


Hitachi Lightning 9900™ Series Architecture Guide



The Only Switched Storage Architecture in the World



Hitachi Data Systems — Built for fast times™

HITACHI
DATA SYSTEMS

Hitachi Lightning 9900™ Series Architecture Guide



The Only Switched Storage Architecture in the World

The Lightning 9900 Series Packaging

Front-end Design of the Hitachi Lightning 9900 Series

Hi-Star Architecture – An Internal Switched Hierarchical Star Network

High Performance Back-end Design

Capacities and RAID Architecture


Hitachi High-Availability Software Improves Application Service Levels

Performance Considerations

The Lightning 9900 Series Software Solutions are the Best in the World

Centralized and Open Systems Management Software Simplifies IT Operations

Professional Services and Support



Hitachi Data Systems — Built for fast times™

HITACHI
DATA SYSTEMS

Table of Contents

Table of Contents	iii
Table of Figures	vii
Table of Tables	viii
Chapter 1 – Introduction	1
The Only Switched-Fabric Storage Architecture in the World	1
Lightning 9900 Series Product Offerings.....	1
An Overview of the Hitachi Lightning 9900 Series Architecture	2
Comparison to the Shared Bus Architecture of the Hitachi Freedom Storage 7700E.....	4
The Lightning 9900 has an advanced availability profile compared to competitive products.....	4
Other Information Sources Available from Hitachi Data Systems	5
Chapter 2 – The Lightning 9900 Series Packaging	7
Model 9960 Packaging	7
Model 9960 Control Frame.....	7
Model 9960 Array Frames	8
Model 9910 Packaging	9
The System Monitoring Network of the Lightning 9900 Series	10
Chapter 3 – Front-end Design of the Hitachi Lightning 9900 Series	11
CHIPs Support Fibre Channel Host and ESCON Host Connectivity and Scalability.....	11
About the Fibre Channel Interface Boards.....	12
Bandwidth of Supported Host Interfaces.....	14
Number of Host Interfaces Supported.....	14
Types of Hosts Supported by Direct Fibre Channel Host Attachment by the Lightning 9900 Series.....	14
Chapter 4 – Hi-Star Architecture – An Internal Switched Hierarchial Star Network	15
The Internal Hi-Star Architecture	15
The Cache Hierarchical Star Network	15
The Cache Switch Provides Non-blocking Switched Access to Cache	16
The Path Bandwidth Between the CSW and CHIPs or ACPs is 6.4GB/sec	16
The Path Bandwidth Between the CSW and the Cache Adapter (CAs) is 3.2GB/sec	16
Advantages of a Switched Architecture Compared to a Shared Bus Architecture	16
The 9960 Provides 32 GB of Fully Addressable Cache	18

Advanced Cache Algorithms of the Lightning 9900 Series	18
Hitachi FlashAccess Allows Datasets to be Permanently Placed in Cache	18
Contiguous Block Allocation for High-speed Sequential Access	18
Read-ahead for High-performance Sequential Reads	18
The Control Memory Hierarchical Star Network	19
Chapter 5 – High Performance Back-end Design	21
High-Performance Back-end Design	21
New Fibre Channel Disk Drives	21
Maximum Configuration of ACPs and Disks	22
The Importance of Command Tag Queuing (CTQ)	22
Chapter 6 – Capacities and RAID Architecture	23
Storage Capacity of the Lightning 9900 Series	23
Advantages of the Lightning 9900 Series RAID Hardware	23
RAID-1+	23
RAID-5 Distributed Parity	23
RAID Rebuild Capability	23
High-density Disk Drive Assemblies (HDDs)	24
Chapter 7 – Hitachi High-Availability Software Improves Application Service Levels	25
An Overview of High Availability (HA) Computing Software.....	25
Software products contribute to High-Availability Computing	25
The Lightning 9900 has an Advanced Availability Profile compared to Competitive Products	25
Open Systems High-Availability Middleware Support	26
1 - Alternate Pathing middleware switches the I/O load in the event of path failure	27
2 - Hitachi Dynamic Link Manager provides path failover and load balancing	27
3 - Host Failover.....	27
4 - Parallel Database Clustering	28
Standard Hot-swap and Redundant Power Supplies	28
Standard Redundant Fans for Cooling.....	29
Unmatched Competitive Advantage in RAID Choices	29
Chapter 8 – Performance Considerations	31
Truly Scalable Performance	31
The Advantages of the Lightning 9900 Series Hi-Star Architecture Over a Shared Bus Architecture.....	31
Internal Switching Overcomes the Limitations of Shared Bus Architectures	32

Test Results	32
Chapter 9 – The Lightning 9900 Series Software Solutions Are the Best in the World	33
Increased IT service levels in availability and performance through Business Continuity Planning (BCP) software and services.....	34
Remote Copy, Data Duplication, Data Migration Defined	35
Remote Copy	35
PiT Copy	35
Data Migration.....	35
Business Goals Served by Hitachi Copy Software	36
Chapter 10 – Centralized and Open Systems Management Software Simplifies IT Operations	37
A powerful and open systems management philosophy	37
HiCommand allows systems management of Hitachi Storage and software through the enterprise’s vendor of choice	37
The components of HiCommand	39
Overview of how the Systems Management Software supports business objectives	39
Hitachi Resource Manager 9000	40
Service Level Management	40
Deployment Management	40
Compliance Management.....	40
SNMP support ensures compliance with all major systems management platforms.....	40
Hitachi Graph-Track takes the guess-work out of performance management.....	40
Hitachi LUN Manager simplifies configuration management and reduces staffing costs	41
Hitachi Virtual Logical Volume Image Manager allows very large to very small virtual volume configuration	42
Hitachi FlashAccess allows data to be “locked and unlocked” in cache on-the-fly	42
Hitachi CruiseControl.....	42
Storage Maximizer	43
Hitachi SANtinel.....	43
Transaction Processing Facility/Multi-Path Lock Facility (TPF/MPLF)	43
Chapter 11 – Professional Services and Support	45
Hitachi Data Systems is consistently ranked Number One in the Industry	45
Professional Services Overview.....	45
Cost of Risk Analysis methodology (CORA).....	45
Data Protection Services	45
Data Protection Services Lab	46

Remote Copy Assessment and Implementation Service	46
Hitachi NanoCopy Implementation Service.....	46
Continuous Availability Service	47
On-line Data Migration Service	47
SAN Services	47
Hitachi Data Systems Interoperability Laboratory Service	48
Enterprise storage and availability management services	48
Software Portfolio Review and Analysis Service	48
Appendix A – System Configurations for the Performance Experiment	49
Glossary of Terms	51

Table of Figures

Figure 1	The 9960 provides for exploding growth and data consolidation while the 9910 provides for managed growth and open system exploitation.	2
Figure 2	The Hitachi Lightning 9960 Series' Hi-Star internal switched-fabric architecture.	3
Figure 3	System block diagram of the 7700E's shared bus architecture.	4
Figure 4	A 9960 Control Frame and its major components.	7
Figure 5	A 9960 Array Frame with its major components.	8
Figure 6	A power supply module of a Lightning 9900 Series system.	9
Figure 7	A 9910 single-cabinet storage subsystem.	10
Figure 8	Fibre Channel and ESCON or FICON adapter boards connect to open systems and mainframe computers.	11
Figure 9	A Fibre Channel host adapter board and its major hardware components.	12
Figure 10	An ESCON host adapter board and its major hardware components.	13
Figure 11	A 9960 system connected to Mainframes and Open Systems hosts.	14
Figure 12	The Cache Hierarchical Star Network (C-HSN) provides fully redundant switched access to cache from all channel adapters.	15
Figure 13	The CSW and CARBs provide non-blocking channel access to all cache.	16
Figure 14	A shared-bus architecture is limited to two simultaneous I/O operations.	17
Figure 15	The Lightning 9900 Series Hi-Star architecture allows for 16 parallel I/Os to cache through four interconnecting cache switches.	17
Figure 16	Separate redundant control memory handles the exchange of control information between processors and cache memory about the status, location, and configuration of data.	19
Figure 17	Functional diagram of a Lightning 9900 Series ACP pair.	21
Figure 18	Functional diagram of an Array Group with FC-AL disks.	22
Figure 19	Diagram of four ACP pairs and 32 FC-AL back-end disks.	22
Figure 20	The Lightning 9900 Series Systems are available with RAID-1+ and RAID-5 arrays.	24
Figure 21	The Lightning 9900 offers a far more complete range of availability protection than competitive products.	26
Figure 22	Hitachi Dynamic Link Manager automatically provides path failover and load balancing for open systems.	28
Figure 23	Test results show the advantages of the 9960 HiStar architecture when compared to a shared bus architecture.	31
Figure 24	Hitachi Data Systems offers a broad range of remote copy, data duplication, and data migration software solutions.	34
Figure 25	HiCommand allows management of virtually all Hitachi hardware and software from the customer's platform of choice.	38
Figure 26	The GUI interface on Hitachi Graph-Track simplifies performance management.	41

Tables

Table 1	Comparison of the Lightning 9900 series models.	2
Table 2	Bandwidth of supported interfaces of the Lightning 9900 Series.	14
Table 3	System Capacities and Number of disk drives for various configurations.	23
Table 4	Summary of Hitachi Freedom Storage software products and the business objectives they serve.	33
Table 5	Hitachi copy products provide Remote Copy, Data Duplication and Data Migration.	36

Introduction

1

The Only Switched-Fabric Storage Architecture in the World

The Hitachi Freedom Storage™ Lightning 9900™ Series internal switched-fabric storage architecture is considered the most advanced design available in the world for open systems and mainframe data. The patented Hi-Star™ architecture powers the most intelligent, highest performance enterprise-class storage product in the industry today.

The most advanced design in storage.

1. Hi-Star™, a **modern switched-fabric storage architecture**, was introduced in 2000 and is the only one available among enterprise class storage systems.
2. The Lightning 9900™ Series leapfrogs competitive products in virtually every category of comparison, especially in internal hardware architecture, availability, performance, and software solutions.
3. **Very high availability** is provided by an unusually robust product design. Full environmental monitoring and hot swap capability is also provided.
4. Front-end support options include virtually all major open systems platforms and S/390® compatible mainframe computers, which provides unsurpassed functionality.
5. The Lightning 9900™ Series offers **scalable performance** through a multiple switched architecture, hardware RAID, real-time operating system (kernel) design, and superior system management tools.
6. **Backup windows, operations efficiency, and disaster recovery are optimized** using leading Hitachi copy, migration, and data duplication software solutions.
7. The Lightning 9900™ Series is fully compatible with all major **UNIX®, Windows NT®, and IBM® S/390 mainframe network and system management tools** in addition to Freedom Storage™ management software.
8. Hitachi offers optional **robust dial-in capability** that provides world class experts to diagnose problems if they occur.

Lightning 9900™ Series Product Offerings

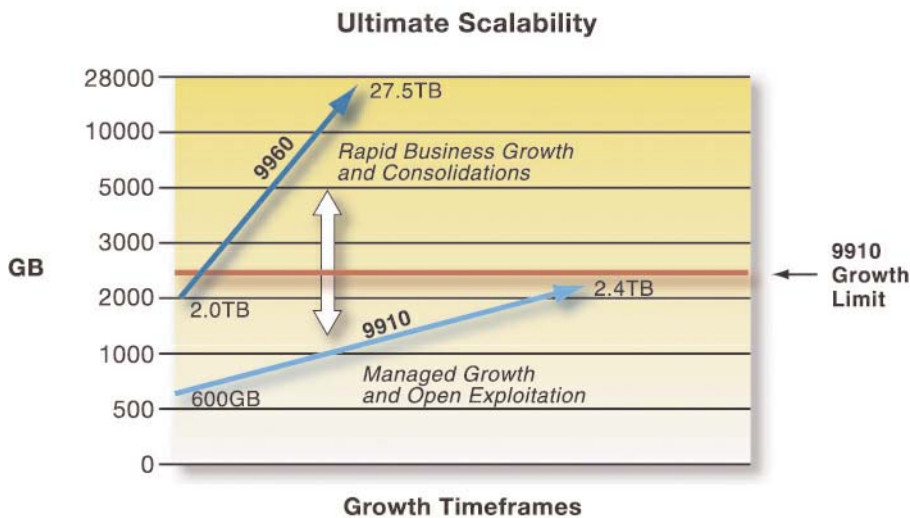
The Lightning 9900™ Series is offered in two models – the 9960 and the 9910. The 9960 is recommended for users requiring large storage capacities of 37+ TB of raw storage capacity. With the same advanced architecture as the Lightning 9960, the 9910 bridges the gap between standard mid-range and premium priced advanced function storage. With capacities up to 3+ TB of storage, the 9910 is recommended for users who require advanced function capabilities of premium storage, yet do not need the actual storage capacity of the 9960. Table 1 compares key differences between the two product offerings in the Lightning 9900™ Series.

Table 1 – Comparison of the Lightning 9900™ Series models.

Product Characteristics	9960	9910
Number of Cabinets	2-7	1
Maximum Cache Capacity	32 GB	16 GB
Number of Fibre Channel/ESCON® Ports	Up to 32	Up to 24
Number of FICON Ports	Up to 16	Up to 12
Number of Disks (HDDs)	Up to 512	Up to 48
Number of Array Groups	Up to 126	Up to 11
Number of ACP pairs	1-4	1
Number of FC-AL Paths to Back-end Disks	Up to 32	Up to 8
Maximum TB of Raw Disk Capacity	37 TB+	3 TB+
Minimum usable RAID-5 Capacity	34.06 GB	34.06 GB
Maximum usable RAID-5 Capacity	27.5 TB	2.4 TB

Selecting the 9910 or the 9960 is dependent on capacity requirements and expected growth of data. The Lightning 9900™ Series therefore offers product alternatives that cover a very broad range of scalability as illustrated in Figure 1.

Figure 1 – The 9960 provides for exploding growth and data consolidation while the 9910 provides for managed growth and open system exploitation.



An Overview of the Hitachi Lightning 9900™ Series Architecture

The Lightning 9900™ Series intelligent storage array is the most powerful enterprise storage system in the industry today. The new architecture used in the Lightning 9900™ Series is specifically designed for the brutal I/O demands of open systems environments. Extremely high internal bandwidths, high-speed back-end design, scalable internal pathways, increased processor speeds, and increased number of processors, larger cache sizes, and a new and improved high-performance RAID-1 (called RAID-1+) sets the Lightning 9900™ Series apart from all other storage systems. Figure 2 illustrates the architecture of the Lightning 9900™ Series.

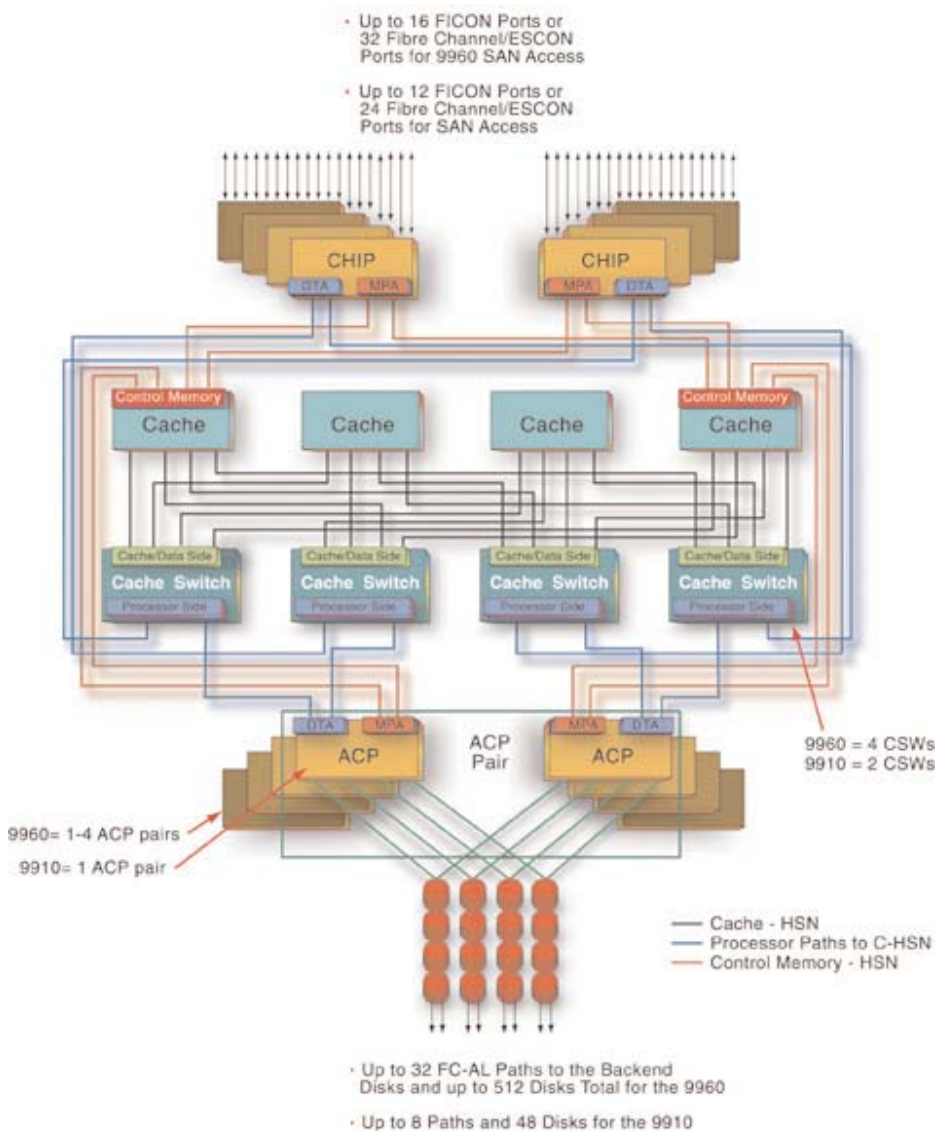


Figure 2 – The Hitachi Lightning 9960™ Series' Hi-Star™ internal switched-fabric architecture.

At the heart of the Lightning 9900™ Series' enormous capacity for throughput lies the Hi-Star™ architecture. This network of interconnecting high-speed links to system components provides an unequalled internal aggregate bandwidth of 6.4GB/sec (gigabytes/sec) in the 9960. This bandwidth is divided equally into 3.2GB/sec for data transfer and 3.2GB/sec for separate internal control information transfer.

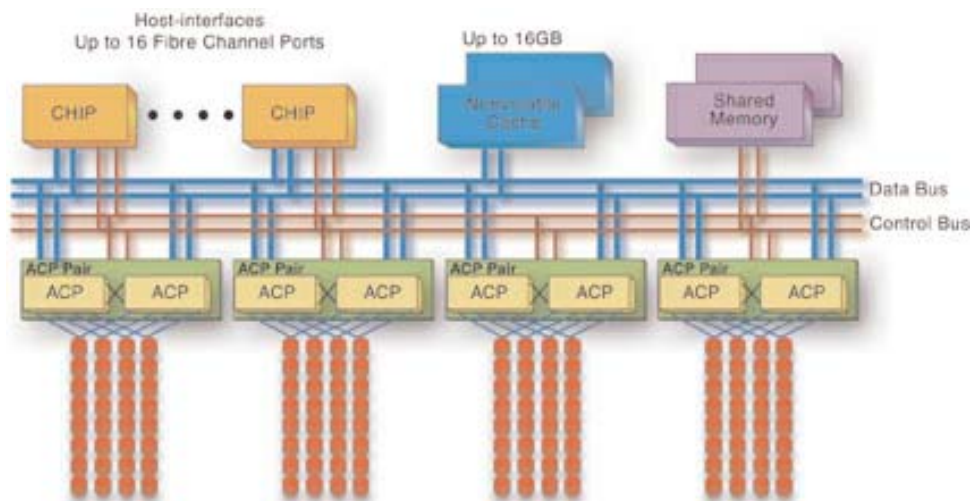
The main components that make up the Lightning 9900™ Series are the Array Control Processor modules (ACPs), Host Adapter modules, called Client Host Interface Processors (CHIPs). These elements enable the unsurpassed bandwidth of the Hi-Star™ network by bringing their own network paths (and bandwidth) to the Cache Hierarchical Star Network (C-HSN) and the Control Memory Hierarchical Star Network (CM-HSN). This means that the full 6.4GB/sec internal bandwidth is available with a fully configured 9960 with all four ACP pairs (eight modules total) and all four CHIP pairs (eight modules total). This is a truly high-performance and cost-effective configuration. With any Lightning 9900™ Series configuration, growing or scaling can be easily completed by adding more CHIP pairs for connectivity or ACP pairs for more capacity or performance. The internal bandwidth increases every time components are added. The bandwidths stated are the rated bandwidths for the paths; the effective rates will vary based on application workload and bandwidth consumption.

Comparison to the Shared Bus Architecture of the Hitachi Freedom Storage™ 7700E

Bus contention is virtually eliminated with the Hi-Star™ Architecture.

Unlike the shared common bus architecture of previous generation storage products, such as the 7700E (as shown in Figure 3 for reference to a shared bus architecture), Hi-Star™ is designed to truly scale in both capacity and performance. This is possible only with the Lightning 9900™ Series' internally switched architecture which, unlike shared bus architectures, does not have to contend for limited bus bandwidth.

Figure 3 – System block diagram of the 7700E's shared bus architecture.



The Lightning 9900™ has an advanced availability profile compared to competitive products

Hitachi is the only major supplier of enterprise class storage systems with a 100% availability guarantee.

The Lightning 9900™ was designed with maximum emphasis on high availability computing for today's most critical enterprises. Figure 15 highlights some of the many advantages in high-availability design that the Lightning 9900™ has over competitive products.

- Active dual-ported disk drives instead of single-port active drives.
- Separate storage for control & configuration tables.
- Mirrored cache instead of single image cache.¹
- Hitachi Data Systems is the leading supplier of storage systems for GDPS production sites.
- NanoCopy™ is a superior scheme for disaster recovery.
- High-availability, custom designed disk drives.

¹ Some competitive products do *not* have dual cache like the Hitachi Lightning 9900™. If there is an unrecoverable error in cache, there is no duplicate backup copy. With these products there is an increased risk of lost data, especially if the loss occurs in the hardware status area. The Lightning 9900™ does not have this problem since all status areas and write data are duplexed. Although a minor concern to some given the low probability of occurrence, the thought of data loss is never comforting to enterprise executives charged with 7x24 operations.

Other Information Sources Available from Hitachi Data Systems®

Additional information on the enterprise storage market, including an in-depth overview of the technically advanced and unsurpassed software solutions available from Hitachi, is available in the companion publication, the *Hitachi Software Solutions Guide*. This document is downloadable in PDF format from www.hds.com. The report contains expert guidelines regarding Hitachi industry leading software solutions for data copy, data migration, and data duplication. It also compares the Hitachi solutions to the major competitive software solutions from EMC and IBM and drills down into specifics.

Since storage today is more than just hardware and software, additional information will be available in a third report titled the *Hitachi Enterprise Solutions Guide*. This report contains information and on the world-class services and solutions available from Hitachi Data Systems® and how these solutions can be combined with Hitachi and other vendor hardware and software for state-of-the-art, simplified storage management for the enterprise.

If further information is required, Hitachi Data Systems® representatives can provide specialized presentations, reports, and expert knowledge on the topics contained in this series of reports.

The Lightning 9900 Series Packaging

2

Model 9960 Packaging

The Lightning 9960™ is available in a two to seven cabinet configuration (as shown on the cover of this guide), depending on connectivity and capacity needs. A Control Frame serves as the control center of each system and manages up to six Array Frames that are connected via high speed Fibre Channel links. A fully configured system is contained in seven cabinets 70.5 inches high, 171.3 inches wide, and 31.5 inches deep.

Model 9960 Control Frame

The Control Frame of the 9960 contains the Service Processor, Client Host Interface Processor modules (CHIPs) with Fibre Channel adapters, ESCON® adapters, and/or FICON adapters, Cache modules, Cache Switch Modules (CSWs), Array Control Processor modules (ACPs), power supplies, and battery modules. An illustration of the Control Frame is shown in Figure 4.





Figure 4 – A 9960 Control Frame and its major components.

Model 9960 Array Frames

Array frames 1-4 of the 9960 contain up to 96 Hard Disk Drive Assemblies (HDDs) and communication interfaces to the Control Frame. Array frames 5-6 contain up to 64 drives for a total of 512 drives for a maximum configured system. An Array Frame is shown in Figure 5.

This design allows for efficient centralized and consolidated storage, which results in reduced management costs and lower overall total cost of ownership.

Figure 5 – A 9960 Array Frame with its major components.



The Lightning 9900™ Series power subsystem consists of redundant power supplies in both Control Frames and Array Frames. The power supplies are both hot pluggable and hot swappable and are N+1 redundant. This means they can be removed or installed during system operation and that if a power supply fails during operation, the failed power can be dynamically removed without any loss to system operation. A Lightning 9900™ Series power supply module is shown in Figure 6.

Cables carry 48-volt main system power from the back of the Control Frame and Array Frame where the power is converted to the required voltages.

Different AC power cord options allow all Lightning 9900™ Series models to connect with the different electrical outlet configurations used around the world.




Figure 6 – A power supply module of a Lightning 9900™ Series system.

Model 9910 Packaging

Although the 9960 has the ability to be upgraded from one to six array frames, the Lightning 9910 is a fixed-configuration, single-cabinet storage system. The 9910 is targeted for IT environments where floorspace is a premium and data growth is limited, yet where premium storage functionality is required. The 9910 looks very much like the Control Frame of a Lightning 9960. The back end contains four HDD arrays housing up to 48 disks in a 70.5 inch high, 27.6 inch wide, and 31.5 inch deep footprint. This is shown in Figure 7.

Figure 7 – A 9910 single-cabinet storage subsystem.



The System Monitoring Network of the Lightning 9900™ Series

The Lightning 9900™ Series systems have both a predictive maintenance and a failure alert system.

The Lightning 9900™ Series systems are designed with a system-monitoring network. The system monitoring cables are connected to all cabinets in the Lightning 9900 Series systems, and report a variety of system parameters, such as component failure, fan speed, power, voltages, and temperature, to ensure trouble-free operation. This network not only provides for efficient predictive maintenance (replacing a component before it fails), but also for the failure alert system to expedite system repair after a component has failed. This information is passed to the Control Frame, where it can be viewed either locally by the Service Processor (SVP) or remotely by the Hitachi Resource Manager 9000™ across a private LAN on a UNIX® or Windows NT® console. Users can view failure information via SNMP or S/390 SIMs. System status and alerts are also sent to Hitachi Data Systems Customer Support Centers via the Hi-Track™ “call-home” software facility.

Front-end Design of the Hitachi Lightning 9900 Series

3

CHIPs Support Fibre Channel Host and ESCON® or FICON Host Connectivity and Scalability

The Lightning 9900™ Series supports both Fibre Channel for open systems and ESCON® or FICON for S/390® compatible platforms as shown in Figure 8.





Figure 8 – Fibre Channel and ESCON® or FICON adapter boards connect to open systems and mainframe computers.

About the Fibre Channel Interface Boards

Each Fibre Channel Interface Board, as shown in Figure 9, has four Interface Processors² for management of connectivity to the host.

Figure 9 – A Fibre Channel host adapter board and its major hardware components.



² The Lightning 9900™ Series uses the Reduced Instruction Set (RISC) Intel® i960 80MHz processor that is especially designed for I/O operations.

There are four ports for each interface board or a total of 32 host interfaces for a fully configured system. An ESCON interface board is shown in Figure 10. A FICON interface board has a similar appearance.





Figure 10 – An ESCON® host adapter board and its major hardware components.

A Lightning 9900™ Series system can connect to mainframes and open systems hosts, thereby consolidating the storage for an enterprise under one roof for consistent performance, availability, and ease of management. This is shown in Figure 12.

Figure 11 – A 9960 system connected to Mainframes and Open Systems hosts.



Bandwidth of Supported Host Interfaces

The bandwidth of interfaces supported by the Lightning 9900™ Series is shown in Table 2.

Table 2 – Bandwidth of supported interfaces of the Lightning 9900™ Series.

Supported Interface	Bandwidth
Fibre Channel Adapters	100MB/sec or 200MB/sec
ESCON® Adapters	17.5MB/sec
FICON Adapters	100MB/sec

Number of Host Interfaces Supported

Up to 32 Fibre Channel, ESCON® or FICON ports are available on the 9960 and up to 24 on the 9910. No card slots are subtracted from anywhere else in the subsystem to accommodate these ports.

Types of Hosts Supported by Direct Fibre Channel Host Attachment by the Lightning 9900™ Series

The Lightning 9900™ Series supports all major open systems hosts and mainframe hosts for unmatched connectivity. The Lightning 9900™ Series supports direct Fibre Channel host attachment to all major open systems platforms as listed below:

- Windows NT 4.0®
- Windows 2000®
- HP-UX®
- Solaris™ (Sun™)
- AIX 32 and 64 bit (IBM®)
- NetWare™ (Novell®)
- Linux® Red Hat
- IRIX® (SGI™)
- Tru64™ UNIX® (Compaq®)
- OpenVMS™ (Compaq®)

Hi-Star Architecture – An Internal Switched Hierarchical Star Network

4

The Internal Hi-Star Architecture

The internal Hi-Star architecture consists of two separate networks, the Cache Hierarchical Star Network (C-HSN) and the Control Memory Hierarchical Star Network (CM-HSN).

The Cache Hierarchical Star Network

The C-HSN is the network used for transferring data to and from the main global data cache. The C-HSN can also be broken down into two component networks: the processor paths to the C-HSN and the actual C-HSN (as shown previously in Figure 2). Figure 12 provides a detailed illustration of the processor paths and the C-HSN.

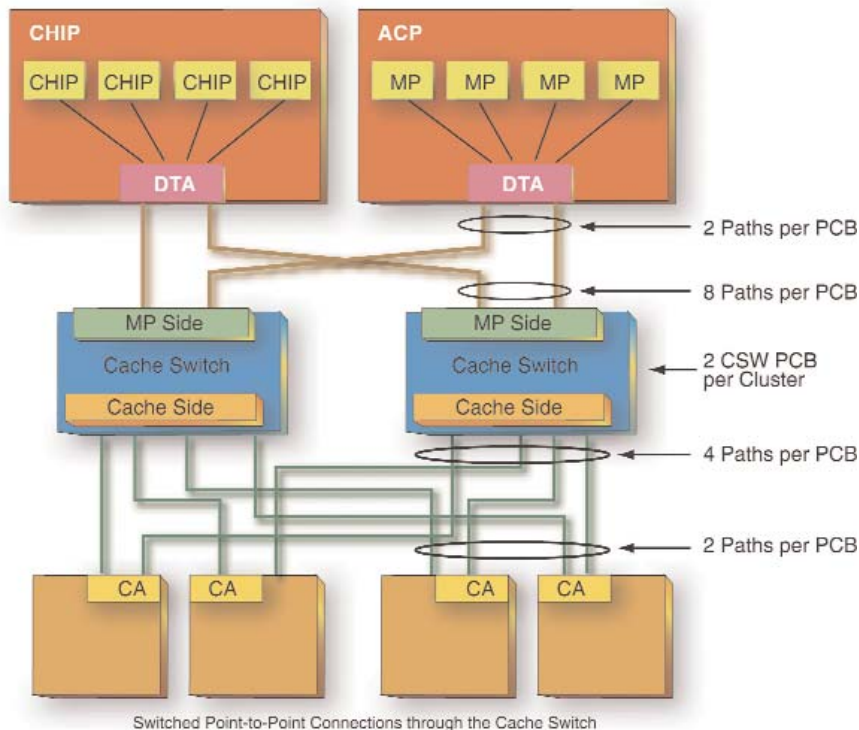


Figure 12 – The Cache Hierarchical Star Network (C-HSN) provides fully redundant switched access to cache from all channel adapters.

C-HSN

16 bit x 100MHz x 16 Paths = 3.2GB
Total C-HSN

200MB per Path x 2 Paths = 400MB/s per PCB

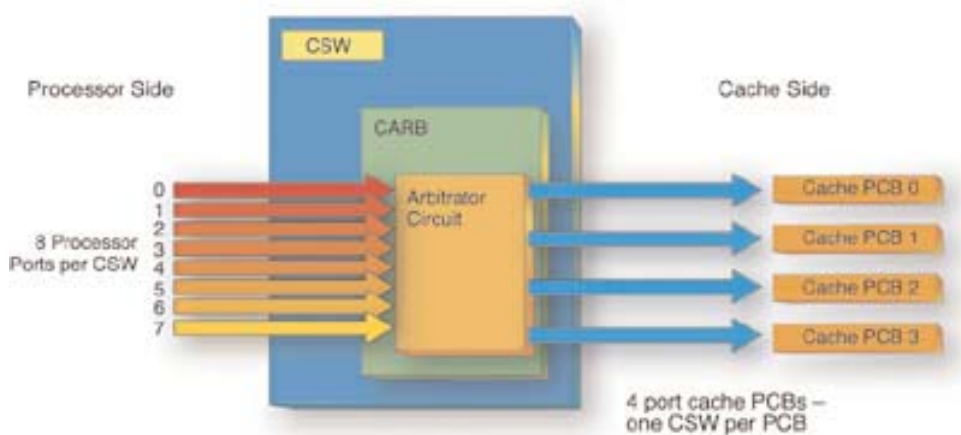
32 Paths x 200MB/s = 6.4 GB/sec to the C-HSN

The CSW is at the heart of the Lightning 9900™ Series.

Figure 13 – The CSW and CARBs provide non-blocking channel access to all cache.

The Cache Switch Provides Non-blocking Switched Access to Cache

At the heart of the C-HSN is the Cache Switch (CSW). Together, these four switches (in the 9960) form the internal star using parallel switch fabric bus (PSFB) of the C-HSN. The CSW is a specially designed crossbar switch that functions as a combination MUX, path arbitrator, and non-blocking network switch. The CSW functions as a MUX by supporting eight paths into the processor side of each switch and four paths to the cache modules (8 to 4 routing). All total, there are 32 paths at the processor side of the fabric network and 16 paths to the cache modules from the cache side of the fabric network. As shown in Figure 13, Cache Memory Arbitrator (CARB) allows the CSW to act as an arbitrator, handling the access and addressing to and from the cache modules. The CSW also manages the non-blocking paths between the processors to and from the cache.



The Path Bandwidth Between the CSW and CHIPS or ACPs is 6.4GB/sec

The paths that connect the Client Host Interface Processors, or CHIPs (through the Data Adapters – DTA) to the CSW and the CSW to the cache modules are each 16 bits wide, plus 2 bits for parity. The bandwidth of each path is 200MB/sec. This is the data transfer rate of a 16-bit wide path (parallel) clocked at 100MHz (16 bits at 100MHz equals 200MB/sec). As mentioned, there are 32 paths from the processor modules of the Lightning 9960™ (ACPs and CHIPs) to the ports of the CSWs. This equals 6.4GB/sec (32 paths at 200MB/sec) of bandwidth to the processor side of the switched fabric (when all 16 component modules are installed).

The Path Bandwidth Between the CSW and the Cache Adapters (CAs) is 3.2GB/sec

The bandwidth between the cache side of the Cache Switch (CSW) and the cache modules is 3.2GB/sec (16 paths at 200MB/sec) total. This bandwidth will always be 3.2GB/sec as long as the four CSWs and the four cache modules are installed and functioning.

Advantages of a Switched Architecture Compared to a Shared Bus Architecture

When considering the advantages of a switched bus architecture to that of a shared bus architecture (Figure 14), the number of operations that can be present on the buses simultaneously is the key difference.




Figure 14 – A shared-bus architecture is limited to two simultaneous I/O operations.

In a shared bus storage system, the number of internal simultaneous processor operations that can occur is equal to the number of buses on the backplane. The shared buses are used by all host channels to communicate between the processors on the host adapters and cache memory. In the case of the shared-bus design shown in Figure 14, there are four separate buses for data and control information. Therefore a maximum of at most four I/Os can be processed simultaneously in this architecture.

Shared-bus architecture handle only two simultaneous I/Os.

On Lightning 9900™ Series system, however, with its internal non-blocking switch topology connecting the processors to cache, that number is 16. This is equal to the total number of paths between the CSW and the cache module ports (Figure 15). This results in significantly improved simultaneous I/O performance of the Lightning 9900™ Series system compared to shared-bus systems as discussed in Chapter 8.

The Lightning 9900™ Series handle 16 simultaneous I/Os.

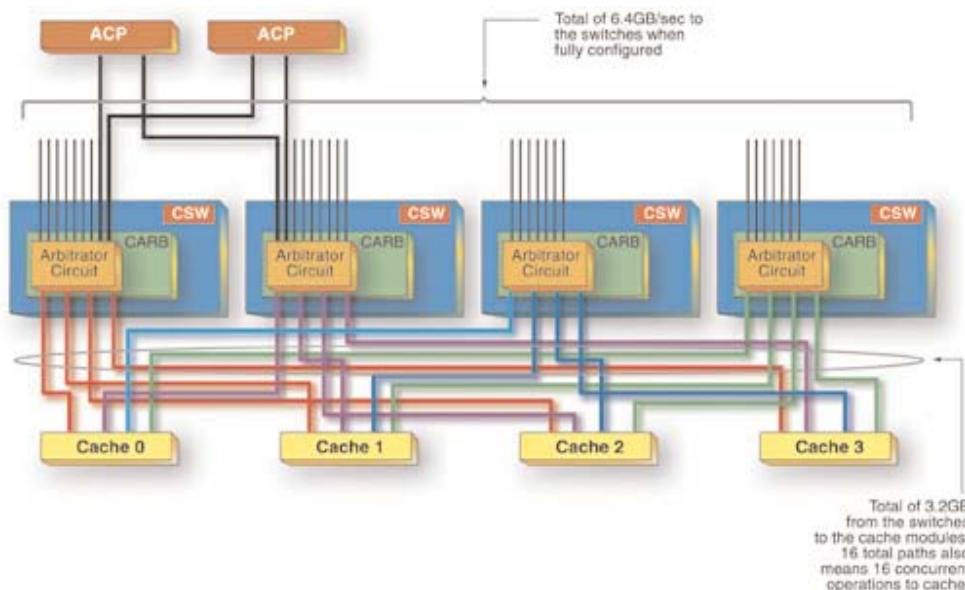


Figure 15 – The Lightning 9900™ Series Hi-Star™ architecture allows for 16 parallel I/Os to cache through four interconnecting cache switches.

A portion of cache memory can be allocated to specific data.

The 9960 Provides 32GB of Fully Addressable Cache

The 9960 supports up to 32GB of data cache, all directly addressable. Separate cache modules (up to 1.5GB) are used for control storage. Competitive systems use their cache for both data and control information, limiting the amount of usable data cache.

Advanced Cache Algorithms of the Lightning 9900™ Series

The Lightning 9900™ Series has a variety of advanced cache algorithms and software solutions that allow exceptional performance.

Hitachi FlashAccess™ Allows Datasets to be Permanently Placed in Cache

Hitachi FlashAccess™ allows users to dynamically “lock and unlock” data into cache in real time. Read and write functions are then performed at cache speeds, with no disk latency delay. With Hitachi FlashAccess™, a portion of cache memory can be allocated to specific data. Administrators can add, delete, or change FlashAccess™ managed data at any time, quickly and easily.

In S/390® environments, defined by the Logical Volume Image (LVI), cache data can be as small as a single track or as large as a full 3390. For increased configuration flexibility, Hitachi FlashAccess™ offers multiple modes of operation. It can be used in conjunction with Hitachi RapidXchange™ to increase the speed of data transfer and, therefore, improve performance of mainframes to open systems data exchange. RapidXchange™ supports both open-to-S/390 and open-to-open high speed data transfers.

Contiguous Block Allocation for High-speed Sequential Access

The Lightning 9900™ Series uses a delayed block allocation mechanism for allocating disk blocks to files. Unlike the UNIX® File System (UFS), where blocks are allocated early, disk blocks in the Lightning 9900™ Series systems are allocated only when the host is ready to write data to disk. The file system attempts to allocate contiguous disk blocks for all the pages to be written to disk. This helps to maximize contiguous block allocation. Contiguous block allocation is highly advantageous for applications, such as Mechanical Design (MCAD), where a relatively small number of large files are read and written by a few clients. Under such conditions, the I/O request stream seen by the server is often highly sequential. By storing file data contiguously, read performance can be substantially improved.

The block allocation mechanism is further tuned for the underlying RAID-5 implementation such that it generates full stripe writes whenever possible. Full stripe writes allow a full stripe of data plus parity information to be written to the RAID array without requiring any data to be read from the stripe. By comparison, partial stripe writes require data to be read from the stripe to generate correct parity.

Read-ahead for High-performance Sequential Reads

Read clustering in the Lightning 9900™ Series is enabled using built-in heuristics to read ahead for every I/O. The heuristics are applied to determine if the data is being accessed sequentially. If so, then the Lightning 9900™ Series reads ahead pages corresponding to that data. Read-ahead helps to ensure that when a client read request is received the requested data will already be stored in the data cache, so the request can be satisfied immediately.

Control Memory Hierarchical Star Network®

The second component of the Hi-Star Architecture™ is the Control Memory Hierarchical Star Network™ (CM-HSN). This is a point-to-point network that handles the exchange of control information between the processors and control memory. The control memory contains information about the status, location, and configuration of the cache, the data in the cache, and the configuration of the Lightning 9900™ Series system (as well as other information related to the operational state of the system). Two control memory areas are mirrored images of each other. This is illustrated in Figure 16. Control data is “data about data” also called “meta-data.” Essentially, control information is handled “out of band” from the data paths, both through a separate memory area and network.

The CM-HSN is a much simpler network design in that every connection is a point-to-point connection. Only the Cache-HSN (data paths) uses a switched-fabric topology for its interconnecting network. The CM-HSN also uses a narrower path and more of them. Figure 16 shows a close-up view of the CM-HSN’s networking topology. Referring back to the diagram in Figure 2, there are two CM-HSN paths connecting the processors to the control memory. However, the diagram in Figure 15 shows four paths per processor module. There are 64 4-bit paths connecting the processors to the control memory. The diagram in Figure 2 shows the 4-bit paths combined into their full 8-bit (plus a parity bit) paths. The zoomed-in view in Figure 16 shows all of the ports to the control memory in its “split” configuration. Again, control memory is mirrored. This path topology is used to facilitate this mirrored architecture.

Either way the bandwidth is calculated in the CM-HSM, the bandwidth is 3.2GB/sec total when all of the processor modules are installed. 4 bits clocked at 100MHz equals 50MB/sec per path multiplied by 64 ports is 3.2GB/sec bandwidth. Or, the full 8-bit path clocked at 100MHz multiplied by the combined 32 ports of the control memory “pair” equals 3.2GB/sec total.




Figure 16 – Separate redundant control memory handles the exchange of control information between processors and cache memory about the status, location, and configuration of data.

High Performance Back-end Design 5

High-speed Back-end Design

New Hitachi high-performance back-end drives are available with the Lightning 9900™ Series. The drives are attached via Fibre Channel – Arbitrated Loops (FC-AL). Eight 100MB/sec FC-AL connect to dual ported drives for each ACP pair. With four configured ACP pairs installed (at 800 MB/sec per ACP pair) that's 3.2GB/sec total back-end bandwidth to the physical drives.

Each FC-AL path is driven with a dedicated processor and a Data Recovery and Regeneration (DRR) circuit. This is in contrast to the Hitachi Freedom 7700E™, where two of the FW-SCSI-2 paths in the back end were shared through a single processor and single DRR circuit. (See also Figure 2). Basically, the number of processors and DRR circuits in the back end of the Lightning 9900 Series has increased to 32 (for the Lightning 9960™), compared to the 7700E's 16. The diagram of a Lightning 9900™ Series ACP pair is shown in Figure 17 and Figure 18. Figure 19 illustrates all four ACP pairs with 32 back-end disks per array group.

The Lightning 9900™ Series uses advanced algorithms for managing performance of ACP pairs.

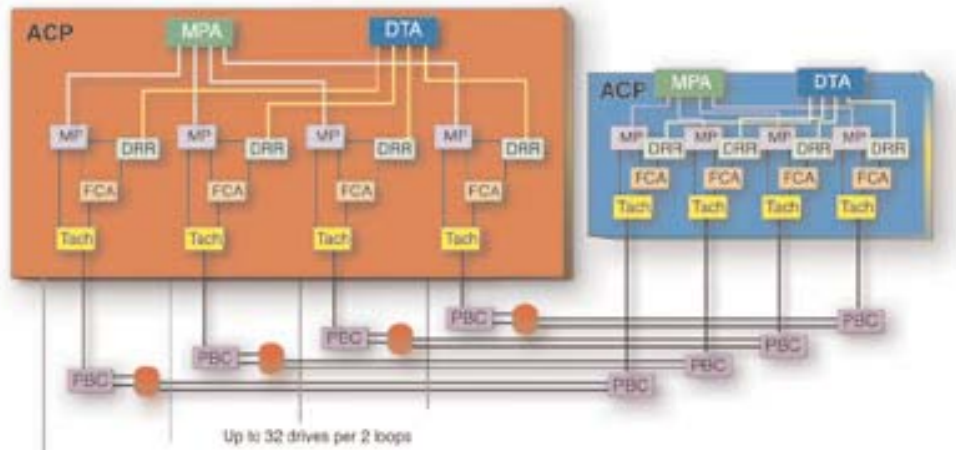


Figure 17 – Functional diagram of a Lightning 9900™ Series ACP pair.


New Fibre Channel Disk Drives

The back-end performance of the Lightning 9900™ Series is enhanced with the introduction of a new dual ported/dual active Fibre Channel disk drive technology. These disk drives are specially designed and built by Hitachi Ltd. for the Lightning 9900™ Series. The disks are in a 3-inch form factor with capacities available in 18GB and 73GB spinning at 10,025RPM. These are all high-density disks (HDDs) spinning at high RPMs. The transfer rates with these new disks are approaching 60MB/sec. Higher sectors per physical track means less head movement and more data transferred per rotation, even when compared to the 7700E 15GB disks, which spin at 12,030RPM.

The Lightning 9900™ Series uses custom Hitachi high availability disk drives.

These new dual-ported/dual-active Fibre Channel disk drives, combined with the technology built into the ACP pair, allow the back end of the Lightning 9900™ Series to use all eight FC-AL paths in an ACP pair for both performance and fault tolerance. However, no one disk or array group can use both paths simultaneously. The ACPs monitor the activity and the utilization of the paths. Based on this information, the ACPs determine the best path to use for accessing a disk. A functional diagram of an Array Group with FC-AL disks is shown in Figure 18.


Figure 18 – Functional diagram of an Array Group with FC-AL disks.



Maximum Configuration of ACPs and Disks

The maximum configuration of four ACPs and 16 Fibre Channel loops is shown in Figure 19.

Figure 19 – Diagram of four ACP pairs and 32 FC-AL back-end disks for the Lightning 9960™.



The Importance of Command Tag Queuing (CTQ)

Another new feature was introduced with the Lightning 9900™ Series that was not employed on the 7700E – Command Tag Queuing (CTQ) to the back-end drives. Command Tag Queuing greatly improves the performance of the Lightning 9900™ Series system back end by offloading much of the seek optimization functions to the disk drives themselves, to allow for more simultaneous back-end I/O operations to occur.

Command Tag Queuing greatly improves performance.

Capacities and RAID Architecture

6

Storage Capacity of the Lightning 9900™ Series

The Lightning 9900™ Series supports an unmatched range of capacities as shown in Table 3.

System Capacities and Number of disk drives for various configurations		
Raw System Configuration	Raw capacity with 18GB drives	Raw capacity with 73GB drives
Lightning 9960 One Array Frame	1.7TBs = 96 disks	7.0TB = 96 disks
Lightning 9960 Six Array Frames	9.4TB = 512 disks	37.4TB = 512 disks
Lightning 9910	.9TB = 48 disks	3.5TB = 48 disks

Table 3 – System capacities and number of disk drives for various configurations.

Advantages of the Lightning 9900™ Series RAID Hardware

Hitachi designs the most advanced RAID (Redundant Array of Independent Disks) controllers in enterprise storage to interface to its disk subsystems. These intelligent controllers provide disk interface and RAID management, offloading these tasks to dedicated processors. Each Lightning 9900™ Series ACP controller supports RAID-1+ and RAID-5 (parity RAID). All disks in the system are defined as part of a RAID array of one type or another. Non-volatile RAM on each controller accelerates RAID functions, particularly disk writes.

Intelligent controllers offload tasks to dedicated processors.

RAID-1+

With the announcement of the Lightning 9900 Series RAID-1+ was introduced. RAID-1+ is a 2 Data + 2 Data disk arrangement (four disks total) in a RAID-1+0 configuration with a twist – a rotating copy.

RAID-1+ is unique to Hitachi.


RAID-5 – Distributed Parity

RAID-5 arrays consist of four disks. Data is striped across RAID-5 arrays in a fashion similar to RAID-1+, but RAID-5 provides fault resilience by keeping parity information on each stripe of data. If a failure occurs, the contents of that block can be recreated by reading back the other blocks in the stripe along with the parity. Parity information is distributed throughout the array to minimize potential bottlenecks in the event of a need to rebuild data from a failed disk. The overhead of RAID-5 is equivalent to one disk drive, regardless of the size of the array.

RAID Rebuild Capability

In the event of a disk failure, RAID-1+ or RAID-5 arrays can be rapidly and automatically rebuilt using available “hot-spare” drives. The Lightning 9900™ also allows logical volumes to be expanded online. Configurations for RAID-1+ and RAID-5 are shown in Figure 20.

Figure 20 – The Lightning 9900™ Series systems are available with RAID-1+ and RAID-5 arrays.



High-density Disk Drive Assemblies (HDDs)

Each drive has the industry unique feature of supporting dual active ports.

The disk drives used in the Hitachi Lightning 9900™ Series are specially designed for high availability and performance by Hitachi Limited. Each drive has the industry unique feature of supporting dual active ports. Whereas merchant market products like the Seagate® drives used in most competitive systems read and write through one port, using the second port as standby, the Hitachi drives read and write from whichever port is available. This provides significantly increased performance.

Hitachi High-Availability Software Improves Application Service Levels

7

An Overview of High-Availability (HA) Computing Software

It is increasingly important for a business to support continued access to global information 24 hours a day, 7 days a week. What's more, careers often depend on the availability of service levels provided by IT to the enterprise.

In computer science, availability refers to the degree to which a system or resource is capable of performing its normal function. Availability is measured in terms of Mean Time Between Failure (MTBF) divided by MTBF plus the Mean Time to Repair (MTTR).

- $AVAILABILITY = MTBF / (MTBF + MTTR)$.

For example, a server which fails on average once every 5,000 hours and takes an average of two hours to diagnose, replace faulty components and reboot, would have an availability rating of $5,000 / (5,000 + 2) = 99.96\%$. This would correspond to a Level 3 rating using the *Scale of 9s*.

Availability refers to the degree to which a system or resource is capable of performing its normal function.

Software products contribute to High-Availability Computing

Many factors can cause unplanned downtime. The Hitachi Freedom Storage™ Lightning 9900™ has been designed to eliminate as many of these factors as possible in both hardware redundancy, on-line replaceable components, and software data copy functions to allow copies of data at other locations either locally or remotely so that processing can continue in the event of an outage.

Backup and restore procedures and products also contribute dramatically to computer system availability by reducing the time to restore operations in the event of an outage.

Hitachi has teamed with industry-leading storage management software providers such as Tantian Technologies® and VERITAS® Software of Mountain View, California to provide world-class high-availability solutions. These solutions contribute to the fact that customers and analysts alike regard Hitachi Freedom Storage™ systems as having the highest availability in the industry.

Customers and analysts alike regard Hitachi Freedom Storage systems as having the highest availability in the industry.

The Lightning 9900™ has an Advanced Availability Profile compared to Competitive Products

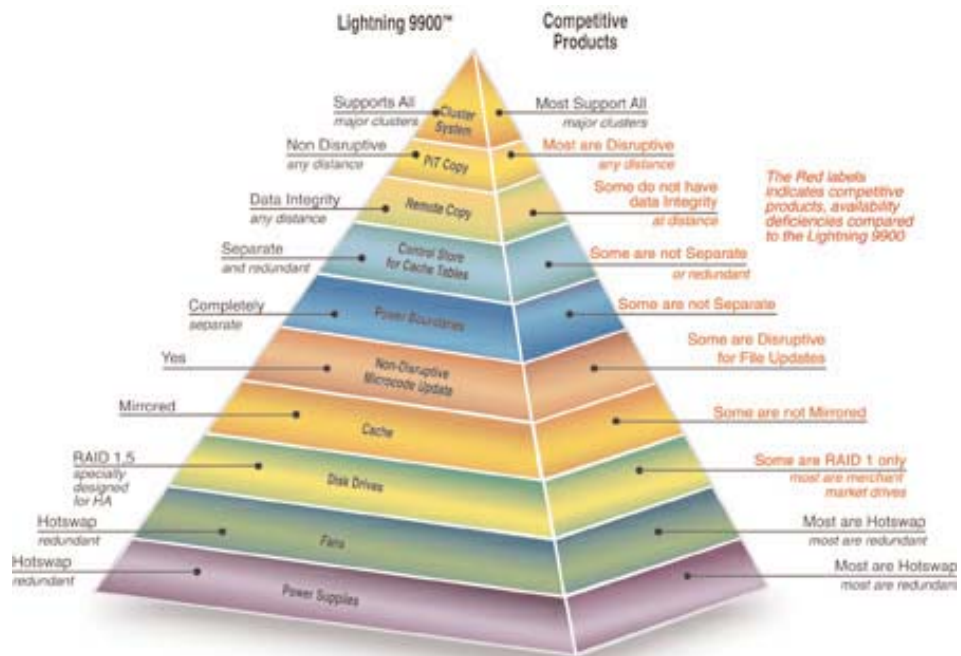
The Lightning 9900™ was designed with maximum emphasis on high-availability computing for today's most critical enterprises. Figure 21 highlights some of the many advantages in high-availability design that the Lightning 9900™ has over other enterprise storage products.

- Active dual-ported disk drives instead of single-port drives or active/passive dual-ported drives.
- Redundant active components throughout the system combined with automatic fail-over architecture.

Hitachi is the only major supplier of enterprise class storage systems with a 100% availability guarantee.

- Dual data paths and dual control paths connecting every component.
- Mirrored cache for all write data instead of single image cache.³
- GDPS support.
- Hitachi TrueCopy combined with ShadowImage™ for disaster recovery.
- Hitachi NanoCopy™ to provide Point-in-Time copies without requiring quiescence of databases or applications.
- High-availability, custom designed disk drives.
- “Phone home” capability to notify the Hitachi support center quickly of any incident that may affect availability.

Figure 21 - The Lightning 9900™ offers a far more complete range of availability protection than competitive products.



Open Systems High-Availability Middleware Support

The recovery process without high-availability middleware involves time-consuming and error-prone manual operations.

Open systems server vendors and third-party software vendors such as VERITAS® software have developed a class of software known as “high-availability middleware” to help reduce downtime by automatically detecting faults and recovering data services on a redundant set of hardware. Without high-availability middleware, time is lost while a fault goes undetected. Once the fault is detected, a diagnose/repair/replace action must take place before data-service recovery can begin. High-availability middleware can begin an automated recovery process immediately on the redundant hardware. The recovery process without high-availability middleware involves time-consuming and error-prone manual operations, which may include resetting the SCSI bus, restarting drivers, reassigning IP addresses, recovering and restarting applications and transactions, and even rebooting.

³ Some competitive products do *not* have dual cache like the Hitachi Lightning 9900™. If there is an unrecoverable error in cache, there is no duplicate backup copy. With these products there is an increased risk of lost data, especially if the loss occurs in the hardware status area. The Lightning 9900™, of course, can never have this problem since all status areas and write data are duplexed.

There are four basic types of high-availability middleware that can reduce downtime in the event of a data path or host failure. The Lightning 9900™ supports each classification of middleware:

1. Alternate Pathing middleware switches the I/O load in the event of path failure

This type of middleware automatically switches the I/O load on a failed primary path to an alternate path on the same host system. The Lightning 9900 support alternate pathing for AIX® 4.2 and above, Compaq True64™ UNIX®, HP® 10.01 and through PVlink, Windows NT/2000® and Sun Solaris™ 2.5 and Sequent DYNIX/ptx®. In addition, VERITAS® Dynamic Path Management™ (DPM) is supported.

Alternate pathing automatically switches the I/O load on a failed primary path.

2. Hitachi Dynamic Link Manager™ provides path failover and load balancing

Hitachi Dynamic Link Manager™ is a family of Hitachi-provided middleware software utilities that are server-based as shown in Figure 22. Hitachi Dynamic Link Manager™ enhances the availability of RAID systems by providing automatic error recovery and path failover from server-to-RAID connection failures. Hitachi Dynamic Link Manager™ provides load balancing in addition to path failover by re-directing I/O activity to the least busy path using complex algorithms.

Hitachi Dynamic Link Manager™ provides load balancing in addition to path failover.

Just because a system is RAID-protected doesn't mean it is protected against connection bus failures, which is why Hitachi Data Link Manager™ is required for true non-stop operations. This product allows systems administrators to take advantage of the multiple paths on a Lightning 9900™ by adding redundant connections between application servers and RAID systems. Hitachi Dynamic Link Manager™ therefore provides increased reliability and performance. Supported platforms include AIX®, Sun Solaris™, Windows NT®, and Windows 2000®.

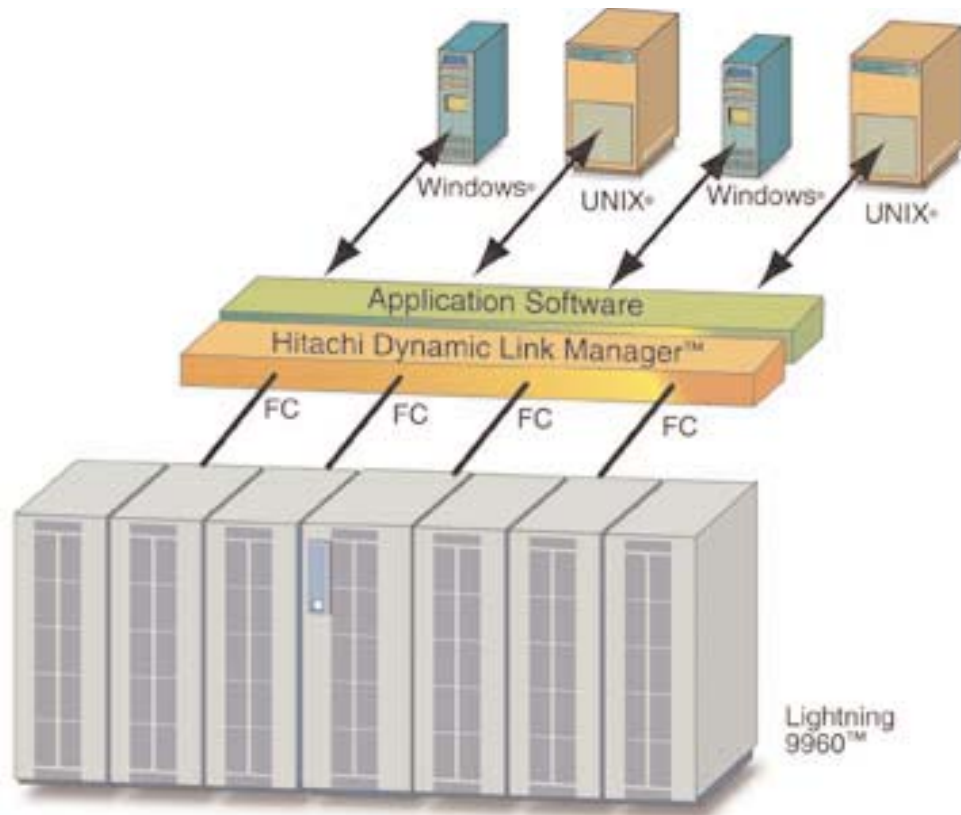
3. Host Failover

Host failover software supports a cluster of host processors in which one of the hosts automatically takes over the workload of any failed host in the cluster. This "take-over" includes the reassignment of networks and peripherals, as well as the restarting of applications. Host clustering can also be used to create fault-tolerant work loads and scale processor capability while sharing network and disk resources.

One of the hosts automatically takes over the workload of any failed host in the cluster.

The Lightning 9900™ supports all major open systems clustering schemes including: Compaq TruCluster®, HP® MC Service Guard, HP® MC Lock Manager, IBM® RS6000 and SP HACMP, Windows NT/2000®, Microsoft® Cluster Server, NCR® UNIX® SVR4 Lifekeeper, IBM DYNIX/ptx ATAP®, and VERITAS ClusterServer™ for Sun Solaris™.

Figure 22 - Hitachi Dynamic Link Manager™ automatically provides path failover and load balancing for open systems.



Clustering middleware supports distributed lock management.

4. Parallel Database Clustering

This type of middleware is a special version of host failover middleware, which supports major parallel database servers like Oracle® Parallel Server, Informix® XPS, and Sybase® MPP. Clustering middleware supports distributed lock management, a feature that enables parallel database software running on separate cluster nodes to share access to the same database. If one host fails, the other hosts can take over its work. Database clusters allow a customer to grow a database incrementally simply by adding additional nodes. With non-parallel database servers, the server has to be replaced or an additional server with another database instance has to be purchased and installed when the capacity of the original system is exceeded.

The Lightning 9900™ system has been certified with MC/Lock Manager for Oracle® Parallel Server and Sun™ PDB for Oracle® Parallel Server.

Standard Hot-swap and Redundant Power Supplies

The Lightning 9900™ Series provides fully redundant power supplies to ensure uninterrupted power and cooling to all chassis in the system – supporting full system power in the event of a power supply failure. In the event of such a failure, the power supply that has failed can be “hot swapped” without disruption to the power subsystem.

Standard Redundant Fans for Cooling

The Lightning 9900™ provides fully redundant fans for uninterrupted cooling to all components of the system. The speed and operation of these fans are monitored by the environmental monitoring system. If a fan should fail, it can be “hot swapped” without disruption to the cooling system.

Unmatched Competitive Advantage in RAID Choices

The flexibility in choice for Lightning 9900™ Series RAID protection for disk drives is unmatched. It is important to have choices in RAID protection since files and data have different characteristics in terms of user required performance, uptime, and rebuild times. RAID-1+ (mirroring) offers the highest performance (read from either disk in the mirror pair) during normal operation. RAID-1+ also offers the highest performance in the event of a failed disk since it is not necessary to read parity (and then data) in the event of a disk failure in other RAID configurations.

The Lightning 9900™ Series also offers the capability to intermix RAID levels within a system, depending on the level of availability and performance required by the data hosted on the Lightning 9900™ Series system. For example, a single Lightning 9900™ Series system can configure both RAID-1+ and RAID-5.

Performance Considerations 8

Truly Scalable Performance

Based on the internal switched-fabric Hi-Star™ architecture, the Lightning 9960™ shows unmatched performance scalability. The more capacity and connectivity added to the 9960, the more total bandwidth increases. As stated previously, the 6.4GB/sec internal bandwidth is reached when all 16 modules are installed and functioning.

When new application workloads are added to a 9960 system with an existing workload, neither the new nor the existing workloads will suffer performance degradations from the sharing of internal resources and bandwidth. The following series of tests proves this point. By measuring the effects of one workload on another by only sharing the controller, we can illustrate this phenomenon.

The Advantages of the Lightning 9900™ Series Hi-Star™ Architecture Over a Shared Bus Architecture

A test was conducted that clearly demonstrates that the 9960 is not susceptible to performance declines typical of a shared bus architecture. The test compares an OLTP and a DSS workload running on both a Lightning 9960™ and a 7700E system. We use the 7700E as an example of the shared bus storage system architecture. In both cases, the same OLTP application workload was run on an HP/9000-K460® server and the same DSS workload was run on two Sun E3000® dual-processor systems. In one case, both processors were connected to a 7700E and then to a 9960 system⁴. There were no differences, therefore, from the application software or the server platform variables. They were the same for both sets of data. The only difference was the storage system in order to illustrate the advantages of the Hi-Star™ switched-fabric architecture. Experimental results are shown in Figure 23.

The more capacity and connectivity added to a Lightning 9900™ Series system, the more total internal bandwidth increases.

Shared bus architectures are susceptible to performance bottlenecks.




Figure 23 – Test results show the advantages of the 9960 Hi-Star™ architecture when compared to a shared bus architecture.

⁴ See Appendix A for configuration details.

*Bottleneck-free
performance with the
Lightning 9900™ Series.*

Internal Switching Overcomes the Limitations of Shared Bus Architectures

It can be seen from Figure 24 that a reduction occurs in the OLTP performance for the shared bus architecture of the 7700E system. This occurs at Point 1, when the DSS application begins at about 45 minutes into the test. It depicts the effect that the DSS application (full table scans) has on the OLTP workload due to the fact that both applications compete for bus bandwidth. In the Lightning 9900™ Series system, with its Hi-Star™ architecture, however, each workload essentially has its own internally routed paths. This means that the main reason for performance decline for the OLTP application i.e., bus contention, does not occur when the DSS application starts. Although there may be some contention at the cache, there are 16 simultaneous operations concurrently happening at the cache. This far exceeds the demands of most application requirements, which can be measured with the limited multiple shared bus architectures on the market today.

Test Results

The DSS application starts performing full table scans while the OLTP application has been running by itself for approximately 45 minutes. The 9960's OLTP application started at a Transaction Per Minute (TPM) rate of about 2350 and sustained this rate during the start of the DSS application and throughout the DSS application's run. The 9960's DSS application started with a transfer rate of about 220MB/sec or over 2.3 million rows per second. There was no reduction in TPM.

In other words, the DSS application did NOT affect the 9960's OLTP application performance as it does on the shared bus controller system.

Also when the 9960's OLTP application completes the DSS application's transfer rate is maintained at 220MB/sec or over 2.3 million rows per second. Again, there is no reduction in MB/sec or rows per second with the Lightning 9900™ Series.

The Lightning 9900 Series Software Solutions Are the Best in the World

9

Customers and analysts regard the Hitachi Freedom Storage Lightning 9900™ Series as the most advanced products in the enterprise storage system market. The **Software Solutions Guide** discusses the advantages of this suite of leading-edge software products that enhances the Lightning 9900™ Series product lines. This report is available from Hitachi Data Systems in PDF form

http://www.hds.com/products/software/solutions_guide.html. Particular emphasis is devoted to the business benefits of this robust suite of software features in the context of three generally accepted categories of benefits that are enabled by enterprise class intelligent storage system software. These three business objectives can be classified as follows:

1. Increased IT service levels in availability and performance for global and non-stop data access.
2. Simplified IT operations for centralized storage and data management.
3. Accelerated IT deployment of new applications and new systems for business agility.

Hitachi Freedom Storage™ software solutions are the most advanced available.


Hitachi Freedom Storage Software Solutions for the Lightning 9900™ Series	Business Benefit		
	Increased IT service levels in availability and performance	Simplified IT operations	Accelerated deployment of new applications and new systems
Synchronous Remote Copy - OpenSystems	✓		
Synchronous Remote Copy - S390®	✓		
Asynchronous Remote Copy - OpenSystems	✓		
Asynchronous Remote Copy - S/390®	✓		
IBM® GDPS™ Compatible Remote Copy	✓		
Software Asynchronous Remote Copy	✓		
Non-disruptive PiT (Point-in-Time) Copy™ - S/390® (Hitachi NanoCopy™)	✓		
Non-disruptive data replication - Open ShadowImage™	✓		✓
Non-disruptive data replication -S/390® ShadowImage™	✓		✓
Data Migration Service to the Lightning 9900-S/390™	✓	✓	✓
Automatic Path Failover - (Dynamic Link Manager™)	✓	✓	
Host Failover and Parallel Database Clustering	✓		
Open backup to mainframe tape using standard backup software	✓	✓	
HARBOR® Backup and HYPERTape	✓	✓	
Restore of single files only - S/390®	✓	✓	
VERITAS® support, volume and file management utilities	✓	✓	
Systems Management - Resource Manager 9000™	✓	✓	✓
S/390® to Open file conversion - RapidXchange™	✓	✓	✓
Open to Open file conversion - RapidXchange™	✓	✓	✓
Open and S/390 channel file transfer - RapidXchange™		✓	✓
LUN Security/Hitachi SANTinel™	✓	✓	✓
Automatic performance tuning -Hitachi CruiseControl™	✓	✓	

Table 4 – Summary of Hitachi Freedom Storage™ software products and the business objectives they serve.

The Hitachi Freedom Storage™ software solutions for Lightning 9900™ Series systems support an enterprise's strategic goal of accessing any information, on any computer, located anywhere, at any time. The many advanced functions available on Lightning 9900 Series™ hardware are initiated, managed, and controlled through these powerful software programs. Table 1 summarizes how Hitachi software solutions can be mapped to the business benefits mentioned above.

The Hitachi Freedom Storage™ software solutions deliver enterprise-wide coverage of on-line data copy/relocation, data access/protection, and storage resource management. Customers have the freedom to choose the precise solution – or combination of solutions – appropriate for their environment. The Hitachi Data Systems Software Solutions are illustrated in Figure 24.

Figure 24 – Hitachi Data Systems offers a broad range of remote copy, data duplication, and data migration software solutions.



Increased IT service levels in availability and performance through Business Continuity Planning (BCP) software and services

In the New Economy, performance and availability go hand in hand. If a Web site is slow in delivering requested information, it might as well be offline. Of course, if a Web site is down for any length of time, the business may not be around for long. Building on Hitachi Data Systems' reputation for "bulletproof" reliability, the Lightning 9900™ Series offer complete redundancy and hot-replaceable components, delivering maximum uptime.

The Lightning 9900 Series delivers unsurpassed uptime and performance.

To this already robust platform, Hitachi Data Systems adds a number of solutions to ensure quick recovery from acts of nature, human errors, application errors, and malicious attacks. In the world of e-commerce, it is no longer enough to have a backup copy of one's data; businesses must also plan for rapid recovery from such outages. Hitachi Data Systems is the only storage vendor that is providing non-disruptive copy solutions that assure fast recovery and full data integrity.

Remote Copy, Data Duplication, Data Migration Defined

It is important to understand the differences in what has become known as the triumvirate of copy software product categories. The jargon of copy software alternatives is made even more confusing when traditional backup methods are considered. Advances in technology have allowed new words such as “real-time,” “point in time” (PiT), and “snapshot” to creep into the language of enterprise class storage. Copy products allow an enterprise to replicate, protect, and share data in dynamic new ways. The three main terms used for copy software are:

The jargon of copy software alternatives is made even more confusing when traditional backup methods are considered.

Remote Copy

A term that refers to the operations procedure of continuously sending updates to a remote geography in order to provide a time-consistent copy of that data. Synchronous Remote Copy is typically used over short distances and careful consideration of performance requirements is required. Asynchronous communication techniques with methods of insuring data sequencing by timestamping are used for longer distances. The purpose of remote copy is to protect the data in the event of a business interruption.

PiT Copy

A process that creates a “static” image of data at a specific time (e.g., backups) is generally referred to as Point-in-Time (PiT) snapshots.

Data Migration

Software that moves data permanently from one storage device to another. This feature is different from data duplication in that at the end of the process there is only one copy of data. The purpose of data migration is to consolidate storage or upgrade to new systems.

Of the literally hundreds of solutions that vendors have devised, only three basic functions are actually being performed:

- Point-in-Time copy (PiT copy)
- Real-time copy (disaster recovery)
- Data relocation or data migration.

There are many reasons for the proliferation of products and techniques, but the primary motivation is to improve on an existing design or technique. If an existing product is insufficient in some form, that becomes the impetus for creating a newer product.

For example, early implementations of copy software would reduce backup windows from hours to minutes by not requiring the application to endure an outage for the duration of the Point-in-Time copy process. The application would still require an outage (or quiesce), but merely for the duration necessary to initiate the process (seconds or minutes). Using these products, however, the backup is not guaranteed and could subsequently fail. This led to the development of such products as Snapshot (StorageTek), ShadowImage™ (Hitachi®), and TimeFinder (EMC®).

Of the literally hundreds of solutions that vendors have devised, only three basic functions are actually being performed.

Business Goals served by Copy Software

Worldwide commerce and industry have become increasingly dependent on IT resources to conduct business. Even a temporary loss of critical applications or data access can cause serious economic impact to a company, and an extended outage can threaten a company’s existence. Regulatory and competitive pressures, coupled with the potential financial impact of unavailable systems, have motivated IT executives to address availability as one of their top priorities. Primarily because of these strict requirements on availability, storage systems have been enhanced over the years to include “copy service” functions. These hardware-assisted functions have allowed customers to minimize outages associated with copying data for backups, maintain disaster recovery sites for critical data, create business intelligence applications that use copies of production data, or simply move data to newer hardware with a minimum of outage times.

Real-time/point-in-time copies heighten both data security and mobility for a variety of critical activities, including protection, testing, disaster recovery, and warehousing. No matter where the data is located, it can be replicated quickly and safely within the same data center or between different data centers. This ability to share data within the enterprise positions a business to respond quickly to competitive pressures and meet the fundamental business goal of business agility. As shown in Table 5, software products provide the functions of data migration, data duplication and remote copy. These functions are critical to perhaps the two most important business objectives of the enterprise: 1) to maintain business continuity in the face of adversity, such as disaster; 2) to rapidly deploy new IT applications for business intelligence, data warehousing/OLAP, for data center relocation, or for new application testing.

Real-time/point-in-time copies heighten both data security and mobility for a variety of critical activities, including protection, testing, disaster recovery, and warehousing.

Table 5 - Hitachi copy products provide Remote Copy, Data Duplication, and Data Migration.

Hitachi Copy Products	Data Migration	Data Duplication	Remote Copy
Hitachi TrueCopy™ Synchronous	✓	✓	✓ - 25 Miles ⁵
Hitachi TrueCopy™ Asynchronous	✓	✓	✓ - Unlimited distance
NanoCopy™		✓	✓
Software Asynchronous Remote Copy™ (HXRC)		✓	✓
Hitachi ShadowImage™		✓	
Hitachi Data Protection Services	✓	✓	✓
Hitachi RapidXchange™		✓	
HARBOR® File Transfer™ (HFT)	✓		

⁵ 25 miles is the current practical distance limitation for direct connection using ESCON® fibre. This is changing with the advent of new networking devices. Synchronous remote copy may also use telephone networks. However the performance of the production application must be considered when using synchronous remote copy over long distances.

Centralized and Open Systems Management Software Simplifies IT Operations

10

A powerful and open systems management philosophy

The Hitachi Freedom Storage® software solutions support an enterprise's strategic goal of helping customers focus on their business issues instead of on deploying enabling technologies by providing powerful centralized management capability. The many advanced functions available on Hitachi Freedom Storage® hardware are initiated, managed, and controlled through the powerful HiCommand™ management framework.


The Hitachi Freedom Storage® software solutions deliver enterprise-wide coverage of on-line data copy/relocation, data access/protection, and storage system resource management. Customers have the freedom to choose the precise solution—or combination of solutions—appropriate for their environment.

The many advanced functions available on Hitachi Freedom Storage hardware are initiated, managed, and controlled through the powerful HiCommand software program.

HiCommand™ allows systems management of Hitachi storage and software through the enterprise's vendor of choice

Hitachi Data Systems® software and hardware solutions are managed through the powerful HiCommand™ management framework. This powerful management tool reduces operations expense and increases business agility while enabling operational excellence. HiCommand™ includes a GUI-based browser, open APIs, Command Line Interface (CLI), server, and host agents that access HDS software management functionality. The program encompasses storage resource management, configuration management and automation, automated data replication and recovery, performance management and optimization, and related functionality. Most importantly HiCommand™ can plug into other management frameworks. This feature provides the utmost in an open architecture and leverages an enterprise's existing investments in software and skills. The HiCommand™ architecture is shown in Figure 25.

Figure 25 - HiCommand™ allows management of virtually all Hitachi hardware and software from the customer's platform of choice.



HiCommand™ enables an enterprise to manage Hitachi storage products and software solutions.

HiCommand™ is an important Hitachi management tool in that it provides a comprehensive storage management software framework that enables an enterprise to manage all supported Hitachi storage products and many Hitachi software solutions:

1. **HiCommand™ for storage management** including alerts, configuration parameters, LUN mapping.
2. **HiCommand™ for SAN management** including WWN Discovery and Display.
3. **HiCommand™ for storage configuration automation** of the Lightning 9900 and Thunder 9200 Configuration Wizard.
4. **HiCommand™ for data archive** with both event and time-based backup.
5. **HiCommand™ for performance management and optimization** that allows correlation of client applications, RDBMS servers, storage capacity, and performance.

HiCommand™ functionality is available as a standalone product from Hitachi and through the Hitachi Data System® independent software vendor partners. Hitachi has long recognized the importance of the ISV vendors and is working with vendors whose offerings are complementary. This is accomplished through both joint development and engineering, joint certification and, where appropriate, marketing agreements. Both Hitachi and their software alliance partners believe this open architecture will best benefit customers by allowing them to manage their Hitachi storage through the enterprise's vendor of choice.

The components of HiCommand™

HiCommand™ consists of three components:

1. The **HiCommand™ Server** resides on any Windows® NT/2000, or Sun Solaris server, and connects to the Lightning 9900™ Series system(s) over a private LAN. A single HiCommand™ server can manage multiple Lightning 9900™ Series storage systems, providing a common point of control that delivers:
 - Out of band storage area networks (SANs), networked attached storage (NAS), or direct attached storage (DAS) connections enable data to run between the server and storage at lightning speeds.
 - Seamless integration between HiCommand™ and software solutions from industry leaders such as VERITAS® Software, Sun Microsystems®, BMC Software®, IBM®, Tivoli®, and Computer Associates® through use of XML-based APIs.
 - Remote and local access control offers a centralized approach to managing storage from any location.
2. A **GUI, browser-based interface** that is easy to learn and allows for simple and efficient storage administration. A browser from anywhere can control the HiCommand™ server. In addition, HiCommand™ also supplies an optional CLI (Command Line Interface).
3. **Host Agents** reside on the customer's application servers and 'push' information back to HiCommand™, including percentage of file utilization of LUNs, HBA WWNs, operating system SCSI addresses; and other useful device specific and file level information.

Overview of how the Systems Management Software supports business objectives

To be successful in today's business environment, companies must attain the highest levels of business agility, with maximum IT system efficiency and performance. The "storage centric" model of computing is widely accepted today. In this model "enterprise data" is the most important element of the IT infrastructure. Intelligent storage and network management software suites like the Hitachi Resource Manager™ therefore are at the heart of the enterprise's ability to achieve high levels of performance and availability to meet business objectives. Explosively growing storage in both centralized and distributed environments must be managed intelligently while ensuring that storage data remains available to all who need it.

The complexities of heterogeneous computing environments make storage management a daunting task. Hitachi Data Systems offers the ideal solution to this difficult challenge, Hitachi Resource Manager™. This comprehensive systems and software management offering brings together Hitachi Graph-Track™, Hitachi Virtual Logical Volume Image Manager™, Hitachi FlashAccess™, and Hitachi LUN Manager™ into one complete package. Designed for the Lightning 9900™ Series, these diverse capabilities provide world class functionality.

Hitachi Graph-Track™ provides the ultimate in reliable, easy-to-use hardware performance and availability management software. Hitachi VLVI Manager™ enhances Lightning 9900™ performance by defining logical volumes in many sizes to fit workloads better and utilize the full capacity of the system. Hitachi FlashAccess™ lets users "lock" Lightning 9900™ data into cache in real time for ultra-high performance. And Hitachi LUN Manager lets users define, configure, and maintain open systems logical units (LUNs) without hardware vendor intervention.

Hitachi Resource Manager™ provides the ultimate in reliable, easy-to-use hardware performance and availability management software.

Hitachi Resource Manager™ 9000

The following is a summary list of the many ways in which the Hitachi Resource Manager™ can simplify IT operations, improve availability and performance service levels for business critical OLTP and DSS applications, and thereby enable the more rapid deployment of new applications by managing data more intelligently.

Service Level Management

- Policy Management
- Security Management
- Automation
- Monitoring/Reporting
- Storage Management

Deployment Management

- Asset Management
- Change Management
- Capacity Planning

Compliance Management

- Event Management
- Performance Management
- Quota Planning
- Accounting Management
- Problem Planning

SNMP support ensures compliance with all major systems management platforms

Hitachi Resource Manager™ supports Simple Network Management Protocol (SNMP). SNMP is the most widely used network management protocol in the marketplace. System Information Messages (SIMs) are converted to SNMP Management Information Blocks (MIBs) for use with Hitachi Resource Manager™ software, or the VERITAS® Storage Manager™, Tivoli®, CA Unicenter TNG, HP® OpenView®. With CA Unicenter TNG, “super MIBs” allow the control of the Lightning 9900™ using CA Unicenter commands. This out-of-band⁶ reporting format ensures that Hitachi Resource Manager™ 9000 can determine the status of your Lightning 9900™ storage system even if the data channel path is down.

Hitachi Graph-Track™ takes the guess-work out of performance management

As the cornerstone of the software utilities of the Hitachi Resource Manager™ 9000 software suite, the Hitachi Graph-Track™ GUI is widely recognized as user friendly. Icons and pull-down menus define and display data, while point-and-click capabilities and on-line help further enhance its ease of use. All Hitachi Graph-Track™ data can be exported in several formats quickly and effortlessly, for use in other data analysis and reporting programs.

Hitachi Graph-Track™ gives users a more reliable and centralized way to manage performance. Available for S/390®, UNIX®, and Windows NT/2000® environments, this unique tool monitors hardware performance and supplies complete system storage

*The Hitachi
Graph-Track™ GUI
is widely recognized as
user friendly.*

⁶ “Out of Band” refers to the path that system management messages travel which is a different path from the data. “In band” messages would compete with data for bandwidth.

information at the touch of a button. The robust GUI-oriented command screen for Hitachi Graph-Track™ is shown in Figure 26. Running on a PC attached to a dedicated LAN, Hitachi Graph-Track™ displays real-time or historical data for all connected Lightning 9900™ systems, helping users identify important peaks in utilization and ongoing trends in processing. With Hitachi Graph-Track™, users can analyze discrete storage activities and determine the precise impact of each activity on system operation. It scrutinizes activity all the way to the logical device level, reporting on channel interface processor usage rates, cache usage rate by function, and logical device utilization. This data pinpoints specific activities that may have a negative impact on operations and lets the user fine-tune the system so that it can be used to maximum advantage.

Hitachi Graph-Track™ also highlights valuable cache information. Cache read/write and read-hit ratios are reported in real time. Each system connected to Hitachi Graph-Track™ can be tuned separately.

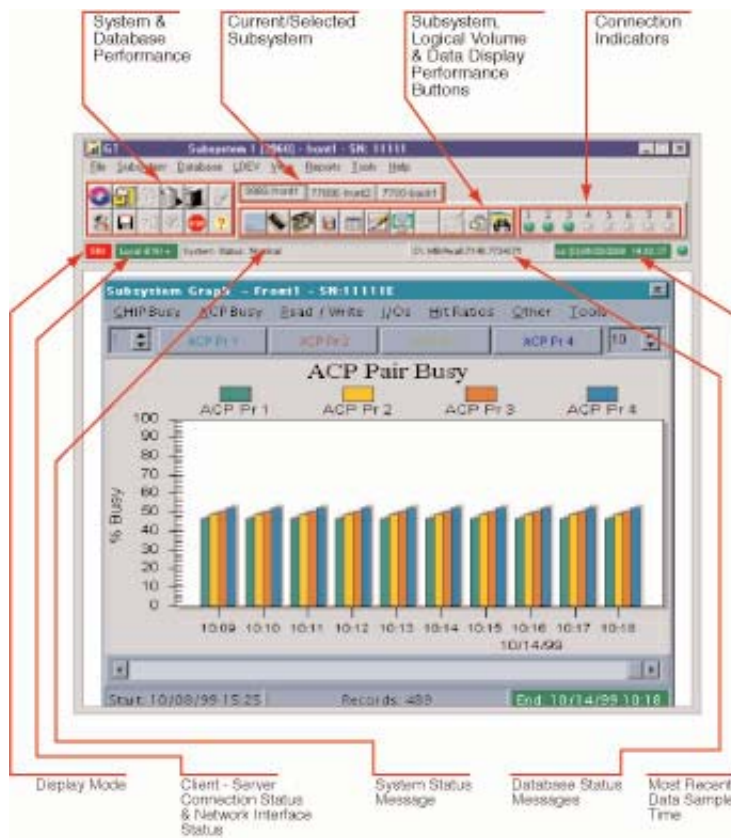


Figure 26 – The GUI interface on Hitachi Graph-Track™ simplifies performance management.

Hitachi LUN Manager simplifies configuration management and reduces staffing costs

Hitachi LUN Manager is an open systems management utility. With Hitachi LUN Manager, open systems Logical Units (LUNs) can be defined, configured, and maintained. There is no more waiting for the hardware vendor to come and make configuration changes. Hitachi LUN Manager includes an easy-to-use, GUI-based interface that allows the easy definition of paths for LUNs, the reconfiguration of LUN-to-port assignments, or the viewing of the Lightning 9900™ remote service information messages. Because the Hitachi LUN Manager can assign multiple paths to a single LUN, support of alternate path failover, path load balancing, and clustered systems is possible. Running on a standard Windows®-based PC connected to the storage subsystems by a dedicated LAN, Hitachi LUN Manager can support up to eight Lightning 9900 systems.

With Hitachi LUN Manager support of alternate path failover, path load balancing, and clustered systems is enhanced.

Hitachi LUN Manager also features logical unit size expansion that dramatically improves LUN flexibility. Up to 36 physical LUNs can be presented to the operating system as a single, large LUN. This capability simplifies storage management because there are fewer LUNs to manage. Two levels of password protection, user and administrator, ensure maximum data security. Users can view only Lightning 9900™ LUN configuration information, while administrators can access all LUN configuration information and functions. Administrators can customize access privileges for individual users, thus providing maximum flexibility and control of Hitachi LUN Manager capabilities.

Hitachi Virtual Logical Volume Image Manager allows very large to very small virtual volume configuration

Hitachi Virtual Logical Volume Image Manager improves performance.

Hitachi Virtual Logical Volume Image Manager is a mainframe software utility that optimizes Lightning 9900™ capacity utilization by allowing users to configure multiple *virtual LVIs* in place of a LVI. Data volumes as small as a single cylinder or as large as a full 3390-9 can be defined. Volume size is determined in cylinder increments. Each virtual LVI requires one physical address, with a maximum of 4096 addresses per Lightning 9900™ system. Different types of LVIs can coexist within an array group with no need to convert them to a common LVI. This feature maximizes array group capacity and enhances configuration flexibility. Hitachi Virtual Logical Volume Image Manager improves performance by reducing logical device contention and operating system queuing. It also boosts remote copy performance by avoiding the need to copy the entire volume.

Hitachi FlashAccess™ allows data to be “locked and unlocked” in cache on-the-fly

Hitachi FlashAccess™ dramatically improves performance.

Hitachi FlashAccess™ is a software utility that allows users to dynamically “lock” and “unlock” data into cache in real time. Read and write functions are then performed at cache speeds, with no disk latency delay. With Hitachi FlashAccess™, a portion of cache memory can be allocated to specific data. Users can add, delete, or change Hitachi FlashAccess™ managed data at any time quickly and easily.

Defined by the Logical Volume Image (LVI) for mainframes or Logical Unit Number (LUN) for open systems, cache data can be as small as a single track or as large as an entire volume. For increased configuration flexibility, Hitachi FlashAccess Manager™ offers multiple modes of operation. It can be used in conjunction with Hitachi RapidXchange™ to increase the speed of data transfer and, therefore, improve performance of mainframe to open systems data exchange. Hitachi RapidXchange™ offers high-speed data transfer among OS/390®, UNIX®, and Windows NT/2000® platforms. Hitachi RapidXchange™ transfers can be open-to-open as well as mainframe-to-open.

Hitachi CruiseControl™

Hitachi CruiseControl™ automatically monitors, analyzes and moves logical volumes to eliminate “hot-spots.”

For open systems and S/390 volumes, Hitachi CruiseControl™ provides automatic performance tuning for Lightning 9900™ Series products by automatically eliminating performance bottlenecks. Utilizing the Lightning 9900™ many high-speed internal paths to optimize data placement CruiseControl™ automatically monitors, analyzes and *moves logical volumes* to eliminate “hot-spots” and provides load balancing to maintain predetermined performance levels. Independent of the type of host attachment, the CruiseControl™ makes recommendations for administrator approval in either automatic or assisted mode. CruiseControl™ replaces time-consuming and sometimes error-prone manual load balancing with simple automated procedures. This ensures long-range optimal performance and reduced cost of ownership.

Storage Maximizer

The Hitachi Storage Maximizer™ is a solution that uses two software products – Hitachi CruiseControl™ and Hitachi FlashAccess™ software for mainframes and open systems. This allows customers to manage the cost of storage, as data is consolidated onto fewer better-managed systems. Storage administrators can “fill it up and max it out,” confident that they can meet performance requirements all the way to 37 TB raw capacity on a Lightning 9900™ using 73 GB drives. Hitachi CruiseControl™ then takes over and eliminates “hot spots,” thereby automating performance to service level objectives, while Hitachi FlashAccess™ locks data into cache, guaranteeing data access at memory, not disk, speed.

With the massive consolidation made possible by Lightning 9900™ Series capabilities, customers will be able to lower not only the original purchase costs, but also get additional savings from lower storage management costs, lower software and maintenance costs, and reduced real estate and environmental costs. This is often the major payback for Lightning 9900™ Series products. Reduced per-megabyte administration costs, by itself, can provide a substantial return on investment.

Hitachi SANTinel™

The Hitachi SANTinel™ controls host access to Lightning 9900™ Series LUNs in open systems and LVIs in S/390 systems for SAN environments. This enables users to restrict server access to LUNs in pre-defined “zones” in order to protect data from unauthorized access.

Transaction Processing Facility/Multi-Path Lock Facility (TPF/MPLF)

The Lightning 9900™ supports the Multi-Path Lock Facility (MPLF) for IBM® highest performance transaction processing software environment - the Transaction Processing Facility (TPF). TPF is used by many of the world’s largest Customer Reservation Systems and by the world’s largest financial institutions. In either native TPF mode or under VM, MPLF provides extremely high performance *record-level locking*, which enables high levels of concurrent data access across multiple channel paths. RAID-1+, RAID-5, 3390-3, and 3390-9 LVIs are supported.

Professional Services and Support

11

Hitachi Data Systems® is consistently ranked Number One in the Industry

In numerous independent surveys on IT services organizations, Hitachi Data Systems® continually wins the highest ratings in terms of overall customer satisfaction. “Service Responsiveness” is the key Hitachi Data Systems® characteristic that allows the world renowned Hitachi Data Systems® service and support organization to ensure that Hitachi products operate at peak performance to complement all hardware and software in the enterprise.

Professional Services Overview

As new technologies gain acceptance, companies must decide on long-term plans and implementation schedules that cause the least disruption to business. It takes time to implement any large-scale technological change. The transition to new network topologies will see the coexistence of distributed and legacy systems, and with SCSI and Fibre Channel on SANs, SWANs, ESCON, and FICON. Hitachi Freedom Storage™ provides the comprehensive connectivity, management, and availability capabilities needed to handle this transition. These built-in product strengths are bolstered by the Hitachi Data Systems® professional services and service and support organizations to insure the optimal operation of hardware, software, and middleware for the enterprise.

Hitachi Data Systems® professional services specialize in infrastructure, hardware, software, and storage management services that provide a vendor-independent view of IT architecture. This provides sharp focus on ways to streamline operations, costs, and interoperability. Hitachi also excels at helping customers chart both the strategies and timelines necessary to remain productive and competitive. Whether an enterprise needs assistance with SANs, Business Continuity Consulting and Implementation, DFSMS performance/capacity issues, migration planning, decisions about platforms and architectures, or maximization of IT investments, Hitachi Data Systems® has the expertise and the resources to guide an enterprise toward the best business solution. A few of the many Hitachi Data Systems® service offerings are highlighted here.

Hitachi Data Systems® professional services provide comprehensive connectivity, management, and availability capabilities.

Cost of Risk Analysis methodology (CORA)

As a first step in data movement assessment, the CORA service is sometimes recommended. CORA is an important component of the Hitachi consulting services. This service can help identify the cost, benefit, and ROI in the enterprise’s computing infrastructure to reduce outage windows and the impact of an outage. Hitachi findings and recommendations will be in accordance with the enterprise’s business drivers and in clear terms that IT executives can take to their executive team. See also Chapter 2.

CORA can help identify the cost, benefit, and ROI in the enterprise’s computing infrastructure.

Data Protection Services

The Hitachi Data Protection Services team is trained in architectural analysis, configuration planning, and enterprise assessment. The team uses structured methodologies that promote consistent results. By applying Hitachi’s world class software solutions, the Hitachi Data Protection Services implement essential copy solutions by moving data in

order to protect it. The Data Protection Services team uses software in all three categories of data movement:

- 1) Data Migration
- 2) Real-Time Copy
- 3) Point-in-Time Copy

The team will perform data relocation and migration tasks that establish rock-solid backup and disaster recovery copy facilities to keep a business running in the event of man-made or natural disaster. Often the need for an enterprise to allocate staff for these tasks is eliminated.

With Data Protection Services, an enterprise can accelerate its return on existing investments, streamline data management processes, eliminate redundant software and costs, and reduce time to recovery.

As a first step in a typical engagement, an Engagement Manager works with the enterprise to determine requirements based on the use of a proven Data Movement questionnaire and the presentation of examples of prior engagements. Disaster recovery, backup, data warehousing, business continuance, testing and development, and data center consolidation are the types of projects that most frequently use this service. Then the team reviews system, network, storage components, and phases of solution delivery with the client. After a clear understanding of requirements, a detailed professional services proposal is presented and reviewed with the enterprise. Hitachi Data Systems uses “best of breed” consulting techniques and has an excellent reference list of satisfied clients that may be contacted prior to beginning of the engagement. Examples of the Hitachi *suite of data protection services* are discussed below:

Data Protection Services Lab

Hitachi Data Systems has designed a laboratory proving ground for developing the tools and expertise an enterprise needs to fully enable the functions of Hitachi data storage systems and software. The Data Protection Services Lab is dedicated to resolving software and hardware issues that affect storage system consolidation, data center productivity, and business continuity.

Remote Copy Assessment and Implementation Service

This service assists the enterprise in implementing a remote copy process for either disaster recovery or for the rapid deployment of new IT systems such as relocation of a data center, population of a geographically dispersed data mart, or testing new applications. Hitachi Data Systems provides expert consultants to assist in assessing processes and procedures to ensure an optimal implementation of the remote copy process. Hitachi also offers services for developing remote testing at a disaster recovery hot site to validate the proper configuration of the remote copy process in terms of hardware, software, and processing method. Enterprise IT professionals will be trained to address remote copy as a strategic tool of the data center to meet the objectives of business continuity, disaster recovery, or rapid deployment of new IT systems.

Hitachi NanoCopy™ Implementation Service

Hitachi NanoCopy™ lets an enterprise make point-in-time copies without disturbing critical applications or causing any disruption to end-user operations. It is the only completely non-disruptive long distance remote copy technique that can be used for Disaster Recovery with confidence. Hitachi Data Systems experts are available to help customize and install this advanced alternative to existing remote copy technologies. The Hitachi service helps the enterprise adapt Hitachi NanoCopy™ to the most rigorous requirements, insures maximum operations efficiency and availability, and shortens enterprise recovery time dramatically in the event of a disaster.

Hitachi Data Systems provides expert consultants to assist in assessing processes and procedures.

Continuous Availability Service

Keeping critical systems available at a level 4 or 5 on the Scale of 9s is no easy task. It involves the best of hardware, software, and operations practice. Hitachi Data Systems® professional services experts review the entire backup and recovery software portfolio, operations procedures, and hardware configurations of the enterprise. Weaknesses that threaten the IT environment are then systematically analyzed and eliminated. Hitachi consultants employ best consulting practices in both open systems and MVS® application backup and recovery (both on-site and off-site) to conduct the assessment. As a first deliverable, a baseline will be created and presented for analysis of potential strengths and weaknesses of the current processes, equipment, software, and practices. Recommendations are then discussed to improve availability and reduce the risk of data loss.

Hitachi consultants employ best consulting practices in both open systems and MVS® application backup and recovery.

On-line Data Migration Service

The Hitachi Data Migration Service helps migrate data from existing systems to newly installed systems while minimizing the impact on mission-critical applications. Hitachi Data Migration Service features a four-phase approach that includes assessment, planning, migration, and post-migration support. The strength of this service lies in two key areas:

- 1) Hitachi Data Systems® professional services methodology/procedure and skills
- 2) The outstanding reliability of Hitachi Freedom Storage™ products

Hitachi Data Migration Service can provide an unprecedented level of data protection and integrity. This is accomplished through a complete solution that moves terabytes of data quickly and efficiently.

In S/390® environments, terabytes of data can be migrated to Hitachi Freedom Storage™ systems from other vendor systems in a matter of hours while applications are on-line and processing remains completely uninterrupted. Hitachi Data Migration Service provides the utmost in availability, allowing users to access data continuously throughout the migration process. Its superior capabilities reduce migration times dramatically, saving considerable expense.

Hitachi Data Migration Service can provide an unprecedented level of data protection and availability during migration.

SAN services

Controlling explosive data growth and the subsequent increases in storage costs is a daunting task. Storage Area Networks (SANs) can help an enterprise achieve increased scalability, availability, and reliability. A SAN is a network of storage systems and servers that enables data to be pooled within an enterprise. SANs enable higher scalability; increased addressing; centralized management of storage systems and backup to either disk or tape; 100MB-per-second Fibre Channel (FC) connectivity for distances of up to 10 kilometers, and high availability.

Hitachi Data Systems offers an entire suite of SAN services to assist clients in planning and implementing SANs.

The main benefits of SAN implementation are:

- Reduces cost of storage ownership through the more efficient use of resources
- Shortens backup windows significantly through parallel backup
- Saves server processing cycles through server free backup
- Reduces the impact of backup on LAN performance
- Enables consolidation of tape backup systems into larger silos
- Permits wide access to distributed critical data on servers and workstations
- Saves on training and personnel costs in highly heterogeneous environments

Hitachi Data Systems® offers an entire suite of professional services that will assist clients in planning and implementing Storage Area Network (SAN) solutions to optimize management and control data across the enterprise. The modular services include:

- The SAN Enterprise Infrastructure Assessment (EIA) Service
- The SAN Strategic and Tactical Design Service
- The SAN Installation Service

Hitachi Data Systems® has a highly experienced team of SAN specialists to aid enterprises in implementing SANs.

Hitachi Data Systems® Interoperability Laboratory Service

Interoperability Laboratory Service provides clients with the knowledge of pre-tested components of a SAN.

The Interoperability Laboratory Service provides clients with the knowledge of pre-tested components of a SAN in various network topologies. Customers, however, will not be limited to the use of these proven interoperable SAN elements. To assure that a wide variety of components will operate effectively in Hitachi Data Systems® SAN solutions, Hitachi Data Systems® has established a new multi-million dollar Interoperability Laboratory at its Santa Clara headquarters to test new SAN elements.

Among the Interoperability Laboratory's earliest projects was the replication and operation of the SAN developed for a major banking customer and a major telecommunications customer company. The experience and methodology gained by the lab's many such customer experiences, enables Hitachi Data Systems® to quickly tailor effective SAN designs to meet wide-ranging customer requirements in a variety of business environments.

Enterprise storage and availability management services

Hitachi Data Systems® offers a variety of consulting services and software utilities to help extend the life of an enterprise's storage investment and define ways to achieve greater information functionality throughout the enterprise. These storage services cover storage management, availability, disk utilization, performance and tuning, data center cabling and configuration, and design/installation of fiber optic components. Hitachi Data Systems® services help the enterprise take full advantage of the enterprise's IT resources, and enhance the continuous availability and integrity of the enterprise's mission-critical applications. In addition, Hitachi Data Systems® Professional Services consultants can work with the enterprise's business continuity planners to exploit the capabilities of the Lightning 9900®.

Software Portfolio Review and Analysis Service

Hitachi Data Systems® uses a software utility to help an enterprise organize its software portfolio.

Hitachi Data Systems® Software Portfolio Review and Analysis Service uses a software utility to help an enterprise organize its software portfolio. Once this is accomplished, the Software Portfolio Review and Analysis Service helps clients develop a management philosophy and improve flexibility in negotiating with vendors. Managing a software portfolio as a business asset involves many important steps and raises critical questions about market and strategy. Together with Hitachi Data Systems® "best of breed" partners this methodology is based on the key disciplines of software asset management practices. This proven methodology focuses on maximizing the return on investment to acquire the greatest possible value.

System Configurations for the Performance Experiment A

Configurations of the HP and Sun Servers and Configurations of the Lightning 9960™ Switched Architecture to a Shared-bus Architecture like most competitive products for the Performance Experiment

In the performance test discussed in Chapter 8, comparing each application resided on its own array groups. The only resource shared by the application is the storage controller (cache, paths, etc.), and two ACP pairs (there are four ACP pairs total). The comparison is between the Hitachi Freedom Storage™ 7700E, representing all shared bus-based architectures, and the Lightning™ 9960 system, representing a switched fabric architecture.

The application workloads consist of two very different types of database environments – an On-line Transaction Processing (OLTP) system and a Decision Support System (DSS). The OLTP workload runs on a HP platform consisting of the following configuration for both storage arrays:

- HP®/9000-K460 dual processor 180MHz, 768MB RAM, 4 Fibre Channel HBAs (FC-AL topology) direct connect to either the Freedom 7700E or the Lightning 9900™ Series system.
- HP-UX® 11.00.
- Oracle® Database Server 8.0.4 with 50MB of SGA (in order to drive I/Os to the storage arrays), 2KB database block size.
- OLTP database with simulated user drivers (similar to the TPC-C benchmark).
- 7700E and Lightning 9960™ system with 1GB cache, 4 Fibre Channel (FC-AL topology) ports, 2 ACP pairs, 6GB 12,030RPM drives, RAID 5 formatted. The following are differences between the 7700E and Lightning 9960™ system for the OLTP workload:
 - 7700E used OPEN-K volumes and a total of 48 disk drives (12 array groups total).
 - Lightning 9960™ system used OPEN-9 volumes (some doubling up of tables occurred because of this), 40 disk drives (10 array groups total), and 18GB 10,025RPM drives.

The DSS workload runs on a Sun platform and consists of the following configuration:

- 2 Sun E3000 dual processor 250MHz, 512MB RAM each, 2 32-bit JNI HBAs (point-to-point fabric topology) Fibre Channel switch (Ancor® MKII 16 ports) connected each.
- Sun® Cluster 2.1.
- Oracle® Parallel Server 8.0.5, 16KB database block size (256KB I/O block size).
- DSS database with one large table – 600037899 rows at 112 bytes each. Performing full parallel table scans with multiple instances.
- 7700E and Lightning 9960™ system with 1GB cache, 4 Fibre Channel (point-to-point fabric topology) ports, 4 ACP pairs, 6GB 12,030RPM drives, RAID-5+ formatted – 64 disks total (16 array groups). The following are

differences between the 7700E and the Lightning 9960™ system for the DSS workload;

- The 7700E used all 6GB drives.
- The Lightning 9960™ system used a mix of 18GB and 47GB drives.

Glossary of Terms

10BaseT

Ethernet with a data transfer rate of 10Mbps/sec.

100BaseT

Also known as Fast Ethernet with a data transfer rate of 100Mbps/sec.

AG

Array Group.

ACP

Array Control Processor. The group of processors on a PCB that performs the physical access of the disk drives across four FC-AL loops.

ACP pair

A combination of two ACPs. Designed in such a manner as to facilitate redundancy and performance.

Alert

A message or log that a computing element generates as the result of an error event collection and analysis. An alert indicates that there is a need to perform some service action, and it can be sent by a variety of methods to operations personnel.

API

Application Programmer's Interface or API. A standardized set of software commands (calls) that can be used to access a particular software program in a consistent and reliable way.

Array frame

Array frames are the three left and three right cabinets of the Lightning 9900™ Series systems housing HDDs (high density disks), power supplies, and batteries.

Array group

The physical arrangement of disk drives independent of RAID level used – 3 Data + 1 Parity for RAID-5+ and 2 Data + 2 Data for RAID-1+.

Asynchronous Remote Copy

The transmission of data between two devices that are not synchronized with a clocking scheme or other technique. The sender can send data at any time and the receiver can accept information when it becomes available. Synchronous

communication is an exactly timed stream of bits when the start of a character is located by using a clocking mechanism such as bipolar encoding. See also Synchronous, Hitachi TrueCopy™.

Availability

In computer science, availability refers to the degree to which a system or resource is capable of performing its normal function. Availability is measured in terms of Mean Time Between Failure (MTBF) divided by MTBF plus the Mean Time to Repair (MTTR). The availability equation is expressed as follows:

$$\text{AVAILABILITY} = \text{MTBF} / (\text{MTBF} + \text{MTTR}).$$

For example, a server that fails on average once every 5,000 hours and takes an average of two hours to diagnose, replace faulty components and reboot, would have an availability rating of $5,000 / (5,000 + 2) = 99.96\%$. This would correspond to a Level 3 rating using the Scale of 9s.

b

Abbreviation for “bit” where 8 “bits” compose a byte.

B

Abbreviation for byte or the equivalent of one character in text.

Back end

In reference to storage arrays, the back end includes the controllers, disk drives, and paths to the disk drives. On the Lightning 9900™ Series these are the ACP pairs, the Fibre Channel loops, and the Fibre Channel disks.

Business Continuity Planning (BCP)

An “umbrella” term covering both disaster recovery planning and business resumption planning. See also: Disaster Recovery.

Business Impact Analysis (BIA)

The process of analyzing all business functions and the effect that a specific disaster may have upon them.

Business Interruption

Any event, whether anticipated (i.e., public service strike) or unanticipated (i.e., blackout) that disrupts the normal course of business operations at a corporate location.

CA

Cache Adapter. The CA resides on the cache boards and connects to the Cache Switch (CSW).

CARB

Cache memory Arbitrator circuit. Hitachi designed logical circuit in the Cache Switch (CSW) used to arbitrate access to cache.

Cache

Cache (pronounced cash) can be either on-chip memory circuits in a microprocessor (e.g. L2 processor cache), a reserved section of main memory (e.g. system or server cache), or an independent, high-speed disk storage device (e.g. a Web cache). Two types of caching are commonly used in personal computers: memory caching and disk caching. Disk caching can dramatically improve the performance of applications, because accessing a byte of data in RAM can be thousands of times faster than accessing a byte on a hard disk. When data is found in the cache, it is called a cache hit, and the effectiveness of a cache is judged by its hit rate.

CHIP

Client Host Interface Processor.

C-HSN

Cache Hierarchical Star Network.

CHT

Channel adapter Tachyon. (For Fibre Channel interfaces).

Client/Server Architecture

Client/Server Architecture is a network architecture in which each computer or process on the network is either a client or a server. Servers are powerful computers or processors dedicated to managing disk drives (file servers), printers (print servers), or network traffic (network servers). Clients are PCs or workstations on which users run applications. Clients rely on servers for resources, such as files, devices, and even processing power.

CMA

Control Memory Adapter.

CM-HSN

Control Memory, Hierarchical Star Network.

Cold site

An alternate facility that is void of any resources or equipment except air-conditioning and raised flooring. Equipment and resources must be installed in such a facility to duplicate the critical business functions of an organization. Cold sites have many variations, depending on their communication facilities, UPS systems, or mobility. See also: Shell site, Recovery site, Alternative site.

Consistency Groups

A software copy technique of suspending multiple volumes at the same time.

Control Frame

The Disk Controller or control frame, is the center cabinet of the Lightning 9900™ Series system, housing the channel switch, cache, power supplies, batteries, and host interface cards.

CORA (Cost of Risk Analysis)

CORA is a unique, software-based risk management system developed initially for the insurance industry by International Security Technology (IST) Inc. of New York City. In an alliance with IST, Hitachi Data Systems® is offering CORA applied to the BCP (Business Continuity Planning) industry as a professional service. See also Business Continuity Planning.

Cost Benefit Analysis

A major benefit of the CORA process is the Cost Benefit Analysis where each action to mitigate risk is evaluated in terms of cost and return on investment. See also CORA.

CPU

CPU is an acronym for Central Processing Unit. Can refer to either a processor chip such as Sun's SPARC™ or Intel's Pentium®, or to a processor chip or chips and support circuitry on a CPU board.

CSW

Cache Switch. A specially designed crossbar switch for use in the Lightning 9900™ Series. They are not Fibre Channel switches.

CTQ

Command-Tag Queuing.

DA

The Data Adapter resides on the CHT and connects the cache boards to the front end.

DASD

An IBM® coined acronym that means Direct Access Storage Device, i.e. a disk drive.

Data availability

Data availability refers to the degree to which a computer system is capable of providing data to its users. See also Availability.

Data Copy

A term that refers to both remote copy, data duplication, and data migration.

DD

Disk Director.

Data Duplication

Software that duplicates data as in remote copy or Point-in-Time (PiT) snapshots. Data duplication is differentiated from data migration in that with data duplication at the end of the process there are two copies of data and with data migration there is only one.

Data Migration

Software that migrates data from one storage device to another. This feature is different from data duplication in that at the end of the process there is only one copy of data.

Disaster Recovery

Disaster Recovery is the profession that plans to avoid disasters and to restore services after a disaster. The level of preparedness to respond to an interruption in services by implementing a disaster recovery plan to restore an organization's critical business functions.

Downtime

A planned or unplanned interruption in system availability. Planned downtime is usually for scheduled system maintenance and unplanned downtime usually includes business interruptions or disasters. See also Business Interruptions, Disaster Recovery, Business Continuity Planning (BCP), and Availability.

DR

Disaster Recovery.

DRR

Data Recovery and Regeneration.

DSS

Decision Support Systems is a term that refers to computer systems used in the areas of business intelligence such as data warehousing, data mining, OLAP, and others.

DTA

Data Adapter.

ECC

Error correction code.

EBCDIC

Abbreviation of Extended Binary-Coded Decimal Interchange Code. Pronounced eb-sih-dik, EBCDIC is an IBM® code for representing characters as numbers. Although it is widely used on large IBM® computers, most other computers, including PCs and Macintoshes®, use ASCII codes. See also ASCII, binary, synchronous, and asynchronous.

Electronic vaulting

The transfer of data to an offsite storage facility via a communication link rather than via portable media. Typically, electronic vaulting is used for batch or journaled updates to critical files to supplement full backups that are taken periodically.

ESA/390®

Enterprise Systems Architecture/390®. Often abbreviated S/390®, ESA/390® is an IBM® architecture for mainframe computers and peripherals. Processor systems that follow this architecture include the IBM ES/9000® family. In 2000, IBM® changed the server series name to eServer zSeries.

ESCON®

Enterprise Systems Connection architecture is an IBM® mainframe ESA/390® computer peripheral interface or connection between two mainframes for data exchange. The I/O interface utilizes ESA/390® logical protocols over a serial interface that configures attached units to a communication fabric. ESCON® is based on networking technology. ESCON® provides direct channel-to-channel connections between mainframe systems over fiber-optic links at distances up to 43 kilometers or 25 miles. ESCON® also provides a way for communication controllers and other devices to share a single channel to a mainframe.

ESCON Director

An I/O interface switch that allows the interconnection of multiple ESCON® interfaces in a distributed-star topology.

Ethernet

A Local Area Network (LAN) protocol developed by Xerox® in cooperation with Digital Equipment and Intel in 1976. Ethernet supports a star or bus topology and supports a data transfer rate of 10 megabits per second or 10 Mbps. The Ethernet specification formed the basis of the IEEE 802.3 standard, which specifies the physical and lower software layers. Ethernet uses the CSMA/CD access method for handling simultaneous demands and is one of the most widely implemented LAN standards. Ethernet is also known as 10BaseT. See also Fast Ethernet, Gigabit Ethernet.

Fabric

A group or network of switches arranged in such a way as to enhance scalability and connectivity.

Failover

Host, host bus adaptor, cable, or controller failover is the routing of all transactions to a second controller when the first controller fails.

Fast Ethernet

Fast Ethernet or 100BaseT, defined by the IEEE 802.3 standards committee, provides a 100 Mbps standard that is compatible with existing 10BaseT installations, preserving the CSMA/CD media access control (MAC) protocol.

Fast write

A write operation at cache speed that does not require immediate transfer of data to a disk drive. The system writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for de-staging (writing to disk). Fast write reduces the time an application must wait for the I/O operation to complete.

FC

Fibre Channel.

FCA

Fibre Channel Adaptor.

FC-AL

Fibre Channel – Arbitrated Loop.

FDN

Hitachi Data Systems' Storage Area Network™ (SAN) solution is known as Freedom Data Networks™ (FDN). FDN provides an open architecture that leverages SAN technology and offers organizations the freedom of choice in deploying data-access and data-sharing capabilities across the enterprise through Hitachi Data System® advanced implementation of Storage Area Networks (SANs). With FDN, customers gain a powerful new tool that enables the consolidation of servers and storage, increased data availability, centralized storage management, and the ability to back up and migrate data without affecting the performance of enterprise networks.

Fibre Channel

Fibre Channel is an ANSI standard designed to provide high-speed data transfers between workstations, servers, desktop computers and peripherals. Fibre channel makes use of a circuit/packet switched topology capable of providing multiple simultaneous point-to-point connections between devices. Fibre Channel is widely deployed in SAN implementations today. Standards for Fibre Channel SANs are being worked on by the Storage and Networking Industry Association (SNIA). The technology has gained interest as a channel for the attachment of storage devices, but has limited popularity as high-speed networks interconnect. Fibre channel can be deployed in point-to-point, arbitrated loop (FC-AL), or switched topologies. Fibre channel nodes

log in with each other and the switch to exchange operating information on attributes and characteristics. This information includes port names and port IDs and is used to establish interoperability parameters.

FICON

An ESA/390® computer peripheral interface.

File Backup

The practice of copying a file that is stored on disk or tape to another disk or tape is referred to as file backup. This is done for protection in case the active file gets damaged. Backup is considered “local copy” as opposed to “remote copy.” See also Remote Copy.

Front End

In reference to storage arrays, the front end is considered to be the interfaces or ports to the “real world,” the processors servicing these ports, and in some cases the cache memory. On the Lightning 9900™ Series systems, the front end consists of CHIP pairs.

Gigabit Ethernet

Provides a standard that supports data transfer at 1000 megabits/sec. Gigabit Ethernet is also called 1000BaseT Category 5 (copper wire) or 1000BaseX (fiber optic). There is a 10,000BaseT version of the Ethernet standard that will be widely available by 2002.

Gigabyte

Technically, a gigabyte is 1024 Megabytes although most disk drive manufacturers today define it as 1000MB. This is a source of great confusion in the storage industry.

GUI

GUI is an acronym that refers to a Graphical User Interface that is the software that controls the screen presented to a user in a computer application.

HARBOR File Level Backup/Recovery with Lightning™ 9900 Agent™

Harbor File Level Backup/Recovery is a software utility that provides for network or channel-based backup/restore of open systems client files to an OS/390® host using a common graphical user interface. Online, non-disruptive backup/recovery of popular database, mail and Enterprise Resource Planning (ERP) applications are optional. Client support is available for Windows®, OS/2®, UNIX®, NetWare™, and VMS host platforms. High-speed data transfer is also available using Hitachi RapidXchange™ and other channel-based technologies.

HARBOR File Transfer

This software utility provides automatic, reliable, and secure data transfer between OS/390® and open systems hosts using high speed ESCON® / FICON™ channels or network communications. HFT can be thought of as a very high speed FTP that does not use a network. High performance is enabled by allowing multiple file transfers to

run concurrently. Reliability is enhanced with retry logic, alternate path routing, and restart from point of failure. Ease of use is provided by a Java-client interface, a system monitor, scheduling support, and remote control of client functions from OS/390®.

HDD

High-density disks. Used in the Lightning 9900™ Series.

HDmS

The Hitachi Data Systems Migration Service™ (HDmS) is a Hitachi Data Systems® professional service that helps users migrate data from existing systems to newly installed systems while minimizing the impact on mission-critical applications. HDmS features a four-phase approach that includes assessment, planning, migration, and post-migration support.

Hitachi CruiseControl™

CruiseControl™ automatically monitors, analyzes and *moves logical volumes* to eliminate “hot spots” within a Lightning 9900 storage system and provides load balancing to maintain predetermined performance levels.

Hitachi Dynamic Link Manager™

Hitachi Dynamic Link Manager™ is a family of software utilities that is server based and enhances RAID systems by providing automatic failover and load balancing from server-to-RAID channel connection failures. This product allows systems administrators to take advantage of the multiple paths on a Lightning 9900™ by adding redundant SCSI connections between data servers and RAID systems. Hitachi Dynamic Link Manager™ therefore provides increased reliability and performance. Supported platforms include AIX®, Sun Solaris™, and Windows NT/2000®.

Hitachi FlashAccess™

FlashAccess™ is a Hitachi feature that allows specified (usually high access) data sets to be “pegged” or permanently placed in cache memory so they are not managed by the data movement algorithms of the Freedom Storage® system. The Hitachi Flash Access™ feature in the Lightning 9900™ can be used for either S/390 or open systems. Hitachi Flash Access™ is a software utility in the Freedom 9000 Resource Manager™ suite that allows the creating, deletion, and monitoring of data managed by the Hitachi Flash Access™ software. See also Freedom 9000 Resource Manager™.

Hitachi Graph-Track™

Graph-Track™ is a software utility in the Freedom 9000 Resource Manager™ suite that allows a robust set of system and network management utilities and provides graphical reports for Lightning 9900™ performance and availability and configuration management.

Hitachi Hi-Star™ Architecture

At the heart of the Lightning 9900™ revolutionary design is the Hitachi Hi-Star™ Architecture, which provides multiple, redundant, non-blocking paths between the

storage ports, multiple cache nodes, and multiple disk Array Control Processors (ACPs). Each path runs at 100MB/s, permitting the Lightning 9900™ to scale up to a total bandwidth of 6.4 GB/s. This is six-to-eight times the internal bandwidth of shared bus based storage systems. For redundancy and performance, Hi-Star™ consists of four cache switches that are cross-connected to four cache modules and two control memory banks. These control memory banks contain addressing and control information and are also cross-connected to front-end storage ports and back-end disk array ports.

Hitachi InstantSplit

Hitachi InstantSplit is an enhancement to ShadowImage™ that provides multiple copies of data against which new applications can be tested or data mining queries can be run even before (newly created) ShadowImage™ pairs have been fully synchronized. With InstantSplit™ and the Lightning 9900™ bottleneck-free performance, applications can be deployed faster, resulting in the faster delivery of products to market and better-targeted sales campaigns. See also ShadowImage™.

Hitachi LUN Manager

Hitachi LUN Manager is a software utility in the Resource Manager 9000™ suite that allows for complete systems management of LUNs. See also LUN and Resource Manager™ 9000.

Hitachi Multiplatform Backup/Restore

Hitachi Multiplatform Backup/Restore is a channel-based backup/restore of Open Systems volumes using standard mainframe utilities, which leverage current investment in hardware, software, skills, and procedures. Hitachi Multiplatform Backup/Restore provides high-performance, high bandwidth capabilities. One copy of software resides on each controller.

Hitachi RapidXchange™

Hitachi RapidXchange™ provides for the file conversion and exchange of data between S/390® and open systems hosts. Hitachi Multiplatform Data Exchange provides file access library software for open systems hosts and runs on the following host servers: HP/UX®, AIX, Sun Solaris®, Compaq Tru64™, UNIX®, Sequent DYNIX/ptx®, SGI IRIX®, NCR UNIX® SVR4, Windows NT®, and Windows 2000®.

Hitachi Multiplatform Resource Sharing

Hitachi Multiplatform Resource Sharing is built right into the Lightning 9900™. With Hitachi Multiplatform Resource Sharing, Lightning 9900™ users can share resources between UNIX®- or Windows NT®/2000-based servers and S/390® mainframe platforms. Sharing resources across heterogeneous platforms lowers total cost of ownership, provides a centralized point for data management, and simplifies the management of heterogeneous systems..

Hitachi Rapid Recovery

The Hitachi Rapid Recovery solution for S/390®-attached storage combines Hitachi NanoCopy™, Hitachi TrueCopy-S/390 Asynchronous™, and Hitachi ShadowImage™. This facility enables data to be copied from one set of S/390®-attached systems to another with complete transaction integrity and without disrupting user applications in any way. Open-systems platform data can be consolidated on the Lightning 9900™ with NanoCopy™ capability to recover their critical data very quickly following an outage.

Hitachi Resource Manager™ 9000

The Hitachi 7000/9000 Resource Manager™ is a comprehensive suite of management software that brings together Hitachi Graph-Track™, Hitachi Virtual Logical Volume Image (VLVI) Manager™, Hitachi FlashAccess Manager™, and Hitachi LUN Manager™ into one complete package.

Hitachi SANTinel™

The Hitachi SANTinel™ software controls host access to Hitachi Freedom Storage 7000E™ or Lightning 9900™ LUNs in open systems or SAN environments.

Hitachi ShadowImage™

The Hitachi ShadowImage™ is a firmware-based software copy utility that uses command-line-interfaces to create up to ten copies of a volume within one Lightning 9900™ system, or up to 20 across multiple Lightning 9900™ systems. Graphic or command-line interfaces control data replication and fast resynchronization of logical volumes. ShadowImage is available for open systems or S/390 environments. ShadowImage also works in concert with Hitachi TrueCopy-S/390 Asynchronous™ to provide additional copies in another system.

Hitachi TrueCopy

Hitachi TrueCopy provides synchronous or asynchronous remote copy capability for open system and S/390 computers. This allows remote copies over virtually unlimited distances. Operating systems that are supported include MVS, HP/UX, AIX, Sun Solaris™, Digital UNIX®, Sequent DYNIX/ptx®, SGI IRIX™, NCR®, UNIX SVR4®, Windows NT/2000®.

Hot site

An alternate facility that has the equipment and resources to recover the business functions affected by the occurrence of a disaster or business interruption. Hot sites may vary in type of facilities offered (such as data processing, communication, or any other critical business functions needing duplication). Location and size of the hot site will be proportional to the equipment and resources needed. Similar terms include backup site; recovery site; recovery center; and alternate processing site. See also Cold site, Warm site, Disaster Recovery, Business Interruption, and Business Continuity Planning.

HSN

Hierarchical Star Network. The technical term used to describe the combined internal networks (the C-HSN and the CM-HSN) of the Lightning 9900™ Series.

Hub

A common connection point for devices in a Fibre Channel network. A hub contains multiple ports. When a Fibre Channel packet of data arrives at one port, it is copied to the other ports so that all storage devices on a SAN can see all packets.

HXRC (Hitachi Extended Remote Copy)

This is Hitachi Data Systems® IBM XRC® compatible host-based software asynchronous remote copy. Using System Data Mover software it guarantees data integrity for dependent write applications.

HYPERtape

An enterprise backup/restore solution that leverages current customer investments. HYPERtape is a three-tier distributed system architecture with central administration and control that supports consolidated and distributed environments. HYPERtape can be used to back up data from any supported host to any system that supports the ftp protocol, including backup to disk for HSM integration or backup to local- or network-attached tape. Over 30 host platforms are covered and 70 library modules are supported. All popular RDBMS programs are supported including Oracle®, SAP® R/3, Informix®, Sybase®, DB2®, Adabas-D RDB, and Microsoft SQL Server®, Exchange®, and Windows NT/2000® registry.

IP

The IP (Internet Protocol) is the underlying protocol for routing packets on the Internet and other TCP/IP-based networks. IP is an internetwork protocol that provides a communication standard that works across different types of linked networks, for example Ethernet, FDDI, or ATM.

Java

Developed by Sun Microsystems®, Java is now a standard software language for developing plug-in applications.

Journaling

A journaling file system keeps track of all changes to files as transactions occur in real time. In the event of unexpected system problems, the file system can be restored to a consistent state by updating a prior copy of the file system for the changes made from the point in time that the copy was made.

LAN

Local area networks or LANs are networks of computers that are geographically close together; this usually means on the same campus. Most LANs are confined to a single building or group of buildings. However, one LAN can be connected to other LANs over any distance via telephone lines, high-speed fibre optic backbones, and radio waves. A system of LANs connected in this way is called a wide-area network (WAN).

Lightning 9900™ Series

The Lightning 9900™ Series was announced in June 2000 (Lightning 9960™) and November (Lightning 9910™). It represents a major advance in enterprise-class storage systems with its Hierarchical Star Network switched internal architecture, which provides for many times more simultaneous transfers to and from the host compared to shared bus architectures.

Logical Unit

The SCSI term for a logical disk drive.

Logical Unit Number

See LUN.

Logical Volume

The storage medium associated with a logical disk drive. A logical volume typically resides on one or more storage devices. A host system sees a logical volume as a physical volume, although it does not correlate directly with a physical disk drive.

LUN

Logical Unit or Logical Unit Numbers. A SCSI term for the field in an identifying message that is used to select a logical unit on a given target.

LPAR

Logical partition is an IBM ESA/390® term for a set of functions that create the programming environment that is defined by the ESA/390® architecture. ESA/390® architecture uses this term when more than one LPAR is established on an ESA/390 server. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is tied to one or more physical processors in a tightly coupled multiprocessor system. Also the LPAR does not depend on an operating system to create the virtual machine environment.

LUSE

LUN Size Expansion feature. This Lightning 9900™ feature allows standard-size LUNs to be combined to create larger LUNs.

LRU

Least Recently Used. A policy for a caching algorithm that chooses to remove the data from cache which has the longest elapsed time since its last access. Least Recently Used algorithms are used in all major caching systems. The Lightning 9900™ Series' LRU scheme keeps a table (in separate non-volatile memory) that chronicles the frequency of use of data in cache memory.

MAN

Metropolitan Area Networks. Networks within a metropolitan area, which might for example, be used for a city government.

MBCP

Master Business Continuity Professional. The highest level of professional certification of the Disaster Recovery Institute (DRI). The master level targets an individual with a minimum of five years of experience as a business continuity/disaster recovery planner. See also DRI (Disaster Recovery Institute), CBCP (Certified Business Continuity Professional), and ABCP (Associate Business Continuity Professional).

MIB

Management Information Base is a set of standards for detailed system information that is reported to a control console for SNMP compliance. Its intent is to provide common parameters for heterogeneous computer systems.

MIPS

Millions of Instructions Per Second (or MIPS) is a rough measure of processor performance within the same class of processor.

Meta-data

Data about data as used by the CM-HSM and the Control Memory Subsystem of the Lightning 9900™ Series.

Mirrored pair

Two disk units or logical units that contain the same data. The operating system software refers to them as one entity and “reads from either” and “writes to both” when RAID-1 is enabled.

Mirroring

A term to describe the process of writing data to two disk volumes, usually to ensure high availability in case one of the disks fails. Mirroring can be hardware or software based.

MP

Micro-Processor.

MPA

Micro-Processor Adapter.

MPLF

The Lightning 9900™ supports the Multiple Path Locking Facility (MPLF) for the IBM® highest performance transaction processing operating system-TPF. In either native TPF mode or under VM, MPLF provides extremely high performance record-level locking so that multiple hosts can read and write to the same file without interfering with each other. See also TPF.

MTBF

Mean Time Between Failure. A commonly used measure of system reliability, usually expressed in hours. Modern disk drives typically have an MTTR of 1 million hours or more.

MTTR

Mean Time To Repair. Includes the time taken to diagnose the failure, replace or repair faulty component(s) and restart the system so it is available to users. See MTBF.

NanoCopy

NanoCopy™ is a feature of the Hitachi Lightning 9900™ product line that enables time-consistent snapshots to be taken without stopping applications to flush in-flight data to disk. Since there is no system impact in taking a NanoCopy™ snapshot, snapshots can be made more frequently for faster recovery in the event of a failure. See also Hitachi TrueCopy™, ShadowImage™.

NDMP

Network Data Management Protocol (NDMP) is a standard protocol for network-based backup of network-attached storage. NDMP hides the unique hardware interfaces from third-party backup software that allows this software to execute on any NDMP compliant system on the network.

Node

See Fibre Channel.

NVM

Non-Volatile Memory is a term used to refer to battery backed up DRAM so that data will not be lost in the event of power failure.

NVRAM

Non-Volatile Random Access Memory such as static RAM will not lose data in the event that power is lost to the memory chips.

OLAP

On-line Analytic Processing. A type of Decision Support System in which real-time analytical software routines are applied to data.

OLTP

On-line Transaction Processing.

Off-site storage facility

A secure location, remote from the primary location, at which backup hardware, software, data files, documents, equipment, or supplies are stored.

On-line system

An interactive computer system supporting users over a network of computer terminals.

Open system

A system whose characteristics comply with standards made available throughout the industry, and therefore can be connected to other systems that comply with the same standards.

Operating System

The operating system is the most important software program that runs on a computer. The operating system (OS) performs basic tasks such as recognizing input from a keyboard, sending output to the display screen, keeping track of files and directories on the disk and controlling peripheral devices such as disk drive and printers or a mouse. The OS acts as a traffic cop and schedules the various programs that the computer executes. The OS is also responsible for security, ensuring that unauthorized users do not access the system. Operating systems can be classified as follows:

- 1) Multi-user – allows two or more users to run programs at the same time.
- 2) Multi-processing – supports running a program on more than one CPU.
- 3) Multi-tasking – allows more than one program to run concurrently.
- 4) Multi-threading – allows different parts of a single program to run concurrently.
- 5) Real Time – Usually a stripped down OS that responds to input instantly.

Out of band

A communication that does not use the same bandwidth that carries data in a system. For example, the control information in the Lightning 9900™ Series systems do not use the same path as data and is therefore referred to as “out of band.”

Parity

A data-checking scheme used in a computer system to ensure the integrity of the data. The RAID implementation uses parity to recreate data if a disk drive fails.

PiT

A Point-in-Time (PiT) copy is a copy of data that is taken at a specific point in time. PiT copies are used in many ways including backups and checkpoints.

POD

Performance On Demand.

Port/Port ID

See Fibre Channel.

RAID

Redundant Array of Independent Disks. RAID is used to increase the reliability of disk arrays by providing redundancy either through complete duplication of the data (RAID-1, i.e., mirroring) or through construction of parity data for each data stripe in the array (RAID-3, -4, -5). RAID-5, which distributes parity information across all disks in an array, is among the most popular means of providing parity RAID since it avoids the bottlenecks of a single parity disk. The Lightning 9900™ algorithms enable performance from RAID-5 that is competitive with some vendor's RAID-1. Some vendors do not offer RAID-5.

RAID Controllers

RAID controllers provide a highly optimized scheme for securely managing RAID configurations on storage systems. Hitachi RAID controllers allow RAID arrays to be expanded online, and support conversion of an array from one RAID level to another.

Recovery Time

The period from the disaster declaration to the recovery of the critical functions.

Remote Copy

Remote Copy refers generically to software or hardware utilities that provide the capability to copy data from one online volume to remote volumes without disruption. Synchronous techniques are used for short distances (typically less than 25 miles), and asynchronous techniques over LAN/WAN/MAN are used at any distance.

Remote Copy Links

This term refers to the links used between storage systems for the movement of data. Today these links are either direct connect ESCON, Fibre Channel or network links (T3, ATM etc). For direct connect ESCON there is a limit of 43km (25 miles). For direct connect Fibre Channel the limit is 10km. However, newer technologies, such as the Nortel OPTERA™ product are allowing direct fibre connect over longer distances.

Risk Management

The discipline that ensures that an organization does not assume an unacceptable level of risk.

SAN

Storage Area Networks (SANs) connect storage systems to servers through Fibre Channel or Ethernet switches. Hitachi's implementation of SAN is known as Freedom Data Networks (FDN). Major benefits of SANs include outboard backup, sharing of resources, pooling, and reduced cost of storage management. Storage Area Networks (SAN) are high-speed subnetworks of shared storage devices. SAN architecture works in a way that makes all storage devices available to all servers on a LAN or WAN. Because stored data does not reside directly on any of a network's servers, server power is utilized for business applications, and network capacity is released to the end user. See also FDN, Fibre Channel.

SCSI

Small Computer System Interface. An intelligent bus-level interface that defines a standard I/O bus and a set of high-level I/O commands. There are currently many flavors of SCSI defined by different bus widths and clock speeds. The seven major variations of SCSI are SCSI 1, SCSI 2 (Fast/Narrow), SCSI 2 (Fast/Wide), Ultra SCSI (Fast/Narrow), Ultra SCSI (Fast/Wide) – also called SCSI 3, Ultra 2 SCSI (Narrow), Ultra 2 SCSI Wide. See also Fibre Channel.

Snapshot

A term that refers to a copy of a file system at a certain point in time. Snapshots are used for backup and recovery.

SNMP

Simple Network Management Protocol. SNMP is a protocol used for communication between simple, server-resident SNMP agents that respond to network administration requests from simple-to-sophisticated SNMP manager tools running on remote workstations.

Solaris

Sun's UNIX® operating system based on System V, release 4.

SPARC

Scalable Processor Architecture. SPARC International's specification for Reduced-Instruction-Set-Computer (RISC) CPUs.

Stripe

In RAID terminology, a stripe is when data is read or written in parallel to or from multiple disks instead of reading or writing all data to one disk. Striping provides much higher performance through its parallel design.

SVP

Service Processor of the Lightning 9900™ Series.

SWAN

Storage Wide Area Networks (SWANs) are interconnected SANs over long distances. They are made possible by Fibre Channel and ESCON extenders.

Switch

In networking terminology, a switch is a computing device that filters and forwards packets between Local Area Network (LAN) segments. Switches operate at the data link layer (layer 2) of the OSI Reference Model and therefore support any packet protocol. A special type of switch called an L4 switch operates at the fourth layer (Transport Layer) of the OSI Reference Model and is responsible for the integrity of data transmissions between LAN segments. LANs that use switches to join segments are called switched LANs or, in the case of Ethernet networks, switched Ethernet LANs. See also L-4 Switch.

Synchronous

Synchronous communications occurs when the transmission of data between two devices is synchronized with a clocking scheme or other technique. The sender and receiver need to synchronize with one another before data is sent. In synchronous communication the bit stream and the clock pulse are synchronized by a special bit transition pattern in the digital signal creating an exactly timed stream of bits from the sending device to the receiving device. An example of such a mechanism is bipolar encoding. Synchronous communication is either character or bit oriented. Character oriented synchronous transmissions are used to send blocks of characters such as found in ASCII (American Standard Code for Information Interchange) files. Bit oriented synchronous communication is used primarily for the transmission of binary data. See also Asynchronous, ASCII, Binary and EBCDIC, HRC, HORC.

Tach

Abbreviation for Tachyon Fibre Channel interface.

TB

A Terabyte (TB) equals 1024 Gigabytes. Many storage vendors today define a terabyte as 1000 GB, causing confusion in the industry.

TCP

Transmission Control Protocol or TCP is a transport layer component of the Internet's TCP/IP protocol suite. It sits above IP in the protocol stack and provides reliable data delivery services over connection-oriented links. TCP uses IP to deliver information across a network and makes up for the deficiency of IP providing a guarantee of reliable delivery services that IP does not. TCP messages and data are encapsulated into IP datagrams and IP delivers them across the network.

VERITAS®

A Mountain View, California software company that develops and supports volume and file management software products for a variety of Unix® and Windows® platforms.

Virtual Logical Volume Image Manager

Virtual Logical Volume Image Manager (VLVI) is a software utility in the Resource Manager 7000/9000™ that allows for configuration of RAID configurations, create, delete, verify, rebuild, tune, and abort operations. See also RAID and the Freedom 9000 Resource Manager 7000/9000™.

Volume

An ESA/390® term for the information recorded on a single disk unit or recording medium. Indirectly, a volume can refer to the unit of recording medium itself. On a non-removable medium storage device such as a disk drive, the terms may also refer, indirectly, to the storage device that is associated with the volume. When a user stores multiple volumes on a single storage medium transparent to the program, the volumes are referred to as logical volumes.

WAN

Wide Area Networks or WANs are networks of computers that are geographically dispersed and connected by radio waves, telephone lines, satellites, or high-speed fibre optic backbones.

Warm site

An alternate-processing site that is only partially equipped (as compared to Hot site which is fully equipped). See also Hot site and Cold site.

Workload

I/O workload refers to the pattern of I/Os presented to the Lightning 9900™ Series system or to a disk drive.

XRC

Extended Remote Copy. IBM's implementation of a software asynchronous Remote Copy technique that preserves data integrity. See also HXRC and Hitachi TrueCopy.

Hitachi Data Systems Worldwide Offices

www.hds.com
www.storage@hds.com

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com

Asia Headquarters

39-09 Tower One
Lippo Centre
89 Queensway
Hong Kong
2525-2385
infoasia@hds.com

Australia/New Zealand Headquarters

11-17 Khartoum Road
North Ryde NSW 2113
Australia
02-9325-3300
info@hds.com.au

Canada Headquarters

380 Saint-Antoine Street West
Suite 7000
Montreal, Quebec H2Y 3X7
Canada
(514) 982-0707
info@hdscanada.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
01753-61-8000
info.eu@hds.com

Latin America Headquarters

750 Central Expressway, MS 3268
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
infolatam@hds.com

U.S. Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
ussalesinfo@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

Hi-Star, Freedom Storage, Lightning 9900, FlashAccess, NanoCopy, Graph-Track, CruiseControl, SANtinel, ShadowImage, Hi-Track, Freedom Data Networks, Dynamic Link Manager, RapidXchange, and iLAB are trademarks of Hitachi Data Systems Corporation.

BETA and HARBOR are trademarks of BETA Systems Software AG.

BROCADE is a trademark of BROCADE Communications Systems, Inc.

Compaq is a registered trademark and Tru64 and TruCluster are trademarks of Compaq Computer Corporation.

Computer Associates and Unicenter TNG are registered trademarks of Computer Associates International, Inc.

EMC and Symmetrix are registered trademarks and TimeFinder is a trademark of EMC Corporation.

Emulex is a trademark of Emulex Corporation.

HP-UX, Hewlett-Packard, and OpenView are registered trademarks of Hewlett-Packard Company.

IBM, ESCON, S/390, OS/390, Sequent, DYNIX/ptx, RS/6000, NetView, MVS, Enterprise System Architecture/390, ESA/390, and ES/9000 are registered trademarks and HACMP/6000, GDPS, and FICON are trademarks of International Business Machines Corporation. AIX is being used under license.

Informix is a registered trademark of Informix Corporation.

Intel is a registered trademark of Intel Corporation.

Linux is a registered trademark of Linus Torvalds.

McDATA is a registered trademark of McDATA Corporation.

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

NCR is a registered trademark of NCR Corporation.

NetWare is a trademark of Novell, Inc.

OPTera is a trademark of Nortel Networks Limited.

Oracle is a registered trademark of Oracle Corporation.

Q-logic is a trademark of Q-logic Corporation.

Red Hat is a registered trademark of Red Hat, Inc.

Seagate is a registered trademark of Seagate Technology, Inc.

SGI is a trademark and IRIX is a registered trademark of Silicon Graphics, Inc.

StorageTek is a trademark of Storage Technology Corporation.

Sybase is a registered trademark of Sybase, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

VERITAS is a registered trademark and ClusterServer is a trademark of VERITAS Software Corporation.

Xerox is a registered trademark of Xerox Corporation.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to <http://www.hds.com/products/systems/9900/licenses.html> or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

©2001, Hitachi Data Systems Corporation
All Rights Reserved.

DISK-387-00 October 2001

Hitachi Data Systems

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050
U.S.A.
Phone: 1+408 970 1000
info@hds.com
www.hds.com

European Offices:

European Headquarters

Hitachi Data Systems
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
Phone: +44 1753 618000
Fax: +44 1753 618550
info.eu@hds.com
www.eu.hds.com

United Kingdom

Hitachi Data Systems
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
Phone: +44 1753 618989
Fax: +44 1753 618543
info@hds.co.uk
www.hds.co.uk

Hitachi Data Systems
Hamilton House
Church Street
Altrincham WA14 4DR
Phone: +44 161 928 4243
Fax: +44 161 928 8461
info@hds.co.uk

Hitachi Data Systems
100 Park Avenue
Aztec West
Almondsbury
Bristol BS32 4TS
Phone: +44 1454 625511
Fax: +44 1454 625502
info@hds.co.uk

Hitachi Data Systems
6220 Bishops Court
Solihull Parkway
Birmingham Business Park
Birmingham B37 7YB
Phone: +44 121 717 4000
Fax: +44 121 717 4001
info@hds.co.uk

Hitachi Data Systems
11 Rutland Square
Edinburgh EH1 2AS
Phone: +44 131 228 4343
Fax: +44 131 228 5573
info@hds.co.uk

Hitachi Data Systems
Tower 42
The International Finance Centre
Level 7, 25 Old Broad Street
London
EC2N 1HN
Phone: +44 20 7877 2052
Fax: +44 20 7877 0688
info@hds.co.uk

Austria

Hitachi Data Systems GmbH
Praterstraße 62-64
1020 Wien
Phone: +43 1 245 82 0
Fax: +43 1 245 82 250
info.at@hds.com

Belgium

Hitachi Data Systems Belgium
S.A./N.V.
Avenue Louise, 283 Box 18
6th Floor
1050 Brussels
Phone: +32 2 643 4550
Fax: +32 2 643 4551

Denmark

Hitachi Data Systems
Egebækgård
Egebækvej 98
2850 Nærum
Phone: +45 45 58 10 00
Fax: +45 45 58 10 01
info.dk@hds.com
www.hds.dk

Finland

Hitachi Data Systems
Tietäjantie 2
02130 Espoo
Finland
Phone: +358 9 251 05 132
Fax: +358 9 251 05 100

France

Hitachi Data Systems
81 rue Marcel Dassault
92514 Boulogne Billancourt Cedex
Phone: +33 1 46 10 14 00
Fax: +33 1 47 61 94 24

Germany

Hitachi Data Systems GmbH
Mozartstraße 2
D-85622 Feldkirchen b. München
Phone: +49 89 9 05 06 510
Fax: +49 89 9 05 06 199
info.de@hds.com
www.hds.de

Hitachi Data Systems GmbH
Industriestraße 5
70565 Stuttgart
Phone: +49 711 783 37 139
Fax: +49 711 783 37 316

Hitachi Data Systems GmbH
Buchschrager Allee 8
D-63303 Dreieich-Buchschlag
Phone: +49 6103 7334 0
Fax: +49 6103 7334 101

Hitachi Data Systems GmbH
Potsdamer Platz 1
D-10785 Berlin
Phone: +49 30 25944 5710
Fax: +49 30 25944 200

Hitachi Data Systems GmbH
Valentinskamp 24
D-20354 Hamburg
Phone: +49 40 31112 900
Fax: +49 40 31112 200

Hitachi Data Systems GmbH
Prinzenallee 7
D-40549 Düsseldorf
Phone: +49 211 52 391 427
Fax: +49 211 52 391 200

Israel

Hitachi Data Systems
Corex Building
27 Maskit St.
Herzlia 46733
Phone: +972 9 970 45 00
Fax: +972 9 970 45 11
info.israel@hds.com

Italy

Hitachi Data Systems
Via T. Gullii, 39
20147 Milan
Phone: +39 02 40 30 92 1
Fax: +39 02 40 30 92 80
info.italy@hds.com
www.it.hds.com

Netherlands

Hitachi Data Systems
Achterweg 29
4181 AD Waardenburg
Phone: +31 418 657 650
Fax: +31 418 657 555
info.nl@hds.com

Norway

Hitachi Data Systems AS
Vollsveien 19
N-1366 Lysaker
Phone: +47 67 51 84 40
Fax: +47 67 51 84 59

Spain

Hitachi Data Systems S.A.
Avda. Diagonal, 640
08017 Barcelona
Phone: +34 933 634 100
Fax: +34 933 634 111

Hitachi Data Systems S.A.
C/Capitán Haya, 38-8ª Planta
28020-Madrid
Phone: +34 91 449 1100
Fax: +34 91 449 1111

Sweden

Hitachi Data Systems AB
Frosundaviks Alle 15
Box 501
SE-169 29 Solna
Phone: +46 8 624 64 00
Fax: +46 8 624 64 65

Hitachi Data Systems AB
Stora Badhusgatan 18-20
SE-411 21 Gothenburg
Phone: +46 31 701 18 00
Fax: +46 31 701 18 19

Switzerland

Hitachi Data Systems
Kriesbachstrasse 3
8600 Dübendorf / ZH
Schweiz
Phone: +41 1 802 64 64
Fax: +41 1 820 39 40
info.ch@hds.com

Hitachi Data Systems
109, Chemin du Pont-du-Centenaire
1228 Plan-les-Ouates
Phone: +41 22 884 86 00
Fax: +41 22 884 86 01
info.ch@hds.com

Hitachi Data Systems
Bernstrasse 131
3052 Zollikofen
Phone: +41 31 910 42 42
Fax: +41 31 911 64 47
info.ch@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

Hi-Star, Freedom Storage, Lightning 9900, FlashAccess, NanoCopy, Graph-Track, CruiseControl, SANtinel, ShadowImage, Hi-Track, Freedom Data Networks, Dynamic Link Manager, RapidXchange, and iLAB are trademarks of Hitachi Data Systems Corporation.

BETA and HARBOR are trademarks of BETA Systems Software AG.

BROCADE is a trademark of BROCADE Communications Systems, Inc.

Compaq is a registered trademark and Tru64 and TruCluster are trademarks of Compaq Computer Corporation.

Computer Associates and Unicenter TNG are registered trademarks of Computer Associates International, Inc.

EMC and Symmetrix are registered trademarks and TimeFinder is a trademark of EMC Corporation.

Emulex is a trademark of Emulex Corporation.

HP-UX, Hewlett-Packard, and OpenView are registered trademarks of Hewlett-Packard Company.

IBM, ESCON, S/390, OS/390, Sequent, DYNIX/ptx, RS/6000, NetView, MVS, Enterprise System Architecture/390, ESA/390, and ES/9000 are registered trademarks and HACMP/6000, GDPS, and FICON are trademarks of International Business Machines Corporation. AIX is being used under license.

Informix is a registered trademark of Informix Corporation.

Intel is a registered trademark of Intel Corporation.

Linux is a registered trademark of Linus Torvalds.

McDATA is a registered trademark of McDATA Corporation.

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

NCR is a registered trademark of NCR Corporation.

NetWare is a trademark of Novell, Inc.

OPTEra is a trademark of Nortel Networks Limited.

Oracle is a registered trademark of Oracle Corporation.

Q-logic is a trademark of Q-logic Corporation.

Red Hat is a registered trademark of Red Hat, Inc.

Seagate is a registered trademark of Seagate Technology, Inc.

SGL is a trademark and IRIX is a registered trademark of Silicon Graphics, Inc.

StorageTek is a trademark of Storage Technology Corporation.

Sybase is a registered trademark of Sybase, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

VERITAS is a registered trademark and ClusterServer is a trademark of VERITAS Software Corporation.

Xerox is a registered trademark of Xerox Corporation.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to <http://www.hds.com/products/systems/9900/licenses.html> or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.