



**Hitachi Freedom Storage™
Lightning 9900™
User and Reference Guide**

© 2003 Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

Notice: No part of this publication may be reproduced or transmitted in any form or by any electronic or mechanical means, including photocopying and recording, or stored in a database or retrieval system for any purpose, without the express written permission of Hitachi Data Systems Corporation.

Hitachi Data Systems reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements, including license agreements. All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.

Trademarks

Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems design mark is a trademark and service mark of Hitachi, Ltd.

Hi-Track is a registered trademark of Hitachi Data Systems Corporation.

Extended Serial Adapter, ExSA, Hitachi Freedom Storage, Hitachi Graph-Track, and Lightning 9900 are trademarks of Hitachi Data Systems Corporation.

APC and Symmetra are trademarks or registered trademarks of American Power Conversion Corporation.

HARBOR is a registered trademark of BETA Systems Software AG.

AIX, DYNIX/ptx, ESCON, FICON, IBM, MVS, MVS/ESA, VM/ESA, and S/390 are registered trademarks or trademarks of International Business Machines Corporation.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

Tantia is a trademark of Tantia Technologies Inc. Tantia Technologies is a wholly owned subsidiary of BETA Systems Software AG of Berlin.

All other brand or product names are or may be registered trademarks, trademarks or service marks of and are used to identify products or services of their respective owners.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

Revision	Date	Description
MK-90RD008-0	July 2000	Initial Release
MK-90RD008-1	November 2000	Revision 1, supersedes and replaces MK-90RD008-0
MK-90RD008-2	March 2001	Revision 2, supersedes and replaces MK-90RD008-1 This document revision applies to 9900 microcode versions 01-12-xx.
MK-90RD008-3	June 2001	Revision 3, supersedes and replaces MK-90RD008-2 This document revision applies to 9900 microcode versions 01-13-xx.
MK-90RD008-4	January 2002	Revision 4, supersedes and replaces MK-90RD008-3 This document revision applies to 9900 microcode versions 01-16-xx.
MK-90RD008-5	February 2002	Revision 5, supersedes and replaces MK-90RD008-4 This document revision applies to 9900 microcode versions 01-17-xx.
MK-90RD008-6	May 2002	Revision 6, supersedes and replaces MK-90RD008-5 This document revision applies to 9900 microcode versions 01-17-yy.
MK-90RD008-7	October 2003	Revision 7, supersedes and replaces MK-90RD008-6 This revision applies to 9900 microcode versions 01-18-67 and higher.

Source Documents for this Revision

- DKC410I/405I Disk Subsystem Maintenance Manual, revision 12.1 (August 2003).
- RAID400_RAID450_Public_Mode_Rev.2a (April 29, 2003).

Changes in this Revision

- Updated description of disk drive and cache upgrades to remove the statement that all upgrades can be made with minimal impact (section 1.1.7).
- Added information on the 146-GB hard disk drive (sections 2.3, 2.4.2, 2.4.3; Table 2.2, Table 2.3, Table 5.8, Table 5.9, Table 5.18).
- Added information on the public system option modes (new section 3.5, new Table 3.2-Table 3.8).

Preface

This document describes the physical, functional, and operational characteristics of the Hitachi Lightning 9900™ subsystem, provides general instructions for operating the 9900 subsystem, and provides the installation and configuration planning information for the 9900 subsystem.

This document assumes that:

- The user has a background in data processing and understands direct-access storage device (DASD) subsystems and their basic functions,
- The user is familiar with the S/390® (mainframe) operating systems and/or open-system platforms supported by the 9900 subsystem, and
- The user is familiar with the equipment used to connect RAID disk array subsystems to the supported host systems.

For further information on Hitachi Data Systems products and services, please contact your Hitachi Data Systems account team, or visit Hitachi Data Systems worldwide web site at <http://www.hds.com>. For specific information on supported host systems and platforms for the Lightning 9900™ subsystem, please refer to the 9900 user documentation for the platform, or contact the Hitachi Data Systems Support Center.

Note: Unless otherwise noted, the term “9900” refers to the entire Hitachi Lightning 9900™ subsystem family, including all models (e.g., 9960, 9910) and all configurations (e.g., all-mainframe, all-open, multiplatform).

Note: The use of Hitachi Data Systems products is governed by the terms of your license agreement(s) with Hitachi Data Systems.

Microcode Level

This document revision applies to 9900 microcode versions 01-18-67 and higher.

COMMENTS

Please send us your comments on this document: doc.comments@hds.com.

Make sure to include the document title, number, and revision.

Please refer to specific page(s) and paragraph(s) whenever possible.

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

Thank you!

Contents

Chapter 1 Overview of the Lightning 9900™ Subsystem

1.1	Key Features of the Lightning 9900™ Subsystem	1
1.1.1	Continuous Data Availability	2
1.1.2	Connectivity	2
1.1.3	S/390® Compatibility and Functionality	3
1.1.4	Open-Systems Compatibility and Functionality	3
1.1.5	Hitachi Freedom NAS™ and Hitachi Freedom SAN™	4
1.1.6	Program Products and Service Offerings	6
1.1.7	Subsystem Scalability	8
1.2	Reliability, Availability, and Serviceability	9

Chapter 2 Subsystem Architecture and Components

2.1	Overview	11
2.2	Components of the Controller Frame	14
2.2.1	Storage Clusters	14
2.2.2	Nonvolatile Shared Memory	15
2.2.3	Nonvolatile Duplex Cache	15
2.2.4	Multiple Data and Control Paths	16
2.2.5	Redundant Power Supplies	16
2.2.6	Client-Host Interface Processors (CHIPs) and Channels	16
2.2.7	Channels	18
2.2.8	Array Control Processors (ACPs)	19
2.3	Array Frame	21
2.3.1	Disk Array Groups	23
2.3.2	Sequential Data Striping	24
2.4	Intermix Configurations	25
2.4.1	RAID-1 & RAID-5 Intermix	25
2.4.2	Hard Disk Drive Intermix	25
2.4.3	Device Emulation Intermix	26
2.5	Service Processor (SVP)	27
2.6	Remote Console PC	27

Chapter 3 Functional and Operational Characteristics

3.1	New 9900 Features and Capabilities	29
3.2	I/O Operations	29
3.3	Cache Management	30
3.3.1	Algorithms for Cache Control	30
3.3.2	Write Pending Rate	30
3.4	Control Unit (CU) Images, LVIs, and LUs	31
3.4.1	CU Images	31
3.4.2	Logical Volume Image (LVIs)	31
3.4.3	Logical Unit (LU) Type	31
3.5	System Option Modes	32

3.6	Open Systems Features and Functions	37
3.6.1	Failover and SNMP Support	37
3.6.2	Share-Everything Architecture	37
3.6.3	SCSI Extended Copy Command Support	37
3.7	Data Management Functions	38
3.7.1	Hitachi TrueCopy (TC)	40
3.7.2	Hitachi TrueCopy - S/390® (TC390)	40
3.7.3	Hitachi ShadowImage (SI)	41
3.7.4	Hitachi ShadowImage - S/390® (SI390)	41
3.7.5	Command Control Interface (CCI)	42
3.7.6	Extended Copy Manager (ECM)	42
3.7.7	Hitachi Extended Remote Copy (HXRC)	43
3.7.8	Hitachi NanoCopy™	44
3.7.9	Data Migration	44
3.7.10	Hitachi RapidXchange (HRX)	45
3.7.11	Hitachi Multiplatform Backup/Restore (HMBR)	45
3.7.12	HARBOR® File-Level Backup/Restore	46
3.7.13	HARBOR® File Transfer	46
3.7.14	HiCommand™	46
3.7.15	LUN Manager	47
3.7.16	LU Size Expansion (LUSE)	47
3.7.17	Virtual LVI/LUN	47
3.7.18	FlashAccess	48
3.7.19	Cache Manager	48
3.7.20	Hitachi SANTinel	49
3.7.21	Hitachi SANTinel - S/390®	49
3.7.22	Prioritized Port and WWN Control (PPC)	50
3.7.23	Hitachi Parallel Access Volume (HPAV)	50
3.7.24	Dynamic Link Manager™ (DLM)	50
3.7.25	LDEV Guard	50
3.7.26	Hitachi CruiseControl	51
3.7.27	Hitachi Graph-Track™	52

Chapter 4 Configuring and Using the 9900 Subsystem

4.1	S/390® Configuration	53
4.1.1	Subsystem IDs (SSIDs)	53
4.2	S/390® Hardware Definition	54
4.2.1	Hardware Definition Using IOCP (MVS, VM, or VSE)	54
4.2.2	Hardware Definition Using HCD (MVS/ESA)	59
4.2.3	Defining the 9900 to VM/ESA® Systems	70
4.2.4	Defining the 9900 to TPF	70
4.3	S/390® Operations	71
4.3.1	Initializing the LVIs	71
4.3.2	Device Operations: ICKDSF	71
4.3.3	MVS Cache Operations	73
4.3.4	VM/ESA® Cache Operations	75
4.3.5	VSE/ESA Cache Operations	75

4.4	Open-Systems Configuration	76
4.4.1	Configuring the Fibre-Channel Ports	77
4.4.2	Virtual LVI/LUN Devices	77
4.4.3	LU Size Expansion (LUSE) Devices	77
4.5	Open Systems Operations	78
4.5.1	Command Tag Queuing	78
4.5.2	Host/Application Failover Support	78
4.5.3	Path Failover Support	79
4.5.4	Remote SIM (R-SIM) Reporting	80
4.5.5	SNMP Remote Subsystem Management	80
4.5.6	NAS and SAN Operations	80

Chapter 5 Planning for Installation and Operation

5.1	User Responsibilities	81
5.2	Electrical Specifications and Requirements for Three-Phase Subsystems	82
5.2.1	Internal Cable Diagram	82
5.2.2	Power Plugs	83
5.2.3	Features	84
5.2.4	Current Rating, Power Plug, Receptacle, and Connector for Three-Phase (60 Hz only)	84
5.2.5	Input Voltage Tolerances	85
5.3	Electrical Specifications and Requirements for Single-Phase Subsystems	86
5.3.1	Internal Cable Diagram	86
5.3.2	Power Plugs	88
5.3.3	Features	91
5.3.4	Current Rating, Power Plug, Receptacle, and Connector for Single-Phase (60 Hz only)	91
5.3.5	Input Voltage Tolerances	92
5.4	Dimensions and Weight	93
5.5	Floor Loading and Cable Routing Requirements	96
5.5.1	Service Clearance Requirements	96
5.5.2	Minimum Subsystem Disk Configuration	101
5.5.3	Floor Load Rating	105
5.5.4	Cable Requirements	109
5.6	Channel Specifications and Requirements	110
5.7	Environmental Specifications and Requirements	111
5.7.1	Temperature and Humidity Requirements	111
5.7.2	Power Consumption and Heat Output Specifications	111
5.7.3	Loudness	113
5.7.4	Air Flow Requirements	113
5.7.5	Vibration and Shock Tolerances	114

Chapter 6	Troubleshooting	
6.1	Troubleshooting	115
6.2	Service Information Messages (SIMs)	116
6.3	Calling the Hitachi Data Systems Support Center	117
Appendix A	Unit Conversions	119
Acronyms and Abbreviations	121

List of Figures

Figure 2.1	Lightning 9900™ HiStar Network (HSN) Architecture	11
Figure 2.2	9960 Subsystem Frames	13
Figure 2.3	9910 Subsystem Frame	13
Figure 2.4	Conceptual ACP Array Domain	19
Figure 2.5	Sample RAID-1 Layout	23
Figure 2.6	Sample RAID-5 Layout (Data Plus Parity Stripe)	24
Figure 2.7	Sample Hard Disk Drive Intermix	25
Figure 2.8	Sample Device Emulation Intermix	26
Figure 3.1	Fibre-Channel Device Addressing	31
Figure 4.1	IOCP Definition for FICON™ Channels (direct connect and via FICON™ switch)	55
Figure 4.2	IOCP Definition for 1024 LVIs (9900 connected to host CPU(s) via ESCD)	56
Figure 4.3	IOCP Definition for 1024 LVIs (9900 directly connected to CPU)	57
Figure 4.4	Master MENU (Step 1)	61
Figure 4.5	Basic HCD Panel (Step 2)	62
Figure 4.6	Define, Modify, or View Configuration Data (Step 3)	62
Figure 4.7	Control Unit List Panel (Step 4)	63
Figure 4.8	Add Control Unit Panel (Step 5)	63
Figure 4.9	Selecting the Operating System (Step 6)	64
Figure 4.10	Control Unit Chpid, CUADD, and Device Address Range Addressing (Step 7)	64
Figure 4.11	Select Processor / Control Unit Panel (Step 8)	65
Figure 4.12	Control Unit List (Step 9)	65
Figure 4.13	I/O Device List Panel (Step 10)	66
Figure 4.14	Add Device Panel (Step 11)	66
Figure 4.15	Device / Processor Definition Panel - Selecting the Processor ID (Step 12)	67
Figure 4.16	Define Device / Processor Panel (Step 13)	67
Figure 4.17	Device / Processor Definition Panel (Step 14)	68
Figure 4.18	Define Device to Operating System Configuration (Step 15)	68
Figure 4.19	Define Device Parameters / Features Panel (Step 16)	69
Figure 4.20	Update Serial Number, Description and VOLSER Panel (Step 18)	69
Figure 4.21	LVI Initialization for MVS: ICKDSF JCL	71
Figure 4.22	Displaying Cache Statistics Using MVS DFSMS	73
Figure 4.23	IDCAMS LISTDATA COUNTS (JCL example)	73
Figure 4.24	Fibre Port-to-LUN Addressing	77
Figure 4.25	Alternate Pathing	79
Figure 5.1	Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (Europe)	82
Figure 5.2	Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (USA)	83
Figure 5.3	Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (Europe)	83
Figure 5.4	Internal Cable Diagram of a Single-Phase 9960 Subsystem	86
Figure 5.5	Internal Cable Diagram of a Single-Phase 9910 Subsystem	87
Figure 5.6	Power Plugs for Single-Phase 9960 Controller (USA)	88
Figure 5.7	Power Plugs for Single-Phase 9910 Subsystem (USA)	88
Figure 5.8	Power Plugs for a Single-Phase 9960 Controller (Europe)	89
Figure 5.9	Power Plugs for a Single-Phase 9910 Subsystem (Europe)	89
Figure 5.10	Power Plugs for Single-Phase 9960 Disk Array Unit (USA)	90

Figure 5.11	Power Plugs for Single-Phase 9960 Disk Array Unit (Europe)	90
Figure 5.15	9960 DKC and DKU Physical Dimensions	93
Figure 5.16	9910 DKC and DKU Physical Dimensions	93
Figure 5.17	9960 Controller Frame Service Clearance and Cutouts (millimeters)	97
Figure 5.18	9960 Controller Frame Service Clearance and Cutouts (inches)	98
Figure 5.19	9960 Disk Array Unit Service Clearance and Cutout (millimeters)	99
Figure 5.20	9960 Disk Array Unit Service Clearance and Cutout (inches)	100
Figure 5.21	9960 Disk Subsystem Minimum Configuration (millimeters)	101
Figure 5.22	9960 Disk Subsystem Minimum Configuration (inches)	102
Figure 5.23	9910 Disk Subsystem All Configurations (millimeters)	103
Figure 5.24	9910 Disk Subsystem All Configurations (inches)	104
Figure 5.25	9960 Disk Subsystem with Controller and 4 Disk Arrays	106
Figure 5.26	9960 Disk Subsystem with Maximum Configuration	107
Figure 6.1	Typical 9900 SIM Showing Reference Code and SIM Type.	116

List of Tables

Table 1.1	Program Products and Service Offerings	6-7
Table 2.1	CHIP and Channel Specifications	17
Table 2.2	ACP Specifications	20
Table 2.3	Disk Drive Specifications	22
Table 3.1	Capacities of Standard LU Types	31
Table 3.2	Common System Option Modes	32
Table 3.3	System Option Modes for Mainframe Connectivity	32
Table 3.4	System Option Modes for Open-System Connectivity	33
Table 3.5	System Option Modes for ShadowImage - S/390® and ShadowImage	34
Table 3.6	System Option Modes for TrueCopy - S/390® Sync & Async	34-35
Table 3.7	System Option Modes for HXRC	36
Table 3.8	System Option Modes for Concurrent Copy (CC)	36
Table 3.9	Data Management Functions for Open-System Users	38
Table 3.10	Data Management Functions for S/390® Users	39
Table 4.1	SSID Requirements	53
Table 4.2	Correspondence between Physical Paths and Channel Interface IDs (CI 1)	58
Table 4.3	Correspondence between Physical Paths and Channel Interface IDs (CI 2)	58
Table 4.4	HCD Definition for 64 LVIs	59
Table 4.5	HCD Definition for 256 LVIs	59
Table 4.6	ICKDSF Commands for 9900 Contrasted to RAMAC	72
Table 4.7	9900 Open-System Platforms and Configuration Guides	76
Table 5.1	9960 Three-Phase Features	84
Table 5.2	Current Rating, Power Plug, Receptacle, and Connector for 3-Phase 9960	84
Table 5.3	Input Voltage Specifications for Three-Phase Power	85
Table 5.4	9900 Single-Phase Features	91
Table 5.5	Current Rating, Power Plug, Receptacle, and Connector for 1-Phase 9900	91

Table 5.6	Input Voltage Specifications for Single-Phase Power	92
Table 5.7	9900 Physical Specifications	94
Table 5.8	9960 Frame and Component Weights	95
Table 5.9	9960 & 9910 Subsystem Weights	95
Table 5.10	Floor Load Rating for 9910 Subsystem	105
Table 5.11	Floor Load Rating for 9960 Controller with 1 Disk Array	105
Table 5.12	Floor Load Rating for 9960 Controller with 4 Disk Arrays	107
Table 5.13	Floor Load Rating for 9960 Controller with Maximum Configuration	108
Table 5.14	Cable Requirements	109
Table 5.15	ESCON [®] and FICON [™] Port Information	110
Table 5.16	Fibre-Channel Port Information	110
Table 5.17	Temperature and Humidity Requirements	111
Table 5.18	9910/9960 Component Power and Heat Output Specifications	112
Table 5.19	9900 Subsystem Power and Heat Output Specifications	112
Table 5.20	Internal Air Flow	113
Table 5.21	Vibration and Shock Tolerances	114
Table 6.1	Troubleshooting	115
Table A.1	Unit Conversions for Standard (U.S.) and Metric Measures	119

Chapter 1 Overview of the Lightning 9900™ Subsystem

1.1 Key Features of the Lightning 9900™ Subsystem

The Hitachi Lightning 9900™ subsystem provides high-speed response, continuous data availability, scalable connectivity, and expandable capacity for both S/390® and open-systems environments. The 9900 subsystem is designed for use in 7×24 data centers that demand high-performance, non-stop operation. The 9900 subsystem is compatible with industry-standard software and supports concurrent attachment to multiple host systems and platforms. The 9900 subsystem employs and improves upon the key characteristics of generations of successful Hitachi disk storage subsystems to achieve high performance and reliability. The advanced components, functions, and features of the Lightning 9900™ subsystem represent an integrated approach to data retrieval and storage management.

The Lightning 9900™ subsystem provides many new benefits and advantages for the user. The 9900 subsystem can operate with multihost applications and host clusters, and is designed to handle very large databases as well as data warehousing and data mining applications that store and retrieve terabytes of data. The Lightning 9900™ provides up to 32 host interface ports and can be configured for all-mainframe, all-open, or multiplatform operations.

- **Instant access to data around the clock:**
 - 100 percent data availability guarantee.
 - No single point of failure.
 - Highly resilient multi-path fibre architecture.
 - Fully redundant, hot-swappable components.
 - Global dynamic hot sparing.
 - Duplexed write cache with battery backup.
 - Hi-Track® “call-home” maintenance system.
 - Non-disruptive microcode updates.
 - RAID-1 and/or RAID-5 array groups within the same subsystem.
- **Unmatched performance and capacity:**
 - Industry’s only internal switched fabric architecture.
 - Multiple point-to-point data and control paths.
 - Up to 6.4-GB/sec internal system bandwidth.
 - Fully addressable 32-GB data cache; separate control cache.
 - Extremely fast and intelligent cache algorithms.
 - Non-disruptive expansion to over 88 TB raw capacity.
 - Simultaneous transfers from up to 32 separate hosts.
 - High-throughput 10K RPM fibre-channel, dual-active disk drives.

- **Extensive connectivity and resource sharing:**
 - Concurrent operation of UNIX[®]-based, Windows NT[®], Windows[®] 2000, Linux[®], NetWare[®], and S/390[®] host systems.
 - Fibre-channel, Fiber Connection (FICON[™]), and Extended Serial Adapter[™] (ESCON[®]) server connections.
 - Optimized for storage-area networks (SANs), fibre-channel switched, fibre-channel arbitrated loop, and point-to-point configurations.

1.1.1 Continuous Data Availability

The Hitachi Lightning 9900[™] is designed for nonstop operation and continuous access to all user data. To achieve nonstop customer operation, the 9900 subsystem accommodates online feature upgrades and online software and hardware maintenance. See section 1.2 for further information on the reliability and availability features of the Lightning 9900[™] subsystem.

1.1.2 Connectivity

The Hitachi Lightning 9900[™] RAID subsystem supports concurrent attachment to S/390[®] mainframe hosts and open-system (UNIX[®]-based and/or PC-server) platforms. The 9900 subsystem can be configured with FICON[™] ports, Extended Serial Adapter[™] (ExSA[™]) ports (compatible with ESCON[®] protocol), and/or fibre-channel ports to support all-mainframe, all-open, and multiplatform configurations.

When FICON[™] channel interfaces are used, the 9900 subsystem can provide up to 16 logical control unit (CU) images and 4096 logical device (LDEV) addresses. Each physical FICON[™] channel interface supports up to 512 logical paths providing a maximum of 8192 logical paths per subsystem. FICON[™] connection provides transfer rates of up to 100 MB/sec (1Gbps).

When ExSA[™] channel interfaces are used, the 9900 subsystem can provide up to 16 logical control unit (CU) images and 4,096 logical device (LDEV) addresses. Each physical ExSA[™] channel interface supports up to 256 logical paths providing a maximum of 8,192 logical paths per subsystem. ExSA[™] connection provides transfer rates of up to 17 MB/sec.

When fibre-channel interfaces are used, the 9900 subsystem can provide up to 32 ports for attachment to UNIX[®]-based and/or PC-server platforms. The type of host platform determines the number of logical units (LUs) that may be connected to each port. Fibre-channel connection provides data transfer rates of up to 200 MB/sec (2 Gbps). The 9900 subsystem supports fibre-channel arbitrated loop (FC-AL) and fabric fibre-channel topologies as well as high-availability (HA) fibre-channel configurations using hubs and switches.

1.1.3 S/390® Compatibility and Functionality

The 9900 subsystem supports 3990 and 2105 controller emulations and can be configured with multiple concurrent logical volume image (LVI) formats, including 3390-1, -2, -3, -3R, -9 and 3380-E, -J, -K. In addition to full System-Managed Storage (SMS) compatibility, the 9900 subsystem also provides the following functionality in the S/390® environment:

- Sequential data striping,
- Cache fast write (CFW) and DASD fast write (DFW),
- Enhanced dynamic cache management,
- Multiple Allegiance support,
- Concurrent Copy (CC) support,
- Enhanced CCW support,
- Priority I/O queuing,
- Parallel Access Volume (PAV) support, and
- Transaction Processing Facility (TPF)/Multi-Path Locking Facility (MPLF) support.

1.1.4 Open-Systems Compatibility and Functionality

The Lightning 9900™ subsystem supports multiple concurrent attachment to a variety of host operating systems (OS). The 9900 supports the following platforms at this time. The type of host platform determines the number of logical units (LUs) that may be connected to each port. Please contact Hitachi Data Systems for the latest information on platform and OS version support. The 9900 is compatible with most fibre-channel host bus adapters (HBAs).

- | | |
|---------------------------|-------------------------------|
| ■ IBM® AIX® OS | ■ Microsoft® Windows NT® OS |
| ■ Sun™ Solaris™ OS | ■ Microsoft® Windows® 2000 OS |
| ■ HP-UX® OS | ■ Novell® NetWare® OS |
| ■ Compaq® Tru64™ UNIX® OS | ■ Red Hat® Linux® OS |
| ■ Sequent® DYNIX/ptx® OS | ■ Compaq® OpenVMS® OS |
| ■ SGI™ IRIX® OS | |

The 9900 subsystem provides enhanced dynamic cache management and supports command tag queuing and multi-initiator I/O. Command tag queuing (see section 4.5.1) enables hosts to issue multiple disk commands to the fibre-channel adapter without having to serialize the operations. The 9900 subsystem operates with industry-standard middleware products providing application/host failover capability, I/O path failover support, and logical volume management. The 9900 subsystem also supports the industry-standard simple network management protocol (SNMP) for remote management from the open-system host.

The 9900 subsystem can be configured with multiple concurrent logical unit (LU) formats (e.g., OPEN-3, -8, -9, -K, -E, -L, -M). The user can also configure custom-size volumes using the Virtual LVI/LUN and LU Size Expansion (LUSE) features of the 9900 subsystem, which are described in the next section.

1.1.5 Hitachi Freedom NAS™ and Hitachi Freedom SAN™

Hitachi Freedom Data Networks™ (FDN) provide an open architecture that offers organizations freedom of choice in deploying data access, protection, and sharing capabilities across the enterprise. Using multiple technologies and solutions such as storage-area networks (SANs) and network-attached storage (NAS), FDN builds, leverages, and augments storage infrastructures, providing access to any data from any computer, anytime and anywhere.

Hitachi Freedom NAS™ and Hitachi Freedom SAN™ solutions are the core offerings behind the FDN approach. They complement the Hitachi Freedom Storage™ subsystems by allowing more flexibility than ever in heterogeneous environments. While SAN architectures respond to high bandwidth needs, NAS addresses the need for rapid file access, especially critical for e-business applications. Hitachi Data Systems offers the best of both.

FDN encompasses storage, switches and hubs, servers/clients, management software, protocols, services, and networks developed by Hitachi, our alliance partners, and third party providers. FDN facilitates consolidation of server and storage resources, data sharing among heterogeneous hosts, centralized resource and data management, superior data security, and increased connectivity.

Hitachi Freedom SAN™. Hitachi Data Systems' SAN solutions give you the freedom to locate storage wherever needed and protect your investment in currently installed components. Made possible by the advent and proliferation of high-speed fibre-channel technology, SANs break the traditional server/storage bond and enable total connectivity. As a result, you can add, remove, or reassign any resource without interfering with ongoing business operations.

The Lightning 9900™ subsystem features unparalleled reliability, a SAN-ready architecture, and support for S/390®, UNIX®, and Windows NT® platforms. Hitachi adds software and services to SAN components to provide functionality such as LAN-free backup, remote copy, and multiplatform data exchange from our Freedom Storage™ software suites.

Hitachi Freedom NAS™. Freedom NAS answers the need for speed with faster file access. Numerous clients can instantly share data with information available on your NAS file server. Freedom NAS is an excellent solution for file/web serving, document/record imaging, streaming media, video design, telco call centers, and manufacturing.

Hitachi Freedom Storage™ subsystems are combined with Network Storage Solutions' NAS file servers to provide Freedom NAS solutions. The modular architecture of Hitachi Freedom Storage™ subsystems provides quick and easy storage expansion. For further information, please refer to the *Hitachi Freedom NAS™ NSS Configuration Guide*, MK-91RD053.

Freedom NAS provides the following benefits for the user:

- Accelerates response times
- Supports rapid deployment of new applications
- Satisfies increasing customer demand
- Enables expanding operations
- Leverages existing storage infrastructure
- Improves service levels
- Reduces I/O bottlenecks
- Minimizes overhead through consolidation and reduced complexity
- Increases availability and reliability
- Eliminates storage islands
- Installs quickly and easily

1.1.6 Program Products and Service Offerings

The Lightning 9900™ subsystem provides many advanced features and functions that increase data accessibility, enable continuous user data access, and deliver enterprise-wide coverage of on-line data copy/relocation, data access/protection, and storage resource management. Hitachi Data Systems' software solutions provide a full complement of industry-leading copy, availability, resource management, and exchange software to support business continuity, database backup/restore, application testing, and data mining.

Table 1.1 Program Products and Service Offerings (continues on the next page)

Function	Description	See Section:
Data Replication:		
Hitachi TrueCopy (TC) Hitachi TrueCopy – S/390® (TC390)	Enables the user to perform remote copy operations between 9900 subsystems (and 7700E and 7700) in different locations. Hitachi TrueCopy provides synchronous and asynchronous copy modes for both S/390® and open-system data. (The 7700 subsystem supports only synchronous remote copy operations.)	3.7.1 3.7.2
Hitachi ShadowImage (SI) Hitachi ShadowImage – S/390® (SI390)	Allows the user to create internal copies of volumes for a wide variety of purposes including application testing and offline backup. Can be used in conjunction with TrueCopy to maintain multiple copies of critical data at both the primary and secondary sites.	3.7.3 3.7.4
Command Control Interface (CCI)	Enables open-system users to perform TrueCopy and ShadowImage operations by issuing commands from the host to the 9900 subsystem. The CCI software supports scripting and provides failover and mutual hot standby functionality in cooperation with host failover products.	3.7.5
Extended Copy Manager (ECM)	Provides server-free backup solutions between the 9900 and backup devices (e.g., tape, disk) in SAN environments. Supports the SCSI Extended Copy command issued from the host server to the 9900 subsystem.	3.7.6
Hitachi Extended Remote Copy (HXRC)	Provides compatibility with the IBM® Extended Remote Copy (XRC) S/390® host software function, which performs server-based asynchronous remote copy operations for mainframe LVIs.	3.7.7
Hitachi NanoCopy™	Enables S/390® users to make Point-in-Time (PiT) copies of production data, without quiescing the application or causing any disruption to end-user operations, for such uses as application testing, business intelligence, and disaster recovery for business continuance.	3.7.8
Data migration (service offering only)	Enables the rapid transfer of data from other disk subsystems onto the 9900 subsystem. Data migration operations can be performed while applications are online using the data which is being transferred.	3.7.9
Backup/Restore and Data Sharing:		
Hitachi RapidXchange (HRX)	Enables users to transfer data between S/390® and open-system platforms using the ExSA™ and/or FICON™ channels, which provides high-speed data transfer without requiring network communication links or tape.	3.7.10
Hitachi Multiplatform Backup/Restore (HMBR)	Allows users to perform mainframe-based volume-level backup and restore operations on the open-system data stored on the multiplatform 9900 subsystem.	3.7.11
HARBOR® File-Level Backup/Restore	Enables users to perform mainframe-based file-level backup/restore operations on the open-system data stored on the multiplatform 9900 subsystem.	3.7.12
HARBOR® File Transfer	Enables users to transfer large data files at ultra-high channel speeds in either direction between open systems and mainframe servers.	3.7.13

Table 1.1 Program Products and Service Offerings (continued)

Function	Description	See Section:
Resource Management:		
HiCommand™	Enables users to manage the 9900 subsystem and perform functions (e.g., LUN Manager, SANtinel) from virtually any location via the HiCommand™ Web Client, command line interface (CLI), and/or third-party application.	3.7.14
LUN Manager	Enables users to configure the 9900 fibre-channel ports for operational environments (e.g., arbitrated-loop (FC-AL) and fabric topologies, host failover support).	3.7.15
LU Size Expansion (LUSE)	Enables open-system users to create expanded LUs which can be up to 36 times larger than standard fixed-size LUs.	3.7.16
Virtual LVI (VLVI) Virtual LUN (VLUN)	Enables users to configure custom-size LVIs and LUs which are smaller than standard-size devices.	3.7.17
FlashAccess (Flsh)	Enables users to store specific high-usage data directly in cache memory to provide virtually immediate data availability.	3.7.18
Cache Manager	Enables users to perform FlashAccess operations from the S/390® host system. FlashAccess allows you to place specific data in cache memory to enable virtually immediate access to this data	3.7.19
Hitachi SANtinel Hitachi SANtinel – S/390®	Allows users to restrict host access to data on the Lightning 9900™ subsystem. Open-system users can restrict host access to LUs based on the host's World Wide Name (WWN). S/390® mainframe users can restrict host access to LVIs based on node IDs and logical partition (LPAR) numbers.	3.7.20 3.7.21
Prioritized Port Control (PPC)	Allows open-system users to designate prioritized ports (e.g., for production servers) and non-prioritized ports (e.g., for development servers) and set thresholds and upper limits for the I/O activity of these ports.	3.7.22
Hitachi Parallel Access Volume (HPAV)	Enables the S/390® host system to issue multiple I/O requests in parallel to single LDEVs in the Lightning 9900™ subsystem. HPAV provides compatibility with the IBM® Workload Manager (WLM) host software function and supports both static and dynamic PAV functionality.	3.7.23
Dynamic Link Manager™	Provides automatic load balancing, path failover, and recovery capabilities in the event of a path failure.	3.7.24
LDEV Guard	Enables the assigning of access permissions (Read/Write, Read-Only, and Protect) to logical volumes in a disk subsystem. 3390-3A, 3390-3B and 3390-3C volumes can be used by both mainframe hosts and open-system hosts. Note: Please check with your Hitachi Data Systems representative for the latest feature availability.	MK-92RD072
Storage Utilities:		
Hitachi CruiseControl	Monitors subsystem and volume activity and performs automatic relocation of volumes to optimize performance.	3.7.26
Hitachi Graph-Track™ (GT)	Provides detailed information on the I/O activity and hardware performance of the 9900 subsystem. Hitachi Graph-Track™ displays real-time and historical data in graphical format, including I/O statistics, cache statistics, and front-end and back-end microprocessor usage.	3.7.27

1.1.7 Subsystem Scalability

The architecture of the 9900 subsystem accommodates scalability to meet a wide range of capacity and performance requirements. The 9960 storage capacity can be increased from a minimum of 54 GB to a maximum of 88 TB of user data. The 9960 nonvolatile cache can be configured from 1 GB to 32 GB. All disk drive and cache upgrades can be performed without interrupting user access to data.

The 9900 subsystem can be configured with the desired number and type of front-end client-host interface processors (CHIPs). The CHIPs are installed in pairs, and each CHIP pair offers up to eight host connections. The 9960 can be configured with four CHIP pairs to provide up to 32 paths to attached host processors. The 9910 supports up to three CHIP pairs and 24 paths.

The ACPs are the back-end processors which transfer data between the disk drives and cache. Each ACP pair is equipped with eight device paths. The 9960 subsystem can be configured with up to four pairs of array control processors (ACPs), providing up to thirty-two concurrent data transfers to and from the disk drives. The 9910 is configured with one ACP pair.

1.2 Reliability, Availability, and Serviceability

The Lightning 9900™ subsystem is not expected to fail in any way that would interrupt user access to data. The 9900 can sustain multiple component failures and still continue to provide full access to all stored user data. **Note:** While access to user data is never compromised, the failure of a key component can degrade performance.

The reliability, availability, and serviceability features of the 9900 subsystem include:

- **Full fault-tolerance.** The 9900 subsystem provides full fault-tolerance capability for all critical components. The disk drives are protected against error and failure by enhanced RAID technologies and dynamic scrubbing and sparing. The 9900 uses component and function redundancy to provide full fault-tolerance for all other subsystem components (microprocessors, control storage, power supplies, etc.). The 9900 has no active single point of component failure and is designed to provide continuous access to all user data.
- **Separate power supply systems.** Each storage cluster is powered by a separate set of power supplies. Each set can provide power for the entire subsystem in the unlikely event of power supply failure. The power supplies of each set can be connected across power boundaries, so that each set can continue to provide power if a power outage occurs. The 9900 can sustain the loss of multiple power supplies and still continue operation.
- **Dynamic scrubbing and sparing for disk drives.** The 9900 uses special diagnostic techniques and dynamic scrubbing to detect and correct disk errors. Dynamic sparing is invoked automatically if needed. The 9960 can be configured with up to sixteen spare disk drives, and any spare disk can back up any other disk of the same capacity, even if the failed disk and spare disk are in different array domains (attached to different ACP pairs).
- **Dynamic duplex cache.** The 9900 cache is divided into two equal segments on separate power boundaries. The 9900 places all write data in both cache segments with one internal write operation, so the data is always duplicated (duplexed) across power boundaries. If one copy of write data is defective or lost, the other copy is immediately destaged to disk. This duplex design ensures full data integrity in the event of a cache or power failure.
- **Remote copy features.** The Hitachi TrueCopy and Hitachi Extended Remote Copy (HXRC) data movement features enable the user to set up and maintain duplicate copies of S/390® and open-system data over extended distances. In the event of a system failure or site disaster, the secondary copy of data can be invoked rapidly, allowing applications to be recovered with guaranteed data integrity.

- **Hi-Track®.** The Hi-Track® maintenance support tool monitors the operation of the 9900 subsystem at all times, collects hardware status and error data, and transmits this data via modem to the Hitachi Data Systems Support Center. The Hitachi Data Systems Support Center analyzes the data and implements corrective action when necessary. In the unlikely event of a component failure, Hi-Track® calls the Hitachi Data Systems Support Center immediately to report the failure without requiring any action on the part of the user. Hi-Track® enables most problems to be identified and fixed prior to actual failure, and the advanced redundancy features enable the subsystem to remain operational even if one or more components fail. **Note:** Hi-Track® does not have access to any user data stored on the 9900 subsystem. The Hi-Track® tool requires a dedicated RJ-11 analog phone line.
- **Nondisruptive service and upgrades.** All hardware upgrades can be performed nondisruptively during normal subsystem operation. All hardware subassemblies can be removed, serviced, repaired, and/or replaced nondisruptively during normal subsystem operation. All microcode upgrades can be performed during normal subsystem operations using the SVP or the alternate path facilities of the host.
- **Error Reporting.** The Lightning 9900™ subsystem reports service information messages (SIMs) to notify users of errors and service requirements. SIMs can also report normal operational changes, such as remote copy pair status change. The SIMs are logged on the 9900 service processor (SVP) and on the Remote Console PC, reported directly to the mainframe and open-system hosts, and reported to Hitachi Data Systems via Hi-Track®.

Chapter 2 Subsystem Architecture and Components

2.1 Overview

Figure 2.1 shows the Hierarchical Star Network (HiStar or HSN) architecture of the Lightning 9900™ RAID subsystem. The “front end” of the 9900 subsystem includes the hardware and software that transfers the host data to and from cache memory, and the “back end” includes the hardware and software that transfers data between cache memory and the disk drives.

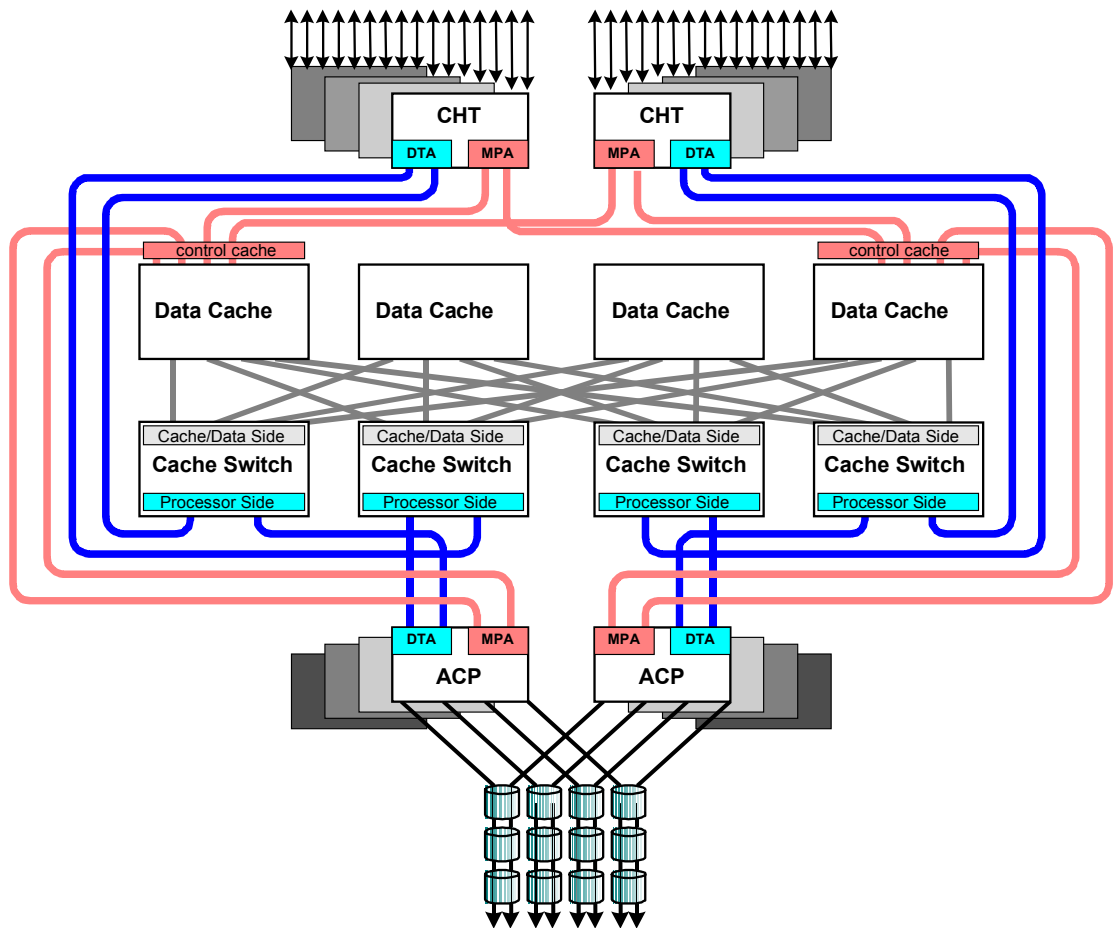


Figure 2.1 Lightning 9900™ HiStar Network (HSN) Architecture

Front End: The 9900 front end is entirely resident in the 9900 controller frame and includes the client-host interface processors (CHIPs) that reside on the channel adapter (CHA or CHT) boards. The CHIPs control the transfer of data to and from the host processors via the fibre-channel, ExSA™, and/or FICON™ channel interfaces and to and from cache memory via independent high-speed paths through the cache switches (CSWs).

- Each channel adapter board (CHA or CHT) can contain two or four CHIPs. The 9960 subsystem supports up to eight CHAs for a maximum of 32 host interfaces, and the 9910 subsystem supports up to six CHAs to provide a maximum of 24 host interfaces.
- The 9960 controller contains four cache switch (CSW) cards, and the 9910 controller contains two CSW cards.
- Cache memory in the 9960 resides on two or four cards depending on features, and each cache card is backed up by a separate battery. The 9910 supports two cache cards.
- Shared memory resides on the first two cache cards and is provided with its own power sources and backup batteries. Shared memory also has independent address and data paths from the channel adapter and disk adapter boards.

Back End: The 9900 back end is controlled by the array control processors (ACPs) that reside on the disk adapter boards in the 9900 controller frame. The ACPs control the transfer of data to and from the disk arrays via high-speed fibre (100 MB/sec or 1 Gbps) and then to and from cache memory via independent high-speed paths through the CSWs.

- The disk adapter board (DKA) contains four ACPs. The 9960 subsystem supports up to eight DKAs for a maximum of 32 ACPs. The 9910 subsystem supports two DKAs for a maximum of eight ACPs.

The 9960 subsystem (see Figure 2.2) includes the following major components:

- One controller frame containing the control and operational components of the subsystem.
- Up to six disk array frames containing the storage components (disk drive arrays) of the subsystem.
- The service processor (SVP) (see section 2.5). The 9900 SVP is located in the controller frame and can only be used by authorized Hitachi Data Systems personnel.
- The Remote Console PC (see section 2.6). The 9900 Remote Console PC can be attached to multiple 9960 and/or 9910 subsystems via the 9900-internal local-area network (LAN).

The 9910 subsystem (see Figure 2.3) includes the following major components:

- One frame containing the controller and disk components of the subsystem.
- The service processor (SVP) (see section 2.5). The 9900 SVP is located in the controller frame and can only be used by authorized Hitachi Data Systems personnel.
- The Remote Console PC (see section 2.6). The 9900 Remote Console PC can be attached to multiple 9960 and/or 9910 subsystems.

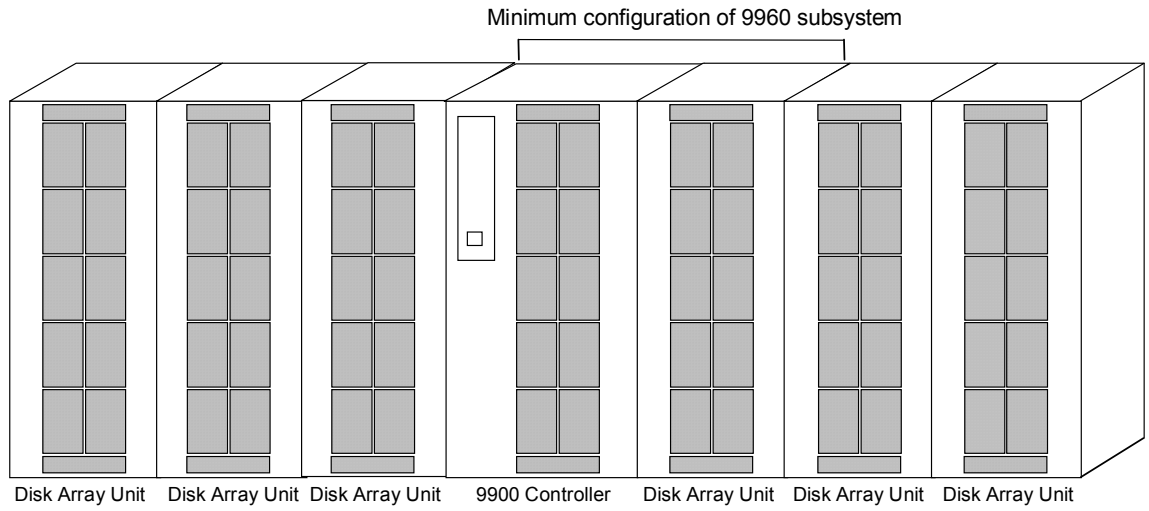


Figure 2.2 9960 Subsystem Frames




Figure 2.3 9910 Subsystem Frame

2.2 Components of the Controller Frame

The 9900 controller frame contains the control and operational components of the subsystem. For the 9910 subsystem, the controller frame also contains the disk array components. The 9900 controller is fully redundant and has no active single point of failure. All controller frame components can be repaired or replaced without interrupting access to user data. The key features and components of the controller frame are:

- Storage clusters (see section 2.2.1),
- Nonvolatile duplex shared memory (see section 2.2.2),
- Nonvolatile duplex cache memory (see section 2.2.3),
- Multiple data and control paths (see section 2.2.4),
- Redundant power supplies (see section 2.2.5),
- CHiPs and channels (FICON™, ExSA™, and/or fibre-channel) (see section 2.2.6),
- ACPs (see section 2.2.8).

2.2.1 Storage Clusters

Each controller frame consists of two redundant controller halves called storage clusters. Each storage cluster contains all physical and logical elements (e.g., power supplies, CHAs, CHiPs, ACPs, cache, control storage) needed to sustain processing within the subsystem. Both storage clusters should be connected to each host using an alternate path scheme, so that if one storage cluster fails, the other storage cluster can continue processing for the entire subsystem.

Each pair of channel adapters is split between clusters to provide full backup for both front-end and back-end microprocessors. Each storage cluster also contains a separate, duplicate copy of cache and shared memory contents. In addition to the high-level redundancy that this type of storage clustering provides, many of the individual components within each storage cluster contain redundant circuits, paths, and/or processors to allow the storage cluster to remain operational even with multiple component failures. Each storage cluster is powered by its own set of power supplies, which can provide power for the entire subsystem in the unlikely event of power supply failure. Because of this redundancy, the Lightning 9900™ subsystem can sustain the loss of multiple power supplies and still continue operation.

Note: The redundancy and backup features of the Lightning 9900™ subsystem eliminate all active single points of failure, no matter how unlikely, to provide an additional level of reliability and data availability.

2.2.2 Nonvolatile Shared Memory

The nonvolatile shared memory contains the cache directory and configuration information for the 9900 subsystem. The path group arrays (e.g., for dynamic path selection) also reside in the shared memory. The shared memory is duplexed, and each side of the duplex resides on the first two cache cards, which are in clusters 1 and 2. Even though the shared memory resides on the cache cards, the shared memory has separate power supplies and separate battery backup. The basic size of the shared memory is 512 MB, and the maximum size is 1.5 GB (for 9960). The size of the shared memory storage is determined by the total cache size and the number of logical devices (LDEVs). Any required increase beyond the base size is automatically shipped and configured during the upgrade process. The shared memory is protected by battery backup.

2.2.3 Nonvolatile Duplex Cache

The 9960 subsystem can be configured with up to 32 GB of cache, and the 9910 can be configured with up to 16 GB of cache. All cache memory in the 9900 is nonvolatile, and each cache card is protected by its own 48-hour battery backup. The cache in the 9900 is divided into two equal areas (called cache A and cache B) on separate cards. Cache A is in cluster 1, and cache B is in cluster 2. The 9900 places all read and write data in cache. Write data is normally written to both cache A and B with one CHIP write operation, so that the data is always duplicated (duplexed) across logic and power boundaries. If one copy of write data is defective or lost, the other copy is immediately destaged to disk. This “duplex cache” design ensures full data integrity in the unlikely event of a cache memory or power-related failure.

Note: Mainframe hosts can specify special attributes (e.g., cache fast write (CFW) command) to write data (typically a sort command) without write duplexing. This data is not duplexed and is usually given a discard command at the end of the sort, so that the data will not be destaged to the disk drives. See section 4.3.3 for further information on S/390® cache operations.

2.2.4 Multiple Data and Control Paths

The 9900 subsystem uses a state-of-the-art architecture called the Hierarchical Star (HiStar) Network (HSN) which utilizes multiple point-to-point data and command paths in order to provide redundancy and improve performance. Each data and command path is independent. The individual paths between the channel or disk adapters and cache are steered by high-speed cache switch cards. The 9900 does not have any common buses, thus eliminating the performance degradation and contention that can occur in a bus architecture. All data stored on the 9900 subsystem is moved into and out of cache via the redundant high-speed paths.

2.2.5 Redundant Power Supplies

Each storage cluster is powered by its own set of redundant power supplies, and each power supply is able to provide power for the entire subsystem, if necessary. Because of this redundancy, the 9900 subsystem can sustain the loss of multiple power supplies and still continue operation. To make use of this capability, the 9900 should be connected either to dual power sources or to different power panels, so if there is a failure on one of the power sources, the 9900 can continue full operations using power from the alternate source.

2.2.6 Client-Host Interface Processors (CHIPs) and Channels

The CHIPs contain the front-end microprocessors which process the channel commands from the host(s) and manage host access to cache. In the S/390[®] environment, the CHIPs perform CKD-to-FBA and FBA-to-CKD conversion for the data in cache. The CHIPs are available in pairs. Depending on the configuration, each CHIP in a pair contains either two or four microprocessors and four buffers which allow data to be transferred between the CHIP and cache. Each CHIP pair is composed of the same type of channel interface (FICON[™], ExSA[™], or fibre-channel). Each ExSA[™] or fibre-channel CHIP pair supports either four or eight simultaneous data transfers to and from cache and four or eight physical connections to the host. Each FICON[™] CHIP pair supports four physical connections to the host. The 9900 can be configured with multiple CHIP pairs to support various interface configurations. Table 2.1 lists the CHIP specifications and configurations and the number of channel connections for each configuration.

Note: The Hitachi CruiseControl and Graph-Track products (see section 3.7) allow users to collect and view usage statistics for the CHIPs in the 9900 subsystem.

Table 2.1 CHIP and Channel Specifications

Parameter	Specification for 9960	Specification for 9910
Number of CHIP pairs	1, 2, 3, or 4	1, 2, or 3
Simultaneous data transfers per CHIP pair:		
S/390®	4 or 8 ExSA™ (serial/ESCON®)	
Open Systems	4 or 8 (fibre-channel)	
Maximum transfer rate:		
FICON™	100 MB/sec (1 Gbps)	
ExSA™ (serial/ ESCON®)	10 or 17 MB/sec	
Fibre	100 or 200 MB/sec (1 or 2 Gbps)	
Physical interfaces per CHIP pair	4 or 8 ExSA™ 4 FICON™ 4 or 8 fibre-channel	
Maximum physical interfaces per subsystem:	32	24
FICON™	0, 4, 8, 12, or 16	0, 4, 8, or 12
ExSA™ (serial/ESCON®)	0, 4, 8, 12, 16, 20, 24, 28 or 32	0, 4, 8, 12, 16, 20, or 24
Fibre-channel	0, 4, 8, 12, 16, 20, 24, 28 or 32	0, 4, 8, 12, 16, 20, or 24
Logical paths per FICON™ port	512	
Logical paths per ExSA™ (ESCON®) port	256	
Maximum logical paths per subsystem	8,192	6144
Maximum LUs per fibre-channel port	256	
Maximum LVI/LUs per subsystem	4,096	

2.2.7 Channels

The Lightning 9900™ subsystem supports all-mainframe, multiplatform, and all-open system operations and offers the following two types of host channel connections:

- **Fiber Connection (FICON™).** The 9960 subsystem supports up to 16 FICON™-channel ports, and the 9910 supports up to 12 FICON™ ports. The FICON™ ports are capable of data transfer speeds of 100 MB/sec (1 Gbps). The 9900 FICON™-channel cards are available with four ports per CHIP pair. The 9900 supports shortwave and longwave non-OFC (non-open FICON™ control) optical interface and multimode optical cables as well as high-availability (HA) FICON™-channel configurations using hubs and switches. When configured with shortwave FICON™ cards, the 9900 subsystem can be located up to 500 meters (2750 feet) from the host(s). When configured with longwave FICON™ cards, the 9900 subsystem can be located up to ten kilometers from the host(s).
- **Extended Serial Adapter™ (ExSA™).** The 9960 subsystem supports a maximum of 32 ExSA™ serial channel interfaces (compatible with ESCON® protocol), and the 9910 supports a maximum of 24 ExSA™ interfaces. The 9900 ExSA™ channel interface cards provide data transfer speeds of up to 17 MB/sec and are available in four or eight ports per CHIP pair. Each ExSA™ channel can be connected to a single processor or logical partition (LPAR) or to serial channel directors. Shared serial channels can be used for dynamic path switching. The 9900 subsystem also supports the ExSA™ Extended Distance Feature (XDF).
- **Fibre-Channel.** The 9960 subsystem supports up to 32 fibre-channel ports, and the 9910 supports up to 24 fibre ports. The fibre ports are capable of data transfer speeds of 100 or 200 MB/sec (1 or 2 Gbps). The 9900 fibre-channel cards are available in either four or eight ports per CHIP pair. The 9900 supports shortwave and longwave non-OFC (non-open fibre control) optical interface and multimode optical cables as well as high-availability (HA) fibre-channel configurations using hubs and switches. When configured with shortwave fibre cards, the 9900 subsystem can be located up to 500 meters (2750 feet) from the open-system host(s). When configured with longwave fibre cards, the 9900 subsystem can be located up to ten kilometers from the open-system host(s).

2.2.8 Array Control Processors (ACPs)

The ACPs, which control the transfer of data between the disk drives and cache, are installed in pairs for redundancy and performance. Figure 2.4 illustrates a conceptual ACP pair domain. The 9960 can be configured with up to four ACP pairs, and the 9910 has one ACP pair. All functions, paths, and disk drives controlled by one ACP pair are called an “array domain.” An array domain can contain a variety of LVI and/or LU configurations.

The disk drives are connected to the ACP pairs by fibre cables using an arbitrated-loop (FC-AL) topology. Each ACP has four microprocessors and four independent fibre backend paths. Each 9960 fibre backend path can access up to 32 disk drives (32 drives \times 4 paths = 128 disk drives per ACP). Each 9910 fibre backend path can access up to 12 disk drives (12 drives \times 4 paths = 48 disk drives per ACP). Each disk drive is dual-ported for performance and redundancy in case of a backend path failure.

Table 2.2 lists the ACP specifications. Each 9960 ACP pair can support a maximum of 128 physical disk drives (in three array frames), including dynamic spare disk drives. Each ACP pair contains eight buffers (one per fibre path), that support data transfer to and from cache. Each disk drive has a dual-port feature and can transfer data via either port. Each of the two paths shared by the disk drive is connected to a separate ACP in the pair to provide alternate path capability. Each ACP pair is capable of eight simultaneous data transfers to or from the disk drives.

Note: The Hitachi CruiseControl and Graph-Track products (see section 3.7) allow users to collect and view usage statistics for the ACPs in the 9900 subsystem.

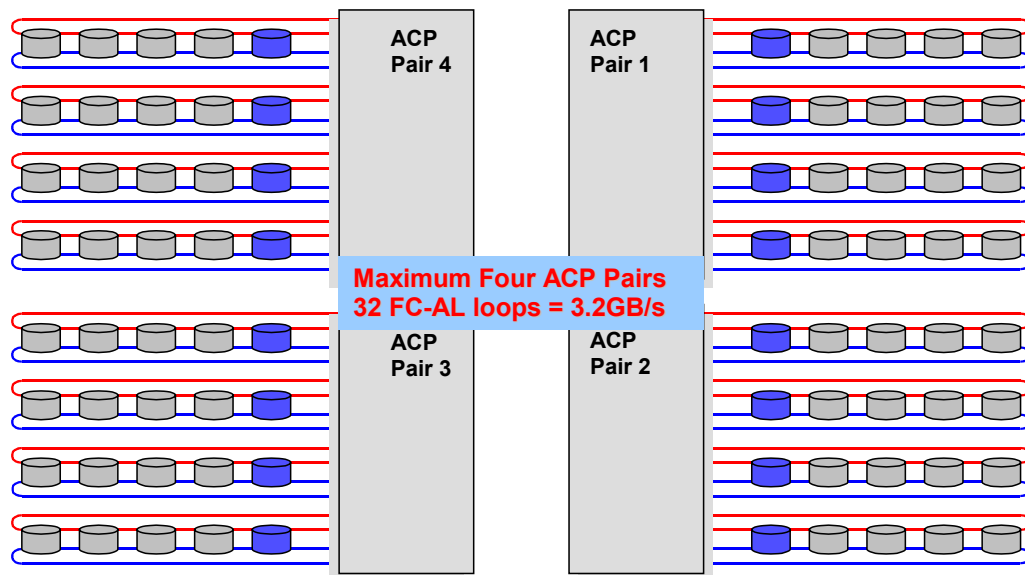


Figure 2.4 Conceptual ACP Array Domain

Table 2.2 ACP Specifications

Description	Specification for 9960	Specification for 9910
Number of ACP pairs	1, 2, 3 or 4	1
Backend paths per ACP pair	8	
Backend paths per subsystem	8, 16, 24 or 32	8
Array group (or parity group) type per ACP pair	RAID-1 and/or RAID-5	
Hard disk drive type per ACP pair ^[1]	18 GB, 47 GB, 73 GB, 146 GB, 180 GB ^[2]	
Logical device emulation type within ACP pair	3380-x, 3390-x, and OPEN-x ^[3]	
Backend array interface type	Fibre-channel arbitrated loop (FC-AL)	
Backend interface transfer rate (burst rate)	100 MB/sec (1 Gbps)	
Maximum concurrent backend operations per ACP pair	8	
Maximum concurrent backend operations per subsystem	32	8
HiStar Network architecture internal bandwidth	1.6 or 3.2 GB/sec	1.6 GB/sec

Notes:

1. All hard disk drives (HDDs) in an array group (also called parity group) must be the same type. Please contact your Hitachi Data Systems representative for the latest information on available HDD types.
2. The 180-GB HDDs should not be intermixed with other HDD types in the same array domain (behind the same ACP pair). See the notes under Table 2.3 for important information on 9900 subsystems which contain 180-GB HDDs.
3. 3390-3 and 3390-3R LVIs cannot be intermixed in the same 9900 subsystem.

2.3 Array Frame

The 9960 array frames contain the physical disk drives, including the disk array groups and the dynamic spare disk drives. Each array frame has dual AC power plugs, which should be attached to two different power sources or power panels. The 9960 can be configured with up to six array frames to provide a storage capacity of up to 88 TB. The 9910 subsystem combines the controller and disk array components in one physical frame.

The 9900 subsystem uses three-inch disk drives with fixed-block-architecture (FBA) format. The currently available disk drives have capacities of 18 GB, 47 GB, 73 GB, 146 GB, and 180 GB. All drives in an array group must have the same capacity. The 18-GB, 47-GB, 73-GB, and 146-GB HDDs can be attached to the same ACP pair. The 180-GB HDDs should not be intermixed with other HDD types behind the same ACP pair. Table 2.3 provides the disk drive specifications.

Each disk drive can be replaced nondisruptively on site. The 9900 utilizes diagnostic techniques and background dynamic scrubbing that detect and correct disk errors. Dynamic sparing is invoked automatically if needed. For both RAID-5 and RAID-1 array groups, any spare disk drive can back up any other disk drive of the same capacity anywhere in the subsystem, even if the failed disk and the spare disk are in different array domains (attached to different ACP pairs). The 9960 can be configured with a minimum of one and a maximum of sixteen spare disk drives. The 9910 can be configured with a minimum of one and a maximum of four spare disk drives. The standard configuration provides one spare drive for type of drive installed in the subsystem. The Hi-Track[®] monitoring and reporting tool detects disk drive failures and notifies the Hitachi Data Systems Support Center automatically, and a service representative is sent to replace the disk drive.

Note: The spare disk drives are used only as replacements and are not included in the storage capacity ratings of the subsystem.

Table 2.3 Disk Drive Specifications

Parameter	Disk Drive Type				
	73 GB	47 GB	18 GB	146 GB	180 GB*
Formatted capacity (GB)	72.91	47.19	18.46	143.7	172.55
Platter diameter	3 inches	3 inches	3 inches	3 inches	3 inches
Physical cylinders (user area)	15,154	12,027	12,027	n/a	24,247
Physical tracks per physical cylinder (number of heads)	24 (DKR1C) 10 (DKR2D)	23 (DKR1B) 16 (DKR1C)	9 (DKR2B) 6 (DKR2C)	10	24
Physical disk platters (user area) (numbers of disks)	12 (DKR1C) 5 (DKR2D)	12 (DKR1B) 8 (DKR1C)	5 (DKR2B) 3 (DKR2C)	5	12
Sector length (byte)	520(512)	520(512)	520(512)	520 (512)	520(512)
Seek time (ms) (Read/Write)	MIN.	0.5/0.7	0.5/0.7	0.5/0.7	0.8 / 1.1
	MAX.	12.0/13.0	12.0/13.0	12.0/13.0	10.0/11.0
	AVE.	5.7/6.5	5.7/6.5	5.2/6.0	4.9/5.4
Revolution speed (rpm)	10,025	10,025	10,025	10,025	7,200
Average latency time (ms)	2.99	2.99	2.99	2.99	4.17
Internal data transfer rate (MB/sec)	33.6 to 56.6	30.2 to 45.6	30.2 to 45.6	57.3 to 99.9	35.3 to 63.5
Max. interface data transfer rate (MB/sec)	100	100	100	100	100

***Caution and notes for the 180-GB HDDs:**

- **Caution:** The 180-GB HDDs should not be intermixed with other HDD types in the same array domain (behind the same ACP pair). If 180-GB HDDs are used, all HDDs in that array domain should also be 180-GB HDDs.
- **Note:** Subsystems which contain the 180-GB HDDs do not have the same availability or serviceability as those subsystems which do not contain 180-GB HDDs. Some offline maintenance may be required.
- **Note:** The 180-GB HDDs are intended for archival use and do not have characteristics that make them suitable as performance devices.

2.3.1 Disk Array Groups

The disk array group is the basic unit of storage capacity for the 9900. Each array group is attached to both ACPs of an ACP pair via eight fibre paths, which enables all disk drives in the array group to be accessed simultaneously by the ACP pair. All disk drives in an array group must have the same logical capacity. Each array frame has two canister mounts, and each canister mount can have up to 48 physical disk drives.

The 9900 supports both RAID-1 and RAID-5 array groups. Figure 2.5 illustrates a sample RAID-1 layout. A RAID-1 array group consists of two pair of disk drives in a mirrored configuration, regardless of disk drive capacity. Data is striped to two drives and mirrored to the other two drives. The stripe consists of two data chunks. The primary and secondary stripes are toggled back and forth across the physical disk drives for high performance. Each data chunk consists of either eight logical tracks (S/390®) or 768 logical blocks (open systems). A failure in a drive causes the corresponding mirrored drive to take over for the failed drive. Although the RAID-5 implementation is appropriate for many applications, the RAID-1 option on the all-open 9900 subsystem is ideal for workloads with low cache-hit ratios.

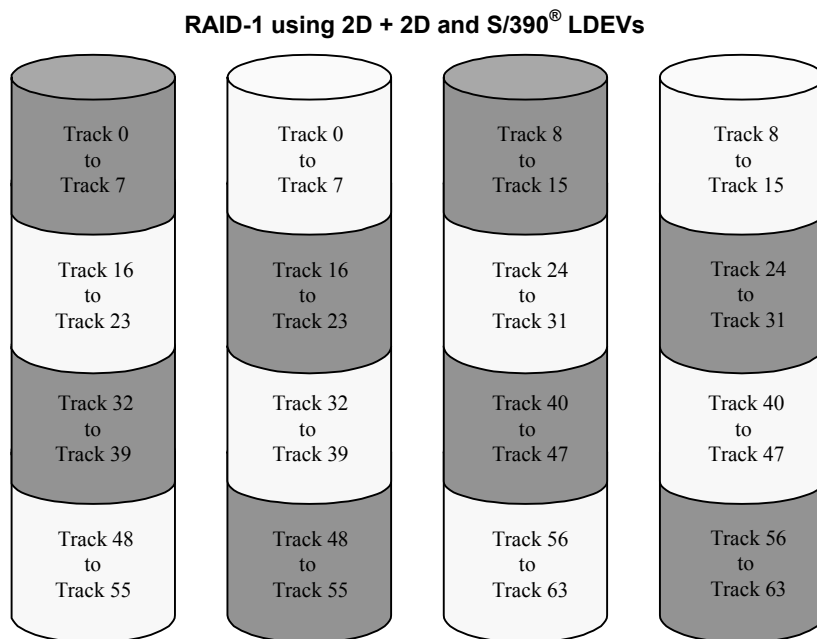


Figure 2.5 Sample RAID-1 Layout

A RAID-5 array group consists of four disk drives. The data is written across the four hard drives in a stripe that has three data chunks and one parity chunk. Each chunk contains either eight logical tracks (S/390®) or 768 logical blocks (open systems). The enhanced RAID-5+ implementation in the 9900 subsystem minimizes the write penalty incurred by standard RAID-5 implementations by keeping write data in cache until an entire stripe can be built and then writing the entire data stripe to the disk drives.

Figure 2.6 illustrates RAID-5 data stripes mapped over four physical drives. Data and parity are striped across each of the disk drives in the array group (hence the term “parity group”). The logical devices (LDEVs) are evenly dispersed in the array group, so that the performance of each LDEV within the array group is the same. Figure 2.6 also shows the parity chunks that are the “Exclusive OR” (EOR) of the data chunks. The parity and data chunks rotate after each stripe. The total data in each stripe is either 24 logical tracks (eight tracks per chunk) for S/390® data, or 2304 blocks (768 blocks per chunk) for open-systems data. Each of these array groups can be configured as either 3380-x, 3390-x, or OPEN-x logical devices. All LDEVs in the array group must be the same format (3380-x, 3390-x, or OPEN-x). For open systems, each LDEV is mapped to a SCSI address, so that it has a TID and logical unit number (LUN).

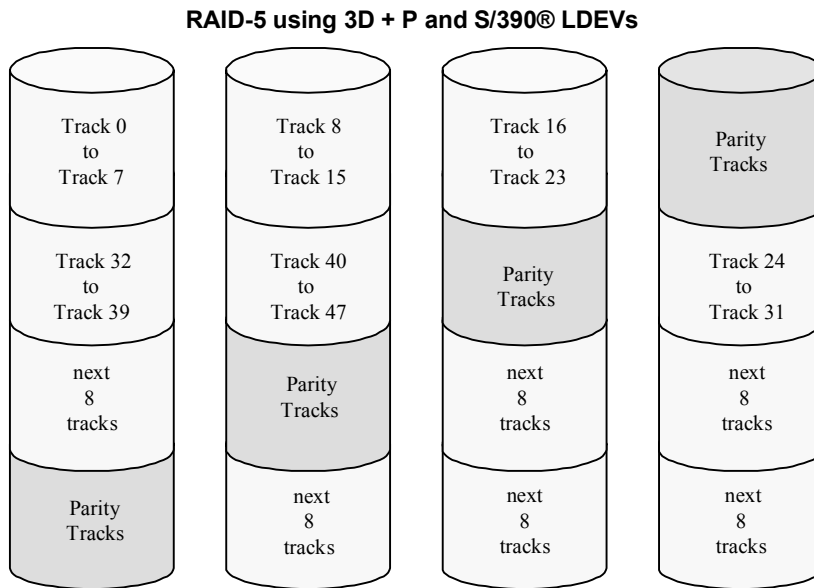


Figure 2.6 Sample RAID-5 Layout (Data Plus Parity Stripe)

Note: The Hitachi CruiseControl and Graph-Track products (see section 3.7) allow users to collect and view detailed usage statistics for the disk array groups in the 9900 subsystem.

2.3.2 Sequential Data Striping

The 9900 subsystem’s enhanced RAID-5+ implementation attempts to keep write data in cache until parity can be generated without referencing old parity or data. This capability to write entire data stripes, which is usually achieved only in sequential processing environments, minimizes the write penalty incurred by standard RAID-5 implementations. The device data and parity tracks are mapped to specific physical disk drive locations within each array group. Therefore, each track of an LDEV occupies the same relative physical location within each array group in the subsystem.

2.4 Intermix Configurations

2.4.1 RAID-1 & RAID-5 Intermix

RAID technology provides full fault-tolerance capability for the disk drives of the 9900 subsystem. The 9900 supports RAID-1, RAID-5, and intermixed RAID-1 and RAID-5 configurations, including intermixed array groups within an array domain. The cache management algorithms (see section 3.3.1) enable the 9900 to stage up to one full RAID stripe of data into cache ahead of the current access to allow subsequent access to be satisfied from cache at host channel transfer speeds.

2.4.2 Hard Disk Drive Intermix

Figure 2.7 illustrates an intermix of hard disk drive types. All hard disk drives in one array group must be of the same capacity and type. The 18-GB, 47-GB, 73-GB, and 146-GB HDDs can be attached to the same ACP pair. The 180-GB HDDs should not be intermixed with other HDD types behind the same ACP pair. See the notes under Table 2.3 for important information on 9900 subsystems which contain 180-GB HDDs.




Figure 2.7 Sample Hard Disk Drive Intermix

2.4.3 Device Emulation Intermix

The 9900 subsystem supports an intermix of different device emulations (e.g., 3390-x LVIs, 3380-x LVIs, OPEN-x LUs) on the same ACP pair. Figure 2.8 illustrates an intermix of device emulation types. The only requirement is that the devices within each array group must have the same type of track geometry or format, as follows:

- 3390-1, -2, -3 or -9 can be intermixed within an array group.
- 3380-E, -J, or -K can be intermixed within an array group.
- OPEN-3, -8, -9, -E, -L, and -M can be intermixed within an array group with the following restrictions:
 - OPEN-L devices can only be configured on array groups of 73, 146, or 180 GB HDDs.
 - OPEN-M devices can only be configured on array groups of 47 or 180 GB HDDs.
- OPEN-K cannot be intermixed with other device types within an array group.

Note: For the latest information on supported LU types and intermix requirements, please contact your Hitachi Data Systems account team.

Note: The 9960 and 9910 subsystems may support different device emulations and intermix configurations.

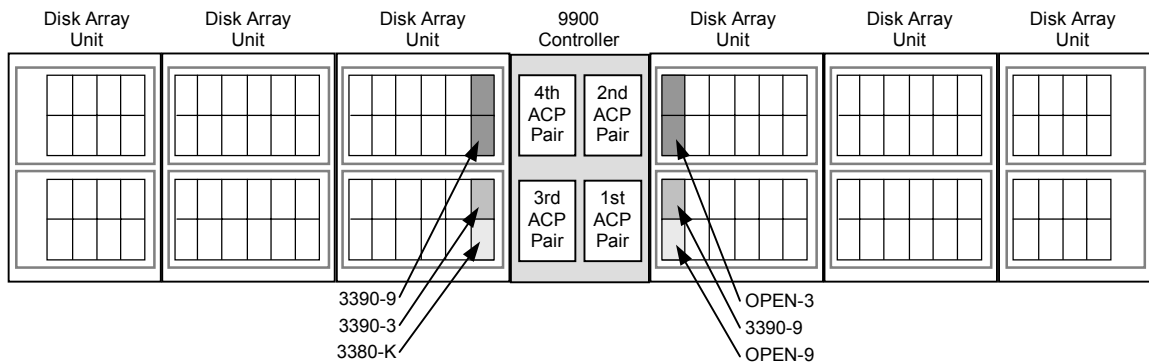


Figure 2.8 Sample Device Emulation Intermix

2.5 Service Processor (SVP)

The Lightning 9900™ subsystem includes a built-in laptop PC called the service processor (SVP). The SVP is integrated into the controller frame and can only be used by authorized Hitachi Data Systems personnel. The SVP enables the Hitachi Data Systems representative to configure, maintain, and upgrade the 9900 subsystem. The SVP also collects performance data for all key components of the 9900 subsystem to enable diagnostic testing and analysis. The Hitachi Graph-Track™ software product (see section 3.7.27) stores the SVP performance data on the Remote Console PC and allows users to view the data in graphical format and export the data for statistical analysis. **Note:** The SVP does not have access to any user data stored on the 9900 subsystem.

2.6 Remote Console PC

The Remote Console PC is LAN-attached to one or more 9900 subsystems via the 9900-internal LAN. The 9900 Remote Console PC supports the Windows® 95, Windows® 98, Windows® 2000, and Windows NT® operating systems to provide a user-friendly interface for the 9900 remote console software products. The remote console software communicates directly with the SVP of each attached subsystem, enabling the user to view subsystem configuration information and issue commands directly to the 9900 subsystems. For further information on the Remote Console PC, please refer to the *Hitachi Lightning 9900™ Remote Console User's Guide* (MK-90RD003).

Chapter 3 Functional and Operational Characteristics

3.1 New 9900 Features and Capabilities

The Hitachi Lightning 9900™ subsystem offers the following new or improved features and capabilities which distinguish the 9900 subsystem from the 7700E subsystem:

- Support for FICON™ channel interface.
- Sixteen (16) logical control unit (CU) images.
- State-of-the-art hard disk drives of 18-GB, 47-GB, 73-GB, 146-GB, and 180-GB capacities.
- Up to 32 GB cache memory for the 9960 subsystem and 16 GB cache for the 9910.
- Up to 32 host interface ports [ExSA™ (ESCON®) and/or fibre-channel].
- Up to 16 host interface ports (FICON™).
- Up to 256 logical paths per ExSA™ (ESCON®) channel interface.
- Up to 512 logical paths per FICON™ channel interface.
- Up to 4096 device addresses.
- Prioritized port control for open-system users.

3.2 I/O Operations

The 9900 I/O operations are classified into three types based on cache usage:

- **Read hit:** For a read I/O, when the requested data is already in cache, the operation is classified as a read hit. The CHIP searches the cache directory, determines that the data is in cache, and immediately transfers the data to the host at the channel transfer rate.
- **Read miss:** For a read I/O, when the requested data is not currently in cache, the operation is classified as a read miss. The CHIP searches the cache directory, determines that the data is not in cache, disconnects from the host, creates space in cache, updates the cache directory, and requests the data from the appropriate ACP pair. The ACP pair stages the appropriate amount of data into cache, depending on the type of read I/O (e.g., sequential).
- **Fast write:** All write I/Os to the 9900 subsystem are fast writes, because all write data is written to cache before being destaged to disk. The data is stored in two cache locations on separate power boundaries in the nonvolatile duplex cache (see section 2.2.3). As soon as the write I/O has been written to cache, the 9900 subsystem notifies the host that the I/O operation is complete, and then destages the data to disk.

3.3 Cache Management

3.3.1 Algorithms for Cache Control

The 9900 subsystem places all read and write data in cache, and 100% of cache memory is available for read operations. The amount of fast-write data in cache is dynamically managed by the cache control algorithms to provide the optimum amount of read and write cache, depending on the workload read and write I/O characteristics.

The algorithms for internal cache control used by the 9900 include the following:

- **Hitachi Data Systems Intelligent Learning Algorithm.** The Hitachi Data Systems Intelligent Learning Algorithm identifies random and sequential data access patterns and selects the amount of data to be “staged” (read from disk into cache). The amount of data staged can be a record, partial track, full track, or even multiple tracks, depending on the data access patterns.
- **Least-recently-used (LRU) algorithm (modified).** When a read hit or write I/O occurs in a nonsequential operation, the least-recently-used (LRU) algorithm marks the cache segment as most recently used and promotes it to the top of the appropriate LRU list. In a sequential write operation, the data is destaged by priority, so the cache segment marked as least-recently used is immediately available for reallocation, since this data is not normally accessed again soon.
- **Sequential prefetch algorithm.** The sequential prefetch algorithm is used for sequential-access commands or access patterns identified as sequential by the Intelligent Learning Algorithm. The sequential prefetch algorithm directs the ACPs to prefetch up to one full RAID stripe (24 tracks) to cache ahead of the current access. This allows subsequent access to the sequential data to be satisfied from cache at host channel transfer speeds.

Note: The 9900 subsystem supports S/390[®] extended count key data (ECKD) commands for specifying cache functions.

3.3.2 Write Pending Rate

The write pending rate is the percent of total cache used for write pending data. The amount of fast-write data stored in cache is dynamically managed by the cache control algorithms to provide the optimum amount of read and write cache based on workload I/O characteristics. Hitachi CruiseControl and Graph-Track allow users to collect and view the write-pending-rate data and other cache statistics for the Lightning 9900™ subsystem.

Note: If the write pending limit is reached, the 9900 sends DASD fast-write delay or retry indications to the host until the appropriate amount of data can be destaged from cache to the disks to make more cache slots available.

3.4 Control Unit (CU) Images, LVIs, and LUs

3.4.1 CU Images

The 9900 subsystem supports the following logical CU images (emulation types): 3990-3, 3990-6E, and 2105. The 9900 subsystem is configured with one logical CU image for each 256 devices (one storage subsystem ID (SSID) for each 64 or 256 devices) to provide a maximum of sixteen CU images per subsystem. The S/390® data management features of the 9900 subsystem may have restrictions on CU image compatibility. FICON™ support requires 2105-F20 emulation. For further information on CU image support, please contact your Hitachi Data Systems account team.

3.4.2 Logical Volume Image (LVIs)

The 9900 subsystem supports the following S/390® LVI types: 3390-1, -2, -3, -3R, -9; 3380-E, -J, -K. The LVI configuration of the subsystem depends on the RAID implementation and physical disk drive capacities. See section 4.1 for further information on LVI configurations.

3.4.3 Logical Unit (LU) Type

The 9900 subsystem currently supports the following LU types: OPEN-3, OPEN-8, OPEN-K, OPEN-9, OPEN-E, OPEN-L, and OPEN-M. Table 3.1 lists the capacities for each standard LU type. The 9900 also allows users to configure custom-size LUs which are smaller than standard LUs as well as size-expanded LUs which are larger than standard LUs. LU Size Expansion (LUSE) volumes can range in size from 3.748 (OPEN-K*2) to 524.448 GB (OPEN-E*36). Each LU is identified by target ID (TID) and LU number (LUN) (see Figure 3.1). Each 9900 fibre-channel port supports addressing capabilities for up to 256 LUNs.

Table 3.1 Capacities of Standard LU Types

LU Type	OPEN-K	OPEN-3	OPEN-8	OPEN-9	OPEN-E	OPEN-L	OPEN-M
Capacity (GB)	1.881	2.461	7.347	7.384	14.568	36.450	47.185




Figure 3.1 Fibre-Channel Device Addressing

3.5 System Option Modes

To provide greater flexibility and enable the 9900 to be tailored to unique customer operating requirements, additional operational parameters, or system option modes, are available for the 9900 subsystem. At installation, the 9900 modes are set to their default values (**OFF**), so make sure to discuss these settings with your Hitachi Data Systems team. The 9900 modes can only be changed by the Hitachi Data Systems representative.

Table 3.2 - Table 3.8 show the public 9900 system option modes. **Note:** This 9900 mode information was current at the time of publication of this document but may change. Please contact your Hitachi Data Systems representative for the latest information on the 9900 System Option Modes.

Table 3.2 Common System Option Modes

Mode	Level	Description and Usage
22	Mandatory	Controlling option for correction/drive copy: Turn mode 22 ON (nondisruptive) after the completion of the micro-program exchange to microcode 01-18-67-00/00 or higher. Do NOT activate this mode unless running the minimum microcode level indicated. Also improves destage schedule to support large-capacity HDD (146-GB and higher). ON: Controlling option for correction/ drive copy is active. OFF: Controlling option for correction/ drive copy is not active.

Table 3.3 System Option Modes for Mainframe Connectivity

Mode	Level	Description and Usage
1	Optional	PDS search assist by learning algorithm. ON: 9900 performs PDS search assist by learning algorithm. OFF: 9900 performs PDS search assist by host indication. Microcode level: 01-10-00-00/10 and higher.
3	Optional	EPW function for ascending record number in a cylinder. ON: EPW can work to ascending record number in a cylinder. OFF: EPW cannot work to ascending record number in a cylinder. Microcode level: 01-10-00-00/10 and higher.
5	Optional	SIM reporting option regarding drive/media SIM for 3380-K. ON: DKC sends drive/media SIM for 3380-K to host. OFF: DKC does not send drive/media SIM for 3380-K to host. Microcode level: 01-10-00-00/10 and higher.
6	Optional	"Suppression of reporting SSB F/M=6X. For VSE/FASTCOPY Copy job" ON: DKC does not report SSB F/M=6X to host. OFF: DKC reports SSB F/M=6X to host. Microcode level: From 01-10-00-00/10 to 01-13-01-00/00.
162	Optional	Tuning option for mainframe performance. Do not set MODE162=ON to the subsystem which has less than 01-18-48-00/00. ON: Tuning for sequential write performance is inactive. OFF: Tuning for sequential write performance is active.

Table 3.4 System Option Modes for Open-System Connectivity

Mode	Level	Description and Usage
57	Optional	SIM notification to open-system host. ON: 9900 reports SIM to open-system host as Sense Key/Sense code. OFF: 9900 does not send SIM to open-system host. SIM will be reported as B/C300. EC=D002 will be logged in SSB.log. Set host mode 04 and System option MODE 56,57 for Sequent. Microcode level: 01-10-00-00/10 and higher.
111	Optional	LUN security option. ON: DKC checks all command for LUN security. OFF: DKC does not check all command for LUN security.
161	Optional	Suppression of high speed micro-program exchange for CHT. ON: 9900 does not perform high speed micro-program for CHT. New micro-program is also written into flash memory on the CHT. OFF: 9900 does perform high speed micro-program for CHT. Only the microprocessor code is loaded. Either a PS OFF/ON or a dummy replace of the CHT board will be required for the subsequent loading of flash memory. Microcode level: 01-11-xx and higher.
185	Mandatory	Solaris + SUN Cluster 3.0 or SDS4.2.1. ON: Mandatory setting when SUN Cluster 3.0 or SDS4.2.1 is connected. OFF: SUN Cluster 3.0 or SDS4.2.1 should never be connected with this setting. Micro-code level: 01-15-39-00/0x and higher.
186	Mandatory	VERITAS Database Editions/Advanced Cluster. ON: Mandatory setting when VERITAS Database Editions/Advanced Cluster is connected. OFF: VERITAS Database Editions/Advanced Cluster should never be connected with this setting. Microcode level: 01-17-94-00/10 and higher.
198	Mandatory	Tru64 5.1a connection. ON: Mandatory setting when Tru64 5.1a is connected. OFF: Tru64 5.1a should never be connected with this setting. Microcode level: 01-16-60-00/00 and higher.
206	Mandatory	An inhibit option regarding Q-Err expansion bit. ON: DKC inhibits AIX Q-Err expansion bit. OFF: DKC does not inhibit AIX Q-Err expansion bit. Microcode level: 01-18-42-00/00 and higher.
213	Mandatory	AIX + HACMP ON: Mandatory setting when AIX + HACMP is connected. OFF: AIX + HACMP should never be connected with this setting. Microcode level: 01-17-94-00/06 and higher.
247	Mandatory	Tru64 (Recognition of LUN). Micro-program has been modified to change Vendor ID and Product ID as response of Inquiry command with system option MODE 247 for Tru64 connection as shown below. Do not set MODE 247=ON if the subsystem is already connected to Tru64 with old microcode. Set MODE 247=ON if the subsystem is to be newly connected to Tru64 with 01-18-92-00/00 and later. Note: The function of MODE 247 is suppressed when MODE 198 =ON. ON: Mandatory setting when Tru64 is connected. OFF: Tru64 should never be connected with this setting. Microcode level: 01-19-92-00/00 and higher.

Table 3.5 System Option Modes for ShadowImage - S/390® and ShadowImage

Mode	Level	Description and Usage
80	Optional	ShadowImage Quick Restore function. ON: 9900 does not perform the ShadowImage Quick Restore function. OFF: 9900 performs the ShadowImage Quick Restore function.
87	Optional	Quick Resync by CCI (RAID Manager). ON: 9900 performs ShadowImage quick resync operation for Resync command from CCI. OFF: 9900 does not perform ShadowImage quick resync operation for Resync command from CCI. Microcode level: 01-13-18-00/07 and higher.
122	Optional	ShadowImage Quick Split and Quick Resync functions. ON: 9900 does not perform ShadowImage quick split and quick resync operations. OFF: 9900 performs ShadowImage quick split and quick resync operations.

Table 3.6 System Option Modes for TrueCopy - S/390® Synchronous & Asynchronous (continues on the next page)

Mode	Level	Description and Usage
20	Optional	Enables TC390 – R-VOL read-only function (RCU only). ON: R-VOL read only function is available. OFF: R-VOL read only function is not available. Microcode level: 01-10-00-00/10 and higher.
36	Optional	TC390 Synchronous – Selects function of CRIT=Y(ALL) or CRIT=Y(PATHS). ON: CRIT=Y(ALL) => equivalent to Fence Level = Data. OFF: CRIT=Y(PATHS) => equivalent to Fence Level = Status. Microcode level: 01-10-00-00/10 and higher.
38	Optional	TC390 – Changes SSB reported against the WRITE I/O to the M-VOL in critical state. ON: Intervention required. OFF: Command rejected (PPRC specification). Microcode level: 01-10-00-00/10 and higher.
49	Optional	TC390 – Changes reporting of SSIDs in response to CQUERY command (which is limited to four SSIDs). When 64 LDEVs per SSID are defined, mode 49 must be ON for TC390, GDPS, and P/DAS operations. When 256 LDEVs per SSID are defined, mode 49 must be OFF. ON: Report first SSID for all (256) devices in the logical CU. OFF: Report SSID specified for each 64 or 256 devices. Microcode level: 01-10-00-00/10 and higher.
64	Optional	TC390 CGROUP – Defines scope of CGROUP command within the 9900. Must be OFF for GDPS. ON: All TC390 volumes in this 9900 subsystem. OFF: TC390 volumes behind the specified LCU pair (main and remote LCUs). Microcode level: 01-10-00-00/10 and higher.

Table 3.6 System Option Modes for TrueCopy - S/390® Synchronous & Asynchronous (continued)

Mode	Level	Description and Usage																		
93	Optional	<p>TC390 Asynchronous graduated delay process for sidefile control.</p> <p>ON: soft delay type. OFF: strong delay type.</p> <table border="1"> <thead> <tr> <th>Amount of sidefile</th> <th>Strong delay type</th> <th>Soft delay type</th> </tr> </thead> <tbody> <tr> <td>threshold – [15-20%] (HWM)</td> <td>100 ms x 1 time</td> <td>20 ms x 1 time</td> </tr> <tr> <td>threshold – [10-15%]</td> <td>200 ms x 1 time</td> <td>40 ms x 1 time</td> </tr> <tr> <td>threshold – [5-10%]</td> <td>300 ms x 1 time</td> <td>60 ms x 1 time</td> </tr> <tr> <td>threshold – [0-5%]</td> <td>400 ms x 1 time</td> <td>80 ms x 1 time</td> </tr> <tr> <td>threshold or higher</td> <td>500 ms x permanent</td> <td>100 ms x permanent</td> </tr> </tbody> </table> <p>Microcode level: 01-13-18-00/07 and higher.</p>	Amount of sidefile	Strong delay type	Soft delay type	threshold – [15-20%] (HWM)	100 ms x 1 time	20 ms x 1 time	threshold – [10-15%]	200 ms x 1 time	40 ms x 1 time	threshold – [5-10%]	300 ms x 1 time	60 ms x 1 time	threshold – [0-5%]	400 ms x 1 time	80 ms x 1 time	threshold or higher	500 ms x permanent	100 ms x permanent
Amount of sidefile	Strong delay type	Soft delay type																		
threshold – [15-20%] (HWM)	100 ms x 1 time	20 ms x 1 time																		
threshold – [10-15%]	200 ms x 1 time	40 ms x 1 time																		
threshold – [5-10%]	300 ms x 1 time	60 ms x 1 time																		
threshold – [0-5%]	400 ms x 1 time	80 ms x 1 time																		
threshold or higher	500 ms x permanent	100 ms x permanent																		
104	Optional	<p>TC390 CGROUP – Selects subsystem default for CGROUP FREEZE option. Applies only to 3990 controller emulation. Note: Mode 104 is invalid if the controller emulation is 2105. For 2105, use the CGROUP option of CESTPATH.</p> <p>ON: FREEZE enabled. OFF: FREEZE disabled.</p> <p>Microcode level: 01-10-00-00/10 and higher.</p>																		
114	Optional	<p>TC390 – Allows dynamic port mode setting (RCP/LCP for serial, Initiator/RCU target for fibre-channel) through PPRC CESTPATH and CDELPATH commands.</p> <p>ON: Set defined port to RCP/LCP mode (serial) or Initiator/RCU target mode (fibre-channel) as needed. OFF: Port must be reconfigured using Remote Console PC (or SVP).</p> <p>Microcode level: 01-10-00-00/10 and higher.</p> <p>Note: For fibre-channel interface, do not use the CESTPATH and CDELPATH commands at the same time as the SCSI path definition function of LUN Manager. The FC interface ports need to be configured as initiator ports or RCU target ports before the CESTPATH and CDELPATH commands are issued.</p> <p>Caution: Before issuing the CESTPATH command, you must make sure that the relevant paths are offline from the host(s) (e.g., configure the Chipid offline, or deactivate the LPAR, or block the port in the ESCD). If any active logical paths still exist, the add path operation will fail because the port mode (LCP/RCP) cannot be changed.</p>																		
118	Optional	<p>HXRC, CC, TC390A – SIM notification when: HXRC+CC sidefile reaches sleep wait threshold (see modes 45, 85, 86, 97). TC390A sidefile reaches high-water mark (HWM = sidefile threshold - 20%) (see mode 93).</p> <p>ON: Generate SIM. OFF: Do not generate SIM.</p> <p>Microcode level: 01-10-00-00/10 and higher.</p>																		

Table 3.7 System Option Modes for HXRC

Mode	Level	Description and Usage
45	Optional	Sleep Wait suppressing option (see modes 61, 85, 86, 97). When Mode 45 is ON and Mode 61 is ON, WRITE I/Os for LDEVs are blocked by the threshold specified by SDM. ON: Sidefile threshold does not activate Sleep Wait timer at the sleep wait threshold. OFF: Sidefile threshold activates Sleep Wait timer at the sleep wait threshold. Microcode level: -10-00-00/10 and higher.
61	Optional	Enables the DONOTBLOCK option of the XADPAIR command. Must be OFF if the operating system does not support the DONOTBLOCK option. ON: DONOTBLOCK option activated. OFF: DONOTBLOCK option ignored. Microcode level: 01-12-18-00/00 and higher.
85,86	Optional	Variable sidefile threshold. Mode 85 ON and mode 86 OFF: Thresholds for Sleep wait/SCP/Puncture = 30/40/50% Mode 85 OFF and mode 86 OFF: Thresholds for Sleep wait/SCP/Puncture = 40/50/60% Mode 85 OFF and mode 86 ON: Thresholds for Sleep wait/SCP/Puncture = 50/60/70% Mode 85 ON and mode 86 ON: Thresholds for Sleep wait/SCP/Puncture = 60/70/80% Microcode level: 01-12-18-00/00 and higher.
97	Optional	Variable Sleep Wait timer duration (see modes 45, 85, 86). ON: Sleep Wait timer duration = 10 ms. OFF: Sleep Wait timer duration = 100 ms. Microcode level: 01-10-00-00/10 and higher.
98	Optional	Selects SCP or session cancel (see modes 45, 85, 86). ON: Forced session cancel. OFF: SCP. Microcode level: 01-10-00-00/10 and higher.
118	Optional	HXRC, CC, TC390A – SIM notification when: HXRC+CC sidefile reaches sleep wait threshold (see modes 45, 85, 86, 97). TC390A sidefile reaches high-water mark (HWM = sidefile threshold - 20%) (see mode 93). ON: Generate SIM. OFF: No SIM generated. Microcode level: 01-10-00-00/10 and higher.

Table 3.8 System Option Modes for Concurrent Copy (CC)

Mode	Level	Description and Usage
118	Optional	HXRC, CC, TC390A – SIM notification when: HXRC+CC sidefile reaches sleep wait threshold (see modes 45, 85, 86, 97). TC390A sidefile reaches high-water mark (HWM = sidefile threshold - 20%) (see mode 93). ON: Generate SIM. OFF: No SIM generated. Microcode level: 01-10-00-00/10 and higher.

3.6 Open Systems Features and Functions

The 9900 subsystem offers many features and functions specifically for the open-systems environment. The 9900 supports multi-initiator I/O configurations in which multiple host systems are attached to the same fibre-channel interface. The 9900 subsystem also supports important open-system functions such as fibre-channel arbitrated-loop (FC-AL) and fabric topologies, command tag queuing, multi-initiator I/O, and industry-standard middleware products which provide application and host failover, I/O path failover, and logical volume management functions. In addition, several program products and services are specifically for open systems. See section 3.7 for more information.

3.6.1 Failover and SNMP Support

The 9900 subsystem supports industry-standard products and functions which provide host and/or application failover, I/O path failover, and logical volume management (LVM), including Hitachi Dynamic Link Manager™, VERITAS® FirstWatch®, VERITAS® Cluster Server, VERITAS® Volume Manager/DMP, Sun Cluster, TruCluster, HP® MC/ServiceGuard, HACMP, Sequent Multi Path, Sequent Cluster Control, VMSCluster, Novell® Cluster Server, Microsoft® Cluster Server. For the latest information on failover and LVM product releases, availability, and compatibility, please contact your Hitachi Data Systems account team.

The 9900 subsystem also supports the industry-standard simple network management protocol (SNMP) for remote subsystem management from the UNIX®/PC server host. SNMP is used to transport management information between the 9900 subsystem and the SNMP manager on the host. The SNMP agent for the 9900 subsystem sends status information to the host(s) when requested by the host or when a significant event occurs.

3.6.2 Share-Everything Architecture

The 9900 subsystem's global cache provides a "share-everything" architecture that enables any fibre-channel port to have access to any LU in the subsystem. In the 9900, each LU can be assigned to multiple fibre-channel ports to provide I/O path failover and/or load balancing (with the appropriate middleware support) without sacrificing cache coherency. The LUN mapping can be performed by the user using the LUN Manager remote console software, or by your Hitachi Data Systems representative (fee-based configuration service).

3.6.3 SCSI Extended Copy Command Support

The Extended Copy Manager (ECM) feature of the 9900 subsystem (see section 3.7.6) supports the SCSI Extended Copy (e-copy) command issued from the host server to the 9900 subsystem. ECM provides a server-free backup solution between the 9900 subsystem and backup devices (e.g., tape, disk) in a storage-area network (SAN) environment, eliminating server CPU and I/O overhead during movement of data and decreasing the time required for backup. The e-copy (backup) operations are performed via fibre-channel interfaces directly to the backup devices. ECM operations are configured using the LUN Manager software on the Remote Console PC. To implement ECM operations, you need a backup application on the host server to issue the Extended Copy commands to the 9900 subsystem.

3.7 Data Management Functions

The 9900 subsystem provides features and functions that increase data availability and improve data management. Table 3.9 and Table 3.10 list the data management features that are currently available for the 9900 subsystem. Please refer to the appropriate user documentation for more details.

Table 3.9 Data Management Functions for Open-System Users

Feature Name	Controlled by:		Licensed Software?	User Document(s)
	Remote Console?	Host OS?		
Replication:				
TrueCopy (section 3.7.1)	Yes	Yes	Yes	MK-91RD051
Hitachi ShadowImage (section 3.7.3)	Yes	Yes	Yes	MK-90RD031
Command Control Interface (CCI) (section 3.7.5)	No	Yes	Yes	MK-90RD003
Extended Copy Manager (ECM) (section 3.7.6)	Yes	Yes	Yes	MK-91RD049
Data Migration (section 3.7.9)	Yes	No	No	Service Offering
Backup/restore and sharing:				
Hitachi RapidXchange (HRX) (section 3.7.10)	No	Yes	Yes	MK-90RD020
Hitachi Multiplatform Backup/Restore (HMBR) (section 3.7.11)	No	Yes	Yes	MK-90RD037
HARBOR® File-Level Backup and Restore (section 3.7.12)	No	No	No	Service Offering
HARBOR® File Transfer (section 3.7.13)	No	Yes	Yes	Service Offering
Resource management:				
HiCommand™ (section 3.7.14)	Yes	No	Yes	Web Client: MK-91HC001 Server: MK-91HC002 Security: MK-91HC003 HiScan: MK-91HC005 Command Line Interface: MK-91HC007
LUN Manager (section 3.7.15),	Yes	No	Yes	MK-91RD049
LUSE (section 3.7.16)	Yes	No	Yes	MK-91RD049
Virtual LUN (section 3.7.17)	Yes	No	Yes	MK-90RD005
FlashAccess (section 3.7.18)	Yes	Yes	Yes	MK-90RD004
Hitachi SANtinel (section 3.7.20)	Yes	No	Yes	MK-91RD049
Prioritized Port Control (section 3.7.22)	Yes	No	Yes	MK-90RD030
Hitachi Dynamic Link Manager™ (section 3.7.24)	No	Yes	Yes	For AIX® systems: MK-91RD056 For Solaris™ systems: MK-91RD057 For Windows® systems: MK-91RD058 For HP-UX® systems: MK-91RD059
LDEV Guard (section 3.7.25)	Yes	Yes	Yes	MK-92RD072
Storage utilities:				
Hitachi CruiseControl (section 3.7.26)	Yes	No	Yes	MK-91RD054
Hitachi Graph-Track (section 3.7.27)	Yes	No	Yes	MK-90RD032
Network solutions:				
Hitachi Freedom NAS™ (section 1.1.5)	N/A	N/A	N/A	MK-91RD053
Hitachi Freedom SAN™ (section 1.1.5)	N/A	N/A	N/A	N/A

Table 3.10 Data Management Functions for S/390® Users

Feature Name	Controlled by:		Licensed Software?	User Document(s)
	Remote Console?	Host OS?		
Replication:				
TrueCopy – S/390® (section 3.7.2)	Yes	Yes	Yes	MK-91RD050
Hitachi ShadowImage – S/390® (section 3.7.4)	Yes	Yes	Yes	MK-90RD012
Hitachi Extended Remote Copy (HXRC) (section 3.7.7)	No	Yes	No	<i>Planning for IBM Remote Copy, SG24-2595; Advanced Copy Services, SC35-0355; DFSMS MVS V1 Remote Copy Guide and Reference, SC35-0169</i>
Hitachi NanoCopy™ (section 3.7.8)	N/A	N/A	N/A	MK-90DD878, Service Offering
Data Migration (section 3.7.9)	Yes	No	No	Service Offering
Backup/restore and sharing:				
Hitachi RapidXchange (HRX) (section 3.7.10)	No	Yes	Yes	MK-90RD020
Resource management:				
Virtual LVI (section 3.7.17)	Yes	No	Yes	MK-90RD005
FlashAccess (section 3.7.18)	Yes	Yes – Cache Manager	Yes	MK-90RD004
Cache Manager (section 3.7.19)	Yes – FlashAccess	Yes	No	MK-91RD045
Hitachi SANtinel – S/390® (section 3.7.21)	Yes	No	Yes	MK-90RD036
Hitachi Parallel Access Volume (HPAV) (section 3.7.23)	Yes	Yes	Yes	MK-91RD047
Storage utilities:				
Hitachi CruiseControl (section 3.7.26)	Yes	No	Yes	MK-91RD054
Hitachi Graph-Track™ (section 3.7.27)	Yes	No	Yes	MK-90RD032

3.7.1 Hitachi TrueCopy (TC)

Hitachi TrueCopy enables open-system users to perform synchronous and/or asynchronous remote copy operations between 9900 subsystems. The user can create, split, and resynchronize LU pairs. TrueCopy also supports a “takeover” command for remote host takeover (with the appropriate middleware support). Once established, TrueCopy operations continue unattended and provide continuous, real-time data backup. Remote copy operations are nondisruptive and allow the primary TrueCopy volumes to remain online to all hosts for both read and write I/O operations. TrueCopy operations can also be performed between 9900, 7700E, and 7700 subsystems (the 7700 supports only synchronous remote copy).

Hitachi TrueCopy supports both serial (ESCON[®]) and fibre-channel interface connections between the main and remote 9900 subsystems. For serial interface connection, TrueCopy operations can be performed across distances of up to 43 km (26.7 miles) using standard ESCON[®] support. For fibre-channel connection, TrueCopy operations can be performed across distances of up to 30 km (18.6 miles) using single-mode longwave optical fibre cables in a switch configuration. Long-distance solutions are provided, based on user requirements and workload characteristics, using approved channel extenders and communication lines.

Note: For further information on Hitachi TrueCopy, please see the *Hitachi Lightning 9900™ Hitachi TrueCopy User and Reference Guide* (MK-91RD051), or contact your Hitachi Data Systems account team.

3.7.2 Hitachi TrueCopy – S/390[®] (TC390)

Hitachi TrueCopy - S/390[®] (TC390) enables S/390[®] users to perform synchronous and asynchronous remote copy operations between 9900 subsystems. Hitachi TrueCopy - S/390[®] can be used to maintain copies of data for backup or duplication purposes. Once established, TC390 operations continue unattended and provide continuous, real-time data backup. Remote copy operations are nondisruptive and allow the primary TrueCopy volumes to remain online to all hosts for both read and write I/O operations.

Hitachi TrueCopy - S/390[®] also supports both serial (ESCON[®]) and fibre-channel interface connections between the main and remote 9900 subsystems. Remote copy operations can also be performed between 9900, 7700E, and 7700 subsystems (with some restrictions, e.g., the 7700 does not support asynchronous remote copy).

Note: For further information on Hitachi TrueCopy - S/390[®], please see the *Hitachi Lightning 9900™ Hitachi TrueCopy - S/390[®] User and Reference Guide* (MK-91RD050), or contact your Hitachi Data Systems account team.

3.7.3 Hitachi ShadowImage (SI)

Hitachi ShadowImage enables open-system users to maintain subsystem-internal copies of LUs for purposes such as data backup or data duplication. The RAID-protected duplicate LUs (up to nine) are created within the same 9900 subsystem as the primary LU at hardware speeds. Once established, ShadowImage operations continue unattended to provide asynchronous internal data backup. ShadowImage operations are nondisruptive; the primary LU of each ShadowImage pair remains available to all hosts for both read and write operations during normal operations. Usability is further enhanced through a resynchronization capability that reduces data duplication requirements and backup time, thereby increasing user productivity. ShadowImage also supports reverse resynchronization for maximum flexibility.

ShadowImage operations can be performed in conjunction with Hitachi TrueCopy operations (see section 3.7.1) to provide multiple copies of critical data at both primary and remote sites. ShadowImage also supports the Virtual LVI/LUN and FlashAccess features of the 9900 subsystem, ensuring that all user data can be duplicated by ShadowImage operations.

Note: For further information on Hitachi ShadowImage, please see the *Hitachi Lightning 9900™ ShadowImage User's Guide (MK-90RD031)*, or contact your Hitachi Data Systems account team.

3.7.4 Hitachi ShadowImage – S/390® (SI390)

Hitachi ShadowImage - S/390® enables S/390® users to create high-performance copies of source LVIs for testing or modification while benefiting from full RAID protection for the ShadowImage copies. The ShadowImage copies can be available to the same or different logical partitions (LPARs) as the original volumes for read and write I/Os. ShadowImage allows the user to create up to three copies of a single source LVI and perform updates in either direction, either from the source LVI to the ShadowImage copy or from the copy back to the source LVI. When used in conjunction with either TrueCopy - S/390® or HXRC, ShadowImage - S/390® enables users to maintain multiple copies of critical data at both primary and remote sites. ShadowImage also supports the Virtual LVI/LUN and FlashAccess features, ensuring that all user data can be duplicated by ShadowImage operations

Note: For further information on Hitachi ShadowImage - S/390®, please see the *Hitachi Lightning 9900™ ShadowImage - S/390® User's Guide (MK-90RD012)*, or contact your Hitachi Data Systems account team.

3.7.5 Command Control Interface (CCI)

Hitachi Command Control Interface (CCI) enables users to perform Hitachi TrueCopy and Hitachi ShadowImage operations on the Lightning 9900™ subsystem by issuing commands from the UNIX®/PC server host to the 9900 subsystem. The CCI software interfaces with the system software and high-availability (HA) software on the UNIX®/PC server host as well as the TrueCopy/ShadowImage software on the 9900 subsystem. The CCI software provides failover and other functions such as backup commands to allow mutual hot standby in cooperation with the failover product on the UNIX®/PC server (e.g., MC/ServiceGuard®, FirstWatch®, HACMP).

CCI also supports a scripting function that allows users to define multiple TrueCopy and/or ShadowImage operations in a script (text) file. Using CCI scripting, you can set up and execute a large number of TrueCopy and/or ShadowImage commands in a short period of time while integrating host-based high-availability control over remote copy operations.

Note: For further information on CCI, please see the *Hitachi Lightning 9900™ Command Control Interface (CCI) User and Reference Guide (MK-90RD011)*, or contact your Hitachi Data Systems account team.

3.7.6 Extended Copy Manager (ECM)

Extended Copy Manager (ECM), together with the backup application on the open-system host server, provides server-free backup solutions between the 9900 subsystem and backup devices such as tape and disk devices. Extended Copy Manager enables non-disruptive backup directly from disk to tape (or disk) in storage-area network (SAN) environments, eliminating server CPU and I/O overhead during movement of data and decreasing the time required for backup. Users can perform copy (backup) operations on the data (in units of block) stored on the multiplatform 9900 subsystem via fibre-channel interfaces directly to the backup devices.

Extended Copy Manager supports the SCSI Extended Copy command issued from the host server to the 9900 subsystem. The 9900 subsystem receives the Extended Copy commands issued by the server, and then copies the data directly to the specified backup device. ECM operations are configured using the LUN Manager software on the Remote Console PC. To implement ECM operations, you need a backup application on the host server to issue the Extended Copy commands to the 9900 subsystem.

Note: For further information on Extended Copy Manager, please see the *Hitachi Lightning 9900™ LUN Manager User's Guide (MK-91RD049)*, or contact your Hitachi Data Systems account team.

3.7.7 Hitachi Extended Remote Copy (HXRC)

The HXRC asynchronous remote copy feature of the 9900 subsystem is functionally compatible with IBM® Extended Remote Copy (XRC). HXRC provides asynchronous remote copy operations for maintaining duplicate copies of S/390® data for data backup purposes. Once established, HXRC operations continue unattended to provide continuous data backup. HXRC operations are nondisruptive and allow the primary HXRC volumes to remain online to the host(s) for both read and write I/O operations. For HXRC operations, there is no distance limit between the primary and remote disk subsystems. HXRC is also compatible with the DFSMS data mover that is common to the XRC environment.

HXRC operations are performed in the same manner as XRC operations. The user issues standard XRC TSO commands from the mainframe host system directly to the 9900 subsystem. The Remote Console PC is not used to perform HXRC operations. HXRC can be used as an alternative to Hitachi TrueCopy - S/390® for mainframe data backup and disaster recovery planning. However, HXRC requires host processor resources that may be significant for volumes with high-write activity. The Data Mover utility may run in either the primary host or the optional remote host.

Note: For 9900-specific information on HXRC (e.g. SVP modes), please see the *Hitachi Lightning 9900™ Hitachi TrueCopy - S/390® User and Reference Guide (MK-91RD050)*, or contact your Hitachi Data Systems account team.

Note: For further information on XRC, please refer to the following IBM® publications: *Planning for IBM Remote Copy (SG24-2595)*, *Advanced Copy Services (SC35-0355)*, and *Remote Copy Administrator's Guide and Reference (SC35-0169)*.

3.7.8 Hitachi NanoCopy™

Hitachi NanoCopy™ is the storage industry's first hardware-based solution which enables customers to make Point-in-Time (PiT) copies without quiescing the application or causing any disruption to end-user operations. NanoCopy™ is based on TC390 Asynchronous (TC390A), which is used to move large amounts of data over any distance with complete data integrity and minimal impact on performance. TC390A can be integrated with third-party channel extender products to address the “access anywhere” goal of data availability. TC390A enables production data to be duplicated via ESCON® or communication lines from a main (primary) site to a remote (secondary) site that can be thousands of miles away.

NanoCopy™ copies data between any number of primary subsystems and any number of secondary subsystems, located any distance from the primary subsystem, without using valuable server processor cycles. The copies may be of any type or amount of data and may be recorded on subsystems anywhere in the world.

NanoCopy™ enables customers to quickly generate copies of production data for such uses as application testing, business intelligence, and disaster recovery for business continuance. For disaster recovery operations, NanoCopy™ will maintain a duplicate of critical data, allowing customers to initiate production at a backup location immediately following an outage. This is the first time an asynchronous hardware-based remote copy solution, with full data integrity, has been offered by any storage vendor.

Hitachi TrueCopy - S/390® Asynchronous with Hitachi NanoCopy™ support is offered as an extension to Hitachi Data Systems' data movement options and software solutions for the Hitachi Lightning 9900™. Hitachi ShadowImage - S/390® can also operate in conjunction with TC390 Synchronous and Asynchronous to provide volume-level backup and additional image copies of data. This delivers an additional level of data integrity to assure consistency across sites and provides flexibility in maintaining volume copies at each site.

Note: For further information on Hitachi NanoCopy™, please contact your Hitachi Data Systems account team.

3.7.9 Data Migration

The Lightning 9900™ subsystem supports data migration operations from other disk array subsystems, including older Hitachi subsystems as well as other vendors' subsystems. Data can be moved to a new location either temporarily or as part of a data relocation process. During normal migration operations, the data being migrated can be online to the host(s) for both read and write I/O operations during data migration operations.

Note: Data migration is available as a Hitachi Data Systems service offering. For further information on data migration, please contact your Hitachi Data Systems account team.

3.7.10 Hitachi RapidXchange (HRX)

Hitachi RapidXchange (HRX) enables the user to transfer data between S/390® and open-system platforms using the ExSA™ channels or fibre channels. HRX enables high-speed data transfer without requiring network communication links or tape. Data transfer is performed via the HRX volumes, which are shared devices that appear to the S/390® host as 3390-3 or 3380-K LVIs and to the open-system host as OPEN-3 or OPEN-K LUs. To provide the greatest platform flexibility for data transfer, the HRX volumes are accessed from the open-system host using SCSI raw device mode.

HRX allows the open-system host to read from and write to S/390® sequential datasets using the HRX volumes. The HRX volumes must be formatted as 3390-3A/B/C or 3380-KA/B/C LVIs. The -A LVIs can be used for open-to-mainframe and/or mainframe-to-open HRX, the -B LVIs are used for mainframe-to-open HRX, and the -C LVIs are used for open-to-mainframe HRX. HRX also supports OPEN-x-HRX devices to provide open-to-open HRX operations for all-open 9900 subsystems.

The HRX software enables the open-system host to read from and write to individual S/390® datasets. The HRX software is installed on the open-system host and includes the File Conversion Utility (FCU) and the File Access Library (FAL). FCU allows the user to set up and perform file conversion operations between S/390® sequential datasets and open-system flat files. The FAL is a library of C-language functions that allows open-system programmers to read from and write to S/390® sequential datasets on the HRX volumes.

Note: For further information on HRX, please see the *Hitachi Lightning 9900™ RapidXchange (HRX) User's Guide* (MK-91RD052), or contact your Hitachi Data Systems account team.

3.7.11 Hitachi Multiplatform Backup/Restore (HMBR)

HMBR allows the user to implement mainframe-based backup procedures and standards for the open-system data stored on the multiplatform 9900 subsystem. HMBR enables standard mainframe backup/restore utilities such as DFDSS, Fast Dump/Restore (FDR), and VSE FASTWRITE to perform volume-level backup and restore operations on OPEN-3 and OPEN-9 LUs. Using these mainframe-based utilities as well as mainframe-based media and high-speed backup devices, the user can use the same procedures and achieve the same standards for both mainframe and open-system backup/restore operations. Before HMBR operations can begin, an offline utility such as ICKDSF must be used to create a volume table of contents (VTOC) to enable the mainframe host to use the OPEN-x LUs as mainframe volumes, which contain a single file. HMBR supports only full-volume backup/restore operations.

Note: For further information on HMBR, please see the *Hitachi Lightning 9900™ Hitachi Multiplatform Backup/Restore User's Guide* (MK-90RD037), or contact your Hitachi Data Systems account team.

3.7.12 HARBOR® File-Level Backup/Restore

Tantia™ Technologies HARBOR® File-Level Backup/Restore features an integrated architecture and includes:

- A host component on MVS,
- Integrated clients for desktops and servers,
- LAN-based distributed storage servers,
- High-speed HRX file-level backup of open-system data, and
- Transparent network support.

Note: For further information on HARBOR® File-Level Backup/Restore, please contact your Hitachi Data Systems account team.

3.7.13 HARBOR® File Transfer

Tantia™ Technologies HARBOR® file transfer adds automation to the process of transferring large data files at ultra-high channel speeds in either direction between open systems and mainframe servers. After automatically breaking large data files into more manageable pieces, HARBOR® File Transfer offers increased transfer speeds by directing data in multiple streams through the Lightning 9900™ storage subsystem.

Note: For further information on HARBOR® File Transfer, please contact your Hitachi Data Systems account team.

3.7.14 HiCommand™

HiCommand™ provides a consistent, easy to use, and easy to configure set of interfaces for managing Hitachi storage products including the Lightning 9900™ subsystem. HiCommand™ provides a web interface for real-time interaction with the storage arrays being managed, as well as a command line interface (CLI) for scripting. HiCommand™ gives storage administrators easier access to the existing Hitachi subsystem configuration, monitoring, and management features such as LUN Manager, SANTinel, TrueCopy, and ShadowImage. **Note:** HiCommand™ 1.x does not support all Hitachi subsystem functions.

HiCommand™ enables users to manage the 9900 subsystem and perform functions from virtually any location via the HiCommand™ Web Client, HiCommand™ command line interface (CLI), and/or third-party application. HiCommand™ displays detailed information on the configuration of the storage arrays added to the HiCommand™ system and allows you to perform important operations such as adding and deleting volume paths, securing logical units (LUs), and managing data replication operations.

Note: For further information on HiCommand™, please refer to the HiCommand™ user documentation (see Table 3.9), or contact your Hitachi Data Systems account team.

3.7.15 LUN Manager

LUN Manager enables users to set and define the port modes for fibre-channel ports and to set the fibre topology (e.g., FC-AL, fabric). Please connect your Hitachi Data Systems account team for further details on this feature.

Note: For further information on LUN Manager, please see the *Hitachi Lightning 9900™ LUN Manager User's Guide* (MK-91RD049), or contact your Hitachi Data Systems account team.

3.7.16 LU Size Expansion (LUSE)

The LUSE (LU Size Expansion) feature allows users to create virtual LUs that are larger than standard OPEN LUs, by expanding the size of a selected LU up to 36 times its normal size. The maximum size depends on the type of configuration. For example, you can expand an OPEN-9 LU to a maximum size of 265 GB (7.3 GB × 36). This capability enables open-system hosts to access the data on the entire 9900 subsystem using fewer logical units. LUSE allows host operating systems that have restrictions on the number of LUs per interface to access larger amounts of data.

Note: For further information on LUSE, please see the *Hitachi Lightning 9900™ LUN Manager User's Guide* (MK-91RD049), or contact your Hitachi Data Systems account team.

3.7.17 Virtual LVI/LUN

Virtual LVI/LUN allows users to convert fixed-size volumes into several smaller variable custom-sized volumes. Using the Remote Console PC, users can configure custom-size volumes by assigning a logical address and a specific number of cylinders/tracks (for S/390® data) or MB (for open-systems data) to each custom LVI/LU.

Virtual LVI/LUN improves data access performance by reducing logical device contention as well as host I/O queue times, which can occur when several frequently accessed files are located on a single volume. Multiple LVI/LU types can be configured within each array group. Virtual LVI/LUN enables the user to more fully utilize the physical storage capacity of the 9900, while reducing the amount of administrative effort required to balance I/O workloads. When Virtual LVI/LUN is used in conjunction with FlashAccess, the user can achieve even better data access performance than when either Virtual LVI/LUN or FlashAccess is used alone.

Note: For further information on Virtual LVI/LUN, please see the *Hitachi Lightning 9900™ Virtual LVI/LUN User's Guide* (MK-90RD005), or contact your Hitachi Data Systems account team.

3.7.18 FlashAccess

FlashAccess allows users to store specific data in cache memory. FlashAccess increases the data access speed for the cache-resident data by enabling read and write I/Os to be performed at front-end host data transfer speeds. The FlashAccess cache areas (called cache extents) are dynamic and can be added and deleted at any time. The 9900 subsystem supports up to 1,024 addressable cache extents.

FlashAccess operations can be performed for open-system LUs (e.g., OPEN-3, -8, -9) as well as S/390® LVIs (e.g., 3390-3/9, 3380-K), including custom-size volumes. Use of FlashAccess in conjunction with the Virtual LVI/LUN feature will achieve better performance improvements than when either of these options is used individually.

Note: For further information on FlashAccess, please see the *Hitachi Lightning 9900™ FlashAccess User's Guide* (MK-90RD004), or contact your Hitachi Data Systems account team.

3.7.19 Cache Manager

Cache Manager enables users to perform FlashAccess operations on S/390® LVIs by issuing commands from the S/390® host system to the 9900 subsystem. For further information on Cache Manager, please see the *Hitachi Lightning 9900™ Cache Manager User's Guide* (MK-91RD045), or contact your Hitachi Data Systems account team.

3.7.20 Hitachi SANtinel

Hitachi SANtinel allows users to restrict LU accessibility to an open-systems host using the host's World Wide Name (WWN). You can set an LU to communicate only with one or more specified WWNs, allowing you to limit access to that LU to specified open-system host(s). This feature prevents other open-systems hosts from either seeing the secured LU or accessing the data contained on it. The Hitachi SANtinel software for the Remote Console PC enables you to configure Hitachi SANtinel operations on the 9900 subsystem.

Hitachi SANtinel can be activated on any installed fibre-channel port, and be turned on or off at the port level. If you disable Hitachi SANtinel on a particular port, that LU will not be restricted to a particular host or group of hosts. If you enable Hitachi SANtinel on a particular port, that port will be restricted to a particular host or group of hosts. You can assign a WWN to as many ports as you want, and you can assign more than one WWN to each port. You can also change the WWN access for any port without disrupting the settings of that port.

Because up to 128 WWNs can access each port and the same WWNs may go to additional ports in the same subsystem, the Hitachi SANtinel software allows you to create LU and WWN groups, so you can more easily manage your 9900 storage subsystem. An LU group allows you to assign specified LUs to a single group name. A WWN group allows you to assign up to 128 WWNs to a single group. A WWN group gives every host in the specified WWN group access to the specified LU or group of LUs.

Note: For further information on Hitachi SANtinel, please see the *Hitachi Lightning 9900™ LUN Manager User's Guide* (MK-91RD049), or contact your Hitachi Data Systems account team.

3.7.21 Hitachi SANtinel – S/390®

Hitachi SANtinel - S/390® allows users to restrict S/390® host access to the logical devices (LDEVs) on the 9900 subsystem. Each LDEV can be set to communicate only with user-selected host(s). Hitachi SANtinel - S/390® prevents other hosts from seeing the secured LDEV and from accessing the data contained on the secured LDEV. The licensed Hitachi SANtinel - S/390® software on the 9900 Remote Console PC displays the Hitachi SANtinel - S/390® information and allows you to perform Hitachi SANtinel - S/390® operations.

Note: For further information on Hitachi SANtinel - S/390®, please see the *Hitachi Lightning 9900™ Hitachi SANtinel - S/390® User's Guide* (MK-90RD036), or contact your Hitachi Data Systems account team.

3.7.22 Prioritized Port and WWN Control (PPC)

Prioritized Port Control (PPC) allows open-system users to designate prioritized ports (e.g., for production servers) and non-prioritized ports (e.g., for development servers) and set thresholds and upper limits for the I/O activity of these ports. PPC enables users to tune the performance of the development server without affecting the production server's performance.

Note: For further information on PPC, please see the *Hitachi Lightning 9900™ Prioritized Port and WWN Control User's Guide* (MK-90RD030), or contact your Hitachi Data Systems account team.

3.7.23 Hitachi Parallel Access Volume (HPAV)

Hitachi Parallel Access Volume (HPAV) enables the S/390® host system to issue multiple I/O requests in parallel to single logical devices (LDEVs) in the Lightning 9900™ subsystem. HPAV can provide substantially faster host access to the S/390® data stored in the 9900 subsystem. The Workload Manager (WLM) host software function enables the S/390® host to utilize the HPAV functionality of the Lightning 9900™ subsystem. The 9900 supports both static and dynamic HPAV functionality.

Note: For further information on HPAV, please see the *Hitachi Lightning 9900™ Hitachi Parallel Access Volume (HPAV) User and Reference Guide* (MK-91RD047), or contact your Hitachi Data Systems account team.

3.7.24 Dynamic Link Manager™ (DLM)

Hitachi Dynamic Link Manager™ (DLM) provides automatic load balancing, path failover, and recovery capabilities in the event of a path failure. Dynamic Link Manager™ helps guarantee that no single path becomes overloaded while others are underutilized.

Note: For further information on DLM, please see the *Hitachi Lightning 9900™ Hitachi Dynamic Link Manager™ User's Guide* for your host platform (see Table 4.7), or contact your Hitachi Data Systems account team.

3.7.25 LDEV Guard

LDEV Guard enables you to assign access permissions (Read/Write, Read-Only, and Protect) to logical volumes in your disk subsystem using a Remote Console PC or a host command.

Note: For further information on LDEV Guard, please see the *Hitachi Lightning 9900™ LDEV User's Guide*, or contact your Hitachi Data Systems account team.

3.7.26 Hitachi CruiseControl

Hitachi CruiseControl enables users to optimize their data storage and retrieval on the 9900 subsystem. Hitachi CruiseControl analyzes detailed information on the usage of 9900 subsystem resources and tunes the 9900 automatically by migrating logical volumes within the subsystem according to detailed user-specified parameters. CruiseControl tuning operations can be used to resolve bottlenecks of activity and optimize volume allocation. CruiseControl operations are completely nondisruptive - the data being migrated can remain online to all hosts for read and write I/O operations throughout the entire volume migration process. CruiseControl also supports manual volume migration operations and estimates performance improvements prior to migration to assist you in tuning the 9900 subsystem for your operational environment.

Hitachi CruiseControl provides the following major benefits for the user:

- Load balancing of subsystem resources.
- Optimizing disk drive access patterns.
- Analysis of subsystem usage using GraphTool (provided with CruiseControl).

Note: For further information on Hitachi CruiseControl, please see the *Hitachi Lightning 9900™ Hitachi CruiseControl User's Guide* (MK-91RD054), or contact your Hitachi Data Systems account team.

3.7.27 Hitachi Graph-Track™

Hitachi Graph-Track™ (GT) allows users to monitor and collect detailed subsystem performance and usage statistics for the 9900 subsystem. GT can monitor as many as 32 subsystems on the 9900-internal LAN. GT monitors the hardware performance, cache usage, and I/O statistics of the attached subsystems and displays real-time and historical data as graphs that highlight key information such as peaks and trends. GT displays the following data for each attached subsystem:

- Subsystem configuration, including controller name, serial number, controller emulation, channel address(s), SSIDs, and cache size.
- LDEV configuration, including total storage capacity and RAID implementation for each array domain; hard disk drive capacity, LDEV type (e.g., 3390-3R, OPEN-3), and LDEV IDs for each array group.
- Subsystem usage, including percent busy versus time for the front-end microprocessors (CHIPs) and back-end microprocessors (ACPs).
- Cache statistics, including percent cache in use and percent write-pending data in cache.
- I/O statistics at the subsystem, array group, and LDEV levels: I/O rates, read/write ratio, read and write hits, backend transfer rates (drive-to-cache and cache-to-drive I/O rates).

In addition to displaying performance and usage data, Hitachi Graph-Track™ manages the collection and storage of the GT data automatically according to user-specified preferences. GT also allows the user to export GT data for use in reports or in other data analysis programs.

Note: For further information on Hitachi Graph-Track™, please see the *Hitachi Lightning 9900™ Hitachi Graph-Track™ User's Guide* (MK-90RD032), or contact your Hitachi Data Systems account team.

Chapter 4 Configuring and Using the 9900 Subsystem

4.1 S/390® Configuration

The first step in 9900 configuration is to define the subsystem to the S/390® host(s). The three basic areas requiring definition are:

- Subsystem ID (SSIDs),
- Hardware definitions, including I/O Configuration Program (IOCP) or Hardware Configuration Definition (HCD), and
- Operating system definitions (HCD or OS commands).

Note: The missing interrupt handler (MIH) value for the 9900 subsystem is 45 seconds without TrueCopy, and 60 seconds when TrueCopy operations are in progress. (The MIH value for data migration operations is 120 seconds.)

4.1.1 Subsystem IDs (SSIDs)

Subsystem IDs (SSIDs) are used for reporting information from the CU (or controller) to the operating system. The SSIDs are assigned by the user and must be unique to all connected host operating environments. Each group of 64 or 256 volumes requires one SSID, so there are one or four SSIDs per CU image. The first (lowest) SSID for each CU image must be divisible by four. The user-specified SSIDs are assigned during subsystem installation, and the 9900 Remote Console PC can also be used to assign and change SSIDs. Table 4.1 lists the SSID requirements.

Table 4.1 SSID Requirements

Controller Emulation	SSID Requirements	LVI Support
3990-3	0004 - 00FD	3380-x, 3390-x, OPEN-3, -8, -9, -K, -E, -L, -M
3990-6E	0004 - FFFD	3390-x, OPEN-3, -8, -9, -K, -E, -L, -M
2105*	0004 - FFFD	3390-x, OPEN-3, -8, -9, -K, -E, -L, -M

***Note:** HPAV operations require that one SSID be set for each set of 256 LDEVs.

4.2 S/390® Hardware Definition

4.2.1 Hardware Definition Using IOCP (MVS, VM, or VSE)

The I/O Configuration Program (IOCP) can be used to define the 9900 subsystem in MVS, VM, and VSE environments (wherever HCD cannot be used). The 9900 subsystem supports up to sixteen logical CU (LCU) images and 4096 LDEVs. Each LCU can hold up to 256 LDEV addresses. An LCU is the same as an IBM® logical sub-system (LSS). The CUADD parameter defines the CU images by number (0-F). The unit type can be 3990 or 2105.

Note: FICON™ support requires 2105-F20 emulation.

The following are cautions when using IOCP or HCD:

- Use FEATURE=SHARE for the devices if multiple LPARs/mainframes can access the volumes.
- 16,384 addresses per physical interface are allowed by MVS with FICON™ channels.
- Only 1024 addresses per physical interface are allowed by MVS with ExSA™ (ESCON®) channels. (This includes PAV base and alias addresses.)

Note: 4096 device addressing requires 16 CU images using CUADD=0 through CUADD=F in the CNTLUNIT statement.

Figure 4.1 shows a sample IOCP definition for a 9900 configured with:

- 2105 ID.
- Four FICON™ channel paths. Two channels paths are connected to a FICON™ switch. Two channel paths are directly connected to the 9900 subsystem.
- Six LCUs (0, 1, 2, 3, 4, 5) with 256 LVIs per control unit.
- Sixty-four (64) base addresses and 128 alias addresses per CU 0, 1, 2, and 3.
- One hundred twenty-eight (128) base addresses and 128 alias addresses per CU 4 and 5.

```

*****
**** Sample FICON CHPID / CNTLUNIT and IODEVICE ****
*****

CHPID PATH=(F9),TYPE=FC,SWITCH=07
CHPID PATH=(FB),TYPE=FC
CHPID PATH=(FD),TYPE=FC
CHPID PATH=(FF),TYPE=FC,SWITCH=07 *

CNTLUNIT CUNUMBR=8100,PATH=(F9,FB,FD,FF),
LINK=(04,**,**,05), *
UNIT=2105,CUADD=0,UNITADD=(00,256) *
CNTLUNIT CUNUMBR=8200,PATH=(F9,FB,FD,FF), *
LINK=(04,**,**,05), *
UNIT=2105,CUADD=1,UNITADD=(00,256) *
CNTLUNIT CUNUMBR=8300,PATH=(F9,FB,FD,FF), *
LINK=(04,**,**,05), *
UNIT=2105,CUADD=2,UNITADD=(00,256) *
CNTLUNIT CUNUMBR=8400,PATH=(F9,FB,FD,FF), *
LINK=(04,**,**,05), *
UNIT=2105,CUADD=3,UNITADD=(00,256) *
CNTLUNIT CUNUMBR=8500,PATH=(F9,FB,FD,FF), *
LINK=(04,**,**,05), *
UNIT=2105,CUADD=4,UNITADD=(00,256) *
CNTLUNIT CUNUMBR=8600,PATH=(F9,FB,FD,FF), *
LINK=(04,**,**,05), *
UNIT=2105,CUADD=5,UNITADD=(00,256) *

IODEVICE ADDRESS=(8100,064),CUNUMBR=(8100),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8180,128),CUNUMBR=(8100),STADET=Y,UNIT=3390A*
,UNITADD=80
IODEVICE ADDRESS=(8200,064),CUNUMBR=(8200),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8280,128),CUNUMBR=(8200),STADET=Y,UNIT=3390A*
,UNITADD=80
IODEVICE ADDRESS=(8300,064),CUNUMBR=(8300),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8380,128),CUNUMBR=(8300),STADET=Y,UNIT=3390A*
,UNITADD=80
IODEVICE ADDRESS=(8400,064),CUNUMBR=(8400),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8480,128),CUNUMBR=(8400),STADET=Y,UNIT=3390A*
,UNITADD=80
IODEVICE ADDRESS=(8500,128),CUNUMBR=(8500),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8580,128),CUNUMBR=(8500),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(8600,128),CUNUMBR=(8600),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8680,128),CUNUMBR=(8600),STADET=Y,UNIT=3390A

```

Figure 4.1 IOCP Definition for FICON™ Channels (direct connect and via FICON™ switch)

Figure 4.2 shows a sample IOCP hardware definition for a 9900 configured with:

- 3990 ID.
- Two (2) LPARs called PROD and TEST sharing 4 ExSA™ (ESCON®) channels connected via 2 ESCDs to the 9980.
- Each switch has ports C0 and C1 attached to the 9980.
- Four (4) LCUs with 256 LVIs per control unit.
- Two (2) cu statements per logical control unit.
- To protect data integrity due to multiple operating systems sharing these volumes, these devices require FEATURE=SHARED.

```

CHPID      PATH=(B0,B1),TYPE=CNC,PARTITION=(PROD,TEST,SHR),SWITCH=01
CHPID      PATH=(B2,B3),TYPE=CNC,PARTITION=(PROD,TEST,SHR),SWITCH=02
CNTLUNIT   CUNUMBR=A000,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=0,UNIT=3990
CNTLUNIT   CUNUMBR=A001,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=0,UNIT=3990
CNTLUNIT   CUNUMBR=A100,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=1,UNIT=3990
CNTLUNIT   CUNUMBR=A101,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=1,UNIT=3990
CNTLUNIT   CUNUMBR=A200,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=2,UNIT=3990
CNTLUNIT   CUNUMBR=A201,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=2,UNIT=3990
CNTLUNIT   CUNUMBR=A300,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=3,UNIT=3990
CNTLUNIT   CUNUMBR=A301,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=3,UNIT=3990
IODEVICE   ADDRESS=(A000,256),CUNUMBR=(A000,A001),UNIT=3390,FEATURE=SHARED
IODEVICE   ADDRESS=(A100,256),CUNUMBR=(A100,A101),UNIT=3390,FEATURE=SHARED
IODEVICE   ADDRESS=(A200,256),CUNUMBR=(A200,A201),UNIT=3390,FEATURE=SHARED
IODEVICE   ADDRESS=(A300,256),CUNUMBR=(A300,A301),UNIT=3390,FEATURE=SHARED

```

Note: 4096 device addressing requires 16 CU images using CUADD=0 through CUADD=F in the CNTLUNIT statement.

Figure 4.2 IOCP Definition for 1024 LVIs (9900 connected to host CPU(s) via ESCD)

Figure 4.3 shows a sample IOCP hardware definition for a 9900 with:

- 2105 ID.
- Eight (8) ExSA™ (ESCON®) channels directly connected to the 9900.
- Four (4) LCUs with 256 LVIs per control unit.
- One (1) cu statement per logical control unit.
- One hundred twenty-eight (128) 3390 base addresses per CU 0 and 1.
- One hundred twenty-eight (128) 3390 alias addresses per CU 0 and 1.
- Sixty-four (64) 3390 base addresses in CU 2.
- One hundred ninety-two (192) 3390 alias addresses in CU 2.
- One hundred twenty-eight (128) 3390 addresses in CU 3.
- Sixty-four (64) 3390 base addresses per CU 3.
- Sixty-four (64) 3390 alias addresses per CU 3.
- To protect data integrity due to multiple operating systems sharing these volumes, these devices require FEATURE=SHARED.

Note: If you maintain separate IOCP definitions files and create your SCDS or IOCDS manually by running the IZP IOCP program, you must define each LCU on a 9900 subsystem using one CNTLUNIT statement in IOCP. While it is possible to define an LCU on a 9900 subsystem using multiple CNTLUNIT statements in IOCP, the resulting input deck cannot be migrated to HCD due to an IBM® restriction allowing only one CNTLUNIT definition.

```

CHPID PATH=(60,61,62,63,64,65,66,67),TYPE=CNC
CNTLUNIT CUNUMBR=8000,PATH=(60,61,62,63,64,65,66,67),          *
        UNITADD=((00,256)),CUADD=0,UNIT=2105
CNTLUNIT CUNUMBR=8100,PATH=(60,61,62,63,64,65,66,67),          *
        UNITADD=((00,256)),CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=8200,PATH=(60,61,62,63,64,65,66,67),          *
        UNITADD=((00,256)),CUADD=2,UNIT=2105
CNTLUNIT CUNUMBR=8300,PATH=(60,61,62,63,64,65,66,67),          *
        UNITADD=((00,256)),CUADD=3,UNIT=2105
IODEVICE ADDRESS=(8000,128),CUNUMBR=(8000),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8080,128),CUNUMBR=(8000),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8100,128),CUNUMBR=(8100),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8180,128),CUNUMBR=(8100),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8200,064),CUNUMBR=(8200),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8240,192),CUNUMBR=(8200),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8300,128),CUNUMBR=(8300),STADET=Y,UNIT=3390,FEATURE=SHARED
IODEVICE ADDRESS=(8380,064),CUNUMBR=(8300),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(83C0,064),CUNUMBR=(8300),STADET=Y,UNIT=3390A,FEATURE=SHARED

```

Figure 4.3 IOCP Definition for 1024 LVIs (9900 directly connected to CPU)

The 9960 subsystem can be configured with up to 32 connectable physical paths to provide up to 32 concurrent host data transfers. The 9910 subsystem can be configured with up to 24 connectable physical paths to provide up to 24 concurrent host data transfers. Since only 16 channel interface IDs are available (due to 16 physical channel interfaces for IBM® systems), the 9900 uses one channel interface ID for each pair of physical paths. For example, link control processors (LCPs) 1A and 1B correspond to channel interface ID 08 (00), and LCPs 1C and 1D correspond to channel interface ID 09 (01). Table 4.2 illustrates the correspondence between physical paths and channel interface IDs on Cluster 1, and Table 4.3 illustrates the same for Cluster 2.

Table 4.2 Correspondence between Physical Paths and Channel Interface IDs (Cluster 1)

Channel Interface ID (used in commands)	08 (00)	09 (01)	0A (02)	0B (03)	0C (04)	0D (05)	0E (06)	0F (07)
LCP ID	1A & 1B	1C & 1D	1E & 1F	1G & 1H	1J & 1K	1L & 1M	1N & 1P	1Q & 1R

Table 4.3 Correspondence between Physical Paths and Channel Interface IDs (Cluster 2)

Channel Interface ID (used in commands)	18 (10)	19 (11)	1A (12)	1B (13)	1C (14)	1D (15)	1E (16)	1F (17)
LCP ID	2A & 2B	2C & 2D	2E & 2F	2G & 2H	2J & 2K	2L & 2M	2N & 2P	2Q & 2R

4.2.2 Hardware Definition Using HCD (MVS/ESA)

The Hardware Configuration Definition (HCD) utility can be used to define the 9900 subsystem in an MVS/ESA environment. The HCD procedures for 3990 and 2105 controller emulation types are described below. FICON™ support requires 2105-F20 emulation.

3990 Controller Emulation. To define a 9900 subsystem with 64 or fewer LVIs, use the same procedure as for an IBM® 3990-6, 3990-6E, or 3990-3 subsystem (see Table 4.4). The hardware definition for a 9900 subsystem with more than 64 LVIs (see Table 4.5) is different than that for an IBM® 3990 subsystem.

Table 4.4 HCD Definition for 64 LVIs

Parameter	Value
Control Frame:	
Control unit number	Specify the control unit number.
Control unit type	3990-6 or 3990-6E (using 3990-6 emulation) 3990-3 (using 3990-3 emulation)
Channel path IDs	Specify ExSA™ or FICON™
Unit address	00 (ExSA™ or FICON™)
Number of units	64
Array Frame:	
Device number	Specify the first device number.
Number of devices	64
Device type	3390
Connected to CUs	Specify the control unit number(s).

Table 4.5 HCD Definition for 256 LVIs

Parameter	Value
Control Frame:	
Control unit number	Specify the control unit number.
Control unit type	NOCHECK* Use UIM 3990 for more than 128 logical paths Use UIM 3990-6 for 128 or fewer logical paths
Channel path IDs	Specify ExSA™ or FICON™
Unit address	00 (ExSA™ or FICON™)
Number of units	256
Array Frame:	
Device number	Specify the first device number.
Number of devices	256
Device type	3390
Connected to CUs	Specify the control unit number(s).

***Note:** The NOCHECK function was introduced by APAR OY62560. Defining the 9900 as a single control unit allows all channel paths to access all DASD devices.

2105 Controller Emulation. To define a 9900 logical control unit (LCU) and the base and alias address range that it will support, please use the following example for HCD.

Note: The following HCD steps correspond to the 2105 IOCP definition shown in Figure 4.3.

Note: The HCD PAV definitions must match the configurations in the 9900 subsystem. If not, error messages will occur when the HOSTs are IPL'd or the devices are varied online.

1. From an ISPF/PDF primary options menu, select the HCD option to display the basic HCD panel (see Figure 4.4). On this panel you must verify the name of the IODF or IODF.WORK I/O definition file to be used.
2. On the basic HCD panel (see Figure 4.5), select the proper I/O definition file, and then select option 1 to display the Define, Modify, or View Configuration Data panel.
3. On the Define, Modify, or View Configuration Data panel (see Figure 4.6), select option 4 to display the Control Unit List panel.
4. On the Control Unit List panel (see Figure 4.7), if a 2105 type of control unit already exists, then an 'Add like' operation can be used by inputting an 'A' next to the 2105 type control unit and pressing the return key. Otherwise press F11 to add a new control unit.
5. On the Add Control Unit panel (see Figure 4.8), input the following new information, or edit the information if preloaded from an 'Add like' operation, and then press enter key:
Control unit number
Control unit type - 2105
Switch information only if a switch exists. Otherwise leave switch and ports blank.
6. On the Select Processor / Control Unit panel (see Figure 4.9), input an S next to the PROC. ID, and then press return key.
7. On the Add Control Unit panel (see Figure 4.10), enter chpids that attach to the control unit, the logical control unit address, the device starting address, and the number of devices supported, and then press return key.
8. Verify that the data is correct on the Select Processor / Control Unit panel (see Figure 4.11), and then press F3.
9. On the Control Unit List panel (see Figure 4.12), add devices to the new Control Unit, input an S next to CU 8000, and then press enter.
10. On the I/O Device List panel (see Figure 4.13), press F11 to add new devices.
11. On the Add Device panel (see Figure 4.14), enter the following, and then press return:
Device number
Number of devices
Device type: 3390, 3390B for HPAV base device, or 3390A for HPAV alias device.
12. On the Device / Processor Definition panel (see Figure 4.15), add this device to a specific Processor/System-ID combination by inputting an S next to the Processor and then pressing the return key.
13. On the Define Device / Processor panel, enter the values shown in Figure 4.16, and press the return key.

14. On the Define Processor / Definition panel (see Figure 4.17), verify that the proper values are displayed, and press the return key.
15. On the Define Device to Operating System Configuration panel, input an S next to the Config ID (see Figure 4.18), and then press the return key.
16. The Define Device Parameters / Features panel displays the default device parameters (see Figure 4.19). **Note:** The WLMPAV parameter defaults to “YES”. Set the desired parameters, and then press return key.
17. This returns to the Define Device to Operating System Configuration Panel. Press F3.
18. The Update Serial Number, Description and VOLSER panel now displays the desired device addresses (see Figure 4.20). To add more control units or device addresses, repeat the previous steps.

```

San Diego OS/390 R2.8 Master MENU
OPTION ==> HC                                SCROLL ==> PAGE

                                USERID - HDS
                                TIME   - 20:23

IS ISMF   - Interactive Storage Management Facility
P  PDF    - ISPF/Program Development Facility
IP IPCS   - Interactive Problem Control Facility
R  RACF   - Resource Access Control Facility
SD SDSF   - System Display and Search Facility
HC HCD    - Hardware Configuration Definition
BMB BMR BLD - BookManager Build (Create Online Documentation)
BMR BMR READ - BookManager Read (Read Online Documentation)
BMI BMR INDX - BookManager Read (Create Bookshelf Index)
SM SMP/E   - SMP/E Dialogs
IC ICSF   - Integrated Cryptographic Service Facility
OS SUPPORT - OS/390 ISPF System Support Options
OU USER   - OS/390 ISPF User Options
S  SORT   - DF/SORT Dialogs
X  EXIT   - Terminate ISPF using list/log defaults

F1=HELP    F2=SPLIT    F3=END      F4=RETURN   F5=RFIND    F6=RCHANGE
F7=UP      F8=DOWN      F9=SWAP    F10=LEFT   F11=RIGHT   F12=RETRIEVE

```

Figure 4.4 Master MENU (Step 1)

```

OS/390 Release 5 HCD
Command ==> _____

Hardware Configuration

Select one of the following.

1_ 1. Define, modify, or view configuration data
   2. Activate or process configuration data
   3. Print or compare configuration data
   4. Create or view graphical configuration report
   5. Migrate configuration data
   6. Maintain I/O definition files
   7. Query supported hardware and installed UIMs
   8. Getting started with this dialog
   9. What's new in this release

For options 1 to 5, specify the name of the IODF to be used.

I/O definition file . . . 'HDS.IODF06.WORK'          +
-----
| (C) Copyright IBM Corp. 1990, 1998. All rights reserved. |
-----
F1=Help  F2=Split  F3=Exit  F4=Prompt  F9=Swap  F12=Cancel

```

Figure 4.5 Basic HCD Panel (Step 2)

```

OS/390 Release 5 HCD
----- Define, Modify, or View Configuration Data -----
|
| Select type of objects to define, modify, or view data.
|
| 4_ 1. Operating system configurations
|     consoles
|     system-defined generics
|     EDTs
|     esoterics
|     user-modified generics
| 2. Switches
|     ports
|     switch configurations
|     port matrix
| 3. Processors
|     partitions
|     channel paths
| 4. Control units
| 5. I/O devices
|
| F1=Help  F2=Split  F3=Exit  F9=Swap  F12=Cancel
|
|-----

```

Figure 4.6 Define, Modify, or View Configuration Data (Step 3)

```

Goto  Filter  Backup  Query  Help
-----
                               Control Unit List                               Row 27 of 40
Command ==> _____ Scroll ==> PAGE

Select one or more control units, then press Enter.  To add, use F11.

/ CU  Type +      Serial-# + Description
- 3107 SCTC      _____
- 3108 SCTC      _____
- 3109 SCTC      _____
- 310A SCTC      _____
- 4000 3990-6    _____
- 4100 3990-6    _____
- 4200 3990-6    _____
- 4300 3990-6    _____
- 5000 3990      _____
- 5001 3990      _____
- 6000 3990      _____
- 6001 3990      _____
- 7000 3990      _____
- 7001 3990      _____

F1=Help    F2=Split    F3=Exit    F4=Prompt    F5=Reset    F7=Backward
F8=Forward F9=Swap     F10=Actions F11=Add     F12=Cancel

```

Figure 4.7 Control Unit List Panel (Step 4)

```

----- Add Control Unit -----
|
| Specify or revise the following values.
|
| Control unit number . . . . . 8000 +
|
| Control unit type . . . . . 2105_____ +
|
| Serial number . . . . . _____
| Description . . . . . new 2105 type for 80xx devices__
|
| Connected to switches . . . _ _ _ _ _ +
| Ports . . . . . _ _ _ _ _ _ _ _ _ _ _ _ +
|
| If connected to a switch, select whether to have CHPIDs/link
| addresses, and unit address range proposed.
|
| Auto-assign . . . . . 2 1. Yes
|                   2. No
|
| F1=Help    F2=Split    F3=Exit    F4=Prompt    F5=Reset    F9=Swap
| F12=Cancel
|
-----

```

Figure 4.8 Add Control Unit Panel (Step 5)

```

Goto Filter Backup Query Help
----- Select Processor / Control Unit -----
Row 1 of 1 More: > |
Command ==> _____ Scroll ==> PAGE |
|
| Select processors to change CU/processor parameters, then press Enter. |
|
| Control unit number . . : 8000      Control unit type . . . : 2105 |
|
|           Log. Addr. -----Channel Path ID . Link Address + ----- |
| / Proc. ID Att. (CUADD) + 1---- 2---- 3---- 4---- 5---- 6---- 7---- 8---- |
| S PROD |
| ***** Bottom of data ***** |
|
|
|
|
|
|
|
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset |
| F6=Previous  F7=Backward   F8=Forward   F9=Swap       F12=Cancel  |
|-----|
| F6=Previous  F7=Backward   F8=Forward   F9=Swap       F12=Cancel  |
|-----|

```

Figure 4.9 Selecting the Operating System (Step 6)

```

Goto Filter Backup Query Help
----- Add Control Unit -----
|
| Specify or revise the following values. |
|
| Control unit number . . : 8000      Type . . . . . : 2105 |
| Processor ID . . . . . : PROD |
|
| Channel path IDs . . . . 60  61  62  63  64  65  66  67  + |
| Link address . . . . . _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ + |
|
| Unit address . . . . . 00  _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ + |
| Number of units . . . . 256  _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
|
| Logical address . . . . 0_ + (same as CUADD) |
|
| Protocol . . . . . _ _ + (D,S or S4) |
| I/O concurrency level . 2 + (1, 2 or 3) |
|
|
|
|
|
|
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F9=Swap |
| F12=Cancel |
|-----|

```

Figure 4.10 Control Unit Chpid, CUADD, and Device Address Range Addressing (Step 7)

```

Goto Filter Backup Query Help
----- Select Processor / Control Unit -----
Row 1 of 1 More: >
Command ==> _____ Scroll ==> PAGE
Select processors to change CU/processor parameters, then press Enter.
Control unit number . . : 8000      Control unit type . . . : 2105

      Log. Addr. -----Channel Path ID . Link Address + -----
/ Proc. ID Att. (CUADD) + 1---- 2---- 3---- 4---- 5---- 6---- 7---- 8----
_ PROD   Yes 0      60   61   62   63   64   65   66   67
***** Bottom of data *****

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
F6=Previous  F7=Backward   F8=Forward   F9=Swap        F12=Cancel
-----

```

Figure 4.11 Select Processor / Control Unit Panel (Step 8)

```

Goto Filter Backup Query Help
-----
Control Unit List
Row 40 of 41
Command ==> _____ Scroll ==> PAGE
Select one or more control units, then press Enter. To add, use F11.

/ CU  Type +      Serial-# + Description
_ 7001 3990
S 8000 2105      _____ add 2105 type for 80xx devices
***** Bottom of data *****

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F7=Backward
F8=Forward   F9=Swap       F10=Actions  F11=Add        F12=Cancel
-----

```

Figure 4.12 Control Unit List (Step 9)

```

Goto Filter Backup Query Help
-----
I/O Device List
Command ==> _____ Scroll ==> PAGE

Select one or more devices, then press Enter. To add, use F11.

Control unit number : 8000      Control unit type . : 2105

-----Device----- --#-- -----Control Unit Numbers + -----
/ Number Type +          PR OS 1--- 2--- 3--- 4--- 5--- 6--- 7--- 8--- Base
***** Bottom of data *****
_____

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F7=Backward
F8=Forward   F9=Swap       F10=Actions  F11=Add       F12=Cancel

```

Figure 4.13 I/O Device List Panel (Step 10)

```

Goto Filter Backup Query Help
----- Add Device -----
|
| Specify or revise the following values.
|
| Device number . . . . . 8000 (0000 - FFFF)
| Number of devices . . . . . 128
| Device type . . . . . 3390B_____ +
|
| Serial number . . . . . _____
| Description . . . . . 3390 Base addresses 8000-807F__
|
| Volume serial number . . . . . _____ (for DASD)
|
| Connected to CUs . . 8000 _____ +
|
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F9=Swap
| F12=Cancel
|
-----

```

Figure 4.14 Add Device Panel (Step 11)


```

.----- Device / Processor Definition -----
|                                                                 |
|                                                                 |
|                                                                 |
| Command ==> _____ Scroll ==> PAGE |
|                                                                 |
| Select processors to change device/processor definitions, then press |
| Enter. |
|                                                                 |
| Device number . . : 8100          Number of devices . . : 128 |
| Device type . . . : 3390B |
|                                                                 |
|                                                                 |
| / Processor ID  UA + Time-Out STADET CHPID + Preferred Explicit Device |
| s PROD          ___ No       Yes   ___   Candidate List             ***** |
| ***** Bottom of data ***** |
|                                                                 |
|                                                                 |
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset |
| F6=Previous  F7=Backward   F8=Forward   F9=Swap       F12=Cancel |
|                                                                 |
.-----

```

Figure 4.15 Device / Processor Definition Panel – Selecting the Processor ID (Step 12)

```

.----- Define Device / Processor -----
|
| Specify or revise the following values.
|
| Device number . . : 8000          Number of devices . . . . : 128
| Device type . . . : 3390B
| Processor ID . . . : PROD
|
| Unit address . . . . . 00 + (Only necessary when different from
|                               the last 2 digits of device number)
| Time-Out . . . . . No (Yes or No)
| STADET . . . . . Yes (Yes or No)
|
| Preferred CHPID . . . . . ___ +
| Explicit device candidate list . No (Yes or No)
|
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F9=Swap
| F12=Cancel
|
.-----

```

Figure 4.16 Define Device / Processor Panel (Step 13)

```

----- Device / Processor Definition -----
|                                                                 Row 1 of 1 |
| Command ==> _____ Scroll ==> PAGE |
|                                                                 |
| Select processors to change device/processor definitions, then press |
| Enter. |
|                                                                 |
| Device number . . : 8000          Number of devices . : 128 |
| Device type . . . : 3390B |
|                                                                 |
|                                                                 Preferred Explicit Device |
| / Processor ID  UA + Time-Out  STADET  CHPID +  Candidate List |
| _ PROD          00   No         Yes      ___      No |
| ***** Bottom of data ***** |
|                                                                 |
|                                                                 |
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset |
| F6=Previous  F7=Backward  F8=Forward   F9=Swap       F12=Cancel |
|                                                                 |
-----

```

Figure 4.17 Device / Processor Definition Panel (Step 14)

```

----- Define Device to Operating System Configuration -----
|                                                                 Row 1 of 1 |
| Command ==> _____ Scroll ==> PAGE |
|                                                                 |
| Select OSs to connect or disconnect devices, then press Enter. |
|                                                                 |
| Device number . . : 8100          Number of devices : 128 |
| Device type . . . : 3390B |
|                                                                 |
| / Config. ID  Type      Description          Defined |
| s PROD       MVS |
| ***** Bottom of data ***** |
|                                                                 |
|                                                                 |
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset |
| F6=Previous  F7=Backward  F8=Forward   F9=Swap       F12=Cancel |
|                                                                 |
-----

```

Figure 4.18 Define Device to Operating System Configuration (Step 15)

```

----- Define Device Parameters / Features -----
|                                                                 Row 1 of 6 |
| Command ==> _____ Scroll ==> PAGE |
|
| Specify or revise the values below.
|
| Configuration ID . . : PROD
| Device number . . . : 8000          Number of devices : 128
| Device type . . . . : 3390B
|
| Parameter/
| Feature   Value   P Req.  Description
| OFFLINE   No      P Req.  Device considered online or offline at IPL
| DYNAMIC   Yes     P Req.  Device supports dynamic configuration
| LOCANY    No      P Req.  UCB can reside in 31 bit storage
| WLMPAV    Yes     P Req.  Device supports work load manager
| SHARED    Yes     P Req.  Device shared with other systems
| SHAREDUP  No      P Req.  Shared when system physically partitioned
| ***** Bottom of data *****
|
| F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
| F7=Backward  F8=Forward    F9=Swap     F12=Cancel
-----

```

Figure 4.19 Define Device Parameters / Features Panel (Step 16)

```

----- Update Serial Number, Description and VOLSER -----
|                                                                 Row 1 of 128 |
| Command ==> _____ Scroll ==> PAGE |
|
| Device number . . . : 8000          Number of devices : 128
| Device type . . . . : 3390B
|
| Specify or revise serial number, description and VOLSER.
|
| Device Number Serial-#  Description          VOLSER
| 8000                _____ 3390 Base addresses 8000-807F  _____
| 8001                _____ 3390 Base addresses 8000-807F  _____
| 8002                _____ 3390 Base addresses 8000-807F  _____
| 8003                _____ 3390 Base addresses 8000-807F  _____
| 8004                _____ 3390 Base addresses 8000-807F  _____
| 8005                _____ 3390 Base addresses 8000-807F  _____
| 8006                _____ 3390 Base addresses 8000-807F  _____
| 8007                _____ 3390 Base addresses 8000-807F  _____
| 8008                _____ 3390 Base addresses 8000-807F  _____
| 8009                _____ 3390 Base addresses 8000-807F  _____
| 800A                _____ 3390 Base addresses 8000-807F  _____
| F1=Help      F2=Split      F3=Exit      F5=Reset      F7=Backward
| F8=Forward    F9=Swap     F12=Cancel
-----

```

Figure 4.20 Update Serial Number, Description and VOLSER Panel (Step 18)

4.2.3 Defining the 9900 to VM/ESA® Systems

64 or Fewer LVIs: To define a 9900 with less than or equal to 64 LVIs to a VM/ESA® system, use the same procedure as for an IBM® 3990-6, 3990-6E, or 3990-3 subsystem. To define a 9900 with more than 64 LVIs to VM/ESA®, enter the LVI address range, storage type, and sharing option for the subsystem as shown below (the address range varies for each installation).

[Address Range] TYPE DASD SHARED YES

More than 64 LVIs: To define a 9900 with more than 64 LVIs to VSE/ESA, use the same procedure as for an IBM® 3990-6, 3990-6E, or 3990-3 subsystem. For 9900 subsystems with more than 64 LVIs, the **ADD cuu:cuu** ECKD statements are the same as for the IBM® 3390.

Caution: There may be APAR requirements at different versions of the VM software. In addition, certain requirements apply when VM is supporting guests (e.g., OS/390®) that can run 2105 native. For further information, please refer to the IBM® documentation.

4.2.4 Defining the 9900 to TPF

The 9900 supports the IBM® Transaction Processing Facility (TPF) and Multi-Path Locking Facility (MPLF) functions in either native mode or under VM. MPLF support requires TPF version 4.1 or higher, and RAID-5+ and 3390-3/3R LVIs are supported. The 9900's TPF/MPLF capability enables high levels of concurrent data access across multiple channel paths. For further information on TPF and MPLF operations, please refer to the following IBM® documentation:

- Storage Subsystem Library, 3390 Transaction Processing Facility Support RPQs, IBM® document number GA32-0134-03.
- Storage Subsystem Library, 3990 Storage Control Reference for Model 6, IBM® document number GA32-0274-03.

Note on 2105 emulation: There are available PTFs to implement exploitation mode for TPF version 4.1. For further information, please refer to the IBM® documentation.

4.3 S/390® Operations

4.3.1 Initializing the LVIs

The 9900 LVIs require only minimal initialization before being brought online. Figure 4.21 shows an MVS ICKDSF JCL example of a minimal init job to write a volume ID (VOLID) and volume table of contents (VTOC).

Note: HPAV base and alias devices require additional definition. For further information, please refer to the *Hitachi Lightning 9900™ Hitachi Parallel Access Volume (HPAV) User and Reference Guide (MK-91RD047)*.

```
// EXAMPLE JOB
// EXEC PGM=ICKDSF
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
INIT UNITADDRESS (XXXX) NOVERIFY VOLID(YYYYYY) -
OWNERID(ZZZZZZ)
/*
```

Note: X = physical install address, Y = new volume ID, Z = volume ID owner.

Figure 4.21 LVI Initialization for MVS: ICKDSF JCL

4.3.2 Device Operations: ICKDSF

The 9900 subsystem supports the ICKDSF media maintenance utility. The ICKDSF utility can also be used to perform service functions, error detection, and media maintenance. Since the 9900 is a RAID device, there are only a few differences in operation from conventional DASD or other RAID devices. Table 4.6 lists ICKDSF commands that are specific to the 9900, as contrasted to RAMAC.

Table 4.6 ICKDSF Commands for 9900 Contrasted to RAMAC

Command	Argument	Subsystem	Return Code
INSPECT	KEEPIT	RAMAC	CC = 12 Invalid parameter(s) for device type.
		9900	CC = 12, F/M = 04 (EC=66BB).
	PRESERVE	RAMAC	CC = 4 Parameter ignored for device type.
		9900	CC = 12, F/M = 04 (EC=66BB) Unable to establish primary and alternate track association for track CCHH=xxxx.
	SKIP	RAMAC	CC = 4 Parameter ignored for device type - skip.
		9900	CC = 12, F/M = 04 (EC=66BB) Primary track CCHH-xxxx found unrecoverable.
	NOPRESERVE, NOSKIP,	RAMAC	CC = 0, ALT information not displayed.
	NOCHECK	9900	CC = 0
ALLTRACKS, ASSIGN, RECLAIM	RAMAC	CC = 12 Invalid parameter(s) for device type.	
	9900	In case of PRESERVE: CC = 12, In case of NO PRESERVE: CC = 0.	
INSTALL	SETMODE (3390)	RAMAC	CC = 0 (but not recommended by IBM).
		9900	CC = 0
	SETMODE (3380)	RAMAC	CC = 12, Invalid parameter(s) for device type.
		9900	CC = 12, Function not supported for nonsynchronous DASD.
ANALYZE		RAMAC	CC = 0
		9900	CC = 0
BUILDX		RAMAC	CC = 0
		9900	CC = 0
REVAL	REFRESH	RAMAC	CC = 12 Device not supported for the specified function.
		9900	CC = 12, F/M = 04 (EC=66BB) Error, not a data check. Processing terminated.
	DATA, NODATA	RAMAC	CC = 0, Data/Nodata parameter not allowed.
		9900	CC=0
CONTROL		RAMAC	CC = 0, ALT information not displayed.
		9900	CC = 0, ALT information not displayed.
INIT		RAMAC	CC = 0, ALT information not displayed.
		9900	CC = 0
REFORMAT		RAMAC	CC = 0, ALT information not displayed.
		9900	CC=0
CPVOLUME		RAMAC	CC = 0, Readcheck parameter not allowed.
		9900	CC=0
AIXVOL		RAMAC	Readcheck parameter not allowed.
		9900	CC = 0

4.3.3 MVS Cache Operations

To display the 9900 cache statistics under MVS DFSMS, use the following operator command: **D SMS, CACHE**. Figure 4.22 shows the cache statistics reported by the 9900. The 9900 reports cache statistics for each SSID in the subsystem. Because the dynamic cache management algorithm has been enhanced, the read and write percentages for the 9900 are displayed as **N/A**. For further information on MVS DFSMS cache reporting, please refer to the IBM® document *DFSMSdfp Storage Administrator Reference (SC28-4920)*.

SSID	DEVS	READ	WRITE	HIT RATIO	FW BYPASSES
0004	15	N/A	N/A	50%	0
0005	11	N/A	N/A	0%	0
0006	11	N/A	N/A	87%	0
0007	10	N/A	N/A	87%	0

 SSID=SUBSYSTEM IDENTIFIER
 DEVS=NUMBER OF MANAGED DEVICES ATTACHED TO SUBSYSTEM
 READ=PERCENT OF DATA ON MANAGED DEVICES ELIGIBLE FOR CACHING
 WRITE=PERCENT OF DATA ON MANAGED DEVICES ELIGIBLE FOR FAST WRITE

 HIT RATIO=PERCENT OF READS WITH CACHE HITS
 FW BYPASSES=NUMBER OF FAST WRITE BYPASSES DUE TO NVS OVERLOAD

Figure 4.22 Displaying Cache Statistics Using MVS DFSMS

The 9900 supports the following MVS cache operations:

- **IDCAMS LISTDATA COUNTS.** When the <subsystem> parameter is used with the LISTDATA command, the user must issue the command once for each SSID to view the entire 9900 image. Figure 4.23 shows a JCL example of the LISTDATA COUNTS command.

```
//LIST JOB. . . . .
//COUNT1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
LISTDATA COUNTS VOLUME (VOL000) UNIT (3390) SUBSYSTEM
LISTDATA COUNTS VOLUME (VOL064) UNIT (3390) SUBSYSTEM
LISTDATA COUNTS VOLUME (VOL128) UNIT (3390) SUBSYSTEM
LISTDATA COUNTS VOLUME (VOL192) UNIT (3390) SUBSYSTEM
/*
```

Figure 4.23 IDCAMS LISTDATA COUNTS (JCL example)

- **Subsystem counter reports.** The cache statistics reflect the logical caching status of the volumes. For the 9900, Hitachi Data Systems recommends that you set the nonvolatile storage (NVS) ON and the DASD fast write (DFW) ON for all logical volumes. This will not affect the way the 9900 caches data for the logical volumes. The default caching status for the 9900 is:

CACHE ON for the subsystem

CACHE ON for all logical volumes

CACHE FAST WRITE ON for the subsystem

NVS OFF for the subsystem

← Change NVS to **ON** for the 9900.

DFW OFF for all volumes

← Change DFW to **ON** for the 9900.

Note: In normal cache replacement, bypass cache, or inhibit cache loading mode, the 9900 performs a special function to determine whether the data access pattern from the host is sequential. If the access pattern is sequential, the 9900 transfers contiguous tracks from the disks to cache ahead of time to improve cache hit rate. Due to this advance track transfer, the 9900 shows the number of tracks transferred from the disks to the cache slot at DASD/CACHE of the SEQUENTIAL in TRANSFER OPERATIONS field in the subsystem counters report, even though the access mode is not sequential.

IDCAMS LISTDATA STATUS. The LISTDATA STATUS command generates status information for a specific device within the subsystem. The 9900 reports two storage sizes:

- **Subsystem storage.** This field shows capacity in bytes of cache. For a 9900 with more than one SSID, the cache is shared among the SSIDs instead of being logically divided. This strategy ensures backup battery power for all cache in the 9900. For the 9900, this field shows three-fourths (75%) of the total cache size.
- **Nonvolatile storage.** This field shows capacity in bytes of random access cache with a backup battery power source. For the 9900, this field shows one-fourth (25%) of the total cache size.

IDCAMS SETCACHE. The 9900 supports the IDCAMS SETCACHE commands, which manage caching for subsystem storage through the use of one command (except for REINITIALIZE). The following SETCACHE commands work for the subsystem storage across multiple SSIDs:

SETCACHE SUBSYSTEM ON|OFF

SETCACHE CACHEFASTWRITE ON|OFF

SETCACHE NVS ON|OFF

SETCACHE DESTAGE

Note: The SETCACHE REINITIALIZE command reinitializes only the logical subsystem specified by the SSID. You must issue the REINITIALIZE command once for each defined SSID.

DEVSERV PATHS. The DEVSERV PATHS command is defined as the number of LVIs that can be specified by an operator (from 1 through 99). To display an entire 9900 subsystem, enter the DEVSERV command for several LVIs, as follows:

DEVSERV PATHS,100,64

DEVSERV PATHS,140,64

DEVSERV PATHS,180,64

DEVSERV PATHS,1C0,64

4.3.4 VM/ESA® Cache Operations

When the 9900 is managed under VM/ESA®, the following SET CACHE commands are effective across multiple SSIDs:

```
SET CACHE SUBSYSTEM ON | OFF
SET NVS SUBSYSTEM ON | OFF
SET CACHEFW SUBSYSTEM ON | OFF
DESTAGE SUBSYSTEM
```

4.3.5 VSE/ESA Cache Operations

When using VSE/ESA to manage the 9900, the following CACHE commands are effective across multiple SSIDs:

```
CACHE SUBSYS=cuu,ON | OFF | STATUS
CACHE SUBSYS=cuu,FAST,ON | OFF
CACHE SUBSYS=cuu,NVS,ON | OFF
CACHE SUBSYS=cuu,REINIT
```

Note: SIMs indicating a drive failure may not be reported to the VSE/ESA console (reference IBM® document GA32-0253). Since the RAID technology and dynamic spare drives ensure non-stop processing, a drive failure may not be noticed by the console operator. If Hi-Track® is not installed, the user should run and read an EREP SIM report on a regular basis. Since all SIMs are also logged the 9900 Remote Console PC, the user can also use the Remote Console PC to monitor the SIMs.

4.4 Open-Systems Configuration

After physical installation of the 9900 subsystem has been completed, the user configures the 9900 subsystem for open-system operations with assistance as needed from the Hitachi Data Systems representative. For specific information and instructions on configuring the 9900 disk devices for open-system operations, please refer to the 9900 configuration guide for the connected platform. Table 4.7 lists the currently supported platforms and the 9900 configuration guides. Please contact your Hitachi Data Systems account team for the latest information on platform and software version support.

Table 4.7 9900 Open-System Platforms and Configuration Guides

Platform	Configuration Guide Document Number
UNIX®-based systems:	
IBM® AIX®	MK-90RD014
HP-UX®	MK-90RD016
Sun™ Solaris™	MK-90RD017
Compaq® Tru64 UNIX® (includes DIGITAL UNIX)	MK-90RD021
SGI™ IRIX®	MK-90RD024
Compaq® OpenVMS®	MK-91RD044
Sequent® DYNIX/ptx®	<i>Please use the IBM® documentation.</i>
PC server systems:	
Windows NT®	MK-90RD015
Windows® 2000	MK-90RD025
Novell® NetWare®	MK-90RD026
Linux®-based systems:	
Red Hat® Linux®	MK-90RD028

4.4.1 Configuring the Fibre-Channel Ports

The LUN Manager remote console software enables users to configure the fibre-channel ports for the connected operating system and operational environment (e.g., FC-AL or fabric). If desired, Hitachi Data Systems can configure the fibre-channel ports as a fee-based service. For further information on LUN Manager, see *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide (MK-91RD049)*, or contact your Hitachi Data Systems account team.

The 9960 subsystem supports a maximum of 32 fibre-channel ports, and the 9910 supports up to 24 fibre-channel ports. Each fibre-channel port is assigned a unique target ID (from 0 to EF). The 9900 subsystem supports up to 256 LUNs per port. Figure 4.24 illustrates fibre port-to-LUN addressing.




Figure 4.24 Fibre Port-to-LUN Addressing

4.4.2 Virtual LVI/LUN Devices

The Virtual LVI/LUN remote console software enables users to configure custom-size LUs which are smaller than standard-size LUs. Open-system users define Virtual LVI/LUN devices by size in MB (minimum device size = 35 MB). S/390® mainframe user define Virtual LVI/LUN devices by number of cylinders.

4.4.3 LU Size Expansion (LUSE) Devices

The LUSE function (included in the LUN Manager remote console software) enables users to configure size-expanded LUs which are from 2 to 36 times larger than standard-size LUs. LUSE devices are identified by the type and number of LDEVs which have been joined to form the single LUSE device. For example, an OPEN-9*36 LUSE device is composed of 36 OPEN-9 LDEVs.

4.5 Open Systems Operations

4.5.1 Command Tag Queuing

The 9900 supports command tag queuing for open-system devices. Command tag queuing enables hosts to issue multiple disk commands to the fibre-channel adapter without having to serialize the operations. Instead of processing and acknowledging each disk I/O sequentially as presented by the applications, the 9900 subsystem processes requests in the most efficient order to minimize head seek operations and disk rotational delay.

Note: The queue depth parameter may need to be adjusted for the 9900 devices. Please refer to the appropriate 9900 configuration guide for queue depth requirements and instructions on changing queue depth and other related system and device parameters (refer to Table 4.7 for a list of the 9900 open-system configuration guides).

4.5.2 Host/Application Failover Support

The 9900 supports many industry-standard products which provide host and/or application failover capabilities (e.g., HP® MC/ServiceGuard, VERITAS® First Watch®, HACMP, Microsoft® Cluster Server, Novell® Cluster Server, Sun Cluster, TruCluster, VMSCluster). For the latest information on failover and LVM product releases, availability, and compatibility, please contact your Hitachi Data Systems account team.

4.5.3 Path Failover Support

The user should plan for path failover (alternate pathing) to ensure the highest data availability. In the open-system environment, alternate pathing can be achieved by host failover and/or I/O path failover software. The 9900 provides up to 32 fibre ports to accommodate alternate pathing for host attachment. Figure 4.25 shows an example of alternate pathing. The LUs can be mapped for access from multiple ports and/or multiple target IDs. The number of connected hosts is limited only by the number of fibre-channel ports installed and the requirement for alternate pathing within each host. If possible, the alternate path(s) should be attached to different channel card(s) than the primary path.

The 9900 subsystem supports industry-standard I/O path failover products, including Hitachi Dynamic Link Manager™, Sequent Multi Path, and VERITAS® Volume Manager/DMP. For the latest information on failover product releases, availability, and compatibility, please contact your Hitachi Data Systems account team.




Figure 4.25 Alternate Pathing

4.5.4 Remote SIM (R-SIM) Reporting

The 9900 subsystem automatically reports all SIMs to the Remote Console PC (if powered on and booted). These R-SIMs contain the same information as the SIMs reported to mainframe hosts, enabling open-system users to monitor 9900 operations from the Remote Console PC. The 9900 remote console software allows the user to view the R-SIMs by date/time or by controller and to manage the R-SIM log file on the Remote Console PC.

4.5.5 SNMP Remote Subsystem Management

The 9900 subsystem supports the industry-standard simple network management protocol (SNMP) for remote subsystem management from the UNIX[®]/PC server host. SNMP is used to transport management information between the 9900 subsystem and the SNMP manager on the host. The SNMP agent for the 9900 subsystem sends status information to the host(s) when requested by the host or when a significant event occurs. Notification of 9900 error conditions is made in real time, providing UNIX[®]/PC server users with the same level of monitoring and support available to S/390[®] mainframe users. The SIM reporting via SNMP enables the user to monitor the 9900 subsystem without having to check the Remote Console PC for R-SIMs.

4.5.6 NAS and SAN Operations

The Hitachi Lightning 9900[™] subsystem supports NAS and SAN configurations. For further information on NAS operations, refer to the *Hitachi Freedom NAS[™] NSS Configuration Guide* (MK-91RD053), or contact your Hitachi Data Systems account team. For further information on SAN operations, please contact your Hitachi Data Systems account team.

Chapter 5 Planning for Installation and Operation

This chapter provides information for planning and preparing a site before and during installation of the Hitachi Lightning 9900™ subsystem. Please read this chapter carefully before beginning your installation planning.

If you would like to use any of the Lightning 9900™ features or software products (e.g., TrueCopy, ShadowImage, HRX, Hitachi Graph-Track), please contact your Hitachi Data Systems account team to obtain the appropriate license(s) and software license key(s).

Note: The general information in this chapter is provided to assist in installation planning and is not intended to be complete. The DKC410 and DKC415 (9960/9910) installation and maintenance documents used by Hitachi Data Systems personnel contain complete specifications. The exact electrical power interfaces and requirements for each site must be determined and verified to meet the applicable local regulations. For further information on site preparation for Lightning 9900™ subsystem installation, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center.

5.1 User Responsibilities

Before the 9900 subsystem arrives for installation, the user must provide the following items to ensure proper installation and configuration:

- Physical space necessary for proper subsystem function and maintenance activity
- Electrical input power
- Connectors and receptacles
- Air conditioning
- Floor ventilation areas (recommended but not required)
- Cable access holes
- RJ-11 analog phone line (for Hi-Track® support)

5.2 Electrical Specifications and Requirements for Three-Phase Subsystems

The Lightning 9960 subsystem supports three-phase power. At this time the 9910 subsystem supports only single-phase power.

5.2.1 Internal Cable Diagram

Figure 5.1 illustrates the internal cable layout of a three-phase 9960 subsystem.

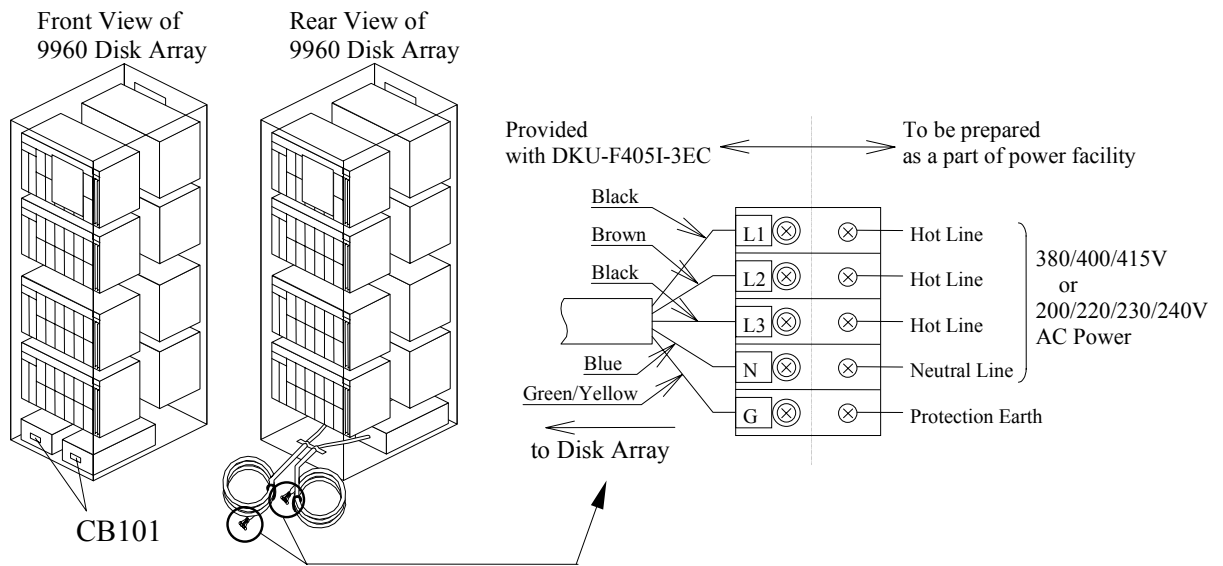


Figure 5.1 Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (Europe)

5.2.2 Power Plugs

Figure 5.2 illustrates the power plugs for a three-phase 9960 disk array unit (USA). Figure 5.3 illustrates the power plugs for a three-phase 9960 disk array unit (Europe).




Figure 5.2 Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (USA)

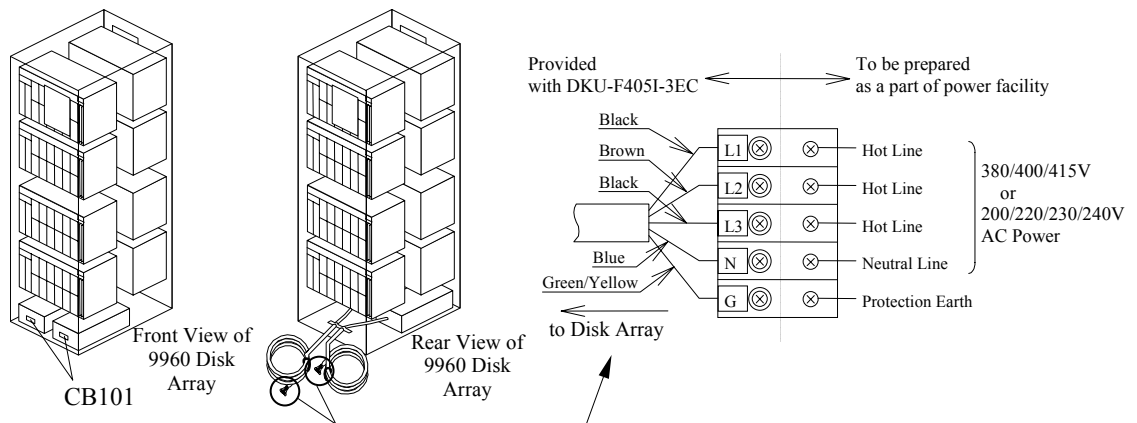


Figure 5.3 Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (Europe)

5.2.3 Features

Table 5.1 lists the features for a three-phase 9960 subsystem.

Table 5.1 9960 Three-Phase Features

Frame	Feature Number	Description	Comments
Controller	N/A	—	If using three-phase power, the Controller Frame receives its input power from the first disk array frame. All Controller single-phase features must be removed to allow three-phase power for the Controller.
Disk Array	DKU-F405I-3PS	AC Box Kit for Three-Phase	Consists of two AC Boxes
Disk Array	DKU-F405I-3EC	Power Cable Kit (Three-Phase Model for Europe)	
Disk Array	DKU-F405I-3UC	Power Cable Kit (Three-Phase Model for USA)	

5.2.4 Current Rating, Power Plug, Receptacle, and Connector for Three-Phase (60 Hz only)

Table 5.2 lists the current rating and power plug, receptacle, and connector requirements for three-phase 60-Hz 9960 subsystems. In a three-phase 9960 subsystem the controller frame (DKC) receives its AC input power from the first disk array frame (DKU) via internal cabling, so that subsystem will not require any customer outlets for the controller frame. The user must supply all power receptacles and connectors for both 60-Hz and 50-Hz subsystems. Russell & Stoll type (R&S) connectors are recommended for 60-Hz systems.

Note: Each disk array frame requires two power connections to ensure power redundancy. It is strongly recommended that the second power source be supplied from a separate power boundary to eliminate source power as a possible single (nonredundant) point of failure.

Table 5.2 Current Rating, Power Plug, Receptacle, and Connector for Three-Phase 9960

Item	9960 DKC	9960 DKU
Hitachi Base Unit	DKC410I-5	DKU405I-14
Circuit Current Rating	(from DKU)	30 A
Hitachi Feature(s) Required	none	DKU-F405I-3PS DKU-F405I-3UC
60-Hz Power Plug (or equiv.) Included with the product.	N/A	R&S 3760PDG/1100
Box-Type Receptacle (or equiv.) (not provided)	N/A	R&S 3754
Inline Connector (or equiv.) (not provided)	N/A	R&S 3934

***Note:** For information on power connection specifications for locations outside the U.S., contact the Hitachi Data Systems Support Center for the specific country.

5.2.5 Input Voltage Tolerances

Table 5.3 lists the input voltage tolerances for the three-phase 9960. Transient voltage conditions must not exceed +15-18% of nominal and must return to a steady-state tolerance within of +6 to -8% of the normal related voltage in 0.5 seconds or less. Line-to-line imbalance voltage must not exceed 2.5%. Nonoperating harmonic contents must not exceed 5%.

Table 5.3 Input Voltage Specifications for Three-Phase Power

Frequency	Input Voltages (AC)	Wiring	Tolerance (%)
60 Hz \pm 0.5 Hz	200V, 208V, or 230V	three-phase three wire + ground	+6% / -8%
50 Hz \pm 0.5 Hz	200V, 220V, 230V, or 240V	three-phase three wire + ground	+6% / -8%
50 Hz \pm 0.5 Hz	380V, 400V, or 415V	three-phase four wire + ground	+6% / -8%

Note: User input requires a 30-amp circuit breaker for three-phase power.

5.3 Electrical Specifications and Requirements for Single-Phase Subsystems

5.3.1 Internal Cable Diagram

Figure 5.4 and Figure 5.5 illustrate the internal cable layout of single-phase 9960 and 9910 subsystems, respectively.

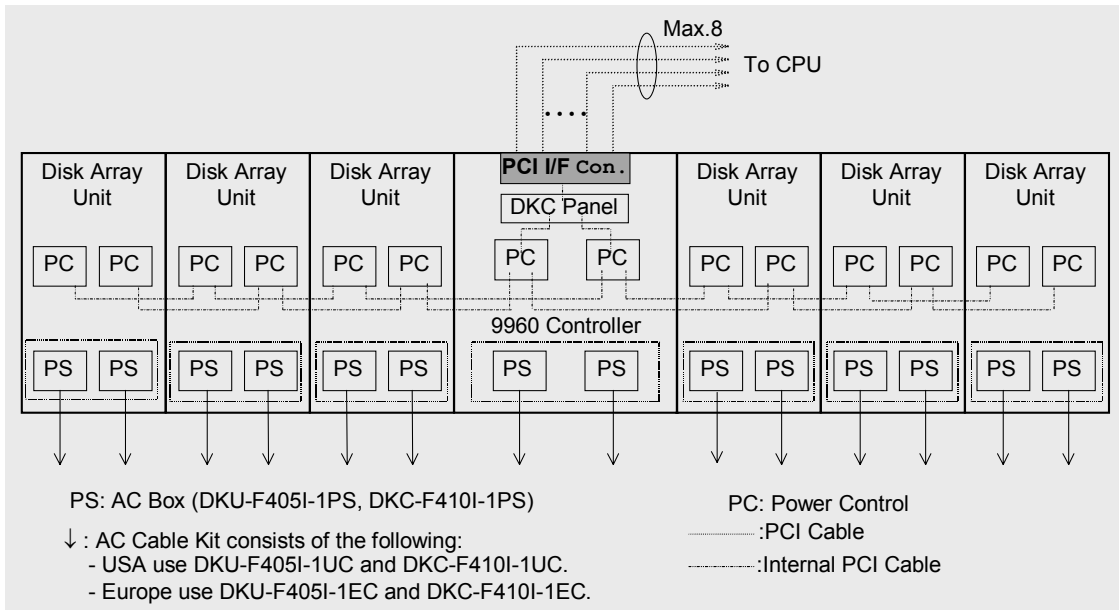


Figure 5.4 Internal Cable Diagram of a Single-Phase 9960 Subsystem

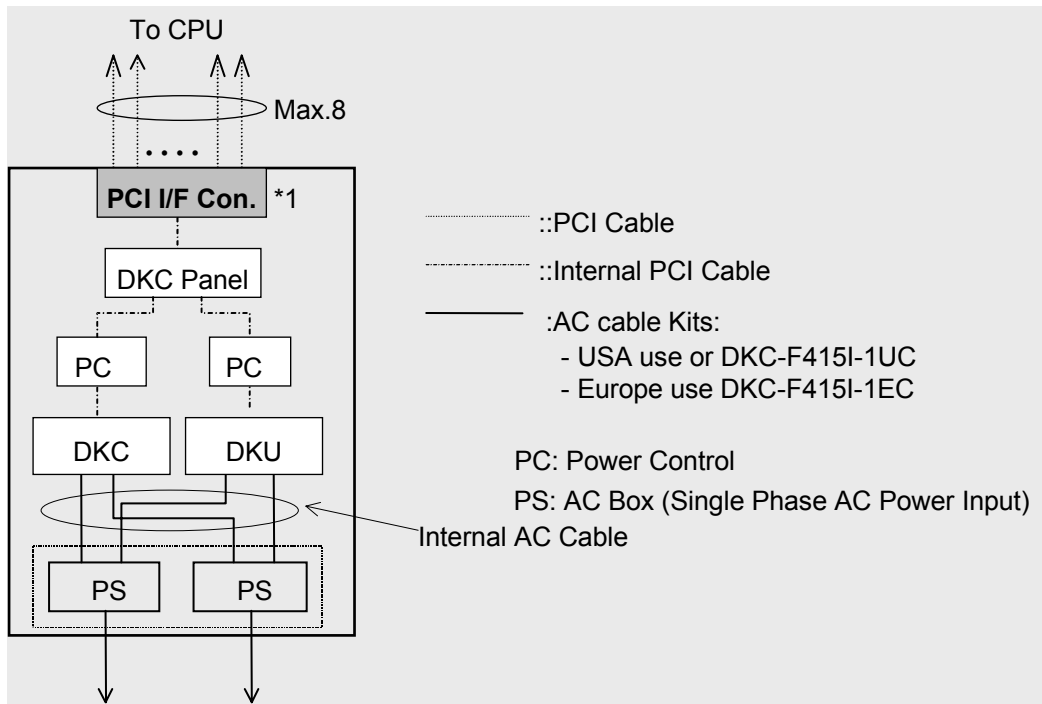


Figure 5.5 Internal Cable Diagram of a Single-Phase 9910 Subsystem

5.3.2 Power Plugs

Figure 5.6 through Figure 5.11 show the power plugs for the 9960 and 9910 subsystems.




Figure 5.6 Power Plugs for Single-Phase 9960 Controller (USA)




Figure 5.7 Power Plugs for Single-Phase 9910 Subsystem (USA)

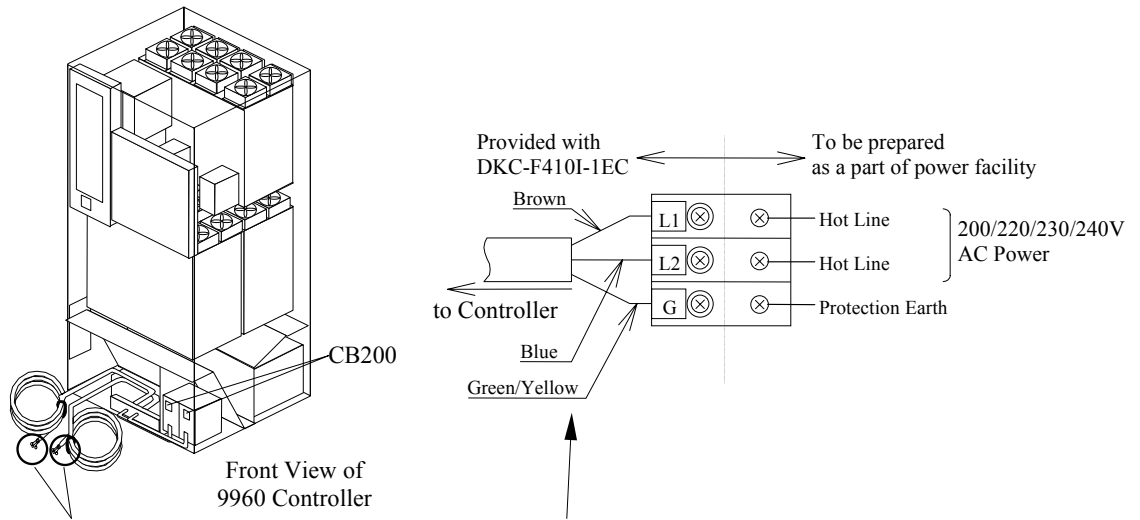


Figure 5.8 Power Plugs for a Single-Phase 9960 Controller (Europe)

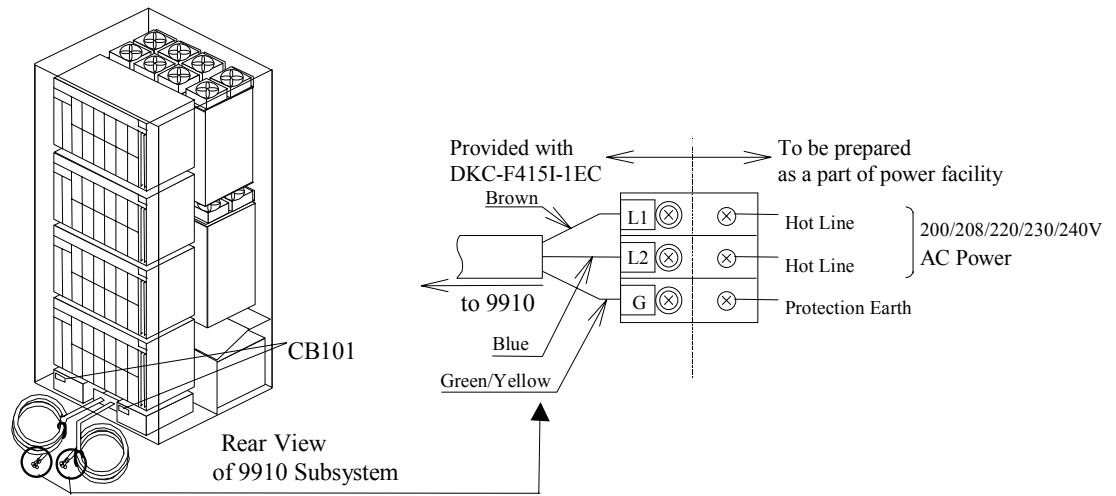


Figure 5.9 Power Plugs for a Single-Phase 9910 Subsystem (Europe)




Figure 5.10 Power Plugs for Single-Phase 9960 Disk Array Unit (USA)

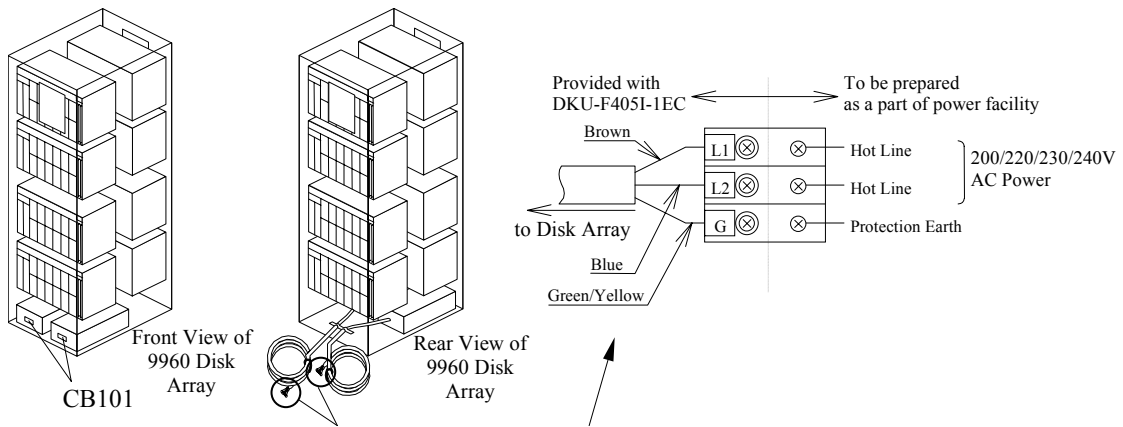


Figure 5.11 Power Plugs for Single-Phase 9960 Disk Array Unit (Europe)

5.3.3 Features

Table 5.4 lists the features for single-phase 9960 and 9910 subsystems.

Table 5.4 9900 Single-Phase Features

Model	Frame	Feature Number	Description
9960	Controller	DKC-F410I-PS	AC Box Kit for Single Phase
		DKC-F410I-EC	Power Cable Kit (Single-Phase Model for Europe)
		DKC-F410I-UC	Power Cable Kit (Single-Phase Model for USA)
	Disk Array	DKU-F405I-1PS	AC Box Kit for Single Phase
		DKU-F405I-1EC	Power Cable Kit (Single-Phase Model for Europe)
		DKU-F405I-1UC	Power Cable Kit (Single-Phase Model for USA)
9910	The entire subsystem consists of a single frame	DKC-F415I-1EC	Power Cable Kit (Single-Phase Model for Europe)
		DKC-F415I-1UC	Power Cable Kit (Single-Phase Model for USA)

5.3.4 Current Rating, Power Plug, Receptacle, and Connector for Single-Phase (60 Hz only)

Table 5.5 lists the current rating and power cable plug, receptacle, and connector requirements (part number or equivalent) for single-phase 60-Hz subsystems. Each disk array frame and controller frame has 2 input power cables with R & S FS 3720 plugs. The customer must supply the outlets for the plugs.

Note: Each disk array frame requires two power connections to ensure power redundancy. It is strongly recommended that the second power source be supplied from a separate power boundary to eliminate source power as a possible single (nonredundant) point of failure.

Table 5.5 Current Rating, Power Plug, Receptacle, and Connector for Single-Phase 9900

Item	9910	9960 DKC	9960 DKU
Hitachi Base Unit	DKC415I-5	DKC410I-5	DKU405I-18
Circuit Current Rating	20 A	20 A	20 A
Hitachi Feature(s) Required	DKC-F415I-1PS DKC-F415I-1UC	DKC-F410I-1PS DKC-F410I-1UC	DKU-F405I-1PS DKU-F405I-1UC
60-Hz Power Plug (or equiv.) Included with the product.	R&S 3720DP/0625	R&S 3720DP/0625	R&S 3720DP/0625
Box-Type Receptacle (or equiv.) (not provided)	R&S 3743	R&S 3743	R&S 3743
Inline Connector (or equiv.) (not provided)	R&S 3913	R&S 3913	R&S 3913

5.3.5 Input Voltage Tolerances

Table 5.6 lists the input voltage tolerances for the single-phase 9900 subsystem. Transient voltage conditions must not exceed +15-18% of nominal and must return to a steady-state tolerance of between +6 and -8% of the normal related voltage in 0.5 seconds or less. Line-to-line imbalance voltage must not exceed 2.5%. Nonoperating harmonic contents must not exceed 5%.

Table 5.6 Input Voltage Specifications for Single-Phase Power

Frequency	Input Voltages (AC)	Wiring	Tolerance(%)
60 Hz \pm 0.5 Hz	200V, 208V or 230V	Single-phase two wire + ground	+6% or -8%
50 Hz \pm 0.5 Hz	200V, 220V, 230V or 240V	Single-phase two wire + ground	+6% or -8%

Note: User input requires a 20-amp circuit breaker for single-phase power.

5.4 Dimensions and Weight

Figure 5.12 and Figure 5.13 illustrate the dimensions and weight of the 9960 and 9910 DKC and DKU. Table 5.7 lists the physical specifications, Table 5.8 lists the frame and component weights, and Table 5.9 lists the subsystem weights.




Figure 5.12 9960 DKC and DKU Physical Dimensions

Figure 5.13 9910 DKC and DKU Physical Dimensions

Table 5.7 9900 Physical Specifications

Item		9960 Controller	9960 Disk Unit *6		9910 Subsystem		
Dimension (mm/in)	Width	750 mm / 29.5 in *1	600 mm / 23.6 in *1		700 mm / 29.5 in *1		
	Depth	800 mm / 31.5 in					
	Height	1,790 mm / 70.5 in					
Floor Area with no service clearance		0.6 sq meters / 6.5 sq feet		0.48 sq meters / 5.16 sq feet		0.56 sq meters / 6 sq feet	
Floor Area with service clearance		1.8 sq meters / 19.4 sq feet		1.44 sq meters / 15.5 sq feet		1.47 sq meters / 15.8 sq feet	
Weight (kg/lb)		412 / 908 *2	480 / 1058 *5	480 / 1058 *3	440 / 970 *4	525 / 1157 *7	506 / 1116 *8
Heat Output (kW)		1.46 *2	3.43 *5	2.32 *3	2.85 *4	3.1 *7	2.81 *8
Power Consumption (kVA)		1.57 *2	3.72 *5	2.41 *3	3.14 *4	3.34 *7	3.06 *8
Air Flow (m ³ /min.)		18 *2 *3		12 *4 *5		15 *7 *8	

Notes:

*1. This includes the thickness of side covers (16 mm/0.6 in × 2).

*2. These values are used when 9960 has 8 GB cache memory, two fibre 8-port adapters, and an additional disk adapter.

*3. These values are used when 9960 Controller has maximum options.

*4. These values apply to a maximum size Disk Unit using 18-GB hard disk drives (HDDs).

*5. These values apply to a maximum size Disk Unit using 47-GB or 73-GB HDDs.

*6. The number of Disk Units per 9960 subsystem is 1-6.

*7. These values are used when 9910 has 16 GB cache, three fibre 8-port adapters, and fully populated with 47- or 73-GB HDDs.

*8. These values are used when 9910 has 16 GB cache memory, three fibre 8-port adapters, and fully populated with 18-GB HDDs.

Table 5.8 9960 Frame and Component Weights

Frame	Component	kg	lb
Controller	- First ACP pair - Cache boards with no memory	402	886
Controller	- Two ACP pairs - 8 GB cache - Two sets of fibre 8-port Adapters	412	908
Controller	- Four ACP pairs - 32 GB cache - Four sets of fibre 8-port Adapters - Additional power supplies	480	1058
Disk Array	Maximum configuration of 18-GB drives	440	970
Disk Array	Maximum configuration of 47-GB, 73-GB, and/or 180-GB disk drives	480	1058
Disk Array	RAID-5/RAID-1 array group (4 x 18-GB HDDs per group)	5.6	12.3
Disk Array	RAID-5/RAID-1 array group (4 x 47, 73, 146, or 180 GB HDDs per group)	7.2	15.9
Disk Array	Spare 18-GB hard drive	1.4	3.1
Disk Array	Spare 47, 73, 146, or 180 GB hard drive	1.8	4.0

Table 5.9 9960 & 9910 Subsystem Weights

Subsystem	Subsystem Configuration	kg*	lb*
9960	Controller with 32 GB cache, 4 sets of 8-port adapters and 1 Disk Array frame fully populated with 47, 73, 146, and/or 180-GB drives	960	2116
	Controller with 32 GB cache, 4 sets of 8-port adapters and 2 Disk Array frames fully populated with 47, 73, 146, and/or 180-GB drives	1440	3174
	Controller with 32 GB cache, 4 sets of 8-port adapters and 3 Disk Array frames fully populated with 47, 73, 146, and/or 180-GB drives	1920	4232
	Controller with 32 GB cache, 4 sets of 8-port adapters and 4 Disk Array frames fully populated with 47, 73, 146, and/or 180-GB drives	2400	5290
	Controller with 32 GB cache, 4 sets of 8-port adapters and 5 Disk Array frames fully populated with 47, 73, 146, and/or 180-GB drives	2880	6348
	Controller with 32 GB cache, 4 sets of 8-port adapters and 6 Disk Array frames fully populated with 47, 73, 146, and/or 180-GB drives	3360	7406
9910	Single frame with 16 GB cache, 3 sets of 8-port adapters, fully populated with 18-GB drives	525	1157
	Single frame with 16 GB cache, 3 sets of 8-port adapters, fully populated with 47-GB, 73-GB, and/or 180-GB drives	506	1116

***Note:** This is maximum weight. Other configurations (e.g. single phase, 4-port adapters) will have less weight.


5.5 Floor Loading and Cable Routing Requirements

5.5.1 Service Clearance Requirements

Figure 5.14 through Figure 5.17 specify the service clearance requirements (a + b) based on the floor load rating and the clearance (c). The following formula can be used to calculate floor loading to ensure that the weight of all equipment to be installed is adequately supported. Total area is defined as machine area plus half the service clearance.

$$\frac{\text{machine weight} + (15 \text{ lb/ft}^2 \times 0.5 \text{ service clearance}) + (10 \text{ lb/ft}^2 \times \text{total area})}{\text{total area}}$$


The additional weight of the raised floor and the weight of the cables is 10 lb/ft² (50 kg/m²) uniformly across the total area used in the calculations. When personnel and equipment traffic occur in the service clearance area, a distributed weight of 15 lb/ft² (75 kg/m²) is allowed. This distributed weight is applied over half of the service clearance area up to a maximum of 30 inches (760 mm) from the machine.



*1: Values in parentheses show allowable range of the floor cutout dimension. The floor cutout should be planned in the center of the DKC. In case that the floor cutout is planned in a right position for the external cable work and it is within the allowable range, the cutout position may be off-center. In this case, check the relation between the positions of the cutout and the opening on the bottom of the frame. If the floor cutout width is planned more than 520 mm, be careful about the restriction of the movable direction because there is a possibility that the caster wheels fall down into the cutout.

*2: These dimensions vary with the floor cutout dimension.

Figure 5.14 9960 Controller Frame Service Clearance and Cutouts (millimeters)



*1: Values in parentheses show allowable range of the floor cutout dimension. The floor cutout should be planned in the center of the DKC. In case that the floor cutout is planned in a right position for the external cable work and it is within the allowable range, the cutout position may be off-center. In this case, check the relation between the positions of the cutout and the opening on the bottom of the frame. If the floor cutout width is planned more than 520 mm, be careful about the restriction of the movable direction because there is a possibility that the caster wheels fall down into the cutout.

*2: These dimensions vary with the floor cutout dimension.

Figure 5.15 9960 Controller Frame Service Clearance and Cutouts (inches)

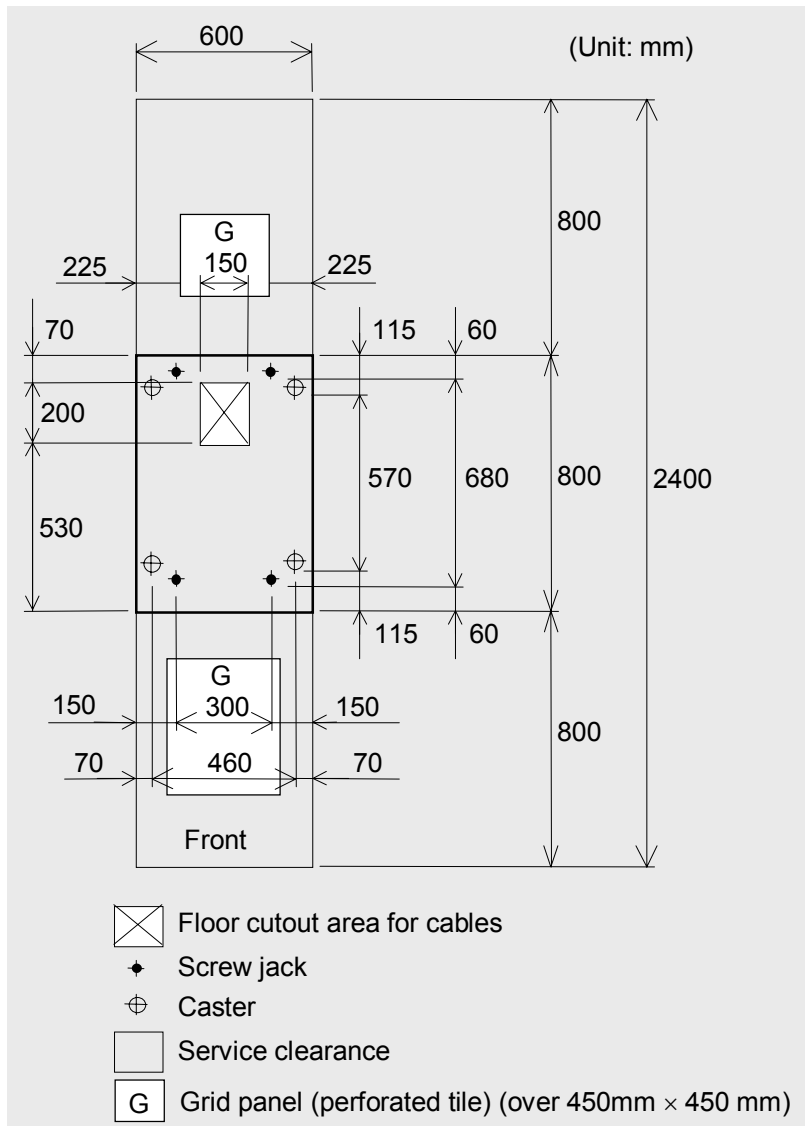


Figure 5.16 9960 Disk Array Unit Service Clearance and Cutout (millimeters)


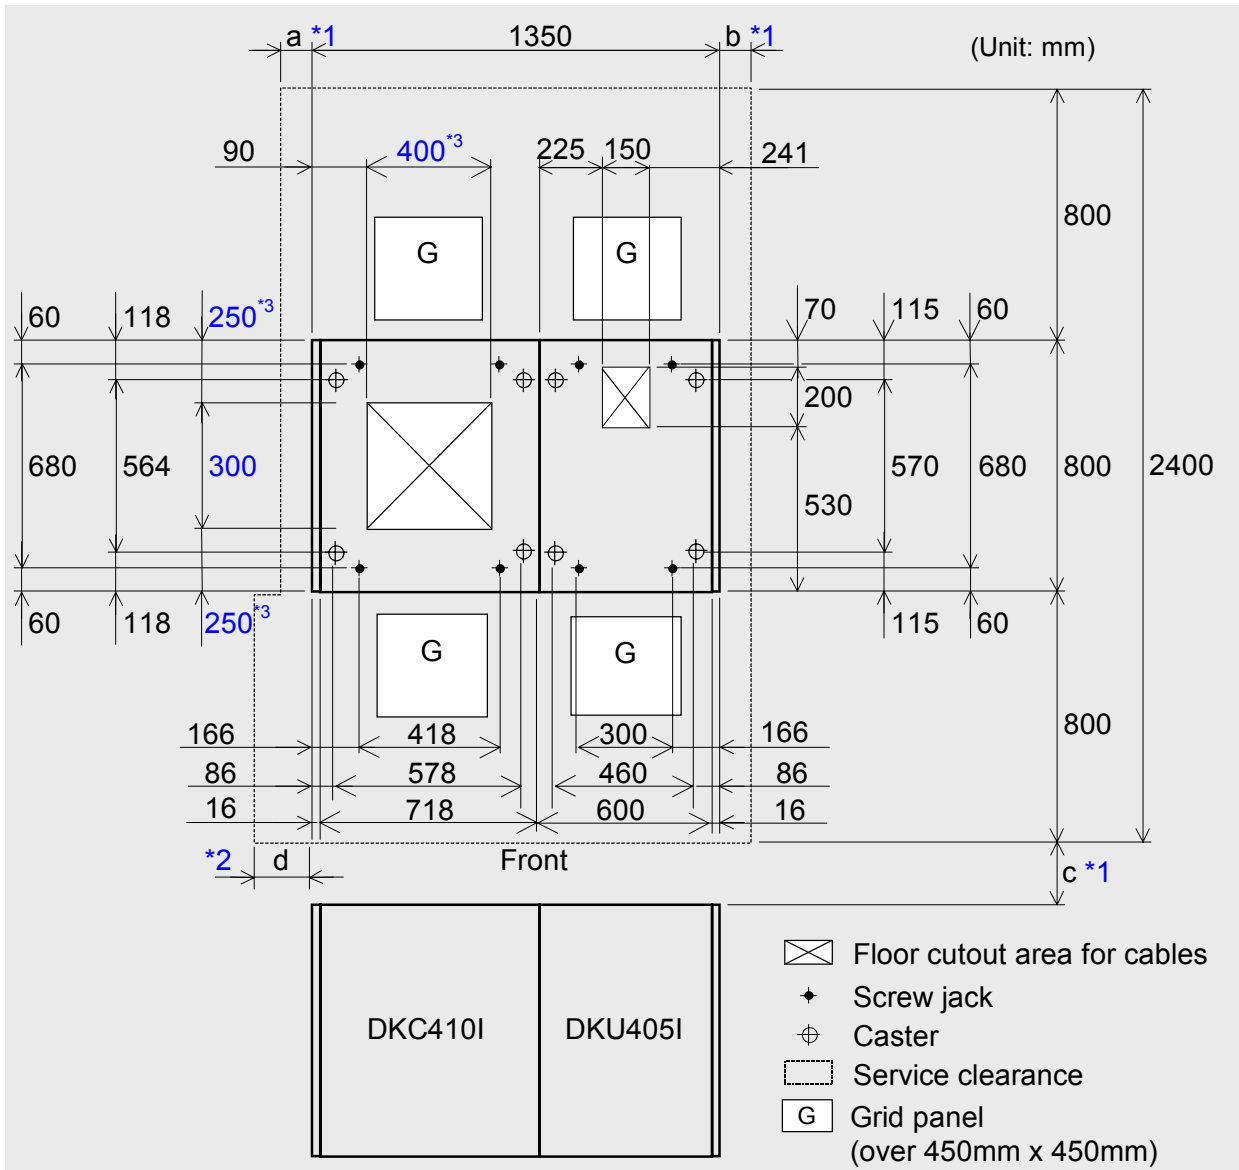


Figure 5.17 9960 Disk Array Unit Service Clearance and Cutout (inches)

5.5.2 Minimum Subsystem Disk Configuration

Figure 5.18 and Figure 5.19 illustrate the 9960 and 9910 subsystem minimum configuration. Figure 5.20 and Figure 5.21 illustrate all configurations of the 9960 and 9910 subsystems.

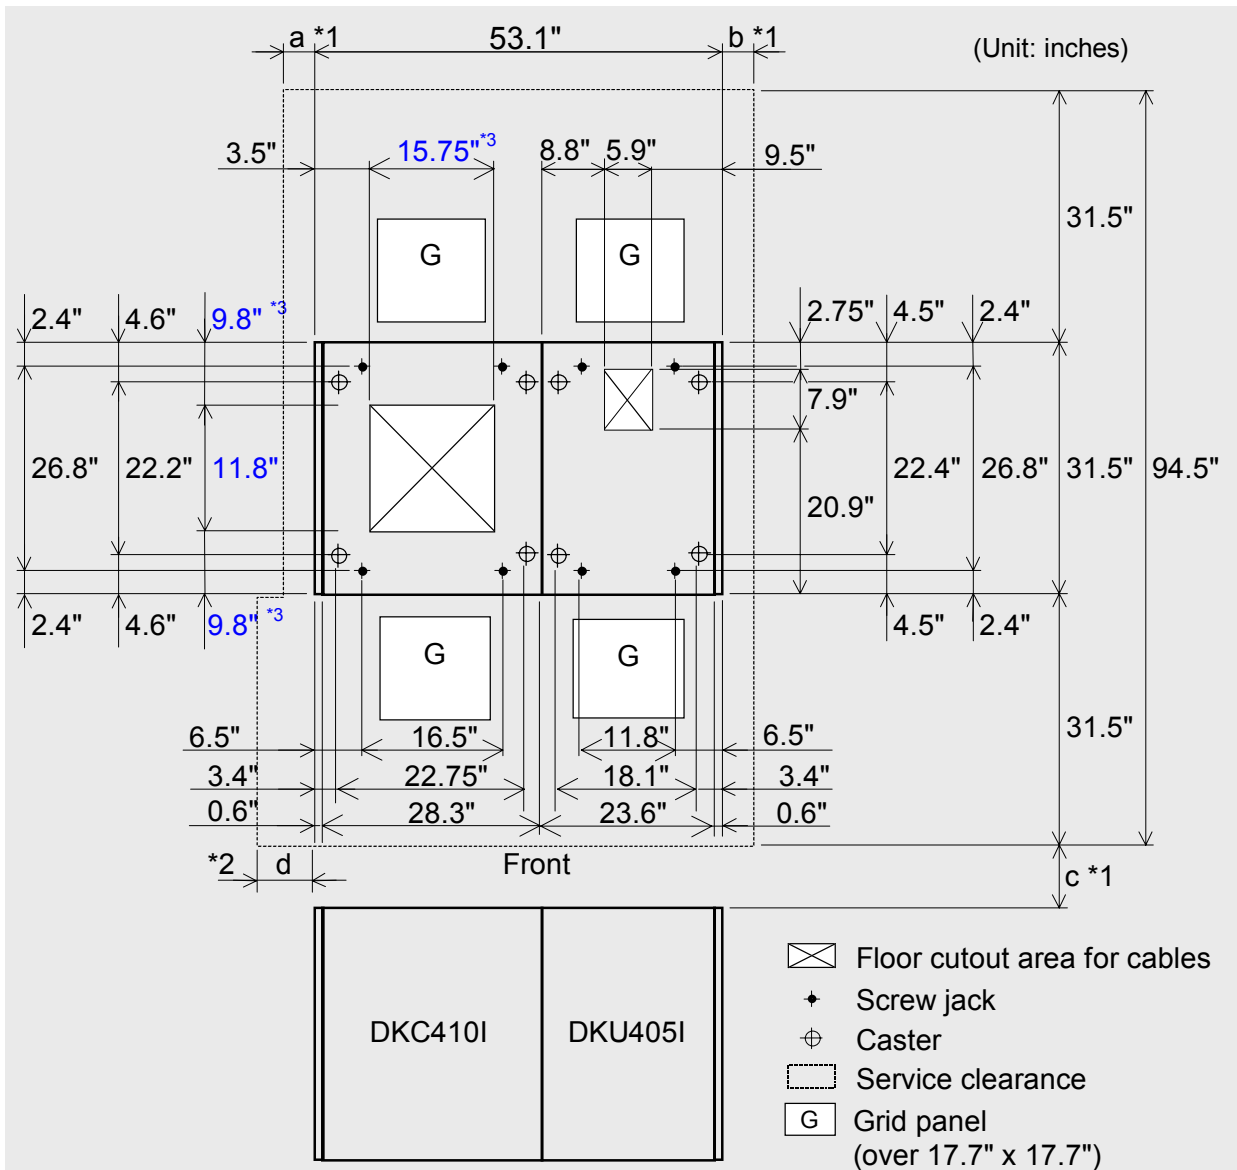


*1: Clearance (a+b) depends on the floor load rating and clearance c (see section 5.5.3).

*2: Clearance (d) is required over 0.28 m so as to open the DKC front door (refer to Figure 5.14). In case that clearance (d) is less than clearance (a), give priority to clearance (a).

*3: Refer to Figure 5.14 for details on the DKC floor cutout.

Figure 5.18 9960 Disk Subsystem Minimum Configuration (millimeters)



*1: Clearance (a+b) depends on the floor load rating and clearance c (see section 5.5.3).

*2: Clearance (d) is required over 0.28 m (11") so as to open the DKC front door (refer to Figure 5.14). In case that clearance (d) is less than clearance (a), give priority to clearance (a).

*3: Refer to Figure 5.14 for details on the DKC floor cutout.

Figure 5.19 9960 Disk Subsystem Minimum Configuration (inches)




Figure 5.20 9910 Disk Subsystem All Configurations (millimeters)




Figure 5.21 9910 Disk Subsystem All Configurations (inches)

5.5.3 Floor Load Rating

Figure 5.22, Figure 5.23, and Table 5.10 through Table 5.13 illustrate the floor loading ratings for various 9960 and 9910 subsystem configurations.

Table 5.10 Floor Load Rating for 9910 Subsystem

Floor Load Rating kg/m ² (lb/ft ²)	Required Clearance (a+b)				
	Clearance (C) in meters (feet)				
	C = 0	C = 0.2 (0.66)	C = 0.4 (1.31)	C = 0.6 (1.97)	C = 1.0 (3.28)
500 (102.4)	0.6 m (2 ft)	0.3 m (1 ft)	0.1 m (.3 ft)	0	0
450 (92.2)	0.8 m (2.6 ft)	0.5 m (1.6 ft)	0.3 m (1 ft)	0.1 m (.3 ft)	0
400 (81.9)	1.0 m (3.3 ft)	0.7 m (2.3 ft)	0.5 m (1.6 ft)	0.3 m (1 ft)	0
350 (71.7)	1.4 m (4.6 ft)	1.0 m (3.3 ft)	0.7 m (2.3 ft)	0.5 m (1.6 ft)	0.2 m (0.66 ft)
300 (61.4)	2.0 m (6.6 ft)	1.5 m (4.9 ft)	1.1 m (3.6)	0.9 m (3 ft)	0.5 m (1.6 ft)

Table 5.11 Floor Load Rating for 9960 Controller with 1 Disk Array

Floor Load Rating kg/m ² (lb/ft ²)	Required Clearance (a+b)				
	Clearance (C) in meters (feet)				
	C = 0	C = 0.2 (0.66)	C = 0.4 (1.31)	C = 0.6 (1.97)	C = 1.0 (3.28)
500 (102.4)	0	0	0	0	0
450 (92.2)	0.2 m (0.66 ft)	0	0	0	0
400 (81.9)	0.5 m (1.6 ft)	0.3 m (1 ft)	0.1 m (.3 ft)	0	0
350 (71.7)	0.9 m (3 ft)	0.7 m (2.3 ft)	0.5 m (1.6 ft)	0.3 m (1 ft)	0
300 (61.4)	1.6 m (5.2 ft)	1.3 m (4.3 ft)	1.0 m (3.3 ft)	0.8 m (2.6 ft)	0.4 m (1.3 ft)

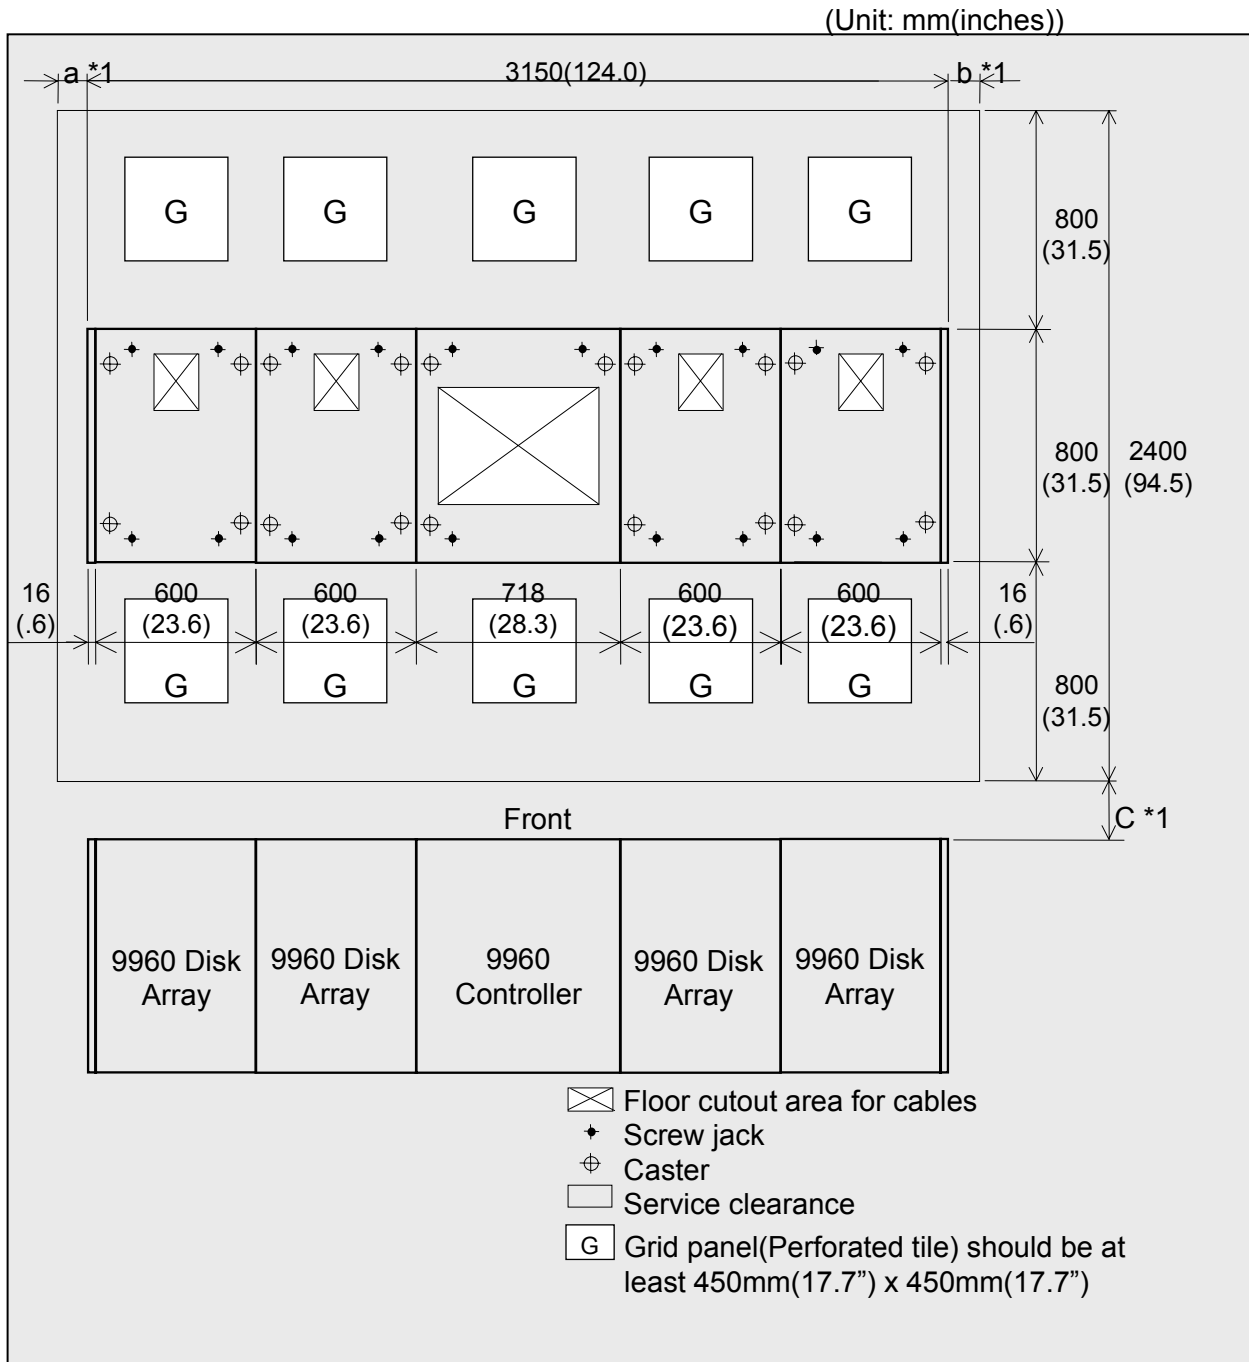


Figure 5.22 9960 Disk Subsystem with Controller and 4 Disk Arrays

Table 5.12 Floor Load Rating for 9960 Controller with 4 Disk Arrays

Floor Load Rating kg/m ² (lb/ft ²)	Required Clearance (a+b)				
	Clearance (C) in meters (feet)				
	C = 0	C = 0.2 (0.66)	C = 0.4 (1.31)	C = 0.6 (1.97)	C = 1.0 (3.28)
500 (102.4)	0.2 m (0.66 ft)	0	0	0	0
450 (92.2)	0.8 m (2.6 ft)	0.3 m (1 ft)	0	0	0
400 (81.9)	1.5 m (4.9 ft)	1.0 m (3.3 ft)	0.5 m (1.6 ft)	0.2 m (0.66 ft)	0
350 (71.7)	2.5 m (8.2 ft)	1.9 m (6.2 ft)	1.4 m (4.6 ft)	1.0 m (3.3 ft)	0.3 m (1 ft)
300 (61.4)	4.1 m (13.5 ft)	3.3 m (10.8 ft)	2.7 m (8.9 ft)	2.1 m (6.9 ft)	1.3 m (4.3 ft)

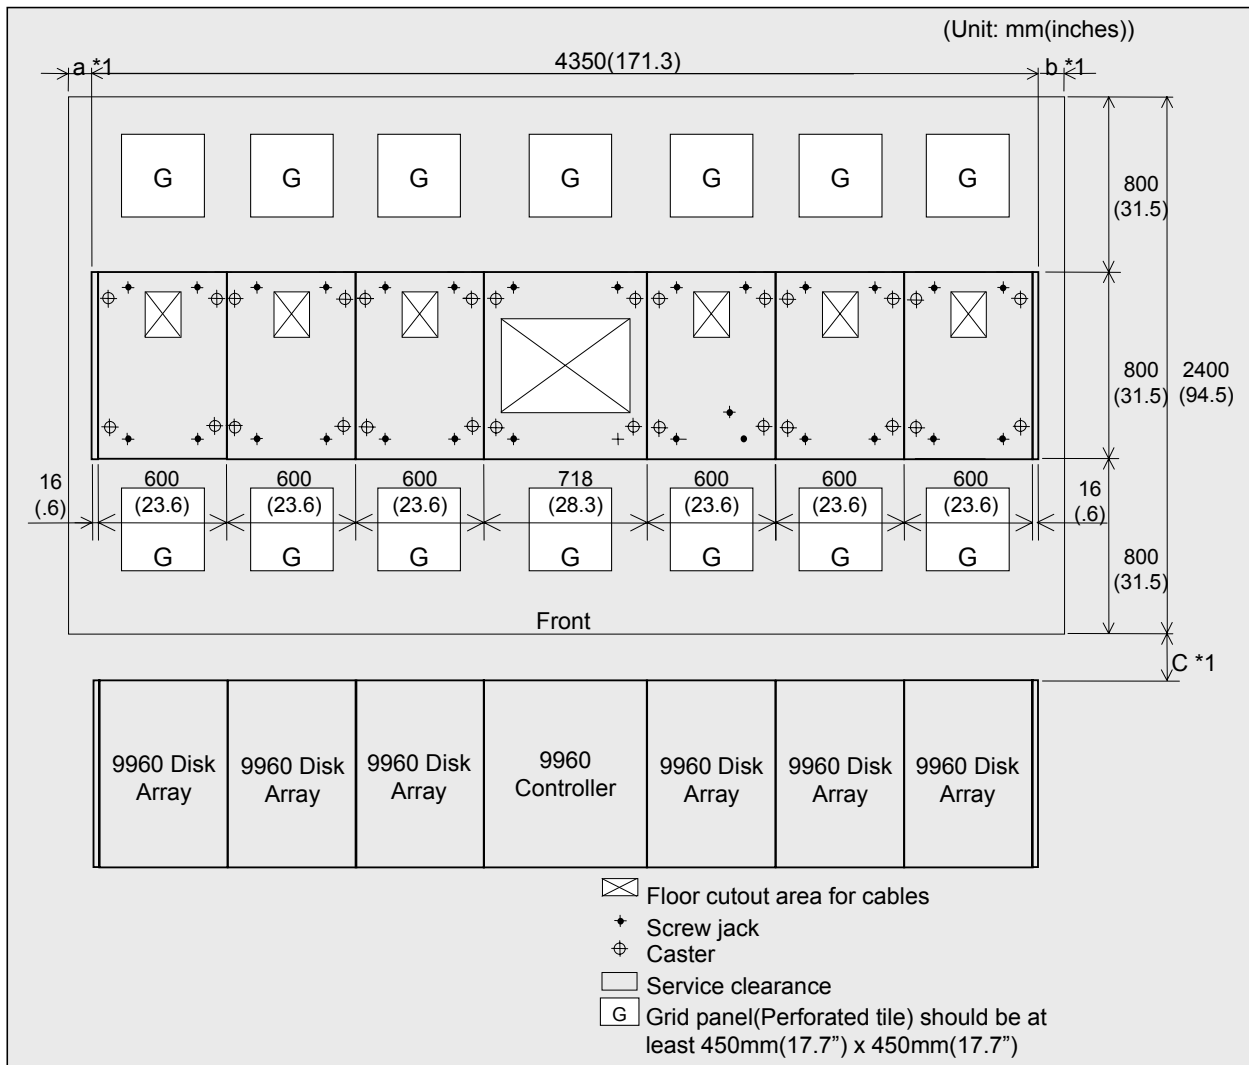


Figure 5.23 9960 Disk Subsystem with Maximum Configuration

Table 5.13 Floor Load Rating for 9960 Controller with Maximum Configuration

Floor Load Rating kg/m ² (lb/ft ²)	Required clearance (a+b)				
	Clearance (C) in meters (feet)				
	C = 0	C = 0.2 (0.66)	C = 0.4 (1.31)	C = 0.6 (1.97)	C = 1.0 (3.28)
500 (102.4)	0.2 m (0.66 ft)	0	0	0	0
450 (92.2)	0.9 m (3 ft)	0.3 m (1 ft)	0	0	0
400 (81.9)	1.8 m (5.9 ft)	1.1 m (3.6 ft)	0.6 m (2 ft)	0.2 m (0.66 ft)	0
350 (71.7)	3.2 m (10.5 ft)	2.4 m (7.9 ft)	1.7 m (5.6 ft)	1.1 m (3.6 ft)	0.3 m (1 ft)
300 (61.4)	5.4 m (17.7 ft)	4.3 m (14.1 ft)	3.4 m (11.2 ft)	2.7 m (8.9 ft)	1.6 m (5.2 ft)

5.5.4 Cable Requirements

Table 5.14 lists the cables required for the 9960 control frame. These cables must be ordered separately, and the quantity depends on the type and number of channels and ports. ExSA™ (ESCON®), FICON™, and fibre-channel cables are available from Hitachi Data Systems.

Table 5.14 Cable Requirements

Cable	Function
PCI cable	Connects 9900 to CPU power control interface.
FICON™ interface cables	Connects mainframe host systems, channel extenders, or FICON™ directors to 9900 ports. Single mode cables: <ul style="list-style-type: none"> ■ Yellow in color with SC- and LC-type connectors. ■ 8-10 micron. Most common is 9-micron singlemode. Multimode cables: <ul style="list-style-type: none"> ■ Orange cables with SC and LC-type connectors. ■ 50/125 micron and 62.5 micron multimode.
ExSA™ (ESCON®) interface cables	Connects mainframe host systems, channel extenders, or ESCON® directors to 9900 ports. Multimode cables: <ul style="list-style-type: none"> ■ Commonly called jumper cables. ■ Use LED light source. ■ Plug directly on CHA cards. ■ Orange cables with black duplex connectors. ■ Contain 2 fibers (transmit and receive). ■ 62.5 micron (up to 3 km per link). ■ 50 micron (up to 2 km per link). Mono/Single mode cables: <ul style="list-style-type: none"> ■ Required on XDF links between ESCDs or IBM® 9036 ESCON® remote control extenders. ■ Use laser light source. ■ Yellow in color with gray duplex connectors. ■ 8-10 micron. Most common is 9 micron.
Fibre cables	Connects open-systems to 9900 ports. Fibre cable types are 50 / 125 micron or 62.5 / 125 micron multimode. SC-type connector is required for 1 Gbps port. LC-type connector is required for 2 Gbps port.
Phone cable with RJ11 connector	Connects phone line to 9900 SVP for Hi-Track®.
10BaseT cable with RJ45 connector	Connects Remote Console PC to 9900 subsystem.
10Base2 cable (RG58) with BNC connector	Connects Remote Console PC to 9900 subsystem, and allows connection to multiple 9900s (up to eight) without using a hub. Requires a transceiver.

5.6 Channel Specifications and Requirements

Table 5.15 lists the ESCON[®] and FICON[™] port specifications for the 9900, and Table 5.16 lists the fibre-channel port specifications for the 9900. Each 9960 subsystem supports up to four channel adapters and 32 interface ports. Each 9910 subsystem supports up to three channel adapters and 24 interface ports. Each adapter consists of two cards.

Note: The 8GSE and 8HSE ports are applicable for Hitachi TrueCopy - S/390[®].

Note: Additional power supplies are needed if three or more channel adapters are installed.

Table 5.15 ESCON[®] and FICON[™] Port Information

Item	ESCON [®] Ports		FICON [™] Ports	
	ExSA [™] 4-Port	ExSA [™] 8-Port	4-Port SW	4-Port LW
Adapter model number	4S	8S	4MS	4ML
Max. data transfer rate (MB/s)	10/17	10/17	100 (1 Gbps)	100 (1 Gbps)
Number of ports/adapter	4	8	4	4
Number of MPs/adapter	4	8	4	4
Host interface	ExSA [™]	ExSA [™]	FICON [™]	FICON [™]
Number of adapter pairs per 9960 subsystem	1 to 4	1 to 4	1 to 4	1 to 4
Number of adapter pairs per 9910 subsystem	1 to 3	1 to 3	1 to 3	1 to 3
Max. interface cable length	43 km	43 km	500 m *1	10 km

*1. Indicates when 50 / 125 μ m multi-mode fiber cable is used. If 62.5 / 125 μ m multi-mode fiber cable is used, maximum length decreases to 300 m.

Table 5.16 Fibre-Channel Port Information

Item	Fibre-Channel Ports				
	4-Port SW	8-Port SW	4-Port LW	8-Port LW	2 Gbps 8-Port SW
Adapter model number	4GS	8GS/GSE	4GL	8GL	8HSE
Max. data transfer rate (MB/s)	100 (1 Gbps)	100 (1 Gbps)	100 (1 Gbps)	100 (1 Gbps)	100/200 (1/2 Gbps)
Number of ports/adapter	4	8	4	8	8
Number of MPs/adapter	4	8	4	8	8
Host interface	FC	FC	FC	FC	FC
Number of adapter pairs per 9960 subsystem	1 or 4	1 to 4	1 to 4	1 to 4	1 to 4
Number of adapter pairs per 9910 subsystem	1 or 3	1 to 3	1 to 3	1 to 3	1 to 3
Max. interface cable length	500 m *1	500 m *1	10 km	10 km	500/300 m *2

*1. Indicates when 50 / 125 μ m multi-mode fiber cable is used. If 62.5 / 125 μ m multi-mode fiber cable is used, maximum length decreases to 300 m.

*2. Indicates when 50 / 125 μ m multi-mode fiber cable is used. If 62.5 / 125 μ m multi-mode fiber cable is used, 500 m (100 MB/s) and 300 m (200 MB/s) decrease to 300 m and 150 m, respectively.

5.7 Environmental Specifications and Requirements

5.7.1 Temperature and Humidity Requirements

Table 5.17 lists the temperature and humidity requirements for the 9900 subsystem. The recommended operational room temperature is 70-75°F (21-24°C). The recommended operational relative humidity is 50% to 55%.

Table 5.17 Temperature and Humidity Requirements

Parameter	Operating ^{*1}		Non-Operating ^{*2}		Shipping & Storage ^{*3}	
	Low	High	Low	High	Low	High
Temperature °F (°C)	60 (16)	90 (32)	14 (-10)	109 (43)	5 (-25)	140 (60)
Relative Humidity (%) ^{*4}	20 - 80		8 - 90		5 - 95	
Max. Wet Bulb °F (°C)	79 (26)		81 (27)		84 (29)	
Temperature Deviation °F (°C) / hour	18 (10)		18 (10)		36 (20)	

Notes:

^{*1}. Environmental specification for operating condition should be satisfied before the disk subsystem is powered on. The maximum temperature of 90°F (32°C) should be strictly satisfied at the air inlet portion of the subsystem. The recommended temperature range is 70-75°F (21-24°C).

^{*2}. This includes both packing and unpacking conditions unless otherwise specified.

^{*3}. During shipping or storage, the product should be packed with factory packing.

^{*4}. No condensation in and around the drive should be observed under any conditions.

5.7.2 Power Consumption and Heat Output Specifications

Table 5.18 lists the power consumption and heat output parameters for the 9900 control frame and array frames. These data generally apply to both 60-Hz and 50-Hz subsystems. Table 5.19 lists the power consumption and heat output parameters for the various configurations of the 9900 subsystem.

Table 5.18 9910/9960 Component Power and Heat Output Specifications

Subsystem	Component	Power (kVA)	Heat (kW)
9960	Basic Controller frame with: - first ACP pair (2 cards total) - no cache or shared memory (2 empty cache cards)	0.95	0.89
	Controller frame with: - two ACP pairs (4 cards total) - 4-GB cache, shared memory - one 8-Port ExSA™ or Fibre adapter (2 cards total)	1.3	1.2
	Controller frame with all options: - four ACP pairs (8 cards total) - 32-GB cache and associated shared memory (4 cards total) - four 8-Port Serial/Fibre adapters (8 cards total) - additional power supplies for >2 channel/fibre adapters	2.41	2.32
	Empty array frame	0.704	0.561
	Disk Array frame with maximum amount of 18-GB hard drives	3.14	2.85
	Disk Array frame with maximum amount of 47-GB or 73-GB hard drives	3.72	3.43
	Disk Array frame with maximum amount of 180-GB hard drives	3.04	2.754
9910/9960	Each array group: RAID-5/RAID-1, 18-GB disk drives (Qty 4)	0.072	0.072
	RAID-5/RAID-1, 47-GB, 73-GB, or 146-GB disk drives (Qty 4)	0.096	0.096
	RAID-5/RAID-1, 180-GB disk drives (Qty 4)	0.068	0.068
	Spare disk drives: 18-GB	0.018	0.018
	47-GB, 73-GB, or 146-GB	0.024	0.024
	180-GB	0.017	0.017

Table 5.19 9900 Subsystem Power and Heat Output Specifications

Subsystem Type	Subsystem Configuration	Power (kVA)	Heat (kW)
9960	Controller with 32 GB cache, 4 sets of 8-port Adapters and 1 fully populated Disk Array frame using 47-GB / 73-GB drives	6.13	5.75
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 2 fully populated Disk Array frames using 47-GB / 73-GB drives	9.85	9.18
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 3 fully populated Disk Array frames using 47-GB / 73-GB drives	13.57	12.61
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 4 fully populated Disk Array frames using 47-GB / 73-GB drives	17.29	16.04
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 5 fully populated Disk Array frames using 47-GB / 73-GB drives	21.01	19.47
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 6 fully populated Disk Array frames using 47-GB / 73-GB drives	24.73	22.9
9910	16 GB cache, 3 sets of 8-port adapters, fully populated with 18-GB drives	3.06	2.81
	16 GB cache, 3 sets of 8-port adapters, fully populated with 47- or 73-GB drives	3.34	3.10

5.7.3 Loudness

The acoustic emission values [loudness in dB(A)] for a maximum 9960 subsystem configuration (one DKC, six DKUs) are:

Front/rear = 64 dB(A)

Both sides = 57 dB(A)

Note: These values were extrapolated from the values for one DKC and one DKU.

5.7.4 Air Flow Requirements

The 9900 subsystem is air cooled. Air must enter the subsystem through the air flow intakes at the sides and bottoms of the frames and must be exhausted out of the tops and sides, so it is very important that the air intakes and outlets remain clear. Hitachi Data Systems recommends that under-floor air cooling has a positive pressure and meets the specifications listed in Table 5.20. For subsystems located at elevations from 3,000 to 7,000 feet (900 to 2,100 meters) above sea level, decrease the maximum air temperature by two degrees for each 1,000 feet (300 meters) above 3,000 feet (900 meters).

Table 5.20 Internal Air Flow

Subsystem Type	Subsystem Configuration	Air Flow (ft ³ /min)	Air Flow (m ³ /min)
9960	Control Frame with 8 GB Cache Memory, two 8-port Fibre Adapter and additional Disk Adapter	636	18
	Control Frame with all options	636	18
	Array Frame, fully configured	424	12
9910	Subsystem fully configured	530	15

5.7.5 Vibration and Shock Tolerances

Table 5.21 lists the vibration and shock tolerance data for the 9900 subsystem. The 9900 can be subjected to vibration and shock up to these limits and still perform normally. The user should consider these requirements if installing the 9900 near large generators located on the floor above or below the 9900 subsystem. Generators or any other source of vibration, if not insulated or shock-mounted, can cause excessive vibration that may affect the subsystem.

Table 5.21 Vibration and Shock Tolerances

Condition	Operating	Nonoperating	Shipping or Storage
Vibration	0.01" (0.25 mm), 5~10 Hz 0.05 g, 10~300 Hz	0.1" (2.5 mm), 5~10 Hz 0.5 g, 10 to 70 Hz 0.002" (0.05 mm), 70~99 Hz 1.0 g, 99 to 300 Hz	0.5 g, 15 min. *1 At four most severe resonance frequencies between 5 to 200 Hz
Shock	—	8 g, 15 ms	Horizontal: Incline Impact 4 ft (1.22 m) *2 Vertical: Rotational Edge 0.3 ft (0.1 m) *3

*1: See ASTM D999-86, Standard Methods for Vibration Testing of Shipping Containers.

*2: See ASTM D880-86, Standard Methods of Incline Impact Test for Shipping Containers.

*3: See ASTM D775-80, Standard Methods for Drop Test for Loaded Boxes.

g = acceleration equivalent to gravity (9.8 m/sec²)

Chapter 6 Troubleshooting

6.1 Troubleshooting

The Hitachi Lightning 9900™ subsystem provides continuous data availability and is not expected to fail in any way that would prevent access to user data. The READY and ENABLE LEDs on the 9900 control panel must be **ON** when the subsystem is operating online.

Table 6.1 lists potential error conditions and provides recommended actions for resolving each condition. If you are unable to resolve an error condition, contact your Hitachi Data Systems representative, or call the Hitachi Data Systems Support Center for assistance (see section 6.3).

Table 6.1 Troubleshooting

Error Condition	Recommended Action
Error message displayed.	Determine the type of error as reported by the subsystem. If possible, remove the cause of the error. If you cannot correct the error condition, call the Hitachi Data Systems Support Center for assistance.
General power failure.	Turn off power to the subsystem (see section 4.2.4) to prevent damage caused by power surges when power is restored. After power has been restored and is no longer fluctuating, power on the 9900 (see section 4.2.4).
Fence message is displayed on the console.	Determine if there is a failed storage path. If so, toggle the RESTART switch, and retry the operation. If the fence message is displayed again, call the Hitachi Data Systems Support Center for assistance.
Pinned track.	<p>A “pinned track” occurs when a disk drive failure temporarily prevents data from being destaged from cache to disk. The next time the data on the pinned track is accessed by the host, the data is destaged to another disk location. The RAID technology and dynamic spare disk drives of the 9900 ensure full recovery from disk drive failures and non-stop access to all user data.</p> <p>Due to the 9900’s architecture, data can remain pinned in cache if it is not accessed by the host following the pinned track condition. Certain 9900 maintenance activities (e.g., microcode updates) require that all pinned tracks be cleared. If a Hitachi Data Systems service representative encounters pinned tracks during these activities, user assistance may be required to access the data (e.g., run DSF Inspect, read and write specific data) to cause it to be destaged to disk.</p>
READY LED does not go on, or there is no power supplied.	Call the Hitachi Data Systems Support Center for assistance. WARNING: Do not open the 9900 control frame or touch any of the controls.
ENABLE LEDs for a cluster do not go on.	Determine if channel I/O operations to that cluster are possible. If so, call the Hitachi Data Systems Support Center to have the LEDs checked. If not, disconnect the channel(s) (see section 4.2.5 for instructions on parallel channels), and/or call the Hitachi Data Systems Support Center.
Emergency (fire, earthquake, flood, etc.)	Pull the emergency power-off (EPO) switch. You must call the Hitachi Data Systems Support Center to have the EPO switch reset.
ALARM LED is on.	If there is an obvious temperature problem in the area, power down the subsystem, lower the room temperature to the specified operating range, and power on the subsystem. If the area temperature is not the obvious cause of the alarm, call the Hitachi Data Systems Support Center for assistance.

6.2 Service Information Messages (SIMs)

The 9900 subsystem generates service information messages (SIMs) to identify normal operations (e.g., TrueCopy pair status change) as well as service requirements and errors or failures. For information on SIMs, please contact the Hitachi Data Systems Support Center (see section 6.3). (The SIM list in this document was removed because it was not complete.)

SIMs can be generated by the CHIP and ACP microprocessors and by the SVP. All SIMs generated by the 9900 are stored on the SVP for use by Hitachi Data Systems personnel, logged in the SYS1.LOGREC dataset of the S/390® host system, logged on the Remote Console PC (the Remote Console PC must be running), and reported via SNMP to the open-system host. The Remote SIM (R-SIM) feature of the 9900 Remote Console PC enables users to view the SIMs reported by the attached 9900 subsystems on the Remote Console PC. Each time a SIM is generated, the amber Message LED on the 9900 control panel turns on and the Remote Console PC displays a warning message, even when the remote console software is not running. The Hi-Track® remote maintenance tool also reports all SIMs to the Hitachi Data Systems Support Center.

SIMs are classified according to severity: service, moderate, serious, or acute. The service and moderate SIMs (lowest severity) do not require immediate attention and are addressed during routine maintenance. The serious and acute SIMs (highest severity) are reported to the S/390® host(s) once every eight hours. **Note:** If a serious or acute-level SIM is reported, the user should call the Hitachi Data Systems Support Center immediately to ensure that the problem is being addressed (see section 6.3).

Figure 6.1 illustrates a typical 32-byte SIM from the 9900 subsystem. SIMs are displayed by reference code (RC) and severity. The six-digit RC, which is composed of bytes 22, 23, and 13, identifies the possible error and determines the severity. The SIM type, located in byte 28, indicates which component experienced the error.




Figure 6.1 Typical 9900 SIM Showing Reference Code and SIM Type

6.3 Calling the Hitachi Data Systems Support Center

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

- The 9900 Remote Console configuration information saved on diskette using the **FDCOPY** function (see the *Hitachi Lightning 9900™ Remote Console User's Guide*),
- The circumstances surrounding the error or failure,
- The exact content of any error messages displayed on the host system(s),
- The error code(s) displayed on the Remote Console PC, and
- The remote service information messages (R-SIMs) logged on the Remote Console PC and the reference codes and severity levels of the recent R-SIMs.

The worldwide Hitachi Data Systems Support Centers are:

- Hitachi Data Systems North America/Latin America
San Diego, California, USA
1-800-348-4357
- Hitachi Data Systems Europe
Buckinghamshire, United Kingdom
011-44-175-361-8000
- Hitachi Data Systems Asia Pacific
North Ryde, Australia
011-61-2-9325-3300

Appendix A Unit Conversions

Table A.1 provides unit conversions for the standard (U.S.) and metric measurement systems.

Table A.1 Unit Conversions for Standard (U.S.) and Metric Measures

From	Multiply By:	To Get:
British thermal units (BTU)	0.251996	Kilocalories (kcal)
British thermal units (BTU)	0.000293018	Kilowatts (kW)
Inches (in)	2.54000508	Centimeters (cm)
Feet (ft)	0.3048006096	Meters (m)
Square feet (ft ²)	0.09290341	Square meters (m ²)
Cubic feet per minute (ft ³ /min)	0.028317016	Cubic meters per minute (m ³ /min)
Pound (lb)	0.4535924277	Kilogram (kg)
Kilocalories (kcal)	3.96832	British thermal units (BTU)
Kilocalories (kcal)	1.16279×10^{-3}	Kilowatts (kW)
Kilowatts (kW)	3412.08	British thermal units (BTU)
Kilowatts (kW)	859.828	Kilocalories (kcal)
Millimeters (mm)	0.03937	Inches (in)
Centimeters (cm)	0.3937	Inches (in)
Meters (m)	39.369996	Inches (in)
Meters (m)	3.280833	Feet (ft)
Square meters (m ²)	10.76387	Square feet (ft ²)
Cubic meters per minute (m ³ /min)	35.314445	Cubic feet per minute (ft ³ /min)
Kilograms (kg)	2.2046	Pounds (lb)
Ton (refrigerated)	12,000	BTUs per hour (BTU/hr)

Acronyms and Abbreviations

ACP	array control processor
ASTM	American Society for Testing and Materials
BS	basic (power) supply
BSA	bus adapter
BTU	British thermal unit
°C	degrees centigrade/Celsius
ca	cache
CFW	cache fast write
CH	channel
CHP	channel processor or channel path
CHIP	client-host interface processor
CHPID	channel path identifier
CKD	count key data
CL	cluster
CPU	central processing unit
CSA	Canadian Standards Association
CSW	cache switch
CU	control unit
CVS	custom volume size
DASD	direct access storage device
dB(A)	decibel (A-weighted)
DFDSS	Data Facility Dataset Services
DFSMS	Data Facility System Managed Storage
DFW	DASD fast write
DKC	disk controller
DKU	disk unit
dr	drive
DSF	Device Support Facilities
DTDS+	Disaster Tolerant Disk System Plus
ECKD	Extended Count Key Data
EOF	end of field
EPO	emergency power-off
EREP	Error Reporting
ESA	Enterprise Systems Architecture
ESCON [®]	Enterprise System Connection (IBM [®] trademark for optical channels)
ExSA [™]	Extended Serial Adapter [™]
FAL	File Access Library (part of the HRX software)
FBA	fixed-block architecture
FC	fibre-channel
FC-AL	fibre-channel arbitrated loop
FCC	Federal Communications Commission
FCU	File Conversion Utility (part of the HRX software)
FDN	Freedom Data Networks

FDR	Fast Dump/Restore
FICON™	Fiber Connection (IBM® trademark for fiber connection technology)
F/M	format/message
FWD	fast wide differential
<i>g</i>	acceleration of gravity (9.8 m/s ²) (unit used for vibration and shock)
Gb	gigabit
GB	gigabyte
Gbps	gigabit per second
GLM	gigabyte link module
GUI	graphical user interface
HACMP	High Availability Cluster Multi-Processing
HBA	host bus adapter
HCD	hardware configuration definition
HCPF	Hitachi Concurrent Processing Facility
HMBR	Hitachi Multiplatform Backup/Restore
HPAV	Hitachi Parallel Access Volume
HRX	Hitachi RapidXchange
HSN	Hierarchical Star Network
HVR	Hitachi Volume Relocation
HXRC	Hitachi Extended Remote Copy
Hz	Hertz
ICKDSF	A DSF command used to perform media maintenance
IDCAMS	access method services (a component of Data Facility Product)
IML	initial microprogram load
in.	inch(es)
IO, I/O	input/output (operation or device)
IOCP	input/output configuration program
JCL	job control language
kB	kilobyte
kcal	kilocalorie
kg	kilogram
km	kilometer
kVA	kilovolt-ampere
kW	kilowatt
LAN	local area network
lb	pound
LD	logical device
LDEV	logical device
LED	light-emitting diode
LPAR	logical partition
LCP	link control processor, local control port
LRU	least recently used
LU	logical unit
LUN	logical unit number, logical unit

LVI	logical volume image (also called device emulation)
LW	long wavelength
m	meter
MB	megabyte
MIH	missing interrupt handler
mm	millimeter
MP	microprocessor
MPLF	Multi-Path Locking Facility
MR	magnetoresistive
ms, msec	millisecond
MVS	Multiple Virtual Storage (including MVS/370, MVS/ESA, MVS/XA)
NAS	network-attached storage
NVS	nonvolatile storage
OEM	original equipment manufacturer
OFC	open fibre control
ORM	online read margin
OS	operating system
P/DAS	PPRC/dynamic address switching (an IBM® S/390® host software function)
PC	personal computer system
PCI	power control interface
PPRC	Peer-to-Peer Remote Copy (an IBM® S/390® host software function)
PS	power supply
R&S	Russell&Stoll
RAB	RAID Advisory Board
RAID	redundant array of independent disks
RAM	random-access memory
RC	reference code
RISC	reduced instruction-set computer
R/W	read/write
S/390®	IBM® System/390® architecture
SAN	storage-area network
SCSI	small computer system interface
sec.	second
seq.	sequential
SGI®	Silicon Graphics, Inc.
SIM	service information message
SMS	System Managed Storage
SNMP	simple network management protocol
SSID	storage subsystem identification
SVP	service processor
SW	short wavelength
TB	terabyte
TC	Hitachi TrueCopy
TCA	Hitachi TrueCopy Asynchronous

TC390	Hitachi TrueCopy - S/390®
TC390A	Hitachi TrueCopy - S/390® Asynchronous
TID	target ID
TPF	Transaction Processing Facility
TSO	Time Sharing Option (an IBM® System/370 operating system option)
UCB	unit control block
UIM	unit information module
UL	Underwriters' Laboratories
µm	micron, micrometer
VA	volt-ampere
VAC	volts AC
VDE	Verband Deutscher Elektrotechniker
VM	Virtual Machine (an IBM® S/390® system control program)
VOLID	volume ID
volser	volume serial number
XRC	Extended Remote Copy (an IBM® S/390® host software function)
VSE	Virtual Storage Extension (an IBM® S/390® operating system)
VTOC	volume table of contents
W	watt
WLM	Workload Manager (an IBM® S/390® host software function)
XA	System/370 Extended Architecture
XDF	Extended Distance Feature (for ExSA channels)