HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Lightning 9900™

User and Reference Guide

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.

Hitachi GRAPH-Track™, Hitachi Freedom Storage™, and Lightning 9900™ are trademarks of Hitachi Data Systems Corporation.

AIX®, ESCON®, and S/390® are registered trademarks of International Business Machines Corporation.

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

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

Notice of Export Controls

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

Document Revision Level

Revision	Date	Description
MK-90RD008-0	July 2000	Initial Release
MK-90RD008-1	November 2000	Revision 1, supersedes and replaces MK-90RD008-0

Source Document Revision Level

The following source documents were used to produce this 9900 user guide:

- *Hitachi DKC410 Disk Subsystem Maintenance Manual*, revision 3.
- *Hitachi DKC415 Disk Subsystem Maintenance Manual*, revision 1.

Changes in this Revision

The major change in this revision is the addition of the configuration information for the 9910 model of the Lightning 9900™ subsystem family.

Preface

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

This document assumes that:

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

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

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

Contents

List of Figures

List of Tables

Chapter 1 Overview of the Lightning 9900™ Subsystem

1.1 Key Features of the Lightning 9900™ Subsystem

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

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

- **Instant access to data around the clock:**
  - 100 percent data availability guarantee
  - No single point of failure
  - Highly resilient multi-path fibre architecture
  - Fully redundant, hot-swappable components
  - Global dynamic hot sparing
  - Duplexed write cache with battery backup
  - Hi-Track® “call-home” maintenance system
  - Non-disruptive microcode updates
  - RAID-1+ and/or RAID-5 arrays within the same system
- **Unmatched performance and capacity:**
  - Industry’s only internal switched fabric architecture
  - Multiple point-to-point data and control paths
  - Up to 6.4-GB/sec internal system bandwidth
  - Fully addressable 32-GB data cache; separate control cache
  - Extremely fast and intelligent cache algorithms
  - Non-disruptive expansion to over 37 TB raw capacity
  - Simultaneous transfers from up to 32 separate hosts
  - High-throughput 10K RPM fibre channel, dual-active disk drives
- **Extensive connectivity and resource sharing:**
  - Concurrent operation of UNIX®-based, Windows® 2000, Windows NT®, Linux®, Novell® NetWare®, and S/390® hosts
  - Fibre-channel and ESCON® server connections
  - Optimized for SANs, fibre-channel switched, fibre-channel arbitrated loop, and point-to-point

1.1.1 Continuous Data Availability

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

1.1.2 Connectivity

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

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

When fibre-channel interfaces are used, the 9900 subsystem can provide up to 32 ports for attachment to UNIX®-based (e.g., AIX®, Solaris®, HP-UX®), Linux®-based (Red Hat® Linux®), and/or PC-server platforms (e.g., Windows NT®, Windows® 2000, NetWare®). Please contact Hitachi Data Systems for the latest information on platform support. 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 100 MB/sec. The 9900 subsystem supports fibre-channel arbitrated loop (FC-AL) and fabric fibre-channel topologies as well as high-availability (HA) fibre-channel configurations using hubs and switches.

1.1.3 S/390® Compatibility and Functionality

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

- Sequential data striping,
- Cache fast write (CFW) and DASD fast write (DFW),
- Enhanced dynamic cache management,
- Concurrent copy, and
- Transaction Processing Facility (TPF)/Multi-Path Locking Facility (MPLF) support.

1.1.4 Open-Systems Compatibility and Functionality

The Lightning 9900™ subsystem supports multiple concurrent attachment to a variety of UNIX®-based hosts (e.g., Solaris®, AIX®, HP-UX®, Tru64® UNIX®, SGI™ IRIX®) as well as PC server hosts (e.g., Windows NT®, Windows® 2000, Novell® NetWare®). The 9900 subsystem also supports the Red Hat Linux operating system. Please contact your Hitachi Data Systems account team for the most recent information on supported platforms and OS versions. The 9900 is compatible with most fibre host bus adapters (HBAs).

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

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

1.1.5 Program Products and Service Offerings

The Lightning 9900™ subsystem provides many advanced data management 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. The **CARE** 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.

- The Hitachi Remote Copy (HRC) and Hitachi Open Remote Copy (HORC) features (see sections 3.6.1 and 3.6.2) enables the user to perform remote copy operations between Hitachi RAID subsystems in different locations. Hitachi Remote Copy provides synchronous and asynchronous copy modes for both mainframe and open-system users.
- The Hitachi ShadowImage feature (see sections 3.6.3 and 3.6.4) allows the user to create internal copies of volumes for a wide variety of purposes including application testing and offline backup. When used in conjunction with Remote Copy, ShadowImage allows the user to maintain multiple copies of critical data at both the main and remote sites.
- The Command Control Interface (CCI) software (see section 3.6.5) enables users to perform Remote Copy and ShadowImage operations by issuing commands from the open-system host to the 9900 subsystem. The CCI software supports scripting and provides failover and mutual hot standby functionality in cooperation with the failover product on the host.
- The Hitachi Extended Remote Copy (HXRC) feature (see section 3.6.6) provides compatibility with the IBM Extended Remote Copy (XRC) host software function, which performs server-based asynchronous remote copy operations for mainframe LVIs.
- The Hitachi Online Data Migration (HODM) feature (see section 3.6.7) allows the rapid transfer of data from other disk subsystems onto the 9900 subsystem. HODM operations can be performed while applications are online using the data which is being transferred.
- The Dynamic Optimizer feature (see section 3.6.8) monitors volume activity and performs automatic relocation of volumes to optimize subsystem performance.
- The FlashAccess feature (see section 3.6.9) enables users to store specific high-usage data directly in cache memory to provide virtually immediate data availability.
- The LUN Manager feature (see section 3.6.10) enables users to configure the 9900 fibre ports for operational environments (e.g., arbitrated-loop (FC-AL) and fabric topologies, host failover support).
- The LU Size Expansion (LUSE) feature (see section 3.6.10) enables open-system users to create expanded LUs which can be up to 36 times larger than standard fixed-size LUs.
- The Virtual LVI/LUN feature (see section 3.6.11) enables users to configure custom-size LVIs and LUs which are smaller than standard-size devices.
- The Zone Allocation Manager (LUN Security) feature (see section 3.6.12) allows open-system users to restrict host access to LUs based on the host's World Wide Name (WWN).

- The LDEV Security feature (see section 3.6.13) allows mainframe users to restrict host access to LVIs based on node IDs and logical partition (LPAR) numbers.
- The Prioritized Port Control (PPC) feature (see section 3.6.14) 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.
- The Hitachi Multiplatform Resource Sharing (HMRS) feature allows the data stored on the 9900 to be shared by mainframe and open-system hosts. The HMRS functions include:
  - The Hitachi Multiplatform Data Exchange (HMDE) feature (see section 3.6.15) enables users to transfer data between S/390® and open-system platforms using the ExSA channels, which provides high-speed data transfer without requiring network communication links or tape.
  - The Hitachi Multiplatform Backup/Restore (HMBR) feature (see section 3.6.16) allows users to perform mainframe-based volume-level backup and restore operations on the open-system data stored on the multiplatform 9900 subsystem.
  - The HARBOR File-Level Backup/Restore feature (see section 3.6.17) enables users to perform mainframe-based file-level backup/restore operations on the open-system data stored on the multiplatform 9900 subsystem.
- The licensed Hitachi GRAPH-Track™ software (see section 3.6.18) provides detailed information on the I/O activity and hardware performance of the 9900 subsystem. Hitachi GRAPH-Track™ displays real-time and historical data in graphical format, including I/O statistics, cache statistics, and front-end and back-end microprocessor usage.

1.1.6 Subsystem Scalability

The architecture of the 9900 subsystem accommodates scalability to meet a wide range of capacity and performance requirements. The 9960 storage capacity can be increased from a minimum of 54 GB to a maximum of 37 TB of user data. The 9960 nonvolatile cache can be configured from 1 GB to 32 GB. All disk drive and cache upgrades can be performed without interrupting user access to data and with minimal impact on subsystem performance.

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

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

1.2 Reliability, Availability, and Serviceability

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

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

- **Full fault-tolerance**. The 9900 subsystem provides full fault-tolerance capability for all critical components. The disk drives are protected against error and failure by enhanced RAID technologies and dynamic scrubbing and sparing. The 9900 uses component and function redundancy to provide full fault-tolerance for all other subsystem components (microprocessors, control storage, power supplies, etc.). The 9900 has no active single point of component failure and is designed to provide continuous access to all user data.
- **Separate power supply systems**. Each storage cluster is powered by a separate set of power supplies. Each set can provide power for the entire subsystem in the unlikely event of power supply failure. The power supplies of each set can be connected across power boundaries, so that each set can continue to provide power if a power outage occurs. The 9900 can sustain the loss of multiple power supplies and still continue operation.
- **Dynamic scrubbing and sparing for disk drives**. The 9900 uses special diagnostic techniques and dynamic scrubbing to detect and correct disk errors. Dynamic sparing is invoked automatically if needed. The 9960 can be configured with up to sixteen spare disk drives, and any spare disk can back up any other disk of the same capacity, even if the failed disk and spare disk are in different array domains (attached to different ACP pairs).

- **Dynamic duplex cache**. The 9900 cache is divided into two equal segments on separate power boundaries. The 9900 places all write data in both cache segments with one internal write operation, so the data is always duplicated (duplexed) across power boundaries. If one copy of write data is defective or lost, the other copy is immediately destaged to disk. This duplex design ensures full data integrity in the event of a cache or power failure.
- **Remote copy features**. The Hitachi Remote Copy (HRC), Hitachi Open Remote Copy (HORC), and Hitachi Extended Remote Copy (HXRC) data movement features enable the user to set up and maintain duplicate copies of S/390® and open-system data over extended distances. In the event of a system failure or site disaster, the secondary copy of data can be invoked rapidly, allowing applications to be recovered with guaranteed data integrity.
- **Hi-Track®.** The Hi-Track® maintenance support tool monitors the operation of the 9900 subsystem at all times, collects hardware status and error data, and transmits this data via modem to the Hitachi Data Systems Support Center. The Hitachi Data Systems Support Center analyzes the data and implements corrective action when necessary. In the unlikely event of a component failure, Hi-Track® calls the Hitachi Data Systems Support Center immediately to report the failure without requiring any action on the part of the user. Hi-Track® enables most problems to be identified and fixed prior to actual failure, and the advanced redundancy features enable the subsystem to remain operational even if one or more components fail. ***Note:*** Hi-Track® does not have access to any user data stored on the 9900 subsystem. The Hi-Track® tool requires a dedicated RJ-11 analog phone line.
- **Nondisruptive service and upgrades**. All hardware upgrades can be performed nondisruptively during normal subsystem operation. All hardware subassemblies can be removed, serviced, repaired, and/or replaced nondisruptively during normal subsystem operation. All microcode upgrades can be performed during normal subsystem operations using the SVP or the alternate path facilities of the host.
- **Error Reporting**. The Lightning 9900™ subsystem reports service information messages (SIMs) to notify users of errors and service requirements. SIMs can also report normal operational changes, such as remote copy pair status change. The SIMs are logged on the 9900 service processor (SVP) and on the Remote Console PC, reported directly to the mainframe and open-system hosts, and reported to Hitachi Data Systems via Hi-Track®.

Chapter 2 Subsystem Architecture and Components

2.1 Overview

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

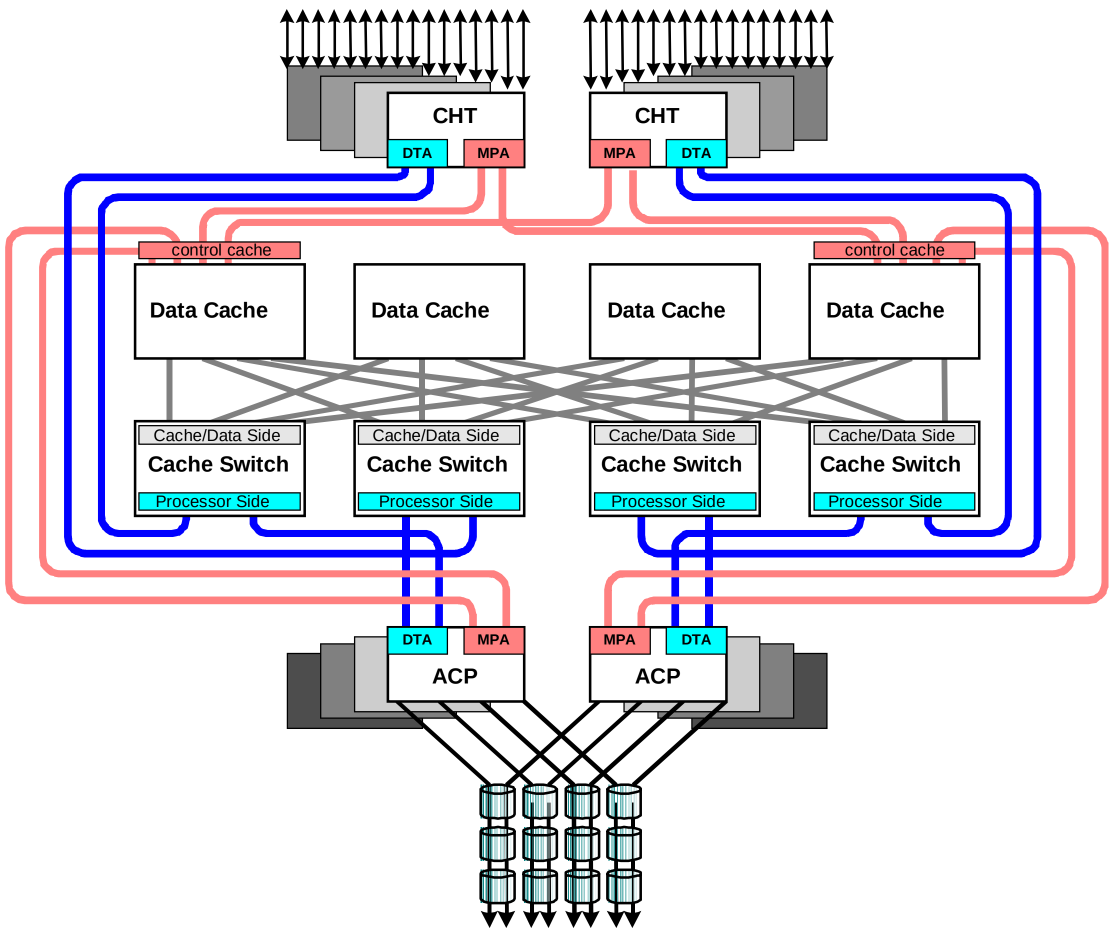

Figure 2.1 Lightning 9900™ HiStar Network (HSN) Architecture

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

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

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

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

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

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

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

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

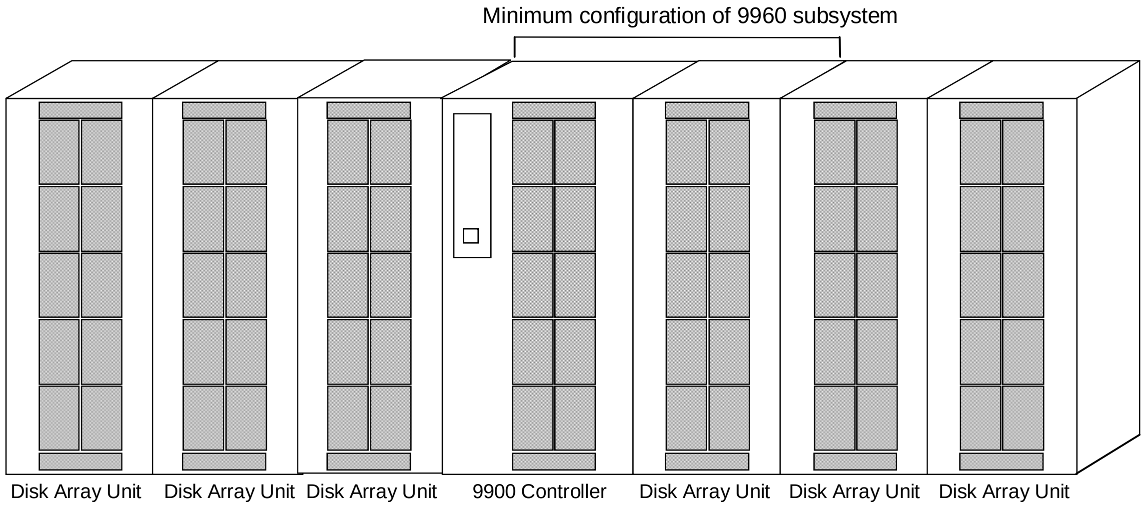

Figure 2.2 9960 Subsystem Frames

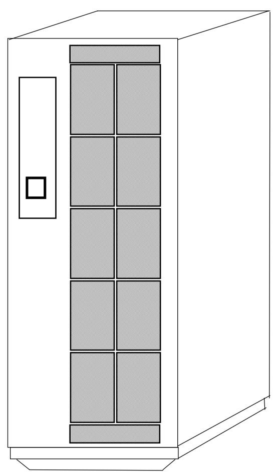

Figure 2.3 9910 Subsystem Frame

2.2 Components of the Controller Frame

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

- Storage clusters (see section 2.2.1),
- Nonvolatile duplex shared memory (see section 2.2.2),
- Nonvolatile duplex cache memory (see section 2.2.3),
- Multiple data and control paths (see section 2.2.4),
- Redundant power supplies (see section 2.2.5),
- CHIPs and channels (ExSA™ and/or fibre channel) (see section 2.2.6), 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, CHAs, CHIPs, ACPs, cache, control storage) needed to sustain processing within the subsystem. Both storage clusters should be connected to each host using an alternate path scheme, so that if one storage cluster fails, the other storage cluster can continue processing for the entire subsystem.

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

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

2.2.2 Nonvolatile Shared Memory

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

2.2.3 Nonvolatile Duplex Cache

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

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

2.2.4 Multiple Data and Control Paths

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

2.2.5 Redundant Power Supplies

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

2.2.6 Client-Host Interface Processors (CHIPs) and Channels

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

Table 2.1 CHIP and Channel Specifications

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

2.2.7 Channels

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

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

2.2.8 Array Control Processors (ACPs)

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

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

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

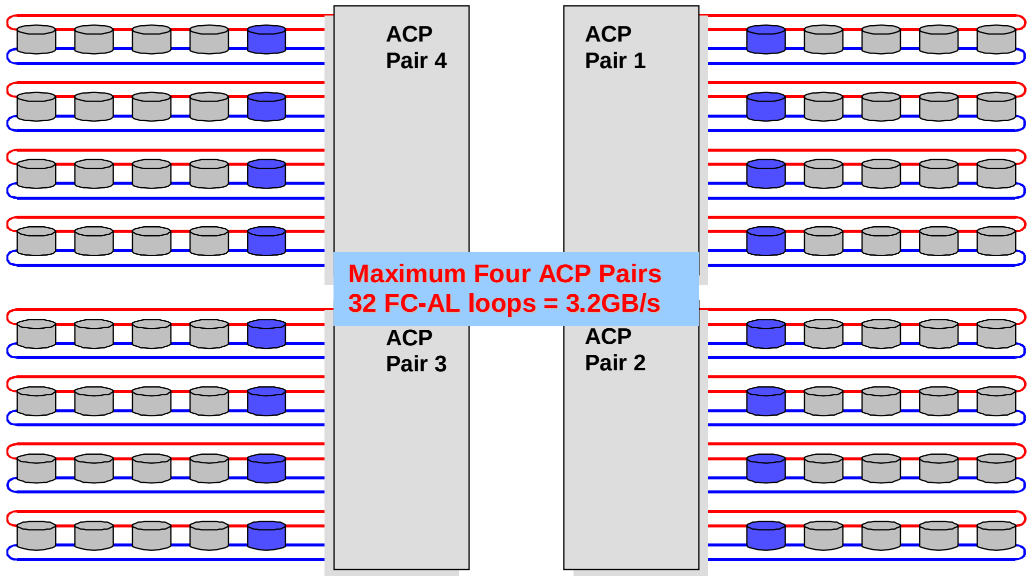

Figure 2.4 Conceptual ACP Array Domain

Table 2.2 ACP Specifications

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

Notes:

1. The entire ACP pair domain must be the same RAID type. The RAID type is set when the ACP pair is installed. RAID-1 and RAID-5 can coexist in the 9900 subsystem, but must be on separate ACP pairs.
2. 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.
3. 3390-3 and 3390-3R LVIs cannot be intermixed in the same 9900 subsystem.

2.3 Array Frame

The 9960 array frames contain the physical disk drives, including the disk array groups and the dynamic spare disk drives. Each array frame has dual AC power plugs, which should be attached to two different power sources or power panels. The 9960 can be configured with up to six array frames to provide a storage capacity of up to 37 TB. The 9910 subsystem combines the controller and disk array components in one physical frame.

The 9900 subsystem uses three-inch disk drives with fixed-block-architecture (FBA) format. The currently available disk drives have capacities of 18 GB, 47 GB, and 73 GB. Disk drives of varying capacities can be attached to the same ACP pair, but all drives in an array group must have the same capacity. Table 2.3 provides the disk drive specifications.

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

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

Table 2.3 Disk Drive Specifications

Parameter		Disk Drive Capacity 73 GB	47 GB	18 GB
Formatted capacity (GB)		72.914	47.191	18.463
Platter diameter		3 inches	3 inches	3 inches
Physical cylinders (user area)		15,154	12,027	12,027
Physical tracks per physical cylinder (user area)		24	23	9
Physical disk platters (user area)		12	12	5
Sector length (byte)		520(512)	520(512)	520(512)
Seek time (ms)	MIN.	0.5/0.7	0.5/0.7	0.5/0.7
(Read/Write)	MAX. AVE.	12.0/13.0 5.7/6.5	12.0/13.0 5.7/6.5	12.0/13.0 5.2/6.0
Revolution speed (rpm)		10,025	10,025	10,025
Mean latency time (ms)		2.99	2.99	2.99
Media transfer rate (MB/s)		33.60 to 56.64	30.20 to 45.60	30.20 to 45.60
Interface transfer rate (MB/s)		100	100	100

2.3.1 Disk Array Groups

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

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

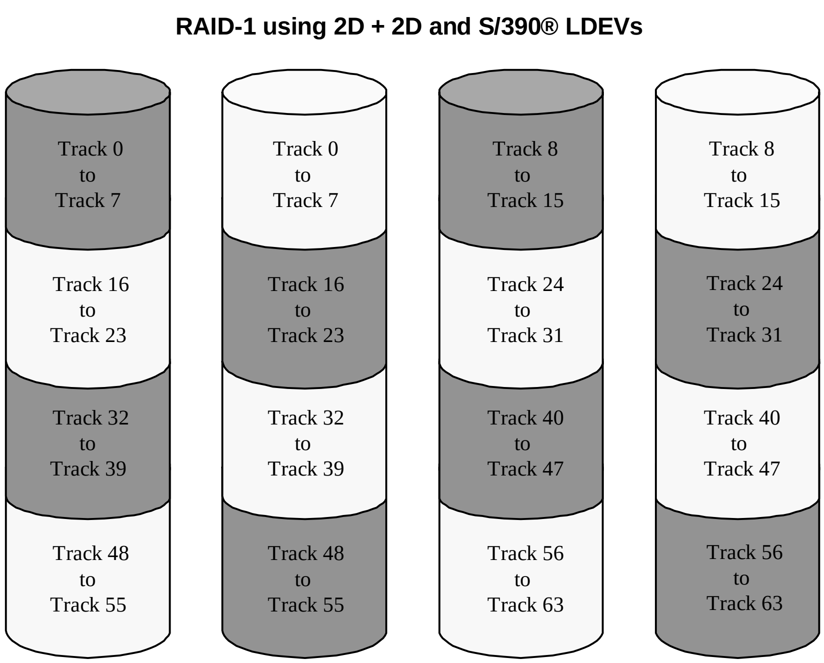

Figure 2.5 Sample RAID-1 Layout

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

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

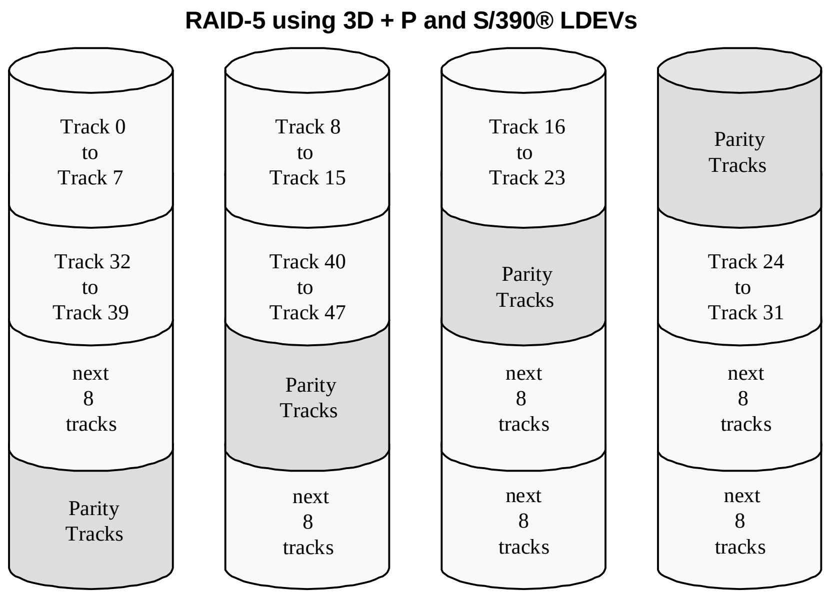

Figure 2.6 Sample RAID-5 Layout (Data Plus Parity Stripe)

2.3.2 Sequential Data Striping

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

2.4 Intermix Configurations

2.4.1 RAID-1 & RAID-5 Intermix

RAID technology provides full fault-tolerance capability for the disk drives of the 9900 subsystem. The 9900 supports RAID-1, RAID-5, and an intermix of RAID-1 and RAID-5 array groups. Figure 2.7 illustrates a RAID-1 and RAID-5 intermix configuration. The cache management algorithms (see section 3.3.1) enable the 9900 to stage up to one full RAID stripe of data into cache ahead of the current access to allow subsequent access to be satisfied from cache at host channel transfer speeds.

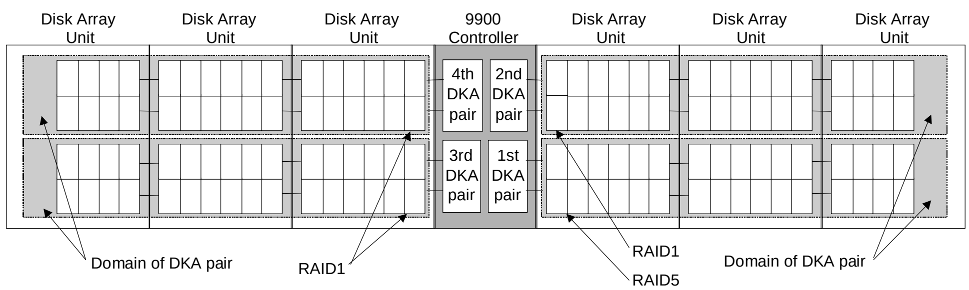

Figure 2.7 Sample RAID-1/RAID-5 Intermix

2.4.2 Hard Disk Drive Intermix

Array groups consisting of each type of hard disk drive can coexist anywhere in the 9900 subsystem. However, all hard disk drives in one array group must be of the same capacity and type. Figure 2.8 illustrates an intermix of hard disk drive types.

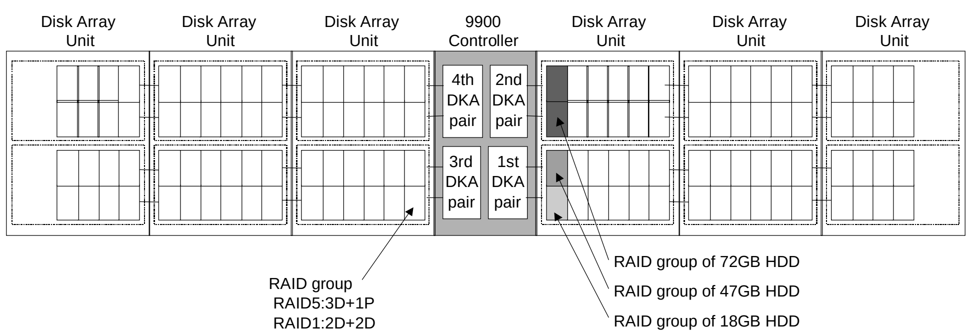

Figure 2.8 Sample Hard Disk Drive Intermix

2.4.3 Device Emulation Intermix

The 9900 subsystem supports an intermix of different device emulations (e.g., 3390-x LVIs, 3380-x LVIs, OPEN-x LUs) on the same ACP pair. Figure 2.9 illustrates an intermix of device emulation types. The only requirement is that the devices within each array group must have the same type of track geometry or format, as follows:

- 3390-1, -2, -3 or -9 can be intermixed within an array group.
- 3380-E, -J, or -K can be intermixed within an array group.
- OPEN-3, -8, and -9 can be intermixed within an array group.
- OPEN-K cannot be intermixed with other device types within an array group.

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

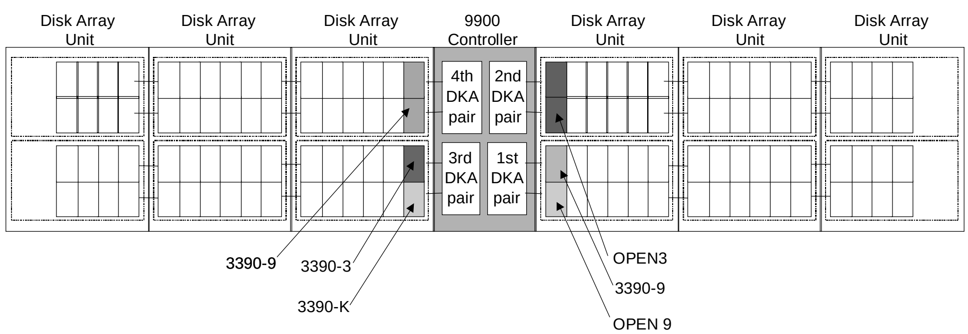

Figure 2.9 Sample Device Emulation Intermix

2.5 Service Processor (SVP)

The Lightning 9900™ subsystem includes a built-in laptop PC called the service processor (SVP). The SVP is integrated into the controller frame and can only be used by authorized Hitachi Data Systems personnel. The SVP enables the Hitachi Data Systems representative to configure, maintain, and upgrade the 9900 subsystem. The SVP also collects performance data for all key components of the 9900 subsystem to enable diagnostic testing and analysis. The Hitachi GRAPH-Track™ software product (see section 3.6.18) stores the SVP performance data on the Remote Console PC and allows users to view the data in graphical format and export the data for statistical analysis. ***Note:*** The SVP does not have access to any user data stored on the 9900 subsystem.

2.6 Remote Console PC

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

Chapter 3 Functional and Operational Characteristics

3.1 New 9900 Features and Capabilities

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

- Sixteen (16) logical control unit images.
- State-of-the-art hard disk drives of 18-GB, 47-GB, and 73-GB capacities.
- Up to 32 GB cache memory for the 9960 subsystem and 16 GB cache for the 9910.
- Up to 256 logical paths per ExSA™ (ESCON®) channel interface.
- Up to 4096 device addresses.
- Prioritized port control for open-system users.

3.2 I/O Operations

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

- **Read hit:** For a read I/O, when the requested data is already in cache, the operation is classified as a read hit. The CHIP searches the cache directory, determines that the data is in cache, and immediately transfers the data to the host at the channel transfer rate.
- **Read miss:** For a read I/O, when the requested data is not currently in cache, the operation is classified as a read miss. The CHIP searches the cache directory, determines that the data is not in cache, disconnects from the host, creates space in cache, updates the cache directory, and requests the data from the appropriate ACP pair. The ACP pair stages the appropriate amount of data into cache, depending on the type of read I/O (e.g., sequential).
- **Fast write:** All write I/Os to the 9900 subsystem are fast writes, because all write data is written to cache before being destaged to disk. The data is stored in two cache locations on separate power boundaries in the dynamic duplex cache (see section 3.3.2). As soon as the write I/O has been written to cache, the 9900 subsystem notifies the host that the I/O operation is complete, and then destages the data to disk.

3.3 Cache Management

3.3.1 Algorithms for Cache Control

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

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

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

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

3.3.2 Dynamic Duplex Cache

The duplex write cache is the area of cache that is dynamically allocated for write operations. The duplex write line (DWL) is the amount of duplex write cache expressed as a percentage of total cache. 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.

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

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

3.4.1 CU Images

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

3.4.2 Logical Volume Image (LVIs)

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

3.4.3 Logical Unit (LU) Type

The 9900 subsystem currently supports the following LU types: OPEN-3, OPEN-8, OPEN-K, OPEN-9, OPEN-E, OPEN-L, and OPEN-M (OPEN-L and OPEN-M are supported only for Windows® 2000 operations). Table 3.1 lists the capacities for each standard LU type. The 9900 also allows users to configure custom-size LUs which are smaller than standard LUs as well as size-expanded LUs which are larger than standard LUs. LU Size Expansion (LUSE) volumes can range in size from 3.748 (OPEN-K*2) to 524.448 GB (OPEN-E*36). Each LU is identified by fibre-channel port ID (target ID) and LU number (LUN) (see Figure 3.1). Each 9900 fibre-channel port supports addressing capabilities for up to 256 LUNs.

Table 3.1 Capacities of Standard LU Types

LU Type	Capacity (GB)
OPEN-K	1.874
OPEN-3	2.461
OPEN-8	7.347
OPEN-9	7.384
OPEN-E	14.567

Note: For information on the OPEN-L and OPEN-M LUs for Windows® 2000 operations, please contact Hitachi Data Systems.

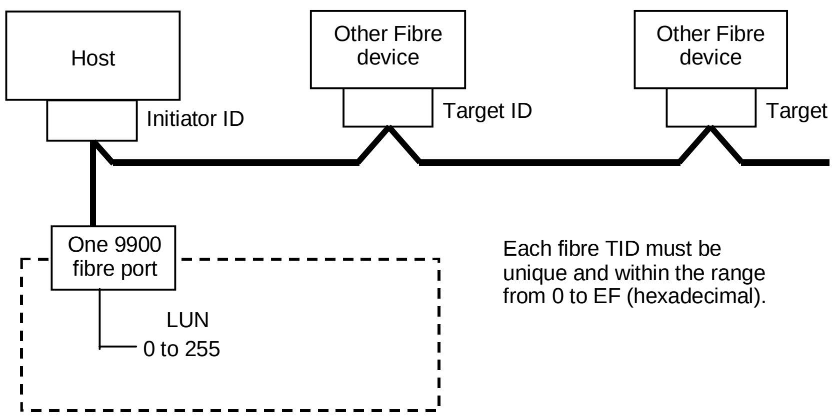

Figure 3.1 Fibre Target IDs and LUNs

3.5 Open Systems Features and Functions

The 9900 subsystem offers many features and functions specifically for the open-systems environment. The 9900 supports multi-initiator I/O configurations in which multiple host systems are attached to the same fibre-channel interface. The 9900 subsystem also supports important open-system functions such as fibre-channel arbitrated-loop (FC-AL) and fabric topologies, command tag queuing, multi-initiator I/O, and industry-standard middleware products which provide application and host failover, I/O path failover, and logical volume management functions. In addition, several program products and services are specifically for open systems. See section 3.6 for more information.

3.5.1 Failover and SNMP Support

The 9900 subsystem supports industry-standard products and functions which provide host and/or application failover, I/O path failover, and logical volume management (LVM) (e.g., HP® MC/ServiceGuard, VERITAS® FirstWatch®, HACMP, HAGEO, IRIS FailSafe™, SGI™ Advanced Cluster Environment (ACE), Microsoft Cluster Server, Novell® High Availability Server (NHAS), Novell® System Fault Tolerance (SFT) III). For the latest information on middleware releases, availability, and compatibility, please contact your Hitachi Data Systems account team.

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

3.5.2 Share-Everything Architecture

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

3.6 Data Management Features and Functions

The 9900 subsystem provides features and functions that increase data availability and improve data management. Table 3.2 lists the data management features that are currently available for the 9900 subsystem. Please review the appropriate user's guide for more details.

3.6.1 Hitachi Remote Copy (HRC)

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

For more information, please see *Hitachi Freedom Storage™ Lightning 9900™ HRC User and Reference Guide* (MK90RD009), or contact your Hitachi Data Systems account team.

3.6.2 Hitachi Open Remote Copy (HORC)

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

For more information, see *Hitachi Freedom Storage™ Lightning 9900™ Open Remote Copy (HORC) User's Guide* (MK90RD010), or contact your Hitachi Data Systems account team.

3.6.3 ShadowImage for S/390® Data (HMRCF)

The Hitachi ShadowImage feature, also called Hitachi Multiple RAID Coupling Feature (HMRCF), 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 ShadowImage copy back to the source LVI. When used in conjunction with either HRC or HXRC, ShadowImage enables users to maintain multiple copies of critical data at both primary and remote sites. ShadowImage also supports the Virtual LVI and FlashAccess features of the 9900 subsystem, ensuring that all user data can be duplicated by ShadowImage operations

For more information, please see *Hitachi Freedom Storage™ Lightning 9900™ ShadowImage (HMRCF) User's Guide* (MK-90RD012), or contact your Hitachi Data Systems account team.

Table 3.2 Data Management Features and Functions

Feature Name	Controlled by Remote Console?	Controlled by Host OS?	Licensed Software Product?	User Document(s)
Hitachi Remote Copy (Synchronous and Asynchronous) (section 3.6.1)	Yes	Yes	Yes	MK-90RD009
Hitachi Extended Remote Copy (HXRC) (section 3.6.6)	No	Yes	No	*Planning for IBM Remote Copy* (SG24-2595), *DFSMS MVS V1 Remote Copy Guide and Reference* (SC35-0169)
Hitachi Open Remote Copy (Synchronous and Asynchronous) (section 3.6.2)	Yes	Yes	Yes	MK-90RD010
Hitachi ShadowImage (HMRCF) (section 3.6.3)	Yes	Yes	Yes	MK-90RD012
Hitachi Open ShadowImage (HOMRCF) (section 3.6.4	Yes	Yes	Yes	MK-90RD031
Hitachi Command Control Interface (section 3.6.5)	No	Yes	Yes	MK-90RD003
Hitachi Online Data Migration (HODM) (section 3.6.6)	Yes	No	No	Service Offering Only
Hitachi Dynamic Optimizer (HIHSM) (section 3.6.8)	Yes	No	Yes	MK-90RD007
FlashAccess (DCR) (section 3.6.6)	Yes	Yes*	Yes	MK-90RD004
Virtual LVI/LUN (CVS) (section 3.6.10)	Yes	No	Yes	MK-90RD005
LUN Manager and LU Size Expansion (section 3.6.10)	Yes	No	Yes	MK-90RD006
Zone Allocation Manager (LUN Security) (section 3.6.12	Yes	No	Yes	MK-90RD006
LDEV Security (section 3.6.13)	Yes	No	Yes	MK-90RD036
Prioritized Port Control (section 3.6.14)	Yes	No	Yes	MK-90RD030
Hitachi Multiplatform Data Exchange (HMDE) (section 3.6.15)	No	Yes	Yes	MK-90RD020
Hitachi Multiplatform Backup/Restore (HMBR) (section 3.6.14)	No	Yes	Yes	MK-90RD037
HARBOR File-Level Backup and Restore (section 3.6.17)	No	No	No	Service Offering Only
Hitachi GRAPH-Track™ (section 3.6.18)	Yes	No	Yes	MK-90RD032

* Contact your Hitachi Data Systems account team for the latest information on availability.

3.6.4 Open ShadowImage (HOMRCF)

The Hitachi Open ShadowImage feature of the Lightning 9900™ subsystem enables open-system users to maintain subsystem-internal copies of LUs for purposes such as data backup or data duplication. The RAID-protected duplicate LUs (up to nine) are created within the same 9900 subsystem as the primary LU at hardware speeds. Once established, Open ShadowImage operations continue unattended to provide asynchronous internal data backup. Open ShadowImage operations are nondisruptive; the primary LU of each Open 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. Hitachi Open ShadowImage also supports reverse resynchronization for maximum flexibility.

Open ShadowImage operations can be performed in conjunction with Hitachi Open Remote Copy (HORC) operations (see section 3.6.2) to provide multiple copies of critical data at both primary and remote sites. Open ShadowImage also supports the Virtual LUN and FlashAccess features of the 9900 subsystem, ensuring that all user data can be duplicated by Open ShadowImage operations.

For further information, please see *Hitachi Freedom Storage™ Lightning 9900™ Open ShadowImage (HOMRCF) User's Guide* (MK90RD031), or contact your Hitachi Data Systems account team.

3.6.5 Command Control Interface (CCI)

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

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

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

HXRC operations are performed in the same manner as XRC operations. The user issues standard XRC TSO commands from the mainframe host system console directly to the 9900 subsystem. The Remote Console PC is not used to perform HXRC operations. HXRC can be used as an alternative to HRC 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. For further information on XRC, please refer to the following IBM publications: *Planning for IBM Remote Copy* (SG24-2595), and *Remote Copy Administrator's Guide and Reference* (SC35-0169).

3.6.7 Hitachi Online Data Migration (HODM)

The HODM feature is used to migrate data onto the 9900 subsystem from other disk array subsystems. HODM supports migration from older Hitachi subsystems as well as from several non-Hitachi subsystems. HODM can be used to move data to a new location either temporarily or as part of a data process. HODM allows the data being migrated to be online to the host(s) for both read and write I/O operations during HODM migration operations.

HODM is available a Hitachi Data Systems service offering. For further details, please contact your Hitachi Data Systems account team.

3.6.8 Dynamic Optimizer (HIHSM)

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

Dynamic Optimizer provides the following major benefits for the user:

- **Load balancing of subsystem resources**. Balancing resource utilization can significantly improve 9900 subsystem performance. The data provided by Dynamic Optimizer enables users to optimize several areas of performance, including front-end and back-end processor usage as well as the allocation of logical devices to physical drives and RAID level.
- **Optimizing disk drive access patterns**. Dynamic Optimizer collects and analyzes detailed information on disk drive access patterns and can migrate volumes to optimize host access to the data stored on the 9900 subsystem. For example, RAID-1 technology may provide better performance than RAID-5 under certain operational conditions, and one disk drive type may provide better performance than another for certain types of access. The Dynamic Optimizer feature enables users to fine-tune the logical volume allocation of the 9900 subsystem to optimize host access to data.

Dynamic Optimizer operations take into account the RAID level and physical HDD performance of each array group, enabling reallocation of logical volumes and optimization with respect to both RAID level and HDD type. The proper combination of RAID level and HDD type for the logical volumes can significantly improve 9900 subsystem performance for the user’s operational environment.

In addition to RAID level and HDD type, Dynamic Optimizer bases its migration plans on logical device usage, back-end path usage, and disk drive access patterns. Dynamic Optimizer also applies detailed user-specified criteria, including maximum disk utilization for each class of installed HDD, the range of data, which parity drives will be included, and the schedule of automatic migration operations.

For further information on the Dynamic Optimizer feature, please see *Hitachi Freedom Storage™ Lightning 9900™ Dynamic Optimizer (HIHSM) User’s Guide* (MK90RD007), or contact your Hitachi Data Systems account team.

3.6.9 FlashAccess (DCR)

The FlashAccess feature of the 9900 subsystem (also called Dynamic Cache Residency, DCR), 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) have the following parameters:

- They are dynamic and can be added and deleted at any time.
- The 9900 supports a maximum of 1024 addressable cache extents per LDEV, and up to 1,024 cache extents per subsystem.
- All write I/Os to FlashAccess data are duplex writes, guaranteeing full data integrity. The data remains fixed in cache until the user manually deletes it. Deletion of FlashAccess extents de-stages any write data to the affected logical device(s).
- FlashAccess operations support both open-system LUs (e.g., OPEN-3, -8, -9) and S/390® LVIs (e.g., 3380-E/J/K, 3390-1/2/3/3R/9). DCR also supports LUSE and Virtual LVI/LUN volumes. Use of FlashAccess in conjunction with Virtual LVI/LUN volumes will achieve better performance improvements than when either of these options is used individually.
  - For S/390® LVIs, FlashAccess cache extents must be defined on contiguous tracks, with a minimum size of one cache slot (or track) and a maximum size of one LDEV. If a user needs an entire LVI in FlashAccess, using a small Virtual LVI volume uses less cache. For more information on Virtual LVI/LUN, please see *Hitachi Freedom Storage™ Lightning 9900™ Virtual LVI/LUN User's Guide* (MK-90RD005).
  - For open-system LUs, FlashAccess cache extents must be defined in logical blocks using logical block addresses (LBAs), with a minimum size of 96 LBAs. However, most users will assign the entire LU for FlashAccess. In this case using a small Virtual LUN device uses less cache.

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

3.6.10 LUN Manager and LU Size Expansion

The LUN Manager feature enables you to set and define the port modes for Fibre channel ports and to set the fibre topology (point to point, FC-AL or fabric). Please connect your Hitachi Data Account team for further details on this feature.

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

3.6.11 Virtual LVI/LUN (CVS)

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

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

For more information, please see *Hitachi Freedom Storage™ Lightning 9900™ Virtual LVI/LUN (CVS) User's Guide* (MK90RD005), or contact your Hitachi Data Systems account team.

For more information, please see *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK90RD006), or contact your Hitachi Data Systems account team.

3.6.12 Zone Allocation Manager (LUN Security)

The Zone Allocation Manager feature of the 9900 subsystem allows you to restrict LU accessibility to an open-systems host using open-systems host's World Wide Name (WWN). You can set a 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 Security software for the 9900 Remote Console PC enables you to configure Zone Allocation Manager operations on the 9900 subsystem.

LUN Security can be activated on any installed fibre channel port, and be turned on or off at the port level. If you disable LUN Security on a particular port, that LU will not be restricted to a particular host or group of hosts. If you enable LUN Security 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 Security software allows you to create LU and WWN groups, so you can more easily manage your 9900 storage subsystem. A 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.

For more information, please see *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK90RD006), or contact your Hitachi Data Systems account team.

3.6.13 LDEV Security

The LDEV Security feature of the Hitachi Freedom Storage™ 7700E subsystem allows you to restrict S/390® host access to the logical devices (LDEVs) on the 7700E subsystem. Each LDEV to can be set to communicate only with user-selected host(s). The LDEV Security feature prevents other hosts from seeing the secured LDEV and from accessing the data contained on the secured LDEV. The licensed LDEV Security software on the 7700E Remote Console PC displays the LDEV Security information and allows you to perform LDEV Security operations.

For more information, please see *Hitachi Freedom Storage™ Lightning 9900™ LDEV Security User's Guide* (MK-90RD036), or contact your Hitachi Data Systems account team.

3.6.14 Prioritized Port Control (PPC)

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

3.6.15 Hitachi Multiplatform Data Exchange (HMDE)

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

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

The HMRS feature and the HMDE software enable the open-system host to read from and write to individual S/390® datasets. The HMDE 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 HMDE volumes. For further information on HMDE, please refer to the *Hitachi Multiplatform Data Exchange (HMDE) User's Guide* (MK-90RD020), or contact your Hitachi Data Systems account team.

3.6.16 Hitachi Multiplatform Backup/Restore (HMBR)

The HMBR feature allows the user to implement mainframe-based backup procedures and standards for the open-system data stored on the multiplatform 9900 subsystem. HMBR enables standard mainframe backup/restore utilities such as DFDSS, Fast Dump/Restore (FDR), and VSE FASTWRITE to perform volume-level backup and restore operations on OPEN-3 and OPEN-9 LUNs. 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.

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

3.6.17 HARBOR File-Level Backup/Restore

The HARBOR File-Level Backup/Restore multiplatform feature of the 9900 subsystem features an integrated architecture and includes:

- A host component on MVS,
- Integrated clients for desktops and servers,
- LAN-based distributed storage servers,
- High-speed HMDE 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.

3.6.18 Hitachi GRAPH-Track

Hitachi GRAPH-Track™ is the performance and usage monitoring utility for the 9900 subsystem. Hitachi GRAPH-Track™ runs on the Remote Console PC and can monitor as many as eight subsystems on the 9900-internal LAN. Hitachi GRAPH-Track™ 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. Hitachi GRAPH-Track™ displays the following data for each attached subsystem:

- Subsystem configuration, including controller name, serial number, controller emulation, channel address(s), cache size, SSIDs, and current duplex write line (DWL).
- 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 (called CHIPs) and back-end microprocessors (called ACPs).
- Cache statistics, including amount of cache in use, and amount of write-pending data in cache.
- I/O statistics at the subsystem, array group, and logical device levels: I/O rates, read/write ratio, read hits and write hits, and 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 storage of the GRAPH-Track data automatically according to user-specified preferences and allows the user to perform the following important functions:

- Starting and stopping subsystem and LDEV data collection for each subsystem.
- Setting the GT data sampling rate for each subsystem (one, three, five, or fifteen minutes).
- Exporting GRAPH-Track data for use in reports or in other data analysis programs.
- Purging (deleting) selected GRAPH-Track data from the GT database.
- Configuring the GRAPH-Track database manager and automatic purge scheduler to control the contents and size of the GT database.

For further information and instructions on using Hitachi GRAPH-Track, please contact your Hitachi Data Systems account team.

Chapter 4 Configuring and Using the 9900 Subsystem

4.1 S/390® Configuration

The first step in 9900 configuration is to define the subsystem to the S/390® host(s). The three basic areas requiring definition are: SSIDs, hardware definitions (including I/O Configuration Program [IOCP] or Hardware Configuration Definition [HCD]), and operating system definitions (HCD or OS commands).

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

4.1.1 Subsystem IDs (SSIDs)

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

Table 4.1 SSID Requirements

Controller Emulation	SSID Requirements	LVI Support
3990-3	0004 - 00fd	3380-x, 3390-x, OPEN-K,3,8,9 LVIs
3990-6 and 3990-6E	0004 - fffd	3390-x, OPEN-3,8,9 LVIs

4.2 S/390® Hardware Definition

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

The I/O Configuration Program (IOCP) can be used to define the 9900 subsystem in MVS, VM, and VSE environments (wherever HCD cannot be used). To define a 9900 subsystem with 64 LVIs or fewer, use the same procedure as for an IBM 3990-3, 3990-6, or 3990-6E subsystem. Figure 4.1 shows IOCP hardware definition for a 9900 configured with ExSA™ (ESCON®) channels and 256 LVIs. Figure 4.2 and Figure 4.3 show IOCP hardware definition for a 9900 configured with 1024 LVIs and four CU images. The CUADD parameter defines the CU images by number (0-3).

Note: The 9900 subsystem supports up to 16 CU Images (CU 0-F) and 4096 LDEVs. Each CU image can hold up to 256 LDEV addresses.

```
CHPID PATH=((2A,2B,2C,2D,3A,3B,3C,3D)),TYPE=CNC
CNTLUNIT CUNUMBR=F00,(PATH=2A,2B,2C,2D),UNIT=3990,UNITADD=((00,256))
CNTLUNIT CUNUMBR=F01,(PATH=3A,3B,3C,3D),UNIT=3990,UNITADD=((00,256))
IODEVICE ADDRESS=(100,256),CUNUMBR=(F00,F01),UNIT=3390,FEATURE=(ALTCTRL,SHARED)
```

Figure 4.1 IOCP Definition for ExSA (ESCON®) Channels and 256 LVIs

```
CHPID      PATH=A0,TYPE=CNC,PARTITION=(LPA1,REC)
CHPID      PATH=A1,TYPE=CNC,PARTITION=(LPA1,REC)
CHPID      PATH=A2,TYPE=CNC,PARTITION=(LPA1,REC)
CHPID      PATH=A3,TYPE=CNC,PARTITION=(LPA1,REC)
CNTLUNIT  CUNUMBR=100,PATH=(A0,A1),UNITADD=((00,256)),CUADD=0,UNIT=3990
CNTLUNIT  CUNUMBR=110,PATH=(A2,A3),UNITADD=((00,256)),CUADD=0,UNIT=3990
CNTLUNIT  CUNUMBR=101,PATH=(A0,A1),UNITADD=((00,256)),CUADD=1,UNIT=3990
CNTLUNIT  CUNUMBR=111,PATH=(A2,A3),UNITADD=((00,256)),CUADD=1,UNIT=3990
CNTLUNIT  CUNUMBR=102,PATH=(A0,A1),UNITADD=((00,256)),CUADD=2,UNIT=3990
CNTLUNIT  CUNUMBR=112,PATH=(A2,A3),UNITADD=((00,256)),CUADD=2,UNIT=3990
CNTLUNIT  CUNUMBR=103,PATH=(A0,A1),UNITADD=((00,256)),CUADD=3,UNIT=3990
CNTLUNIT  CUNUMBR=113,PATH=(A2,A3),UNITADD=((00,256)),CUADD=3,UNIT=3990
IODEVICE   ADDRESS=(000,256),CUNUMBR=(100,110),UNIT=3390
IODEVICE   ADDRESS=(100,256),CUNUMBR=(101,111),UNIT=3390
IODEVICE   ADDRESS=(200,256),CUNUMBR=(102,112),UNIT=3390
IODEVICE   ADDRESS=(300,256),CUNUMBR=(103,113),UNIT=3390
```

Figure 4.2 IOCP Definition for 1024 LVIs (9900 connected directly to host CPU)

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

```
CHPID     PATH=B0,TYPE=CNC,PARTITION=(LPA1,REC),SWITCH=01
CHPID     PATH=B1,TYPE=CNC,PARTITION=(LPA1,REC),SWITCH=01
CHPID     PATH=B2,TYPE=CNC,PARTITION=(LPA1,REC),SWITCH=01
CHPID     PATH=B3,TYPE=CNC,PARTITION=(LPA1,REC),SWITCH=01
CNTLUNIT  CUNUMBR=0200,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=0,UNIT=3990
CNTLUNIT  CUNUMBR=0210,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C2,C3),CUADD=0,UNIT=3990
CNTLUNIT  CUNUMBR=0201,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=1,UNIT=3990
CNTLUNIT  CUNUMBR=0211,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C2,C3),CUADD=1,UNIT=3990
CNTLUNIT  CUNUMBR=0202,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=2,UNIT=3990
CNTLUNIT  CUNUMBR=0212,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C2,C3),CUADD=2,UNIT=3990
CNTLUNIT  CUNUMBR=0203,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=3,UNIT=3990
CNTLUNIT  CUNUMBR=0213,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C2,C3),CUADD=3,UNIT=3990
IODEVICE  ADDRESS=(000,256),CUNUMBR=(200,210),UNIT=3390
IODEVICE  ADDRESS=(100,256),CUNUMBR=(201,211),UNIT=3390
IODEVICE  ADDRESS=(200,256),CUNUMBR=(202,212),UNIT=3390
IODEVICE  ADDRESS=(300,256),CUNUMBR=(203,213),UNIT=3390
```

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

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

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

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

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

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

Figure 4.4 shows an IOCP definition for a 9900 using 16 channels (16 physical paths), four CU images, and 1024 devices. The 9900 is configured using four CU macro statements for each logical CU image. For example, 16 CU macro statements are used for connection to four CU images with 16 physical paths to the host. This configuration enables access to the devices over eight different channel cards and also provides alternate CUs in case CU busy messages are encountered.

```
CU       100           104           101           105
CU       108           10C           109           10D
      ----------    ----------    ----------    ----------
CHPID   60    61      62    63      64    65      66    67
LCP     1A    1C      1E    1G      1J    1L      1N    1Q
        |     |       |     |       |     |       |     |
                                CLUSTER1
CU       102           106           103           107
CU       10A           10E           10B           10F
      ----------    ----------    ----------    ----------
CHPID   68    69      6A    6B      7A    7B      7C    7D
LCP     2A    2C      2E    2G      2J    2L      2N    2Q
        |     |       |     |       |     |       |     |
                                CLUSTER2
CHPID Path=60,type=cnc,partition=(prod,maint,shr)
CHPID Path=61,type=cnc,partition=(prod,maint,shr)
CHPID Path=62,type=cnc,partition=(prod,maint,shr)
CHPID Path=63,type=cnc,partition=(prod,maint,shr)
CHPID Path=64,type=cnc,partition=(maint,prod,shr)
CHPID Path=65,type=cnc,partition=(maint,prod,shr)
CHPID Path=66,type=cnc,partition=(maint,prod,shr)
CHPID Path=67,type=cnc,partition=(maint,prod,shr)
CHPID Path=68,type=cnc,partition=(prod,maint,shr)
CHPID Path=69,type=cnc,partition=(prod,maint,shr)
CHPID Path=6A,type=cnc,partition=(prod,maint,shr)
CHPID Path=6B,type=cnc,partition=(prod,maint,shr)
CHPID Path=7A,type=cnc,partition=(maint,prod,shr)
CHPID Path=7B,type=cnc,partition=(maint,prod,shr)
CHPID Path=7C,type=cnc,partition=(maint,prod,shr)
CHPID Path=7D,type=cnc,partition=(maint,prod,shr)
cntlunit cunumber=100,path=(60,61),unit=3990,unitadd=(00,256),cuadd=0
cntlunit cunumber=101,path=(64,65),unit=3990,unitadd=(00,256),cuadd=0
cntlunit cunumber=102,path=(68,69),unit=3990,unitadd=(00,256),cuadd=0
cntlunit cunumber=103,path=(7A,7B),unit=3990,unitadd=(00,256),cuadd=0
cntlunit cunumber=104,path=(62,63),unit=3990,unitadd=(00,256),cuadd=1
cntlunit cunumber=105,path=(66,67),unit=3990,unitadd=(00,256),cuadd=1
cntlunit cunumber=106,path=(6A,6B),unit=3990,unitadd=(00,256),cuadd=1
cntlunit cunumber=107,path=(7C,7D),unit=3990,unitadd=(00,256),cuadd=1
cntlunit cunumber=108,path=(60,61),unit=3990,unitadd=(00,256),cuadd=2
cntlunit cunumber=109,path=(64,65),unit=3990,unitadd=(00,256),cuadd=2
cntlunit cunumber=10A,path=(68,69),unit=3990,unitadd=(00,256),cuadd=2
cntlunit cunumber=10B,path=(7A,7B),unit=3990,unitadd=(00,256),cuadd=2
cntlunit cunumber=10C,path=(62,63),unit=3990,unitadd=(00,256),cuadd=3
cntlunit cunumber=10D,path=(66,67),unit=3990,unitadd=(00,256),cuadd=3
cntlunit cunumber=10E,path=(6A,6B),unit=3990,unitadd=(00,256),cuadd=3
cntlunit cunumber=10F,path=(7C,7D),unit=3990,unitadd=(00,256),cuadd=3
iodevice address=(1000,256),cunumber=(100,101,102,103),unit=3390
iodevice address=(1100,256),cunumber=(104,105,106,107),unit=3390
iodevice address=(1200,256),cunumber=(108,109,10A,10B),unit=3390
iodevice address=(1300,256),cunumber=(10C,10D,10E,10F),unit=3390
```

This IOCP definition maps to:

Devices	Partitions	CHPIDs
1000-10ff	maint,prod	60, 61, 64, 65, 68, 69, 7A, 7B
1100-11FF	maint,prod	62, 63, 66, 67, 6A, 6B, 7C, 7D
1200-12ff	maint,prod	60, 61, 64, 65, 68, 69, 7A, 7B
1300-13FF	maint,prod	62, 63, 66, 67, 6A, 6B, 7C, 7D

Figure 4.4 IOCP Definition for 16 Channels, 4 CUs, and 1024 LVIs

4.2.2 Hardware Definition Using HCD (MVS/ESA)

The Hardware Configuration Definition (HCD) utility can be used to define the 9900 subsystem in an MVS/ESA environment. To define a 9900 with 64 or fewer LVIs, use the same procedure as for an IBM 3990-6, 3990-6E, or 3990-3 subsystem (see Table 4.4). The hardware definition for a 9900 subsystem with more than 64 LVIs (see Table 4.5) is different than that for an IBM 3990 subsystem.

Table 4.4 HCD Definition for 64 LVIs

Parameter	Value
Control Frame:	
Control unit number	Specify the control unit number.
Control unit type	3990-6 or 3990-6E (using 3990-6 emulation), 3990-3 (using 3990-3 emulation).
Channel path IDs	Specify the ExSA.
Unit address	00 (ExSA)
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 the ExSA.
Unit address	00
Number of units	256
Array Frame:	
Device number	Specify the first device number.
Number of devices	256
Device type	3390
Connected to CUs	Specify the control unit number(s).

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

4.2.3 Defining the 9900 to VM/ESA

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

[Address Range] TYPE DASD SHARED YES

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

4.2.4 Defining the 9900 to TPF

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

4.3 S/390® Operations

4.3.1 Initializing the LVIs

The 9900 LVIs require only minimal initialization before being brought online. Figure 4.5 shows an MVS ICKDSF JCL example of a minimal init job to write a volume ID (VOLID) and volume table of contents (VTOC).

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

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

Figure 4.5 LVI Initialization for MVS: ICKDSF JCL

4.3.2 Device Operations: ICKDSF

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

Table 4.6 ICKDSF Commands for 9900 Contrasted to RAMAC

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

4.3.3 MVS Cache Operations

To display the 9900 cache statistics under MVS DFSMS, use the following operator command: **D SMS, CACHE**. Figure 4.6 shows the cache statistics reported by the 9900. The 9900 reports cache statistics for each SSID in the subsystem. Because the dynamic cache management algorithm has been enhanced, the read and write percentages for the 9900 are displayed as **N/A**. For further information on MVS DFSMS cache reporting, please refer to the IBM document *DFSMSdfp Storage Administrator Reference* (SC28-4920).

```
SSID          DEVS     READ     WRITE    HIT RATIO                 FW BYPASSES

0004          15       N/A      N/A      50%             0

0005          11       N/A      N/A      0%              0

0006          11       N/A      N/A      87%             0

0007          10       N/A      N/A      87%             0
****************************************************
SSID=SUBSYSTEM IDENTIFIER
DEVS=NUMBER OF MANAGED DEVICES ATTACHED TO SUBSYSTEM
READ=PERCENT OF DATA ON MANAGED DEVICES ELIGIBLE FOR CACHING
WRITE=PERCENT OF DATA ON MANAGED DEVICES ELIGIBLE FOR FAST WRITE

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

Figure 4.6 Displaying Cache Statistics Using MVS DFSMS

The 9900 supports the following MVS cache operations:

- **IDCAMS LISTDATA COUNTS**. When the <**subsystem**> parameter is used with the LISTDATA command, the user must issue the command once for each SSID to view the entire 9900 image. Figure 4.7 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 4.7 IDCAMS LISTDATA COUNTS (JCL example)

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

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

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

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

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

SETCACHE SUBSYSTEM ON|OFF
SETCACHE CACHEFASTWRITE ON|OFF
SETCACHE NVS ON|OFF
SETCACHE DESTAGE

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

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

DEVSERV PATHS,100,64
DEVSERV PATHS,140,64
DEVSERV PATHS,180,64
DEVSERV PATHS,1C0,64

4.3.4 VM/ESA Cache Operations

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

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

4.3.5 VSE/ESA Cache Operations

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

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

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

4.4 Open-Systems Configuration

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

Table 4.7 9900 Open-System Configuration Guides

Title	Document Number
IBM® AIX® Configuration Guide	MK-90RD014
HP-UX® Configuration Guide	MK-90RD016
Sun® Solaris® Configuration Guide	MK-90RD017
Compaq® Tru64 UNIX® Configuration Guide	MK-90RD021
Windows NT® Configuration Guide	MK-90RD015
Windows® 2000 Configuration Guide	MK-90RD025
Novell® NetWare® Configuration Guide	MK-90RD026
Red Hat® Linux® Configuration Guide	MK-90RD028

4.4.1 Configuring the Fibre-Channel Ports

The LUN Manager remote console software enables users to configure the fibre-channel ports for the connected operating system and operational environment (e.g., FC-AL or fabric). If desired, Hitachi Data Systems can configure the fibre-channel ports as a fee-based service. For further information on LUN Manager, see *Hitachi Freedom Storage™ Lightning 9900™ LUN Manager User's Guide* (MK-90RD006), or contact your Hitachi Data Systems account team.

The 9960 subsystem supports a maximum of 32 fibre-channel ports, and the 9910 supports up to 24 fibre-channel ports. Each fibre-channel port is assigned a unique target ID (from 0 to EF). The 9900 subsystem supports up to 256 LUNs per port. Figure 4.8 illustrates fibre port-to-LUN addressing.

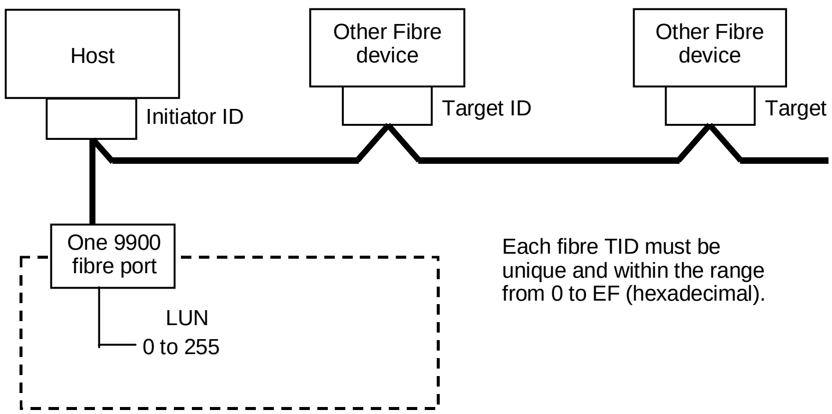

Figure 4.8 Fibre Port-to-LUN Addressing

4.4.2 Custom-Volume-Size (CVS) Devices

The Virtual LVI/LUN remote console software enables users to configure custom-size LUs (CVS devices) which are smaller than standard-size LUs. CVS devices which are created from OPEN-x LUs are defined by size in MB (minimum CVS device size = 35 MB). (CVS devices which are created from mainframe LVIs are defined by number of cylinders.)

4.4.3 LU Size Expansion (LUSE) Devices

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

4.5 Open Systems Operations

4.5.1 Command Tag Queuing

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

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

4.5.2 Host/Application Failover Support

The 9900 supports many industry-standard middleware products which provide host and/or application failover capabilities (e.g., HP® MC/ServiceGuard, VERITAS® First Watch®, HACMP, HAGEO, IRIS FailSafe™, SGI™ Advanced Cluster Environment (ACE), Microsoft Cluster Server, Novell® High Availability Server (NHAS), Novell® System Fault Tolerance (SFT) III). Please consult the appropriate 9900 configuration guide for further information on failover support.

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 middleware. The 9900 provides up to 32 fibre ports to accommodate alternate pathing for host attachment. Figure 4.9 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.

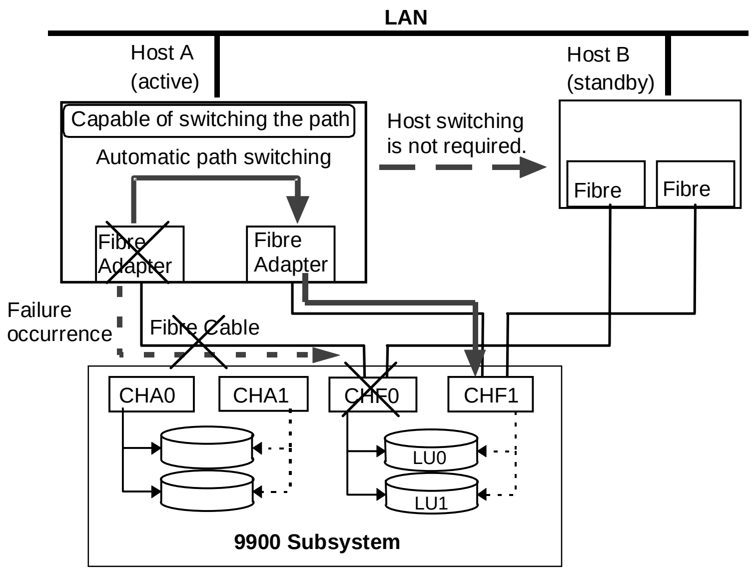

Figure 4.9 Alternate Pathing

4.5.4 Remote SIM (R-SIM) Reporting

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

4.5.5 SNMP Remote Subsystem Management

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

Chapter 5 Planning for Installation and Operation

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

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

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

5.1 User Responsibilities

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

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

5.2 Electrical Specifications and Requirements for Single-Phase Subsystems

5.2.1 Internal Cable Diagram

Figure 5.1 and Figure 5.2 illustrate the internal cable layout of single-phase 9960 and 9910 subsystems, respectively.

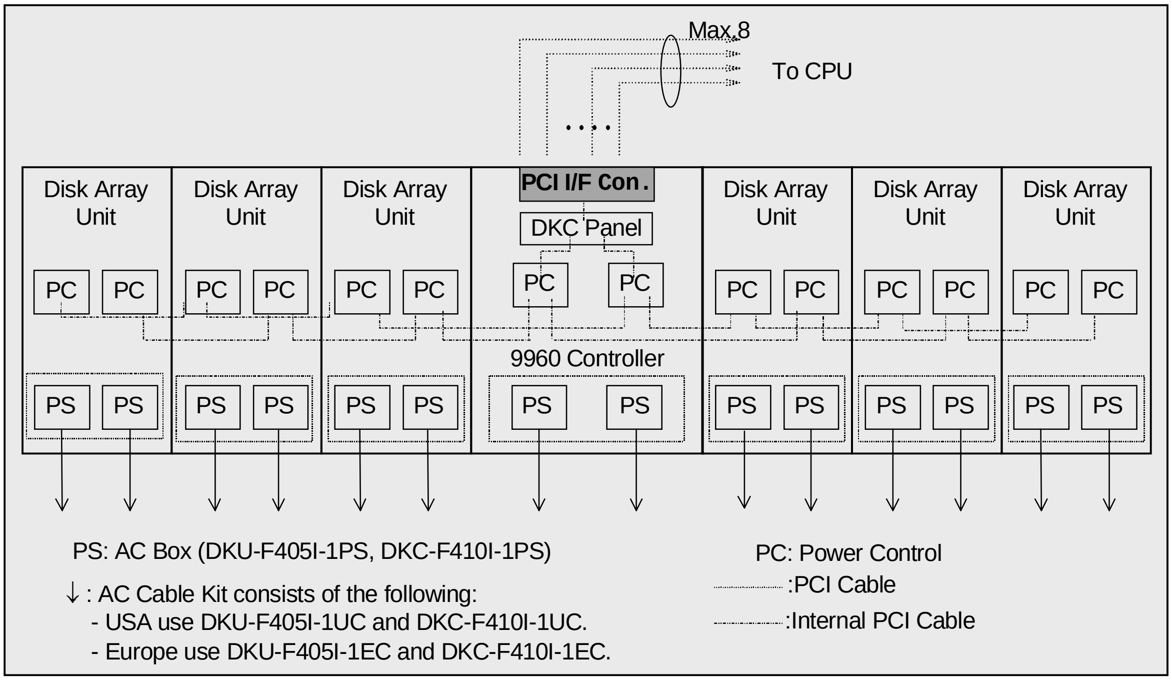

Figure 5.1 Internal Cable Diagram of a Single-Phase 9960 Subsystem

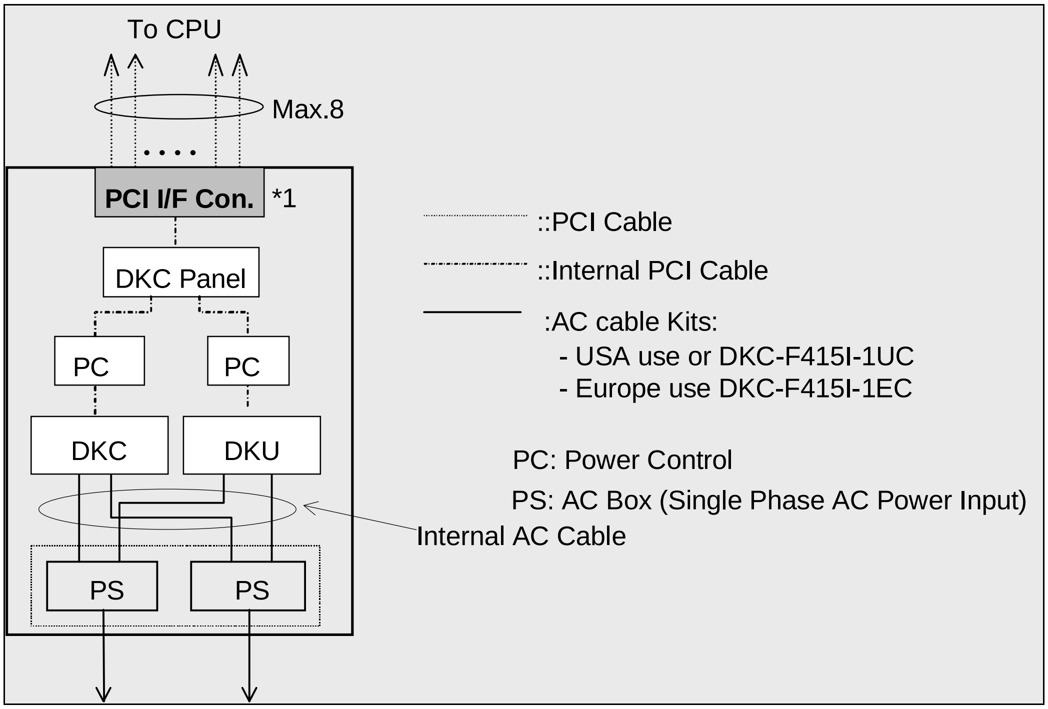

Figure 5.2 Internal Cable Diagram of a Single-Phase 9910 Subsystem

5.2.2 Power Plugs

Figure 5.3 through Figure 5.8 show the power plugs for the 9960 and 9910 subsystems.

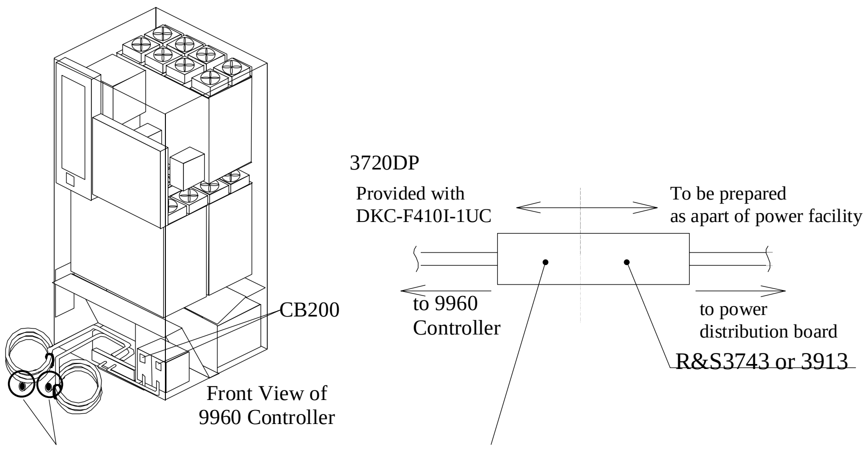

Figure 5.3 Power Plugs for Single-Phase 9960 Controller (USA)

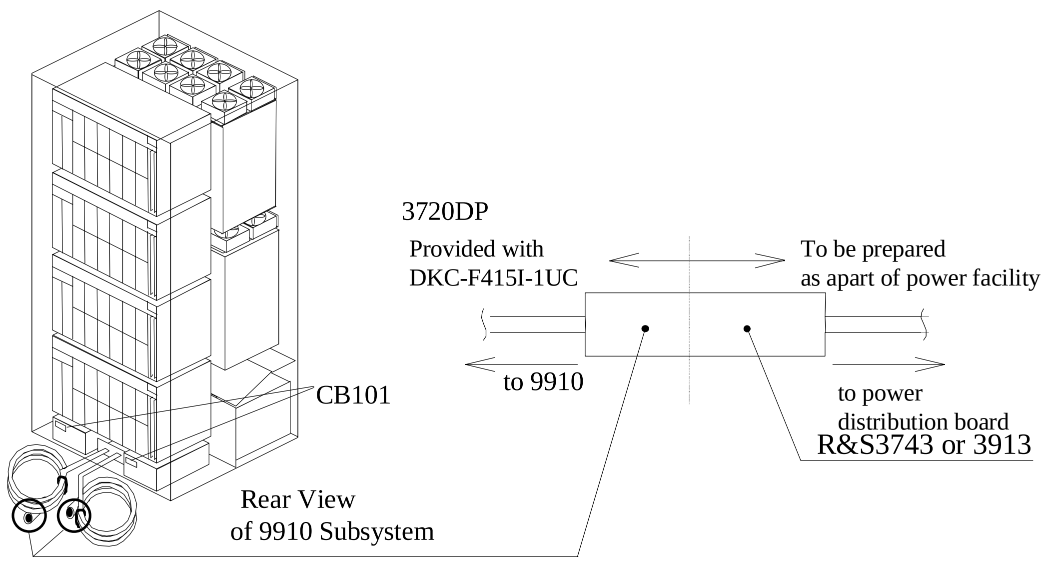

Figure 5.4 Power Plugs for Single-Phase 9910 Subsystem (USA)

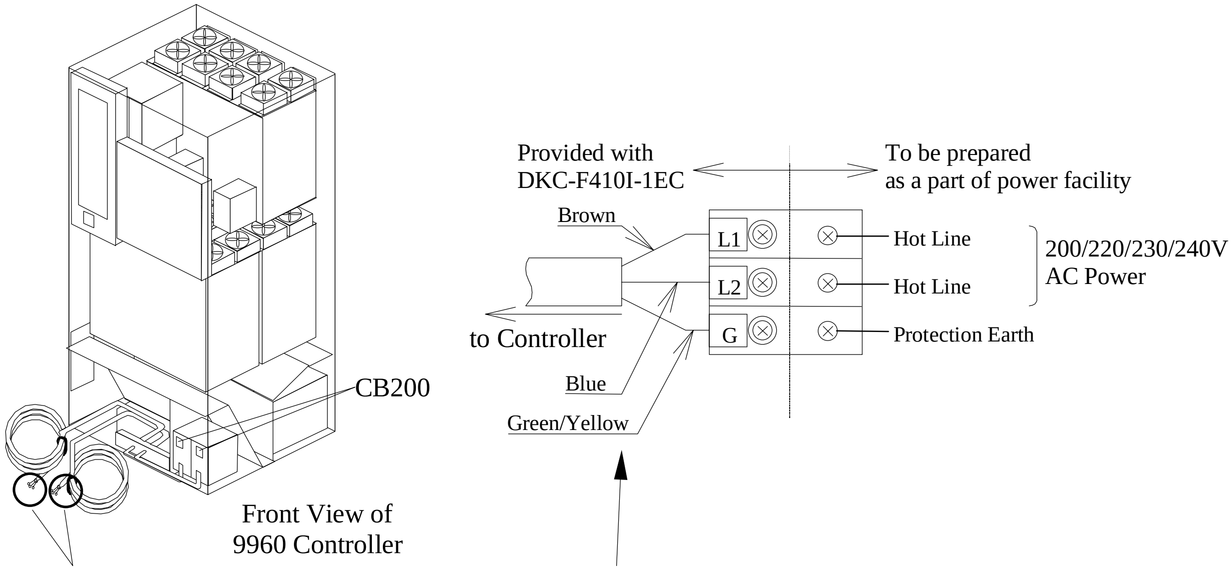

Figure 5.5 Power Plugs for a Single-Phase 9960 Controller (Europe)

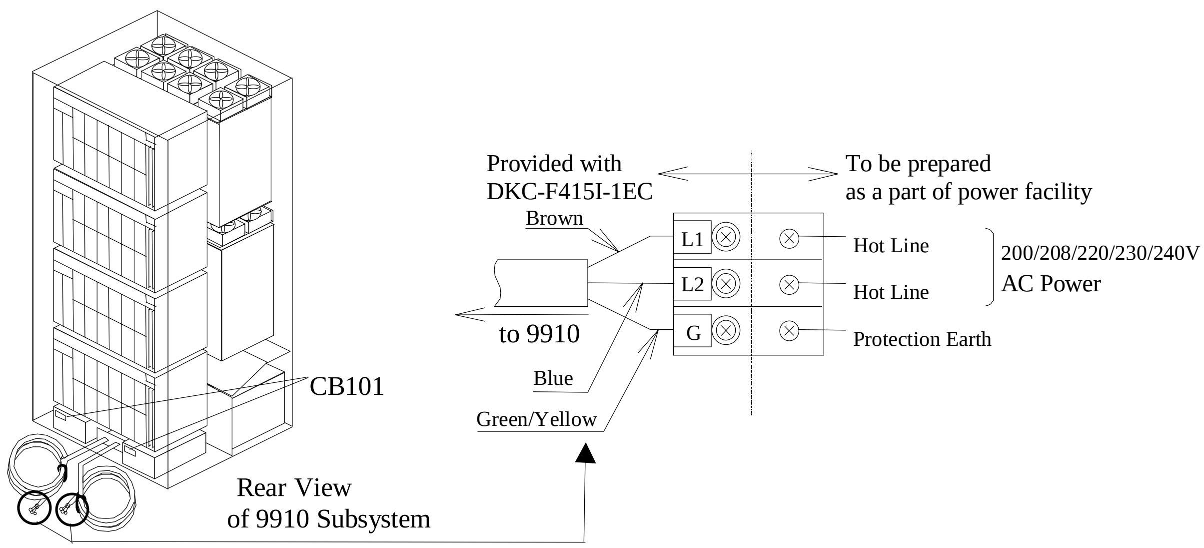

Figure 5.6 Power Plugs for a Single-Phase 9910 Subsystem (Europe)

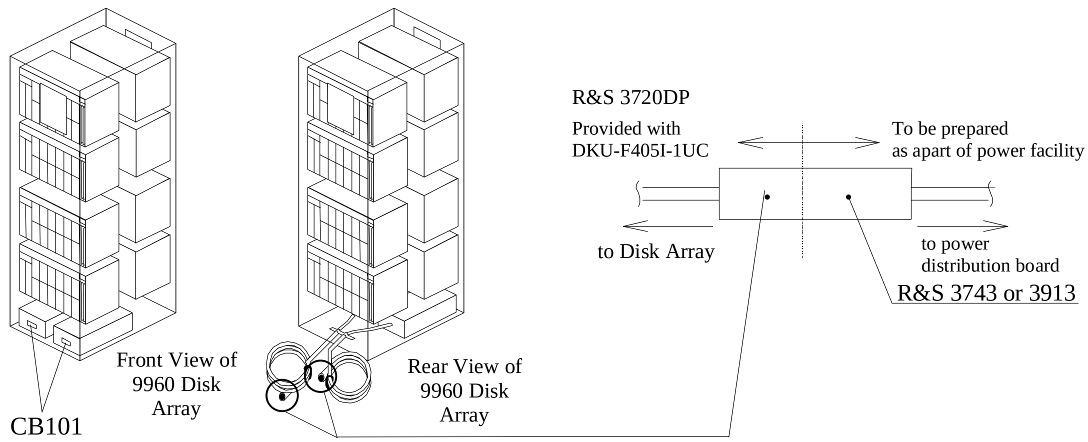

Figure 5.7 Power Plugs for Single-Phase 9960 Disk Array Unit (USA)

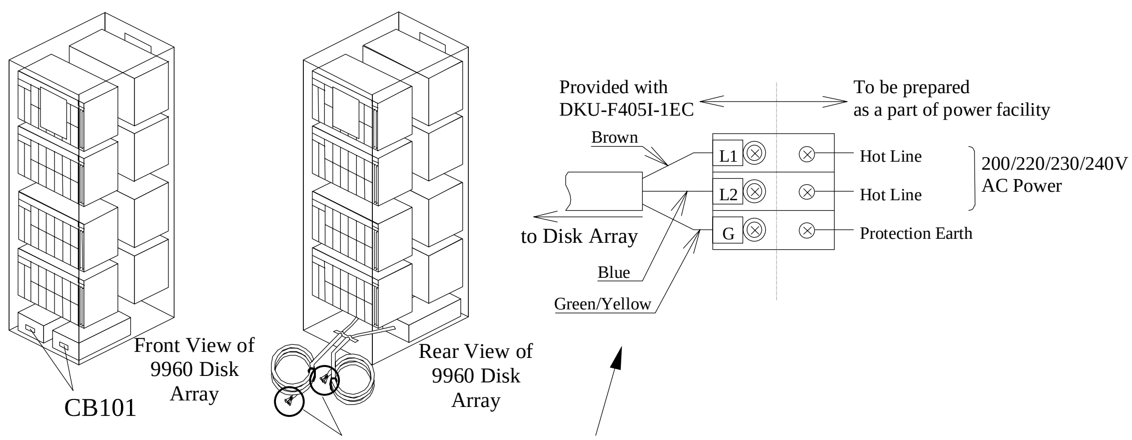

Figure 5.8 Power Plugs for Single-Phase 9960 Disk Array Unit (Europe)

5.2.3 Features

Table 5.1 lists the features for single-phase 9960 and 9910 subsystems.

Table 5.1 9900 Single-Phase Features

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

5.2.4 Connector, Receptacle Part Numbers and Electrical Ratings

Table 5.2 lists the connector and receptacle part numbers and electrical ratings recommended for single-phase 60-Hz input power. Each disk array frame and controller frame has 2 input power cables with R & S FS 3720 plugs. The customer must supply the outlets for the plugs.

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

Table 5.2 Input Power Connector, Receptical and Ratings for Single-Phase 9900

Item	Device	Quantity	Plugs	Connector	Receptacle	Electrical Rating
9960	Controller frame	2	R&S FS 3720	R&S 3913	R&S 3743	60HZ 200V, 208V, 230 V 20 Amps single phase
	Disk Array frame					
9910	Subsystem					

5.2.5 Input Voltage Tolerances

Table 5.3 lists the input voltage tolerances for the single-phase 9900 subsystem. Transient voltage conditions must not exceed +15-18% of nominal and must return to a steady-state tolerance of between +6 and -8% of the normal related voltage in 0.5 seconds or less. Line-to-line imbalance voltage must not exceed 2.5%. Nonoperating harmonic contents must not exceed 5%.

Table 5.3 Input Voltage Specifications for Single-Phase Power

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

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

5.3 Electrical Specifications and Requirements for Three-Phase Subsystems

5.3.1 Internal Cable Diagram

Figure 5.9 illustrates the internal cable layout of a three-phase 9960 subsystem.

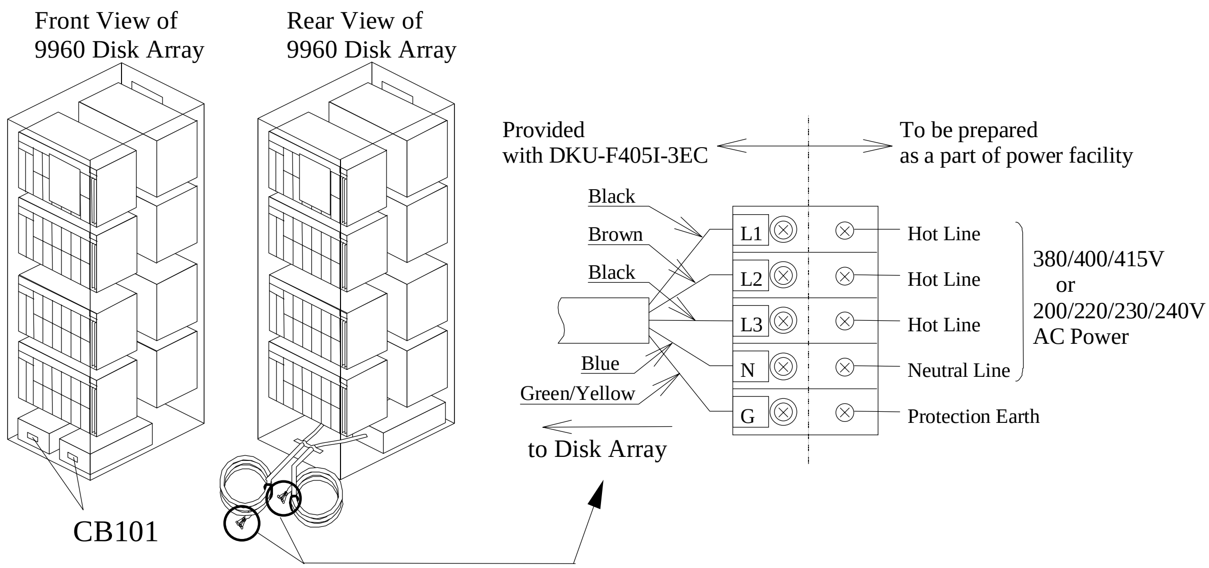

Figure 5.9 Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (Europe)

5.3.2 Power Plugs

Figure 5.10 illustrates the power plugs for a three-phase 9960 disk array unit (USA). Figure 5.11 illustrates the power plugs for a three-phase 9960 disk array unit (Europe).

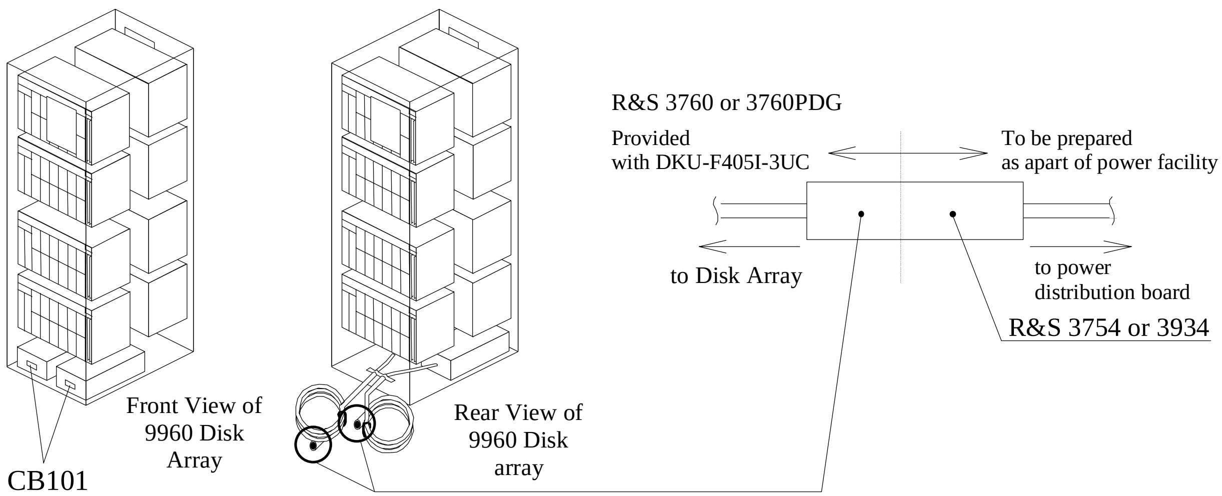

Figure 5.10 Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (USA)

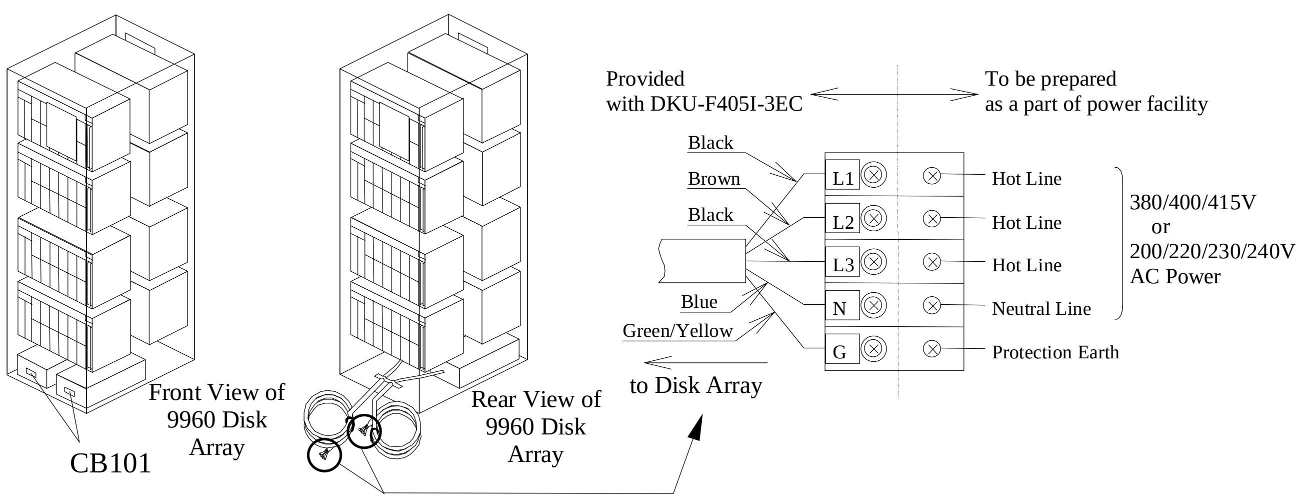

Figure 5.11 Diagram of Power Plugs for Three-Phase 9960 Disk Array Unit (Europe)

5.3.3 Features

Table 5.4 lists the features for a three-phase 9960 subsystem.

Table 5.4 9960 Three-Phase Features

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

5.3.4 Connector and Receptacle Part Numbers and Electrical Ratings

Table 5.5 lists the connector and receptacle part numbers and electrical ratings recommended for three-phase 60-Hz input power. In a three-phase 9960 subsystem the controller frame receives its AC input power from the first Disk Array Frame via internal cabling, so that subsystem will not require any customer outlets for the controller frame. The user must supply all power receptacles and connectors for both 60-Hz and 50-Hz subsystems. Russell & Stoll type (R&S) connectors are recommended for 60-Hz systems.

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

Table 5.5 Input Power Connector and Ratings for Three-Phase 9960 Model

Device	Quantity	Plugs	Connector	Receptacle	Electrical Rating
Controller	0	N/A	N/A	N/A	60 Hz 200V,208V, 230 V
Disk Unit (US)*	2	R&S FS 3760	R&S 3934	R&S 3754	60HZ 200V, 208V, 230 V 30 Amps three-phase

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

5.3.5 Input Voltage Tolerances

Table 5.6 lists the input voltage tolerances for the three-phase 9900. 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 5.6 Input Voltage Specifications for Three-Phase Power

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

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

5.4 Dimensions and Weight

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

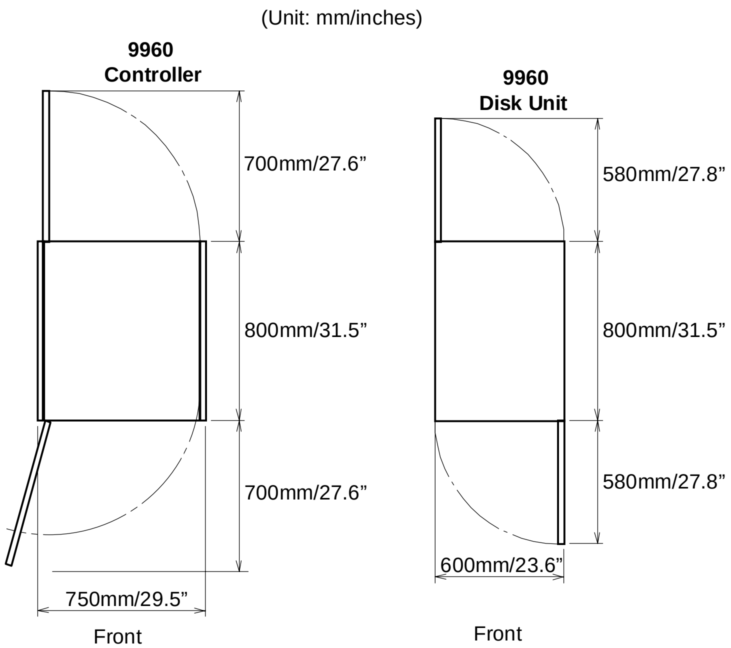

Figure 5.12 9960 DKC and DKU Physical Dimensions

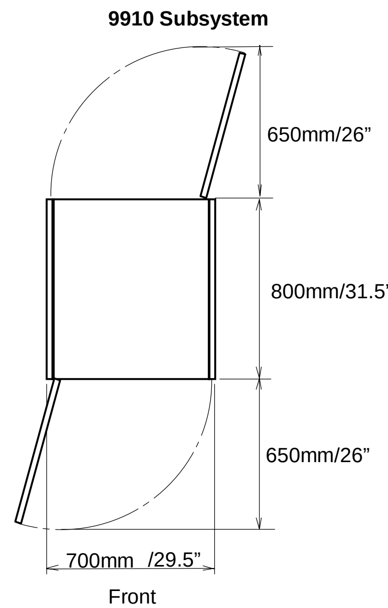

Figure 5.13 9910 DKC and DKU Physical Dimensions

Table 5.7 9900 Physical Specifications

Item		9960 Controller		9960 Disk Unit *6		9910 Subsystem	
Dimension (mm/in)	Width	750 mm / 29.5 in *1		600 mm / 23.6 in *1		700 mm / 29.5 in *1	
	Depth	800 mm / 31.5 in					
	Height	1,790 mm / 70.5 in					
Floor Area with no service clearance		0.6 sq meters / 6.5 sq feet		0.48 sq meters / 5.16 sq feet		0.56 sq meters / 6 sq feet	
Floor Area with service clearance		1.8 sq meters / 19.4 sq feet		1.44 sq meters / 15.5 sq feet		1.47 sq meters / 15.8 sq feet	
Weight (kg)/lb		412 / 908 *2	480 / 1058 *5	480 / 1058 *3	440 / 970 *4	525 / 1157 *7	506 / 1116 *8
Heat Output (kW)		1.46 *2	3.43 *5	2.32 *3	2.85 *4	3.1 *7	2.81 *8
Power Consumption (kVA)		1.57 *2	3.72 *5	2.41 *3	3.14 *4	3.34 *7	3.06 *8
Air Flow (m3/min.)		18 *2 *3		12 *4 *5		15 *7 *8	

Notes:

*1. This includes the thickness of side covers (16mm/.6 in × 2).
*2. These values are used when 9960 has 8 GB cache memory, two fibre 8-port adapters, and an additional disk adapter.
*3. These values are used when 9960 Controller has maximum options.
*4. These values apply to a maximum size Disk Unit using 18-GB hard disk drives (HDDs).
*5. These values apply to a maximum size Disk Unit using 47-GB or 73-GB HDDs.
*6. The number of Disk Units per 9960 subsystem is 1-6.
*7. These values are used when 9910 has 16 GB cache, three fibre 8-port adapters, and fully populated with 47- or 73-GB HDDs.
*8. These values are used when 9910 has 16 GB cache memory, three fibre 8-port adapters, and fully populated with 18-GB HDDs.

Table 5.8 9960 Frame and Component Weights

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

Table 5.9 9960 & 9910 Subsystem Weights

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

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

5.5 Floor Loading and Cable Routing Requirements

5.5.1 Service Clearance Requirements

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

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

The additional weight of the raised floor and the weight of the cables is 10 lb/ft^2 (50 kg/m^2) 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/ft^2 (75 kg/m^2) is allowed. This distributed weight is applied over half of the service clearance area up to a maximum of 30 inches (760 mm) from the machine.

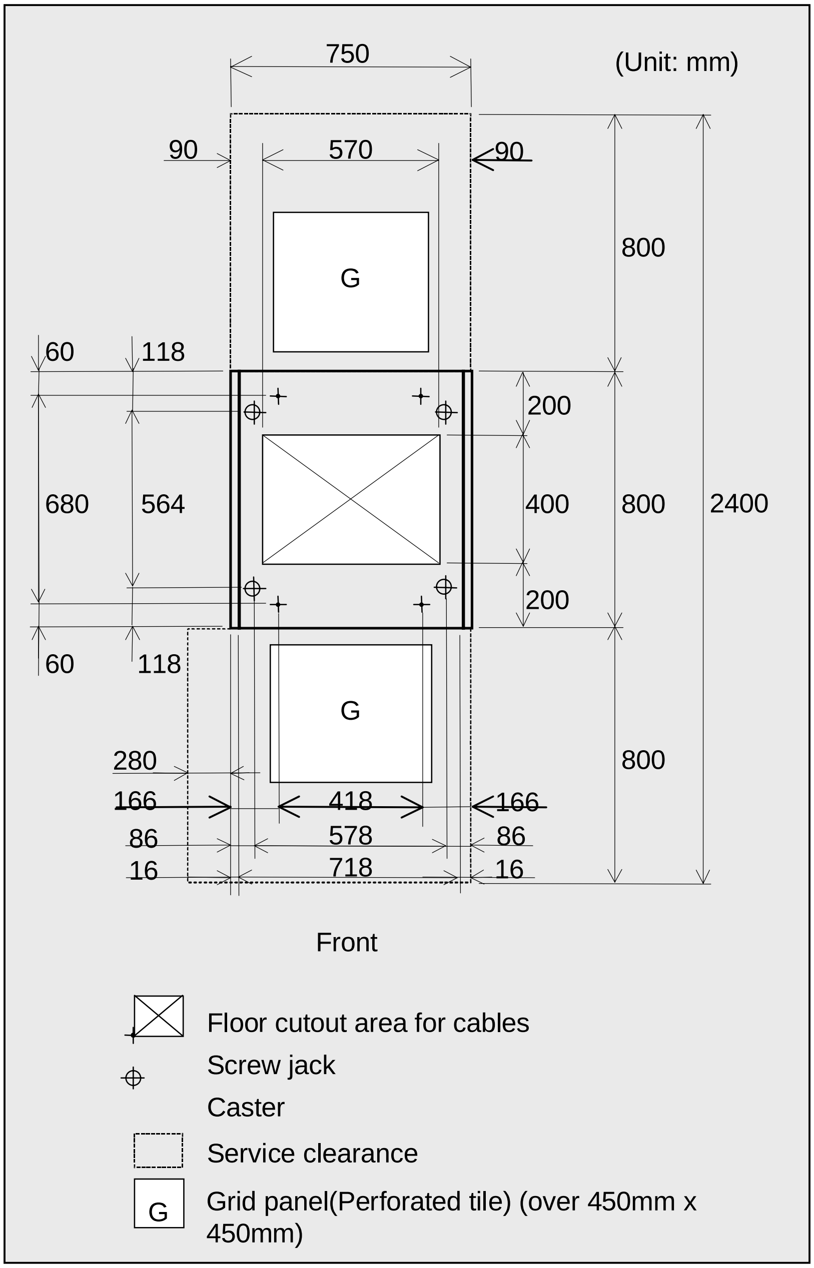

Figure 5.14 9960 Controller Frame Service Clearance and Cutouts (millimeters)

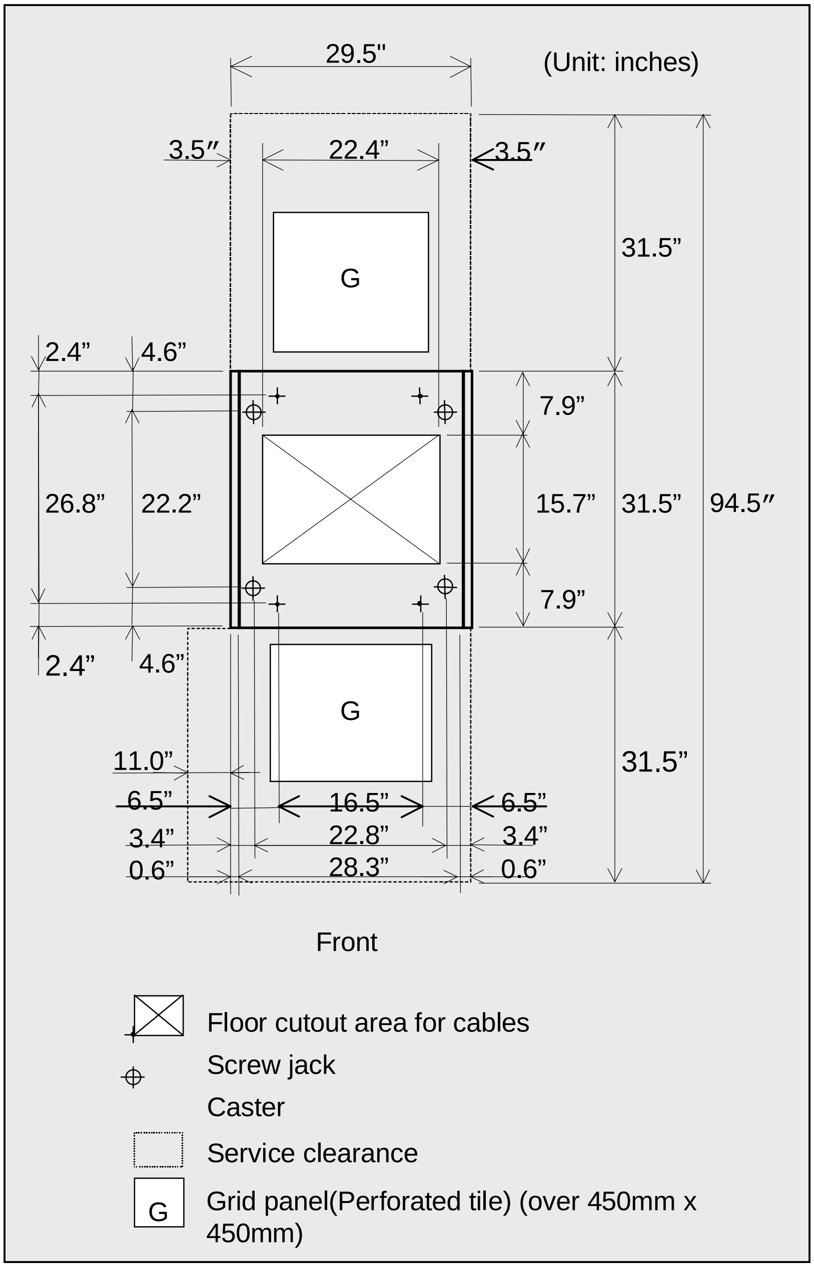

Figure 5.15 9960 Controller Frame Service Clearance and Cutouts (inches)

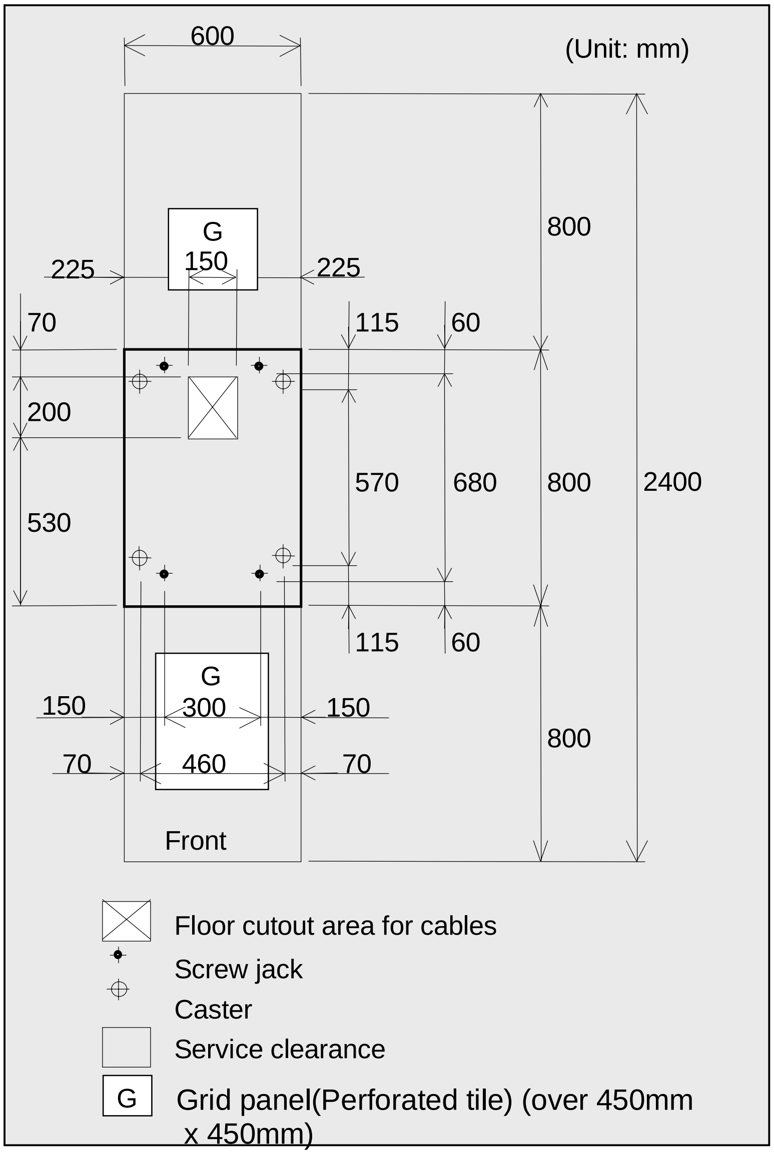

Figure 5.16 9960 Disk Array Unit Service Clearance and Cutout (millimeters)

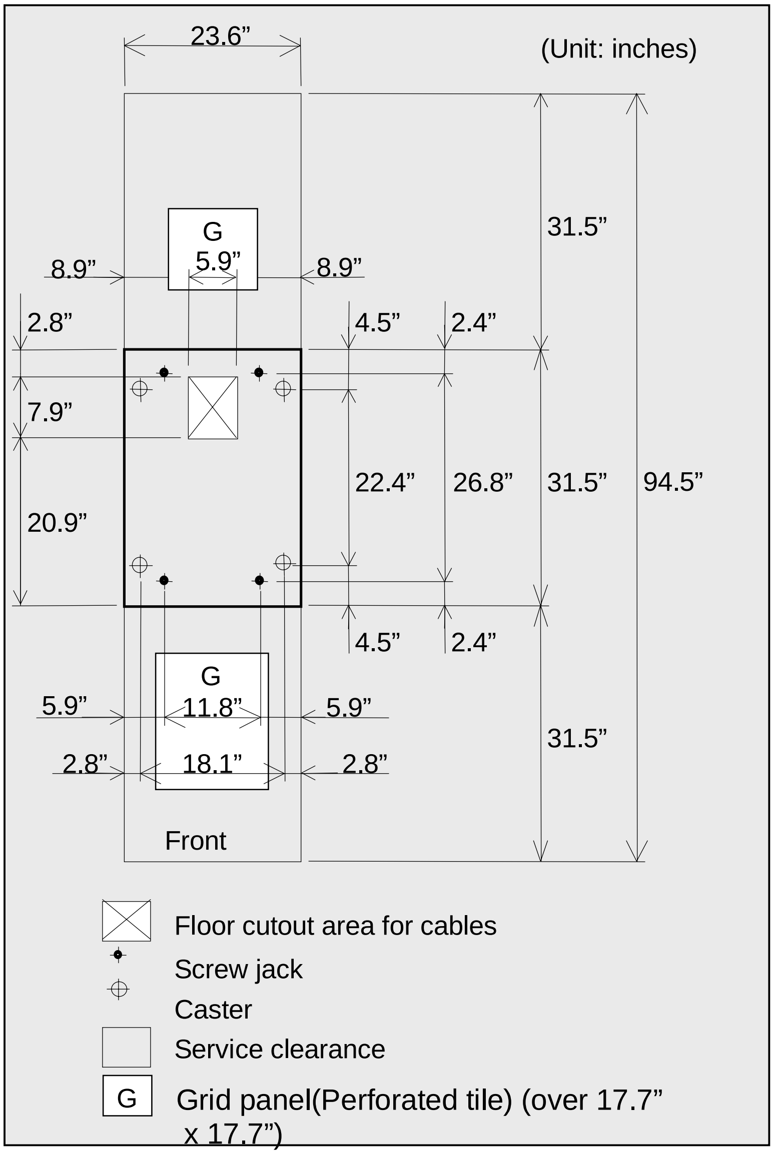

Figure 5.17 9960 Disk Array Unit Service Clearance and Cutout (inches)

5.5.2 Minimum Subsystem Disk Configuration

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

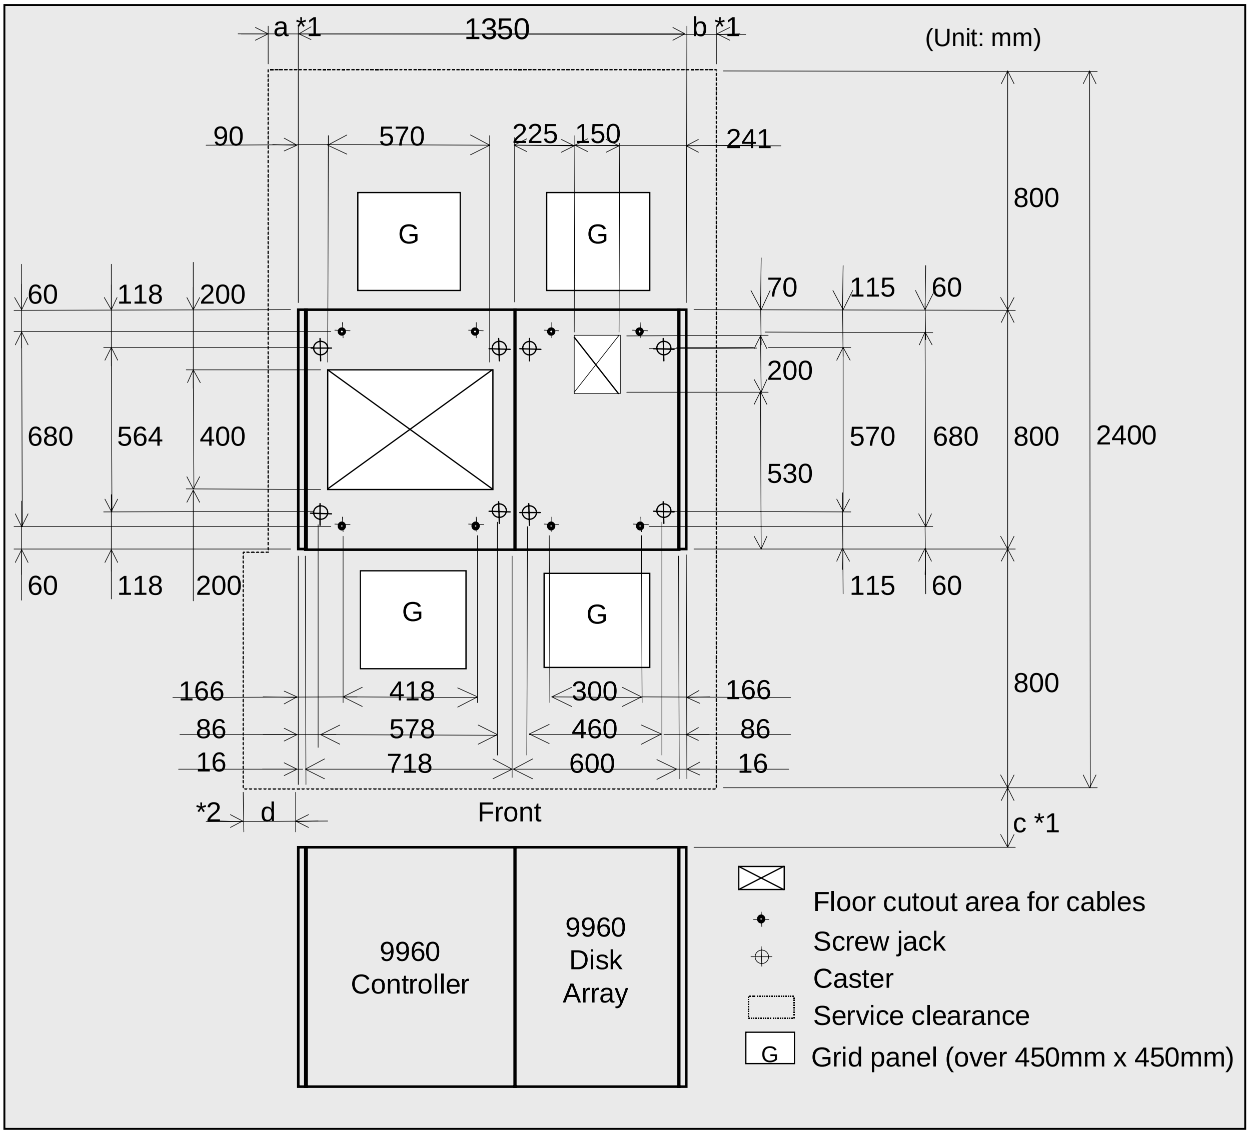

Figure 5.18 9960 Disk Subsystem Minimum Configuration (millimeters)

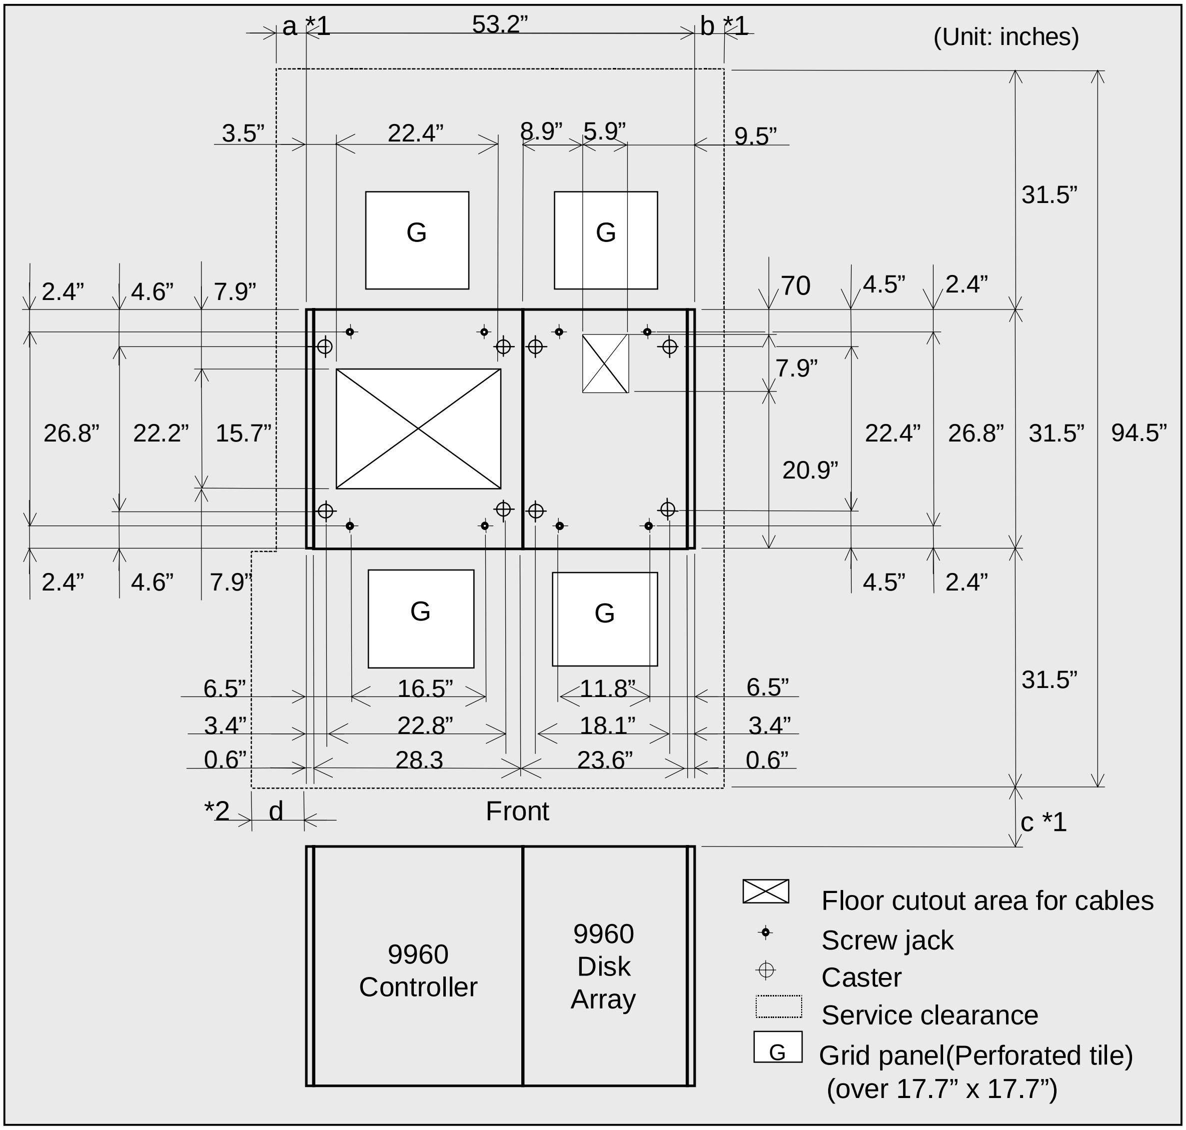

Figure 5.19 9960 Disk Subsystem Minimum Configuration (inches)

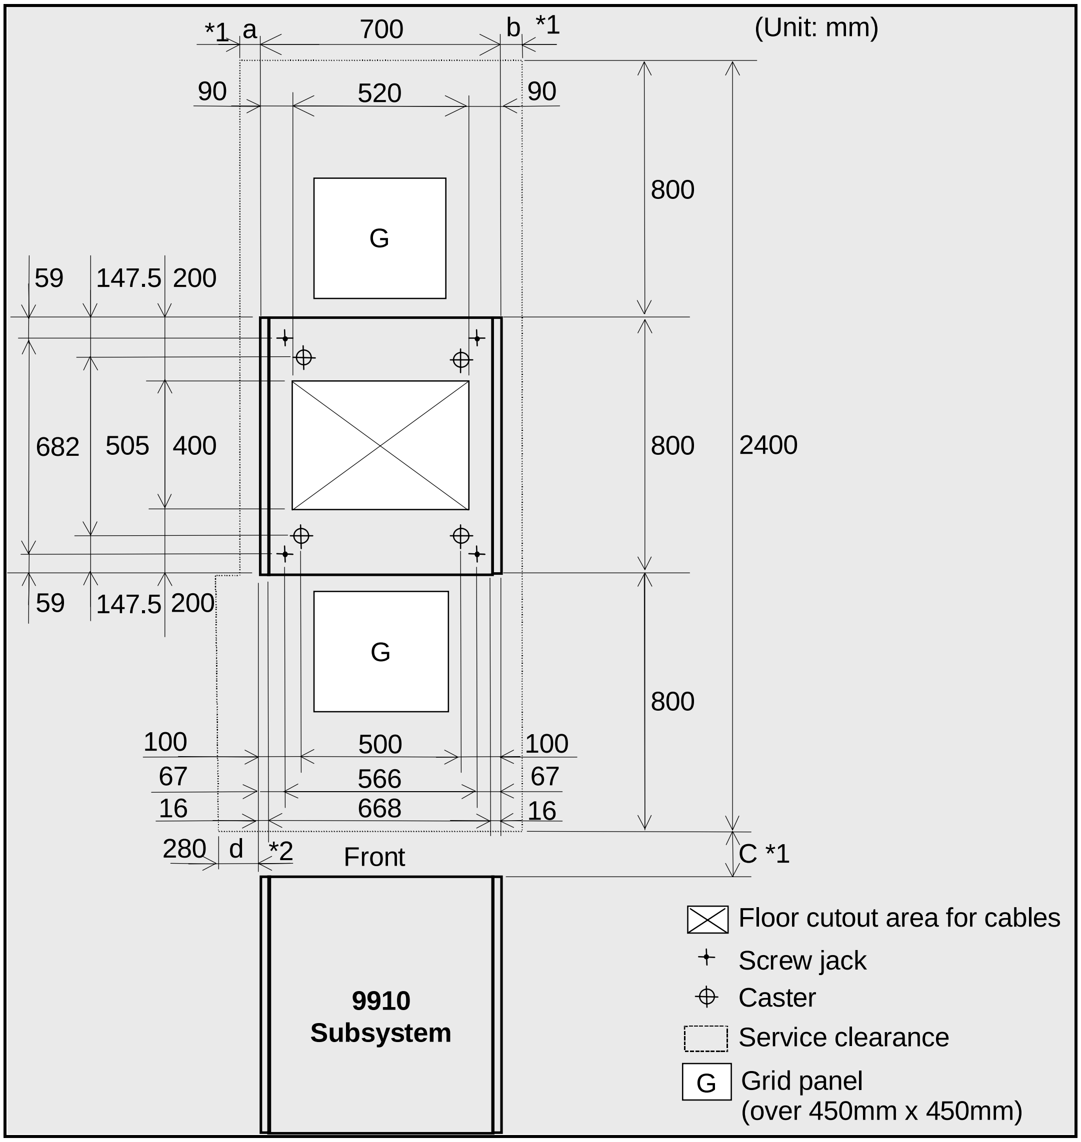

Figure 5.20 9910 Disk Subsystem All Configurations (millimeters)

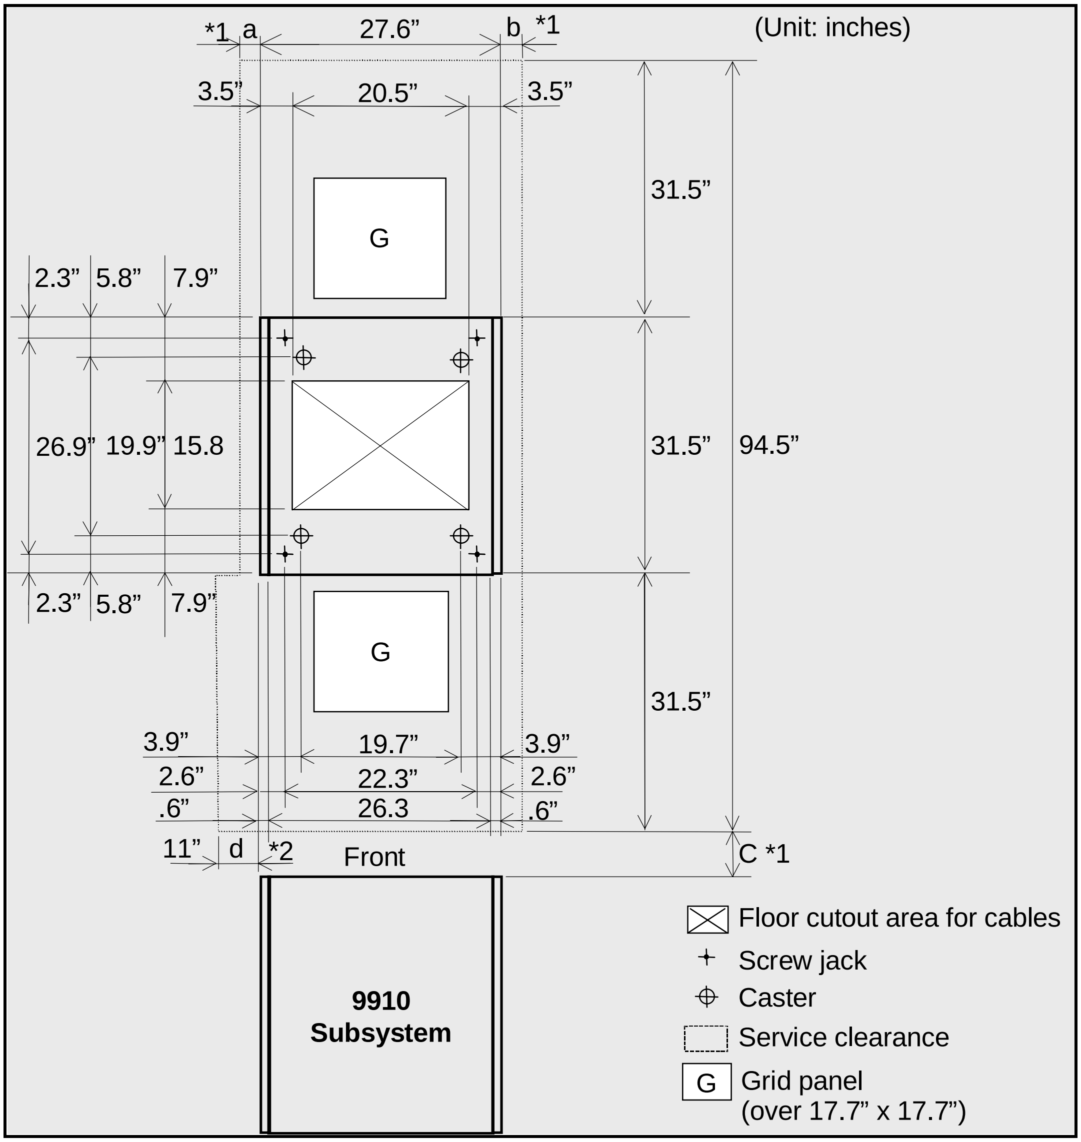

Figure 5.21 9910 Disk Subsystem All Configurations (inches)

5.5.3 Floor Load Rating

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

Table 5.10 Floor Load Rating for 9910 Subsystem

Floor load Rating Kg/m²(lb/ft²)	Required clearance (a+b) Clearance(C) in meters(feet)				
	C=0	C=0.2(.66)	C=0.4(1.31)	C=0.6(1.97)	C=1.0(3.28)
500(102.4)	0.6 m(2 ft)	0.3 m(1 ft)	0.1 m(.3 ft)	0	0
450(92.2)	0.8 m(2.6 ft)	0.5 m(1.6 ft)	0.3 m(1 ft)	0.1 m(.3 ft)	0
400(81.9)	1.0 m(3.3 ft)	0.7 m(2.3 ft)	0.5 m(1.6 ft)	0.3 m(1 ft)	0
350(71.7)	1.4 m(4.6 ft)	1.0 m(3.3 ft)	0.7 m(2.3 ft)	0.5 m(1.6 ft)	0.2 m(0.66 ft)
300(61.4)	2.0 m(6.6 ft)	1.5 m(4.9 ft)	1.1 m(3.6)	0.9 m(3 ft)	0.5 m(1.6 ft)

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

Floor load Rating Kg/m²(lb/ft²)	Required clearance (a+b) Clearance(C) in meters(feet)				
	C=0	C=0.2(.66)	C=0.4(1.31)	C=0.6(1.97)	C=1.0(3.28)
500(102.4)	0	0	0	0	0
450(92.2)	0.2 m(0.66 ft)	0	0	0	0
400(81.9)	0.5 m(1.6 ft)	0.3 m(1 ft)	0.1 m(.3 ft)	0	0
350(71.7)	0.9 m(3 ft)	0.7 m(2.3 ft)	0.5 m(1.6 ft)	0.3 m(1 ft)	0
300(61.4)	1.6 m(5.2 ft)	1.3 m(4.3 ft)	1.0 m(3.3 ft)	0.8 m(2.6 ft)	0.4 m(1.3 ft)

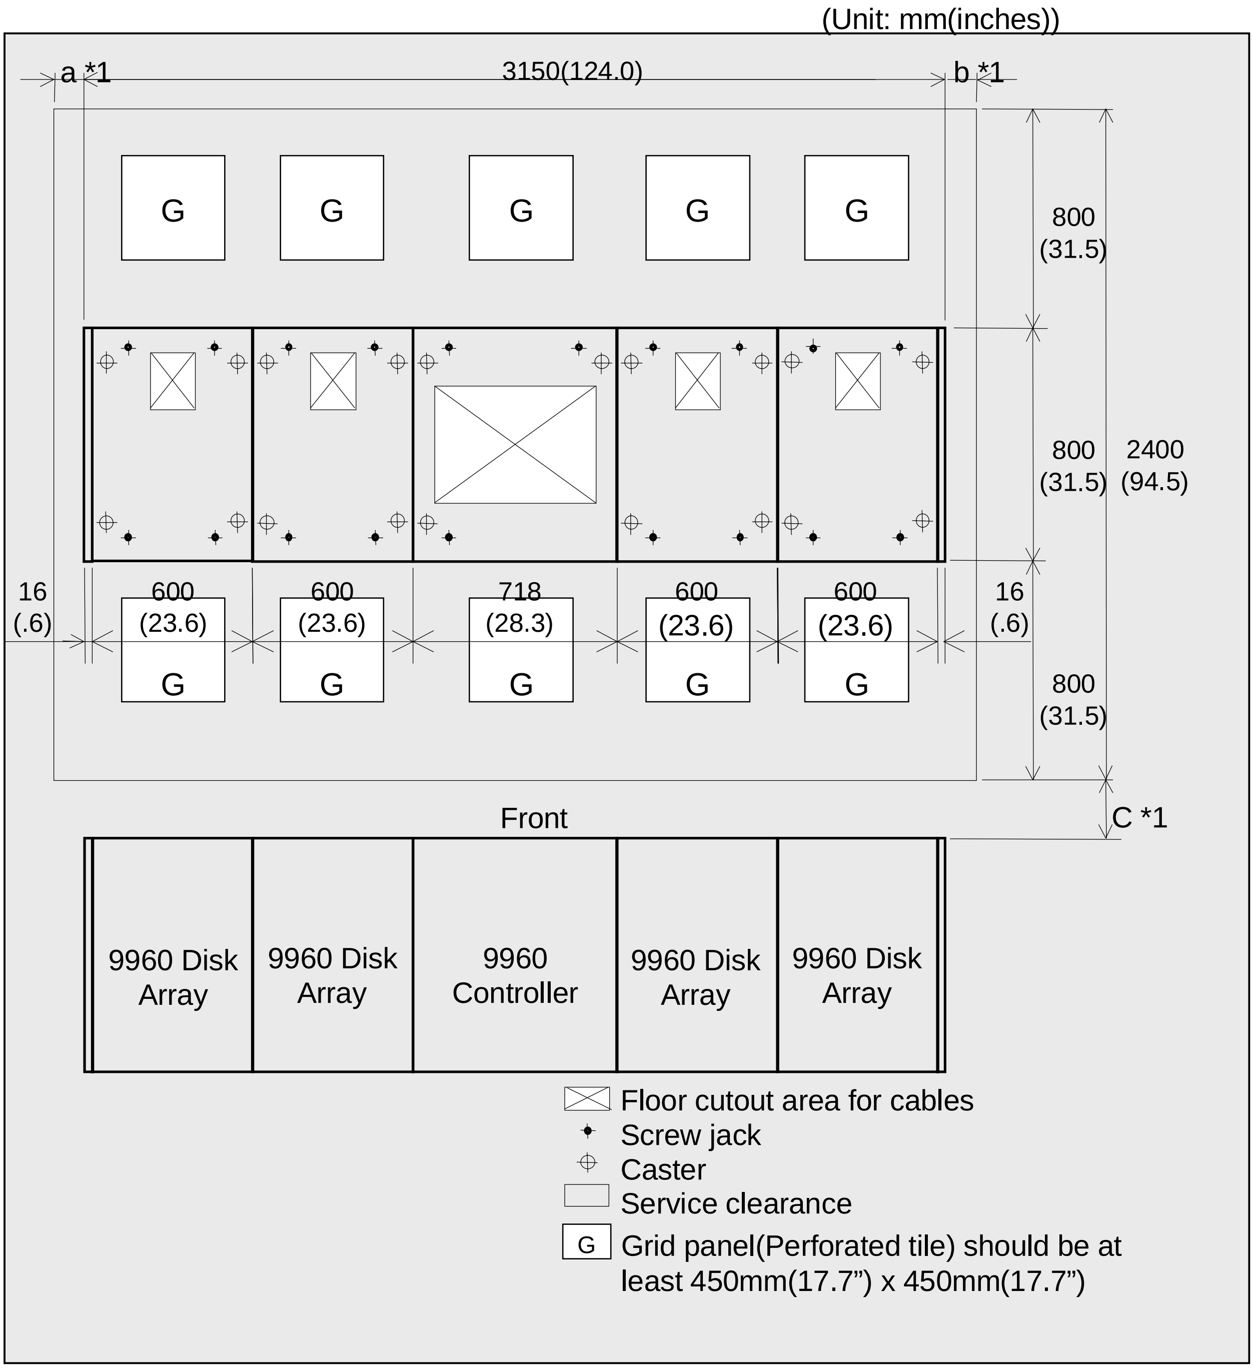

Figure 5.22 9960 Disk Subsystem with Controller and 4 Disk Arrays

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

Floor load Rating Kg/m2(lb/ft2)	Required clearance (a+b)				
	Clearance(C) in meters(feet)				
	C=0	C=0.2(.66)	C=0.4(1.31)	C=0.6(1.97)	C=1.0(3.28)
500(102.4)	0.2 m(0.66 ft)	0	0	0	0
450(92.2)	0.8 m(2.6 ft)	0.3 m(1 ft)	0	0	0
400(81.9)	1.5 m(4.9 ft)	1.0 m(3.3 ft)	0.5 m(1.6 ft)	0.2 m(0.66 ft)	0
350(71.7)	2.5 m(8.2 ft)	1.9 m(6.2 ft)	1.4 m(4.6 ft)	1.0 m(3.3 ft)	0.3 m(1 ft)
300(61.4)	4.1 m(13.5 ft)	3.3 m(10.8 ft)	2.7 m(8.9 ft)	2.1 m(6.9 ft)	1.3 m(4.3 ft)

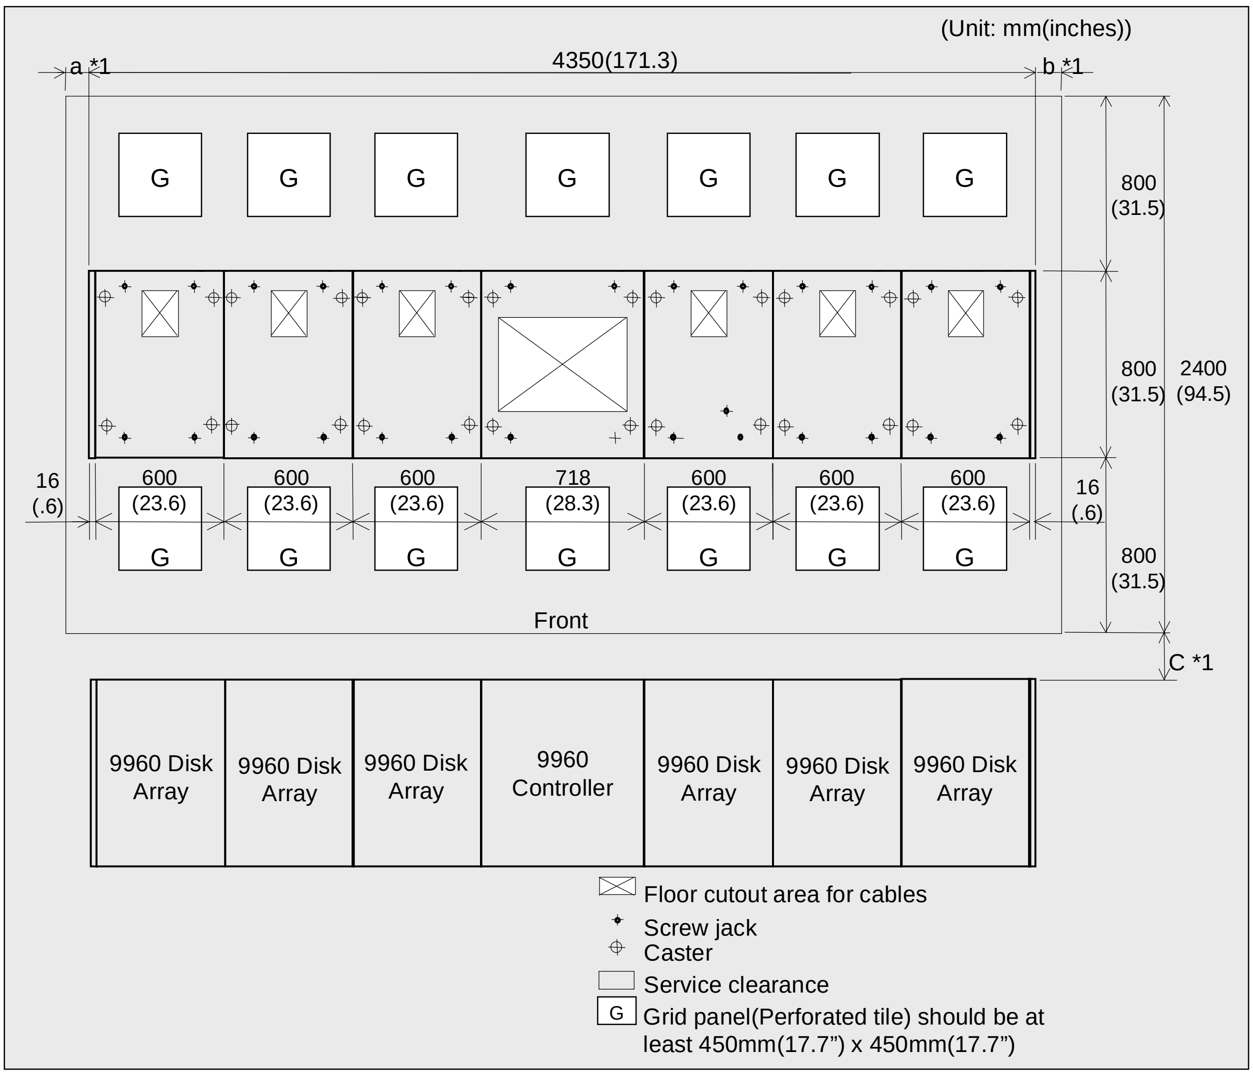

Figure 5.23 9960 Disk Subsystem with Maximum Configuration

Table 5.13 Floor Load Rating for 9960 Controller with Maximum Configuration

Floor load Rating kg/m²(lb/ft²)	Required clearance (a+b) Clearance(C) in meters(feet)				
	C=0	C=0.2(.66)	C=0.4(1.31)	C=0.6(1.97)	C=1.0(3.28)
500(102.4)	0.2 m(0.66 ft)	0	0	0	0
450(92.2)	0.9 m(3 ft)	0.3 m(1 ft)	0	0	0
400(81.9)	1.8 m(5.9 ft)	1.1 m(3.6 ft)	0.6 m(2 ft)	0.2 m(0.66 ft)	0
350(71.7)	3.2 m(10.5 ft)	2.4 m(7.9 ft)	1.7 m(5.6 ft)	1.1 m(3.6 ft)	0.3 m(1 ft)
300(61.4)	5.4 m(17.7 ft)	4.3 m(14.1 ft)	3.4 m(11.2 ft)	2.7 m(8.9 ft)	1.6 m(5.2 ft)

5.5.4 Cable Requirements

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

Table 5.14 Cable Requirements

Cable	Function
PCI cable	Connects 9900 to CPU power control interface.
ESCON® interface cables	Connects mainframe Host Systems, channel extenders or ESCON® Directors to 9900 ports. Multimode cables: • 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 2km per link) Mono/Single mode cables: • required on XDF links between ESCON® directors or IBM 9036 ESCON® remote control extenders. • Use laser light source • Yellow in color with gray duplex connectors • 8-10 micron. However, the most common is 9 micron.
Fibre Cables	Connects Open Systems to 9900 ports. Fibre cable type is 50 micron multimode.
Phone cable with RJ11 connector.	Connects phone line to 9900 SVP for Hi-Track.
10BaseT cable with RJ45 connector.	Connects remote console PC to 9900 subsystem.
10Base2 cable (RG58) with BNC connector.	Connects Remote console to 9900 subsystem and allows connection to multiple 9900's (up to 8) without using a hub. Requires a transceiver.

5.6 Channel Specifications and Requirements

Table 5.15 lists the port specifications for each 9900 channel option. Each 9960 subsystem supports up to four channel adapters and 32 interface ports. Each 9910 subsystem supports up to three channel adapters and 24 interface ports. Each adapter consists of two cards.

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

Table 5.15 Port Information

Item	ExSA™ 4-Port	ExSA™ 8-Port	Fibre 4-Port Short Wavelength	Fibre 8-Port Short Wavelength	Fibre 4-Port Long Wavelength	Fibre 8-Port Long Wavelength
Adapter model number	4S	8S	4GS	8GS	4GL	8GL
Max. data transfer rate (MB/s)	10/17	10/17	100	100	100	100
Number of ports/adapter	4	8	4	8	4	8
Number of MPs/adapter	4	8	4	8	4	8
Host interface	ExSA™	ExSA™	Fibre	Fibre	Fibre	Fibre
Number of adapter pairs per 9960 subsystem	1 to 4	1 to 4	1 or 4	1 to 4	1 to 4	1 to 4
Number of adapter pairs per 9910 subsystem	1 to 3	1 to 3	1 or 3	1 to 3	1 to 3	1 to 3
Max. interface cable length	43 km	43 km	500 m	500 m	10 km	10 km

5.7 Environmental Specifications and Requirements

5.7.1 Temperature and Humidity Requirements

Table 5.16 lists the temperature and humidity requirements for the 9900 subsystem. The recommended operational room temperature is 70–75°F (21–24°C). The recommended operational relative humidity is 50% to 55%.

Table 5.16 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	50(10)		50(10)		68(20)	

Notes:

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

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

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

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

5.7.2 Power Consumption and Heat Output Specifications

Table 5.17 lists the power consumption and heat output parameters for the 9900 control frame and array frames. These data generally apply to both 60-Hz and 50-Hz subsystems. Table 5.18 lists the power consumption and heat output parameters for the various configurations of the 9900 subsystem.

Table 5.17 9910/9960 Component Power and Heat Output Specifications

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

Table 5.18 9900 Subsystem Power and Heat Output Specifications

Subsystem Type	Subsystem Configuration	Power (kVA)	Heat (kW)
9960	Controller with 32 GB cache, 4 sets of 8-port Adapters and 1 fully populated Disk Array frame using 47 GB / 73 GB drives	6.13	5.75
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 2 fully populated Disk Array frames using 47 GB / 73 GB drives	9.85	9.18
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 3 fully populated Disk Array frames using 47 GB / 73 GB drives	13.57	12.61
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 4 fully populated Disk Array frames using 47 GB / 73 GB drives	17.29	16.04
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 5 fully populated Disk Array frames using 47 GB / 73 GB drives	21.01	19.47
	Controller with 32 GB cache, 4 sets of 8-port Adapters and 6 fully populated Disk Array frames using 47 GB / 73 GB drives	24.73	22.9
9910	16 GB cache, 3 sets of 8 port adapters, fully populated with 18 GB drives	3.06	2.81
	16 GB cache, 3 sets of 8 port adapters, fully populated with 147 GB or 73 GB drives	3.34	3.10

5.7.3 Air Flow Requirements

The 9900 subsystem is air cooled. Air must enter the subsystem through the air flow intakes at the sides and bottoms of the frames and must be exhausted out of the tops and sides, so it is very important that the air intakes and outlets remain clear. Hitachi Data Systems recommends that under-floor air cooling has a positive pressure and meets the specifications listed in Table 5.19. For subsystems located at elevations from 3,000 to 7,000 feet (900 to 2,100 meters) above sea level, decrease the maximum air temperature by two degrees for each 1,000 feet (300 meters) above 3,000 feet (900 meters).

Table 5.19 Internal Air Flow

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

5.7.4 Vibration and Shock Tolerances

Table 5.20 lists the vibration and shock tolerance data for the 9900 subsystem. The 9900 can be subjected to vibration and shock up to these limits and still perform normally. The user should consider these requirements if installing the 9900 near large generators located on the floor above or below the 9900 subsystem. Generators or any other source of vibration, if not insulated or shock-mounted, can cause excessive vibration that may affect the subsystem.

Table 5.20 Vibration and Shock Tolerances

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

Chapter 6 Troubleshooting

6.1 Troubleshooting

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

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

Table 6.1 Troubleshooting

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

6.2 Service Information Messages (SIMs)

The 9900 subsystem generates service information messages (SIMs) to identify normal operations (e.g., HRC pair status change) as well as service requirements and errors or failures. Table 6.2 lists the basic SIMs generated by the 9900 subsystem. SIMs can be generated by the CHIP and ACP microprocessors and by the SVP. All SIMs issued by the 9900 are stored on the SVP for use by Hitachi Data Systems personnel, logged in the SYS1.LOGREC dataset of the S/390® host system, logged on the Remote Console PC (the Remote Console PC must be running), and reported via SNMP to the open-system host. The Remote SIM (R-SIM) feature of the 9900 Remote Console PC enables users to view the SIMs reported by the attached 9900 subsystems on the Remote Console PC. Each time a SIM is generated, the amber Message LED on the 9900 control panel turns on and the Remote Console PC displays a warning message, even when the remote console software is not running. The Hi-Track® remote maintenance tool also reports all SIMs to the Hitachi Data Systems Support Center.

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

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

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
SIM SSB	00	90	10	00	00	00	8F	E0	44	10	00	04	00	80	04	0C	69	00	00	00	00	02	30	70	05	10	42	C0	F1	00	02

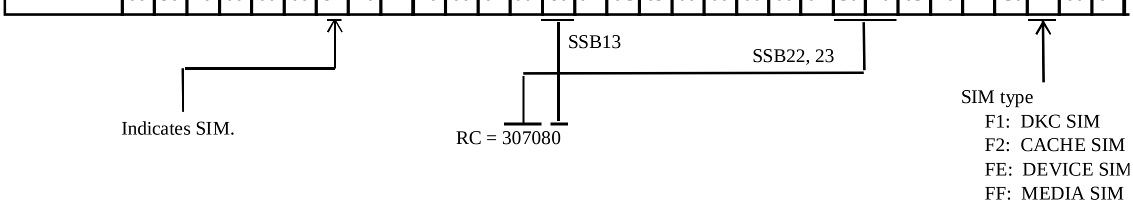

Figure 6.1 Typical 9900 SIM Showing Reference Code and SIM Type

Table 6.2 lists the basic SIMs generated by the 9900 subsystem. For information on feature-specific SIMs, please see the *Hitachi Freedom Storage™ Lightning 9900™ Remote Console User's Guide* (MK-90RD003). For further information and assistance with SIMs, please call the Hitachi Data Systems Support Center.

Table 6.2 SIM Reference Codes (continues on the next page)

Error		REF CODE			SIM	Error Level	Host Report?
		36	37	2D	3C		
LCP/FCP error	LCM Hardware error	21	3Z	XY	F1	Moderate	Yes
	ADP permanent error	21	70	X0	F1	Moderate	Yes
	ADP temporary error	21	71	XY	F1	Service	No
	ADP blocking	21	72	XY	F1	Moderate	Yes
	AL_PA value conflict	21	90	XY	F1	Service	No
CHA Processor error	CHK1A threshold over	30	70	XY	F1	Service	No
	CHK1B threshold over	30	71	XY	F1	Service	No
	CHK3 threshold over	30	72	XY	F1	Service	No
	Processor blocking	30	73	XY	F1	Moderate	Yes
	FM threshold over	30	74	XY	F1	Service	No
	FM error	30	75	XY	F1	Moderate	Yes
	Incorrect SUM value of FM	30	76	XY	F1	(*4)	No
	LDEV blockade (Effect of processor blockade)	30	90	XY	F1	Serious	Yes
	CHF PCB exchange impossible (No PERFORMANCE FIBRE)	31	97	XY	F1	Service	No
	CHF PCB exchange impossible (No supported Micro-Program)	31	98	XY	F1	Service	No
	P/S OFF impossible	38	8F	00	F1	Moderate	No
	P/S OFF impossible (Device reserved)	38	9F	00	F1	Moderate	No
	Undefined Package is mounted	39	90	XY	F1	Moderate	No
	V-R or serial number is inconsistent	39	91	XY	F1	Moderate	No
	Replace failed	39	93	XY	F1	Moderate	No
	Micro-program version up	39	94	X0	F1	Service	No
	Micro-program version up impossible	39	95	XY	F1	Service	No
CHA Processor error	CHF PCB exchange impossible (No PERFORMANCE FIBRE)	39	97	XY	F1	Service	No
	CHF PCB exchange impossible (no supported microprogram)	39	98	XY	F1	Service	No
	CHF PCB is exchanged to PERFORMANCE FIBRE	39	99	XY	F1	Service	No
	Warning of SM DISABLE (MPA detected)	39	9A	XY	F1	Moderate	No

Table 6.2 SIM Reference Codes (continued)

Error		REF CODE 36	37	2D	SIM 3C	Error Level	Host Report?
CHA Processor error	SMA slave error	39	9B	XY	F1	Moderate	No
	MPA slave error	39	9C	XY	F1	Moderate	No
	Injustice DC voltage control	39	9D	X0	F1	Moderate	No
	Injustice CE MODE	39	9E	X0	F1	Moderate	No
	CHT PCB exchange impossible (4GL ↔ 4GS)	39	9F	X0	F1	Moderate	No
DKA Processor error	CHK1A threshold over	31	70	XY	F1	Service	No
	CHK1B threshold over	31	71	XY	F1	Service	No
	CHK3 threshold over	31	72	XY	F1	Service	No
	Processor blocking	31	73	XY	F1	Moderate	Yes
	FM threshold over	31	74	XY	F1	Service	No
	FM error	31	75	XY	F1	Moderate	Yes
	Incorrect SUM value of FM	31	76	XY	F1	(*1)	No
	LDEV blockade (Effect of processor blockade)	31	90	XY	F1	Serious	Yes
	P/S OFF impossible	3C	8F	00	F1	Moderate	No
	P/S OFF impossible (Device reserved)	3C	9F	00	F1	Moderate	No
	Undefined Package is mounted	3D	90	XY	F1	Moderate	No
	V-R or serial number is inconsistent	3D	91	XY	F1	Moderate	No
	Replace failed	3D	93	XY	F1	Moderate	No
	Micro-program version up	3D	94	X0	F1	Service	No
	Micro-program version up impossible	3D	95	XY	F1	Service	No
	Warning of SM DISABLE (MPA detected)	3D	9A	XY	F1	Moderate	No
	SMA slave error	3D	9B	XY	F1	Moderate	No
	MPA slave error	3D	9C	XY	F1	Moderate	No
	Injustice DC voltage control	3D	9D	X0	F1	Moderate	No
	Injustice CE MODE	3D	9E	X0	F1	Moderate	No
	FSW LED BUS test	3D	AZ	XY	F1	Service	No
	Tachyon error	3D	B0	XY	F1	Service	No
SHSN error	Logical path blockade (CHA Processor)	32	8Z	XX	F1	Moderate	No
	Logical path blockade (DKA Processor)	33	8Z	XX	F1	Moderate	No
CHA, CHK2 (processor)	RCHA temporary error	3F	84	X0	F1	Service	No
	RCHA blocking	3F	85	X0	F1	Serious	Yes

Table 6.2 SIM Reference Codes (continued)

Error		REF CODE 36	37	2D	SIM 3C	Error Level	Host Report?
CHS, CHK2 (processor)	SCP path temporary error	3F	86	XY	F1	Service	No
	SCP path blockade	3F	87	XY	F1	Moderate	Yes
DKA CHK2 (processor)	Pinned slot	CF	4X	XX	F1	Moderate	No
	DRR temporary error	CF	80	XY	F1	Service	No
	FCA temporary error	CF	81	XY	F1	Service	No
	DRR blocking	CF	82	XY	F1	Moderate	Yes
	FCA blocking (0-3)	CF	83	XY	F1	Moderate	Yes
	Fibre port blocking (Effect of PATH INLINE failed)	CF	84	XY	F1	Moderate	No
	LDEV blockade (Effect of FCA blockade)	CF	90	XY	F1	Serious	Yes
Cache data error	Pinned slot	FF	4X	XX	F2	Moderate	No
Cache error	Area is volatized	FF	CD	0X	F2	Service	No
	Package is volatized	FF	CE	0X	F2	Service	No
	Module group is volatized	FF	CF	YX	F2	Service	No
	Correctable error	FF	F0	YX	F2	Service	No
	Cache temporary error	FF	F1	YX	F2	Service	Yes
	Module group blocking	FF	F2	YX	F2	Moderate	Yes
	Package blocking	FF	F3	0X	F2	Moderate	Yes
	Area blocking	FF	F4	0X	F2	Serious	Yes
	Both area failed	FF	F5	0X	F2	Moderate	No
	Injustice DC voltage control	FF	F6	0X	F1	Moderate	No
	Injustice CE MODE	FF	F7	0X	F1	Moderate	No
Shared Memory Error	Loss of duplicated information	FF	DE	XX	F1	Service	No
	One side Area is volatized	FF	DF	0X	F1	Service	No
	Correctable error	FF	E0	YX	F1	Service	No
	Uncorrectable error	FF	E1	0X	F1	Service	Yes
	Area blocking	FF	E2	0X	F1	Serious	Yes
	Real memory size inconsistent	FF	E3	XY	F1	Serious	No
	Replace failed	FF	E4	0X	F1	Serious	No
	Both side invalid	FF	E5	00	F1	Acute	No
	Configuration information compare error	FF	E6	00	F1	Acute	No
	Shared memory is volatized.	FF	E7	00	F1	Serious	No
	Configuration unmatch	FF	E8	00	F1	Acute	No
	MAIN DIFF lost	FF	E9	00	F1	Serious	No

Table 6.2 SIM Reference Codes (continued)

Error		REF CODE 36	37	2D	SIM 3C	Error Level	Host Report?
	CHK3 threshold over	FF	EC	0X	F1	Service	No
	Area temporary blocking	FF	EE	0X	F1	Service	Yes
	Rebooted without volatilization after an instantaneous down	FF	EF	00	F1	Service	No
Drive error (normal R/W)	Drive port temporary error (Drive path: Boundary 0)	DF	6X	XX	FE	Service	No
	Drive port temporary error (Drive path: Boundary 1)	DF	7X	XX	FE	Service	No
	Drive temporary error	EF	AX	XX	FE	Service	No
	Drive media error	43	4X	XX	FF	Service	No
	Drive port blockade (Drive path: Boundary 0)	DF	8X	XX	FE	Moderate	Yes *2
	Drive port blockade (Drive path: Boundary 1)	DF	9X	XX	FE	Moderate	Yes *2
	LDEV blockade (Drive path: Boundary 0 Effect of SCSI blockade)	DF	AX	XX	FE	Serious	Yes *2
	LDEV blockade (Drive path: Boundary 1) Effect of SCSI blockade)	DF	BX	XX	FE	Serious	Yes *2
	Drive blockade (drive)	EF	1X	XX	FE	Serious	Yes *2
	Drive blockade (Effect of Dynamic sparing normal end)	EF	2X	XX	FE	Service	Yes *2
	Side file 40% over	49	9X	XX	FE	Service	Yes *3
	LDEV blockade (Effect of drive blockade)	EF	9X	XX	FE	Serious	Yes *2
	Drive blockade (media)	43	CX	XX	FF	Serious	Yes *2
	Correction copy start	45	1X	XX	FE	Service	Yes *2
	Correction copy normal end	45	2X	XX	FE	Service	Yes *2
	Correction copy abnormal end	45	3X	XX	FE	Serious	Yes *2
	Correction copy discontinued	45	4X	XX	FE	Service	No
	Correction copy warning end* (With blockade LDEV or some error)	45	5X	XX	FE	Service	Yes *2
	Dynamic sparing start (Drive copy)	46	1X	XX	FE	Service	Yes *2
	Dynamic sparing normal end (Drive copy)	46	2X	XX	FE	Service	Yes *2
	Dynamic sparing abnormal end (Drive copy)	46	3X	XX	FE	Moderate	Yes *2

Table 6.2 SIM Reference Codes (continued)

Error		REF CODE			SIM	Error Level	Host Report?
		36	37	2D	3C		
	Dynamic sparing discontinued	46	4X	XX	FE	Service	No
	Dynamic sparing warning end* (With blockade LDEV or some error) (Drive copy)	46	5X	XX	FE	Service	Yes *2
	Pinned slot	EF	4X	XX	FE	Moderate	No
	FIBRE temporary error (Drive path: Boundary 0)	50	0X	XX	FE	Service	No
	FIBRE temporary error (Drive path: Boundary 1)	50	1X	XX	FE	Service	No
	Drive temporary error	50	2X	XX	FE	Service	No
	Drive media error	50	3X	XX	FF	Service	No
CSW error	Injustice DC voltage control	FF	20	X0	F1	Moderate	No
	Injustice CE MODE	FF	21	X0	F1	Moderate	No
SVP Interface Error (CHA Side)	Ethernet error for SVP	14	00	X0	F1	Moderate	Yes
	SIM transfer failure to SVP	14	01	X0	F1	Serious	Yes
SVP Interface Error (DKA Side)	Ethernet error for SVP	15	00	X0	F1	Moderate	Yes
	SIM transfer failure to SVP	15	01	X0	F1	Serious	Yes
Power Error	HDU power off	AC	50	XY	F1	Moderate	Yes
	HDU power recovered	AC	51	XY	F1	Service	No
	DKC was set to power error mode	AC	60	00	F1	Moderate	No
	DKC was released from power error mode	AC	61	00	F1	Service	No
	When DKC was set to power error mode, urgent destaging start succeeded.	AC	62	00	F1	Service	No
	When DKC was set to power error mode, urgent destaging start failed.	AC	63	00	F1	Moderate	No

*1: Warning end: “1” indicates that LDEV blockade status is detected, and copy process for the LDEV is skipped; “2” indicates that some failure is detected, and the copy process is complete.
*2: When the drive is 3880 emulation mode, the SIM is not reported to host.
*3: The SIM is not reported to SVP.
*4: If this is a ROM error, the error level is “Moderate”. Otherwise, it is “Service”

6.3 Calling the Hitachi Data Systems Support Center

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including the circumstances surrounding the error or failure, the exact content of any messages displayed, and the severity levels and reference codes of the SIMs. The worldwide Hitachi Data Systems Support Centers are:

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

Appendix A Acronyms and Abbreviations

ACP	array control processor
BS	basic (power) supply
BSA	bus adapter
BTU	British thermal unit
°C	degrees centigrade/Celsius
ca	cache
CFW	cache fast write
CH	channel
CHP	channel processor or channel path
CHIP	client-host interface processor
CHPID	channel path identifier
CKD	count key data
CL	cluster
CPU	central processing unit
CSA	Canadian Standards Association
CU	control unit
CVS	custom volume size (also called Virtual LVI/LUN)
DASD	direct access storage device
DCR	dynamic cache residency (also called FlashAccess)
DFDSS	Data Facility Dataset Services
DFSMS	Data Facility System Managed Storage
DFW	DASD fast write
DKC	disk controller
dr	drive
DSF	Device Support Facilities
DTDS+	Disaster Tolerant Disk System Plus
DW	duplex write
DWL	duplex write line
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 HMDE 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 HMDE software)
FDR	Fast Dump/Restore

F/M	format/message
FWD	fast wide differential
GB	gigabytes
GLM	gigabyte link module
GUI	graphical user interface
HACMP	High Availability Cluster Multi-Processing
HAGEO	High Availability Geographic
HARC	Hitachi Asynchronous Remote Copy (also called HRCA)
HCD	hardware configuration definition
HCPF	Hitachi Concurrent Processing Facility
HIHSM	Hitachi Internal Hierarchical Storage Manager
HMBR	Hitachi Multiplatform Backup/Restore
HMDE	Hitachi Multiplatform Data Exchange
HMRCF	Hitachi Multiple RAID Coupling Feature (also called ShadowImage)
HMRS	Hitachi Multiplatform Resource Sharing
HODM	Hitachi Online Data Migration
HOMRCF	Hitachi Open Multiple RAID Coupling Facility (also called Open ShadowImage)
HORCA	Hitachi Open Remote Copy Asynchronous
HORC	Hitachi Open Remote Copy
HRC	Hitachi Remote Copy
HRCA	Hitachi Asynchronous Remote Copy
HSN	Hierarchical Star Network
HVR	Hitachi Volume Relocation
HXRC	Hitachi Extended Remote Copy
Hz	Hertz
ICKDSF	A DSF command used to perform media maintenance
IDCAMS	access method services (a component of Data Facility Product)
IML	initial microprogram load
in.	inch(es)
IO, I/O	input/output (operation or device)
IOCP	input/output configuration program
JCL	job control language
kB	kilobyte
kcal	kilocalorie
kg	kilogram
km	kilometer
kVA	kilovolt-ampere
kW	kilowatt
LAN	local area network
lb	pound
LD	logical device
LDEV	logical device
LED	light-emitting diode

LPAR	logical partition
LCP	link control processor, local control port
LRU	least recently used
LUN	logical unit, logical unit number
LVI	logical volume image (also called device emulation)
m	meters
MB	megabytes
MIH	missing interrupt handler
mm	millimeters
MP	microprocessor
MPLF	Multi-Path Locking Facility
MR	magnetoresistive
ms, msec	milliseconds
MVS	Multiple Virtual Storage (including MVS/370, MVS/ESA, MVS/XA)
NVS	nonvolatile storage
OEM	original equipment manufacturer
OFC	open fibre control
ORM	online read margin
OS	operating system
P/DAS	PPRC/dynamic address switching
PC	personal computer system
PCI	power control interface
PPRC	Peer-to-Peer Remote Copy
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
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
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
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
VSE	Virtual Storage Extension (an IBM S/390® operating system)
VTOC	volume table of contents
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 (US) and metric measurement systems.

Table B.1 Unit Conversions for Standard (US) 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)