary data volumes if two or more primary data volumes are selected in the Pair Operation window.
 - **Increment:** UR increments LUNs of the resulting secondary data volumes one by one. For example, if the LUNs of the three primary data volumes are *001*, *002*, and *003*, the LUNs of the resulting secondary data volumes will be *011*, *012*, and *013* if you select the LUN *011* with the **S-VOL** drop-down list.
 - **Input Next:** UR users can specify a LUN for each primary data volume. For example, if two primary data volumes are selected, the UR user must select **Input Next**, specify the secondary data volume for the first primary data volume with **S-VOL**, and then click **Set**. Next, the user must specify the secondary data volume for the second primary data volume with S-VOL, and then click **Set**.
 - **Relative to P-VOL:** UR calculates the difference between LUNs of two adjacent primary data volumes, and then determines LUNs of secondary data volumes.

For example, if the LUNs of the three primary data volumes are *001*, *005*, and *006*, the LUNs of the resulting secondary data volumes will be *002*, *006*, and *007* if you select the LUN *002* with the **S-VOL** drop-down list.

Note: If you select and right-click more than one volume in the **Preview** list and then select **Modify**, the Paircreate window appears and the **Select Other S-VOL(s)** drop-down list displays **No Change**. If you want to change the setting with this drop-down list, select another item from the drop-down list.

- **Mirror:** Assigns a master journal group to the primary data volume, assigns a restore journal group to the secondary data volume, and also assigns a mirror ID to the volume pair.
 - **M-JNL:** Assigns a master journal group to the primary data volume.
Note: **M-JNL** does not display journal group numbers used by Universal Replicator for z/OS®.
 - **Mirror ID:** Assigns a mirror ID to the volume pair.

- **R-JNL**: Assigns a restore journal group to the secondary data volume.
Note: **R-JNL** does not display journal group numbers used by Universal Replicator for z/OS®.
- **CT Group**: Allows you to specify a consistency group.
Note: This option cannot be changed after a pair has been created.
- **DKC**: Allows you to select the serial number with the number of LDKC and the controller ID with the model name of the secondary storage system. This option also allows you to specify the path type (i.e., channel type).
Note: The secondary storage system must be the same for all pairs being created during one operation.
- **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 will not start after the volume pair is created. The primary storage system starts copying of update data as needed.
Caution: The user must ensure that the primary data volume and secondary data volume are already identical when using **None**.
 - **Delta**: No initial copy operation will 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.
- **Priority**: Allows you to specify priority (scheduling order) of the initial copy operations (1-256). The default setting is 32.
Note: If the time out error occurs during the Paircreate operation, the copy operation may not be executed correctly in the order that you have set with the **Priority** parameter. The time out error may be caused by the CU configuration or a remote copy path error. Review the error, release the pair with the error, and then retry the Paircreate operation.
- **Error Level**: Allows you to specify 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.
- **M-JNL Information**: Displays information about the master journal.
 - **Current Mirror(s)**: Indicates the number of mirrors registered in the master journal.
 - **Total Mirror(s)**: Indicates the sum of the following:
 - (1) The number of mirrors registered in the master journal.
 - (2) The number of mirrors to be added by the Paircreate window

(3) The number of mirrors for volumes in the **Preview** list in the Pair Operation window.

- **Set:** Applies the settings to the **Preview** list in the Pair Operation window (Figure 4-2).
- **Cancel:** Discards the settings.

Displaying Detailed Information about a Pair of Data Volumes

The Detailed Information window allows you to view detailed information about a pair of data volumes.

To display detailed information about a pair of data volumes, follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the Pair Operation window is displayed (see section Pair Operation Window).
2. In the tree of the Pair Operation window, select and right-click a pair.
Note: You cannot select more than one pair.
3. From the pop-up menu, select Detailed Information.
4. In the Detailed Information window (Figure 7-4), view detailed information about the pair.
5. After you finish viewing the information, select **Close** to close the Detailed Information window.

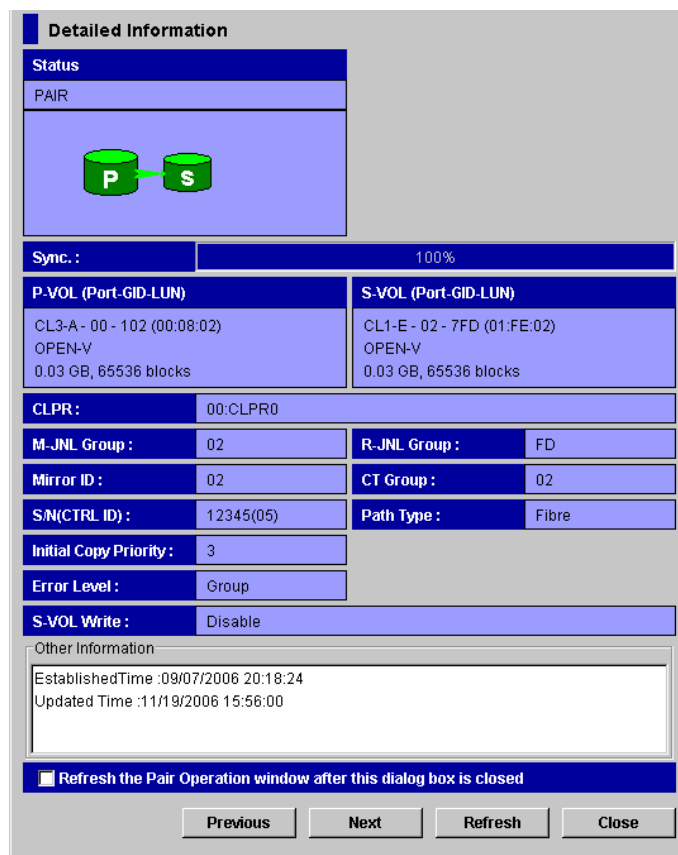


Figure 7-4 Detailed Information Window

The Detailed Information window displays the following:

- **Status:** Indicates the status of the pair. If the pair is split (or suspended), **Status** also displays the suspend type. If the pair is waiting for initial copy, **Status** also displays the word (**Queuing**).
- **Sync.:** If the volume in the local storage system is a primary data volume, **Sync.** displays progress of an initial copy operation.

If the volume in the local storage system is a secondary data volume, **Sync.** displays information in the following ways:

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

Caution: 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 window (Figure 7-4) displays the phrase "Initial copy failed".

- In the following cases, this column will be blank.

This column is blank when the volume in the local storage system is neither a primary data volume nor a secondary data volume.

When the status of the volume pair is **HOLD** or **HLDE**.

- **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 or an iSCSI target. If the primary data volume exists in the local 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 (i.e., 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.

Note: If a 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*.

Note: If the primary data volume is a volume of TagmaStore USP/NSC, "00" is displayed as for the LDKC number.

- **S-VOL:** Indicates the secondary 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 secondary data volume; the GID is a group number for a host group or an iSCSI target. If the secondary data volume exists in the local storage system, the first line also displays the CLPR number and the CLPR name. If the secondary data volume is an LUSE volume, the LUN is the LDEV number of the top LDEV (i.e., 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.

Note: If a 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*.

Note: If the secondary data volume is a volume of TagmaStore USP/NSC, "00" is displayed as the LDKC number.

- **CLPR:** Indicates the CLPR number and the CLPR name of the volume in the local 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 secondary storage system. The control ID is enclosed by parentheses.
- **Path Type:** Indicates the channel type of the path interface between the storage systems (Fiber, serial).

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 secondary 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 window closes.
- **Previous:** Displays the pair status information for the previous pair in the list (the pair in the row above).
Note: The Display Filter settings (see section Filtering Information in the List in the Pair Operation Window) can affect how **Previous** or **Next** is recognized.
Note: The list displays a maximum of 1,024 rows at once. The **Previous** and **Next** buttons on the Detailed Information window can only be used for the currently displayed 1,024 rows.
- **Next:** Displays the pair status information for the next pair in the list (the pair in the row below)
Note: The Display Filter settings (see noted above) can affect how **Previous** or **Next** is recognized.
Note: The list displays a maximum of 1,024 rows at once. The **Previous** and **Next** buttons on the Detailed Information window can only be used for the currently displayed 1,024 rows.
- **Refresh:** Updates the pair status information.
- **Close:** Closes the Detailed Information window.

Changing Options for a Data Volume Pair

The Change Pair Option window allows you to change options for a pair of data volumes.

To change the pair options, follow the procedure below. Note that only the administrator of the primary storage system can perform the following operation. The administrator of the secondary storage system cannot perform the following operation:

1. Ensure that the Storage Navigator main window is in Modify mode.

For detailed information about how to do this, please refer to the *Storage Navigator User's Guide*.

2. Ensure that the Pair Operation window is displayed (see section Pair Operation Window).

3. In the list, select and right-click the desired pair(s).

Caution: You can only change options for pairs in **PAIR**, **PSUS**, or **PSUE**, **HOLD**, or **HLDE** status. If you select a pair in another status, you cannot go on to the subsequent steps.

4. From the pop-up menu, select **Change Pair Option**.

The Change Pair Option window appears (Figure 7-5).

5. In the Change Pair Option window, change pair options and select **Set**.

Note: If two or more pairs are selected in the list, the drop-down list on the window may display blank for option values. If you do not want to change the current option value for each of the pairs, ensure that the drop-down list displays blank. If you select a value other than blank from the drop-down list, the selected value will be set to each of the pairs.

6. See the **Preview** list in the Pair Operation window 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**.

- If you want to cancel a setting, select and right-click the setting and then select **Cancel**.

7. Select **Apply** to apply the changes that you have made.

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

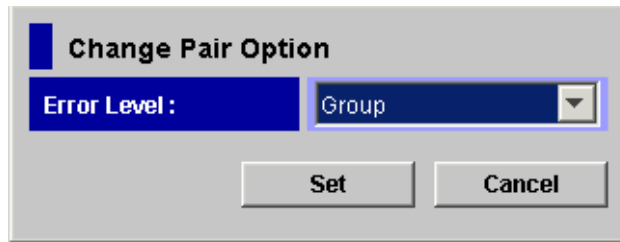


Figure 7-5 Change Pair Option Window

The Change Pair Option window displays the following items that can be configured:

- **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.



When the Change Pair Option window appears, the window displays the current option value.

- **Set:** Applies the settings to the **Preview** list in the Pair Operation window (Figure 4-2).
- **Cancel:** Discards the settings.

Splitting a Pair of Data Volumes (Pairsplit-r)

After the initial copy operation finishes, you can use the Pairsplit-r window to split the pair of data volumes.

To split one or more pairs, follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Pair Operation window is displayed (see section Pair Operation Window).
3. In the tree, select a port  or a host group (.
4. In the list, select and right-click the pair that you want to split.
The pair status must be COPY or PAIR.
5. From the pop-up menu, select **Pairsplit-r**.
The Pairsplit-r window appears (Figure 7-6).
6. In the Pairsplit-r window, select the desired pairsplit options for the pair(s).
7. Select Set to close the Pairsplit-r window.
8. See the **Preview** list in the Pair Operation window 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**.
 - If you want to cancel a setting, select and right-click the setting and then select **Cancel**.
9. Select **Apply** to split pair(s).
Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.
10. In the list of the **Pair Operations** window, verify that the pair has been split successfully. If the pair has been split, the status of the pair is **PSUS**.
Note: To restore a pair that has already been split, use the **Pairresync** command (see the next section)

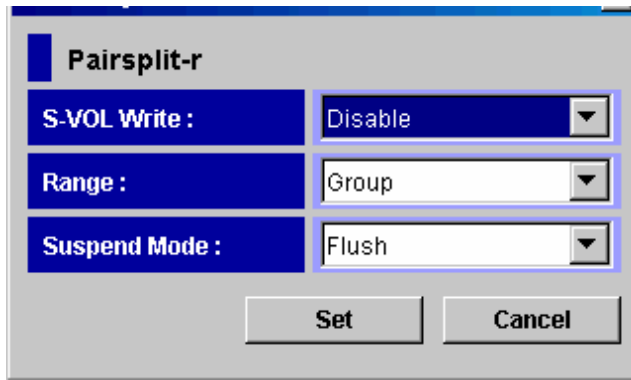


Figure 7-6 Pairsplit-r Window

The Pairsplit-r window displays the following:

- **S-VOL Write:** Allows you to specify whether to permit hosts to write data to the secondary data volume. The default is **Disable** (i.e., do not permit):
 - **Disable:** Hosts cannot write data to the secondary data volume while the pair is split.
 - **Enable:** Hosts can write data to the secondary data volume while the pair is split. This option is available only when the selected volume is a primary data volume.
- **Range:** Allows you to specify the split range. The default is **LU** if two or more pairs in the same consistency group are selected. The default is **Group** if not.
 - **LU:** Only the specified pair(s) 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 frequency of host I/Os. The suspend pair operations should be performed under light I/O load conditions.
 - **Group:** All pairs in the same consistency group(s) as the selected pair(s) 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 **Preview** 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 **Preview** 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 secondary data volume. The default is **Flush**:

- **Flush:** When you split the pair, update data will be copied to the secondary data volume.

When the secondary storage system receives a request for splitting a pair, all the journal data (i.e., update data) that the pair retains will be written to the secondary data volume. After that, the status of the pair will change from **Suspending** to **PSUS** if the pair does not retain any journal data for a certain period of time.



- **Purge:** When you split the pair, update data will not be copied to the secondary data volume. If you restore the pair later, the update data will be copied to the secondary data volume.
- **Set:** Applies the settings to the **Preview** list in the Pair Operation window (Figure 4-2).
- **Cancel:** Discards the settings.

Restoring a Pair of Data Volumes (Pairresync)

The Pairresync window (Figure 7-7) allows you to restore a pair that has been split. In addition, the window allows you to recover a mirror and start data copying of UR pair in **HOLD** status in order to perform delta resync operation.

Note: If the primary or secondary storage system is powered off and its backup batteries are fully discharged while pairs are suspended, the M-VOL/R-VOL bitmaps will not be retained. In this unlikely case, the primary/secondary storage system will mark all cylinders/tracks 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 administrator of the primary storage system can follow the procedure below to restore one or more pairs that have been split, or to perform delta resync operation. Note that the administrator of the secondary storage system cannot perform the following operation.

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Pair Operation window is displayed (see section Pair Operation Window).
3. If any pair was suspended due to an error condition (use the Detailed Information window to view the suspend type), make sure that the error condition has been removed. The primary storage system will not resume the pair(s) until the error condition has been removed.
4. In the tree, select a port  or a host group .
5. In the list, select and right-click the pair that you want to restore.
When you select the pair in **PSUS** or **PSUE** status, the split pair will recover.
The result of this operation differs according to the status of the selected mirror.
 - When you select the pair in PSUS or PSUE status, the split pair will recover.
 - When you select the pair in HOLD status, delta resync operation will be performed.
 - When you select the pair in **HLDE** status, the status will change to **HOLD**.
6. From the pop-up menu, select **Pairresync**.
The Pairresync window appears (Figure 7-7).
7. In the Pairresync window, select the desired options for the pair(s).
8. Select Set to close the Pairresync window.

9. See the **Preview** list in the Pair Operation window 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**.
 - If you want to cancel a setting, select and right-click the setting and then select **Cancel**.

10. Select **Apply** to restore pair(s).

Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.

11. In the list of the **Pair Operations** window, verify that the pair(s) has been restored successfully.

If the pair has been restored successfully, the status of the pair is **COPY** or **PAIR**.

- If delta resync operation has been performed successfully, the status of the pair is **COPY** or **PAIR**.
- If the operation to change pair status has been performed successfully, the status of the pair changes to **HOLD**.

12. To monitor the progress of restoring pair(s), click **File, Refresh** on the menu bar of the Storage Navigator main window to update the information in the list, or use the Detailed Information window (see section Displaying Detailed Information about a Pair of Data Volumes) to monitor the detailed status of each pair.

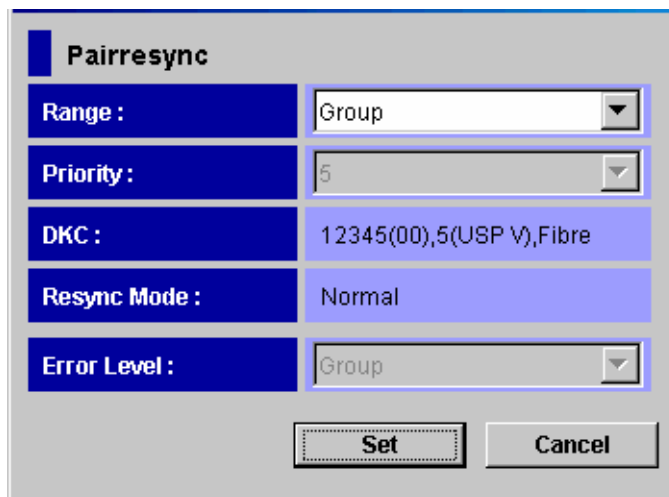


Figure 7-7 Pairresync Window

The Pairresync window displays the following:

- **Range:** Allows you to specify the restore range. The default is **LU** if two or more pairs in the same consistency group are selected. The default is **Group** if not. Note that the default setting is **Group** when the pair status is **HOLD** or **HLDE**, and in that case you cannot change the default setting.

- **LU:** Only the specified pair(s) will be restored.
- **Group:** All pairs in the same consistency group(s) as the selected pair(s) 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 **Preview** 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 **Preview** 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 **Range** is **Group**, you cannot change the **Priority** option.

- **DKC:** Indicates the followings about a remote storage system.

- The serial number with the LDKC number.
- The controller ID with the model name.
- The path type.

- **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. For information about the requirements to perform delta resync operation, see section Requirements for Performing Delta Resync Operation.
- **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 **Range** is **Group**, you cannot change the **Error Level** option.

- **Set:** Applies the settings to the **Preview** list in the Pair Operation window (Figure 4-2).

- **Cancel:** Discards the settings.



Releasing a Pair of Data Volumes (Pairsplit-S)

The Pairsplit-S window (Figure 7-8) allows you to release a pair of data volumes.

When one journal group uses two mirror IDs, the result of the operation differs as follows according to the status of the mirror of specified UR pair.

- When you specify a mirror of the UR pair in **HOLD** or **HLDE** status, only the UR pair of the specified mirror will be deleted.
- When you specify a mirror of the UR pair that is not in **HOLD** or **HLDE** status, UR pairs of both mirrors (including the mirror which you do not specify) will be deleted.

To release one or more pair(s), follow the procedure below. Not only the primary storage system administrator but also the secondary storage system administrator can perform this operation:

1. Ensure that the Storage Navigator main window is in Modify mode.
For detailed information about how to do this, please refer to *Storage Navigator User's Guide*.
2. Ensure that the Pair Operation window is displayed (see section Pair Operation Window).
3. In the tree, select a port  or a host group .
4. In the list, select and right-click the pair that you want to release.
5. From the pop-up menu, select **Pairsplit-S**.
The Pairsplit-S window appears (Figure 7-8).
6. In the Pairsplit-S window, select the desired options for the pair(s).
7. Select Set to close the Pairsplit-S window.
8. See the **Preview** list in the Pair Operation window 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**.
 - If you want to cancel a setting, select and right-click the setting and then select **Cancel**.
9. Select **Apply** to release pair(s).
Note: If an error occurs, the rightmost column of the **Preview** list displays the error code. To view detailed information about the error, right-click the error code and then select **Error Detail**. An error message appears and gives you detailed information about the error.
10. In the list of the **Pair Operations** window, verify that the pair(s) has been released successfully. If the pair has been released, the status of the pair is **SMPL**.

To monitor the progress of releasing pair(s), click **File, Refresh** on the menu bar of the Storage Navigator main window to update the information in the list, or use the Detailed Information window to monitor the detailed status of each pair.

Note: To restore a pair which was released from the secondary storage system, first release the pair from the primary storage system, and then restore the pair using the appropriate initial copy option.

Important: If a pair deletion is performed, it is possible that the deletion takes effect on the secondary data volume later than it takes effect on the primary data volume. Therefore, an error may occur if you reference the status of a pair or manipulate a pair immediately after the pair is deleted by CCI. If such an error occurs, please check the status of the secondary data volume and then retry referencing the pair status or manipulating the pair.

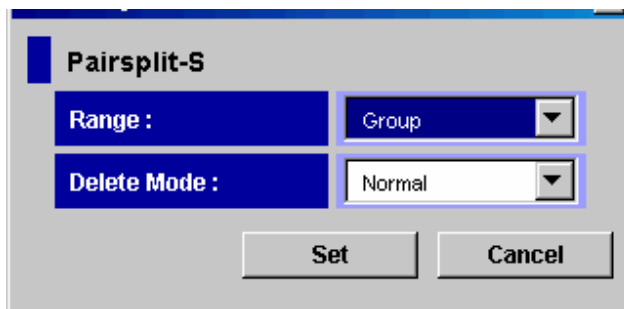


Figure 7-8 Pairsplit-S Window

The Pairsplit-S window displays the following:

- **Range:** Allows you to specify the release range. The default is **LU** if two or more pairs in the same consistency group are selected. The default is **Group** if not. Also, if the pair status is **SMPL**, **Range** is set to **LU**. If the pair status is **Deleting** or **Suspending**, **Range** is set to **Group**.

- **LU:** Only the specified pair(s) will be released.

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 frequency of host I/Os. The pair operations should be performed under light I/O load conditions.

- **Group:** All pairs in the same consistency group(s) as the selected pair(s) will be released.

Caution: Do not use this option when releasing pairs at the secondary storage system during disaster recovery.

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

- The **Preview** 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 **Preview** 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 release the pair(s) forcibly. When the status of the pair(s) to be released is **SMPL** or **Deleting**, the default setting is **Force**. Otherwise, the default setting is **Normal**:
 - **Force:** The pair(s) will forcibly be released even if the primary storage system is unable to communicate with the secondary storage system. This option may be used to free a host waiting for device-end from a primary storage system that cannot communicate with its the secondary storage system, thus allowing host operations to continue.
If you specify **Force** when the pair status is other than **SMPL**, **Range** will be set to **Group**.
 - **Normal:** The pair(s) will be released only if the primary storage system is able to change the pair status of the primary and secondary data volumes to SMPL.
- **Set:** Applies the settings to the **Preview** list in the Pair Operation window (Figure 4-2).
- **Cancel:** Discards the settings.

Creating a Point-in-Time Copy of Data Volumes

Universal Replicator enables you to make Point-in-Time (PiT) duplicates of groups of volumes. The term "Point-in-Time (PiT)" refers to an act of obtaining a copy of data immediately when you want to obtain it. The Universal Replicator can be used to create a PiT copy of data volumes belonging to a journal group.

To produce a PiT duplicate of data volumes belonging to a journal group:

1. Stop all write I/Os from hosts to primary data volumes in the journal group. Stop all host updates to all Universal Replicator primary data volumes in the group.
2. After all write I/Os to primary data volume have stopped, split the mirror (i.e., a pair of the master and restore journal group. See section Splitting a Mirror (Pairsplit-r: Suspending a copy operation)). When splitting it, use the Pairsplit-r window and then specify Group for Range and specify Flush for Suspend Mode.
3. When the status of all data volume pairs in the journal group has changed to **PSUS**, the duplicate set of volumes is complete. If desired, you can restart your business application at the primary site to resume business activities.

Recovering a Pinned Track

If a hardware failure occurs when the storage system is in operation, some data in the cache memory of the storage system may not be written to disk drives and may remain in the cache memory. Data that remains in cache memory due to such a reason is referred to as "pinned track".

Recovering a Pinned Track on a UR Data Volume

If a pinned track occurs on a Universal Replicator primary or secondary data volume, the primary storage system will suspend the data volume pair (SIM reference code = D41x, D42x, DB1x, DB2x). Use the following procedure to ensure full data integrity of the data volume pair while recovering the pinned track:

1. Connect to the primary storage system.
2. Release the data volume pair that contains the volume with the pinned track (see section Releasing a Pair of Data Volumes (Pairsplit-S)).
3. If the volume is offline (e.g., secondary data volume has pinned track), vary the volume online.

4. Perform your usual procedure for recovering data from a pinned track. Refer to the pinned track recovery procedures for your operating system, or contact your Hitachi representative for assistance in recovering the pinned track.
5. If the volume was previously offline (e.g., secondary data volume. See step 3), make sure to vary the volume offline again.
6. Use the **Paircreate** command in the Pair Operation window (see section Pair Operation Window) to recover the data volume pair. Make sure to select the **Entire** initial copy option in the Paircreate window.

Recovering a Pinned Track on a UR Master Data Volume

If a pinned track occurs on a Universal Replicator master or restore journal volume, the procedure for pinned track recovery is as follows:

1. Identify the journal group that contains the journal volume where a pinned track occurs. After that, split the mirror where the journal group belongs (see section Splitting a Mirror (Pairsplit-r: Suspending a copy operation)).
2. Delete the journal volume containing the pinned track from the journal group (see section Deleting Journal Volumes from a Journal Group).
3. Format the journal volume by using the LDEV format function of VLL. For detailed instructions on how to format the volume, please refer to the *Virtual LVI/LUN and Volume Shredder User's Guide*.
4. Register the formatted volume (or another volume that can be used instead of the formatted volume) in the journal group (see section Registering Journal Volumes in a Journal Group).
 - If you do not need to resume copying, the procedure is now finished.
 - If you want to resume copying, please continue to restore the mirror to which the journal group belongs (see section Restoring a Mirror (Pairresync: Resuming a copy operation))

Usage Monitor Operations

This chapter describes using the Usage Monitor window to follow the collection and display of performance monitoring statistics.

- Reviewing the Usage Monitor Window
- Starting and Stopping Usage Monitoring
- Displaying the Usage Monitor Graph
- Saving Monitoring Data in Text Files
- Viewing Pair Operation History
- Saving Operation History to a Text File

Reviewing the Usage Monitor Window

The Usage Monitoring window enables you to collect I/O statistics for all volumes to be monitored on the connected storage system. The I/O statistics that are displayed in the window are collected by Performance monitor. The Usage Monitoring is able to display the data of remote copy operations. At each data sampling time interval that is specified from Performance Monitor, the usage monitoring window collects 1,440 pieces of data for each volume and up to 1,440 pieces of data can be collected for entire volumes to be monitored.

Starting and Stopping Usage Monitoring

Starting and stopping usage monitoring can be set using Performance Monitor. For detail about starting and stopping usage monitoring, please refer to *Performance Manager User's Guide*.

Displaying the Usage Monitor Graph

When usage monitoring is running, the Usage Monitor window can display user-selected remote copy I/O statistics in real time.

The usage monitor graph plots the user-selected I/O statistics (up to 65 data points) on an x-y graph. The x-axis displays time. The y-axis displays the number of I/Os during the last sampling period. The legend (right side of the graph) indicates the data being displayed. A value on the y-axis varies according to the maximum value of the statistical data that is displaying. If the value on the y-axis exceeds 10,000,000, the value is displayed in exponential notation (e.g., $1E7 = 1 \times 10^7 = 10,000,000$; $2E8 = 2 \times 10^8 = 200,000,000$).

To display the usage monitor graph:

1. Make sure that usage monitoring is running (Monitoring Switch = Enable). The usage monitor graph can only be displayed when monitoring is on.
2. Right-click the graph area of the Usage Monitor window, and select Display Item to open the Display Item window (see Figure 8-1).
3. Select an appropriate radio button in **Select Volume**, following the instructions below:
 - If you want to display I/O statistics for all LDEVs to be monitored in the storage system, select **ALL Volumes**.
 - If you want to display I/O statistics for a specific journal group, select **JNL Group**, and then enter a journal group number (00-FF).
Note: If you specify an unused journal group, the graph is not displayed.
 - If you want to display I/O statistics for a specific LU, select **Volume**, and then specify the desired LU by selecting a port (CL1-A to CLG-R) and entering the **GID** (00-FE) and **LUN** (00-7FF).
Note: If you specify the un-mounted volume, the graph is not displayed.
Important: When a graph displays, the LDKC number, the CU number and the LDEV number appear above the graph. If the 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*.
4. In the **Monitor Data** box, select the I/O statistics data that you want to display on the graph. You must select at least one box. Table 8-1 describes the I/O statistics data.
5. Select **Set** to close the Display Item window. The Usage Monitor window now displays a graph showing the selected I/O statistics data for the selected LU(s).

To enlarge the displayed graph, right-click the graph, and select the **Large Size** command. To return the graph to normal size, right-click the graph, and select the **Normal Size** command.

Note: The length of a host group name or an iSCSI target name at the upper portion of the graph is limited to a maximum of eight characters when the graph is in normal size. To display a full host group name, please enlarge the graph.

Important: To stop displaying the usage monitor graph, right-click the graph, and select **Close**. To stop displaying all graphs, select **Close All**. The usage monitor graph closes automatically in the following cases:

- When you select another tab,
- When you select another program product,
- When you exit the Storage Navigator software,

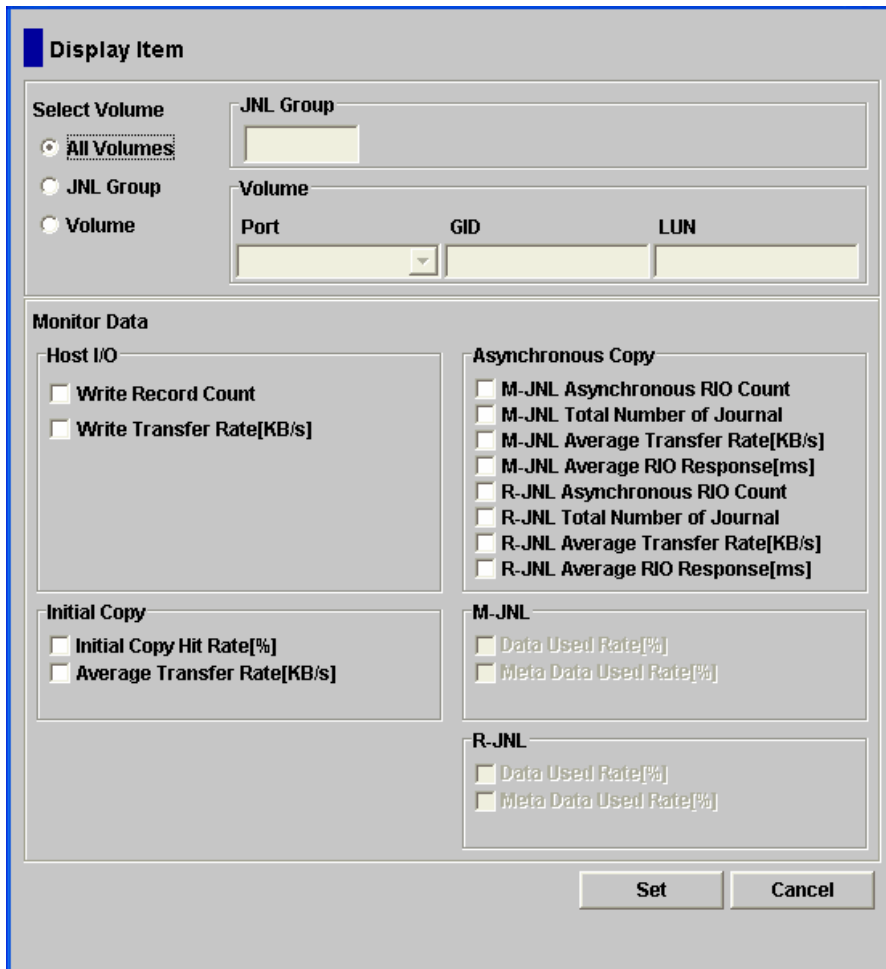


Figure 8-1 Display Item Window

Table 8-1 Remote Copy I/O Statistics

Statistic	Description
Host I/O	
Write Record Count	The number of write I/Os per second
Write Transfer Rate	The amount of data that are written per second. The unit is kilobytes per second.
Initial Copy	
Initial Copy Hit Rate	The initial copy hit rate. The unit is percent.
Average Transfer Rate	The average transfer rate for initial copy operations. The unit is kilobytes per second.
Asynchronous Copy	
M-JNL Asynchronous RIO count	The number of asynchronous remote I/Os per second at the primary storage system..
M-JNL Total Number of Journal	The number of journals at the primary storage system.
M-JNL Average Transfer Rate	The average transfer rate for journals in the primary storage system. The unit is kilobytes per second.
M-JNL Average RIO Response	The remote I/O process time on the primary storage system. The unit is milliseconds.
R-JNL Asynchronous RIO count	The number of asynchronous remote I/Os per second at the secondary storage system..
R-JNL Total Number of Journal	The number of journals at the secondary storage system.
R-JNL Average Transfer Rate	The average transfer rate for journals in the secondary storage system. The unit is kilobytes per second.
R-JNL Average RIO Response	The remote I/O process time on the secondary storage system. The unit is milliseconds.
M-JNL	
Data Used Rate	Data usage rate for master journals. The unit is percent.
Meta Data Used Rate	Metadata usage rate for master journals. The unit is percent.
R-JNL	
Data Used Rate	Data usage rate for restore journals. The unit is percent.
Meta Data Used Rate	Metadata usage rate for restore journals. The unit is percent.

Saving Monitoring Data in Text Files

To save monitoring data in text files, use the Export Tool of Performance Monitor. For information and instructions on using the Export Tool, please refer to the *Performance Manager User's Guide*.

Viewing Pair Operation History

The **History** window (see Figure 4-10) enables you to view history of operations for data volume pairs. For example, you can check the date and time when data volume pairs are created or released. You can also save operation history in a text file.

This chapter explains how to do the following:

Display the History window (see section History Window).

Save operation history in a text file (see the next section).

Saving Operation History to a Text File

The export function enables you to save operation history into a CSV file.

Note: You can use the export function only while the client file access is enabled. When the client file access is not enabled, the **Export** button is not displayed on the History window. For the further information of the client file access settings, please refer to the *Storage Navigator User's Guide*.

Note: If the number of operations exceeds approximately 10,000, the exported text file will become too large to be saved in a floppy disk. If this happens, please save the exported text file in a large-capacity disk.

The following is an example of a exported text file.

```
2007/02/22 09:57:54,Paircreate Complete,01,01,00:02:01,00:02:04,--,000:01:08
2007/02/22 09:56:46,Paircreate Start,01,01,00:02:01,00:02:04,--,
2007/02/22 09:56:46,Pair definition,01,01,00:02:01,00:02:04,--,
```

(1) (2) (3) (4) (5) (6) (7) (8)

- (1) Data and time when the operation finished.
- (2) Operation.
- (3) Journal group number
- (4) Mirror ID
- (5) LDKEC number, CU number and LDEV number of the volume in the local storage system (LDKC : CU : LDEV).
- (6) LDKEC number, CU number and LDEV number of the volume in the remote storage system (LDKC : CU : LDEV).
- (7) Hyphens(--) are displayed.
- (8) Time taken for copying (from the start to the end of copying).

This is displayed only when the operation (refer to (2)) is Paircreate Complete or Pairresync Complete

Figure 8-2 Example of a Exported File

To save operation history in a file:

1. In the History window (Figure 4-10), select **Export**.
A message appears and asks whether you want to save operation history in a file.
2. Click **Yes** to close the message.
A window appears and prompts you to save a file.
3. Enter the name of the file and then click **Save**.
A message appears when UR finishes saving the file.
4. Click **OK** to close the message.

Performing Disaster Recovery Operations

This chapter describes how to determine the best approach for disaster recovery. The user should anticipate and plan for all types of failures and disasters.

- Preparing for Disaster Recovery
- Recovering from a Disaster in a 3DC Cascading Configuration
- Recovering from Primary Site Disaster in 3DC Multi-Target Configuration

Preparing for Disaster Recovery

The type of disaster and the status of the Universal Replicator volume pairs will determine the best approach for disaster recovery. Unfortunately, some disasters are not so “orderly” and involve intermittent or gradual failures occurring over a longer period of time. The major steps in preparing for disaster recovery are:

1. Identify the journal groups and data volumes that contain important files and data for disaster recovery.
2. Install the Storage Navigator PC and Universal Replicator hardware and software, and establish Universal Replicator operations for the journal groups and data volumes identified in step (1).
3. Establish file and database recovery procedures. These procedures should already be established for recovering data volumes that become inaccessible due to some failure.
4. Install and configure host failover software between the primary and secondary sites.

File and Database Recovery Procedures

When the primary or secondary storage system suspends a Universal Replicator pair due to a disaster, the secondary data volume may contain in-process data. A data set could be open, or transactions may not have completed. Therefore, you need to establish file recovery procedures. These procedures should be the same as those used for recovering data volume that becomes inaccessible due to control unit failure.

Universal Replicator does not provide any procedure for detecting and retrieving lost updates. To detect and recreate lost updates, you must check other current information (e.g., database log file) that was active at the primary site when the disaster occurred. Since this detection/retrieval process can take a while, your disaster recovery scenario should be designed so that detection/retrieval of lost updates is performed after the application has been started at the secondary site.

You should prepare for file and database recovery by using:

Files for file recovery (e.g., database log files which have been verified as current).

Important: Remote copy and disaster recovery procedures are inherently complex. Consult your Hitachi account team on recovery procedures.

Note: See section Recovering a Pinned Track on a UR Data Volume for information on recovering a pinned track on a Universal Replicator volume.

Switching Operations to the Secondary Site

If a disaster or failure occurs at the primary site, the first disaster recovery activity is to switch your operations to the secondary site. Executing the **horctakeover** command of CCI (Command Control Interface) on Universal Replicator pairs enables your business tasks to be taken over to the secondary site, using secondary data volumes in the secondary site. The **horctakeover** command checks the pair status of secondary data volumes and splits journal groups that will become a Universal Replicator pair. This ensures 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. If the **horctakeover** command executes successfully, your business tasks will be taken over to (or your business application will be started at) the secondary site, using secondary data volumes. For detailed information about how to use CCI, please refer to *Hitachi Command Control Interface User and Reference Guide*.

Transferring Operations Back to the Primary Site

Once the disaster recovery procedure is finished and your business applications are running at the secondary site, the next activity is to restore the primary site and make arrangements for copying data from the secondary site back to the primary site. The following procedure explains how to use CCI to copy data from the secondary site to the primary site:

1. Restore the primary storage system and remote copy connections, bring up the host server(s) at the primary site, and make sure that all Universal Replicator components are fully operational.
2. At the primary site, locate primary data volumes whose pair status is COPY or PAIR, and then locate corresponding secondary data volumes whose pair status is SSWS. If such volume pairs are found, issue a request for splitting the pairs to the primary data volumes.
3. At the primary site, locate primary data volumes whose pair status is not SMPL, and then locate corresponding secondary data volumes whose pair status is SMPL. If such volume pairs are found, issue a request for releasing the pairs to the primary data volumes.
4. At the primary site, locate data volume pairs whose pair status is SMPL, and then issue a request for releasing the pairs to the secondary data volumes.
5. Execute the **pairresync -swaps** command on secondary data volumes whose pair status is SSWS (**pairresync** is the CCI command for resynchronizing pair and **-swaps** is a swap option). This reverses primary data volumes and secondary data volumes to resynchronize pairs.

Caution: When you execute the **pairresync -swaps** command, you can use the **-d** option to specify a data volume. However, if the restore journal group where the data volume belongs is in **Active**, **Halting** or **Stopping** status, the **pairresync -swaps** command will be rejected. For detailed information about journal group status, see section Displaying Detailed Information about a Journal Group.

6. Create pairs, specifying secondary data volumes whose pair status is SMPL as primary data volumes. This creates pairs in which primary data volumes and secondary data volumes are reversed.
7. Verify that pair status of all secondary data volumes (which were originally primary data volumes) changes from COPY to PAIR. If the pair status is changed to PAIR, initial copy operations are finished and consistency is maintained.

The above procedure enables copying of data from the secondary site to the primary site. Data in the secondary site will be reflected on the primary site.

For detailed information about how to use CCI, please refer to *Hitachi Command Control Interface User and Reference Guide*.

For detailed information about status of data volumes, refer to section Pair Status.

Resuming Normal Operations at the Primary Site

Once the Universal Replicator volume pairs have been established in the reverse direction, you are ready to resume normal operations at the primary site. The following procedure explains how to resume normal operations at the primary site by using CCI. Remember that the Universal Replicator terminology is now reversed: the original primary data volumes are now secondary data volumes, and the original secondary data volumes are now primary data volumes.

1. At the primary and secondary sites, make sure that all Universal Replicator components are fully operational and are free from failures.
2. Make sure that pair status of primary and secondary data volumes in all Universal Replicator pairs is "PAIR". This indicates that the Universal Replicator initial copy operations are complete and consistency is maintained.
3. Stop the applications at the secondary site.
4. Issue a request for splitting pairs to master journal groups (which were originally restore journal groups); please use the Flush option when splitting pairs. If an error occurs when splitting pairs, please remove the error cause and go back to step 1 after resuming your business task at the secondary site.

5. If no error occurs in step 4, wait until suspension finishes. After suspension finishes, check whether there is a secondary data volume (which is originally a primary data volume) whose pair status is other than PSUS. If such a pair exists, please remove the error cause and go back to step 1 after resuming your business task at the secondary site.
6. If there is no secondary data volume (which is originally a primary data volume) whose pair status is other than PSUS, data in primary data volumes are the same as data in secondary data volumes, and the secondary data volume (which are originally primary data volumes) are usable. Please resume applications at the primary site.
7. Execute the **pairresync -swaps** command (**pairresync** is the CCI command for resynchronizing pair and **-swaps** is a swap option). This reverses primary data volumes and secondary data volumes to resynchronize pairs.

Caution: When you execute the **pairresync -swaps** command, you can use the **-d** option to specify a data volume. However, if the restore journal group where the data volume belongs is in **Active**, **Halting** or **Stopping** status, the **pairresync -swaps** command will be rejected. For detailed information about journal group status, see section Displaying Detailed Information about a Journal Group.

For detailed information about how to use CCI, please refer to *Hitachi Command Control Interface User and Reference Guide*.

For detailed information about status of data volumes, refer to section Pair Status.

Recovering from a Disaster in a 3DC Cascading Configuration

If a disaster or failure occurs in the primary site in a 3DC cascading configuration, you can use TrueCopy Synchronous secondary data volumes in the intermediate site to enable the intermediate site to take over your business tasks. If a disaster or failure occurs in the primary site, first you must transfer your business tasks from the primary site to the intermediate site.

To transfer your business tasks to the secondary site:

1. Check consistency of the TrueCopy Synchronous S-VOL in the intermediate site.
2. Use Command Control Interface (CCI) to execute the horctakeover command onto the TrueCopy Synchronous pair that exists between the primary site and the intermediate site.
3. Wait until the TrueCopy Synchronous S-VOL in the intermediate site becomes suspended (in SSWS status) or changes into a TrueCopy Synchronous P-VOL.

After the S-VOL becomes suspended or changes into P-VOL, you can use the volume to resume your business tasks.

For detailed information about how to use Command Control Interface (CCI), please refer to *Hitachi Command Control Interface User and Reference Guide*.

When a disaster or failure occurs, the above procedure enables you to resume your business tasks by using the TrueCopy Synchronous volume in the intermediate site. The TrueCopy Synchronous volume in the intermediate site will contain the same data that are stored in the TrueCopy Synchronous P-VOL in the primary site.

For detailed information about pair status of data volumes, see section Pair Status.

Recovering from Primary Site Disaster in 3DC Multi-Target Configuration

If a disaster or failure occurs only in the primary site in 3DC multi-target system, you can use secondary volume in the TrueCopy Synchronous secondary site to enable the secondary site to take over your business tasks. After transferring your business tasks to the secondary site, you need to remove the failure from the primary site, so that you will be able to transfer your business tasks back to the primary site.

To transfer your business tasks to the TrueCopy Synchronous secondary site, follow the procedure below. Command Control Interface (CCI) is used in this procedure:

1. Use CCI to release the UR pair.
2. Use CCI to execute the horctakeover command on the TrueCopy Synchronous pair.
3. Use the TrueCopy Synchronous secondary volume to resume your business tasks.
4. Check the execution result of the horctakeover command, and then do either of the following:
 - a. If the primary/secondary relationship of the TrueCopy Synchronous pair is not reversed (i.e., the copy direction is not reversed as illustrated in the upper-right picture in Figure 9-1), use CCI to create a UR pair. The primary data volume of the UR pair must be the TrueCopy Synchronous secondary volume. The secondary data volume of the UR pair must be the secondary data volume in the UR secondary site.
 - b. If the primary/secondary relationship of the TrueCopy Synchronous pair is reversed (i.e., the copy direction is reversed as illustrated in the lower-right picture in Figure 9-1), recreate the UR pair that were released at step 1 of this procedure. When you finish recreating the UR pair, the system configuration changes from 3DC multi-target configuration to 3DC cascading configuration, and the former TrueCopy Synchronous secondary volume changes to the primary volume.

Notes:

- If you followed step a above, please remove failures in the primary site and other locations, and then change the system into 3DC multi-target configuration (see section Changing to 3DC Multi-target Configuration after Recovering from Primary Site Failures), and follow the instructions in section Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC multi-target configuration), in order to transfer business tasks back to the primary site.
- If you followed step b above, please remove failures in the primary site and other locations, and then follow the instructions in section Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC Cascading Configuration), in order to transfer business tasks back to the primary site.




Figure 9-1 Recovery from Disaster in the Primary Site

Recovering from Primary Site Failures (When Delta Resync Operation is Performed)

If a disaster or failure occurs only in the primary site in 3DC multi-target system, you can use secondary volume in the TrueCopy Synchronous secondary site to enable the secondary site to take over your business tasks. After transferring your business tasks to the secondary site, you need to remove the failure from the primary site, so that you will be able to transfer your business tasks back to the primary site.

To transfer your business tasks to the TrueCopy Synchronous secondary site, follow the procedure below. Command Control Interface (CCI) is used in this procedure:

1. Use CCI to execute the horctakeover command on the TrueCopy Synchronous pair.
2. Use the TrueCopy Synchronous S-VOL to resume your business tasks.

3. Check the execution result of the horctakeover command.
4. In the either cases when the primary/secondary relationship of the TrueCopy Synchronous pair is not reversed (i.e., the copy direction is not reversed as illustrated in the upper-right picture in Figure 9-2), or when the primary/secondary relationship of the TrueCopy Synchronous pair is reversed (i.e., the copy direction is reversed as illustrated in the lower-right picture in Figure 9-2), use CCI to perform delta resync operation on the volume in the TrueCopy Synchronous secondary site. When you perform delta resync operation on the volume in the TrueCopy Synchronous secondary site, the status of the data volumes of UR pair changes as explained in the following table.

Table 9-1 Changes of UR Pair Status by Delta Resync Operation (When Recovering from Failures in the Primary Site)

UR Pair	Pair Status before Delta Resync Operation		Pair Status after Delta Resync Operation	
	Primary Data Volume	Secondary Data Volume	Primary Data Volume	Secondary Data Volume
UR pair between TrueCopy Synchronous primary site and UR secondary site	PAIR, PSUS, or PSUE	PAIR, PSUS, or PSUE	HOLD	HOLD
UR pair between TrueCopy Synchronous secondary site and UR secondary site	HOLD	HOLD	PAIR or COPY	PAIR or COPY

Note: In order to transfer business tasks back to the primary site, please remove failures in the primary site and other locations, and then follow the instructions in section Transferring Business Tasks from TrueCopy Secondary to Primary Site (in 3DC Cascading Configuration). For the latest information on the availability of the delta resync function, please contact your Hitachi Data Systems account.




Figure 9-2 Recovery from Disaster in the Primary Site (When Delta Resync Operation is Performed)



Troubleshooting

Table 9-2 provides general troubleshooting information for Universal Replicator. Table 9-3 provides troubleshooting information for logical paths between storage systems. Table 9-4 and Table 9-5 provide troubleshooting information for suspended Universal Replicator pairs.

See section Checking Universal Replicator Error Codes for a description of the Universal Replicator error messages displayed on the Storage Navigator PC. See Checking Service Information Messages (SIMs) for a brief description of the USP V service information messages (SIMs).

Table 9-2 General Troubleshooting Information for Universal Replicator

Error	Corrective Action
The Storage Navigator computer hangs, or Universal Replicator operations do not function properly.	Make sure that the problem is not being caused by the computer or Ethernet hardware or software, and restart the computer. Restarting the Storage Navigator computer does not affect Universal Replicator operations in progress. Make sure that all Universal Replicator requirements and restrictions are met. See Preparing for Universal Replicator Operations. Make sure that the primary storage system, the secondary storage system and remote copy connections are powered on, and NVS and cache are fully operational. Check all the entered values and parameters to make sure that you entered the correct information on the Storage Navigator computer (e.g., secondary storage system S/N, path parameters, primary data volume IDs, and secondary data volume IDs). If you are using Performance Monitor, please refrain from using it.
An initiator or RCP channel-enable LED indicator (on the USP V control window) is off or flashing.	Please call the Hitachi Support Center for assistance.
Status of pairs and/or logical paths is not displayed correctly.	Check whether you selected correct items or options in Universal Replicator windows, such as the DKC Operation window, the Pair Operation window, and the Display Filter window.
A Universal Replicator error message is displayed on the Storage Navigator computer.	Remove the error cause, and then retry the Universal Replicator operation.
The status of a logical path is not normal.	Check the path status in the DKC Status window, and see Table 9-3.
The pair status is <i>suspended</i> .	Check the pair status in the Detailed Information window, and see Table 9-4.

Error	Corrective Action
Paircreate or pairresync operation resulted in a timeout error.	<p>If the timeout error was caused by a hardware failure, a SIM will be generated. If this occurs, call the Hitachi Support Center, and retry Universal Replicator operations after the problem is solved.</p> <p>If no SIM was generated, wait for a while (5 or 6 minutes), then check the pair status of the pair(s) being created or resumed. If the pair status changed correctly, the failed operation completed after the timeout message was issued. If the pair status did not change as expected, heavy workload might have prevented the Universal Replicator operation from being completed. In this case, retry the Universal Replicator operation when the storage system has a lighter workload.</p> <p>If a communication error occurs between the Storage Navigator computer and the SVP, please refer to the <i>Storage Navigator User's Guide</i> for instructions.</p>
There is a pinned track on a Universal Replicator volume.	See Recovering a Pinned Track on a UR Data Volume for instructions.

Table 9-3 Troubleshooting Problems with Status of Logical Paths

Path Status	Description	Corrective Action
Initialization Failed	The link initialization procedure for the remote storage system failed.	<p>Make sure that the local and remote storage systems are connected correctly via cables.</p> <p>Make sure that you entered the correct serial number (S/N) of the remote storage system and the correct path parameters (e.g., local storage system port, remote storage system port, controller ID).</p> <p>Make sure that the topology settings of the local and remote storage system ports are correct.</p>
Communication Time Out	Communication between the local and remote storage systems timed out.	<p>Make sure the remote storage system is powered on, and NVS and cache are fully functional.</p> <p>Make sure that the network relay devices (cables, connectors, switches, ESCDs, extender devices, communication lines, and all other devices connected to the extenders) are properly configured and functional.</p> <p>Delete the failed path. You may need to change the minimum paths setting or delete the remote storage system in order to delete the path. After deletion finishes, add the path or remote storage system.</p>
Resource Shortage (Port) or, Resource Shortage (Pair-Port)	The storage system rejected the request for establishing a logical path because all resources in the local and remote storage systems are being used for other connections.	<p>Delete the failed path, and also delete all paths and remote storage systems not currently in use. The local storage system can be connected to up to 64 remote storage systems with up to eight paths to each remote storage system.</p> <p>Make sure that all ports on the local and remote storage systems are properly configured: ordinary target ports or RCU target ports for hosts, RCU target ports for remote storage systems, initiator ports for local storage systems.</p> <p>If necessary, connect to the remote storage system to delete paths or remote storage systems and reconfigure ports, then reconnect to the local storage system. After that, add the path and remote storage system again.</p>

Path Status	Description	Corrective Action
Serial Number Mismatch	The serial number (S/N) of the remote storage system does not match the specified serial number (S/N).	Make sure that you entered the correct remote storage system S/N and path parameters (e.g., local storage system port, remote storage system port, and controller ID). After that, delete the failed path. You may need to change the minimum paths setting or delete the remote storage system in order to delete the path. After deletion finishes, add the path and the remote storage system. Make sure that you entered the correct controller ID, and also make sure that the topology settings of the local and remote storage system ports are correct.
Invalid Port	The specified port is not configured as an initiator port, or this path already exists.	make sure that the topology settings of the local and remote storage system ports are correct. Make sure that you entered the correct remote storage system S/N and path parameters (e.g., local storage system port, remote storage system port, controller ID). After that, delete the failed path. You may need to change the minimum paths setting or delete the remote storage system in order to delete the path. After deletion finishes, add the path and the remote storage system again.
Port number invalid	The specified port number is invalid. The cable is not connected to the specified port.	Delete the failed path. Check the specified port number or the cable connection and then retry.
Communication error	A timeout error has occurred in the Fiber path between the local and remote storage systems.	Delete the failed path and retry.
Pair-Port Number Mismatch	The specified port in the remote storage system is physically disconnected from the local storage system.	Make sure that you specified the correct remote storage system port. Correct the port number if necessary. Make sure that the local and the remote storage systems are connected via cables to each other. If the storage systems are not connected via cables, connect them via cables and then retry the operation. Make sure that the topology settings of ports on the local and remote storage systems are correct.
Pair-Port Type Mismatch	The specified remote storage system port is not configured as an RCU target port.	Make sure that the remote storage system port is configured as an RCU target port.
Communication Failed	The local storage system is connected to the remote storage system successfully, but logical communication timeout occurred.	Make sure that the remote storage system port and network relay devices (e.g., cables and switches) are properly configured and functioning.
Logical Blockade	A failure occurs at a local storage system port.	Repair the port status on the local storage system.
	A failure occurs at a remote storage system port.	Repair the port status on the remote storage system. After that, delete the path by using the Delete Path or Delete DKC command, and then add the path and the remote storage system again by using the Add Path or Add DKC command.
	A failure occurs in the path relay equipment.	Repair the path relay equipment. After that, delete the path by using the Delete Path or Delete DKC command, and then add the path and the remote storage system again by using the Add Path or Add DKC command.

Path Status	Description	Corrective Action
	The connection cable is physically broken.	Replace the broken cable. After that, delete the path by using the Delete Path or Delete DKC command, and then add the path and the remote storage system again by using the Add Path or Add DKC command.
Program Error	Program error is detected.	Delete the path by using the Delete Path or Delete DKC command, and then add the path and the remote storage system again by using the Add Path or Add DKC command.
<blank>	This path was not established.	Delete the failed path. You may need to change the minimum paths setting or delete the remote storage system in order to delete the path. After deletion finishes, add the path and the remote storage system by using the Add Path or Add DKC command.

Table 9-4 Troubleshooting Suspended Universal Replicator 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 secondary storage system. The suspend type for the secondary data volume is <i>by MCU</i> .	Clear the error condition at the secondary storage system or secondary data volume. If you need to access the secondary data volume, release the pair from the primary storage system. If data in the secondary data volume has been changed, release the pair from the primary storage system and then re-create the pair by using the Paircreate window. If data in the secondary 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 secondary storage system or detected an I/O error during update copy. In this case, the suspend type for the secondary data volume is usually <i>by MCU</i> .	Check the path status on the DKC Status window (see Table 9-3). Clear any error conditions at the secondary storage system and the secondary data volume. If you need to access the secondary data volume, release the pair from the secondary storage system. If data in the secondary data volume has been changed, release the pair from the primary storage system and then re-create the pair by using the Paircreate window. If data in the secondary data volume has not been changed, restore the pair from the primary storage system.
PSUE, MCU IMPL	Primary data volume, secondary 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 secondary data volume.

Suspend Type	Applies to	Description	Corrective Action
PSUE, Initial Copy Failed	Primary data volume, secondary data volume	The primary storage system suspended this pair during the initial copy operation. The data on the secondary data volume is not identical to the data on the primary data volume.	Release the pair from the primary storage system. Clear all error conditions at the primary storage system, primary data volume, secondary storage system, and secondary data volume. Restart the initial copy operation by using the Paircreate window.
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.

Table 9-5 provides troubleshooting instructions for Universal Replicator pair suspension. Hardware failures which affect the cache memory and the shared memory of the primary or secondary storage system may also cause the Universal Replicator pairs to be suspended.

Table 9-5 Resolving Universal Replicator Pair Suspension

Classification	Causes of Suspension	SIM	Recovery procedure
Primary storage system hardware or secondary storage system hardware	<p>Hardware redundancy has been lost due to some blockade condition. As a result, one of the following could not complete: primary-secondary storage system communication, journal creation, copy operation, restore operation, staging process, or de-staging process.</p> <p>Journals cannot be retained because some portion of the cache memory or the shared memory has been blocked due to hardware failure.</p> <p>The primary storage system failed to create and transfer journals due to unrecoverable hardware failure.</p> <p>The secondary storage system failed to receive and restore journals due to unrecoverable hardware failure.</p> <p>The drive parity group was in correction-access status while the Universal Replicator pair was in COPY status.</p>	DC0x DC1x DC2x	<p>According to SIM, remove the hardware blockade or failure.</p> <p>Restore the failed volume pairs (Pairresync).</p> <p>If a failure occurs during execution of the CCI horctakeover command, secondary volumes in SSWS pair status may remain in the master journal group. If these volumes remain, execute the pairresync -swaps command on the secondary volumes whose pair status is SSWS (pairresync is the CCI command for resynchronizing pair and -swaps is a swap option). This operation changes all volumes in the master journal group to primary volumes. After this operation, restore the volume pairs (Pairresync).</p>
Communication between the primary and secondary storage systems	<p>Communication between the storage systems failed because the secondary storage system or network relay devices were not running.</p> <p>Journal volumes remained full even after the timeout period elapsed.</p>	DC0x DC1x	<p>Remove the failure from the primary and secondary storage systems or the network relay devices.</p> <p>If necessary, increase resources as needed (e.g., the amount of cache, the number of paths between primary and secondary storage systems, the parity groups for journal volumes, etc.).</p> <p>Restore the failed pairs (Pairresync).</p>

Classification	Causes of Suspension	SIM	Recovery procedure
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	Release failed pairs (pairsplit-S). If necessary, increase resources as needed (e.g., the amount of cache, the number of paths between primary and secondary 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.

Checking Universal Replicator Error Codes

Storage Navigator computers display an error message when an error occurs during Universal Replicator operations. The error message describes the error and displays an error code consisting of four digits. The error message may also include a USP V SVP error code. If you need to call the Hitachi Support Center for assistance, please report the error code(s). Please refer to *Storage Navigator Messages* for a list of error codes displayed on the Storage Navigator computers.

Checking Service Information Messages (SIMs)

The USP V storage system generates SIMs (service information messages) when it is necessary to notify the user of possible service requirements for the storage system. SIMs can be generated by the channel and storage path microprocessors of the storage system and by the service processor (SVP) of the storage system. The SVP reports all SIMs related to Universal Replicator processing. Each time an SIM is generated, the amber **Message** LED on the USP V control window (under the **Ready** and **Alarm** LEDs) turns on as an additional alert for the user.

SIMs are classified into four categories according to severity for reporting and logging purposes: *service*, *moderate*, *serious*, or *acute*. All SIMs are recorded on the USP V SVP and displayed on Storage Navigator computers. The user can find SIMs in the **Status** window of Storage Navigator. SIMs generated by the primary storage system will include the device ID of the primary data volume (byte 13), and SIMs generated by the secondary storage system will include the device ID of the secondary data volume (byte 13).

Note: For further information on SIM reporting, please contact your Hitachi representative or the Hitachi Support Center.

If SNMP is installed and operational for the USP V storage system, each SIM will result in an SNMP trap being sent to the appropriate host(s). For further information on SNMP operations, please refer to the *Storage Navigator User's Guide*, or call the Hitachi Support Center for assistance.

Figure 9-3 shows a typical 32-byte SIM from the Hitachi USP V storage system. The six-digit reference code (RC) consists of bytes 22, 23, and 13. The RC identifies the possible error and determines the severity. The SIM type (byte 28) indicates the component which experienced the error. When the value of byte 22 is 21, the SIM is a control unit SIM. When the value of byte 22 is Dx (where x is an arbitrary character), the SIM is a device SIM.

Table 9-6 illustrates severity of SIMs. The table also shows SVP log files.




Figure 9-3 Typical SIM Showing Reference Code and SIM Type

Table 9-6 Typical SIM Showing Reference Code and SIM Type

Reference Code		Severity	Description	SIM Generated by	SVP Log File
Byte 22	Byte 23				
21	80	Moderate	Logical path is blocked due to a failure.	Primary/secondary storage system	SIM Log
21	81	Service	Logical path is restored.	Primary/secondary storage system	SSB Log
21	82	Moderate	Line failure is reported by the extender.	Primary/secondary storage system	SIM Log
DC	0X	Serious	A volume being used by primary data volume has been suspended. Path recovery is impossible.	Primary storage system	SIM Log
DC	1X	Serious	A volume being used by primary data volume has been suspended. A failure has been detected in the primary storage system.	Primary storage system	SIM Log
DC	2X	Serious	A volume being used by primary data volume has been suspended. A failure has been detected in the secondary storage system.	Primary storage system	SIM Log
DC	4X	Serious	A volume being used by primary data volume has been suspended. A data volume pair has been suspended at the secondary storage system.	Primary storage system	SIM Log
DC	5X	Serious	A volume being used by primary data volume has been suspended. A data volume pair has been released at the secondary storage system.	Primary storage system	SIM Log
DC	6X	Serious	A volume being used by secondary data volume has been suspended. Path recovery is impossible.	Primary storage system	SIM Log
DC	7X	Serious	A volume being used by secondary data volume has been suspended. A failure has been detected in the secondary storage system.	Secondary storage system	SIM Log

Calling the Hitachi Data Systems Support Center

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any error messages displayed on the host system(s).
- The exact content of any error messages displayed by Storage Navigator.
- The Storage Navigator configuration information (use the FD Dump Tool).
- The service information messages (SIMs), including reference codes and severity levels, displayed by Storage Navigator.

The Hitachi Data Systems customer support staff is available 24 hours/day, seven days a week. If you need technical support, please call:

- United States: (800) 446-0744
- Outside the United States: (858) 547-4526



Acronyms and Abbreviations

CCI	Command Control Interface
CU	control unit
ESCON [®]	Enterprise System Connection (IBM trademark for optical channels)
HBA	host bus adapter
HRX	Hitachi RapidXchange
I/O	input/output
IBM	International Business Machines Corporation
LAN	local-area network
LCP	local control port
LED	light-emitting diode
LUN	logical unit (also called device emulation or device type)
Mb/s	megabits per second
MCU	main control unit
NVS	nonvolatile storage
PC	personal computer system
PCB	printed circuit board
PSUE	pair suspended-error
PSUS	pair suspended-split
P-VOL	primary volume
RAID	redundant array of independent disks
RAID-1/-5	specific RAID architectures
RCP	remote control port
RCU	remote control unit
RIO	remote IO
SIM	service information message
SMPL	simplex
S/N	serial number (also abbreviated as s#)
SSID	storage storage system identification
S-VOL	secondary volume
SVP	service processor
sync	synchronous
s#	serial number (also abbreviated as S/N)

VOL volume

Index

C

components, 2-5
control units (CUs), 2-8

D

disaster recovery
file/DB recovery, 9-2

P

Performance Monitor, 3-45
pinned track, 7-28, 7-29
point-in-time copy, 7-28
power off/on requirements, 5-21

R

read and write I/O operations, 2-20

S

Support Center, Hitachi Data Systems, xviii, 9

