


Hitachi Freedom Storage™ Lightning 9900™ V Series

Architecture Guide



The direction you can trust

HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Lightning 9900™ V Series

Architecture Guide



Introduction

Hitachi Freedom Storage Lightning
9900 V Series Packaging

Front-end Design of the Hitachi
Freedom Storage Lightning
9900 V Series

Hi-Star™ Architecture – An
Internal Switched
Hierarchical Star Network

High-performance Back-end
Design

Capacities and RAID
Architecture

Hitachi Freedom Storage
Lightning 9900 V Series
Availability is Highest
in the Industry


Unmatched Performance
and Scalability

Hitachi Freedom Storage
Software Solutions
Are Best-of-Breed

The HiCommand™ Management
Framework and Hitachi Resource
Manager™ System Management
Software

Optimizing Performance, Availability,
and Cost with the HiReturn™
Investment Analysis Tool Kit

Global Solutions Services



The direction you can trust

HITACHI
DATA SYSTEMS

Table of Contents

Table of Contents	iii
Table of Figures	vii
Table of Tables	viii
Chapter 1 – Introduction	1
Redefining the Industry Standard, Once Again	1
Lightning 9900 V Series Product Offerings	2
An Overview of the Lightning 9900 V Series Architecture	3
Limitation of Competitive Shared-bus Architectures	4
The Lightning 9900 V Series Systems Have an Advanced Availability Profile Compared to Competitive Products.....	4
Other Information Sources Available from Hitachi Data Systems	5
Reduced TCO	5
Second-generation “Virtual” Switched Storage.....	6
Chapter 2 – Hitachi Freedom Storage Lightning 9900 V Series Packaging	9
The Hitachi Freedom Storage Lightning 9980V System Packaging	9
Lightning 9980V System Control Frame	9
Lightning 9980V System Array Frames	10
Hitachi Freedom Storage Lightning 9970V System Packaging	11
The System Monitoring Network of the Lightning 9900 V Series.....	12
Chapter 3 – Front-end Design of the Hitachi Freedom Storage Lightning 9900 V Series	13
CHiPs Support Fibre Channel Host and ESCON or FICON Host Connectivity and Scalability.....	13
About the Fibre Channel Interface Boards	14
Bandwidth of Supported Host Interfaces	16
Number of Host Interfaces Supported	16
Types of Hosts Supported with Direct Fibre Channel Host Attachment by the Lightning 9900 V Series Systems	16
Chapter 4 – Hi-Star Architecture – An Internal Switched Hierarchical Star Network	17
The Internal Hi-Star Architecture	17
The Cache Hierarchical Star Network	17
The Cache Switch Provides Non-blocking Switched Access to Cache	18
The Lightning 9980V System Provides 64GB of Fully Addressable Cache	18

Advanced Cache Algorithms of the Lightning 9900 V Series.....	18
Hitachi FlashAccess Software Allows Data Sets to be Permanently Placed in Cache	18
Read-ahead for High-performance Sequential Reads	18
Control Memory Hierarchical Star Network	18
Chapter 5 – High-performance Back-end Design	21
High-speed All Fibre Channel Back-end Design	21
Maximum Configuration of ACPs and Disks	22
New Fibre Channel Disk Drives	22
The Importance of Command Tag Queuing (CTQ)	22
Chapter 6 – Capacities and RAID Architecture	23
Storage Capacity of the Hitachi Freedom Storage Lightning 9900 V Series	23
Advantages of the Lightning 9900 V Series RAID Hardware	23
RAID-1+	23
RAID-5 – Distributed Parity	23
RAID Rebuild Capability	23
High-density Disk Drive (HDD) Assemblies.....	24
Chapter 7 – Hitachi Freedom Storage Lightning 9900 V Series Availability Is Highest in the Industry	25
An Overview of High-availability (HA) Computing Software	25
Software Products Contribute to High-availability Computing	25
The Lightning 9900 V Series Has an Advanced Availability Profile Compared to Competitive Products	25
Open Systems High-availability Middleware Support	26
Alternate Pathing Middleware Switches the I/O Load in the Event of Path Failure	27
Hitachi Dynamic Link Manager Software Provides Path Failover and Load Balancing	27
Host Failover	27
Parallel Database Clustering	28
Standard Hot Swap and Redundant Power Supplies	28
Standard Redundant Fans for Cooling.....	29
Competitive Advantage in RAID Choices	29
Chapter 8 – Unmatched Performance and Scalability	31
Truly Scalable Performance	31
Internal Switching Overcomes the Limitations of Shared-bus Architectures	31
There Is NO Match for the Performance of the Lightning 9980V System	31

Chapter 9 – Hitachi Freedom Storage Software Solutions Are <i>Best of Breed</i>	33
Hitachi Freedom Storage Software Solutions Simplifies Information Management, Protects Data Assets, and Optimizes Storage Investments	33
Storage Area Management Suite	33
Business Continuity Suite	35
Backup and Recovery Suite	36
Performance Enhancement Suite	36
Data Movement Suite	37
Chapter 10 – The HiCommand Management Framework and Hitachi Resource Manager System Management Software	39
A Powerful and Open Systems Management Philosophy	39
HiCommand Management Framework Allows Systems Management of Hitachi Freedom Storage and Software through the Enterprise’s Vendor of Choice	39
HiCommand Management Framework Provides an Open Framework for Storage Management	39
Open, Standards-based Approach	40
CIM Benefits to Customers	40
Overcome the Challenges Created by De-coupling Storage Resources from Specific Server(s) in a SAN	42
View and Operate Networked Storage as a Holistic Resource	42
Avoid Being "Locked" into a Proprietary Management Environment	42
Overview of How the Systems Management Software Supports Business Objectives	43
Hitachi Resource Manager Provides Powerful Device-level Storage Management Capability	43
Out-of-band Management	43
SNMP Support Ensures Compliance with All Major Systems-management Platforms	44
Java Agent	44
Host Storage Domains	44
Hitachi LUN Manager and LUN Expansion Manager Software Simplifies Configuration Management and Reduces Staffing Costs	44
Hitachi SANTinel Software	45
Hitachi Virtual LUN and LUN Expansion Software	46
Hitachi FlashAccess Software Allows Data to be “Locked and Unlocked” in Cache On-the-fly	44
Remote Console – Storage Navigator	47
Hitachi Graph-Track Software Takes the Guesswork Out of Performance Management	47
Hitachi Virtual Logical Volume Image Manager Software for S/390 Environments Allows Very Large to Very Small Virtual-Volume Configuration	48
Cache Manager Host Agent	48

Chapter 11 – Optimizing Performance, Availability, and Cost with the HiReturn Investment Analysis Tool Kit	49
HiReturn Investment Analysis Tool Kit	49
The HiReturn Methodology Is Implemented in Three Phases	49
Phase 1 - Preparation	49
Phase 2 - On-site, One-day Assessment	49
Infrastructure Review.....	49
Assimilation.....	49
Findings.....	49
Phase 3 - Final Report	50
The On-site Assessment Produces the Needed Facts to Build a Business Case	50
Results and Reports Quantify Both IT Benefits and Business Benefits	50
Chapter 12 – Global Solutions Services	51
Hitachi Data Systems Is Consistently Ranked Number One in the Industry.....	51
Global Solutions Services Overview	51
Storage Management Solutions	51
The HiCommand Device Manager Enablement Service	51
Basic Implementation Service for InterSAN PATHLINE.....	51
Data Protection Solutions	52
Data Replication Planning and Design Service	52
Implementation Service for Hitachi TrueCopy Software.....	52
Integration Service for Hitachi TrueCopy Software	52
Implementation Service for Hitachi ShadowImage Software	52
Integration Service for Hitachi ShadowImage Software	52
Hitachi SplitSecond Solutions Rapid Recovery Service for Microsoft Exchange	53
Hitachi SplitSecond Solutions Rapid Recovery Service for Microsoft SQL Server	53
Storage Infrastructure Solutions.....	53
Configuration Management Service	53
Fiber Design and Implementation Service	53
Infrastructure Planning Service for FICON	54
Storage Consolidation Planning and Design Service.....	54
Storage Consolidation Implementation Service	54
Open Systems Data Migration Planning Service	54
Open Systems Data Migration Implementation Service	54
SAN Services	54
Glossary of Terms	55

Table of Figures

Figure 1	The Lightning 9980V system provides for exploding growth and data hyper-consolidation while the Lightning 9970V system provides for managed growth and open systems exploitation.	2
Figure 2	The Lightning 9900 V Series is based on an internal switched architecture.	3
Figure 3	A shared-bus architecture is limited to two simultaneous I/O operations.	4
Figure 4	Actual average deployed capacity of the Lightning 9900 V Series systems exceeds that of shared-bus systems by a factor of 2.5x.	6
Figure 5	Today's heterogeneous IT environment demands virtualization of storage.	6
Figure 6	Heterogeneous solutions today.	7
Figure 7	The Lightning 9900 V Series virtual Fibre Channel ports and Host Storage Domains.	8
Figure 8	A Lightning 9980V system Control Frame and its major components.	9
Figure 9	A Lightning 9980V system Array Frame with its major components.	10
Figure 10	Power supply modules for the Lightning 9900 V Series systems.	11
Figure 11	A Lightning 9970V system single-cabinet storage system.	12
Figure 12	Fibre Channel and ESCON or FICON adapter boards connect to open systems and mainframe computers.	13
Figure 13	A 2Gbit/sec Fibre Channel host adapter board and its major hardware components.	14
Figure 14	An ESCON host adapter board and its major hardware components.	15
Figure 15	The Lightning 9900 V Series systems support both mainframe and open system hosts.	16
Figure 16	The C-HSN of a Lightning 9970V system provides fully redundant switched access to cache from all channel adapters.	17
Figure 17	Separate redundant control memory handles the exchange of control information, such as status, location, and configuration of data, between processors and cache memory.	19
Figure 18	Functional diagram of an Array Group with FC-AL disks.	21
Figure 19	Diagram of four ACP pairs and 32 FC-AL back-end disks for the Lightning 9980V system.	21
Figure 20	The Lightning 9900 V Series systems are available with RAID-1+ and RAID-5 arrays.	24
Figure 21	The Lightning 9900 V Series offers a far more complete range of availability protection than competitive products.	26
Figure 22	Dynamic Link Manager software automatically provides path failover and load balancing for open systems.	28
Figure 23	Test results show that the advantages of the Lightning 9980V system maintain user response time while workloads dramatically increase.	32
Figure 24	Hitachi Data Systems offers a broad range of remote copy, data duplication, and data migration software solutions which are profiled in the Hitach Freedom Storage Software Solutions Guide.	35
Figure 25	HiCommand Device Manager software and Tuning Manager software allows management of virtually all Hitachi Freedom Storage hardware and software from the customer's platform of choice.	40
Figure 26	The HiCommand Management Framework is based on the CIM and SOAP industry standard.	42
Figure 27	Host Storage Domains reduce costs through fewer physical connections.	45
Figure 28	SANTinel software uses pre-defined zones to protect data from unauthorized access.	46
Figure 29	The GUI interface on Graph-Track software simplifies performance management.	48

Tables

Table 1	Comparison of Lightning 9900 V Series models.....	2
Table 2	Bandwidth of supported interfaces of the Lightning 9900 V Series.....	16
Table 3	System capacities and number of disk drives for various configurations.	23
Table 4	Lightning 9980V system bandwidth is ten times that of competition.....	31
Table 5	Summary of Hitachi Freedom Storage software products.....	34
Table 6	Summary of Hitachi Resource Manager solutions and the products they support.....	43

Introduction

1

Redefining the Industry Standard, Once Again

Hitachi Data Systems, in presenting the Hitachi Freedom Storage™ Lightning 9900™ V Series internal switched storage architecture with virtual Fibre Channel switches, again redefines the industry standard for “best-of-breed” storage systems. It is the most advanced design available in the world for open systems and mainframe data, and it is now two generations ahead of the competition. The patented Hi-Star™ architecture, combined with industry-unique “virtual” Fibre Channel ports, powers the most intelligent, highest performance, enterprise-class storage product in the industry today.

Hitachi Data Systems is committed to enabling its customers to build world-class, profit-oriented solutions on top of the Hitachi TrueNorth™ vision and strategic direction. With its state-of-the-art components, TrueNorth enables customers to optimize the value of information within an extremely agile, easy-to-manage and highly-resilient systems environment. This, in turn, leads to competitive advantage for the enterprise by fostering the close cooperation of employees, owners, and suppliers through integrated, bulletproof information systems. TrueNorth allows the enterprise to simplify, protect, and optimize its information infrastructure, thereby reducing total cost of ownership (TCO) and ensuring a more rapid return on investment (ROI). The TrueNorth vision and strategic direction, along with an associated product road map, anticipates customer needs. The TrueNorth vision views the future storage infrastructure as a synergistic combination of management software and powerful, intelligent storage systems that deliver set-and-forget management, complemented by flexible capacity, performance, and connectivity.

1. Hi-Star, a **modern switched storage architecture**, was introduced in 2000 and is the only architecture of its type available among enterprise-class storage systems. It was upgraded in May 2002 to almost twice the internal processing power and bandwidth.
2. The Lightning 9900 V Series **leapfrogs competitive products** in every category of comparison, including internal hardware architecture, availability, performance, and embedded 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 IBM® S/390® and zOS™ compatible mainframe computers, which provides unsurpassed functionality.
5. The Lightning 9900 V Series dramatically increases **front-end connectivity** by making available multiple “virtual” Fibre Channel ports associated with each of 32 physical Fibre Channel front-end ports. This enables truly large-scale data consolidation.
6. The Lightning 9900 V Series offers **exceptional performance**, with no decline in user response time in tests simulating 7200 OLTP users running in parallel with two decision-support applications performing full table scans at 6.6 million rows per second.
7. **Backup windows, operations efficiency, and disaster recovery are optimized** using leading Hitachi Data Systems copy, migration, and data duplication software solutions.
8. The many **advanced system management functions** of the Lightning 9900 V Series can be initiated, managed, and controlled by the powerful HiCommand™ Management Framework.
9. The Lightning 9900 V Series is fully compatible with all major **UNIX®, Microsoft® Windows NT®, Windows® 2000 and S/390 mainframe network and system management tools**.
10. Hitachi offers Hi-Track® service, with an optional **robust dial-in capability** that provides world-class experts to diagnose problems if they occur.

Two generations ahead of the competition.

Industry-unique “virtual” Fibre Channel ports.

Ten times the bandwidth of competitive products.

Lightning 9900 V Series Product Offerings

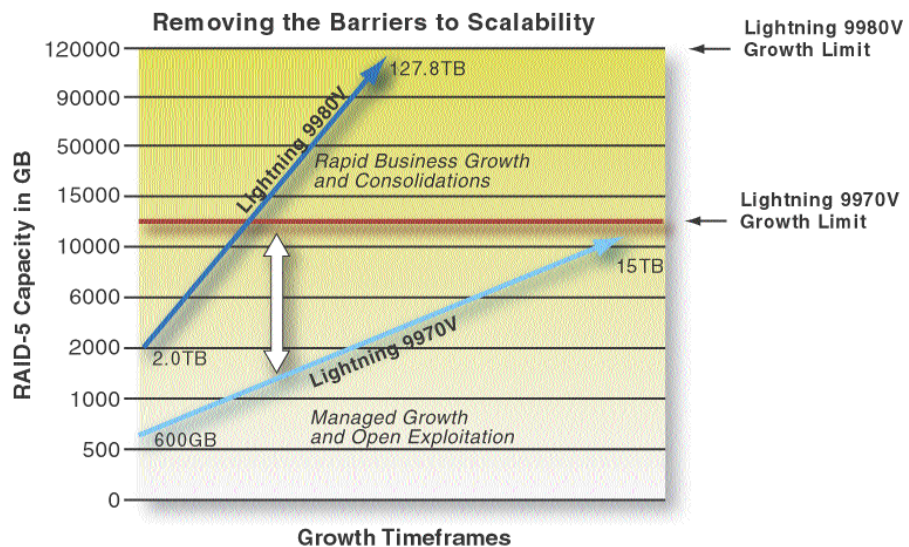
The Lightning 9900 V Series offers two models – the Lightning 9980V™ and the Lightning 9970V™. The Lightning 9980V system is recommended for users requiring large storage capacities of 149TB+ of raw storage capacity. With the same advanced architecture as the Lightning 9980V system, the Lightning 9970V system bridges the gap between standard midrange and premium-priced advanced-function storage, with capacities up to 18TB+ of storage. The Lightning 9970V system is recommended for users who require advanced function capabilities of premium storage, yet do not need the actual storage capacity of the Lightning 9980V system. Table 1 compares key features of the two new Lightning 9900 V Series product offerings.

Table 1 – Comparison of Lightning 9900 V Series models.

Product Characteristics	Lightning 9900 V Series	
	Lightning 9980V system	Lightning 9970V system
Number of Cabinets	2-5	1
Maximum Cache Capacity	64GB	32GB
Number of Fibre Channel Ports	Up to 64	Up to 48
Number of ESCON® Ports ¹	Up to 48	Up to 24
Number of FICON™ Ports	Up to 32	Up to 24
Number of Disks (HDDs)	Up to 1024	Up to 128
Number of Array Groups	Up to 254	Up to 31
Number of ACP pairs	1-4	1-2
Number of FC-AL Paths to Back-end Disks	Up to 32	Up to 16
Maximum TB of Raw Disk Capacity	149TB+	18TB+
Minimum usable RAID-5 Capacity	73GB	73GB
Maximum usable RAID-5 Capacity	127.8TB	15TB

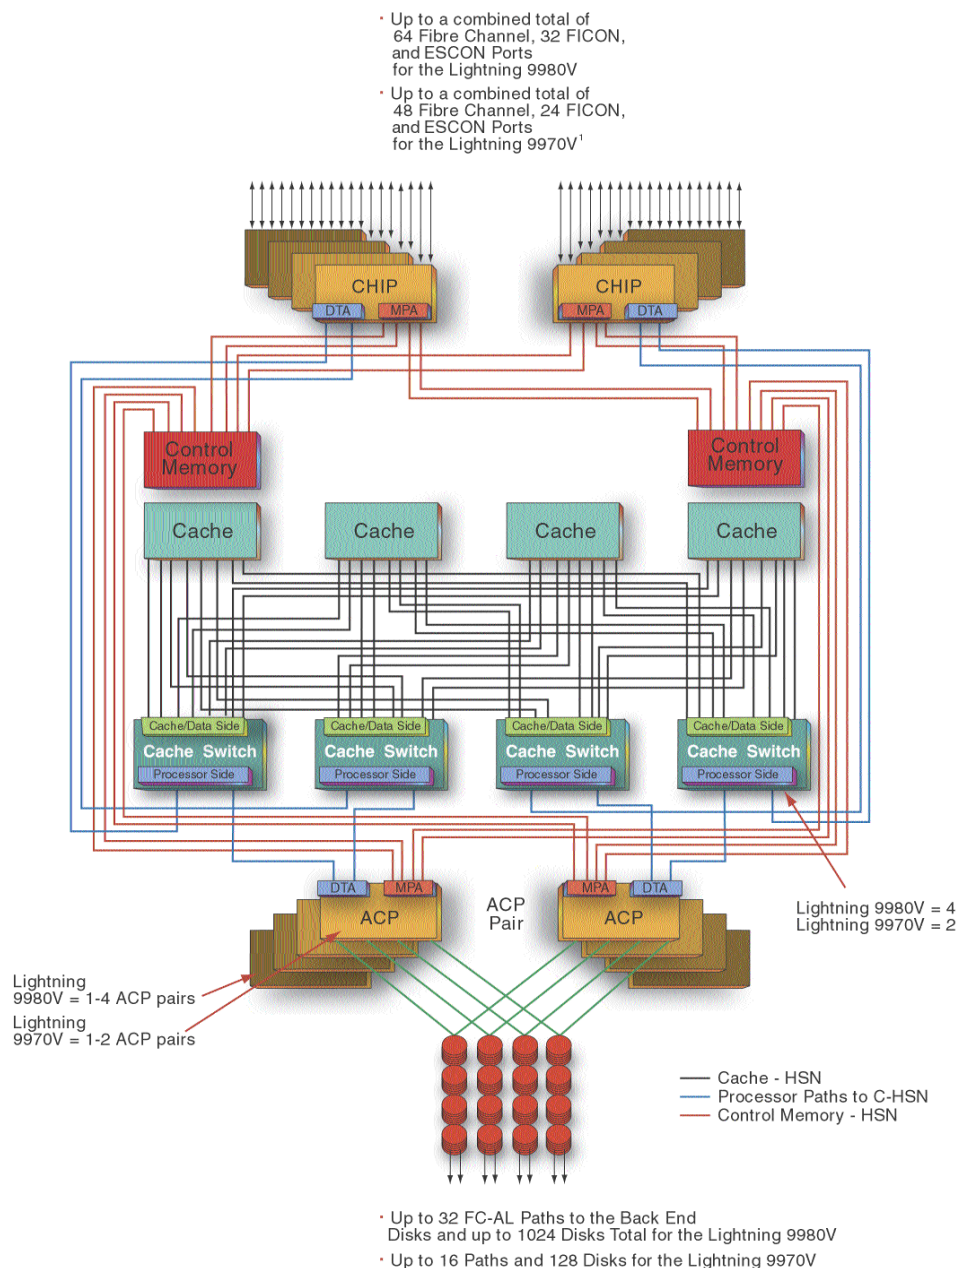
Selecting the Lightning 9970V system or Lightning 9980V system depends on capacity requirements, performance requirements, and expected growth of data. The Lightning 9900 V Series therefore offers product alternatives that cover a very broad range of scalability, as illustrated in Figure 1.

Figure 1 – The Lightning 9980V system provides for exploding growth and data hyper-consolidation while the Lightning 9970V system provides for managed growth and open systems exploitation.



An Overview of the Lightning 9900 V Series Architecture

The Lightning 9900 V Series intelligent storage systems are the most powerful enterprise systems in the industry today and two generations ahead of the competition. The new architecture used in the Lightning 9900 V Series has almost twice the internal processing power and bandwidth of our earlier enterprise systems. Figure 2 illustrates the architecture of the Lightning 9900 V Series.



At the heart of the Lightning 9900 V Series is enormous capacity for throughput, thanks to Hi-Star architecture. This network of interconnecting high-speed links to system components provides an unequalled internal aggregate bandwidth of 15.9GB/sec (gigabytes/sec) in the Lightning 9980V system. This bandwidth is divided into 10.6GB/sec for data transfer and 5.3GB/sec for separate internal control information transfer. This superhigh bandwidth is approximately 10 times that of competitive products.

The Lightning 9980V system has almost twice the processing power of our earlier enterprise systems.

Figure 2 – The Lightning 9900 V Series is based on an internal switched architecture.

The Lightning 9980V system has an internal bandwidth of 15.9GB/second.

The main components that make up the Lightning 9900 V Series systems are the Array Control Processor modules (ACPs) and Host Adapter modules, called Channel 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 15.9GB/sec internal bandwidth is available with a fully configured Lightning 9980V system 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 V 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.

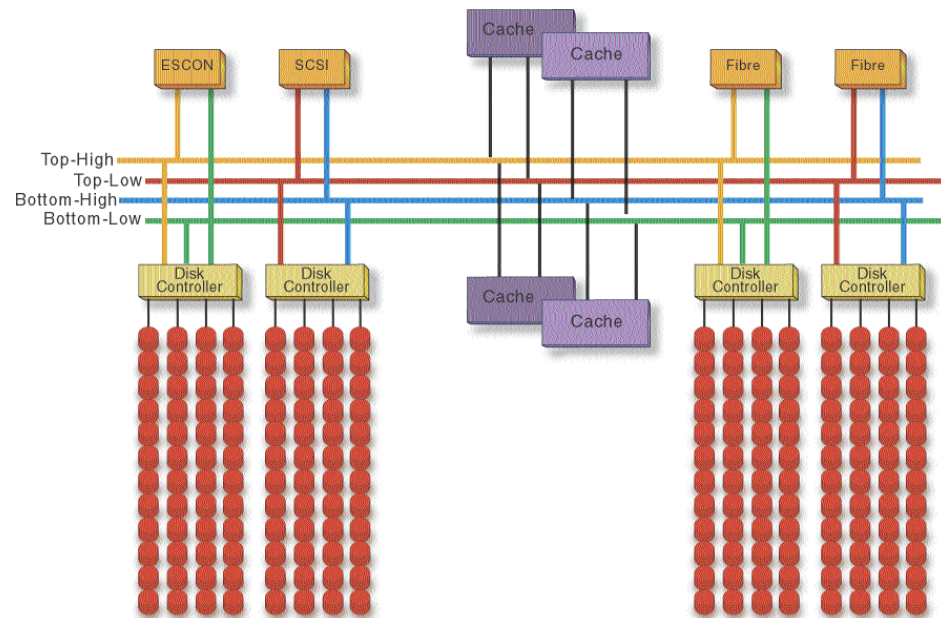
Although two ACP pairs are allowed for the Lightning 9970V system, more than one pair is rarely needed except for scientific applications requiring massive bandwidth of raw data from disk.

Bus contention is virtually eliminated with the Hi-Star architecture.

Limitation of Competitive Shared-bus Architectures

Unlike the shared-bus architecture of some competitive products (as shown in Figure 3), Hi-Star internally switched architecture is designed to truly scale in both capacity and performance and does not have to contend for limited bus bandwidth.

Figure 3 – A shared-bus architecture is limited to two simultaneous I/O operations.



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

The Lightning 9900 V Series Systems Have an Advanced Availability Profile Compared to Competitive Products

The Lightning 9900 V Series was designed with maximum emphasis on high-availability computing for today's most critical enterprises. Chapter 7 discusses some of the many advantages in high-availability design that the Lightning 9900 V Series has over competitive products including:

- System microcode which can be upgraded nondisruptively
- Active/active dual-ported disk drives instead of single-port active drives

- Separate storage for control and configuration tables
- Mirrored cache instead of single image cache²
- Offered by Hitachi Data Systems, the leading supplier of storage systems for geographically GDPS™ production sites
- NanoCopy™ technology which offers a scheme for disaster recovery
- High-availability, custom-designed disk drives

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 Data Systems, is available in the companion publication, the ***Hitachi Freedom Storage Software Solutions Guide***. This document is downloadable in PDF format from www.hds.com. The report contains expert guidelines regarding industry-leading software solutions suites from Hitachi Data Systems for storage management, business continuity, performance enhancement, backup and recovery, and data transfer. It also compares the Hitachi Data Systems solutions to the major competitive software solutions and drills down into specifics.

Since enterprise storage features are increasingly required in the modular storage market, additional information on the Hitachi Freedom Storage Thunder 9500 V™ Series systems is available in a separate report entitled the ***Hitachi Freedom Storage Thunder 9500 V Series Architecture Guide***. The Thunder 9500 V Series are modular (midrange) RAID storage systems designed for use in heterogeneous open systems computing environments and targeted at the e-commerce, Web serving, and data warehousing applications. The Thunder 9500 V Series storage systems are the only modular products that truly deliver large, enterprise-level features with a midrange package and price.

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.

Other reports, presentations, and expert knowledge are also available.

Reduced TCO


Bandwidth contention in a shared-bus architecture serves as a limit to data hyper-consolidation initiatives and limits the enterprise's ability to reduce TCO. Figure 4 illustrates that a recent survey showed that actual deployed average capacity of Lightning 9980V systems and Lightning 9970V systems exceeded that of fixed-bus systems by a factor of 2.5x. Since the Lightning 9980V systems and Lightning 9970V systems nearly double internal bandwidth of our earlier enterprise-level systems, it is clear that obstacles to massive consolidation projects will be removed. This is especially true when the Hitachi industry-unique "virtual" Fibre Channel port feature of the Lightning 9900 V Series is considered.

The major objective of massive consolidation is to keep up with exploding storage capacities without increasing head count and deploying staff to address the application backlog. With the Lightning 9900 V Series, TCO can be reduced not only through virtualization but also by:

- Reduced network complexity with fewer switches, connections
- Simplified and automated management tools and software
- Lower licensing, maintenance charges, and environmental costs

² Some competitive products do *not* have dual cache like the Hitachi Freedom Storage Lightning 9900 V Series. 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 V Series 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 24/7 operations.


Figure 4 – Actual average deployed capacity of the Lightning 9900 V Series systems exceeds that of shared-bus systems by a factor of 2.5x.



Second-generation “Virtual” Switched Storage

An important key to lower TCO is the Lightning 9900 V Series “virtual” Fibre Channel ports. Virtualization is becoming increasingly important in today’s heterogeneous storage environment, which is made up of these storage architectures: direct attached storage (DAS), storage area networks (SANs), and network attached storage (NAS), as shown in Figure 5. This increases the need to centralize and efficiently manage all heterogeneous storage with uniform enterprise policies and procedures.

Figure 5 – Today’s heterogeneous IT environment demands virtualization of storage.



As shown in Figure 6, the heterogeneous solutions of today require the separate configuration of physical ports for each host platform type supported, particularly in high availability path failover and clustered environments. This both increases expense through unnecessary extra ports and limits heterogeneous consolidation, resulting in higher TCO.

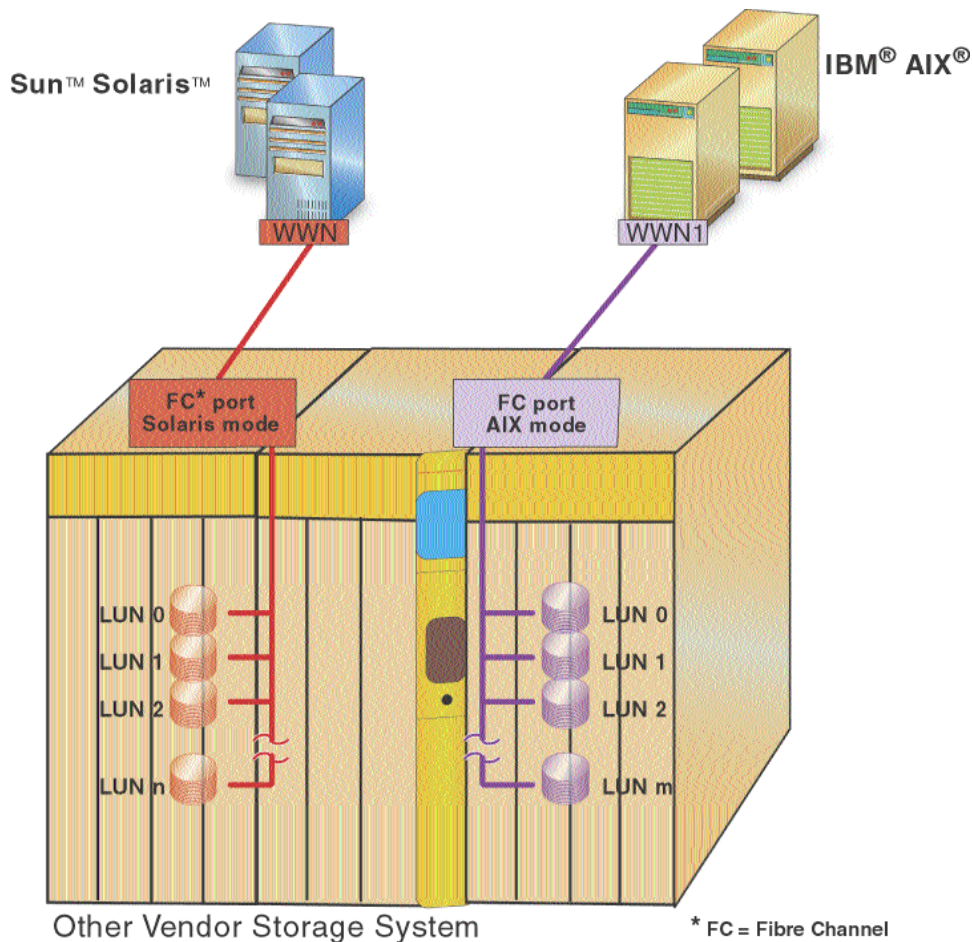



Figure 6 – Heterogeneous solutions today.

The Lightning 9900 V Series provides “virtual” Fibre Channel ports that are logically managed by intelligent Fibre Channel controller cards. Each physical Fibre Channel port may have multiple Host Storage Domains. A Host Storage Domain supports a “logical” Fibre Channel port (Fig. 7) each with its own set of LUNs. Hosts are matched to their assigned HSD based on a unique World Wide Names (WWNs) identifier. LUN security provided by Hitachi SANtinel™ software fences access to LUNs based upon host WWNs. Under this innovative software structure, each of multiple host systems may have its own unique LUN 0, from which it can boot, within its own Host Storage Domain.

Figure 7 – The Lightning 9900 V Series virtual Fibre Channel ports and Host Storage Domains.



Hitachi Freedom Storage™ Lightning 9900™ V Series Packaging

2

The Hitachi Freedom Storage Lightning 9980V™ System Packaging

The Lightning 9980V system is available in a two-to-five 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 four Array Frames that are connected via high-speed Fibre Channel links. A fully configured system is contained in five cabinets 73.2 inches high, 148.9 inches wide, and 31.5 inches deep.

Lightning 9980V System Control Frame

The Control Frame of the Lightning 9980V system 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 8.





Figure 8 – A Lightning 9980V system Control Frame and its major components.

Lightning 9980V System Array Frames

Each array frame of the Lightning 9980V system contains up to 256 Hard Disk Drive (HDD) modules and communication interfaces to the Control Frame, for a total of 1024 drives for a maximum configured system. An Array Frame is shown in Figure 9.

This design allows for efficient centralized and consolidated storage, which results in reduced management costs and lower overall TCO.

Figure 9 – A Lightning 9980V system Array Frame with its major components.



The Lightning 9900 V Series power systems consist 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, if a power supply fails during operation, the failed power can be dynamically removed without any loss to system operation. The Lightning 9900 V Series power supply modules are shown in Figure 10.

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 V Series models to connect with the different electrical outlet configurations used around the world.





Figure 10 – Power supply modules for the Lightning 9900 V Series systems.

Hitachi Freedom Storage Lightning 9970V™ System Packaging

Although the Lightning 9980V system has the ability to be upgraded from one to four array frames, the Lightning 9970V™ system is a fixed-configuration, single-cabinet storage system. The Lightning 9970V system is targeted for IT environments where floorspace is at a premium and data growth is limited, yet where premium storage functionality is required. From the front, the Lightning 9970V system looks very much like the Control Frame of a Lightning 9980V model. However, the back end contains HDD arrays housing up to 128 disks in a 73.2 inch high, 30.8 inch wide, and 31.5 inch deep footprint. This is shown in Figure 11.

Figure 11 – A Lightning 9970V system single-cabinet storage system.



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

The System Monitoring Network of the Lightning 9900 V Series

Lightning 9900 V Series systems are designed with a system-monitoring network. The system-monitoring cables are connected to all cabinets in Lightning 9900 V Series systems, and they 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™ storage package across a private LAN on a UNIX® or Microsoft® Windows NT® console. Users can view failure information via SNMP or IBM® 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 Freedom Storage™ Lightning 9900™ V Series

3

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

The Lightning 9900™ V Series supports both Fibre Channel for open systems and ESCON® or FICON™ for IBM® S/390® and zOS™ compatible platforms as shown in Figure 12.

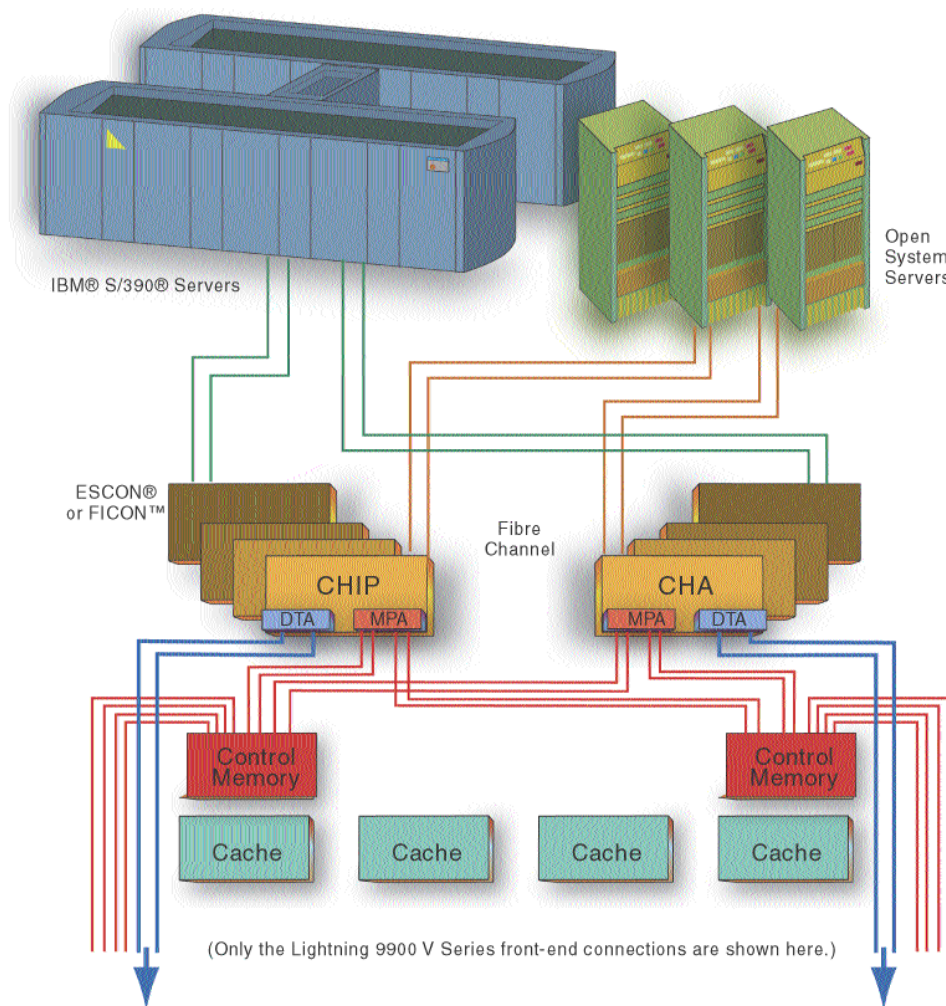



Figure 12 – Fibre Channel and ESCON or FICON adapter boards connect to open systems and mainframe computers.

About the Fibre Channel Interface Boards

Each Fibre Channel Interface Board, as shown in Figure 13, has four Interface Processors for management of connectivity to the host. Each port supports a transfer rate of 2Gbit/sec (gigabits/second) in conformance with the Fibre Channel standard. An optional eight-port Fibre Channel Interface Board is also available.

Figure 13 – A 2Gbit/sec Fibre Channel host adapter board and its major hardware components.



With the Lightning 9900 V Series industry-leading implementation of “virtual” Fibre Channel ports, each interface card is equivalent to multiple interface cards in competitive systems. Each interface card also contains Oracle® checksum dedicated hardware logic.

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 14. A FICON interface board has a similar appearance. The FICON interface card is available with 1Gbit/sec or 2Gbit/sec transfer rates.





Figure 14 – An ESCON host adapter board and its major hardware components.

A Lightning 9900 V Series system can connect to both mainframe and open system hosts, thereby consolidating the storage for an enterprise under one roof for consistent performance, availability, and ease of management. This is shown in Figure 15.

Figure 15 – The Lightning 9900 V Series systems support both mainframe and open system hosts.



Bandwidth of Supported Host Interfaces

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

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

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

Number of Host Interfaces Supported

Up to 64 Fibre Channel, 48 ESCON, or 32 FICON ports are available on the Lightning 9980V™ system and up to 24 on the Lightning 9970V™ system. Only 2 ACP pairs are sequenced on the Lightning 9980V system for 48 ESCON ports.

Types of Hosts Supported with Direct Fibre Channel Host Attachment by the Lightning 9900 V Series Systems

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

- Microsoft® Windows NT®
- Microsoft Windows® 2000
- HP-UX®
- Solaris™ (Sun™)
- AIX® 32 and 64 bit (IBM®)
- NetWare® (Novell®)
- Linux® (United Linux®, Red Hat®, SuSE®, Turbolinux®)
- IRIX® (SGI™)
- Tru64™ UNIX® (HP®)
- OpenVMS™ (HP®)
- DYNIX/ptx® (IBM®)
- Mainframes (MVS®, VSE®, VM®, OS/390®, zOS®, Linux®)

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 16 provides a detailed illustration of the processor paths and the C-HSN of a Hitachi Freedom Storage™ Lightning 9900™ V Series Lightning 9970V™ system.

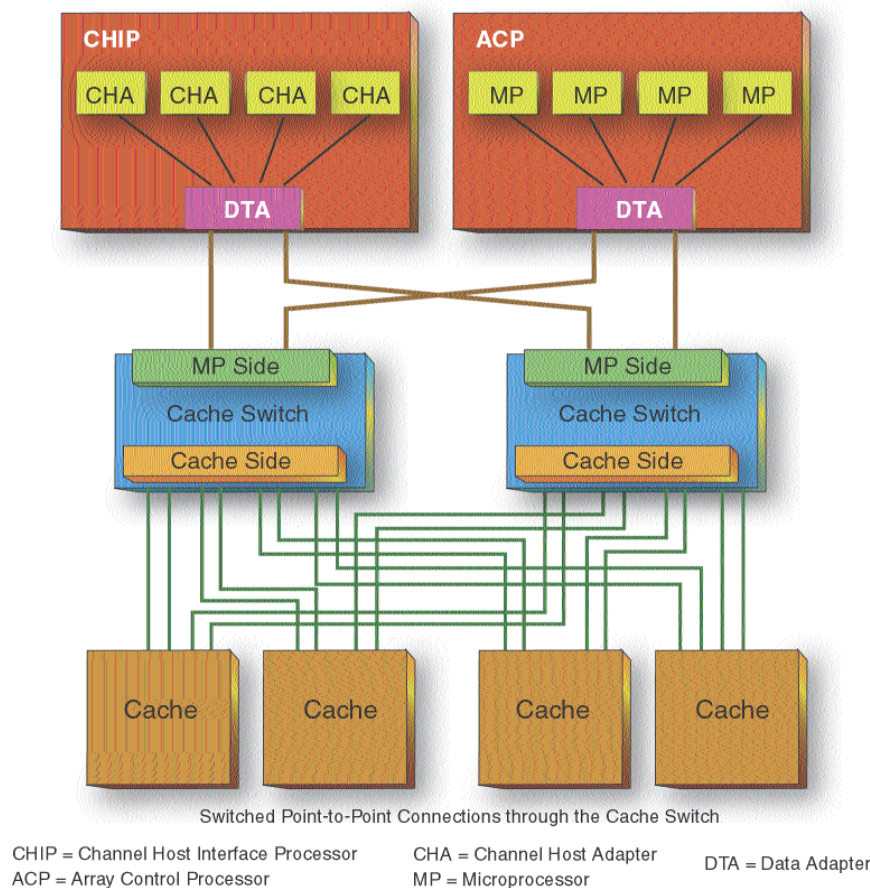


Figure 16 – The C-HSN of a Lightning 9970V system provides fully redundant switched access to cache from all channel adapters.

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

The Cache Switch Provides Non-blocking Switched Access to Cache

At the heart of the Lightning 9900 V Series are four Cache Switches (CSWs). Together, these four switches (in the Lightning 9980V™ system) use a parallel switch fabric bus (PSFB). 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 eight paths to the cache modules. All total, there are 32 paths at the processor side of the fabric network and 32 paths to the cache modules from the cache side of the fabric network.

The Lightning 9980V System Provides 64GB of Fully Addressable Cache

The Lightning 9980V system supports up to 64GB of data cache, all directly addressable. Separate cache modules (up to 3.0GB) are used for control storage. Competitive systems use their cache for both data and control information, limiting the amount of usable data cache and creating performance limitations for certain workloads.

Advanced Cache Algorithms of the Lightning 9900 V Series

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

Hitachi FlashAccess™ Software Allows Data Sets to be Permanently Placed in Cache

Hitachi FlashAccess software 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. FlashAccess software, a portion of cache memory can be allocated to specific data. Administrators can add, delete, or change FlashAccess software managed data at any time, quickly and easily.

In IBM® 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, FlashAccess software offers multiple modes of operation. It can be used in conjunction with Hitachi RapidXchange™ software to increase the speed of data transfer and, therefore, improve performance of mainframes to open systems data exchange. RapidXchange software supports both open-to-S/390 and open-to-open high-speed data transfers.

Read-ahead for High-performance Sequential Reads

Read clustering in the Lightning 9900 V 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 V 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 V Series system (as well as other information related to the operational state of the system). Two control memory areas are

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

mirrored images of each other. This is illustrated in Figure 17. Control data is “data about data,” also called “metadata.” 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 17 shows a close-up view of the CM-HSN’s networking topology. Referring back to the diagram in Figure 2, Chapter 1, there are two CM-HSN paths connecting the processors to the control memory. However, the diagram in Figure 16 shows four paths per processor module. There are up to 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 17 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.

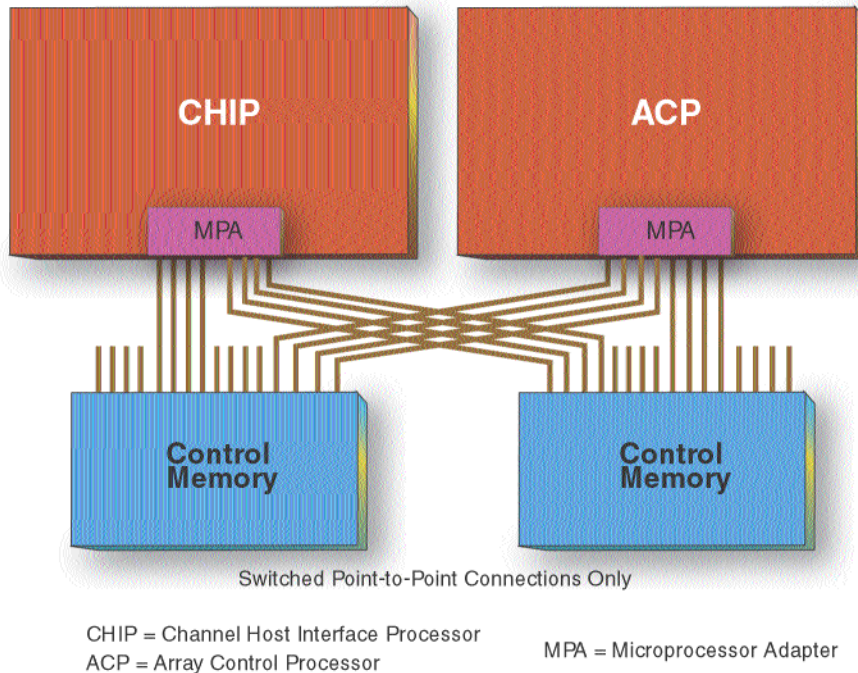


Figure 17 – Separate redundant control memory handles the exchange of control information, such as status, location, and configuration of data, between processors and cache memory.

High-performance Back-end Design 5

High-speed All Fibre Channel Back-end Design

Special high-performance back-end drives are available with the Hitachi Freedom Storage™ Lightning 9900™ V Series. All the drives are 3.5" Low Profile (LP) 1"-height form factor with capacities available in 36GB, 73GB, and 146GB drives. The 36GB drives spin at 15,000RPM, and the 73GB and 146GB drives spin at 10,000RPM.

Each FC-AL path is driven with a dedicated processor. The diagram of a Lightning 9900 V Series ACP pair is shown in Figure 18. The Lightning 9900 V Series uses advanced algorithms for managing performance of ACP pairs.




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

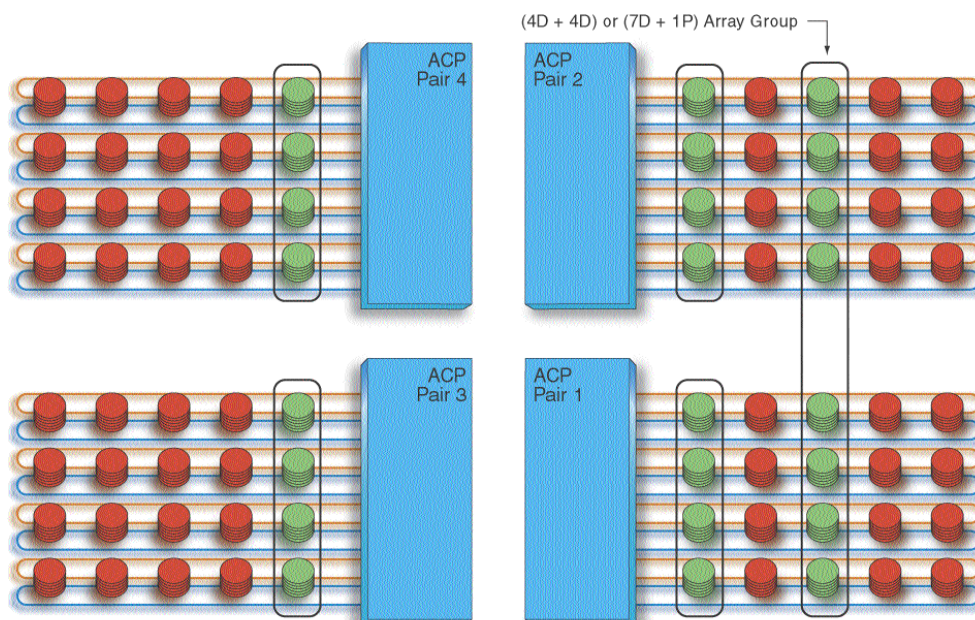


Figure 19 – Diagram of four ACP pairs and 32 FC-AL back-end disks for the Lightning 9980V™ system.

Maximum Configuration of ACPs and Disks

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

New Fibre Channel Disk Drives

The back-end of the Lightning 9900 V Series uses dual ported/dual active Fibre Channel disk drives. All the drives are 3.5" Low Profile (LP) 1"-height form factor with capacities available in 36GB, 73GB, and 146GB drives. The 36GB drives spin at 15,000 RPM, and the 73GB and 146GB drives spin at 10,000RPM. These dual-ported/dual-active Fibre Channel disk drives, combined with the technology built into the ACP pair, allow the back end of Lightning 9900 V Series systems to use all eight FC-AL paths in an ACP pair for both performance and fault tolerance. 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.

Command Tag Queuing greatly improves performance.

The Importance of Command Tag Queuing (CTQ)

Another new feature introduced with the Lightning 9900 V Series is Command Tag Queuing (CTQ) to the back-end drives. CTQ greatly improves the performance of 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.

Capacities and RAID Architecture

6

Storage Capacity of the Hitachi Freedom Storage™ Lightning 9900™ V Series

The Lightning 9900™ V 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 36GB drives	Raw capacity with 73GB drives	Raw capacity with 146GB drives
Lightning 9980V™ System One Array Frame	9.2TB = 256 disks	18.7TB = 256 disks	37.2TB = 256 disks
Lightning 9980V System Four Array Frames	36.6TB = 1024 disks	74.7TB = 1024 disks	149TB = 1024 disks
Lightning 9970V™ System	4.6TB = 128 disks	9.3TB = 128 disks	18.6TB = 128 disks

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

Advantages of the Lightning 9900 V Series RAID Hardware

Hitachi designs the most advanced RAID (Redundant Array of Independent Disks) controllers in enterprise storage to interface to its disk systems. These intelligent controllers provide disk interface and RAID management, offloading these tasks to dedicated processors. Each Lightning 9900 V Series ACP controller supports RAID-1+ and RAID-5 (parity RAID). All user data disks in the system are defined as part of a RAID array of one type or another.

Intelligent controllers offload tasks to dedicated processors.

RAID-1+

For the Lightning 9900 V Series, RAID-1+ is available in 2 Data + 2 Data and 4 Data + 4 Data disk arrangements in a RAID-1+0 configuration.

RAID-1+ is unique to Hitachi.


RAID-5 – Distributed Parity

RAID-5 disk arrangements for the Lightning 9900 V Series consist of four disks (3 Data and 1 Parity) of eight disks (7 Data and 1 Parity). 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 V Series also allows logical volumes to be expanded online. The 2D +2D RAID-1+ configuration and 3D + 1P and 7D + 1P RAID-5 configurations are shown in Figure 20.

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



High-density Disk Drive (HDD) Assemblies

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

The disk drives used in the Lightning 9900 V Series are specially designed for high availability and performance. Each drive has the industry unique feature of supporting dual active ports. This provides significantly increased performance.

Hitachi Freedom Storage™ Lightning 9900™ V Series Availability Is Highest in the Industry

7

An Overview of High-availability (HA) Computing Software

It is increasingly important for a business to support continued access to global information 24/7. 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 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*.

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™ V Series has been designed to eliminate as many of these factors as possible in hardware redundancy, online replaceable components, and software data copy functions. This allows 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 VERITAS® Software, 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 V Series Has an Advanced Availability Profile Compared to Competitive Products

The Lightning 9900 V Series was designed with maximum emphasis on high-availability computing for today's most critical enterprises, including:

- System microcode that can be upgraded nondisruptively³
- Active/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 failover architecture

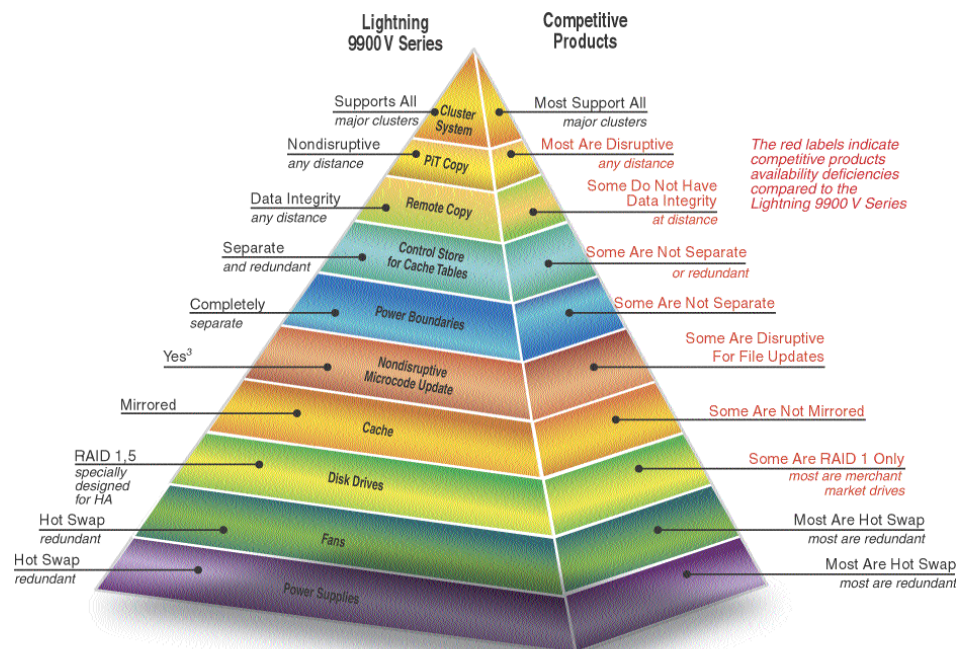
³ With the exception of Seagate® HDD microcode.

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

- Dual data paths and dual control paths connecting every component
- Mirrored cache for all write data instead of single image cache⁴
- IBM® Geographically Dispersed Parallel Sysplex™ (GDPS™) support
- Hitachi TrueCopy™ software combined with Hitachi ShadowImage™ software for disaster recovery
- NanoCopy™ technology to provide point-in-time copies without requiring quiescence of databases or applications
- High-availability, custom-designed disk drives
- Hi-Track® “call-home” capability to notify the Hitachi Data Systems support center quickly of any incident that may affect availability

Figure 21 highlights some of the many advantages in high-availability design that the Lightning 9900 V Series has over other enterprise storage products.

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



Open Systems High-availability Middleware Support

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.

⁴ Some competitive products do *not* have dual cache like the Lightning 9900 V Series systems. 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 V Series systems do not have this problem since all status areas and write data are duplexed.

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

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 V Series supports each classification of middleware:

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 V Series systems support alternate pathing for IBM AIX®, HP® True64™ UNIX®, HP-UX® and through PVlink, Microsoft® Windows NT®/Windows® 2000, Sun™ Solaris™, and IBM DYNIX/ptx®.

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

Hitachi Dynamic Link Manager™ Software Provides Path Failover and Load Balancing

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

Dynamic Link Manager software 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 Dynamic Link Manager software is required for true nonstop operations. This product allows systems administrators to take advantage of the multiple paths on a Lightning 9900 V Series system by adding redundant connections between application servers and RAID systems. Dynamic Link Manager software therefore provides increased reliability and performance. Supported platforms include IBM AIX, Sun Solaris, Windows NT, Windows 2000, and LINUX® (July 2003).


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 "takeover" includes the reassignment of networks and peripherals, as well as the restarting of applications. Host clustering can also be used to create fault-tolerant workloads and scale processor capabilities while sharing network and disk resources.

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

Lightning 9900 V Series systems supports all major open systems clustering schemes including: HP TruCluster™, HP MC/ServiceGuard™, HP MC/Lock Manager, IBM RS6000® and SP HACMP™, Windows NT/Windows 2000, Microsoft Cluster Server, NCR UNIX SVR4 Lifekeeper, IBM DYNIX/ptx ATAP, VERITAS Cluster Server™ for Sun Solaris and Red Hat® Advanced Server 2.1 and SuSE Enterprise Server 7.0 for LINUX®.

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



Parallel Database Clustering

Clustering middleware supports distributed lock management.

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 V Series systems have 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 V 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 system.

Standard Redundant Fans for Cooling

The Lightning 9900 V Series 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.

Competitive Advantage in RAID Choices

The flexibility in choice for Lightning 9900 V 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 V 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 V Series system. For example, a single Lightning 9900 V Series system can configure both RAID-1+ and RAID-5.

With the Lightning 9900 V Series flexible LUN configuration feature, the user can select LUN sizes as small as 46MB or as large as almost 2TB (using LUSE).

Unmatched Performance and Scalability

8

Truly Scalable Performance

Based on the internal switched Hi-Star™ architecture, the Hitachi Freedom Storage™ Lightning 9900™ V Series shows unmatched performance scalability. This chapter discusses the ability to dramatically add more online transaction processing (OLTP) users and mixed decision-support and OLTP workloads without a user response time decrease. The more capacity and connectivity added to a Lightning 9900 V Series system, the more total the bandwidth increases. As Table 4 shows, total aggregate bandwidth is ten times that of competition.

When new application workloads are added to a Lightning 9980V™ system with an existing workload, neither the new nor the existing workloads suffer performance degradations from the sharing of internal resources and bandwidth. The following series of tests proves this point. By measuring user-response time while workloads are added to the system, we can illustrate this phenomenon.

Model Type	Control Bandwidth	Data Bandwidth	Aggregate Bandwidth	User-useable Throughput
Lightning 9970V™ system Single Cabinet	2.6GB/sec	5.3GB/sec	7.9GB/sec	5.3GB/sec
Lightning 9980V™ system Multi-Cabinet	5.3GB/sec	10.6GB/sec	15.9GB/sec	10.6GB/sec
Competition	1.6GB/sec	1.6GB/sec	1.6GB/sec	1.6GB/sec

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

Table 4 – Lightning 9980V system bandwidth is ten times that of competition.

Internal Switching Overcomes the Limitations of Shared-bus Architectures

It can be seen from Figure 24 that very little user-response time reduction occurs in the OLTP performance of the Lightning 9980V system. This is because the Lightning 9900 V Series system, with its Hi-Star architecture, essentially has its own internally routed path for each workload. This means that the main reason for performance decline of a competitive shared-bus architecture, i.e. bus contention, does not occur with Lightning 9900 V Series systems. Although there may be some contention in cache, many simultaneous operations concurrently happen at the cache, so it far exceeds the demands of most applications.

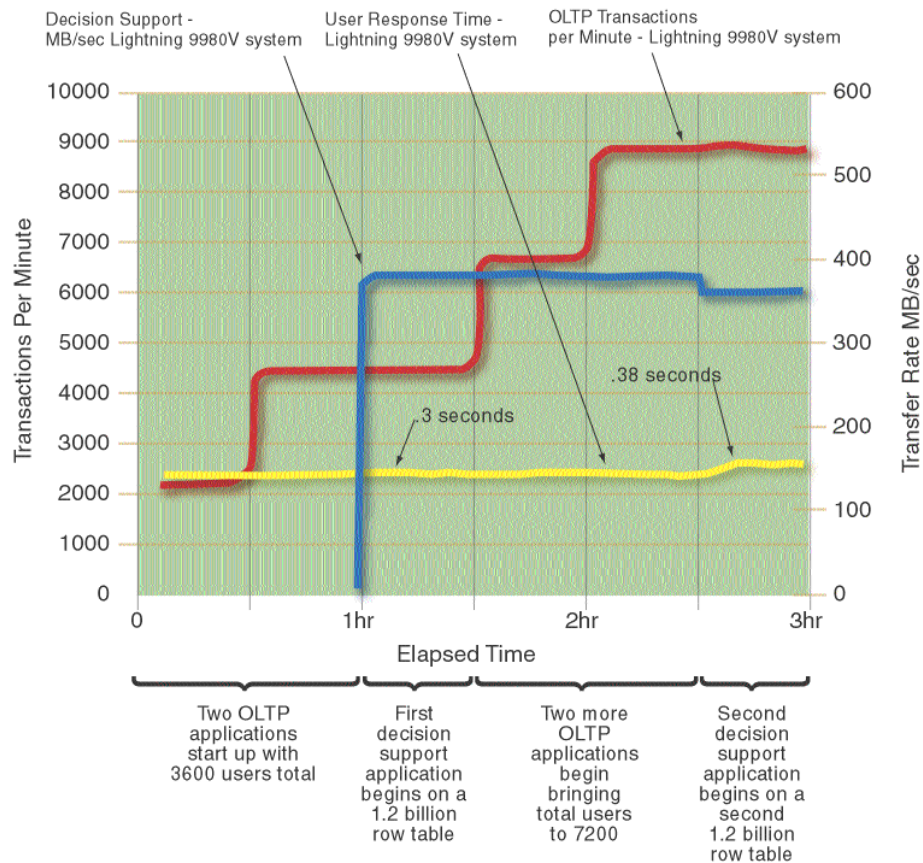
Bottleneck-free performance with the Lightning 9900 V Series.

There Is NO Match for the Performance of the Lightning 9980V System

As shown in Figure 23, during the first 30 minutes of the test 1800 simulated users were connected to the first OLTP server, and the transactions per minute (TPM) rate and user-response time reached steady state at about 2200TPM and 3.0 seconds respectively. OLTP user response time is the delay time between when a simulated user issues a transaction request and when the request responds with a commit. After 30 minutes, a second OLTP server was connected with an additional 1800 simulated users, bringing the total number of users to 3600.

After one hour, the first decision support system (DSS) workload was started. The DSS workload began a series of full table scans, counting records from a 1.2 billion row table. The transfer rate for the DSS system was 380MB/sec or 3.4 million rows per second. We

Figure 23 – Test results show that the advantages of the Lightning 9980V system maintain user response time while workloads dramatically increase.



overshot the targeted transfer rate for the DSS system by approximately 50MB/sec. This was managed by using Hitachi Priority Access software to limit the throughput for a particular application. At this point, the two OLTP workloads were still able to sustain their 2200TPM each, or 4400TPM aggregate. The three combined workloads were allowed to run for approximately 30 minutes together to obtain a steady state.

From 90 minutes to 150 minutes of the test, two more OLTP servers were added to the workload of the Lightning 9980V system, bringing the total number of users to 7200 and the TPM to 8800, a linear increase from the time the first OLTP server was started with a TPM of 2200.

From 150 minutes to 180 minutes of the test, the second and final DSS system started its table scan activity. This mixed workload of six total servers (four OLTP and two DSS) caused only a slight (2.6 percent) reduction in the first DSS system's throughput from a transfer rate of 380MB/sec to 370MB/sec. However, the two DSS systems were performing full tables scans at an aggregate rate of about 740MB/sec or 6.6 million rows per second each. During this time the total OLTP transaction rate was maintained at 8800TPM; however, the average user-response time increased insignificantly to 0.38 seconds per transaction. The 0.08 second increase was considered noise at these levels of measurement, but is recorded here for discussion.

After three hours, the Lightning 9980V system was under a heavy mixed-performance workload characterized by a mix of sequential reads at high bandwidths and random reads and writes. Total processor utilization for the entire system was only 30 to 35 percent.

Hitachi Freedom Storage™ Software Solutions Are *Best-of-Breed*

9

Customers and analysts regard the Hitachi Freedom Storage™ Lightning 9900™ V Series as the most advanced products in the enterprise storage system market. The ***Hitachi Freedom Storage Software Solutions Guide*** (Figure 24) discusses the advantages of the suite of leading-edge software products that enhances the Lightning 9900 V Series product lines. This report is available from Hitachi Data Systems in PDF form at: <http://www.hds.com>. Particular emphasis is devoted to the business benefits of these robust suite of software in the context of massive consolidation projects that are for the first time enabled by the switched architecture and virtual Fibre Channel ports of the Lightning 9900 V Series systems' advanced design. HiCommand™ Management Framework software enables a dramatically lower total cost of ownership (TCO) for the enterprise due to its ability to provide centralized, worldwide control of Hitachi Freedom Storage Software Solutions from a centralized location.

Hitachi Freedom Storage™ Software Solutions are the most advanced available.

Hitachi Freedom Storage Software Solutions Simplifies Information Management, Protects Data Assets, and Optimizes Storage Investments.

With Hitachi Freedom Storage Software Solutions, customers have the ability to choose the precise solution or a combination of solutions that are appropriate for their environment. Software solutions are combined into five Suites with each Suite targeting specific business values as discussed below. In keeping with the Hitachi commitment to open solutions, Hitachi Freedom Storage Software Solutions include not only software solutions provided by Hitachi Data Systems, but also solutions provided by TrueNorth™ Solutions Alliance partners.

Storage Area Management Suite

The Storage Area Management Suite includes five powerful software packages that dramatically simplify storage management complexity, reduce operating costs and improve return on investment (ROI). The HiCommand Management Framework enables users to construct the most suitable management environment. Hitachi Data Systems is working with the Storage Networking Industry Association's (SNIA) Storage Management Initiative to deliver an interoperable storage management framework based on the Common Information Model (CIM) and the Simple Object Access Protocol (SOAP) industry standards. (See also Chapter 10). By deploying the Sun™ StorEdge™ Resource Management Suite with the HiCommand Management Framework, users can develop products that seamlessly integrate into the framework to improve resource utilization, optimize storage network planning, enhance data continuity solutions, and exceed existing service levels. Resource Manager software is a package of management software utilities, including Hitachi Graph-Track™ software, Hitachi FlashAccess™ software, Hitachi Virtual Logical Volume Image (VLVI) Manager, Hitachi LUN Manager software, and Hitachi SANtinel™ software. The Resource Manager software package eases complex storage management processes and allows users to handle oncoming performance and capacity requirements proactively. (See also Chapter 10). VERITAS Volume Manager removes the physical limitations of disk storage, so users can configure, share, and manage storage in enterprise environments without interrupting data availability. InterSAN PATHLINE provides comprehensive,

The Storage Area Management Suite simplifies storage management complexity, reduces costs and improves ROI

Table 5 – Summary of Hitachi Freedom Storage software products.

Software Suite	Hitachi Freedom Storage™ Product Support			OS Support	
	Thunder 9531V™ Thunder 9532V™ Thunder 9533V™	Thunder 9570V™	Lightning 9900™ V	Open	IBM® S/390®
Storage Area Management Suite					
HiCommand™ Management Framework	✓	✓	✓	✓	
- HiCommand Device Manager	✓	✓	✓	✓	
- HiCommand Tuning Manager	✓	✓	✓	✓	
Hitachi Resource Manager	✓	✓	✓	✓	✓
- Hitachi Graph-Track™			✓	✓	✓
- Hitachi FlashAccess	✓	✓	✓	✓	✓
- Hitachi Virtual Logical Volume (VLVI) Manager			✓		✓
- Hitachi LUN Manager	✓	✓	✓	✓	
- Hitachi SANTinel	✓	✓	✓	✓	✓
Hitachi Command Control Interface	✓	✓	✓	✓	✓
Hitachi StorEdge Resource Management Suite	✓	✓	✓	✓	✓
VERITAS Volume Manager	✓	✓	✓	✓	
InterSAN PATHLINE	✓	✓	✓	✓	✓
Business Continuity Suite					
Hitachi TrueCopy™		✓	✓	✓	✓
NanoCopy™			✓		✓
Hitachi ShadowImage™	✓	✓	✓	✓	✓
Hitachi Extended Remote Copy (HXRC)			✓		✓
Hi-Track®	✓	✓	✓		✓
Hitachi Dynamic Link Manager™	✓	✓	✓	✓	
Geographically Dispersed Parallel Sysplex™ (IBM)			✓		✓
Hitachi SplitSecond™ Solutions	✓	✓	✓	✓	✓
Data Protection Services	✓	✓	✓	✓	✓
Backup and Recovery Suite					
Hitachi Multiplatform Backup/Restore			✓		✓
Hitachi e-Copy			✓	✓	
VERITAS NetBackup™	✓	✓	✓	✓	
TANTIA Technologies® HARBOR® Backup			✓		✓
CommVault™ Galaxy™	✓	✓	✓	✓	
Performance Enhancement Suite					
Hitachi CruiseControl™			✓	✓	✓
Hitachi Priority Access			✓	✓	✓
Hitachi FlashAccess	✓	✓	✓	✓	✓
Hitachi Parallel Access Volume (PAV)			✓		✓
Multiple Allegiance (MA)			✓		✓
Hitachi Graph-Track			✓	✓	✓
Hitachi Dynamic Link Manager	✓	✓	✓	✓	
Hitachi Performance Maximizer			✓	✓	✓
- Hitachi Performance Monitor			✓	✓	✓
- Hitachi CruiseControl™			✓	✓	✓
- Hitachi Priority Access			✓	✓	✓
Data Movement Suite					
Hitachi RapidXchange™			✓		✓
Data Protection Services	✓	✓	✓	✓	✓
TANTIA Technologies HARBOR File Transfer			✓		✓

policy-driven automated management across the entire life cycle of shared storage services, ranging from planning, implementation, and provisioning, to monitoring.

Business Continuity Suite

A well-planned business continuity strategy is essential if an organization is to survive planned or unplanned downtime. The Business Continuity Suite includes nine software solutions that enable users to protect data assets by ensuring continuous access to information. TrueCopy software addresses the challenge of non-stop information processing by replicating information between Lightning 9900 V Series systems. Copies may be replicated between systems contained within one data center or to remote models in geographically dispersed locations—with minimal impact on performance. For S/390 environments, Hitachi NanoCopy technology offers the industry's first automated solution to make copies over long distances with zero outage time. For the testing and rapid deployment of new applications, ShadowImage software rapidly replicates LUNs within Lightning 9900 V Series storage systems without disrupting or affecting production

The Business Continuity Suite protects data assets by ensuring continuous access to information




Figure 24 – Hitachi Data Systems offers a broad range of remote copy, data duplication, and data migration software solutions which are profiled in the Hitachi Freedom Storage Software Solutions Guide.

information and performance levels. The Hi-Track "call home" predictive maintenance system is based on intelligent software that monitors a Lightning 9900 V Series system and notifies the Hitachi Data Systems support center for preventative maintenance. Dynamic Link Manager software solution increases information availability by ensuring complete utilization of all paths between a server and a Lightning 9900 V Series system. HXRC⁵ allows users to create and share server-based remote copies between Lightning 9900 V Series systems and an IBM Enterprise Storage Server™ system. Hitachi Data Systems has demonstrated proven TrueCopy software support for GDPS, an IBM service offering for systems failover, workload balancing, and data mirroring. In addition, Hitachi Data Systems offers a number of services to help users plan, integrate, and deploy a GDPS environment based on Lightning 9900 V Series systems. The SplitSecond Solutions package provides a communications layer to seamlessly integrate a customer application with ShadowImage software and TrueCopy software to ensure that essential information can survive a rolling disaster with full integrity. But what good is the world's best software without first-rate consulting services to go with it? Data Protection Services are offered by the Global Solutions Services team to analyze and implement all software solutions in the Business Continuity Suite.

Backup and Recovery Suite

Backup/Restore is the largest cost element in storage ownership and the most important to high-availability computing. In today's global IT environment, it is no longer enough to have a backup copy of one's data.

The Backup and Recovery Suite consists of eight powerful software packages that allow businesses to perform copy, backup and restore in the shortest possible time and with minimum disruption to information availability. Hitachi Multiplatform Backup/Restore software delivers high-performance volume-level backup of open systems data to a S/390 host thereby leveraging mainframe based storage management procedures. The Hitachi e-Copy software feature enables nondisruptive backups to be made directly from disk to tape in a SAN environment without incurring server and I/O overhead during actual data movement. Hitachi Data Systems has teamed up with industry leaders to deliver first-rate solutions that simplify the copy, backup and recovery processes. Through these strategic alliances, users can seamlessly and cost-effectively deploy, implement, and optimize backup tasks. CommVault™ Galaxy™, Tantia Technologies® HARBOR® Backup, and VERITAS NetBackup™ provide choices to help safeguard information using a solution that best fits user requirements.

Performance Enhancement Suite

For today's enterprise, instant access to mission-critical information is a must. The Performance Enhancement Suite optimizes performance for business critical online applications with eight powerful software packages. Hitachi CruiseControl software automatically monitors, analyzes, and balances the Lightning 9900 V Series system while offering tuning recommendations for approval in either automatic or assisted mode. This automated software tool replaces traditional, time-consuming, and sometimes error-prone, manual load-balancing processes with streamlined procedures for improved performance efficiency and overall response time. To further maximize system performance, Hitachi Priority Access software allows users to define host access priority into a Lightning 9900 V Series system, capitalizing on this series' port virtualization technology. Combining Priority Access software with Host Storage Domains allows host pooling at the logical unit level, optimizing connectivity and reducing networking infrastructure. For Lightning 9900 V

The Backup and Recovery Suite protects information with sophisticated copy, backup and restore functions

⁵ HXRC is a cross license technology shared by Hitachi Data Systems and IBM.

Series' systems, Priority Access software and Host Storage Domains guarantee that the most crucial applications get the highest level of service. FlashAccess allows users to quickly and easily lock and unlock information into cache in real time. A user can confidently "fill it up and max it out" to meet performance requirements using the maximum raw capacity of the Lightning 9900 V Series system. Hitachi Parallel Access Volume (HPAV) software permits multiple applications running on an S/390 server to access the same information simultaneously. The Multiple Allegiance (MA) software solution extends the HPAV software capability to applications running on multiple S/390 servers. Together, HPAV software and MA software reduce queuing, which results in significantly decreased batch times and swifter responses in today's high-transaction environments.

Graph-Track software is a unique and effective tool that has the scope to monitor all aspects of hardware performance from a centralized utility. Through a GUI, information can be exported to multiple formats quickly and easily for analysis and reporting programs. Graph-Track software displays real-time or historical information for all connected Lightning 9900 V Series and Thunder 9500 V Series storage systems. Identifying utilization peaks and ongoing trends positions users to make rapid and informed decisions as needed to stay competitive. Dynamic Link Manager software automates I/O load balancing, path failover, and recovery capabilities in the event a single path breaks down. In a unique balancing act, Dynamic Link Manager software ensures that no single path becomes congested and overworked while another is underutilized. By automatically allocating data to an alternate path, the threat of application failure is removed. The Performance Maximizer package includes Hitachi Performance Monitor and CruiseControl software for intelligent performance tuning. The third component, Priority Access software, works with port virtualization and VERITAS Volume Manager to provide the highest level of service access available today.

The Performance Enhancement Suite optimizes performance for business critical online applications

Data Movement Suite

Information needs to be exchanged regularly for such tasks as data synchronization, populating data warehousing systems, or application development and testing. Typically, data is exchanged in either direction, between a mainframe and open systems. However, the ever-diminishing mainframe batch window, coupled with increasing volumes of information, often impedes this process. The Data Movement Suite includes three software packages that eliminate many of the constraints related to exchanging information in heterogeneous environments. RapidXchange software saves time as it delivers high performance and reliable data sharing among heterogeneous host platforms. Swift information transfer is achieved without placing additional burdens on the network infrastructure or tape transport equipment. By improving exchange processes between mainframe and open systems, the threat of over-stressing the network is reduced. When combined with FlashAccess software, RapidXchange software delivers information even faster to keep pace with widespread applications and their demands.

Multi-vendor file transfer capabilities are demonstrated in the Hitachi Data Systems strategic alliance with Tantia Technologies. HARBOR File Transfer adds automation to large data file transfers at ultrahigh channel speeds in either direction between open systems and mainframe servers. After splintering large files into more manageable pieces, HARBOR File Transfer and RapidXchange software direct the data in multiple streams through Lightning 9900 V Series systems, making information available and accessible at mega rates. With TrueCopy software, the Hitachi Data Systems team of experts ensures continuous access to information while they carry out a client's plan to restructure operations or consolidate decentralized systems into an existing datacenter. Data Protection Services incorporates components of the Business Continuity Suite for applications with industry-leading Hitachi Freedom Storage systems. This combination of services, software, and storage provides a complete solution to help implement the processes involved in duplicating and transferring data to both local and remote locations.

The Data Movement Suite optimizes the transformation of information into knowledge.

The HiCommand™ Management Framework and Hitachi Resource Manager™ System Management Software

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 capabilities. 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 also deliver enterprise-wide coverage of online 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.

HiCommand Management Framework Allows Systems Management of Hitachi Freedom Storage and Software through the Enterprise's Vendor of Choice

Hitachi Freedom Storage software and hardware solutions are managed through the powerful HiCommand Management Framework, which was created to reduce operations expense and increase business agility while enabling operational excellence. Through use of standards-based application programming interfaces (APIs), Hitachi has ensured that any independent software or hardware vendor (ISV or IHV) can develop products that seamlessly integrate into the framework. The HiCommand Management Framework encompasses storage resource management, configuration management and automation, automated data replication and recovery, performance management and optimization, and related functionality. Beginning with the HiCommand Device Manager and HiCommand Tuning Manager modules, the framework will expand to provide a full suite of storage management applications.

HiCommand Device Manager, the first module of the HiCommand Management Framework, consists of three components: the Device Manager server, a Web-based graphical user interface (GUI) browser, and Device Manager host agents that access software management functionality. Most importantly, Device Manager software 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.

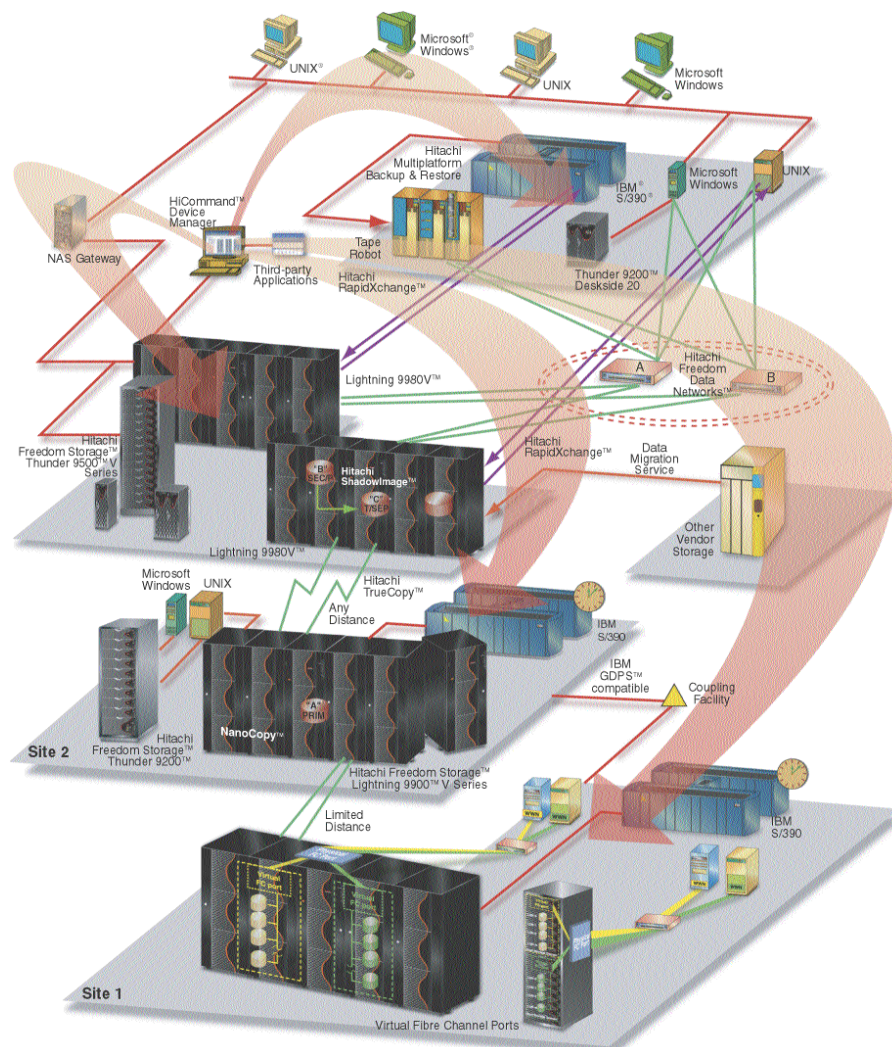
HiCommand Tuning Manager software maps, monitors, analyzes, and reviews an entire storage network's resources, all from a consolidated single screen. Intuitive and Web-based, Tuning Manager software provides the necessary data to improve the administration-to-storage resource ratio, maintain IT service- and operating-level agreements, employ predictive trending to forecast future storage requirements, generate storage reports for current usage analysis, and set real-time capacity and performance metrics. The ability of Device Manager software and Tuning Manager software to globally manage all Hitachi Freedom Storage hardware and software is shown in Figure 25.

HiCommand Management Framework Provides an Open Framework for Storage Management

With the advent of SAN and NAS, customers have greater flexibility and control over

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

Figure 25 – HiCommand Device Manager software and Tuning Manager software allows management of virtually all Hitachi Freedom Storage hardware and software from the customer's platform of choice.



their storage environments than they have ever had in the past. However, most storage managers are faced with administering a myriad of existing components and new, complex storage topologies, which lack a clear and consolidated approach to management. It is no longer good enough to provide basic services, such as LUN management, device monitoring and backup; now, it is critical to provide integrated management of all aspects of storage across a heterogeneous environment made up of multiple devices, such as SAN switches, complex storage systems, and storage appliances. Customers have also invested heavily in "point" solutions and systems management applications for their current environments. However, it is not sufficient for a vendor to provide a replacement for existing technologies; the vendor must provide a complete solution by allowing customers to leverage their existing investments, whether they are in hardware, software, or people.

Hitachi Data Systems recognizes that no single vendor can meet all the storage requirements for today's enterprises. That's why Hitachi Data Systems provides an open management platform, based on industry-standard management protocols, along with high-value-add functional components, such as policy-based automation and heterogeneous device management. This allows customers to incorporate multiple vendor offerings to create a truly complete solution for managing and monitoring their storage environments. HiCommand Management Framework supports a modular, building-block approach to designing and addressing business-specific storage management requirements. As a result, this framework enables customers to build the most suitable management environment

based on a combination of Hitachi products and best-of-breed storage hardware and software. Likewise, Device Manager software is customizable to integrate easily with existing infrastructures and accommodate future growth plans.

The Device Manager module simplifies management functions and enables integrated administration of multi-vendor storage environments. Furthermore, standard storage management operations come pre-loaded in the policy-based automation module. Device Manager also provides definition tools and an easy-to-use GUI interface for automating more complicated storage management processes, resulting in lowered TCO.

HiCommand Device Manager enables an enterprise to manage Hitachi Freedom Storage products and software solutions.

Open, Standards-based Approach

At the foundation of HiCommand Management Framework is a message bus (as shown in Figure 26) that allows individual functional components to "plug and play." Functional components are incorporated into the overall solution based on a usage of the Common Information Model (CIM), and the Simple Object Access Protocol (SOAP) provides a standard envelope for a messaging bus that enables simple plug-in of functional modules. This model provides a clear set of information content for management and monitoring of the entire storage and system environment. Using this object-model-based approach, solutions can readily communicate with all network-attached devices and perform a variety of configuration, administration, and other functions, allowing for a complete, integrated hardware and software management environment. Adopting this model is the key to overall interoperability in the storage and systems management environments. If a solution component adheres to this model, then its management or incorporation will be relatively straightforward.

Another industry standard, eXtensible Markup Language (XML), is fast becoming the norm for data communication within an integrated environment. Using XML to express the messages within the CIM model enables not only a common understanding of the required content, but also a clear and universal understanding of the message format and flexible application integration. Hitachi is creating an open and extensible architecture using these two foundational standards.

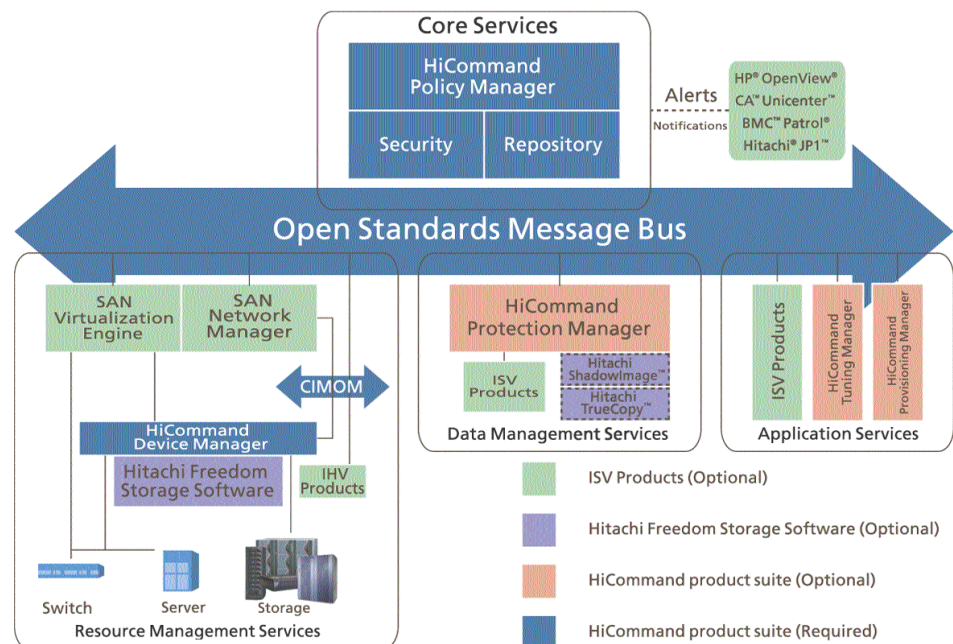
XML is fast becoming the norm for data communication within an integrated environment.

Hitachi is creating an open and extensible architecture using these three foundational standards. A more thorough discussion of the HiCommand Management Framework initiative and the open industry standards can be found at http://www.hds.com/pdf/PER1116-00_WhtPpr.pdf

CIM Benefits to Customers

CIM answers customers' demands for interoperability, creating more functional management applications and enabling customers and integrators to construct multi-vendor solutions more easily. Vendor adoption of CIM becomes a critical success factor for a product's inclusion in a heterogeneous storage management environment. Without CIM, vendors would need to swap and maintain an increasing tangle of APIs. CIM provides a

Figure 26 – The HiCommand Management Framework is based on the CIM and SOAP industry standards.



common model with which all vendors can work. A CIM-enabled storage management environment will deliver significant benefits to customers, including the ability to:

1. Overcome the Challenges Created by De-coupling Storage Resources from Specific Server(s) in a SAN

Neither a single vendor solution nor individual point solutions can meet customers' long-term comprehensive storage management and interoperability requirements. Customers need assurance that products deployed today can be expanded and enhanced to meet emerging interoperability requirements by adherence to open industry standards, such as CIM.

2. View and Operate Networked Storage as a Holistic Resource

CIM provides a common, normalized, yet extensible view of these resources, their relationships, and the operations that can be performed against them.

By spanning the management of all components in the storage stack, from the application to the physical storage medium, CIM allows a management application to be as focused or all-encompassing as it needs to be, all within a single standard.

3. Avoid Being "Locked" into a Proprietary Management Environment

Storage management vendors supporting heterogeneous, interoperable environments must differentiate their value based on products, services, and solutions, and not because their customers are locked into distinct management frameworks incompatible with other SAN elements. Customers who choose new and innovative companies that can create competitive or complementary functionality (built on existing management infrastructure and components based on the CIM standard) realize investment protection, greater choice, purchasing leverage, and additional options.

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. 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. Intelligent storage and network management software suites like the Hitachi Resource Manager software package, therefore, are at the heart of the enterprise's ability to achieve high levels of performance and availability to meet business objectives.

Hitachi Resource Manager Provides Powerful Device-level Storage Management Capability

As shown in Table 6, this comprehensive systems and software management offering brings together 11 powerful software solutions, with those shown in **bold** being standard.

	Product Support		O/S Support	
	Thunder 9500™ V Series	Lightning 9900™ V Series	Open	IBM® S/390®
Hitachi Resource Manager™				
Remote Console – Storage Manager		✓	✓	✓
Hitachi Graph-Track™		✓	✓	✓
SNMP Agent	✓	✓	✓	✓
Sun™ Java™ Agent	✓	✓	✓	✓
Host Storage Domains	✓	✓	✓	
Hitachi LUN Manager / LUN Expansion	✓	✓	✓	
Hitachi SANTinel™	✓	✓	✓	✓
Hitachi Virtual LUN / LUN Expansion	✓	✓	✓	
Hitachi VLVI Manager		✓		✓
Hitachi FlashAccess™	✓	✓	✓	✓
Cache Management Host Agent		✓		✓

Table 6 - Summary of Hitachi Resource Manager solutions and the products they support.

Shipped with every Hitachi Freedom Storage™ Thunder 9500™ V Series and Lightning 9900 V Series systems, these integrated software packages are used by storage managers to display system configuration, create user name/password security for administrators, set up RAID groups, allocate LUNs, expand LUNs, and format storage. Storage administrators who need a flexible, easy-to-use configuration and error-management tool will find Resource Manager software package is an ideal choice. Users can monitor and manage their storage systems through a graphical user interface.

The Resource Manager software package can also help optimize performance by providing valuable resource utilization information, such as I/O activity, cache usage, and availability status/event notification. Administrators can use Password Protection to authorize users. This provides protection against unauthorized access to the Management mode of Resource Manager software, and stores the password list at the array. Reliable error management and reporting offers a big breakthrough for IT troubleshooters who seek rapid-fire detection and notification to maintain productivity.

The following is a description of the Resource Manager modules and capabilities applicable to Thunder 9500 V Series and Lightning 9900 V Series systems as summarized in Table 6.

1. Out-of-band Management

Two LANs can be attached to the Thunder 9500 V Series systems: an internal LAN (private LAN), which is used to connect the service processors of multiple

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

Out-of-band management ensures data security, data accessibility, and data manageability.

systems, and the user's intranet (public LAN), which allows users to access the Resource Manager software package.

2. SNMP Support Ensures Compliance with All Major Systems-management Platforms

The Resource Manager software package 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 the Resource Manager software package, or the VERITAS® Storage Manager™, IBM Tivoli®, CA™ Unicenter® or HP® OpenView®. With CA Unicenter, “super MIBs” allows the control of the Lightning 9900 V Series or Thunder 9500 V Series systems using CA Unicenter commands. This out-of-band⁶ reporting format ensures that Resource Manager software can determine the status of your Hitachi Freedom Storage system even if the data channel path is down.

3. Java™ Agent

The Hitachi Web Utility runs on Web browsers, such as Internet Explorer® and Netscape® Navigator®. These programs run under the Microsoft® Windows and Sun™ Solaris™ operating systems and provide a user-friendly interface for the Lightning 9900 V Series and Thunder 9500 V Series systems Remote Console – Storage Navigator functions, which are all based on the use of Java agents.

4. Host Storage Domains

The Lightning 9900 V Series and Thunder 9500 V Series systems provide for “virtual” Fibre Channel ports that are logically managed by intelligent Fibre Channel controller cards. Each physical Fibre Channel port has multiple Host Storage Domains or HSDs. An HSD supports a “logical” Fibre Channel port (see Figure 27) each with its own set of logical units (LUNs). Hosts are matched to their assigned HSD based on a unique World Wide Names (WWNs) identifier. LUN security is provided by Hitachi SANTinel™ software, which fences access to LUNs based upon host WWNs. Under this innovative software structure, one LUN is reserved for the command and control device.

5. Hitachi LUN Manager and LUN Expansion Manager Software Simplifies Configuration Management and Reduces Staffing Costs

Hitachi LUN Manager software is an open systems management utility. With LUN Manager software, open systems LUNs can be defined, configured, and maintained. There is no more waiting for the hardware vendor to come and make configuration changes. LUN Manager software includes an easy-to-use, GUI-based interface that allows the definition of paths for LUNs, the reconfiguration of LUN-to-port assignments, or the viewing of the Hitachi Freedom Storage remote service information messages. Because LUN Manager software 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 systems by a dedicated LAN, LUN Manager software can support multiple systems.

Hitachi LUN Expansion Manager software features logical unit size expansion that dramatically improves LUN flexibility. Multiple LUNs can be presented to

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.




Figure 27 – Host Storage Domains reduce costs through fewer physical connections.

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 Hitachi Freedom Storage 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 these powerful management capabilities.

6. Hitachi SANTinel Software

SANTinel software controls host access to Hitachi Freedom Storage systems LUNs in open systems, multiplatform, or SAN environments. This enables users to restrict server access to only those LUNs for which they are authorized, as shown in Figure 28.

SANTinel software for open systems, running on the Hitachi Freedom Storage systems, allows administrators to define multiple HSDs for each physical Fibre Channel port on the system. An HSD consists of a virtual (logical) Fibre Channel port and an associated set of LUNs visible only via that logical Fibre Channel port.

Connections from host servers arriving at the physical Fibre Channel port are routed to the logical Fibre Channel port within the appropriate HSD based upon the Fibre Channel WWN of the host.

The storage administrator creates the HSD, sets the host mode of the logical Fibre Channel port within that HSD, and maps the LUNs to appear within that HSD. Then the storage administrator authorizes a particular host to "see" an HSD on that physical port by assigning the Fibre Channel WWN of that host to the HSD in question.

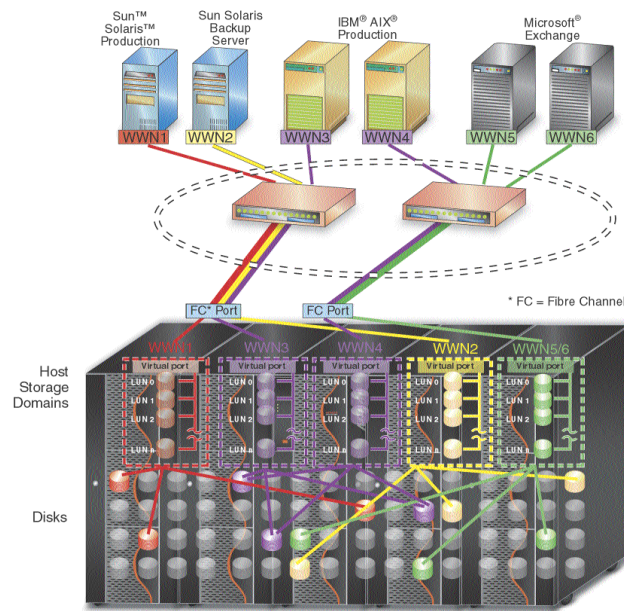
Because the WWN is used to route connections to the appropriate HSD, the

WWN may only be assigned to a single HSD on each Fibre Channel port.

SANTinel software may be used to authorize multiple hosts to access a particular HSD. SANTinel software may be further used to permit or deny a particular host to access individual LUNs within the HSD, based upon the WWN of the host.

SANTinel software may be used to create up to 128 HSDs per physical Fibre Channel port. Each HSD may contain up to 256 LUNs, with the maximum number of LUNs contained within all HSDs on a given Fibre Channel port reaching 512.

Figure 28 – SANTinel software uses pre-defined zones to protect data from unauthorized access.



7. Hitachi Virtual LUN and LUN Expansion Software

Virtual LUN software allows administrators to convert fixed-size volumes into several smaller, variable custom-sized volumes. Using the Resource Manager software, users can configure a specific number of megabytes to each custom logical volume image LUN. Virtual LUN software improves data access performance by reducing logical-device contention as well as host I/O-queue times, which can occur when several frequently accessed files are located on a single volume. Multiple LUN types can be configured within each RAID group. Virtual LUN software enables administrators to more fully utilize the physical storage capacity of the Hitachi Freedom Storage systems, while reducing the amount of administrative effort required to balance I/O workloads.

When Virtual LUN software is used in conjunction with FlashAccess software, the user can achieve even better data access performance than when either Virtual LUN software or FlashAccess software is used alone.

The LUN Expansion Manager allows administrators to create virtual LUNs that are larger by expanding the size of a selected LUN. The maximum size is 2TB.

8. Hitachi FlashAccess™ Software Allows Data to be “Locked and Unlocked” in Cache On-the-fly

The Hitachi FlashAccess software utility 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 FlashAccess software, a portion of cache memory can be allocated to specific data. Users can add, delete, or change FlashAccess software managed data at any time, quickly and easily.

Defined by the 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, FlashAccess software offers multiple modes of operation.

9. Remote Console – Storage Navigator

The Lightning 9900 V Series Remote Console – Storage Navigator is provided as a Sun Java® applet program, which can execute on any machine that supports a Java Virtual Machine (JVM). The Remote Console – Storage Navigator PC hosts the Remote Console – Storage Navigator Java-applet program and is attached to the Lightning 9900 V Series system(s) via a TCP/IP local-area network (LAN). When a Remote Console – Storage Navigator accesses and logs into the desired service processor (SVP), the Remote Console applet is downloaded from the SVP to the Remote and runs on the Web browser of the Remote Console PC. In this way, the Remote Console communicates with the attached Lightning 9900 V Series systems via a TCP/IP network.

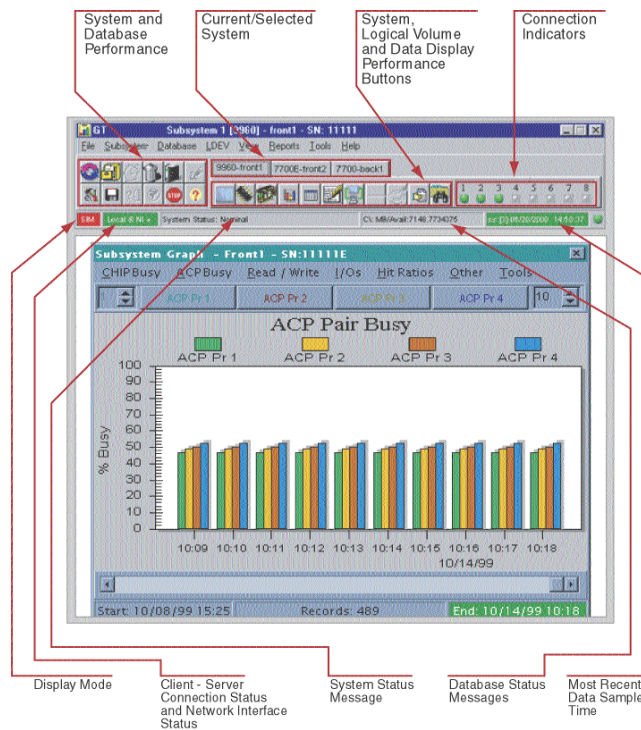
Two LANs can be attached to the Lightning 9900 V Series and Thunder 9500 V Series systems: an internal LAN (private LAN), which is used to connect the SVPs of multiple systems, and the user's intranet (public LAN), which allows you to access the Remote Console – Storage Navigator functions from individual Remote Console – Storage Navigator PCs. The Remote Console – Storage Navigator communicates directly with the SVP of each attached system to obtain system configuration and status information and send user-requested commands to the system. The Lightning 9900 V Series systems' Remote Console – Storage Navigator Java-applet program is downloaded to the Remote Console – Storage Navigator (Web client) from the SVP (Web server) each time the Remote Console – Storage Navigator is connected to the SVP. The Remote Console – Storage Navigator Java-applet program runs on Web browsers, such as Internet Explorer® and Netscape® Navigator®, which run under the Microsoft® Windows NT®, Windows® 2000, and Sun™ Solaris™ operating systems to provide a user-friendly interface for the Lightning 9900 V Series Remote Console – Storage Navigator functions.

10. Hitachi Graph-Track™ Software Takes the Guesswork Out of Performance Management

As the cornerstone of the software utilities of the Resource Manager 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 online help further enhance its ease of use. All Graph-Track data can be exported in several formats quickly and effortlessly, for use in other data analysis and reporting programs.

Graph-Track software gives users a more reliable and centralized way to manage performance. Available for S/390, UNIX, Windows NT, and Windows 2000 environments, this unique tool monitors hardware performance and supplies complete system-storage information at the touch of a button. The robust GUI-oriented command screen for Graph-Track is shown in Figure 29. Running on a PC attached to a dedicated LAN, Graph-Track software displays real-time or historical data for all connected Lightning 9900 V Series systems, helping users identify important peaks in utilization and ongoing trends in processing. With Graph-Track software, users can analyze discrete storage activities and determine the precise impact of each activity on system operations. It

Figure 29 – The GUI interface on Graph-Track software simplifies performance management.



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.

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

11. Hitachi Virtual Logical Volume Image Manager Software for S/390 Environments Allows Very Large to Very Small Virtual-volume Configuration

Hitachi Virtual Logical Volume Image (VLVI) Manager is a mainframe software utility that optimizes Lightning 9900 V Series system capacity utilization by allowing users to configure multiple *virtual LVs* in place of an 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 VLVI requires one physical address, with a maximum of 4096 addresses per Lightning 9900 V Series 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 VLVI Manager software 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.

12. Cache Manager Host Agent

The Cache Manager Host Agent enables mainframe users to perform FlashAccess software operations on S/390 LVIs by issuing commands from the S/390 host system to the Lightning 9900 V Series system.

Hitachi VLVI Manager software improves performance.

Optimizing Performance, Availability, and Cost with the HiReturn™ Investment Analysis Tool Kit

11

HiReturn™ Investment Analysis Tool Kit

Most enterprises today require that financial investment analysis be prepared prior to IT capital expenditures. Identifying hardware and software as "best-of-breed" is no longer sufficient justification for a major investment. Networked storage infrastructures and consolidated storage systems have generally appeared to be cost-effective IT strategies, but in order to commit to significant capital expenditures, today's decision-makers seek detailed, real-world information on operating expenses. With this in mind, Hitachi Data Systems provides the HiReturn™ investment analysis tool kit so that customers can make informed capital expenditure IT decisions.

HiReturn is a collection of return on investment (ROI) tools, methodologies, best practices, industry empirical data, and one-on-one coaching that allows customers to calculate the costs and benefits of implementing Hitachi Freedom Storage™ technology and solutions. With these advanced analytical methods, Hitachi Data Systems delivers a pre-sales, no-cost service that provides the objective validation needed for the storage system and solution capital expenditure decision.

Cold, hard facts create a solid empirical basis for financial purchase decisions involving Hitachi Freedom Storage networked infrastructure consolidation solutions. That's why the HiReturn tool kit can play such a valuable role in the effort to simplify, protect, and, optimize the storage environment that supports the business objectives of the enterprise.

Since 1997, Hitachi Data Systems has documented 29 ROI conditions, or "use cases," that demonstrate hard- and soft-dollar savings. Conservative, nominal, and aggressive modeling parameters are discussed within these "use cases." These cases and the ROI calculation methods behind them are based on a proven cost-determination methodology.⁷

*HiReturn™ quantifies
ROI for Hitachi
Freedom Storage™
solutions and systems*

*The HiReturn tool kit
helps develop a solid
business case*

The HiReturn Methodology Is Implemented in Three Phases:

Phase 1 – Preparation.

A conference call or brief on-site meeting is conducted prior to the engagement to review the approach, goals, and scope.

Phase 2 – On-site, One-day Assessment.

Delivery is provided in three modules:

Infrastructure review. Includes collection of facts on the site's current storage infrastructure, financial data, and IT requirements.

Assimilation. The HiReturn team reviews the information, formulates recommendations, and develops a strategic plan.

Findings. The customer and the HiReturn team confirm and corroborate collected data. Recommendations, financial impact, and strategic plans are reviewed and discussed.

⁷ To download our white paper covering the 29 use cases and cost-determination methods for a SAN Return on Investment, please go to http://www.hds.com/pdf/3SAN_ROI_Wht_Ppr.pdf

Phase 3 – Final report. The HiReturn team provides a business Financial Impact Report that summarizes the ROI findings of the assessment and outlines the suggested strategy, along with other key recommendations.

The On-site Assessment Produces the Needed Facts to Build a Business Case

The HiReturn analysis is conducted at the customer's site. Hitachi Data Systems consultants spend time prior to the assessment gathering data and facilitating cost and technology workshops to capture and model the customer's current and future IT environment. The following elements need to be defined for use in the ROI determination process:

- Cost of fully burdened labor rates for storage administrators
- Depreciation rate, Internal Rate of Return and/or lease terms
- Cost of data center space (\$/sq. ft or m²)
- Corporate tax rate (marginal tax rate)
- Current data center floor space with current storage systems and equipment
- Cost of power (\$/kWh)
- Power and BTU ratings of current storage systems
- Estimates of \$/GB and annual storage erosion rate
- Number and average cost of servers (typically Microsoft® Windows NT® servers) used exclusively for storage access (NFS, CIFS mounted servers)
- Hourly cost of scheduled and unscheduled downtime (opportunity cost/loss)
- Current hardware maintenance costs for storage networks and elements
- Current software maintenance/license costs for storage networks and elements

To arrange a no-cost HiReturn analysis, please contact your Hitachi Data Systems representative.

Results and Reports Quantify Both IT Benefits and Business Benefits

As a result of the on site meetings and data collection, the HiReturn team produces a Financial Impact Report that summarizes ROI assessment findings. Overall client assumptions, ROI tool output, and business and technical conclusions are provided. The report may also introduce alternative technology and solutions that are likely to produce a superior ROI result and meet enterprise business requirements. The ROI impact is defined in hard- and soft-dollar savings that can be used by you to rationalize and justify future capital expenditures.

Global Solutions Services

12

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 Support Services and Global Solution Services organizations to ensure that Hitachi products operate at peak performance to complement all hardware and software in the enterprise.

Global Solutions Services Overview

The Hitachi Data Systems Solutions Services consultants offer a powerful, one-stop source for storage services to help customers simplify enterprise management. Highly experienced consultants work closely with data center personnel to define a total solution that makes the most of existing IT investments and provides for a clean transition to new computing environments.

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. Hitachi built-in product strengths are bolstered by the Global Solution Services group to ensure the optimal operation of hardware, software, and middleware for the enterprise.

Global Solutions Services specializes in three basic categories of service solutions 1) Storage Management Solutions, 2) Data Protection Solutions and 3) Storage Infrastructure Solutions. Each of the services in these three categories focus on solving customer business problems and providing project ROI whether it is through streamlining operations, reducing costs, or ensuring system interoperability. The Global Solutions Services group 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 service offerings are highlighted here.

***Hitachi Data Systems
Global Solutions Services
provides comprehensive
connectivity,
management, and
availability capabilities.***

Storage Management Solutions

The HiCommand™ Device Manager Enablement Service

The HiCommand™ Device Manager Enablement Service (Enablement Service) is designed to help users understand the basic concepts pertaining to the management of Hitachi Freedom Storage™ systems using Device Manager software. In a workshop approach, the Enablement Service provides the basic skills required to install and configure the Device Manager server, Web-based GUI browser, and HiScan™ host agents. In addition, demonstrations of basic storage configuration and management tasks are performed in order to illustrate Device Manager functionality.

Basic Implementation Service for InterSAN® PATHLINE®

The Basic Implementation Service for InterSAN® PATHLINE® provides planning, training, and installation assistance to get PATHLINE storage area management software configured and working in production environments. Hitachi Data Systems

Global Solutions Services team start by uncovering the complexity of an individual organization's storage environment in order to define a baseline deployment. Using baseline objectives, the consultants then plan and implement a PATHLINE management environment. In the course of the engagement, product overview training is provided and training of key administrators on the newly customized PATHLINE configuration concludes the service.

Data Protection Solutions

Data Replication Planning and Design Service

The Data Replication Planning and Design Service provides a high-level design of a distance-replication environment, as well as a detailed analysis of workload and performance characteristics that help support expensive bandwidth decisions. Using Hitachi Data Systems data replication best practices, the Global Solutions Services team produces a detailed study of the environment's current workload to recommend bandwidth necessary to support a current and future distance-data-replication environment.

Implementation Service for Hitachi TrueCopy™ Software

Implementation Service for Hitachi TrueCopy™ Software provides technical assistance with configuration and control of Hitachi Freedom Storage™ system-based, remote data mirroring TrueCopy software. The service is available in both open systems and mainframe environments and assists with TrueCopy configurations and basic control for data volumes already identified for remote replication.

Integration Service for Hitachi TrueCopy™ Software

Integration Service for Hitachi TrueCopy™ Software provides assistance for configuring and integrating Hitachi Freedom Storage™ controller-based, remote data mirroring TrueCopy software to support customer data replication requirements, including backup and recovery, testing and development, business continuity, and disaster recovery. The service is available in both open systems and mainframe environments and is individually tailored to address unique customer requirements.

Implementation Service for Hitachi ShadowImage™ Software

Implementation Service for Hitachi ShadowImage™ Software provides assistance with the configuration and operational automation of Hitachi Freedom Storage™ system-based, data mirroring software in both open systems and mainframe environments. Hitachi Data Systems Global Solutions Services assist customers with the planning and execution of ShadowImage software volume configurations and scripted automation of split-mirror backup operations.

Integration Service for Hitachi ShadowImage™ Software

Integration Service for Hitachi ShadowImage™ Software tailors the configuration and integration of Hitachi Freedom Storage™ system-based, data mirroring software with customer backup and recovery application requirements in both open systems and mainframe environments. Hitachi Data Systems Global Solutions Services assist customers with the planning and execution of ShadowImage software volume configuration, scripted automation of split-mirror backup operations, and the integration of split-mirror operations with explicit application quiesce and recovery mechanisms.

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

Hitachi SplitSecond™ Solutions Rapid Recovery Service for Microsoft® Exchange

The Hitachi SplitSecond™ Solutions Rapid Recovery Service for Microsoft® Exchange employs hardware and software in a service package that is designed to solve business problems that can be avoided by the rapid recovery for mission-critical applications or databases in the event of a failure. SplitSecond Solutions software utilizes the point-in-time (PiT) disk copy capability of Hitachi ShadowImage™ software and Command Control Interface (CCI) scripts. By incorporating proven technology, ShadowImage software, and the Hitachi Data Systems Global Solutions Services experience and knowledge of operating systems and applications, the SplitSecond Solutions for Exchange service package delivers a solution that integrates Hitachi Freedom Storage™ Lightning 9900™ V Series and Lightning 9900 Series systems with customer applications.

Hitachi SplitSecond Solutions® Rapid Recovery Service for Microsoft® SQL Server®

Hitachi SplitSecond™ Solutions Rapid Recovery Service for Microsoft® SQL Server™ employs hardware and software in a service package that is designed to solve business problems that can be avoided by the rapid recovery of mission-critical applications and databases in the event of a failure. SplitSecond Solutions software uses the point-in-time (PiT) copy capability of Hitachi ShadowImage™ software and Command Control Interface (CCI) scripts. By incorporating proven technology, ShadowImage software and the Hitachi Data systems Global Solutions Services experience and knowledge of operating systems and applications, the SplitSecond Solutions for SQL Server service package delivers a solution that integrates the Hitachi Freedom Storage™ Lightning 9900™ V Series and Lightning 9900 Series systems with customers applications.

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

Storage Infrastructure Solutions**Configuration Management Service**

The Configuration Management Service is designed to help management and staff improve the documentation and management a data center's raised floor environment. The Configuration Management Service increases the usability and accuracy of configuration information through a wide variety of on-line and printed reports which improve the access of information by data center support personnel. Through this service, Hitachi Data Systems provides the powerful OBTAIN 24/7™ database system that is designed to document IBM® S/390® and open systems equipment and configurations, fiber optic infrastructures, as well as other power-related equipment. The built-in strengths of Hitachi Data System's products are enhanced further by the Global Solution Services group to ensure the optimal operation of the customer's enterprise architecture.

Fiber Design and Implementation Service

The Fiber Design and Implementation Service is designed to help customers implement and/or manage data center fiber optic infrastructures for ESCON®, FICON™, open systems, and SANs. This service can be used for complex efforts, such as the implementation of a new fiber optic infrastructure, or for more simple and frequent tasks, such as periodic change-add-deletes to an existing fiber optic infrastructure. The Fiber Design and Implementation Service includes infrastructure analysis and design, provisioning of fiber optic trunks, patch jumpers, patch panels, cabinets, conveyance trays, as well as other necessary equipment, documentation, and maintenance services. Through this service, Hitachi Data Systems also provides copper-based connectivity solutions for open systems and SAN environments.

Infrastructure Planning Service for FICON™

The Hitachi Data Systems Infrastructure Planning Service for FICON™ is an on-site, implementation planning and training start up service designed to accelerate the installation of a FICON solution. This service also helps protect current data center infrastructure investments by leveraging the current, physical environment and by projecting long-term growth requirements associated with the FICON migration. Hitachi Data Systems will help identify critical customer concerns, provide answers to technology questions, assist with budgetary estimates and make FICON infrastructure design recommendations. Once the service is complete, customers will not only have the technical information necessary to implement a FICON infrastructure, but will also have a complete understanding of the associated infrastructure costs related to the migration.

Storage Consolidation Planning and Design Service

The Storage Consolidation Planning and Design Service provides analysis of existing storage environments, performance characteristics, and capacity requirements. The analysis is used to create the implementation plan for migrations to a consolidated storage and storage area network (SAN) environment. The objective of this service is to ensure successful planning and design of the consolidated storage environment.

Storage Consolidation Implementation Service

The Storage Consolidation Implementation Service provides storage consolidation implementation services in environments using Hitachi Freedom Storage™ systems and software. The objective of this service is to ensure a successful storage consolidation of customer data according to set plans created in accordance with the Storage Consolidation Planning and Design Service.

Open Systems Data Migration Planning Service

The Open Systems Data Migration Service takes an inventory of a customer's existing storage environment and matches this to the customer's capacity plan for Hitachi Freedom Storage™ systems. This analysis is then used to create the plan for the migration of the customer's existing data to Hitachi Freedom Storage Lightning 9900™ V Series or Thunder 9500™ V Series systems. The objective of this service is the successful planning of the customer's data migration to Hitachi Freedom Storage systems.

Open Systems Data Migration Implementation Service

The Open Systems Data Migration Implementation Service provides for the actual migration of existing customer data onto Hitachi Freedom Storage™ systems according to the course of action developed during this service's prerequisite - the Open Systems Data Migration Planning Service. The objective of this service is to ensure a successful migration of customer data from the customer's current systems to a Hitachi Freedom Storage environment.

SAN Services

Hitachi Data Systems offers an entire suite of infrastructure solutions that will assist clients in planning and implementing SANs to optimize the management and control of data across the enterprise. The modular services include:

- SAN Assessment Service
- SAN Design Service
- Basic SAN Implementation Service
- Advanced SAN Implementation Service
- SAN Health Check Service

The Global Solutions Services group has a highly experienced team of SAN specialists to aid enterprises in implementing SANs and maintaining them in optimal condition.

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 Interface or API is a set of calls that allow software developers to interface to a specific program.

Array Frame

Array frames are the (up to three) left and (up to three) right cabinets of the Hitachi Freedom Storage™ Lightning 9980V™ systems housing HDDs (high density disks), power supplies, and batteries.

Array Group

The physical arrangement of disk drives independent of RAID level used.

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 Remote Copy and 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.

Byte

A unit of data that is eight binary digits long and represents one text character.

Back End

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

Bit

Eight bits compose a byte. This term is spelled out in compound abbreviations (Mbit/sec and Gbit/sec) to avoid confusion with abbreviation for byte - B (as in MB or GB).

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. A 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.

CHA

Channel Host Adapter.

CHIP

Channel 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 processor 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.

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 Hitachi Freedom Storage™ Lightning 9980V™ system, housing the channel switch, cache, power supplies, batteries, and host interface cards.

CPU

CPU is an acronym for Central Processing Unit. It can refer to either a processor chip, such as the Sun™ SPARC™ or the Intel® 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 Hitachi Freedom Storage™ Lightning 9900™ V Series. These 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.

DAS

Direct Attached Storage.

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 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

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. It also indicates 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.

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 10Mbit/sec. 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.

eXtended SCSI Copy (e-Copy)

Extended SCSI is a command in the SCSI 3 specification that allows for the transfer of data without the involvement of the server. The Hitachi Freedom Storage™ Lightning 9900™ V Series implements this command, which is used by programs such as VERITAS® NetBackup™ to enable serverless and LANfree backup.

Fabric

A fabric is one of three Fibre Channel topologies. In a Fabric topology, Node Ports (N_Ports) are connected to Fabric Ports (F_Ports) on a switch. See also Switch.

Failover

Host, host bus adapter, 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 100Mbit/sec 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 Adapter.

FC-AL

Fibre Channel Arbitrated Loop is the most dominant of the three topologies of Fibre Channel. Loops are a cost-effective way of connecting up to 127 ports in a network without the need for a switch. See also Fabric and Switch.

Freedom Data Networks™

Hitachi Data Systems groups its storage area network (SAN) and network attached storage (NAS) solutions under the label Freedom Data Networks. Freedom Data Networks provides an open architecture that leverages SAN and NAS technology and offers organizations freedom of choice in deploying data-access and data-sharing capabilities across the enterprise through Hitachi Data Systems advanced implementations: Hitachi Freedom SAN™ and Hitachi Freedom NAS™. With Freedom Data Networks, 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 among 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 Networking Industry Association (SNIA). The technology has attracted interest as a channel for the attachment of storage devices, but it has a limited popularity as a 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™

FICON is an IBM® trademark for Fibre Connection technology (ESCON® over Fibre Channel).

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 Hitachi Freedom Storage™ Lightning 9900™ V Series systems, the front end consists of CHIP pairs.

GB

Abbreviation for gigabyte.

Gigabit Ethernet

Provides a standard that supports data transfer at 1000Mbit/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 early 2003.

GUI

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

HARBOR® File Level Backup/Recovery with Lightning 9900™ V Series Agent

Harbor File Level Backup/Recovery (HBR) is a software utility offered by Tantia Technologies® that provides for network or channel-based backup/restore of open systems client files to an IBM® OS/390® host using a common graphical user interface. Online, nondisruptive backup/recovery of popular database, mail, and enterprise resource planning (ERP) applications are optional. Client support is available for Microsoft® Windows®, IBM OS/2®, UNIX®, Novell® NetWare®, and HP® 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. It 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 Sun™ Java™ client interface, a system monitor, scheduling support, and remote control of client functions from OS/390.

HDD

High-density disks. Used in the Hitachi Freedom Storage™ Lightning 9900™ V Series systems.

HDmS

The Hitachi Data Systems Migration Service (HDmS) is a Hitachi Data Systems Global Solutions 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 V Series storage system and provides load balancing to maintain predetermined performance levels.

Hitachi Dynamic Link Manager™

Dynamic Link Manager is a server-based family of software utilities that 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 V Series system by adding redundant SCSI connections between data servers and RAID systems. Dynamic Link Manager therefore provides increased reliability and performance. Supported platforms include IBM AIX®, Sun Solaris™, and Windows NT®/Windows 2000.

Hitachi FlashAccess™

FlashAccess software 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 Hitachi Freedom Storage™ system. The FlashAccess feature in the Hitachi Freedom Storage™ Lightning 9900™ V Series systems can be used for either IBM® S/390® or open systems. Hitachi FlashAccess is a software utility in the Hitachi Resource Manager™ suite that allows the creation, deletion, and monitoring of data managed by the FlashAccess software. See also Resource Manager.

Hitachi Graph-Track™

Graph-Track is a software utility in the Hitachi Resource Manager suite that enables a robust set of system and network management utilities and provides graphical reports for Lightning 9900 V Series performance, availability, and configuration management.

Hi-Star™ Architecture

At the heart of the Lightning 9900 V Series revolutionary design is Hi-Star architecture, which provides multiple, redundant, non-blocking paths between the storage ports, multiple cache nodes, and multiple disk Array Control Processors (ACPs).

Hitachi LUN Manager

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

Hitachi Multiplatform Backup/Restore

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

Hitachi Multiplatform Resource Sharing

Multiplatform Resource Sharing is built right into the Lightning 9900 V Series systems. With Multiplatform Resource Sharing, Lightning 9900 V Series systems users can share resources between UNIX® or Microsoft® Windows NT®/Windows® 2000 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 Priority Access

Priority Access allocates bandwidth on the basis of quality of service (QoS) requirements either at the physical port level or at the Host Storage Domain level within a physical port. See also Host Storage Domains.

Hitachi Rapid Recovery

The Hitachi Rapid Recovery solution for IBM® S/390®-attached storage combines NanoCopy™, Hitachi TrueCopy™ Asynchronous for S/390, 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 Hitachi Freedom Storage™ Lightning 9900™ V Series systems with NanoCopy capability to recover their critical data very quickly following an outage.

Hitachi RapidXchange™

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

Hitachi Resource Manager™

Resource Manager is a comprehensive package 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™

SANtinel software controls host access to Hitachi Freedom Storage (including Lightning 9900 V Series systems) LUNs in open systems or SAN environments.

Hitachi ShadowImage™

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™ V Series system, or up to 20 across multiple Lightning 9900™ V Series 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 asynchronous for S/390 to provide additional copies in another system.

Hitachi TrueCopy™

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 IBM MVS®, HP/UX, AIX, Sun Solaris, Digital UNIX, Sequent DYNIX/ptx, SGI IRIX, NCR, UNIX SVR4, and Windows NT/Windows 2000.

Host Storage Domain

One or more Host Storage Domains may be defined on each Fibre Channel port on the Lightning 9900 V Series. Each Domain has its own logical (or virtual) Fibre Channel port and hosts, which are matched to their assigned Domain based on their WWNs. See also LUN, Fibre Channel, and WWN.

HSD

See Host Storage Domain (HSD).

HSN

Hierarchical Star Network. The technical term used to describe the combined internal networks (the C-HSN and the CM-HSN) of the Hitachi Freedom Storage™ Lightning 9900™ V 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 IBM® XRC® compatible, host-based software offered by Hitachi Data Systems for 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®/Windows® 2000 registry.

InterSan® Pathline®

InterSAN Inc.'s Storage Area Management (SAM) software product PATHLINE is the first application-based SAM software that aligns the management of storage networks with business policies and objectives. PATHLINE helps customers efficiently manage large, multiprotocol, multivendor storage area networks (SANs) by taking an applications-oriented view that masks the complexity of the underlying storage-networking infrastructure.

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.

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 fiber optic backbones, and radio waves. A system of LANs connected in this way is called a wide-area network (WAN).

Lightning 9900™ V Series

The Hitachi Freedom Storage™ Lightning 9900™ V Series was announced in May 2002 (Lightning 9980V™ and Lightning 9970V™ models). 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 Fibre Channel Port

The Lightning 9900 V Series Fibre Channel adapter cards each contain four physical Fibre Channel ports. Each physical port can be subdivided into multiple logical Fibre Channel ports. Each logical Fibre Channel port supports a Host Storage Domain (HSD), which is a collection of LUNs matched to assigned hosts based on WWNs.

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

LPAR or 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 V Series 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 V 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.

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.

Metadata

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

MPLF

The Lightning 9900™ V Series 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 Freedom Storage™ Lightning 9900™ V Series 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™, Hitachi ShadowImage™.

NAS

Network Attached Storage or NAS servers are a special class of server that allows files to be stored over networks using the UNIX® or Microsoft® Windows® remote file system standards.

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.

Obtain 24/7™

Hitachi Data Systems is a licensed reseller for OBTAIN 24/7 software, which is owned and supported by Knowledge Flow Corporation. OBTAIN 24/7 software is a powerful database system designed to document IBM® S/390® and open systems equipment and configurations, fiber optic infrastructures, as well as other power-related equipment.

OLTP

Online Transaction Processing.

Online System

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

Open System

A system with characteristics that comply with standards made available throughout the industry, which 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 Hitachi Freedom Storage™ Lightning 9900™ V 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.

QoS

Quality of Service refers to service level agreement (SLAs) between users and the IT organization. See also SLA.

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 Hitachi Freedom Storage™ Lightning 9900™ V Series algorithms enable performance from RAID-5 that is competitive with some vendors' RAID-1.

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 Hitachi SAN™ and falls under the larger Hitachi Freedom Data Networks™ umbrella. Major benefits of SANs include outboard backup, sharing of resources, pooling, and reduced cost of storage management. SANs 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 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.

Serverless Backup

Using the eXtended SCSI command, and products like VERITAS® NetBackup™, the Hitachi Freedom Storage™ Lightning 9900™ V Series offers serverless and LANfree backup.

SLA

Service level agreements or SLAs are agreements regarding level of service between user departments and the IT department. SLAs refer to all aspects of IT service, including availability, performance, and repair.

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

The Sun™ UNIX® operating system based on System V, release 4.

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 V Series.

SWAN

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

Switch

A switch is a network device that examines and forwards packets between LAN segments.

Synchronous

Synchronous communications occur 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 those 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 Remote Copy and Hitachi TrueCopy™.

TCO

Total cost of ownership or TCO is a computer industry financial method of identifying the cost of operating computer equipment. TCO includes depreciated capital costs, manpower expense, power costs, communication costs, overhead, etc.

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 Microsoft® Windows® platforms.

Virtual Logical Volume Image Manager

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

Volume

An IBM® 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 fiber optic backbones.

Workload

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

WWN

World Wide Name(s) or WWN(s) refer to an eight-byte identifier assigned to each product that can be used as a port on a Fibre Channel network. The WWN is stored in nonvolatile memory and is frequently stamped on the surface of the product or used as a serial number. It applies to all HBAs, switches, or storage controller cards that interface to a Fibre Channel network.

XRC

Extended Remote Copy. The IBM implementation of a software asynchronous remote copy technique that preserves data integrity. See also HXRC and Hitachi TrueCopy™.

zOS™

IBM's latest mainframe operating system.

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

Suite 3301-6, Shell Tower
Times Square, 1 Matheson Street
Causeway Bay
Hong Kong
2525-2385
infoasia@hds.com

Australia/New Zealand Headquarters

Hitachi Data Systems
Level 3
82 Waterloo Rd.
North Ryde, NSW 2113
Australia
+ 61-2-9325-3300
info.australia@hds.com

Canada Headquarters

2550 Victoria Park Avenue
Suite 601
Toronto, Ontario M2J 5A9
Canada
(416) 494-4114
www.hds.com
Canada.Sales@hds.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
+44 (0) 1753-618000
info.eu@hds.com
www.eu.hds.com

Latin America Headquarters

750 Central Expressway, MS 3268
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
infolatin@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 a service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

Hi-Track is a registered service mark and Freedom Storage, Lightning 9900, HiCommand, Resource Manager, ShadowImage, Dynamic Link Manager, RapidXchange, SANtinel, Graph-Track, FlashAccess, NanoCopy, HiReturn, Thunder 9500, Lightning 9970V, Lightning 9980V, TrueCopy, TrueNorth, CruiseControl, SplitSecond, HiScan and Hi-Star are trademarks of Hitachi Data Systems Corporation.

All other trademarks, trade names, and service marks used herein are the property of their respective owners.

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/software/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.

©2003, Hitachi Data Systems Corporation
DISK-437-02 January 2003