

THEORY OF OPERATION SECTION

Contents

THEORY01-10	1. RAID Architecture Overview
THEORY01-10	1.1 Outline of RAID Systems
THEORY01-30	1.2 Comparison of RAID levels
THEORY02-10	2. Hardware Specifications
THEORY02-10	2.1 General
THEORY02-20	2.1.1 Features
THEORY02-40	2.2 Architecture
THEORY02-40	2.2.1 Outline
THEORY02-60	2.2.2 Hardware Architecture
THEORY02-110	2.2.3 Hardware Component
THEORY02-240	2.3 Subsystem Specifications
THEORY02-250	2.4 Physical Specifications
THEORY02-250	2.4.1 DKC515I Physical Specifications
THEORY02-260	2.5 Power Specifications
THEORY02-260	2.5.1 Subsystem Power Specifications
THEORY02-261	2.5.2 Steady Current
THEORY02-270	2.5.3 Notes of Power Supply
THEORY02-310	2.6 Environmental Specifications
THEORY02-320	2.7 Layout
THEORY02-320	2.7.1 DKC515I Layout
THEORY02-330	2.7.2 Twin Rack Configuration of Disk Subsystem
THEORY03-10	3. Internal Operation
THEORY03-10	3.1 Hardware Block Diagram
THEORY03-20	3.2 Software Organization
THEORY03-40	3.3 Data Formats
THEORY03-80	3.4 Cache Management
THEORY03-170	3.5 Destaging Operations
THEORY03-200	3.6 Operations Performed when Drive Errors Occur
THEORY03-230	3.7 Inter Mix of Drives and Emulation types
THEORY03-230	3.7.1 Drives to be Connected
THEORY03-240	3.7.2 Emulation Device Type
THEORY03-290	3.8 16384 logical addresses

THEORY03-300	3.9 LDEV Formatting
THEORY03-300	3.9.1 Outlines of a high-speed LDEV Formatting
THEORY03-310	3.9.2 Estimation of a high-speed LDEV Formatting Time
THEORY03-310	3.9.2.1 DKxxx-JxxxFC/KxxxFC
THEORY03-311	3.9.3 Notes when logical volume format is executed
THEORY03-320	3.10 High-Speed Online Replacement of Microprogram
THEORY03-320	3.10.1 Outline
THEORY03-330	3.10.2 Processing Time
THEORY03-340	3.11 HRC
THEORY03-340	3.11.1 HRC Components
THEORY03-430	3.11.2 HRC Software Requirements
THEORY03-440	3.11.3 HRC Hardware Requirements
THEORY03-590	3.11.4 HRC Theory of Operations
THEORY03-630	3.11.5 HRC Control Operations
THEORY03-760	3.11.6 Managing HRC Environment
THEORY03-810	3.11.7 HRC Error Recovery
THEORY03-930	3.12 HMRCF & HOMRCF
THEORY03-930	3.12.1 Overview
THEORY03-1130	3.12.2 Construction of HMRCF & HOMRCF
THEORY03-1140	3.12.3 Status transition
THEORY03-1150	3.12.4 Interface
THEORY03-1160	3.12.5 Cascade function
THEORY03-1180	3.12.6 Reverse-RESYNC
THEORY03-1260	3.12.7 ShadowImage-FlashCopy (R) Option function
THEORY03-1290	3.12.8 ShadowImage-FlashCopy (R) version2 Option function
THEORY03-1310	3.12.9 Micro-program Exchange
THEORY03-1320	3.12.10 Notes on powering off
THEORY03-1321	3.12.11 Copy-on-Write Snapshot option
THEORY03-1330	3.13 TPF
THEORY03-1330	3.13.1 An outline of TPF
THEORY03-1360	3.13.2 TPF Support Requirement
THEORY03-1370	3.13.3 TPF trouble shooting method
THEORY03-1380	3.13.4 The differences of DASD-TPF(MPLF) vs DASD-MVS
THEORY03-1420	3.13.5 Notices for HRC-option setting
THEORY03-1430	3.14 HXRC
THEORY03-1430	3.14.1 Outline of HXRC
THEORY03-1440	3.14.2 HXRC Support Requirements
THEORY03-1440	3.14.2.1 OS
THEORY03-1450	3.14.2.2 Hardware
THEORY03-1460	3.14.2.3 Micro-program
THEORY03-1480	3.14.2.4 HXRC recommendations
THEORY03-1490	3.14.3 Online Maintenance while Concurrent Copy (CC)/HXRC in use

THEORY03-1510	3.15 HRC Asynchronous
THEORY03-1510	3.15.1 Components
THEORY03-1530	3.15.2 Consistency Group
THEORY03-1570	3.15.3 HRC Asynchronous Theory of Operations
THEORY03-1620	3.15.4 MCU-RCU Communications to Maintain Asynchronous Copies
THEORY03-1630	3.15.5 Failure Detected by the RCU
THEORY03-1640	3.15.6 Inflow Control for Sidefiles
THEORY03-1650	3.15.7 HRC Asynchronous Control Operations
THEORY03-1650	3.15.7.1 DKC Options
THEORY03-1660	3.15.7.2 Registering/Monitoring/Deleting the Consistency Group
THEORY03-1700	3.15.7.3 Pair Status
THEORY03-1740	3.15.7.4 Controlling Volume Pairs
THEORY03-1840	3.15.7.5 Monitoring subsystem Statistics
THEORY03-1850	3.15.7.6 On-line Micro-program exchange procedure
THEORY03-1860	3.15.8 Management/Recovery Procedures
THEORY03-1860	3.15.8.1 Managing HRC Asynchronous Subsystems
THEORY03-1900	3.15.8.2 Recovering from Pair Suspended
THEORY03-1910	3.15.8.3 Disaster Recovery - Switching to the Secondary Subsystem
THEORY03-1940	3.15.8.4 Disaster Recovery - Switching Back to the Primary Subsystem
THEORY03-1960	3.16 HIHSM(Hitachi internal Hierarchical Storage Management)
THEORY03-1960	3.16.1 HIHSM Overview
THEORY03-1970	3.16.2 Hardware requirements
THEORY03-1970	3.16.3 Software requirements
THEORY03-1980	3.16.4 Monitor function
THEORY03-2000	3.16.5 Estimate function
THEORY03-2010	3.16.6 Volume moving (migration) function
THEORY03-2100	3.16.7 Decision of volume moving (migration)
THEORY03-2120	3.16.8 Automating Volume Migration function
THEORY03-2180	3.17 Data Assurance at a Time When a Power Failure Occurs
THEORY03-2180	3.17.1 Outline

THEORY03-2220	3.18 HPAV
THEORY03-2220	3.18.1 Overview
THEORY03-2220	3.18.1.1 Overview of HPAV
THEORY03-2250	3.18.1.2 How to Obtain Optimum Results from HPAV
THEORY03-2260	3.18.2 Preparing for HPAV Operations
THEORY03-2260	3.18.2.1 System Requirements
THEORY03-2270	3.18.2.2 Preparations at the Host Computer
THEORY03-2280	3.18.2.2.1 Generation Definition of Base Devices and Alias Addresses
THEORY03-2290	3.18.2.2.2 Setting the WLM Operation Mode
THEORY03-2300	3.19 FICON
THEORY03-2300	3.19.1 Introduction
THEORY03-2310	3.19.2 Environment
THEORY03-2320	3.19.3 DKC515I FICON specification
THEORY03-2330	3.19.4 Configuration
THEORY03-2360	3.19.5 The operation procedure
THEORY03-2370	3.20 Universal Volume Manager (UVM)
THEORY03-2370	3.20.1 Overview
THEORY03-2390	3.20.2 Procedure of using external volumes
THEORY03-2390	3.20.2.1 Prepare in the external subsystem a volume to be used for UVM
THEORY03-2400	3.20.2.2 Change the port attribute to External
THEORY03-2400	3.20.2.3 Connect the external subsystem to the external port
THEORY03-2400	3.20.2.4 Search for the external subsystem from the UVM operation panel (Discovery)
THEORY03-2410	3.20.2.5 Map an external volume
THEORY03-2430	3.20.2.6 Format volume (MF volume only)
THEORY03-2430	3.20.2.7 Define the host path (Open volume only)
THEORY03-2440	3.20.2.8 Other settings
THEORY03-2470	3.20.3 Cross-system Copy
THEORY03-2470	3.20.3.1 Overview
THEORY03-2500	3.21 Save and Recovery of the SM Control Information in the case of PS OFF/ON

THEORY03-2510	3.22 Universal Replicator
THEORY03-2510	3.22.1 URz Components
THEORY03-2530	3.22.1.1 TagmaStore™ USP Series Storage Subsystems
THEORY03-2540	3.22.1.2 Main and Remote Control Units (Primary Subsystems and Secondary Subsystems)
THEORY03-2550	3.22.1.3 Journal Group
THEORY03-2560	3.22.1.4 Data Volume Pair
THEORY03-2570	3.22.1.5 Journal Volume
THEORY03-2610	3.22.1.6 Remote Copy Connections
THEORY03-2620	3.22.1.7 Initiator Ports and RCU Target Ports
THEORY03-2630	3.22.1.8 URz Remote Console Software
THEORY03-2640	3.22.1.9 Host I/O Time-Stamping Function
THEORY03-2650	3.22.1.10 Error Reporting Communications (ERC)
THEORY03-2660	3.22.2 Remote Copy Operations
THEORY03-2670	3.22.2.1 Initial Copy Operation
THEORY03-2680	3.22.2.2 Update Copy Operation
THEORY03-2690	3.22.2.3 Read and Write I/O Operations During URz Volumes
THEORY03-2700	3.22.2.4 Secondary Data Volume Write Option
THEORY03-2710	3.22.2.5 Secondary Data Volume Read Option
THEORY03-2720	3.22.2.6 Difference Management
THEORY03-2730	3.22.3 Journal Processing
THEORY03-2740	3.22.3.1 Creating and Storing Journals at the Primary Subsystem
THEORY03-2750	3.22.3.2 Copying Journals to the Secondary Subsystem
THEORY03-2760	3.22.3.3 Storing Journal at the Secondary Subsystem
THEORY03-2770	3.22.3.4 Selecting and Restoring Journal at the Secondary Subsystem
THEORY03-2790	3.22.3.5 Types of Journal
THEORY03-2800	3.22.4 UR operation
THEORY03-2800	3.22.4.1 Pair operation
THEORY03-2930	3.22.4.2 USAGE/HISTORY
THEORY03-2960	3.22.4.3 Option
THEORY03-2980	3.22.5 Maintenance features and procedure
THEORY03-2980	3.22.5.1 Maintenance
THEORY03-2990	3.22.5.2 PS OFF/ON Process
THEORY03-3000	3.22.5.3 Power failure
THEORY03-3010	3.22.6 Cautions on software
THEORY03-3010	3.22.6.1 Error recovery
THEORY03-3030	3.22.7 Disaster Recovery Operations
THEORY03-3030	3.22.7.1 Preparation for Disaster Recovery
THEORY03-3040	3.22.7.2 Sense information is transferred between sites
THEORY03-3050	3.22.7.3 File and Database Recovery Procedures
THEORY03-3060	3.22.7.4 Switching Operations to the Secondary Site
THEORY03-3070	3.22.7.5 Transferring Operations Back to the Primary Site
THEORY03-3080	3.22.7.6 Resuming Normal Operations at the Primary Site

THEORY03-3090	3.23 CUIR
THEORY03-3090	3.23.1 Overview
THEORY03-3100	3.23.2 Preparing for Operations
THEORY03-3100	3.23.2.1 System Requirements
THEORY03-3120	3.23.2.2 Extended SAID
THEORY03-3130	3.23.2.3 CUIR Installation procedures
THEORY03-3140	3.23.3 Operations of CUIR
THEORY03-3150	3.23.4 Notes
THEORY03-3170	3.24 Cautions when Stopping the Subsystem
THEORY03-3170	3.24.1 Cautions when Stopping the Subsystem
THEORY03-3180	3.24.2 Operation when Turning Off the Distribution Board Breaker
THEORY03-3190	3.24.3 Cautions when Turning Off the Battery Switch (Clearing the Management Information)
THEORY03-3250	3.25 Hyper PAV
THEORY03-3250	3.25.1 Overview of Hyper PAV
THEORY03-3250	3.25.2 Hyper PAV function setting procedure
THEORY04-10	4. Power-on Sequences
THEORY04-10	4.1 IMPL Sequence
THEORY04-30	4.2 Drive Power-on Sequence
THEORY04-40	4.3 Planned Stop
THEORY05-10	5. Appendixes
THEORY05-10	5.1 Physical - Logical Device Matrixes
THEORY05-90	5.2 Commands
THEORY05-130	5.3 Comparison of pair status on SVP, Web Console, Raid Manager

1. RAID Architecture Overview

The objectives of the RAID technology are the low cost, high reliability, and high I/O performance of disk storage devices. To achieve these objectives, this subsystem supports levels 1, 5 and 6 of RAID technologies (in this section, part of level 3 RAID technology is explained to make the outline of RAID5 more understandable). The features of the levels of RAID technologies are described below.

1.1 Outline of RAID Systems


The concept of disk array was announced in 1987 by the research group of University of California at Berkeley.

The research group called the disk array RAID (Redundant Array of Inexpensive Disks: A disk subsystem that has redundancy by employing multiple inexpensive and small disk drives), classified the RAID systems into five levels, that is, RAID 1 to RAID 5, and added RAID 0 and RAID 6 later. Since the DKC515I disk subsystem supports RAID 1, RAID 5, and RAID 6, the method, advantage, and disadvantage of each of them are explained below.

Table 1.1-1 Outline of RAID Systems

Level	Configuration		Characteristics
RAID 1 (2D+2D) configuration	The diagram shows a Data block row with blocks A, B, C, D, E, F, and an ellipsis. Below this is a DKC (Data Key Controller) box. Under the DKC are two RAID pairs. Each RAID pair consists of a Primary drive and a Secondary drive. The first RAID pair contains blocks A and A' (mirrored), and the second RAID pair contains blocks B and B' (mirrored). Vertical ellipses indicate that blocks C, D, E, and F are also mirrored. The entire setup is labeled as a Parity group.	Outline	Mirror disks (duplicated writing) Two disk drives, primary and secondary disk drives, compose a RAID pair (mirroring pair) and the identical data is written to the primary and secondary disk drives. Further, data is scattered on the two RAID pairs.
		Advantage	RAID 1 is highly usable and reliable because of the duplicated data. It has higher performance than ordinary RAID 1 (when it consists of two disk drives) because it consists of the two RAID pairs.
		Disadvantage	A disk capacity twice as large as user data capacity is required.
RAID 1 (4D+4D) configuration (Two concatenation of RAID1 (2D+2D))	The diagram shows a Data block row with blocks A, B, C, D, E, F, G, H, and an ellipsis. Below this is a DKC box. Under the DKC are four RAID pairs, each with a Primary drive and a Secondary drive. The first RAID pair contains blocks A and A', the second contains B and B', the third contains C and C', and the fourth contains D and D'. Vertical ellipses indicate that blocks E, F, G, and H are also mirrored. A note states: 'Note: A RAID pair consists of two disk drives.' The entire setup is labeled as a Parity group.	Outline	Mirror disks (duplicated writing) The two parity groups of RAID 1 (2D+2D) are concatenated and data is scattered on them. In the each RAID pair, data is written in duplicate.
		Advantage	This configuration is highly usable and reliable because of the duplicated data. It has higher performance than the 2D+2D configuration because it consists of the four RAID pairs.
		Disadvantage	A disk capacity twice as large as user data capacity is required.

RAID5	Data block A B C D E F DKC A E : : B F : : C P1 : : P0 D : : Data disks + Parity disk Note: There are two configurations of RAID 5: 3D+1P configuration (four disk drives) and 7D+1P configuration (eight disk drives). The above diagram shows the 3D+1P configuration. In the 7D+1P configuration, data is arranged in the same way. <td> Outline Data is written to multiple disks successively in units of block (or blocks). Parity data is generated from data of multiple blocks and written to optional disk. Advantage RAID 5 fits the transaction operation mainly uses small size random access because each disk can receive I/O instructions independently. It can provide high reliability and usability at a comparatively low cost by virtue of the parity data. Disadvantage Write penalty of RAID 5 is larger than that of RAID 1 because pre-update data and pre-update parity data must be read internally because the parity data is updated when data is updated. </td>	Outline Data is written to multiple disks successively in units of block (or blocks). Parity data is generated from data of multiple blocks and written to optional disk. Advantage RAID 5 fits the transaction operation mainly uses small size random access because each disk can receive I/O instructions independently. It can provide high reliability and usability at a comparatively low cost by virtue of the parity data. Disadvantage Write penalty of RAID 5 is larger than that of RAID 1 because pre-update data and pre-update parity data must be read internally because the parity data is updated when data is updated.
RAID6	Data block A B C D E F DKC A E : : B F : : C P1 : : P0 Q1 : : Q0 D : : Data disks + Parity disks P and Q Note: RAID 6 (6D+2P) configuration practically consists of eight disk drives. In the above diagram, three disk drives are omitted. <td> Outline Data blocks are scattered to multiple disks in the same way as RAID 5 and two parity disks, P and Q, are set in each row. Therefore, data can be assured even when failures occur in up to two disk drives in a parity group. Advantage RAID 6 is far more reliable than RAID 1 and RAID 5 because it can restore data even when failures occur in up to two disks in a parity group. Disadvantage Because the parity data P and Q must be updated when data is updated, RAID 6 is imposed write penalty heavier than that on RAID 5, performance of the random writing is lower than that of RAID 5 in the case where the number of drives makes a bottleneck. </td>	Outline Data blocks are scattered to multiple disks in the same way as RAID 5 and two parity disks, P and Q, are set in each row. Therefore, data can be assured even when failures occur in up to two disk drives in a parity group. Advantage RAID 6 is far more reliable than RAID 1 and RAID 5 because it can restore data even when failures occur in up to two disks in a parity group. Disadvantage Because the parity data P and Q must be updated when data is updated, RAID 6 is imposed write penalty heavier than that on RAID 5, performance of the random writing is lower than that of RAID 5 in the case where the number of drives makes a bottleneck.

RAID5 concatenation	Data block  DKC Parity group Note: The above-mentioned figure is four concatenation configuration, but it is the same in the case of two concatenation.	Outline Advantage Disadvantage	In the case of RAID5 (7D+1P), two or four parity groups (eight drives) are concatenated, and the data is distributed and arranged in 16 drives or 32 drives. Two concatenation: V05 support; Four concatenation: V06 support. When the parity group becomes a performance bottleneck, the performance improvement can be attempted because it is configured with twice and four times the number of drives in comparison with RAID5 (7D+1P). The influence level when two drives are blocked is large because twice and four times LDEVs are arranged in comparison with RAID5 (7D+1P). However, the probability that the read of the single block in the parity group becomes impossible due to the failure is the same as that of RAID5 (7D+1P).
----------------------------	--	---	--

1.2 Comparison of RAID levels

(1) Space efficiency

RAID level	Space efficiency (User area/Disk capacity)	Remarks
RAID1 2D+2D	50.0%	Because of the mirroring
RAID1 4D+4D	50.0%	Because of the mirroring
RAID5 3D+1P	75.0%	Ratio of the number of parity disks to the number of data disks The space efficiency of the 6D+2P is the same as that of the 3D+1P. Two concatenation and four concatenation of 7D+1P are also the same.
RAID5 7D+1P	87.5%	
RAID6 6D+2P	75.0%	

(2) Comparison of performance limits of parity groups (When supposing the marginal efficiency of the RAID 1 (2D+2D) to be 100%)

RAID level	Random and sequential reading	Sequential writing	Random writing
RAID1 2D+2D	100%	100%	100%
RAID1 4D+4D	200%	200%	200%
RAID5 3D+1P	100%	150%	50%
RAID5 7D+1P	200%	350%	100%
RAID6 6D+2P	200%	300%	66.7% (The efficiency is lowered by 33% compared with the 7D+1P.)
Remarks	Proportionate to the number of HDDs	Proportionate to the number of data HDDs	See the explanation below.

- In the case of two concatenation and four concatenation RAID5 (7D+1P), it becomes the value twice and four times the above-mentioned.
- The reason why the efficiency is lowered by 33% in the case of RAID 6 in comparison with RAID 5 (7D+1P) is as follows.

When RAID 5 executes random writing, it issues a total of four IOs, that is, reading of old data, reading of old parity data, writing of new data, and writing of new parity data to disk drives.

In the case of RAID 6, on the other hand, it issues a total of six IOs, that is, reading of old data, reading of old parity data (P), reading of old parity data (Q), writing of new data, writing of new parity data (P), and writing of new parity data (Q) to disk drives.

The number of IOs that RAID 5 issues is four, whereas those that RAID 6 issues is six; the latter is 1.5 times as many as the former. Therefore, the random writing performance of RAID 6 is lowered by 33% in comparison with RAID 5.

However, unless RAID 6 is in an environment in which the number of drives makes a bottleneck, the write penalty is absorbed by the cache memory, so that the performance is not lowered.

(3) Reliability

RAID level	Conditions of data assurance
RAID1 2D+2D	When a failure occurs in one of the mirroring pair of disk drives, data can be restored through use of data of the other disk drives.
RAID1 4D+4D	When failures occur in both of the mirroring pair of disk drives, an LDEV blockade is caused.
RAID5 3D+1P	When a failure occurs in one disk drive in a parity group, data can be restored through use of the parity data.
RAID5 7D+1P	When failures occur in two disk drives, an LDEV blockade is caused.
RAID6 6D+2P	When the failure(s) occur(s) in one or two disk drive(s) in a parity group, data can be restored through use of the parity data. When three disk drives fail, an LDEV blockade is caused.

In the case of RAID 6, data can be assured when up to two drives in a parity group fail, as explained above. Therefore, RAID 6 is the most reliable in the RAID levels.

2. Hardware Specifications


2.1 General

The DKC515I is a high performance and large capacity disk subsystem at the high end that follows the architecture of the DKC460I and has improved Hi-Star Net Architecture and a speeded up microprocessor.


DKC515I consists of disk control frame (primary rack) which can install 120 HDDs, and disk array frame (secondary rack) which can install 120 HDDs and a subsystem is a flexible configuration, which goes from 0 to 240 HDDs (minimum and maximum).

DKC515I is only a single-phase AC power model and this model is connectable to both mainframe systems and open systems.

Single Rack configuration



Twin Rack configuration



2.1.1 Features

(1) Scalability

The DKC515I is capable of providing disk subsystems that fit customer's needs by having the channel adapters, cache memories, and disk drives installed according to subsystem specifications.

- Number of installed channel options: 1 set.
- Capacity of cache memory: 4 GB to 32 GB
- Number of disk drives: Up to 240 HDDs/8 disk paths.

(2) High-performance

- The DKC515I supports two kinds of high-speed disk drives with a speed of 10 kmin-1 or 15 kmin-1
- It realizes high-speed data transfer between the MIX and HDDs at a rate of 2Gbps with the fibre channel (FC-AL).
- Performance of the microprocessor installed in the MIX is two times as fast as the DKC460I.
- A bandwidth of 17GB/s is realized through the further improvement of Hi-Star Net Architecture (HSN) which gave satisfactory results to DKC460I.

(3) Large Capacity

- The DKC515I supports three kinds of disk drives with various capacities: 72GB, 146GB and 300GB.
- It can control up to 16,384 logical volumes and up to 240 disk drives, so that it realizes a physical disk capacity of approximately 69.2TB per subsystem.

(4) Connectivity

The DKC515I supports OS's for various UNIX servers, PC servers, and mainframes, so that it conforms to heterogeneous system environment in which those various OS's coexist.

Show the platform that can be connected in the following table.

Mainframe		Open system	
Maker	OS version	Maker	OS version
Hitachi	VOS3/FS, VOS3/AS	HP	HP-UX (11.0, 11I, 11.i v2)
IBM	OS/390		Tru64 (5.1B) Open VMS(V7.3)
	MVS/ESA, MVS/XA	Sun	Solaris (8, 9)
	VM/ESA, VSE/ESA	IBM	AIX (5.1, 5.2)
	Z/OS & Z/OS.e, Z/VM		AIX (4.3.3, 5L)
	Redhat 7.2 for S/390 & Z-series	—	Windows (2000, 2003)
		NOVELL	NetWare (6.0)
		SGI	IRIX64 (6.5.16)
		REDHAT	Linux (3.0)

Host interfaces supported by the DKC515I are shown below. They can mixture within the subsystem.

- Mainframe: Serial channel (ESCON) and fibre channel (FICON)
- Open system: Fibre channel and Networks (NAS)

(5) High reliability

- The DKC515I supports RAID6 (6D+2P), RAID5 (3D+1P/7D+1P) and RAID1 (2D+2D). Two concatenation and four concatenation configurations are possible for RAID5 (7D+1P). Two concatenation: V05 support; Four concatenation: V06 support.
- Main components are implemented with a duplex or redundant configuration, so even when single point of the component failure has occurred, the subsystem can continue non-stop operation.

(6) Non-disruptive Service and Upgrades

- Main components can be added, removed, and replaced without shutting down a device while the subsystem is in operation.
- A service processor (SVP) mounted on the DKC monitors the running condition of the subsystem. Connecting the SVP with a service center enables remote maintenance.
- The microcode can be upgraded without shutting down the subsystem.


2.2 Architecture

2.2.1 Outline

Primary rack is composed of DKC Box and two Disk Chassis which 60 disk drives can be installed in and the Secondary rack is composed of two Disk Chassis which 60 disk drives can be installed in.

DKC can control a maximum of 240 HDDs when connected with four Disk Chassis.

An outline of frame components of the DKC515I disk subsystem is shown below.



*1: This includes the thickness of side covers (12.5mm × 2)

(1) Disk Controller

The DKC consists of the DKC Box in which the MIX-PCB (channel adapters & disk adapters), BASE-PCB (cache memories & shared memories & CSW) are installed, and power supplies and Battery Boxes that supply power to the components above. Most of main components have been implemented with a duplex or redundant configuration, and achieves non-stop operation against failure of single point of the components. In addition, components can be replaced and added, and the microcode can be upgraded while the subsystem is in operation. The control unit is equipped with an SVP, which is used to service the subsystem, monitor its running condition, and analyze faults. Connecting the SVP with a service center enables remote maintenance of the subsystem.

(2) Disk Chassis

The Disk Chassis consists of the four HDD Boxes in each of which 30 disk drives can be installed, cooling fans, and power supplies that supply power to the components above. It can have up to 60 disk drives installed.

This disk unit is configured from duplex power supply system and cooling fans with redundancy introduced. In addition, the disk drives achieve non-stop operation against failure of one piece of components by employing the RAID1 or RAID5 or RAID6 level. Components can be replaced and added while the subsystem is in operation.

2.2.2 Hardware Architecture

The normal configuration of DKC515I is divided into the power supply section, DKC section, and DKU section in which disk drives are installed.


The Disk-less configuration without DKU section is divided into the power supply section and DKC section.

The Power Supply section consists of AC-DC power supplies and Battery Boxes.


The DKC section consists of channel adapters (CHAs), disk adapters (DKAs), cache memories (CACHEs), shared memories (SMs) and disk units (DKUs). Each component is connected with the cache paths and SM paths.

For the normal configuration, the DKC section and DKU section are connected with the disk paths.

The normal configuration is shown below.



The Disk-less configuration is shown below.




(1) Power Supply

The DKC515I adopts the duplicated structure for the AC power reception as before, so that it continues its operation when a trouble occurs in one of the power systems. When the AC power supply is shut off owing to a power failure, the subsystem advances to the backup process in which the cache and shared memories are backed up for 36hours.

In the DKC515I, the AC power supplied is converted into 56V (or 48V) DC power by the AC-DC Power Supply and supplied to each component inside the subsystem. Each component has a DC-DC converter and generates necessary voltage from the 56V DC to supply it to itself.

(2) SM Path


In the DKC515I, the shared memory was installed in the exclusive slot on the BASE PCB.



The reliability of data stored in the shared memory is enhanced through the duplication of the data separated into the clusters 1 and 2.

(3) Cache Path

In the DKC515I, the number of cache paths is doubled and furthermore the transfer rate per path is made higher. Performance of accessing the cache memory varies depending on the configuration of the CACHE and the CSW.




Hi-Star Net performance according to subsystem configuration

- Cache (BASE) + CSW (BASE) ----- max. 17GB/s

(4) Disk Path

In the DKC515I, the number of ports per MIX-PCB is increased to four, the transfer rate per fibre port is made higher to 2Gbps.



2.2.3 Hardware Component

(1) Channel Adapter

The channel adapter (CHA) controls data transfer between the upper host and the cache memory. The DKC515I provides various kinds of CHAs, which support the mainframe, SAN (Storage Area Network), and NAS (Network Attached Storage) interface, as options to be added in units of pair.

Mainframe Connection

	ESCON	Mainframe Fibre 8port	
		Short Wave	Long Wave
Option name	16S/16SR	8MS/8MSR	8ML/8MLR
Host interface	ESCON	FICON	FICON
Data transfer rate (MB/s)	17	100/200	100/200
Number of options installed	1	1	1
Number of ports / Option	16	8	8
Number of ports / Subsystem	16	8	8
Maximum cable length	3km	500m/300m (*1)	10km


	Mainframe Fibre 16port				
	Short Wave		Long Wave		
Option name	16MSR	16MFSR	16MLR	16MFLR	16MFL4R
Host interface	FICON	FICON	FICON	FICON	FICON
Data transfer rate (MB/s)	100/200	100/200/400	100/200	100/200/400	100/200/400
Number of options installed	1	1	1	1	1
Number of ports / Option	16	16	16	16	16
Number of ports / Subsystem	16	16	16	16	16
Maximum cable length	500m/300m (*1)	500m/300m/150m (*1)	10km	10km	4km

Open System Connection

		Fibre (*3)			
		16port	8port	16port	32port
Option name		MIX2R	8HSR	16HS/16HSR	32HS/32HSR
Host interface		FCP	FCP	FCP	FCP
Data transfer rate (MB/s)		100/200	100/200	100/200	100/200
Number of options installed		1	1	1	1
Number of ports / Option		16	8	16	32
Number of ports / Subsystem		16	8	16	32
Maximum cable length (*3)	Short Wave	500m/300m (*1)			
	Long Wave	10km			

		Fibre (*3)			NAS (*4)	iSCSI (*4)
		8port	16port		32port	8port
Option name		8FS2R	MIX4R	16FS2R	32FSR/ 32FS2R	8NS/ 8NSR
Host interface		FCP	FCP		FCP	Gigabit Ethernet
Data transfer rate (MB/s)		100/200/400	100/200/400		100/200/400	100
Number of options installed		1	1	1	1	1
Number of ports / Option		8	16	16	32	8
Number of ports / Subsystem		8	16	16	32	8
Maximum cable length (*3)	Short Wave	500m/300m/150m (*1)				500m/275m (*2)
	Long Wave	4km: 1FL4R use, 10km: 1FLR use				—

- *1: Indicates when 50/125mm multi-mode fiber cable is used. If 62.5/125mm multi-mode fiber cable is used, 500m (100MB/s), 300m (200MB/s) and 150m (400MB/s) are decreased to 300m, 150m and 75m respectively.
- *2: Indicates when 50/125mm multi-mode fiber cable is used. If 62.5/125mm multi-mode fiber cable is used, maximum length is decreased to 275m.
- *3: The CHA for the fibre channel connection can conform to any one of the short and long wavelengths concerning each port depending on a selection of a transceiver installed on a port on a PCB. Since a transceiver conforming to the short wavelength is installed on a port on the each CHA as the standard, an optional transceiver conforming to the long wavelength is required separately when changing the standard port to that conforming to the long wavelength (DKC-F510I-1HL/1HLR/1FLR/1FL4R).
- *4: The CHA for connection with NAS and iSCSI supports the short wavelength only.




(2) Disk Adapter

The disk adapter (DKA) controls data transfer between the disk drive and cache memory. In DKC515I, the option only of DKA is not support. However, the new component (MIX-PCB) that consists of both channel adapter functionality and disk adapter functionality is provided.

The number of the ports per MIX-PCB is four and the transfer rate per port is 2Gbps.

Model number	DKC-F515I-MIX2R	DKC-F515I-MIX4R
Number of PCBs/subsystem	2	2
Configuration of PCBs	WP525-B × 1 & SH343-B × 4	WP526-A × 1 & SH343-B × 4
Number of microprocessor/PCB	4	4
Performance of microprocessor	400MHz	400MHz
Performance of fibre port	2Gbps	2Gbps
Number of Fibre port/PCB	4	4
Maximum number of disk paths/subsystem	8	8
Maximum number of disk drives/FCAL	60	60




(3) Cache Memory

This is the memory located in the middle of CHA and DKA, and it is used in order to perform I/O processing asynchronously with the reading and writing to a disk drive.

The cache capacity can be enlarged from the minimum of 4 GB to the maximum of 64 GB by the 4GB or 8GB.

Since the DKC has the two sets of cache memories onto which data to be written on the disk is written doubly, the data will not be lost even if a cache failure occurs.

Memory Chip	Number of Additional	Cache Memory Capacity	
		Minimum	Maximum
512Mbit DRAM (DKC-F510I-C4G/C4GR)	4GB (4DIMM)	4GB	32GB
1Gbit DRAM (DKC-F510I-C8GR)	8GB (4DIMM)	8GB	64GB




(4) Shared Memory

This is a memory that stores configuration data and information for controlling the cache memory and disk drives. The memory can be accessed commonly from the CHA and DKA. The shared memory capacity has a close relation with capacities of the installed cache memories and disks (number of LDEV) and can be enlarged from the minimum of 1GB to the maximum of 8GB by the 1GB or 2GB.

In the DKC515I, the shared memory was installed in the exclusive slot on the BASE-PCB. Since the DKC has the two sets of shared memories to duplicate data, the data will not be lost even if a SM failure occurs.

Memory Chip	Number of Additional	Shared Memory Capacity	
		Minimum	Maximum
512Mbit DRAM (DKC-F510I-S1G/S1GR)	1GB (2DIMM)	1GB	6GB
1Gbit DRAM (DKC-F510I-S2GR)	2GB (2DIMM)	2GB	8GB



(5) Disk Drive


The DKC515I supports seven types of disk drives whose capacities are 72GB/10kmin⁻¹, 72GB/15kmin⁻¹, 146GB/10kmin⁻¹, 146GB/15kmin⁻¹, 300GB/10kmin⁻¹, 300GB/15kmin⁻¹ and 400GB/10kmin⁻¹ respectively and each of them has 2Gbps fibre channel interface.

Item		72GB/15kmin ⁻¹	146GB/10kmin ⁻¹	146GB/15kmin ⁻¹
Disk Drive Model Name	Seagate	DKS2C-K72FC ^{*1}	DKS2D-J146FC ^{*9}	DKS2D-K146FC ^{*16}
		DKS2D-K72FC ^{*2}	DKS2E-J146FD ^{*10}	DKS2E-K146FC ^{*17}
		DKS2E-K72FC ^{*3}	DKS2G-J146FD ^{*11}	DKS2F-K146FC ^{*18}
		DKS2F-K72FD ^{*4}	—	DKS2G-K146FD ^{*19}
		DKS2G-K72FD ^{*5}		
	HGST	DKR2F-K72FC ^{*6}	DKR2E-J146FC ^{*12}	DKR2F-K146FC ^{*20}
		DKR2G-K72FC ^{*7}	DKR2F-J146FC ^{*13}	DKR2G-K146FC ^{*21}
		DKR2J-K72FD ^{*8}	DKR2G-J146FC ^{*14}	DKR2J-K146FD ^{*22}
	DKR2J-J146FD ^{*15}			
Capacity/HDD		71.50GB	143.76GB	143.76GB
Number of heads		8 ^{*1/4} 4 ^{*2/2} 3 ^{*3/3} 4 ^{*4/4} 5 ^{*5/2} 6 ^{*6/2} 7 ^{*7/4} 8 ^{*8}	4 ^{*9/4} 4 ^{*10/4} 4 ^{*11/10} 10 ^{*12/5} 13 ^{*13/4} 14 ^{*14/4} 15 ^{*15}	8 ^{*16/4} 4 ^{*17/3} 3 ^{*18/4} 4 ^{*19/10} 20 ^{*20/4} 21 ^{*21/4} 22 ^{*22}
Number of disks		4 ^{*1/2} 2 ^{*2/1} 3 ^{*3/2} 4 ^{*4/2} 5 ^{*5/3} 6 ^{*6/1} 7 ^{*7/2} 8 ^{*8}	2 ^{*9/2} 10 ^{*10/2} 11 ^{*11/5} 5 ^{*12/3} 13 ^{*13/2} 14 ^{*14/2} 15 ^{*15}	4 ^{*16/2} 17 ^{*17/2} 18 ^{*18/2} 19 ^{*19/5} 20 ^{*20/2} 21 ^{*21/2} 22 ^{*22}
Seek Time (ms) (Read/Write)	minimum	0.4/0.6	0.65/0.85	0.4/0.6
	average	3.8/4.2	4.9/5.5	3.8/4.1
	maximum	6.8/7.6	10.0/11.0	6.8/7.6
Average latency time (ms)		2.01	2.99	2.01
Revolution speed (min ⁻¹)		15,000	10,000 ^{*9,*13/} 15,000 ^{*14,*15}	15,000
Interface data transfer rate (MB/S)		200	200	200
Internal data transfer rate (MB/S)		74.5 to 111.4	57.3 to 99.9	76.13 to 113.78

Item		300GB/10kmin ⁻¹	72GB/10kmin ⁻¹	300GB/15kmin ⁻¹
Disk Drive Model Name	Seagate	DKS2D-J300FC ^{*23}	—	DKS2E-K300FC ^{*32}
		DKS2E-J300FC ^{*24}	—	DKS2F-K300FC ^{*33}
		DKS2G-J300FC ^{*25}		DKS2G-K300FC ^{*34}
	HGST	DKR2F-J300FC ^{*26}	DKR2F-J72FC ^{*29}	DKR2G-K300FC ^{*35}
		DKR2G-J300FC ^{*27}	DKR2G-J72FD ^{*30}	DKR2H-K300FC ^{*36}
		DKR2J-J300FC ^{*28}	DKR2J-J72FD ^{*31}	DKR2J-K300FC ^{*37}
Capacity/HDD		288.20GB	72.91GB	288.20GB
Number of heads		8 ^{*23/4} *24/4 ^{*25/} 10 ^{*26/8} *27/4 ^{*28}	3 ^{*29/4} *30/4 ^{*31}	8 ^{*32/6} *33/4 ^{*34/} 8 ^{*35/8} *36/4 ^{*37}
Number of disks		4 ^{*23/2} *24/2 ^{*25/} 5 ^{*26/4} *27/2 ^{*28}	2 ^{*29/2} *30/2 ^{*31}	4 ^{*32/3} *33/2 ^{*34/} 4 ^{*35/4} *36/2 ^{*37}
Seek Time (ms) (Read/Write)	minimum	0.65/0.85	0.4/0.45	0.2/0.44
	average	4.9/5.5	4.7/5.1	3.6/4.1
	maximum	10.0/11.0	10.0/11.0	6.8/7.4
Average latency time (ms)		2.99	2.99	2.01
Revolution speed (min ⁻¹)		10,000 ^{*23.*26/} 15,000 ^{*27.*28}	10,000 ^{*29/} 15,000 ^{*30.*31}	15,000
Interface data transfer rate (MB/S)		200	200	200
Internal data transfer rate (MB/S)		58.75 to 114.13	72.7 to 134.4	108 to 180

Note: min⁻¹ = r.p.m.

Item		400GB/10kmin ⁻¹
Disk Drive Model Name	Seagate	DKS2E-J400FC ^{*38}
		DKS2G-J400FD ^{*39}
	HGST	—
		—
Capacity/HDD		393.85GB
Number of heads		8 ^{*38} /6 ^{*39}
Number of disks		4 ^{*38} /3 ^{*39}
Seek Time (ms) (Read/Write)	minimum	0.35/0.52
	average	3.9/4.4
	maximum	8.1/8.7
Average latency time (ms)		2.98
Revolution speed (min ⁻¹)		10,000
Interface data transfer rate (MB/S)		200
Internal data transfer rate (MB/S)		90.6 to 151.4




In the canister that is mounted with the HDD above, the DC-DC converter, which converts 56V DC to 5V and 12V DC used by the HDD, is installed.

(6) Battery

The battery to be installing in DKC515I adopts the nickel metal hydride battery which consists of kind material to the environment.

The batteries are connected with the control section (cache memories, shared memories). When the batteries appropriate to the subsystem configuration are installed, the subsystem can continue its operation with the power supplied from the batteries when the AC input power is shut off owing to a power failure and duration of the power stoppage is not longer than 20m seconds. When duration of the power stoppage is longer than 20m seconds, the subsystem executes the backup process.

One minute UPS and De-stage function is not supported in DKC515I.



*1: When the power is recovered from a failure while the backup power is supplied from the battery, the subsystem operates depending on the status of the Auto-Power-On jumper position.

ENABLE position: The subsystem is powered on automatically.

DISABLE position: The subsystem is powered on through an operation of the Power ON/OFF switch.

The batteries must be added according to the subsystem configuration. Required numbers of the additional battery options (DKC-F515I-LGABR) are shown in the following table.

	Additional conditions	Required number of DKC-F515I-LGABR
1	Add one of these components when one of the following two conditions are met: • Cache memory (DKC-F510I-C4G/C4GR/C8GR) quantity are 4 sets or greater. • Shared memory (DKC-F510I-S1G/S1GR/S2GR) quantity are 4 sets or greater.	1 set added
2	Add one of these components when the following condition is met. • Cache memory (DKC-F510I-C4G/C4GR/C8GR) quantity are 8 sets.	1 set added

CAUTION:

Supplying power to cache memory and shared memory and Battery charge are not executed until POWER ON procedure of Disk Subsystem is performed, even if AC power is being received.

When Disk Subsystem goes down without POWER-OFF procedure due to a power failure, data in the cache memory and the shared memory are maintained by the battery even though the power failure is recovered.

If maximum battery backup time, 36 hours as a full charge of 24 hours, is exceeded, the data will be volatilized.

Please execute POWER-ON procedure of Disk Subsystem before the battery backup time is exceeded.

(7) SVP

A notebook PC on the market was used as an SVP for the former subsystem, whereas an exclusive SVP for RAID was developed for the DKC515I.

By virtue of the newly developed SVP, a model change of the PC became unnecessary and the loads formerly required for the structure review and evaluation of the SVP microprogram and maintenance tools are reduced sharply. Furthermore, 3.5-inch disk drive adopted as the built-in disk drive makes reliability of the SVP enhanced.

The SVP High Reliability Support Kit is not supported.

Specification of custom SVP

Item	Specifications
OS	Windows XP Professional
CPU	Celeron 2.0GHz
Internal Memory	512MB
Key board	— *1
Display	— *1
Disk Drive	20GB or more (3.5" HDD)
LAN-1	On-Board 10Base-T/ 100Base-TX
LAN-2	On-Board 10Base-T/ 100Base-TX
Modem	PCMCIA (Modem Card not include) *2
Device	FDD: 1.44MB
	CD-ROM
Serial Port	RS232C
Other	USB × 2

*1: Since the SVP has neither a display nor a keyboard, the service personnel must bring CE Laptop PC that satisfy the following specifications with them and connect it to the SVP for exclusive use to perform installation or maintenance of the subsystem.

*2: Since the SVP has no Modem card, the Modem card must be prepared separately when connecting the subsystem to the service center via the telephone line.
In addition, the recommendation modem card is provided by the model name of DKC-F510I-MDMR.

Specification of CE Laptop PC

Item	Specifications
OS	Windows9x /NT4.0 /2000*3 WindowsXP
CPU	Pentium/Celeron 300MHz or upper
Memory	Windows9x: 64MB or upper WindowsNT /2000 /XP: 128MB or upper
Disk Drive	Available hard disk space: 100MB or upper
Display	800x600 (SVGA) or higher-resolution 1024X768 (XGA) Recommendation
CD-ROM	Need
LAN	Ethernet 10Base-T/100Base-TX


*3: Remote Desktop Connection Software is required.

(8) Disk Chassis


In the Disk Chassis installed in the RACK frame, one HDU contains 15 HDDs are installed.

The HDDs in the Disk Chassis are connected to the MIX via the FSWs with the FCAL

As to the Disk Chassis installed in the RACK, the FSW is installed in it as the standard and the HDDs are controlled by four fibre channel ports in the MIX pair.



Connections inside the HDU (15 HDDs/FCAL)



2.3 Subsystem Specifications

Subsystem specifications are shown in the following table.

Table 2.3-1 Subsystem specifications

Item			Specifications
Subsystem	Number of disk drives	Minimum	0
		Maximum	240
	RAID level		RAID6/RAID5/RAID1
	RAID group configuration	RAID6	6D+2P
		RAID5	3D+1P/7D+1P
		RAID1	2D+2D
	Maximum number of spare disk drives		16 (*1)
	Maximum number of volumes		16,384
	Subsystem capacity (Physical capacity)	400GB HDD	94.5 TB
		300GB HDD	69.2 TB
146GB HDD		34.5 TB	
72GB HDD		17.1 TB	
Internal path	Architecture		Hi-Star Net
	Maximum Bandwidth	Cache Path	8.5GB/s
		SM Path	3.6GB/s
Memory	Cache memory capacity (Increment size)		4GB ~ 32GB (4GB)
Device I/F	DKC-DKU interface		Fibre (FC-AL) / Dual Port
	Data transfer rate (MB/S)		Max. 200
	Maximum number of HDD/FC-AL		60
	Maximum number of MIX-PCB (DKA)		1
Channel I/F	Support channel option	Mainframe	Serial channel: 16S/16SR 1-2Gbps Fibre channel: 8MS/8MSR/8ML/8MLR/16MSR/16MLR 1-4Gbps Fibre channel: 16MFSR/16MFLR/16MFL4R
		Open system	1-2Gbps Fiber Short Wavelength (*2): 8HSR/16HS/16HSR/32HS/32HSR 1-4Gbps Fiber Short Wavelength (*3): 32FSR/8FS2R/16FS2R/32FS2R 1.25Gbps NAS Short Wavelength: 8NS/8NSR 1.25Gbps iSCSI Short Wavelength: 8ISR
	Data transfer rate (MB/s)	Serial channel	17
		MF Fibre channel	100 / 200 / 400
		Fibre channel	100 / 200 / 400
		NAS/iSCSI channel	100
	Maximum number of CHA options		1
Power	AC Input	1 Phase	60Hz: 200V, 208V or 230V 50Hz: 200V, 220V, 230V or 240V
Dimension	W × D × H (mm)	19-inch Rack	610 (*4) × 925 × 1,920
Non stop maintenance	Control PCB		Support
	CM/SM memory module		Support
	Power Supply, Fan		Support
	Battery		Support
	Microcode		Support
	Disk drive		Support

*1: Available as spare or data disks.

*2: If SFP Transceiver of the fibre port on CHA is exchanged for DKC-F510I-1HL/1HLR, it can be used as Long Wavelength.

*3: If SFP Transceiver of the fibre port on CHA is exchanged for DKC-F510I-1FLR/1FL4R, it can be used as Long Wavelength.

*4: This includes the thickness of side covers. (12.5mm × 2)

2.4 Physical Specifications

2.4.1 DKC515I Physical Specifications

DKC515I physical specifications are shown in the following table.

Table 2.4.1-1 DKC515I physical specifications

No.	Model Number	Weight (kg)	Heat Output (kW)	Power Consumption (kVA)	Dimension (mm)			Air Flow (m ³ /min.)
					Width	Depth	Height	
1	DKC515I-5/5R	82.0	0.638	0.658	480	648	703 ^{*1}	7.0
2	DKC-F515I-FRMR	180	—	—	610	925	1920	—
3	DKC-F515I-B2R	66.0	0.100	0.103	480	648	522	8.0
4	DKC-F515I-BCHR	4.3	—	—	—	—	—	—
5	DKC-F515I-BUHR	3.0	—	—	—	—	—	—
6	DKC-F515I-BCPR	4.3	—	—	—	—	—	—
7	DKC-F515I-BUPR	3.0	—	—	—	—	—	—
8	DKC-F515I-FIPR	0.4	—	—	—	—	—	—
9	DKC-F515I-BCSR	4.3	—	—	—	—	—	—
10	DKC-F515I-BUSR	3.0	—	—	—	—	—	—
11	DKC-F515I-MIX2R	5.0	0.370	0.381	—	—	—	—
12	DKC-F515I-MIX4R	5.0	0.380	0.392	—	—	—	—
13	DKC-F515I-LGABR	8.0	0.044	0.045	—	—	—	—
14	DKC-F510I-C4G/C4GR	0.2	0.014	0.014	—	—	—	—
15	DKC-F510I-S1G/S1GR	0.06	0.011	0.012	—	—	—	—
16	DKC-F510I-C8GR	0.2	0.014	0.014	—	—	—	—
17	DKC-F510I-S2GR	0.06	0.014	0.014	—	—	—	—
18	DKC-F515I-UC0R	2.5	—	—	—	—	—	—
19	DKC-F515I-UC1R	1.5	—	—	—	—	—	—
20	DKC-F515I-UC4R	1.5	—	—	—	—	—	—
21	DKC-F515I-EXC0R	4.3	—	—	—	—	—	—
22	DKC-F515I-EXC2R	4.3	—	—	—	—	—	—
23	DKC-F510I-16S/16SR	5.4	0.313	0.323	—	—	—	—
24	DKC-F510I-8MS/8MSR	5.7	0.346	0.356	—	—	—	—
25	DKC-F510I-8ML/8MLR	5.7	0.346	0.356	—	—	—	—
26	DKC-F510I-16MSR	5.6	0.35	0.361	—	—	—	—
27	DKC-F510I-16MLR	5.6	0.35	0.361	—	—	—	—
28	DKC-F510I-8HSR	4.6	0.273	0.281	—	—	—	—
29	DKC-F510I-16HS/16HSR	5.0	0.287	0.296	—	—	—	—
30	DKC-F510I-32HS/32HSR	5.8	0.382	0.394	—	—	—	—
31	DKC-F510I-32FSR	5.8	0.382	0.394	—	—	—	—

(To be continued)

(Continued from the preceding page)

No.	Model Number	Weight (kg)	Heat Output (kW)	Power Consumption (kVA)	Dimension (mm)			Air Flow (m ³ /min.)
					Width	Depth	Height	
32	DKC-F510I-8FS2R	4.6	0.273	0.281	—	—	—	—
33	DKC-F510I-16FS2R	5.0	0.287	0.296	—	—	—	—
34	DKC-F510I-32FS2R	5.8	0.382	0.394	—	—	—	—
35	DKC-F510I-16MFLR	5.7	0.36	0.371	—	—	—	—
36	DKC-F510I-16MFSR	5.7	0.36	0.371	—	—	—	—
37	DKC-F510I-16MFL4R	5.7	0.36	0.371	—	—	—	—
38	DKC-F510I-1HL/1HLR	0.02	—	—	—	—	—	—
39	DKC-F510I-1HS/1HSR	0.02	—	—	—	—	—	—
40	DKC-F510I-1FLR	0.02	—	—	—	—	—	—
41	DKC-F510I-1FSR	0.02	—	—	—	—	—	—
42	DKC-F510I-1FL4R	0.02	—	—	—	—	—	—
43	DKC-F510I-8NS/8NSR	6.5	0.446	0.46	—	—	—	—
44	DKC-F510I-8ISR	4.2	0.198	0.204	—	—	—	—
45	DKC-F510I-MDMR	0.04	0.006	0.006	—	—	—	—
46	DKU-F505I-72KS	1.1	0.022	0.024	—	—	—	—
47	DKU-F505I-72KSR	1.0	0.022	0.024	—	—	—	—
	DKU-F505I-72K1R							
48	DKU-F505I-146JS/146JSR	1.0	0.023	0.025	—	—	—	—
	DKU-F505I-146J1R							
49	DKU-F505I-146KSR	1.0	0.023	0.025	—	—	—	—
	DKU-F505I-146K1R							
50	DKU-F505I-300JSR	1.0	0.023	0.025	—	—	—	—
	DKU-F505I-300J1R							
51	DKU-F505I-300KSR	1.0	0.023	0.025	—	—	—	—
52	DKU-F505I-400JSR	1.0	0.023	0.025	—	—	—	—

*1: This includes the thickness of SVP (43.2mm).

2.5 Power Specifications

2.5.1 Subsystem Power Specifications

The inrush current and leakage current of DKC515I are shown in the following table.


Rack	AC INPUT Location No.	Input Power	Inrush Current			Leakage Current
			1st (0-p)	2nd (0-p)	1st (0-p) Time (-25%)	
Primary Rack	00	1-phase/20A	116.8A	38.1A	0.2ms	2.2mA
	01	1-phase/20A	35.6A	33.0A	0.4ms	1.7mA
	10,11	1-phase/20A respectively	35.6A	33.0A	0.3ms	1.2mA
Secondary Rack	20,21,30,31	1-phase/20A respectively	25.4A	24.1A	0.2ms	1.0mA

2.5.2 Steady Current


The steady current of DKC515I are shown in the following table.

This current isn't MAX current. If you need Max. current, you have to calculate by Physical Specifications sheet.

Frame	Input Power	Volt Freq. 50/60Hz	Steady Current (A rms)
DKC515I	1-pase/30A (LOGIC BOX) (SVP)	200V	2.4
			0.9
	1-pase/30A (HDD BOX)	200V	7.9




*1: HDS: HDS prepares the AC Power Cables and PDUs.
 HP: HP prepares the AC Power Cables.
 PDUs are included in option (DKC-F515I-FRMR).



*1: HDS: HDS prepares the AC Power Cables and PDUs.
 HP: HP prepares the AC Power Cables.
 PDUs are included in option (DKC-F515I-FRMR).



2.5.3 Notes of Power Supply

(1) Power Specifications of PDU.



(2) Adoption of the Breakers

As for DKC515I, a AC power is supplied to each PDU from the breaker of 20A of single phases.


When Primary Rack is usedWhen Secondary Rack is used

(3) Connection to the Power Facility

AC input to a subsystem has duplex structure per PDU. When installing a subsystem, don't mistake connection of AC cable which connects PDU and power distribution panel.

If connection of AC cable is mistaken, a system failure is caused when one of AC inputs intercept.

When AC input line is connected to direct Power Facility




*1: The output of two PDUs supplies electric power to the whole DKC515I.

*2: When PDP1 breaks, since the output of two PDUs cannot supply the whole DKC515I, it causes a system failure.

PDP: Power Distribution Panel

When AC input line is connected to UPS



PDP: Power Distribution Panel

2.6 Environmental Specifications


The environmental specifications are shown in the following table.

Item	Condition		
	Operating ^{*1}	Non-operation ^{*2}	Shipping & Storage ^{*3}
Temperature (°C)	16 to 32	-10 to 43	-25 to 60
Relative Humidity (%) ^{*4}	20 to 80	8 to 90	5 to 95
Max. Wet Bulb (°C)	26	27	29
Temperature Deviation (°C/hour)	10	10	20
Vibration ^{*5}	5 to 10Hz: 0.25mm 10 to 300Hz: 0.49m/s ²	5 to 10Hz: 2.5mm 10 to 70Hz: 4.9m/s ² 70 to 99Hz: 0.05mm 99 to 300Hz: 9.8m/s ²	4.9m/s ² , 15min. At four most severe resonance between 5 to 200Hz ^{*6}
Shock	—	78.4m/s ² , 15ms	Horizontal: Incline Impact 1.22m/s ^{*7} Vertical: Rotational Edge 0.1m ^{*8}
Acoustic level ^{*9}	65dB	—	—
Earthquake resistance (m/s ²)	up to 2.5 (*10)	—	—



- *1: Environmental specification for operating condition should be satisfied before the disk subsystem is powered on. Maximum temperature of 32°C should be strictly satisfied at air inlet portion.
Recommended temperature range is 21 to 24°C.
- *2: Non-operating condition includes both packing and unpacking conditions unless otherwise specified.
- *3: On shipping/storage condition, the product should be packed with factory packing.
- *4: No condensation in and around the drive should be observed under any conditions.
- *5: The above specifications of vibration apply to all three axes.
- *6: See ASTM D999-01 The Methods for Vibration Testing of Shipping Containers.
- *7: See ASTM D5277-92 Test Method for Performing Programmed Horizontal Impacts Using an Inclined Impact Tester.
- *8: See ASTM D6055-96 Test Methods for Mechanical Handling of Unitized Loads and Large Shipping Cases and Crates.
- *9: Measurement Condition: The point 1m far from floor and surface of the product.
- *10: Time is 5 seconds or less in case of the testing with device resonance point (4~5Hz).

2.7 Layout

2.7.1 DKC515I Layout



2.7.2 Twin Rack Configuration of Disk Subsystem

TOP VIEWFRONT VIEW

3. Internal Operation

3.1 Hardware Block Diagram

The normal configuration of DKC515I is divided into the power supply section, DKC section, and DKU section in which disk drives are installed.

The Disk-less configuration without DKU section that has only the external volumes (*1) is divided into the power supply section and DKC section.


*1: Refer to 3.20 Universal Volume Manager (UVM).

The Power Supply section consists of AC-DC power supplies and Battery Boxes.


The DKC section consists of channel adapters (CHAs), disk adapters (DKAs), cache memories (CACHEs), shared memories (SMs) and disk units (DKUs). Each component is connected with the cache paths and SM paths.

For the normal configuration, the DKC section and DKU section are connected with the disk paths.

The normal configuration is shown below.



The Disk-less configuration is shown below.



3.2 Software Organization




Fig. 3.2-1 Software Organization

Real Time OS:

A basic OS for controlling the RISC processor. Its primary tasks are to control and switch between the main tasks and SVP communication tasks.

Main tasks:

Made up of DKC control tasks (CHA Prog, CHF Prog, DMP, DSP) and the SC kernel tasks that supervise the DKC control tasks. They switch the control tasks by making use of the SC kernel's task switching facility.

SVP communication task:

Controls the communication with the SVP.

LCP (Link Control Program):

Controls the optical channel links.

HTP (Hyper Transfer Program):

Controls the FICON channel links.

CHA (Prog):

Is a channel command control layer that processes channel commands and controls cache and data transfer operations. It is located in the CHA. CHA Prog is recognized by the logical volume number and logical block number.

CHF (Prog):

Is an open channel command control layer that processes open channel commands and controls cache and data transfer operations. It is located in the CHF. CHF Prog is recognized by the logical volume number and logical block number.

DMP (Disk Master Program):

RAID control functions. DMP is located in the DKA. DMP is recognized by the logical volume number and logical block number.

DSP (Disk Slave Program):

Is a Fibre drive control layer and provides Fibre control, drive data transfer control, and parity control functions. It is located in the DKA. DSP is recognized by the physical volume number and LBA number.

Shared memory:

Stores the shared information about the subsystem and the cache control information (director names). This type of information is used for the exclusive control of the subsystem. Like CACHE, shared memory is controlled as two areas of memory and is fully non-volatile (time sustained during power failure dependant on configuration, de-stage or backup mode). The size of shared memory must be 20 MB for 1 GB of cache.

SM PATH (Shared Memory Access Path):

Access Path from the processors of CHA, CHF, DKA, PCB to Shared Memory.

3.3 Data Formats

(1) Data Conversion Overview

Since the disk subsystem uses SCSI drives, data in the CKD format are converted to the FBA format on an interface before being written on the drives. The data format is shown in Fig. 3.3-1.

CKD-to-FBA conversion is carried out by the CHA. Data is stored in cache (in the DKC) in the FBA format. Consequently, the drive need not be aware of the data format when transferring to and receiving data from cache.

Each field of the CKD-format record is left-justified and the data is controlled in units of 528-byte subblocks (because data is transferred in 16-byte units). Each field is provided with data integrity code (LRC). An address integrity code (LA: logical address) is appended to the end of each subblock. A count area (C area) is always placed at the beginning of the subblock.

Four subblocks make up a single block. The first subblock of a block is provided with T information (record position information).

If a record proves not to fit in a subblock during CKD-to-FBA conversion, a field is split into the next subblock when it is recorded. If a record does not fill a subblock, the subblock is padded with 00s, from the end of the last field to the LA.

On a physical drive, data is recorded data fields in 520-byte units (physical data format). The format of the LA in the subblock in cache is shown in Fig. 3.3-1. The last 8 bytes of the LA area are padding data which is insignificant (the reason for this is because data is transferred to cache in 16 byte units). When data is transferred to a drive from cache, the last 8 bytes of each LA area are discarded and 520 bytes are transferred.




Fig. 3.3-1 Data Format

(2) Block format




Fig. 3.3-2 Block Format

The RAID system records T information for each block of 4 subblocks as positional information that is used during record search. This unit of data is called a block.

1 block = 4 subblocks = 2 KB

The T information is 16 bytes long. However, only two bytes have meaning and the remaining 14 byte positions are padded with 0s. The reason for this is the same as that for the LA area. Unlike the LA, the insignificant bytes are also stored on the drive as are.

As seen from Fig. 3.3-2, the T information points to the closest count area in its block in the form of an SN (segment number). The drive computes the block number from the sector number with the SET SECT and searches the T information for the target block. From the T information, the drive computes the location of the closest count area and starts processing the block at the count area. This means that the information plays the role of the AM of the conventional disk storage.

(3) Data integrity provided




Fig. 3.3-3 Outline of Data Integrity

In the DKC and DKU system, a data integrity code is appended to the data being transferred at each component as shown in Fig. 3.3-3. Since data is striped onto two or more disk devices and the address integrity code is also appended. The data integrity codes are appended by hardware and the address integrity codes by software.

3.4 Cache Management

Since the DKC requires no through operation, its cache system is implemented by two memory areas called cache A and cache B so that write data can be duplexed. To prevent data loss due to power failures, cache is made non-volatile by being fully battery-backed (48 hours). This dispenses with the need for the conventional NVS.

The minimum unit of cache is the segment. Cache is destaged in segment units. Emulation Disk type at one or four segments make up one slot. The read and write slots are always controlled in pair. Cache data is enqueued and dequeued usually in slot units. In real practice, the segments of the same slot are not always stored in a contiguous area in cache, but are stored in discreet areas. These segments are controlled using CACHE-SLCB and CACHE-SGCB so that the segments belonging to the same slot are seemingly stored in a contiguous area in cache.




Fig. 3.4-1 Cache Data Structure

For increased directory search efficiency, a single virtual device (VDEV) is divided into 16-slot groups which are controlled using VDEV-GRPP and CACHE-GRPT.


1 cache segment = 24 blocks = 96 subblocks = 48 KB
 29 blocks = 116 subblocks = 58 KB
 32 blocks = 128 subblocks = 64 KB

1 slot = 1 stripe = 1 segment = 48 KB = OPEN-X (where X is other than OPEN-V)
 1 segment = 58 KB = 3390-X mainframe volume
 1 segment = 64 KB = OPEN-V on RAID450
 4 segments = 256 KB = OPEN-V on RAID500

The directories VDEV-GRPP, CACHE-GRPT, CACHE-SLCB, and CACHE-SGCB are used to identify the cache hit and miss conditions. These control tables are stored in the shared memory.

In addition to the cache hit and miss control, the shared memory is used to classify and control the data in cache according to its attributes. Queues are something like boxes that are used to classify data according to its attributes.

Basically, queues are controlled in slot units (some queues are controlled in segment units). Like SLCB-SGCB, queues are controlled using a queue control table so that queue data of the seemingly same attribute can be controlled as a single data group. These control tables are briefly described below.

(1) Cache control tables (directories)**LDEV-DIR (Logical DEV-directory):**

Contains the shared memory addresses of VDEV-GRPPs for an LDEV. LDEV-DIR is located in the local memory in the CHA.

VDEV-GRPP (Virtual DEV-group Pointer):

Contains the shared memory addresses of the GRPTs associated with the group numbers in the VDEV.

GRPT (Group Table):

A table that contains the shared memory address of the SLCBs for 16 slots in the group. Slots are grouped to facilitate slot search and to reduce the space for the directory area.

SLCB (Slot Control Block):

Contains the shared memory addresses of the starting and ending SGCBs in the slot. One or more SGCBs are chained. The SLCB also stores slot status and points to the queue that is connected to the slot. The state transitions of clean and dirty queues occur in slot units. The processing tasks reserve and release cache areas in this unit.

SGCB (Segment Control Block):

Contains the control information about a cache segment. It contains the cache address of the segment. It is used to control the staged subblock bit map, dirty subblock bitmap, and other information. The state transitions of only free queues occur in segment units.

(2) Cache control table access method (hit/miss identification procedure)




Fig. 3.4-2 Outline of Cache Control Table Access

1. The current VDEV-GRPP is referenced through the LDEV-DIR to determine the hit/miss condition of the VDEV-groups.
2. If a VDEV-group hits, CACHE-GRPT is referenced to determine the hit/miss condition of the slots.
3. If a slot hits, CACHE-SLCB is referenced to determine the hit/miss condition of the segments.
4. If a segment hits, CACHE-SGCB is referenced to access the data in cache.

If a search miss occurs during the searches from 1. through 4., the target data causes a cache miss.

Definition of VDEV number

Since the host processor recognizes addresses only by LDEV, it is unaware of the device address of the parity device. Accordingly, the RAID system is provided with a VDEV address which identifies the parity device associated with an LDEV. Since VDEVs are used to control data devices and parity devices systematically, their address can be computed using the following formulas:

Data VDEV number = LDEV number

Parity VDEV number = 1024 + LDEV number

From the above formulas, the VDEV number ranges from 0 to 2047.

(3) Queue structures

The DKC and DKU uses 10 types of queues to control data in cache segments according to its attributes. These queues are explained below.

- CACHE-GRPT free queue

This queue is used to control segments that are currently not used by CACHE-GRPT (free segments) on an FIFO (First-In, First-Out) basis. When a new table is added to CACHE-GRPT, the segment that is located by the head pointer of the queue is used.

- CACHE-SLCB free queue

This queue is used to control segments that are currently not used by CACHE-SLCB (free segments) on an FIFO basis. When a new slot is added to CACHE-SLCB, the segment that is located by the head pointer of the queue is used.

- CACHE-SGCB free queue

This queue is used to control segments that are currently not used by CACHE-SGCB (free segments) on an FIFO basis. When a new segment is added to CACHE-SGCB, the segment that is located by the head pointer of the queue is used.

- Clean queue

This queue is used to control the segments that are reflected on the drive on an LRU basis.

- Bind queue

This queue is defined when the bind mode is specified and used to control the segments of the bind attribute on an LRU basis.

- Error queue

This queue controls the segments that are no longer reflected on the drive due to some error (pinned data) on an LRU basis.

- Parity in-creation queue

This queue controls the slots (segments) that are creating parity on an LRU basis.

- DFW queue (host dirty queue)

This queue controls the segments that are not reflected on the drive in the DFW mode on an LRU basis.

- CFW queue (host dirty queue)

This queue controls the segments that are not reflected on the drive in the CFW mode on an LRU basis.

- PDEV queue (physical dirty queue)

This queue controls the data (segments) that are not reflected on the drive and that occur after a parity is generated. Data is destaged from this queue onto the physical DEV. There are 32 PDEV queues per physical DEV.

The control table for these queues is located in the shared memory and points to the head and tail segments of the queues.

(4) Queue state transitions

Fig. 3.4-3 shows the state transitions of the queues used in. A brief description of the queue state transitions follows.

- State transition from a free queue

When a read miss occurs, the pertinent segment is staged and enqueued to a clean queue.

When a write miss occurs, the pertinent segment is temporarily staged and enqueued to a host dirty queue.

- State transition from a clean queue

When a write hit occurs, the segment is enqueued to a host dirty queue. Transition from clean to free queues is performed on an LRU basis.

- State transition from a host dirty queue

The host dirty queue contains data that reflects no parity. When parity generation is started, a state transition occurs to the parity in-creation queue.

- State transition from the parity in-creation queue

The parity in-creation queue contains parity in-creation data. When parity generation is completed, a transition to a physical dirty queue occurs.

- State transition from a physical dirty queue

When a write hit occurs in the data segment that is enqueued in a physical dirty queue, the segment is enqueued into the host dirty queue again. When destaging of the data segment is completed, the segment is enqueued into a queue (destaging of data segments occur asynchronously on an LRU basis).





Fig. 3.4-3 Queue Segment State Transition Diagram

(5) Cache usage in the read mode



The cache area to be used for destaging read data is determined depending on whether the result of evaluating the following expression is odd or even:

$$(\text{CYL\#} \times 15 + \text{HD\#}) / 16$$

The read data is destaged into area A if the result is even and into area B if the result is odd.

Fig. 3.4-4 Cache Usage in the Read Mode

Read data is not duplexed and its destaging cache area is determined by the formula shown in Fig. 3.4-4. Staging is performed not only on the segments containing the pertinent block but also on the subsequent segments up to the end of track (for increased hit ratio).

Consequently, one track equivalence of data is prefetched starting at the target block. This formula is introduced so that the cache activity ratios for areas A and B are even. The staged cache area is called the cache area and the other area NVS area.

(6) Cache usage in the write mode




Fig. 3.4-5 Cache Usage in the Write Mode

This system handles write data (new data) and read data (old data) in separate segments as shown in Fig. 3.4-5 (not overwritten as in the conventional systems), whereby compensating for the write penalty.

- (1) If the write data in question causes a cache miss, the data from the block containing the target record up to the end of the track is staged into a read data slot.
- (2) In parallel with step (1), the write data is transferred when the block in question is established in the read data slot.
- (3) The parity data for the block in question is checked for a hit or miss condition and, if a cache miss condition is detected, the old parity is staged into a read parity slot.
- (4) When all data necessary for generating new parity is established, it is transferred to the DRR circuit in the DKA.
- (5) When the new parity is completed, the DRR transfers it into the write parity slots for cache A and cache B (the new parity is handled in the same manner as the write data).

The reason for writing the write data into both cache areas is that data will be lost if a cache error occurs when it is not yet written on the disk.

Although two cache areas are used as explained above, the read data (including parity) is staged into either cache A or cache B simply by duplexing only the write data (including parity) (in the same manner as in the read mode).

(7) CFW-inhibited write-operation (with Cache single-side error)

The non RAID-type Disk systems write data directly onto disk storage in the form of cache through, without performing a DFW, when a cache error occurs. In this system, cache must always be passed, which fact disables the through operation. Consequently, the write data is duplexed, and a CFW-inhibited write-operation is performed; that is, when one cache subsystem goes down, the end of processing status is not reported until the data write in the other cache subsystem is completed. This process is called CFW-inhibited write-operation.

The control information necessary for controlling cache is stored in the shared memory.

(8) Shared memory




Fig. 3.4-6 Outline of Shared Memory

This system has two areas of cache memory, shown in Fig. 3.4-6, as it has two areas of cache memory. One part of its internal data is fully duplexed (this serves as the role of the conventional ECM). The other part of the shared memory area contains the control information about the corresponding cache area (shared memory area A for cache A and shared memory area B for cache B). If an error occurs on one side of shared memory (A or B), the corresponding cache area becomes inoperative (equivalent to a cache error).

Like cache, shared memory is made non-volatile (the time dependant on configuration, de-stage or backup mode) to prevent data loss in case of power failures.

3.5 Destaging Operations

(1) Cache management in the destage mode (RAID5)

Destaging onto a drive is deferred until parity generation is completed. Data and parity slot transitions in the destage mode occur as shown in Fig. 3.5-1.




Fig. 3.5-1 Cache Operation in the Destage Mode

Write data is stored in write segments before parity is generated but stored in read segments after parity is generated. When drive data is stored, therefore, the data from the read segment is transferred.

- (2) Cache management in the destage mode (RAID1)
Data slot is destaged to primary/secondary drive.




Fig. 3.5-2 RAID1 asynchronous destage

- ① Destage to primary drive.
- ② Destage to secondary drive.
- ③ The data read segment in the NVS area is released.

(3) Blocked data write

The purpose of blocked data write is to reduce the number of accesses to the drive during destaging, whereby increasing the subsystem performance. There are three modes of blocked data write: single-stripe blocking, multiple-stripe blocking, and drive blocking. These modes are briefly explained below.

- Single-stripe blocking

Two or more dirty segments in a stripe are combined into a single dirty data block.

Contiguous dirty blocks are placed in a single area. If an unloaded block exists between dirty blocks, the system destages the dirty blocks separately at the unloaded block. If a clean block exists between dirty blocks, the system destages the blocks including the clean block.

- Multiple-stripe blocking

The sequence of stripes in a parity group are blocked to reduce the number of write penalties. This mode is useful for sequential data transfer.

- Drive blocking

In the drive blocking mode, blocks to be destaged are written in a block with a single drive command if they are contiguous when viewed from a physical drive to shorten the drive's latency time.

The single- and multiple-stripe blocking modes are also called in-cache blocking modes. The DMP determines which mode to use. The drive blocking mode is identified by the DSP.

3.6 Operations Performed when Drive Errors Occur


(1) I/O operations performed when drive errors occur

This system can recover target data using parity data and data stored on normal disk storage even when it cannot read data due to errors occurring on physical drives. This feature ensures non-disruptive processing of applications in case of drive errors. This system can also continue processing for the same reason in case errors occur on physical drives while processing write requests.


Fig. 3.6-1 shows the outline of data read processing in case a drive error occurs.

Request for reading data B

(i) Normal time



(ii) When a disk error occurs



A, B, C . . . : Data (A=A', B=B', C=C')

P : Parity data

Fig. 3.6-1 Outline of Data Read Processing

(2) Data integrity feature and drive errors

This system uses spare disk drives and reconfigures any drives that are blocked due to errors or drives whose error count exceeds a specified limit value using spare disks. (Drives belonging to a Parity Group with no LDEVs defined are not reconfigured.)

Since this processing is executed on the host in the background, this system can continue to accept I/O requests. The data saved on spare disks are copied into the original location after the error drives are replaced with new ones.

1. Dynamic sparing

This system keeps track of the number of errors that occurred, for each drive, when it executes normal read or write processing. If the number of errors occurring on a certain drive exceeds a predetermined value, this system considers that the drive is likely to cause unrecoverable errors and automatically copies data from that drive to a spare disk. This function is called dynamic sparing. In RAID1 method, this system is same as RAID5 dynamic sparing.




Fig. 3.6-2 Outline of the Dynamic Sparing Function

2. Correction copy

When this system cannot read or write data from or to a drive due to an error occurring on that drive, it regenerates the original data for that drive using data from the other drives and the parity data, and copies it onto a spare disk. In RAID1 method, this system copies data from the another drive to a spare disk.

In the case of RAID 6, the correction copy can be made to up to two disk drives in a parity group.




Fig. 3.6-3 Outline of the Correction Copy Function

3. Allowable number of copying operations

Table 3.6-1 Allowable number of copying operations

RAID level	Allowable number of copying operations
RAID1	Either the dynamic sparing or correction copy can be executed within a RAID pair.
RAID5	Either the dynamic sparing or correction copy can be executed within a parity group.
RAID6	The dynamic sparing and/or correction copy can be executed up to a total of twice within a parity group.

3.7 Inter Mix of Drives and Emulation types

3.7.1 Drives to be Connected

The models of disk units which are connectable with the RAID500 disk subsystem and the specifications of each disk unit are shown in Table 2.3-1 ([THEORY02-240](#)).

The RAID500 disk subsystem can connect up to 1152 disk drives mentioned above, though the number of connectable disk drives varies with the emulation types and the RAID configuration. These will be explained in detail in Section 3.7.2.

SVP displays each drive model as the following table.

Disk drive model	SVP screen
DKS2C-K72FC	DKS2C-K072FC
DKS2D-K72FC	DKS2D-K072FC
DKS2E-K72FC	DKS2E-K072FC
DKS2F-K72FD	DKS2F-K072FC
DKS2G-K72FD	DKS2G-K072FC
DKR2F-K72FC	DKR2F-K072FC
DKR2G-K72FC	DKR2G-K072FC
DKR2J-K72FD	DKR2J-K072FC
DKS2D-K146FC	DKS2D-K146FC
DKS2E-K146FC	DKS2E-K146FC
DKS2F-K146FC	DKS2F-K146FC
DKS2G-K146FD	DKS2G-K146FC
DKR2F-K146FC	DKR2F-K146FC
DKR2G-K146FC	DKR2G-K146FC
DKR2J-K146FD	DKR2J-K146FC
DKR2E-J146FC	DKR2E-J146FC
DKR2F-J146FC	DKR2F-J146FC
DKR2G-J146FC	DKR2G-J146FC
DKR2J-J146FD	DKR2J-J146FC
DKS2D-J146FC	DKS2D-J146FC
DKS2E-J146FC	DKS2E-J146FC
DKS2G-J146FD	DKS2G-J146FC
DKR2F-J300FC	DKR2F-J300FC
DKR2G-J300FC	DKR2G-J300FC
DKR2J-J300FC	DKR2J-J300FC
DKS2D-J300FC	DKS2D-J300FC
DKS2E-J300FC	DKS2E-J300FC
DKS2G-J300FC	DKS2G-J300FC
DKS2E-K300FC	DKS2E-K300FC
DKS2F-K300FC	DKS2F-K300FC
DKS2G-K300FC	DKS2G-K300FC
DKR2G-K300FC	DKR2G-K300FC
DKR2H-K300FC	DKR2H-K300FC
DKR2J-K300FC	DKR2J-K300FC
DKS2E-J400FC	DKS2E-J400FC
DKS2G-J400FC	DKS2G-J400FC

Since DKC micro code version “05-04-4x”, when installed “Disc drive models (same capacity and same rotation speed)” are intermixed in same ECC, Recommendation setting on SVP is following.

Disk drive model	Recommendation setting
DKS2C-K72FC DKS2D-K72FC DKS2E-K72FC DKS2F-K72FD DKS2G-K72FD DKR2F-K72FC DKR2G-K72FC DKR2J-K72FD	DKR2J-K072FC
DKR2E-J146FC DKR2F-J146FC DKR2G-J146FC DKR2J-J146FD DKS2D-J146FC DKS2E-J146FC DKS2G-J146FD	DKR2J-J146FC
DKS2D-K146FC DKS2E-K146FC DKS2F-K146FC DKS2G-K146FD DKR2F-K146FC DKR2G-K146FC DKR2J-K146FD	DKR2J-K146FC
DKR2F-J300FC DKR2G-J300FC DKR2J-J300FC DKS2D-J300FC DKS2E-J300FC DKS2G-J300FC	DKR2J-J300FC
DKS2E-K300FC DKS2F-K300FC DKS2G-K300FC DKR2G-K300FC DKR2H-K300FC DKR2J-K300FC	DKR2J-K300FC
DKS2E-J400FC DKS2G-J400FC	DKS2G-J400FC

When the HDD is replaced, the HDD which has the compatibility is following.

Before replacing	After replacing
DKR2x-K072FC	DKR2x-K072FC DKS2x-K072FC
DKS2x-K072FC	DKR2x-K072FC DKS2x-K072FC
DKR2x-J146FC	DKR2x-J146FC DKS2x-J146FC
DKS2x-J146FC	DKR2x-J146FC DKS2x-J146FC
DKR2x-K146FC	DKR2x-K146FC DKS2x-K146FC
DKS2x-K146FC	DKR2x-K146FC DKS2x-K146FC
DKR2x-J300FC	DKR2x-J300FC DKS2x-J300FC
DKS2x-J300FC	DKR2x-J300FC DKS2x-J300FC
DKR2x-K300FC	DKR2x-K300FC DKS2x-K300FC
DKS2x-K300FC	DKR2x-K300FC DKS2x-K300FC
DKS2x-J400FC	DKS2x-J400FC

x: C, D, E,... etc

Blank Sheet

Blank Sheet

3.7.2 Emulation Device Type

Refer to 3.5 Volume Specification in OPENPLATFORM SECTION about OPEN Volume Type.
The emulation types of disk controller and disk units of the RAID500 are shown in Tables 3.7.2-1 to 3.7.2-5.

Table 3.7.2-1 List of RAID500 Model number

Model Number	Disk drive model	RAID Level
DKU-F505I-72KS/72KSR	DKS2C-K72FC	RAID5(3D+1P,7D+1P)
	DKS2D-K72FC	/RAID1(2D+2D)
	DKS2E-K72FC	/RAID6(6D+2P)
	DKS2F-K72FD	
	DKS2G-K72FD	
	DKR2F-K72FC	
	DKR2G-K72FC	
	DKR2J-K72FD	
DKU-F505I-72JSR	DKR2E-J72FC	
	DKR2F-J72FD	
	DKR2G-J72FD	
	DKR2J-J72FD	
DKU-F505I-146JS/146JSR	DKR2E-J146FC	
	DKR2F-J146FC	
	DKR2G-J146FC	
	DKR2J-J146FD	
	DKS2D-J146FC	
	DKS2E-J146FC	
	DKS2G-J146FD	
DKU-F505I-146KSR	DKS2D-K146FC	
	DKS2E-K146FC	
	DKS2F-K146FC	
	DKS2G-K146FD	
	DKR2F-K146FC	
	DKR2G-K146FC	
	DKR2J-K146FD	
DKU-F505I-300JSR	DKR2F-J300FC	
	DKR2G-J300FC	
	DKR2J-J300FC	
	DKS2D-J300FC	
	DKS2E-J300FC	
	DKS2G-J300FC	
DKU-F505I-300KSR	DKS2E-K300FC	
	DKS2F-K300FC	
	DKS2G-K300FC	
	DKR2G-K300FC	
	DKR2H-K300FC	
	DKR2J-K300FC	
DKU-F505I-400JSR	DKS2E-J400FC	
	DKS2G-J400FC	

Note: As for RAID1, the two concatenation of a parity groups is possible (8HDDs).

In this case the number of volume become doubled.

Two concatenation and four concatenation (16HDDs and 32HDDs) of the RAID Group are possible for RAID5 (7D+1P).

In this case, the number of volumes becomes twice and four times.

Two connection: V05 support; Four connection: V06 support.

Table 3.7.2.2 List of RAID500 Emulation Types for RAID5(3D+1P)

Item		Emulation contents			
Emulation Type	DKC	3990-6/3990-6E/I-2105		I-2105	
	DKU	3390-9	3390-3/3R	3390-L*1	3390-M
Storage capacity (G byte/volume)		8.510	2.838	27.8	55.689
Number of volumes / Parity groups	DKU-F505I-72KS/72KSR/72K2R	23	71	7	3
	DKU-F505I-72JSR	24	73	7	3
	DKU-F505I-146JS/146JSR/146KSR/146K2R	48	144	14	7
	DKU-F505I-300JSR/300KSR/300K2R	96	250	29	14
	DKU-F505I-400JSR	132	250	40	20
Maximum number of Parity groups	DKU-F505I-72KS/72KSR/72K2R	59	59	59	59
	DKU-F505I-72JSR	59	59	59	59
	DKU-F505I-146JS/146JSR/146KSR/146K2R	59	59	59	59
	DKU-F505I-300JSR/300KSR/300K2R	59	59	59	59
	DKU-F505I-400JSR	59	59	59	59
Maximum number of Volumes	DKU-F505I-72KS/72KSR/72K2R	1357	4189	413	177
	DKU-F505I-72JSR	1416	4307	413	177
	DKU-F505I-146JS/146JSR/146KSR/146K2R	2832	8496	826	413
	DKU-F505I-300JSR/300KSR/300K2R	5664	14750	1711	826
	DKU-F505I-400JSR	7788	14750	2360	1180
Subsystem capacity (user area) (G byte)	DKU-F505I-72KS/72KSR/72K2R	Min	196	201	195
		Max	11548	11888	11481
	DKU-F505I-72JSR	Min	204	207	195
		Max	12050	12223	11481
	DKU-F505I-146JS/146JSR/146KSR/146K2R	Min	408	409	389
		Max	24100	24112	22963
	DKU-F505I-300JSR/300KSR/300K2R	Min	817	710	806
		Max	48201	41861	47566
	DKU-F505I-400JSR	Min	502	167	1640
		Max	1666	572	5410

Note: The DKC emulation type 3390-6, 3390-6E and I-2105 support only the emulation type 3390*.

The Emulation type 3390-3 and 3390-3R cannot be intermixed in the subsystem.

Table 3.7.2.3 List of RAID500 Emulation Types for RAID5(7D+1P)

Item		Emulation contents			
Emulation Type	DKC	3990-6/3990-6E/I-2105		I-2105	
	DKU	3390-9	3390-3/3R	3390-L*1	3390-M
Storage capacity (G byte/volume)		8.510	2.838	27.8	55.689
Number of volumes / Parity groups	DKU-F505I-72KS/72KSR/72K2R	55	167	17	8
	DKU-F505I-72JSR	57	171	17	8
	DKU-F505I-146JS/146JSR/146KSR/146K2R	112	337	34	17
	DKU-F505I-300JSR/300KSR/300K2R	225	500	69	34
	DKU-F505I-400JSR	308	500	94	47
Maximum number of Parity groups	DKU-F505I-72KS/72KSR/72K2R	28	28	28	28
	DKU-F505I-72JSR	28	28	28	28
	DKU-F505I-146JS/146JSR/146KSR/146K2R	28	28	28	28
	DKU-F505I-300JSR/300KSR/300K2R	28	28	28	28
	DKU-F505I-400JSR	28	28	28	28
Maximum number of Volumes	DKU-F505I-72KS/72KSR/72K2R	1540	4676	476	224
	DKU-F505I-72JSR	1596	4788	476	224
	DKU-F505I-146JS/146JSR/146KSR/146K2R	3136	9436	952	476
	DKU-F505I-300JSR/300KSR/300K2R	6300	14000	1932	952
	DKU-F505I-400JSR	8624	14000	2632	1316
Subsystem capacity (user area) (G byte)	DKU-F505I-72KS/72KSR/72K2R	Min	468	474	473
		Max	13105	13270	13233
	DKU-F505I-72JSR	Min	485	485	473
		Max	13582	13588	13233
	DKU-F505I-146JS/146JSR/146KSR/146K2R	Min	953	956	945
		Max	26687	26779	26466
	DKU-F505I-300JSR/300KSR/300K2R	Min	1915	1419	1918
		Max	53613	39732	53710
	DKU-F505I-400JSR	Min	238	79	778
		Max	3983	1345	13138

Note: The DKC emulation type 3390-6, 3390-6E and I-2105 support only the emulation type 3390*.

The Emulation type 3390-3 and 3390-3R cannot be intermixed in the subsystem.

Table 3.7.2.4 List of RAID500 Emulation Types for RAID1(2D+2D)

Item		Emulation contents			
Emulation Type	DKC	3990-6/3990-6E/I-2105		I-2105	
	DKU	3390-9	3390-3/3R	3390-L*1	3390-M
Storage capacity (G byte/volume)		8.510	2.838	27.8	55.689
Number of volumes / Parity groups	DKU-F505I-72KS/72KSR/72K2R	15	47	4	2
	DKU-F505I-72JSR	16	48	4	2
	DKU-F505I-146JS/146JSR/146KSR/146K2R	32	96	9	4
	DKU-F505I-300JSR/300KSR/300K2R	64	193	19	9
	DKU-F505I-400JSR	88	250	26	13
Maximum number of Parity groups	DKU-F505I-72KS/72KSR/72K2R	59	59	59	59
	DKU-F505I-72JSR	59	59	59	59
	DKU-F505I-146JS/146JSR/146KSR/146K2R	59	59	59	59
	DKU-F505I-300JSR/300KSR/300K2R	59	59	59	59
	DKU-F505I-400JSR	59	59	59	59
Maximum number of Volumes	DKU-F505I-72KS/72KSR/72K2R	885	2773	236	118
	DKU-F505I-72JSR	944	2832	236	118
	DKU-F505I-146JS/146JSR/146KSR/146K2R	1888	5664	531	236
	DKU-F505I-300JSR/300KSR/300K2R	3776	11387	1121	531
	DKU-F505I-400JSR	5192	14750	1534	767
Subsystem capacity (user area) (G byte)	DKU-F505I-72KS/72KSR/72K2R	Min	128	133	111
		Max	7531	7870	6561
	DKU-F505I-72JSR	Min	136	136	111
		Max	8033	8037	6561
	DKU-F505I-146JS/146JSR/146KSR/146K2R	Min	272	272	250
		Max	16067	16074	14762
	DKU-F505I-300JSR/300KSR/300K2R	Min	545	548	528
		Max	32134	32316	31164
	DKU-F505I-400JSR	Min	502	167	1640
		Max	1086	379	3091
					6203

Note: The DKC emulation type 3390-6, 3390-6E and I-2105 support only the emulation type 3390*.

The Emulation type 3390-3 and 3390-3R cannot be intermixed in the subsystem.

Table 3.7.2.5 List of RAID500 Emulation Types for RAID6(6D+2P)

Item		Emulation contents			
Emulation Type	DKC	3990-6/3990-6E/I-2105		I-2105	
	DKU	3390-9	3390-3/3R	3390-L*1	3390-M
Storage capacity (G byte/volume)		8.510	2.838	27.8	55.689
Number of volumes / Parity groups	DKU-F505I-72KS/72KSR/72K2R	47	143	14	7
	DKU-F505I-72JSR	48	146	14	7
	DKU-F505I-146JS/146JSR/146KSR/146K2R	96	289	29	14
	DKU-F505I-300JSR/300KSR/300K2R	193	500	59	29
	DKU-F505I-400JSR	264	500	80	40
Maximum number of Parity groups	DKU-F505I-72KS/72KSR/72K2R	28	28	28	28
	DKU-F505I-72JSR	28	28	28	28
	DKU-F505I-146JS/146JSR/146KSR/146K2R	28	28	28	28
	DKU-F505I-300JSR/300KSR/300K2R	28	28	28	28
	DKU-F505I-400JSR	28	28	28	28
Maximum number of Volumes	DKU-F505I-72KS/72KSR/72K2R	1316	4004	392	196
	DKU-F505I-72JSR	1344	4088	392	196
	DKU-F505I-146JS/146JSR/146KSR/146K2R	2688	8092	812	392
	DKU-F505I-300JSR/300KSR/300K2R	5404	14000	1652	812
	DKU-F505I-400JSR	7392	14000	2240	1120
Subsystem capacity (user area) (G byte)	DKU-F505I-72KS/72KSR/72K2R	Min	400	406	389
		Max	11199	11363	10898
	DKU-F505I-72JSR	Min	408	414	389
		Max	11437	11602	10898
	DKU-F505I-146JS/146JSR/146KSR/146K2R	Min	817	820	806
		Max	22875	22965	22574
	DKU-F505I-300JSR/300KSR/300K2R	Min	1642	1419	1640
		Max	45988	39732	45926
	DKU-F505I-400JSR	Min	238	79	778
		Max	3404	1152	10820

Note: The DKC emulation type 3390-6, 3390-6E and I-2105 support only the emulation type 3390*.

The Emulation type 3390-3 and 3390-3R cannot be intermixed in the subsystem.

Specifications for coexistence of elements

Table 3.7.2-6 shows permitted coexistence of RAID levels, HDD types, and emulation types respectively.

Table 3.7.2-6 Specifications for Coexistence of Elements

Item	Specification	Remarks
Coexistence of RAID levels	RAID1/RAID5/RAID6 can exist in DKC.	
Coexistence of numbers of HDDs composing ECC group (coexistence of configurations, 3D+1P and 6D+2P)	- The numbers of HDDs inside an ECC group are 4 HDDs or 8HDDs and they are applicable to RAID5 (7D+1P, 3D+1P), RAID1 (2D+2D) and RAID6 (6D+2P). - All HDD types support RAID5 (7D+1P, 3D+1P) and RAID1 (2D+2D) configuration. 7D+1P, 3D+1P, 2D+2D and 6D+2P can coexist within units of one MIX PCB pair.	
Coexistence of HDD types	- HDD types can coexist in each ECC group. - The specification for selecting the spare HDD can be common in the subsystem if there exists a spare HDD having the same capacity as that of the data HDD, though the HDD types coexist.	
Coexistence of emulation types	- Emulation types can coexist in each ECC group. - LDEV ID addressing must be of the same emulation type (3380-3, 3390-M/L/9/3, or 3390-3R) for every 32-address boundary.	An emulation of the same system should be set for every 32 addressings due to generation restriction by the host.

3.8 16384 logical addresses

The host connection interface specification are outlined in Tables 3.8-1 and 3.8-2.

Table 3.8-1 List of Allowable Maximum Values of Host Connection Interface Items on the DKC Side

	ESCON channel	Fibre channel
Maximum number of CUs	64	
Maximum number of SSIDs	256	
Maximum number of LDEVs	16384	

Table 3.8-2 Allowable Range of Host Connection Interface Items on DKC Side

	ESCON channel	Fibre channel
CU address	0 to 3F *1	
SSID	0004 to FFFD	
Number of logical volumes	1 to 16384	

*1: Number of CUs connectable to the one ESCON channel (CHPID) is 16. The CU addresses in the interface with a host are 00 to 0F for the ESCON channel.

Notice: If you use PPRC command and specify 0xFFXX as SSID of MCU and RCU, the command may be rejected. Please specify 0x0004 0xFEFF as SSID of MCU and RCU.

Detailed numbers of logical paths of the main frame fibre and serial channels are shown in Table 3.8-3.

Table 3.8-3 List of Numbers of Connectable Logical Paths

Item	Fibre channel	Serial channel
Number of channel ports	8 ^{*2} 16 ^{*3}	16
Max. number of logical paths per CU	2,048	512
Max. number of logical paths per port	65,536 ^{*4}	512 ^{*5}
Max. number of logical paths per channel adapter	65,536	1,024
Max. number of logical paths per system	131,072	2,048

*2: When four ports per channel adapter are installed

*3: When eight ports per channel adapter are installed

*4: The maximum number of paths for connection to a host per fibre channel port is 1024 (1024 host paths × 64 CUs = 65536 logical paths).

*5: The number of logical paths per ESCON port is up to 512 (32 host paths × 16CUs = 512 logical paths).

But, the number of connecting devices per ESCON channel port is limited up to 1024 devices.

Therefore if 256 devices are defined by CU, the number of logical paths is normally 128 (32 host paths × 4CUs = 128 logical paths).

3.9 LDEV Formatting

3.9.1 Outlines of a high-speed LDEV Formatting

DKC can LDEV-Format to two or more ECC at the same time by providing the HDD with the LDEV formatting function.

Item No.	Item	Contents
1	SVP operation	The operation is performed by selecting functions from the Maintenance menu.
2	Display of execution status	Display of the execution progress in the SVP message box (%)
3	Execution result	Normal/abnormal LDEV: Same indications as the conventional ones are displayed. Normal/abnormal PDEV: STATUS is displayed.
4	Recovery action when a failure occurs	Same as the conventional one. However, a retry is to be executed in units of ECC. (Because the LDEV-FMT is terminated abnormally in units of ECC when a failure occurs in the HDD.)
5	Operation of the SVP which is a high-speed LDEV-FMT object	When an LDEV-FMT of more than one ECC is instructed, the high-speed processing is performed.
6	PS/OFF or powering off	The LDEV formatting is suspended. No automatic restart is executed.
7	SVP PC powering off during execution of an LDEV-FMT	After the SVP is rebooted, the indication before the PC powering off is displayed in succession.
8	Execution of a high-speed LDEV-FMT in the status that the spare is saved	ECC of HDD which the spare is saved fails the high-speed LDEV-FMT, and changes to a low-speed format. (Because the low-speed format is executed after the high-speed format is completed, the format time becomes long.) After the high-speed LDEV-FMT is completed, execute the copy back of HDD which the spare is saved from SIM log and restore it.

3.9.2 Estimation of a high-speed LDEV Formatting Time

3.9.2.1 DKxxx-JxxxFC/KxxxFC

The format time of DKxxx-JxxxFC/KxxxFC doesn't depend on number of ECC, and be decided by capacity and the rotational speed of HDD.

Formatting time is indicated as follows.

HDD Capacity/rotation Speed	Formatting Time	Time Out Value ^{*2}
400GB/10krpm	approx. 120min	180min
300GB/10krpm 300GB/15krpm	approx. 90min ^{*1}	160min
146GB/10krpm	approx. 90min	150min
146GB/15krpm	approx. 45min	80min
72GB/10krpm	approx. 40min	75min
72GB/15krpm	approx. 25min	55min

*1: LDEV-Formatting time has been shorted because of the oncoming generation.

*2: The progress rate on SVP is displays as "99%" during the "Formatting Time" and the "Time Out Value".

3.9.3 Notes when logical volume format is executed

Execute the logical volume format when the drive use rate in the parity group that executes the logical volume format is a proper value (40% or less).


There is a possibility that the format processing fails by the response delay to Host and the logical volume format processing delay due to the DKC internal resource competition with Host I/O when the logical volume format is executed with the drive use rate exceeding the limit value (80% or more), which is in the overloaded status.

3.10 High-Speed Online Replacement of Microprogram

The microprograms are stored in the shared memory and transferred in a batch. Thus the number of times of the transfer from the SVP to the DKC via the LAN is reduced and the online microprogram replacement is speeded up by recovering two or more processors at the same time.

3.10.1 Outline

A microprogram storage area is reserved in the shared memory. The microprogram by which each processor operates is stored in it, and data is written into the flash memory from the shared memory by rebooting the processor. The reason for executing the writing during reboot processing is that it is intended to unify the microprogram writing processing executed in the PCB replacement and cold replacement (not supported at present). In addition, when the microprograms are stored in the shared memory, processors can execute the microprogram writing processing at the same time and the processing time can be shortened substantially.



Two or more processors execute the processing at the same time.

[Processing sequence]

- ① The SVP transfers the microprograms (DKCMAIN, RAMBOOT, LCP/LCDG, HTP/FCDG, PCDG, NBIOS) to the shared memory.
- ② The SVP executes a reboot instruction for each processor of the first group.
- ③ In the reboot processing, the microprograms are transferred from the shared memory to the flash memory and the writing processing is executed.
- ④ The SVP monitors the status of the processors executing the reboot processing and writing processing to the flash memory and waits until the status changes to normal.
- ⑤ The SVP executes the processing of steps ② to ④ repeatedly for each group.

3.10.2 Processing Time

The estimated time of microprogram replacement is shown below. The time below is required for the process for the programs DKCMAIN, RAMBOOT LCDG, HTP, FCDG, PCDG, NASBIOS and LCP. The SVP and DKU are processed in the conventional processing sequences and the time required for the sequences are not included in the processing time.


Processing time required for microprogram replacement
RAID500 method (Estimated time for 4 groups)
20 min. / 64 processors

3.11 HRC

3.11.1 HRC Components

There are two modes of the update copy, a synchronous copy and asynchronous copy. Description of the asynchronous is described in chapter 3.15 HRC Asynchronous.

(1) HRC Components



*: Serial Interface for MCU-RCU paths is not Supported.

Fig. 3.11.1-1 HRC Components for Serial Interface Connection




Fig. 3.11.1-2 HRC Components for Fibre-Channel Interface Connection

(a) HRC Volume Pair

An **HRC volume pair** consists of two logical volumes, an M-VOL and an R-VOL, in different DKC515I subsystems.

An **M-VOL** (main volume) is a primary volume. It can be read or written by I/O operations from host processors.

An **R-VOL** (remote volume) is a secondary or a mirrored volume. Under control of the DKC515I subsystems, contents of an M-VOL and updates from host processors are copied to an R-VOL. Read or write I/O operations from host processors to R-VOLs are rejected.

Note : R-VOL Read Only function

HRC has R-VOL Read Only function to accept read commands to R-VOL of suspended pairs of HRC.

R-VOL Read Only function becomes effective with SVP system option setting for RCU of HRC.

With this function, RCU accepts all RD commands including CTL/SNS commands and WR command to cylinder zero, head zero, record three of R-VOL. (It is necessary to change VOLSER of the volume.)

And, when it is combined with another system option, RCU accepts all RD commands including CTL/SNS commands and WR command to all tracks, and all records in cylinder zero of R-VOL. (It is necessary to change VOLSER and VTOC of the Volume.)

The RCU rejects some PPRC commands such as ADDPAIR to the R-VOL nevertheless the status of the R-VOL looks 'Simplex'. They must be controlled by system administration.

With this function, RCU displays the status of the R-VOL as 'Simplex' instead of 'Suspended'. It is necessary to accept I/O to R-VOL.

MCU copies cylinder zero of the pair at RESYNC copy unconditionally, besides the ordinary RESYNC copy.

With this function, if DKC Emulation type is 2105/2107, CSUSPEND command to R-VOL of suspended Pair of HRC is rejected.

The M-VOLs of the HRC volume pairs and the R-VOLs of other HRC volume pairs can be intermixed in one DKC515I subsystem.

Note: Do not use M-VOLs or R-VOLs from hosts that have different CU emulation types (2105/2107 and 3990) at the same time. If you use the M-VOLs or R-VOLs from the 2105/2107 and 3990 hosts simultaneously, an MIH message might be reported to the 3990 host.

Note: When 3390-L/M volume is used as M/R-VOL, CU emulation type of package used for connection between MCU and RCU should be other than 3990.

(b) MCU and RCU

An **MCU** (main disk control unit) and an **RCU** (remote disk control unit) are disk control units in the DKC515I subsystems to which the M-VOLs and the R-VOLs are connected respectively.

An MCU controls I/O operations from host processors to the M-VOLs and copy activities between the M-VOLs and the R-VOLs. An MCU also provides functions to manage HRC status and configuration.

An RCU executes write operations directed by the MCU. The manner to execute write operations is almost same as that of I/O operations from host processors. An RCU also provides a part of functions to manage HRC status and configuration.

Note that an MCU/RCU is defined on each HRC volume pair basis. One disk control unit can operate as an MCU to control the M-VOLs and an RCU to control the R-VOLs.

Note: When serial interface connection is used for the remote copy operation, the controller emulation of the connected port of the MCU and RCU can be different. However, when MCU and RCU are connected to the same host, the controller emulation of the connected port of the MCU and RCU must be the same (3990-3/6/6E or 2105/2107).

(c) Remote Copy Connections

There are two kinds of Serial interface (ESCON/ACONARC) and Fibre channel interface of connection form.

At least two independent remote copy connections should be established between an MCU and an RCU.

(d) RCP

An **RCP** (remote control port) is a serial interface port to which an RCU is connected. Any serial interface port of the DKC515I subsystems can be configured as an RCP.

When an MCU communicates with an RCU through ESCON interface protocol, the RCP plays the role of a host processor channel. The RCP supports ESCON dynamic connection. A serial interface port of the RCU to which the MCU is connected can be connected to host processor channels by using dynamic switching capability provided by ESCON directors.

However an RCP can not communicate with host processor channel. Channel interface paths must be connected to other serial interface ports.

(e) SVP and Web Console

An **SVP** provides functions to set up , modify and display HRC configuration and status.

A **Web Console** is a personal computer compatible with the PC/AT. It should be connected to DKC515I subsystems with an Ethernet network (TCP/IP). Several DKC515I subsystems can be connected with one Ethernet network.

For Web Console, Hitachi provides only two software components, an HRC application program and dynamic link library. Both of them require Microsoft Windows operating system. A personal computer, Ethernet materials and other software products are not provided by Hitachi.

(f) Error Reporting Communications

Error reporting communication is a communication means between host processors. An MCU generates the sense information when it fails in keeping synchronization of HRC volume pair. The sense information causes the corresponding message to be displayed on the host processor console. For the reference during disaster recovery at the secondary (recovery) site, this console message should be transferred to the secondary site through the error reporting communication.

The error reporting communications may be configured by using channel-to-channel communications, Netview technology or other interconnect technologies, depending on installation. Hitachi does not provide any product for error reporting communications.

(g) PPRC Support

HRC provides a host processor interface compatible with IBM PPRC. TSO commands, DSF commands and disaster recovery PTFs provided for PPRC can be used for HRC.

(h) XRC Time-Stamping Capability (For HRC Asynchronous)

In case of N-to-1 configuration, the XRC time-stamping capability requires to be installed in the primary host system. MVS/DFP 3.2.0 or higher level is required.

In order to get benefit of time-stamping capability during copy-back process(pair establishment from the secondary to the primary subsystem), the XRC time-stamping capability recommends to be installed in the secondary system.

If the primary system(and the secondary system)consists of several CPU complexes, SYSPLEX timer must be installed for common time reference.

(i) Consistency group

HRC asynchronous ensures update-sequence-consistency across serveral volume pairs. It also provides some group-based operations. A set of volume pairs treated by such group-based functions is called a consistency group.

HRC asynchronous supports a maximum of 128 consistency groups. Every HRC asynchronous volume pair belongs to the one consistency group.

(j) Initiator Port

An **Initiator Port** (remote control port) is a Fibre Channel interface port to which an RCU is connected. Any Fibre Channel interface port of the DKC515I subsystems can be configured as an Initiator Port.

But, as for the channel port of the host computer, it can't communicate. A path from the host computer must be connected with other Fibre Channel interface ports.

(k) RCU Target Port

An **RCU Target Port** (remote control port) is a Fibre Channel interface port to which an MCU is connected. Any Fibre Channel interface port of the DKC515I subsystems can be configured as an RCU Target Port.

It can be connected with the channel of the host computer by the Fibre Channel switch.

HRC operations from an SVP or a Web Console and the corresponding TSO commands are shown in Table 3.11.1-1. Before using TSO commands or DSF commands for PPRC, the serial interface ports to which the RCU(s) will be connected must be set to the RCP mode. Table 3.11.1-2 shows the value of the SAID (system adapter ID) parameters required for CESTPATH command. For full description on TSO commands or DSF commands for PPRC, refer to the appropriate manuals published by IBM corporation.

Table 3.11.1-1 HRC operations and corresponding TSO commands for PPRC

Function	HRC operations	TSO commands
Registering an RCU and establishing remote copy connections	Add RCU	CESTPATH (note)
Adding or removing remote copy connection(s)	Edit Path	CESTPATH
Deleting an RCU registration	Delete RCU	CDELPATH
Registering consistency groups	Add Group	—
Deleting consistency group registration	Delete Group	—
Establishing an HRC volume pair	Add Pair	CESTPAIR MODE (COPY)
Suspending an HRC volume pair	Suspend Pair	CSUSPEND
Disestablishing an HRC volume pair	Delete Pair	CDELPAIR
Recovering an HRC volume pair from suspended condition	Resume Pair	CESTPAIR MODE (RESYNC)
Controlling HRC volume groups	—	CGROUP

Note: Required Parameters

(How to set up LINK PARAMETER for CESTPATH command)

LINK PARAMETER

a	a	a	a	b	b	c	c
---	---	---	---	---	---	---	---

aaaa: SAID (refer to Table 3.11.1-2)

bb : destination address

cc : CUI# of RCU

Table 3.11.1-2 SAID (system adapter ID) required for CESTPATH command (1/2)

FRONT CL1

Package Location	Port	SAID	Package Location	Port	SAID	Package Location	Port	SAID	Package Location	Port	SAID
1E (Basic) *1	CL1-A	X'0000'	1G (Add2)	CL1-J	X'0008'	1K (Add4)	CL9-N	X'008C'	1B (Add6)	CL9-E	X'0084'
	CL3-A	X'0020'		CL3-J	X'0028'		CLB-N	X'00AC'		CLB-E	X'00A4'
	CL5-A	X'0040'		CL5-J	X'0048'		CLD-N	X'00CC'		CLD-E	X'00C4'
	CL7-A	X'0060'		CL7-J	X'0068'		CLF-N	X'00EC'		CLF-E	X'00E4'
	CL1-B	X'0001'		CL1-K	X'0009'		CL9-P	X'008D'		CL9-F	X'0085'
	CL3-B	X'0021'		CL3-K	X'0029'		CLB-P	X'00AD'		CLB-F	X'00A5'
	CL5-B	X'0041'		CL5-K	X'0049'		CLD-P	X'00CD'		CLD-F	X'00C5'
	CL7-B	X'0061'		CL7-K	X'0069'		CLF-P	X'00ED'		CLF-F	X'00E5'
	CL1-C	X'0002'		CL1-L	X'000A'		CL9-Q	X'008E'		CL9-G	X'0086'
	CL3-C	X'0022'		CL3-L	X'002A'		CLB-Q	X'00AE'		CLB-G	X'00A6'
	CL5-C	X'0042'		CL5-L	X'004A'		CLD-Q	X'00CE'		CLD-G	X'00C6'
	CL7-C	X'0062'		CL7-L	X'006A'		CLF-Q	X'00EE'		CLF-G	X'00E6'
	CL1-D	X'0003'		CL1-M	X'000B'		CL9-R	X'008F'		CL9-H	X'0087'
	CL3-D	X'0023'		CL3-M	X'002B'		CLB-R	X'00AF'		CLB-H	X'00A7'
	CL5-D	X'0043'		CL5-M	X'004B'		CLD-R	X'00CF'		CLD-H	X'00C7'
	CL7-D	X'0063'		CL7-M	X'006B'		CLF-R	X'00EF'		CLF-H	X'00E7'
1F (Add1) *2	CL1-E	X'0004'	1H (Add3)	CL1-N	X'000C'	1L (Add5)	CL9-J	X'0088'	1A (Add7)	CL9-A	X'0080'
	CL3-E	X'0024'		CL3-N	X'002C'		CLB-J	X'00A8'		CLB-A	X'00A0'
	CL5-E	X'0044'		CL5-N	X'004C'		CLD-J	X'00C8'		CLD-A	X'00C0'
	CL7-E	X'0064'		CL7-N	X'006C'		CLF-J	X'00E8'		CLF-A	X'00E0'
	CL1-F	X'0005'		CL1-P	X'000D'		CL9-K	X'0089'		CL9-B	X'0081'
	CL3-F	X'0025'		CL3-P	X'002D'		CLB-K	X'00A9'		CLB-B	X'00A1'
	CL5-F	X'0045'		CL5-P	X'004D'		CLD-K	X'00C9'		CLD-B	X'00C1'
	CL7-F	X'0065'		CL7-P	X'006D'		CLF-K	X'00E9'		CLF-B	X'00E1'
	CL1-G	X'0006'		CL1-Q	X'000E'		CL9-L	X'008A'		CL9-C	X'0082'
	CL3-G	X'0026'		CL3-Q	X'002E'		CLB-L	X'00AA'		CLB-C	X'00A2'
	CL5-G	X'0046'		CL5-Q	X'004E'		CLD-L	X'00CA'		CLD-C	X'00C2'
	CL7-G	X'0066'		CL7-Q	X'006E'		CLF-L	X'00EA'		CLF-C	X'00E2'
	CL1-H	X'0007'		CL1-R	X'000F'		CL9-M	X'008B'		CL9-D	X'0083'
	CL3-H	X'0027'		CL3-R	X'002F'		CLB-M	X'00AB'		CLB-D	X'00A3'
	CL5-H	X'0047'		CL5-R	X'004F'		CLD-M	X'00CB'		CLD-D	X'00C3'
	CL7-H	X'0067'		CL7-R	X'006F'		CLF-M	X'00EB'		CLF-D	X'00E3'

*1: DKC515I is installed Location 1A.

*2: DKC515I is installed Location 1B.

Table 3.11.1-2 SAID (system adapter ID) required for CESTPATH command (2/2)

REAR CL2

Package Location	Port	SAID	Package Location	Port	SAID	Package Location	Port	SAID	Package Location	Port	SAID
2Q (Basic) *3	CL2-A	X'0010'	2T (Add2)	CL2-J	X'0018'	2W (Add4)	CLA-N	X'009C'	2N (Add6)	CLA-E	X'0094'
	CL4-A	X'0030'		CL4-J	X'0038'		CLC-N	X'00BC'		CLC-E	X'00B4'
	CL6-A	X'0050'		CL6-J	X'0058'		CLE-N	X'00DC'		CLE-E	X'00D4'
	CL8-A	X'0070'		CL8-J	X'0078'		CLG-N	X'00FC'		CLG-E	X'00F4'
	CL2-B	X'0011'		CL2-K	X'0019'		CLA-P	X'009D'		CLA-F	X'0095'
	CL4-B	X'0031'		CL4-K	X'0039'		CLC-P	X'00BD'		CLC-F	X'00B5'
	CL6-B	X'0051'		CL6-K	X'0059'		CLE-P	X'00DD'		CLE-F	X'00D5'
	CL8-B	X'0071'		CL8-K	X'0079'		CLG-P	X'00FD'		CLG-F	X'00F5'
	CL2-C	X'0012'		CL2-L	X'001A'		CLA-Q	X'009E'		CLA-G	X'0096'
	CL4-C	X'0032'		CL4-L	X'003A'		CLC-Q	X'00BE'		CLC-G	X'00B6'
	CL6-C	X'0052'		CL6-L	X'005A'		CLE-Q	X'00DE'		CLE-G	X'00D6'
	CL8-C	X'0072'		CL8-L	X'007A'		CLG-Q	X'00FE'		CLG-G	X'00F6'
	CL2-D	X'0013'		CL2-M	X'001B'		CLA-R	X'009F'		CLA-H	X'0097'
	CL4-D	X'0033'		CL4-M	X'003B'		CLC-R	X'00BF'		CLC-H	X'00B7'
	CL6-D	X'0053'		CL6-M	X'005B'		CLE-R	X'00DF'		CLE-H	X'00D7'
	CL8-D	X'0073'		CL8-M	X'007B'		CLG-R	X'00FF'		CLG-H	X'00F7'
2R (Add1) *4	CL2-E	X'0014'	2U (Add3)	CL2-N	X'001C'	2X (Add5)	CLA-J	X'0098'	2M (Add7)	CLA-A	X'0090'
	CL4-E	X'0034'		CL4-N	X'003C'		CLC-J	X'00B8'		CLC-A	X'00B0'
	CL6-E	X'0054'		CL6-N	X'005C'		CLE-J	X'00D8'		CLE-A	X'00D0'
	CL8-E	X'0074'		CL8-N	X'007C'		CLG-J	X'00F8'		CLG-A	X'00F0'
	CL2-F	X'0015'		CL2-P	X'001D'		CLA-K	X'0099'		CLA-B	X'0091'
	CL4-F	X'0035'		CL4-P	X'003D'		CLC-K	X'00B9'		CLC-B	X'00B1'
	CL6-F	X'0055'		CL6-P	X'005D'		CLE-K	X'00D9'		CLE-B	X'00D1'
	CL8-F	X'0075'		CL8-P	X'007D'		CLG-K	X'00F9'		CLG-B	X'00F1'
	CL2-G	X'0016'		CL2-Q	X'001E'		CLA-L	X'009A'		CLA-C	X'0092'
	CL4-G	X'0036'		CL4-Q	X'003E'		CLC-L	X'00BA'		CLC-C	X'00B2'
	CL6-G	X'0056'		CL6-Q	X'005E'		CLE-L	X'00DA'		CLE-C	X'00D2'
	CL8-G	X'0076'		CL8-Q	X'007E'		CLG-L	X'00FA'		CLG-C	X'00F2'
	CL2-H	X'0017'		CL2-R	X'001F'		CLA-M	X'009B'		CLA-D	X'0093'
	CL4-H	X'0037'		CL4-R	X'003F'		CLC-M	X'00BB'		CLC-D	X'00B3'
	CL6-H	X'0057'		CL6-R	X'005F'		CLE-M	X'00DB'		CLE-D	X'00D3'
	CL8-H	X'0077'		CL8-R	X'007F'		CLG-M	X'00FB'		CLG-D	X'00F3'

*3: DKC515I is installed Location 2F.

*4: DKC515I is installed Location 2E.

- (l) DKC emulation type = 2105/2107

The RESETHP option of the CESTPATH command reset host's I/Os, so you have to stop them in advance.

3.11.2 HRC Software Requirements

Minimum level for HRC is MVS/DFP 3.2.0 + PTF or VM/ESA 2.1.0 + PTF.

- Optional error recovery procedure (ERP) functions - MVS/DFP 3.2.0 or above.
- ICKDSF R16 + PTF functions - VM/ESA 2.1.0 or above.

3.11.3 HRC Hardware Requirements

(1) HRC Supported models

Refer to Specifications of “THEORY OF OPERATION SECTION” for the Support models.

- Emulation type of an MCU and RCU can be different.
- Emulation type of an M-VOL and R-VOL must be same.
- CVS/DCR is able to define on the M-VOL and R-VOL.
- The 3990-6E or 2105/2107 DKC emulation supports only 3390 DKU emulation.
- The R-VOL must have the same track sizes, and the same or larger volume capacities, as the M-VOL. (note)
- When T-VOL of HMRCF is established as M-VOL of HRC, the T-VOL must be “Split” state. If not so, an equipment check error will occur in establishing HRC pair.

Note:

When executing an HRC between volumes with different capacities, take notice of the following.

- ① State of the subsystem when executing the HRC
All M-VOL data including the VTOC is physically copied by executing the HRC. Accordingly, the R-VOL is recognized as having the same capacity as that of the M-VOL in the HRC execution.
 - ② Expansion of the VTOC
It is needed to expand the VTOC to make the R-VOL simplex and be accessed as having the normal capacity by the host. The VTOC is expanded by issuing the ICKDSF REFORMAT (REFVTOC) command.
- * The system environment is required to support the REFVTOC parameter. This parameter can be executed only by the ICKDSF which supports the PPRC function.

(2) Web Station PC Requirements

An HRC application software and dynamic link library require Microsoft Windows.

(3) Distance between MCU and RCU

(a) Serial interface connection

An MCU and an RCU must be connected with serial interface (ESCON) cables. Only multi mode ESCON cables whose length is up to 3km can be connected to the DKC515I subsystems. In order to locate disk subsystems more than 3km apart, IBM 9032/9033 ESCON directors (ESCDs) or 9036 ESCON repeaters are required.

IBM 9032/9033 ESCON director supports an extended distance facility (XDF). The XDF uses single mode ESCON cables of which length is up to 20km. IBM 9036 ESCON repeater supports single mode - to - single mode connection or single mode - to - multi mode connection. In order to locate disk control units more than 9km apart, the XDF connections provided by the ESCON directors or ESCON repeaters are required.

Maximum distance between disk subsystems is 43km.

Furthermore, the distance between an MCU and an RCU becomes unrestricted by connecting a Channel Extender (CX5000 or ULTRANET manufactured by CNT). (refer to [THEORY03-500](#) Appendix B)




Fig 3.11.3-1 Distance between Disk subsystems of Serial interface connection

(b) Fibre channel interface connection

You must connect MCU and RCU with Optical Fibre cable.

With ShortWave (Optical Multi Mode), the longest cable is 500 m. The longest cable is 10 km with LongWave (Optical Single Mode).

By connecting Switch, the longest cable for ShortWave is 1.5 km, and 30 km for LongWave. But the Switch can be connected with a maximum of two steps.

Channel Extender Connects MCU and RCU with no distance restriction.

In case of a direct connection between MCU and RCU, each Fibre channel port topology must be “Fabric:Off and FC-AL”.

In case of via FC-Switch connection between MCU and RCU, each Fibre channel port topology must be set the same as for the closest FC-Switch's topology.

(Eg.) “Fabric:On and FC-AL” or “Fabric:On and Point-to-Point” or “Fabric:Off and Point-to-Point”




Fig. 3.11.3-2 Distance between Disk subsystems of Fibre channel interface connection

(4) Recommendation of MIH time and HRC configuration

Recommendation of MIH time is 60 sec. for HRC. In addition that, MIH time needs to be set with consideration of the following factors.

- The number of pair volumes
- Cable length between MCU and RCU
- Volume status (Initial copy status)
- Maintenance operation pending

When HRC configuration is implemented, consider the following.

- The logical paths between MCU and RCU need to be established independent of the logical paths between Host and RCU.
- The maximum logical paths between MCU and RCU should not exceed the maximum allowable per subsystem.
- When configuring HRC paths, host I/O and RCP port configuration should be considered.
- The Only important volumes had better be established HRC pair.
- Under the situation where the number of paths between MCU and RCU is decreased due to path failure, the traffic of the remaining paths is increased. It may cause some I/O time-outs to RCU operations and excess time-outs might cause suspension of HRC pairs.

When HRC asynchronous configuration is implemented, take note of the following in addition to the statements above.

- The cache of the MCU should be double that of the RCU's

When HRC asynchronous N-to-1 configuration is implemented, take note of the following. Failure to do so may cause suspension of HRC pairs.

- The host write I/O rate should be balanced in all MCUs.
- The cache amount of all MCUs should be equal.
- The cache amount of RCU should be equal to the total cache of all MCUs.
- The logical number of paths between MCU and RCU should be the same for all MCUs.

On the use of HRC Async, HORC Async and HXRC, take notice of the following.

Failure to do so may cause suspension of pairs.

- HRC Async, HORC Async and HXRC should not be used together.

Note: HRC with FICON

If you use FICON paths between CHL and MCU, you should consider the following.

Table 3.11.3-1 HRC with FICON

		CHL-MCU	
		FICON	ESCON
MCU-RCU	ESCON	Not supported *	Supported
	Open Fiber	Supported	Supported

*: The link bandwidth of FICON is greater than that of ESCON. Considering this performance balance of ESCON path and FICON path, if FICON path is used between Channel and MCU, the path between MCU and RCU should be OPEN-FC link, not ESCON link.

(5) Requirements for Track format, Cache, NVS, DASD Fast Write

Track format Cache, NVS, DASD Fast Write must satisfy following conditions for both of M-VOL and R-VOL of Remote Copy.

		HRC	
		M-VOL	R-VOL
Track format	Track format of record zero	Standard format	
	Key length of record zero	Zero	
	Data length of record zero	Eight	
	CCHH of record zero	Identical to physical cylinder address and physical head address of a track	
	CCHH of each user records	Unique in a track	
Cache		Depends on HRC pair option*	Depends on HRC pair option*
NVS		Depends on HRC pair option*	Depends on HRC pair option*
DASD Fast Write		Depends on HRC pair option*	Depends on HRC pair option*

*: Required if “DFW to R-VOL” option is set.

(6) HRC available CU image

CU#0 ~ CU#31 (In hexadecimal, CU#1F)

(7) Connections between DKC515 and DKC410 or DKC460 Subsystems

Usable configurations depending on combinations of models of MCU and RCU to be connected are shown in the following table.

Descriptions following the (1) and (2) in the table express extents of the usable ports and CU#'s respectively.

Also, descriptions on the left and right of a slash (/) express the extents of the usable models of the MCU and RCU respectively.

MCU \ RCU	DKC410I (01-19-81-xx/xx)	DKC460I (21-12-00-xx/xx)	DKC515I (50-03-xx-xx/xx) (50-05-xx-xx/xx (*1))
DKC410I (01-19-81-xx/xx)	(1) 1A ~ 2R/1A ~ 2R (2) 00 ~ 0F/00 ~ 0F	(1) 1A ~ 2R/1A ~ 4R (2) 00 ~ 0F/00 ~ 1F	(1) 1A ~ 2R/1A ~ 4R (2) 00 ~ 0F/00 ~ 1F
DKC460I (21-12-00-xx/xx)	(1) 1A ~ 4R/1A ~ 2R (2) 00 ~ 1F/00 ~ 0F	(1) 1A ~ 4R/1A ~ 4R (2) 00 ~ 1F/00 ~ 1F	(1) 1A ~ 4R/1A ~ GR (2) 00 ~ 1F/00 ~ 3F
DKC515I (50-03-xx-xx/xx) (50-05-xx-xx/xx (*1))	(1) 1A ~ 4R/1A ~ 2R (2) 00 ~ 1F/00 ~ 0F	(1) 1A ~ GR/1A ~ 4R (2) 00 ~ 3F/00 ~ 1F	(1) 1A ~ GR/1A ~ GR (2) 00 ~ 3F/00 ~ 3F

*1: You need the following microcode version or later when you would like to use TrueCopy-MF Async N:1 configuration.

Especially when you use DKC510I as MCU DKC, RCU DKC must be DKC510I from a performance balance view point.

Note: In the case that you would like to perform microcode version down to 50-04-xx or lower version when you use TrueCopy-MF Async N:1 configuration, you must remove TrueCopy-MF Async N:1 configuration before microcode version down.

AppendixA: HRC Installation check list

Table 3.11.3-3 HRC Installation Check List

No.	Item	Check
1	MCU/RCU emulation type must be correct.	
2	M_VOL/R_VOL emulation type must be correct.	
3	OS version must be as followed. • Optional error recovery procedure (ERP) functions - MVS/DFP 3.2.0 or above. • ICKDSF R16 + PTF functions - VM/ESA 2.1.0 or above.	
4	RCP port must be set.	
5	ESCON cable between MCU and RCU must be connected.	
6	ESCON cable test between MCU and RCU must be executed.	

AppendixB: Guide for HRC via EXTENDER(CHANNELink(CX5000), ULTRANET)

1. Preliminary remarks

This is the guide for HRC operation via Computer Network Technology's (CNT's) CHANNELink and UltraNet Storage Director products. These two CNT products provide channel extension (store and forward functionality) between a Hitachi MCU and RCU used in an HRC configuration. The channel extension provided by the CNT products allows the ESCON connection between the MCU(s) and the RCU(s) to be greater than the ESCON standard 43 km apart but still provide near native data transfer between the MCU and RCU.

2. Summary of Specification

Table 3.11.3-4 Summary of Specification

		CHANNELink(CX5000)		ULTRANET	
Copy Mode		Synchronous	Asynchronous	Synchronous	Asynchronous
u-code Version	Extender	CHANNELink Release 4.3 with SSDX at version 4.B or later		UltraNet Release 2.2 with SSDX at version 4.B or later	
Link Types(between local and remote Extenders)		T3		ATM T3	
Maximum Number of physical ESCON paths(between DKC515I and EXTENDER)		3 or 4 ESCON Interfaces *1		8 ESCON Interfaces *2	
Combination of MCU and RCU		The following combination are available. (for example, refer to Fig 3.12.3-3, 3.12.3-4, 3.12.3-5) MCU:CU#n <=> RCU:CU#m (0<=n<=31, 0<=m<=31) [restriction] A CU image of the RCU can not be shared by more than one CU images of the MCU(s) that are connected to the same ESCON port(s) of the local extender.(Fig 3.11.3-6) If the CU image of the RCU requires to be shared, the CU images of the MCU(s) must be connected to the different ESCON port(s) as shown in Fig 3.11.3-7.			

*1: if the CHANNELink has 11 slots, the maximum number of ESCON interfaces is 3. If the CHANNELink has 13 slots, then the maximum is 4 ESCON interfaces.

*2: The stated maximum number of ESCON interfaces in an UltraNet Storage Director is 8.

Notes:

- One ESCON port of Extender supports up to 254 volumes. For HRC installations with requirements for greater than 254 pairs, additional ESCON interfaces are required in both the local and remote channel extenders.
- HRC via Extender does not support P/DAS function.

Figures 3.11.3-3 to 3.11.3-7 on the following pages describe various valid and invalid HRC with channel extension configurations.




Fig 3.11.3-3 MCU:CU#0 <--> RCU:CU#0 (valid configuration)




Fig 3.11.3-4 MCU:CU #0 <--> RCU:CU #1, MCU:CU #1 <--> RCU:CU #0 (valid configuration)




Fig 3.11.3-5 MCU:CU #0 <--> RCU:CU #0, MCU:CU #0 <--> RCU:CU #1 (valid configuration)




Fig 3.11.3-6 MCU:CU #0 <--> RCU:CU #1, MCU:CU #1 <--> RCU:CU #1
(invalid configuration)

(2 CU Images of MCU and 1 CU Image of RCU using same ESCON port)

The above diagram is an invalid configuration as 2 CU Images of an MCU are configured to utilize 1 RCU CU Image via the same ESCON port. Unique remote ESCON ports are required to enable concurrent access to a RCU Image.




Fig 3.11.3-7 MCU:CU #0 <--> RCU:CU #1 , MCU:CU #1 <--> RCU:CU #1
(valid configuration)

(2 CU Images of MCU and 1 CU Image of RCU using different ESCON ports)

The above diagram is a valid configuration as 2 CU Images of an MCU are configured to utilize 1 RCU CU Image via different remote Extender ESCON ports.

3. Operation for HRC via Extender

3.1 Planning the configuration of HRC via Extender

Please determine the number of Extenders and the number of ESCON paths from the configuration of HRC.

There are some restrictions in planning the configuration of HRC via Extender.

[restriction]

- One ESCON port of an Extender supports up to 254 volumes. For HRC connections between more than 254 pairs, add the additional physical ports as required.

- Some combination of MCU and RCU are not available. Please refer to page

[THEORY03-500 ~ THEORY03-530.](#)

3.2 Preliminary arrangement on Extender.

(1) Change the configuration information in Extender.

There is a necessity that configuration information is set according to the configuration of HRC. Before setting HRC via Extender, please request the modification of configuration information of Extender from CNT. The information that should be given to CNT prior to installation is listed below.

- CU# of RCU

- R-VOL#

- Indication whether or not an ESCD (ESCON director) will be included in the configuration. *3


- ESCON link address (if using ESCD)

- A request to set serial number(s) of Local Extender *4

Notes:

*3: If an ESCD (ESCON director) is put between an MCU & local Extender, or between an RCU & remote Extender, there occurs a necessity to have Extender configuration information set by CNT specially. (According to CNT's specification, it is allowed to locate one ESCD in master and remote site, respectively, and define all link ports as 'Dynamic Connection'.)

*4: The Extender can be set any serial numbers to every ESCON ports independently. At ADDRUCU (/CESTPATH) operation, the serial number of Local Extender should be set as the serial number of RCU. So, in advance, the same serial number should be set to ESCON physical paths that make the logical paths to same RCU. For an example, refer to Fig. 3.11.3-8. When 2 RCUs are connected via Extender, different serial numbers should be set. But such configurations have not certified yet. Please use a MCU and a RCU.



These parameters are required in ADD RCU operation of DKC515I from SVP.

RCU		Path		
Serial number	SSID	PORT#	Link Address	Logical Address
10924	0040	1A	00	02
10924	0040	2A	00	02

Without connecting CNT, this parameter is 30128.

Fig 3.11.3-8 reference for notice *4

2.2 Operation of DKC515I

(1) Setting of Special mode (Mode 21=ON)

There is a necessary that Mode21 is set, which is a special mode for HRC via Extender. Specifically, when the MCU restarts the WRFBA (HRC special command) CCW chain, it will be an incomplete domain pattern from the beginning.

(2) ADD RCU

At ADDRUCU (/CESTPATH) operation, input the serial number (S#) of Local Extender as the serial number of RCU. The procedure to seek the Local Extender's serial number is as follows.

<Procedure to seek Extender's S#>

To begin with, please connect MCU to RCU via CX5000 directly with no ESCD (ESCON director).

- Please open the icon "MAINTENANCE" of SVP (MCU).
- Please click the button "LCP/MCP Path" in General Status Display.
- Please select "Physical Path Status".
- Please notice the information of RCP port to connect to local CNT. The last five digits of "SEQNUMBER" correspond to the target manufacturing number to input.

3. Configuration Setup Example

(1) MCU/RCU Example Pair Configuration

Table 3.11.3-5 Example M-VOL and R-VOL Pair Configuration Matrix

Table 3.11.3-5 Pair configuration matrix

MCU/ CU Image	M-VOL#	SSID		RCU/ CU Image	R-VOL#	SSID
CU#0	0x00-0x3F	0x0004	----->	CU#0	0x00-0x3F	0x0040
CU#1	0x00-0x3F	0x0008	----->	CU#2	0x00-0x3F	0x0048
	0x40-0x7F	0x0009	----->	CU#3	0x00-0x3F	0x004C
CU#3	0x00-0x3D	0x0010	----->		0x40-0x7D	0x004D

(2) Configuration of HRC via Extender

The following figure (Fig. 3.11.3-9) includes a configuration that matches the configuration that is described in Table 3.11.3-5.




Fig 3.11.3-9 configuration of HRC via Extender

(3) ADD RCU parameter

These parameters are required in ADD RCU operation of DKC515I from SVP in case of example in this section.

Table 3.11.3-6 ADD RCU parameter

RCU		Path		
Serial number	SSID	PORT#	Link Address	Logical Address
10924	0040	1A	00	00
10924	0048	1A	00	02
10924	004C	1A	00	03
10924	004D	2A	00	03

Without connecting EXTENDER, this parameter is 30128.

4. Notes

- (1) When one of HRC paths between the remote EXTENDER and the RCU is blocked, the blocked path will not be recovered voluntarily. And then, "PATH STATUS" of this path may be abnormal(-) status, otherwise the HRC will suspend due to any failure of paths(SSB F/M=8F, EC=C8Ax). If C8Ax status is indicated, then the failure path is determined by the information as follows.

- CL1/2 : (it is known by MP# detecting SSB)
- LPN in CL1/2 : (SSB log byte42 means LPN in CL1/2)

In these cases, please confirm the connection of the physical path and follow the following procedure to recover the path.

[Recover procedure]

- (a) Please delete the failed HRC path by way of "EDIT PATH" menu.
- (b) Please edit the path again to select "EDIT PATH" menu.

- (2) When using the procedure of switching the M-VOL and the R-VOL for disaster recovery of HRC via Extender, the configuration of Local/Remote Extender should be set again because the Local Extender and Remote Extender are reversed. And, when a CU image of the original MCU and plural CU image of the original RCU are connecting through the same physical path, the above mentioned disaster recovery procedure is inapplicable. In this case, please recover the combination of 1 CU Image of the new MCU and 1 CU Image of the new RCU according to the procedure switching the M-VOL and the R-VOL. After that, please try to recover the remaining CU Images one by one.

3.11.4 HRC Theory of Operations

(1) HRC Copy Activities

HRC executes two kinds of copy activities, initial copy and update copy.




Fig. 3.11.4-1 HRC Copy Activities

(a) Initial Copy

Responding to an Establish HRC Volume Pair operation from an SVP/remote console or an ESTPAIR PPRC command, HRC begins initial copy. Data field of record zero and following records on all tracks, except for alternate and CE tracks, are copied from M-VOL to R-VOL. The initial copy operation is performed in ascending order of cylinder numbers.

“No copy” can be specified as a parameter to the initial copy. When “no copy” is specified, HRC will complete an Establish HRC Volume Pair operation without copying any data. An operator or a system administrator should be responsible for ensuring that data on the M-VOL and the R-VOL is already identical.

“Only out-of-sync cylinders” can also be specified as a parameter to the initial copy. This parameter is used to recover (re-establish) HRC volume pair from suspended condition. After suspending HRC volume pair, the MCU maintains a cylinder basis bit map which indicates the cylinders updated by I/O operations from the host processors. When this parameter is specified, HRC will copy only cylinders indicated by the bit map.

(b) Controlling Initial Copy

Number of tracks copied by one initial copy activity can be specified by an SVP/remote console or an ESTPAIR PPRC command.

Number of volume pairs for which the initial copy are concurrently executed and priority of each volume pair can be specified from an SVP/remote console.

(c) Update Copy

Responding to the write I/O operations from the host processors, HRC copies the records updated by the write I/O operation to the R-VOL.

The update copy is a synchronous remote copy. An MCU starts the update copy after responding only channel-end status to the host processor channel, and sends device-end status after completing the update copy. The MCU will start the update copy when it receives:

- The last write command in the current domain specified by preceding locate record command;
- A write command for which track switch to the next track is required;
- Each write command without being preceded by locate record command.

If many consecutive records are updated by single CCW chain which does not use locate record command, the third condition above may cause the significant impact on performance.

(d) Update Copy for Cache Fast Write Data

Cache fast write (CFW) data does not always have to be copied because CFW is used for temporary files, such as sort work data sets. These temporary files are not always necessary for disaster recovery.

In order to reduce update copy activities, HRC supports a parameter which specifies whether CFW data should be copied or not.

(e) Special Write Command for Initial Copy and Update Copy

In order to reduce overhead by the copy activities, HRC uses a special write command which is allowed only for copy activities between the DKC515I subsystems. The single write command transfers control parameters and an FBA formatted data which includes consecutive updated records in a track. It reduces interlocks on ESCON interface protocol and overhead required for converting FBA-to-CKD format and CKD-to-FBA format.

(2) HRC Read I/O Operations

Responding to read I/O operations, an MCU transfers the requested records from an M-VOL to a host processor. Even if reading records from the M-VOL is failed, the R-VOL is not automatically read for recovery. The redundancy of the M-VOL itself provided by RAID5 or RAID1 technique would recover the failure.

(3) HRC Volume Pair Status

All volumes in a DKC515I subsystem are in one of the states shown in Table 3.11.4-1.

Status of the M-VOLs or the R-VOLs are kept by the MCU and the RCU respectively. The MCU is responsible to keep status of the R-VOLs identical to status of the M-VOLs. However, in the case of communication failure between the MCU and the RCU, they could be different.

From an Web Console or by using an appropriate command for IBM PPRC, status of M-VOLs or status of R-VOLs can be obtained from the MCU or the RCU respectively.

Table 3.11.4-1 HRC Volume Status

Status	Description
Simplex	This volume does not belong to HRC volume pair. When the initial copy is started by an Add Pair operation, the volume is changed to “pending duplex” state.
Pending Duplex	The initial copy is in progress. Data on HRC volume pair is not fully identical. When completing the initial copy, the volume will be changed to “duplex” state.
Duplex	Volumes in HRC volume pair are synchronized. All updates from the host processors to the M-VOL are duplicated to the R-VOL.
Suspended	Volumes in HRC volume pair are not synchronized. • When the MCU can not keep synchronization between HRC volume pair due to, for example, failure on the update copy, the MCU will put the M-VOL and the R-VOL in this state. • When the MCU or the RCU accepts a Suspend operation from an SVP/remote console, the M-VOL and the R-VOL will be put in this state. • When the RCU accepts the Delete Pair operation from the SVP/remote console, the MCU will detect the operation and put the M-VOL in this state.

Table 3.11.4-2 HRC Volume Status - Sub-status of Suspended Volume

Cause of Suspension	Description
M-VOL by Operator	The Suspend operation with “M-VOL failure” option was issued to the M-VOL. This cause of suspension is defined only for the M-VOLs.
R-VOL by Operator	The Suspend operation with “R-VOL” option was issued to the M-VOL or the R-VOL. This cause of suspension is defined for both the M-VOLs and the R-VOLs.
by MCU	The RCU received a request to suspend the R-VOL from an MCU. This cause of suspension is defined for only the R-VOLs.
by RCU	The MCU detected an error condition of the RCU which caused HRC volume pair to be suspended. This cause of suspension is defined only for the M-VOLs.
Delete Pair to RCU	The MCU detected that the R-VOL had been changed to “simplex” state by the Delete Pair operation. This cause of suspension is defined only for the M-VOLs.
R-VOL Failure	The MCU detected an error condition on the communication between the RCU or I/O error on the update copy. This cause of suspension is defined only for the M-VOLs. The cause of suspension of the R-VOLs are usually set to “by MCU” in this situation.
MCU IMPL	The MCU could not find valid control information in its non-volatile memory during its IMPL procedure. This situation may occur after the power supply failure.
Initial Copy Failed	The volume pair was suspended before completing the initial copy. Even if no write I/O has been issued after being suspended, the data in the R-VOL is not completely identical to the M-VOL.

3.11.5 HRC Control Operations

This section describes HRC control operations from a Web Console.

(1) Add RCU Operation

(a) Serial interface connection

An Add RCU operation makes an MCU register the specified disk control unit as an RCU and establish the logical paths to the RCU. This operation also provides a function to modify the Remote Copy options which will be applied to all Remote Copy volume pairs in this subsystem.

To register the RCU, the following parameters are required:

RCU S#	Serial number of the RCU.
SSID	SSID (subsystem identifier) of the RCU.
Num. of Path	Number of logical paths which should be established to the RCU on the remote copy connections.

Path parameters to specify one logical path are shown below. Up to four sets of the path parameters, as many as the “Num. of Paths” parameter specifies, must be specified. The description of path parameters is similar to the channel path definitions in IOCDS (I/O configuration dataset). In the IOCDS, a logical path is described with a sub-channel number, a link destination address and a logical address of the control unit image. A “Port” parameter, instead of the sub-channel number, is used to specify the serial interface port in the DKC515I subsystem.

Port	The serial interface port of the DKC515I subsystem where the logical path begins from. Before this operation, the serial interface port must be set to the RCP mode.
Link Adr	The link destination address. Similar to the logical path definitions in the IOCDS, this is the destination port address on the ESCD which is set to provide the dynamic connection capability. If the remote copy connection does not include the dynamic connection, x'00' must be specified.
Logical Adr	The logical control unit address of the control unit image. If the RCU is configured as an SPSD (single path storage director), either x'00' or x'01' must be specified. Otherwise, x'00' must be specified.

Fig. 3.11.5-1 shows an example for the RCU and the path parameters.




Fig. 3.11.5-1 Add RCU Operation Parameters Example

(b) Fibre channel interface connection

The following parameters are necessary to register RCU as a Fibre Channel connection.

Port Type	Serial: Serial channel interface is used for the connection of MCU and RCU. Fibre: Fiber channel interface is used for the connection of MCU and RCU.
Controller ID	But, Set it up with '02' when RCU is RAID400, and set it up with '03' when RCU is RAID450. Default is '04' fixed.
RIO MIH Time	A data transfer complete waiting time to RCU from MCU. Usual: 15[Sec]. Avail. range: 10[Sec] ~ 100[Sec]
MCU Port	An Initiator port of the DKC515I subsystem which set up a logic pass. You must set up a Fibre Channel interface port in Initiator port before this operation.
RCU Port	The Fibre Channel interface port of the place of the connection. You must specify a RCU target port.




Fig. 3.11.5-2 Add RCU Operation

The following parameters modify the Remote Copy options which will be applied to all Remote Copy volume pairs in this subsystem.

- Minimum Paths** When the MCU blocks the logical path due to communication failure, if the number of remaining paths becomes less than the number specified by this parameter, the MCU will suspend all of the Remote Copy volume pairs. The default value is set to “1”. If the installation requirements prefers the subsystem I/O performance to the continuation of Remote Copy, value between “2” and the number of the established logical paths can be specified.
When you use asynchronous pairs of fibre channel interface, please set the same value for the minimum paths of all RCU.
- Maximum Initial Copy Activities** It specifies how many HRC initial copies can be simultaneously executed by the MCU. If more Remote Copy volume pairs are specified by an Add Pair operation, the MCU will execute the initial copy for as many volumes as specified by this parameter. The initial copy for other volumes is delayed until one of the initial copies is completed. This parameter can control the performance impact caused by the initial copy activity.
Note: Default value of this parameter is “4”.
- Incident of RCU** This parameter specifies whether the link incident record generated by the RCU is to be sent or not to the host processor connected to the MCU.
- PPRC supported by HOST** If “Yes” is specified, the MCU will generate the sense information which is compatible with IBM PPRC when the HRC volume pair is suspended. If “No” is specified, the MCU will generate only service information messages. Even if the SSB (F/M=FB) is specified by the Suspend Pair Operation, the x'FB' sense information will not be reported to the HOST.
- Service SIM of Remote Copy** If “Report” is specified, the Remote Copy Service SIM will be reported to the HOST. If “Yes” is specified in PPRC supported by HOST option, DEV_SIM of HRC will not be reported. If “Not Report” is specified, the Remote Copy Service SIM reporting will be suppressed. Refer to “SIM Reference Codes Detected by the Processor for Remote Copy” in SIM-RC SECTION.

Note that these parameters will be applied to ALL RCUs registered to the MCU. If different parameters are specified, the last parameter will be applied.

(2) Edit Path Operation

An Edit Path operation makes the MCU add/delete the logical path to/from the registered RCU.

To add a logical path, the same path parameters as an Add RCU operation are required. The added logical path will be automatically used to execute the copy activities.

When deleting a logical path, pay attention to the number of remaining logical paths. If it becomes less than the number specified by “Minimum Paths”, Remote Copy volume pair could be suspended.

(3) RCU Option Operation

An RCU Option operation modifies the Remote Copy options described in “3.11.5(1) Add RCU operation”.

(4) Delete RCU Operation

A Delete RCU operation makes the MCU delete the specified RCU from RCU registration. All logical paths to the specified RCU will be removed.

If some volumes connected to the specified RCU are active R-VOLs, this operation will be rejected. All R-VOLs must be deleted by a Delete Pair operation before a Delete RCU operation.

(5) RCU Status Operation

An RCU Status operation makes the MCU display the status of RCU registration. It also provides the current status, time of registration and time of changing status for each logical path.

The current status of each logical path is defined as follows:

Normal	This logical path has been successfully established and can be used for the Remote Copy activities.
Initialization Failed	The link initialization procedure between the RCU is failed. It occurred due to Missing physical path connection between MCU and RCU, or connecting MCU with HOST as RCU.
Resource Shortage (RCU)	Establish Logical Path link control function has been rejected by the RCU. All logical path resources in the RCU might be used for other connections.
Serial Number Mismatch	The serial number of the control unit which is connected to this logical path does not match to the serial number specified by "RCU S#" parameter.
Invalid Port	The serial interface port specified by "Port" parameter is not in the RCP mode.

(6) Add Pair Operation

An Add Pair operation makes the MCU establish a new Remote Copy volume pair. It also provides function to modify the Remote Copy options which will be applied to the selected Remote Copy volume pair. Up to 8192 Remote Copy volume pairs can be established in one DKC460I subsystem.

To establish Remote Copy volume pair, following parameters are required:

RCU	The disk control unit which controls the R-VOL of this Remote Copy volume pair. It must be selected from RCUs which have already been registered by Add RCU operations.
R-VOL	Device number of the R-VOL.
Priority	Priority (scheduling order) of the initial copy for this volume pair. When the initial copy for one volume pair has been terminated, the MCU selects and start the initial copy for another volume pair which has the lowest value of this parameter. For the Add Pair operations, the value “1” through “256” can be specified. For establishing HRC volume pair by TSO command or DSF command for PPRC, “0” is implicitly applied to. “0” is the highest priority, “256” is the lowest, and default value for the Add Pair operation is “32”. For the volume pairs to which the priority has been specified, the MCU prioritizes the volume pairs in the arrival order of the Add Pair operations or TSO/DSF commands. If the MCU are performing the initial copy for the number of volume pairs, as much as the value of “maximum initial copy activities”, and accepts further Add Pair operation, the MCU does not start other initial copy until one of the copy being performed will be completed. Note: When a time out occurs in this operation, a schedule may not be done as the priority parameter. The cause of the time-out is thought the problem of the configuration of DKC or Remote-copy connection path. Confirm configuration. After that, cancel a pair, and re-establish a pair.
Operation Mode	It specifies what kind of remote copy capability should be applied to this volume pair. “Remote Dual Copy” means HRC respectively.
Initial Copy	It specifies what kind of initial copy activity should be executed for this HRC volume pair. The kind of the initial copy can be selected out of: - “Entire Volume” specifies that all cylinders excluding the alternate cylinder and the CE cylinders should be copied. - “Only Out-of-Sync Cylinders” specifies that only cylinders which have been updated during this HRC volume pair is in “suspended” state - “None” specifies that the initial copy does not need to be executed. The synchronization between volume pair must have been ensured by the operator.

Remote Copy option parameters which will be applied to this Remote Copy volume pair are as follows:

- Initial Copy Pace** It specifies how many tracks should be copied at once by the initial copy. “15 Tracks” or “3 Tracks” can be specified. When “15 Tracks” is selected, elapsed time to complete the initial copy becomes shorter, however, the subsystem I/O performance during the initial copy could become worse. This parameter is valid only for HRC volume pair. Note: The default value of this parameter is “15”.
- DFW to R-VOL** It specifies whether the DFW capability of the R-VOL is required or not. If “DFW required” is specified, the HRC volume pair will be suspended when the RCU can not execute the DFW due to, for example, cache failure. If the installation requirements prefers the continuation of HRC to the subsystem I/O performance, “DFW not required” is recommended. This parameter is valid only for HRC volume pair.
- CFW Data** It specifies whether the records updated by CFW should be copied to the R-VOL or not. “Only M-VOL”, which means that CFW updates are not copied, is recommended because CFW data is not always necessary for disaster recovery.
- M-VOL Fence Level** It specifies by what conditions the M-VOL will be fenced (the MCU will reject the write I/O operations to the M-VOL).
- “R-VOL Data”: The M-VOL will be fenced when the MCU can not successfully execute the update copy.
 - “R-VOL Status”: The M-VOL will be fenced when the MCU can not put the R-VOL into “suspended” state. If status of the R-VOL is successfully changed to “suspended”, the subsequent write I/O operations to the M-VOL will be permitted.
 - “Never”: The M-VOL will never be fenced. The subsequent write I/O operations after the HRC volume pair has been suspended will be permitted. This parameter is valid only for HRC volume pairs.

(7) Delete Pair Operation

A Delete Pair operation makes the specified Remote Copy volume pair being terminated. It can be operated on either the MCU or the RCU.

- When operated on the MCU, both the M-VOL and the R-VOL will be put into the “simplex” state.
- When operated on the RCU, only the R-VOL will be put into the “simplex” state. The M-VOL will be suspended when the MCU detects this operation. To complete deleting this volume pair, the MCU requires another Delete Pair operation.

When the MCU accepts this operation and it can not communicate with the RCU, this operation will be rejected. “Delete Pair by Force” option can make the MCU complete this operation, even if it can not communicate with the RCU.

For the purpose of the recovery operation simply, “Delete All Pairs” option is provided in the delete pair operation. This option is need to use “Delete Pair by Force” option together, and specifies that the all volume pairs in the same RCU (CU Image) should be deleted. In the case of the delete operation at the RCU, specifies that the all volume pairs in the same serial number of the MCU and the same CU image of the MCU should be deleted.

(8) Suspend Pair Operation

A Suspend Pair operation makes the MCU or the RCU suspend the specified Remote Copy volume pair.

The option parameters for this operation are as follows:

SSB (F/M=FB)	The MCU and the RCU will generate sense information to notify the suspension of this volume pair to the attached host processors. This option is valid only for HRC volume pairs.
M-VOL Failure	The subsequent write I/O operations to the M-VOL will be rejected regardless of the fence level parameter. This option can be selected only when operating on the MCU. This option is valid for only HRC volume pairs.
R-VOL	For HRC volume pairs. This option can be accepted by the MCU and the RCU.

(9) Pair Option Operation

A Pair Option operation modifies the Remote Copy option parameters which has been applied to the selected Remote Copy volume pair. Refer to “3.11.5(6) Add Pair Operation” for the option parameters.

(10) Pair Status Operation

A Pair Status operation makes the MCU or the RCU display the result of the Add Pair operation or the Pair Status operation to the specified Remote Copy volume pair, along with the following information:

Initial Copy Complete	When this Remote Copy volume pair is in “pending duplex” state, it indicates how many cylinders have been successfully copied by the initial copy. When this Remote Copy volume pair is in “suspended” state, it indicates how many cylinders are currently identical between this Remote Copy volume pair. This information is provided only by the MCU.
Pair Status	It indicates the status of the M-VOL or the R-VOL. Definition of the volume states is described in “3.11.4(3) HRC Volume Pair Status”.
Last Update	Indicates the time stamp when the volume pair status has been updated. Note that the time stamp value is obtained from an internal clock in the DKC515I subsystem.
Pair Established	It indicates the time stamp when the volume pair has been established by an Add Pair operation. Note that the time stamp value is obtained from an internal clock in the DKC515I subsystem.

(11) Resume Pair Operation

A Resume Pair operation restarts the suspended Remote Copy volume pair. It also provides function to modify the Remote Copy options which will be applied to the selected Remote Copy volume pair.

“Out-of-Sync Cylinders” are recorded in the form of cylinder-bit-map allocated in SM (shared memory) of the DKC515I. If the MCU is powered off and the cylinder-bit-map is not retained due to the battery being discharged, the MCU resumes the initial copy as follows:

- (a) For the HRC volume pair in “pending duplex” state, the initial copy is automatically resumed. Then all cylinders of this volume will be copied.
- (b) For the HRC volume pair in “suspended” state, then all cylinders of this volume will be copied responding to the Resume Pair operation.

(12) Port Operation

(a) Serial interface connection

All serial interface ports in the DKC515I subsystem are initially set to the LCP mode, to which the host processor channels can be connected. At least two serial interface ports, one port from each storage cluster, must be set to the RCP mode for remote copy connections.

A Port operation makes the DKC515I subsystem change the operating mode of the specified serial interface port(s).

Before changing the operating mode from the LCP mode to the RCP mode, all channel paths to the specified port must be removed using host processor console or ESCD commands.

Before changing the operating mode from the RCP mode to the LCP mode, all RCUs which are connected through the specified port must be deleted by a Delete RCU operation.

Note: The Define Configuration & Installation operation also provide the function to set the operating mode of each serial interface mode.

(b) Fibre channel interface connection

You must set up the connection port of MCU and RCU prior to the path formation in Initiator port or the RCU target port from the usual target port.

Port topology of Initiator port and the RCU target port must be set up as follows.

- Direct connection : Fabric = OFF, FC-AL
- A connection via Switch : Fabric = ON, FC-AL or Point to point
- A connection via CN2000 : Fabric = OFF, Point to point

In case of a direct connection between MCU and RCU, each Fibre channel port topology must be the same as "Fabric:Off and FC-AL".

In case of via FC-Switch connection between MCU and RCU, each Fibre channel port topology must be set suitable for the closest FC-Switch's topology.

(Eg.) "Fabric:On and FC-AL" or "Fabric:On and Point-to-Point" or "Fabric:Off and Point-to-Point"

(13) Remote Copy function Switch Operation

The Specification of a present function switches is displayed.

Moreover, the function switches which want to be specified/released can be set.

Number of function switches is 64.

This function is only SVP operation.

The function allocated in each switch is as follows.

00 ~ 06: Reserved.

07: Do not report SSB of F/M = FB to HOST.

08 ~ 10: Reserved.

11: When “15 Tracks” is selected in the Fibre Remote Copy option parameter, a maximum of 32 tracks are copied by the initial copy.

12: When Add Pair or Resume Pair is operated with the HOST command, only M-VOL is set up in the CFW Data option.

Note: When ‘Only M-Vol’ to CFW Option at the time of the TC pairs created, cannot be used the data set that updated in CFW in M-Vol for R-Vol.

When use this data set in R-Vol, please format data set after deleted TC pairs.

13: Unused.

14: Reserved.

15: The path failure threshold values are changed, when the both function switch 17 and 15 are ON.

16: Reserved.

17: When the count of the LIPs, RSCNs, or time-outs that occur in the fibre path exceeds the respective threshold value, the logical path of the port concerned is detached owing to a failure.

18: In the case where the switch 17 is turned on, the count of link failures that occur in the fibre path is also a factor that results in a detachment caused by an excess of a threshold value.

19 ~ 20: Unused.

21: When Add Pair or Resume Pair is operated with the HOST command to DKC emulation type ‘2105’ or ‘2107’, PACE = 3 Tracks is set up in the Initial Copy Pace option.

22: Unused.

23: I/O multiplicity per CU changes from 4 (default) to 8.

Note: This change is only available when System Option Mode 499 is set to ON.

24 ~ 32: Unused.

33: In the cases that a PDCM function of McData ES3232 is being used on MCU-RCU path, when response of LOGIN response is late, path status of MCU is changed to non-normal.

34 ~ 63: Unused.

(14) Connection composition

Connection composition examples are shown below:


In case of a direct connection between MCU and RCU, each Fibre channel port topology must be the same as “Fabric:Off and FC-AL”.

In case of via FC-Switch connection between MCU and RCU, each Fibre channel port topology must be set suitable for the closest FC-Switch’s topology.

(Eg.) “Fabric:On and FC-AL” or “Fabric:On and Point-to-Point” or “Fabric:Off and Point-to-Point”


(a) Direct connection

*1: Fabric OFF



(b) Switch connection

*2: Fabric ON



(c) Extender connection

*3: Fabric ON




Fig. 3.11.5-3 Connection composition examples

3.11.6 Managing HRC Environment

(1) Setting Up HRC Volume Pairs

(a) Sequence of Operations

Sequence of operations to establish the HRC volume pairs are shown below.

Table 3.11.6-1 Operations to Set Up HRC Volume Pairs

Step		Operation	
		SVP *	Others
1	Set appropriate serial interface ports to the RCP mode.	Port	
2	Establish logical paths between the DKC515I HRC subsystems	Add RCU	Before this step, remote copy connections must be established between DKC515I subsystems.
3	Ensure that the R-VOLs are offline from host processors		If necessary, perform the following system command. <In case of MVS system> • VARY OFFLINE <In case of VM system> • VARY OFFLINE from guest OS • VARY PATH OFFLINE from VM
4	Establish HRC volume pairs.	Add Pair	

* : Operations from the SVP/remote console attached to the MCU.

Several volume pairs can be specified within one Add Pair Operation. After completing an Add Pair operation, another Add Pair operation can be executed to establish another HRC volume pairs.

Be sure to vary the R-VOLs offline from the attached host processors before executing the Add Pair operation. The RCU will reject the write I/O operations to the R-VOLs once the Add RCU operation has been accepted.

(b) Considering HRC Parameters

Setting of the “fence level” parameter to the Add Pair operation and the “PPRC supported by host” and “Service SIM of Remote Copy” option to the Add RCU operation depends on your disaster recovery planning. Refer to “3.11.7(1) Preparing for Disaster Recovery” for these parameters.

Setting of the “CFW data” and “DFW to R-VOL” parameters to the Add Pair operation and the “minimum paths” parameter to the Add RCU operation depends on your performance requirement to the DKC515I subsystem at the primary site. Refer to “3.11.5(6) Add Pair operation” and “3.11.5(1) Add RCU operation” for these parameters.

Setting of the “maximum initial copy activities” parameter to the Add RCU operation and the “priority” and the “initial copy pace” parameters can control performance effect from the initial copy activities. Refer to “3.11.6(1)(c) Controlling Initial Copy Activities” for more detailed description.

Refer to “3.11.5(1) Add RCU operation” and “3.11.5(6) Add Pair operation” for other parameters.

(c) Controlling Initial Copy Activities

To control performance effect from the initial copy activities, the “maximum initial copy activities” parameter and the “priority” and the “copy pace” parameters can be specified:

- The “maximum initial copy activities” parameter controls the number of volumes for which the initial copy are concurrently executed;
- The “priority” parameter specifies the executing order of the migration copy on volume pair basis;
- The “copy pace” parameter specifies how many tracks should be copied by each initial copy activity.

Refer to the following example for the “maximum initial copy activities” and the “priority” parameters.

Example

Conditions:

- The Add Pair operation specifies that devices 00~05 should be M-VOLs.
- “Maximum initial copy activities” is set to “4” (this is the default value).
- “Priority” parameters for devices 00~05 are set to “3”, “5”, “5”, “1”, “4”, and “2” respectively.

Under the above conditions, the MCU will perform the initial copy:

- for devices 00, 03, 04 and 05 immediately.
- for device 01 when one of the initial copy has been terminated.
- for device 02 when the initial copy for the second device has been terminated.

(2) Suspending and Resuming the HRC Volume Pairs

This section describes the operations to suspend or resume the HRC volume pair, which are necessary for the following sections in this chapter.

The Suspend Pair operation with the “R-VOL” option parameters can suspend the specified HRC volume pairs while the M-VOLs are still accessed from the attached host processors. The “SSB” option should not be selected to prevent the sense information from being generated.

To resume the suspended HRC volume pairs, the Resume Pair operation must be executed.

Refer to “3.11.5(8) Suspend Pair Operation” and “3.11.5(6) Add Pair Operation” for more detailed description.

(3) Managing Power On/Off of HRC Components

(a) Cutting Power to HRC component

Cutting power to the RCU or the ESCDs on the remote copy connections, or other equivalent events which make the MCU unable to communicate with the RCU should be controlled in order not to affect the remote copy activities. If the MCU detects these events when it intends to communicate with the RCU, it would suspend all HRC volume pairs.

To avoid this problem, the applications on the primary host processors must be terminated or all HRC volume pairs must be suspended or terminated, before performing these events.

Refer to “3.11.6(2) Suspending and Resuming the HRC Volume Pairs” for the operations to suspend and resume the HRC volume pairs.

(b) Power Control Interface at the Secondary Site

In the secondary site, It is not recommended to use the power control interface which remotely cuts the power to the RCU or the ESCD on the remote copy connections in order to avoid the situation described in “3.11.6(3)(a) Cutting Power to HRC components”.

(c) Power-on-sequence

The RCU and the ESCDs on the remote copy connections must become operable before the MCU accepts to first write I/O operation to the M-VOLs.

After the power-on-reset sequence of the MCU, It communicates with the RCU in order to confirm the status of the R-VOLs. If it is not possible, the MCU retries the confirmation until it is successfully completed or the MCU accepts the first write I/O operations to the M-VOLs.

If the MCU accepts the first write I/O operation before completing the confirmation, the MCU will suspend the HRC volume pair. This situation is critical because the status of the R-VOL can not be changed, that is, remains “duplex” state.

(4) Executing ICKDSF to HRC Volume Pairs

The updates by the channel programs which specify “diagnostic authorization” or “device support authorization” are not reflected to the R-VOL. ICKDSF commands which issue the write I/O operations to the M-VOL must be controlled. The HRC volume pairs must be suspended or terminated before performing ICKDSF commands.

Refer to “3.11.6(2) Suspending and Resuming the HRC Volume Pairs” for the operations to suspend and resume the HRC volume pairs.

3.11.7 HRC Error Recovery

(1) Preparing for Disaster Recovery

(a) Considering Fence Level Parameter

Table 3.11.7-1 shows how the fence level parameter of the Add Pair operation has an effect on the write I/O operations to the M-VOL after the HRC volume pair has been suspended. You should select one of the fence level considering the “degree of the currency” of the R-VOL required by your disaster recovery planning. The SVP or remote console, which is connected to either the MCU or the RCU, can display the fence level parameter which has been set to the HRC volume pairs.

Table 3.11.7-1 Effect of the Fence Level Parameter

Failure		Subsequent write I/O operations to the M-VOL will be ...		
		“Data”	“Status”	“Never”
1)	The update copy has failed,	Rejected *	—	—
2)	(1) & however the status of the R-VOL could have been successfully changed to “suspended” state.	Rejected *	accepted	accepted
3)	(1) & furthermore the status of the R-VOL could not have been changed to “suspended” state.	Rejected *	Rejected *	accepted

*: Sense bytes includes “command reject” and x'0F' of format/message.

Note: “Data” and “Status” has an effect when an HRC volume pair of “duplex” state is suspended. For HRC volume pairs which are in “pending duplex” state, subsequent write I/O operations will not be rejected regardless of Fence Level parameter.

1) Fence Level = "Data"

The data of the R-VOL is always identical to the M-VOL if once the HRC volume pair has been successfully synchronized. You can reduce the time to analyze whether the R-VOL is current or not in your disaster recovery procedures.

However, this parameter will make the M-VOL not accessible from your applications whenever the HRC copy activity has failed. Therefore you should specify this parameter to the most critical volumes for your disaster recovery planning.

Most of the database system supports duplexing the critical files, for example log files of DB2, for its file recovering capability. It is recommended to locate the duplexed files on the volumes in the physically separated DKC515I subsystems, and establish HRC volume pairs for each volumes by using physically separated remote copy connections.

Note 1: If the failure has occurred before completing the initial copy, the R-VOL can not be used for disaster recovery because the data of the R-VOL is not fully consistent yet. You can become aware of this situation with referring status of the R-VOL in your disaster recovery procedures. Refer to "3.11.7(2)(b) Analyzing the Currency of R-VOLs" for more detailed description.

Note 2: Only the difference between the HRC volume pair must be the last update from the host processor. HRC is a synchronous remote copy. The MCU reports a "unit check" if it detects the failure on the write I/O operation including the update copy to the R-VOL. Therefore, the operating system and the application program does not regard the last (failed) I/O operation as successfully completed.

This parameter is functionally equivalent to "CRIT=YES" parameter for IBM PPRC.

2) Fence Level = Never

The subsequent write I/O operations to the M-VOL will be accepted even if the HRC volume pair has been suspended. Therefore the contents of the R-VOL can become “older” (behind the currency of corresponding M-VOL) if the application program continue updating the M-VOL. Furthermore, the status of the R-VOL which will be obtained from the RCU can not be in a “suspended” state.

To use this parameter, your disaster recovery planning must satisfy the following requirements:

- The currency of the R-VOL should be decided by referring the error message which might have been transferred through the error reporting communications or analyzing the R-VOL itself with other files which are confirmed to be current.
- The data of the R-VOL should be recovered by using other files which are ensured to be current.

This parameter is functionally equivalent to “CRIT=NO” parameter for IBM PPRC.

3) Fence Level = Status

The level of this parameter is between “Data” and “Never”. Only when the status of the R-VOL can be ensured, the subsequent write I/O operations to the M-VOL will be permitted. Therefore the disaster recovery procedure of deciding the currency of the R-VOL can be reduced.

(b) Transferring the Sense Information through Error Reporting Communications

When the HRC volume pair is suspended, the MCU generates the sense information which notifies the host processor of the failure. This will help in deciding the currency of the R-VOLs in the disaster recovery procedures by transferring the sense information, or the system console message caused by the sense information, with the system time stamp information.

The sense information can be selected out of:

- x'FB' of format/message. The sense information is compatible with IBM PPRC and result on a corresponding system console message, for example IEA491E of MVS, if the operating system supports it.
- service information message whose reference code means that the HRC volume pair has been suspended.

Note: The first version of HRC is not completely certified under the operating system which does not support IBM PPRC. Therefore the x 'FB' sense information must be selected.

The error reporting communications are essential if you use the fence level of "Status" or "Never".

(c) File Recovery Procedures Depending on Installations

HRC is a synchronous remote copy. All updates to the M-VOLs are copied to their R-VOLs before completing each channel program of the write I/O operations. When the HRC volume pairs have been suspended or the MCU has become inoperable due to a disaster, therefore, many data "in progress" could remain in the R-VOLs. That is, some data set might be still opened, or some transactions might not be committed yet. All breakdown cases should be previously considered.

Therefore, even if you have selected the fence level of "Data" for all HRC volume pairs, you should establish the file or volume recovery procedures. The situation which should be assumed is similar to that where the volumes have become not accessible due to the disk controller failure in non-remote copied environment.

If you use the fence level of “Status” or “Never”, the suspended R-VOLs could become “ancient” compared to other volumes. This situation might cause a data inconsistency problem among several volumes.

You should prepare, in your disaster recovery, for recovering some files or some volumes which have become “ancient” by using:

- files for file recovery, for example DB2 log files, which have been confirmed to be current. To ensure the currency of these files, it is recommended to use the fence level of “Data” for these critical volumes.
- the sense information with the system time stamp which have been transferred through the error reporting communications.
- full consistent file or volume backups, if the sense information and the system time stamp can not be used.

(d) CSUSPEND/QUIESCE for IBM PPRC

PPRC recommends to customers to establish their disaster recovery planning where the CSUSPEND/QUIESCE TSO command is programmed to be issued responding to the IEA491E system console messages. This procedure intentionally suspend the remaining volume pairs when some volume pairs have been suspended due to a disaster.

The CSUSPEND/QUIESCE TSO command will be supported as the enhancement to HRC.


(e) All SIM of the HRC clear option

For the purpose of the restrain to report the HRC SIM which should be generated at the disaster and the recovery operations during the OS IPL, “Clear SIM” button is provided on HRC Main Control Screen at the SVP and the RMC. This function is able to use for the disaster recovery operation, and specifies that the all Remote Copy SIM in the subsystem should be deleted.

(2) Disaster Recovery Procedures - Switching to the Recovery System

(a) Summary

① Primary system and MCU becomes inoperable due to disaster.



(b) Analyzing the Currency of R-VOLs (Step 1)

1) Analyzing Status of the R-VOLs and Fence Level Parameter

Table 3.11.7-2 Currency of the R-VOLs

Status of R-VOL	Fence Level for this HRC volume pair		
	Data	Status	Never
Simplex	To be confirmed	To be confirmed	To be confirmed
Pending Duplex	Inconsistent	Inconsistent	Inconsistent
Duplex	Current	Current	To be analyzed
Suspended (Initial Copy Failed)	Inconsistent	Inconsistent	Inconsistent
Suspended (by other reason)	Current	Suspected	Suspected

Table 3.11.7-2 shows how to analyze the currency of the R-VOLs referring the status of the R-VOLs and the fence level parameter which have been specified when establishing the HRC volume pairs.

The status of the R-VOLs must be obtained from the RCU in your disaster recovery procedures.

The fence level parameter must be previously field since it cannot be obtained From RCU.

The meaning of the results or further actions shown in each column of Table 3.11.7-2 are as follows:

To be confirmed This volume does not belong to any HRC volume pair. If you have certainly established the HRC volume pair for this volume and you have never deleted it, you should regard this volume as inconsistent.

Inconsistent The data on this volume is inconsistent because not all cylinders have successfully been copied to this volume yet. You can not use this volume for the applications unless this volume is initialized (or successfully copied from the M-VOL at later time).

Current The data on this volume is completely synchronized with the corresponding M-VOL.

To be analyzed The currency on this volume can not determined. To determine the currency, further analysis described in (2) of this section should be performed.

Suspected

The data on this volume must be “older”, behind the currency of corresponding M-VOL. You should restore the consistency of this volume at least, and the currency of this volume if required. The system time information which might have been transferred through the error reporting communications or time of suspension obtained from the Pair Status operation will help you decide the last time when this volume was current.

2) Further Analysis by Referring to Other Information

The M-VOLs, to which the fence level parameter has been set to “Never”, will accept the subsequent write I/O operations regardless of the result of communication to change the R-VOL into the “suspended” state. Therefore, the status of the R-VOL should be analyzed by referring to the following information:

- The sense information through the error reporting communications. If the sense information which denote the suspension of this volume is found, you can return to Table 3.11.7-2 with assumption of the “suspended” state.
- The status of the M-VOL obtained from the MCU, if possible. You should return to Table 3.11.7-2 with assumption of the same status as the M-VOL and fence level of “Status”.
- The other related files, for example DB2 log files, which have been confirmed to be current.

(c) Terminating HRC Volume Pairs (Step 2)





The “Delete Pair” operation to the RCU terminates the specified HRC volume pairs. These R-VOLs will be changed to “simplex” state. Specified “Delete Pair by Force” option and “Delete All Pairs” option the all volume pairs in the same serial number of the MCU and the same CU image of the MCU should be deleted. Refer to “3.11.5(7) Delete Pair Operation”.


(d) Vary all R-VOLs online (Step 3)

In the case of the OS IPL, execute the “Clear SIM” operation at the SVP or the RMC before OS IPL.

(3) Disaster Recovery Procedures - Returning to the Primary Site

(a) Summary





(b) Terminating the HRC Settings Remaining in the MCU (Step2)

After the DKC515I subsystem becomes operable, the remaining registration of the HRC volume pairs and the RCU should be deleted by performing the Delete Pair operation and Delete RCU operation respectively.

To complete the Delete Pair operation, the “delete pair unconditionally” option is required because the original R-VOLs do not belong presently to any HRC volume pairs. The MCU will change the specified M-VOLs into “simplex” state without checking the current status of the corresponding R-VOLs.

Specified “Delete Pair by Force” option and “Delete All Pairs” option the all volume pairs in the same RCU should be deleted.

Note that the status of M-VOLs may be “Suspended (Delete Pair to RCU)” because of Delete Pair operation issued to the RCU in step 2 of “3.11.7(2) Disaster Recovery Procedures - Switching to the Recovery System”. It is normal condition in this situation.

Before performing the Delete RCU operation, all HRC volume pairs must be deleted.

If you want to use same remote copy connections for step 3, the serial interface ports which have been set to the RCP mode should be changed to the LCP mode by the Port operation.

(c) Establish HRC with the Reverse Direction (Step3)

The HRC volume pair should be established with the reverse direction to synchronize the original M-VOLs with the original R-VOLs. The procedures for this step are same as those described in “3.11.6(1) Setting Up HRC Volume Pairs”. Note that the DKC515I subsystems in the original primary site and the recovery site are treated as the RCUs/R-VOLs and the MCUs/M-VOLs respectively.

Do not select “only out-of-sync cylinders” or “none” parameter to the Add Pair operations. The volumes in the original primary site are now behind the volumes in the recovery site. Furthermore the updates to the volumes in the recovery site have not been accumulated in cylinder bit map.

(d) Terminate Applications and HRC Settings at the Recovery Site (Step 4~6)

HRC settings with the reverse direction must be deleted after halting the applications in the recovery site (step 4) and confirming that all HRC volume pairs are in “duplex” state (step 5).

Specified “Delete Pair by Force” option and “Delete All Pairs” option the all volume pairs in the same RCU should be deleted.

If you want to use same remote copy connections for step 7, the serial interface ports which have been set to the RCP mode should be changed to the LCP mode by the Port operation.

(e) Establish HRC Pair with the Original Direction and Start Applications (Step 7)

The HRC volume pair should be established with the original direction to synchronize the original M-VOLs with the original R-VOLs. The procedures for this step are same as those described in “3.11.6(1) Setting Up HRC Volume Pairs”.

Do not select “only out-of-sync cylinders” or “none” parameter to the Add Pair operations. The volumes in the original primary site are now behind the volumes in the recovery site. Furthermore the updates to the volumes in the recovery site have not been accumulated in cylinder bit map.

3.12 HMRCF & HOMRCF

3.12.1 Overview

(1) Main object

- 1) Reduce Backup time.
- 2) Easy testing before system upgrade with the data whose applications are actually used on the system.

(2) Function Outline

- 1) Making duplicated volumes.
- 2) There is no conflict on volume because the duplicated volumes are on another physical storage subsystem.
- 3) Three destination volumes can be with one master volume.
Those three pairs can be split independently.
- 4) HMRCF (Hitachi Multiple RAID Coupling Feature) can be controlled by PPRC Command interface.
- 5) HOMRCF (Hitachi Open Multiple RAID Coupling Feature) can be controlled from RAID Manager.

Table 3.12.1-1 Outline of HMRCF & HOMRCF

No.	Items	Specification
1	Coupling object	One logical volume (LDEV)
2	Support emulation type of LDEV	HMRCF: 3390-3/3A/3B/3C/3R/9/9A/9B/9C/9L/LA/LB/LC/M/MA/MB/MC, 3380-3/3A/3B/3C/K/KA/KB/KC/F, NF80-K/KA/KB/KC/F, Emulation types and CVSs of them HOMRCF: OPEN-3/8/9/E/L/V, emulation types and CVSs (except OPEN-L emulation type) and LUSE of them (*1)
3	Requirement for create a pair	(1) Pair LDEVs have to be a same track format and same capacity. (2) Pair LDEVs have to exist in a same subsystem. (3) It is not possible to share a destination volume at same time.
		• When forming a pair using volumes of the NAS system.
		Master Volume Destination Volume Support
		NAS System Volume User Volume Supported
		User Volume NAS System Volume Not Supported
		NAS System Volume NAS System Volume Not Supported
4	Support of CVS (Customized Volume Size)	HMRCF : Supported HOMRCF : Supported
5	Combination of RAID level between master volume and destination volume	RAID1(2D+2D)←→RAID1(2D+2D) RAID5(3D+1P or 7D+1P)←→RAID5(3D+1P or 7D+1P) RAID5(3D+1P or 7D+1P)←→RAID1(2D+2D) RAID6(6D+2P)←→RAID6(6D+2P) RAID6(6D+2P)←→RAID1(2D+2D) RAID6(6D+2P)←→RAID5(3D+1P or 7D+1P)
6	Data protection	There is a parity protection for both master volume and destination volume.
7	RESYNC pattern	HMRCF/HOMRCF supports 2 types of RESYNC pattern. From Master Volume data to Destination volume and from Destination Volume to Master Volume
8	Time for transition from Duplex to Split.	3 min./VOL(3390-3) without IO(Depend on the number of pairs and the load of DKC)
9	When the destination volume can be accessed from HOST.	The destination volume can be accessed at only Split status.
10	Cooperation with HRC/XRC	HMRCF : Supported The master volume of HMRCF can be an M-VOL or R-VOL of HRC. • HRC volumes are shared with Universal Replicator for Mainframe volumes can't be HMRCF pair volumes.

*1: "0" is added to the emulation type of the V-VOLs (ex. OPEN-0V).

When you create a Copy-on-Write Snapshot pair, specify the volume whose emulation type is displayed with "0" like OPEN-0V as the S-VOL.

11	Cooperation with HORC	HOMRCF : Supported The master volume of HOMRCF can be an M-VOL or R-VOL of HORC. HORC volumes are shared with Universal Replicator volumes can't be HOMRCF pair volumes.
	Cooperation with HORC (only for HOMRCF)	Supported - The M-VOL or R-VOL of HORC can be a primary volume of HOMRCF. - The secondary volume of HOMRCF can be an M-VOL of HORC.
	Cooperation with HMBR (only for HOMRCF)	Supported - A primary volume of HOMRCF can be accessed from HOST like a usual volume. - A secondary volume of HOMRCF can not be accessed from HOST if the pair is not in SPLIT status. In the SPLIT status, a secondary volume of HOMRCF can access from HOST like a usual volume.
12	Cooperation with HIHSM	Supported - The source, destination or RESERVE volume of HIHSM can not be the primary, secondary or RESERVE volume of HMRCF/HOMRCF. - The primary, secondary or RESERVE volume of HMRCF/HOMRCF can be the source volume of HIHSM. But if HMRCF/HOMRCF P-VOL or RootVOL already has 3 pairs, it can not be the source volume of HIHSM. - When HMRCF/HOMRCF pair which is combined with HIHSM is split, the migration of HIHSM is canceled.
13	Cooperation with Hi-Copy	Not supported Hi-Copy pair volumes (primary/secondary) can't be HOMRCF pair volumes (primary/secondary).
14	Cooperation with Universal Replicator (only for HOMRCF)	Supported - Universal Replicator volumes (primary / secondary) can be a primary volume of HOMRCF. - The secondary volume of HOMRCF can't be Universal Replicator volumes (primary / secondary). - The journal volume of Universal Replicator can't be a primary or secondary volume of HOMRCF. - Universal Replicator volumes are shared with HORC volumes can't be HOMRCF pair volumes.
15	Cooperation with Universal Replicator for Mainframe (only for HMRCF)	Supported - Universal Replicator for Mainframe volumes (primary / secondary) can be a primary volume of HMRCF. - The secondary volume of HMRCF can't be Universal Replicator for Mainframe volumes (primary / secondary). - The journal volume of Universal Replicator for Mainframe can't be a primary or secondary volume of HOMRCF. - Universal Replicator for Mainframe volumes are shared with HRC volumes can't be HMRCF pair volumes.

16	ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2 Option function	The ShadowImage-FlashCopy (R) Option function provides the same function as IBM Flash Copy. You can operate the ShadowImage-FlashCopy (R) Option function by using the PPRC TSO and DFSMSdss. The ShadowImage-FlashCopy (R) Option function forms a pair by virtually or physically copying S-VOL data to the T-VOL. (A pair formed by means of the ShadowImage-FlashCopy (R) Option function is especially called relationship.) Concerning ShadowImage-FlashCopy (R), locations of extents of a copy source and a copy destination must be the same. Concerning ShadowImage-FlashCopy (R) version2, locations of extents different from each other can be specified. ShadowImage-FlashCopy (R) forms one relationship for each unit of volume even when the two or more extents are specified; ShadowImage-FlashCopy (R) version2 forms one relationship for each unit of specified extent. When the relationship is established, a host can execute a reading/writing from/to T-VOL data that is a virtual or physical copy of S-VOL data. When establishing the relationship, the user can specify an extent for copying (to be referred to as extent). The relationship is canceled at the time when a copying of data in the extent is completed. Note: ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2 do not start at the same time. For the functions of ShadowImage-FlashCopy (R) Option and ShadowImage-FlashCopy (R) version2 Option, refer to the chapters, "ShadowImage-FlashCopy (R) Option Function" and "ShadowImage-FlashCopy (R) version2 Option Function" in the "ShadowImage- Mainframe User's Guide."
----	--	---

17	At-Time Split Function (HOMRCF)	The HOMRCF At-Time Split function applies to HOMRCF pairs that belong to a consistency group, and enables you to create Source volumes of all Target volumes in the same consistency group, at the time when the pairsplit command is executed using the Command Control Interface (CCI) software from the UNIX[®]/PC server host to the 9900V subsystem. Note: For further information on Command Control Interface, please refer to the Hitachi Lightning 9900™ V Series and Lightning 9900™ Command Control Interface (CCI) User and Reference Guide (MK-90RD011). An HOMRCF consistency group is a user-defined set of HOMRCF volume pairs used for the At-Time Split function. Users can defined a consistency group by using CCI on the UNIX[®]/PC server host. HOMRCF consistency groups also correspond to the groups registered in the CCI configuration definition file. HOMRCF consistency groups have the following restrictions: • You can configure up to 128 consistency groups in a subsystem with HOMRCF consistency group. • A number (0~127) is assigned to each consistency group. You can specify a consistency group number when you create HOMRCF pairs. If you do not specify a number, then the 9900V subsystem assigns a number automatically. • You can define up to 1,024 HOMRCF pairs in a consistency group. However, for LUSE volumes that contain n LDEVs, you should count as n pairs. Note: For further information on LUSE volumes, please refer to the Hitachi Lightning 9900™ V Series LUN Expansion (LUSE)/Virtual LVI/LUN (VLL) User's Guide (MK-92RD104). • HOMRCF pair and HOMRCF pair cannot share with same consistency group. • To configure HOMRCF consistency groups, you can only use the CCI software. However, to confirm the HOMRCF consistency group numbers, you can also use the Storage Navigator.
----	---------------------------------	---

18	At-Time Split Function (HMRCF)	To use this function, utility of host is necessary. The HMRCF At-Time Split function applies to HMRCF pairs that belong to a consistency group, and enables you to create Source volumes of all Target volumes in the same consistency group, at the time that is specified by PPRC command. An HMRCF consistency group is a defined with Storage Navigator by user before creating pair. • You can configure up to 128 consistency groups in a subsystem with HMRCF consistency groups. • A number (0~127) is assigned to each consistency group • You can specify a consistency group number when pair is created. • To configure HMRCF consistency group numbers, you can also use PPRC command or the Storage Navigator. • You can define up to 1,024 HMRCF pairs in a consistency group. • HMRCF pair and HOMRCF pair cannot share with same consistency group.
19	The maximum number of pairs	The maximum number of pairs is as follows. The maximum number of pairs is the sum of the numbers of pairs of the HMRCF, HOMRCF, HIHSM and Hi-Copy.

(1) In the case of the open system

A number of formable pairs varies depending on the state of addition of shared memories for the differential tables, a type and a capacity (size specified as CVS) of an emulation to be used.

The shared memories for the differential tables can be added.

When the shared memories for the differential tables are not added, the number of the differential tables is 13,652.

When the shared memories for the differential tables are added, the number of the differential tables is 30,718.

In the both cases above, total number of the formable pairs is 8192 or less.

Note: Since number of volumes usable as pairs is up to 16383 when a command device is installed, the maximum number of pairs is 8191 in the case of a configuration of PVOL:SVOL = 1:1.

The number of sheets of differential tables a pair used according to the emulation type and the capacity is calculated by using the following expression.

(a) For the emulation type of the open system (OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-L)

Differential tables per pair (α) = $\uparrow ((X) \div 48[\text{KB}] + (Y) \times 15) \div (Z) \uparrow$

(X): User-specified size [KB]

(Y): Number of control cylinders [Cyl] (See Table 3.12.1-2.)

(Z): Number of slots managed by one differential table. (1,916 × 32)

$\uparrow \uparrow$: rounding up decimal digits

For example, differential tables per pair for the emulation type OPEN-3 is

$(2,403,360 \div 48 + 8 \times 15) \div (1,916 \times 32) = \uparrow 0.81860 \uparrow = 1$

(b) For the emulation type of the open system (OPEN-V)

Differential tables per pair (α) = $\uparrow ((X) \div 256[\text{KB}]) \div (Z) \uparrow$

(X): User-specified size [KB]

(Z): Number of slots managed by one differential table. (1,916 × 32)

$\uparrow \uparrow$: rounding up decimal digits

For example, differential tables per pair for the emulation type OPEN-V is

$(3,019,898,880 \div 256) \div (1,916 \times 32) = \uparrow 192.40084 \uparrow = 193$

Table 3.12.1-2 The Number of the Control Cylinders According to the Emulation Types (for Open System)

Emulation type	Number of control cylinders (CYL)
OPEN-V	0KB (0Cyl)
OPEN-3	5,760KB (8Cyl)
OPEN-8	19,440KB (27Cyl)
OPEN-9	19,440KB (27Cyl)
OPEN-E	13,680KB (19Cyl)
OPEN-L	5,040KB (7Cyl)

Note: VLL functions are not applicable to OPEN-L volumes.

When making pairs of volumes of various emulation types coexist, the number of formable pairs is determined on the following conditions.

- In the case of a configuration of 1LU = 1 volume (when no LUSE is used)
Number of formable pairs: The maximum number of pairs that satisfies a condition,
 $\Sigma(\alpha) \leq (\beta)$

$\Sigma(\alpha)$: The sum total of numbers of differential tables per pair (α 's)

(β) : Number of differential tables in a subsystem concerned

When the shared memories for the differential tables are not added:

$(\beta) = 13652$.

When the shared memories for the differential tables are added:

$(\beta) = 30718$.

For example, when 10 pairs of OPEN-3 and 20 pairs of OPEN-L (non-CVS) are formed, when the shared memories for the differential tables are not added, the sum total of numbers of the differential tables per pair ($\Sigma(\alpha)$) is calculated as $(1 \times 10) + (13 \times 20) = 270$.

It is shown 10 pairs of OPEN-3 and 20 pairs of OPEN-L (non-CVS) can be formed judging from a comparison, $270 \leq 13652$.

- In the case of a configuration in which an LU consists of two or more volumes
(when the LUSEs are used)
Number of formable pairs: The maximum number of pairs that satisfies a condition,
$$\Sigma(\alpha) \leq (\beta)$$
$$\Sigma(\alpha)$$
: The sum total of numbers of differential tables used by volumes composing an LUSE.
(Each volume composing an LUSE uses the differential tables necessary for a mate volume to a pair and one pair.)
$$\Sigma(\beta)$$
: $\Sigma(\alpha)$ (The sum total of numbers of the differential tables for all pairs in the subsystem concerned)
$$(\beta)$$
: Number of differential tables in a subsystem concerned

For example, when forming 10 pairs of LUSEs each consists of three volumes of OPEN-3 and one volume of OPEN-E (non-CVS) and 20 pairs of LUSEs each consists of five volumes of OPEN-L, $\Sigma(\alpha)$ is calculated as $((1 \times 3) + (5 \times 1)) \times 10 + (13 \times 5 \times 20) = 1380$.

It is shown that 10 pairs of LUSEs each consists of three volumes of OPEN-3 and 20 pairs of LUSEs each consists of five volumes of OPEN-L (20G) can be formed judging from a comparison, $730 \leq 13652$.

(2) In the case of the MF system

A number of formable pairs varies depending on the state of addition of shared memories for the differential tables, a type and a capacity (size specified as CVS) of an emulation to be used.

The shared memories for the differential tables can be added in two stages.

When the shared memories for the differential tables are not added, the number of the differential tables is 13,652.

When the shared memories for the differential tables are added, the number of the differential tables is 30,718.

In the both cases above, total number of the formable pairs is 8192 or less.

The number of sheets of differential tables a pair used according to the emulation type and the capacity is calculated by using the following expression.

Differential tables per pair (α) = $\uparrow ((X) + (Y)) \times 15 \div (Z) \uparrow$

(X): User-specified size [Cyl]

(Y): Number of control cylinders [Cyl] (See Table 3.12.1-3.)

(Z): Number of slots managed by one differential table. (1,916 × 32)

$\uparrow \uparrow$: rounding up decimal digits

For example, differential tables per pair for the emulation type 3390-3 is

$(3,339 + 6) \times 15 \div (1,916 \times 32) = \uparrow 0.81836 \uparrow = 1$

Table 3.12.1-3 The Number of the Control Cylinders According to the Emulation Types (for Mainframe System) (1/2)

Emulation type	Number of control cylinders (CYL)
3380-F	22
3380-3	7
3380-K	7
NF80-F	22
NF80-K	7
3390-3	6
3390-3R	6
3390-9	25
3390-L	23
3380-KA	7
3380-KB	7
3380-KC	7
NF80-KA	7
NF80-KB	7
NF80-KC	7
3390-M	53
3390-MA	53
3390-MB	53
3390-MC	53

Table 3.12.1-3 The Number of the Control Cylinders According to the Emulation Types (for Mainframe System) (2/2)

Emulation type	Number of control cylinders (CYL)
3390-3A	6
3390-3B	6
3390-3C	6
3380-3A	7
3380-3B	7
3380-3C	7
3390-9A	25
3390-9B	25
3390-9C	25
3390-LA	23
3390-LB	23
3390-LC	23

When making pairs of volumes of various emulation types coexist, the number of formable pairs is determined on the following conditions.

Number of formable pairs: The maximum number of pairs that satisfies a condition, $\Sigma(\alpha) \leq (\beta)$

$\Sigma(\alpha)$: The sum total of numbers of differential tables per pair (α 's)

(β) : Number of differential tables in a subsystem concerned

When the shared memories for the differential tables are not added:

$(\beta) = 13652$.

When the shared memories for the differential tables are added:

$(\beta) = 30718$.

For example, when 10 pairs of 3390-3 and 20 pairs of 3390-L are formed, when the shared memories for the differential tables are not added, the sum total of numbers of the differential tables per pair ($\Sigma(\alpha)$) is calculated as $(1 \times 10) + (9 \times 20) = 190$.

It is shown that 10 pairs of 3390-3 and 20 pairs of 3390-L can be formed judging from a comparison, $190 \leq 13652$.

- (3) Case of the MF system (Forming of the relationship of ShadowImage-FlashCopy (R) version2 for each volume)

The number of sheets of differential tables a pair used according to the emulation type and the capacity is calculated by using the following expression.

(a) Source

Source differential tables per pair (α) = $\uparrow ((X) + (Y)) \times 15 \div (Z) \uparrow$

(X): User-specified size [Cyl]

(Y): Number of control cylinders [Cyl] (See Table 3.12.1-3A.)

(Z): Number of slots managed by one differential table. (1,916 × 32)

$\uparrow \uparrow$: rounding up decimal digits

(b) Target

Target differential tables per pair (α) = $\uparrow ((X) + (Y)) \times 15 \div (Z) \uparrow$

(X): User-specified size [Cyl]

(Y): Number of control cylinders [Cyl] (See Table 3.12.1-3A.)

(Z): Number of slots managed by one differential table. (1,916 × 32)

$\uparrow \uparrow$: Rounding up decimal digits

For example, differential tables per pair for the emulation type 3390-3 is sum of source differential tables and target differential tables.

(i) Source differential tables per pair

$$(3,339 + 6) \times 15 \div (1,916 \times 32) = \uparrow 0.81836 \uparrow = 1$$

(ii) Target differential tables per pair

$$(3,339 + 6) \times 15 \div (1,916 \times 32) = \uparrow 0.81836 \uparrow = 1$$

Table 3.12.1-3A The Number of the Control Cylinders According to the Emulation Types (for Mainframe System)

Emulation type	Number of control cylinders (CYL)
3380-3	7
3390-3	6
3390-3R	6
3390-9	25
3390-L	23
3390-M	53

When forming the relationships with volumes with emulation types and capacities different from each other, the number of relationships that can be formed is decided as shown below.

Number of relationships that can be formed:

Maximum number of relationships that satisfies the following inequality. $\Sigma (\alpha) \leq (\beta)$

$\Sigma (\alpha)$: The sum of numbers of differential tables ((α)'s) per relationship.

(β): Number of all the differential tables in the subsystem concerned

When the shared memories for the differential tables are not added:

(β) = 13,652

When the shared memories for the differential tables are added:

(β) = 30,718

Blank Sheet

- (4) Case of the MF system (Forming of the relationship of ShadowImage-FlashCopy (R) version2 for each data set)

When forming the relationship for each extent, resources as many as the number of differential tables per relationship are required if there is no duplicate extent in the same volume. If there is a duplicate extent, the differential tables as many as the number of differential tables per relationship multiplied by a number of the duplicates are required.

The number of relationships that can be formed is decided as shown below.

The number of relationships that can be formed: Maximum number of relationships that satisfies the following inequalities

$\Sigma (\alpha) \leq (\beta)$, and $\Sigma (\gamma) \leq 32,768$ (Maximum number of managed relationships)

$\Sigma (\alpha)$: The sum of numbers of required differential tables ((α)'s)

(β) : Number of differential tables in the subsystem concerned

When the shared memories for the differential tables are not added:

$(\beta) = 13,652$

When the shared memories for the differential tables are added:

$(\beta) = 30,718$


$\Sigma (\gamma)$: The sum of numbers of relationships ((γ)'s)

In regard to the 3390-3/3R, the maximum number of the differential tables available as source volumes is 16. When the 16 differential tables are already used for the volume and an extent of data of the relationship already established duplicates with all the differential tables, a relationship cannot be formed.

In the case of the 3390-9, it can use 48 differential tables.

In the case of the 3390-L, it can use 144 differential tables.

Example: The following figure shows a case where seven relationships of the 3390-3 and three relationships of the 3390-L (32,760 CYL's) are formed.



The sum of numbers of the differential tables per relationship ($\Sigma (\alpha)$) is:

$$\text{The sum of numbers of differential tables of A to H} = (3 + 1) + 1 + 1 + 1 + 1 + 1 + (9 \times 3) + 9 + 9 = 54$$

Requirements for the number of relationships that can be formed are:

$$\Sigma (\alpha) \leq (\beta), \text{ and } \Sigma (\gamma) \leq 32,768$$

$$54 \leq 13,652 \text{ (When the shared memories for the differential tables are not added) or}$$

$$54 \leq 30,718 \text{ (When the shared memories for the differential tables are added)}$$

$$10 \leq 32,768$$

As shown above, the requirements for the number of relationships that can be formed are satisfied. Therefore, the relationships in the configuration shown above can be formed.

For the forming of relationship for each data set of ShadowImage-FlashCopy (R) version2, refer to the chapter, “ShadowImage-FlashCopy (R) version2 Option Function” in the “ShadowImage-Mainframe User’s Guide.”

In case of HMRCF/HOMRCF volume is selected for HIHSM volume

	P-VOL	S-VOL	RootVOL	NodeVOL	LeafVOL	ReserveVOL
Source VOL	Possible ^{*1}	Possible	Possible ^{*1}	Possible ^{*2}	Command Reject	Possible
Target VOL	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject
ReserveVOL	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject

*1: It is impossible if HMRCF/HOMRCF P-VOL or RootVOL already has 3 pairs.

*2: It is impossible if HMRCF/HOMRCF NodeVOL is already paired with 2 LeafVOLs.
If you want to execute migration of HIHSM, you need to delete the pair and reset ReserveVOL of HMRCF/HOMRCF.

- (5) Differences between ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2
Table 3.12.1-3B shows differences between ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2.

Table 3.12.1-3B Differences between ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2

#	Item	ShadowImage-FlashCopy (R)	ShadowImage-FlashCopy (R) Version2
1	Required program products and memory	ShadowImage-Mainframe program product ShadowImage-FlashCopy (R) program product	ShadowImage-Mainframe program product ShadowImage-FlashCopy (R) version2 program product The shared memory is required at the location 1.
2	Sizes of S-VOL (source volume) and T-VOL (target volume) at the time when copying each volume	$S-VOL = T-VOL$	$S-VOL \leq T-VOL$
3	S-VOL/T-VOL LSS	Relationship can be formed with the same LSS's.	Relationship can be formed with the same/different LSS's.
4	Extent of relationship	An extent can be specified for each volume or it can be specified when locations of a copy source and a copy destination are the same.	An extent can be specified for each volume or it can be specified when locations of a copy source and a copy destination are the same or different.
5	Unit of relationship management	The relationships are managed for each volume. Even if two or more data sets are specified for the same volume, one relationship is formed.	The relationships are managed for each extent that is specified. They are managed for each volume when the volume is specified. When the two or more data sets are specified for the same volume, an independent relationship is formed for the each data set.
6	Number of relationships of copy source volume	One relationship can be formed for one volume.	Up to 16 relationships can be formed for one extent.
7	Establishment of relationships of ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2	The relationship can be formed with a HMRCF volume in the Simplex status. Besides, the relationship can be formed using an S-VOL and T-VOL in the Split or Duplex status as copy sources. An S-VOL that already has three T-VOLs of HMRCF cannot be a copy source of ShadowImage-FlashCopy (R) version2.	The relationship can be formed with a HMRCF volume in the Simplex status. Besides, the relationship can be formed using an S-VOL in the Split or Duplex status as a copy source. Even with an S-VOL that already has three T-VOLs of the HMRCF, up to 16 relationships can be formed using if as a copy source of ShadowImage-FlashCopy (R) version2.
8	Use with the other copy solution	The ShadowImage-FlashCopy (R) relationship can share volumes with the other copy solution such as TrueCopy-Mainframe (TC-MF), Extended Remote Copy (XRC), Universal Replicator-Mainframe (UR-MF), and ConcurrentCopy (CC).	The ShadowImage-FlashCopy (R) version2 relationship can share volumes with HMRCF, TC-MF, XRC, CC, and UR-MF only. However, IBM copy solution can share volumes with "the other copy solutions" shown to the left.

In case of HMRCF/HOMRCF volume is selected for HIHSM volume

	P-VOL	S-VOL	RootVOL	NodeVOL	LeafVOL	ReserveVOL
Source VOL	Possible (*1)	Possible	Possible (*1)	Possible (*2)	Command Reject	Possible
Target VOL	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject
ReserveVOL	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject	Command Reject

*1: It is impossible if HMRCF/HOMRCF P-VOL or RootVOL already has 3 pairs.

*2: It is impossible if HMRCF/HOMRCF NodeVOL is already paired with 2 LeafVOLs.

If you want to execute migration of HIHSM, you need to delete the pair and reset ReserveVOL of HMRCF/HOMRCF.

In case of HIHSM volume is selected for HMRCF/HOMRCF volume

	Source VOL	Target VOL	ReserveVOL
P-VOL	Command Reject	Command Reject	Command Reject
S-VOL	Command Reject	Command Reject	Command Reject
RootVOL	Command Reject	Command Reject	Command Reject
NodeVOL	Command Reject	Command Reject	Command Reject
LeafVOL	Command Reject	Command Reject	Command Reject
ReserveVOL	Command Reject	Command Reject	Command Reject

If you want to add pair of HMRCF/HOMRCF, you need to cancel HIHSM migration.

⚠ CAUTION


Copy process is done asynchronously with HOST i/o according to differential bit map. Differential bit map is recorded on shared memory. So if shared memory is lost by offline micro exchange or volatile PSON etc., DKC lost differential bit map.

In these cases DKC treat as whole volume area has differential data, so copy process will take longer time than usual. And if the pair is SPLIT-PEND status, the pair become SUSPEND status because lost of differential bit map.

Primary volumes and secondary volumes of HMRCF/HOMRCF pairs should be placed on many RAID groups separately. And HMRCF/HOMRCF pairs which are operated at the same time should be placed in other RAID groups. HMRCF/HOMRCF pairs which are concentrated at very few RAID groups may influence HOST I/O performance.

If DKC is busy, increase Cache, DKA and RAID groups. And secondary volumes of HMRCF/HOMRCF pairs should be placed in the increased RAID groups.

HMRCF/HOMRCF pairs in very busy DKC may influence HOST I/O performance.



3.12.2 Construction of HMRCF & HOMRCF

- HMRCF & HOMRCF can be controlled from SVP, Remote Console and HOST.
- DKA with LA exchange has to exist in the subsystem.




Fig. 3.12.2-1 Construction of HMRCF & HOMRCF

3.12.4 Interface

(1) Outline

HMRCF & HOMRCF support a command set to control HMRCF & HOMRCF functions. This command set is a common interface in a subsystem. So the commands from different HOSTs are translated to the HMRCF & HOMRCF command at each command process.




Fig. 3.12.4-1 Outline of HMRCF & HOMRCF IF

Notice: It is necessary to define Command Device before using RAID Manager on OPEN HOST.

Do not define Command Device on a heavy-load path.

(2) HMRCF & HOMRCF operation


Table 3.12.4-1 HMRCF & HOMRCF operation

No.	Command	Operation
1	Duplex	Creates a pair and start initial copy
2	Split	Splits the pair
3	RESYNC	Resumes the pair and start differential copy
4	Simple	Deletes the pair
5	Suspend	Suspends the pair action
6	Status Check	Requires the status information
7	Reserve	Marks and Unmarks the volume for a candidate of destination volume

3.12.5 Cascade function

Cascade function makes a pair with an existed Target volume as a new Source volume. See the figure below.

This function is available for HOMRCF only.



No.	Content	Specification
1	Pair structure	A Target volume of L1 pair (=Node volume) can be a Source volume of L2 pair.
2	Number of copies	Root : Node = 1 : 3 Node : Leaf = 1 : 2 Root : (Node + Leaf) = 1 : 9
3	Split pair condition	L2 pair is able to execute split pair request only when the L1 pair is in the split status.
4	Delete pair condition	- No conditions. - When L1 pair is deleted, L2 pair becomes L1 pair.
5	Combination with HORC	Possible. However, Node volume and Leaf volume are treated as a target volume.
6	Combination with HIHSM	Possible. However, Leaf volume cannot be moved.

- Name of volume type
 - Source volume of the first pair : Root Volume
 - Target volume of the first pair : Node Volume
 - Source volume of the 2nd pair : Node Volume
 - Target volume of the 2nd pair : Leaf Volume
- Name of pair
 - The first pair (A pair of root volume is source volume) : L1 pair
 - The second pair (A pair of node volume is source volume) : L2 pair
- Name of pair chain
 - A chain of L1 pair and L2 pair with a node volume: stream

3.12.6 Reverse-RESYNC

(1) Reverse-RESYNC Function/Quick Restore Function

The Reverse-RESYNC function is an extension of the RESYNC function of the MRCF.

The Quick Restore function is a similar function with Reverse-RESYNC, but it speeds up the operation.

When a pair in the Split status is requested to perform the Reverse-RESYNC, the differential data between the target volume and the source volume is copied to the source volume from the target volume.

When a pair in the Split status is requested to perform the Quick Restore, a volume map in DKC is changed to swap contents of Source volume and Target volume without copying the Source volume data to the Target volume. The Source volume and the Target volume are resynchronized when update copy operations are performed for pairs in the Duplex status.

Note on RAID Level and DCR swap:

The Quick Restore operation changes locations of the data for primary volumes and secondary volumes and location of DCR of HMRCF/HOMRCF pairs. Therefore, the operation may change RAID levels and HDD types of the volumes. For example, if the primary volume is RAID1 and the secondary volume is RAID5, Quick Restore operation changes the primary volume to RAID5 and the secondary to RAID1. RAID6 volume is also similar.

If you want to go back to the previous state, follow the actions below:

- step1 : Stop HOST I/O to the pair
- step2 : Split the pair
- step3 : Perform Quick Restore for the pair
- step4 : Restart HOST I/O to the pair

Due to the replacement of DCR setting locations, you must operate 1 or 2 shown below.

1. Set the same DCR location for Source volume and Target volume.
2. Reset the DCR settings of Source volume and Target volume before Quick Restore, and set DCR of Source volume and Target volume after the pair transits to the Duplex status by Quick Restore.

Unless you perform the operation above, I/O performance to the same data may be down for the change of the locations of cache-resident area after Quick Restore.




Fig. 3.12.6-1 Normal RESYNC Process




Fig. 3.12.6-2 Reverse RESYNC Process

(2) Specifications

No.	Item	Description
1	RESYNC copy pattern	- The data of the target volume is copied to the source volume. - The copy pattern can be selected by specifying a unit of operation. Specified operation unit : SVP/RMC: In units of pair operation at a time RAID manager: In units of command
2	Copy range	In the case of the Reverse-Copy and Quick Restore in the Split status, a range for merging the writing into the source and target volumes
3	Copy format	Same format as that of a copy in the Duplex status
4	Applicable LDEV type	- HMRCF : 3390-3/3R/9/L/M/,3380-3/K/F,NF80-K/F,Emulation types and CVSs of them - HOMRCF : OPEN-3/8/9/E/L/V,emulation types and CVSs (except OPEN-L emulation type) and LUSE of them
5	Host access during copying	(1) In the case of the main frame volume - Source volume: Reading and writing disabled - Target volume: Reading and writing disabled (2) In the case of the open volume - Source volume: Writing disabled - Target volume: Reading and writing disabled Note: The reason why the source volume is not disabled to read is to make the volume recognizable by the host and it does not mean that the data is assured.
6	Specification method	SVP/ Storage Navigator: Add a specification for the RESYNC pattern onto the Pair Resync screen.
7	Conditions of command reception	- The pair concerned is in the Split status. - In the case of Quick Restore, the pair must not be combined CVS Volume and Normal Volume. - Another pair sharing the source volume is in the Suspend or Split status. → If this condition is not satisfied, the CMD RJT takes place. - When the Reverse-Resync or Quick Restore is being executed by another pair which is sharing the source volume, it is impossible to change the pair status of the pair concerned. (However, the pair deletion and pair suspension requests are excluded.) - The source volume of the pair concerned has no pair of the HRC/HORC or in the Suspend status. (See Item No.14 in this table.)
8	Status display during copying	- SVP/Storage Navigator HMRCF : RESYNC -R HOMRCF : COPY(RS-R) The display of the attribute, source or target, is not changed. - RAID manager Pair status display: RCPY The display of the attribute, source or target, is not changed.

9	Condition after normal end of the copy operation	- The pair concerned enters the Duplex status. - The conditions of the host access after the status transition are shown below. - Main frame volume - Source volume: Reading and writing enabled - Target volume: Reading and writing disabled - Open volume - Source volume: Reading and writing enabled - Target volume: Writing disabled
10	Impacts on another pair	In another pair sharing the source volume, the part actually copied becomes the difference after executing this function. Example: Pair of the other target volumes in the 1:3 configuration
11	Operation when the copying terminates abnormally	- The pair concerned enters the Suspend status. - The source volume of the pair concerned is enabled to read and write. → Data is not assured. The target volume of the pair concerned is disabled to read nor write in the case of the main frame volume and disabled to write in the case of the open volume. - The status of a pair sharing the source volume is not changed.
12	Operation when a suspension request is received during copying	Same as above
13	Relation to the cascade function	The Reverse-RESYNC and Quick Restore cannot be executed for the L2 pair.
14	Relation to the HRC/HORC	- In the case where “M-volume of the HRC/HORC” = “Source volume of the MRCF” → The Reserve-Resync and Quick Restore cannot be executed. (Command Reject) - In the case where “R-volume of the HRC/HORC” = “Source volume of the MRCF” → The Reserve-Resync and Quick Restore cannot be executed. (Command Reject) - In the case where “Target volume of the MRCF” = “M-VOL of the HRC/HORC” → The Reserve-Resync and Quick Restore cannot be executed. (Command Reject) - A pair of the HRC/HORC cannot be created with the volume of the MRCF executing the Reserve-Resync or Quick Restore. (Command Reject)

15	Relation to the Universal Replicator / Universal Replicator for Mainframe	• In the case where “Primary volume of the Universal Replicator / Universal Replicator for Mainframe” = “Source volume of the MRCF” “Pair status of the Universal Replicator / Universal Replicator for Mainframe” = “Suspend” → The Reserve-Resync and Quick Restore can be executed. “Pair status of the Universal Replicator / Universal Replicator for Mainframe” ≠ “Suspend” → The Reserve-Resync and Quick Restore cannot be executed. (Command Reject) • In the case where “Secondary volume of the Universal Replicator / Universal Replicator for Mainframe” = “Source volume of the MRCF” “Pair status of the Universal Replicator / Universal Replicator for Mainframe” = “Suspend” → The Reserve-Resync and Quick Restore can be executed. “Pair status of the Universal Replicator / Universal Replicator for Mainframe” ≠ “Suspend” → The Reserve-Resync and Quick Restore cannot be executed. (Command Reject)
----	---	---

(3) Action to be taken when the pair is suspended during the Reverse-RESYNC


The recovery procedure to be used when the pair executing the Reverse-RESYNC is suspended owing to some problem or is explicitly transferred to the Suspend status by a command from the SVP/remote console/RAID manager is explained below.

(a) Case 1: A case where the Suspend status can be recovered without recovering the LDEV concerned

This is equivalent to a case where the pair encounters an event that copying cannot be continued owing to a detection of pinned data or a staging time- out.

Or, it is equivalent to a case where the pair is explicitly transferred to the Suspend status by a command.

<<Recovery procedure>>



- (b) Case 2: A case where the Suspend status cannot be recovered unless the LDEV concerned is recovered


This is equivalent to a case that the LDEV is blocked.

To recover the blockade of the LDEV, an LDEV formatting or LDEV recovery is required. Both of them cannot be executed in the state that the MRCF pair is created. (A guard works against it.) Therefore, delete the pair once, recover the LDEV, and then create the pair once again.

However, in the pending state, caution must be taken because the data of the source volume is copied to the target volume if the pair is simply created again. Recover the blockade following the procedure below.

The following procedure is applicable just to a restoration of the source volume using the target volume. The following procedure does not include a procedure for directly restoring the source volume when the target volume is blocked.


- Case 2-1: A case where the source volume is blocked
<<Recovery procedure>>



- Case 2-2: A case where the target volume is blocked

A recovery procedure for restoring data of the target volume is added because the copy source of the Reverse-RESYNC cannot be accessed.

<<Recovery procedure>>



3.12.7 ShadowImage-FlashCopy (R) Option function

Notice:

ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2 do not start at the same time. A relationship of ShadowImage-FlashCopy (R) version2 cannot be formed in the state in which a relationship of ShadowImage-FlashCopy (R) exists. When forming a relationship of ShadowImage-FlashCopy (R) version2, all the relationships of ShadowImage-FlashCopy (R) must be dissolved.

Operations that can be done for the ShadowImage-FlashCopy (R) Option pair are shown below.

Operation	Pair status		
	Simplex	F-COPY	The others
Delete Pair	See Table 2.4.	Yes	See Table 2.4.
Relationship definition	Yes	No	No
The others	See Table 2.4.	No	See Table 2.4.

The S-VOL in the F-COPY status and the T-VOL in the status other than F-COPY can be shared, so the pair in the second layer (L2=Layer2) can be formed under the pair in the first layer (L1=Layer1). (Fig. 3.12.7-1)

(When the L1 pair is the F-COPY pair, the L2 pair can not be formed. The L1 pair and the L2 pair which are other than the F-COPY pair can not be formed. The L3 (L3=Layer3) pair can not be formed.)




Fig. 3.12.7-1 HMRCF Extended Configuration Formed by Means of ShadowImage-FlashCopy (R) Option

Relation between the L1 pair statuses and operability of the L2 pair is shown below.

Relation between the L1 pair statuses and operations of the L2 pair

L1 pair status	Operation of L2 pair						Relationship definition
	Add Pair	Split Pair	Resync Pair	Reverse Resync/ Quick Restore	Suspend	Delete	
Pending	—	—	—	—	—	OK	NG
Duplex	—	—	—	—	—	OK	NG
SP-Pend	—	—	—	—	—	OK	NG
V-Split	—	—	—	—	—	OK	NG
Split	—	—	—	—	—	OK	OK
Resync	—	—	—	—	—	OK	NG
Reverse Resync / Quick restore	—	—	—	—	—	OK	NG
Suspend	—	—	—	—	—	OK	NG
F-COPY	NG	NG	NG	NG	NG	NG	NG

“—” (dash) means that the combination has no relationship to F-COPY pair.
(The operation can not be performed.)

Relation between the L2 pair statuses and operability of the L1 pair is shown below.

Relation between the L2 pair statuses and operations of the L1 pair

L2 pair status	Operation of L1 pair						Relationship definition
	Add Pair	Split Pair	Resync Pair	Reverse Resync/ Quick Restore	Suspend	Delete	
Pending	—	—	—	—	—	OK	NG
Duplex	—	—	—	—	—	OK	NG
SP-Pend	—	—	—	—	—	OK	NG
V-Split	—	—	—	—	—	OK	NG
Split	—	—	—	—	—	OK	NG
Resync	—	—	—	—	—	OK	NG
Reverse Resync / Quick restore	—	—	—	—	—	OK	NG
Suspend	—	—	—	—	—	OK	NG
F-COPY	NG	NG	NG	NG	OK	OK	NG

“—” (dash) means that the combination has no relationship to F-COPY pair.
(The operation can not be performed.)

3.12.8 ShadowImage-FlashCopy (R) version2 Option function

When there is no shared memory at Location 1, be sure to add it before installing ShadowImage-FlashCopy (R) version2. The addition of the shared memory is to be done by service personnel sent for by a user.

Notice:

ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2 do not start at the same time. A relationship of ShadowImage-FlashCopy (R) version2 cannot be formed in the state in which a relationship of ShadowImage-FlashCopy (R) exists. When forming a relationship of ShadowImage-FlashCopy (R) version2, all the relationships of ShadowImage-FlashCopy (R) must be dissolved.

(1) ShadowImage-FlashCopy (R) version2 Option functions

ShadowImage-FlashCopy (R) version2 provides a function to copy data in high-speed similarly to ShadowImage-FlashCopy (R).

ShadowImage-FlashCopy (R) version2 forms a pair by copying data of a copy source (source volume) virtually or physically to a copy destination (target volume).

A pair formed by means of ShadowImage-FlashCopy (R) version2 is called a “relationship.”

Once a relationship of ShadowImage-FlashCopy (R) version2 is established, a host can execute a reading/writing from/to target data that is a virtual or physical copy of source volume data.

When making a copy of each data set, a relationship of only the specified data set (an extent of the CCHH). This extent of data to be copied is called an “extent.” The minimum unit of the extent is a track.

(2) Use of ShadowImage-FlashCopy (R) version2 Option together with the other function

A ShadowImage-FlashCopy (R) version2 pair can be formed using an HMRCF volume in the Simplex status. Besides the pair can be formed also using a P-VOL in the Split or Duplex status as a copy source.

Table 3.12.8-1 Possibility of Volume Sharing by ShadowImage-FlashCopy (R) version2 and Other Copy Solutions

	Possibility of coexistence with ShadowImage-FlashCopy (R) version2	
	ShadowImage-FlashCopy (R) version2 S-VOL	ShadowImage-FlashCopy (R) version2 T-VOL
MRCF PVOL	Possible	Impossible
MRCF SVOL	Impossible	Impossible
XRC PVOL	Possible	Impossible
XRC SVOL	Impossible	Impossible
TC-MF M-VOL	Possible	Impossible
TC-MF R-VOL	Possible	Impossible
UR-MF PVOL	Possible	Impossible
UR-MF SVOL	Impossible	Impossible
CC PVOL	Possible	Impossible
CC SVOL	Impossible	Impossible
HIHSM	Impossible	Impossible

Note: Even if a volume can be shared by ShadowImage-FlashCopy (R) version2 and another copy solution, there may be a case where restrictions are placed on the pair status. For details of the restriction, refer to the section, “ShadowImage-FlashCopy (R) version2 Option Function” in the “ShadowImage-Mainframe User’s Guide.”

(3) Relationship expansion

The relationship expansion function increases the maximum number of the relationships ShadowImage-FlashCopy (R) version2 can establish. Normally, ShadowImage-FlashCopy (R) version2 can establish up to 32,768 relationships, but if you use the relation expansion function, ShadowImage-FlashCopy (R) version2 will be able to establish up to 1,048,575 relationships.

For details, refer to the section, “ShadowImage-FlashCopy (R) version2 Option Function” in the “ShadowImage-Mainframe User’s Guide.”

Caution: In order to enable or disable the relationship expansion function, you must delete all ShadowImage-FlashCopy (R) version2 relationships and Copy-on-Snapshot pairs beforehand.

3.12.9 Micro-program Exchange

(1) Off-line Micro-program Exchange

- ① The existence of relationship of ShadowImage-FlashCopy (R) or ShadowImage-FlashCopy (R) version2 option is checked on the WebConsole ShadowImage –S/390 screen on SVP. Refer ShadowImage –S/390 User's Guide4.2.1 The Volume List Box.
 - ①-1: In the case of existing no relationship of ShadowImage-FlashCopy(R) option. go to ②,
 - ①-2: In the case of existing relationship of ShadowImage-FlashCopy (R) option.
In the case of ShadowImage-FlashCopy (R), watch progress of the copying made by ShadowImage-FlashCopy (R) Option for a while, and then go to Step a) or b).
In the case of ShadowImage-FlashCopy (R) version2, go to Step b) because progress of the copying (expressed in percentage) cannot be displayed.
 - a) When the copy is likely to completed within permission time.
Wait for the completion of the copy.
 - b) When the copy is unlikely to completed within permission time.
Request to delete all relationship to user. However, notify user of information no longer being guaranteed T-VOL by deleting the relationship under copy.
- ② Perform micro-program exchange operation.
- ③ If required, establish the relationship of ShadowImage-FlashCopy (R) or ShadowImage-FlashCopy (R) version2 option again.

Note 1: If step 1 above is not performed, all relationship of ShadowImage-FlashCopy (R) or ShadowImage-FlashCopy (R) version2 option is delete forcibly and all T-VOL of established relationship of ShadowImage-FlashCopy (R) or ShadowImage-FlashCopy (R) version2 option is in blockade state with generating SIM (47E600).

If T-VOL is external volume, the subsystem might stand up with normal T-VOL, without blocked T-VOL.

In this case, the data of T-VOL isn't guaranteed, so you need operate either following one.

- Dataset of T-VOL is deleted
- Initialization of volume is performed

Note 2: Request user to delete all relationship of ShadowImage-FlashCopy(R) option beforehand at the time of performing off-line micro exchange.

3.12.10 Notes on powering off

When performing a powering off, take notice of the following.

Item	Note	Reason
1	(MRCF) Take care that the time required for the copying becomes longer. Make a schedule taking the copying time into consideration.	If data in the shared memory has volatilized when the next powering on is performed, the following phenomena occur. - When the pair is in the Pending or Resync status, the data, from which a copying has been completed before the powering off, is also treated as data to be copied again. - Even if no I/O has been issued, the rate of data identity does not reach 100% when the pair status is changed to Duplex. - The data that has become the one to be copied again is copied to the secondary volume after the pair status is changed to Duplex. - When the pair is in the Duplex status, the data, from which a copying has been completed before the powering off, is also treated as data to be copied again. - The rate of data identity will be 0%. - The copying of the data, which has become the one to be copied again, is performed in the state in which the pair is in the Duplex status. - When the pair is in the Split status, the whole volume will be a differential between the two volumes. - The rate of data identity will be 0%. - Data of the whole volume is copied when a resynchronization is performed.
2	(MRCF) As to a pair in the Split transitional status (SP-Pend, V-Split), complete the copying of it and put it in the Split status.	If data in the shared memory has volatilized when the next powering on is performed, the following phenomenon occurs. When the pair is Split transitional status (SP-Pend, V-Split), it is changed to Suspend.
3	(ShadowImage-FlashCopy (R) and ShadowImage-FlashCopy (R) version2 Option function) Perform a powering off of the subsystem after the copying is completed.	If data in the shared memory has volatilized when the next powering on is performed, the following phenomena occur. - The relationship is dissolved. - The secondary volume is detached. If T-VOL is external volume, the subsystem might stand up with normal T-VOL, without blocked T-VOL. In this case, the data of T-VOL isn't guaranteed, so you need operate either following one. - Dataset of T-VOL is deleted - Initialization of volume is performed
4	(HOMRCF At-Time Split function) Perform PS OFF when by split operation by At-Time Split function, after change status of all pairs belonging to consistency group is completed. (HMRCF At-Time Split function) Perform PS OFF when by pair operation of copy group specification, or split operation by split time registration, after change status of all pairs belonging to consistency group is completed.	Pair by which change status is not carried out even if it carries out PS ON may occur.

3.12.11 Copy-on-Write Snapshot option

Before installing the Copy-on-Write Snapshot option, when there is no shared memory in the location 1 of Basic PCB and the locations 1 to 4 of the Option PCB, it is necessary to surely add it. In the case of adding the shared memory, a user contacts the service personnel and performs it.


(1) Copy-on-Write Snapshot function

It is a product to copy and manage the data in the disk subsystem as well as the ShadowImage function. The Copy-on-Write Snapshot forms the pair of which the logical volume is made as the primary volume (hereafter, indicated as P-VOL), and the virtual volume (hereafter, indicated as V-VOL) is made as the secondary volume (hereafter, indicated as S-VOL).

Because the S-VOL of the Snapshot pair is the V-VOL with no reality, the S-VOL does not actually consume the capacity of the disk subsystem.

As for the Snapshot pair, only the part to be updated is copied to the pool VOL among the data of the P-VOL. Therefore, the capacity used by the entire disk subsystem can be reduced.

It is possible to access the S-VOL of the Snapshot pair. In this case, the data of the P-VOL is referred to via the S-VOL. Therefore, the load concentrates on the parity group of the P-VOL. Also, if the P-VOL becomes a failure, it cannot access the S-VOL.



(2) Pool VOL

A pool is an area to store the snapshot data acquired by the Copy-on-Write Snapshot.

The pool is configured with two or more pool-VOLs, and the snapshot data is actually stored in the pool VOL.

As a result of writing the data in the volume of the Snapshot pair, when the amount of the pool in use exceeds the capacity of the pool, the Snapshot pair becomes PSUE (status of the failure occurrence), and cannot acquire the snapshot data.

(3) Shared memories

It is necessary to secure the capacity of 512Mbytes × 2 to 2Gbytes × 2 in the option PCB for the V-VOL management area (VFS SYSAREA) where the pool is managed.

3.13 TPF

3.13.1 An outline of TPF

TPF stands for Transaction Processing Facility.

TPF is one of operating systems (OS) for mainframes mainly used for airline customer reservation systems.

To correspond to TPF, DKC must support logical exclusive lock facility and extended cache facility.

The former is a function which is called MPLF (Multi-Path Lock Facility) and the latter is a function which is called RC (Record Cache).

DKC has implemented a special version of microprogram which supports the MPLF and RC functions of TPF feature(RPQ#8B0178), described in IBM public manuals;

- (a) IBM3990 Transaction Processing Facility support RPQs (GA32-0134-03)
- (b) IBM3990 Storage Control Reference for Model 6 (GA32-0274-03)

(1) An outline of MPLF

This facility provides a means, using a DKC, to control concurrent usage of resources in host systems via use of logical locks. A logical lock may be defined for the control of a shared resource, where the sharing of that resource must be controlled. Each shared resource has its own name called Lock Name. Every Lock Name controls multiple lock states (2 to 16).

The following figure shows the outline of I/O sequence which uses MPLF.

DKC recognizes up to 16 MPLF users. In this figure, user A and user B are shown. These users may belong to the same HOST or different HOSTs. Each user must indicate MPLP (Multi-Path Lock Partition) to use MPLF. MPLP is a means of logically subdividing the MPLs (Multi-Path Locks) for a user set. The maximum number of MPLP is four. Each MPLP has numbered from 1 to 4. The process to get permission to use MPLF is called CONNECT.

The connected user executes the SET LOCK STATE process using Lock Name. The MPL corresponding to specified Lock Name is assigned to the user. This assignment is canceled by the UNLOCK process. HOSTs can share the DASD without contradiction by using this MPLF.




Fig. 3.13.1-1 An outline of MPLF

(2) An outline of RC

RC has the following two features:

- (a) Record Mode Chain
- (b) Record Caching

The following explains these features.

(a) Record Mode Chain

Record Mode Chain consists of the following 4 command chains:

- 1) Mainline Processing (Read)
- 2) Mainline Processing (Write)
- 3) Capture
- 4) Restore

To execute Record Mode Chain, a subsystem must be initialized for Record Caching, and Record Mode must be allowed for the addressed device. Under these conditions, Record Mode Chain works when Record Mode Chain is indicated in the Extended Global Attributes of Define Extent command. Otherwise, the chain is processed in a standard mode.

A Mainline Processing chain consists of a Define Extent command, a Locate Record command, and a single Read Data or Write Update Data command.

A Capture chain consists of a Define Extent command followed by a Seek command and multiple Read Count, Key, and Data commands.

A Restore chain consists of a Define Extent command, a Locate Record command, and multiple Write Update Data commands.

(b) Record Caching

Record Caching is a naming contract with Track Caching used in a standard model. At the completion of first initialization, all caches are allocated to Track Slot as a standard model. Record Cache will be allocated if Set Cache Allocation Parameters Order is issued.

3.13.2 TPF Support Requirement

(1) OS

TPF Ver.4.1.

(2) Hardware

The following table shows subsystem hardware specification for TPF support.

Table 3.13.2-1 TPF Support Hardware Specification

Item	Description
Number of MODs	Max. 16384/box
Number of LCUs/Box	Max. 64
Number of SSIDs/LCU	1
Cache/SM capacity	(Refer to INST07-20)
RAID level	5 or 1
Emulation type (1) LCU (2) Device	3990-6 or 2105-F20 3390-3 or 3390-9
Number or Multi-Path Locks	16384/LCU (when up to 16 LCUs) 4096/LCU (when 17 ~ 64 LCUs)
Option features; (1) CVS (2) DCR (3) HRC (4) HMRCF (5) HIHSM (6) HMDE (7) HMBR (8) Destage Mode (Power Lost Mode)	Available Available Available Available Available *1 (Not Available) (Not Available) (Not Available) *2

*1: HIHSM supports only a monitor function.

*2: Power Lost Mode supports only a “Memory Backup Mode” function.

3.13.3 TPF trouble shooting method

Basically TPF environment and MVS (as a standard operation system) are same in troubleshooting.

A example order is below;

- (a) Collect system error information by Syslog ...etc.
- (b) Collect DKC error information by SVP dump operation.
 - Normal dump which contains TPF dump data as well. (Refer [SVP02-540](#))
- (c) Send the above data to T.S.D.

3.13.4 The differences of DASD-TPF(MPLF) vs DASD-MVS

(1) Data-exclusive method by MPLF function


MVS environments

- (a) Logical volume(Device) is the unit of data-exclusive between several CPUs.
- (b) "Device" is owned by one CPU during CPU processing(accessing), and "Device-busy" status is reported to another CPU's accesses.
- (c) "Device-end" status is used to notify the waiting CPUs
- (d) when the device becomes free.

TPF environments


- (a) Logical "Lock" is used for this purpose, instead of logical volume(device) of MVS.
 - (b) Most Read/Write CCWs have a unique: Prefix-CCW(Set Lock) to own the target lock. And only when the request-lock is granted to, its CCW continues the following Read/Write processes.
 - DSB="4C" is for granted / DSB="0C" is for NOT-granted(wait).
 - (c) "Attention" status is used to notify the waiting CPUs when the lock becomes free.
 - (d) The relationship between Lock and Dataset is completely free.
- Usually TPF users(customers) have their own definitions.

MVS environments



- ① Reserve/Read&Write Access by CPU-A (Successful).
- ② CPU-B's trial is rejected by Device-busy (Failed).
- ③ Terminate its process and release the volume.
- ④ Free(Device-end) will be sent. CPU-B can use this volume.

TPF environments



- ① Set Lock/Read&Write process *1 by CPU-A (Successful).
- ② CPU-B's trial is rejected by Not-granted (failed).
- ③ Terminate with Unlock, by CPU-A.
- ④ Free(Attention) will be sent. *2. CPU-B can use this Dataset.

*1:Typical CCW chain:

- Set lock State(x27/order(x30));
- Read Subsystem Data(x3E);
- TIC(to be continued if granted)
- (ordinary CCW chain)

*2:This report's path/Address is usually different from above②.

Fig. 3.13.4-1 Environments of DASD-TPF and DASD-MVS

(2) No path-group

MVS environments

- (a) Each CPU establishes the Path-group on every DASD Online-device, using all the connected paths.
- (b) Channel and DASD (Control Unit) rotate the I/O service path to meet each occasion within this group.
- (c) “Device-end” status can be reported through any-path of this group.

TPF environments

- (a) TPF OS/CPU does not establish this Path-Group, even if the configuration has multiple-paths for DASD.
- (b) But the Channel rotates the I/O request-path, within the connected paths. (Like old MVS way.)
- (c) “Attention” report is restricted to one “Connect-Path” which has been defined during IPL (Vary-online) procedure.

(3) Connect Path/Device

- (a) TPF system issues “Connect order” to define ;
 - Lock tables on each DASD control-unit,
 - Report path & Device for Attention interrupt.
- (b) This order is code(x33) of Perform Subsystem Function (x27) command.
- (c) This order is issued during the IPL process of each CPU.
- (d) CPU (channel) only has the capability to change this path and device definition.

Table 3.13.4-1 Order-list of Perform Subsystem Function (x27) command

Order	Meaning	Function
x10	Commit	RC
x11	Discard	
x18	Prepare for Read Subsystem Data	
x19	Destage Modified Tracks	
x1B	Set Special Intercept Condition	
x20	Set Cache Allocation Parameters	
x21	Suspend/Resume Function	
x22	Prepare to Read Lock Data	MPLF
x30	Set Lock State	
x31	Purge Lock	
x32	Unlock	
x33	Connect	
x34	Disconnect	

For details, please see the following IBM RPQ manual;
 “IBM 3990 Transaction Processing Facility Support RPQs” (GA32-0134-03)

(4) Channel Re-drive function

MVS environments

- (a) In general, the Channel selects the most proper path (in the Path-group) for each I/O request.
- (b) In MVS environments, there is not this kind of function.

TPF environments

- (a) In TPF environments,
To keep a fast IO response & IO request-order(fast-in, fast-out),
this kind of special function has been introduced. (Thus is our conjecture.)
- (b) By the channel-monitor data,
Once IO request is rejected with Control-unit busy by DASD,
Sub-channel repeats a reconnect-trial to DASD (with some short interval)
until (1) the sub-channel gets into DASD or (2) it reaches the trial-count threshold (in this case the IO request is registered to some waiting Queue in the channel.)
- (c) And once the IO request is accepted by DASD control-unit,
next Control-unit judges this would be accepted or not using “Lock” status.

(5) Fixed Record-size

Dataset structure in DASD

In general, TPF software makes the following logical structure in DASDs.




Fig. 3.13.4-2 Logical structure in DASDs

Table 3.13.4-2 Pool records Classification

Logical (usable) size	Physical size
381 Bytes	384 Bytes
1055 Bytes	1056 Bytes
4095 Bytes	4096 Bytes

Only three lengths exist for pool records.

Table 3.13.4-3 More detailed classification

381 Record	1055 Record	4095 Record
SLT (Small, Long Term)	LLT (Large, Long Term)	4LT (4KB, Long Term)
SST (Small, Short Term)	LST (Large, Short Term)	4ST (4KB, Short Term)
SDP (Small, Long Term, Duplicated)	LDP (Large, Long Term, Duplicated)	4DP (4KB, Long Term, Duplicated)

(6) Prime/Dupe MODs pairs

- (a) To improve Data-integrity of DASD, TPF system often makes the Data-duplications on different two DASD subsystems.
- (b) The following figure shows one example of these pairs. Prime MOD(module)s and Dupe MODs are always located on each side of subsystem(spread to all subsystems).




Fig. 3.13.4-3 Prime/Dupe MODs pairs

(7) Data Copy procedures

The Copy procedures are taken for the following purposes:

- (a) To make a pair (To copy data from Prime MOD to Dupe);
- (b) To recover the failed data (To copy the remaining data to the re-init MOD).

There are two ways to make a pair.

- (a) AFC (All File Copy),
 - (b) AMOD copy.
- (a): In this copy process, the destination-drive of the copy keeps “Offline” status, and just after the completion of this copy, the source-drive becomes “Offline” and the destination-drive changes to “Online”. From the view-point of TPF software, there is only one MOD, independent of copy process.
 - (b): In this copy process, both source-drive and destination-drive stay “Online”. TPF software can distinguish both drives, even in the copy process.

3.13.5 Notices for HRC-option setting

<SVP operation>

(1) RCU Option

We strongly recommend you to select “No” in the “PPRC support by host” column of the “RCU Option” window.

We strongly recommend you to select the “Not Report” in the “Service SIM of Remote Copy” column of the “RCU Option” window.

(2) Add Pair

We strongly recommend you to select the “Copy to R-VOL” in the “CFW Data” column of the “Add Pair” window.

(3) Suspend Pair

We strongly recommend you to select the “Disable” in the “SSB(F/M=FB)” column of the “Suspend Pair” window.

<Host (TPF-OS) consideration>

In MVS-OS world, DKC with HRC expects (requires) Customers to extend I/O Patrol Timer to prevent many MIHs from reporting.


Also in TPF-OS world, same consideration is required, so please discuss with your Customer to find the opportunity to extend “Stalled Module Queue” timer over 5 seconds.

3.14 HXRC

3.14.1 Outline of HXRC

HXRC is not supported in Ver.1.

HXRC (Hitachi eXtended Remote Copy) function provides for data replication at distance in order to recover for disaster.



- System Data Mover defines a XRC pair session. (①)
- When a write command is issued to the primary volume from application site, the primary DKC makes Sidefile data (replication) on cache memory. (②)
- System Data Mover reads Sidefile data non-synchronously at distance, and writes it to the secondary volume. (③, ④)

Fig.3.14.1-1 An outline of HXRC

3.14.2 HXRC Support Requirements

3.14.2.1 OS

(1) OS level

- (a) MVS/ESA 4.3.0 or upper.
- (b) DFSMS/MVS 1.1.0 or upper.

<Restriction of SMS 1.4 environment>

The Maximum number of HXRC pair per CU image is up to 128 under SMS 1.4 environment. CCA (Channel Connection Address) can be specified to 128 logical devices per CU image such as '00' - '7F'. CCA addresses '80' - 'FF' for HXRC may be rejected by System Data Mover.

(2) Conditions of HXRC in use

- (a) The following conditions must be satisfied by OS before starting the HXRC function,

- CACHE ON
- NVS ON

and must be the CACHE ON status on DKC.

When CACHE OFF/NVS OFF commands are issued by OS or Cache malfunctions (includes 'Ref code=FFEE: Area temporary blocking) occur, the HXRC function is stopped.

(b) I/O Patrol Value

- (I) Without CHL Extender
 - Current patrol time(more than 30sec).
- (II) With CHL Extender
 - More than 70sec.

(c) Session ID

- Up to 64Session ID's can be utilized per CU for Concurrent Copy and HXRC.
- Up to 64Session ID's can be utilized per CU for HXRC.
- Up to 16Session ID's can be utilized per VOL for Concurrent Copy and HXRC.
- Only 1Session ID can be utilized per VOL for HXRC.

(d) HRC/HODM

- HXRC cannot be used for the same volumes using the HODM pair volumes.
- HODM cannot be used for the same volumes using the HXRC pair volumes.
- HXRC must be the volumes after deleting the HODM pair volumes.

3.14.2.2 Hardware

(1) HXRC Support Hardware Specification.

Table 3.14.2.2-1 HXRC Support Hardware Specification

CU Type	3990-6/6E, 2105/2107(*1)
DEV Type	3390-3/3R/9/L, 3380-3(*2), 3390-M(*3)
DKC model	Primary: RAID450 Recommended Secondary: RAID450/RAID400/RAID300/RAID200HA/ DKC80/DKC90
RAID level	RAID5/RAID1
Channel	ESCON

*1: Do not intermix of DKC emulation type '2105/2107' and 3990-6/6E in the same DKC.

If you change DKC emulation type '2105/2107', the following operation.

- Delete All CC/XRC pairs
- Change DKC emulation type 2105/2107 of All CHE PK
- RESUME CC/XRC pairs

*2: 3380-3 is supported only when the DKC emulation type is '2105/2107'.

*3: If you want to use DKU emulation type 3390-M, PTF 'UA18053: Support XRC volume SIZE up to 65520 CYL' adaptation is necessary.

(2) CACHE SIZE

Cache capacity should be doubled from the current cache size.

(The amount of Sidefile data may occupy up to 60% of total cache capacity.)

3.14.2.3 Micro-program

- (1) HXRC supports from the 1st version of Main Frame Micro-program.
- (2) CNT extender version 4.9 or upper level code is recommended.

(3) Device Blocking Function and Load Balancing Control

DKC does not block Write I/Os for the logical device which is specified the DONOTBLOCK option not to affect performance impact for application programs.

<Requirements>

The following conditions need to activate the DONOT BLOCK option.

For Operating system

- The operating system should support the DONOT BLOCK option.

For RAID system

- Set XRC option DONOT BLOCK = Enable for the DONOT BLOCK option.

DKC performs current load balancing control, if DONOT BLOCK = Disable (default).

- Should be holding DONOT BLOCK = Disable (default), if the operating system does not support the function.

3.14.2.4 HXRC recommendations

- (1) Recovery site CPU is the most ideal location for Data Mover.
- (2) Data Mover's path should be utilized only to read Sidefile.
- (3) Subsystem configurations
 - Cache capacity : Should be doubled from current cache size.
 - Confirmation for Number of channel paths for system data mover (SDM)
 - Confirmation for Work loads for the subsystem
- (4) Utility device for primary volume
 - Should be prepared for each XRC session.
 - A low activity device should be selected as a Utility Device
 - Utility Device should be specified at the 1st time before establishing pair volumes.
- (5) System Data Mover (SDM)
 - Confirmation for PTF levels
 - No record found problem. : APAR # OW30183, OW33680
 - Necessary tuning for SDM data set : Capacity, Geometry of the data set
- (6) DB2
 - Broken VSAM index file problem : APAR # II08859
- (7) Others
 - CPU MIPS : Enough for HXRC environment
 - LINE CAPACITY : Enough for HXRC environment with channel extender.
- (8) HXRC PP option
 - If the ANTA5107E(RC=9014 REAS=604 OR REAS=608) console message is displayed during the XADDPAIR operation for HXRC pairs the operation might be unsuccessful.
 - In this case, you may check the HXRC PP Option Installed. If HXRC PP Option not install Please install HXRC PP option.
 - Hitachi – Extended Remote Copy PP Option effects only 2105/2107 dkctype.
 - You can use Hitachi – Extended Remote Copy for 3990 dkctype without this PP option.
- (9) HXRC with FICON

Table 3.14.2.4-1 HXRC and FICON configuration

		Record set transfer path (System Data Mover - DKC)	
		FICON	ESCON
Application site	ESCON	Supported	Supported
(System and DKC)	FICON	Supported	Not recommended (*1)

*1: If the path of Application site is FICON, System Data Mover (SDM) path should be also FICON in order to balance the performance of Application path and SDM path.

- (10) TOD setting or updating of Synchronization information
In the case that the amount of Sidefiles reach to the threshold, XRC pair may be suspended due to operator's manual operation of "TOD setting" or updating of "Synchronization information".

3.14.3 Online Maintenance while Concurrent Copy(CC)/HXRC in use

(1) Availability of Installation and DE-installation.

Component	Maintenance Type	During initial copy		Established		Suspend	
		Primary	Secondary	Primary	Secondary	Primary	Secondary
HDD canister	Installation	*	x	*	x	*	x
	De-installation	*	x	*	x	*	x
Cache PCB	Installation	*	x	*	x	*	x
	De-installation	*	x	*	x	*	x
CHA	Installation	x	x	x	x	x	x
	De-installation	x	x	x	x	x	x
DKA	Installation	x	x	x	x	x	x
	De-installation	x	x	x	x	x	x

x: Maintenance is available.

*: Maintenance is available but it should take place when workload is low. The following are recommendations.

When a maintenance operation is needed while CC/HXRC is being used, I/O's for CC/HXRC pair volumes or CC/HXRC itself should be stopped before the maintenance operation.

If the maintenance operation must be done while CC/HXRC is being used, you must confirm that the usage of Sidefile monitor is less than 20% of total Cache capacity before you start the maintenance operation. Only when the usage of Sidefile monitor is less than 20% of total Cache capacity, you can proceed the maintenance operation.

Refer to "Monitoring" in the SVP SECTION about Sidefile monitor.

- Select the [Monitor] button in the SVP main panel to start the monitoring feature.
- From the menu in the 'Monitor' panel, select [Monitor]-[Open...].
- Select 'Cache' from [Object] and 'Cache Sidefile' from [Item] in the "Select Monitor Item" panel. After that, select [=>] button and then the [OK] button.

(2) Availability of the System tuning .

When the following System tuning operation is needed while CC/HXRC is being used, CC/HXRC should be stopped before the System tuning operation.

- It is impossible to change the DKC No, SSID, or DKC Emulation type by System tuning operation while CC/HXRC is being used.
- When the DRV emulation type of CC/HXRC pair volumes are 3390-3 or 3390-3R, it is impossible to change the emulation type between 3390-3 and 3390-3R by CHANGE EMULATION operation while CC/HXRC is being used.

(3) Availability of the Replacement.

Component	Maintenance Type	During initial copy		Established		Suspend	
		Primary	Secondary	Primary	Secondary	Primary	Secondary
Logical Device	Blockade	**	**	**	**	**	**
	Recovery	**	**	**	**	**	**
	Format	**	**	**	**	**	**
	Verify	x	x	x	x	x	x
HDD canister	Replace	x	x	x	x	x	x
Cache PCB	Replace	*	x	*	x	*	x
CHA	Replace	x	x	x	x	x	x
DKA	Replace	x	x	x	x	x	x
LTM PCB	Replace	x	x	x	x	x	x

x: Maintenance is available

*: Maintenance is available but it should take place when workload is low. The following are recommendations.

When a maintenance operation is needed while CC/HXRC is being used, I/O's for CC/HXRC pair volumes or CC/HXRC itself should be stopped before the maintenance operation.

If the maintenance operation must be done while CC/HXRC is being used, you must confirm that the usage of Sidefile monitor is less than 20% of total Cache capacity before you start maintenance operation. Only when the usage of Sidefile monitor is less than 20% of total Cache capacity, you can proceed the maintenance operation.

Refer to "Monitoring" in the SVP SECTION about Sidefile monitor.

- Select the [Monitor] button in the SVP main panel to start the monitoring feature.
- From the menu in the 'Monitor' panel, select [Monitor]-[Open...].
- Select 'Cache' from [Object] and 'Cache Sidefile' from [Item] in the "Select Monitor Item" panel. After that, select [=>] button and then the [OK] button.

** : When a maintenance operation is needed while CC/HXRC is being used, CC/HXRC should be stopped before the maintenance operation.

3.15 HRC Asynchronous

3.15.1 Components

Asynchronous mode is one of the update copy modes of the HRC volume pairs. The HRC asynchronous subsystem consists of the same components as HRC synchronous with the following exceptions:

- A set of the HRC volume pairs named consistency group is introduced.
- Only 1-to-1 and n-to-1 ($n \leq 4$) configuration is supported
- For n-to-1 configuration, the XRC time-stamping capability is required.
- A communicating facility to transfer error information from the primary system to the secondary system is not required.

Note: In this document, the term n-to-m means that n-MCUs and m-RCUs are connected to each other to establish the HRC volume pairs. N and m is number of control units of physical unit bases, not of control unit image bases.




Fig. 3.15.1-1 HRC Asynchronous Subsystem Components

(1) MCU and RCU

- Both MCU and RCU must be RAID400 or upper model.
- Maximum configuration is 4-to-1. For n-to-1 ($n > 1$) configuration, XRC time-stamping capability is required.
- To use XRC time-stamping capability, control unit emulation must be 2105/2107 or 3990-6 basic or enhanced mode.
- Maximum number of control unit image pairings (established by Add RCU or ESTPATH) is 128.

(2) Consistency Group

- HRC asynchronous ensures update-sequence-consistency across several volume pairs. It also provides some group-based operations. A set of volume pairs treated by such group-based functions is called a consistency group.
- HRC asynchronous supports 128 consistency groups at maximum. Every HRC asynchronous volume pair belongs to one consistency group.

(3) PPRC Support

- Although HRC asynchronous is not fully compatible to PPRC, it can be controlled and monitored with PPRC host facilities, PPRC TSO commands, ICKDSF PPRC commands and some console messages. For this purpose, MVS/DFP 3.2.0 or higher level and ICKDSF release 16 or upper are available.
- Only fundamental facilities are available. Neither P/DAS SWAP nor CGROUP is supported.
- If the primary system (and the secondary system) consists of several CPU complexes, SYSPLEX timer must be installed for the common time reference.

(4) XRC Time-Stamping Capability

- In case of N-to-1 configuration, the XRC time-stamping capability requires to be installed in the primary host system. MVS/DFP 3.2.0 or higher level is required.
- In order to get benefit of time-stamping capability during copy-back process (pair establishment from the secondary to the primary subsystem), the XRC time-stamping capability recommends to be installed in the secondary system.
- If the primary system (and the secondary system) consists of several CPU complexes, SYSPLEX timer must be installed for common time reference.

3.15.2 Consistency Group

(1) HRC Asynchronous Volume Pairs and Consistency Group

- Every HRC asynchronous volume pair belongs to one consistency group.
- When establishing the HRC asynchronous volume pair, an operator specifies the consistency group number that the volume pair will belong to with a new parameter of Add Pair and ESTPATH.
- The consistency group must be registered prior to pair establishment.

(2) Functions of Consistency Group Basis

(a) Ensuring Update Sequence Consistency

- The updated records are copied to the corresponding R-VOLs in the same order as the M-VOLs have been updated by the primary host systems.
- The update sequence consistency is ensured within a consistency group. The updated records of the different consistency groups may be copied in the different order from the original.

(b) Suspending Volumes Pairs (Error Level)

- When one R-VOLs is not updated correctly due to the failure, all the HRC asynchronous volume pairs will be suspended with keeping update sequence consistency.
- When establishing the HRC asynchronous volume pair, an operator specifies whether other volume pairs will be suspended together or not against the failure of the volume pair. A new parameter called Error Level is defined for this purpose.

Error Level = Group	When this volume pair is suspended due to the failure, all volume pairs in the same consistency group will be suspended together.
Error Level = Volume	Even if this volume pair is suspended due to the failure, other volume pairs in the same consistency group will not be suspended, as long as the failure prevent.

(c) Providing the Consistency Time

The latest time stamp value of the update that has been successfully copied to the R-VOL is called a **consistency time**. The consistency time is a group basis indication. It means that all the updates performed before or at the consistency time have been successfully copied to the R-VOLs in the consistency group.

- The consistency time can be displayed with the following operations issued to the R-VOL of duplex or suspended state. If 'LOCAL' is specified for timer type, Consistency time is not displayed.
 - Web Console Pair Status panel
 - PPRC CQUERY command (only at suspended state.)
- At the R-VOLs in duplex state, the consistency time is a ticking value. Any R-VOL displays the consistency time in that instance. It can be used for feeling how long the R-VOLs are behind the M-VOLs.
- Whenever the volume pair is suspended, the consistency time of the R-VOL is frozen.
 - ① If the update sequence consistency between the R-VOL and other R-VOLs in the consistency group is ensured, the R-VOL indicates the latest consistency time of the consistency group.
 - ② Otherwise, the R-VOL indicates the latest time stamp value of the update that has been successfully copied to the R-VOL. It may be older than other R-VOLs because the consistency time of the consistency group is still ticking.
- If the R-VOL is in suspended state, the supplementary status that indicates whether the consistency time is of the consistency group (case (a) above) or the R-VOL (case (b) above) is also displayed.

(d) Consistency Group Basis Operations

In order to make the disaster/failure recovery procedure simple, the following consistency group basis operations are provided.

- Operations at the RCU

- ① Deleting all suspended volume pairs except for inconsistent volume pairs
- ② Deleting all volume pairs regardless of the consistency among them
- ③ Suspending all volume pairs

- Operations at the MCU

- ① Suspending all volume pairs
- ② Deleting all volume pairs behind this unit (except for M-VOLs behind other MCU)
- ③ Resuming all suspended volume pairs behind this unit (except for M-VOLs behind other MCU)

(3) Configuration of the Consistency Group

(a) Disposition of the Volume Pairs

- All R-VOLs that belong to the same consistency group must be located behind one RCU.
- The M-VOLs that belong to in the same consistency group can be located behind up to 4 different MCUs.
- Up to 128 consistency groups can be established within one pair of MCU and RCU. The RCU sup-ports up to 128 consistency groups.
- The R-VOLs of the different consistency groups can be located behind the different RCU.
- Up to 8,192 volume pairs can belong to one consistency group.

(b) Primary Host Systems and Consistency Group

1) Primary host systems and timer type

- Every update I/O to the M-VOL of the same consistency group must be time-stamped by using common timer facility. Table 3.15.2-1 shows the relationship between the primary host system and available timer resource.
- When registering the consistency group, an operator must specify which timer resource should be used for the consistency group based functions.

Table 3.15.2-1 Primary host systems and timer resource

Primary Host System	Timer Resource	Notes on Configuration
MVS with the XRC time-stamping capability	System timer (CPU TOD clock)	• N-to-1 ($N \leq 4$) configuration is possible. • SYSPLEX timer must be installed if the primary system consists of the several CPU complexes.
Other main frame host systems	Local timer (MCU internal clock)	• Only 1-to-1 configuration allowed. • The consistency time is not displayed.
Open host systems		

2) Restrictions and notes on the primary host systems

- The primary host systems can not access the volume pairs of the same consistency group unless they have the common timer reference.
- The M-VOLs updated by the same primary host system can belong to the same or different consistency group if the M-VOLs have no requirement on update sequence consistency (i.e. they are updated independently of each other.) However it is recommended for them to belong to the different consistency groups because, for example, they might be suspended together resulting from the failure
- For the same reason, the independent M-VOLs accessed by the independent primary host systems is recommended to belong to the different consistency groups, even if they can use common timer reference.

3.15.3 HRC Asynchronous Theory of Operations

(1) Update Copy

The updates from the primary host systems and additional control information are queued in the cache storage of the MCU, and sent to the RCU independent of host I/O processes. The RCU stores the data and control information into the provisional spaces allocated in the cache storage. According to the time-stamp and the sequence information, the RCU promotes the updates in the provisional spaces the formal data of the R-VOLs in the same order as they have been performed at the MCU.

(a) Receiving Time-stamp Information

In case of the Timer Type of System specified, the MCU receives the time-stamp information as follows:

- When directed to establish the HRC asynchronous volume pair, the MCU reports the state-change-interrupt (SCI) to all the attached host systems. The host system issues a series of sense group commands to recognize what status of the device has changed. The MCU generates the response as if the device became a member of an XRC session. This response activates the XRC time-stamping capability if installed in the host systems.
- Once activated, MVS IOS routine attaches the time-stamp information (contents of time-of-day clock) to each I/O operation to read and write the device. The time-stamp information indicates when the corresponding update has been issued at the primary host system. It is transferred to the MCU at the beginning of each I/O operation.

(b) Creating Recordset

- When accepting the updates from the primary host systems, the MCU creates a set of information called a **recordset**. A recordset includes:
 - ① updated record
 - ② time-stamp information received
 - ③ sequence number
 - ④ record locations (device, cylinder, track and record number) and record length
- The **sequence number** is the number of recordsets the MCU has created for the consistency group. That is, all recordsets in each MCU and each consistency group are independently numbered.
- The recordset information other than the updated records is stored and queued into the exclusive spaces allocated in the cache storage.
- The updated records are stored as the host-dirty data and do not occupy the exclusive space until the following events happen before the recordset is sent to the RCU
 - The same record is updated again, or
 - The host-dirty status is removed by the de-staging process.

When the above mentioned event happens, the MCU moves the updated records into an exclusive space, called **Sidefile**, in the cache storage.

(c) Sending Recordset to the RCU

- The MCU sends the recordset in a similar manner as HRC synchronous. That is, the MCU and RCP port act as the host processor channel and issue I/O operation, called Remote I/O(RIO), to the RCU.
- The RIO transfers the time-stamp information, the sequence number, the record locations and length, and the updated records in the FBA format (not in the CKD format) by one channel command, like HRC synchronous. However the parameter length and detailed specification of this channel command is different from HRC synchronous. Therefore the micro code of the store-and-forward type channel extender (i.e. Channelink and UltraNet of CNT corp.) should be upgraded to support the command.
- Unlike the recordset offloading of the XRC Data Mover, the MCU sends the recordset with directly specifying the device address of the R-VOL. The RIO independently activates each R-VOL. Furthermore, the MCU may send several recordsets by one RIO even if their sequence numbers are not contiguous to each other. Therefore the recordsets are usually sent in different order from their arrivals to the MCU.

(d) Storing Recordset into Sidefile Space

- The RCU stores the received recordsets into the spaces exclusively allocated in the cache storage. The exclusive space is called a Sidefile. The updated records in the Sidefile are not treated as the formal data. That is, the host I/O processes and the de-staging processes do not access the records in the Sidefile at this time.
- The records in the Sidefile will be promoted the formal data later. A term settle means to promote the records in the Sidefile.
- The RCU also allocates exclusive spaces in the cache storage so that the Sidefiles form a queue.
 - The RCU makes a queue per MCU and per consistency group.
 - This queue is not of the FIFO fashion. Each entry of the queue is previously assigned to each sequence number. The arrived recordsets are queued into the corresponding entries with indexed by their sequence number. The entries for the recordsets that do not arrive yet are left empty. As a result, the RCU lines up the recordsets in the order of their sequence number.

(e) Selecting Recordset and Deciding Consistency Time

- The RCU selects the recordset to be settled with the following algorithm.
 - ① Checks if there are the valid entries at the top of all the queues in the consistency group. If one of them is empty, the RCU waits for the entry.
 - ② When all the top entries are filled with the valid Sidefiles, the RCU selects one entry that has the smallest time-stamp. It can be settled.
 - ③ Repeats step 1 and 2.
- Fig. 3.15.3-1 shows an example. All the top entries are filled with the recordsets of S11/T3, S21/T2, S31/T1 and S41/T5. The RCU selects the recordset of S31/T1 to be settled because the T1 is the smallest time-stamp. Then the top entry S31/T1 is removed from the MCU3's queue. The S32 becomes the top but it is empty. The next recordset S11/T2 will be selected when recordset S32 arrives and its time-stamp is smaller than T2




Fig. 3.15.3-1 Example of Selecting Recordset at the RCU

(f) Settling Recordset

- The recordset selected by the algorithm described in section(e) will be marked as host-dirty and treated as the formal data after that time. The time-stamp value of the recordset is promoted to the consistency time.
- The RCU settles the updated records in the recordset as follows:
 - If the corresponding track is not in cache storage (track-miss), the cache directory of the Sidefile is changed to be the formal data. No data is moved.
 - If the corresponding track is in cache (track-hit), the updated records in the recordset are copied to the existing cached-track and the cache space for the Sidefile is released.

3.15.4 MCU-RCU Communications to Maintain Asynchronous Copies

(1) Dummy Recordset

The RCU needs to receive the recordset continuously from all the MCUs even if the MCU does not have to create the new recordset.

- The MCU creates and sends a dummy recordset when it has received no update I/O in a second. The dummy recordset contains only the sequence number and the time-stamp information. Contents of the time-stamp information is generated by incrementing the largest time-stamp that the RCU has (the MCU reads it from the RCU before creating the dummy recordset) by one.
- The RCU receives the dummy recordset and puts it into the queue. It can help other recordsets being selected.

Another purpose of the dummy recordset is have the RCU be aware of the disaster. If the RCU can not receive any recordset in the predetermined duration (this time can be specified by Maximum Copy Delay Time parameter), the RCU regards such situation as the disaster and suspends all the HRC asynchronous volume pairs.

Due to these reasons, the MCU and RCU always need to communicate with each other once the HRC asynchronous volume pair established and not all the pairs have been suspended.

(2) Change-Status Recordset

When the volume pairs are suspended or deleted due to operations or failure, the update sequence consistency must be ensured. In order to meet this requirement, the negotiation on changing volume pair status is made by means of recordset. The recordset used for this purpose also has only the sequence number and the time-stamp information.

(3) PS OFF Notification Recordset

At the power-off sequence, the MCU creates and sends the recordset to notify of the power-off event. The recordset used for this purpose also has only the sequence number and the time-stamp information.

3.15.5 Failure Detected by the RCU

(1) Pair Suspend by the RCU

Table 3.15.5-1 shows the failures detected by the RCU and the volume pairs to be suspended due to the failure.

Table 3.15.5-1 Volume Pairs to be Suspended by the RCU-detected Failure

Failures	Volume pairs to be suspended
The RCU could not settle the pending recordset or could not communicate with the MCU before the maximum copy delay time expired.	All pairs in the consistency group
The RCU could not receive the recordset successfully due to the hardware failure.	All pairs in the consistency group, or only the affected pair (depending on the failure)
The RCU detected the logical error while selecting the recordset to be settled.	All pairs in the consistency group
The RCU could not settle the recordset due to the hardware failure, the track condition, or the logical error.	All pairs in the consistency group, or only the affected pair (depending on the failure)

(2) Pair Suspended and Re-synchronization

For the HRC asynchronous volume pairs, both the MCU and RCU maintain the bit map for pair re-synchronization. When/after the volume pair(s) suspended, the cylinders that contain the following records are marked in the bit map as modified(to-be-copied later):

- The recordsets that have been created by the MCU but not sent the RCU yet. After marking the cylinders as modified, the recordsets are discarded.
- The recordset that have reached at the RCU but not settled yet. After marking the cylinders as modified, the recordsets are discarded.
- The records updated by the primary system after the volume pair(s) suspended

At the beginning of the pair re-synchronization, the contents of the RCU's bit map are sent to the MCU and merged into the MCU's bit map. The MCU performs the initial copy for the pair re-synchronization according to its bit map. That is, the cylinders that contain the lost recordset are re-synchronized at this time.

3.15.6 Inflow Control for Sidefiles

- As described in section 3.15.4, both the MCU and RCU create the Sidefiles for storing the recordsets. Since the Sidefile is an exclusive space in the cache storage, both the MCU and RCU perform the inflow control to prevent the subsystem overload.
- Both the MCU and RCU use the threshold value specified with the Web Station PC/SVP panel.

(1) Inflow Control by MCU

- When the amount of Sidefiles reaches at the threshold, the MCU responds to the update I/Os from the primary system with the state-change-pending (SCP) or channel-command-retry request.
- If no recordset has been sent to the RCU after the specified time duration, the MCU will suspend all the volume pairs and reset the SCP condition in order to avoid the system being hung up.

Note: In the case that the amount of Sidefiles reach to the threshold, Async pair may be suspended due to operator's manual operation of "TOD setting" or updating of "Synchronization information".

(2) Inflow Control by RCU

- When the amount of Sidefiles reached at the threshold, the RCU responds to the command that transfers the recordset from the MCU with the channel-command-retry request. Only the recordset of the sequence number necessary to continue settling the pending recordsets is accepted.
- If the recordset has not been settled after the specified time duration, the RCU will suspend all the volume pairs and reset the channel-command-retry condition in order to avoid the MCU being hung up.

3.15.7 HRC Asynchronous Control Operations

This chapter describes the Web Console operations for HRC asynchronous.

3.15.7.1 DKC Options

(1) Async Option - Modifying HRC Options on Physical Unit Basis

- Async Option panel provides the function to modify asynchronous options.
- These options are effective to entire physical control unit (i.e. all M-VOLs and R-VOLs behind the control unit.)
- These options can be modified when no asynchronous volume pair is established.
- These options may be modified before/after performing Port, Add RCU, and Add Group operations.

Table 3.15.7.1-1 Async Options on Physical Control Unit Basis

Option Name	Description
Pending Update Data Rate	• It specifies the amount of cache storage in percent that allows to be used for storing recordset (Sidefile). • When the amount of cache storage for the recordset reaches the specified threshold, the MCU and RCU invokes its own inflow control as follows: – The MCU responds to the update I/Os from the primary system with the state-change-pending (SCP) or channel-command-retry request. – The RCU responds to the command that transfers recordset from the MCU with the channel-command- retry request. However the specific recordset that will help the RCU settle the pending recordset is still accepted. • Any percent between 30% and 70% can be specified with a unit of 10%. The default is 50%.
Offloading Timer	• It specifies how long the MCU can continue the inflow control described above. • The MCU stops the inflow control and suspends all asynchronous volume pairs after the specified time expires unless no recordset has been offloaded to the RCU. • Every minute between 1 to 20 and “None” can be specified. If “None” is specified, the MCU will immediately become suspend, when the Sidefile threshold is exceeded.

Notice: If Pending Update Data Rate is modified when asynchronous volume pair is established, the host I/O timeout may occur.

3.15.7.2 Registering/Monitoring/Deleting the Consistency Group

(1) Add Group - Registering Consistency Groups

- The consistency group can be registered with Web Console attached to the MCU.
- The consistency group must be registered prior to the volume pair establishment.
- The consistency group has its own attributes and parameters, consistency group number, timer type, and others. They are specified when registered.
- The consistency group is registered in the RCU too. However it is not necessary to be specified. When the volume pair is established, the MCU directs the RCU to register the consistency group.

(a) Consistency Group Number

- The consistency group number is described with one digit of hexadecimal character.
- The volume pair control operations require the consistency group number as the parameter. The pair status displayed by Pair Status operation of Web Console and CQUERY command also includes the consistency group number.

(b) Timer Type

- The timer type must be specified out of System, Local and None when the consistency group is registered.

Table 3.15.7.2-1 Timer Type Attributes

Timer Type	Meaning
System	The system timer (CPU TOD clock) provided by the XRC time-stamping capability is used for controlling this consistency group.
Local	The local timer (<i>internal</i> TOD clock of this MCU) is used to for controlling this consistency group.
None	The system timer (CPU TOD clock) provided by the XRC time-stamping capability is used for controlling this consistency group. The R-VOLs in this consistency group can be located behind the different RCUs. <i>However the update sequence consistency across the RCUs is not ensured.</i> This timer type should be selected only when volume pairs are established from the original secondary to the original primary volumes (<i>copy back</i>).

- Table 3.15.7.2-2 shows the related configuration and timer type to be specified.

Table 3.15.7.2-2 Timer Types to Be Specified

Configuration			Timer type to be specified	
Host system	XRC time-stamping capability	MCU-to-RCU	For <i>P-to-S</i> copy (original direction)	For <i>S-to-P</i> copy (copy back)
Main frame	Installed	N-to-1 (n > 1)	System	None
		1-to-1	System	System
	Not installed	1-to-1	Local	Local
Open systems	(unavailable)	1-to-1	Local	Local

(c) Timeout Parameters

- The following parameters can be specified to modify the expiration time for the timeout event of the consistency group basis.

① Maximum copy delay time

Table 3.15.7.2-3 Maximum copy delay time to Be Specified

Name in the RMC panel	Time Out: Write Pending [min]
Available range	3 min. to 15 min. or "None"(no time out event occurs)
Default	5 min.
Description	It specifies the maximum delay allowed for asynchronous copy. Based on this parameter, the RCU will suspend all R-VOLs if following time out event occurs: • The RCU has received the updated data but it can not be settled in the specified time. • The RCU has had no communication from the MCU until the specified time expires

Note. The RCU stores the updated data received into a provisional space of cache storage, and will make it available for use later. The term settle means to making it available for use.

When the User of SE use ASYNC HRC function in N-to-1 configuration, they must reset the maximum copy delay allowed with the notice of as follows.

If not take, it may cause suspension of HRC pairs.

- (1) Execute ASYNC HRC with the maximum delay allowed = 'NONE'.
- (2) You can recognize the current copy delay with the difference the Time Stamp of HOST I/O and the Consistency Time of "Group Status" in Web Console.
- (3) Execute "Suspend Pair" for all pairs on the CT Group by Web Console, and reset the maximum delay allowed with over the current copy delay. If it is longer than maximum time (15 min.), reduce the HOST I/O rate, or you should leave it "NONE".
- (4) Restart ASYNC HRC with "Resume Pair" for all pairs on the CT Group by Web Console.

② Maximum RCU-ready-wait time

Table 3.15.7.2-4 Maximum RCU-ready-wait time to Be Specified

Name in the RMC panel	Time Out: RCU Ready [min]
Available range	1 min. to 10 min or "None"(no wait)
Default	5 min.
Description	During the power-on-reset procedure, the MCU intends to communicate with the RCU and will suspend all the volume pairs if it can not communicate until the specified time expires.

(2) Delete Group - Deleting Consistency Group Registration

- If no volume pair belongs to the consistency group, the group can be deleted.
- This operation is available only at the MCU. The registration to the RCU is automatically deleted, when the last volume pair in this group is deleted.
- In N-to-1 (N > 1) configuration, deleting the consistency group does not affect the consistency group that has been registered in another MCU.

(3) Group Option - Modifying HRC Options on Group Basis

- This operation allows MCU to delete the consistency group currently registered.
- The Group Option can be operated only when no volume pair belongs to this group.

(4) Group Status - Displaying Consistency Group Status

- The options and the working status can be displayed on this panel.

Table 3.15.7.2-5 Consistency Group Status

Item	Contents	Displayed by:	
		MCU	RCU
Consistency group number	Consistency group number in one digit of hexadecimal character.	Yes	Yes
RCU serial number/SSID	Serial number and SSID of the RCU that belongs to this group.	Yes	No
Volume list	List of volumes that belong to this group and are behind this control unit.	Yes *1	Yes
Consistency time	Current consistency time of this group.	Yes *2	Yes
Timer type	Specified timer type of this group.	Yes	Yes
SEQCHK *3	At least one volume pair of this group has the SEQCHK status.	Yes *2	Yes
Maximum copy delay time	Specified maximum copy delay time.	Yes	Yes
Maximum RCU-ready-wait time	Specified maximum RCU-ready-wait time.	Yes	No

*1: In N-to-1 (N > 1) configuration, it does not include the M-VOLs behind other MCU.

*2: The consistency time and SEQCHK status are decided by the RCU. The MCU displays these items after reading them from the RCU. If the MCU can not communicate due to communication failure, the latest contents are not displayed. Therefore these items should not be used for disaster recovery.

*3: SEQCHK is one of the R-VOL statuses.

3.15.7.3 Pair Status

- Suspending and Deleting states are newly defined to indicate the status in transition.
- Suspended by MCU powered-off is added as the caused of suspension.
- To indicate whether the update sequence consistency is kept or not, the subsidiary pair status (Group or Volume) and SEQCHK indicator are newly defined.

(1) Status in Transition - Suspending and Deleting

When suspending or deleting the HRC asynchronous volume pairs, the MCU and RCU intend to process all pending recordset before changing pair status. It takes longer time for the asynchronous volume pairs to change to suspended or simplex state than the synchronous volume pairs. Therefore adding to the conventional pair statuses (simplex, pending, duplex, suspended), two new statuses are introduced.

(a) Definitions and conditions of transition

Suspending This volume pair is in transition from duplex or pending to suspended state. When cause of suspension (failure or operation) is detected, all affected volume pairs change to suspending state. After completing suspension, they will automatically change to suspended state.

Deleting This volume pair is in transition from duplex, pending or suspended to simplex state. When accepting delete pair operation, all affected volume pairs change to deleting state. After completing delete pair operation, they will automatically change to simplex state.

(b) Indication of pair status

- Suspending and deleting statuses can be indicated only on Web Station PC (or SVP) main control and pair status panels.
- For main frame host systems, these states are not indicated. Status in transition is treated as follows:
 - In case of operations (suspend pair or delete pair), status is not changed until transition completes. After completing status transition, affected volume pairs are changed to suspended or simplex state.
 - In case of failure, affected volume pairs are changed to suspended state when cause of suspension is detected.
 - In any cases, the MCU or RCU report the state change interrupt (SCI) after completing status transition. IEA491E or IEA494I console messages appear on the system console at this time

(2) Suspended by MCU Powered-Off

When the MCU is being powered off, the MCU suspends all volume pairs behind it. The suspended volume pairs will automatically return to their original state (duplex or pending) when the MCU is powered-on again. During this suspension, the suspended R-VOLs indicate by MCU powered-off as the cause of suspension.

(a) Conditions of transition

- During power-off sequence, the MCU notifies the RCU of the power-off event. The RCU changes all related R-VOLs to suspended states and sets by MCU powered-off as the cause of suspension.
 - Only R-VOLs which are in duplex or pending state are affected. The cause of suspension of the R-VOL that is already suspended is not changed.
 - Only R-VOLs of which corresponding M-VOLs are behind the MCU are affected. The volume pairs between other MCU are not affected.
- During power-on sequence, the MCU notifies the RCU of the power-on event. The RCU changes all R-VOLs in suspended by MCU powered-off state to original (duplex or pending) state.

(b) Indication of pair status

- Only the RCU can indicate this cause of suspension.
- This cause of suspension is displayed as follows:
 - Web Console Main Control panel: OFF in the Sub column
 - Web Console Pair Status panel: Suspended (MCU PS OFF)
 - PPRC CQUERY command: SUSPENDED(5) with the specific indicator in the serial number field

(3) Consistency Status - Group or Volume

(a) Definition of the status

Group

Update sequence consistency between this R-VOL and other R-VOLs in this consistency group is ensured. This R-VOL can be used for disaster recovery at the secondary system after deleting the HRC volume pair.

This status is indicated when:

- All volume pairs in this consistency group have been suspended due to the failure that affected the consistency group (not the specific volume pair.)
- The volume pair having Error Level = Group has been suspended due to the failure.
- This volume pair has been suspended by suspend pair operation with Group parameter.

Volume

Probably this volume pair has been suspended alone. Update sequence consistency between this R-VOL and other R-VOLs in this consistency group is not ensured. This R-VOL can not be used for disaster recovery at the secondary system. This status is indicated when:

- This volume pair has Error Level = Volume and has been suspended due to the failure that did not affect entire group.
- This volume pair has been suspended by suspend pair operation with Volume parameter.

(b) Indication of the status

- Only the RCU can display this cause of suspension.
- This status has the meaning described above only when the R-VOL is in suspended state.
- This status is displayed as follows:
 - Web Station PC (or SVP) Main Control panel: GRP or VOL in the Sub column
 - Web Station PC (or SVP) Pair Status panel: GRP or VOL as “Suspended by:” item
 - PPRC CQUERY command: SUSPENDED(x) with the specific indicator in the serial number field

(4) Alert for Non-Time-stamped Updates - SEQCHK

SEQCHK indicates that the volume pair has accepted some updates without time-stamp attached while the consistency group is specified to use system timer.

(a) Set/reset conditions

When the RCU is settling the updated data, the RCU turns this indicator on if the updated data is not time-stamped, otherwise turns this indicator off.

(b) Indication of this indicator

- Only the RCU can display this indicator.
- This status can be indicated unless the volume is in simplex status.
- This status is displayed as follows:
 - Web Console Main Control panel: SEQ (on) or blank (off) in the SEQ column
 - Web Console Pair Status panel: SEQCHK (on) or blank (off)
 - PPRC CQUERY command: The specific indicator in the serial number field

3.15.7.4 Controlling Volume Pairs

- Basic operations to control/monitor the HRC asynchronous volume pairs are same as the HRC synchronous pairs. That is, at first the volume pairs to be operated must be selected on the main control panel, and then the operation and its parameters/options should be specified.
- For the HRC asynchronous volume pairs, some group basis operations are supported. Several volume pairs within the consistency group to which the selected volume belongs can be deleted, suspended, resumed by the group basis operations. For these operations, any volume pair within the consistency group may be selected.

(1) *Add Pair* - Establishing HRC Volume Pairs

- This operation has the MCU establish the HRC volume pair(s). Selected volumes will be the M-VOLs and the R-VOLs should be specified by the parameters.
- For asynchronous volume pairs, the MCU-RCU logical connection must be established and the consistency group must be registered prior to this operation.

(a) Basic parameters and pair options

The basic parameters and the pair options of Add Pair are described in and respectively. “Changed” column of these tables indicates the difference between HRC asynchronous and conventional HRC.

Table 3.15.7.4-1 Add Pair - Basic Parameters

Item	Parameters	Changed	Description
Pair configuration	- R-VOL - RCU	No	Specifies the serial number and SSID of the RCU and the logical device number of the R-VOL
Initial copy control	- Initial Copy - Priority	No	Specifies whether the initial copy for this volume pair is necessary (“Entire Volume”) or not (“None”) and its priority
Update copy mode	- Copy Mode	Yes	Specifies the update copy mode for this volume pair out of Synchronous and Asynchronous. For the asynchronous volume pair, specifies the consistency group number to which this volume pair will belong.

Table 3.15.7.4-2 Add Pair/Resume Pair - Option Parameters

Option Name	Changed	For:	Description
Initial copy pace	No	Any	Specifies the initial copy pace.
Fence Level	No	Synch.	Specifies the fence level out of “Data”, “Status” and “Never” for synchronous volume pairs. This option is valid for the HRC synchronous volume pairs.
CFW data	No	Any	Specifies whether the CFW updates should be copied to the R-VOL or not.
DFW to R-VOL	No	HODM	Specifies the volume pair should be suspended when DFW is disabled/failed at the RCU. This option is valid only for HODM volume pairs.
Error Level	Yes (new)	Asynch.	Specifies whether all the volume pairs in the same consistency group should be suspended together or not when this volume pair is suspended. Group All the volume pairs group should be suspended together. Volume Only this volume pair may be suspended. This option is valid only for the HRC asynchronous volume pairs and the default is Group.
Pair Resume	Yes (new)	Asynch.	Specifies whether all the suspended volume pairs, which belong to the same consistency group and whose M-VOLs are behind this MCU, should be resumed together. Group All the volume pairs should be resumed together. Volume Only this volume pair may be resumed. This option is valid only for Resume Pair operation of the HRC asynchronous volume pairs and the default is Group (currently Volume).

(b) End conditions

The Resume Pair operation with the Group option gets the normal end condition before the initial copy for each volume pair begins. If some unusual conditions (pinned tracks, correction-access status of the parity group, etc.) prevents the volume pair from being re-established, the volume pair would be still in suspended status. Therefore the result of this operation should be confirmed by referring to the actual pair status or the console messages that indicate pair status change.

(2) *Delete Pair* - Deleting HRC Volume Pairs

- In order to make the disaster recovery operations at the secondary subsystem simple, a new option that specifies the volume pairs to be deleted is supported (Delete option).
- The pair status of the asynchronous volume pairs will be Deleting when the operation completes, and then Simplex after the internal process completes.

(a) Volume pairs to be deleted

Table 3.15.7.4-3 Delete Pair - Delete Option

Option Name	Operable at:		Description
	MCU	RCU	
Consistent Volumes	No	Yes	Specifies that the volume pairs that meet the following conditions should be deleted. — Belongs to the same consistency group as the selected volume pair and; — Is in Suspended status and; — Has the consistency status of Group. Other volume pairs than described above are not deleted. This is the default when operated at the RCU.
Group	Yes	Yes	Specifies that the all volume pairs (but volume pairs behind other MCU when operated at the MCU) in the same consistency group should be deleted.
Volume	Yes	Yes	Specifies that only selected volume pairs should be deleted. This is the default when operated at the MCU.

- This option is valid only when the asynchronous volume pair is selected.
- “Consistent Volumes” and “Group” can be specified when only one volume pair is selected.

Note: Async pair using with [Suspend Range: Volume]

If there are “Duplex” status pairs and “not-Duplex” status pairs in one CT group, unexpected pair suspend may occur for your pair operation (Suspend/Delete/Resume) under high I/O stress condition * (Eg. about 30% ratio of Sidefile).

*: You can know high I/O stress condition with about 30% ratio of Sidefile (if you can see Sidefile ratio) or high frequency of host I/O (if you can not see Sidefile ratio).

(b) Force option for asynchronous volume pairs

The Force option specifies that the volume pair should be deleted even if the MCU and RCU can not communicate with each other. For the asynchronous volume pairs, this option is still effective but can be specified with the Delete option of Group.

(c)

1) Group option for R-VOL pairs

If you delete the R-VOL pairs without the Delete option of Group, the pair status of the pairs will be Simplex after settling all pending recordsets when the operation is accepted by the MCU/RCU. On the other hand, if you delete the R-VOL pairs with the Delete option of Group, the pair status of the pairs will be Simplex when the state of no recordset for the pair continue during some time after settling all pending recordsets when the operation is accepted by the RCU. Therefore, if you delete the R-VOL pairs with the Delete option of Group, you need the following procedures.

- Stop I/Os to the volume pairs, and
- Delete the pairs.

After the operation is accepted by the RCU, if the state of no recordset for the pair does not continue for some time, the pair is forcibly suspended. This suspend is due to R-VOL failure (not user-requested). This status is temporary, and will change to Simplex finally. In this case it is possible that all pending recordsets are not settled.

2) End condition/Delete operation

The Delete Pair operation for the asynchronous volume pairs gets the normal end condition when the operation is accepted by MCU/RCU. The pair status of the volume pairs to be deleted are Deleting at this time. And then the pair status will be Simplex when the condition above mentioned in 1) is concluded.

Therefore the result of this operation should be confirmed by referring to the actual pair status or the console messages that indicate pair status change.

(d) Consistency time/status after deleted

Once the volume pair deleted, the consistency time/status of the volume pair will be reset. Therefore the consistency time/status should be memorized prior to this operation.

(3) *Suspend Pair* - Suspending HRC Volume Pairs

- A new option that specifies the volume pairs to be suspended is supported (Suspend option).
- The pair status of asynchronous volume pairs will be Suspending when the operation completes, and then Suspended after the internal process completes.

(a) Parameters and options

Table 3.15.7.4-4

Suspend Pair - Parameters and Options

Item Name	Operable at:		Description
	MCU	RCU	
SSB (F/M=FB)	Yes	Yes	Specifies that IEA494E console message should be generated.
Suspend Kind	-	-	Only “R-VOL” can be selected for the asynchronous volume pairs.
Suspend			Specifies the volume pairs to be suspended.
Group	Yes	Yes	Specifies that all volume pairs in the same consistency group as the selected volume pair should be suspended together.
Volume	Yes	Yes	Specifies that only the selected volume pair should be suspended.
Pending Update			Specifies how the pending recordset should be treated.
Drain	Yes	Yes	Specifies that the volume pair(s) should be suspended after all pending recordset are settled. Refer to (3) in this section for notes on this option.
Purge	Yes	Yes	Specified that that volume pair(s) should be suspended even if pending recordset remain. The pending recordset may be purged. If the MCU/RCU discard the pending recordset, the MCU/RCU marks the cylinders that contain discarded recordset as modified in its shared memory. The marked cylinders will be copied when the Resume Pair operated.

Note: Async pair using with [Suspend Range: Volume]

If there are “Duplex” status pairs and “not-Duplex” status pairs in one CT group, unexpected pair suspend may occur for your pair operation (Suspend/Delete/Resume) under high I/O stress condition* (Eg. about 30% ratio of Sidefile).

*: You can know high I/O stress condition with about 30% ratio of Sidefile (if you can see Sidefile ratio) or high frequency of host I/O (if you can not see Sidefile ratio).

(b) Update sequence consistency

Regardless of the Pending Update parameter, the update sequence consistency across the volume pairs to be suspended is ensured. However, if some volume pairs have been suspended with the Volume option and other related volume pairs are updated after that, the update sequence consistency is not ensured.

(c) Notes on Drain option

- If the Purge option specified, the volume pairs are suspend when the RCU accepts this operation. On the other hand, if the Drain option is specified, the volume pairs will be suspended when the RCU completes the steps described bellow.
 - ① Accepts this operation and;
 - ② Has finished settling all the pending recordsets and;
 - ③ Completes the negotiation with all MCUs (reports ready-for-suspension to all the MCUs and receives their acknowledgements) without further recordset generated.
- Therefore, the procedures to get the R-VOLs whose contents are frozen at specific point in time relative to the application, are as follows
 - ① Quiesce the application (Quiesce all update activity to the volume pairs) and;
 - ② Perform the Suspend Pair operation with Drain option and;
 - ③ Confirm that all volume pairs complete to be suspended and;
 - ④ Restart the application.
- After the operation is accepted by the RCU, if the state of no recordset for the pair does not continue for some time, the pair is forcibly suspended. This suspend is due to R-VOL failure (not user-requested). In this case it is possible that all pending recordsets are not settled.

(d) End conditions

The Suspend Pair operation for the asynchronous volume pairs gets the normal end condition when the operation is accepted by the MCU/RCU. The pair status of the volume pairs to be suspended are Suspending at this time. And then the pair status will be Suspended after the internal process (negotiation between the MCU and RCU) completes. Therefore the result of this operation should be confirmed by referring to the actual pair status or the console messages that indicate pair status change.

(4) *Resume Pair* - Resuming HRC Volume Pairs

The Pair Resume option is newly supported. It specifies whether all suspended volume pairs, which belong to the same consistency group and whose M-VOLs are behind this MCU, should be resumed together or not. Refer to Table 3.15.7.4-2 for this option.

(5) *Pair Option* - Modifying HRC Options on Volume Pair Basis

Refer to Table 3.15.7.4-2 for the options on volume pair basis. The Error Level option can be changed regardless of the pair status (even if the volume pair is in Suspended state), but become effective at the next time of suspension.

(a) Modified items on Pair Status panel

Table 3.15.7.4-4 Pair Status - Modified Items on Pair Status Panel

Items	Indicated at:		Description
	MCU	RCU	
M-VOL and R-VOL	Yes	Yes	Indicates control unit image number and device number in form of “c:dd”.
Pair Synchronized	Yes	Yes	Indicates the amount (%) of cylinders that are marked as modified in the bitmap for pair re-synchronization. Refer to (2) below for more detailed information.
Pair Status			Adding to the contents for the synchronous HRC volume pairs, the following are added for the asynchronous HRC volume pairs.
Suspending	Yes	Yes	Indicates that the volume pair is in transition state to Suspended.
Deleting	Yes	Yes	Indicates that the volume pair is in transition state to Simplex.
Suspended[MCU PS OFF]	No	Yes	Indicates that the volume pair has been suspended due to the power-off event of the MCU.

(b) The Pair Synchronized indicator

- For the asynchronous volume pairs, both MCU and RCU maintain the bitmap for pair re-synchronization. The indicated percentage is calculated as follows.
 - The M-VOL of Pending state indicates the remaining cylinders to be copied for pair re-synchronization.
 - The R-VOL of Suspended state indicates the cylinders that contain the recordsets lost at the RCU (reached at the RCU but can not be settled before suspension).
 - The M-VOL of Suspended state indicates the cylinders that contain;
 - ① The tracks that have not copied yet by the initial copy and;
 - ② The records updated by the primary system after suspension and;
 - ③ The recordsets lost at the MCU (created in the MCU but can not be sent to the RCU before suspension) and;
 - ④ The recordsets lost at the RCU.

The last item (d) is included as long as the MCU can get the information from the RCU. Otherwise, only items (a)-(c) are included. In this case, the item (d) will be included at the beginning of the pair resynchronization.

- This percentage is always calculated based on the total cylinders of the M-VOL (even if the R-VOL is larger than M-VOL).

(c) Added items on Pair Status panel

Table 3.15.7.4-5 Pair Status - Added Items on Pair Status Panel

Items	Indicated at:		Description
	MCU	RCU	
C/T Group	Yes	Yes	Indicates the consistency group number to which this volume pair belongs.
C/T Type	Yes	Yes	Indicates the timer type that has been specified to control the consistency group. The displayed content is out of System, Local and None.
C/T	Yes (Note)	Yes	Indicates the consistency time of this volume pair in form of “mm/dd/yyyy hh:mm:ss.uuuuuu”. (uuuuu: micro seconds)
SEQCHK	Yes (Note)	Yes	Indicates SEQCHK if the volume pair is in SEQCHK status. Otherwise it blanks.
Suspended by	No	Yes	Indicates the consistency status of the R-VOL when this volume pair is in Suspended status. Group The consistency between this R-VOL and other R-VOLs in the consistency group is ensured. Volume The consistency between this R-VOL and other R-VOLs in the consistency group is not ensured. The contents of this R-VOL may be behind other R-VOLs.

Note: Only the RCU maintains the latest result and can display it. The result at the MCU may be behind the actual result because the MCU get the result from the RCU before displaying. Therefore the result at the MCU can be used for monitoring the normal activities, but can not for the disaster recovery at the RCU.

(d) Added items

Table 3.15.7.4-6 Pair Status - Added Items on Pair Status Panel

Items	Indicated at:		Description
	MCU	RCU	
Grp (Lv)	Yes	Yes	Grp Indicates the consistency group number to which this volume pair belongs. (Lv) Indicates the Error Level of this volume pair, “Grp” for group or “Vol” for volume.
Sub	No	Yes	Indicates the consistency status or the others when the volume pair is in “suspended” state. GRP This R-VOL has the consistency status of Group. VOL This R-VOL has the consistency status of Volume. OFF This volume pair has been suspended due to the power-off event of the MCU.
Seq	No	Yes	Indicates SEQCHK if the volume pair is in SEQCHK status. Otherwise it blanks.

3.15.7.5 Monitoring subsystem Statistics

(1) *Usage* - Displaying Remote I/O Statistics Information

Table 3.15.7.5-1 Usage - RIO Statistics in Async Copy Category

Item name	Unit	Description
Async IO count	—	Indicates the total number of RIO (Remote I/O) activities completed in a specified interval.
Total number of recordset	—	Indicates the total number of recordset sent to the RCU in a specified interval.
RCU command retries	—	Indicates the total number of RIO command retries requested by the RCU in a specified interval.
MCU command retries	—	Indicates the total number of channel command retries performed by the MCU in order to avoid cache slot conflict between the host I/O and RIO processes.
Average transfer rate	KB/sec.	Indicates the average data transfer rate of the recordset sent to the RCU in a specified interval.
Average RIO response	msec.	Indicates the average RIO response time in a specified interval.

(2) *Information/Monitor* - Displaying Subsystem Resource Usage

- The item named “Async Write Pending Data” is newly supported. It is displayed on the Information/Monitor panel of the SVP and recorded into monitor.dat file.
- “Async Write Pending Data” indicates the total amount of cache space in percent that store the pending recordset (Sidefile).
- The conventional item “Sidefile” indicates the total usage of the Sidefile for HXRC, Concurrent Copy and HRC asynchronous.

3.15.7.6 On-line Micro-program exchange procedure

(1) Version up procedure

- Usually, exchange the Micro-program of MCU first.
- If the 'Copy Back function' is running, exchange the Micro-program of RCU first.

(2) Version down procedure

- If in the 'N-to-1' configuration, set back to '1-to-1' configuration.
- Execute 'Delete- Group' for all CT groups on the subsystem.

3.15.8 Management/Recovery Procedures

3.15.8.1 Managing HRC Asynchronous Subsystems

(1) Checking on SEQCHK Status

- SEQCHK status would be indicated when the asynchronous volume pair in the consistency group with the timer type of System specified accepts the non-time-stamped updates from the primary system.
- SEQCHK status does not affect copy activities of HRC asynchronous and will be removed when the next time-stamped update is successfully copied to the R-VOL. However if the disaster happens before the next time-stamped update, the update sequence consistency between the R-VOL and other R-VOLs in the consistency group is not ensured. Therefore to make the disaster recovery more certain, the source of SEQCHK status should be detected and removed.
- Source of SEQCHK status to be suspected is as follows:
 - An application may issue the update I/Os with bypassing MVS standard I/O procedure.
 - An XRC time-stamping capability may not active at the primary system.
 - An Operating system of the primary system may not support the time-stamping capability.

(2) Checking on the Consistency Time

The consistency time is indicated as a part of pair status of the HRC asynchronous volume pairs. While the primary system continues to update the M-VOLs, the difference between the current time and the consistency time indicates how long the R-VOLs are behind the M-VOLs. The updates to the M-VOLs during this duration may be lost when the disaster happens

- If the disaster recovery design can not accept this delay, performance improvement by adding the remote copy resources (paths, cache amount, etc.) and/or reducing unnecessary I/O workload should be considered to get shorter delay.
- If this delay is close to the time specified by the Maximum Copy Delay Time parameter, the timeout failure may occur and the affected volume pair may be suspended due to the I/O workload fluctuations. The performance improvement described above and/or increasing this parameter setting should be considered.

(3) Planned Outage of HRC Asynchronous Components

The MCU requires communication with the RCU even if it receives no update I/Os from the primary system. Suspending all the duplex pairs is necessary prior to the planned outage of the RCU.

(a) Planned outage of the MCU

- No special procedure is required for the asynchronous volume pairs. The MCU automatically suspends all volume pairs in Duplex/Pending state during the power-off sequence, and will remove the Suspended state at the power-on-reset sequence.

Note: Perform planned outage operation of the MCU after the Sidefile data becomes a zero. The Sidefile ratio of the subsystem is shown on the Monitor dialog box.

- Note that the volume pairs whose M-VOLs are behind the MCU will not be suspended. In N-to-1 configuration, the volume pairs behind other MCU(s) continue to be copied. If the full-consistent volume set requires to be kept during the planned outage, the procedures are as follows:

- ① Quiesce the applications.
- ② Perform Suspend Pair operation with Group and Purge (or Drain) option at the RCU.
- ③ Perform planned outage operation of the MCU(s).
- ④ After getting all MCUs ready to resume, perform Resume Pair operation at all MCUs.

(b) Planned outage of the RCU

- Suspending all volume pairs is required prior to the planned outage according to the followings:

- ① Suspend all volume pairs;
 - 1) Perform Suspended Pair operation at each MCU and;
 - 2) Confirm volume status with Pair Status operation.
- ② Perform planned outage operation of the RCU.
- ③ Get the RCU ready to resume.
- ④ Remove Suspended status from all volume pairs;
 - 1) Perform Resume Pair operation at each MCU and;
 - 2) Confirm volume status with Pair Status operation.

Note: If step 1 above is not performed, the MCU detects the communication failure and suspends all affected volume pairs, then generates SIM(s) and console message(s) that indicates the failure.

(c) Planned outage of the components on the remote copy connection

The same restriction should be considered and the same procedures should be performed as the RCU for the components (channel extender, ESCON director, etc.) on the remote copy connection.

- Suspending all volume pairs is required prior to the planned outage according to the followings:

- ① Suspend all volume pairs;
 - 1) Perform Suspended Pair operation at each MCU and;
 - 2) Confirm volume pair status with Pair Status operation.
- ② Perform planned outage operation of the component.
- ③ Get the component ready to resume.
- ④ Remove Suspended status from all volume pairs;
 - 1) Perform Resume Pair operation at each MCU and;
 - 2) Confirm volume status with Pair Status operation.

Note: If step 1 above is not performed, the MCU detects the communication failure and suspends all affected volume pairs, then generates SIM(s) and console message(s) that indicates the failure.

(d) Planned outage of both MCU and RCU

- The MCU(s) must become not-ready first, and the RCU must be back before the MCU(s).
 - ① Perform planned outage operation of the MCU(s).
 - ② Perform planned outage of operation of the RCU.
 - ③ Get the RCU ready to resume.
 - ④ Get the MCU(s) ready to resume.
- Note that the MCU make all volume pairs suspended if it can not communicate with the RCU during the power-on-reset sequence. Therefore step 4 above should be started after completing step 3. If it is difficult to control due to some installation requirements (Eg. Power-Control-Interface setting), consider using the Maximum RCU-ready-wait Time parameter.

(4) ICKDSF on the Asynchronous Volume Pairs

- ICKDSF activities involve write I/Os with device support authorization or diagnostic authorization instead of normal authorization. Since the MCU does not duplicate write I/Os with device support or diagnostic authorization, the HRC volume pairs must be suspended before running ICKDSF. The procedures to do are as follows:

- ① Perform Suspend Pair operation with Volume option at the MCU to suspend volume pair(s) for which ICKDSF will be performed.
- ② Perform ICKDSF.
- ③ Perform Resume Pair operation at the MCU.

(5) Micro-program Exchange

No special procedure is required for the asynchronous volume pairs for the on-line micro-program exchange and the MCU's off-line micro-program exchange.

In the case of HRC Asynchronous components, Suspending all volume pairs is required prior to the RCU's off-line micro-program exchange.

- ① Suspend all volume pairs;
 - 1) Perform Suspended Pair operation at each MCU and;
 - 2) Confirm volume status with Pair Status operation.
- ② Perform micro-program exchange operation of the RCU.
- ③ Remove Suspended status from all volume pairs;
 - 1) Perform Resume Pair operation at each MCU and;
 - 2) Confirm volume status with Pair Status operation.

Note: If step 1 above is not performed, the MCU detects the communication failure and suspends all affected volume pairs, then generates SIM(s) and console message(s) that indicates the failure.

In step 3, if both micro-program exchange operation is performed, perform Resume Pair operation after waiting 10 minutes after RCU's micro-program exchange is completed.

3.15.8.2 Recovering from Pair Suspended

(1) Recovering from Pair Suspended - Suspended by the MCU

The cause of suspension and the recovery procedures are basically the same as HRC synchronous with exceptions that:

- ① The cache storage/shared memory of the one side (not both sides) may cause the HRC asynchronous volume pairs to be suspended.
- ② The MCU requires communicating with the RCU in power-on-reset sequence and suspends all the HRC asynchronous volume pairs if it can not do so.
- ③ The MCU suspends all the HRC asynchronous volume pairs if no recordset can be sent in the specified time period.

(2) Recovering from Pair Suspended - Suspended by the RCU

Adding to directed by the MCU, the RCU of HRC synchronous may detect the cause of suspension and suspended the failed volume pair(s) on its own initiative.

3.15.8.3 Disaster Recovery - Switching to the Secondary Subsystem

(1) Switching Procedures to the Secondary Subsystem

Basic procedures to switch to the secondary subsystem are as follows:

- ① Check pair status and memorize the consistency time of the R-VOLs.
- ② Make the R-VOLs simplex by performing Delete Pair operation.
- ③ Confirm that the R-VOLs have successfully been changed to simplex.
- ④ Perform IPL of the secondary system.
- ⑤ Perform the application restart procedure at the secondary system

(2) Suspending Volume Pairs due to the Disaster

If the RCU receives no communication with the MCU after the time specified by Maximum Copy Delay Time expires, the RCU regards this timeout event as a disaster and suspends all the HRC asynchronous volume pairs.

- Pair status of the R-VOLs normally become Suspended - Group.
- The latest time-stamp of the recordset that has been successfully settled is frozen and indicated as the consistency time of the R-VOLs.

(3) Checking Volume Pair Status

- The R-VOLs of Suspended - Volume should not be used for the disaster recovery. Therefore Error Level of Group should be selected if the volume is essential to the disaster recovery.
- Delete Pair operation with Group option can make only the R-VOLs of Suspended - Group status. Therefore an operator seems to be able to skip this checking. However the volumes of Simplex status can not be distinguished from others after the R-VOLs in Suspended - Group status become Simplex by Delete Pair operation. Therefore at least the volumes in Simplex must be checked at this time.

Table 3.15.8.3-1 Checking R-VOL Status When Switching to the Secondary System

Pair status		Usable for recovery?	Description
Suspended	Group	Yes *1	Since the update sequence consistency across these R-VOLs is ensured at point in time indicated by the consistency time, these R-VOLs can be used for the disaster recovery at the secondary system. Note that some updates performed after the consistency time at the primary system may be lost.
Suspended	Volume	No	Since contents of this R-VOL may be behind other R-VOLs in the consistency group, this R-VOL should not be used for the disaster recovery if this volume requires the update sequence consistency with other volumes. Suspected reason for this status are as follows: — This R-VOL has been suspended due to the failure or Suspend Pair operation prior to the disaster. — This R-VOL was in pending status when the disaster happened. — The Error Level of Volume has been specified for this R-VOL and this R-VOL was suspended by the first symptom of the disaster.
Duplex		No	This status does not usually take place in the disaster recovery procedure. This R-VOLs should not be used for the disaster recovery. <i>It should be especially noted that the volumes in simplex status can not be distinguished from others after deleted by Delete Pair operation.</i>
Pending		No	
Simplex		No	

*1: If SEQCHK status should be decided by the RCU using the XRC time-stamping capability, this R-VOL shall not be used for disaster recovery.

(4) Memorizing the Consistency Time

The consistency time of the R-VOLs should be memorized. It could help the disaster recovery retrieve the lost updates.

- All the R-VOLs of Suspended - Group status indicate the same consistency time. Therefore any R-VOL of this status can represent the consistency time of the consistency group.
- Once Delete Pair operation performed, the R-VOL never indicates the consistency time again. Therefore the consistency time should be memorized before Delete Pair operation.

(5) Deleting Volume Pairs

By Delete Pair operation, the R-VOLs that are used for the disaster recovery should be changed to Simplex status.

- Group option of Delete Pair operation can make all the R-VOLs of Suspended - Group status Simplex. It would be helpful in reducing number of operations.
- Group option of Delete Pair operation does not change the pair status of the R-VOLs if they are other than Suspended - Group. It can prevent the inconsistent volumes from being used for the disaster recovery.

(6) Recovering the Lost Updates

HRC asynchronous provides no factory standard procedure to retrieve the lost updates.

- In order to detect and recreate lost updates, it is necessary for customers to check other current information, for example, data base journal log file that had been active at the primary system when disaster occurred. Note that the journal log file entries of most DBMS may be related to time-of-day clock information and the source of the consistency time is time-of-day clock of the primary system (when the timer type of System specified.).
- However such a detection/retrieval would take long time to do. The customers' disaster recovery scenario recommends to be designed to enable such a detection/retrieval after the application has been started at the secondary system.
- Maximum Copy Delay Time parameter of Add Group operation can control the maximum time duration during which the updates may be lost.

3.15.8.4 Disaster Recovery - Switching Back to the Primary Subsystem

This section describes the procedures to switch back from the secondary system to the primary system. The basic concept and procedures are the same as HRC synchronous. After volume pairs establishment in the opposite direction, planned switching-back and volume pair establishment in the original direction are performed.

That is, the original R-VOLs are working as Simplex and the application is running on the secondary system by using the original R-VOLs. The original HRC asynchronous configuration/status should remain in the original MCU. In this section, the original MCU/M-VOLs and RCU/R-VOLs are called the primary subsystem and secondary subsystem respectively.

(1) Switching Back Procedures to the Primary Subsystem

- (a) Make the primary subsystem and the communication facilities on the remote copy connection operable. Note that all the M-VOLs in the primary subsystem may be suspended since the original R-VOLs are now in Simplex status.
- (b) Remove entire HRC asynchronous configuration remained in the primary subsystem:
 - ① Make all the volume pairs Simplex by Delete Pair operation.
 - ② Remove registration of the consistency group by Delete Group operation.
 - ③ Remove MCU-RCU logical paths by Delete RCU operation.
 - ④ Change the serial interface port to LCP mode.
- (c) Establish HRC asynchronous in the opposite direction:
 - ① Change the operating mode of the communication facilities in the opposite direction.
 - ② Change the serial interface port of the secondary subsystem to RCP mode.
 - ③ Register the consistency group at the secondary subsystem by Add Group operation.
 - ④ Establish MCU-RCU path in the opposite direction at the secondary subsystem by Add RCU operation. The primary subsystem is now defined as the RCU.
 - ⑤ Establish the HRC asynchronous volume pairs in the opposite direction. The volumes in primary and secondary subsystem are now defined as the R-VOLs and M-VOLs respectively.
Note that Initial Copy of Entire Volume option must be specified.
- (d) Quiesce the application at the secondary system.
- (e) Confirm all the HRC asynchronous volume pairs are in Duplex status.

(f) Remove entire HRC asynchronous configuration at the secondary subsystem:

- ① Make all the volume pairs Simplex by Delete Pair operation.
- ② Remove registration of the consistency group by Delete Group operation.
- ③ Remove MCU-RCU logical paths by Delete RCU operation.
- ④ Change the serial interface port to LCP mode.

(g) Establish HRC asynchronous in the original direction

- ① Change the operating mode of the communication facilities in opposite direction.
 - ② Change the serial interface port of the primary subsystem to RCP mode.
 - ③ Register the consistency group at the primary subsystem by Add Group operation.
 - ④ Establish MCU-RCU path in opposite direction at the secondary subsystem by Add RCU operation. The secondary subsystem is now defined as the RCU.
 - ⑤ Establish the HRC asynchronous volume pairs in the original direction. The volumes in primary and secondary subsystem are now defined as the M-VOLs and R-VOLs respectively.
- Note that Initial Copy of No Copy option may be specified.

(h) Make the primary subsystem online from the primary system and restart the application at the primary system.

Note: Since the CNT channel extenders have the operating mode, they must be re-configured to change copy direction. The boxes (or nodes) to which the current MCU and RCU are connected must be set as channel-mode and device-mode respectively.

(2) Setting the Consistency Group for Switching Back

- For N-to-1 ($N \geq 2$) configuration, timer type of “None” must be specified.
 - None enables the consistency group across up to 4 RCUs to be established.
 - However the update sequence consistency across the RCUs is not ensured.
- In 1-to-1 configuration, the update sequence consistency across all R-VOLs is ensured regardless of the timer type. However System is recommended if the secondary system can use the XRC time-stamping capability in order for the consistency time to be indicated.

Table 3.15.8.4-1 Timer Type for Switching Back

Original Configuration		Timer Time for Switching Back	
Number of MCU-RCU	Timer Type	XRC time-stamping capability at the secondary system?	
		Yes	No
N-to-1 ($N \geq 2$)	System	None	None
1-to-1	System	System	None
	Local	Local	Local

3.16 HIHSM(Hitachi internal Hierarchical Storage Management)

3.16.1 HIHSM Overview

This document describes the function of HIHSM (Hitachi internal Hierarchical Storage Management) that is one of program products.

RAID system can be constructed by several types of physical drives and three types of RAID levels (RAID1, RAID5 and RAID6).

This combination of the type of physical drive and the type of RAID level provides a system that cost and performance are optimized to user environment. However, it is difficult to get information about actual operation of physical drives in the RAID system unlike other disk subsystems.

- (1) HIHSM provides solutions of the problem and supports decision of users to determine system construction as described below.
 - (a) Load balancing of system resources
Unbalance of utilization of system resources makes performance worse. HIHSM supports decision of optimized allocation of logical volumes to physical drives.
 - (b) Migration of logical volumes optimized to access patterns to physical drives
For instance, RAID5/RAID6 are suitable to sequential access, and RAID1 of high performance drive is suitable to random access that is required small response time. HIHSM shows types of access pattern to physical drives clearly, and supports migration of logical volumes to suit the access pattern.
- (2) HIHSM consists of following subfunction to achieve above purposes. Users can refer to utilization of system resources monitored by monitor function, decide reallocation plan by using estimate function, and reallocate the logical volumes by volume moving (migration) function.
 - (a) Monitor function(*1) monitors and shows utilization of system resources.
 - (b) Estimate function estimates utilization of parity groups after migration of logical volumes.
 - (c) Volume moving (migration) function moves logical volumes to specified parity groups.
 - (d) Automating Migration function makes a migration plan from information that users preset, and moves the logical volumes by the migration plan automatically.

*1: The license of Performance Monitor is required.

For details, refer to Performance Management User's Guide of Program products.

3.16.2 Hardware requirements

More SM may be needed if DKC does not have enough SM. Please refer at INSTALLATION SECTION.

3.16.3 Software requirements

The Program Product, Performance Monitor is required.

3.16.4 Monitor function

The Performance Monitor Program is required to perform this function.

(1) How to start and stop.

Monitoring starts in DKC by direction from the Monitoring Option window.

Monitor function monitors ratios of utilization of system resources described below.

(a) CHP utilization and DKP utilization ratio

(b) Starnet utilization ratio

(c) DRR utilization ratio

(d) Parity group utilization: Disk utilization of parity groups.

The used time of physical drives in a parity groups.

(e) Parity group utilization ratio of each logical volume: The used time of physical drives of synchronous and asynchronous access on each logical volume, averaged by the number of physical drives in the parity group.

Parity group utilization means the sum of utilization of each logical volume in the parity group.

Directions for an stop from the Monitoring Option window to finish the collection of the Monitoring information.

(2) How to collect

Monitoring information is acquired automatically until directions for an acquisition end are issued after directions for an acquisition start of the Monitoring information are issued on the Monitoring Option Window.

When acquired Monitoring information is collected in SVP, it learns to refer to it with a [Physical] tab.

It is done automatically collecting opportunity Monitoring information every 15 minutes.

The collected Monitoring information up to 3 months is stored in the hard disk of SVP.

When a monitor function is continued for more than 3 months, it is erased in order from the old information, and new Monitoring information is cumulated.

(3) How to view

The Monitoring information which accumulated in SVP can be referred to with a [Physical] tab.

The period of the Monitoring information which accumulated can be specified by the operation of the part Term.(It is optional in the unit for 15 minutes.)

Information about the Monitoring information of the period specified in the part Term is indicated in the Table part and the Graph part.

The utilization of the resource (average and maximum value) chosen in the Tree part is indicated in the Table part.

The utilization of the element chosen by the Table part is graphed and indicated by the Graph part.

3.16.5 Estimate function

The estimate function estimates changes of parity group utilization and parity group utilization of each logical volume after migration of the logical volume to specified parity group. The estimate function estimates the changes from the monitored information.

3.16.6 Volume moving (migration) function

Volume moving (migration) function moves data in logical volume (source volume) to physical location of another logical volume (destination volume). Users specify the volumes in HIHSM utility window.




Fig. 3.16.6-1 Before moving




Fig. 3.16.6-2 After moving

(1) Volume moving function Overview

(a) Source volumes:

The following volumes cannot be used as source volumes:

- Volumes which are set as command devices (devices reserved for use by the host)
- Volumes which are used by XRC
- Volumes which are used by CC (Concurrent Copy)
- Volumes which have FlashAccess (also called DCR) data stored in cache
- Volumes which are in an abnormal or inaccessible condition (e.g., pinned track, fenced)
- Journal Volumes which are used by Universal Replicator

If the status of volumes that form HRC pairs is suspended, the volumes can be used as source volumes. If you delete an HRC pair from an MCU, the status of the M-VOL and the R-VOL changes to simplex so that the volumes can be used as source volumes. If you delete an HRC pair from an RCU, the status of the M-VOL changes to suspended and the status of the R-VOL changes to simplex so that the volumes can be used as source volumes.

If migration source volume is HRC volume, migration target volume must be an internal volume. Migration target volume can not be an external volume.

If the status of volumes that forms HORC pairs is PSUS or PSUE, the volumes can be used as source volumes. If not, the volumes cannot be used as source volumes. If you delete an HORC pair from an MCU, the status of the P-VOL and the S-VOL changes to SMPL so that the volumes can be used as source volumes. If you delete an HORC pair from an RCU, the status of the P-VOL changes to PSUS and the status of the S-VOL changes to SMPL, so that the volumes can be used as source volumes.

If migration source volume is HORC volume, migration target must be an internal volume. Migration target volume can not be an external volume.

— Volumes which configure the Universal Replicator for Mainframe pairs

If the status of the volumes is Pending duplex or Duplex, they cannot be source volumes. Also, if the volumes which configure the Universal Replicator for Mainframe pairs are used as source volumes, external volumes cannot be specified as target volumes.

— Volumes which configure the Universal Replicator pairs

If the status of the volumes is COPY or PAIR, they cannot be source volumes. Also, if the volumes which configure the Universal Replicator pairs are used as source volumes, external volumes cannot be specified as target volumes.

For volumes that form an HMRCF pair or an HOMRCF pair, it depends on the status or configuration of the pair whether the volumes can be used as source volumes, as explained below:

- If the status of the pair is not SP-Pend/V-Split, the volumes can be used as source volumes. If the status of the pair is SP-Pend/V-Split, the volumes cannot be used as source volumes.
- The table below explains whether volumes that do not form a cascade pair can be used as source volumes:

Table 3.16.6-1 Whether volumes that do not form a cascade pair can be used as source volumes

If the pair is configured as follows	Can P-VOLs be used as source volumes?	Can S-VOLs be used as source volumes?
If the ratio of P-VOLs to S-VOLs is 1:1	Yes	Yes
If the ratio of P-VOLs to S-VOLs is 1:2	Yes	Yes
If the ratio of P-VOLs to S-VOLs is 1:3	No	Yes

- The table below explains whether volumes that form a cascade pair can be used as source volumes:

Table 3.16.6-2 Whether volumes that form a cascade pair can be used as source volumes

If the pair is configured as follows	Can P-VOLs be used as source volumes?	Can S-VOLs be used as source volumes?
If the pair is an L1 pair and the ratio of P-VOLs to S-VOLs is 1:1	Yes	Yes
If the pair is an L1 pair and the ratio of P-VOLs to S-VOLs is 1:2	Yes	Yes
If the pair is an L1 pair and the ratio of P-VOLs to S-VOLs is 1:3	No	Yes
If the pair is an L2 pair and the ratio of P-VOLs to S-VOLs is 1:1	Yes	No
If the pair is an L2 pair and the ratio of P-VOLs to S-VOLs is 1:2	No	No

Caution: If any of the following operations is performed on a source volume, the volume migration process stops:

- XRC operation
- CC operation
- HRC/HORC operation that changes the volume status to something other than suspended
- ShadowImage (HMRCF/HOMRCF) operation that changes the volume status to SP-Pend/V-Split.
- Universal Replicator for Mainframe operation or Universal Replicator operation which seems to make the volumes the status of COPY

(b) Target volumes:

Target volumes must be reserved prior to migration. The HIHSM remote console software allows you to reserve volumes as HIHSM target volumes.

Hosts cannot access HIHSM-reserved volumes.

The following volumes cannot be reserved as target volumes:

- Logical Unit Size Expansion (LUSE) volumes
- Volumes which are set as command devices (devices reserved for use by the host)
- Volumes which are assigned to Hitachi ShadowImage (HMRCF/HOMRCF) or Hitachi Remote Copy (HRC/HORC) pairs
- Volumes which are used by XRC
- Volumes which are used by CC (Concurrent Copy)
- Volumes which are reserved for ShadowImage operations
- Volumes which have FlashAccess (also called DCR) data stored in cache
- Volumes which are used by ENAS system
- Volumes which are in an abnormal or inaccessible condition (e.g., pinned track, fenced)
- Onlined volume
- Volumes which are used by Universal Replicator (Primary Data Volume, Secondary Data Volume and Journal Volume)
- Volumes which are set as Read Only or Protect from Volume Retention Manager
- Volumes which are set as T-VOL/R-VOL Disabled from Volume Security
- Volumes which are set as Read Only, Protect or S-VOL Disable from Data Retention Utility

(c) Specifying Volumes:

Source volumes and Target volumes are specified by LDEV number.

An open volume is also specified by one LDEV number. If the volumes are set as LUSE, you can specify by one LDEV number. HIHSM will migrate the LDEV which is one of a LUSE and its access is higher.

(d) Moving of multi volumes:

Moving of volumes can be performed by repeating instruction about each volume.

The number of formable copy pairs in HIHSM is calculated as the number of the differential tables used in HMRCF, HOMRCF and Hi-Copy to form pairs is subtracted from the total number of the differential tables, and it is up to 8,192 at the maximum. Refer to 2.7HMRCF function and HOMRCF function to calculate the number of the differential tables used in HMRCF, HOMRCF and Hi-Copy.

The number of formable pairs varies depending on the state of the emulation type and capacity of volumes.

The number of sheets of differential tables a pair used according to the emulation type and the capacity is calculated by using the following expression.

(i) For the main system

Differential tables per pair (α) = $\uparrow ((X) + (Y)) \times 15 \div (Z) \uparrow$

(X): User-specified size [Cyl]

(Y): Number of control cylinders [Cyl] (See Table 3.16.6-3.)

(Z): Number of slots managed by one differential table. (1916 × 32)

$\uparrow \uparrow$: rounding up decimal digits

For example, differential tables per pair for the emulation type 3390-3 is

$$(3,339 + 6) \times 15 \div (1,916 \times 32) = \uparrow 0.81836 \uparrow = 1$$

(ii) For the emulation type of the open system (OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-L)

Differential tables per pair (α) = $\uparrow ((X) \div 48[\text{KB}] + (Y) \times 15) \div (Z) \uparrow$

(X): User-specified size [KB]

(Y): Number of control cylinders [Cyl] (See Table 3.16.6-4.)

(Z): Number of slots managed by one differential table. (1916 × 32)

$\uparrow \uparrow$: rounding up decimal digits

For example, differential tables per pair for the emulation type OPEN-3 is

$$(2,403,360 \div 48 + 8 \times 15) \div (1,916 \times 32) = \uparrow 0.81860 \uparrow = 1$$

(iii) For the emulation type of the open system (OPEN-V)

Differential tables per pair (α) = $\uparrow ((X) \div 256[\text{KB}]) \div (Z) \uparrow$

(X): User-specified size [KB]

(Z): Number of slots managed by one differential table. $(1,916 \times 32)$

$\uparrow \uparrow$: rounding up decimal digits

For example, differential tables per pair for the emulation type OPEN-V is

$(3,019,898,880 \div 256) \div (1,916 \times 32) = \uparrow 192.40084 \uparrow = 193$

Table 3.16.6-3 The Number of the Control Cylinders According to the Emulation Types (for Mainframe System) (1/2)

Emulation type	Number of control cylinders (CYL)
3380-F	22
3380-3	7
3380-K	7
NF80-F	22
NF80-K	7
3390-3	6
3390-3R	6
3390-9	25
3390-L	23
3380-KA	7
3380-KB	7
3380-KC	7
NF80-KA	7
NF80-KB	7
NF80-KC	7
3380-M	53
3390-MA	53
3390-MB	53
3390-MC	53

Table 3.16.6-3 The Number of the Control Cylinders According to the Emulation Types (for Mainframe System) (2/2)

Emulation type	Number of control cylinders (CYL)
3390-3A	6
3390-3B	6
3390-3C	6
3380-3A	7
3380-3B	7
3380-3C	7
3390-9A	25
3390-9B	25
3390-9C	25
3390-LA	23
3390-LB	23
3390-LC	23

Table 3.16.6-4 The Number of the Control Cylinders According to the Emulation Types (for Open System)

Emulation type	Number of control cylinders (CYL)
OPEN-V	0KB (0Cyl)
OPEN-3	5,760KB (8Cyl)
OPEN-8	19,440KB (27Cyl)
OPEN-9	19,440KB (27Cyl)
OPEN-E	13,680KB (19Cyl)
OPEN-L	5,040KB (7Cyl)

Note: VLL functions are not applicable to OPEN-L volumes.

When forming pairs of volumes with various emulation types and capacity, the number of formable pairs is determined by the following conditional formulas.

The condition of the number of formable pairs: The maximum number of pairs that satisfies $\Sigma(\alpha) \leq (\beta)$

$\Sigma(\alpha)$: The sum total of the number of the differential tables per pair (α)

(β) : The number of the differential tables in the subsystem concerned

When the shared memories for the differential tables are not added, $(\beta) = 13,652$.

When the shared memories for the differential tables are added, $(\beta) = 30,718$.

For example, when 10 pairs of 3390-3 and 20 pairs of 3390-L (32,760CYL) are formed, the sum total of the number of the differential tables per pair, ($\Sigma(\alpha)$), is calculated as $(1 \times 10) + (9 \times 20) = 190$.

Because $\Sigma(\alpha) \leq (\beta)$, $190 \leq 13,652$ (when the memories of the differential tables are not added) or $190 \leq 30,718$ (when the memories of the differential tables are added) satisfies the condition of the number of formable pairs. Accordingly, 10 pairs of 3390-3 and 20 pairs of 3390-L (32,760CYL) can be formed.

(e) Abort moving:

Users can direct to abort the instructed moving before completion.

With aborting, the data in the destination volume is not guaranteed.

Users can direct to abort by each LDEV.

(f) Notice when the DKC is maintenance:

Volume migration may be failed if you start the following action:

(i) replace SM/Cache/HDD

(ii) install/deinstall SM/Cache/HDD

(2) Conditions for moving

Data moving is performed when all conditions about source volume and destination volume described below are satisfied.

(a) Both of the volumes have same emulation type.

(b) Both of the volumes have same size.

(c) There is no PIN data in the source volume.

(d) Both of the volumes are not blockade.

(e) Both of the volumes in same DKC.

(f) The volumes are not instructed to move already and not waiting to move.

(g) The volumes are not combination of CVS Volume and Normal Volume.

(3) Viewing History

Users can see the history of volume moving (migration).

3.16.7 Decision of volume moving (migration)

HIHSM supports decision of users about disk system performance tuning by logical volume moving (migration). This section describes usage and points to notice about monitor function.

(1) Inspecting utilization of system resources

First of all, using monitoring function, a user investigates whether there exists overloaded resources, or imbalance of resource utilization. Then the user tunes resource utilization in the manner described in the following clause.

Note: Due to average system resource utilization, there will be such a case as a portion of system performance will be negatively effected although total performance of a system will be improved. For example, if there exists RAID groups A and B of utilization 20% and 90% respectively, and if the utilization will become 55% and 55% if a logical volume residing in parity group B moves parity group A. Then response time of I/Os to parity group A will be increased while response time of I/Os and throughput to parity group B will be improved.

(2) Tuning Starnet utilization

Since Starnet are common resources in RAID500, migration of logical volumes does not improve system performance.

(3) Tuning CHP utilization

Migration of logical volumes does not improve CHP performance. Therefore if CHPs are overloaded on an average, the user should consider installation of new CHAs. And if the utilization of CHPs are imbalance, the user should consider that channel paths connecting to a CHA, which includes overloaded CHPs, is reconfigured into the connection to another CHA which includes CHPs of lower utilization.

(4) Tuning DKP/DRR utilization

If utilization of DRRs or DKPs is in high average, the user should consider installing new DKAs and HDDs. After installation of DKAs and HDDs, logical volumes which had high traffic of write access (especially of sequential write access) should be migrated to a parity group in newly installed HDDs.

If utilization of each DKP are imbalance, the user should consider migration of logical volumes from current parity group deployed under a DKA pair of high utilization to one under a DKA pair of lower utilization. The estimate function cannot simulate the DKP utilization. Therefore this method should be applied under the prospect of large improvement. For example, there would be least improvement in case of slight difference of utilization of each DKP, or if DRRs are comparatively in high utilization.

(5) Tuning RAID group utilization

If parity groups are in high utilization, the user should consider installing new HDDs. After installation of HDDs, logical volume which had high traffic of I/Os should be migrated to a parity group in newly installed HDDs.

If utilization of each parity group is imbalanced, the user should consider migration of logical volumes from the current parity group showing high utilization to the one showing lower utilization.

These methods should be applied with a view to large improvement. There would be least improvement in a case of slight difference of utilization of each parity group, or if DRRs or DKPs are already comparatively highly utilized.

When errors exist in the system, utilization of system resource can increase or be unbalanced.

3.16.8 Automating Volume Migration function

(1) Automating Volume Migration function Overview


Automating Volume Migration function provides a typical tuning method of Volume Migration based on parameters given by users, and performs tuning plan automatically.

(2) Overview of tuning

Tuning by Volume Migration can be proceeded by performing the following steps repeatedly.

- (a) Monitoring information
- (b) Analyzing information
- (c) Making volume migration plan (decision of volume migration)
- (d) Moving volume (migration)
- (a') Monitoring information again to confirm condition and effect of the performed tuning.

Preset function reduces users' work by providing a typical tuning method of Volume Migration based on parameters given by users, and performs tuning plan automatically. When (b) Analyzing and (c) Making plan are done by users, it can perform fine tuning.



(3) Process of tuning by Automating Volume Migration function


(a) Basic process flow of tuning by Automating Volume Migration function

Preset function performs the following two processes.

- Making Plan: Detecting volumes that have problems by monitored information (disk utilization) and making plan of volume Migration
- Migration: Moving volume according to the plan.

These two processes are repeated by the cycle described above, the length of the cycle depends on users. In the preset function, this cycle is supposed from one day to one week or several weeks at most.


Users can specify the term where the monitored information is used to make the plan focusing on concerns of users. For example, users can specify the highest load term in a day or in a week as the referred term for making plan.



(b) Confirming and remaking of plan by users

During the period from the planning by preset function until migration, users can see the plan and delete it if needed.

During the same period as above, users can remake the plan manually with changing parameters.



(4) Making plan

(a) Tuning based on disk utilization ratio


HIHSM monitors utilization ratios of various system resources, but the preset function uses disk utilization ratio to make a migration plan to solve disk neck as a typical tuning method. Users could refer to other information if needed.

Users also specify parameters for tuning based on disk utilization ratio. Users should set these parameters for their system (preset function provides default value for these parameters roughly).

(b) Hierarchy of parity groups and management by class

Parity groups in DKC have hierarchy by drive type and RAID type. HIHSM provides a function to optimize the usage of this hierarchy. Preset function manages this hierarchy as class (parity group set).

Parity groups are divided into classes. The classes are ordered from high level (high performance) to low level (low performance). This classification is decided by performance of physical drive type of each parity group.



A method of migration plan which the preset function performs is described below.


Preset function uses the monitored information in the term that is the data to be referred to make the plan.

(b-1) Management by maximum limit of disk utilization ratio

A maximum limit of disk utilization (parity group utilization) is specified to each parity group. Users should specify this limit for their systems. (HIHSM uses the default value but it is desirable that user sets the desired value.)


For the parity groups that exceed this limit, preset function makes plan of moving volumes from this parity group to another parity group in higher class.

This avoids physical disk neck and provides load balancing of disk utilization.

**(b-2) Selection of volumes to migration**

In the migration plan described above, volumes of larger disk utilization are selected to be moved from the parity group that exceeds the limit to higher class. Moving larger utilization volumes to higher performance class is expected to make large tuning effect. And larger utilization means larger amount of access from host, and this also makes large tuning effect.

When the reserved (empty) volumes run short in the high class, volumes of smaller disk utilization are selected to be moved from higher class to lower class to make reserve volumes.



Preset function provides some criteria for you to select them, average of disk utilization, average of highest Nth value of disk utilization in the referred term, and a value considering sequential/random access pattern.

(b-3) Specifying maximum limit of disk utilization ratio by users

When the same maximum limit of disk utilization ratio described above is given to each class (parity group) classified by drive types, performance of each physical drive type makes performance of each class directly.

When the users specify the limit to each class with bias, users can make the difference of performance of classes larger or smaller.

Users can specify parity groups to fixed parity groups in which volumes users do not want to move automatically. Preset function does not make a migration plan about fixed parity groups and volumes in the fixed parity groups.

(b-4) Notice for making plan

HIHSM can make these plans only on the following conditions:

- HIHSM can estimate the disk utilization ratio against all migrated parity groups.
- The disk utilization rate of all migrated parity groups should not be over the maximum rate. If the rate of one parity group is over the maximum rate, HIHSM could not make a plan.

(b-5) Notice for reference term

HIHSM could not use old information before the last volume migration in order to reduce the influence of performance by volume migration.

Therefore, HIHSM sometimes fails in making a plan by lack of information.

(5) Moving (migration) by preset function

Preset function performs moving (migration) process once a day at time specified by users.

If there is a migration plan made by preset function at the time, volumes are moved by the plan.

Users can specify the limit of moving to avoid the overload by moving (data copy) process.

If the disk utilization ratio of parity groups in which moving is started exceeds the moving process limit, the moving in the parity group will be aborted.

Users can specify the time limit of moving. If a plan does not complete in the time limit, the remaining plan will be executed in the following day. If the new plan will be made until the next migration time, those remaining plans will be deleted before making the plan.

3.17 Data Assurance at a Time When a Power Failure Occurs

3.17.1 Outline

Operation that is done when a power failure occurs is as follows.

It is possible to retain data in the cache by means of the memory backup. When the power stoppage of the DKU occurs, the LDEV is detached temporarily and after the power supply is recovered, the LDEV status is returned to normal automatically.

Blank Sheet

Blank Sheet

Blank Sheet

3.18 HPAV

3.18.1 Overview

3.18.1.1 Overview of HPAV

HPAV (Hitachi Parallel Access Volume) enables a host computer to issue multiple I/O requests in parallel to each device in the disk subsystem. Usually, host computers are able to issue only a single I/O request to a single device. When a host computer issues one I/O request to a device, the host computer is unable to issue another I/O request to that device. However, HPAV enables you to assign one or more aliases to a single device so that the host computer is able to issue multiple I/O requests. In this way, HPAV provides the host computer with substantially faster access to data in the disk subsystem.

When you assign aliases to a device, you specify the addresses of unused LDEVs (logical devices or logical volumes) in the disk subsystem. The specified addresses are used as alias addresses.

Throughout this manual, the term base device refers to a device to which aliases will be assigned. Also, the term alias device refers to an alias.

HPAV operates in either of the following ways: static PAV and dynamic PAV. These are described next:

■ Static PAV

When static PAV is used, the number of aliases for each base device remains unchanged even if the number of I/O requests to each device changes. As explained later, when dynamic PAV is used, the number of aliases for a base device is likely to increase as the number of I/O requests to the device increases; this means the number of aliases for other base devices may decrease.

However, when static PAV is used, the number of aliases remains as specified by the Remote Console user or SVP operation user.

Before you assign aliases to base devices, you should consider whether I/O requests will converge on some of the base devices. We recommend that you assign more aliases on base devices on which I/O requests are expectedly converge. Otherwise, HPAV might not be able to provide much faster access to data in the disk subsystem.

The following figure gives an example of static PAV. In this figure, each of the three base devices (numbered 10, 11, and 12, respectively) has two aliases assigned. I/O requests converge on the base device #10, but the number of aliases for each base device remains unchanged.




Fig. 3.18.1.1-1 Static PAV

■ Dynamic PAV

When dynamic PAV is used, the number of aliases for a base device may change as the number of I/O requests to the device changes. If I/O requests converge on some of the base devices, the number of aliases may increase for these base devices but may decrease for the other base devices. Dynamic PAV can balance workloads on base devices and optimize the speed for accessing data in the disk subsystem.

The following figure gives an example of dynamic PAV. In this example, each of the three base devices (#10, #11, and #12) was originally assigned two aliases. As I/O requests converge on #10, the number of aliases for #10 increases to four. For the base devices #11 and #12, the number of aliases decreases to one.

Dynamic PAV requires the Workload Manager (WLM), a special function provided by the operating system at the host computer. For details, see sections 3.18.2.1 and 3.18.2.2.




Fig. 3.18.1.1-2 Dynamic PAV

3.18.1.2 How to Obtain Optimum Results from HPAV

To obtain good results from HPAV, you should be aware of the following:

- The best results can be obtained if the number of aliases is “Number of available channel paths minus 1”. If the number of aliases is specified this way, I/O operations can use all the channel paths, and thus the best results can be obtained.
- HPAV may not produce good results when many channel paths are used. If all the channel paths are used, no good results can be expected.
- HPAV lets you assign unused devices for use as aliases. If you assign most of the unused devices for use as aliases, only a small amount of free devices are available. It is recommended that you think about adding more disks in the future when you determine the number of aliases to be assigned.

If we assume that there are 256 devices and we assign the same number of alias devices to each base devices, the number of base devices and alias devices is calculated as explained in Table 3.18.1.2-1. The recommended ratio of base devices to alias devices is 1:3.

If you can expect the types of jobs to be passed to base devices, or if you can expect how many accesses should be made to each base device, you should determine the number of aliases for each base device so that it meets the requirements for each base device.

Table 3.18.1.2-1 The ratio of base devices to aliases

Ratio (base devices : alias devices)	The number of base devices	The number of alias devices
1:3 (recommended)	64	192
1:1	128	128

- Good results cannot be expected on devices that are always shared and used by multiple host computers.
- If dynamic PAV can be used in all the systems, good results can be expected if you assign 8 to 16 aliases to each CU (control unit).

3.18.2 Preparing for HPAV Operations

3.18.2.1 System Requirements

To be able to run, HPAV requires the following operating systems to be installed on the host computer:

- For static PAV
 - OS/390 V1R3 & DFSMS/MVS 1.3 with PTF
 - VM/ESA 2.4.0
- For dynamic PAV
 - OS/390 V2R7 with PTF & DFSMS/MVS 1.5

Note: To perform operations with HPAV, you must have administrator access privileges. Users who do not have administrator access privileges can only view HPAV information.

The following restrictions apply when using HPAV.

Table 3.18.2.1-1 Restrictions that apply when using HPAV

No.	Item	Specifications
1	DKC emulation type	I-2105/2107
2	DKU emulation type	3390-3, 3390-3R, 3390-9, 3390-L, 3390-M, 3380-3
3	Number of aliases that can be assigned to a single base device	Up to 255
4	Alias device numbers that can be used	When you set aliases for base devices, you can use the device numbers of unused devices as the alias device numbers. When you set aliases for base devices, you must be aware that the alias devices and the base devices must belong to the same CU.
5	Device functions that can concurrently be used with HPAV	• CVS (Customized Volume Size) • DCR (Dynamic Cache Residence) • LDEV Security
6	Device functions that cannot concurrently be used with HPAV	• HMBR (Hitachi Multiplatform Backup/Restore) • HMDE
7	Copy functions that can concurrently be used with HPAV	• Concurrent Copy/XRC – Do not intermix DKC emulation type of 2105/2107 and 3990-6/6E in the MCU. • HRC (Hitachi Remote Copy) / HMRCF (Hitachi Multiple RAID Coupling Feature) / HIHSM (Hitachi Internal Hierarchical Storage Manager) – You can be intermixed DKC emulation type of 2105/2107 and 3990-6/6E in the same DKC.

3.18.2.2 Preparations at the Host Computer

This section briefly describes arrangements that should be made at the host computer. For detailed information, see the documentation for MVS.

3.18.2.2.1 Generation Definition of Base Devices and Alias Addresses

The address mapping between base devices and the corresponding alias devices must be defined at generation.

The address mapping between base devices and alias devices at the host computer should match the corresponding address mapping at the DKC side. If it does not match, a serious failure might occur during data processing.

The following gives an example of mapping between base devices and aliases devices:

(A) x 00-x1F:Base	(B) x 00-x3F:Base	(C) x 00-x7F:Alias	(D) x 00-x3F:Alias
x 20-xFF:Alias	x 40-x7F:Alias	x 80-xFF:Base	x 40-x7F:Base
	x 80-xBF:Base		x 80-xBF:Alias
	x C0-xFF:Alias		x C0-xFF:Base

Note: The recommended ratio of base devices to aliases is 1:3, if each base device is assumed to be assigned the same number of aliases.

3.18.2.2.2 Setting the WLM Operation Mode

If you want to use dynamic PAV, you must set the WLM (Workload Manager) operation mode to *goal mode*. WLM manages workloads on MVS and can use two operation modes, which include goal mode. In goal mode, WLM manages the system to fulfill the performance goal that was specified before the system began to operate.

You should be aware that static PAV is used instead of dynamic PAV if compatibility mode is used instead of goal mode.

For details on the WLM operation modes, see the documentation for MVS.

3.19 FICON

3.19.1 Introduction

FICON is the new mainframe architecture, which is FC-SB-2 protocol based on Fiber channel physical layer protocol (FC-PH) and it is approved by ANSI.

The specification of FICON is below.

- Full duplex data transfer
- Multiple concurrent I/O operations on channel
- High bandwidth data transfer (100MB/s, 200MB/s, 400MB/s)
- Fewer control unit interfaces
- Pipelined CCW execution

Note: When you select 2107 DKC Emulation, power-on and POR (Power off/on) of SOM493 are required in advance.

3.19.2 Environment

If you use FICON, the environment shown below is required.

Table 3.19.2-1

Items	Contents	
CPU	z9, z990, z900, G5/G6, z800	
DKC Emulation type	2105-F20, 2107	
FICON Director	MCDATA	ED-5000
		ES-3232
		ES-4700
		ED-6064
		ED-6140
		ED-10000 (i10K)
	McDATA (CNT)	FC9000
		UMD
	Brocade	Silkworm12000
		Silkworm24000

3.19.3 DKC515I FICON specification

Table 3.19.3-1 shows the specification of DKC515I FICON.

Table 3.19.3-1 DKC515I FICON specification

Items		Contents
Range of CU address		0 to 3F (*4)
Number of logical volumes		1 to 16384
Number of connectable channel port		8 (*1) 16 (*2)
Maximum number of logical paths per CU		2,048
Maximum number of logical paths per port		65,536 (*3)
Maximum number of logical paths per CHA		65,536
Maximum number of logical paths per subsystem		131,072
Support DKC emulation type		2105-F20, 2107
Support fiber channel	Bandwidth	1Gbps / 2Gbps / 4Gbps
	Cable and connector	LC-Duplex
	Mode	Single Mode Fiber/Multi Mode Fiber

*1: In case of 4 ports/CHA.

*2: In case of 8 ports/CHA.

*3: The maximum number of connectable hosts per FICON port is 1024.
(1024 host paths x 64 CUs = 65,536 logical paths.)

*4: Number of CUs connectable to the one FICON channel (CHPID) is 64.
The CU addresses in the interface with a host are 00 to 3F for the FICON channel.
When the number of CUs per FICON channel (CHPID) exceeds the limitation, there is a possibility that HOST OS IPL fails.

3.19.4 Configuration

(1) Topology

The pattern of connection between DKC515I FICON and channel are such as below.

- Point to point connection
- Switched point to point connection
- Non cascading connection
- Cascading connection




Fig. 3.19.4-1 FICON Topology

(2) Switched point to point configuration

FICON director (FICON switch) specifications are below.

Table 3.19.4-1 FICON director specification

FICON Director		Bandwidth	Connector
Vendor	Model		
McDATA	ED-5000	1Gbps	SC-Duplex
	ES-3232	1/2Gbps	LC-Duplex
	ES-4700	1/2Gbps	LC-Duplex
	ED-6064	1/2Gbps	LC-Duplex
	ED-6140	1/2Gbps	LC-Duplex
	ED-10000 (i10K)	1/2Gbps	LC-Duplex
McDATA (CNT)	FC-9000	1Gbps	SC-Duplex
		1/2Gbps	LC-Duplex
	UMD	1/2Gbps	LC-Duplex
Brocade	Silkworm12000	1/2Gbps	LC-Duplex
	Silkworm24000	1/2Gbps	LC-Duplex

Note:

1. Above switches are all that RSD certified connectivity between DKC515I FICON and each switch.
2. Do not change the Operation Speed mode of switch during I/O operation.

(3) FICON Cascade configuration

Cascaded FICON director specifications are below.

Table 3.19.4-2 Cascaded FICON Director specification

FICON Director		Necessary Feature
Vendor	Model	
McDATA	ES-3232 ES-4700 ED-6064 ED-6140 ED-10000 (i10K)	(1) SANteglity (2) FICON Management Server
McDATA (CNT)	FC-9000 UMD	(1) Fabric security License
Brocade	Silkworm12000 Silkworm24000	(1) Secure Fabric OS License (2) Advanced Zoning License

Note:

1. FICON cascade configuration is only supported by zSeries servers.
2. The FICON directors used by the cascade connection must be the same vendor.

(4) Configuration of FICON and ESCON intermixed

(a) Migration from ESCON to FICON

FICON/ESCON intermix within the same path group is allowed for only migration from ESCON to FICON.

(b) FICON and ESCON intermixed configuration within the same subsystem

If the generation of ESCON and that of FICON are separated, each system can use the same volume, without disturbing each other.

3.19.5 The operation procedure

(1) Notice about Speed Auto Negotiation

A stable physical environment (fully mated connectors, no cable flexing, no transient noise sources, etc.) is expected on Speed Auto Negotiation. Otherwise, Speed Auto Negotiation may settle to not fastest speed but an optimum speed.

To change into fastest speed, check whether it is the stable physical environment, and execute either of the following operations.

Confirm the Link speed from the 'Physical path status' window of DKC after executing each operation.

- DKC PS OFF/ON
- Dummy replace the package including the FICON ports.
- Remove and insert the FICON cable which is connected to the FICON port in DKC. (*1)
- Block and Unblock the associated outbound switch/director port. (*1) (*2)

*1: Execute this after deleting the logical paths from the host with the "CHPID OFFLINE" operation. If this operation is not executed, Incident log may be reported.

*2: Alternate method using switch/director configuration.

<Operating procedure from switch control window (Ex: in the EFCM)>

- (a) Block the associated outbound switch/director port to the CHA interface that is currently "Negotiated" to not fastest speed (Ex: 1Gb/s).
- (b) Change port speed setting from "Negotiate mode" to "Fastest speed fix mode (Ex: 2Gb/s mode)", then from "Fastest speed fix mode (Ex: 2Gb/s mode)" back to "Negotiate mode" in the switch/director port configuration window.
- (c) Unblock the switch/director port.
- (d) Observe the "Online" and "Fastest speed (Ex: 2Gb/s)" link establish without errors on the switch/director port status window.

(2) About the change work of the package

The change by "CHA Type Change Operation" is possible when there is no vacant slot for the addition at the time of the package change.

CHA before the change	CHA after the change
8 port adapter (8MS/8MSR/8ML/8MLR)	16 port adapter (16MSR/16MLR)
8 port adapter (8MS/8MSR/8ML/8MLR)	16 port adapter (16MFSR/16MFLR/16MFL4R)
16 port adapter (16MSR/16MLR)	16 port adapter (16MFSR/16MFLR/16MFL4R)
16 port adapter (16MFSR/16MFLR/16MFL4R)	16 port adapter (16MSR/16MLR)
16 port adapter (16MSR/16MLR) *	8 port adapter (8MS/8MSR/8ML/8MLR)
16 port adapter (16MFSR/16MFLR/16MFL4R) *	8 port adapter (8MS/8MSR/8ML/8MLR)

* The emulation type of CHA should not be mixed.

Refer to "CHA type change operation" of "SVP Section" for the operation procedure.

(3) Configuring Alternate paths

In mainframe systems, the concept of alternate paths is available for the purpose of avoiding system down.

We recommend you to configure alternate paths based on the following priorities.

TagmaStore USP supports following mainframe fibre packages.

- Mainframe fibre 8 port adapter
- Mainframe fibre 16 port adapter

The internal structure and characteristics of each package are as shown in Table 3.19.5-1.

Table 3.19.5-1 Internal structures and characteristics of TagmaStore USP mainframe packages

Item	Mainframe fibre 8 port adapter	Mainframe fibre 16 port adapter
Internal structure		
Characteristics	• Can access 2CHPs from 1Port (↔) • 1HTP controls 1Port • 2HTP/1MHUB/4CHP	• Can access 4CHPs from 1Port (↔) • 1HTP controls 2Port • 2HTP/1MHUB/4CHP

HTP: Processor used for controlling FICON protocols

MHUB: LSI for mainframe control

CHP: Processor used for controlling channels

Port 1A/3A/5A/7A/1B/3B/5B/7B: Port IDs in Cluster 1/BASIC slot

The mainframe fibre 16 port adapter can control four processors (CHP) from four ports. If you use different MHUB ports (For example, Port 1A and 1B in the figure), one port can use four processors exclusively. Therefore, the throughput performance of one path is better than using the same MHUB ports (For example, Port 1A and 5A).

In addition to the packages described above, TagmaStore USP provides power redundancy using the cluster configurations.

Considering the structures and performance, we recommend you to set paths based on the following priorities when configuring alternate paths.

Priority 1: Set paths to clusters

Priority 2: Set paths to packages

Priority 3: Set paths to MHUBs

Priority 4: Set paths to HTPs (Mainframe fibre 16 port adapter only)

3.20 Universal Volume Manager (UVM)

3.20.1 Overview

Universal Volume Manager is a program product that realizes the virtualization of the storage subsystem. Universal Volume Manager enables you to operate multiple storage subsystems including RAID500 subsystem as if they are all in one storage subsystem. As Universal Volume Manager realizes the virtualization of the storage subsystem, the system administrator can manage the different kinds of multiple storage subsystems as one storage subsystem.

Once you connect the RAID500 subsystem and another kind of storage subsystem using Universal Volume Manager, the system administrator can also manage another storage subsystem using RAID500 Storage Navigator. For example, the system administrator can set the path from a host computer to the volume of DF600 subsystem using LUN Management of RAID500.

In addition to the function of Universal Volume Manager, the Cross-system Copy function enables you to easily make a backup copy of data stored in the RAID500 subsystem to another storage subsystem. It is also easy to restore the backed up copy of data to RAID500 subsystem.

In this manual, the source RAID500 subsystem is called “local subsystem”, and the connected storage subsystem is called “external subsystem”. The volume managed in the local subsystem is called “internal volume”, and the volume in the external subsystem is called “external volume”.

Features of Universal Volume Manager are as follows:

By mapping an external volume as an internal volume using Universal Volume Manager, it becomes possible to operate the external volume using RAID500 Storage Navigator as if it is a volume in the RAID500 subsystem.

“Mapping” means assigning the CU:LDEV numbers of the internal volumes to the external volumes. By assigning the numbers of the internal volumes to the external volumes, the system administrator will be able to operate not only internal volumes of RAID500 subsystem but also external volumes using RAID500 Storage Navigator.

Universal Volume Manager has an option feature, called Cross-system Copy.

If you use Cross-system Copy with Universal Volume Manager, you can make copy bidirectional between external volume and internal volume.




Figure 3.20.1-1

Figure 3.20.1-1 shows the idea of connection between a RAID500 subsystem and an external storage subsystem which are connected by the Universal Volume Manager function. In Figure 3.20.1-1, the external subsystem is connected to the external port of the RAID500 subsystem via a switch (or hub) using the fibre-channel interface. The external port is a kind of port attribute, which is used for Universal Volume Manager. In Figure 3.20.1-1, the external volumes are mapped as RAID500 volumes.

3.20.2 Procedure of using external volumes

- (1) Prepare in the external subsystem a volume to be used for UVM.
- (2) Change the port attribute to External.
- (3) Connect the external subsystem to the external port.
- (4) Search for the external subsystem from the UVM operation panel (Discovery).
- (5) Map an external volume.
 - (a) Register to an external volume group
 - (b) Select the IO Suppression mode
 - (c) Select the cache mode
- (6) Format the volume (MF volume only)
- (7) Define the host path (Open volume only)
- (8) Other settings

3.20.2.1 Prepare in the external subsystem a volume to be used for UVM

Prepare a volume to be used for UVM in the external subsystem connected to the RAID500.

Note: The volume in the external subsystem should be about 38 MB (77,760 Blocks) ~ 2 TB (4,294,967,296 Blocks). If the capacity of the external volume is 2 TB or larger, you can use up to 2 TB of the volume.

3.20.2.2 Change the port attribute to External

The port used for Universal Volume Manager needs to be set as the external port. When the external subsystem is connected to the external port of the local subsystem, you can view the information on the external subsystem from the RAID500 Storage Navigator computer. The external subsystem cannot be connected to the ports other than the external port.

In order to set the port attribute to external, you need to release the paths set to the port. The attribute of the port where the paths are set cannot be changed to external. Before starting the Universal Volume Manager operations, you need to know the ports whose attributes can be changed to external.

Note: The ports whose attributes are set for remote copy software (eg., RCU target, initiator) or the other features cannot be used as external ports for Universal Volume Manager. Change the port attribute to external if the port attribute is set to other than external.

3.20.2.3 Connect the external subsystem to the external port

Insert a Fibre cable into the external port from the external subsystem.

3.20.2.4 Search for the external subsystem from the UVM operation panel (Discovery)

Search for the external subsystem from the UVM operation panel (LU Operation tab) in the StorageNavigator.

You cannot the external volume from the RAID500 by just discovering the external subsystem. You need to perform the Add LU operation (mapping) shown in the next section.

3.20.2.5 Map an external volume

When you connect the external subsystem to the external port, volumes in the external subsystem (external volumes) can be mapped as volumes in the local subsystem (internal volumes). Confirm which volumes in which external subsystem you need to map as internal volumes.

Only one external volume can be mapped as one internal volume. The maximum number of external volumes, which can be mapped, is 1,024 per port.

When the external volume is more than 2 TB (4,294,967,296 Blocks), you can access the data stored in the field up to 2 TB. You cannot access the data that is stored in the field over 2 TB. The external volumes of about 38 MB (77,760 Blocks) or smaller cannot be mapped. When you want to set the OPEN-V emulation type, 47 MB (96,000 Blocks) or smaller external volume cannot be mapped.

(1) Register to an external volume group

When you map an external volume as an internal volume, you need to register the external volume to an external volume group.

The user can classify the external volumes, which is set by Universal Volume Manager, into the groups according to the use. The group is called external volume group (ExG). For instance, you can register multiple volumes in one external subsystem to one external volume group. Or, even though the data you want to manage in a lump is stored in volumes in the different external subsystems, you can register the volumes in one external volume group and manage them in block.

You need to assign numbers from 1 to 16,384 to external LU groups. A maximum of 16,384 external volume groups can be created. A maximum number of volumes, which can be registered in one external group, is 256.

(2) Select the external volume attribute

When you map an external volume as an internal volume, you set the attributes of the external volume. The attributes of an external volume can be set using Add LU panel of Universal Volume Manager.

(a) I/O Suppression mode (IO Suppression mode: Disable or Enable)

You can set when you map the volume, if you allow the I/O from the hosts to the mapped external volume or not.

When you select Disable for the I/O Suppression mode, you can use the mapped external volume as if it is a volume inside the local RAID500 subsystem from the host. The volume that has the Disable attribute for the I/O Suppression mode cannot be used for the Cross-system Copy operation.

When you select Enable for the I/O Suppression mode, the mapped volume can be used only for the Cross-system Copy operation.

(b) I/O cache mode (Cache mode: Disable or Enable)

You can set if the operation uses the cache or not when the host I/O is demanded. If you select Enable, the host I/O once goes to the cache and then to the volume. If you select Disable, the host I/O comes directly to the volume not coming through the cache.

When the IO Suppression mode is Enabled, the cache mode is Disabled.

(3) Emulation type

You can set the emulation type of the mapped volume. However, if you are planning to use the mapped volume for the Cross-system Copy operation, you need to select the OPEN-V emulation type.

If you select the emulation type other than OPEN-V, you need to know that the volume requires some specific area provided for the data of RAID500 management information. This means that the volume capacity after the mapping becomes less than the actual external volume capacity for the area (volume) provided for the data management information.

3.20.2.6 Format volume (MF volume only)

If you set the MF emulation type (3380-x/3390-x) for the external volume which is mapped as an internal volume, you need to format the LDEV before using the volume.

Use the VLL feature to format the LDEV.

3.20.2.7 Define the host path (Open volume only)

If you set the Open emulation type (Open-x) for the external volume which is mapped as an internal volume, define a path between the host and the internal subsystem. Use the LUNM feature to define a path.

3.20.2.8 Other settings

(1) Define alternate path to the external volume

When you map an external volume as an internal volume, the path(s) will be set from the internal volume to the external volume. When two paths, which can be set, from the two different clusters to the external volume exist, the two paths are set at the time of the mapping. When two settable paths do not exist, one path is set at the time of the mapping.

You can set up to eight paths to the external volume including the automatically set paths. Among the paths to the external volume, a path that is given the highest priority is called a primary path. Paths other than the primary path are also called alternative paths.

When the external volume is mapped as the internal volume using Universal Volume Manager, the host I/O operations to the external volume are enabled normally using the path set in the mapping operation. However, the path is automatically switched to the alternate path when the path set in mapping operation cannot be used due to, for instance, maintenance operation in the subsystem, or a failure in the channel processor. Because the path is switched to the alternate path, you can continue performing the I/O operation to the external volume that is mapped by Universal Volume Manager as usual even though an error occurred in the original path.

If the alternate path is not set, host I/O, a copying operation with Cross-system Copy is aborted when a maintenance operation is performed for the subsystem or a trouble such as a failure in the channel processor occurs.

It is recommended to set the alternate paths for safer operation.

To set the alternate path, use the path setting function of Universal Volume Manager.

Figure 3.20.2.8-1 illustrates an example of setting an alternate path. In Figure 3.20.2.8-1, external subsystem ports, “WWN A” and “WWN B”, are connected to “CL1-A” and “CL2-A” respectively which are set to the external ports in the RAID500 subsystem. You need to specify the port of a different cluster in the RAID500 subsystem for the alternate path as “CL1” port and “CL2” port are specified in Figure 3.20.2.8-1.




Figure 3.20.2.8-1

Figure 3.20.2.8-2 also illustrates an example of setting an alternate path. In Figure 3.20.2.8-2, two ports are specified in the RAID500 subsystem, and connected to the ports in the external subsystem via the switch. In this case, two ports of different clusters are specified in the RAID500 subsystem. Therefore, the setting of the alternate path is enabled.

In Figure 3.20.2.8-3, two paths are also set between the internal volume and the external volume. However, one port is specified in the RAID500 subsystem, and two ports are specified in the external subsystems via the switch. Since two ports of different clusters need to be set in the RAID500 subsystem for alternate path settings in Universal Volume Manager, we do not recommend the setting shown in Figure 3.20.2.8-3.




Figure 3.20.2.8-2




Figure 3.20.2.8-3

3.20.3 Cross-system Copy

3.20.3.1 Overview

After you have connected a subsystem other than RAID500 subsystem to the RAID500 subsystem using Universal Volume Manager, Cross-system Copy enables you to copy the data stored in the RAID500 subsystem to a storage subsystem other than RAID500. It also allows you to copy the data stored in a storage subsystem other than RAID500 to the RAID500 subsystem.

If you can copy data between a RAID500 subsystem and a storage subsystem other than RAID500, you will be able to, for example, gather the data from the multiple storage subsystems to one RAID500 and compile them. And then, you will be able to distribute the compiled data back to the multiple storage subsystems in result.

Cross-system Copy also enables you to back up the data in a RAID500 subsystem to other storage subsystem easily. You can easily restore the backed up data to the RAID500 subsystem.

Features of Cross-system Copy are as follows.

After the mapping of an external volume as an internal volume using Universal Volume Manager, Cross-system Copy enables you to copy data from a storage subsystem other than RAID500 (external volume) to a RAID500 volume using Cross-system Copy.

For instance, you can easily migrate the data to a RAID500 when you remove the external subsystem, by copying all the data in the external volumes to the RAID500 volumes.

After the mapping of an external volume as an internal volume using Universal Volume Manager, Cross-system Copy enables you to create a copy of the RAID500 volume in the external volume. The system administrator will be able to make copies of data in the RAID500 to the external volumes as if the copies are created in the internal volumes.

Figure 3.20.3.1-1 shows the concept of copying the data in the external volumes to the RAID500 volumes. In Figure 3.20.3.1-1, “A”, the volume in the external subsystem is mapped as an internal volume in the RAID500 subsystem. In addition, “A”, the data in the volume, is copied to the RAID500 volume using the copy feature of Cross-system Copy.




Figure 3.20.3.1-1

Figure 3.20.3.1-2 shows the concept of copying the data in the RAID500 volumes to the external volumes. In Figure 3.20.3.1-2, “B”, the volume in the RAID500 subsystem is copied to the mapped external volume using Cross-system Copy.




Figure 3.20.3.1-2

3.21 Save and Recovery of the SM Control Information in the case of PS OFF/ON

In the DKC515I disk subsystem, the control information on SM is saved in HDD of SVP at the time of PS OFF, and it recovers from SVP to SM when the volatilization PS is ON. The processing of TrueCopy or ShadowImage can be restarted from the status immediately before PS OFF by this function even in the case of the volatilization PS ON.

This function saves the SM information at the time of PS OFF, and recovers in the case of the volatilization PS ON.

The PP function that this function supports is shown below:

- (1) ShadowImage
- (2) ShadowImage OS/390
- (3) Volume Migration
- (4) Compatible Mirroring for IBM® FlashCopy®
- (5) Compatible Mirroring for IBM® FlashCopy® Version 2
- (6) TrueCopy
- (7) TrueCopy for Mainframe
- (8) Universal Replicator
- (9) Universal Replicator for Mainframe
- (10) Copy-on-Write Snapshot

Save is executed only at the time of PS OFF, and it recovers when the volatilization PS is ON only if the effective save data exists in SVP. It does not operate in the following cases:

- (1) Define Configuration
- (2) System Tuning
- (3) Offline microprogram exchange
- (4) Online microprogram exchange
- (5) Power supply destage

It is necessary to set the local mode to use this function.

Note: When this function is applied, the PS OFF/ON time may extend for about 5 to 30 minutes.

The PS OFF extension time depends on the save recovery processing on the number of executable processors (it is usually 16, but it is 6 at the minimum for the NS model, and it may be less if the processor is blocked).

When save is failed, it reports SIM.

If this SIM is detected, PS should be ON as promptly as possible before the battery volatilizes.

If save fails and the volatilization PS is ON, ShadowImage or the TrueCopy pair becomes all copies or suspended. That is, it becomes the processing same as the existing volatilization PS ON (before this function is supported). Refer to the basic specification described in the device explanation of this manual or the user's guide for the event that occurs at the time of the volatilization PS ON.

3.22 Universal Replicator

This Chapter explains Universal Replicator (hereafter referred to as UR).

3.22.1 URz Components

URz operations involve the TagmaStore™ USP subsystems at the primary and secondary sites, the physical communications paths between these subsystems, and the TagmaStore™ USP URz remote console software. URz copies the original online data at the primary site to the offline backup volumes at the secondary site via the dedicated fibre-channel remote copy connections using a journal volume. You can operate the URz software with the user-friendly GUI environment using the TagmaStore™ USP URz remote console software. Note: Host failover software is required for effective disaster recovery with URz.

For management of URz journal groups that consist of journal volumes located in multiple subsystems, host I/O time stamping function (provided by MVS DFSMSdfp) is a requisite functional item. An error reporting communications (ERC) feature is essential for URz to be able to recover data lost in a disaster.

Figure 3.22-1 shows the URz components and their functions:




Figure 3.22-1 URz components

Table 3.22-2 shows the plural secondary subsystems connection configuration of URz. By connecting one primary subsystem with more than one secondary subsystem, you can create a volume pair that has a one-to-one relationship for each journal group.




Figure 3.22-2 Connection Configuration of Plural Secondary Subsystems

This URz components describes:

- TagmaStore™ USP RAID storage subsystem (see section 3.22.1.1)
- Main and remote control units (primary subsystems and secondary subsystems) (see section 3.22.1.2)
- Journal group (see section 3.22.1.3)
- Data volume pair (see section 3.22.1.4)
- Journal volume (see section 3.22.1.5)
- Remote copy connections (see section 3.22.1.6)
- Initiator ports and RCU target ports (see section 3.22.1.7)
- TagmaStore™ USP URz remote console software (see section 3.22.1.8)
- Host I/O time stamping function (see section 3.22.1.9)
- Error reporting communications (ERC) (see section 3.22.1.10)

3.22.1.1 TagmaStore™ USP Series Storage Subsystems

URz operations involve the TagmaStore™ USP subsystems at the primary and secondary sites. The primary subsystem consists of the main control unit (primary subsystem) and SVP. The secondary subsystem consists of the remote control unit (secondary subsystem) and SVP.

3.22.1.2 Main and Remote Control Units (Primary Subsystems and Secondary Subsystems)

The main control unit (primary subsystem) and remote control unit (secondary subsystem) control URz operations:

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

3.22.1.3 Journal Group

Journal group consists of two or more data volumes and journal volumes. It is a feature that allows you to sort multiple data volumes and journal volumes into collective units to tailor URz to meet your unique business needs. The journal group in the primary subsystem is referred to as the master journal group. The journal group in the secondary subsystem is referred to as the restore journal group. The data volumes in the master journal group are also called the primary data volumes. The journal volumes in the master journal group are called the master journal volumes. The data volumes in the restore journal group are similarly called the secondary data volumes. The journal volumes in the restore journal group are called the restore journal volumes.

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.

3.22.1.4 Data Volume Pair

URz performs remote copy operations for data volume pairs created by the user. Each URz pair consists of one primary data volume and one secondary data volume which can be located in different subsystems. The URz primary data volumes are the primary volumes (LDEVs) which contain the original data, and the URz secondary data volumes are the secondary volumes (LDEVs) which contain the backup or duplicate data. During normal URz operations, the primary data volume remains available to all hosts at all times for read and write I/O operations. During normal URz operations, the secondary subsystem 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 3.22.2.4).

3.22.1.5 Journal Volume

When URz is used, updates to primary data volumes can be stored in other volumes, which are called journal volumes. The updates (which is 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 subsystem and the secondary subsystem fails temporarily, remote copy operations can continue after the communication path is recovered.
- If data transfer from hosts to the primary subsystem is temporarily faster than data transfer between the primary subsystem and the secondary subsystem, remote copy operations between the primary subsystem and the secondary subsystem 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 subsystem is faster for a relatively long period of time than data transfer between the primary subsystem and the secondary subsystem.

(1) The Number of Journal Volumes

One journal group can contain up to 16 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.

(2) Specifications of Journal Volumes

■ Types of logical units (LUs):

The following DKU emulation types are allowed for journal volumes:

Table 3.22.1.5-1 Emulation Types for Journal Volumes

Emulation Category	Supported Emulation Types
DKU (drive)	• OPEN-V • All mainframe volumes that can be used with TagmaStore™ USP Note: Status of mainframe volumes cannot be referenced.

■ 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 TagmaStore™ USP. Journal volumes also support all physical volumes that are supported by TagmaStore™ USP.

■ Support for option programs:

Cache Residency Manager volumes can be used for journal volumes.

Caution: Volumes containing a VMA (volume management area) cannot be used as journal volumes.

(3) Restrictions on Journal Volumes

■ 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 deleted)
- When all data volume pairs in the journal group are suspended.
- When processing for changing the status of a data volume pair (for example, deletion 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 (4)) 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 panel.

■ 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 deleted)
- When all data volume pairs in the journal group are suspended.

You can use Storage Navigator computers to delete journal volumes.

Caution: After you delete a journal volume that is a mainframe volume, you must format the volume (LDEV) by using Virtual LVI/LUN (VLL). Unless you format the volume, data in the volume will not be guaranteed.

■ 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.

(4) 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 3.22.3.1-1. The journal data that the metadata manages is stored in the journal data area.

3.22.1.6 Remote Copy Connections

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

When fibre-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 longwave interface connections are required. When fibre-channel interface (single-mode longwave) connections are used, two switches are required for distances greater than 10 km (6.2 miles), and distances up to 30 km (18.6 miles) are supported.

3.22.1.7 Initiator Ports and RCU Target Ports

The initiator port and the RCU target port are required at both the primary subsystem and secondary subsystem. The initiator port at the primary subsystem is connected to the RCU target port at the secondary subsystem via the fibre channel interface. The initiator port at the secondary subsystem is connected to the RCU target port at the primary subsystem. The initiator port at the secondary subsystem issues a “read journal” command to the primary subsystem, and then the RCU target port at the primary subsystem sends journal data to the secondary subsystem in response to the “read journal” command.

Any fibre-channel interface port of the TagmaStore™ USP 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 fibre-channel interface port other than the initiator port.

3.22.1.8 URz Remote Console Software

TagmaStore™ USP Storage Navigator Java applet program product includes URz for the TagmaStore™ USP subsystem. The TagmaStore™ USP Storage Navigator software communicates with the SVP of each TagmaStore™ USP subsystem via defined TCP/IP connections.

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

3.22.1.9 Host I/O Time-Stamping Function

If you plan to establish URz journal groups, the I/O time-stamping function must be installed on the host processor at the primary site. The I/O time-stamp, which is provided by MVS DFSMSdfp, is the same time-stamp that is used by IBM® XRC pairs. The I/O time-stamping function should also be installed on the host processor at the secondary site, so that time-stamps can be used when copying data in the reverse direction.

Note: If the system at the primary and/or secondary site consists of several CPU complexes, a SYSPLEX timer is required to provide a common time reference for the I/O time-stamping function.

3.22.1.10 Error Reporting Communications (ERC)

Error reporting communications (ERC), which transfers information between host processors at the primary and secondary sites, is a critical component of any disaster recovery effort. You can configure ERC using channel-to-channel communications, NetView technology, or other interconnect technologies, depending on your installation requirements and standards. Neither URz nor the URz remote console software provides ERC between the primary and secondary sites.

When URz is used as a data migration tool, ERC is recommended but is not required. When URz is used as a disaster recovery tool, ERC is required to ensure effective disaster recovery operations. When a URz pair is suspended due to an error condition, the primary subsystem generates sense information which results in an IEA491E system console message. This information should be transferred to the primary site via the ERC for effective disaster detection and recovery.

3.22.2 Remote Copy Operations

Figure 3.22-3 illustrates the two types of URz remote copy operations: initial copy and update copy.




Figure 3.22-3 Remote copy operations

This section describes the following topics that are related to remote copy operations with URz:

- Initial copy operation (see section 3.22.2.1)
- Update copy operation (see section 3.22.2.2)
- Read and write I/O operations for URz volumes (see section 3.22.2.3)
- Secondary data volume write option (see section 3.22.2.4)
- Secondary data volume read option (see section 3.22.2.5)
- Difference management (see section 3.22.2.6)

3.22.2.1 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 subsystem to the secondary subsystem, and then restores the base-journal to the secondary data volume.

If the journal-obtain operation starts at the primary data volume, the primary subsystem 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 subsystem to the secondary subsystem after the secondary subsystem issues a read-journal command. After a base-journal is copied to the secondary subsystem, 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 subsystem is stored to the secondary subsystem journal volume according to the read command from the secondary subsystem. 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.

3.22.2.2 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 subsystem, and then restored to the secondary data volume.

The primary subsystem 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 subsystem issues “read journal” commands, update journals will be copied from the primary subsystem to the secondary subsystem asynchronously with completion of write I/Os by the host. Update journals that are copied to the secondary subsystem will be stored in journal volumes in the journal group that the secondary data volume belongs to. The secondary subsystem 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.

3.22.2.3 Read and Write I/O Operations During URz Volumes

When a primary subsystem receives a read I/O for a URz primary data volume, the primary subsystem 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 subsystem does not read the URz secondary data volume for recovery.

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

The primary subsystem completes the primary data volume write operations independently of the update copy operations at the secondary data volume. The secondary subsystem 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 subsystem reports a unit check and does not create the journal data for this operation. If the update copy operation fails, the secondary subsystem suspends either the affected pair or all URz pairs in the journal group, depending on the type of failure. When the suspended URz pair or journal group is resumed (Resume Pair), the primary subsystem and secondary subsystem negotiate the resynchronization of the pair(s).

During normal URz operations, the secondary subsystem does not allow URz secondary data volumes to be online (mounted), and therefore hosts cannot read from and write to secondary data volumes. The URz secondary data volume write enable option allows write access to a secondary data volume while the pair is split (see section 3.22.2.4). The secondary data volume write option can only be enabled when you split the pair from the primary subsystem.

3.22.2.4 Secondary Data Volume Write Option

For additional flexibility, URz provides a secondary data volume write option (S-Vol. Write) which enables write I/O to the secondary data volume of a split URz pair. The secondary data volume write option can be selected by the user during the Suspend Pair 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 subsystem. When you resync a split URz pair which has the secondary data volume write option enabled, the secondary subsystem sends the secondary data volume track bitmap to the primary subsystem, and the primary subsystem 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.

3.22.2.5 Secondary Data Volume Read Option

For additional flexibility, URz offers a special secondary data volume read option. The Hitachi representative enables the secondary data volume read option on the secondary subsystem (mode 20). The secondary data volume read option allows you to read a URz secondary data volume only while the pair is suspended, that is, without having to delete the pair. The secondary subsystem will allow you to change only the VOLSER of the suspended secondary data volume, so that the secondary data volume can be online to the same host as the primary data volume while the pair is suspended. All other write I/Os will be rejected by the secondary subsystem. The primary subsystem copies the VOLSER of the primary data volume back onto the secondary data volume when the pair is resumed. When the secondary data volume read option is not enabled and/or the pair is not suspended, the secondary subsystem rejects all read and write I/Os to a URz secondary data volume.

3.22.2.6 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 (Resume Pair), the primary subsystem merges the primary data volume and secondary data volume bitmaps, and the differential data is copied to the secondary data volume.

3.22.3 Journal Processing

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

- Creating and storing journals at the primary subsystem (see section 3.22.3.1)
- Copying journals to the secondary subsystem (see section 3.22.3.2)
- Storing journals at the secondary subsystem (see section 3.22.3.3)
- Selecting and restoring journals at the secondary subsystem (see section 3.22.3.4)
- Types of journals (see section 3.22.3.5)

3.22.3.1 Creating and Storing Journals at the Primary Subsystem

When a primary subsystem performs an update (host-requested write I/O) on a URz primary data volume, the primary subsystem creates a journal data to be transferred to secondary subsystem. 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 3.22.3.1-1). When base-journal is obtained, only metadata information is created and stored in UR cache or the journal volume.

Table 3.22.3.1-1 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 subsystem has created for each journal group. The journal data is transferred to the secondary subsystem asynchronously with the host I/O. The secondary subsystem updates the secondary data volume in the same order as the primary data volume according to the sequence number information in the journal.

3.22.3.2 Copying Journals to the Secondary Subsystem

When a primary subsystem receives a read journal command from a secondary subsystem, the primary subsystem sends the journal data to the secondary subsystem. The secondary subsystem's initiator ports act as host processor channels and issue special I/O operations, called remote I/Os (RIOs), to the primary subsystem. The RIO transfers the journal data in FBA format using a single channel command. The primary subsystem 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 subsystem in a different order than the journal data were created at the primary subsystem. The secondary subsystem 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 subsystem-to-secondary subsystem link resources.

3.22.3.3 Storing Journal at the Secondary Subsystem

A secondary subsystem receives the journal data that is transferred from a primary subsystem 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 subsystem 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 subsystem. This is true even if the primary subsystem and secondary subsystem are connected via a channel extender product.

3.22.3.4 Selecting and Restoring Journal at the Secondary Subsystem

The secondary subsystem selects journal data to be promoted to formal data (or “restored”) as follows:

1. The secondary subsystem 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 subsystem. If the number is 1, the journal data arrived at the secondary subsystem. If the number is 0, the journal data has not arrived yet. The secondary subsystem determines whether the journal data should be settled or not according to this number. If the journal data has not arrived yet, the secondary subsystem waits for the journal data.
2. When the top of queue in the journal group indicates the journal data arrival, the secondary subsystem selects the journal data which has the lowest sequence number, and then settles this journal data.
3. The secondary subsystem repeats steps 1. and 2. to select and settle the journal data.

Figure 3.22-4 illustrates the journal data selection and settling at the secondary subsystem. This diagram shows that journal data S1 arrives at the secondary subsystem because the management information indicates 1. The secondary subsystem 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 subsystem waits journal data S2. When journal data S2 arrives, the secondary subsystem selects S2 as the next journal data to be settled. The journal data selected by the secondary subsystem is marked as “host-dirty” and treated as formal data.




Figure 3.22-4 Selecting and Settling Journal at the Secondary Subsystem

The secondary subsystem 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.

3.22.3.5 Types of Journal

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

3.22.4 UR operation

3.22.4.1 Pair operation

The following figure illustrates an UR pair configuration. In the configuration, UR pairs belong to the journal group. Each journal group and UR pair have an attribute and status. Each attribute and status is described in the following subsections.




Figure 3.22-5 Relationship between pairs and journal groups

(1) Journal group

Prior to pair creation, the journal volume needs to be registered, and the journal group needs to be defined. You can delete journal groups that are not in use. To define journal groups, you need to register journal volumes. On the other hand, to delete journal groups, you need to delete all journal volumes.

Table 3.22.4.1-1 Journal group attributes

#	Journal group attribute	Meaning
1	Initial	No UR pair is set in the journal group to which journal volumes are registered.
2	Master	The journal group to which the logical volume storing the original data (PVOL) belongs. The attribute of the journal group in Master is called “M journal”.
3	Restore	The journal group to which the logical volume storing the duplicated data (SVOL) belongs. The attribute of the journal group in Restore is called “R journal”.

Table 3.22.4.1-2 Journal group status

#	Journal group status	Meaning
1	–	UR pairs are not set in the target group
2	Active	Base/update copy is in progress in the target group
3	Halting	Base/update copy is halted in the target group
4	Stopping	The target group is being suspended, or being deleted
5	Stop	The target group is suspended. (All UR pairs in the target group are suspended)

Table 3.22.4.1-3 Journal group operations

#	Operation	Specified attribute	Description
1	Register journal volume	MJNL/RJNL	✓ Register a logical volume, to which no path is defined, as a JNL volume. ✓ Can register the additional volume to the existing journal group only when the UR remote copy is stopped (Journal group status is STOP). ✓ When the journal group state is ACTIVE, the journal volume can be registered since ver.50-05-00-00. (However, only the journal area increases, and after the journal group suspend and resync, the meta data area increases). ✓ Can register UR pair to the journal group to which the journal volume is registered.
2	Delete journal volume	MJNL/RJNL	✓ Delete the logical volume registered as a journal volume from the journal volume. ✓ Can delete it only when UR remote copy is stopped (Journal group status is STOP or the group attribute is Initial) if there are two or more journal volumes in the target journal group. ✓ Can delete it only when the target group does not have any UR pair (Group attribute is Initial) if the number of journal volumes in the target journal group is one ✓ No UR pair can be registered to such a journal group from which all journal volumes are deleted. ✓ When the journal volume of MF emulation type is deleted, a volume concerned blockades because LDEV format is necessary since version 50-06-00-00. When the same volume is re-used as a journal volume, format is not necessary for automatic blockage restoration.

(2) UR pairs

UR performs a remote copy operation for the logical volume pair set by the user. Based on pair operations, pair attributes and pair statuses are added to the logical volumes. You can perform the following remote copy operations for UR pairs. The figure below illustrates how the UR pair status changes due to each operation.

Table 3.22.4.1-4 Pair attributes


#	Pair attribute	Description
1	SMPL	Target volume is not assigned to an UR pair.
2	PVOL	Primary volume. Data volume to which the original data is stored.
3	SVOL	Secondary volume. Data volume to which backup or duplicated data is stored.

Table 3.22.4.1-5 Pair status

#	Pair status	Description
1	SMPL	Target volume is not assigned to an UR pair.
2	COPY	Base copy is in progress and data of the PVOL and SVOL of the UR pair do not match completely. When their data match, the status changes to PAIR.
3	PAIR	Base copy is completed, and data of the PVOL and SVOL match completely.
4	PSUS/SSUS	Copy operation is suspended in the UR pair.
5	PSUE	An error is detected in the DKC, and the copy in the UR pair is stopped (Suspended).
6	PFUL	The journal usage in the journal volume exceeded the threshold. The copy operation is continued.
7	PFUS	The capacity of the stored journal exceeded the journal volume capacity, and the copy operation is suspended in DKC.
8	SSWS	Data can be written to the SVOL in which Takeover is in progress.
9	Suspending	The status of the UR pair is being changed to Suspend.
10	Deleting	The UR pair is deleted, and the status is being changed to SMPL.

Table 3.22.4.1-6 Pair operations

#	Operation	Specified attribute	Description
1	Pair create	MJNL	Register the logical volume to which a path is defined as an UR pair. There are two types of copy instruction, "All copy" and "NO copy". "All copy" performs a base copy, and "NO copy" does not perform a base copy.
2	Pair suspend	MJNL/RJNL	Change the status of the UR pair, which is performing the base/update copy, to the suspend status.
3	Pair delete	MJNL/RJNL	Delete the already registered UR pair.
4	Pair resync	MJNL/RJNL	Change the pair status of the UR, in which the copy operation is suspended, to the pair resume status. (RJNL can be specified only when swapping is specified)
5	Takeover	RJNL	Swap MJNL and RJNL (reverse the source and the target) and resync the pair.



(#x) : Correspond to the pair operation # in Table 3.22.4.1-6.

——→ : Status transition due to pair operation

-----→ : Status transition in DKC

△ : Status transition is in progress. DKC internal processing is in progress.

○ : Normal status. Status is changed to this when operation completes.

Figure 3.22-6 Pair attributes

(3) Pair create

Logical volume to which a path is defined is registered to the target journal group as an UR pair (data volume). The specification of a data volume and UR pair in a journal group is as follows.

Table 3.22.4.1-7 Pair create

#	Item	Mainframe	Open platform
1	Pair create timing	✓ Can register the pair only when the specified journal group exists, and no Suspending/Deleting pair is included in that journal group	✓ Can register the pair only when the specified journal group exists, and no Suspending/Deleting pair is included in that journal group

● Pair create option

When creating a pair, you can set a pair create option. The following table shows the options and applications that can be specified.

Table 3.22.4.1-8 Pair create options

#	Option	Feature overview	OP	MF	RM	BC
1	Copy type	You can choose to or not to perform a base copy. The following copy types are available. ✓ All copy: Perform base copy when creating a pair ✓ No copy: Not perform base copy when creating a pair	○	○	○	○
2	Priority	Specify which pair you want to perform a base copy first when creating multiple pairs at a time. ✓ Setting: 1-255	—	○	—	—
3	Error level	Set error levels. You can choose either of the following levels. ✓ Group: Even if an error which affects only the specified volume occurs, all pairs in the journal group will suspend due to the error. ✓ Volume: When an error which affects only the specified volume occurs, only that pair suspends due to the error. However, when an error which affects the entire journal group occurs, all pairs in the journal group are suspended.	○	○	—	○
4	SVOL ONLINE Check	Check whether SVOL is ONLINE or not, and if it is ONLINE, you cannot create a pair. You can choose either of the following. ✓ No Check: Create a pair without checking whether SVOL is ONLINE ✓ SVOL ONLINE: Check if SVOL is ONLINE.	—	—	—	○
5	CFW data	Choose whether MCU copies the CFW data to SVOL. You can choose either of the following. ✓ Copy to SVOL: MCU copies the CFW data to SVOL. ✓ PVOL only: MCU does not copy the CFW data to SVOL.	—	○	—	—

Legend) OP: Instruction from StorageNavigator (OPEN), MF: Instruction from StorageNavigator (MF)
RM: Instruction from RaidManager, BC: Instruction from BC Manager

(4) Pair suspend

The copy operation of an UR pair is stopped, and the pair status is changed to PSUS. When all UR pairs in a journal group are suspended, the status of the journal group is changed to STOP. The specifications of the suspend operation are as follows.

- ✓ This operation is performed when the target volume is in the COPY/PAIR status and all volumes in the journal group are in the status other than Suspending/Deleting status.
- ✓ The pair suspend operation can be performed from PVOL/SVOL. The processing of the suspend operation is the same for both PVOL/SVOL instructions.
- ✓ When you perform the pair suspend operation, you can specify the Pend Update and the suspend range. The table below shows the relationships of Pend Update and the suspend range.

Table 3.22.4.1-9 Pair suspend

#	Pend Update	Feature overview	Volume	Group
1	Flush	✓ When the suspend operation is received, the pending data in MCU/RCU is reflected. ✓ When the operation is performed while the host I/O is being processed, the contents of the PVOL and SVOP are not the same. (In the operation after the host I/O is stopped, the contents of the PVOL and SVOL are the same)	○	○
2	Purge	✓ When the suspend operation is received, the difference of the pending data in MCU/RCU is recorded to the differential bitmaps. (Since the data is not reflected, the contents of the PVOL and SVOL are not the same) ✓ The status can be changed to Suspend in a short time.	—	○

● Pair suspend option

When you perform a pair suspend operation, you can specify a pair suspend option. The following table shows the options and applications you can specify.

Table 3.22.4.1-10 Pair suspend option



#	Option	Feature overview	OP	MF	RM	BC
1	Pend Update	You can choose a Pend Update type when performing a pair suspend. (Note that only Flush is available for RaidManager) ✓ Flush: The Pend data is processed in the Flush mode. ✓ Purge: The Pend data is processed in the Purge mode.	○	○	–	○
2	Range	You can specify the suspend range. (However, when the Purge mode is selected in Pend Update, only “Group” can be selected) ✓ Volume: Only the specified UR pairs are suspended. ✓ Group: All pairs in the specified journal group are suspended.	○	○	○	○
3	SVOL write	You can choose to perform Write/Read to the SVOL when performing a pair suspend. ✓ Disable: Cannot Read/Write data to the suspended SVOL. ✓ Enable: Can Read/Write data to the suspended SVOL.	○	–	○	○

(5) Pair delete

The UR pair is deleted, and the pair volumes are restored to non-UR logical volumes. When all UR pairs in the target journal group are deleted, the attribute of the journal group is changed to Initial. The following table shows the specifications of the pair delete operation.

- ✓ This operation is performed when the target volume is in the normal status (COPY/PAIR/PSUS(PSUE)) and all volumes in the journal group are in the status other than Suspending/Deleting.
- ✓ If you perform the pair delete operation for an UR pair after stopping the host I/O, the data in the PVOL and that in the SVOL are the same. (When this operation is performed while the host I/O is in progress, the data are not the same.)
- ✓ The pair delete operation can be performed from PVOL/SVOL. Note that the feature depends on the specified volume attribute. The following table shows the specified volume attributes and the features.

Table 3.22.4.1-11 Pair delete

#	Specified volume	Feature overview
1	PVOL specified	✓ The statuses of both PVOL and SVOL are changed to SMPL, and the UR pair will be deleted. 
2	SVOL specified	✓ Only the status of SVOL changes to SMPL, and the PVOL status is not changed. If the PVOL is in the COPY/PAIR status, the pair will be suspended due to a failure (PSUE). When the operation is performed to the suspended pair, the PVOL status is not changed.  ✓ The UR pair of only the PVOL after the pair delete operation cannot be resynchronized. Therefore, you need to delete the pair using the PVOL-specified pair delete operation.

● Pair delete option

When you perform a pair delete operation, you can specify a pair delete option. The following table shows the options and applications you can specify.

Table 3.22.4.1-12 Pair delete option

#	Option	Feature overview	OP	MF	RM	BC
1	Range	You can specify the following pair delete range. ✓ Volume: Only the specified UR pairs are deleted. ✓ Group: All pairs in the specified journal group are deleted.	○	○	○	○
2	Specified volume attribute	You can specify the attribute of the volume to which the pair delete is performed. PVOL specified: Normal pair delete operation is performed. SVOL specified: Pair is deleted only in SVOL.	○	○	○	—
3	ForceDelete	You can specify ForceDelete which deletes pairs regardless of the pair status. The overview of the ForceDelete feature is as follows. ✓ ForceDelete deletes only the pairs in the specified volume without recognizing the status of the paired volume. Therefore, the operations in the target (paired) volume is not guaranteed. ✓ Only “Group” range is available. (Note that only “Volume” range is available for SMPL volumes) ✓ After the ForceDelete operation, the data in PVOL and that in SVOL are not the same.	○	○	—	—

(6) Pair resync (Pair resume)

The copy operation of the suspended UR pair is resumed. When a copy operation of one pair is resumed in the journal group in which the copy operation is stopped (STOP), the journal group status will be changed to Active. The specifications of pair resync operation are as follows.

- ✓ This operation is performed when the target UR pair is in the PSUS/PSUE status and the journal group to which the target UR belongs does not contain UR pair in the Suspending/Deleting status. However, this operation cannot be performed to the PVOL which is paired with the SVOL in the SSWS status.
- ✓ The operations can be specified in PVOL. However, the Swap instruction can be specified in SVOL optionally.

● Pair resync option

When you perform a pair resync operation, you can specify a pair resync option. In addition, when the range of the operation is “volume”, you can change the pair create option for the target volume. The following table shows the options and applications you can specify.

Table 3.22.4.1-13 Pair resync option

#	Option	Feature overview	OP	MF	RM	BC
1	Range	You can specify the following pair resync range. ✓ Volume: Only the specified UR pairs are resynchronized. ✓ Group: All pairs in the specified journal group are resynchronized.	○	○	○	○
2	Priority	To resynchronize multiple pairs at a time, set the priority of the base copy operation. ✓ Range: 1-255	—	○	—	—
3	Error level	Set the error level in the case of a failure. (Enabled only when the range “volume”) ✓ Group: Even if a failure affects only the specified volume, all pairs in the journal group will be suspended due to the failure. ✓ Volume: When a failure affects only the specified volume, only the pair will be suspended due to the failure. However, when a failure affects the entire journal group, all pairs in the journal group will be suspended due to the failure.	○	○	—	○
4	Swap	This option is specified in the SVOL, the volumes are swapped (SVOL→PVOL/PVOL→SVOL). The differential data is resynchronized from the swapped PVOL to SVOL. See “3.22.4.1(7) takeover” for the swapping feature”.	—	—	○	○

(7) Takeover (Reverse Resync)

Takeover swaps PVOL and SVOL (switch the source and the target) and resynchronizes the pair. The specifications of the Takeover operation are as follows.

- ✓ Only RaidManager can perform this operation.
- ✓ This operation is performed to SVOL. It can be performed even if the PVOL (i.e. DKC in the MCU) is lost due to disaster etc.
- ✓ This operation is performed when the pair status of the SVOL is PAIR or PSUS/SSWS and when the target RCU does not contain UR pairs in the Suspending/Deleting or COPY status. The following table shows when the Takeover command is issued.

Table 3.22.4.1-14 Takeover

#	Pair status	Feature overview
1	PAIR	(1) When a Takeover command is issued to the RCU, the Flush suspend is performed for the specified group. When it is issued, the SVOL status changes to SSWS, and it becomes a volume that can be read/written. (2) After the pair is suspended, a pair resync swap is performed to swap the PVOL and the SVOL, and an initial copy operation is performed.
2	PSUS/PSUE	(1) When a Takeover command is issued to RCU, the Flush suspend is not performed for the specified group. Only the SVOL status is changed to SSWS. (2) After the suspend operation completes, a pair resync swap is performed to swap the PVOL and the SVOL, and an initial copy is performed.

(8) JNL group status and display of pair status by Raid Manager

We summarize the relationship between display of JNL group status and display of pair status regarding Raid Manager(RM) compared with JAVA display as the following.

① Operation suspend

#	JAVA Pair status	RM Pair status	RM (M-JNL) JNLGroup status	RM (R-JNL) JNLGroup status	Remarks
1	SMPL	SMPL	—	—	—
2	COPY/PAIR	COPY/PAIR	PJNN	SJNN	—
3	Suspending				—
4	PSUS	PSUS/SSUS	PJSN	SJSN	—

② JNL utilization threshold over (JNL utilization is over 80% (but...not failure))

#	JAVA Pair status	RM Pair status	RM (M-JNL) JNLGroup status	RM (R-JNL) JNLGroup status	Remarks
1	SMPL	SMPL	—	—	—
2	COPY/PAIR	PFUL	PJNF	SJNF	—
3	Suspending	—	—	—	Pair status will not change to “Suspend”.
4	PSUS	PSUS/SSUS	—	—	Same as above

③ Failure suspend (except for JNL puncture)

#	JAVA Pair status	RM Pair status	RM (M-JNL) JNLGroup status	RM (R-JNL) JNLGroup status	Remarks
1	SMPL	SMPL	—	—	—
2	COPY/PAIR	COPY/PAIR	PJNN	SJNN	—
3	Suspending	PSUE	PJSE	SJSE	—
4	PSUE	PSUE			—

④ JNL puncture suspend

#	JAVA Pair status	RM Pair status	RM (M-JNL) JNLGroup status	RM (R-JNL) JNLGroup status	Remarks
1	SMPL	SMPL	—	—	—
2	COPY/PAIR	COPY/PAIR	PJNN	SJNN	—
3	Suspending	PFUS	PJSF	SJSF	—
4	PSUS	PFUS			—

3.22.4.2 USAGE/HISTORY

(1) USAGE

This feature enables you to view the UR information (Frequency of I/O for UR pair, transfer rate of journals between MCU/RCU, usage of journals in MCU/RCU etc.) using SVP and WEBCONSOLE. The specifications are as follows.

Table 3.22.4.2-1 USAGE specification

#	Item	Specification etc.
1	Sampling interval	1-15 min. (1 min.)
2	Samplings	1440
3	Unit of sampling/display	• For each LU • For each JNL group • For each system
4	Sampling item	See Table 3.22.4.2-2
5	Others	You can save the monitor data in the text format using the Export Tool of Performance Monitor.

Table 3.22.4.2-2 Sampled items

#		Category	Statistics	Description
1	MCU	Host I/O	Read Record Count	The number of Read I/Os per second
2			Read Hit Record Count	The number of Read Hit I/Os per second
3			Write Record Count	The number of Write I/Os per second
4			Write Hit Record Count	The number of Write Hit I/Os per second
5			Read Transfer Rate	The amount of data that are read per second (KB/sec.)
6			Write Transfer Rate	The amount of data that are written per second (KB/sec.)
7		Initial copy	Initial copy HIT rate	Initial copy hit rate (%)
8			Average Transfer Rate	The average transfer rate for initial copy operations (KB/sec.)
9	MCU	Async copy	M-JNL Asynchronous RIO count	Number of async RIOs per second in MCU
10			M-JNL Total Number of Journal	The number of journals at MCU
11			M-JNL Average Transfer Rate	The average transfer rate for journals at MCU (KB/sec.)
12			M-JNL Average RIO Response	The remote I/O process time at MCU (milliseconds)
13	RCU	Async copy	R-JNL Asynchronous RIO count	The number of asynchronous remote I/Os per second at the RCU
14			R-JNL Total Number of Journal	The number of journals at the RCU
15			R-JNL Average Transfer Rate	The average transfer rate for journals at RCU (KB/sec.)
16			R-JNL Average RIO Response	The remote I/O process time at RCU (milliseconds)
17	MCU	Journal group	Data Used Rate	Data usage rate for journals at MCU (%)
18			Meta Data Used Rate	Metadata usage rate for journals at MCU (%)
19	RCU	Journal group	Data Used Rate	Data usage rate for journals at RCU (%)
20			Meta Data Used Rate	Metadata usage rate for journals at RCU (%)

(2) HISTORY

This feature enables you to view the history of operations for data volume pairs (Operations performed in the past) using SVP and WEBCONSOLE. The specifications are as follows.

Table 3.22.4.2-3 HISTORY specifications

#	Item	Spec etc.
1	Displayed info	Date and time of the operation, contents of the operation (See #3, Journal group number, Mirror ID, Data volume, Target volume)
2	Samplings	65535, or for one week
3	Operation	• Pair create • Pair delete • Pair recovery • Pair split etc.
4	Others	The snapshot function enables you to save operation history to a text file.

3.22.4.3 Option

The following table shows UR settings. For pair settings, see 3.22.4.1. The following system option and journal group options are supported.

Table 3.22.4.3-1 List of options

#	Affected	Setting	Description	Supported version
1	System	Max number of initial copy VOLs	✓ Number of volumes to which initial copy operations are performed at a time. The setting value is 1-128. ✓ Default is 64.	
2	Journal group	Inflow Control	✓ Indicates whether to restrict inflow of update I/Os to the PVOL (whether to delay a response to the hosts). ✓ If you select “Yes”, the inflow control is performed according to the condition of #3.	
3		Data Overflow Watch	✓ Indicates the time (in seconds) for monitoring whether metadata and journal data are of journal volume are full. ✓ When either of the areas becomes full, the target group will be suspended due to a failure after this specified time. ✓ The setting value is 0-600 sec.	
4		Copy pace	✓ Initial copy pace (speed). ✓ The setting value is “Low”, “Medium” or “High”.	
5		Path Watch time	✓ Time for monitoring a path from the time it is blocked until it is suspended due to a failure. ✓ The setting value is 1-60 min. ✓ You can 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.	50-04-28-00/00 (v04+1)
6		Forward Path Watch time	✓ If you select “Yes”, the Path Watch time (#5) value of the master journal group will be forwarded to the restore journal group.	50-04-28-00/00 (v04+1)
7		Use of Cache	✓ If you select “USE”, journal data will be stored into the cache and don't destage journal volumes in the restore journal group as long as the cache memory is not filled. The JNL restore performance improves.	50-04-28-00/00 (v04+1)
8		Speed of Line	✓ JNL copy performance pace (speed). ✓ The setting value is “256Mbps”, “100Mbps” or “50Mbps”.	50-04-28-00/00 (v04+1)

Table 3.22.4.3-2 Link failure monitoring mode

Mode	Feature digest	Feature overview
448	Mode for suspending the pair immediately after UR path is disconnected	Turn this mode ON when you want to suspend the pair due to a failure immediately after a communication failure is detected between UR M-R. ON: In MCU, the pair is suspended due to a failure immediately after the RDJNL from RCU is stopped. In RCU, the pair is suspended due to a failure immediately after the RDJNL fails. OFF: In MCU, the pair is suspended due to a failure when the RDJNL from RCU is stopped for a certain period of time. In RCU, the pair is suspended due to a failure when the RDJNL fails for a certain period of time. This mode runs only when Mode 449 is OFF.
449	Mode for prohibiting UR path watch	Turn this mode ON when you want to prevent communication failures between UR M-R from being detected. On is default since ver.50-05-00-00. ON: In MCU, the RDJNL stop check from RCU is prevented. In RCU, the RDJNL failure monitoring is prevented. OFF: Communication failures between M-R are detected.

3.22.5 Maintenance features and procedure

3.22.5.1 Maintenance

The following table shows limitations/impact on the maintenance operations regarding the UR operations. Other limitations/impact are based on those in RAID500.

Table 3.22.5.1-1 Limitations/restrictions on maintenance

#	Item	Limitation/impact	Note
1	CM/SM replacement in MCU	Pair in the initial copy status will suspend due to a failure.	
2	Microprogram exchange in RCU	The pair may be suspended in MCU due to a failure. (Cause: Link failure, journal full)	Alternate paths are set between MCU and RCU. If microprogram is not replaced using these alternate paths at a time, the pair will not be suspended due to a link failure.
3	CHT/CHI replacement in MCU/RCU	The pair may be suspended in MCU/RCU due to a failure. (Cause: Link failure, journal full)	Alternate paths are set between MCU and RCU. If microprogram is not replaced using these alternate paths at a time, the pair will not be suspended due to a link failure.

3.22.5.2 PS OFF/ON Process

Before you power off/on (PS OFF/ON) the DKC, we recommend that you suspend all pairs (all JNL groups) by specifying Flush in advance as shown below. When you stop the host I/O and suspend (Flush) the pairs in advance, the contents of PVOL and SVOL match. Consequently, the operation can be continued using the SVOL data in R-DKC even if you fail to power on the M-DKC for some reason. The recommended procedure is as follows.

- (1) Stop the host I/O.
- (2) Issue the suspend (Flush) request for M-JNL, and change the statuses of both M-JNL and R-JNL to suspend.
- (3) Power off the M-DKC.
- (4) Power off the R-DKC.
- (5) Power on the R-DKC.
- (6) Power on the M-DKC.
- (7) Resynchronize the pair upon the pair resync request.
- (8) Resume the host I/O.

If you do not perform the procedure above, operations are performed as follows. The numbers in the table below show the target of PS OFF/ON, and the order of PS OFF/ON. “-” shows that it is not the target of PS OFF/ON.

Table 3.22.5.2-1 PS OFF/ON

	Part and Order of PS OFF/ON				Operation
	M-DKC		R-DKC		
	OFF	ON	OFF	ON	
Case 1	1	2	—	—	[PS OFF/ON only M-DKC] No problem.
Case 2	—	—	1	2	[PS OFF/ON only R-DKC] Pair may be suspended due to a failure in MCU. As a result, R-JNL detects the failure suspension when PS ON is performed, and the pair may be also suspended. Such a failure suspension is caused because the M-JNL is full or the read journal is stopped.
Case 3	1	3	2	4	[PS OFF/ON both M-DKC and R-DKC] PS ON the M-DKC first. If it takes time until R-DKC PS ON (4) after M-DKC PS ON (3), the same phenomenon as Case 2 will occur.
Case 4	1	4	2	3	[PS OFF/ON both M-DKC and R-DKC] The procedure is the recommended one, except that host I/O stop and pair suspension are excluded. When PS ON is performed, the pair status before PS OFF is maintained.
Case 5	2	3	1	4	[PS OFF/ON both M-DKC and R-DKC] PS OFF R-DKC first. If it takes time until M-DKC PS OFF (2) after R-DKC PS OFF (1), the same phenomenon as Case 2 will occur. If it takes time until R-DKC PS ON(4) after M-DKC PS ON (3), the same phenomenon as Case 2 will occur.
Case 6	2	4	1	3	[PS OFF/ON both M-DKC and R-DKC] If it takes time until M-DKC PS OFF (2) after R-DKC PS OFF (1), the same phenomenon as Case 2 will occur.

3.22.5.3 Power failure

When a power failure occurs, UR pairs will be in the following status. The status change is the same in both the memory backup mode and the destage mode. The destage mode is supported in the 2005/7E version.

Table 3.22.5.3-1 Power failures

#	Item	Pair in PS ON	Recovery
1	SM/CM non-volatile	Suspended due to a failure	Resynchronize to the target group
2	SM/CM volatile	Journal group information and pair information will be lost.	Register journal group and pair

3.22.6 Cautions on software

3.22.6.1 Error recovery

The following subsections describes failure detection, error report, and recovery in UR. When a failure occurs in UR, Group or Volume (according to the error level setting) will be suspended.

(1) Error report

The following table shows the error report in the case of a failure.

Table 3.22.6.1-1 Error report

Item	Type	SSB(F/M)	SIM
Path	Link failure	–	0x2180-XX
Pair	Failure suspend	F/M=0xFB Not reported to host	Serious SIM Not reported to host

(2) Failure detection and recovery

Failure detection and recovery in MCU are described in the following table.

Table 3.22.6.1-2 Failure detection and recovery

Part	Description	Error	Recovery
Path failure	Link failure of MCU->RCU	SIM=0x2180	Recover the link
	Link failure of MCU<-RCU	When Read JNL is stopped for a certain period of time, the target Volume/Group is suspended, and F/M=0xFB, SIM is generated. (see Table 3.22.4.3-1 and Table 3.22.4.3-2.)	Recover the link Resync after the failure recovered
Data volume failure	PDEV failure	Target Volume/Group is suspended F/M=0xFB, SIM Update I/O is managed as differential data	Resync after the failure recovered
Journal failure	PDEV failure	Target Group is suspended F/M=0xFB, SIM Update I/O is managed as differential data	Resync after the failure recovered
Journal full	Metadata area or journal data area of JNL volume is not sufficient for a certain period of time	Target Group is suspended F/M=0xFB, SIM Update I/O is managed as differential data	Resync after the failure recovered

(3) Failure detection in RCU

Failure detection and recovery in RCU are described in the following table.

Table 3.22.6.1-3 Failure detection in RCU

Part	Description	Error	Recovery
Path failure	Link failure of MCU<-RCU	SIM=0x2180When Read JNL is stopped for a certain period of time, the target Volume/Group is suspended, and F/M=0xFB, SIM is generated. (see Table 3.22.4.3-1 and Table 3.22.4.3-2.)	Recover the link Resync after the failure recovered
Data volume failure	PDEV failure	Target Volume/Group is suspended F/M=0xFB, SIM	Resync after the failure recovered
Journal failure	PDEV failure	Target Group is suspended F/M=0xFB, SIM	Resync after the failure recovered

3.22.7 Disaster Recovery Operations

3.22.7.1 Preparation for Disaster Recovery

The type of disaster and the status of the URz 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 user should anticipate and plan for all types of failures and disasters.

The major steps in preparing for disaster recovery are:

1. Identify the journal groups and data volumes that contain important files and data (e.g. DB2 log files, master catalogs, key user catalogs, and system control datasets) for disaster recovery.
2. Install the Storage Navigator PC and URz 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 error reporting communications (ERC) between the primary and secondary sites.

3.22.7.2 Sense information is transferred between sites

When the primary subsystem (or secondary subsystem for URz) suspends a URz pair due to an error condition, the primary subsystem or secondary subsystem sends sense information with unit check status to the appropriate host(s). This sense information is used during disaster recovery. You must transfer the sense information to the secondary site via the error reporting communications (ERC).

3.22.7.3 File and Database Recovery Procedures

When the primary or secondary subsystem suspends a URz 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.

URz 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. Note that the journal log file entries of most DBMS have the same system TOD clock information that is used for the I/O time-stamps (when timer type = system). The URz group consistency time can be extremely useful when performing this detection and retrieval. 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).
- The sense information with system time stamp which will be transferred via ERC.

3.22.7.4 Switching Operations to the Secondary Site

If a disaster or failure occurs at the primary site, the first disaster recovery activity is to use Business Continuity Manager to switch your operations to the remote backup site.

The basic procedures for switching operations to the remote backup site are:

1. Check whether the restore journal group includes a secondary data volume whose pair status is Pending duplex or Suspend (equivalent to SUSPOP in Business Continuity Manager).
If such a pair exists, consistency in the secondary data volume is dubious, and recovery with guaranteed consistency is impossible. In this case, if you want to use the secondary data volume, you must delete the pair.
2. If such a pair does not exist, use Business Continuity Manager to execute the YKSUSPND REVERSE option on the restore journal group (YKSUSPND is a command for splitting a pair).
If an error occurs, consistency in the secondary data volume is dubious, and recovery with guaranteed consistency is impossible. In this case, if you want to use the secondary data volume, you must delete the pair.
3. If no error occurs in step 2, wait until the splitting finishes. When the splitting finishes, the secondary data volume becomes usable with maintained consistency.
4. When the splitting finishes, use Business Continuity Manager to execute the YKRESYNC REVERSE option on the restore journal group (YKRESYNC is a command for restoring a pair). This option attempts to restore the pair and reverse the primary/secondary relationship.
5. Check whether there is a pair whose pair status of the restore journal group is Suspend (equivalent to SWAPPING in Business Continuity Manager).
If such a pair does not exist, the pair is successfully restored and the copy direction is reversed, and then copying of data from the secondary site to the primary site will start.

If the YKSUSPND command finishes successfully and the splitting ends successfully, you can resume business tasks (i.e., you can start business applications) by using secondary data volumes in the secondary site. Also, if the primary subsystem, the secondary subsystem, and remote copy connections are free from failure and fully operational, the restoring of the pair will finish successfully, and then copying of data from the secondary site to the primary site will start.

3.22.7.5 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 Business Continuity Manager to copy data from the secondary site to the primary site:

1. Restore the primary subsystem and remote copy connections, and make sure that all URz components are fully operational.
2. At the primary site, locate primary data volumes whose pair status is Pending duplex or Duplex, and then locate corresponding secondary data volumes whose pair status is Suspend, which is equivalent to SWAPPING in Business Continuity Manager terminology. 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 Simplex, and then locate corresponding secondary data volumes whose pair status is Simplex. If such volume pairs are found, issue a request for deleting the pairs to the primary data volumes.
4. At the primary site, locate data volume pairs whose pair status is Simplex, and then use Business Continuity Manager to execute YKRECOVER on the secondary data volume (YKRECOVER is a command for deleting a pair).
5. Execute the YKRESYNC REVERSE option on secondary data volumes whose pair status is Suspend, which is equivalent to SWAPPING in Business Continuity Manager terminology (YKRESYNC is the Business Continuity Manager command for resynchronizing pair). This reverses primary data volumes and secondary data volumes to resynchronize pairs.
6. Create pairs, specifying secondary data volumes whose pair status is Simplex 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 Pending Duplex to Duplex. If the pair status is changed to Duplex, 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.

3.22.7.6 Resuming Normal Operations at the Primary Site

Once the URz 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 Business Continuity Manager. Remember that the URz 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 URz components are fully operational and are free from failures.
2. Make sure that pair status of primary and secondary data volumes in all URz pairs is “Duplex”. This indicates that the URz 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 Business Continuity Manager to execute the YKSUSPND FLUSH SVOL PERMIT option on the master journal group (which was originally the restore journal group); YKSUSPND is a command for 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 Suspend (equivalent to SUSPOP with Business Continuity Manager). 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 Suspend (equivalent to SUSPOP with Business Continuity Manager), 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 YKSUSPND REVERSE command on the restore journal group (which were originally master journal group); YKSUSPND is a Business Continuity Manager command and REVERSE is an option. Wait until suspension completes.
8. After suspension completes, execute the Business Continuity Manager YKRESYNC REVERSE command on the restore journal group (which were originally master journal group). This reverses primary data volumes and secondary data volumes to resynchronize pairs and restores copy direction to its original direction.

3.23 CUIR

3.23.1 Overview

CUIR (Control Unit Initiated Reconfiguration) can perform “Vary Path Offline” or “Vary Path Online” instead of an operator, when replacing CHA PCB.

The work of replacement of CHA PCB is mitigated by this function.

3.23.2 Preparing for Operations

3.23.2.1 System Requirements

To be able to run, CUIR requires the following operating systems to be installed on the host computers connected to the PCB.

- z/OS V1.6 or later with PTF

The following restrictions apply when using CUIR.

Table 3.23.2.1-1 Restrictions that apply when using CUIR

No.	Item	Specifications
1	DKC emulation type	I-2105, 2107
2	CHA PCB type	FICON (DKC-F510I-8MS/8ML or DKC-F510I-16MSR/16MLR)
3	SVP operation	Only the replacement to same PCB type can be used. The addition (or removing) of PCB cannot be used.
4	Installation	Updating to extended SAIDs is required. (The power supply of equipment needs to be turned off temporarily)
5	Downgrade of microprogram	When changing to the microprogram which does not support the function, PSOFF is required in order to restore previous SAID.
6	Connection of a host and a PCB port	(1) The same LPAR is not physically patched to same PCB port via two or more switches. (See 3.23.4 (1)) (2) When performing 'PCB Replace', all online volumes must have the alternative host's path. If this requirement is not met, a host will suspend CUIR. (3) Only the paths to z/OS must be online, and other paths must not be online. (4) Don't change the online state of a path from a host after starting CUIR processing until it completes. This requirement is applied not only to target PCB but to all the mainframe paths in a subsystem. (box) It is applied also to the secondary-volume of a copy function. (TrueCopy/ShadowImage) (5) If the unstable path of connection is in a subsystem, CUIR processing may not be completed normally. In such a case, please use conventional "Replace" without CUIR.
7	Use of XRC	Don't use XRC simultaneously with CUIR.
8	Number of established logical paths	Up to 256 logical paths per LCU.
9	Device Reserve	The utility which performs prolonged device reservation (which exceeds 60 seconds) should not be running. (ex. DFDSS)

(To be continued)

(Continued from the preceding page)

No.	Item	Specifications
10	PCB status	Each MP on FICON PCB belongs to the group which processes only 16K specified volumes. CUIR is not available when all MPs of one certain group blockade. (See 3.23.4 (4)) In this case, please use the conventional 'Replace' without CUIR.
11	Host status	The host has to be running normally throughout CUIR processing. For example, if an operator inputs the system command "QUIESCE" of z/OS, CUIR will not be completed normally.
12	LED lamp of port	If there is a port which its LED lamp has switched off, remove the cable before execution of CUIR processing and insert it after completion of CUIR.
13	De-installation of device	When a CE returns an arbitrary volume of VIRTUAL-LVI to a space (VOL TO SPACE), there is a case where the host does not recognize the change. In such a case, CUIR will not be completed normally.
14	Universal Replicator	When the JOURNAL volume has the I/O definition (generation) on a host, CUIR is not completed normally.
15	After completing CUIR processing	• If it completes normally, confirming completion of RESUME by the message on all hosts is desired. • If it completes abnormally, be sure to perform the procedure of 'Recover all the paths in the Quiesce state' (See TRBL05-167).

3.23.2.2 Extended SAID

In the CUIR, each physical port uses the following extended SAIDs.

Table 3.23.2.2-1 Extended SAID for DKC515I

CL1		CL2	
1A	1B	2F	2E
—	0050	—	00D0
—	0051	—	00D1
—	0054	—	00D4
—	0055	—	00D5
—	0058	—	00D8
—	0059	—	00D9
—	005C	—	00DC
—	005D	—	00DD

3.23.2.3 CUIR Installation procedures

- (1) Exchanging microprogram
 - If the microprogram is not supporting CUIR, exchange for a new program
- (2) Updating to extended SAIDs
 - Perform “Vary Path Offline” for all paths from the hosts
 - Turn ON MODE493 (option for extension of SAID)
 - Turn ON after turning OFF the power switch of DKC
 - Perform “Vary Path Online” for all paths from the hosts
- (3) Validation of CUIR
 - Turn ON MODE494 (option for validating CUIR)
- (4) Confirming extended SAIDs
 - Check the extended SAIDs (Refer to Table 3.23.2.2-1) using ICKDSF/ANALYZE (Refer to the manual of z/OS about the details of ICKDSF/ANALYZE)
- (5) Device not recognized from the host

If CUIR processing encounters device conditions (Table 3.23.2.1-1 No.13, No.14), it will be abnormally ended by the following messages.

IOS283I	C.U.I.R. VARY PATH (XXXX, xx) REJECTED, PATH NOT OPERATIONAL
IOS290I	C.U.I.R. REQUEST UNSUCCESSFUL

In order to avoid this error, please input the following host command to those devices before execution of CUIR REPLACE. This command should be issued from all the hosts with the definition (generation) of those devices. And all the paths of those devices must be specified.

V PATH (XXXX, xx), OFFLINE

3.23.3 Operations of CUIR

Perform the following procedure, after checking CPU host's normal run.


- (1) Confirm that the MODE494 (option for validating CUIR) is ON.
- (2) Check that heavy load does not exist on target PCB.
- (3) Start the replacement of CHA PCB.
- (4) When the prompt "ONL3948i" of CUIR QUIESCE is displayed, answer permission of use
Warning is displayed to guarantee the identity of the paths after exchange of PCB.
- (5) If removal of PCB is requested, perform it.
- (6) If insertion of PCB is requested, perform it.
- (7) Wait for the link of all paths to become effective before answering the prompt "ONL3950i" of CUIR RESUME.
Warning is displayed to guarantee the identity of the paths before exchange of PCB.
- (8) Push "OK" button, and so CUIR RESUME processing starts.
- (9) The message which shows completion of the replacement of PCB is displayed.

If an error occurred, refer to [TRBL05-161](#).

3.23.4 Notes

(1) Narrowing down of host paths

FICON CUIR doesn't support the configuration shown below. If a LPAR uses multiple channels and the host paths are connected from these channels to a port of CU configuration via switch, CUIR is not available. Please do not build such the configuration of paths physically. Response of the host as expected is unreceivable only by the logical blockade (CHP OFF) of one of Path-1 or Path-2.



(2) Guarantee of path configuration

At CUIR operation, if the path configuration when performing CUIR Resume was different from the path configuration when CUIR Quiesce, CUIR Resume procedure is not performed correctly.

Therefore, when inserting the new package and connecting the cable to the package, please make it the same path configuration as previous one. CU checks the path configuration described above before performing starting CUIR Resume procedure.

If CU detects the path configuration difference, an error message appears on SVP and the CUIR procedure is suspended.

If the path cable drawn out from the CU port is inserted again, the path state will usually be recovered. Rarely, a host may not issue the command for recovery. Since CUIR cannot recognize the path as a target when a path state is not normal, CUIR may not complete normally.

(3) Combinations with PPs

Basically, there is no restriction in combination with other P.P.

However, when using CUIR operation, to change SAID, PS OFF/ON procedure is required.

You should be careful about it.

When using PPRC, the unique SAIDs have already assigned for the ports between MCU and RCU.

Therefore, CUIR support does not affect on it.

Regarding XRC, the path from Data Mover is often vary offline status.

In this case, CUIR function does not work on the path.

Therefore, if you forward the package replacement, XRC link path may be down.

This is the same case as usual.

You should do offline procedure the path of Data Mover before executing package replacement.

(4) PCB Status

When the PCB for replacement blockades, the conventional replacement procedure should be used. CUIR can be used even if some of MPs on target PCB blockade.

If the status of MP changes while performing CUIR, it will not complete normally like the above (3).

3.24 Cautions when Stopping the Subsystem

Powering off the subsystem greatly affects the operation when it is started next time.

For this reason, you should pay attention to the following points for the operation when you stop the subsystem.

3.24.1 Cautions when Stopping the Subsystem

If the battery backup time (36 hours ^(*)) is exceeded with the subsystem powered off, the management information stored in the memory will be cleared.

If the management information is cleared, some considerations are required when you use the device next time depending on the function installed in the device. That is, (1) the customer or the SE may have to set the management information again or follow the cautions related to the function; or (2) the service personnel may have to replace the battery that has overdischarged.

Therefore, you should propose the operation that even after the equipment is powered off, the user will power it on according to the battery backup specification (which enables a backup of 36 hours ^(*) by a charge for 24 hours ^(*) or more), in order to maintain the management information stored in the memory and prevent the battery from overdischarging.

The user may turn off the breakers during the legal inspection, etc. If this occurs, they can do so as long as they follow the battery backup specification (which enables a backup of 36 hours ^(*) by a charge for 24 hours ^(*) or more).

After stopping the subsystem intentionally, when the breakers in the distribution board are turned off exceeding battery backup time, it is proposed to set operation mode (Set local mode 677 and local mode 460 as “ON” together.) to restrain SM battery backup to prevent battery being wasted. For the cautions when setting this mode, refer to “3.24.3 Cautions when Turning Off the Battery Switch (Clearing the Management Information)”.

Table 3.24.1-1 How the Subsystem Operates after It Is Powered Off

After the subsystem is powered off	Operation in the subsystem	When the subsystem is started
AC input is supplied (breaker ON) or AC input is disconnected (breaker OFF)	This establishes the battery backup mode, in which the built-in battery backs up the data in the memory for up to 36 hours ^(*1) . If the power-OFF time exceeds the battery backup time, the data in the memory volatilizes and the battery deteriorates because of the overdischarge.	<u>Within the backup time</u> The subsystem can be started from the same status in which it is powered off. Charging the built-in battery to the full requires up to 24 hours. <u>Exceeding the backup time</u> The management information is cleared, so that it may require some cautions (*2), depending on the various functions installed. It may become necessary to replace the overcharged battery.

*1: When there is no battery failure and charging is done for 24 hours or more.

*2: For the cautions by each function, see 3.24.3, "Cautions when Turning Off the Battery Switch (Clearing the Management Information)".

3.24.2 Operation when Turning Off the Distribution Board Breaker

<Cautions>

When turning off the battery switch (see [LOC06-40](#)), you need to pay attention to the following points.

Turning off the battery switch carelessly will cause a serious failure leading to the data lost.

Before you turn off the battery switch, check that the “PS Off” of the device is completed normally by confirming that the LED on the DC Power Supply is off and the Power Off event log of the SVP is collected (see “3.12.2 Power OFF Procedure” in the installation section).

When you fail to confirm that the LED on the DC Power Supply is off and the Power Off event log is collected, stop the work and ask the customer to restart the device, and then check that the “PS Off” is completed normally after the request.

<Cautions>

1. See the Maintenance screen of the SVP ([SVP03-10](#)) and check that the battery is normal.
2. Check that there is no discharge from the battery when the breaker is on or the power supply is on for 24 hours or more.
3. To confirm that the subsystem is normally powered off, check that the LED on the DC Power Supply is off and the Power Off event log is collected according to “3.12.2 Power OFF Procedure” in the installation section.
4. If you fail to confirm that the LED on the DC Power Supply is off and the Power Off event log is collected in step 3, you must not turn off the battery switch. Before you turn off the battery switch, you must restart the device and confirm that the “PS Off” is completed normally, the LED on the DC Power Supply is off, and the Power Off event is collected.

If it cannot be powered on according to the battery backup specification (which enables a backup of 36 hours by a charge for 24 hours or more) after the subsystem is powered off, request the customer to perform the following operation thoroughly.

1. Request the customer to check that the power off procedure has been completed (the READY, MESSAGE, and ALARM lamps are off) before the breaker is off. The service personnel should check the Power Off event log shown in the cautions above.
2. If the device stops exceeding the backup time, request the customer to instruct the service personnel to turn off the battery switch beforehand in order to prevent the battery from overdischarging. (It is unnecessary to turn off the battery switch if the operation mode (local mode 677=ON) is set to restrain SM battery backup beforehand.)
3. Turning off the battery switch makes it impossible to maintain the management information stored in the memory. Inform that the cautions must be followed when the device is started next time depending on various functions installed (such as resident Cache and remote copy). For the cautions, see 3.24.3, “Cautions when Turning Off the Battery Switch (Clearing the Management Information)”.
4. When powering on the device, you should request the customer to instruct the service personnel to turn on the battery switch beforehand and ask the customer to set the management information again as required.

3.24.3 Cautions when Turning Off the Battery Switch (Clearing the Management Information)

Table 3.24.3-1 shows the cautions to be followed before you turn off the battery switch (clearing the management information) regarding each function and the cautions to be followed after you turn on the battery switch and power on the device. When you turn on the battery switch, you must pay attention to the time required and the procedure until each function recovers after the device is powered on.

Even when the operation mode (local mode 677=ON) that restrains the battery backup is set, cautions similar to the case of the battery switch-off need to follow.

For the following functions, however, the management information is not cleared because the local mode setting makes the management information saved to the HDD of the SVP when the power supply is turned off and makes it recovered from the SVP to the SM when the volatile power supply is turned on (for details, see 3.21, the THEORY section).

- (1) ShadowImage
- (2) ShadowImage OS/390
- (3) Volume Migration
- (4) Compatible Mirroring for IBM® FlashCopy®
- (5) Compatible Mirroring for IBM® FlashCopy® Version 2
- (6) TrueCopy
- (7) TrueCopy for Mainframe
- (8) Universal Replicator
- (9) Universal Replicator for Mainframe
- (10) Copy-on-Write Snapshot

If data saving fails and the volatile power supply is turned on, the management information is cleared.

You are requested to ask our SE about details of the work procedure and the estimate of the time for processing because they differ depending on the device environment and the operation status of each function.

This work does not occur on the functions not listed in the Table 3.24.3-1.

**Table 3.24.3-1 Cautions when Turning Off the Battery Switch
(Clearing the Management Information)**

No.	Function name	Precautions	Cautions after the device is powered on
1	[Open remote copy function] TrueCopy TrueCopy Asynchronous [M/F open remote copy function] TrueCopy for Mainframe TrueCopy Asynchronous for Mainframe	If the local mode setting is Off, the copying time after turning on the power supply varies depending on the pair status when the device is powered off. Change the pair status to DUPLEX and turn off the power supply. For how to turn off/on the power supply using this function, follow the procedures in the User's Guides shown below. For details, ask the SE before starting the work. [Related documents] <u>TrueCopy User's Guide</u> 5.10 Powering Off/On TrueCopy Components <u>TrueCopy for Mainframe User's Guide</u> 5.11 Powering Off/On TrueCopy for z/OS® Components	If the local mode setting is Off, <u>the copying time after turning on the power supply varies depending on the pair status when the device is powered off.</u> (1) If the pair state is SUSPEND and the power supply is off, the difference information is cleared, so that it takes more time required because the pair resync processing performed after turning on the power supply will be the same as the formation copy processing (copying all volumes). (2) If the pair status is DUPLEX and the power supply is turned off, the data copy shown in (1) does not occur when the power supply is turned on. The pair automatically becomes the DUPLEX status. (3) When the power supply is turned off in the PENDING status, copying restarts automatically when turning on the power supply. However, you need to pay attention to the time required because the difference information is cleared and all of the data having been copied before the power supply is turned off will become the target of the formation copy.
2	[Open data replication function] ShadowImage [M/F data replication function] ShadowImage for Mainframe	If the local mode setting is Off, copying all volumes will become necessary irrespective of the pair status when the device is powered off. For how to turn off/on the power supply using this function, ask the SE and follow the regular work procedure.	If the local mode setting is Off, <u>copying all volumes will become necessary irrespective of the pair status when the power supply is turned off. Because the difference information between the primary and secondary volumes is cleared, you need to pay attention to the time required for the status recovery.</u> (1) If the pair status is SP-PEND or SPLIT and the power supply is turned off, you should pay attention that it takes more time because the pair resync processing performed after turning on the power supply will be the way of copying all volumes. If the pair status is SP-PEND, the pair will be suspended when the power supply is turned on. (2) If the pair status is DUPLEX, PENDING, or RESYNC and if the power supply is turned on, automatic copying restarts. However, it takes more time to copy because of copying all volumes.

(To be continued.)

(Continued from the preceding page)

No.	Function name	Precautions	Cautions after the device is powered on
3	[Data replication function, flash copy] Compatible Mirroring for IBM® FlashCopy® Compatible Mirroring for IBM® FlashCopy®V2	If the local mode setting is Off, release <u>the relationship of the primary and secondary volumes beforehand</u> . For FlashCopy V2, in particular, it is a data set copy and therefore regular user data must exist in only the secondary volume. If the secondary volume is blocked under this status, there may be a serious effect on the application using the particular user data. Therefore, be sure to follow the regular procedure. For how to turn off/on the power supply using this function, ask the SE and follow the regular work procedure.	If the local mode setting is Off, <u>set the relationship again after the device is powered on</u> . <u>The secondary volume is blocked.</u> The secondary volume is blocked if the power supply is turned off when there is a relationship between the primary and secondary volumes. To recover the blocked secondary volume, you need to set the volume format and relationship again.
4	[Hierarchy control function] Volume Migration	If the local mode setting is Off, power <u>off the device after the volume transfer is completed</u> . <u>A part of the hierarchy control monitor information is lost.</u> The monitor information, covering the period from the last time when the monitor information required for hierarchy control is taken until the power supply is turned off, is lost. Collect the monitor information before the power supply is turned off. [Related documents] <u>Performance Manager User's Guide</u> 2.7 Overview of Performance Monitor Operations (see "2.7.1 Note".) For details about how to turn off/on the power supply using this function, ask the SE and follow the regular work procedure.	If the local mode setting is Off, <u>pay attention to the time required until the volume transfer is completed</u> . If the power supply is turned off before the volume transfer is completed, the data, which has already been copied before the power supply is turned off, is handled as the target of copy again so that copying covers all the volumes. The above work does not occur if volume transfer is already completed.

(To be continued.)

(Continued from the preceding page)

No.	Function name	Precautions	Cautions after the device is powered on
5	[Open universal replicator function] UniversalReplicator [M/F universal replicator function] UniversalReplicator - M/F	If the local mode setting is Off, <u>the copying time after turning on the power supply varies depending on the pair status when the device is powered off.</u> It is recommended to turn off the power supply when the pair status is DUPLEX. Turning off the power supply under the DUPLEX status automatically makes it in the DUPLEX status when the subsystem is started next time. For how to turn off/on the power supply using this function, follow the procedures in the User's Guides shown below. For details, ask the SE before starting the work. [Related documents] <u>UniversalReplicator User's Guide</u> Appendix A. Power Management for Disk Subsystem and Network Relay Devices <u>UniversalReplicator Mainframe User's Guide</u> Appendix A. Power Management for Disk Subsystem and Network Relay Devices	If the local mode setting is Off, when the pair status is changed to something other than DUPLEX before the power supply is turned off, the difference information on pairs is initialized so that the difference covers everything. This increases the time required until the status becomes DUPLEX. (1) If the pair is DUPLEX and the power supply is turned off, the data copy shown in (2) does not occur when the power supply is turned on. As a result, the DUPLEX status remains. (2) If the pair status is SUSPEND and the power supply is off, the difference information is cleared, so that it will take more time required because the pair resync processing performed after turning on the power supply will be the same as the formation copy processing (copying all volumes). (3) When the power supply is turned off in the PENDING status, copying restarts automatically when turning on the power supply. However, you need to pay attention to the time required because the difference information is cleared and all of the data having been copied before the power supply is turned off will become the target of the formation copy.

(To be continued.)

(Continued from the preceding page)

No.	Function name	Precautions	Cautions after the device is powered on
6	[External subsystem connection function] Universal Volume Manager	When you turn off the power supply and start the subsystem next time, you should use the following order. Beginning with the TagmaStore NSC and the external subsystem (hereinafter called the NSC and the outside), use the following procedure to turn off the power supply. (1) When stopping both the NSC and the outside Stop the I/O for the external subsystem and power off the TagmaStore NSC, and then power off the external subsystem. (2) When stopping either the TagmaStore NSC or the external subsystem <u>2.1 When powering off only the TagmaStore NSC</u> Stop the I/O for the external subsystem, and then power off the TagmaStore NSC. <u>2.2 When powering off only the external subsystem</u> Stop the I/O for the external subsystem, execute the Disconnect Subsystem command (to destage the external subsystem) of Universal Volume Manager, and then power off the external subsystem. [Related documents] <u>Universal Volume Manager User's Guide</u> 2.7 Turning On or Off Power Supply of Subsystem	<u>When powering on the NSC and the outside, you should use the following order.</u> (1) <u>When powering on both the NSC and the outside</u> Power on the outside, check that READY is displayed, and then power on the NSC. (2) When powering on either the NSC or the outside <u>2.1 Powering On Only the NSC</u> Check that the external subsystem shows READY status and then power on the NSC. <u>2.2 Powering On Only the Outside</u> Power on the external subsystem and then execute the Check Path & Restore command of Universal Volume Manage.
7	[External connection copy] Cross-System Copy	When you turn off the power supply and start the subsystem next time, see No. 8. For copying, there is no particular thing to pay attention because copying does not cover all volumes but proceeds continuously.	See No. 8.

(To be continued.)

(Continued from the preceding page)

No.	Function name	Precautions	Cautions after the device is powered on
8	[Snap shot creation function] Copy-On Write Snapshot	If the local mode setting is Off, the pool information and pair information having been set before the power supply is turned off will disappear irrespective of the pair status when the device is powered off. <u>After starting the subsystem next time, you need to recover the pool and create Snapshot pairs again.</u> [Related documents] <u>Copy-on-Write Snapshot User's Guide</u> 4.4.1 Note on Switching Off the Power Supply For how to turn off/on the power supply using this function, ask the SE and follow the regular work procedure.	If the local mode setting is Off, <u>you need to recover the pool and create Snapshot pairs again.</u>
9	[Resident cache function] Cache Residency Manager for Mainframe Cache Residency Manager	There are no particular precautions, but be careful of the following case: If pre-staging has not been completed, a Cache error occurs at the first read access to the resident region after the power supply is turned on. Although data is read from the drive, the following access result in Cache hits.	<u>A Cache error occurs only at the first access.</u> Because the resident data in the Cache memory is cleared, a Cache error occurs at the first read access to the resident region after the device is powered on. Although data is read from the drive, the following access result in Cache hits.

3.25 Hyper PAV

3.25.1 Overview of Hyper PAV

Hyper PAV was evolved from static PAV and dynamic PAV.

If you use Hyper PAV, alias devices that were assigned to a base device are shared with all base devices in the same CU. Thereby, alias devices do not need to be shifted by dynamic PAV.

Moreover, you can use more devices for a base device than when you use PAV because you can reduce devices that are assigned for an alias device.

You can specify the type of PAV (PAV or Hyper PAV) to use for each host computer. Therefore, an alias device may accept both I/O requests that are issued when you use the PAV and Hyper PAV.

Hyper PAV is supported by version 50-09-5x or later.

3.25.2 Hyper PAV function setting procedure

This appendix describes the procedures for installing and uninstalling Hyper PAV.

This appendix also describes the point to be checked when you change the setting to enable or disable Hyper PAV on the host computer and when you restart TagmaStore NSC while using Hyper PAV.

■ Installing Hyper PAV

To install Hyper PAV:

1. Upgrade the host software to support Hyper PAV, that is, z/OS 18 or later, or z/OS 1.6 with PTF or later.
2. If you have used TagmaStore NSC already, upgrade TagmaStore NSC to the microcode version supporting Hyper PAV (50-09-5x or later). If you want to upgrade TagmaStore NSC, contact the Hitachi Technical Support Center.
(Refer to “MICRO-FC SECTION” for upgrade of the microcode.)
3. Install the Compatible PAV Storage Navigator software.
(Refer to “Storage Navigator™ User’s Guide” “Compatible PAV for z/OS® User’s Guide” for upgrade of the microcode.)
4. Install the Compatible Hyper PAV Storage Navigator software.
5. Assign aliases.
If you introduce TagmaStore NSC to new, and you have already assigned aliases to the base volumes for Hyper PAV, please skip this step.
(Refer to “Compatible PAV for z/OS® User’s Guide” “Chapter 3 Preparing for Compatible PAV Operations” and “Chapter 4 Performing Compatible PAV Operations” for assign aliases.)
6. On the host computer, enable Hyper PAV.
(Refer to “Compatible PAV for z/OS® User’s Guide” “2.1.6 Setting of Enabling Compatible Hyper PAV” for enable/disable Hyper PAV.)

7. Issue the DEVSERV QPAV command from the host, and make sure that the aliases are displayed as ones for Hyper PAV. If aliases can not be displayed as for Hyper PAV, execute either operation of the following, and then, issue the command and check the display again.
 - If the host accesses only to the corresponding TagmaStore NSC, disable Hyper PAV on the host computer, and then, enable Hyper PAV again .
 - If the host accesses to other disk subsystems that use Hyper PAV, turn VARY off, and turn CHP off of the host path, and then, turn CHP on, and turn VARY on again.(Refer to “Compatible PAV for z/OS® User’s Guide” “5.1 Additional MVS Commands” “DEVSERV QPAV” “DS QP,xxxx,HPAV” for the alias display.)
8. If you also use Cross-OS File Exchange on the host computer, execute the following steps after executing step 1 to step 7.
 - (1) Turn VARY off of the host path of Cross-OS File Exchange volume.
 - (2) Turn VARY on of the host path of Cross-OS File Exchange volume.
9. Procedure is end.

■ Check Point for Restarting TagmaStore NSC While Using Hyper PAV

If you restart TagmaStore NSC while using Hyper PAV, please issue the DEVSERV QPAV command from the host after restart TagmaStore NSC, and make sure that the aliases are displayed as ones for Hyper PAV. If aliases can not be displayed as for Hyper PAV, execute the following operations, and then, check again.

- Turn VARY off, and then turn CHP off of the host path.
- Turn CHP on, and then turn VARY on of the host path.

(Refer to “Compatible PAV for z/OS® User’s Guide” “5.1 Additional MVS Commands” “DEVSERV QPAV” “DS QP,xxxx,HPAV” for the alias display.)

■ Uninstalling Hyper PAV

To uninstall Hyper PAV:

1. On the host computer, disable Hyper PAV.
(Refer to “Compatible PAV for z/OS® User’s Guide” “2.1.6 Setting of Enabling Compatible Hyper PAV” for enable/disable Hyper PAV.)
2. Uninstall the Hyper PAV Storage Navigator software.
(For information on uninstall of Storage Navigator software, please refer to Storage Navigator™ User’s Guide.)
3. Execute the DEVSERV (DS) command as shown below from the host to a device per CU:
DS QD,xxx,VALIDATE (XXX = device number)
4. Procedure is end.

To uninstall Hyper PAV if other disk subsystem which the host accesses to uses Hyper PAV and Cross-OS File Exchange:

1. Turn VARY off, and then turn CHP off of the host path.
2. Uninstall the Hyper PAV Storage Navigator software.
3. Turn CHP on, and then turn VARY on of the host path.
4. Procedure is end.

4. Power-on Sequences

4.1 IMPL Sequence

The IMPL sequence, which is executed when power is turned on, is comprises of the following four modules:

(1) Boot loader

The boot loader performs the minimum necessary amount of initializations after a ROM boot. Subsequently, the boot loader expands the local memory loader from the flash memory into the local memory and the local memory loader is executed.

(2) Local memory loader

The local memory loader loads the Real Time OS load modules into the local memory and the Real Time OS is executed.

(3) Real Time OS

Real Time OS is a root task that initializes the tables in the local memory that are used for intertask communications. Real Time OS also tests the hardware resources.

(4) DKC task

When the DKC task is created, it executes initialization routines. Initialization routines initialize the most part of the environment that the DKC task uses. When the environment is established so that the DKC task can start scanning, the DKC task notifies the SVP of a power event log. Subsequently, the DKC task turns on the power for the physical drives and, when the logical drives become ready, The DKC task notifies the host processor of an NRTR. The control flow of IMPL processing is shown in Fig. 4.1-1.




Fig. 4.1-1 IMPL Sequence

4.2 Drive Power-on Sequence

An overcurrent condition will occur if two or more drives are started at the same time. To preclude the overcurrent condition, DKUs should be started at the power supply level, 6HDDs (MAX) at a time, at approximately 10 second intervals.

When the logical devices become ready as the result of the startup of the physical drives, the host processor is notified to that effect.





Fig. 4.2-1 Drive Power-on Sequence

4.3 Planned Stop

When a power-off is specified by a maintenance personnel, this subsystem checks for termination of tasks that are blocked or running on all logical devices. When all the tasks are terminated, this subsystem disables the CHL and executes emergency destaging. If a track for which destaging fails (pinned track) occurs, this subsystem stores the pin information in shared memory. Subsequently, this subsystem saves the pin information, which is used as hand-over information, in flash memory, sends Power Event Log to the SVP, and notifies the hardware of the grant to turn off the power.

The hardware turns off main power when power-off grants for all processors are presented.



5. Appendixes

5.1 Physical-Logical Device Matrixes

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (1/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDU-00	HDD00-00	00/00	01-01	01-01
	HDD00-01	00/01	01-02	01-01
	HDD00-02	00/02	01-03	01-03
	HDD00-03	00/03	01-04	01-03
	HDD00-04	00/04	01-05	01-05
	HDD00-05	00/05	01-06	01-05
	HDD00-06	00/06	01-07	01-07
	HDD00-07	00/07	01-08	01-07
	HDD00-08	00/08	01-09	01-09
	HDD00-09	00/09	01-10	01-09
	HDD00-0A	00/0A	01-11	01-11
	HDD00-0B	00/0B	01-12	01-11
	HDD00-0C	00/0C	01-13	01-13
	HDD00-0D	00/0D	01-14	01-13
	HDD00-0E	00/0E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (2/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDU-01	HDD01-00	01/00	01-01	01-01
	HDD01-01	01/01	01-02	01-01
	HDD01-02	01/02	01-03	01-03
	HDD01-03	01/03	01-04	01-03
	HDD01-04	01/04	01-05	01-05
	HDD01-05	01/05	01-06	01-05
	HDD01-06	01/06	01-07	01-07
	HDD01-07	01/07	01-08	01-07
	HDD01-08	01/08	01-09	01-09
	HDD01-09	01/09	01-10	01-09
	HDD01-0A	01/0A	01-11	01-11
	HDD01-0B	01/0B	01-12	01-11
	HDD01-0C	01/0C	01-13	01-13
	HDD01-0D	01/0D	01-14	01-13
	HDD01-0E	01/0E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (3/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-02	HDD02-00	02/00	01-01	01-01
	HDD02-01	02/01	01-02	01-01
	HDD02-02	02/02	01-03	01-03
	HDD02-03	02/03	01-04	01-03
	HDD02-04	02/04	01-05	01-05
	HDD02-05	02/05	01-06	01-05
	HDD02-06	02/06	01-07	01-07
	HDD02-07	02/07	01-08	01-07
	HDD02-08	02/08	01-09	01-09
	HDD02-09	02/09	01-10	01-09
	HDD02-0A	02/0A	01-11	01-11
	HDD02-0B	02/0B	01-12	01-11
	HDD02-0C	02/0C	01-13	01-13
	HDD02-0D	02/0D	01-14	01-13
	HDD02-0E	02/0E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (4/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-03	HDD03-00	03/00	01-01	01-01
	HDD03-01	03/01	01-02	01-01
	HDD03-02	03/02	01-03	01-03
	HDD03-03	03/03	01-04	01-03
	HDD03-04	03/04	01-05	01-05
	HDD03-05	03/05	01-06	01-05
	HDD03-06	03/06	01-07	01-07
	HDD03-07	03/07	01-08	01-07
	HDD03-08	03/08	01-09	01-09
	HDD03-09	03/09	01-10	01-09
	HDD03-0A	03/0A	01-11	01-11
	HDD03-0B	03/0B	01-12	01-11
	HDD03-0C	03/0C	01-13	01-13
	HDD03-0D	03/0D	01-14	01-13
	HDD03-0E	03/0E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (5/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-10	HDD10-00	00/10	02-01	02-01
	HDD10-01	00/11	02-02	02-01
	HDD10-02	00/12	02-03	02-03
	HDD10-03	00/13	02-04	02-03
	HDD10-04	00/14	02-05	02-05
	HDD10-05	00/15	02-06	02-05
	HDD10-06	00/16	02-07	02-07
	HDD10-07	00/17	02-08	02-07
	HDD10-08	00/18	02-09	02-09
	HDD10-09	00/19	02-10	02-09
	HDD10-0A	00/1A	02-11	02-11
	HDD10-0B	00/1B	02-12	02-11
	HDD10-0C	00/1C	02-13	02-13
	HDD10-0D	00/1D	02-14	02-13
	HDD10-0E	00/1E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (6/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-11	HDDR11-00	01/10	02-01	02-01
	HDDR11-01	01/11	02-02	02-01
	HDDR11-02	01/12	02-03	02-03
	HDDR11-03	01/13	02-04	02-03
	HDDR11-04	01/14	02-05	02-05
	HDDR11-05	01/15	02-06	02-05
	HDDR11-06	01/16	02-07	02-07
	HDDR11-07	01/17	02-08	02-07
	HDDR11-08	01/18	02-09	02-09
	HDDR11-09	01/19	02-10	02-09
	HDDR11-0A	01/1A	02-11	02-11
	HDDR11-0B	01/1B	02-12	02-11
	HDDR11-0C	01/1C	02-13	02-13
	HDDR11-0D	01/1D	02-14	02-13
	HDDR11-0E	01/1E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (7/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDU-12	HDDR12-00	02/10	02-01	02-01
	HDDR12-01	02/11	02-02	02-01
	HDDR12-02	02/12	02-03	02-03
	HDDR12-03	02/13	02-04	02-03
	HDDR12-04	02/14	02-05	02-05
	HDDR12-05	02/15	02-06	02-05
	HDDR12-06	02/16	02-07	02-07
	HDDR12-07	02/17	02-08	02-07
	HDDR12-08	02/18	02-09	02-09
	HDDR12-09	02/19	02-10	02-09
	HDDR12-0A	02/1A	02-11	02-11
	HDDR12-0B	02/1B	02-12	02-11
	HDDR12-0C	02/1C	02-13	02-13
	HDDR12-0D	02/1D	02-14	02-13
	HDDR12-0E	02/1E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (8/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDU-13	HDDR13-00	03/10	02-01	02-01
	HDDR13-01	03/11	02-02	02-01
	HDDR13-02	03/12	02-03	02-03
	HDDR13-03	03/13	02-04	02-03
	HDDR13-04	03/14	02-05	02-05
	HDDR13-05	03/15	02-06	02-05
	HDDR13-06	03/16	02-07	02-07
	HDDR13-07	03/17	02-08	02-07
	HDDR13-08	03/18	02-09	02-09
	HDDR13-09	03/19	02-10	02-09
	HDDR13-0A	03/1A	02-11	02-11
	HDDR13-0B	03/1B	02-12	02-11
	HDDR13-0C	03/1C	02-13	02-13
	HDDR13-0D	03/1D	02-14	02-13
	HDDR13-0E	03/1E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (9/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-20	HDD20-00	00/20	03-01	03-01
	HDD20-01	00/21	03-02	03-01
	HDD20-02	00/22	03-03	03-03
	HDD20-03	00/23	03-04	03-03
	HDD20-04	00/24	03-05	03-05
	HDD20-05	00/25	03-06	03-05
	HDD20-06	00/26	03-07	03-07
	HDD20-07	00/27	03-08	03-07
	HDD20-08	00/28	03-09	03-09
	HDD20-09	00/29	03-10	03-09
	HDD20-0A	00/2A	03-11	03-11
	HDD20-0B	00/2B	03-12	03-11
	HDD20-0C	00/2C	03-13	03-13
	HDD20-0D	00/2D	03-14	03-13
	HDD20-0E	00/2E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (10/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-21	HDD21-00	01/20	03-01	03-01
	HDD21-01	01/21	03-02	03-01
	HDD21-02	01/22	03-03	03-03
	HDD21-03	01/23	03-04	03-03
	HDD21-04	01/24	03-05	03-05
	HDD21-05	01/25	03-06	03-05
	HDD21-06	01/26	03-07	03-07
	HDD21-07	01/27	03-08	03-07
	HDD21-08	01/28	03-09	03-09
	HDD21-09	01/29	03-10	03-09
	HDD21-0A	01/2A	03-11	03-11
	HDD21-0B	01/2B	03-12	03-11
	HDD21-0C	01/2C	03-13	03-13
	HDD21-0D	01/2D	03-14	03-13
	HDD21-0E	01/2E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (11/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-22	HDD22-00	02/20	03-01	03-01
	HDD22-01	02/21	03-02	03-01
	HDD22-02	02/22	03-03	03-03
	HDD22-03	02/23	03-04	03-03
	HDD22-04	02/24	03-05	03-05
	HDD22-05	02/25	03-06	03-05
	HDD22-06	02/26	03-07	03-07
	HDD22-07	02/27	03-08	03-07
	HDD22-08	02/28	03-09	03-09
	HDD22-09	02/29	03-10	03-09
	HDD22-0A	02/2A	03-11	03-11
	HDD22-0B	02/2B	03-12	03-11
	HDD22-0C	02/2C	03-13	03-13
	HDD22-0D	02/2D	03-14	03-13
	HDD22-0E	02/2E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (12/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-23	HDD23-00	03/20	03-01	03-01
	HDD23-01	03/21	03-02	03-01
	HDD23-02	03/22	03-03	03-03
	HDD23-03	03/23	03-04	03-03
	HDD23-04	03/24	03-05	03-05
	HDD23-05	03/25	03-06	03-05
	HDD23-06	03/26	03-07	03-07
	HDD23-07	03/27	03-08	03-07
	HDD23-08	03/28	03-09	03-09
	HDD23-09	03/29	03-10	03-09
	HDD23-0A	03/2A	03-11	03-11
	HDD23-0B	03/2B	03-12	03-11
	HDD23-0C	03/2C	03-13	03-13
	HDD23-0D	03/2D	03-14	03-13
	HDD23-0E	03/2E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (13/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-30	HDD30-00	00/30	04-01	04-01
	HDD30-01	00/31	04-02	04-01
	HDD30-02	00/32	04-03	04-03
	HDD30-03	00/33	04-04	04-03
	HDD30-04	00/34	04-05	04-05
	HDD30-05	00/35	04-06	04-05
	HDD30-06	00/36	04-07	04-07
	HDD30-07	00/37	04-08	04-07
	HDD30-08	00/38	04-09	04-09
	HDD30-09	00/39	04-10	04-09
	HDD30-0A	00/3A	04-11	04-11
	HDD30-0B	00/3B	04-12	04-11
	HDD30-0C	00/3C	04-13	04-13
	HDD30-0D	00/3D	04-14	04-13
	HDD30-0E	00/3E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (14/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-31	HDD31-00	01/30	04-01	04-01
	HDD31-01	01/31	04-02	04-01
	HDD31-02	01/32	04-03	04-03
	HDD31-03	01/33	04-04	04-03
	HDD31-04	01/34	04-05	04-05
	HDD31-05	01/35	04-06	04-05
	HDD31-06	01/36	04-07	04-07
	HDD31-07	01/37	04-08	04-07
	HDD31-08	01/38	04-09	04-09
	HDD31-09	01/39	04-10	04-09
	HDD31-0A	01/3A	04-11	04-11
	HDD31-0B	01/3B	04-12	04-11
	HDD31-0C	01/3C	04-13	04-13
	HDD31-0D	01/3D	04-14	04-13
	HDD31-0E	01/3E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (15/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-32	HDD32-00	02/30	04-01	04-01
	HDD32-01	02/31	04-02	04-01
	HDD32-02	02/32	04-03	04-03
	HDD32-03	02/33	04-04	04-03
	HDD32-04	02/34	04-05	04-05
	HDD32-05	02/35	04-06	04-05
	HDD32-06	02/36	04-07	04-07
	HDD32-07	02/37	04-08	04-07
	HDD32-08	02/38	04-09	04-09
	HDD32-09	02/39	04-10	04-09
	HDD32-0A	02/3A	04-11	04-11
	HDD32-0B	02/3B	04-12	04-11
	HDD32-0C	02/3C	04-13	04-13
	HDD32-0D	02/3D	04-14	04-13
	HDD32-0E	02/3E	Spare	Spare

RELATIONSHIP BETWEEN DISK DRIVE# AND PARITY GROUP# (16/16)

HDD BOX Number	Disk Drive Number	C# / R#	Parity Group Number (RAID5 3D+1P) (RAID1 2D+2D)	Parity Group Number (RAID5 7D+1P) (RAID6 6D+2P)
HDDU-33	HDD33-00	03/30	04-01	04-01
	HDD33-01	03/31	04-02	04-01
	HDD33-02	03/32	04-03	04-03
	HDD33-03	03/33	04-04	04-03
	HDD33-04	03/34	04-05	04-05
	HDD33-05	03/35	04-06	04-05
	HDD33-06	03/36	04-07	04-07
	HDD33-07	03/37	04-08	04-07
	HDD33-08	03/38	04-09	04-09
	HDD33-09	03/39	04-10	04-09
	HDD33-0A	03/3A	04-11	04-11
	HDD33-0B	03/3B	04-12	04-11
	HDD33-0C	03/3C	04-13	04-13
	HDD33-0D	03/3D	04-14	04-13
	HDD33-0E	03/3E	Spare	Spare

5.2 Commands

These subsystem commands are classified into the following eight categories:

(1) Read commands

The read commands transfer the readout data from devices to channels.

(2) Write commands

The write commands write the transfer data from channels to devices.

(3) Search commands

The search commands follow a control command and logically search for the target data.

(4) Control commands

The control commands include the SEEK command that positions cylinder and head positions. The SET SECTOR command that executes latency time processing, the LOCATE RECORD command that specifies the operation of the ECKD command, the SET FILE MASK command that defines the permissible ranges for the WRITE and SEEK operations, and the DEFINE EXTENT command that defines the permissible ranges for the WRITE and SEEK operations and that defines the cache access mode.

(5) Sense commands

The sense commands transfer sense bytes and device specifications.

(6) Path control commands

The path control commands enable and disable the exclusive control of devices.

(7) TEST I/O command

The TEST I/O command transfers the specified device and its path state to a given channel in the form of DSBs.

(8) Subsystem commands

The subsystem commands include the commands which define cache control information in the DKCs and the commands which transfer cache-related information to channels.

Table 5.2 -1 Command Summary (1/3)

Command Name			Command Code	
			Single Track	Multitrack
Read commands	READ INITIAL PROGRAM LOAD (RD IPL)		02	—
	READ HOME ADDRESS (RD HA)		1A	9A
	READ RECORD ZERO (RD R0)		16	96
	READ COUNT,KEY,DATA (RD CKD)		1E	8E
	READ KEY,DATA (RD KD)		0E	86
	READ DATA (RD D)		06	92
	READ COUNT (RD C)		12	—
	READ MULTIPLE COUNT,KEY AND DATA (RD MCKD)		5E	—
	READ TRACK (RD TRK)		DE	—
	READ SPECIAL HOME ADDRESS (RD SP HA)		0A	—
WRITE commands	WRITE HOME ADDRESS (WR HA)		19	—
	WRITE RECORD ZERO (WR R0)		15	—
	WRITE COUNT,KEY,DATA (WR CKD)		1D	—
	WRITE COUNT,KEY,DATA NEXT TRACK (WR CKD NT)		9D	—
	ERASE (ERS)		11	—
	WRITE KEY AND DATA (WR KD)		0D	—
	WRITE UPDATE KEY AND DATA (WR UP KD)		8D	—
	WRITE DATA (WR D)		05	—
	WRITE UPDATE DATA (WR UP D)		85	—
	WRITE SPECIAL HOME ADDRESS (WR SP HA)		09	—
SEARCH commands	SEARCH HOME ADDRESS (SCH HA EQ)		39	B9
	SEARCH ID EQUAL (SCH ID EQ)		31	B1
	SEARCH ID HIGH (SCH ID HI)		51	D1
	SEARCH ID HIGH OR EQUAL (SCH ID HE)		71	F1
	SEARCH KEY EQUAL (SCH KEY EQ)		29	A9
	SEARCH KEY HIGH (SCH KEY HI)		49	C9
	SEARCH KEY HIGH OR EQUAL (SCH KEYD HE)		69	E9

Table 5.2 -1 Command Summary (2/3)

Command Name			Command Code	
			Single Track	Multitrack
CONTROL commands	DEFINE EXTENT	(DEF EXT)	63	—
	LOCATE RECORD	(LOCATE)	47	—
	LOCATE RECORD EXTENDED	(LOCATE EXT)	4B	—
	SEEK	(SK)	07	—
	SEEK CYLINDER	(SK CYL)	0B	—
	SEEK HEAD	(SK HD)	1B	—
	RECALIBRATE	(RECAL)	13	—
	SET SECTOR	(SET SECT)	23	—
	SET FILE MASK	(SET FM)	1F	—
	READ SECTOR	(RD SECT)	22	—
	SPACE COUNT	(SPC)	0F	—
	NO OPERATION	(NOP)	03	—
	RESTORE	(REST)	17	—
	DIAGNOSTIC CONTROL	(DIAG CTL)	F3	—
SENSE commands	SENSE	(SNS)	04	—
	READ AND RESET BUFFERED LOG	(RRBL)	A4	—
	SENSE IDENTIFICATION	(SNS ID)	E4	—
	READ DEVICE CHARACTERISTICS	(RD CHR)	64	—
	DIAGNOSTIC SENSE/READ	(DIAG SNS/RD)	C4	—
PATH CONTROL commands	DEVICE RESERVE	(RSV)	B4	—
	DEVICE RELEASE	(RLS)	94	—
	UNCONDITIONAL RESERVE	(UNCON RSV)	14	—
	SET PATH GROUP ID	(SET PI)	AF	—
	SENSE SET PATH GROUP ID	(SNS PI)	34	—
	SUSPEND MULTIPATH RECONNECTION	(SUSP MPR)	5B	—
	RESET ALLEGIANCE	(RST ALG)	44	—
TST I/O	TEST I/O	(TIO)	00	—
TIC	TRANSFER IN CHANNEL	(TIC)	X8	—

Table 5.2 -1 Command Summary (3/3)

Command Name		Command Code	
		Single Track	Multitrack
SUBSYSTEM commands	SET SUBSYSTEM MODE (SET SUB MD)	87	—
	PERFORM SUBSYSTEM FUNCTION (PERF SUB FUNC)	27	—
	READ SUBSYSTEM DATA (RD SUB DATA)	3E	—
	SENSE SUBSYSTEM STATUS (SNS SUB STS)	54	—
	READ MESSAGE ID (RD MSG IDL)	4E	—

Note:

- Command Reject, format 0, and message 1 are issued for the commands that are not listed in this table.
- TEST I/O is a CPU instruction and cannot be specified directly. However, it appears as a command to the interface.
- TIC is a type of command but runs only on a channel. It will never be visible to the interface.

5.3 Comparison of pair status on SVP, Web Console, Raid Manager

Table.5.3-1 Comparison of pair status on SVP, Web Console, Raid Manager

NO	Event	Status on Raid Manager	Status on SVP, Web Console
1	Simplex Volume	P-VOL: SMPL S-VOL: SMPL	P-VOL: SMPL S-VOL: SMPL
2	Copying LUSE Volume Partly completed (SYNC only)	P-VOL: PDUB S-VOL: PDUB	P-VOL: PDUB S-VOL: PDUB
3	Copying Volume	P-VOL: COPY S-VOL: COPY	P-VOL: COPY S-VOL: COPY
4	Pair volume	P-VOL: PAIR S-VOL: PAIR	P-VOL: PAIR S-VOL: PAIR
5	Pairsplit operation to P-VOL	P-VOL: PSUS S-VOL: SSUS	P-VOL: PSUS (S-VOL by operator) S-VOL: PSUS (S-VOL by operator)
6	Pairsplit operation to S-VOL	P-VOL: PSUS S-VOL: PSUS	P-VOL: PSUS (S-VOL by operator) S-VOL: PSUS (S-VOL by operator)
7	Pairsplit -P operation *1 (P-VOL failure, SYNC only)	P-VOL: PSUS S-VOL: SSUS	P-VOL: PSUS (P-VOL by operator) S-VOL: PSUS (by MCU)
8	Pairsplit -R operation *1	P-VOL: PSUS S-VOL: SMPL	P-VOL: PSUS (Delete pair to RCU) S-VOL: SMPL
9	P-VOL Suspend (failure)	P-VOL: PSUE S-VOL: SSUS	P-VOL: PSUE (S-VOL failure) S-VOL: PSUE (S-VOL failure)
10	S-VOL Suspend (failure)	P-VOL: PSUE S-VOL: PSUE	P-VOL: PSUE (S-VOL failure) S-VOL: PSUE (S-VOL failure)
11	PS ON failure	P-VOL: PSUE S-VOL: —	P-VOL: PSUE (MCU IMPL) S-VOL: —
12	Copy failure (P-VOL failure)	P-VOL: PSUE S-VOL: SSUS	P-VOL: PSUE (Initial copy failed) S-VOL: PSUE (Initial copy failed)
13	Copy failure (S-VOL failure)	P-VOL: PSUE S-VOL: PSUE	P-VOL: PSUE (Initial copy failed) S-VOL: PSUE (Initial copy failed)
14	Suspending volume (ASYNCR only)	P-VOL: COPY or PAIR or PSUE S-VOL: COPY or PAIR or PSUE	P-VOL: Suspending S-VOL: Suspending
15	Deleting volume (ASYNCR only)	P-VOL: COPY or PAIR or SUS S-VOL: COPY or PAIR or SUS	P-VOL: Deleting S-VOL: Deleting
16	RCU accepted the notification of MCU's P/S-OFF	P-VOL: — S-VOL: SSUS	P-VOL: — S-VOL: PSUE (MCU P/S OFF)
17	Sidefile overload (under margin, ASYNCR only)	P-VOL: PFUL S-VOL: PAIR	P-VOL: PAIR S-VOL: PAIR
18	Sidefile overload Suspend (over margin, ASYNCR only)	P-VOL: PFUS S-VOL: PFUS	P-VOL: PSUS (Sidefile Overflow) S-VOL: PSUS (Sidefile Overflow)

*1: Operation on Raid Manager