Hitachi Thunder 9200™ Series
Architecture Guide

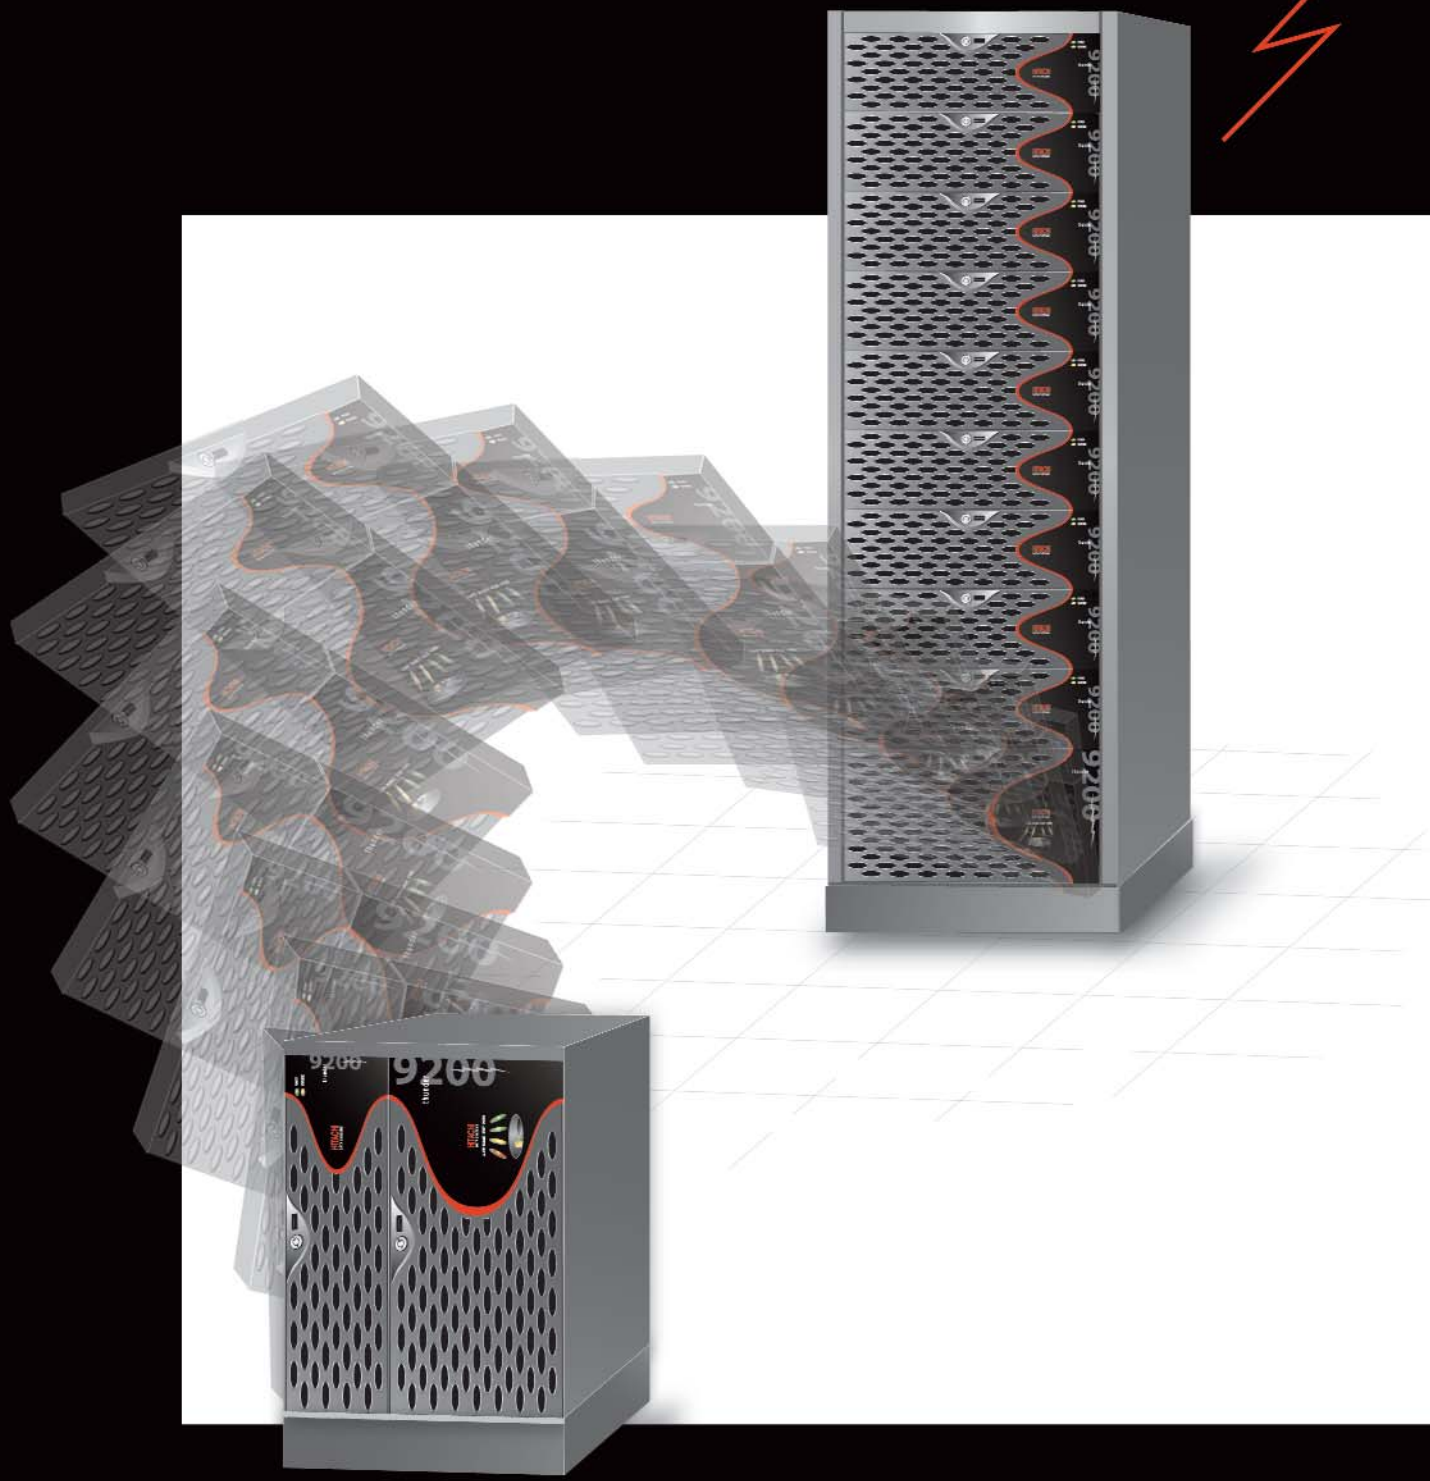

Enterprise-level features with a Midrange Package and Price

HITACHI
DATA SYSTEMS

Hitachi Thunder 9200™ Series
Architecture Guide

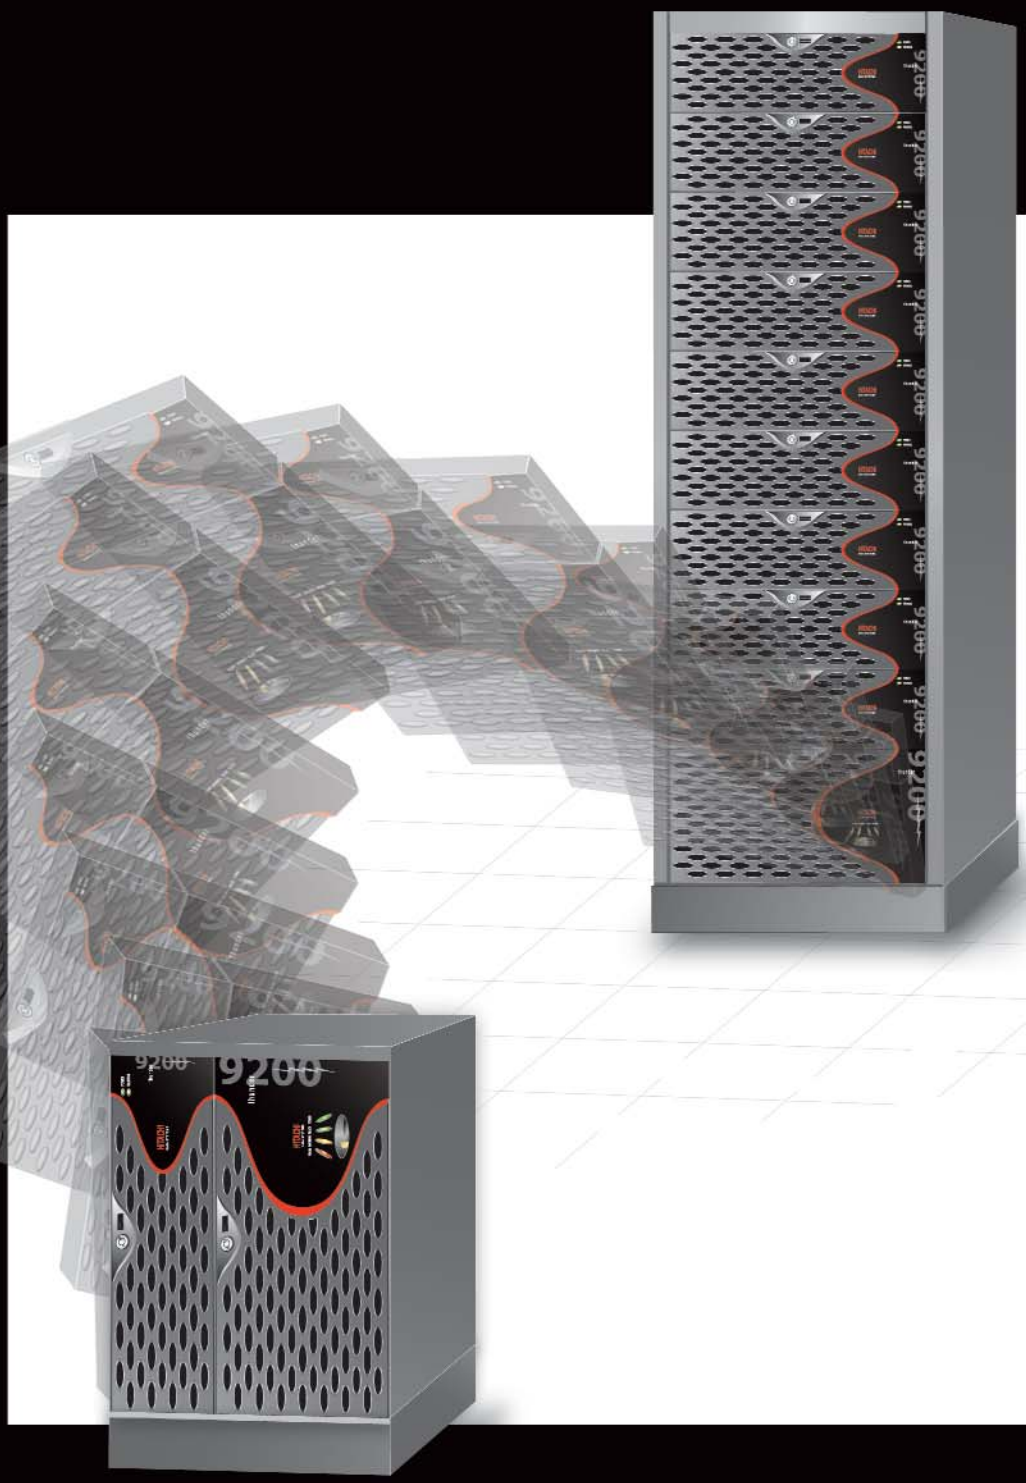

Enterprise-level features with a Midrange Package and Price

HITACHI
DATA SYSTEMS

Table of Contents

Table of Figures

Tables

Is Hitachi Freedom Storage™ Thunder 9200™ Right for Your Business?

1

The Hitachi Freedom Information Infrastructure "Bullet Proofs" Data

Business data must be continuously available AND always protected from loss.

There is little doubt that information and information security will continue to be at the center of the global economy for the foreseeable future. To ensure continued viability and success, both large and small businesses must be able to creatively and cost effectively capture, access, make sense of, and take action on information from the front office to the back office. Not only must information be continuously available but also it must be protected from loss.

To achieve these goals in managing critical business information, Hitachi Data Systems has created an Intelligent Information Infrastructure. This Infrastructure allows businesses of all sizes to build world-class profit oriented solutions on top of a "bullet proof" vault of business data that is both continuously available AND always protected from loss. To achieve this goal, Hitachi designed the competitively advanced "all Fibre" Hitachi Freedom Storage™ 9000™ family of products and wrapped them in advanced function software, SAN solutions, and other advanced solutions and services

There are four design goals for the Hitachi Intelligent Information Infrastructure. These goals are illustrated in Figure 1.

1. **Accelerate Information Delivery through Connectivity**

   Information should be made available to the entire enterprise and beyond, reaching employees, partners, suppliers, and customers. To achieve optimum performance and corporate agility, the information store must be central to IT strategy and infrastructure design.

2. **Exceed Availability Service Levels**

   Continuous availability of data must be "designed-in" to an IT infrastructure. This is because, however resilient the hardware is, the complexity of the software layers means that interruptions in data availability are inevitable. The capability to recover swiftly is not merely a goal – it is a commercial imperative.

3. **Reduce Complexity and Management Costs through Consolidation**

   Since all future IT architectures will be centered on information, information must be logically centralized and physically consolidated as far as is practical. The enterprise IT infrastructure must be powerful, scalable, available, flexible, and resilient. This can only be achieved through infrastructure design.

4. **Aggressively Scale Storage Capability**

   Ability to scale is central to business agility.

   The ability to scale up as a business grows or to re-deploy computing resources as business changes is central to an effective competitive posture for the enterprise. Hitachi's Intelligent Information Infrastructure provides customers, partners and employees with data when and where they need it. The advanced Hitachi architecture also provides an adaptable foundation upon which new applications can be constantly added and adapted thereby greatly reducing the risk of IT limiting business agility, performance, and success.

With Hitachi Data Systems® Intelligent Information Infrastructure world class architecture all four design objectives are met for an optimal and totally secure information-centric foundation on which to build.

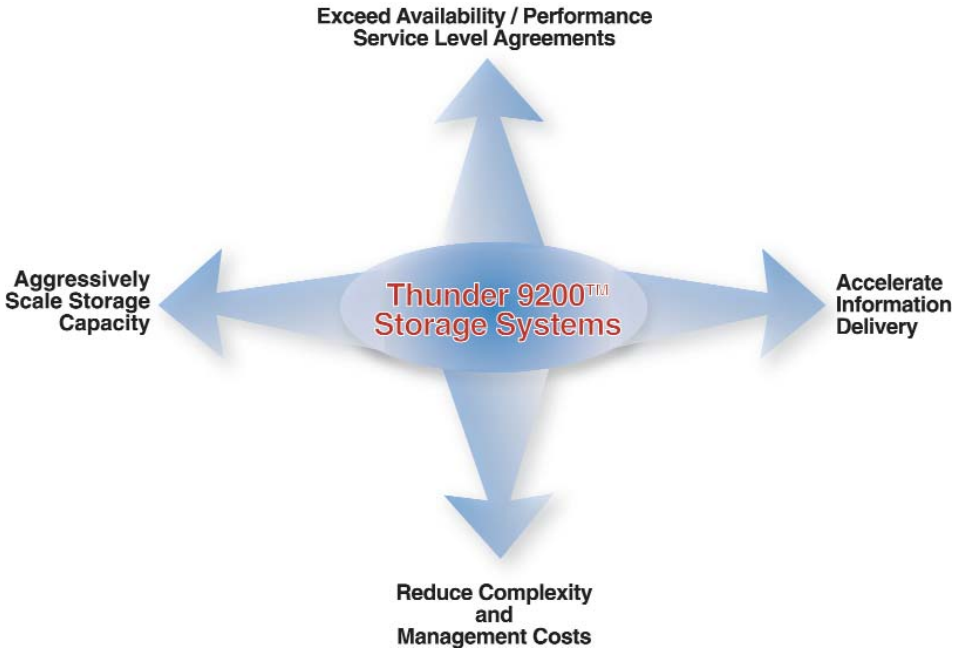

Figure 1 – The Four Dimensions of Excellence.

Hitachi Freedom Storage™ Products are at the Center of the Freedom Information Infrastructure

Just like information is at the center of our customer's business, the Hitachi Freedom Storage™ 9000 "all Fibre" family of products is at the center of this IT Architecture. The Hitachi Freedom Storage™ 9000 products are divided into two Series, the Lightning 9900™ Series and the Thunder 9200™ storage systems. In general the Lightning 9900™ Series is positioned for the high capacity set of rich advanced function software higher cost per megabyte "Premium" segment of the storage market. The Thunder 9200™ is aimed at the "Value" segment and is the only midrange product that truly delivers significant enterprise-level features with a midrange package and price. This is illustrated in Figure 2.

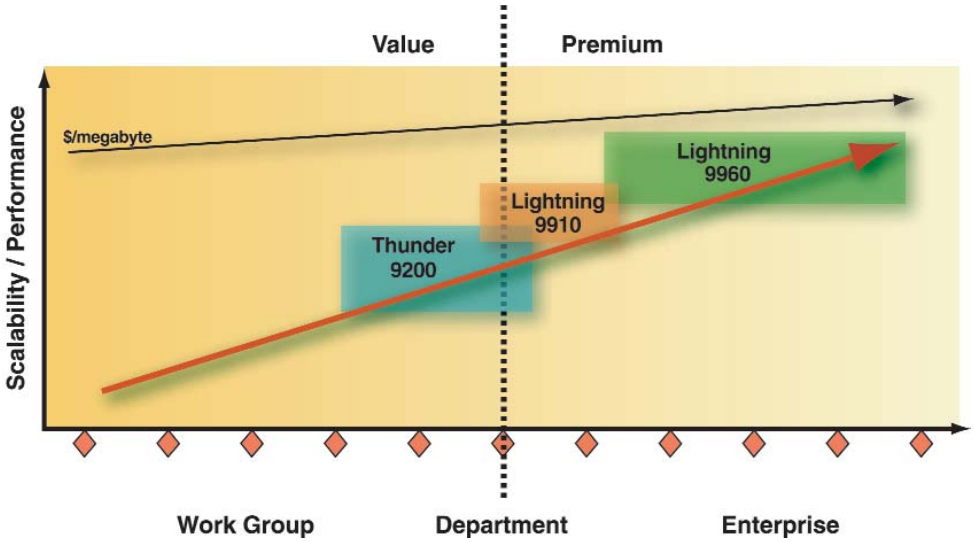

Figure 2 – Hitachi Freedom Storage™ products span all storage market segments.

Lightning 9900™ Series

Lightning 9900™ is the most powerful enterprise storage system in the industry.

The Lightning 9900™ Series intelligent storage system is the most powerful enterprise storage system in the industry today. The new and unique "internally switched" architecture used in the Lightning 9900™ Series is specifically designed for the high I/O and connectivity demands high-end open systems of S/390® and environments as well as SAN. Extremely high internal bandwidths, high-speed back-end design, scalable internal pathways, increased processor speeds, and increased number of processors, larger cache sizes, and a new and improved high-performance RAID-1 (called RAID-1+) sets the Lightning 9900™ Series apart from all other storage systems.

The Lightning 9900™ Series is the best choice for any of four circumstances:

- When S/390® connectivity or data sharing of open systems platforms and S/390® platforms is needed now or in the future.
- When more than four front-end Fibre Channel or two front-end SCSI attachments are needed.
- When storage capacity requirements are greater than 17.6 TB for high density 180 GB 7,200 rpm disks or greater than 7.2 TB for high performance 72 GB 10,000 rpm disks.
- When any advanced software not supported by the Thunder 9200™ is needed. (See also Table 2 and Chapter 2).

Thunder 9200™

Thunder 9200™ is the only midrange product that delievers enterprise-level features with a midrange package and price.

The Thunder 9200™ Series is a midrange RAID storage system designed for use in heterogeneous open systems computing environments and targeted at the e-commerce, web serving, and data warehousing applications. Thunder 9200™ is the only midrange product that truly delivers large enterprise level features with a midrange package and price.

The Thunder 9200™ is the best choice for any of four circumstances:

- When S/390® connectivity or data sharing of open systems platforms and S/390® platforms is not required.
- When four front-end Fibre Channel or two front-end SCSI attachments provide adequate bandwidth and connectivity.
- When storage capacities requirements of 17.6 TB for high density 180 GB 7,200 rpm disks or 7.2 TB for high performance 72 GB 10,000 rpm disks are adequate.
- When only the advanced software supported by the Lightning 9900™ is needed. (See also Table 2 and Chapter 2.)

A quick comparison of the major features of the Lightning 9900™ Series and the Thunder 9200™ is shown in Table 1 to illustrate the differences and similarities.

Table 1 – Comparison of Hitachi Freedom Storage® products.

	Lightning 9900™ Series Model 9960	Lightning 9900™ Series Model 9910	Thunder 9200™	Thunder 9200™	Thunder 9200™
Configuration	Control Frame and up to 6 Array Frames	Single, Integrated Control & Array Frame	Rackmount	Deskside 10	Deskside 20
Max. Raw Storage	83.3 TB	8.3 TB	17.6 TB	1.8 TB	3.5 TB
Max. Data Cache	32 GB	16 GB	4 GB	4 GB	4 GB
Max. Control Memory	1.5 GB	1.5 GB	NA	NA	NA
Max. Host Interfaces	32 Fibre Channel or ESCON	24 Fibre Channel or ESCON	4 Fibre Channel or 2 Ultra – 2 Wide SCSI	4 Fibre Channel or 2 Ultra – 2 Wide SCSI	4 Fibre Channel or 2 Ultra – 2 Wide SCSI
Max. Disks	512	48	100	10	20
Max. Hot Spares	16	4	5	5	5
Disks	18, 73, 180 GB	18, 73, 180 GB	18, 36, 72, 180 GB	18, 36, 72, 180 GB	18, 36,72, 180 GB
Host Support	OS/390, UNIX®, Linux®, Windows® NT/2000®, NetWare®, OpenVMS®	OS/390, UNIX®, Linux®, Windows® NT/2000®, NetWare®, OpenVMS®	UNIX®, Linux®, Windows® NT/2000®, NetWare®, + 7 others	UNIX®, Linux®, Windows® NT/2000®, NetWare®, + 7 others	UNIX®, Linux®, Windows® NT/2000®, NetWare®, + 7 others

An Overview of the Thunder 9200™ Architecture

The Thunder 9200™ architecture is fully fault tolerant.

The Thunder 9200™ offers Fibre Channel connectivity to both host and disk, though Ultra SCSI interfaces to the host are also supported. Each Thunder 9200™ model supports up to two disk controllers in keeping with the systems complete fault tolerant design. The controllers are linked with twin 235 MB/sec links, which is about five times as fast as competitive offerings. This ensures controller failover in the case of component failure for uninterrupted operation.

Each of the two fault tolerant controllers comes standard with two 100 MB/sec Fibre Channel host ports (four total) or optional 2 – 2 Gb/sec Fibre Channel host port per controller (four total). For added flexibility, the Thunder 9200™ also supports Ultra2 Wide SCSI (LVD/single ended) and Ultra Wide SCSI (HVD) connectivity to the host.

Each controller supports up to 2 GB of cache (4 GB maximum per system). RAID levels 0, 1, 5 and 0+1 are supported and can be intermixed with the Thunder 9200™. Drive types can also be intermixed within the modules, but must be the same within RAID groups. On the back-end, the Thunder 9200™ supports four Fibre Channel loops (two per controller) in conjunction with its dual-ported drives. The Thunder 9200™ high level architecture is shown in Figure 3.

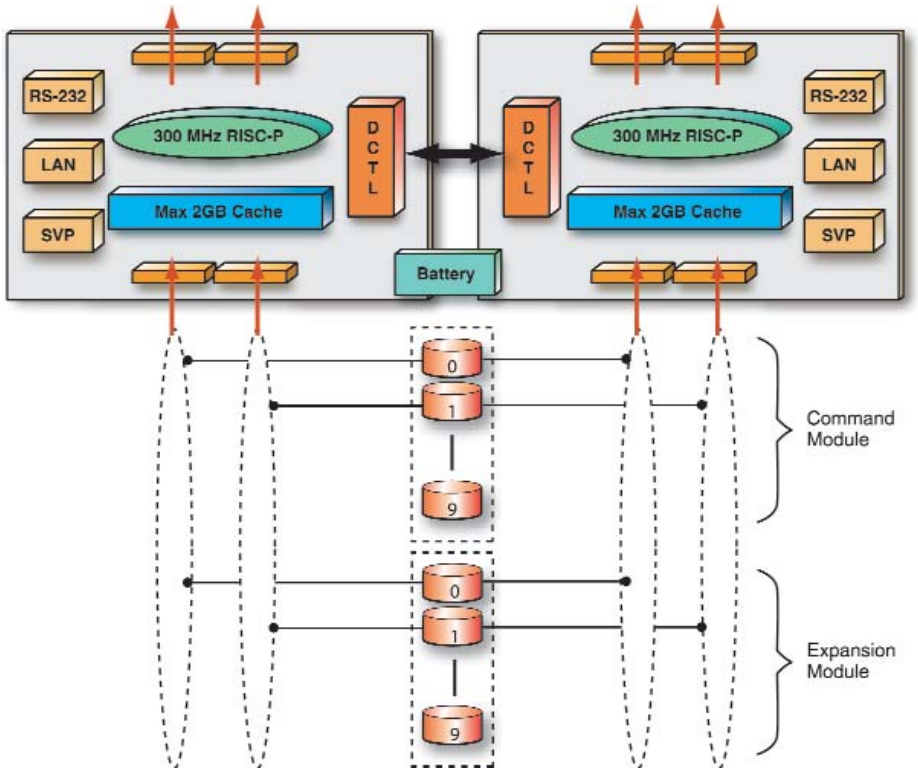

Figure 3 – System block diagram of the Thunder 9200™ fault tolerant architecture.

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 for both the Lightning 9900™ Series and the Thunder 9200™ Series, is available in two companion publications, the ***Hitachi Software Solutions Guide*** and the ***Hitachi Lightning 9900™ Series Architecture Guide*** (Figure 4).

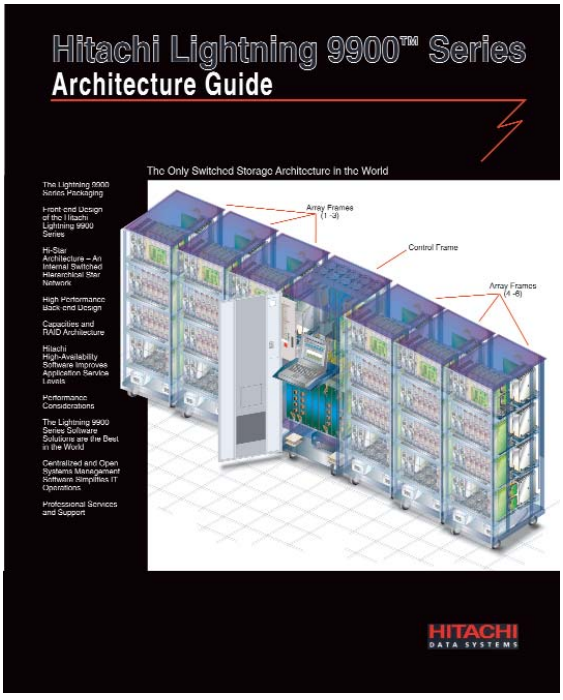

Figure 4 – The Lightning 9900™ Series Architecture Guide.

These documents are downloadable in PDF format from www.hds.com. The reports contain an in-depth assessment of Hitachi Freedom Storage™ hardware and software products. They also compare the Hitachi solutions to the major competitive software solutions and get down into specifics.

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.

Hitachi Surrounds the Thunder 9200™ with Software, SAN, and Service Solutions

2

The Most Advanced Storage Software Solutions in the World

Advanced Hitachi software enables key business benefits.

Customers and analysts regard the Hitachi Freedom Storage Lightning 9900™ Series and Thunder 9200™ storage systems as the most advanced products in their respective storage system markets. Recently Hitachi has also won the respect of many analysts as having taken the leadership in advanced function storage software that supports all Freedom Storage™ products. The ***Software Solutions Guide*** discusses the advantages of this suite of leading-edge software products that enhances both the Lightning 9900™ and Thunder 9200™ Series product lines. This guide is available from Hitachi Data Systems® in PDF form at http://www.hds.com/.

In the Guide's text and illustrations, particular emphasis is devoted to the business benefits of Hitachi's robust suite of advanced function software in the context of three generally accepted categories of business benefits that are enabled by enterprise class intelligent storage system software. These three business objectives can be classified as follows:

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

Thunder 9200™ enables an enterprise to access any information, on any computer, located anywhere, at any time.

The Hitachi Freedom Storage™ Software Solutions for both Lightning 9900™ Series and Thunder 9200™ support an enterprise's strategic goal of accessing any information, on any computer, located anywhere, at any time. The many advanced functions available on Hitachi Freedom Storage™ hardware products are initiated, managed, and controlled through these powerful software programs. Table 2 summarizes how Hitachi software solutions can be mapped to the business benefits mentioned above and which are supported by both the Lightning 9900™ Series and Thunder 9200™ and which are supported on the Lightning 9900™ Series only.

The Hitachi Freedom Storage™ Software Solutions deliver enterprise-wide coverage of on-line data copy/relocation, data access/protection, and storage resource management whether for the midrange or enterprise market segments. Customers have the freedom to choose the precise solution – or combination of solutions – appropriate for their environment after, of course, choosing the appropriate storage systems platform.

Table 2 - Summary of Hitachi Freedom Storage™ software products and the business objectives they serve.

Hitachi Freedom Storage Software Solutions for BOTH the Thunder 9200™ and the Lightning 9900™ Series	Business Benefit		
	Increased IT service levels in availability and performance	Simplified IT operations	Accelerated deployment of new applications and new systems
Hitachi TrueCopy - OpenSystems	✔		
Non-disruptive data replication - Open ShadowImage™	✔		✔
Automatic Path Failover - (Dynamic Link Manager™)	✔	✔	
Host Failover and Parallel Database Clustering	✔		
LUN Security/SANtinel™	✔	✔	✔
VERITAS® support, volume and file management utilities	✔	✔	
Systems Management - HiCommand™	✔	✔	✔
Systems Management - Resource Manager™	✔	✔	✔
Hitachi Freedom Storage Software Solutions for the Lightning 9900™ Series ONLY	**Business Benefit**		
	Increased IT service levels in availability and performance	Simplified IT operations	Accelerated deployment of new applications and new systems
Hitachi TrueCopy - OpenSystems	✔		
Hitachi TrueCopy - S390®	✔		
IBM® GDPS™ Compatible Remote Copy	✔		
Software Asynchronous Remote Copy	✔		
Non-disruptive PiT (Point-in-Time) Copy™ - S/390® (NanoCopy™)	✔		
Non-disruptive data replication -S/390® ShadowImage™	✔		✔
Data migration Service to the Lightning 9900-S/390®	✔	✔	✔
Open backup to mainframe tape using standard backup software	✔	✔	
HARBOR® Backup and HYPERtape	✔	✔	
Restore of single files only - S/390®	✔	✔	
S/390® to Open file conversion - RapidXchange™	✔	✔	✔
Open and S/390 channel file transfer - RapidXchange™		✔	✔
Open to Open file conversion - RapidXchange™		✔	✔
Automatic performance tuning - CruiseControl™	✔	✔	

The Thunder 9200™ Delivers Non-Stop Operations – Not Just Disaster Recovery

Hitachi Data Systems® focuses on minimizing planned interruptions.

Whereas the paradigm for high availability computing is based on avoiding unplanned interruptions to data availability, Hitachi Data Systems® consulting services focuses not only on preventing these unplanned interruptions but also on minimizing planned interruptions, such as scheduled downtime for servicing and backup windows. Using a unique, Six-Sigma, phased approach, Hitachi delivers continuous data availability by applying ISO 9001 type disciplines to threat assessment and risk management. In other words, the focus of Hitachi is on *continuous business* instead of *business continuity.* This subtle yet important paradigm shift is illustrated in Figure 5.

Over the past 13 years, the professional standards for risk assessment and threat management have been centered in practices and writings of the Disaster Recovery Institute (DRI). The DRI knowledge base serves as the industry's best practices standard for contingency planning for the threats imposed by disasters. A discussion of the DRI can be found in Appendix A of this report.

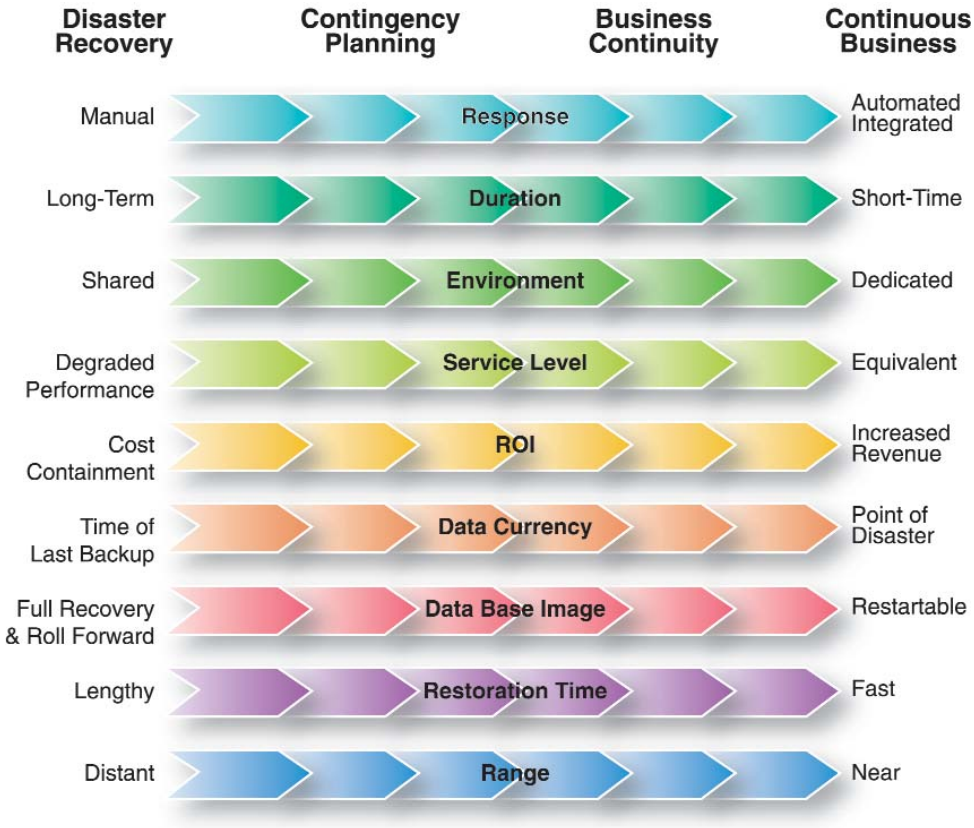

Figure 5 – Hitachi changes the paradigm of high availability computing from business continuity to continuous business.

Hitachi Data Systems Applies Six-Sigma Techniques to Achieve High Data Availability

The Hitachi Continuous Business Paradigm is based on the Six-Sigma standard of quality of operations.

Hitachi Data Systems takes a unique approach to high-availability computing in the storage industry. This approach is based on the Six-Sigma standard of quality of operations. First pioneered by Motorola® in the 1980s, Six Sigma refers to six standard deviations above the mean in sampling of product and process quality. It sets a target of 99.9997 percent defect-free operations for an enterprise, whether for product quality, invoicing accuracy, communications, or any other business process. In short, Six Sigma means getting everything right. Six Sigma can be thought of as a process that is similar to those of ISO 9001 in today's modern enterprise. Using Six-Sigma techniques, the Hitachi solution requires thorough dedication to all quality processes that affect data availability. Implementing the Hitachi Data Systems® solution ensures that an enterprise can easily replicate, backup, and manage all information vital to the enterprise business processes.

The Six-Sigma approach reinforces Hitachi's reputation for non-stop data availability in the industry, in terms of products, services, and networking. Not only are the company's software solutions for high-availability computing accepted as best-of-breed by industry analysts and customers alike, but the Hitachi Data Systems® tool set for infrastructure assessment and high-availability computing planning is unsurpassed.

The Hitachi Data Systems® paradigm of Continuous Business begins by redefining the Scale of 9s and looking at the Scale of 9s from an entirely new perspective. By including both planned and unplanned downtime in the availability equation, Hitachi Data Systems personnel are able to uncover new opportunities for the enterprise to reduce all interruptions – whether they are planned or not. This unique perspective is key to understanding the Hitachi Freedom Storage™ Software Solutions.

Where Customers are Today

SHARE has established a de facto standard for classifying the preparedness of an enterprise to survive an interruption.

The IBM® users group SHARE has established what has become a de facto standard for classifying the preparedness of an enterprise to survive an interruption. Although the SHARE tier description scale was developed within the perspective of disaster recovery (DR) planning, it is important to note that the Hitachi Continuous Business Solution not only encompasses DR, but goes way beyond this traditional emphasis on high-consequence/low-likelihood events. Hitachi expands the focus of interruptions to moderate-to-high-consequence/high-likelihood events that relate to more frequently encountered operations interruptions, such as backup problems, viruses, hackers, software defects, and hardware failures.

SHARE has defined seven tiers (0-6), with tier 6 being the classification for the highest level of availability on the Scale of 9s. There are very few enterprises in tier 6, however, and the majority (90 percent) of enterprises rank themselves in tier 2, as shown in Figure 6. It is often useful for an IT professional or a business manager who is not involved directly in disaster recovery to review Figure 6 and evaluate where his or her enterprise fits into a SHARE category. This can serve as a valid starting point in understanding the financial tradeoff between moving into a higher tier or risking data loss to the extent shown in the event of a disaster. Although it will vary from enterprise to enterprise, the process to move up the scale to tier 6 is exponentially costly.

Figure 6 – Percent of enterprises today operating within SHARE Disaster Recovery Tiers 0 through 6.

The Importance of Six-Sigma Computing

Planned downtime is often an order of magnitude greater than unplanned downtime in terms of real hours when data is not available.

There is no doubt that unplanned downtime has an entirely different impact on IT operations than planned downtime. However, a dramatic new picture can be seen when availability is looked at based on the combination of both planned and unplanned downtime per year. Planned downtime can vary from enterprise to enterprise, but on average it is somewhere between 96 hours per year (8 hours per month) and 1,152 hours per year (96 hours per month). Planned downtime is time spent on backups, database reorganizations, bringing new software on-line, time for vendor support, etc. Table 3 illustrates two important points. First, whereas the emphasis has traditionally been on unplanned downtime, planned downtime is often an order of magnitude greater in terms of real hours when data is not available. Second, when both planned and unplanned downtime are accounted for, data availability is dramatically lower on the Scale of 9s, as can be seen in the right-hand column of Table 3.

Availability Based Only on Unscheduled Downtime			Availability Based on Scheduled and Unscheduled Downtime			
9s Scale	Unscheduled Downtime/yr.	% Availability	Scheduled Downtime/yr.[1]	Total Downtime	% Availability	9s Scale
1	876 hrs.	90	1152 hrs. (96/mo)	2028 hrs.	76.84	0.77
2	87.6 hrs.	99.0	1008 hrs. (84/mo)	1095.6 hrs.	87.49	0.87
3	8.76 hrs.	99.9	864 hrs. (72/mo)	872.8 hrs.	90.03	0.99
4	52.56 min.	99.99	720 hrs. (60/mo)	770.8 hrs.	91.20	1.20
5	5.25 min.	99.999	576 hrs. (48/mo)	576.1 hrs.	93.42	1.34
6	31.54 sec.	99.9999	432 hrs. (36/mo)	432 hrs.	95.06	1.51
7	3.15 sec.	99.99999	288 hrs. (24/mo)	288 hrs.	96.71	1.67
8	.31 sec.	99.999999	144 hrs. (12/mo)	144 hrs.	98.35	1.83
9	.03 sec.	99.9999999	96 hrs. (8/mo)	96 hrs.	98.90	1.89

Table 3 – Through the Six-Sigma approach, Hitachi attacks both planned and unplanned downtime.

The Business Benefits of the Hitachi Continuous Business Paradigm

Hitachi Data Systems® Continuous Business Paradigm, related services and software solutions provide more key benefits than conventional approaches. These include:

- Conventional hot-site solution alternative
- Dedicated, on-demand, managed rapid-recovery service
- Dedicated testing configuration
- Concurrent change activity support
- Reduced technical complexities
- Reduced communications expense
- Flexible, close secondary facilities at over 500 locations with Hitachi Data Systems® partners
- Secure, hardened, and reliable data center facilities
- Comprehensive, highest quality network capability

Continuous Business Planning is Composed of Four Related, but Standalone Phases

1. **Workshop**—In the first phase of Continuous Business Planning, Hitachi Data Systems assembles a cross-functional team to understand the customer processing environment and business requirements for continuous business and data availability. During this three-day interactive workshop phase, the Hitachi Data Systems® team engages the customer with a detailed set of questions outlined in the Continuous Business Planning Questionnaire (Appendix B). The workshop also uses worksheets that allow for a detailed review of the important characteristics of continuous business in the areas of:

Hitachi Data Systems engages the customer with the Continuous Business Planning Questionnaire.

   - Facilities
   - Hardware

[1] In this example, planned downtime is spread over the Scale of 9s chart as shown to reflect the fact that high availability shops tend to have less planned downtime in addition to less unplanned downtime.

- System software
- Communications
- Support software
- Applications
- Business drivers, service levels, and investment practices
- Management practices

After completion of the worksheet, all potential actions for Six-Sigma improvement are then subjected to a "Likert" grid analysis that is used to determine the relative effectiveness of implementation. In a "Likert" grid analysis, each item is displayed on an "x, y" grid. The "x" axis of the grid ranks the item in five categories of importance from "Little Importance" to "Very Important." The "y" axis of the grid ranks the item from "Not Effective" to "Very Effective." Using this technique, it is possible to focus on the most important items to implement first those that will have the highest benefit (effectiveness) to continuous business objectives.

2. **Assessment Phase**—This phase is a fee-based consulting engagement that defines the business impact of interruptions, identifies mitigation alternatives, and estimates high-level costs. Hitachi Data Systems® has adopted the Cost-of-Risk Analysis (CORA) software tool to analyze the cost-benefit of Hitachi high-availability solutions for its client. CORA makes it easy for Hitachi Data Systems and the customer to identify and configure optimum solutions, and to size the business case for these demanding business requirements. The CORA cost-benefit analysis is refined further in the Design Phase.
3. **Design Phase**—The third phase is also a fee-based consulting engagement that develops detailed management, configuration, network, and facility requirements to meet agreed to continuous business objectives, such as SAN Strategic and Tactical Design Source. Based on this information, a more accurate CORA cost-benefit analysis is conducted. CORA is discussed (with an example) later in this chapter.
4. **Implementation Phase**—The fourth phase is a planning and installation consulting engagement to implement the required technology—servers, storage systems, software solutions, SANs, etc.

Cost of Risk Analysis (CORA)

CORA is an expert, software-based risk management program.

CORA is an expert, software-based risk management program that helps identify the cost, benefit, and returns on investments in high-availability computing infrastructures. Developed initially for the insurance industry by New York City-based International Security Technology, Inc., CORA draws on three decades of experience in risk management. Use of the CORA tool with Hitachi Data Systems software solutions discussed in this Guide will reduce interruption windows and the impact of an interruption on business operations. The findings and recommendations will be in accordance with enterprise business drivers and in terms that users can take to their executive team.

The CORA software tool has five unique features:

1. CORA interfaces with an organization's risk information management system (RIMS) and can import basic identifying information about individual operating units or facilities. This ensures that CORA's risk management financial data will track the organization's "book values" for individual facilities.

2. Risk experts define and store risk parameters in CORA as "risk rule" files. These rules include nominal risk parameters and the rules for adjusting these nominal parameters to reflect local risk factors at individual facilities.
3. Line operating personnel collect field data easily by supplying descriptive information about the local risk environment and level of activity at individual objects. CORA makes this practical by printing out tailor-made data collection forms for each of the facilities. This feature conserves the costly time of the risk experts, and it has the added benefit of involving the line personnel in the risk management process.
4. CORA uses the expert risk rules to adjust the nominal risk parameters automatically to reflect local conditions at individual facilities (similar to the way in which risk management experts would act). CORA then calculates for each facility the expected losses (dollars/year), the single occurrence losses (dollars/loss), and the Return On Investment (ROI) of security measures, insurance policies, and business resumption plans—in place or proposed.
5. CORA's Financial Simulator consolidates financial results, expected losses, losses averted, and security costs onto a projected income and expense statement, and then calculates the overall financial impact of risk on the organization for each of the defined security strategies. The Risk Manager can construct, save, and recall alternate analysis setups.

CORA has a number of other features that enhance its value. The risk analyses, ROI calculations and Financial Simulations can be exported as Excel spreadsheets for further processing, graphing, or incorporation into external report documents. All CORA data files include a standard header block that stores the name of the person who entered the data as well as a time-date stamp. This information appears on the printed reports.

Actions to Mitigate the Risk and the Cost Benefit Analysis

CORA is based on a concept called ALE or Annualized Loss Expectancy.

A major benefit of the CORA process is the Cost Benefit Analysis, where each action to mitigate risk is evaluated in terms of cost and ROI. CORA is based on a concept called ALE or Annualized Loss Expectancy. This is the term for risk loss expressed at an annual rate. It is the product of the threat occurrence rate, the loss potential of the function or asset, and the vulnerability of the function or asset to the threat.

In the example below, these three factors are 1/10 years, $20,000, and 50 percent, respectively. In this case, ALE is calculated as follows:

ALE = Threat Occurrence Rate x Loss Potential x Vulnerability Factor

ALE = 0.10 x $20,000 x 0.50 = $1,000/year

If the Single Occurrence Loss (SOL), therefore, of a function or asset is $20,000 and the occurrence rate of the threat is once in ten years, the associated ALE is $1,000/year. Figure 6 shows an example of a Cost Benefit Analysis for various items using the ALE concept described above. The Hitachi Data Systems CORA service will work with clients to determine the ROI of alternative risk management investments in a fashion similar to that shown below:

Figure 7 – Example of a Cost Benefit Analysis.

COST - BENEFIT ANALYSIS

	ALE	ALE Reduction	ROI	PV - ALE
Baseline Conditions	$81.840	Discount Rate =	0.75%	0.067
Security measures. . .				
Inert gas fire extinguishment	$81,826	$214	-95%	
Fire Detection & Alarm	$81,787	$53	-90%	
Personel ID / Access Control	$80,200	$21,802	1010%	
Burglar Detection & Alarm	$79,233	$2,607	438%	
Automatic Sprinkle	$30,893	$341	-74%	
Guards	$59,144	$22,898	-02%	
Housekeeping for fire safety	$31,743	$97	-11%	

In the example, only two security measures, Personal ID / Access Control and Burglar Detection & Alarm, have positive ROIs.

Hitachi Software Solutions and Services Provide a World-Class Selection of Risk Reduction Alternatives for the Enterprise

Duplicating and replicating data within IT architecture are critical to achieve Six-Sigma objectives.

The Hitachi Data Systems solutions include a broad selection of hardware, software, SAN, networking, management practices, business focus, and management software selections that can be implemented to reduce risk of planned and unplanned interruptions. Duplicating and replicating data within a solid worldwide IT architecture based on best-of-breed products and infrastructure services are critical techniques for achieving Six-Sigma objectives. All actions that can move, add, or change existing practices and infrastructure components are considered to reduce the business impact of both expected and unexpected interruptions.

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

Remote Copy—Refers to the mirroring of data, hopefully to provide time-and transaction-consistent remote copy of that data. The purpose of remote copy is to protect the data in the event of a business interruption.

Data Duplication—Software that duplicates data as in remote copy or PiT snapshots. Data duplication differs from data migration in that with data duplication, at the end of the process there are two copies of data and with data migration there is only one.

Data Migration—Software that migrates data from one storage device to another. This feature differs from data duplication in that at the end of the process there is only one copy of data. The purpose of data migration is to reduce operational complexity and costs for storage subsystem upgrades or equipment refurbishment.

There are many copy software alternatives today, especially when traditional backup methods of copying data sets are considered. This is why it is important to choose only from best-of-breed solutions when addressing continuous business objectives. Hitachi Data Systems provides a full range of copy software functionality that has a proven track record of allowing customers to minimize interruptions, create business intelligence applications that use copies of production data, and move data to newer hardware to avoid the possibility of future interruptions. These are discussed in detail in the next chapter.

Thunder 9200™ SAN Solutions and the Hitachi Freedom Data Networks™ Frameworks

Freedom Data Networks™ dramatically reduces planned downtime.

A major potential benefit of Storage Area Networks (SANs), based on the Thunder 9200™ storage systems, is that LAN-free and server-free backup reduce dramatically the amount of planned downtime needed for data center operations. Backup windows are chronic major issues on the minds of storage executives as discussed earlier in this chapter. Table 3 shows that scheduled downtime for backups and maintenance can create an order of magnitude more downtime than unplanned downtime due to equipment failures and disasters for the typical enterprise.

Hitachi Freedom Data Networks™ provides an open architecture that offers organizations freedom of choice in deploying data access, protection and sharing capabilities across the enterprise. Using multiple technologies, protocols, and solutions, such as SAN, NAS, and IP, Freedom Data Networks enables customers to design and build, leverage, and augment storage infrastructures – providing access to any data from any computer, anytime and anywhere.

Hitachi Freedom NAS™ and Hitachi Freedom SAN™ solutions are the core offerings behind the Freedom Data Networks approach. They complement the Hitachi Freedom Storage™ 9000 All-Fibre Family by allowing more flexibility than ever in heterogeneous environments. While SAN architectures respond to high bandwidth needs, NAS addresses the need for rapid file access, especially critical for e-business applications. HDS offers the best of both.

The Freedom Data Networks architecture exploits advances in servers, storage systems, interconnection devices, network protocols, and network configurations. However, the Freedom Data Networks architecture goes beyond SANs, providing the overall structure for solutions that allow customers to manage their data without being tied to proprietary technology that limits their business options. This gives the enterprise the option to locate storage either inside or outside the data center, wherever it makes the greatest business sense to do so. Moreover, it gives operations personnel the ability to manage a wide variety of server, interconnection, and storage platforms under the Freedom Data Network® methodologies. This provides them flexibility in establishing open system configurations and protecting investment in currently installed system components.

Within the Freedom Data Networks architecture, SANs provide high-speed Fibre channel networks for connecting multiple, multi-vendor servers to a pool of multi-vendor storage devices distributed throughout the enterprise. SANs offered by Hitachi Data Systems® will support the open-systems standards being developed by the Storage Networking Industry Association (SNIA).

Thunder 9200™ Packaging

3

The Thunder 9200 Packaging is Completely Modular and Expandable

With the Thunder 9200, single-image storage systems as large as 18.0 TB can be created. Command Modules (controller plus 2-10 disks) combined with Expansion Modules (1-10-disk add-ons) form building blocks for three packaging configurations to meet capacity needs from the department to the enterprise.

A Command Module with 2-10 disks comprises the smallest element of scalability – the Deskside 10 configuration that supports capacities up to 1.8 Terabytes (10 disk drives). A Deskside 10 unit is shown in Figure 8.

Thunder 9200™ is based on Command and Expansion Modules

Figure 8 – The Deskside 10 System is the smallest unit of capacity and scalability.

When an Expansion Module of 1-10 disks is added to a Command Module, the next larger package size of the Thunder 9200™ can be created. The Deskside 20 package supports capacities up to 3.6 Terabytes (20 disk drives) and is shown in Figure 9.

Figure 9 – The Deskside 20 System contains one Command module and one Expansion module.

A Deskside system can be removed from its cabinet and mounted on rails to start a Rackmount system.

The Rackmount System is Ideal for Growing Enterprises

Total scalability and building block architecture is at the heart of the design philosophy of the Thunder 9200™ storage systems. As illustrated on the cover of this Guide, the Command Module and first Expansion Modules of the Deskside 10 and 20 packages form the base of a Rackmount Package when turned sideways into a standard 19 inch wide rack mount cabinet. A fully populated Rackmount System is shown in Figure 10 with a maximum of 100 disk drives and 18 Terabytes of RAID-0 storage capacity.

Figure 10 – A fully configured Rackmount System contains one Command Module and nine Expansion Modules.

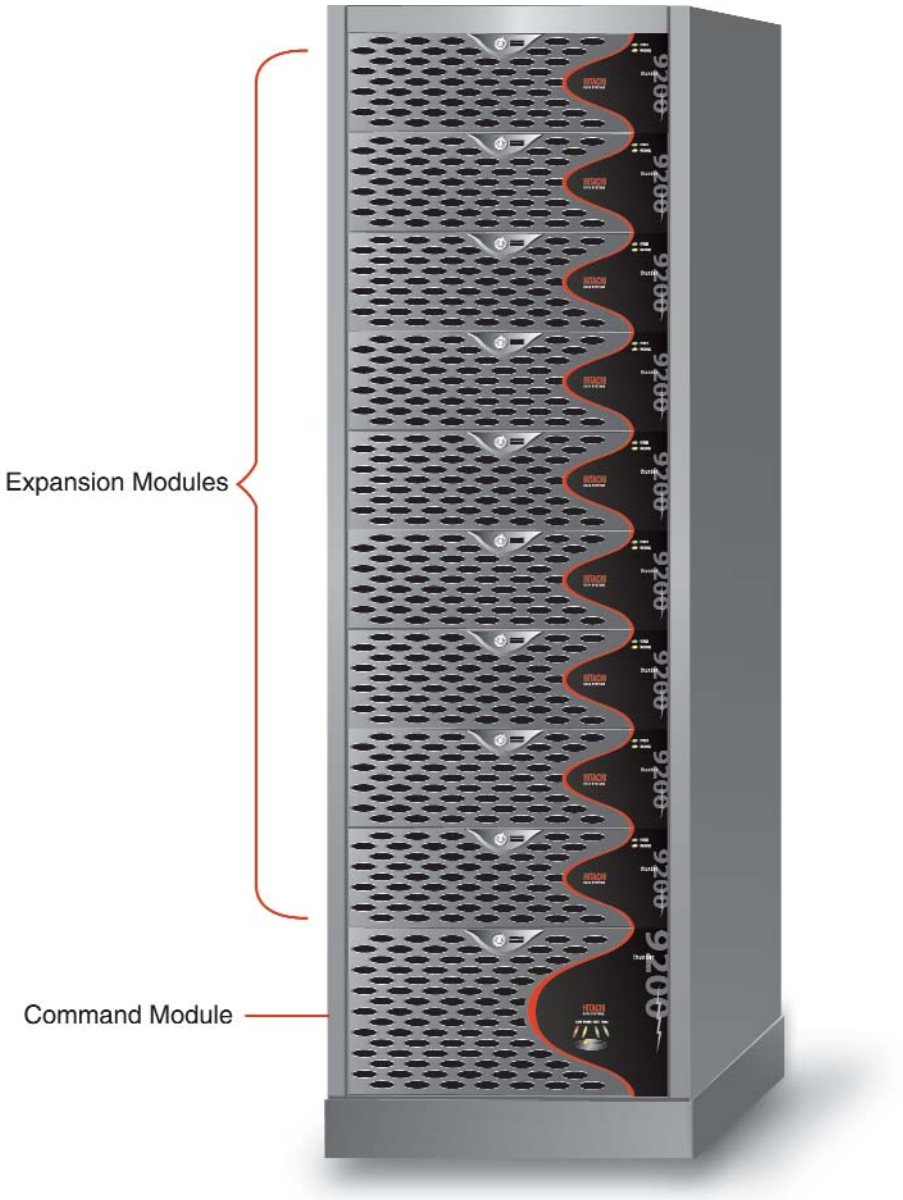

Command Module Design

All elements of the Command Module are fully redundant and efficiently packaged into a highly modular enclosure.

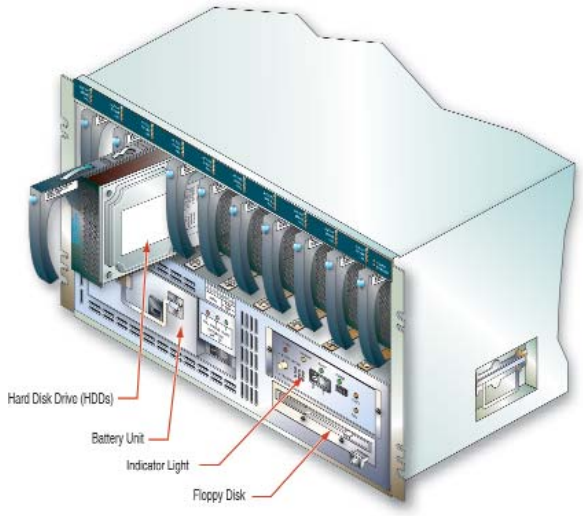

Figure 11 – Front view of a Thunder 9200™ Command Module.

Figure 11 shows the front view of the Command Module containing 10 Hard Disk Drive Modules (HDDs), the battery unit to protect data in cache in the event of a power failure, indicator lights, and a floppy disk for installing microcode. There is a floppy disk in the Command Module since microcode can also be downloaded through LAN or RS232 connections on the controller board as shown in Figure 3 (Chapter 1). This design is consistent with the complete fault tolerant philosophy of the Thunder 9200.

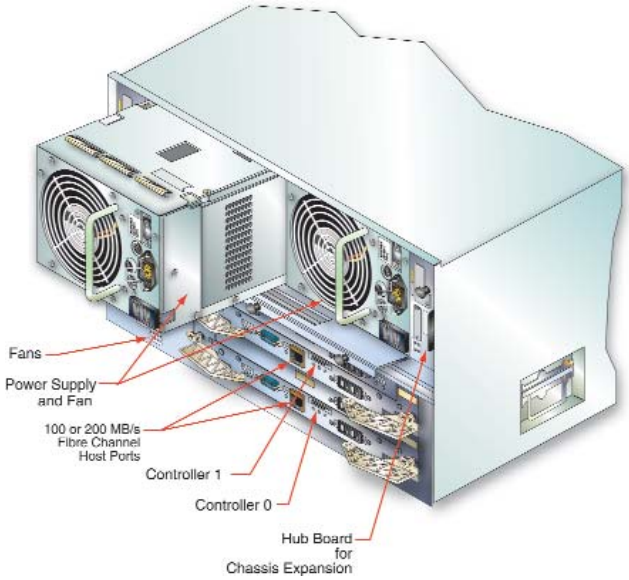

Figure 12 – Rear view of a Thunder 9200 Command Module.

As shown in Figure 12, the back of the Command Module contains redundant power supply/fan units, redundant controller boards (optional) and boards which interconnect with Expansion Modules for additional storage capacity. A controller board and power supply/fan module are shown in Figures 13 and 14.

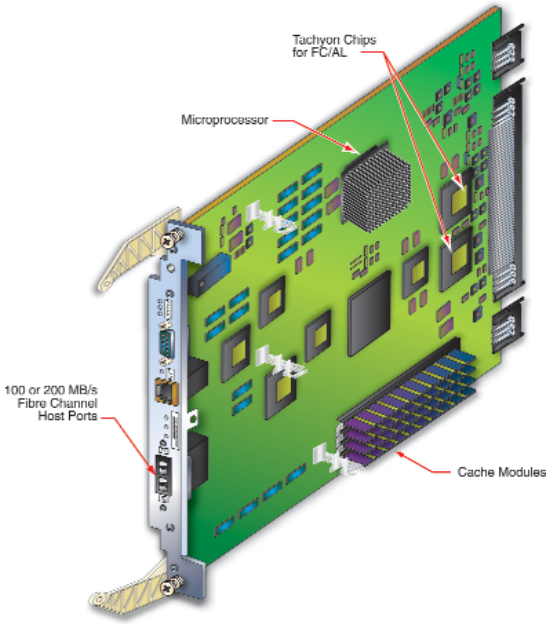

Figure 13 – One of two fault tolerant controllers that form the heart of the Command Module.

The controller board illustrates the "all Fibre" design of the Thunder 9200 with two standard 100 or 200 MB/sec Fibre Channel host ports built into the board and the Tachyon Chips for Fibre Channel Arbitrated Loops FC/AL back-end loops to support back-end disks. Although Fibre Channel front-end interfaces are standard, alternative interfaces are available and discussed in the next chapter.

Figure 14 – One of two fault tolerant power/supply fan modules.

Expansion Module Design

Expansion Modules contain 10 HDDs each and stack on top of the Command Module for efficiency. An Expansion Module is shown in Figure 15 and a Hard Disk Drive is shown in Figure 16.

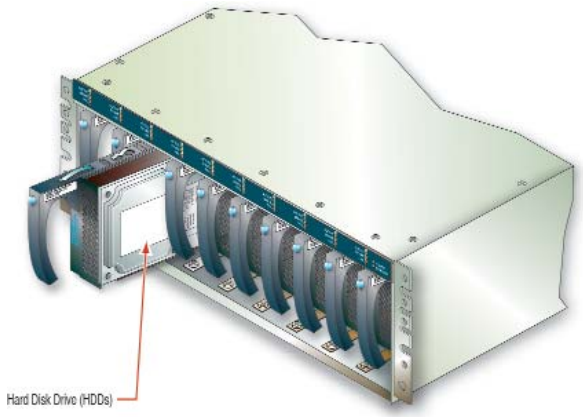

Figure 15 – A Thunder 9200 Expansion Module with 10 HDDs.

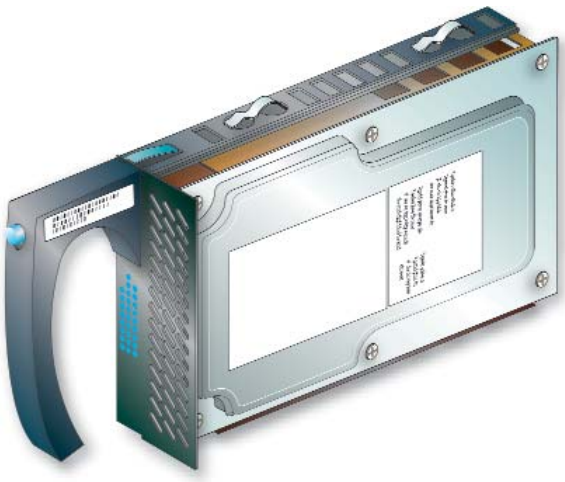

Figure 16 – A hot-swappable Hard Disk Drive Assembl;y (HDD).

NAS and SAN Enabler Packaging Options

4

The Importance of SAN and NAS Networked Storage

As today's information-centric businesses expand operations and implement data-intensive applications, IT managers seek new ways to quickly deploy storage and storage infrastructures without disrupting business operations. To this end, networked storage plays a key role, delivering numerous benefits, from shared pooling and better utilization of storage resources to improved performance, better security, and higher availability. Now, it's possible to have it all in one solution.

NAS Complements SAN

The Hitachi Freedom Storage Thunder 9200 SAN Enabler packaging option responds to the need for high bandwidth connectivity between servers and storage nodes and for greatly improved central management of storage resources.

But for many applications especially Web, design, and medical - the concern is not bandwidth. It's file access response time. With Hitachi Freedom NAS, information is immediately available on the Hitachi Freedom NAS™ file server packaging option, and instantly accessible to numerous clients, allowing them to share data and integrate computers throughout the enterprise network. The modular design of the Thunder 9200 SAN Enabler and Freedom NAS packaging options provides quick, easy, and inexpensive expansion of storage.

Thunder 9200 SAN Enabler (SE) Packaged for Quick Deployment

The Thunder 9200 SAN Enabler (SE) provides state-of-the-art storage area networking benefits.

Hitachi Data Systems offers the ultimate turnkey networked storage solution. Thunder 9200 SAN Enabler (SE), combining the Hitachi Freedom Storage™ Thunder 9200 with two Brocade® SilkWorm® fabric switches in one rackmounted configuration, pre-wired and ready-to-go. This option gives the high-availability and high-performance advantages of Hitachi Freedom Storage, plus the benefits of a switched Fibre Channel networking architecture developed by Brocade Communications Systems, Inc. This networked fabric design provides the high level of intelligence and scalability that successful enterprise deployments require. This solution fulfills a wide variety of success criteria for IT departments that face growth, consolidation, and response-time issues. This option is illustrated in Figure 17.

Hitachi Freedom NAS™ (Network Attached Storage)

Freedom NAS can optimize applications requiring high performance file serving/file sharing.

Hitachi Freedom NAS™ is the latest component of Hitachi Freedom Data Networks. It is designed specifically to provides a coherent, intelligent information infrastructure to support new applications, satisfy increasing customer demand, and enable expanding operations. NAS applications are often deployed for File / Web serving, document / record imaging, streaming media, video design, telco call centers and computer aided design and manufacturing (CAD/CAM). The benefits of NAS include faster data response times for selected applications, reduced server overhead by moving file I/O to the NAS server, high availability, and operating system independence. Hitachi Freedom NAS is a state-of–the-art way to respond to Web or departmental requests for immediate access to information:

Configured with the NSS open-architecture µSTOR-II™, Hitachi Data Systems high-availability storage is transformed into an incredibly fast and reliable, low-cost Web NAS server/storage system. This Thunder 9200 packaging option is shown in Figure 18.

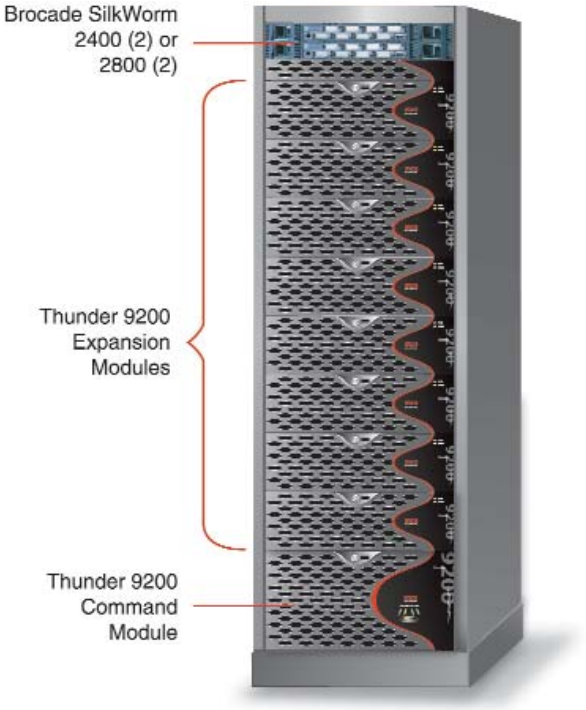

Figure 17 – A Thunder 9200 Rackmount System with SAN Enabler (SE) packaging option.

Figure 18 – The Hitachi Freedom NAS packaging option supports both the Thunder 9200 Rackmount and Deskside System options.

High Performance Front-end Design

5

Fibre Channel and/or SCSI Host Connectivity and Scalability

Fibre Channel front-end interfaces are standard.

The Thunder 9200 supports both Fibre Channel and SCSI interfaces for open systems connectivity. For connectivity to S/390® mainframe platforms, the Lightning 9900™ Series ESCON® or FICON are required. Each of the two fault tolerant Thunder 9200™ controllers comes standard with two 100 MB/sec Fibre Channel host ports (four total) or optional two - 2 Gb/sec Fibre Channel host port per controller (four total). For added flexibility, the Thunder 9200 also supports Ultra2 Wide SCSI (LVD/single-ended) and Ultra Wide SCSI (HVD) connectivity to the host. These optional interfaces are shown in Figure 19 and the bandwidth of supported interfaces is shown in Table 4.

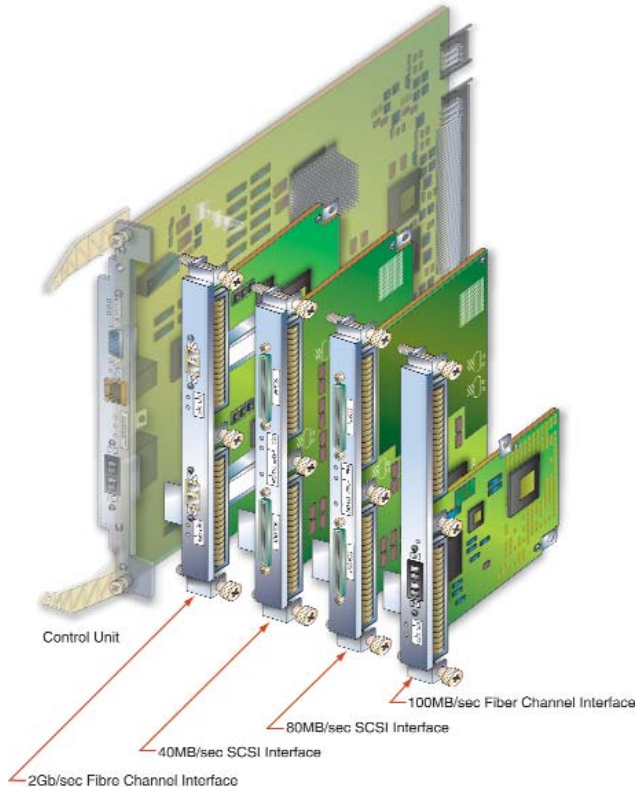

Figure 19 – A Thunder 9200™ controller board with optional Fibre Channel or SCSI interfaces.

Supported Interface	Bandwidth	Number of Ports
Fibre Channel Adapters	100 MB/sec or 2 Gb/sec	4 Maximum
SCSI Adapters	40/MB or .80 MB/sec	2 Maximum

Table 4 – Bandwidth of supported interfaces of the Thunder 9200.

Non-volatile Fault Tolerant Cache Memory Accelerates Front-end Performance

Thunder 9200 cache design is superior to all midrange competitors.

Each controller supports up to 2 GB of cache (4 GB maximum per command module) to accelerate information flow to and from the connected hosts. System microcode maintains cache coherency between the fault tolerant caches using high-speed 235 MB/sec links, which is about five times as fast as competitive offerings. This design is very important to maintain consistent performance. First, a large non-volatile cache can receive large writes without having to de-stage to disk and can terminate the operation to the host in the fastest possible time. Some applications write large temporary database files in the process of joining two large tables that exceed the size of system memory. For these and other types of decision support (DSS) applications, the Thunder 9200 is ideal for this reason.

A second performance benefit of a large cache design is that more frequently accessed data can be kept in cache using Hitachi FlashAccess™ thereby increasing the possibility that it will be read from cache and not disk if re-referenced by the application.

The third performance implication of the mirrored cache and the high-speed inter-connect design of the Thunder 9200 can be seen from the perspective of rapid updating of the mirror in cache to maintain matching copies of data (cache coherency). Since this is a critical path item in system operation, the faster the interconnect, the faster overall system performance regardless of the "read write" mix.

Hitachi FlashAccess™ 9200 Accelerates Front-end Performance Even Further

Using FlashAccess™ 9200 is like having a solid state disk.

Optional FlashAccess 9200 software enables storage managers to increase the speed of database applications and lock selected logical units (LUNs) into cache. This is ideal for frequently accessed data such as database indices etc. The result is ultra high performance with a 100-percent hit rate to the selected LUNs locked in cache as they are accessed at memory speed rather than disk I/O speed. Using FlashAccess 9200 for a data set is like placing data on a solid state disk. This is illustrated in Figure 20.

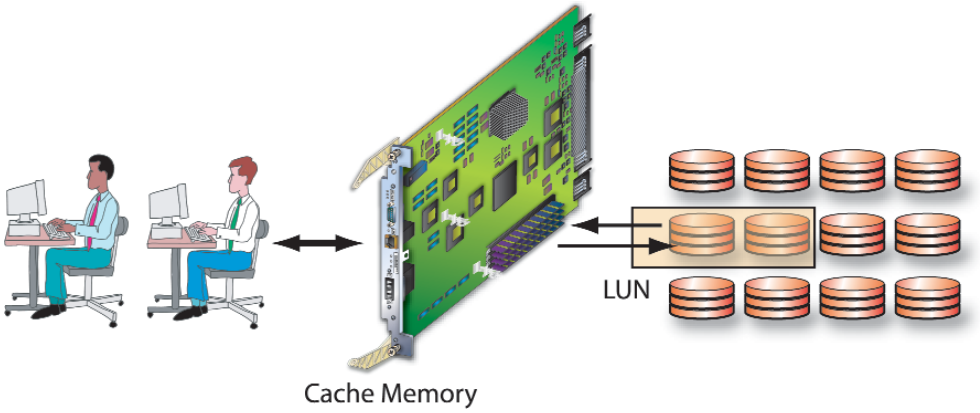

Figure 20 – Hitachi FlashAccess™ 9200 maximizes front-end performance by placing data permanently in cache.

Read-ahead for High-performance Sequential Reads

Read-ahead ensures that requested data will already be in cache.

Read clustering in the Thunder 9200 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 Thunder 9200 "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 cache, so the request can be satisfied immediately for the highest possible front-end performance.

Types of Hosts Supported by the Thunder 9200™ Storage System

The Thunder 9200™ supports all major open systems hosts and mainframe hosts for unmatched heterogeneous open system host connectivity. The Thunder 9200 supports direct Fibre Channel and/or SCSI host attachment to platforms as listed below:

- Windows NT® 4.0
- Windows® 2000
- HP-UX®
- Solaris™ (Sun™)
- AIX® 32 and 64 bit (IBM®)
- NetWare® (Novell®)
- Linux® (Red Hat)
- IRIX® (SGI™)
- Tru64™ UNIX® (Compaq®)
- OpenVMS™ (Compaq®)
- DYNIX/ptx® (IBM® - formerly Sequent)
- NCR® SVR4

"All Fibre" Back-end Design, System Capacity, and RAID

6

High-Performance Back-end Fibre Channel Arbitrated Loop (FC/AL) Design

"All Fibre" back-end design.

State-of-the-art high-performance back-end drives are available with the Thunder 9200™ to build either high capacity or high-performance system configurations. Drive capacities can be mixed within a given configuration within certain restrictions. All drives are attached via two Fibre Channel Arbitrated Loops (FC-AL) per controller with each loop operating at 100 MB/sec. This gives an unsurpassed back-end bandwidth of 400 MB/sec.

Each FC-AL path is driven with a dedicated processor and a Data Recovery and Regeneration (DRR) circuit. This is in contrast to some competitive products where back-end paths are shared through a single processor and single DRR circuit.

State-of-the-Art Fibre Channel Disk Drives

Configure high-performance or high capacity systems.

The Thunder 9200 supports four types of Fibre Channel disk drives in an "all Fibre" back-end.

Two choices are offered to configure the highest performance systems with low capacity high-spin drives with the minimum number of GB per actuator.

- 18 GB 15,000 RPM drives
- 36 GB 10,000 or 15,000 RPM drives

Two choices are offered to configure the highest capacity systems within high capacity drives.

- 72 GB 10,000 RPM drives
- 180 GB 7,200 RPM drives

Storage Capacities of the Thunder 9200™

The Thunder 9200™ supports an unmatched range of capacities among midrange products as shown in Table 5.

System capacities and number of disk drives for various configurations in Terabytes (TB)				
Raw System Configuration	Raw capacity w/ 18 GB drives	Raw capacity w/ 36 GB drives	Raw capacity w/ 72 GB drives	Raw capacity w/ 180 GB drives
Thunder 9200™ Rackmount	1.8 TB=100 disks	3.6 TB= 100 disks	7.2 TB=100 disks	17.6 TB=100 disks
Thunder 9200™ Deskside 10	.18 TB=10 disks	.36 TB= 10 disks	.7 TB=10 disks	1.8 TB=10 disks
Thunder 9200™ Deskside 20	.36 TB=20 disks	.7 TB= 20 disks	1.4 TB=20 disks	3.5 TB=20 disks

Table 5 – Thunder 9200™ system capacities and number of disk drives for various configurations.

Command Tag Queuing (CTQ)

CTQ greatly enhances performance.

All Thunder 9200 Fibre Channel drives support Command Tag Queuing (CTQ) commands issued by the controller. The Thunder 9200 allows up to 128 commands to be queued per port for a maximum of 256 commands per controller. CTQ allows writes to be "held" in cache and ordered in sequence as the data is arrayed physically on the disk to minimize head seeking. Suspended writes are then written to disk with the head traveling in the same direction and not back and forth across the platter. This feature greatly enhances performance.

Advantages of the Thunder 9200 RAID Hardware

The Thunder 9200 supports RAID 0, 1, 5, and 0+1. Hitachi is known to design the most advanced RAID (Redundant Array of Independent Disks) controllers in enterprise and midrange storage to interface to its disk subsystems. These intelligent controllers provide disk interface and RAID management, offloading these tasks to dedicated processors. All disks in the system are defined as part of a RAID array of one type or another and non-volatile RAM on each controller accelerates RAID functions, particularly disk writes.

Configurations for all Thunder 9200™ Series are shown in Figure 21.

Figure 21 – The Thunder 9200 supports RAID-0, 1, 5 and 0+1.

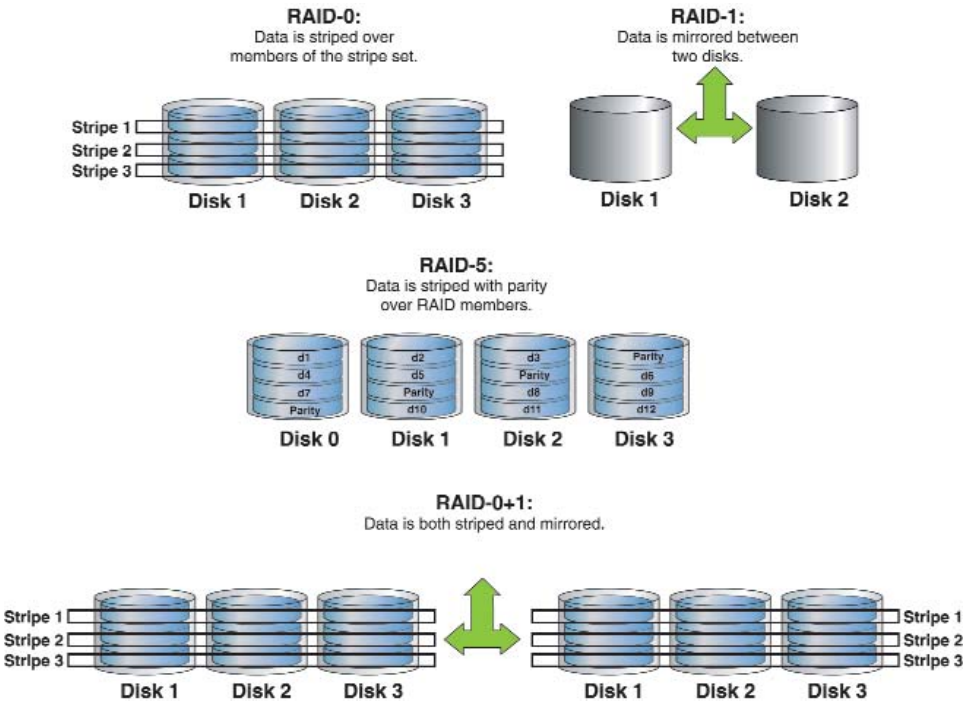

RAID-0

Data is striped across all members of an array. This RAID type can be optimal for large file read and write workloads where the lowest levels of data protection is required.

RAID-1

Data is read from whichever drive is available first but written to both.

With RAID 1 all data is mirrored across to separate disks. This RAID choice is the highest performance and highest availability option but the most expensive. Data is read from whichever drive is available first but written to both.

RAID-5 – Distributed Parity

The overhead of RAID-5 is equivalent to one disk drive.

In RAID-5 arrays, data is striped across the array members in a fashion similar to RAID 0, 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 rotated 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 or the specific RAID configuration. Through the use of unique algorithms, Hitachi RAID-5 arrays perform as well as other vendors' RAID-1 arrays. This feature is discussed in greater detail in Chapter 7. Thunder 9200™ also provides the unique ability to configure RAID-5 arrays with up to 16 drives thereby greatly reducing the cost of RAID protection. With the Thunder 9200, up to 94% of RAID-5 capacity is usable, in contrast to most competitor's 80% usable storage.

RAID-0+1

Also known as RAID-10, RAID-0+1 arrays combine the best of both RAID-0 and RAID-1 by mirroring two striped arrays.

Definable RAID Rebuild Priorities

Users can choose between Slow, Medium, or Fast rebuild priorities.

In the event of a disk failure, RAID-1+ or RAID-5 arrays can be rapidly and automatically rebuilt using up to five available "global hot spare" drives. Users can choose between Slow, Medium, or Fast rebuild priorities depending on the time of the failure and other operational considerations.

Up to Five Global Hot Spares

Data on the failing disk is automatically copied to one of the global hot spares.

Designed to ensure the integrity of data, the Thunder 9200™ allows users to specify up to five global hot spares from among any of the storage units in the storage system. In the unlikely event that a storage unit reaches an error threshold, but is still able to transmit data correctly, the data on the failing disk is automatically copied to one of the global hot spares. If the storage unit fails completely before the copy is produced and the user has defined the RAID group as either RAID-1 (mirrored), RAID-0+1 (mirrored and striped), or RAID-5 (rotating parity group), a global hot spare is selected and the data is rebuilt.

Logical Unit Expansion (LUSE)

The Thunder 9200 also allows logical volumes to be expanded online.

RAID Levels Can be Intermixed for Optimal Data Protection and Performance

All Thunder 9200 RAID choices can be intermixed concurrently within the Thunder 9200 system. Different RAID levels can be selected to best fit application data availability and performance requirements.

The Highest Availability Midrange Product

7

Thunder 9200™ Fault Tolerant and High-Availability Features

Thunder 9200™ offers an array of premium features not usually offered in midrange products.

The Thunder 9200™ offers all the fault tolerance and high-availability features expected in a midrange class storage subsystem like fully redundant and hot-swappable components, battery backed up cache, host failover, alternate pathing, and cluster support. However, the Thunder 9200™ offers an array of additional premium availability features that are usually only available in higher-end systems. Some of these extra features supported by the Thunder 9200™ are remote copy, "point-in-time" copy (a traditionally "enterprise class" function), microcode sparing, non-disruptive microcode upgrades, automatic monitoring and call-home (Hi-Track®), and several premium data integrity and availability tures.

High-Availability (HA) is Really a Dimension of High Performance

When data is not available users have "zero" performance. It can't get any lower! As we discussed in Chapter 2 (See Table 3), it does not matter whether downtime is planned or unplanned, it is still downtime and "zero" performance for users. As it becomes increasingly important for a business to support continued access to global information 24 hours a day, 7 days a week, careers often depend on the availability of service levels provided by IT to the enterprise or by the system administrator to the department, workgroup or branch office.

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 storage unit 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*.

Software Products Contribute to High-Availability Computing

Thunder 9200 software allows processing to continue in the event of an outage.

Many factors can cause unplanned downtime. The Hitachi Thunder 9200™ Series has been designed to eliminate as many of these factors as possible in both hardware redundancy, on-line replaceable components, and software data copy functions to allow copies of data at other locations either locally or remotely so that processing can continue in the event of an outage.

Backup and restore procedures and products also contribute dramatically to computer system availability by reducing the time to restore operations in the event of an outage. Hitachi has teamed with industry-leading storage management software providers such VERITAS® Software of Mountain View, California to provide world-class high-availability solutions. These solutions contribute to the fact that customers and analysts alike regard Hitachi Freedom Storage™ systems as having the highest availability in the industry. Hitachi TrueCopy 9200 and ShadowImage™ 9200 are discussed in more detail in the next chapter.

The Thunder 9200™ Has an Advanced Availability Profile Compared to Competitive Products

The Thunder 9200 was designed with maximum emphasis on high-availability computing for today's most critical enterprises. Figure 22 highlights some of the many advantages in high-availability design that the Thunder 9200™ has over other enterprise storage products.

Figure 22 – The Thunder 9200 offers higher availability than competitive products.

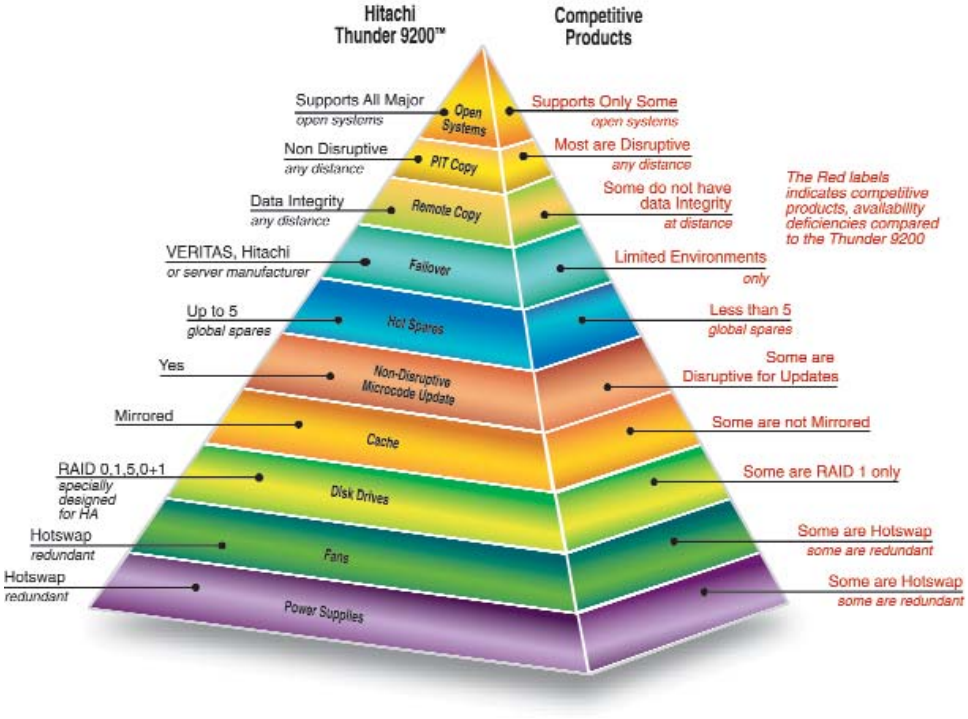

Thunder 9200 offers enterprise-level features with a midrange package and price.

- Unique Hi-Track® phone home preventive maintenance software.
- Redundant active components throughout the system.
- Dual data paths and dual control paths connecting every component.
- Mirrored cache for all write data instead of single image cache.
- Quadruple redundant microcode sparing.
- Hitachi TrueCopy 9200 for disaster recovery.
- Automatic fail-over architecture.
- ShadowImage™ 9200 executes logical backups at faster speeds.

Concurrent Maintenance

Hitachi can perform maintenance activities concurrently while continuing to provide users with access to information. Hardware components, such as disk drives, controllers in dual controller configuration, power supplies, fans, batteries, disk drives can be replaced non-disruptively.

[2] Without requiring quiescence of databases or applications.

Microcode Sparing

The Thunder 9200 stores up to five copies of the systems microcode – the active one plus four additional copies. The reason for this is to ensure that multiple copies of the microcode are available in the event of disk failure.

Non-Disruptive Microcode Upgrades

The Thunder 9200 supports both non-disruptive microcode enhancements and non-disruptive LUN expansion. In addition, scaling up capacity can be done non-disruptively, at least from a hardware standpoint. Disk expansion modules can be added to the base module and recognized by the Thunder 9200 without rebooting the system, provided that functionality is supported by the operating system in use.

Automatic Failure Monitoring and Reporting

Hi-Track® is Hitachi Data Systems monitoring feature that allows the Thunder 9200 to self-monitor its own components and provide notification of soft errors via SNMP alerts. It also has a phone-home capability that enables the Thunder 9200 to automatically initiate a data transmission to the HDS support center when it detects a failure, thus enabling the support team to take corrective action before system integrity is compromised. System status is also available down to the unit or component level from any location via web browser.

Disk expansion modules can be added and recognized by the Thunder 9200 without rebooting.

Cache Reliability

The Thunder 9200 offers important data protection:

- Mirrored—all writes are duplicated for data protection
- Battery backup (minimum, 48 hrs; maximum, more than 15 days)
- Optional integrated UPS for ultimate protection
- Data Consistency checking in memory

On-line V erify, Data Assurance Code, and Error Correction Code

The Thunder 9200 is designed to minimize the potential for a fatal data block read error via a robust, system-based, self-analysis program that continually checks the storage environment for potential data integrity issues. During idle periods, the online program reads and checks data blocks to ensure full data integrity throughout the Thunder 9200 system. If a data block exhibits the potential for a fatal error, the Thunder 9200 system automatically reassigns the data to a problem-free data block. If the Thunder 9200 determines that there are too many potential errors on a given disk drive, the dynamic spare drive is enabled and the data is safely copied to the spare drive. To maintain data integrity, a Data Assurance Code is added at the end of each data block, and the relative block number is confirmed. The Data Assurance Code includes cyclic redundancy checking (CRC) for guaranteed safe and secure data placement. CRC is compared on read and write commands to further ensure data integrity.

During idle periods, the on-line program reads and checks data blocks to ensure full data integrity.

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

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

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.

There are four basic types of high-availability middleware that can reduce downtime in the event of a data path or host failure. The Thunder 9200™ supports each classification of middleware:

1. Alternate Pathing Middleware Switches the I/O Load in the Event of Path Failure

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

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 supports alternate pathing for AIX® 4.2 and above, Compaq True64™ UNIX®, HP® 10.01 and through PVlink, Windows NT/2000® and Sun Solaris™ 2.5 and Sequent DYNIX/ptx®. In addition, VERITAS® Dynamic Path Management™ (DPM) is supported.

2 . Hitachi Dynamic Link Manager™ Provides Path Failover

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

Hitachi Dynamic Link Manager™ is a family of Hitachi-provided middleware software utilities that are server-based as shown in Figure 23. Hitachi Dynamic Link Manager™ enhances the availability of RAID systems by providing automatic error recovery and path failover from server-to-RAID connection failures.

Just because a system is RAID-protected doesn't mean it is protected against connection bus failures, which is why Hitachi Data Link Manager™ is required for true non-stop operations. This product allows systems administrators to take advantage of the multiple paths on a Thunder 9200™ by adding redundant connections between application servers and RAID systems. Hitachi Dynamic Link Manager™ therefore provides increased reliability and performance. Supported platforms include AIX®, Sun Solaris™, Windows NT®, and Windows® 2000.

3. Host Failover

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

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

The Thunder 9200 supports all major open systems clustering schemes including: Compaq TruCluster®, HP® MC Service Guard, HP® MC Lock Manager, IBM® RS6000 and SP HACMP, Windows NT/2000®, Microsoft® Cluster Server, NCR® UNIX® SVR4 Lifekeeper, IBM DYNIX/ptx ATAP®, and VERITAS ClusterServer™ for Sun Solaris™.

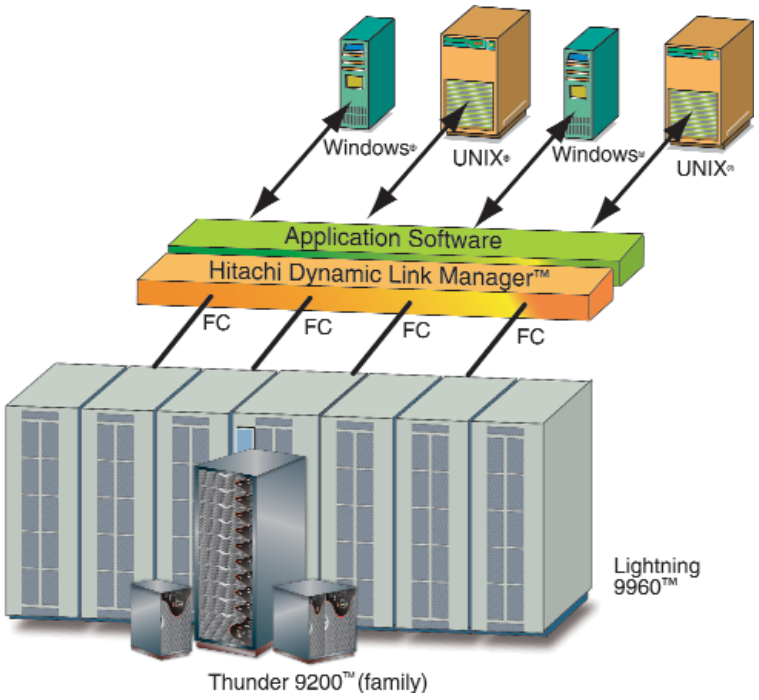

Figure 23 - Hitachi Dynamic Link Manager™ automatically provides path failover and load balancing for open systems connected to either the Thunder 9200 or Lightning 9900.

4. 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.

Standard Hot-swap and Redundant Power Supplies

The Thunder 9200 provides fully redundant power supplies to ensure uninterrupted power and cooling to all chassis in the system – supporting full system power in the event of a power supply failure. In the event of such a failure, the power supply that has failed can be "hot swapped" without disruption to the power subsystem.

Standard Redundant Fans for Cooling

The Thunder 9200 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.

Optimizing Performance, Availability, and Cost for the Thunder 9200 Target Workloads

8

Understanding the Target Workload is Critical to Optimizing Objectives

Rules of thumb can be applied to optimize the cost/performance/ availability equation.

Before determining which Thunder 9200™ configuration is optimal for a particular application, it is important to understand the workload and the computing environment. In this way general rules of thumb can be applied to optimize the cost/performance/availability equation. Although we attempt to simulate a variety of environments, it is impossible in reality to know in advance each and every application workload or the mix that will be generated by all users of the system. The items of most importance are as follows:

- What is the customer trying to accomplish?
  The more data you have, the easier it is to architect a solution.
- What are the current workload statistics?
  If none, then try to characterize the workload based on several models.
- What server platform(s) – CPUs, memory, HBAs, OS levels (including patches) will to be atached?
- What database(s) is being used, if any?
- What is the network infrastructure and average number of users?
- What are the high-availability requirements – Clusters, DMP, etc.?
- What are the capacity requirements?

Possibly the most important selection criteria for performance are the type of drives to deploy and the type of RAID to deploy for that application.

Rules of Thumb for Drive Selection

As discussed in Chapter 5, two choices are offered to configure the highest performance systems using low-capacity high-spin drives to minimize the number of GB per actuator:

18GB 15,000 RPM drives

36GB 10,000 or 15,000 RPM drives

Configure high-performance systems with high-spin drives and RAID-1 or RAID-0+1.

Two choices are offered to configure the highest capacity systems using high-capacity drives:

72 GB 10,000 RPM drives

180GB 7,200 RPM drives

Rules of Thumb for RAID selection

The selection of RAID is critical for each application to be supported. The flexibility of the Thunder 9200™ in allowing the intermixing of RAID configurations and disk types within a particular system allows high-capacity and high-performance applications to coexist in the same installation. General rules of thumb regarding the tradeoffs of RAID selection are shown in Table 6.

Table 6 – RAID Level Comparisons.

	RAID-0	RAID-1	RAID-0+1	RAID-5
Description	Data Striping NO DATA PROTECTION	Disk Mirroring	Data Striping and Mirroring	Data Striping with distributed parity
Minimum # Disks	2	2	4	3
Maximum # Disks	16	16	16	16
Benefit	High Performance for most workload models	Data Protection through redundancy	High Performance with data redundancy	The best balance of cost, reliability, and performance
Disadvantages	NO DATA PROTECTION	Higher cost per number of physical disk	Higher cost per number of physical disk	Performance penalty for high percentage of random writes

Hitachi's Uniquely Fast RAID-5 is Often the Perfect Tradeoff for Midrange Applications

Many storage providers claim that their RAID-1 provides the best information protection and performance. The Thunder 9200 offers superior RAID-5 performance, which has a performance that approaches that of competitors' RAID-1 performance but with RAID-5 value. Using an innovative combination of compatible technologies and management techniques, Hitachi Data Systems brings remarkable performance to the Thunder 9200 RAID-5 implementation, reducing both operating costs and management complexity. For mission-critical applications, the Thunder 9200 RAID-5 offers a real choice for cost-effective, high-performance, fully protected storage. In addition, with the Thunder 9200, ultimate RAID-5 flexibility and performance can be achieved through the ability to configure RAID-5 using up to 16 drives, with rotating parity.

Advanced Function Software and Services

9

Thunder 9200 Supports Enterprise-level Software in a Midrange Package and Price

The Hitachi Freedom Storage™ software solutions for Thunder 9200™ systems support an enterprise's strategic goal of accessing any information, on any computer, located anywhere, at any time. The many advanced functions available on Thunder 9200™ hardware are initiated, managed, and controlled through these powerful software programs.

The Hitachi Freedom Storage™ software solutions deliver enterprise-wide coverage of on-line data copy/relocation, data access/protection, and storage resource management. Customers have the freedom to choose the precise solution – or combination of solutions – appropriate for their environment. The Hitachi Data Systems Software Solutions supported on Thunder 9200™ Series and the business benefits they support are shown in Table 7 and Figure 24.

Hitachi Freedom Storage Software Solutions for the Thunder 9200™	Business Benefit		
	Increased IT service levels in availability and performance	Simplified IT operations	Accelerated deployment of new applications and new systems
Hitachi TrueCopy 9200 - OpenSystems	✔		
Non-disruptive Data replication - ShadowImage™ 9200	✔		✔
Automatic Path Failover - (Dynamic Link Manager™)	✔	✔	
Host Failover and Parallel Database Clustering	✔		
Hitachi SANtinel™ (LUN Security)	✔	✔	✔
VERITAS® support, volume and file management utilities	✔	✔	
Systems Management - HiCommand™ 9200	✔	✔	✔
Systems Management - Resource Manager™ 9200	✔	✔	✔

Table 7 – Hitachi Data Systems Software Solutions supported on Thunder 9200™.

Figure 24 –The Hitachi Software Solutions Guide.

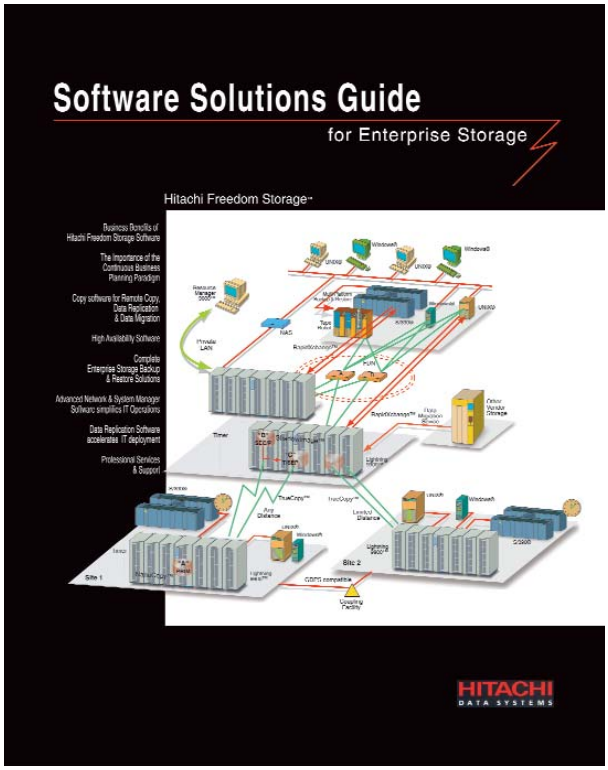

Remote Copy, Data Duplication, Data Migration Defined

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

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

Remote Copy

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

PiT Copy

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

Data Migration

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

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

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

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

There are many reasons for the proliferation of products and techniques, but the primary motivation is to improve on an existing design or technique. If an existing product is insufficient in some form, that becomes the impetus for creating a newer product. For example, early implementations of copy software would reduce backup windows from hours to minutes by not requiring the application to endure an outage for the duration of the Point-in-Time copy process. The application would still require an outage (or quiesce), but merely for the duration necessary to initiate the process (seconds or minutes). Using these products, however, the backup is not guaranteed and could subsequently fail. This led to the development of such products as Snapshot (StorageTek), ShadowImage™ (Hitachi®), and TimeFinder (EMC®).

Business Goals Served by Hitachi Copy Software

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

Real-time/point-in-time copies heighten both data security and mobility for a variety of critical activities, including protection, testing, disaster recovery, and warehousing. No matter where the data is located, it can be replicated quickly and safely within the same data center or between different data centers. This ability to share data within the enterprise positions a business to respond quickly to competitive pressures and meet the fundamental business goal of business agility. As shown in Table 8, software products provide the functions of data migration, data duplication and remote copy. These functions are critical to perhaps the two most important business objectives of the enterprise: 1) to maintain business continuity in the face of adversity, such as disaster; 2) to rapidly deploy new IT applications for business intelligence, data warehousing/OLAP, for data center relocation, or for new application testing.

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

Table 8 - Hitachi copy products provide Remote Copy, Data Duplication, and Data Migration.

Hitachi Copy Products	Data Migration	Data Duplication	Remote Copy
Hitachi TrueCopy Synchronous – 9200 and 9900	✔	✔	✔- 25 Miles[3]
Hitachi TrueCopy™ Asynchronous – 9200 and 9900	✔	✔	✔- Unlimited distance
NanoCopy™ – 9900 Only		✔	✔
Software Asynchronous Remote Copy™ (HXRC) – 9900 Only		✔	✔
Hitachi ShadowImage™ – 9200 and 9900		✔	
Hitachi Data Protection Services – 9200 and 9900	✔	✔	✔
Hitachi RapidXchange™ – 9900 Only		✔	
HARBOR® File Transfer™ (HFT) – 9900 Only	✔		

Hitachi ShadowImage™ 9200

Among the powerful products created to complement the robustness of the Thunder 9200, ShadowImage 9200™ enables open systems information to be replicated in innovative ways to meet ever-changing business challenges. The Thunder 9200 point-in-time (PiT) copy facility, also used on the higher-end Lightning 9900™, is called ShadowImage™ 9200. It is now possible to execute logical backups at faster speeds and with less effort than previously possible. In addition backups can be easily configured across your Storage Area Network (SANs), and can be managed from a central location.

Hitachi TrueCopy 9200

Hitachi TrueCopy 9200 provides both synchronous or asynchronous remote copy capabilities for open system computers. This allows remote copies over virtually unlimited distances. Operating systems that are supported include HP/UX, AIX, Sun Solaris™, Digital UNIX®, Sequent DYNIX/ptx®, SGI IRIX™, NCR®, UNIX SVR4®, Windows NT/2000®. Hitachi TrueCopy is illustrated in Figure 25.

[3] 25 miles is the current practical distance limitation for direct connection using ESCON® fibre. This is changing with the advent of new networking devices. Synchronous remote copy may also use telephone networks. However the performance of the production application must be considered when using synchronous remote copy over long distances.

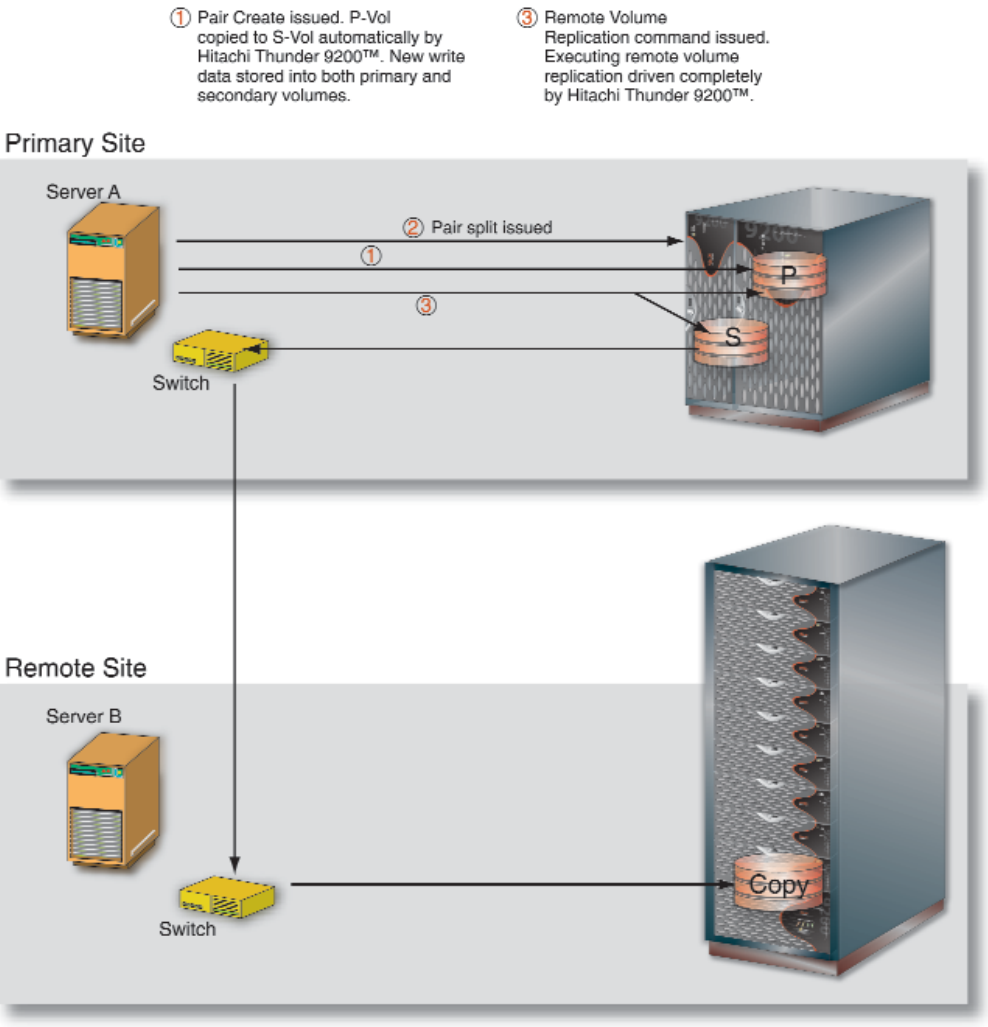

Figure 25 - Hitachi TrueCopy 9200 copies data to remote sites for disaster recovery.

Hitachi Dynamic Link Manager™

Hitachi Dynamic Link Manager™ is a family of software utilities that is server-based and enhances RAID systems by providing automatic failover and load balancing from server-to-RAID channel connection failures. This product allows systems administrators to take advantage of the multiple paths on a Lightning 9900™ or Thunder 9200™ by adding redundant SCSI connections between data servers and RAID systems. Hitachi Dynamic Link Manager™ therefore provides increased reliability and performance. Supported platforms include AIX®, Sun Solaris™, and Windows NT/2000®.

Hitachi SANtinel™

The Hitachi SANtinel™ software controls host access to Hitachi Lightning 9900 or Thunder 9200 LUNs in SAN environments.

VERITAS® Foundation Suite

The VERITAS® Foundation Suite™ supports both the Hitachi Lightning 9900™ or Thunder 9200™ Series. The suite combines the VERITAS® Volume Manager and VERITAS® File System™ to simplify storage management while ensuring high availability for critical systems. VERITAS Volume Manager™ is an easy-to-use, on-line disk and storage

management tool for enterprise computing environments. VERITAS File System™ is an enterprise-class, journaling file system that facilitates high performance, quick recovery, and easy scalability. The VERITAS Foundation Suite™ provides a strong foundation, not only for managing current storage requirements, but also via its open architecture for leveraging future technologies and opportunities. Downtime, whether scheduled or unscheduled, is costly. The VERITAS Foundation Suite™ helps decrease the total cost of operations and brings greater stability to vital information systems.

Hitachi Resource Manager™ 9200

Resource Manager 9200™ is shipped with every Thunder 9200™ Series system and is used by storage managers to display system configuration, create user name and password security, set up RAID groups, allocate LUNs, expand LUNs and format storage. Resource Manager 9200 is discussed in greater detail in the next chapter.

Reduced Complexity and Management Costs

10

A Powerful and Open Systems Management Philosophy

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

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

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

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

Hitachi Data Systems® software and hardware solutions are managed through the powerful HiCommand™ management framework. This powerful management tool reduces operations expense and increases business agility while enabling operational excellence. HiCommand™ includes a GUI-based browser, open APIs, Command Line Interface (CLI), server, and host agents that access HDS software management functionality. The program encompasses storage resource management, configuration management and automation, automated data replication and recovery, performance management and optimization, and related functionality. Most importantly HiCommand™ can plug into other management frameworks. This feature provides the utmost in an open architecture and leverages an enterprise's existing investments in software and skills. The HiCommand™ architecture is shown in Figure 26.

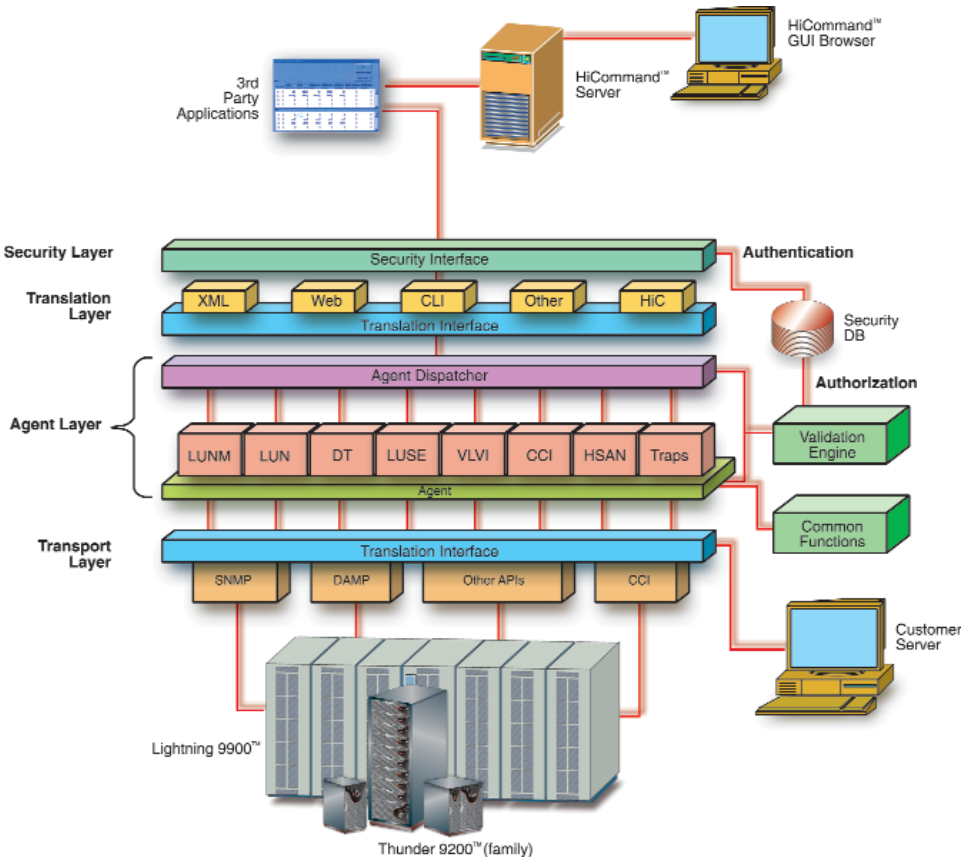

Figure 26 – HiCommand allows management of virtually all Hitachi hardware and software from the customer's platform of choice.

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

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

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

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

The Components of HiCommand

HiCommand consists of three components:

1. The **HiCommand™ Server** resides on any Windows NT®/Windows 2000®, or Sun Solaris server, and connects to the Lightning 9900™ Series system(s) and Thunder 9200 over a private LAN. A single HiCommand server can manage multiple Lightning 9900, Thunder 9200 and Sun StoreEdge™ 9900 Series storage systems, providing a common point of control that delivers:
   - Out of band storage area networks (SANs), networked attached storage (NAS), or direct attached storage (DAS) connections enable data to run between the server and storage at lightning speeds.
   - Seamless integration between HiCommand™ and software solutions from industry leaders such as VERITAS® Software, Sun Microsystems®, BMC Software®, IBM®, Tivoli®, and Computer Associates® through use of XML-based APIs.
   - Remote and local access control offers a centralized approach to managing storage from any location.
2. A **GUI, browser-based interface** that is easy to learn and allows for simple and efficient storage administration. A browser from anywhere can control the HiCommand server. In addition, HiCommand also supplies an optional CLI (Command Line Interface).
3. **Host Agents** reside on the customer's application servers and "push" information back to HiCommand, including percentage of file utilization of LUNs, HBA WWNs, operating system SCSI addresses, and other useful device specific and file level information.

Out of band storage area networks (SANs) enable data to run between the server and storage at lightning speeds.

Overview of How the Systems Management Software Supports Business Objectives

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

Hitachi Resource Manager™ 9200 provides the ultimate in reliable, easy-to-use hardware performance and availability management software.

Hitachi Resource Manager™ 9200

The complexities of heterogeneous computing environments make storage management a daunting task. Hitachi Data Systems offers the ideal solution to this difficult challenge, Hitachi Resource Manager™ 9200. This comprehensive systems and software management offering brings together Hitachi FlashAccess™, SNMP support, and Hitachi LUN Manager into one complete package. Designed for the Thunder 9200, these diverse capabilities provide world class functionality.

Shipped with every Thunder 9200 system, 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™ 9200 an ideal choice. Users can monitor and manage their storage systems through a graphical user interface.

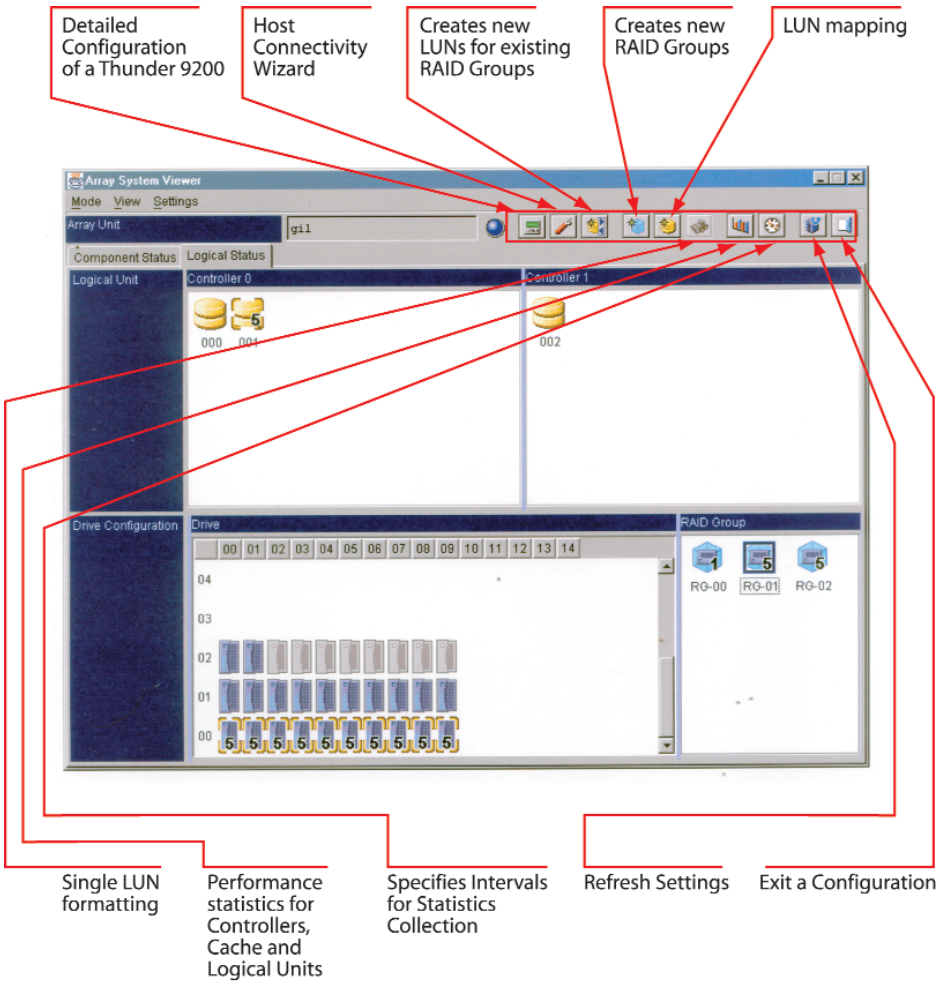

Figure 27 – The GUI interface on the Thunder 9200 Resource Manager simplifies storage and performance management.

Resource Manager 9200 helps 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 up to 20 users. This provides protection against unauthorized access to the Management mode of Resource Manager 9200, 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.

Hitachi Resource Manager 9200 Simplifies Operations

The following is a summary list of the many ways in which the Resource Manager 9200 can simplify IT operations, improve availability and performance service levels for business critical OLTP and DSS applications, and thereby enable the more rapid deployment of new applications by managing data more intelligently.

Service Level Management

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

Deployment Management

- Asset Management
- Change Management
- Capacity Planning

Compliance Management

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

Hitachi Hi-Track® and Web-enabled Monitoring via any Browser

Hitachi Hi-Track® provides the ultimate in enterprise-level error monitoring and call home preventive maintenance software. The World Wide Web capabilities of Hitachi Hi-Track® deliver the power to monitor storage – anywhere, from any computer, at any time. Users gain point-and-click access to worldwide storage, in addition to vivid visual aids to help with error management decision-making and the ability to monitor at the unit or component level via any Web browser.

Hitachi Resource Manager 9200 Simplifies Configuration Management and Reduces Staffing Costs

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

Hitachi Resource Manager 9200 is an open systems management utility. With Resource Manager 9200, open systems Logical Units (LUNs) can be defined, configured, and maintained. There is no more waiting for the hardware vendor to come and make configuration changes. Resource Manager 9200 includes an easy-to-use, GUI-based interface that allows the easy definition of paths for LUNs, the reconfiguration of LUN-to-port assignments, or the viewing of the Thunder 9200 remote service information messages. Because the Resource Manager 9200 can assign multiple paths to a single LUN, support of alternate path failover, and clustered systems is possible. Running on a standard workstation connected to the storage subsystems by a dedicated LAN, Resource Manager 9200 can support up to 1024 Thunder 9200 systems.

Resource Manager 9200 also features logical unit size expansion that dramatically improves LUN flexibility. 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 Thunder 9200 LUN configuration information, while administrators can access all LUN configuration information and functions. Administrators can customize access privileges for individual users, thus providing maximum flexibility and control of Hitachi LUN Manager capabilities.

Hitachi FlashAccess™ 9200 Allows Data to be "Locked and Unlocked" in Cache on-the-fly

Hitachi FlashAccess™ dramatically improves performance.

Hitachi FlashAccess™ 9200 is a software utility that allows users to dynamically "lock" and "unlock" data into cache. Read and write functions are then performed at cache speeds, with no disk latency delay. With FlashAccess 9200, a portion of cache memory can be allocated to specific data. Defined by Logical Unit Number (LUN) for open systems, cache data can be as small as a single track or as large as 900 MB.

Hitachi SANtinel™

The Hitachi SANtinel™ software controls host access to Thunder 9200 LUNs in SAN environments.

Out-of-band SAN Management

Traditionally, network managers have relied on in-band management, where management traffic, such as SNMP commands, is mixed with the data and sent along the primary data path. This causes contention for network bandwidth and makes management and traffic information unavailable if the network is down. SANs cannot afford the negative fallout of in-band management. Out-of-band SAN management is the answer.

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

The Thunder 9200™ system fits squarely into the out-of-band management architecture for storage area networks, which allows all status activities and management traffic to occur off the data path. This ensures data security, data accessibility, and data manageability, even when a data path is down, without compromising network performance – a critical business advantage.

SNMP Compliance

In order to participate in a managed SAN environment, all hubs, switches, directors, and storage devices must be capable of supporting SNMP. Developed to solicit status and set operating parameters for basic availability management in a network, SNMP is an industry-standard IP-based protocol supported by multi-vendor platforms. The Thunder 9200 is fully compliant with SNMP. This means that you can integrate the Thunder 9200 into existing VERITAS®, CA Unicenter TNG®, IBM Tivoli®, HP® OpenView®, and Sun® Management Center network environments. The Thunder 9200 SNMP compliance makes it truly SAN-enabled.

Professional Services and Support

11

Hitachi Data Systems® is Consistently Ranked Number One in the Industry

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

Professional Services Overview

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

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

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

Cost of Risk Analysis Methodology (CORA)

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

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

Data Protection Services

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

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

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

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

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

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

Data Protection Services Lab

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

Remote Copy Assessment and Implementation Service

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

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

Continuous Availability Service

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

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

Recommendations are then discussed to improve availability and reduce the risk of data loss.

On-line Data Migration Service

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

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

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

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

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

SAN Services

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

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

The main benefits of SAN implementation are:

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

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

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

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

Hitachi Data Systems Interoperability Laboratory Service

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

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

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

Enterprise Storage and Availability Management Services

Hitachi Data Systems offers a variety of consulting services and software utilities to help extend the life of an enterprise's storage investment and define ways to achieve greater information functionality throughout the enterprise. These storage services cover storage management, availability, disk utilization, performance and tuning, data center cabling and configuration, and design/installation of fiber optic components. Hitachi Data Systems services help the enterprise take full advantage of the enterprise's IT resources, and enhance the continuous availability and integrity of the enterprise's mission-critical applications. In addition, Hitachi Data Systems Professional Services consultants can work with the enterprise's business continuity planners to exploit the capabilities of the Lightning 9900 and the Thunder 9200 storage systems.

Software Portfolio Review and Analysis Service

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

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

Disaster Recovery Institute A

The Disaster Recovery Institute (DRI) was formed as a non-profit organization in 1988 to provide a base of common knowledge in contingency planning, a rapidly growing industry. Today DRI administers the industry's only global certification program for qualified business continuity/disaster recovery (DR) planner professionals. The Professional Practices for Business Continuity Planners, and the DRI knowledge base, serve as the industry's best practices standard. DRI's acclaimed training courses educate and inform business continuity and disaster recovery planners worldwide.

The DRI knowledge base serves as the industry's best practices standard.

The DRI certifies professionals at three levels as shown in Table 9.

Level	Professional Certification	Years experience as a business continuity / DR planner
Entry	Associate Business Continuity Planner (ABCP)	Less than two years
Intermediate	Certified Business Continuity Professional (CBCP)	Two years minimum
Advanced	Master Business Continuity Professional (MBCP)	Five years minimum

Table 9 – Professional certification levels of the Disaster Recovery Institute (DRI).

The Certified Business Continuity Planner professional, or CBCP, was formerly known as the Disaster Recovery Planner (CDRP). Since the DRI is a not-for-profit corporation, it promotes both credibility and professionalism in the DR industry. CBCP is the intermediate level of certification by the Disaster Recovery Institute, and professionals certified at this level must meet a minimum of two years of experience as a business continuity or disaster recovery planner. The entry level of professional certification is the Associate Business Continuity Planner or ABCP, and the advanced level is the Master Business Continuity Professional or MBCP. In addition to the experience requirements, a passing grade at an extensive certification examination is required.

An excellent source of information on disaster recovery and business continuity is the ***Disaster Recovery Journal (DRJ)***, which today has over 50,000 subscribers and is over 1000 pages in length. The ***DRJ*** has an informative and useful Web site at http://www.drj.com/ that offers many products such as:

The Disaster Recovery Journal is an excellent source of information on disaster recovery and business continuity.

- CD-ROMs
- Reference books
- Disaster recovery videos on floods, hurricanes, earthquakes, and lessons learned
- Business continuity training videos

The ***DRJ*** also sponsors two annual conferences that attract 2,500 disaster recovery professionals from all over the world. This makes the conferences the largest in the industry. In addition to conferences and products, the ***DRJ*** has an on-line subscription service for a monthly fee and is an excellent resource for business continuity information.

Hitachi Continuous Business Planning Questionnaire

B

1. Storage

1.1. Storage System Performance

1.1.1. How confident are you in the ability of your disk storage system configuration to handle a surge in I/O loads without severely impacting application performance?

1.1.2. Is your current disk storage system configuration delivering the type of data access performance that you expect?

1.1.3. Is your current tape storage system configuration delivering the type of data access performance that you expect?

1.1.4. How would you rate your ability to centrally manage multi-platform (S/390, UNIX®, Windows NT®) storage performance?

1.2. Scalability

1.2.1. Have you been able to add storage system capacity without disrupting service and negatively impacting application and data throughput performance?

1.2.2. Do you have difficulty in meeting storage growth needs due to lag time of storage acquisition? Have you considered storage-on-demand?

1.2.3. How would you characterize your data by platform (OS/390, UNIX®, Windows NT®): constant, growing, or shrinking?

1.2.4. What is your projected growth rate of DASD storage?

1.2.5. What processes are in place to monitor and evaluate new storage technologies?

1.3. Ease of Storage Maintenance

1.3.1. Please describe the *level of difficulty* in performing routine maintenance (microcode updates, adding cache or HDDs, etc.) and the impact on service and the need for scheduling downtime.

1.4. Storage RAS

1.4.1. As measured by Mean-Time-To-Recovery (MTTR), how would you rate the serviceability of your storage systems?

1.4.2. How reliable do you feel your storage systems are, and how do they perform over time (as expressed in MTBF, and in Mean-Time-To-Data-Loss - MTDL = length of the expected continuous span of time over which data stored on RAID can be correctly retrieved)?

1.4.3. Are you satisfied with the current levels of availability of your storage systems?

1.4.4. How confident are you that the storage configuration has the ability to provide timely, continuous access to reliable data under abnormal conditions (internal/external failures, environmental failures, etc.)?

1.4.5. Do you currently have an infrastructure that provides instantaneous failover in the event of a hardware outage?

1.4.6. What is the current maximum tolerable outage threshold for downtime, scheduled or unscheduled, that your business can tolerate?

1.5. Storage Connectivity

1.5.1. What is your primary connection architecture for storage now (ESCON, SCSI, Fiber, FICON, FC-AL)? What plans exist for the future? Is there an architectural preference for the disk I/O?

1.5.2. Based on your current storage connectivity, what is the level of difficulty in being able to install new storage systems that can easily interface with multiple servers, with robust connectivity, across your platforms?

1.6. Data Backup

1.6.1. What is your current methodology (procedures and processes) for taking backups of data?

1.6.2. Do you periodically test to ensure that data on the *backup tapes* is good?

Do you currently have concerns or issues regarding your tape systems?

1.6.3. For point-in-time copies (backups), can you tolerate an application outage and, if so, for how long?

1.7. Data Recovery

1.7.1. Do you maintain an alternate internal site? If yes, has your primary data been mirrored to this site?

1.7.2. For business recovery purposes, please rank (from 1 to 5) what is most important to you:

__ No impact to performance at primary site

__ Data integrity/consistency

__ Ease of use

__ Cost of the solution

__ Distance limitations

1.7.3. How confident are you that you can recover your business applications at the alternate site:

From a copy of production data at primary site?

Without restoring your data from backup tapes?

1.8. Data Access Management

1.8.1. How do you currently manage the movement of data among your S/390, UNIX, and Windows NT platforms? (For example, moving S/390 data to UNIX in support of data warehousing, etc.)

1.9. Storage Manageability

1.9.1. What tools (Legato, VERITAS, etc.) do you use for volume management? Are these tools integrated with Enterprise Resource Management tools (CA, Tivoli, BMC, SMS®, etc.)?

2. Support Software

2.1. Transaction Workload Balancing

2.1.1. Currently, what method of transaction routing are you using for your on-line transaction managers, static or dynamic? Why?

2.1.2. Have you had requirements to make use of cloning (replication) your database managers across your system images?

2.2. Data Sharing

2.2.1. Have you implemented *database data sharing? (Check which you have applied.)*

__ None

__ S/390 platform only

__ Across all platforms

2.2.2. What were the requirements for and what is the extent of data sharing in your organization?

2.3. Maintenance/Service Strategy

2.3.1. Describe your current maintenance methodology. Does it lean more towards one of fix-on-fail or the traditional application of preventive maintenance?

2.3.2. Do you install preventive maintenance on a regular basis in an effort to avoid problems, or do you only install service when required to resolve a specific problem?

2.4. Quality of Testing

2.4.1. Have you configured a testing environment for testing software procedures, policies, and new releases of software? Describe this environment.

2.4.2. Do the test cases accurately reflect the production environment? (If so, what process do you have in place?)

Are the test cases kept up to date?

2.5. Performance

2.5.1. Are you satisfied with current levels of transaction performance from your on-line transaction managers, and is performance meeting the stated SLAs?

2.5.2. Is the overhead inherent in data sharing currently impacting performance of your critical applications? Are you measuring this overhead?

2.6. Data Synchronization

2.6.1. How do you approach data synchronization across your database subsystems/platforms?

2.7. Change Management

2.7.1. How do you implement and propagate software changes across all systems (OS/390, CICS, IMS, and DB2)?

2.7.2. Do your critical applications require the latest software functionality as quickly as possible?

2.8. Problem Management

2.8.1. Have you recently experienced software problems that impacted application availability? If so, were fixes readily available, or were these new software problems?

3. Management Practices

3.1. Quality Management

3.1.1. How is your quality management helping you produce quality products and services for your customer?

3.1.2. How have you implemented measurement and performance of your processes and service level standards?

3.1.3. How have you documented all your processes?

3.2. Change Management

3.2.1. How have you implemented a formal change management process?

How do you link it to problem management?

3.3. Problem Management

3.3.1. How do you report successes or problems with changes?

3.4. System Life Cycle

3.4.1. What process have you implemented to ensure proper design, testing, and implementation of all new and maintained business applications?

3.5. Vendor Management

3.5.1. How are you measuring your current support by your vendors?

3.5.2. How are you communicating the evaluation of that support to your vendors?

3.6. Operations Management

3.6.1. How is your operations management team managing continuous business?

4. FACILITIES

4.1. Physical Security

4.1.1. What formal physical security processes and procedures are in place?

4.1.2. Do these processes cover all sites to include any alternate failover or recovery site?

4.1.3. What formal education and awareness process is in place?

4.1.4. What reporting process is in place to identify irregularities?

4.2 Logical Security

4.2.1. What formal logical security processes and procedures are in place?

4.2.2. Does the process include review of new technologies and tools?

4.2.3. Do these processes cover all sites to include any alternate failover or recovery site?

4.2.4. What formal education and awareness process is in place?

4.2.5. What reporting process is in place to identify irregularities?

4.3. Environmentals

4.3.1 What environmental planning, monitoring, and maintenance process is in place?

4.3.2 Does it include review of new technologies?

4.3.3 Is there an environmental review process in place for the continuous business methodology?

4.4 Power

4.4.1. What is your strategy for protection of the power to the continuous operation of the business? (For example, do you employ dual power feeds from different grids, diesel generators, etc.)

4.4.2. What contingency plan is in place for continuous operation if all the mitigation strategies fail?

4.5. Real Estate Structure

4.5.1. Is your facility prepared for an outage, community disaster, or regional disaster?

4.5.2. What contingency plan is in place for continuous business?

4.6 Regional Location

4.6.1. Does your facility fall within areas prone to calamitous events, such as earthquakes, tornadoes, hurricanes, snow/ice storms, nuclear reactors, fuels/chemical or hazardous manufacturing, or hazardous shipping lanes, civil unrest, prisons, etc.?

5. Communications

5.1 Network management

5.1.1. Is your network IP only, SNA over IP, or a convergence of SNA/IP?

5.1.2. Which network transport are you using for SNA and IP integration: DLSW, Frame Relay, or Enterprise Extender?

5.1.3. To what extent do you currently use APPN within your network?

Is it pure APPN, pure Subarea, or mixed APPN/Subarea?

5.1.4. Have you employed a Communication Management Configuration (CMC)?

5.1.5. If a CMC configuration is in place, do you utilize multiple CMC hosts, allowing for quick recovery of resources?

5.1.6. Are consistent naming conventions used within the SNA and TCP/IP networks?

5.1.7. Are symbolic symbols and wild cards used to ease management of the network subsystems?

5.1.8. Do you feel that your current network topology is well documented?

5.1.9. Are automated network discovery tools, such as Netsleuth or other similar tools utilized?

5.1.10. Do you utilize network-monitoring tools, and how effective are they?

5.1.11. What level of network automation do you employ?

5.1.12. What is your network technology adoption style? Is it early, stable technology only, or cutting edge?

5.1.13. What processes are used to monitor and analyze new technologies?

5.2. Network Reliability/Availability/Serviceability/Failover

5.2.1. How do you feel about the current levels of availability of your network and the degree to which the network has resiliency against failures?

5.2.2. Do your servers on the Wide Area Network (WAN) have multiple redundant connections through multiple gateways?

5.2.3. Do you employ a backbone network that will allow continuous availability if one of its components, such as a concentrator or link, fails?

5.2.4. If a component, such as a concentrator or link fails, will your backbone network still allow for continuous availability?

5.2.5. Do you employ major redundant lines, thus allowing for diversion of network traffic to another path?

5.2.6. Are your user workstations at user locations (which are usually interconnected through a LAN) configured to provide fault tolerance via redundant connections?

5.2.7. What is the level of implementation of High Performance Routing (HPR) in your network?

5.2.8. Have you implemented Persistent Sessions?

5.2.9. Have you implemented Multi-node Persistent Sessions?

5.2.10. Is there a single point of control for ownership of network resources?

5.2.11. Is Virtual IP implemented (VIPA)?

5.2.12. If VIPA is implemented, what type of takeover is being utilized, static or dynamic?

5.2.13. Are redundant routers employed?

5.2.14. For Web access, does your ISP provider have dual access paths to avoid loss to the Internet backbone, or do you have links to separate ISPs, thus allowing for redundant access to the Internet backbone?

5.3. Network Reliability/Failover

5.3.1. Where do statistics tell you that most network failures occur?

5.3.2. Do you actively manage with statistics?

5.3.3. How or where do you obtain this information?

5.3.4. Do you feel that adequate failover processes are in place?
If so, what percentage of these failover processes are automated?

5.3.5. Are manual processes in place for continuous network operation if the automation routines for network failover fail?

5.4. Network Capacity/Performance

5.4.1. Do you measure your network performance and how do you report this?

5.4.2. What level of capacity on demand is available for your network components?

5.4.3. Are you achieving your performance goals for your network as specified in your Service Level Agreement (SLA)?

5.5. Workload Balancing

5.5.1. Have you implemented any form of workload balancing in the network, such at VTAM Generic Resources or user-written programs?

5.5.2. Have you implemented any form of IP workload balancing within your TCP/IP network?

5.5.3. Have you implemented any load balancing solutions within your IP network, such as DNS/WLM, Network Dispatcher, NAT, or Cisco's MNLB?

5.6. Network Connectivity

5.6.1. Detail your network connections to remote site recovery site, if applicable.

5.6.2. Are there existing fiber pairs between your buildings or between the floors of your building?

5.6.3. Do you use dark fiber (optical fiber dedicated to the customer)?

5.6.4. Do you have any optical networking in place or in the planning stages, such as DWDM or extended System Area Network (GeoSAN)?

5.6.5. Characterize the number and locations of remote offices or partner offices. Document the network capabilities between the sites.

5.6.6. How might your WAN technology be used to exploit remote copy and remote site replication?

5.7. Modification of Single Points of Failure

5.7.1. What review processes do you use to identify single points of failure?

6. Server Clustering

6.1 Server Clustering Performance

6.1.1. What utilities are used to measure the performance of each server?

6.1.2. Are you achieving your performance goals as outlined in your Service Level Agreements (SLA)?

6.2. Server Clustering Capacity

6.2.1. How many system images are supported by each server?

6.2.2. How much of your computer resources are being used? (CPU, Processor Storage, I/O)

6.2.3. Is spare memory available in your current environment (in the form of spare memory chips)?

6.2.4. Do you have the capability to move resources dynamically from one partition to another?

6.2.5. Do you have enough capacity to support future growth?

6.2.6. What are the expected growth rates?

6.2.7. When will the demands on current resources impact service levels?

6.3 Server Clustering Reliability/Availability/Serviceability (RAS)

6.3.1. Does each server have a dual power feed?

6.3.2. Does each server have redundant power supplies?

6.3.3. Have you identified any elements that lack redundancy?

6.3.4. Have you identified any elements that are potential single points of failure?

6.3.5. How often are your servers unavailable due to scheduled or unscheduled outages?

6.3.6. If currently partitioned, can processor storage be reconfigured dynamically?

6.3.7. Can maintenance and upgrades be performed concurrently on each server?

6.4. Server Clustering Security

6.4.1. How is access to the "hardware management console" controlled?

6.4.2. Are there formal processes in place for creating/deleting security user access profiles for each server platform?

6.5. Server Clustering Connectivity

6.5.1. What type of connectivity is used between the servers?

Glossary of Terms

10BaseT

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

100BaseT

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

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.

Alternate Pathing

Alternate pathing, path failover support, or I/O path switching is provided by middleware that resides on the server and automatically switches the I/O load on a failed primary path to an alternate path on the same host system. This feature is supported by both Lightning 9900™ Series and Thunder 9200™ systems.

API

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

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, 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:

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.

b

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

B

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

Back-end

In reference to storage arrays, the back-end includes the controllers, disk drives, and paths to the disk drives. On the Lightning 9200™ Series these are the disk controller chips, the Fibre Channel loops, and the Fibre Channel disks.

Business Continuity Planning (BCP)

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

Business Impact Analysis (BIA)

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

Business Interruption

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

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.

Client/Server Architecture

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

Cluster Failover

Both the Lightning 9900 Series and Thunder 9200 systems software supports the major popular open systems cluster computing environments so that one host

automatically and seamlessly takes over the work of a failed server in a cluster by reassigning networks and peripherals and assuring cache coherency without workload interruption.

Cold Site

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

Command Module

A basic building block of the Thunder 9200 consisting of a controller and ten disks. Together with Expansion Modules containing ten disk drives each, storage systems as large as 17.6 TB can be created using 180GB drives. See also Expansion Module.

CORA (Cost of Risk Analysis)

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

Cost Benefit Analysis

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

CPU

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

CTL

Controller board.

CTQ

Command-Tag Queuing. The Thunder 9200 supports CTQ. Up to 128 commands can be queued per port, for a maximum of 256 per Controller when using Fibre Channel host interfaces. This allows the controller to use algorithms in the drives for best organizing and prioritizing the access to the drives to minimize head seeking and to optimize performance.

Data Availability

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

Data Copy

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

Data Duplication

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

Data Migration

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

Disaster Recovery

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

Downtime

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

DR

Disaster Recovery.

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.

ECC

Error correction code.

Electronic vaulting

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

Ethernet

A Local Area Network (LAN) protocol developed by Xerox® in cooperation with Digital Equipment and Intel in 1976. Ethernet supports a star or bus topology and supports a data transfer rate of 10 megabits per second or 10 Mbps. The Ethernet

specification formed the basis of the IEEE 802.3 standard, which specifies the physical and lower software layers. Ethernet uses the CSMA/CD access method for handling simultaneous demands and is one of the most widely implemented LAN standards. Ethernet is also known as 10BaseT. See also Fast Ethernet, Gigabit Ethernet.

Expansion Module

A basic building block of the Thunder 9200 with a capacity up to 1.8 TB using ten 180 GB drives. Together with Command Modules containing a controller and ten disk drives each, storage systems as large as 17.6 TB can be created using 180GB drives. See also Command Module.

Fabric

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

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 100 Mbps standard that is compatible with existing 10BaseT installations, preserving the CSMA/CD media access control (MAC) protocol.

Fast write

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

FC

Fibre Channel.

FC-AL

Fibre Channel – Arbitrated Loop.

Fibre Channel

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

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

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.

Freedom Data Networks (FDN)

Freedom Data Networks is Hitachi Data framework for network storage solutions. Freedom Data Networks provides an open architecture that leverages networking technologies and prtocols, offering organizations the freedom of choice in deploying data management, protection, access and sharing capabilities across the enterprise. Current offerings enable SAN, NAS, and IP storage. With Freedom Data Networks, customers gain a powerful new tool that enables the consolidation of servers and storage, the formation of large pools of sharable storage, increased data availability, seamless scalability, centralized storage management, and the ability to backup and migrate data without affecting the performance of enterprise networks.

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 Thunder 9200™ systems, the front end consists of Host Interface Cards and/or the interface logic on the controller board.

Gigabit Ethernet

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

Gigabyte

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

GUI

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

HDD

High-density disks. Used in the Lightning 9900™ Series and the Thunder 9200™ storage.

HDmS

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

Hitachi Dynamic Link Manager™

Hitachi Dynamic Link Manager™ is a family of software utilities that is server-based and enhances RAID systems by providing automatic failover and load balancing from server-to-RAID channel connection failures. This product allows systems administrators to take advantage of the multiple paths on a Lightning 9900™ or Thunder 9200 by adding redundant SCSI connections between data servers and RAID systems. Hitachi Dynamic Link Manager™ therefore provides increased reliability and performance. Supported platforms include AIX®, Sun Solaris™, and Windows NT®/2000.

Hitachi FlashAccess™

FlashAccess™ is a Hitachi feature that allows specified (usually high access) data sets to be "pegged" or permanently placed in cache memory so they are not managed by the data movement algorithms of the Freedom Storage® system. The Hitachi Flash Access™ feature in the Lightning 9900 can be used for either S/390 or open systems. Similarly the Hitachi Flash Access™ feature in the Thunder 9200 can be used for supported open systems platforms. Hitachi Flash Access™ is a software utility in the Freedom 9000 Resource Manager™ for the Lightning 9900 and the Resource Manager™ 9200 that allows the creating, deletion, and monitoring of data managed by the Hitachi Flash Access™ software. See also Hitachi Resource Manager™.

Hitachi Graph-Track™

Graph-Track™ is a software utility in the Resource Manager™ suite that allows a robust set of system and network management utilities and provides graphical reports for performance and availability and configuration management. This feature is only supported on Lightning 9900™ series systems.

Hitachi LUN Manager

Hitachi LUN Manager is a software utility in the Resource Manager™ 9000 and in the Resource Manager 9200™ suites that allows for complete systems management of LUNs. See also LUN, Resource Manager, Resource Manager 9200.

Hitachi SANtinel™

The Hitachi SANtinel™ software controls host access to Hitachi Freedom Storage Thunder 9200, and Lightning 9900™ LUNs in open systems or SAN environments.

Hitachi ShadowImage™

The Hitachi ShadowImage™ is a firmware-based software copy utility that uses command-line-interfaces to create multiple copies of a volume for both Lightning 9900 systems and Thunder 9200 systems. Graphic or command-line interfaces control data replication and fast resynchronization of logical volumes.

Hitachi TrueCopy

Hitachi TrueCopy provides synchronous or asynchronous remote copy capability for open system and S/390 computers. Available for both Lightning 9900 Series and Thunder 9200 systems Hitachi TrueCopy allows remote copies over virtually unlimited distances. Operating systems that are supported include MVS (Lightning 9900 Series only), HP/UX, AIX, Sun Solaris™, Digital UNIX®, Sequent DYNIX/ptx®, SGI IRIX™, NCR®, UNIX SVR4®, Windows NT®/2000.

Host Interface Cards

The Thunder 9200 supports four types of front end host interfaces: 100 MB/sec Fibre Channel, 2 Gb/sec Fibre Channel, 40 MB/sec SCSI, and 80 MB/sec SCSI.

Hot Site

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

Hot Spare

A hot spare is an extra disk that is kept running (hot) but not populated with data until another disk in the array fails. If a disk fails in an array, the data is mapped by copying (RAID-1) or by rebuilding from parity (RAID-5) at priorities selectable by the user, This feature is supported by both Lightning 9900 Series and Thunder 9200 systems.

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.

I/O Path Switching

See alternate pathing.

IP

The IP (Internet Protocol) is the underlying protocol for routing packets on the Internet and other TCP/IP-based networks. IP is an internetwork protocol that

provides a communication standard that works across different types of linked networks, for example Ethernet, FDDI, or ATM.

Java™

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

Journaling

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

LAN

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

Lightning 9900 Series

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

Logical Unit

The SCSI term for a logical disk drive.

Logical Unit Number

See LUN.

Logical Volume

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

LUN

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

LUSE

LUN Size Expansion feature. This Lightning 9900 and Thunder 9200 feature allows standard-size LUNs to be combined to create larger LUNs.

LRU

Least Recently Used. A policy for a caching algorithm that chooses to remove the data from cache which has the longest elapsed time since its last access. Least Recently Used algorithms are used in all major caching systems. The Lightning 9900 Series and Thunder 9200 LRU scheme keeps a table (in separate non-volatile memory) that chronicles the frequency of use of data in cache memory.

MAN

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

MBCP

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

MIB

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

MIPS

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

Meta-data

Meta-data or data about data, is often an excellent candidate for permanent placement in cache (using Hitachi FlashAccess) for either the Lightning 9900 Series or Thunder 9200.

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.

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.

Multi-platform Connectivity

A term that refers to the ability of a storage device to connect to computer systems with different operating systems (platforms). The Thunder 9200 provides multi-platform connectivity to UNIX®, Windows NT®/2000, and Linux® platforms.

NAS

Network Attached Storage.

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. NDMP is supported by both the Lightning 9900 Series and Thunder 9200.

Node

See Fibre Channel.

NVM

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

NVRAM

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

OLAP

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

OLTP

On-line Transaction Processing.

Off-site Storage Facility

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

On-line System

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

Open System

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

Operating System

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

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

Out of Band

A communication that does not use the same bandwidth that carries data in a system. For example, the control information in the Lightning 9900 Series systems and the Thunder 9200 systems do not use the same path as data and is therefore referred to as "out of band." This optimizes SAN performance.

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.

Path Failover

See alternate pathing.

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.

Platform

A term that refers to a computing device and its operation system, e.g., a UNIX® platform, Windows 2000® platform, etc..

POD

Performance On Demand.

Port/Port ID

See Fibre Channel.

RAID

Redundant Array of Independent Disks. RAID is used to increase the reliability of disk arrays by providing redundancy either through complete duplication of the data (RAID-1, i.e., mirroring) or through construction of parity data for each data stripe in the array (RAID-3, -4, -5). RAID-5, which distributes parity information across all disks in an array, is among the most popular means of providing parity RAID since it avoids the bottlenecks of a single parity disk. Algorithms for both the Lightning 9900™ Series systems and the Thunder 9200 systems enable performance from RAID-5 that is competitive with some vendor's RAID-1. Some vendors do not offer RAID-5.

RAID Controllers

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

Recovery Time

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

Remote Copy

Remote Copy refers generically to software or hardware utilities that provide the capability to copy data from one on-line 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.

RISC

Reduced Instruction Set Computing (RISC) refers to a type of microprocessor that executes a limited set of instructions. RISC processors are often optimized for specific tasks, such as storage controllers, as in the case of the Thunder 9200.

Risk Management

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

SAN

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

SCSI

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

Snapshot

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

SNMP

Simple Network Management Protocol. SNMP is a protocol used for communication between simple, server-resident SNMP agents that respond to network administration requests from simple-to-sophisticated SNMP management tools running on remote workstations. This feature is supported by both Lightning 9900 Series systems and Thunder 9200.

Solaris

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

SPARC

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

Stripe

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

SVP

Service Processor. This term applies to both the Lightning 9900 Series and Thunder 9200 systems.

SWAN

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

Switch

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

Synchronous Remote Copy

The transmission of data between two devices that are synchronized with a clocking scheme or other technique. 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 Asynchronous Remote Copy, Hitachi TrueCopy.

TB

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

TCP

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

Thunder 9200

The Thunder 9200 was announced in January 2001 in both rackmount and deskside packages. It represents the first storage system with enterprise-level hardware and software features with midrange package and pricing.

VERITAS®

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

Volume

An ESA/390® term for the information recorded on a single disk unit or recording medium. Indirectly, a volume can refer to the unit of recording medium itself. On a non-removable medium storage device such as a disk drive, the terms may also refer,

indirectly, to the storage device that is associated with the volume. When a user stores multiple volumes on a single storage medium transparent to the program, the volumes are referred to as logical volumes.

WAN

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

Warm Site

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

Workload

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

Hitachi Data Systems Worldwide Offices

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

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

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

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

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

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

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

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

DISK-414-00 January 2002

Hitachi Data Systems

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

European Offices:

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

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

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

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

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

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

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

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

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

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

Finland
Hitachi Data Systems
Tietäjäntie 2
02130 Espoo
Finland
Phone: +358 9 251 05 133
Fax: +358 9 251 05 100

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

DISK-414-EU-00 January 2002