



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

© 2002 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-92RD100-P	March 2002	Preliminary Release
MK-92RD100-P1	April 2002	Revision 1 of Preliminary Release
MK-92RD100-0	April 2002	Initial Release

Source Documents for this Revision

The following source documents were used to produce this user and reference guide:

- DKC465I Disk Subsystem Hardware Specifications, draft for rev 2, April 2002.

Changes in this Revision

- Added support for Cache Manager (section 3.6.19).

Referenced Documents

For a listing of the 9900V user documentation, please see section 3.6.

Preface

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

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 open-system platforms and/or S/390® (mainframe) operating systems supported by the Lightning 9900™ V Series 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 online at <http://www.hds.com>. For specific information on supported host systems and platforms for the Lightning 9900™ V Series subsystem, please refer to the Lightning 9900™ V Series user documentation for the platform, or contact the Hitachi Data Systems Support Center.

Note: Unless otherwise noted, the term “9900V” refers to the entire Hitachi Lightning 9900™ V Series subsystem family, including all models (e.g., 9980V, 9970V) and all configurations (e.g., all-open, all-mainframe, multiplatform).

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

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™ V Series Subsystem

1.1	Key Features of the Lightning 9900™ V Series Subsystem	1
1.1.1	Continuous Data Availability	2
1.1.2	Connectivity	3
1.1.3	S/390® Compatibility and Functionality.....	3
1.1.4	Open-Systems Compatibility and Functionality.....	4
1.1.5	Hitachi Freedom NAS™ and Hitachi Freedom SAN™	5
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.....	15
2.2.1	Storage Clusters.....	15
2.2.2	Nonvolatile Shared Memory	16
2.2.3	Nonvolatile Duplex Cache	16
2.2.4	Multiple Data and Control Paths	17
2.2.5	Redundant Power Supplies	17
2.2.6	Channel Adapters and CHiPs.....	18
2.2.7	Channels	19
2.2.8	Array Control Processors (ACPs).....	20
2.3	Array Frame.....	22
2.3.1	Disk Array Groups.....	24
2.3.2	Sequential Data Striping.....	25
2.4	Intermix Configurations.....	26
2.4.1	RAID-1 & RAID-5 Intermix.....	26
2.4.2	Hard Disk Drive Intermix.....	26
2.4.3	Device Emulation Intermix	27
2.5	Service Processor (SVP)	28
2.6	Remote Console	29

Chapter 3 Functional and Operational Characteristics

3.1	New Features and Capabilities of the Lightning 9900™ V Series	31
3.2	I/O Operations	31
3.3	Cache Management	32
3.3.1	Algorithms for Cache Control	32
3.3.2	Write Pending Rate.....	32
3.4	Control Unit (CU) Images, LVIs, and LUs	33
3.4.1	CU Images.....	33
3.4.2	Logical Volume Image (LVIs).....	33
3.4.3	Logical Unit (LU) Type	34
3.5	Open Systems Features and Functions	35
3.5.1	Failover and SNMP Support.....	35
3.5.2	Share-Everything Architecture	35

3.6	Data Management Functions.....	36
3.6.1	Hitachi TrueCopy	38
3.6.2	Hitachi TrueCopy - S/390®	38
3.6.3	Hitachi ShadowImage	39
3.6.4	Hitachi ShadowImage - S/390®	39
3.6.5	Command Control Interface (CCI)	40
3.6.6	Hitachi Extended Remote Copy (HXRC).....	40
3.6.7	Hitachi NanoCopy™.....	41
3.6.8	Data Migration	41
3.6.9	Hitachi RapidXchange (HRX).....	42
3.6.10	Hitachi Multiplatform Backup/Restore (HMBR).....	42
3.6.11	HARBOR® File-Level Backup/Restore.....	43
3.6.12	HARBOR® File Transfer	43
3.6.13	HiCommand™.....	43
3.6.14	LUN Manager	44
3.6.15	Hitachi SANTinel	44
3.6.16	LUN Expansion (LUSE)	45
3.6.17	Virtual LVI/LUN.....	45
3.6.18	FlashAccess.....	46
3.6.19	Priority Access	46
3.6.20	Hitachi Parallel Access Volume (HPAV)	46
3.6.21	Hitachi Dynamic Link Manager™ (HDLM).....	47
3.6.22	Hitachi Performance Monitor.....	47
3.6.23	Hitachi CruiseControl	47
3.6.24	Hitachi Graph-Track™	48

Chapter 4 Configuring and Using the 9900V Subsystem

4.1	S/390® Configuration	49
4.1.1	Subsystem IDs (SSIDs)	49
4.2	S/390® Hardware Definition	50
4.2.1	Hardware Definition Using IOCP (MVS, VM, or VSE).....	50
4.2.2	Hardware Definition Using HCD (MVS/ESA)	55
4.2.3	Defining the 9900V to VM/ESA® Systems	66
4.2.4	Defining the 9900V to TPF	66
4.3	S/390® Operations.....	67
4.3.1	Initializing the LVIs.....	67
4.3.2	Device Operations: ICKDSF	67
4.3.3	MVS Cache Operations.....	69
4.3.4	VM/ESA® Cache Operations	71
4.3.5	VSE/ESA Cache Operations	71
4.4	Open-Systems Configuration.....	72
4.4.1	Configuring the Fibre-Channel Ports	73
4.4.2	Virtual LVI/LUN Devices.....	74
4.4.3	LUN Expansion (LUSE) Devices	74
4.5	Open-Systems Operations	75
4.5.1	Command Tag Queuing.....	75
4.5.2	Host/Application Failover Support	75
4.5.3	Path Failover Support	76
4.5.4	SIM Reporting	77
4.5.5	SNMP Remote Subsystem Management	77

4.5.6	NAS and SAN Operations	77
4.6	Control Panel	78
4.6.1	Emergency Power-Off (EPO)	80

Chapter 5 Planning for Installation and Operation

5.1	User Responsibilities and Safety Precautions.....	83
5.1.1	Safety Precautions.....	83
5.2	Dimensions, Physical Specifications, and Weight.....	84
5.3	Service Clearance, Floor Cutout, and Floor Load Rating Requirements	87
5.4	Electrical Specifications and Requirements for Three-Phase Subsystems	93
5.4.1	Power Plugs for Three-Phase	93
5.4.2	Features for Three-Phase	96
5.4.3	Current Rating, Power Plugs, Receptacles, and Connectors for Three-Phase (60 Hz only)	97
5.4.4	Input Voltage Tolerances for Three-Phase	98
5.4.5	Cable Dimensions for 50-Hz Three-Phase Subsystems	98
5.5	Electrical Specifications and Requirements for Single-Phase Subsystems	99
5.5.1	Power Plugs for Single-Phase	99
5.5.2	Features for Single-Phase	103
5.5.3	Current Rating, Power Plugs, Receptacles, and Connectors for Single-Phase (60 Hz only)	103
5.5.4	Input Voltage Tolerances for Single-Phase	104
5.5.5	Cable Dimensions for 50-Hz Single-Phase Subsystems	104
5.6	Cable Requirements	105
5.6.1	Device Interface Cable.....	106
5.7	Channel Specifications and Requirements	107
5.8	Environmental Specifications and Requirements	108
5.8.1	Temperature and Humidity Requirements.....	108
5.8.2	Power Consumption and Heat Output Specifications	109
5.8.3	Loudness	112
5.8.4	Air Flow Requirements.....	112
5.8.5	Vibration and Shock Tolerances	113

Chapter 6 Troubleshooting

6.1	Troubleshooting	115
6.2	Calling the Hitachi Data Systems Support Center	116
6.3	Service Information Messages (SIMs).....	117

Appendix A	Acronyms and Abbreviations	119
-------------------	---	------------

Appendix B	Unit Conversions	123
-------------------	-------------------------------	------------

List of Figures

Figure 2.1	Lightning 9980V HiStar Network (HSN) Architecture	12
Figure 2.2	9980V Subsystem Frames	13
Figure 2.3	9970V Subsystem Frame	14
Figure 2.4	Conceptual 9980V ACP Array Domain	20
Figure 2.5	Sample RAID-1 Layout	24
Figure 2.6	Sample RAID-5 Layout (Data Plus Parity Stripe).....	25
Figure 2.7	Sample Hard Disk Drive Intermix	26
Figure 2.8	Sample Device Emulation Intermix.....	27
Figure 2.9	Example of Remote Console PC and SVP Configuration.....	29
Figure 3.1	Fibre-Channel Device Addressing	34
Figure 4.1	IOCP Definition for FICON™ Channels (direct connect, via FICON™ switch)	51
Figure 4.2	IOCP Definition for 1024 LVIs (9900V connected to host CPU(s) via ESCD)	52
Figure 4.3	IOCP Definition for 1024 LVIs (9900V directly connected to CPU).....	53
Figure 4.4	Master MENU (Step 1).....	57
Figure 4.5	Basic HCD Panel (Step 2).....	58
Figure 4.6	Define, Modify, or View Configuration Data (Step 3).....	58
Figure 4.7	Control Unit List Panel (Step 4).....	59
Figure 4.8	Add Control Unit Panel (Step 5)	59
Figure 4.9	Selecting the Operating System (Step 6)	60
Figure 4.10	Control Unit Chpid, CUADD, and Device Address Range Addressing (Step 7)	60
Figure 4.11	Select Processor / Control Unit Panel (Step 8).....	61
Figure 4.12	Control Unit List (Step 9)	61
Figure 4.13	I/O Device List Panel (Step 10)	62
Figure 4.14	Add Device Panel (Step 11).....	62
Figure 4.15	Device / Processor Definition Panel - Selecting the Processor ID (Step 12)	63
Figure 4.16	Define Device / Processor Panel (Step 13).....	63
Figure 4.17	Device / Processor Definition Panel (Step 14)	64
Figure 4.18	Define Device to Operating System Configuration (Step 15)	64
Figure 4.19	Define Device Parameters / Features Panel (Step 16)	65
Figure 4.20	Update Serial Number, Description and VOLSER Panel (Step 18).....	65
Figure 4.21	LVI Initialization for MVS: ICKDSF JCL	67
Figure 4.22	Displaying Cache Statistics Using MVS DFSMS.....	69
Figure 4.23	IDCAMS LISTDATA COUNTS (JCL example).....	69
Figure 4.24	Fibre Port-to-LUN Addressing	73
Figure 4.25	Alternate Pathing	76
Figure 4.26	9980V Control Panel	78
Figure 4.27	Emergency Power-Off (EPO).....	80
Figure 5.1	Physical Overview of 9980V Subsystem	81
Figure 5.2	Physical Overview of the 9970V Subsystem	82
Figure 5.3	9980V DKC and DKU Physical Dimensions.....	84
Figure 5.4	9970V Physical Dimensions	85
Figure 5.5	9980V DKC Service Clearance and Floor Cutout.....	88
Figure 5.6	9980V DKU Service Clearance and Floor Cutout.....	89
Figure 5.7	9980V Subsystem Service Clearance and Floor Cutouts - Min Configuration.....	90

Figure 5.8	9980V Subsystem Service Clearance and Floor Cutouts - Max Configuration ...	91
Figure 5.9	9970V Subsystem Service Clearance and Floor Cutout - All Configurations	92
Figure 5.10	Power Plugs for Three-Phase 9980V Disk Array Unit (Europe).....	93
Figure 5.11	Power Plugs for Three-Phase 9980V Disk Array Unit (USA).....	94
Figure 5.12	Power Plugs for Three-Phase 9970V Subsystem (Europe)	94
Figure 5.13	Power Plugs for Three-Phase 9970V Subsystem (USA)	95
Figure 5.14	Cable Dimensions for 50-Hz Three-Phase Subsystems	98
Figure 5.15	Power Plugs for Single-Phase 9980V Controller (Europe)	100
Figure 5.16	Power Plugs for Single-Phase 9980V Controller (USA)	100
Figure 5.17	Power Plugs for a Single-Phase 9980V Disk Array Unit (Europe)	101
Figure 5.18	Power Plugs for a Single-Phase 9980V Disk Array Unit (USA)	101
Figure 5.19	Power Plugs for a Single-Phase 9970V Subsystem (Europe).....	102
Figure 5.20	Power Plugs for a Single-Phase 9970V Subsystem (USA).....	102
Figure 5.21	Cable Dimensions for 50-Hz Single-Phase Subsystems	104
Figure 5.22	9980V Subsystem Layout and Device Interface Cable Options	106
Figure 6.1	Typical 9900V SIM Showing Reference Code and SIM Type.....	117

List of Tables

Table 1.1	Program Products and Service Offerings.....	6-7
Table 2.1	Channel Adapter Specifications	18
Table 2.2	ACP Specifications.....	21
Table 2.3	Disk Drive Specifications	23
Table 3.1	Device Numbers for Each CU.....	33
Table 3.2	Capacities of Standard LU Types.....	34
Table 3.3	Data Management Functions for Open-System Users	36
Table 3.4	Data Management Functions for S/390® Users.....	37
Table 4.1	SSID Requirements.....	49
Table 4.2	Correspondence between Physical Paths and Channel Interface IDs (Cl 1)	54
Table 4.3	Correspondence between Physical Paths and Channel Interface IDs (Cl 2)	54
Table 4.4	HCD Definition for 64 LVIs.....	55
Table 4.5	HCD Definition for 256 LVIs	55
Table 4.6	ICKDSF Commands for 9900V Contrasted to RAMAC.....	68
Table 4.7	9900V Open-System Platforms and Configuration Guides	72
Table 4.8	9980V Control Panel.....	79
Table 5.1	9980V Physical Specifications.....	86
Table 5.2	9970V Physical Specifications.....	86
Table 5.3	Floor Load Rating and Required Clearances for 9980V Min Configuration.....	90
Table 5.4	Floor Load Rating and Required Clearances for 9980V Max Configuration	91
Table 5.5	Floor Load Rating and Required Clearances for 9970V Subsystem.....	92
Table 5.6	9980V and 9970V Three-Phase Features	96
Table 5.7	Current Rating, Power Plug, Receptacle, Connector for 3-Phase 9900V	97

Table 5.8	Input Voltage Specifications for Three-Phase AC Input	98
Table 5.9	Cable Dimensions for 50-Hz Three-Phase Subsystems.....	98
Table 5.10	9900V Single-Phase Features	103
Table 5.11	Current Rating, Power Plug, Receptacle, Connector for 1-Phase 9900V	103
Table 5.12	Input Voltage Specifications for Single-Phase Power.....	104
Table 5.13	Cable Dimensions for 50-Hz Single-Phase Subsystems.....	104
Table 5.14	Cable Requirements	105
Table 5.15	Open-Systems Channel Specifications	107
Table 5.16	Mainframe Channel Specifications.....	107
Table 5.17	Temperature and Humidity Requirements	108
Table 5.18	9980V DKC Component Power and Heat Output Specifications.....	109
Table 5.19	9980V DKU Component Power and Heat Output Specifications.....	110
Table 5.20	9970V Component Power and Heat Output Specifications	111
Table 5.21	Internal Air Flow	112
Table 5.22	Vibration and Shock Tolerances	113
Table 6.1	Troubleshooting	115
Table B.1	Unit Conversions for Standard (U.S.) and Metric Measures.....	123

Chapter 1 Overview of the Lightning 9900™ V Series Subsystem

1.1 Key Features of the Lightning 9900™ V Series Subsystem

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

The Lightning 9900™ V Series 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™ V Series can be configured for all-open, all-mainframe, and multiplatform operations.

The Lightning 9900™ V Series subsystem provides many benefits and advantages as well as advanced new features for the user, including double or more scalability from the 9900 in both capacity and performance. The 9900V subsystem supports an intermix of fibre-channel, ESCON®, and FICON™ host attachment for heterogeneous environments, and will introduce an enhanced design which provides high performance, high reliability, and high scalability in network-attached storage (NAS) and storage-area network (SAN) environments. The 9900V's new Java™-based user interface makes accessing, configuring, managing, and using the 9900V subsystem easier than ever. The Hitachi HiCommand™ licensed software product also supports the new 9900V subsystem for maximum flexibility in configuration and management. The Lightning 9900™ V Series subsystem is designed to meet customers' evolving and increasing needs for data storage and retrieval in the 21st century.

- **Instant access to data around the clock:**
 - 100 percent data availability guarantee with no single point of failure.
 - Highly resilient multi-path fibre architecture.
 - Fully redundant, hot-swappable components and non-disruptive subsystem microcode updates.
 - Global dynamic hot sparing.
 - Duplexed write cache with battery backup.
 - Hi-Track® “call-home” maintenance system.
 - RAID-1 and/or RAID-5 array groups within the same subsystem.

- **Unmatched performance and capacity:**
 - Multiple point-to-point data and control paths.
 - Up to 10-GB/sec internal system bandwidth.
 - Fully addressable 64-GB data cache; separate 3-GB control cache.
 - Extremely fast and intelligent cache algorithms.
 - Non-disruptive expansion to over 74 TB raw capacity.
 - Simultaneous transfers from up to 32 separate hosts.
 - Up to 1024 high-throughput (10K or 15K rpm) fibre-channel, dual-active disk drives.
- **Extensive connectivity and resource sharing:**
 - Concurrent operation of UNIX®, Windows®, and S/390® host systems.
 - Fibre-channel, FICON™, and Extended Serial Adapter™ (ESCON®) server connections.
 - Fibre-channel switched, arbitrated loop, and point-to-point configurations.

1.1.1 Continuous Data Availability

The Hitachi Lightning 9900™ V Series subsystem is designed for nonstop operation and continuous access to all user data. To achieve nonstop customer operation, the 9900V 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™ V Series subsystem.

1.1.2 Connectivity

The Hitachi Lightning 9900™ V Series RAID subsystem supports concurrent attachment to S/390® mainframe hosts and open-system (UNIX®-based and/or PC-server) platforms. The 9900V subsystem can be configured with the following port types to support all-open, all-mainframe, and multiplatform configurations:

- **Fibre-channel.** When fibre-channel interfaces are used, the 9900V 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 9900V 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.
- **Extended Serial Adapter™ (ExSA™)** (compatible with ESCON® protocol). When ExSA™ channel interfaces are used, the 9900V can provide up to 32 logical control unit (CU) images and 8192 logical device (LDEV) addresses. Each physical ExSA™ channel interface supports up to 256 logical paths providing a maximum of 8192 logical paths per subsystem. ExSA™ connection provides transfer rates of up to 17 MB/sec.
- **FICON™.** When FICON™ channel interfaces are used, the 9900V subsystem can provide up to 32 logical control unit (CU) images and 8192 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). Please contact your Hitachi Data Systems account team for information on the availability date of FICON™ support.

1.1.3 S/390® Compatibility and Functionality

The 9900V subsystem supports 3990-6, 3990-6E, and 2105 controller emulations and can be configured with multiple concurrent logical volume image (LVI) formats, including 3390-3, 3390-3R, 3390-9, and larger. In addition to full System-Managed Storage (SMS) compatibility, the 9900V 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™ V Series subsystem supports multiple concurrent attachment to a variety of host operating systems (OS) and is compatible with most fibre-channel host bus adapters (HBAs). The number of logical units (LUs) that may be connected to each port is determined by the type of host platform being attached. The 9900V supports the following platforms at this time. Please contact Hitachi Data Systems for the latest information on platform, OS version, and HBA support.

- IBM® AIX® OS
- Sun® Solaris® OS
- HP-UX® OS
- Microsoft® Windows NT® OS
- Microsoft® Windows® 2000 OS

Note: Future support is planned for additional platforms. Please contact your Hitachi Data Systems account team for the latest information on platform and OS version support.

The 9900V 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 9900V subsystem operates with industry-standard middleware products providing application/host failover capability, I/O path failover support, and logical volume management. The 9900V subsystem also supports the industry-standard simple network management protocol (SNMP) for remote subsystem management from the open-system host.

The 9900V subsystem can be configured with multiple concurrent logical unit (LU) formats (e.g., OPEN-3, -9, -E, -L). The user can also configure custom-size volumes using the Virtual LVI/LUN and LUN Expansion (LUSE) features of the 9900V subsystem (see section 3.6).

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 using, protecting and leveraging their data across the enterprise. Using multiple technologies and solutions including SAN and 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 storage 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™ V Series 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™ V Series 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 Hitachi TrueCopy – S/390®	Enables the user to perform remote copy operations between 9900V (and 9900) subsystems in different locations. Hitachi TrueCopy provides synchronous and asynchronous copy modes for S/390® and open-system data.	3.6.1 3.6.2
Hitachi ShadowImage Hitachi ShadowImage – S/390®	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.	3.6.3 3.6.4
Command Control Interface (CCI)	Enables open-system users to perform TrueCopy and ShadowImage operations by issuing commands from the host to the 9900V subsystem. The CCI software supports scripting and provides failover and mutual hot standby functionality in cooperation with host failover products.	3.6.5
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.	3.6.6
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.	3.6.7
Data migration (service offering only)	Enables the rapid transfer of data from other disk subsystems onto the 9900V subsystem. Data migration operations can be performed while applications are online using the data which is being transferred.	3.6.8
Data Sharing and Backup/Restore:		
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.	3.6.9
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 9900V subsystem.	3.6.10
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 9900V subsystem.	3.6.11
HARBOR® File Transfer	Enables users to transfer large data files at ultra-high channel speeds in either direction between open systems and mainframe servers.	3.6.12

Table 1.1 Program Products and Service Offerings (continued)

Function	Description	See Section:
Resource Management:		
HiCommand™	Enables users to manage the 9900V subsystem and perform functions (e.g., LUN Manager, LUN Security) from virtually any location via the HiCommand™ Web Client, command line interface (CLI), and/or third-party application.	3.6.13
LUN Manager	Enables users to configure the 9900V fibre-channel ports for operational environments (e.g., arbitrated-loop (FC-AL) and fabric topologies, host failover support).	3.6.14
Hitachi SANtinel (LUN Security)	Allows users to restrict host access to data on the Lightning 9900™ V Series subsystem. Open-system users can restrict host access to LUs based on the host's World Wide Name (WWN).	3.6.15
LUN Expansion (LUSE)	Enables open-system users to create expanded LUs which can be up to 36 times larger than standard fixed-size LUs.	3.6.16
Virtual LVI (VLVI) Virtual LUN (VLUN)	Enables users to configure custom-size LVIs and LUs which are smaller than standard-size devices.	3.6.17
FlashAccess (Flash)	Enables users to store specific high-usage data directly in cache memory to provide virtually immediate data availability.	3.6.18
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	3.6.19
Priority Access	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.	3.6.20
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™ V Series subsystem. HPAV provides compatibility with the IBM® Workload Manager (WLM) host software function and supports both static and dynamic PAV functionality.	3.6.21
Hitachi Dynamic Link Manager™	Provides automatic load balancing, path failover, and recovery capabilities in the event of a path failure.	3.6.22
Storage Utilities:		
Hitachi Performance Monitor	Performs detailed monitoring of subsystem and volume activity.	3.6.23
Hitachi CruiseControl	Performs automatic relocation of volumes to optimize performance.	3.6.24
Hitachi Graph-Track™ (GT)	Provides detailed information on the I/O activity and hardware performance of the 9900V 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.	3.6.25

1.1.7 Subsystem Scalability

The architecture of the 9900V subsystem accommodates scalability to meet a wide range of capacity and performance requirements. The 9980V storage capacity can be increased from a minimum of 108 GB net or 144 GB raw (one four-drive RAID-5 parity group, 36-GB HDDs) to a maximum of 55 TB net or 74 TB raw (254 four-drive RAID-5 parity groups, 73-GB HDDs). The 9980V nonvolatile cache can be configured from 4 GB to 64 GB in increments of 4 GB. The 9970V cache can be configured from 2 GB to 32 GB in increments of 2 GB. All disk drive and cache upgrades can be performed without interrupting user access to data and with minimal impact on subsystem performance.

The 9900V subsystem can be configured with the desired number and type of front-end client-host interface processors (CHIPs). The CHIPs reside on the channel adapters, which are installed in pairs. Each channel adapter pair offers eight host connections. The 9980V can be configured with up to four channel adapter pairs to provide up to 32 paths to attached host processors. The 9970V subsystem supports up to three channel adapter 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 9980V 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 9970V can be configured with one or two pairs of ACPs, providing up to sixteen concurrent data transfers to and from the disk drives.

The 9970V subsystem can support up to a combined total of four channel adapter pair features and ACP pair features. Thus if three channel adapter features with a total of 24 host interfaces are configured, one ACP pair must be configured. If two channel adapter features with a total of 16 host interfaces are configured, one or two ACP pairs may be configured. Due to the very high bandwidth of the ACP pairs in the 9970V, it is anticipated that one ACP pair will be sufficient for most 9970V applications.

1.2 Reliability, Availability, and Serviceability

The Lightning 9900™ V Series subsystem is not expected to fail in any way that would interrupt user access to data. The 9900V 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 9900V subsystem include:

- **Full fault-tolerance.** The 9900V subsystem provides full fault-tolerance capability for all critical components. The subsystem is protected against disk drive error and failure by enhanced RAID technologies and dynamic scrubbing and sparing. The 9900V uses component and function redundancy to provide full fault-tolerance for all other subsystem components (microprocessors, control storage, power supplies, etc.). The 9900V 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 9900V can sustain the loss of multiple power supplies and still continue operation.
- **Uninterruptable power supply (UPS).** The 9900V supports connection to a UPS to provide additional battery backup capability (see section 5.5.5).
- **Dynamic scrubbing and sparing for disk drives.** The 9900V uses special diagnostic techniques and dynamic scrubbing to detect and correct disk errors. Dynamic sparing is invoked automatically if needed. The 9980V can be configured with up to sixteen spare disk drives, and any spare disk can back up any other disk of the same speed (RPMs) and the same or less capacity, even if the failed disk and spare disk are in different array domains (attached to different ACP pairs).
- **Dynamic duplex cache.** The 9900V cache is divided into two equal segments on separate power boundaries. The 9900V 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 9900V subsystem at all times, collects hardware status and error data, and transmits this data via modem to the Hitachi Data Systems Support Center. The Support Center analyzes the data and implements corrective action as needed. 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 9900V 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 subsystem microcode upgrades can be performed during normal operations using the SVP or the alternate path facilities of the host.

*With one exception: In a 9970V with two ACP pairs, the first 64 drives are connected to the first ACP pair, and the second 64 drives are connected to the second ACP pair. It is also possible to attach both sets of 64 disk drives to a single ACP pair. If you want to add a second ACP pair to an existing single-ACP-pair 9970V, a service interruption will be required to disconnect any disk drives in the second set of 64 drive locations from the first ACP pair, and to reconnect them instead to the newly installed second ACP pair.
- **Error Reporting.** The Lightning 9900™ V Series 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 9900V service processor (SVP), 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

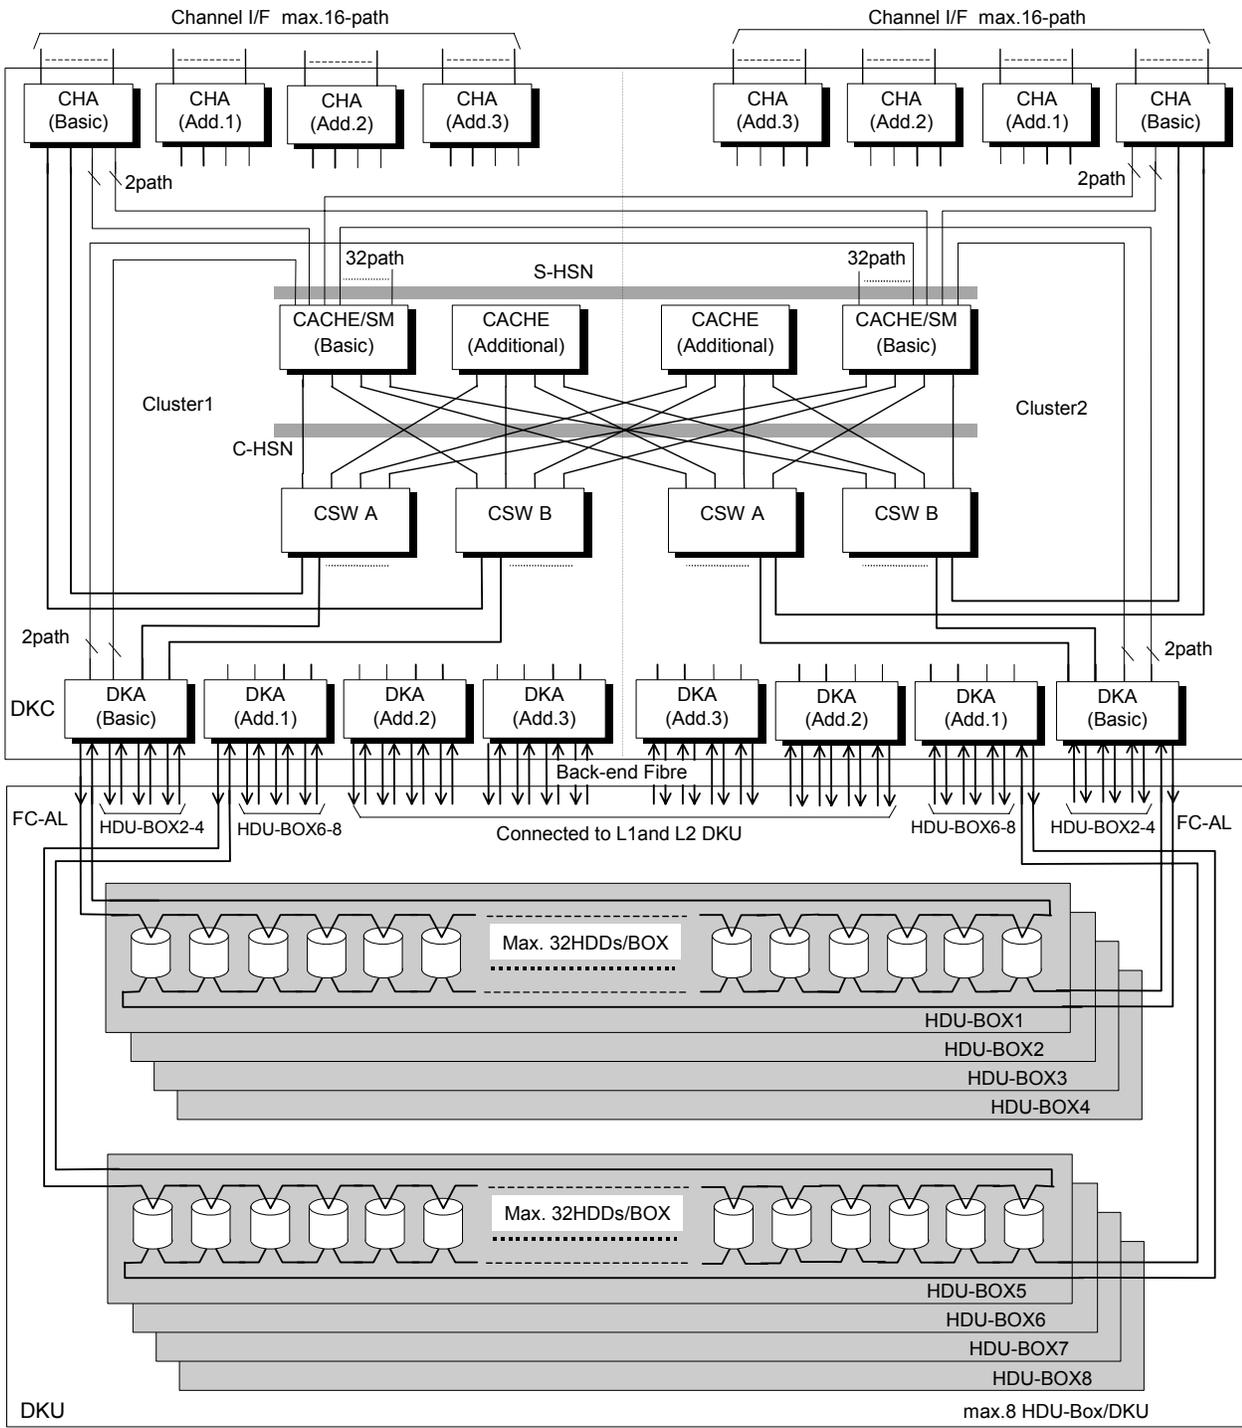
Figure 2.1 shows the Hierarchical Star Network (HiStar or HSN) architecture of the Lightning 9900™ V Series RAID subsystem. The “front end” of the 9900V subsystem includes the hardware and software that transfers the host data to and from cache memory, and the “back end” includes the hardware and software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

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

- Each channel adapter board contains four CHIPs, and the channel adapter boards are installed in pairs, for a total of eight host interfaces per feature. The 9980V subsystem supports up to four pairs of channel adapter boards (four sets of eight CHIPs) for a maximum of 32 host interfaces, and the 9970V subsystem supports up to three pairs of channel adapter boards to provide a maximum of 24 host interfaces.
- The 9980V controller contains four cache switch (CSW) cards, and the 9970V controller contains two CSW cards.
- Cache memory in the 9980V resides on four cache cards, and cache memory in the 9970V resides on two cards. Cache memory is backed up by battery. An additional battery must be configured on 9980V subsystems with over 32 GB of cache memory.
- Shared memory (minimum 2 GB) 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 9900V back end is controlled by the array control processors (ACPs) that reside on the disk adapter boards in the 9900V 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, and the DKAs are installed in pairs. The 9980V subsystem supports from one to four DKA pairs for a maximum of 32 ACPs. The 9970V subsystem supports one or two DKA pairs for a maximum of sixteen ACPs.



S-HSN: Shared Memory Hierarchical Star Net
 C-HSN: Cache Memory Hierarchical Star Net
 CSW: Cache Switch

Max.4 DKU/Subsystem
 In the maximum configuration of the subsystem,
 up to 64 HDDs can be connected through one FC-AL.

Figure 2.1 Lightning 9980V HiStar Network (HSN) Architecture

The 9980V subsystem (see Figure 2.2) includes the following major components:

- One controller frame containing the control and operational components of the subsystem.
- Up to four disk array frames containing the storage components (disk drive arrays) of the subsystem.
- The service processor (SVP) (see section 2.5). The 9900V SVP is located in the controller frame and can only be used by authorized Hitachi Data Systems personnel. The SVP provides the 9900V Remote Console functionality (see section 2.6).

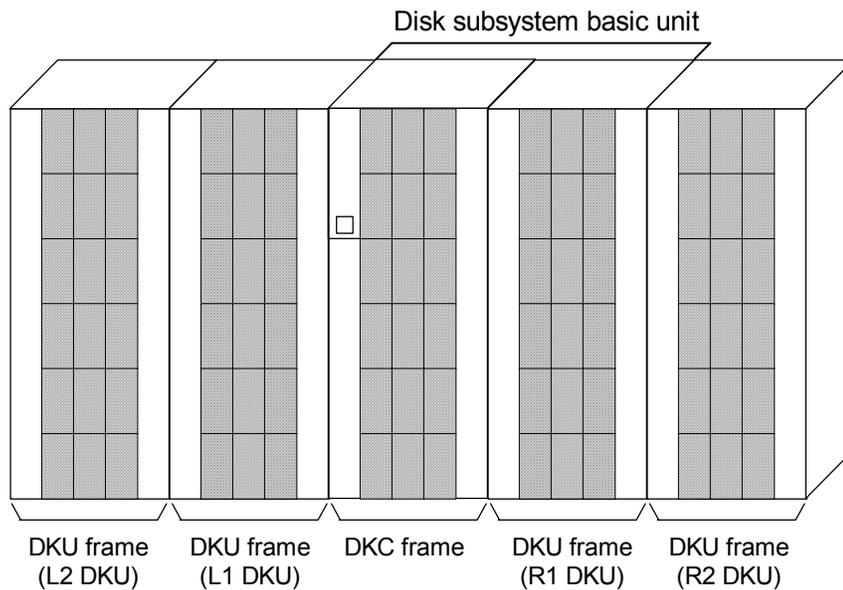


Figure 2.2 9980V Subsystem Frames

The 9970V 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 9900V SVP is located in the controller frame and can only be used by authorized Hitachi Data Systems personnel. The SVP provides the 9900V Remote Console functionality (see section 2.6).

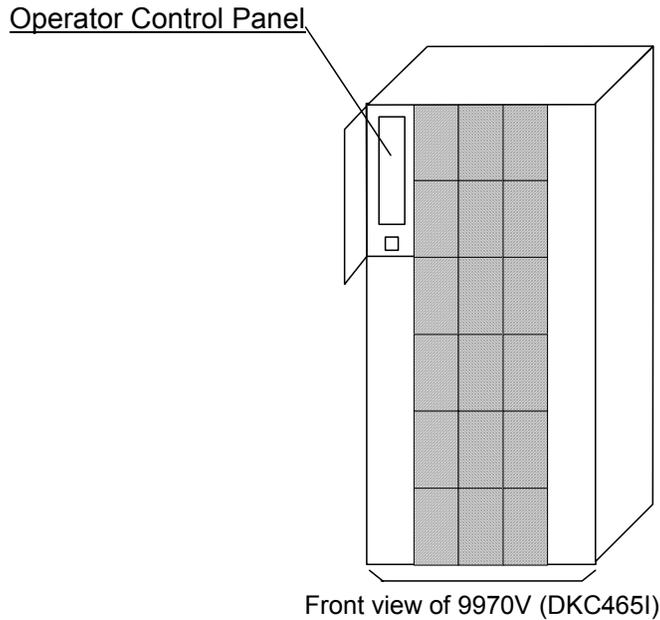


Figure 2.3 9970V Subsystem Frame

2.2 Components of the Controller Frame

The 9900V controller frame contains the control and operational components of the subsystem. For the 9970V subsystem, the controller frame also contains the disk array components. The 9900V 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),
- Channel adapters (ExSA™, fibre-channel, FICON™) (see section 2.2.6),
- Channels (see section 2.2.7), and
- 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, channel adapters, 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™ V Series subsystem can sustain the loss of multiple power supplies and still continue operation. In addition, the 9900V supports connection to a UPS to provide additional battery backup capability (see section 5.5.5).

Note: The redundancy and backup features of the Lightning 9900™ V Series 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 9900V 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 seven-day battery backup. The basic size of the shared memory is 2 GB, and the maximum size is 3 GB. 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 9980V subsystem can be configured with up to 64 GB of cache (increments of 4 GB), and the 9970V can be configured with up to 32 GB of cache (increments of 2 GB). All cache memory in the 9900V is nonvolatile and is protected by 48-hour battery backup. The cache in the 9900V 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 9900V 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 sortwork data) 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 9900V 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 9900V 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 9900V 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 9900V subsystem can sustain the loss of multiple power supplies and still continue operation. To make use of this capability, the 9900V 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 9900V can continue full operations using power from the alternate source.

The 9900V also supports connection to a UPS to provide additional battery backup capability.

2.2.6 Channel Adapters and CHiPs

The channel adapter boards contain the front-end microprocessors (CHiPs) 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. Channel adapter boards are installed in pairs. Each channel adapter board has four channel interfaces, which can all transfer at once, independently. Each channel adapter board pair is composed of one type of channel interface (e.g., ExSA™, FICON™, or fibre-channel). Fibre-channel adapters and FICON™-channel adapters are available in both shortwave (multimode) and longwave (single-mode) versions. The 9900V can be configured with multiple channel adapter pairs to support various interface configurations. Table 2.1 lists the channel adapter specifications and configurations and the number of channel connections for each configuration.

Note: Please contact your Hitachi Data Systems account team for information on the availability date of FICON™ channel interface support.

Note: Hitachi Performance Monitor and Hitachi Graph-Track (see section 3.6) allow users to collect and view usage statistics for the CHiPs in the 9900V subsystem.

Table 2.1 Channel Adapter Specifications

Parameter	Specification for 9980V	Specification for 9970V
Number of channel adapter pairs	1, 2, 3, or 4	1, 2, or 3
Simultaneous data transfers per CHA pair: S/390® Open Systems	8 ExSA™ (ESCON®) 8 fibre-channel	
Maximum transfer rate: ExSA™ (serial/ ESCON®) FICON™ Fibre-Channel	17 MB/sec 100 MB/sec (1 Gbps) 200 MB/sec (2 Gbps)	
Physical interfaces per CHA pair	8 ExSA™ (ESCON®) 8 FICON™ 8 fibre-channel	
Maximum physical interfaces per subsystem: ExSA™ (serial/ESCON®) FICON™ Fibre-channel	32 0, 8, 16, 24, or 32 0, 8, 16, 24, or 32 0, 8, 16, 24, or 32	24 0, 8, 16, or 24 0, 8, 16, or 24 0, 8, 16, or 24
Logical paths per FICON™ port	512	
Logical paths per ExSA™ (ESCON®) port	256	
Maximum logical paths per subsystem	8192	6144
Maximum LUs per fibre-channel port	512	
Maximum LVIs/LUs per subsystem	8192	8192

2.2.7 Channels

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

- **Fibre-Channel.** The 9980V subsystem supports up to 32 fibre-channel ports, and the 9970V supports up to 24 fibre ports. The fibre ports are capable of data transfer speeds of 200 MB/sec (2 Gbps). The 9900V fibre-channel cards have eight ports per pair of channel adapter boards. The 9900V supports shortwave (multimode) and longwave (single-mode) versions of fibre-channel adapters. When configured with shortwave fibre-channel cards, the 9900V subsystem can be located up to 500 meters (2750 feet) from the open-system host(s). When configured with longwave fibre-channel cards, the 9900V subsystem can be located up to ten kilometers from the open-system host(s).
 - **Extended Serial Adapter™ (ExSA™).** The 9980V subsystem supports a maximum of 32 ExSA™ serial channel interfaces (compatible with ESCON® protocol), and the 9970V supports a maximum of 24 ExSA™ interfaces. The 9900V ExSA™ channel interface cards provide data transfer speeds of up to 17 MB/sec and have a total of eight ports per pair of channel adapter boards. Each ExSA™ channel can be directly connected to a CHPID or to a serial channel director. Shared serial channels can be used for dynamic path switching. The 9900V also supports the ExSA™ Extended Distance Feature (XDF).
 - **FICON™.** The 9980V subsystem supports* up to 32 FICON™ ports, and the 9970V supports up to 24 FICON™ ports. FICON™ ports are capable of data transfer speeds of up to 100 MB/sec (1 Gbps). FICON™ features, available in both shortwave (multimode) and longwave (single mode) versions, have a total of 8 FICON™ host interfaces per pair of FICON™ channel adapter cards. The 9900V supports shortwave (multimode) and longwave (single-mode) versions of FICON™ channel adapters. When configured with shortwave FICON™-channel cards, the 9900V subsystem can be located up to 500 meters (2750 feet) from the open-system host(s). When configured with longwave FICON™-channel cards, the 9900V subsystem can be located up to ten kilometers from the open-system host(s).
- * **Note:** Please contact your Hitachi Data Systems account team for information on the availability date of FICON™ channel interface support.

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 9980V can be configured with up to four ACP pairs, and the 9970V can be configured with one or two ACP pairs. 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. ACPs are installed in pairs. Each dual-ported fibre-channel disk drive is connected via its two ports to each ACP in a pair via separate physical paths for reasons of performance and redundancy. Each 9980V fibre backend path can access up to 64 disk drives (64 drives \times 4 pairs of paths = 256 disk drives per ACP pair). Each 9970V fibre backend path can access up to 32 disk drives (32 drives \times 4 pairs of paths = 128 disk drives per ACP pair). In a 9970V with two ACP pairs, each fibre backend path can access up to 16 disk drives (16 drives \times 4 pairs of paths = 64 disk drives per ACP pair).

Table 2.2 lists the ACP specifications. Each 9980V ACP pair can support a maximum of 256 physical disk drives (in two 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: Hitachi Performance Monitor and Hitachi Graph-Track (see section 3.6) allow users to collect and view usage statistics for the ACPs in the 9900V subsystem.

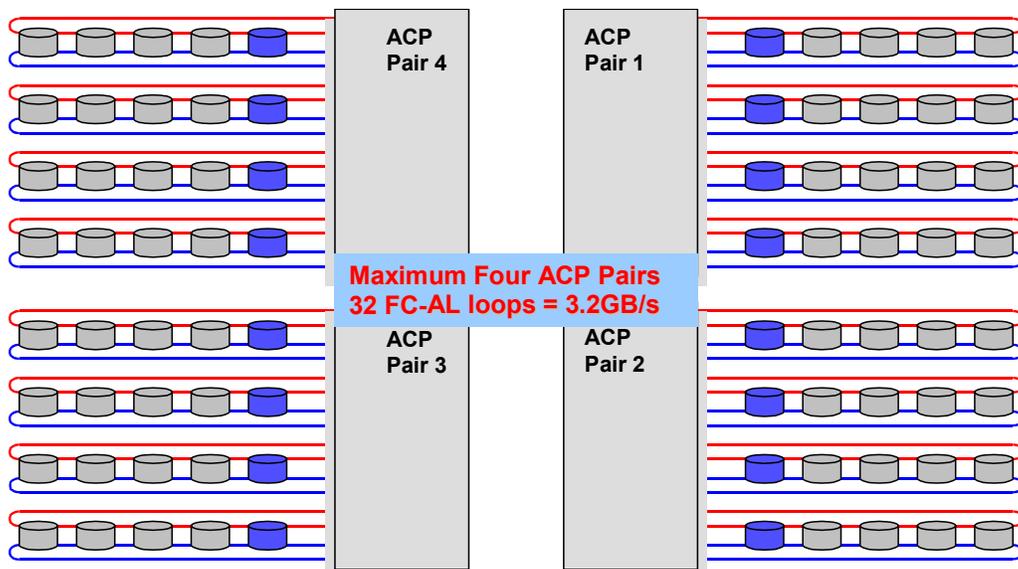


Figure 2.4 Conceptual 9980V ACP Array Domain

Table 2.2 ACP Specifications

Description	Specification for 9980V	Specification for 9970V
Number of ACP pairs	1, 2, 3 or 4	1 or 2
Backend paths per ACP pair	8	
Backend paths per subsystem	8, 16, 24 or 32	8 or 16
Array group (or parity group) type per ACP pair	RAID-1 and/or RAID-5	
Hard disk drive type per ACP pair ^[1]	36 GB, 73 GB	
Logical device emulation type within ACP pair	3390-x and OPEN-x ^[2]	
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 or 16
HiStar Network architecture internal bandwidth	10 GB/sec	5 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. 3390-3 and 3390-3R LVIs cannot be intermixed in the same 9900V subsystem.

2.3 Array Frame

The 9980V 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 9900V also supports connection to a UPS to provide additional battery backup capability. The 9980V can be configured with up to four array frames to provide a raw storage capacity of up to 74 TB. The 9970V subsystem combines the controller and disk array components in one physical frame to provide a raw storage capacity of up to 9 TB. When configured in four-drive RAID-5 parity groups, $\frac{3}{4}$ of the raw capacity is available to store user data, and $\frac{1}{4}$ of the raw capacity is used for parity data.

The 9900V subsystem uses disk drives with fixed-block-architecture (FBA) format. The currently available disk drives have capacities of 36 GB and 73 GB. All drives in an array group must have the same rotation speed and the same capacity. The 36-GB and 73-GB HDDs can be attached to the same ACP pair. Table 2.3 provides the disk drive specifications.

Each disk drive can be replaced nondisruptively on site. The 9900V 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 rotation speed and the same or lower 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 9980V can be configured with a minimum of one and a maximum of sixteen spare disk drives. The 9970V can be configured with a minimum of one and a maximum of four spare disk drives. The standard configuration provides one spare drive for each 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 Capacity	
	73 GB	36 GB
Formatted capacity (GB)	72.91	35.76
Platter diameter	3 inches	2.5 inches
Physical tracks per physical cylinder (user area) (number of heads)	10	8
Physical disk platters (user area) (numbers of disks)	5	4
Sector length (byte)	520(512)	520(512)
Seek time (ms) (Read/Write)	MIN.	0.5/0.7
	MAX.	11.0/12.0
	AVE.	4.9/5.7
Revolution speed (rpm)	10,025	14,904
Average latency time (ms)	2.99	2.01
Internal data transfer rate (MB/sec)	44.2 to 74.0	68.5 to 88.3
Maximum interface data transfer rate (MB/sec)	100	100

2.3.1 Disk Array Groups

The disk array group is the basic unit of storage capacity for the 9900V. 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 128 physical disk drives.

The 9900V 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 9900V subsystem is ideal for workloads with low cache-hit ratios.

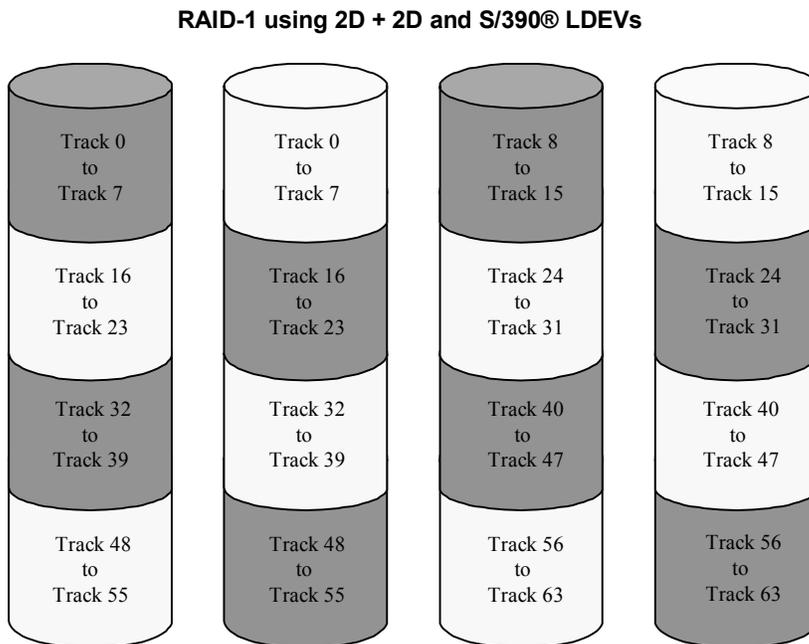


Figure 2.5 Sample RAID-1 Layout

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 9900V 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 3390-x or OPEN-x logical devices. All LDEVs in the array group must be the same format (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).

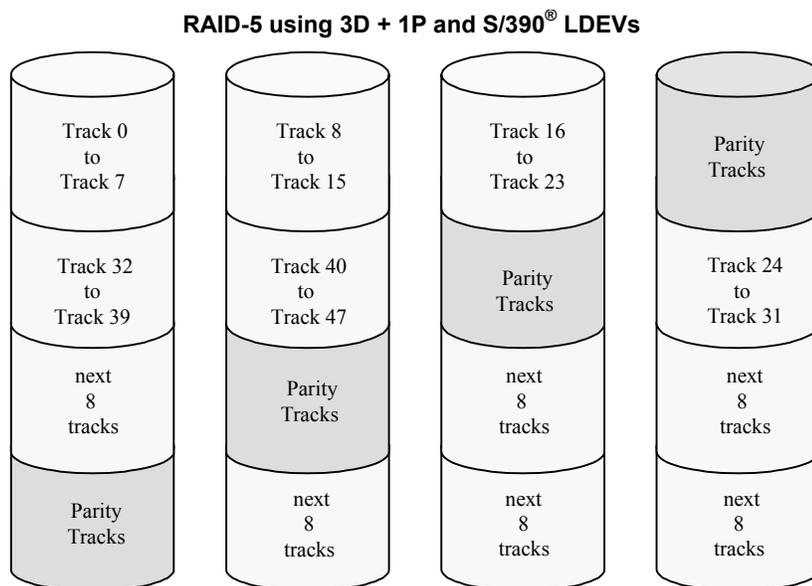


Figure 2.6 Sample RAID-5 Layout (Data Plus Parity Stripe)

Note: Hitachi Performance Monitor and Hitachi Graph-Track (see section 3.6) allow users to collect and view detailed usage statistics for the disk array groups in the 9900V subsystem.

2.3.2 Sequential Data Striping

The 9900V 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 9900V subsystem. The 9900V 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 9900V 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 (parity group) must be of the same capacity and type. The 36-GB and 73-GB HDDs can be attached to the same ACP pair.

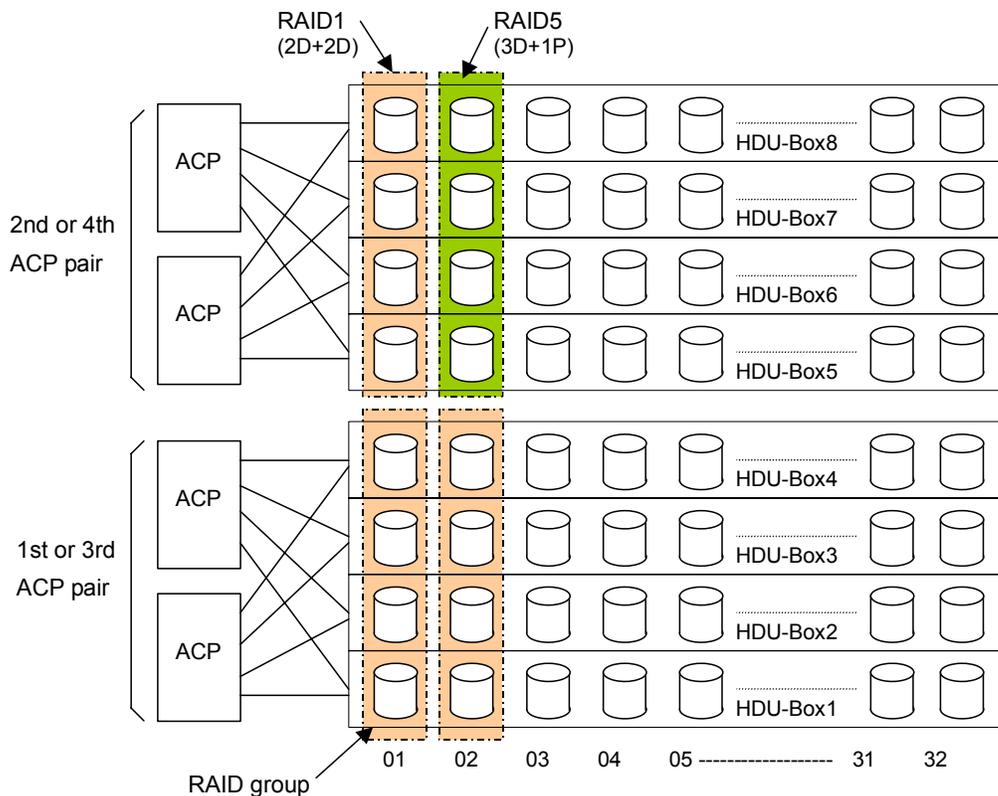


Figure 2.7 Sample Hard Disk Drive Intermix

2.4.3 Device Emulation Intermix

The 9900V subsystem supports an intermix of different device emulations (e.g., 3390-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:

- 3390-3 and 3390-9 can be intermixed within an array group.
- OPEN-3, -9, -E, and -L can be intermixed within an array group with the following restriction:
 - OPEN-L devices can only be configured on array groups of 73-GB HDDs.

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

Note: The 9980V and 9970V subsystems may support different device emulations and intermix configurations.

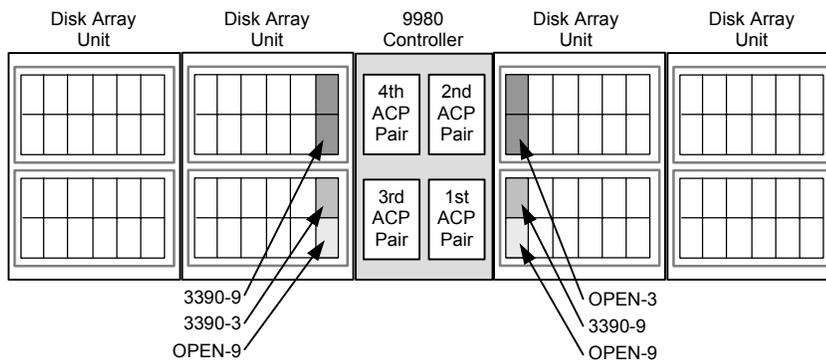


Figure 2.8 Sample Device Emulation Intermix

2.5 Service Processor (SVP)

The Lightning 9900™ V Series 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 9900V subsystem. The SVP also collects performance data for all key components of the 9900V subsystem to enable diagnostic testing and analysis. **Note:** The SVP does not have access to any user data stored on the 9900V subsystem.

The 9900V Remote Console functionality is provided by the SVP.

The Hitachi Graph-Track™ software product (see section 3.6.25) stores the SVP performance data and allows users to view the data in graphical format and export the data for statistical analysis.

2.6 Remote Console

The Hitachi Lightning 9900™ V Remote Console is provided as a Java applet program which can execute on any machine that supports a Java Virtual Machine (JVM). The Remote Console PC hosts the Remote Console Java applet program and is attached to the 9900V subsystem(s) via a TCP/IP local-area network (LAN). When a Remote Console accesses and logs into the desired SVP, the Remote Console applet is downloaded from the SVP to the Remote Console, runs on the web browser of the Remote Console PC, and communicates with the attached 9900V subsystems via a TCP/IP network. Figure 2.9 shows an example of Remote Console and SVP configuration.

Two LANs can be attached to the 9900V: the 9900V internal LAN (private LAN), which is used to connect the SVP(s) of multiple subsystems, and the user's intranet (public LAN), which allows you to access the Remote Console functions from individual Remote Console PCs. The Remote Console communicates directly with the service processor (SVP) of each attached subsystem to obtain subsystem configuration and status information and send user-requested commands to the subsystem. The 9900V Remote Console Java applet program is downloaded to the Remote Console (web client) from the SVP (web server) each time the Remote Console is connected to the SVP. The Remote Console Java applet program runs on Web browsers, such as Internet Explorer and Netscape Navigator, which run under the Windows® and Solaris® operating systems to provide a user-friendly interface for the 9900V Remote Console functions.

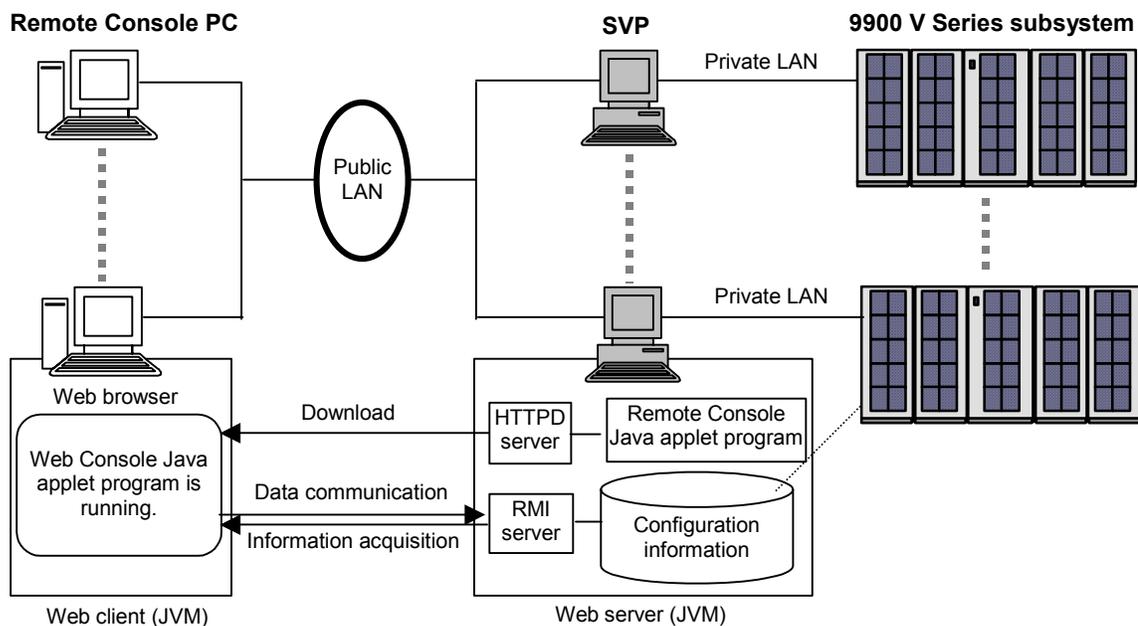


Figure 2.9 Example of Remote Console PC and SVP Configuration

Chapter 3 Functional and Operational Characteristics

3.1 New Features and Capabilities of the Lightning 9900™ V Series

The Hitachi Lightning 9900™ V Series subsystem offers the following new or improved features and capabilities which distinguish this model from the 9900 (9960/9910) subsystem:

- Thirty-two (32) logical control unit (CU) images (16 for 9900).
- Up to 8192 device addresses (4096 for 9900).
- Up to 64 GB cache memory for the 9980V subsystem (32 GB for 9960) and 32 GB for the 9970V (16 GB for 9910).
- Up to 1024 HDDs (512 HDDs in the 9900).
- Up to 512 LUNs per fibre-channel port (256 for 9900).
- Internal bandwidth increased by over three times relative to 9900.
- State-of-the-art hard disk drives of 36-GB and 73-GB capacities.

3.2 I/O Operations

The 9900V 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 9900V 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 9900V 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 9900V 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 9900V 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 9900V 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 Performance Monitor and Hitachi Graph-Track allow users to collect and view the write-pending-rate data and other cache statistics for the Lightning 9900™ V Series subsystem.

Note: If the write pending limit is reached, the 9900V 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 9900V subsystem supports the following logical CU images (emulation types): 3990-6, 3990-6E, and 2105. The 9900V subsystem is configured with one logical CU image for each 256 devices (one storage subsystem ID (SSID) for each 64 or 256 LDEVs) to provide a maximum of 32 CU images per subsystem. The S/390[®] data management features of the 9900V subsystem may have restrictions on CU image compatibility. FICON[™] channel interface 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 9900V subsystem supports the following S/390[®] LVI types: 3390-3, -3R, and -9. The LVI configuration of the subsystem depends on the RAID implementation and physical disk drive capacities. The 9900V LDEVs are accessed using a combination of CU number (0-1F) and device number (00-FF), as shown in Table 3.1. See section 4.1 for further information on LVI configurations.

Table 3.1 Device Numbers for Each CU

CU Number	Device Numbers		CU Number	Device Numbers
0	0:00 to 0:FF		8	8:00 to 8:FF
1	1:00 to 1:FF		9	9:00 to 9:FF
2	2:00 to 2:FF		A	A:00 to A:FF
3	3:00 to 3:FF		B	B:00 to B:FF
4	4:00 to 4:FF		C	C:00 to C:FF
5	5:00 to 5:FF		D	D:00 to D:FF
6	6:00 to 6:FF		E	E:00 to E:FF
7	7:00 to 7:FF		F	F:00 to F:FF

3.4.3 Logical Unit (LU) Type

The 9900V subsystem currently supports the following LU types: OPEN-3, OPEN-9, OPEN-E, and OPEN-L. Table 3.2 lists the capacities for each standard LU type.

Table 3.2 Capacities of Standard LU Types

LU Type	OPEN-3	OPEN-9	OPEN-E	OPEN-L
Capacity (GB)	2.461	7.384	14.568	36.450

The 9900V 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. LUN Expansion (LUSE) volumes can range in size from 4.92 GB (OPEN-3*2) to a maximum of 1312 GB (OPEN-L*36). Each LU is identified by target ID (TID) and LU number (LUN) (see Figure 3.1). Each 9900V fibre-channel port supports addressing capabilities for up to 256 LUNs when not using LUN Security and up to 512 LUNs when using LUN Security.

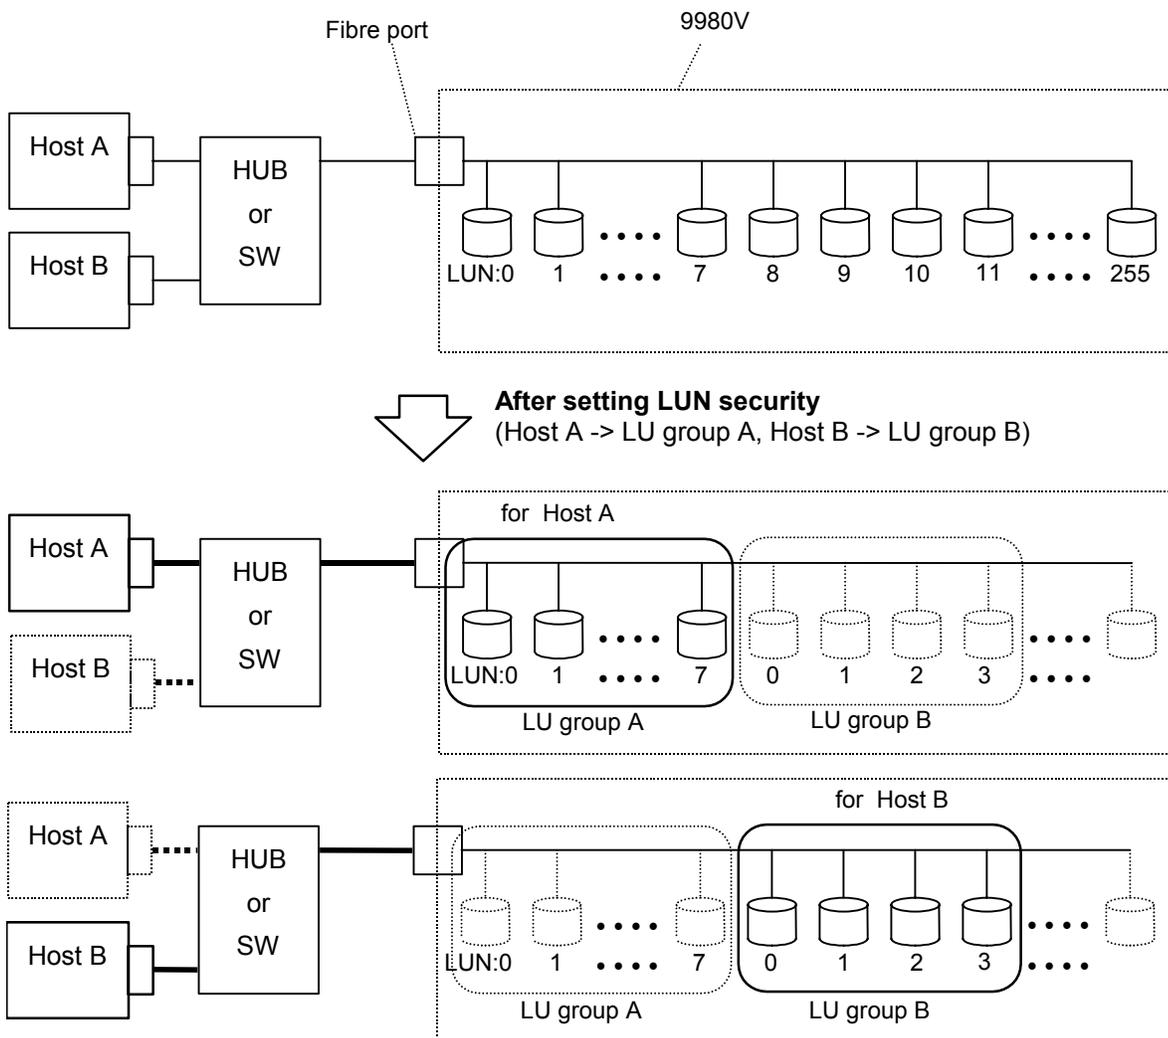


Figure 3.1 Fibre-Channel Device Addressing

3.5 Open Systems Features and Functions

The 9900V subsystem offers many features and functions specifically for the open-systems environment. The 9900V supports multi-initiator I/O configurations in which multiple host systems are attached to the same fibre-channel interface. The 9900V 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 3.6 for more information.

3.5.1 Failover and SNMP Support

The 9900V 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® Cluster Server, Sun Cluster, VERITAS® Volume Manager/DMP, HP® MC/ServiceGuard, HACMP, 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 9900V 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 9900V subsystem and the SNMP manager on the host. The SNMP agent for the 9900V subsystem sends status information to the host(s) when requested by the host or when a significant event occurs.

3.5.2 Share-Everything Architecture

The 9900V 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 9900V, 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 (this is a fee-based configuration service).

3.6 Data Management Functions

The 9900V subsystem provides features and functions that increase data availability and improve data management. Table 3.3 and Table 3.4 list the data management features that are currently available for the 9900V subsystem. Please refer to the appropriate user documentation for more details.

Table 3.3 Data Management Functions for Open-System Users

Feature Name	Controlled by:		Licensed Software?	User Document(s)
	Remote Console?	Host OS?		
Replication and migration:				
Hitachi TrueCopy (section 3.6.1)	Yes	Yes	Yes	MK-92RD108
Hitachi ShadowImage (section 3.6.3)	Yes	Yes	Yes	MK-92RD110
Command Control Interface (CCI) (section 3.6.5)	No	Yes	Yes	MK-90RD011
Data Migration (section 3.6.8)	Yes	No	No	Service Offering
Backup/restore and sharing:				
Hitachi RapidXchange (HRX) (section 3.6.9)	No	Yes	Yes	MK-91RD052
Hitachi Multiplatform Backup/Restore (HMBR) (section 3.6.10)	No	Yes	Yes	MK-92RD136
HARBOR® File-Level Backup and Restore (section 3.6.11)	No	No	No	Service Offering
HARBOR® File Transfer (section 3.6.12)	No	Yes	Yes	Service Offering
Resource management:				
HiCommand™ (section 3.6.13)	Yes	No	Yes	Web Client: MK-91HC001 Server: MK-91HC002 Security: MK-91HC003 HiScan: MK-91HC005 CLI: MK-91HC007
LUN Manager (section 3.6.14)	Yes	No	Yes	MK-92RD105
Hitachi SANtinel (LUN Security) (section 3.6.15)	Yes	No	Yes	MK-92RD105
LUN Expansion (LUSE) (section 3.6.16)	Yes	No	Yes	MK-92RD104
Virtual LUN (section 3.6.17)	Yes	No	Yes	MK-92RD104
FlashAccess (section 3.6.18)	Yes	No	Yes	MK-92RD102
Priority Access (section 3.6.20)	Yes	No	Yes	MK-92RD117
Hitachi Dynamic Link Manager™ (section 3.6.22)	No	Yes	Yes	For AIX®: MK-92DLM111 For HP-UX®: MK-92DLM112 For Solaris®: MK-92DLM114 For Windows®: MK-92DLM129
Storage utilities:				
Hitachi Performance Monitor (section 3.6.23)	Yes	No	Yes	MK-92RD133
Hitachi CruiseControl (section 3.6.24)	Yes	No	Yes	MK-92RD106
Hitachi Graph-Track (section 3.6.25)	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 3.4 Data Management Functions for S/390® Users

Feature Name	Controlled by:		Licensed Software?	User Document(s)
	Remote Console?	Host OS?		
Replication and migration:				
Hitachi TrueCopy – S/390® (section 3.6.2)	Yes	Yes	Yes	MK-92RD107
Hitachi ShadowImage – S/390® (section 3.6.4)	Yes	Yes	Yes	MK-92RD109
Hitachi Extended Remote Copy (HXRC) (section 3.6.6)	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 3.6.7)	N/A	N/A	N/A	MK-90DD878, Service Offering
Data Migration (section 3.6.8)	Yes	No	No	Service Offering
Backup/restore and sharing:				
Hitachi RapidXchange (HRX) (section 3.6.9)	No	Yes	Yes	MK-90RD020
Resource management:				
Virtual LVI (section 3.6.17)	Yes	No	Yes	MK-92RD104
FlashAccess (section 3.6.18)	Yes	Yes – Cache Manager	Yes	MK-92RD102
Cache Manager (section 3.6.19)	Yes – FlashAccess	Yes	No	MK-92RD103
Hitachi Parallel Access Volume (HPAV) (section 3.6.21)	Yes	Yes	Yes	MK-92RD118
Storage utilities:				
Hitachi Performance Monitor (section 3.6.23)	Yes	No	Yes	MK-92RD133
Hitachi CruiseControl (section 3.6.24)	Yes	No	Yes	MK-92RD106
Hitachi Graph-Track™ (section 3.6.25)	Yes	No	Yes	MK-90RD032

3.6.1 Hitachi TrueCopy

Hitachi TrueCopy enables open-system users to perform synchronous and/or asynchronous remote copy operations between 9900V 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 be performed between 9900V subsystems and between 9900V and 9900 subsystems.

Hitachi TrueCopy supports both fibre-channel and serial (ESCON®)* interface connections between the main and remote 9900V subsystems. 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. For serial interface connection, TrueCopy operations can be performed across distances of up to 43 km (26.7 miles) using standard ESCON® support. Long-distance solutions are provided, based on user requirements and workload characteristics, using approved channel extenders and communication lines.

* **Note:** Please contact your Hitachi Data Systems account team for information on the availability date of TrueCopy over serial (ESCON®) interface.

Note: For further information on Hitachi TrueCopy, please see the *Hitachi Lightning 9900™ V Series Hitachi TrueCopy User and Reference Guide* (MK-92RD108), or contact your Hitachi Data Systems account team.

3.6.2 Hitachi TrueCopy – S/390®

Hitachi TrueCopy - S/390® enables S/390® users to perform synchronous and asynchronous remote copy operations between 9900V subsystems. Hitachi TrueCopy - S/390® can be used to maintain copies of data for backup or duplication purposes. Once established, TrueCopy - S/390® 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 9900V subsystems. Remote copy operations can be performed between 9900V subsystems and between 9900V and 9900 subsystems.

Note: For further information on Hitachi TrueCopy - S/390®, please see the *Hitachi Lightning 9900™ V Series Hitachi TrueCopy - S/390® User and Reference Guide* (MK-92RD107), or contact your Hitachi Data Systems account team.

3.6.3 Hitachi ShadowImage

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 9900V 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 3.6.1) 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 9900V 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™ V Series ShadowImage User's Guide* (MK-92RD110), or contact your Hitachi Data Systems account team.

3.6.4 Hitachi ShadowImage – S/390®

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™ V Series ShadowImage - S/390® User's Guide* (MK-92RD109), or contact your Hitachi Data Systems account team.

3.6.5 Command Control Interface (CCI)

Hitachi Command Control Interface (CCI) enables users to perform Hitachi TrueCopy and Hitachi ShadowImage operations on the Lightning 9900™ V Series subsystem by issuing commands from the UNIX®/PC server host to the 9900V 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 9900V 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™ V Series Command Control Interface (CCI) User and Reference Guide (MK-90RD011)*, or contact your Hitachi Data Systems account team.

3.6.6 Hitachi Extended Remote Copy (HXRC)

The HXRC asynchronous remote copy feature of the 9900V 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 9900V subsystem. The Remote Console software 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 9900V-specific information on HXRC (e.g. SVP modes), please see the *Hitachi Lightning 9900™ V Series Hitachi TrueCopy - S/390® User and Reference Guide (MK-92RD107)*, 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)*.

3.6.7 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 Hitachi TrueCopy - S/390® Asynchronous, which is used to move large amounts of data over any distance with complete data integrity and minimal impact on performance. Hitachi TrueCopy - S/390® Asynchronous can be integrated with third-party channel extender products to address the "access anywhere" goal of data availability. Hitachi TrueCopy - S/390® Asynchronous 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™ V Series subsystem. Hitachi ShadowImage - S/390® can also operate in conjunction with Hitachi TrueCopy - S/390® 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.

3.6.8 Data Migration

The Lightning 9900™ V Series 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.

3.6.9 Hitachi RapidXchange (HRX)

Hitachi RapidXchange (HRX) enables the user to transfer data between S/390® and open-system platforms using the ExSA™ 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 LVIs and to the open-system host as OPEN-3 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 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 9900V 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.

3.6.10 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 9900V 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™ V Series Hitachi Multiplatform Backup/Restore User's Guide* (MK-92RD136), or contact your Hitachi Data Systems account team.

3.6.11 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.

3.6.12 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™ V Series storage subsystem.

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

3.6.13 HiCommand™

HiCommand™ provides a consistent, easy to use, and easy to configure set of interfaces for managing Hitachi storage products including the Lightning 9900™ V Series 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, LUN security, TrueCopy, and ShadowImage. **Note:** HiCommand™ 1.x does not support all Hitachi subsystem functions.

HiCommand™ enables users to manage the 9900V 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 3.3), or contact your Hitachi Data Systems account team.

3.6.14 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™ V Series LUN Manager and LUN Security User's Guide* (MK-92RD105), or contact your Hitachi Data Systems account team.

3.6.15 Hitachi SANTinel (LUN Security)

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 LUN Manager and LUN Security Remote Console software enables you to configure Hitachi SANTinel operations on the 9900V subsystem.

Hitachi SANTinel can be activated on any installed fibre-channel port, and can be turned on or off at the port level. 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 LUN Manager and LUN Security software allows you to create LU and WWN groups, so you can more easily manage your 9900V 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 and the LUN Security software, please see the *Hitachi Lightning 9900™ V Series LUN Manager and LUN Security User's Guide* (MK-92RD105), or contact your Hitachi Data Systems account team.

3.6.16 LUN Expansion (LUSE)

LUN Expansion 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 9900V 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™ V Series LUN Expansion and Virtual LVI/LUN User's Guide* (MK-92RD104), or contact your Hitachi Data Systems account team.

3.6.17 Virtual LVI/LUN

Virtual LVI/LUN allows users to convert fixed-size volumes into several smaller variable custom-sized volumes. Using the Remote Console software, 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 9900V, 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™ V Series LUN Expansion and Virtual LVI/LUN User's Guide* (MK-92RD104), or contact your Hitachi Data Systems account team.

3.6.18 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 9900V subsystem supports up to 1024 addressable cache extents.

FlashAccess operations can be performed for open-system LUs (e.g., OPEN-3, OPEN-9) as well as S/390® LVIs (e.g., 3390-3, 3390-9), 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™ V Series FlashAccess User's Guide* (MK-92RD102), or contact your Hitachi Data Systems account team.

3.6.19 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 9900V subsystem. For further information on Cache Manager, please see the *Hitachi Freedom Storage™ Cache Manager User's Guide* (MK-92RD103), or contact your Hitachi Data Systems account team.

3.6.20 Priority Access

Priority Access 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. Priority Access enables users to tune the performance of the development server without affecting production server performance.

Note: For further information on Priority Access, please see the *Hitachi Lightning 9900™ V Series Priority Access User's Guide* (MK-92RD117), or contact your Hitachi Data Systems account team.

3.6.21 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™ V Series subsystem. HPAV can provide substantially faster host access to the S/390® data stored in the 9900V subsystem. The Workload Manager (WLM) host software function enables the S/390® host to utilize the HPAV functionality of the Lightning 9900™ V Series subsystem. The 9900V supports both static and dynamic HPAV functionality.

Note: For further information on HPAV, please see the *Hitachi Lightning 9900™ V Series Hitachi Parallel Access Volume (HPAV) User and Reference Guide* (MK-92RD118), or contact your Hitachi Data Systems account team.

3.6.22 Hitachi Dynamic Link Manager™ (HDLM)

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

Note: For further information on HDLM, please see the *Hitachi Lightning 9900™ V Series Hitachi Dynamic Link Manager™ User's Guide* for your host platform (see Table 3.3), or contact your Hitachi Data Systems account team.

3.6.23 Hitachi Performance Monitor

Hitachi Performance Monitor provides detailed monitoring and collection of 9900V subsystem usage and performance statistics. Hitachi CruiseControl and Priority Access use the data collected by Performance Monitor to identify and resolve bottlenecks of activity.

Note: For further information on Hitachi Performance Monitor, please see the *Hitachi Lightning 9900™ V Series Hitachi Performance Monitor User's Guide (MK-92RD133)*, or contact your Hitachi Data Systems account team.

3.6.24 Hitachi CruiseControl

Hitachi CruiseControl enables users to optimize their data storage and retrieval on the 9900V subsystem. Hitachi CruiseControl analyzes detailed information on the usage of 9900V subsystem resources and tunes the 9900V 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 9900V 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™ V Series Hitachi CruiseControl User's Guide (MK-92RD106)*, or contact your Hitachi Data Systems account team.

3.6.25 Hitachi Graph-Track™

Hitachi Graph-Track™ (GT) allows users to monitor and collect detailed subsystem performance and usage statistics for the 9900V subsystem. GT can monitor as many as 32 subsystems on the 9900V-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 Freedom Storage™ Hitachi Graph-Track™ User's Guide* (MK-90RD032), or contact your Hitachi Data Systems account team.

Chapter 4 Configuring and Using the 9900V Subsystem

4.1 S/390® Configuration

The first step in configuring the Lightning 9900™ V Series subsystem 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 9900V 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 user-specified SSIDs are assigned during subsystem installation. Table 4.1 lists the SSID requirements.

Table 4.1 SSID Requirements

Controller Emulation	SSID Requirements	LVI Support
3990-6, 3990-6E	0004 - FFFD	3390-x, OPEN-3, -9, -E, -L
2105*	0004 - FFFD	3390-x, OPEN-3, -9, -E, -L

***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 9900V subsystem in MVS, VM, and VSE environments (wherever HCD cannot be used). The 9900V subsystem supports up to thirty-two logical CU (LCU) images and 8192 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™ channel interface 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 9900V 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 9900V 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 9900V configured with:

- 3990 ID.
- Two (2) LPARs called PROD and TEST sharing 4 ExSA™ (ESCON®) channels connected via 2 ESCDs to the 9980V.
- Each switch has ports C0 and C1 attached to the 9980V.
- Four (4) logical control units 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 (9900V connected to host CPU(s) via ESCD)

Figure 4.3 shows a sample IOCP hardware definition for a 9980V with:

- 2105 ID.
- Eight (8) ExSA™ (ESCON®) channels directly connected to the 9980V.
- Four (4) logical control units 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 9900V subsystem using one CNTLUNIT statement in IOCP. While it is possible to define an LCU on a 9900V 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 (9900V directly connected to CPU)

The 9980V subsystem can be configured with up to 32 connectable physical paths to provide up to 32 concurrent host data transfers. The 9970V 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 9900V 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 9900V subsystem in an MVS/ESA environment. The HCD procedures for 3990 and 2105 controller emulation types are described below. FICON™ interface requires 2105-F20 emulation.

3990 Controller Emulation. To define a 9900V subsystem with 64 or fewer LVIs, use the same procedure as for an IBM® 3990-6/6E subsystem (see Table 4.4). The hardware definition for a 9900V 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)
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 9900V as a single control unit allows all channel paths to access all DASD devices.

2105 Controller Emulation. To define a 9980V 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 9900V 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 -----
Row 1 of 1
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

/ Processor ID  UA + Time-Out  STADET  CHPID + Preferred Explicit Device
_ PROD         00   No         Yes     ___   No         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 4.17 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 9900V to VM/ESA® Systems

64 or Fewer LVIs: To define a 9900V with less than or equal to 64 LVIs to a VM/ESA® system, use the same procedure as for an IBM® 3990-6E subsystem. To define a 9900V 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 9900V with more than 64 LVIs to VSE/ESA, use the same procedure as for an IBM® 3990-6E subsystem. For 9900V 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 9900V to TPF

The 9900V 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 9900V'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 9900V 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™ V Series Hitachi Parallel Access Volume (HPAV) User and Reference Guide (MK-92RD118)*.

```
// 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 9900V 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 9900V 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 9900V, as contrasted to RAMAC.

Table 4.6 ICKDSF Commands for 9900V Contrasted to RAMAC

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

4.3.3 MVS Cache Operations

To display the 9900V cache statistics under MVS DFSMS, use the following operator command: **D SMS, CACHE**. Figure 4.22 shows the cache statistics reported by the 9900V. The 9900V 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 9900V 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 9900V 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 9900V 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 9900V, 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 9900V caches data for the logical volumes. The default caching status for the 9900V 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 9900V.

DFW OFF for all volumes

← Change DFW to **ON** for the 9900V.

Note: In normal cache replacement, bypass cache, or inhibit cache loading mode, the 9900V performs a special function to determine whether the data access pattern from the host is sequential. If the access pattern is sequential, the 9900V transfers contiguous tracks from the disks to cache ahead of time to improve cache hit rate. Due to this advance track transfer, the 9900V 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 9900V reports two storage sizes:

- **Subsystem storage.** This field shows capacity in bytes of cache. For a 9900V 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 9900V. For the 9900V, 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 9900V, this field shows one-fourth (25%) of the total cache size.

IDCAMS SETCACHE. The 9900V 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 9900V 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 9900V 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 9900V, 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 on the SVP, the user can also use the Remote Console to monitor the SIMs.

4.4 Open-Systems Configuration

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

Table 4.7 9900V Open-System Platforms and Configuration Guides

Platform	Configuration Guide Document Number
UNIX®-based platforms:	
IBM® AIX®	MK-92RD119
HP-UX®	MK-92RD122
Sun® Solaris®	MK-92RD123
PC server platforms:	
Windows NT®	MK-92RD120
Windows® 2000	MK-92RD121

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 Lightning 9900™ V Series LUN Manager and LUN Security User's Guide* (MK-92RD105), or contact your Hitachi Data Systems account team.

The 9980V subsystem supports a maximum of 32 fibre-channel ports, and the 9970V supports up to 24 fibre-channel ports. Each fibre-channel port is assigned a unique target ID (from 0 to EF). Each 9900V fibre-channel port supports addressing capabilities for up to a total of 512 LUNs across up to 128 host storage domains, each with its own LUN 0 and its own host mode. Each host storage domain may have up to 256 LUNs. Multiple host storage domains are supported using LUN Security. Figure 4.24 illustrates fibre port-to-LUN addressing.

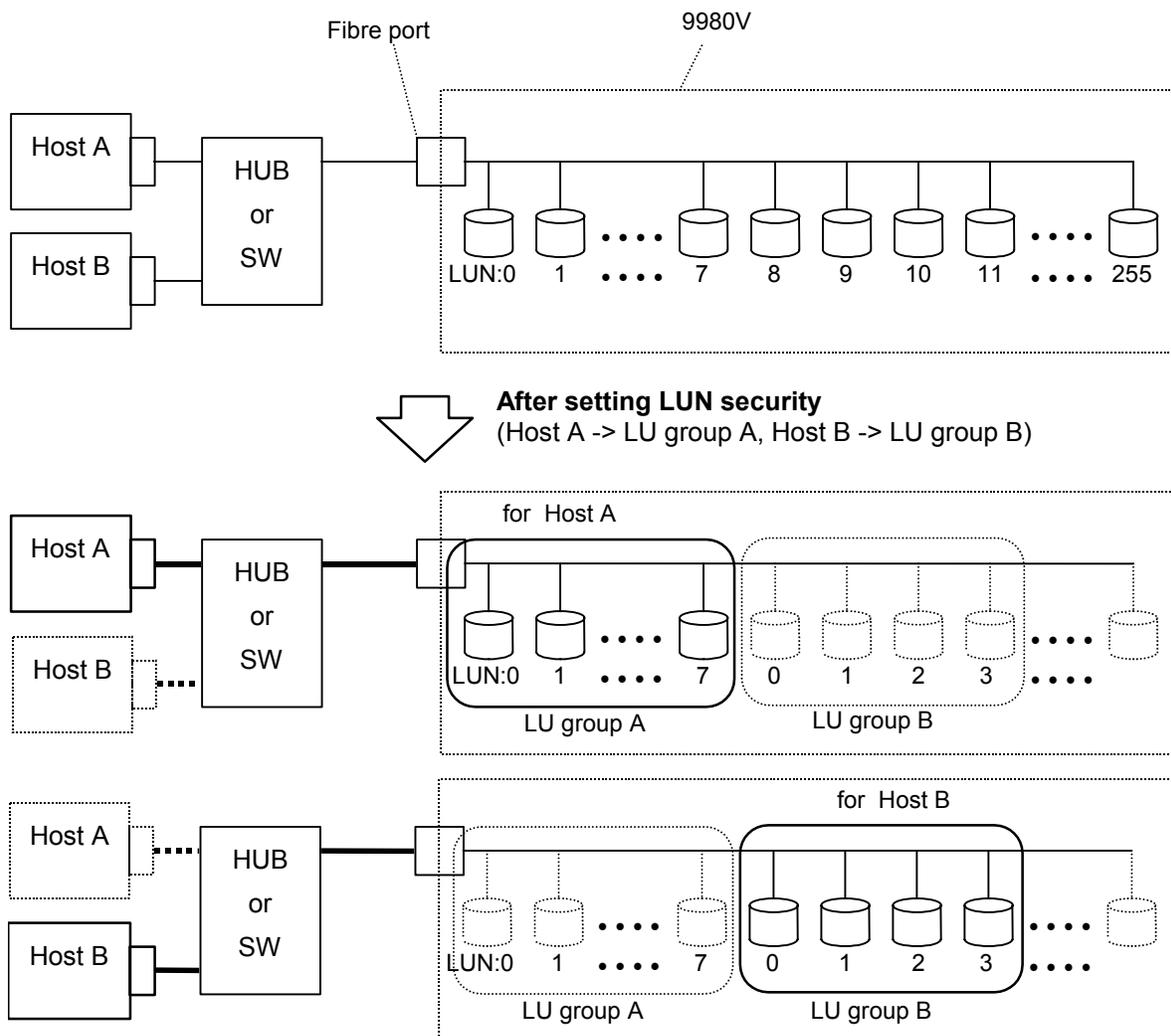


Figure 4.24 Fibre Port-to-LUN Addressing

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 LUN Expansion (LUSE) Devices

The LUSE 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 LUs which have been joined to form the LUSE device. For example, an OPEN-9*36 LUSE device is composed of 36 OPEN-9 LUs.

4.5 Open-Systems Operations

4.5.1 Command Tag Queuing

The 9900V 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 9900V 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 9900V devices. Please refer to the appropriate 9900V 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 9900V open-system configuration guides).

4.5.2 Host/Application Failover Support

The 9900V supports many industry-standard products which provide host and/or application failover capabilities (e.g., Hitachi Dynamic Link Manager™, VERITAS® Cluster Server, Sun® Cluster, VERITAS® Volume Manager/DMP, HP® MC/ServiceGuard, HACMP, Microsoft® Cluster Server). 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 9900V 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 9900V subsystem supports industry-standard I/O path failover products such as Hitachi Dynamic Link Manager™ and VERITAS® Volume Manager/DMP. Hitachi Dynamic Link Manager™ provides load balancing in addition to path failover. For the latest information on failover product releases, availability, and compatibility, please contact your Hitachi Data Systems account team.

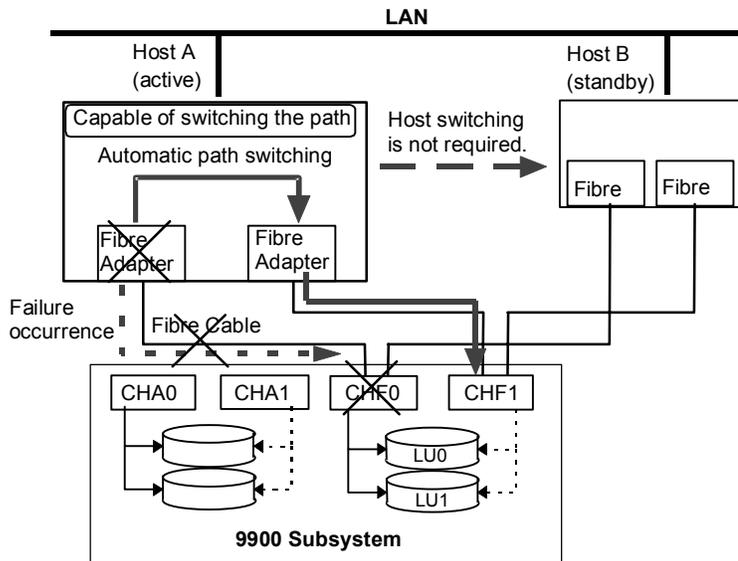


Figure 4.25 Alternate Pathing

4.5.4 SIM Reporting

The 9900V subsystem logs all SIMs on the SVP. When the user accesses the 9900V subsystem using the Remote Console software, the SIM log is displayed. This enables open-system users to monitor 9900V operations from any Remote Console PC. The Remote Console software allows the user to view the SIMs by date/time or by controller.

4.5.5 SNMP Remote Subsystem Management

The 9900V 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 9900V subsystem and the SNMP manager on the host. The SNMP agent on the 9900V sends status information to the host(s) when requested by a host or when a significant event occurs. Notification of 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 users to monitor the 9900V subsystem without having to check the Remote Console for SIMs.

4.5.6 NAS and SAN Operations

The Hitachi Lightning 9900™ V Series subsystem supports NAS and SAN configurations. For further information on 9900V NAS and SAN operations, please see section 1.1.5, refer to the *Hitachi Freedom NAS™ NSS Configuration Guide (MK-91RD053)*, or contact your Hitachi Data Systems account team.

4.6 Control Panel

Figure 4.26 shows the 9980V operator control panel. The 9970V control panel is the same except that the interfaces are A-M (not A-R). Table 4.8 describes the items on the operator control panel. To open the control panel cover, push and release on the point marked PUSH.

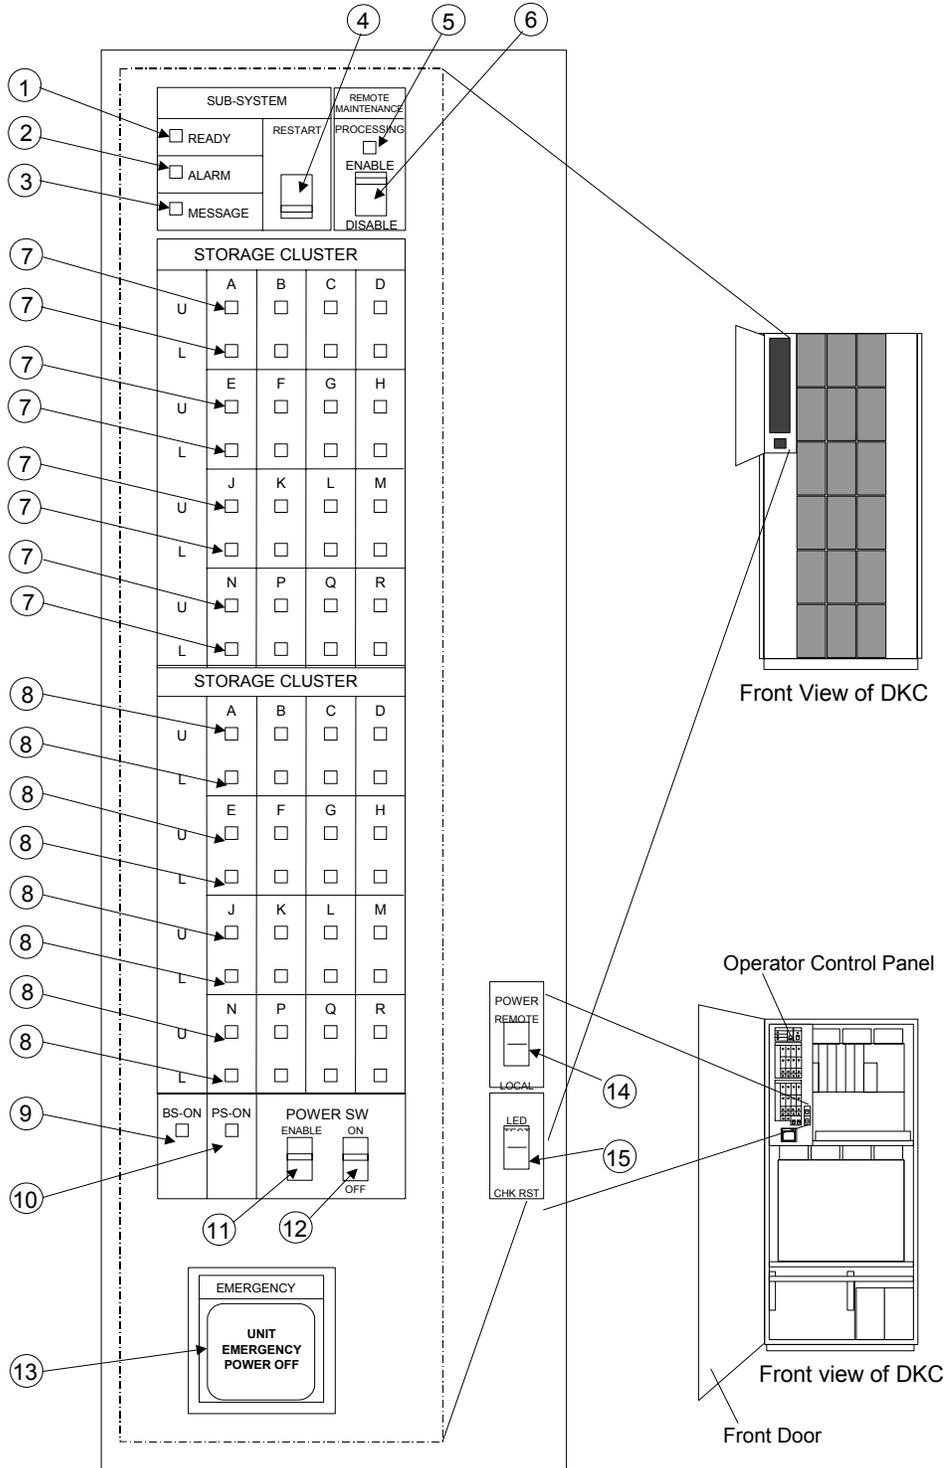


Figure 4.26 9980V Control Panel

Table 4.8 9980V Control Panel

Item	Name	Type	Description
1	SUBSYSTEM READY	LED (Green)	When lit, indicates that input/output operation on the channel interface is possible. Applies to both storage clusters.
2	SUBSYSTEM ALARM	LED (Red)	When lit, indicates that low DC voltage, high DC current, abnormally high temperature, or a failure has occurred. Applies to both storage clusters.
3	SUBSYSTEM MESSAGE	LED (Amber)	On: Indicates that a SIM (message) was generated from either of the clusters. Applies to both storage clusters. Blinking: Indicates that the SVP failure has occurred.
4	SUBSYSTEM RESTART	Switch	Used to unfence a fenced drive path and to release the Write Inhibit command. Applies to both storage clusters.
5	REMOTE MAINTENANCE PROCESSING	LED (Amber)	When lit, indicates that remote maintenance activity is in process. If remote maintenance is not in use, this LED is not lit. Applies to both storage clusters.
6	REMOTE MAINTENANCE ENABLE/DISABLE	Switch	Used for remote maintenance. While executing remote maintenance (the REMOTE MAINTENANCE PROCESSING LED in item 5 is blinking), when switching from ENABLE to DISABLE, remote maintenance is interrupted. If the remote maintenance function is not used, this switch is ineffective. Applies to both storage clusters.
7	STORAGE CLUSTER 1 CHANNEL A-R* ENABLE U: Upper L: Lower	LED (Green)	On: Indicates some of the logical paths are established. Fast blinking (Only Serial Channel): Indicates that the corresponding channel paths are executing the channel command. Slow blinking: Indicates none of the logical path is established. Off: Indicates that the corresponding channel paths are not enabled. When the 16-port CHA is installed, LED of the L side is effective.
8	STORAGE CLUSTER 2 CHANNEL A-R* ENABLE U: Upper L: Lower	LED (Green)	On: Indicates some of the logical paths are established. Fast blinking (Only Serial Channel): Indicates that the corresponding channel paths are executing the channel command. Slow blinking: Indicates none of the logical path is established. Off: Indicates that the corresponding channel paths are not enabled. When the 16-port CHA is installed, LED of the L side is effective.
9	BS-ON	LED (Amber)	Indicates input power is available.
10	PS-ON	LED (Green)	Indicates that subsystem is powered on. Applies to both storage clusters.
11	POWER SW ENABLE	Switch	Used to enable the POWER SW ON/OFF switch. To enable the POWER SW ON/OFF switch, turn the POWER SW ENABLE switch to the ENABLE position.
12	POWER SW ON / OFF	Switch	Used to power subsystem on/off. This switch is valid when REMOTE/LOCAL switch (item 14) is set to LOCAL. Applies to both storage clusters.
13	EMERGENCY POWER OFF (EPO)	1-Way Locking Switch	Used to shut down power to both storage clusters in an emergency.
14	REMOTE/LOCAL	Switch	REMOTE position: Subsystem is powered on/off by instructions from CPU. LOCAL position: Subsystem is powered on/off via the POWER SW ON / OFF switch. Applies to both storage clusters.
15	LED TEST/CHK RESET	Switch	LED TEST position: The LEDs on Control Panel go on. CHK RESET position: The PS ALARM and TH ALARM are reset.

* The 9970V subsystem has channel interfaces A-M.

4.6.1 Emergency Power-Off (EPO)

The disk subsystem EMERGENCY POWER OFF (EPO) switch is located on the operator control panel (see Figure 4.27). Use this switch only in case of an emergency.

To power off the disk subsystem in case of an emergency:

1. Open the operator control panel cover by pushing on the point marked PUSH.
2. Pull the emergency power-off switch (item 13 in Figure 4.26) up and then out towards you as illustrated on the switch.
3. Call the technical support center. The EPO switch must be reset by service personnel before the disk subsystem can be powered on again.

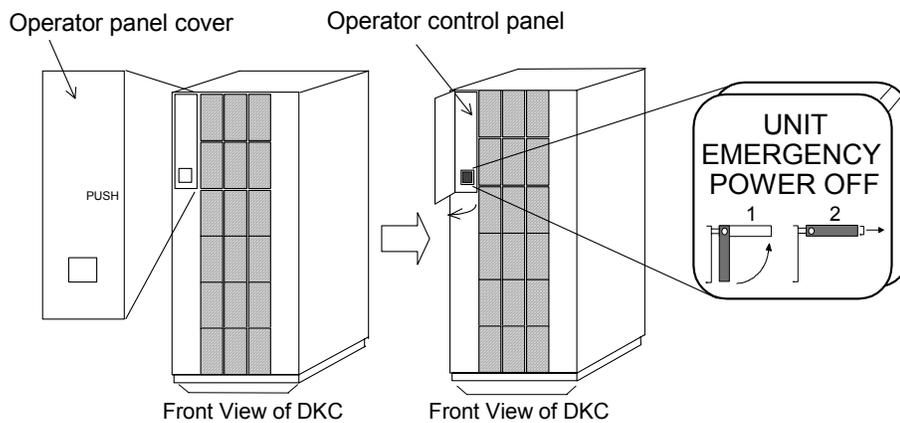


Figure 4.27 Emergency Power-Off (EPO)

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™ V Series subsystem. Please read this chapter carefully before beginning your installation planning. Figure 5.1 shows a physical overview of the 9980V subsystem. Figure 5.2 shows a physical overview of the 9970V subsystem.

If you would like to use any of the Lightning 9900™ V Series features or software products (e.g., TrueCopy, ShadowImage), 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 DKC460I/DKU455I and DKC465I (9980V and 9970V) 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™ V Series subsystem installation, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center.

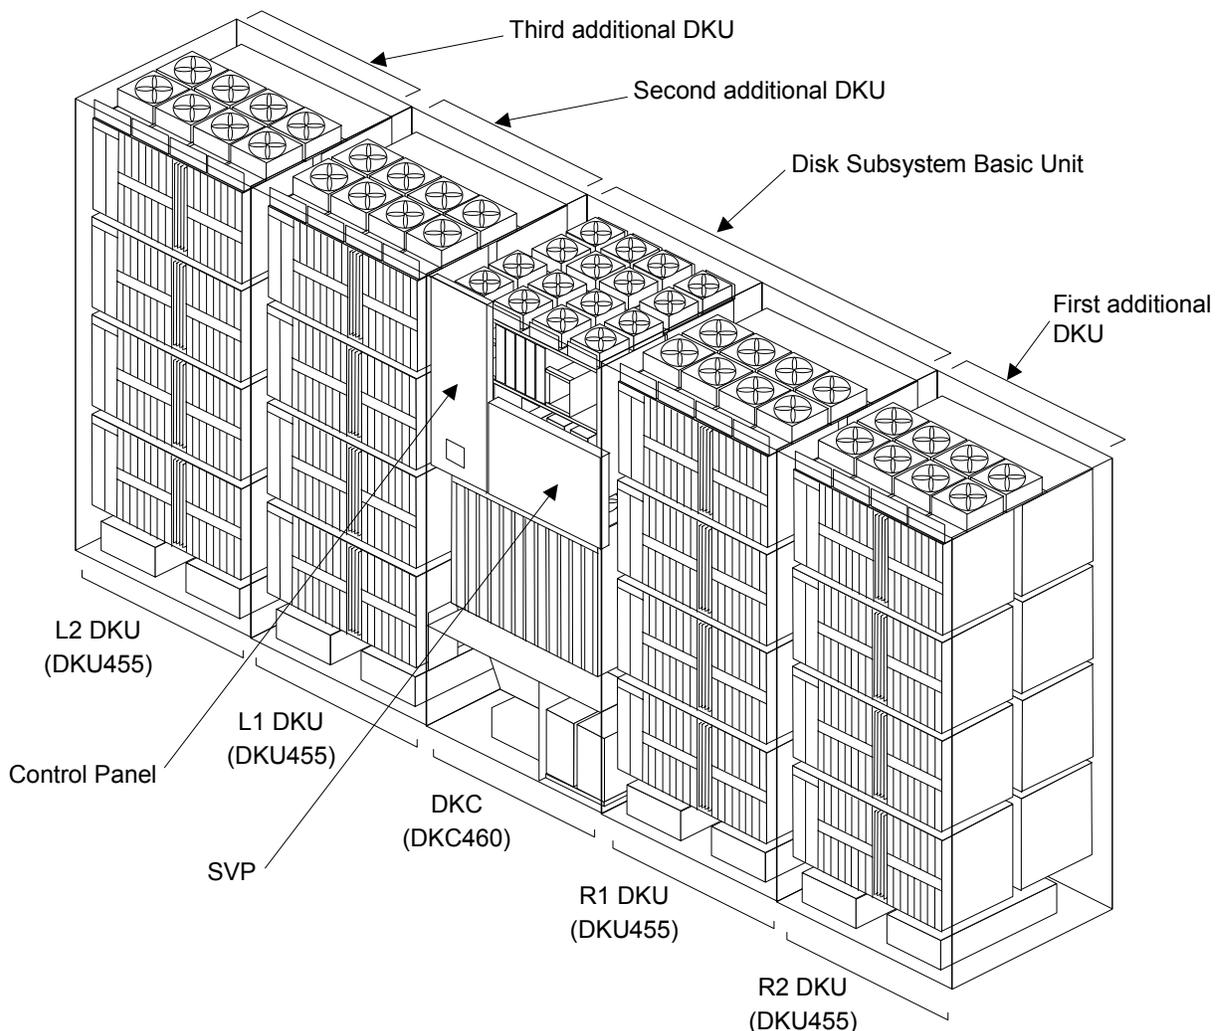


Figure 5.1 Physical Overview of 9980V Subsystem

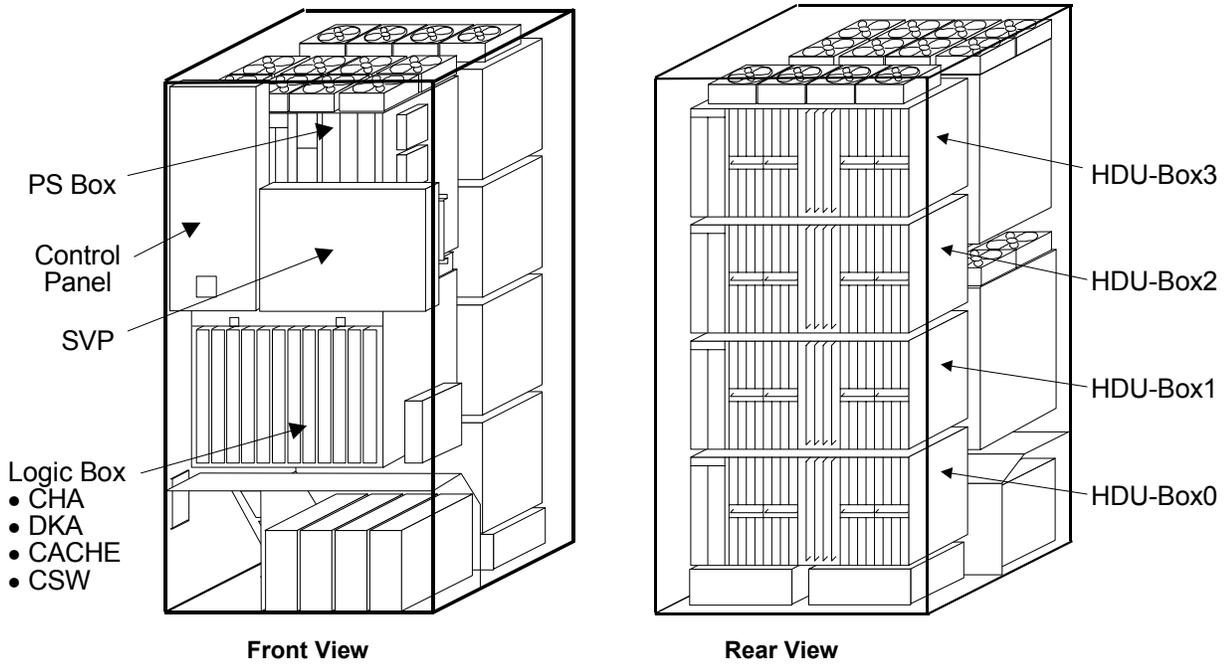


Figure 5.2 Physical Overview of the 9970V Subsystem

5.1 User Responsibilities and Safety Precautions

Before the 9900V 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)

5.1.1 Safety Precautions

For safe operation of the 9980V and 9970V disk subsystems, please observe the following precautions:

- Use the subsystems with the front and rear doors closed. The doors are designed for safety and protection from noise, static electricity, and EMI emissions.
- Make sure that all front and rear doors are closed before operating the subsystem. The only exceptions are during the power-up or power-down processes. Perform only the procedures described in this manual when the front and rear doors must be opened for operation.
- Before performing power-down or power-up, make sure that the disk subsystem is not undergoing any maintenance and is not being used online.
- Do not place objects on top of the frames, as this is where air is exhausted. This interferes with the flow of cooling air.
- For troubleshooting, perform only the instructions described in this manual. If you need further information, please contact Hitachi Data Systems maintenance personnel.
- In case of a problem with the subsystem, please report the exact circumstances surrounding the problem and provide as much detail as possible to expedite problem isolation and resolution.

5.2 Dimensions, Physical Specifications, and Weight

Figure 5.3 and Figure 5.4 show the physical dimensions of the 9980V and 9970V subsystems. Table 5.1 and Table 5.2 lists the physical specifications of the 9980V and 9970V subsystems.

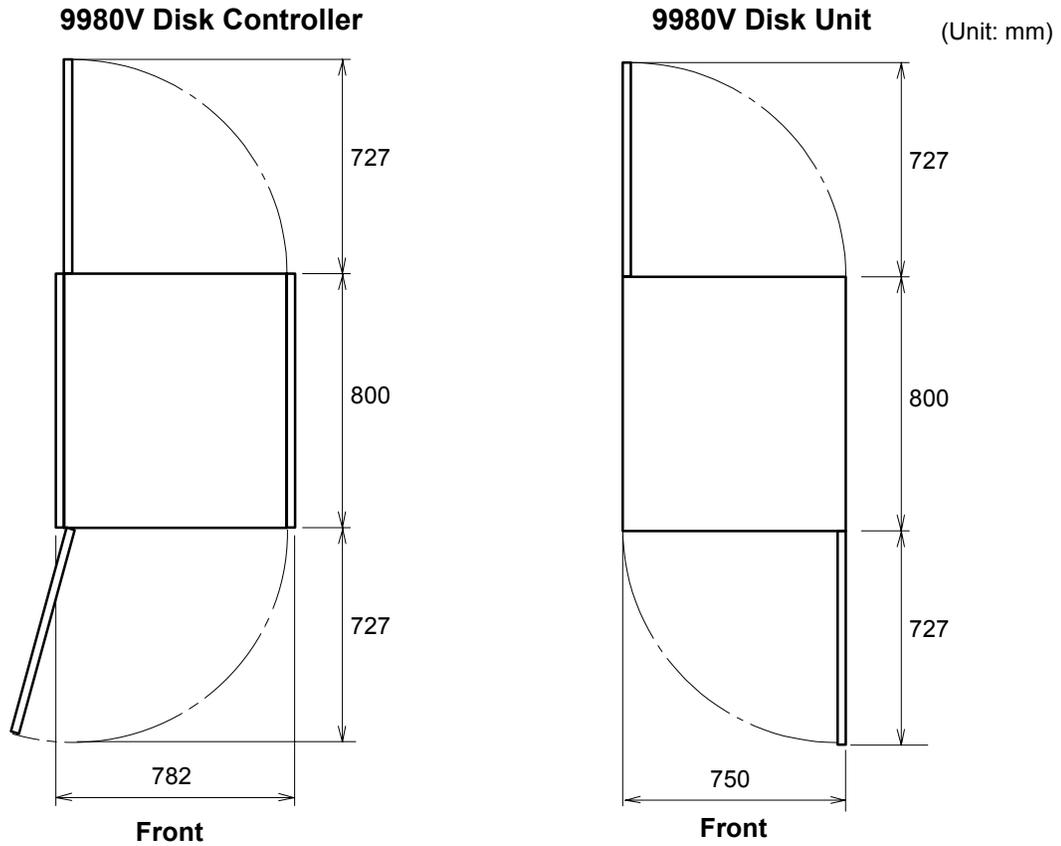


Figure 5.3 9980V DKC and DKU Physical Dimensions

9970V Subsystem

(Unit: mm)

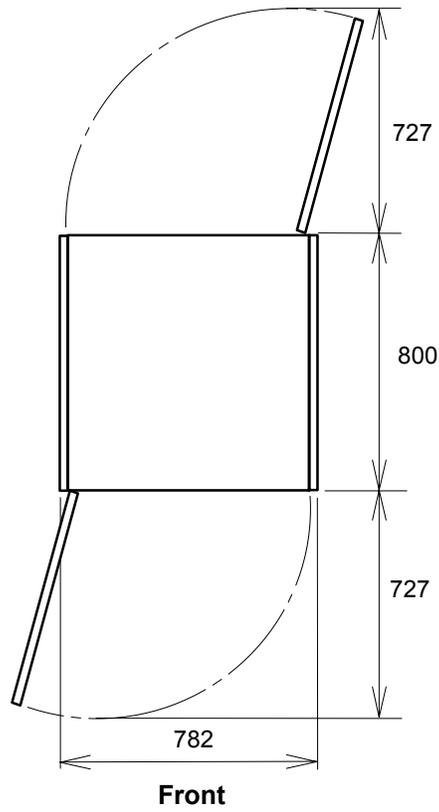


Figure 5.4 9970V Physical Dimensions

Table 5.1 9980V Physical Specifications

Item		9980V Controller (DKC460I)			9980V Disk Unit (DKC455I)	
Weight (kg)		499 ^{*1}	575 ^{*2}	575 ^{*3}	752 ^{*4}	752 ^{*5}
Heat Output (kW)		1.57 ^{*1}	2.68 ^{*2}	3.08 ^{*3}	6.14 ^{*4}	6.40 ^{*5}
Power Consumption (kVA)		1.65 ^{*1}	2.83 ^{*2}	3.27 ^{*3}	6.68 ^{*4}	6.94 ^{*5}
Air Flow (m ³ /min.)		18			32	
Dimensions (mm)	Width	782 ^{*6}			750	
	Depth	800			800	
	Height	1860			1860	

Notes:

- *1: These values are used when DKC460I has 16 GB cache memory, two fibre 8-port CHA options, and two DKA pairs.
- *2: These values are used when DKC460I has 64 GB cache memory, four fibre 8-port CHA options, and four DKA pairs.
- *3: These values are used when DKC460I has 64 GB cache memory, four MF fibre 8-port CHA options, and four DKA pairs.
- *4: These values are used when DKU455I is fully configured with 36-GB (15K rpm) HDDs.
- *5: These values are used when DKU455I is fully configured with 73-GB (10K rpm) HDDs.
- *6: This includes the thickness of side covers (16 mm × 2).

Table 5.2 9970V Physical Specifications

Item		Configured with 1 DKA Pair		Configured with 2 DKA Pairs	
		36 GB	73 GB	36 GB	73 GB
Weight (kg)		770 ^{*1}	770 ^{*2}	770 ^{*1}	770 ^{*2}
Heat Output (kW)		4.89 ^{*1}	5.02 ^{*2}	4.81 ^{*1}	4.94 ^{*2}
Power Consumption (kVA)		5.28 ^{*1}	5.41 ^{*2}	5.19 ^{*1}	5.32 ^{*2}
Air Flow (m ³ /min.)		25			
Dimensions (mm)	Width	782 ^{*3}			
	Depth	800			
	Height	1860			

Notes:

- *1: These values are used when DKC465I is fully configured with all CHA options and 36-GB (15K rpm) HDDs.
- *2: These values are used when DKC465I is fully configured with all CHA options and 73-GB (10K rpm) HDDs.
- *3: This includes the thickness of side covers (16 mm × 2).

5.3 Service Clearance, Floor Cutout, and Floor Load Rating Requirements

This section specifies the 9980V and 9970V service clearance requirements (a + b) based on the floor load rating and the clearance (c).

- Figure 5.5 shows the service clearance and floor cutout for the 9980V disk controller (DKC) frame. Figure 5.6 shows the service clearance and floor cutout for the 9980V disk unit (DKU) frame.
- Figure 5.7 shows the service clearance and floor cutouts for the 9980V subsystem minimum configuration (one DKC, one DKU). Table 5.3 shows the floor load rating requirements for the 9980V subsystem minimum configuration.
- Figure 5.8 shows the service clearance and floor cutouts for the 9980V subsystem maximum configuration (one DKC, four DKUs). Table 5.4 shows the floor load rating requirements for the 9980V subsystem maximum configuration.
- Figure 5.9 shows the service clearance and floor cutouts for the 9970V subsystem (all configurations). Table 5.5 shows the floor load rating requirements for the 9970V subsystem.

Note: The service clearance is required for service work. Do not use this space for storage of any article to prevent damage.

Note: Actual clearances for installation should be decided after consulting with construction specialist responsible for installation building, as clearances could vary depending on the size/layout of the system and building conditions.

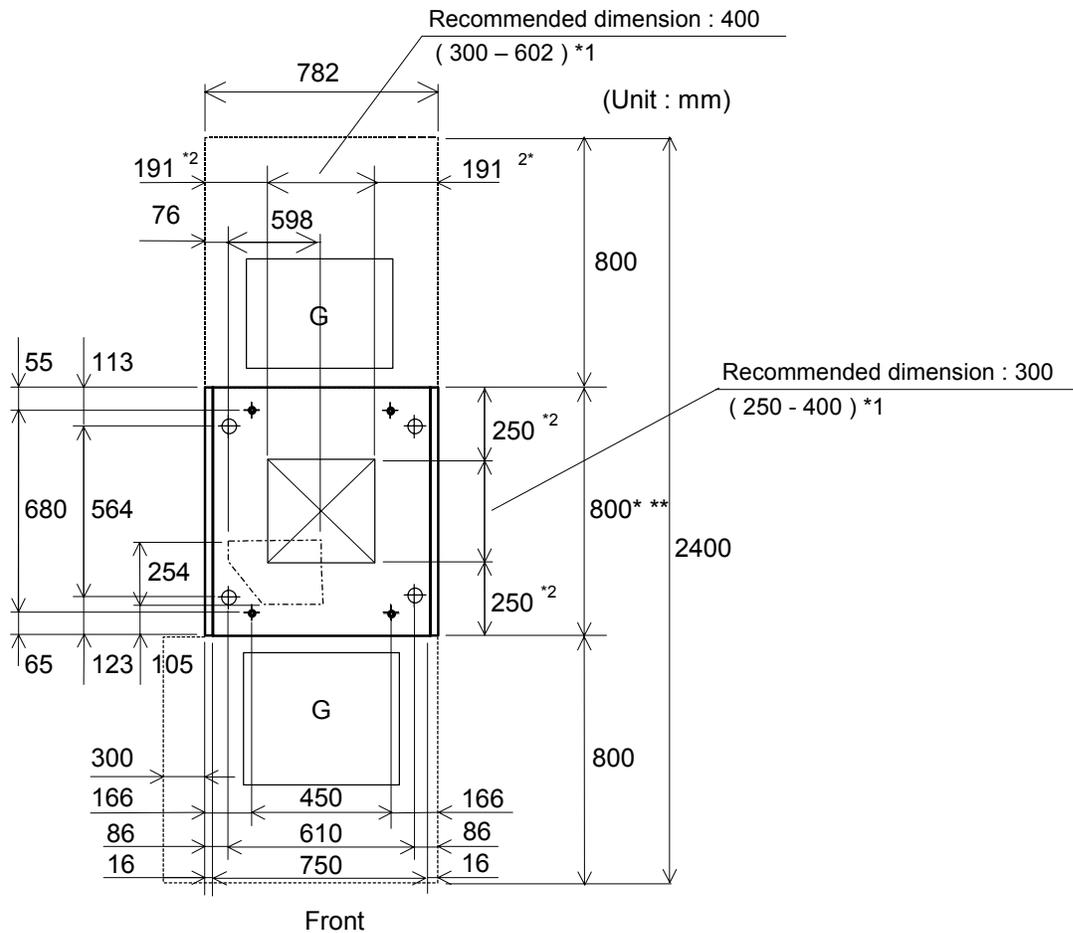
Note: When various configurations of subsystems are arranged in a row, use the clearance values based on the maximum subsystem configuration.

Note: For efficient maintenance operations, it is recommended that clearance (c) be made as large as possible.

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 760 mm (30 inches) from the machine.



- Floor cutout area for cables
- Screw jack
- Caster
- Service clearance
- Grid panel (over 450 mm × 450 mm)
- Opening on the bottom of the frame (for external cable entry)

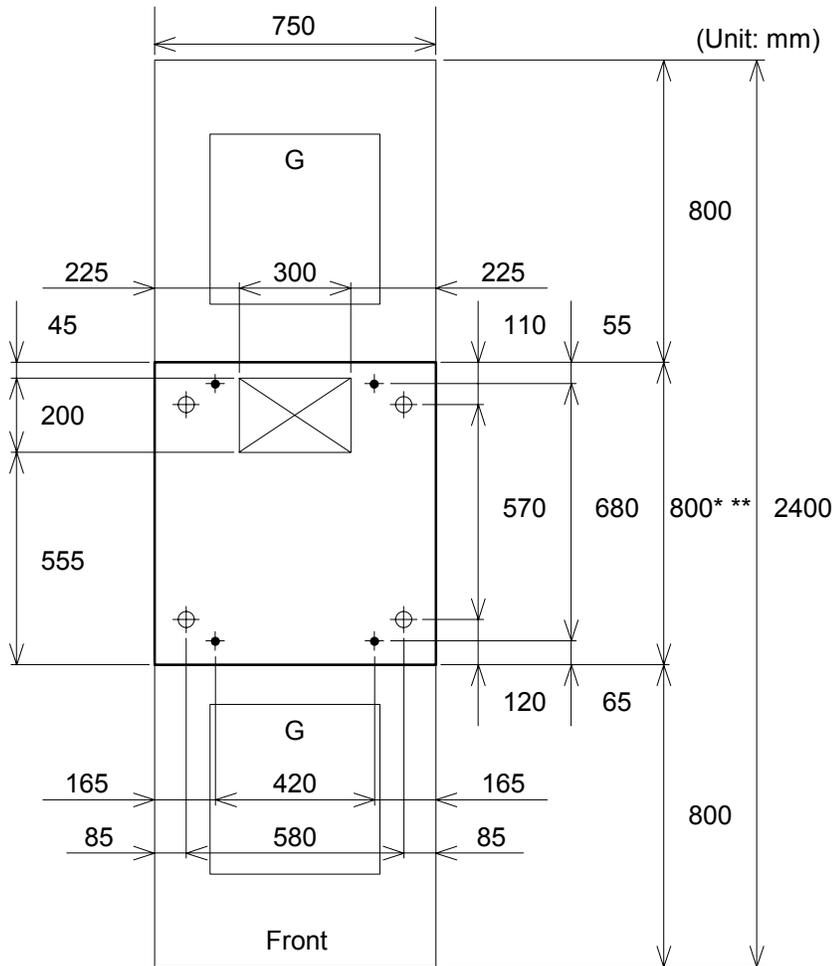
*1 Values in parentheses show the allowable range of the floor cutout dimension.
 The floor cutout should be in the center of the DKC. In case that the floor cutout is in the right position for the external cable work and is within the allowable range, the cutout position may be off-center. In this case, check the relation between the positions of the cutout and the opening on the bottom of the frame. If the floor cutout width is more than 552 mm, be careful about the restriction of the movable direction so that the caster wheels do not fall down into the cutout.

*2 These dimensions vary with the floor cutout dimension.

* The thickness of the door is different in the FRONT (35 mm) than in the REAR (25 mm).

** Overhang of the MOSAIC (LOUVER) of the FRONT DOOR (7 mm) is not included.

Figure 5.5 9980V DKC Service Clearance and Floor Cutout

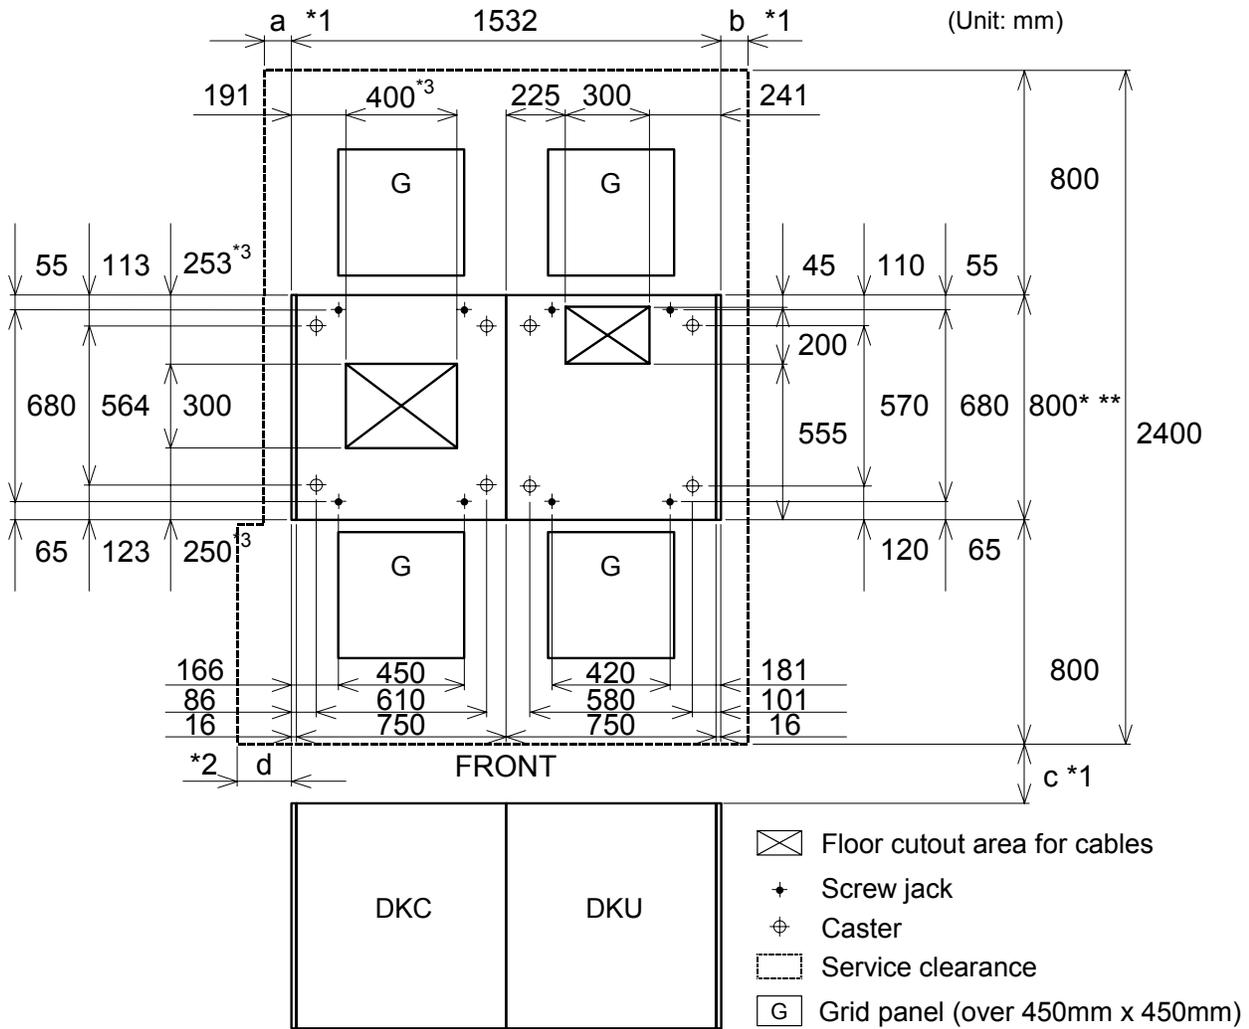


-  Floor cutout area for cables
-  Screw jack
-  Caster
-  Service clearance
-  Grid panel (over 450mm x 450mm)

*: The thickness of the door is different in the FRONT (35 mm) than in the REAR (25 mm).

** : Overhang of the MOSAIC (LOUVER) of the FRONT DOOR (7 mm) is not included.

Figure 5.6 9980V DKU Service Clearance and Floor Cutout

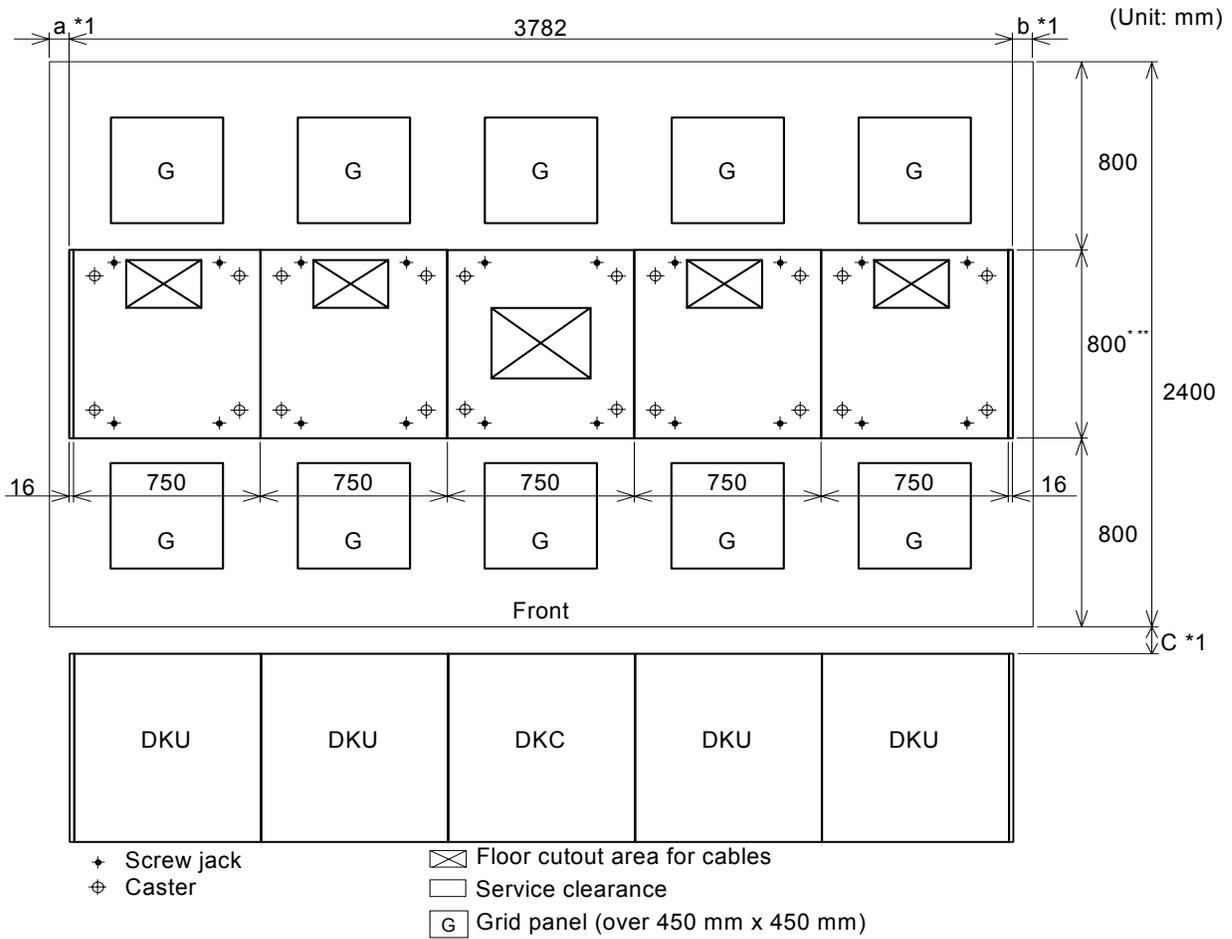


- *1 Clearance (a+b) depends on the floor load rating and clearance ©. Floor load rating and required clearances are in Table 5.3.
- *2 Clearance (d) must be over 0.3 m so as to open the DKC front door (refer to Figure 5.5). In case that clearance (d) is less than clearance (a), give priority to clearance (a).
- *3 See Figure 5.5 for details on the DKC floor cutout.
- * The thickness of the door is different in the FRONT (35 mm) than in the REAR (25 mm).
- ** Overhang of the MOSAIC (LOUVER) of the FRONT DOOR (7 mm) is not included.

Figure 5.7 9980V Subsystem Service Clearance and Floor Cutouts – Minimum Configuration

Table 5.3 Floor Load Rating and Required Clearances for 9980V Minimum Configuration

Floor Load Rating kg/m ² (lb/ft ²)	Required Clearance (a+b) in meters				
	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	0.4	0.2	0	0
450 (92.2)	0.9	0.7	0.4	0.3	0
400 (81.9)	1.4	1.0	0.8	0.6	0.3
350 (71.7)	2.0	1.6	1.3	1.0	0.7
300 (61.4)	3.0	2.5	2.1	1.8	1.3



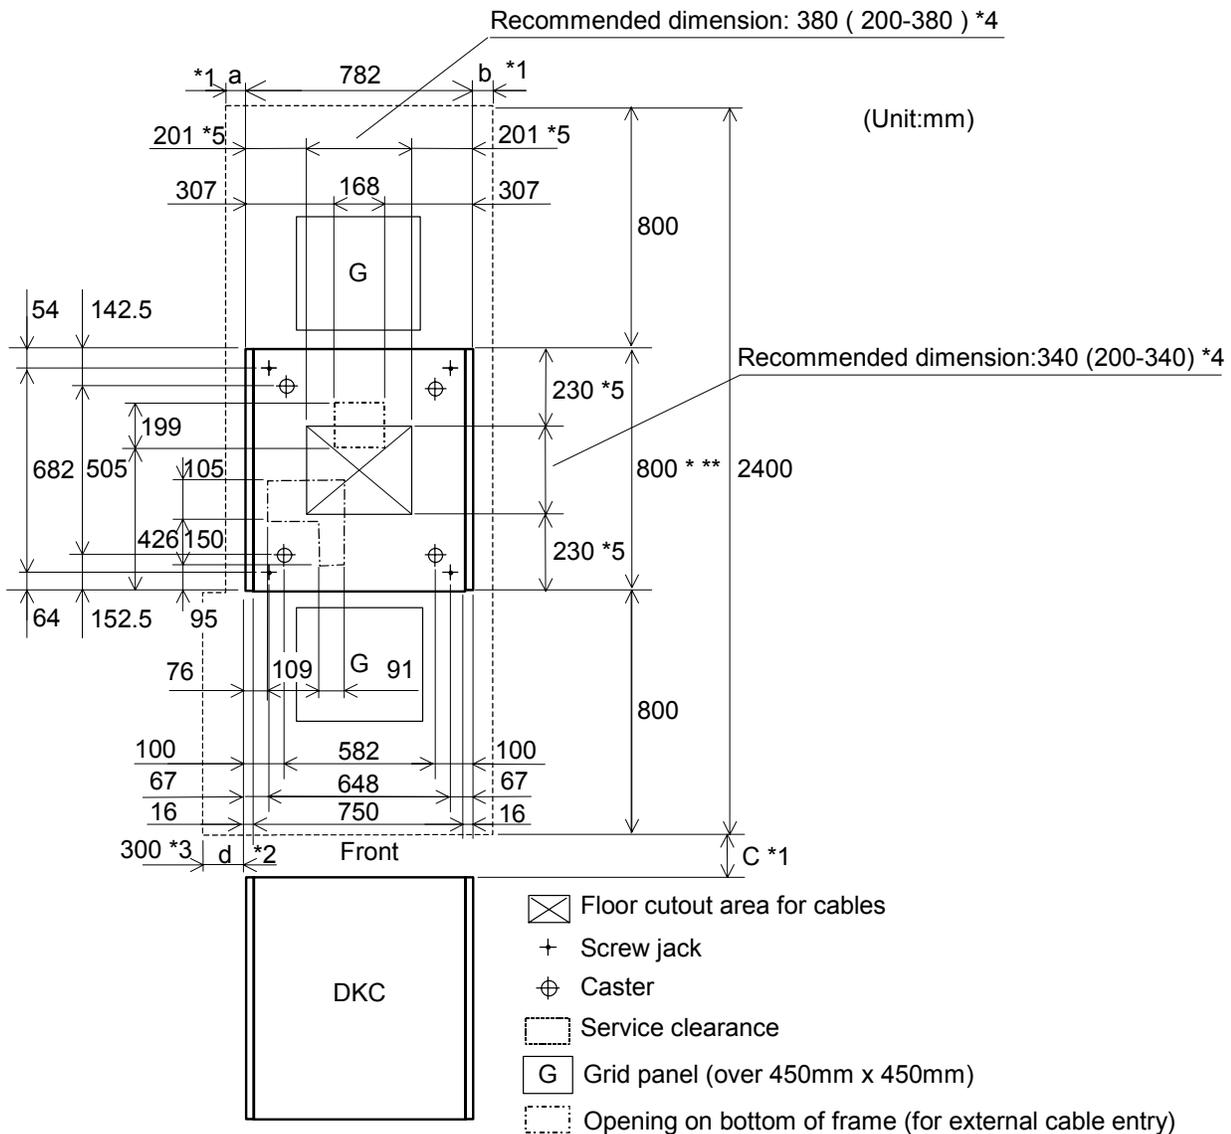
* The thickness of the door is different in the FRONT (35 mm) than in the REAR (25 mm).

** Overhang of the MOSAIC (LOUVER) of the FRONT DOOR (7 mm) is not included.

Figure 5.8 9980V Subsystem Service Clearance and Floor Cutouts – Maximum Configuration

Table 5.4 Floor Load Rating and Required Clearances for 9980V 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)	1.9	1.2	0.8	0.3	0
450 (92.2)	2.7	2.0	1.4	1.0	0.2
400 (81.9)	3.9	3.1	2.4	1.8	1.0
350 (71.7)	5.6	4.6	3.7	3.1	2.0
300 (61.4)	8.3	6.9	5.9	5.0	3.7



*: The thickness of the door is different in the FRONT (35 mm) than in the REAR (25 mm).

**:: Overhang of the MOSAIC (LOUVER) of the FRONT DOOR (7 mm) is not included.

Figure 5.9 9970V Subsystem Service Clearance and Floor Cutout – All Configurations

Table 5.5 Floor Load Rating and Required Clearances for 9970V 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)

5.4 Electrical Specifications and Requirements for Three-Phase Subsystems

The Lightning 9980V and 9970V subsystems support three-phase and single-phase power. This section applies to three-phase subsystems. For electrical specifications and requirements for single-phase subsystems, please see section 5.5.

5.4.1 Power Plugs for Three-Phase

- Figure 5.10 illustrates the power plugs for a three-phase 9980V disk array unit (Europe).
- Figure 5.11 illustrates the power plugs for a three-phase 9980V disk array unit (USA).
- Figure 5.12 illustrates the power plugs for a three-phase 9970V subsystem (Europe).
- Figure 5.13 illustrates the power plugs for a three-phase 9970V subsystem (USA).

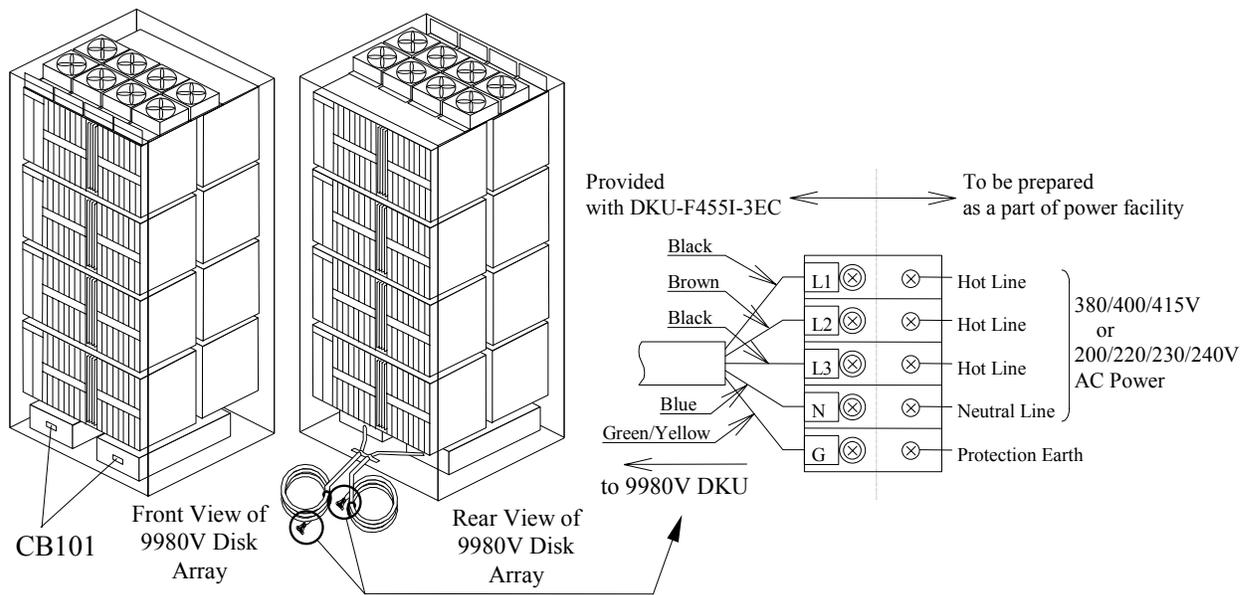


Figure 5.10 Power Plugs for Three-Phase 9980V Disk Array Unit (Europe)

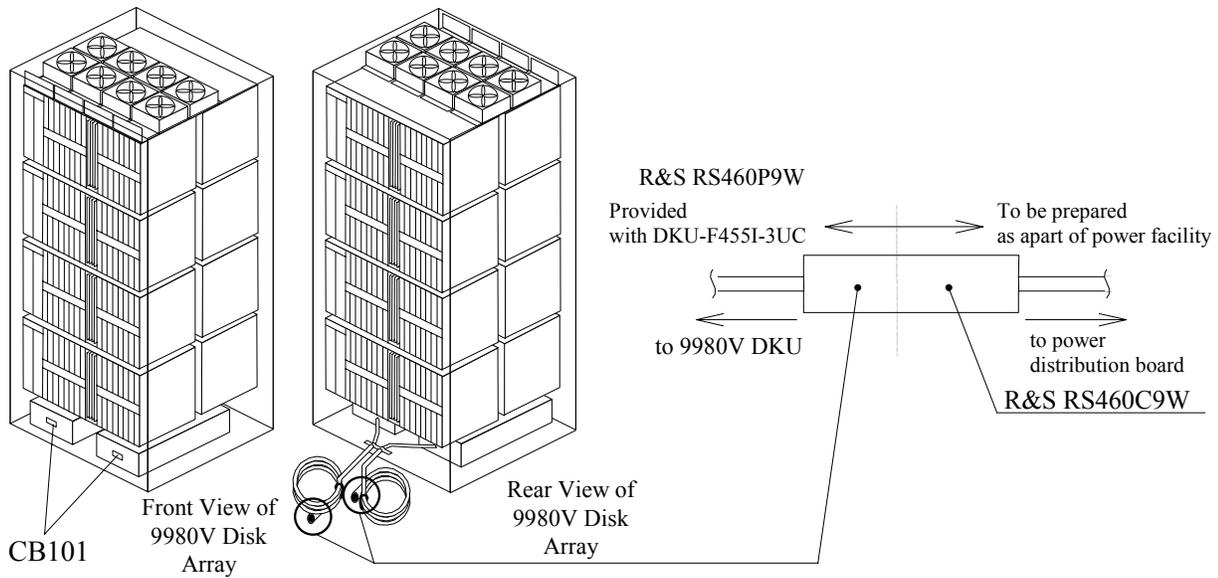


Figure 5.11 Power Plugs for Three-Phase 9980V Disk Array Unit (USA)

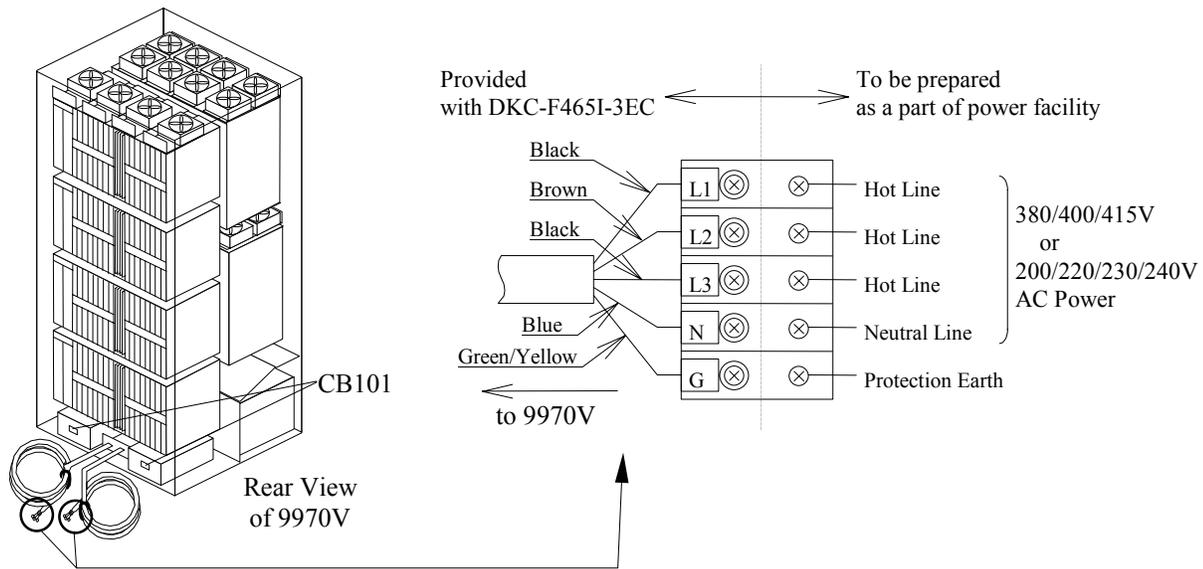


Figure 5.12 Power Plugs for Three-Phase 9970V Subsystem (Europe)

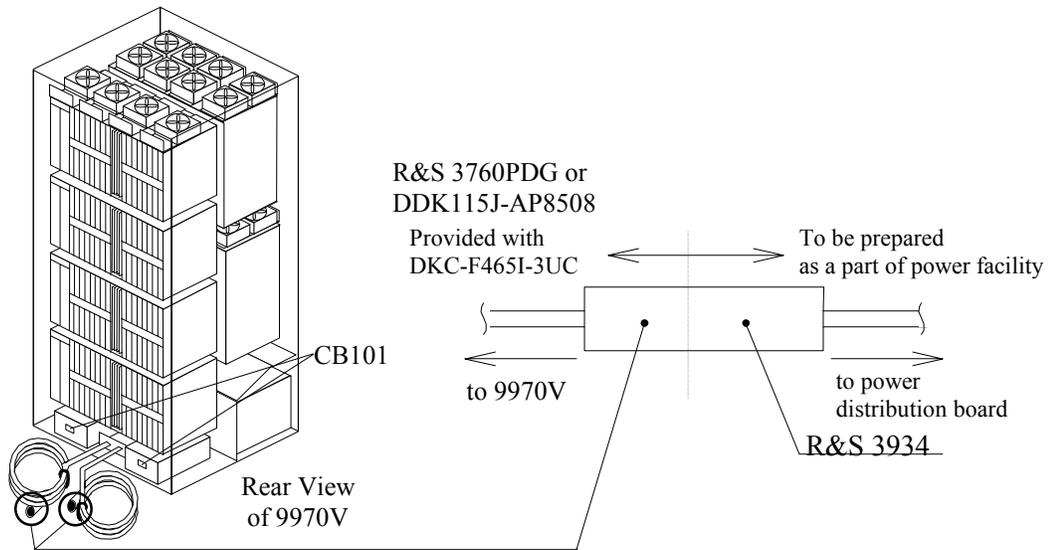


Figure 5.13 Power Plugs for Three-Phase 9970V Subsystem (USA)

5.4.2 Features for Three-Phase

Table 5.6 lists the features for three-phase 9980V and 9970V subsystems.

Table 5.6 9980V and 9970V Three-Phase Features

Model	Frame	Feature Number	Description	Comments
9980V	Controller	DKC-F460I-3PS	Breaker Box Kit for Three-Phase	For 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-F455I-3PS	AC Box Kit for Three-Phase	Consists of two AC boxes.
	Disk Array	DKU-F455I-3EC	Power Cable Kit for Three-Phase	50 Hz for Europe
	Disk Array	DKU-F455I-3UC	Power Cable Kit for Three-Phase	60 Hz for USA
9970V	The 9970V consists of a single frame.	DKC-F465I-3PS	AC Box Kit for Three-Phase	Consists of two AC boxes.
		DKU-F465I-3EC	Power Cable Kit for Three-Phase	50 Hz for Europe
		DKU-F465I-3UC	Power Cable Kit for Three-Phase	60 Hz for USA

5.4.3 Current Rating, Power Plugs, Receptacles, and Connectors for Three-Phase (60 Hz only)

Table 5.7 lists the current rating and AC power cable plug, receptacle, and connector requirements (part number or equivalent) for three-phase 60-Hz subsystems. In a three-phase 9980V subsystem the controller frame (DKC) receives its AC input power from the first disk array frame (DKU) via internal cabling, so the subsystem will not require any customer outlets for the controller frame. The user must supply all power receptacles and connectors for 60-Hz subsystems. Russell & Stoll (R&S) type connectors (or Hubbell or Leviton) are recommended for 60-Hz systems.

Note: Each 9980V disk array frame requires two power connections for 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.

Note: If you are planning to provide power using 3/4-inch diameter flexible conduit, consider using box-type receptacle connectors. In some cases, inline connectors may be used with 3/4-inch flexible conduit if the appropriate adapter is available from the connector vendor.

Table 5.7 Current Rating, Power Plug, Receptacle, and Connector for Three-Phase 9900V

Item	9970V	9980V DKC	9980V DKU
Hitachi Base Unit	DKC465I-5	DKC460I-5	DKU455I-18
Circuit Current Rating	30 A	(from DKU)	60 A
Hitachi Feature(s) Required	DKC-F465I-3PS DKC-F465I-3UC	DKC-F460I-3PS	DKU-F455I-3PS DKU-F455I-3UC
60-Hz Power Plug (or equiv.) Included with the product.	R&S 3760PDG, or DDK 115J-AP8508	N/A	R&S RS460P9W
Box-Type Receptacle (or equiv.) (not provided)	R&S 3754	N/A	R&S RS460R9W, or Hubbell HBL460R9W, or Leviton 460R9W
Back Box for Receptacle Note: Use R&S back box with R&S receptacle, Hubbell back box with Hubbell receptacle, etc.	included	N/A	R&S JB6-B125 plus R&S AA6L 20-degree-angle adapter (or Hubbell BB601W, Leviton BX60-V)
Inline Connector (or equiv.) (not provided)	R&S 3934	N/A	R&S RS460C9W (or Hubbell HBL460C9W, or Leviton 460C9W)

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

5.4.4 Input Voltage Tolerances for Three-Phase

Table 5.8 lists the input voltage tolerances for the three-phase 9980V. 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 5.8 Input Voltage Specifications for Three-Phase AC Input

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: For 9980V DKU, distribution board with circuit breaker or equivalent is rated 60 amps. For 9970V DKC, distribution board with circuit breaker or equivalent is rated 30 amps. This unit does not apply to IT power system.

5.4.5 Cable Dimensions for 50-Hz Three-Phase Subsystems

Table 5.9 and Figure 5.14 show the data required for 50-Hz three-phase cable installations.

Table 5.9 Cable Dimensions for 50-Hz Three-Phase Subsystems

Model	Power Cable			Terminal		Screw Size
	Outer Sheath Overall Diameter	Insulator Outer Diameter	Electric Wire Cross-Section Area	"Internal Diameter"	"External Diameter"	
	A	B	C	D	E	
DKU-F455I-3EC	24.0-30.5 mm	6.6 mm	10.0 mm ²	6.4 mm	12.0 mm	M6
DKC-F465I-3EC	18.0-24.5 mm	5.2 mm	6.0 mm ²	6.4 mm	12.0 mm	M6

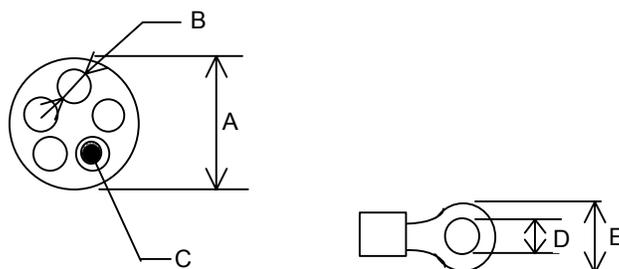


Figure 5.14 Cable Dimensions for 50-Hz Three-Phase Subsystems

5.5 Electrical Specifications and Requirements for Single-Phase Subsystems

5.5.1 Power Plugs for Single-Phase

- Figure 5.15 illustrates the power plugs for a single-phase 9980V controller (Europe).
- Figure 5.16 illustrates the power plugs for a single-phase 9980V controller (USA).
- Figure 5.17 illustrates the power plugs for a three-phase 9980V disk array unit (Europe).
- Figure 5.18 illustrates the power plugs for a three-phase 9980V disk array unit (USA).
- Figure 5.19 illustrates the power plugs for a three-phase 9970V subsystem (Europe).
- Figure 5.20 illustrates the power plugs for a three-phase 9970V subsystem (USA).

For information on UPS configurations, please contact your Hitachi Data Systems team.

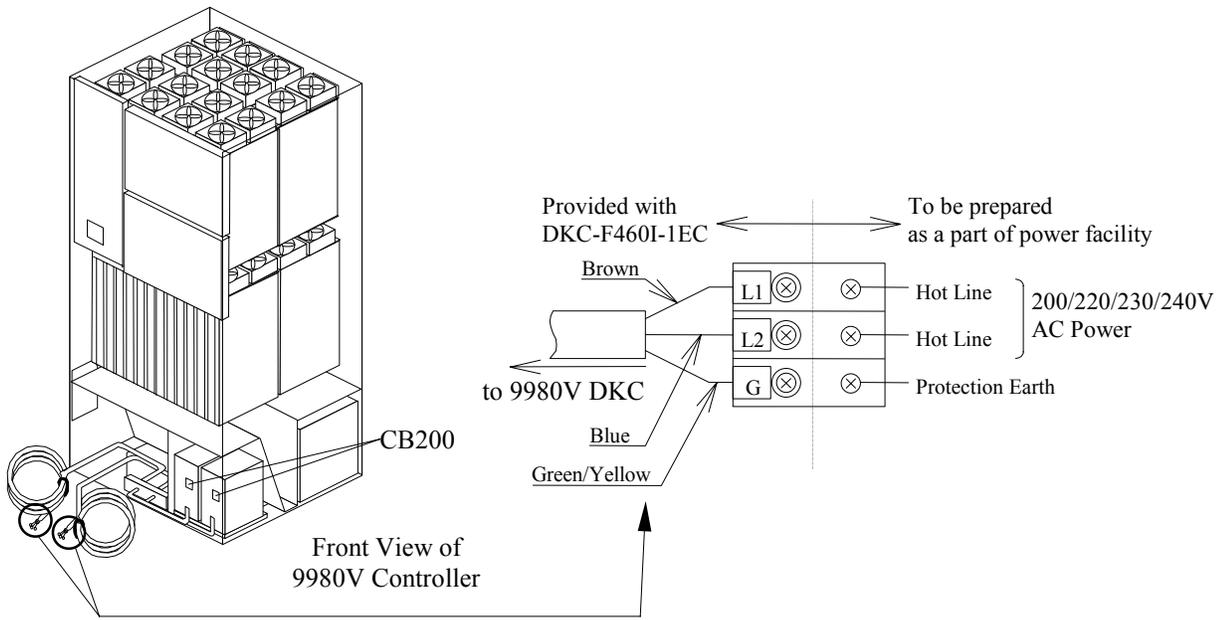


Figure 5.15 Power Plugs for Single-Phase 9980V Controller (Europe)

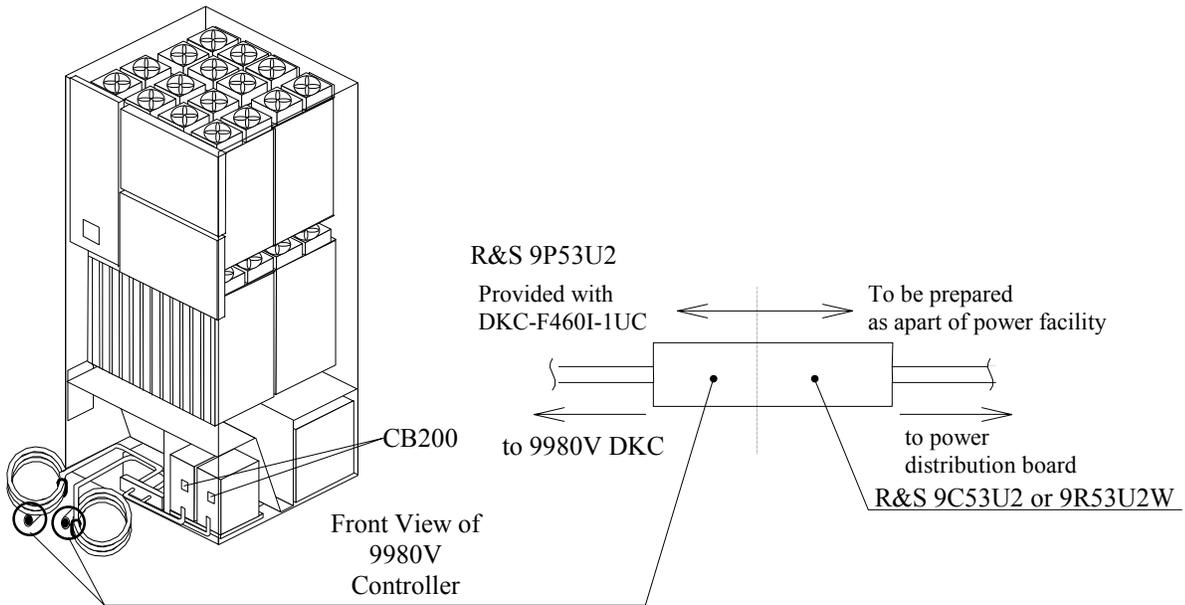


Figure 5.16 Power Plugs for Single-Phase 9980V Controller (USA)

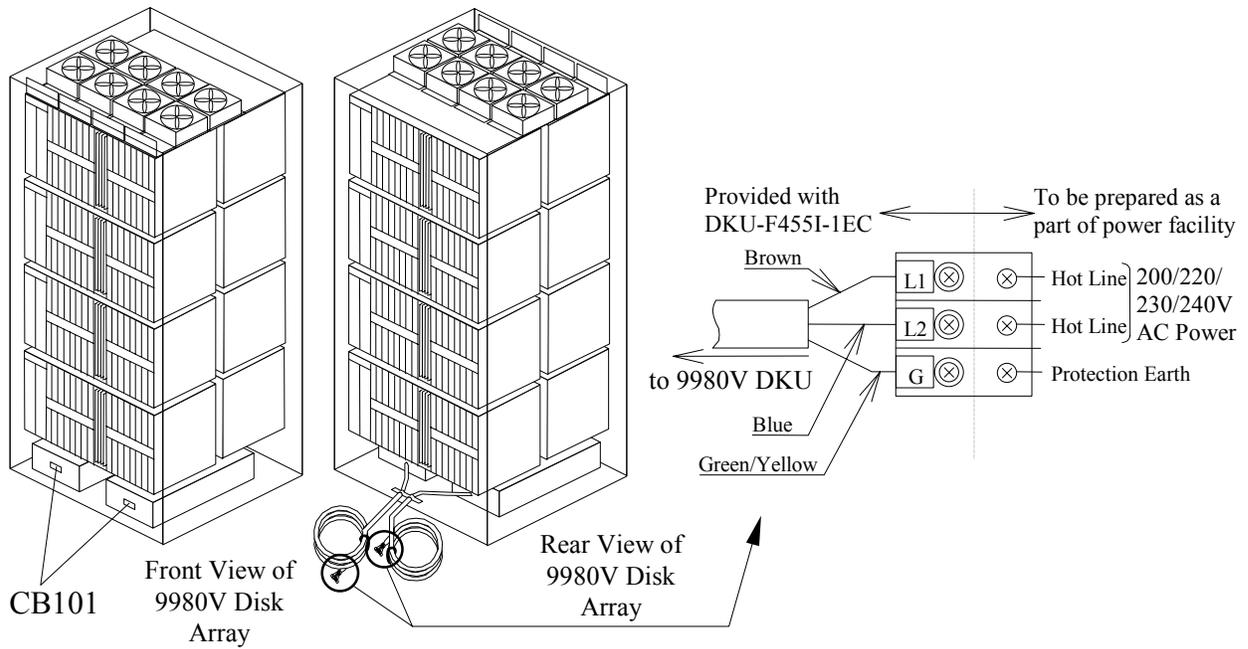


Figure 5.17 Power Plugs for a Single-Phase 9980V Disk Array Unit (Europe)

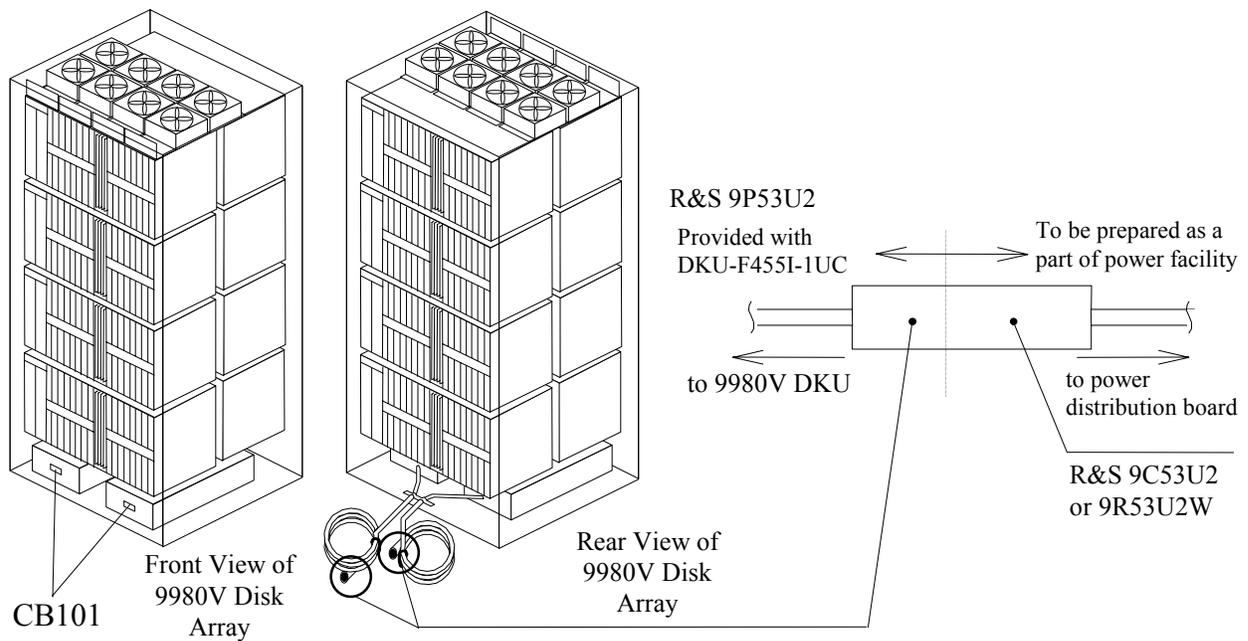


Figure 5.18 Power Plugs for a Single-Phase 9980V Disk Array Unit (USA)

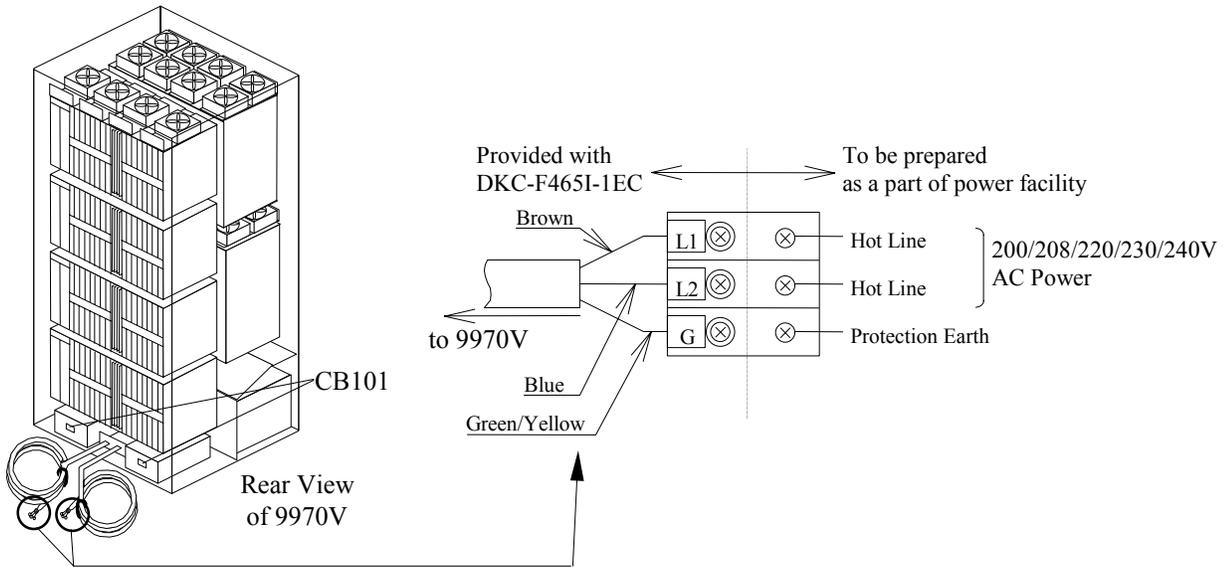


Figure 5.19 Power Plugs for a Single-Phase 9970V Subsystem (Europe)

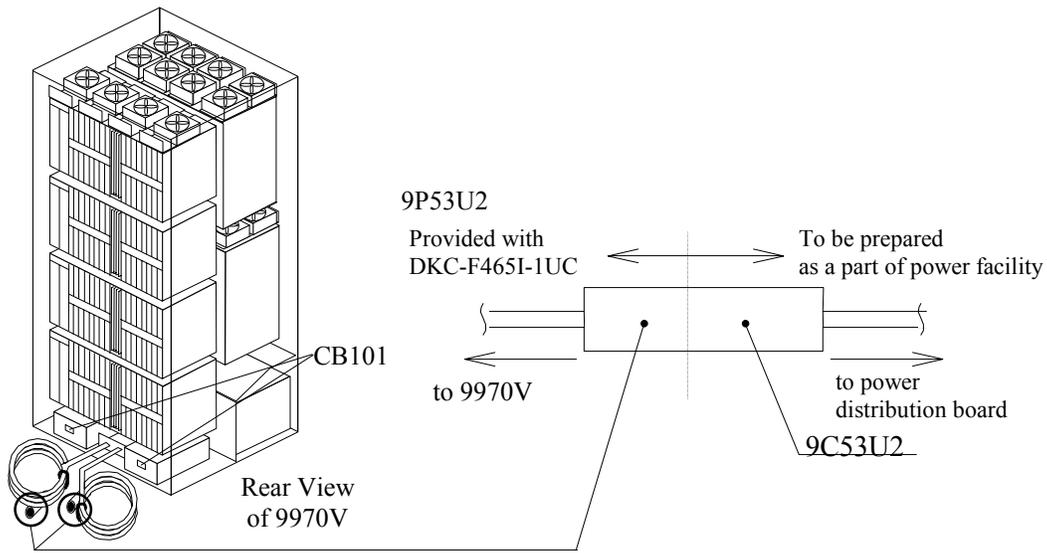


Figure 5.20 Power Plugs for a Single-Phase 9970V Subsystem (USA)

5.5.2 Features for Single-Phase

Table 5.10 lists the features for single-phase 9980V and 9970V subsystems.

Table 5.10 9900V Single-Phase Features

Model	Frame	Feature Number	Description
9980V	Controller	DKC-F460I-PS	AC Box Kit for Single Phase
		DKC-F460I-EC	Power Cable Kit for Single Phase, 50 Hz for Europe
		DKC-F460I-UC	Power Cable Kit for Single Phase, 60 Hz for USA
	Disk Array	DKU-F455I-1PS	AC Box Kit for Single Phase
		DKU-F455I-1EC	Power Cable Kit for Single Phase, 50 Hz for Europe
		DKU-F455I-1UC	Power Cable Kit for Single Phase, 60 Hz for USA
9970V	The 9970V consists of a single frame.	DKC-F465I-1EC	Power Cable Kit for Single Phase, 50 Hz for Europe
		DKC-F465I-1UC	Power Cable Kit for Single Phase, 60 Hz for USA

5.5.3 Current Rating, Power Plugs, Receptacles, and Connectors for Single-Phase (60 Hz only)

Table 5.11 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 (DKU) and controller frame (DKC) has two input power cables with R&S FS 3720 plugs. The user must supply the outlets for the plugs.

Note: Each 9980V disk array frame requires two power connections for 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.

Note: If you are planning to provide power using 3/4-inch diameter flexible conduit, consider using box-type receptacle connectors. In some cases, inline connectors may be used with 3/4-inch flexible conduit if the appropriate adapter is available from the connector vendor.

Table 5.11 Current Rating, Power Plug, Receptacle, and Connector for Single-Phase 9900V

Item	9970V	9980V DKC	9980V DKU
Hitachi Base Unit	DKC465I-5	DKC460I-5	DKU455I-18
Circuit Current Rating	50 A	40 A	50 A
Hitachi Feature(s) Required	DKC-F465I-1PS DKC-F465I-1UC	DKC-F460I-1PS DKC-F460I-1UC	DKU-F455I-1PS DKU-F455I-1UC
60-Hz Power Plug (or equiv.) Included with the product.	R&S 9P53U2	R&S 9P53U2	R&S 9P53U2
Box-Type Receptacle (or equiv.) (not provided)	R&S 9R53U2W	R&S 9R53U2W	R&S 9R53U2W
Back Box for Receptacle	R&S 3781A	R&S 3781A	R&S 3781A
Inline Connector (or equiv.) (not provided)	R&S 9C53U2	R&S 9C53U2	R&S 9C53U2

5.5.4 Input Voltage Tolerances for Single-Phase

Table 5.12 lists the input voltage tolerances for the single-phase 9980V and 9970V subsystems. 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 5.12 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: For 9980V DKC, distribution board with circuit breaker or equivalent is rated 40 amps. For 9980V DKU, distribution board with circuit breaker or equivalent is rated 50 amps. For 9970V DKC, distribution board with circuit breaker or equivalent is rated 50 amps.

5.5.5 Cable Dimensions for 50-Hz Single-Phase Subsystems

Table 5.13 and Figure 5.21 show the data required for 50-Hz single-phase cable installations.

Table 5.13 Cable Dimensions for 50-Hz Single-Phase Subsystems

Model	Power Cable			Terminal		Screw Size
	Outer Sheath Overall Diameter	Insulator Outer Diameter	Electric Wire Cross-Section Area	"Internal Diameter"	"External Diameter"	
	A	B	C	D	E	
DKU-F455I-1EC	20.0-25.5 mm	6.6 mm	10.0 mm ²	6.4 mm	12.0 mm	M6
DKC-F460I-1EC	20.0-25.5 mm	6.6 mm	10.0 mm ²	6.4 mm	12.0 mm	M6
DKC-F465I-1EC	20.0-25.5 mm	6.6 mm	10.0 mm ²	6.4 mm	12.0 mm	M6

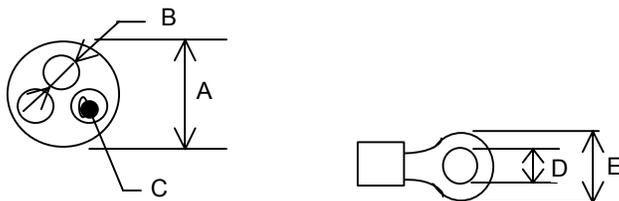


Figure 5.21 Cable Dimensions for 50-Hz Single-Phase Subsystems

5.6 Cable Requirements

Table 5.14 lists the cables required for the 9980V control frame. These cables must be ordered separately, and the quantity depends on the type and number of channels and ports. ExSA™ (ESCON®), FICON™, fibre-channel cables are available from Hitachi Data Systems.

Table 5.14 Cable Requirements

Cable	Function/Description
PCI cable	Connects 9900V to CPU power control interface.
FICON™ interface cables	Connects mainframe host systems, channel extenders, or FICON™ directors to 9900V 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 9900V 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-system host to 9900V ports. Fibre cable types are 50 / 125 micron or 62.5 / 125 micron multimode. SC-type (standard) connector is required for 1 Gbps port. LC-type (little) connector is required for 2 Gbps port.
Phone cable with RJ11 connector	Connects phone line to 9900V SVP for Hi-Track®.
10/100 BaseT (Cat 5) cable with RJ45 connector	Connects Remote Console PC to 9900V subsystem. Can also be used for connecting multiple 9900Vs together (daisy-chain).
10Base2 cable (RG58) with BNC connector	Connects Remote Console PC to 9900V subsystem, and allows connection to multiple 9900Vs (up to 8) without using a hub. Requires a transceiver.

5.6.1 Device Interface Cable

Figure 5.22 shows the relationship between the 9980V subsystem layout and the device interface cable options.

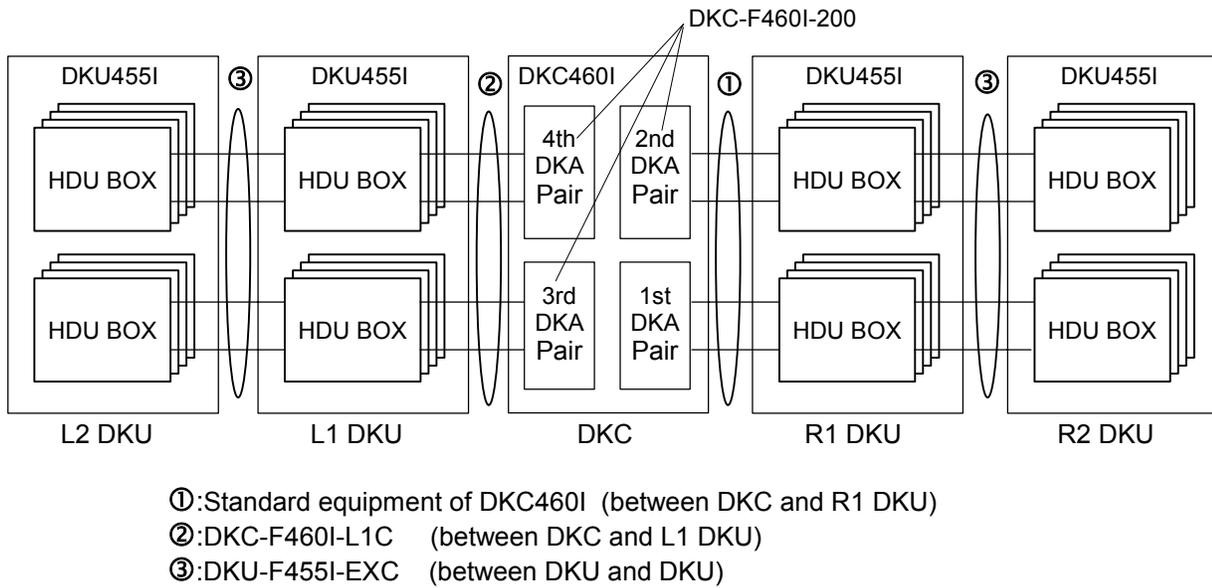


Figure 5.22 9980V Subsystem Layout and Device Interface Cable Options

5.7 Channel Specifications and Requirements

Table 5.15 lists the specifications for each 9900V open-systems channel option. Table 5.16 lists the specifications for each 9900V mainframe channel option. Each 9980V subsystem supports up to eight channel adapters and 32 interface ports. Each 9970V subsystem supports up to six channel adapters and 24 interface ports. Each channel adapter feature consists of a pair of cards.

Note: The 8HSE port is applicable for Hitachi TrueCopy - S/390®.

Note: Additional power supply is needed on the 9980V if the total number of channel adapter options and disk adapter options is four or more.

Table 5.15 Open-Systems Channel Specifications

Item	Fibre Short Wave 2 Gbps 8-port	Fibre Long Wave 2 Gbps 8-port
Option name	8HSE	8HLE
Host interface	FCP	FCP
Data transfer rate (MB/s)	100 / 200	100 / 200
Number of options installed	1 / 2 / 3 / 4	1 / 2 / 3 / 4
Number of PCBs installed	2 / 4 / 6 / 8	2 / 4 / 6 / 8
Number of ports / subsystem	8 / 16 / 24 / 32	8 / 16 / 24 / 32
Maximum cable length	500 m / 300 m *2	10 km

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

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

Table 5.16 Mainframe Channel Specifications

Item	Serial 8-port	Mainframe Fibre	
		Short-Wave 8-Port	Long-Wave 8-Port
Option name	8S	8MS	8ML
Host interface	ESCON®	FICON™	FICON™
Data transfer rate (MB/s)	17	100	100
Number of options installed	1 / 2 / 3 / 4	1 / 2 / 3 / 4	1 / 2 / 3 / 4
Number of PCBs installed	2 / 4 / 6 / 8	2 / 4 / 6 / 8	2 / 4 / 6 / 8
Number of ports / subsystem	8 / 16 / 24 / 32	8 / 16 / 24 / 32	8 / 16 / 24 / 32
Maximum cable length	3 km	500 m *1	10 km

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

5.8 Environmental Specifications and Requirements

5.8.1 Temperature and Humidity Requirements

Table 5.17 lists the temperature and humidity requirements for the 9900V 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. Non-operating condition 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 or around the drive should be observed under any conditions.

5.8.2 Power Consumption and Heat Output Specifications

Table 5.18 lists the power consumption and heat output parameters for the 9980V DKC components. Table 5.20 lists the power consumption and heat output parameters for the 9980V DKU components. Table 5.20 lists the power consumption and heat output parameters for the 9970V components. These data generally apply to both 60-Hz and 50-Hz subsystems.

Table 5.18 9980V DKC Component Power and Heat Output Specifications

Model number	Weight (kg)	Heat Output (kW)	Power Consumption (kVA)	Dimensions (mm)			Air Flow (m ³ /min.)
				Width	Depth	Height	
DKC460I-5	466.0	0.73	0.77	782 *1	800	1860	18
DKC-F460I-1PS	10.0	—	—	—	—	—	—
DKC-F460I-1EC	7.5	—	—	—	—	—	—
DKC-F460I-1UC	10.5	—	—	—	—	—	—
DKC-F460I-3PS	7.0	—	—	—	—	—	—
DKC-F460I-200	3.6	0.21	0.22	—	—	—	—
DKC-F460I-2048	0.2	0.007	0.008	—	—	—	—
DKC-F460I-S512	0.05	0.004	0.004	—	—	—	—
DKC-F460I-8S	4.2	0.28	0.29	—	—	—	—
DKC-F460I-8MS	4.2	0.31	0.33	—	—	—	—
DKC-F460I-8ML	4.2	0.31	0.33	—	—	—	—
DKC-F460I-8HSE	4.2	0.21	0.22	—	—	—	—
DKC-F460I-8HLE	4.2	0.21	0.22	—	—	—	—
DKC-F460I-SVP	5.5	0.07	0.07	—	—	—	—
DKC-F460I-UPS	1.2	0.01	0.01	—	—	—	—
DKC-F460I-80	8.8	—	—	—	—	—	—
DKC-F460I-41	5.6	0.06	0.06	—	—	—	—
DKC-F460I-42	42.2	0.09	0.10	—	—	—	—
DKC-F460I-L1C	3.5	—	—	—	—	—	—
DKC-F460I-U405R	20.0	0.06	0.07	—	—	—	—
DKC-F460I-U405L	19.0	—	—	—	—	—	—
DKC-F460I-U405	15.0	—	—	—	—	—	—
DKC-F460I-18	0.6	0.002	0.003	—	—	—	—

Note. *1: This includes the thickness of side covers (16 mm × 2).

Table 5.19 9980V DKU Component Power and Heat Output Specifications

Model number	Weight (kg)	Heat Output (kW)	Power Consumption (kVA)	Dimensions (mm)			Air Flow (m ³ /min.)
				Width	Depth	Height	
DKU455I-18	460.0	0.51	0.54	750	800	1860	32
DKU-F455I-1PS	20.0	—	—	—	—	—	—
DKU-F455I-1EC	7.0	—	—	—	—	—	—
DKU-F455I-1UC	10.0	—	—	—	—	—	—
DKU-F455I-3PS	20.0	—	—	—	—	—	—
DKU-F455I-3EC	11.0	—	—	—	—	—	—
DKU-F455I-3UC	16.0	—	—	—	—	—	—
DKU-F455I-EXC	3.0	—	—	—	—	—	—
DKU-F455I-36K1	1.0	0.022	0.024	—	—	—	—
DKU-F455I-36K4	4.0	0.088	0.096	—	—	—	—
DKU-F455I-72J1	1.0	0.023	0.025	—	—	—	—
DKU-F455I-72J4	4.0	0.092	0.100	—	—	—	—

Table 5.20 9970V Component Power and Heat Output Specifications

Model number	Weight (kg)	Heat Output (kW)	Power Consumption (kVA)	Dimensions (mm)			Air Flow (m ³ /min.)
				Width	Depth	Height	
DKC465I-5	586.0	0.52	0.55	782 *3	800	1860	25
DKC-F465I-1PS	15.0	—	—	—	—	—	—
DKC-F465I-1EC	7.5	—	—	—	—	—	—
DKC-F465I-1UC	10.5	—	—	—	—	—	—
DKC-F465I-3PS	12.0	—	—	—	—	—	—
DKC-F465I-3EC	6.0	—	—	—	—	—	—
DKC-F465I-3UC	10.0	—	—	—	—	—	—
DKC-F465I-FSW	2.0	0.20	0.21	—	—	—	—
DKC-F465I-FSW2	2.0	0.22	0.23	—	—	—	—
DKC-F465I-100	3.4	0.13	0.13	—	—	—	—
DKC-F460I-200 *1	3.6	0.21	0.22	—	—	—	—
DKC-F460I-2048 *1	0.2	0.007	0.008	—	—	—	—
DKC-F460I-S512 *1	0.05	0.004	0.004	—	—	—	—
DKC-F460I-SVP *1	14.5	0.07	0.07	—	—	—	—
DKC-F460I-UPS *1	1.2	0.01	0.01	—	—	—	—
DKC-F460I-8S *1	4.2	0.28	0.29	—	—	—	—
DKC-F460I-8MS *1	4.2	0.31	0.33	—	—	—	—
DKC-F460I-8ML *1	4.2	0.31	0.33	—	—	—	—
DKC-F460I-8HSE *1	4.2	0.21	0.22	—	—	—	—
DKC-F460I-8HLE *1	4.2	0.21	0.22	—	—	—	—
DKC-F460I-18 *1	0.6	0.002	0.003	—	—	—	—
DKU-F455I-36K1 *2	1.0	0.022	0.024	—	—	—	—
DKU-F455I-36K4 *2	4.0	0.088	0.096	—	—	—	—
DKU-F455I-72J1 *2	1.0	0.023	0.025	—	—	—	—
DKU-F455I-72J4 *2	4.0	0.092	0.100	—	—	—	—

Note. *1: They are common to the options installed in DKC460I-5.

Note. *2: They are common to the options installed in DKU455I-18.

Note. *3: This includes the thickness of side covers (16 mm × 2).

5.8.3 Loudness

The acoustic emission values [loudness in dB(A)] for a maximum 9980V subsystem configuration (one DKC, four DKUs) are:

Front/rear = 60 dB(A)

Both sides = 60 dB(A)

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

5.8.4 Air Flow Requirements

The 9900V 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 top, 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 5.21. For subsystems located at elevations from 3000 to 7000 feet (900 to 2100 meters) above sea level, decrease the maximum air temperature by two degrees for each 1000 feet (300 meters) above 3000 feet (900 meters).

Note: The air flow requirements for 9900V are greater than those for the 9900 subsystem.

Table 5.21 Internal Air Flow

Subsystem Type	Subsystem Configuration	Air Flow (m ³ /min)	Air Flow (ft ³ /min)
9980V	Controller Frame (all configurations)	18	636
	Array Frame (all configurations)	32	1131
9970V	All configurations	25	883

5.8.5 Vibration and Shock Tolerances

Table 5.22 lists the vibration and shock tolerance data for the 9900V subsystem. The 9900V can be subjected to vibration and shock up to these limits and still perform normally. The user should consider these requirements if installing the 9900V near large generators located on the floor above or below the 9900V subsystem. Generators or any other source of vibration, if not insulated or shock-mounted, can cause excessive vibration that may affect the subsystem.

Table 5.22 Vibration and Shock Tolerances

Condition	Operating	Nonoperating	Shipping or Storage
Vibration *1	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. *2 At four most severe resonance frequencies between 5 and 200 Hz
Shock	—	8 g, 15 ms	Horizontal: Incline Impact 4 ft (1.22 m) *3 Vertical: Rotational Edge 0.3 ft (0.1 m) *4

*1. The vibration specifications apply to all three axes.

*2. See ASTM D999-91 Standard Methods for Vibration Testing of Shipping Containers.

*3. See ASTM D5277-92 Standard Test Methods for Performing Programmed Horizontal Impacts Using an Inclined Impact Tester.

*4. See ASTM D1083-91 Standard Test Methods for Mechanical Handling of Unitized Loads and Large Shipping Cases and Crates.

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

Chapter 6 Troubleshooting

6.1 Troubleshooting

The Hitachi Lightning 9900™ V Series 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 9900V 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 6.2).

Table 6.1 Troubleshooting

Error Condition	Recommended Action
Error message displayed.	Determine the type of error (refer to the SIM codes in section 6.2). 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.	Call the Hitachi Data Systems Support Center for assistance. WARNING: Do not open the 9900V control frame or touch any of the controls.
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.
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 9900V 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) and/or call the Hitachi Data Systems Support Center.
Emergency (fire, earthquake, flood, etc.)	Pull the emergency power-off (EPO) switch (see section 4.6.1). 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 (call the Hitachi Data Systems Support Center for assistance), lower the room temperature to the specified operating range, and power on the subsystem (call the Hitachi Data Systems Support Center for assistance). If the area temperature is not the obvious cause of the alarm, call the Hitachi Data Systems Support Center for assistance.

6.2 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 9900V Remote Console configuration information saved on diskette using the FD Dump tool (see the *Hitachi Lightning 9900™ V Series Hitachi Remote Console - Storage Navigator 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, and
- The service information messages (SIMs) displayed on the Remote Console and the reference codes and severity levels of the recent 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

6.3 Service Information Messages (SIMs)

The 9900V 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. **Note:** For assistance with SIMs, please call the Hitachi Data Systems Support Center.

SIMs can be generated by the CHIP and ACP microprocessors and by the SVP. All SIMs generated by the 9900V are stored on the SVP for use by Hitachi Data Systems personnel, logged in the SYS1.LOGREC dataset of the S/390® host system, displayed by the Remote Console software, and reported via SNMP to the open-system host. The SIM display on the 9900V Remote Console PC enables users to remotely view the SIMs reported by the attached 9900V subsystems. Each time a SIM is generated, the amber Message LED on the 9900V control panel turns on. 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 6.2).

Figure 6.1 illustrates a typical 32-byte SIM from the 9900V 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.

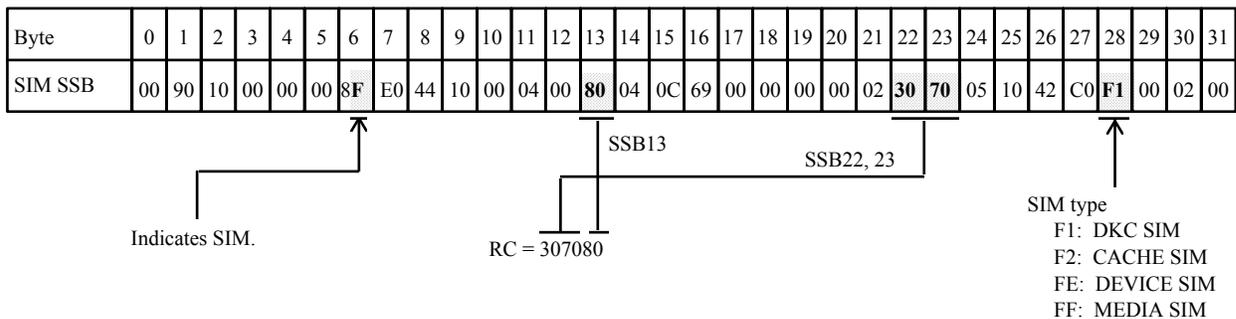


Figure 6.1 Typical 9900V SIM Showing Reference Code and SIM Type

Appendix A 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
CHA	channel adapter
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
DKA	disk adapter
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, Gb/s	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
HDLM	Hitachi Dynamic Link Manager™
HMBR	Hitachi Multiplatform Backup/Restore
HPAV	Hitachi Parallel Access Volume
HRX	Hitachi RapidXchange
HSN	Hierarchical Star Network
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
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
UPS	uninterruptable power supply
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)

Appendix B Unit Conversions

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

Table B.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)
Degrees Fahrenheit (°F)	subtract 32, then multiply result by 0.555556	Degrees centigrade (°C) °C = (°F - 32) × 0.555556
Degrees centigrade (°C)	multiply by 1.8, then add 32 to result	Degrees Fahrenheit (°F) °F = (°C × 1.8) + 32
Degrees Fahrenheit per hour (°F/hour)	0.555555	Degrees centigrade per hour (°C/hour)
Degrees centigrade per hour (°C/hour)	1.8	Degrees Fahrenheit per hour (°F/hour)

