Hitachi Freedom Storage™

Software

Solutions Guide

Best of Breed Software for Centralized Storage Management

HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™

Software

Solutions Guide

Best of Breed Software for Centralized Storage Management

Storage Area Management Solutions

Business Continuity Solutions

Backup and Recovery Solutions

Data Movement Solutions

Professional Services and Support

Continuous Business Planning Questionnaire

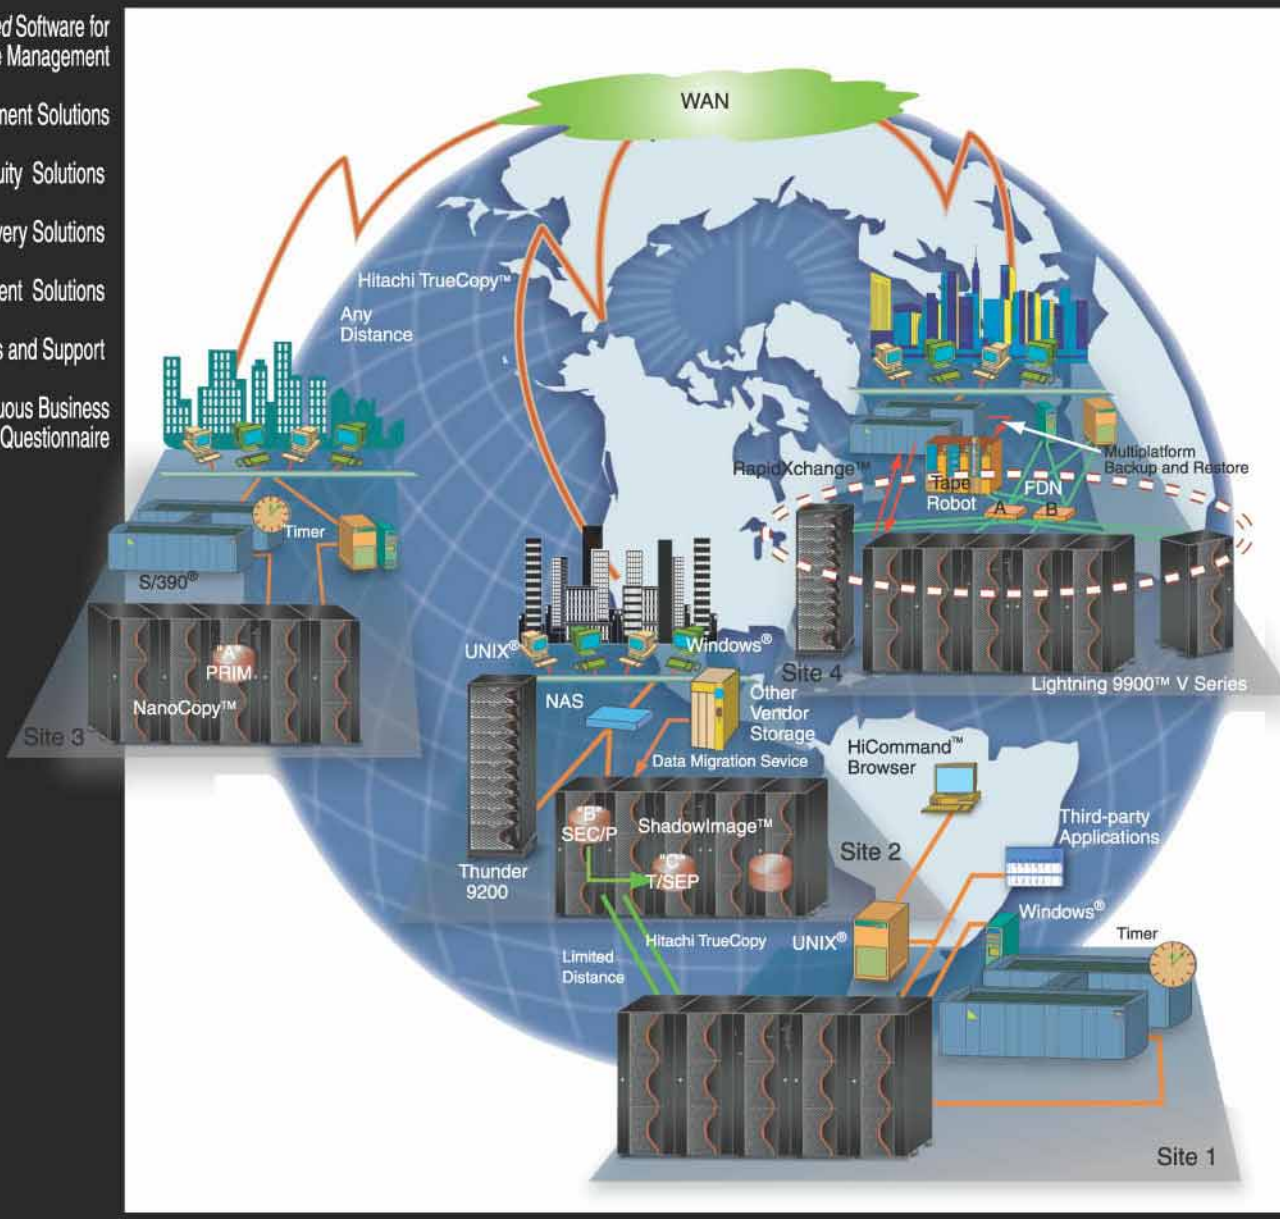

Best of Breed Software for Centralized Storage Management

HITACHI
DATA SYSTEMS

Table of Contents

Table of Figures

Table of Tables

Best of Breed Software for Centralized Storage Management

1

Hitachi TrueNorth™ – A Vision for Open Information Management

The TrueNorth vision and direction is for policy based automated management of storage investments, continuous availability, and lower cost.

As the information sources driving businesses expand, organizations must manage the escalation in the physical quantity of storage, the reliance on universal access to information, the growing complexity of their storage environments, and the adoption of emerging technologies. Few businesses have unlimited resources or time to address these challenges.

Today, companies need to consider a new storage management strategy based on high-performance, intelligent systems and sophisticated software that enable total manageability of existing investments, while maximizing uptime and lowering the total cost of ownership. Hitachi TrueNorth provides a vision and direction, along with an associated product roadmap, for an open information management framework that meets these challenges. With TrueNorth, we view the future storage infrastructure as a synergistic combination of management software and powerful, intelligent storage systems that will deliver always-on, set-and-forget management, complemented by flexible provisioning capabilities.

As first steps to achieving the TrueNorth vision, Hitachi Data Systems is working with its partners to expand the HiCommand™ storage area management framework. Open and collaborative by design, HiCommand software adheres to accepted industry standards and provides policy-based automation tools. To power this software we introduce the Hitachi Freedom Storage™ Lightning 9900™ V systems, our second-generation systems based on an internal switched architecture, delivering the industry's highest bandwidth, availability, and scalability. HiCommand and the Lightning 9900 V Series, combined with Hitachi Freedom Storage software solutions, are central to the Hitachi TrueNorth vision and will allow organizations to simplify, protect, and optimize their information infrastructures, reducing TCO and ensuring a more rapid ROI.

Hitachi Freedom Storage Software Solutions are the most advanced available.

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

Hitachi Freedom Storage Software Solutions support an enterprise's strategic goal of accessing any information, on any computer, located anywhere, at any time. The many advanced functions available on Lightning 9900 Series, Lightning 9900 V Series, and Thunder 9200 hardware are initiated, managed, and controlled through these powerful software programs. Table 1 summarizes how Hitachi Freedom Storage Software Solutions apply to open systems and mainframe environments.

Table 1 - Summary of Hitachi Freedom Storage software products.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390®
Storage Area Management Solutions					
Hitachi HiCommand™	✔	✔	✔	✔	
Hitachi Resource Manager™	✔	✔	✔	✔	✔
Hitachi Performance Enhancement Solutions	✔	✔	✔	✔	✔
Hitachi Parallel Access Volume		✔	✔		✔
VERITAS® Volume, File and SAN Management	✔	✔	✔	✔	
Business Continuity Solutions					
Hitachi TrueCopy	✔	✔	✔	✔	✔
Hitachi ShadowImage™	✔	✔	✔	✔	✔
Hitachi Extended Remote Copy (HXRC)		✔	✔		✔
Hi-Track®		✔	✔		✔
Hitachi Dynamic Link Manager™	✔	✔	✔	✔	
Alternate Pathing, Host Failover, and Parallel Database Clustering	✔	✔	✔	✔	✔
Backup and Recovery Solutions					
Hitachi Multiplatform Backup/Restore		✔	✔		✔
Tantia Technologies® HARBOR® Backup/Recovery and HYPERtape		✔	✔		✔
Non-disruptive database backup	✔	✔	✔	✔	✔
Non-disruptive application backup	✔	✔	✔	✔	✔
Data Movement Solutions					
Hitachi RapidXchange™		✔	✔		✔
Hitachi Data Migration Services		✔	✔		✔
Tantia Technologies HARBOR File Transfer		✔	✔		✔
Hitachi e-Copy			✔	✔	

The Business Value of Hitachi Freedom Storage Software Solutions

The HiCommand management framework delivers the ability to achieve centralized worldwide control of the storage environment. Customers have the freedom to choose the precise solution – or combination of solutions – appropriate for their environment in four important areas, with each area targeted to have specific business value. In keeping with the Hitachi commitment to open solutions, Hitachi Freedom Storage Software Solutions include not only software solutions provided by Hitachi Data Systems, but also solutions provided by software vendors.

Storage Area Management Solutions – Hitachi Data Systems software and hardware solutions are managed through the powerful HiCommand management framework. This management tool substantially lowers total cost of ownership (TCO) by enabling storage hyper-consolidation AND the intelligent remote management of critical information. The business value of centralized global storage management is to increase business agility through heightened operational excellence and greatly reduced TCO. The HiCommand management framework is illustrated in Figure 1 and its ability to drive down TCO through centralized worldwide control is illustrated in Figure 2.

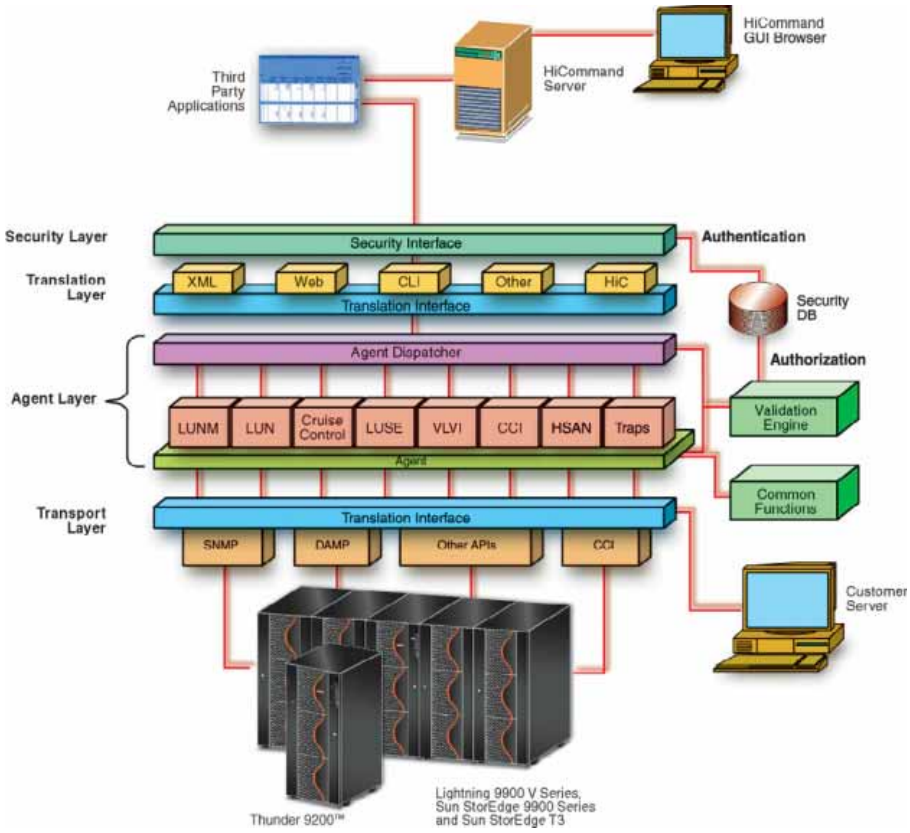

Figure 1 - HiCommand allows management of multi-vendor hardware and software from the customer's platform of choice.

At the heart of Hitachi Freedom Storage Software Solutions are the Hitachi Resource Manager and **Performance Enhancement Solutions**. Hitachi Resource Manager simplifies IT operations, improves information availability, and enables truly centralized worldwide control of critical enterprise information. The Performance Enhancement Solutions Suite increase performance service levels for business-critical OLTP and DSS applications, thereby enabling more rapid deployment of new applications by managing data more intelligently. For the Lightning 9900 V Series, both Resource Manager and the Performance Enhancement Solutions Suite have been enhanced to support the new virtual Fibre Channel Port technology and Host Storage Domains with the powerful Hitachi SANtinel™ and Priority Access programs.

Business Continuity Solutions – Hitachi Data Systems has an industry-unique commitment to the paradigm of Continuous Business. Continuous Business refers to an enterprise's ability to minimize system downtime, whether it is planned or unplanned. This is distinct from Business Continuity, which refers to the ability to recover from unplanned downtime. Building on the Hitachi Data Systems reputation for bulletproof reliability, the Lightning 9900 Series, the Lightning 9900 V Series, and the Thunder 9200 systems offer complete redundancy and hot-replaceable components, delivering maximum uptime. To these already robust platforms, Hitachi Data Systems adds Business Continuity solutions to ensure quick recovery from unplanned downtime resulting from acts of nature, human errors, application errors, and malicious attacks. Hitachi Data Systems is the only storage vendor that is providing non-disruptive copy solutions that ensure fast recovery and full data integrity. For the Lightning 9900 V Series, Hitachi TrueCopy performance has been improved to provide track instead of cylinder transfers and *resynch* time as been cut in half. For ShadowImage, copy performance has improved with an increase of more than 40 percent in copy transfer rate and twice the number of pairs (4096).

Figure 2 - HiCommand allows global systems management of Hitachi Freedom Storage hardware and software solutions.

Backup and Recovery Solutions – Backup/Restore is the largest cost element in storage ownership and the most important to high-availability computing. In today's global IT environment, it is no longer enough to have a backup copy of one's data; businesses must be able to perform backup and restore in the shortest possible time and with minimum disruption to information availability. Hitachi Data Systems therefore combines its own leading-edge Business Continuity software solutions with those of industry-leading software partners to address customers' backup/restore requirements with solutions that reduce TCO. Hitachi Multiplatform Backup/Restore delivers high-performance volume-level backup of open systems data. In addition to the HARBOR Backup/Recovery (HBR) and HYPERtape Backup/Restore software packages, Hitachi Data Systems offers several features and services to achieve this goal. Hitachi ShadowImage and Hitachi TrueCopy for S/390 can be used to create a snapshot copy for near-instantaneous backup.

All Hitachi Freedom Storage backup/restore solutions are designed to meet the following customer requirements:

- Maximum reliability
- Centralized administration
- File-level backup/restore
- Online backup of popular databases, applications, and messaging services
- Out-of-the-box agents for popular servers such as Oracle® Server and BMC Software™
- Multiplatform backup to a mainframe or open systems server
- High-performance movement of data
- Ability to "snapshot" the backup source in UNIX®, Windows NT®/Windows® 2000 or MVS®
- Ability to simulate or test the recovery scenario.

Data Movement Solutions – Sharing and moving large amounts of data is the key to business success in today's global information processing environments, which require that corporate data be readily available to all worldwide users, regardless of their location or the nature of their computing environment. With core business systems residing on a mix of mainframe and open systems, sharing and moving data among multiple heterogeneous systems is key to fulfilling this requirement. Data from heterogeneous systems needs to be exchanged regularly, either for data synchronization between discrete applications or, increasingly, for populating data warehousing systems.

Hitachi RapidXchange offers high-performance and high-reliability data sharing between heterogeneous host platforms such as OS/390®, UNIX, or Windows. With Hitachi RapidXchange, high-speed data transfer is achieved without putting additional data load on the network infrastructure or tape transport equipment. Running with Hitachi RapidXchange, HARBOR File Transfer adds automation to the process of transferring large data files at high channel speeds in either direction between open systems and mainframe servers. It can be thought of as an ultra-fast FTP transfer, except data is transferred over an ESCON® link instead of Ethernet. HARBOR File Transfer enables UNIX and Windows applications, such as enterprise resource planning and decision support systems, to work with up-to-the-minute data extracted from S/390 environments. The Hitachi Data Migration Service helps migrate data from existing systems to newly installed systems while minimizing the impact on mission-critical applications. For the Lightning 9900 V Series, the Hitachi e-Copy feature provides server-free transfer of data between devices whether it is disk to disk or disk to tape.

Other Information Sources Available from Hitachi Data Systems

Two other publications that discuss the core competencies upon which the Hitachi TrueNorth vision and strategic direction are based are available from Hitachi Data Systems. ***The Lightning 9900 V Series Architecture Guide*** and the ***Thunder 9200 Architecture Guide*** describe these leading-edge products' key features and business values relating to respective product classes. These documents are downloadable in PDF format from http://www.hds.com.

If further information is required on any of the topics covered in these three guides, Hitachi Data Systems representatives can provide specialized presentations, reports, and expert knowledge.

Storage Area Management Solutions

2

HiCommand: Hitachi's Framework For Open Storage Management

Storage management is a complex problem that must support a myriad of existing components and new technologies. With the advent of SAN and NAS, customers have greater flexibility and control over their storage environments than they have ever had in the past. However, this has also dramatically increased their storage management burden. These new, complex storage topologies lack a clear and consolidated approach to management. It is no longer good enough to provide basic services such as LUN management, device monitoring and backup; now, it is critical to provide integrated management of all aspects of storage across a heterogeneous environment made up of multiple devices, such as SAN switches, complex storage systems and storage appliances.

Customers have also invested heavily in "point" solutions and systems management applications for their current environments. It is not sufficient for a vendor to provide a replacement for existing technologies; the vendor must provide a complete solution by allowing customers to leverage their existing investments, whether they are in hardware, software or people.

No single vendor can meet all the storage requirements for today's enterprises.

Hitachi provides an open platform, based on industry-standard management protocols, along with high value-add functional components such as policy-based automation and heterogeneous device management, allowing customers to incorporate multiple vendor offerings to create a truly complete solution for managing and monitoring their storage environments.

As shown previously in Figure 1, HiCommand supports a modular, building-block approach to designing and addressing business-specific storage management requirements. As a result, Hicommand enables customers to build the most suitable management environment based on a combination of Hitachi products and best-of-breed storage hardware and software. HiCommand is customizable to integrate easily with existing infrastructures and accommodate future growth plans.

HiCommand's policy-based automation module (Figure 3) simplifies management functions and enables integrated administration of multi-vendor storage environments. Furthermore, standard storage management operations come pre-loaded in the policy-based automation module. HiCommand also provides definition tools and an easy-to-use GUI interface for automating more complicated storage management processes, resulting in lowered total cost of ownership.

Figure 3 – A Logical View of HiCommand Storage Management Architecture

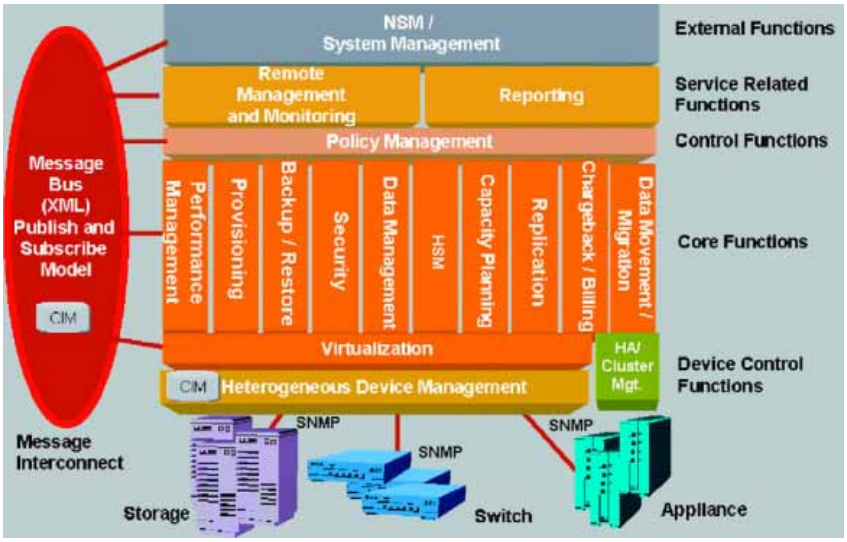

The CIM Model and XML Provide for an Open, Standards-Based Approach

The foundation of Hitachi's architecture is a message bus that allows individual functional components to "plug and play". Functional components are incorporated into the overall solution based on a usage of CIM that provides a clear set of information content for management and monitoring of the entire storage and system environment. Using this object-model-based approach, solutions can readily communicate with all network-attached devices and perform a variety of configuration, administrative and other functions, allowing for a complete, integrated hardware and software management environment. Adopting this model is the key to overall interoperability in the storage and systems management environments. If a solution component adheres to this model, then its management or incorporation will be relatively straightforward. Another key standard, XML, is fast becoming the norm for data communication within an integrated environment. Using XML to express the messages within the CIM model enables not only a common understanding of the required content, but also a clear and universal understanding of the message format. Hitachi is creating an open and extensible architecture using these two foundational standards.

By building to an industry-accepted specification, the Common Information Model1 (CIM), Hitachi and other vendors provide a clear set of information content for management and monitoring of the entire storage and system environment. Using this model-based approach, storage management solutions can readily communicate not only with devices but also with other functions, allowing for a complete, integrated hardware and software management environment. The adoption of this model is the key to overall interoperability in the storage and systems management environment. If a solution component adheres to this model, then its integration into a CIM-based management environment will be relatively straightforward.

Another key standard for interoperability is XML (eXtensible Markup Language). This is fast becoming the de-facto standard for data communication within an integrated environment. Using the CIM model in conjunction with XML to express messages ensures a common understanding of these messages, and provides for flexible application integration. Hitachi is creating an open and extensible architecture using these two foundational standards. The HiCommand Device Manager is the first deliverable component of this comprehensive framework.

HiCommand™ Allows Global Systems Management of Multi-vendor Storage Systems

At the heart of the Hitachi Storage Area Management Solutions are HiCommand™, Hitachi Resource Manager™ and the Hitachi Performance Enhancement Solutions. Hitachi Resource Manager simplifies IT operations, improves information availability and enables truly centralized worldwide control of critical enterprise information. Hitachi Performance Enhancement Solutions increases performance service levels for business-critical OLTP and DSS applications, thereby enabling more rapid deployment of new applications by managing data more intelligently. For the Hitachi Freedom Storage™ Lightning 9900™ V Series, both Hitachi Resource Manager and Hitachi Performance Enhancement Solutions have been enhanced to support the new virtual Fibre Channel port technology and Host Storage Domains in combination with the powerful Hitachi SANtinel™ and Hitachi Priority Access programs.

Hitachi Freedom Storage software and hardware solutions are administered through the powerful HiCommand™ management framework, which substantially lowers total cost of ownership (TCO) by enabling storage massive consolidation AND the intelligent remote management of critical information. The business value of centralized global storage management is to increase business agility through increased operational excellence and greatly reduced TCO.

Hitachi Data Systems has long recognized the importance of ISVs and is working with vendors whose offerings are complementary. This is accomplished through the HiCommand open APIs and through joint development and engineering, joint certification, and where appropriate, marketing agreements. At this time, there are over 75 vendors supporting the HiCommand open API. Both Hitachi Data Systems and its software strategic alliance members believe this open architecture will best benefit customers by allowing them to manage their 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 Hitachi Freedom Storage™ Lightning 9900™ V Series system(s) and Hitachi Freedom Storage Thunder 9200™ system(s) over a private LAN. A single HiCommand server can manage multiple Lightning 9900 Series, Lightning 9900 V Series, Thunder 9200, as well as Sun™ StorEdge™ 9900 Series and T3 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 Sun Microsystems®, BMC Software™, IBM®, 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 is easy to learn for simple and efficient storage administration.The HiCommand GUI browser (Figure 4) is intuitive and easy to use. A browser from anywhere can control the HiCommand server. In addition, Hitachi Data Systems also supplies an optional Hitachi Command Control Interface (CCI), which enables users to perform Hitachi TrueCopy and Hitachi ShadowImage™ operations by issuing commands from the UNIX®/PC server host to the Lightning 9900 V Series system. The CCI software interfaces with the

HiCommand enables an enterprise to manage multi-vendor systems and software solutions.

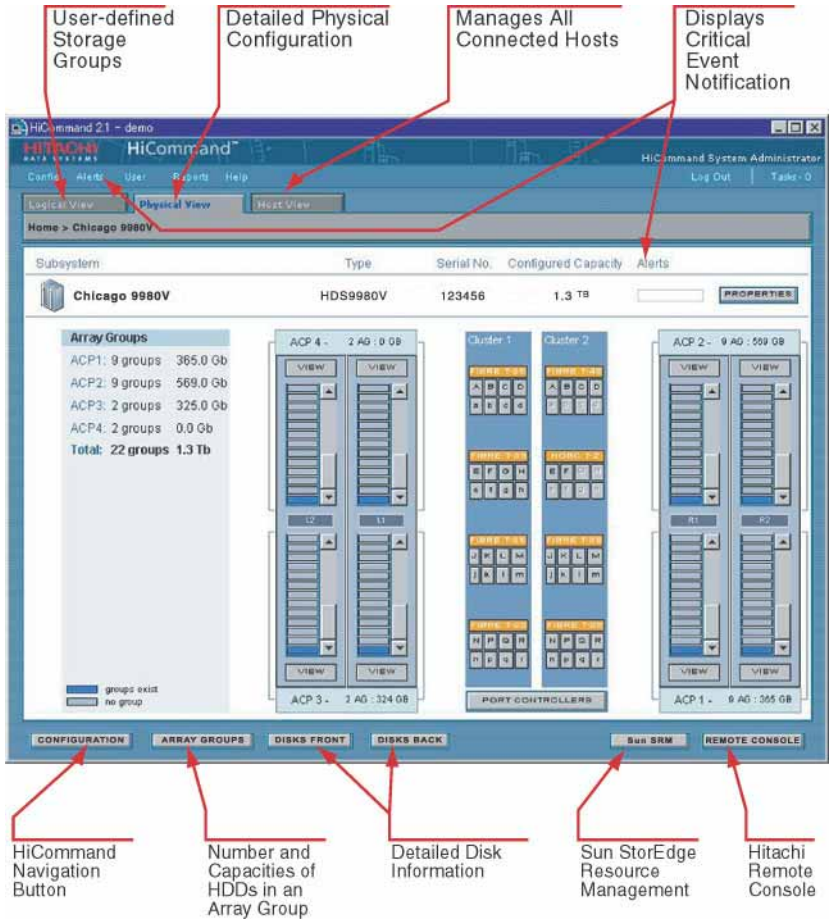

Figure 4 - The HiCommand GUI browser is intuitive and easy to use.

system software and high-availability (HA) software on the UNIX/PC server host as well as the Hitachi TrueCopy/ShadowImage software on the Lightning 9900 V Series system. The CCI software provides failover and other functions such as backup commands to allow mutual hot standby in cooperation with the failover product on the UNIX/PC server. CCI also supports a scripting function that allows users to define multiple Hitachi TrueCopy and/or ShadowImage operations in a script (text) file. Using CCI scripting, administrators can set up and execute a large number of Hitachi TrueCopy and/or ShadowImage commands in a short period of time while integrating host-based high-availability control over remote copy operations.

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

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.

How HiCommand 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. Today there are three major storage architectures – DAS, NAS and SAN – and many types of servers. 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 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 is Modularly Designed

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. As shown in Table 2, this comprehensive systems and software management offering brings together 11 powerful software solutions, with those shown in **bold** being standard for either the open or S/390®-compatible suites.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390®
Hitachi Resource Manager™					
Remote Console - Storage Navigator	✔	✔	✔	✔	✔
Hitachi Graph-Track™		✔	✔	✔	✔
SNMP Agent	✔	✔	✔	✔	✔
Java™ Agent	✔	✔	✔	✔	✔
Host Storage Domains			✔	✔	
LUN Manager / LUN Expansion	✔	✔	✔	✔	
Hitachi SANtinel™	✔	✔	✔	✔	✔
Virtual LUN / LUN Expansion			✔	✔	
VLVI			✔		✔
Hitachi FlashAccess™	✔	✔	✔	✔	✔
Cache Management Host Agent		✔	✔		✔

Table 2 - Summary of Hitachi Resource Manager solutions and the products they support.

Shipped with every Thunder 9200, Lightning 9900 Series, or Lightning 9900 V Series 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 an ideal choice. Users can monitor and manage their storage systems through a graphical user interface.

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

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

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

Service Level Management

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

Deployment Management

- Asset Management
- Change Management
- Capacity Planning

Compliance Management

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

1. Remote Console – Storage Navigator

The Lightning 9900 V Series Remote Console – Storage Navigator is provided as a Java applet program, which can execute on any machine that supports a Java Virtual Machine (JVM). The Remote Console – Storage Navigator PC hosts the Storage Navigator Java applet program and is attached to the Lightning 9900 V Series system(s) via a TCP/IP local area network (LAN). When a Storage Navigator accesses and logs into the desired SVP, the Remote Console applet is downloaded from the SVP to the Remote and runs on the Web browser of the Remote Console PC. In this way, the Remote Console communicates with the attached Lightning 9900 V Series or Thunder 9200 systems via a TCP/IP network.

Two LANs can be attached to the Lightning 9900 V Series or Thunder 9200 systems: an internal LAN (private LAN), which is used to connect the service processors (SVPs) of multiple subsystems, and the user's intranet (public LAN), which allows you to access the Storage Navigator functions from individual Storage Navigator PCs. The Remote Console – Storage Navigator communicates directly with the service processor (SVP) of each attached subsystem to obtain subsystem configuration and status information and send user-requested commands to the subsystem. The Lightning 9900 V Series or Thunder 9200 system's Storage Navigator Java applet program is downloaded to the Storage Navigator (Web client) from the SVP (Web server) each time the Storage Navigator is connected to the SVP. The Storage Navigator Java applet program runs on Web browsers, such as Internet Explorer and Netscape Navigator®, which run under the Windows and Solaris operating systems to provide a user-friendly interface for the Lightning 9900 V Series or Thunder 9200 Storage Navigator functions.

2. Hitachi Graph-Track takes the guesswork out of performance management

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

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

Hitachi Graph-Track gives users a more reliable and centralized way to manage performance. Available for S/390, UNIX, and Windows NT/2000 environments, this unique tool monitors hardware performance and supplies complete system storage information at the touch of a button. The robust GUI-oriented command screen for Hitachi Graph-Track is shown in Figure 5. Running

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

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

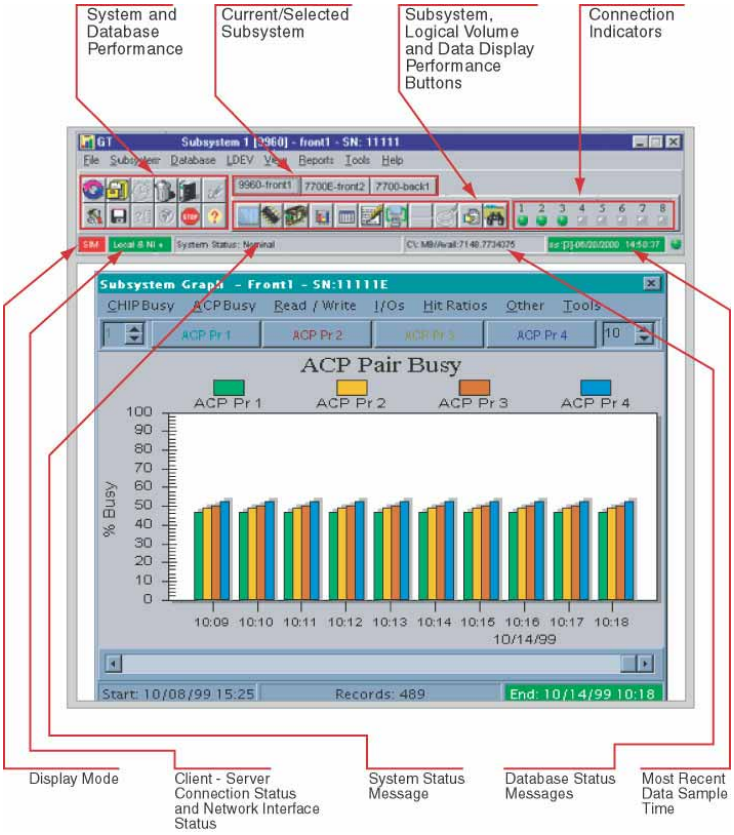

Figure 5 – The GUI interface on Hitachi Graph-Track simplifies performance management.

3. SNMP support ensures compliance with all major systems management platforms

Hitachi Resource Manager supports Simple Network Management Protocol (SNMP). SNMP is the most widely used network management protocol in the marketplace. System Information Messages (SIMs) are converted to SNMP Management Information Blocks (MIBs) for use with Hitachi Resource Manager software, or the VERITAS® Storage Manager™, Tivoli, CA Unicenter TNG® or HP® OpenView®. With CA Unicenter TNG, "super MIBs" allows the control of the Lightning 9900 V Series or Thunder 9200 systems using CA Unicenter commands. This out-of-band[1] reporting format ensures that Hitachi Resource Manager can determine the status of your Hitachi Freedom Storage system even if the data channel path is down.

[1] "Out of Band" refers to the path that system management messages travel, which is a different path (Figure 3) from the data. "In-band" messages would compete with data for bandwidth.

4. Java™ agent

The Hitachi Resource Manager Java agent allows the Remote Console – Storage Navigator Java applet program to run on Web browsers, such as Internet Explorer and Netscape Navigator. These programs run under the Windows and Solaris operating systems and provide a user-friendly interface for the Lightning 9900 Series, Lightning 9900 V Series, or Thunder 9200 Remote Console – Storage Navigator functions.

5. Host Storage Domains

The Lightning 9900 V Series provides for "virtual" Fibre Channel ports that are logically managed by intelligent Fibre Channel controller cards. Each physical FC port has multiple Host Storage Domains or HSDs. An HSD supports a "logical" FC port (Fig. 6) each with its own set of LUNs. Hosts are matched to their assigned HSD based on a unique World Wide Names (WWNs) identifier. LUN security is provided by Hitachi SANtinel, which fences access to LUNs based upon host WWNs. Under this innovative software structure, one LUN is reserved for the command and control device.

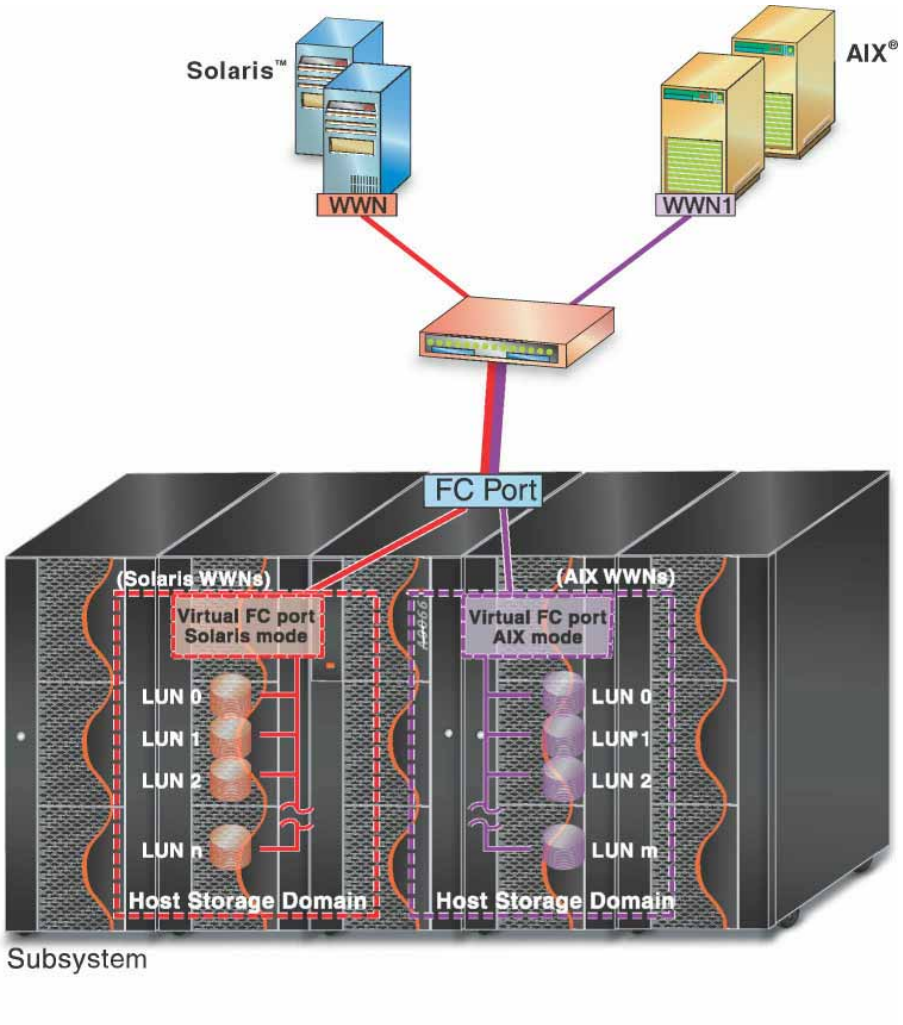

Figure 6 – Host Storage Domains reduce costs through fewer physical connections.

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

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

7. Hitachi SANtinel

Hitachi SANtinel enables users to restrict server access to only those LUNs for which they are authorized.

The Hitachi SANtinel™ controls host access to Hitachi Freedom Storage LUNs in open systems, multi-platform, or SAN environments. This enables users to restrict server access to only those LUNs for which they are authorized, as shown in Figure 7.

SANtinel for open systems, running on the Lightning 9900 V Series systems, allows administrators to define multiple Host Storage Domains for each physical Fibre Channel port on the system. A Host Storage Domain consists of a virtual (logical) Fibre Channel port, and an associated set of LUNs visible only via that logical Fibre Channel port.

Connections from host servers arriving at the physical Fibre Channel port are routed to the logical Fiber Channel port within the appropriate Host Storage Domain based upon the Fibre Channel World Wide Name (WWN) of the host.

Hitachi Virtual Logical Volume Image Manager improves performance.

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

Because the WWN is used to route connections to the appropriate Host Storage Domain, the WWN may only be assigned to a single Host Storage Domain on each Fibre Channel port.

SANtinel may be used to authorize multiple hosts to access a particular Host Storage Domain. Within that Host Storage Domain, SANtinel may be further used to permit or deny a particular host to access individual LUNs within the Host Storage Domain, based upon the WWN of the host.

SANtinel may be used to create up to 128 Host Storage Domains per physical Fibre Channel port. Each Host Storage Domain may contain up to 256 LUNs, with the maximum number of LUNs contained within all Host Storage Domains on a given Fibre Channel port reaching 512.

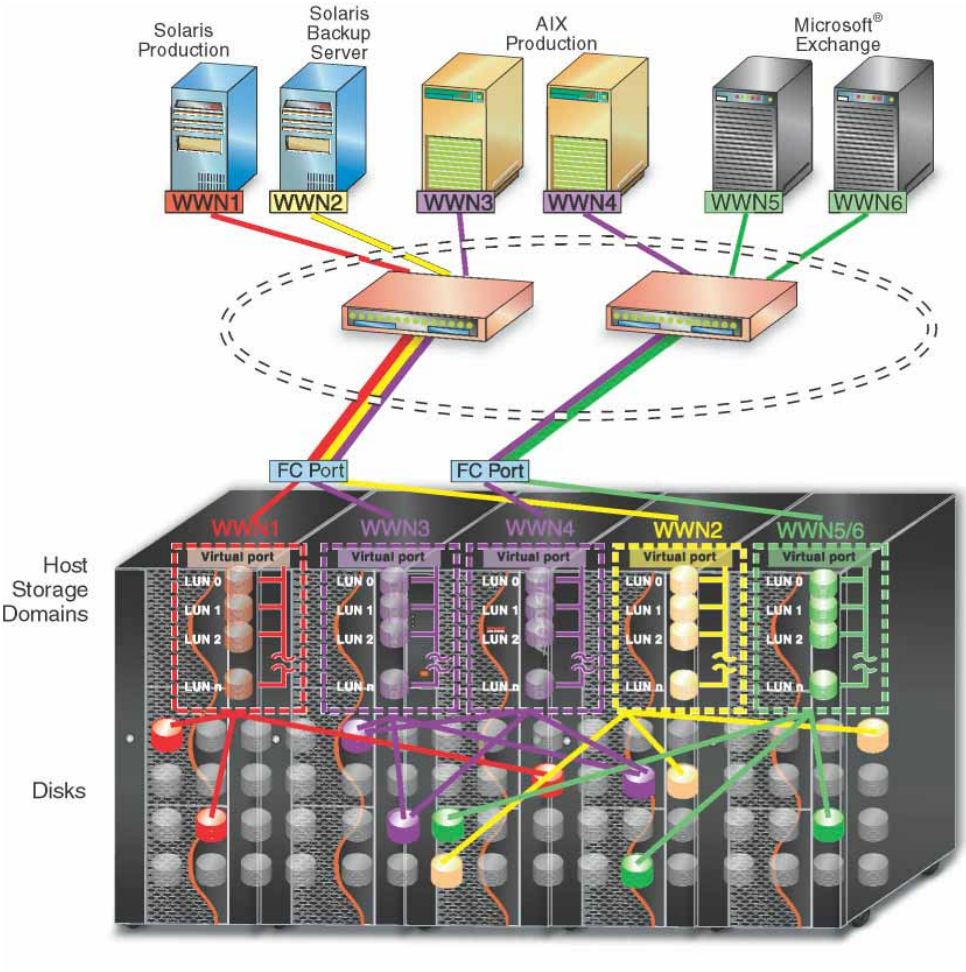

Figure 7 – SANtinel uses pre-defined zones to protect data from unauthorized access.

In a similar fashion to open systems, Hitachi SANtinel for S/390 Environments allows administrators to restrict S/390 host access to the logical devices (LDEVs) on the Lightning 9900 V Series systems. Each LDEV to can be set to communicate only with user-selected host(s). Hitachi SANtinel for S/390 prevents other hosts from seeing the secured LDEV and from accessing the data contained on the secured LDEV. The licensed SANtinel for S/390 Storage Navigator software displays the SANtinel for S/390 information and allows administrators to perform SANtinel for S/390 operations.

8. Virtual LUN and LUN Expansion

Virtual LUN allows administrators to convert fixed-size volumes into several smaller variable custom-sized volumes. Using the Remote Console – Storage Navigator software, users can configure custom-size volumes by assigning a logical address and a specific number of cylinders/tracks (for S/390 data) or MB (for open-systems data) to each custom logical volume image (LVI)/LUN. Virtual LUN improves data access performance by reducing logical device contention as well as host I/O queue times, which can occur when several frequently accessed files are located on a single volume. Multiple LUN types can be configured within each array group. Virtual LUN enables administrators to more fully utilize the physical storage capacity of the Lightning 9900 V Series systems, while reducing the amount of administrative effort required to balance I/O workloads.

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

The LUN Expansion Manager allows administrators to create virtual LUNs that are larger than standard OPEN LUNs, by expanding the size of a selected LUN up to 36 times its normal size. The maximum size depends on the type of configuration. For example, an administrator can expand an OPEN-9 LUN to a maximum size of 265GB (7.3GB x 36). This capability enables open-systems hosts to access the data on the entire Lightning 9900 V Series system using fewer logical units. LUN Expansion Manager allows host operating systems that have restrictions on the number of LUNs per interface to access larger amounts of data.

Hitachi FlashAccess dramatically improves performance.

9. Hitachi Virtual Logical Volume Image LVI Manager for S390 Environments allows very large to very small virtual volume configuration

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

10. Hitachi FlashAccess allows data to be "locked and unlocked" in cache on-the-fly

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

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

11. Cache Manager Host Agent

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

Hitachi Performance Enhancement Solutions

Hitachi Performance Enhancement Solutions increase performance service levels for business-critical OLTP and DSS applications, thereby enabling more rapid deployment of new applications by managing data more intelligently. For the Lightning 9900 V Series, both Hitachi Resource Manager and Hitachi Performance Enhancement Software have been enhanced to support the new virtual Fibre Channel Port technology and Host Storage Domains with the powerful SANtinel and Hitachi Priority Access programs.

Table 3 - Summary of Hitachi Performance Enhancement Solutions Solutions and the products they support.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390
Hitachi Performance Enhancement Solutions					
Hitachi Performance Monitor		✔	✔	✔	✔
Hitachi CruiseControl™		✔	✔	✔	
Hitachi Priority Access			✔	✔	
TPF/MPLF		✔	✔	✔	

1. Hitachi Performance Monitor

Hitachi Performance Monitor is a powerful software utility that provides a variety of performance measurements for system administrators to optimize performance of their Hitachi Freedom Storage systems. It is included at no extra charge with the purchase of either Hitachi CruiseControl™ or Hitachi Priority Access.

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

2. Hitachi CruiseControl

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

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

With the massive consolidation made possible by Lightning 9980V capabilities, customers will be able to not only lower the original purchase costs, but also get additional savings from lower storage management costs, lower software and maintenance costs, and reduced real estate and environmental costs. This is often the major payback for Hitachi Freedom Storage products. Reduced per-megabyte administration costs, by itself, can provide a substantial return on investment.

3. Hitachi Priority Access

Hitachi Priority Access allows open-systems users to designate prioritized ports (e.g., for production servers) and non-prioritized ports (e.g., for development servers) and set thresholds and upper limits for the I/O activity of these ports. This enables users to tune the performance of the development server without affecting the production server's performance as illustrated in Figure 8.

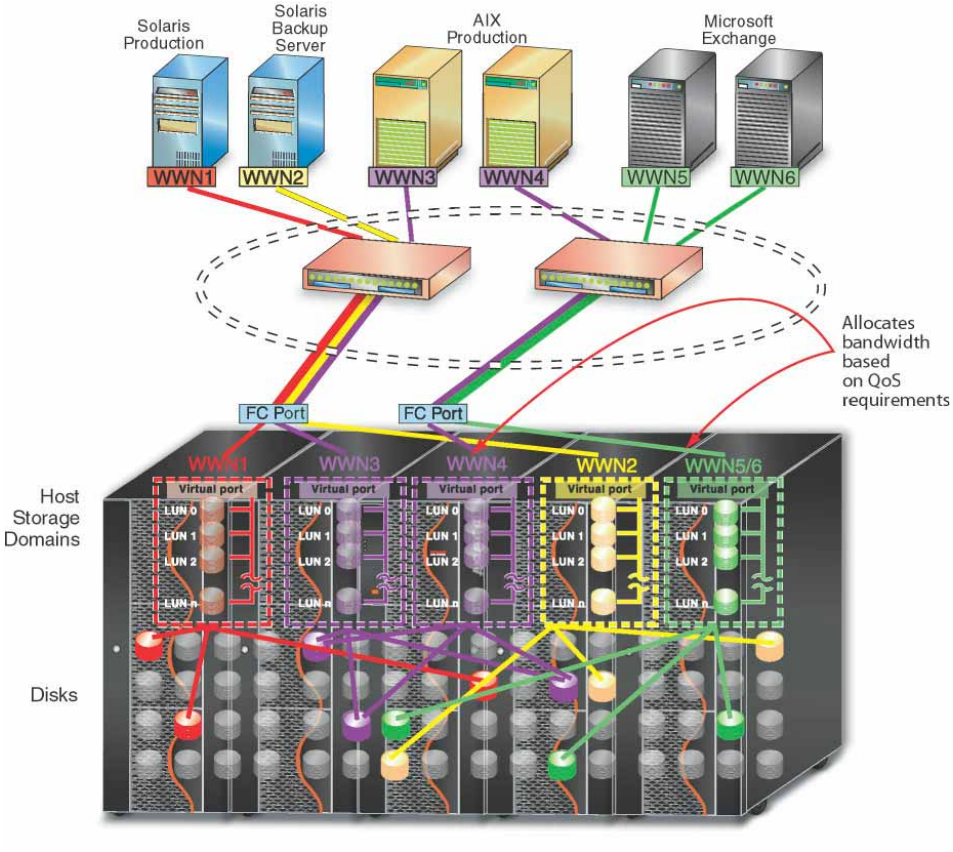

Figure 8 – Priority Access allows administrators to allocate bandwidth in line with QoS.

4. Transaction Processing Facility/Multi-path Lock Facility (TPF/MPLF)

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

5. Hitachi Parallel Access Volume (HPAV) and Multiple Allegiance (MA)

Hitachi Parallel Access Volume (HPAV) enables the S/390 host system to issue multiple I/O requests in parallel to single logical devices (LDEVs) in the Lightning 9900 V Series system. HPAV can provide substantially faster host access to the S/390 data stored in the Lightning 9900 V Series system. The Workload Manager (WLM) host software function enables the S/390 host to utilize the HPAV functionality of the Lightning 9900 V Series system. The Lightning 9900

V Series supports both static and dynamic HPAV functionality. Multiple Allegiance (MA) extends this capability to applications running on multiple S/390 servers. Together, HPAV and MA reduce queuing, which results in significantly decreased batch times and lightning-fast responses in today's high-transaction environments.

The VERITAS Foundation Suite helps decrease the total cost of operations.

VERITAS Software Simplifies Management and Increases Availability

The VERITAS Foundation 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, online 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.

VERITAS Volume Manager increases availability and lowers costs.

1. VERITAS Volume Manager

VERITAS Volume Manager increases availability by creating redundant storage solutions for critical data. At the same time, it helps lower storage costs by making the most of the storage already in-house. "Hot relocation" capabilities automatically construct a mirror on free space in the event of the failure of device. This capability is ideal for those applications that require a non-hardware RAID solution (that is, software-based RAID-0 or RAID-5). Hot spare devices are identified for use in case there is a disk failure, but if they are unavailable, any available free space is used to maintain redundancy. If the server hardware and operating system support hot swapping, data can be swapped to the new device while the system is running.

Instead of adding to overall storage costs, the VERITAS Foundation Suite manages storage more efficiently for overall cost savings and improves the performance of existing storage. By providing the flexibility and scalability to add and manage new storage easily, it allows users to take a proactive approach to storage management. VERITAS Volume Manager enables the effective use of current hardware by identifying available segments or partitions on storage devices and combining them into one or more virtual disks or logical volumes online, without interruption.

This feature also allows the striping of data across multiple physical volumes. Striping reduces the time spent by the application waiting for disk head movements and generates faster overall system performance. In addition, data can be mirrored between multiple RAID devices, and online load balancing between RAID logical units, along with the reconfiguration of RAID logical units can be performed without interrupting access to data.

VERITAS Volume Manager performs regular maintenance tasks online while maintaining system and data availability. These tasks include data backup and defragmentation, storage allocation resizing, storage additions, faulty or slow device replacement, and device reconfiguration for I/O load balancing or new application support. Availability is further maintained via the Dynamic MultiPathing (DMP) capabilities of VERITAS Volume Manager. If a path fails or becomes unavailable, VERITAS Volume Manager automatically reroutes data to the open path(s).

VERITAS File System reduces fragmentation and requires fewer indirect pointers to files, reducing seek time.

2. VERITAS File System

VERITAS File System works to reduce outages due to file system "panics." If it detects a disk or file system error, VERITAS File System isolates the problem without causing a

failure on the whole file system. If an outage occurs, VERITAS File System can restore maximum access to data in seconds. All writes to the file system are recorded in a journal or intent log until they are written to index or data blocks. During the restart after failure, VERITAS File System replays the log record and completes any uncompleted writes within seconds. Its online snapshot, supported by common backup packages, enables full backup without data lockout.

When time is money, it pays to ensure optimized application performance, especially for high I/O applications. Performance can be significantly increased with the features of VERITAS File System and VERITAS Volume Manager. By allocating storage in extents (groups of contiguous blocks) rather than in fixed block sizes, VERITAS File System reduces fragmentation and requires fewer indirect pointers to files. This reduces seek time due to disk latency. Reading/writing to larger extents rather than to smaller blocks requires fewer I/O operations. VERITAS File System also facilitates fast performance by automatically tuning I/O to match VERITAS Volume Manager data layout. VERITAS Volume Manager dramatically improves overall I/O throughput on disk controllers that have two separate I/O paths by using both paths simultaneously. By configuring the VERITAS File System to do this automatically, an immediate performance improvement in existing applications that are I/O intensive will result. If there are storage locations with very high access patterns, VERITAS Volume Manager will establish a "sparse mirror." A "sparse mirror" is a mirror of a small part of the overall data. This reduces contention on read requests by making another copy available. If the sparse mirror resides on storage with very fast access characteristics, such as the Lightning 9900 V Series systems, then the access provided is even faster. In addition, Hitachi FlashAccess software for open systems and S/390 platforms for the Lightning 9900 V Series systems provides for dynamic cache residency, which can further enhance performance.

3. VERITAS Storage Administrator Improves System Manageability

The VERITAS Foundation Suite and Hitachi Freedom Storage allow the centralization, automation, and integration of critical resources.

When an enterprise has limited time and resources to monitor storage, centralization of data management can provide maximum benefits to operations efficiency and reduced cost efforts. The VERITAS Foundation Suite and Hitachi Freedom Storage allow the centralization, automation, and integration of critical resources – whether local or distributed, online or offline. VERITAS Foundation Suite is easy to use, reducing the training costs and the time spent completing administrative tasks.

Simplified remote storage management is possible through VERITAS Storage Administrator, a Java-based graphical user interface (GUI) that ships with VERITAS Foundation Suite. The GUI helps administrators defragment storage, identify and group devices, implement mirroring and striping, combine segments from multiple devices into logical volumes, and resize or move volumes on physical disks. All of this occurs while storage remains available and users remain online. Systems managers can monitor I/O traffic to identify bottlenecks and inefficiencies with the VERITAS Foundation Suite command-line interface. VERITAS Volume Manager further contributes to manageability in storage solution by making it easy to move data between nodes in a SAN environment. Finally, the VERITAS File System allows a storage administrator to group large collections of files for commercial applications, eliminate media size limitations, and easily manage space requirements for a large number of user accounts.

4. VERITAS SANPoint Control™

VERITAS SANPoint Control is the data storage industry's most advanced, easy-to-use, centralized management tool for automatic discovery, visualization, and zoning administration of SAN-connected devices. Leveraging the VERITAS breakthrough V3™ technology and existing industry standards, VERITAS SANPoint Control provides the first drag-and-drop zoning capabilities in a heterogeneous storage area network (SAN). Building

on the unique storage visualization capabilities of VERITAS Volume Manager, VERITAS SANPoint Control enables a single point of management from LUN to interconnect and to SAN-connected hosts. Host-, device- or fabric-based maps provide customizable intuitive views into your SAN as well as real-time events to diagnose and avoid outages. By masking the administrative complexity of the SAN environment, VERITAS SANPoint Control allows the potential of SANs to be realized today.

Business Continuity Solutions

3

Hitachi Data Systems has an industry unique commitment to the paradigm of Continuous Business. Continuous Business refers to an enterprise's ability to minimize system downtime whether it is planned or unplanned. This is distinct from Business Continuity, which refers to the ability to recover from unplanned downtime. Building on the Hitachi Data Systems reputation for bulletproof reliability, the Hitachi Freedom Storage™ Lightning 9900™ Series, the Lightning 9900™ V Series and the Thunder 9200™ systems offer complete redundancy and hot-replaceable components, delivering maximum uptime. To these already robust platforms, Hitachi Data Systems adds Business Continuity software solutions to ensure quick recovery from unplanned downtime resulting from acts of nature, human errors, application errors, and malicious attacks. Hitachi Data Systems is the only storage vendor that is providing non-disruptive copy solutions that ensure fast recovery and full data integrity. For the Lightning 9900 V Series, Hitachi TrueCopy performance has been improved to provide track instead of cylinder transfers and *resync* time as been cut in half. For Hitachi ShadowImage™, copy performance has improved with a copy transfer rate increase of more than 40 percent and twice the number of pairs (4096).

Increased IT Service Levels in Availability and Performance through Business Continuity Software Solutions and Services

The Lightning 9900 V Series delivers unsurpassed uptime and performance.

In today's global IT environment, performance and availability go hand in hand. If a Web site is slow in delivering requested information, it might as well be offline. Of course, if a Web site is down for any length of time, the business may not be around for long. That's why the Lightning 9900 V Series and the Thunder 9200 systems offer complete redundancy and hot-replaceable components, delivering maximum uptime.

High availability terminology is important to understand.

It is important to understand the terminology used in discussions of Business Continuity and Disaster Recovery since both are similar terms used in discussions of high availability computing. Continuous Business refers to an enterprise's ability to minimize system downtime whether it is planned or unplanned. This is distinct from Business Continuity, which refers to the ability to recover from unplanned downtime. Business Continuity Planning software and services is an "umbrella" term covering both disaster recovery planning and business resumption from a business interruption.

A business interruption is any event, whether anticipated (i.e., public service strike) or unanticipated (i.e., blackout) which disrupts the normal course of business operations at an enterprise location. A disaster is any event that creates an inability for an organization to provide critical business functions for an undetermined period of time. While a disaster is a business interruption, not all business interruptions are disasters.

Disasters can be subdivided into three categories:

Man-made Disasters

- Power outages
- Fires
- Explosions
- Accidents
- Bombings
- Sabotage
- Building collapse
- Burst pipes
- Environmental contamination

Natural Disasters

- Power outages
- Fires
- Earthquakes
- Tornadoes and wind storms
- Volcanic eruptions
- Hurricanes, monsoons, and typhoons

Man-made Political Disasters

- Riots and civil disturbances
- Terrorists attacks

Not all business interruptions threaten the very existence of the enterprise. However, some industries are far more sensitive to computer system related business interruptions than others.

The "Scale of 9s" Is an Accepted Measure of Computer System Availability

What is the value of one minute of unplanned downtime to you?

In the field of high availability engineering, a "***Scale of 9s"*** was first used by respected technologist Jim Gray (formerly of Tandem Computers®) to classify computer system availability into ten levels based on the number of "nines" in the availability equation. As shown in Table 4, a Level 2 computer system is available 99 percent of the time in an average year with an expectancy of 1 percent unplanned downtime per year. This translates to an expectancy of 87.6 hours of unplanned downtime, which is far greater than a Level 7 system with unplanned downtime of three seconds per year. Typically fully fault tolerant computer products, such as the Hitachi Freedom Storage Lightning 9900 V Series systems, exhibit availability of Level 5. Triple redundant computers have availability ratings of 6 or 7. Most competitive products are believed to have availability ratings of Level 4. A Level 4 rating would result in an order of magnitude more unplanned downtime per year for competitive products compared to Hitachi Freedom Storage products.

Table 4 – The "Scale of 9s" is an accepted measure of availability.

Level	% Availability	% Downtime	Downtime
1	90%	10%	876 hrs.
2	99.0%	1%	87.6 hrs.
3	99.9%	.1%	8.76 hrs.
4	99.99%	.01%	52.56 min.
5	99.999%	.001%	5.25 min.
6	99.9999%	.0001%	31.54 sec.
7	99.99999%	.00001%	3.15 sec.
8	99.999999%	.000001%	.31 sec.
9	99.9999999%	.0000001%	.03 sec.

The Importance of the Continuous Business Paradigm to Nonstop Operations

Whereas the paradigm for high-availability computing is based on avoiding unplanned interruptions to data availability, Hitachi Data Systems 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 Data Systems delivers continuous data availability by applying ISO 9001 type disciplines to threat assessment and risk management. The process begins with a careful analysis of all downtime variables via the Hitachi Data Systems Continuous Business Planning Questionaire as shown in Appendix A. In other words, the focus of Hitachi Data Systems is on *continuous business* instead of *business continuity.* This subtle yet important paradigm shift is illustrated in Figure 9.

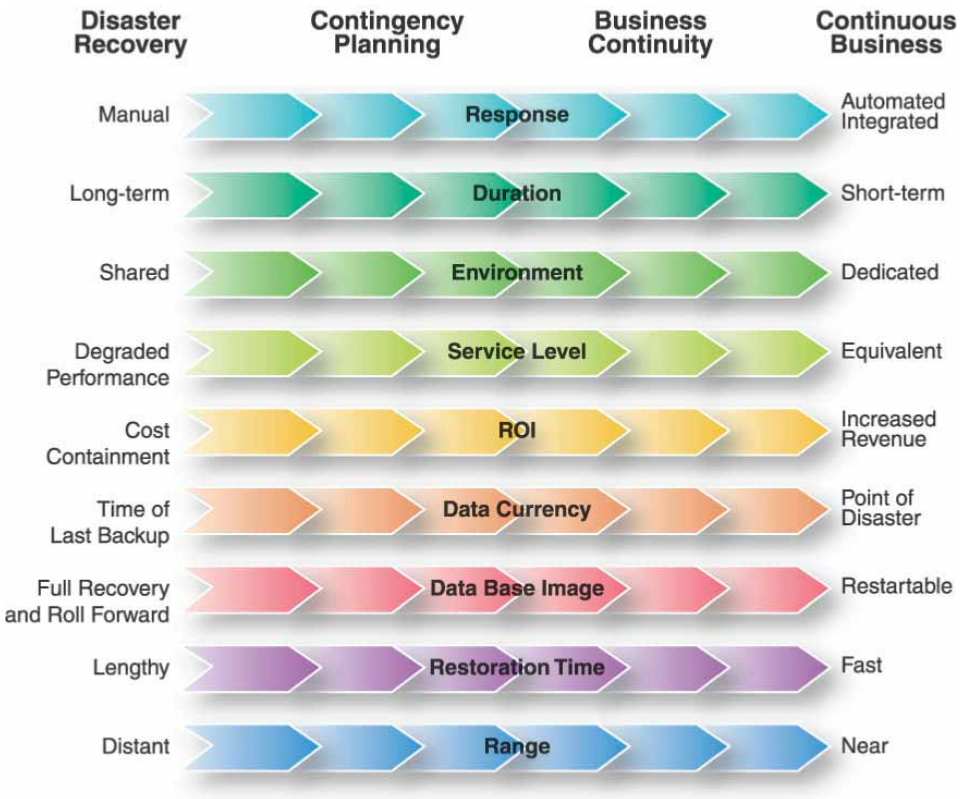

Figure 9 – Hitachi Data Systems changes the paradigm of high-availability computing from business continuity to continuous business.

Hitachi 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 Data Systems 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, back up, and manage all information vital to the enterprise business processes.

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

The Six-Sigma approach reinforces the Hitachi Data Systems reputation for nonstop

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. Hitachi looks at the Scale of 9s from an entirely new perspective. By including both planned and unplanned downtime in the availability equation, Hitachi 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 Hitachi Software Solutions.

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 5 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 5.

Table 5 – Through the Six-Sigma approach, Hitachi Data Systems attacks both planned and unplanned downtime.

	Availability Based Only on Unscheduled Downtime		Availability Based on Scheduled and Unscheduled Downtime			
9s Scale	**Unscheduled Downtime/yr.**	**Percent Availability**	**Scheduled Downtime/yr.[2]**	**Total Downtime**	**Percent Availability**	**9s Scale**
1	876 hrs.	90.0	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	0.31 sec.	99.999999	144 hrs. (12/mo)	144 hrs.	98.35	1.83
9	0.03 sec.	99.9999999	96 hrs. (8/mo)	96 hrs.	98.90	1.89

[2] 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.

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

Hitachi Business Continuity 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 techniques for achieving Six-Sigma objectives.

The Hitachi Data Systems Business Continuity Software Solutions are intended to 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. These Solutions are shown in Table 6. 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.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390®
Business Continuity Solutions					
Hitachi TrueCopy	✔	✔	✔	✔	✔
Hitachi ShadowImage™	✔	✔	✔	✔	✔
Hitachi Extended Remote Copy		✔	✔		✔
Hi-Track®		✔	✔		✔
Hitachi Dynamic Link Manager™	✔	✔	✔	✔	
Alternate Pathing, Host Failover and Parallel Database Clustering	✔	✔	✔	✔	✔

Table 6 – Hitachi Data Systems Business Continuity Solutions.

It is important to understand the terminology of Hitachi Data Systems Business Continuity Software Solutions as they relate to generic storage terms for copy software. The jargon of copy software alternatives is made even more confusing when traditional backup methods are considered. Advances in technology have allowed new words such as "real-time," "point in time" (PiT) and "snapshot" to creep into the language of enterprise class storage. Copy products allow an enterprise to replicate, protect, and share data in dynamic new ways. The three main terms used for copy software are:

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

1. 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 necessary before implementaion. eg. IBM® Peer-to-Peer Remote Copy (PPRC), Hitachi TrueCopy, and EMC® SRDF™ Synchronous. Asynchronous Remote Copy uses methods of ensuring data sequencing by

timestamping and is used for longer distances. The purpose of remote copy is to protect the data in the event of a business interruption to enable business continuity. eg. Hitachi TrueCopy, HXRC/XRC™, GDPS™ and NanoCopy™.

2. 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. eg. Snapshot (StorageTek™), ShadowImage (Hitachi), and TimeFinder™ (EMC).

3. Data Migration

Software that moves data permanently from one storage device to another. This feature is different from data duplication in that at the end of the process there is only one copy of data. The purpose of data migration is to consolidate storage or upgrade to new systems. For example, see Hitachi RapidXchange (Chapter 5).

As shown in Table 7, 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, data center relocation, or new application testing.

Table 7 - Hitachi copy products provide remote copy, data duplication, and data migration capabilities.

Hitachi Copy Products	Data Migration	Data Duplication	Remote Copy
Hitachi TrueCopy basic – Thunder 9200, Lightning 9900, 9900 V Series	✔	✔	✔- 25 Miles[3]
Hitachi TrueCopy asynchronous extension – Lightning 9900, 9900 V Series	✔	✔	✔- Unlimited distance
Remote Volume Replication – Thunder 9200	✔	✔	✔
NanoCopy – Lightning 9900, 9900 V Series Only		✔	✔
HXRC asynchronous remote copy – Lightning 9900, 9900 V Series Only		✔	✔
Hitachi ShadowImage™ – Thunder 9200, Lightning 9900, 9900 V Series		✔	
Hitachi Data Protection Services – Thunder 9200, Lightning 9900, 9900 V Series	✔	✔	✔
Hitachi RapidXchange™ – Lightning 9900, 9900 V Series Only		✔	
HARBOR® File Transfer (HFT) – Lightning 9900, 9900 V Series Only	✔		

Hitachi TrueCopy Solutions

Hitachi TrueCopy can be used to move data locally, between multiple Lightning 9900, 9900 V Series systems, or between multiple Thunder 9200 systems (as shown in Table 8) within a data center, or remotely in geographically separated data centers. There is minimal impact on performance, and no disruption of system operations. Applications continue to run smoothly throughout the remote copy process.

[3] While 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.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390
Hitachi TrueCopy					
Hitachi TrueCopy basic	✔	✔	✔	✔	✔
Hitachi TrueCopy asynchronous		✔	✔		✔
NanoCopy		✔	✔		✔

Table 8 – Summary of Hitachi TrueCopy Solutions and the products they support.

Hitachi remote copy solutions deliver the choice customers need, regardless of computing environment. These solutions accommodate all major platforms for both local or remote copy and provide fully consistent data images for disaster recovery for the following supported operating systems:

- 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
- OS/390®

Hitachi Data Systems is the only vendor, other than IBM, demonstrating proven remote copy support for GDPS.

Hitachi Data Systems is the only vendor other than IBM, demonstrating proven remote copy support for Geographically Dispersed Parallel Sysplex™ (GDPS), an IBM service offering for system failover, workload balancing, and data mirroring on systems spread across two or more sites up to 40 kilometers (25 miles) apart. Hitachi Data Systems provides a number of services to help customers plan, integrate, and deploy a GDPS environment based on Hitachi Freedom Storage systems. This is illustrated in Figure 15, later in this chapter, after a discussion of the importance of timestamping.

Hitachi TrueCopy Basic (Synchronous)

Synchronous remote copy provides the highest degree of data currency and integrity but can operate only at limited distances.

Synchronous remote copy provides the highest degree of data currency and integrity, but can operate only at limited distances. In synchronous communications, the sender and receiver need to synchronize with one another before data is sent. This generally limits the communication between sender and receiver to short distances, typically less than 40km.

Storage system-based synchronous solutions include IBM PPRC, EMC SRDF and Hitachi TrueCopy basic as shown in Figures 10 and 11.

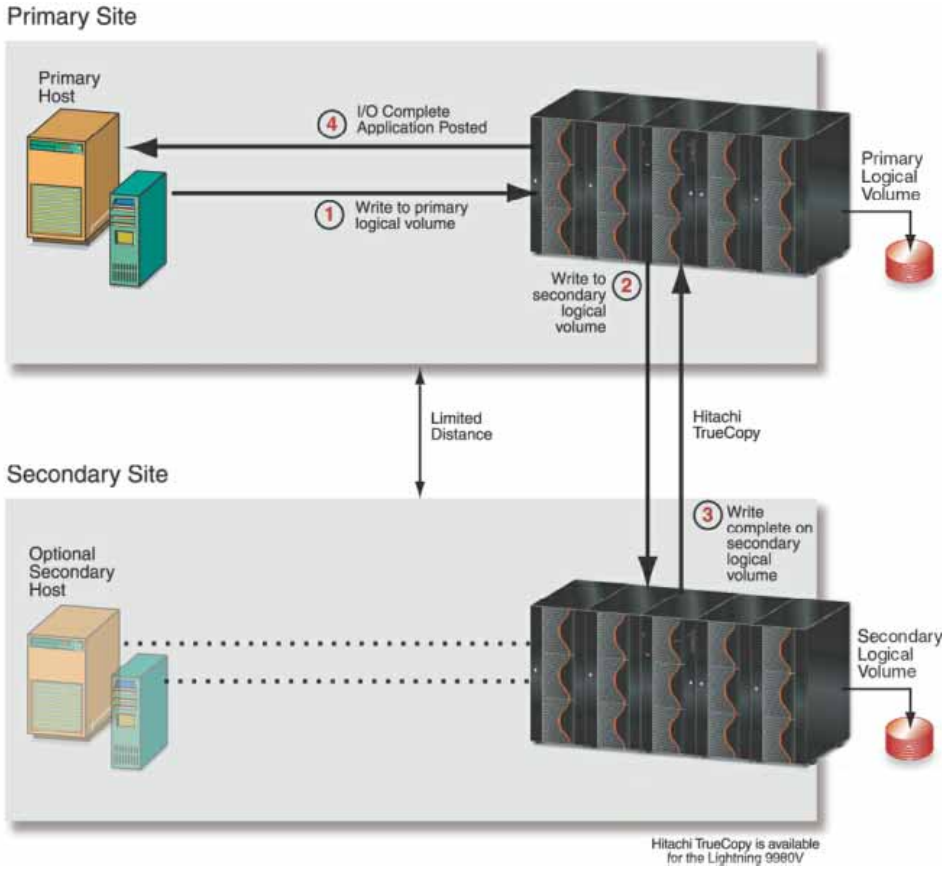

Figure 10 - Hitachi TrueCopy basic for open systems environments.

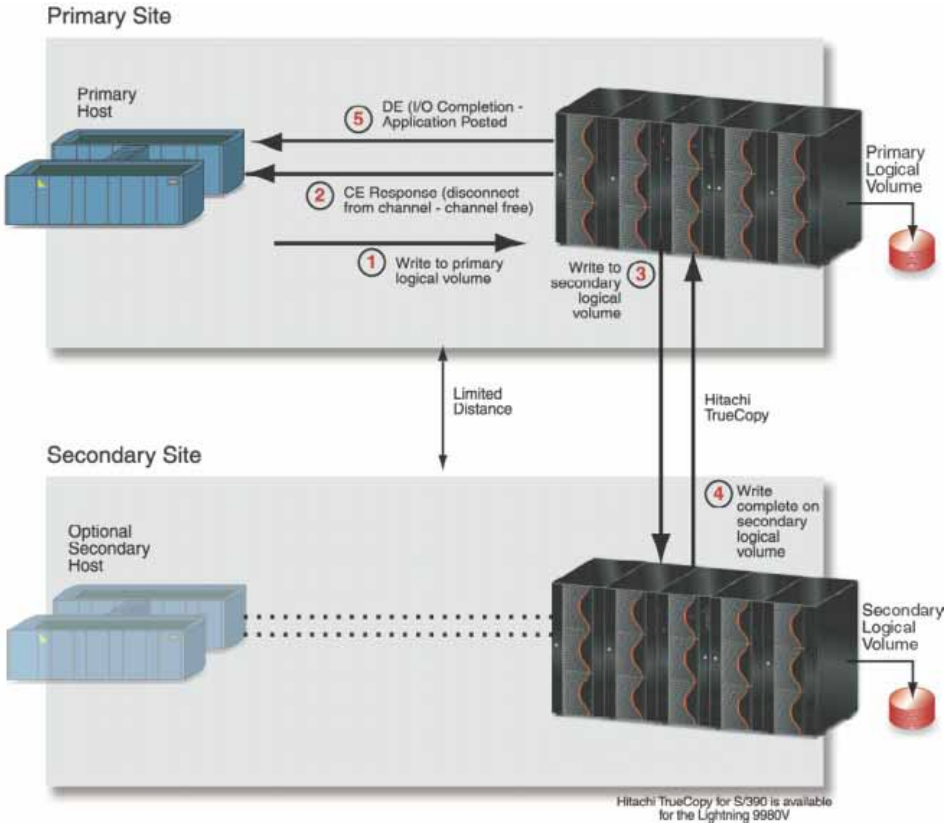

Figure 11 - Hitachi TrueCopy basic for S/390 environments.

Synchronous remote copy, despite being the most popular form of remote copy technology currently deployed, has a performance deficiency as I/O response times increases with the distance between the primary and secondary sites. This is due to what is known as "propagation delay," or the fact that, on average, the speed of light takes one nanosecond to travel one foot. This can add up when many miles are involved and synchronous copy techniques are effectively limited to distances of 25 miles or 40 kilometers (10 kilometers for Fibre Channel) depending on performance requirements. These limitations have been solved by Hitachi and are discussed later in this chapter.

Hitachi TrueCopy Asynchronous

Asynchronous remote copy is not "real time" like synchronous remote copy, but it can operate at any distance and has the data and transaction integrity necessary for disaster recovery.

Hitachi TrueCopy asynchronous extension for remote copy is not "real time" like synchronous remote copy, but it can operate at any distance and has the data and transaction integrity necessary for disaster recovery. This is very important to many businesses since many causes of U.S. computer outages affect large geographic areas so that synchronous remote copy at ESCON distances does not provide real protection. One out of four U.S. computer outages is caused by storms, hurricanes, or floods that span an area greater than that protected by limited-distance synchronous remote copy techniques. Hitachi Data Systems offers a more complete selection of asynchronous copy techniques than any other vendor and offers a number of solutions to address today's 24/7 anywhere business continuity requirements. Most importantly, other asynchronous remote copies do not have transaction integrity.

In open systems environments, timestamps from the Lightning 9900 V Series or system timer are used for data synchronization; in S/390 configurations the host timer or Sysplex Timer® is used. Utilizing reliable timestamps, along with other information created by the primary system(s), allows the Lightning 9900 V Series to transmit updates directly to the secondary system (without any host intervention) and have the secondary system apply the same consistency group algorithms normally provided by the System Data Mover. This approach allows the enterprise to achieve data integrity in a significantly simpler and less costly hardware-based solution with no impact on server or application performance. Figure 12 shows how the Lightning 9900 V Series systems with Hitachi TrueCopy asynchronous extension utilizes timestamps to provide the correct sequencing of "writes" to the secondary volume at the remote site[4].

[4] Lack of individually timestamped I/Os can force an unplanned recovery situation.

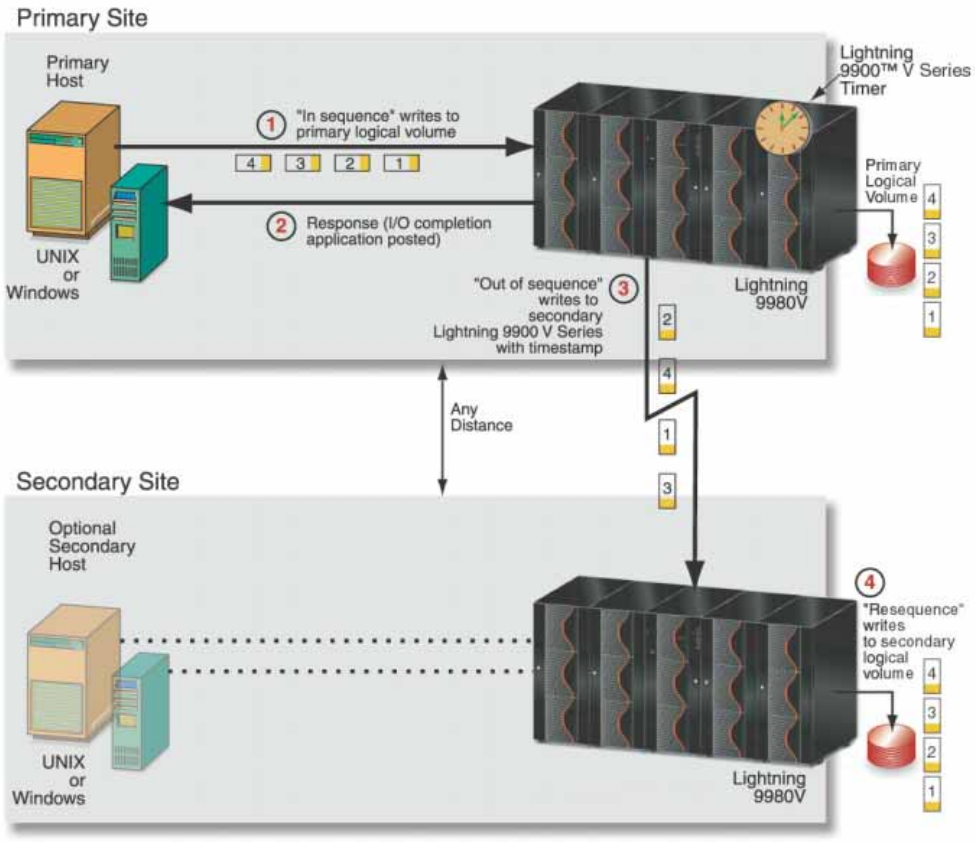

Figure 12 - Hitachi TrueCopy asynchronous extension for open systems.

In the open systems example of Hitachi TrueCopy asynchronous extension in Figure 13 and the S/390 example in Figure 10, the primary control units are asynchronously connected to control units on the secondary system using reliable inter-system timestamps. Writes (timestamped 1 to 4 in the Figures) are sent asynchronously (i.e., "out of sequence") to the secondary system. These writes are then buffered in the secondary system's cache, queued, sorted by the "timestamp" and written to the correspondent volumes in the same sequence as issued by the primary system over the remote link checking sequence numbers embedded in records to ensure no records are missing. In this way, I/O consistency is maintained.

There are six major characteristics of Hitachi TrueCopy asynchronous extension.

1. Data integrity is guaranteed for dependent write applications.
2. Excellent performance for both long and short distance requirements, due to its asynchronous nature.
3. In S/390 up to four primary systems are supported and in open systems a one-to-one configuration is supported per consistency group.
4. The maximum distance between the units is 25 miles or 43 kilometers with ESCON connections or 10 kilometers with Fibre Channel links; for greater distances channel extenders can be used.
5. Control is performed from a GUI remote console interface, via a command line interface for open systems (command control software), or via either host PPRC TSO commands or GUI remote console for S/390.
6. Synchronous, asynchronous, open, and mainframe may be intermixed within the same Lightning 9900 V Series system.

Hitachi TrueCopy asynchronous extension in mainframe environments is an excellent product for data relocation and migration activity as well as PiT and real-time copy. Hitachi TrueCopy asynchronous extension is currently supported in a one-to-one configuration for open systems and in a four-to-one ratio for mainframes per consistency group as shown in Figure 13[5]. Given that the Lightning 9900 V Series systems scale up to 74TB+ raw capacity, it is unlikely that for consistency purposes a customer in open systems would outgrow a one-to-one configuration. In MVS® environments, NanoCopy is available for larger configurations as is discussed later in this Chapter.

Hitachi TrueCopy asynchronous extension in mainframe environments is an excellent product for data relocation.

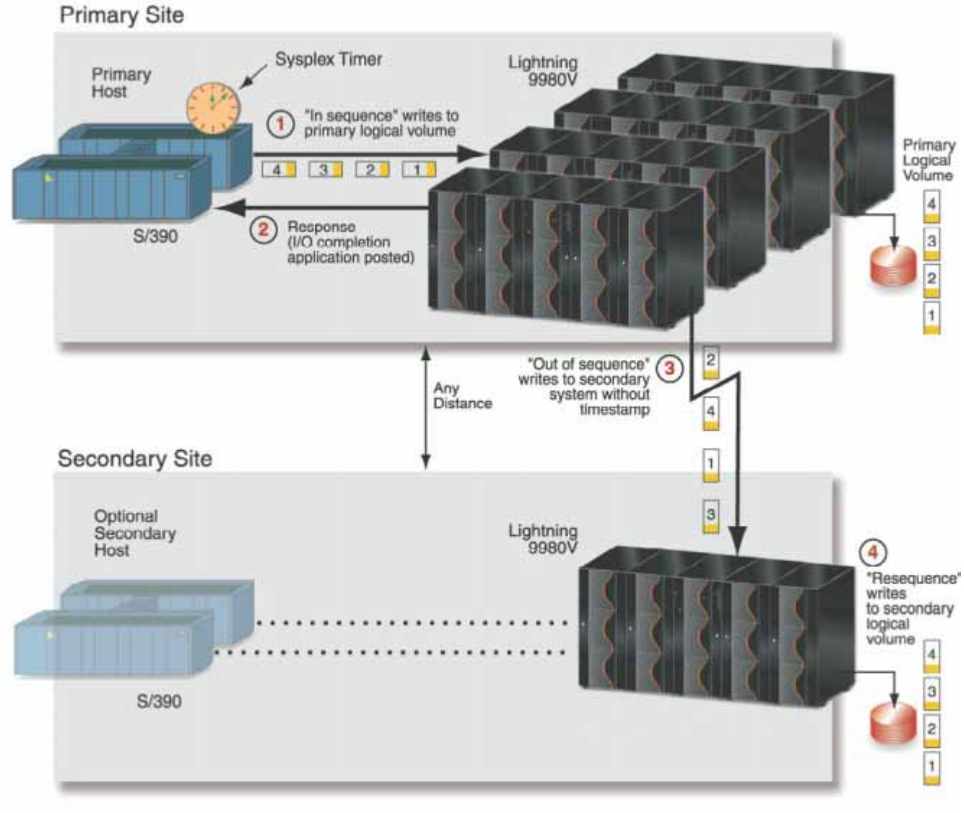

Figure 13- Hitachi TrueCopy asynchronous extension for S/390 environments.

As shown in Figure 14, some competitive products use a technique based on sending changed tracks to the secondary system instead of timestamped I/Os. Users can either specify an "invalid tracks" value, which is the number of changed tracks that can accumulate before that changed data is sent over to the secondary system in full-track images. The problem with this technique is that applying changed tracks (and not individually timestamped I/Os) cannot preserve the original sequence of writes and therefore should not be used for real-time disaster protection or business continuance. Or, in another approach, a table of the changed tracks is accumulated over some period of time (for example, a 24-hour period). At this point, the changed tracks are sent to the secondary location, and when all changed data has been sent, the primary location is quiesced (outage incurred), at which point the two sites are again suspended and the primary site/applications can be restarted. This latter technique is analogous to disk vaulting since it is essentially a point-in-time copy requiring an outage.

[5] This limitation is only for data that has to be consistent across multiple systems (i.e. Application A and B could run on separate one-to-one configurations if they were separate entities for recovery purposes).

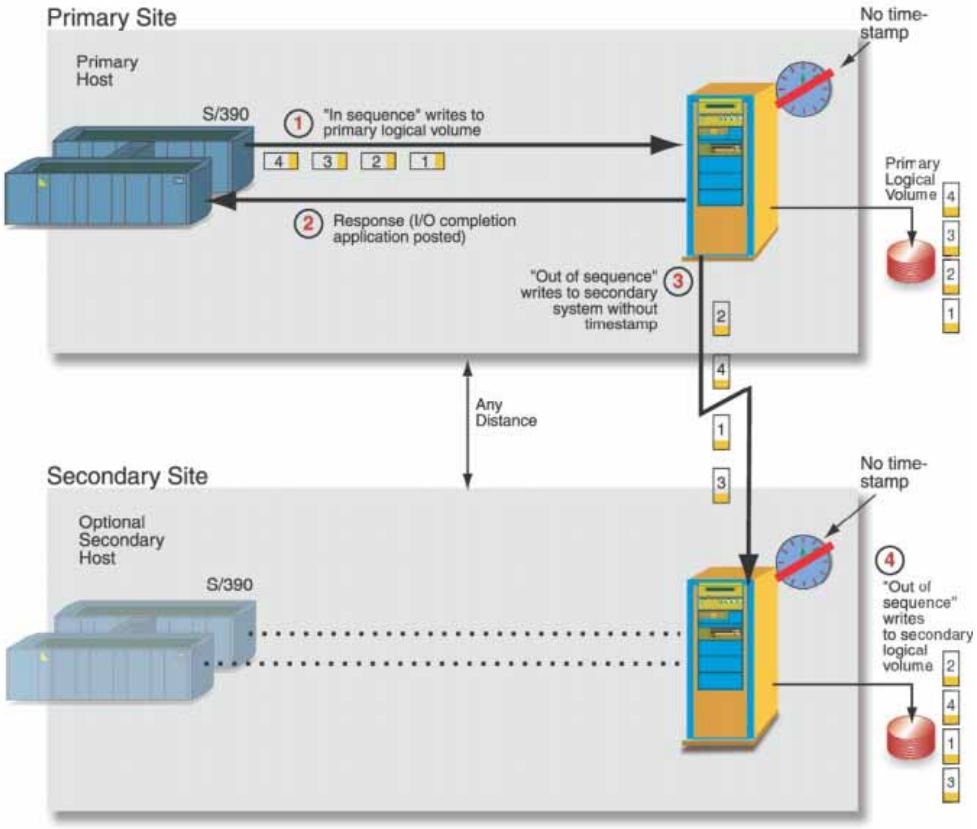

Figure 14 - Not using individually timestamped I/Os can force an unplanned recovery situation.

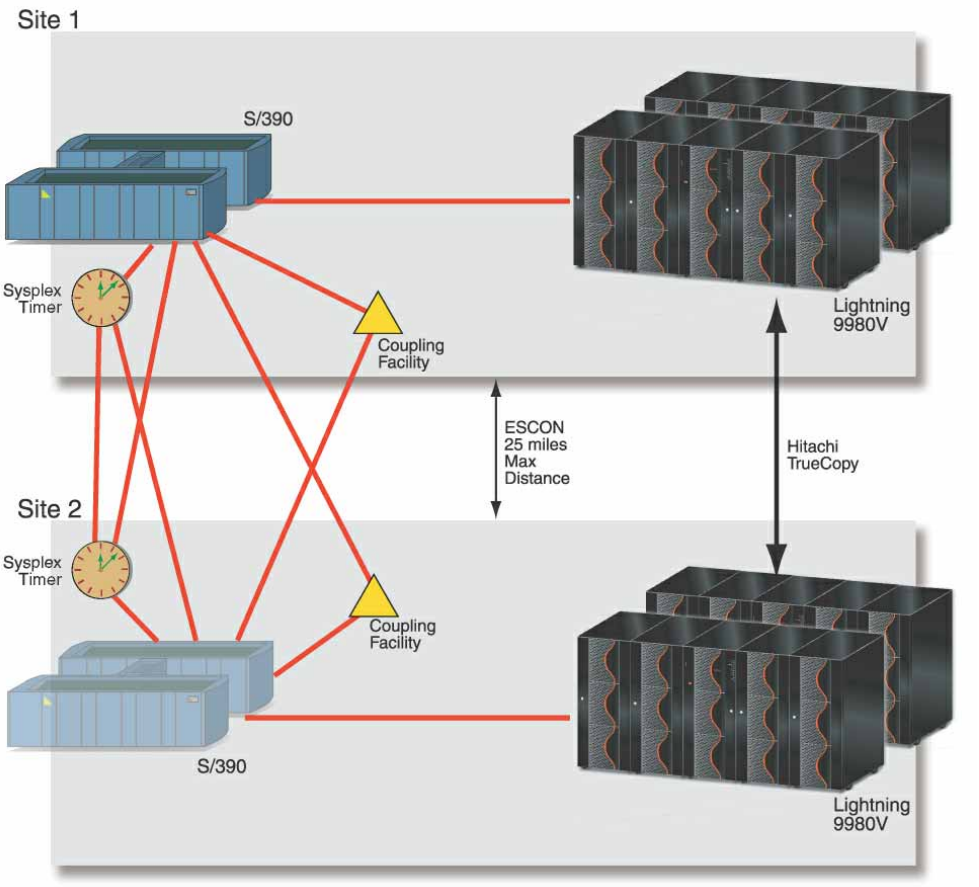

Figure 15 - Hitachi TrueCopy is compatible with IBM GDPS clusters.

NanoCopy

In 1999, Hitachi Data Systems announced the patented NanoCopy feature for the Lightning 9900 Series and in May 2002 for the Lightning 9900 V Series. With these functions, Hitachi leapfrogged competitors by delivering a solution that removes some of the drawbacks of synchronous remote copy (such as maximum distance and data integrity across multiple systems) without the processor performance overhead usually associated with asynchronous remote copy implementations like XRC.

Hitachi TrueCopy asynchronous for open systems and Hitachi TrueCopy asynchronous for S/390 environments are best described as "XRC in a box." The System Data Mover component of XRC contains a number of consistency group algorithms that allow it to maintain data integrity. These algorithms have now been implemented in the microcode of the secondary Lightning 9900 Series and Lightning 9900 V Series systems.

Working with Hitachi TrueCopy asynchronous, Hitachi NanoCopy extension provides for the world's only non-disruptive I/O-consistent, and time-consistent long-distance remote copy.

Working with Hitachi TrueCopy asynchronous, NanoCopy provides the world's only non-disruptive I/O- and time-consistent long-distance remote copy that can be used safely for Disaster Recovery and PiT copy. NanoCopy software allows "any-to-any" remote copy. This means NanoCopy software provides remote copy with integrity for any number of primary storage systems connected to a single Sysplex Timer to any number of remote image systems anywhere in the world, at any time.

Building on Hitachi TrueCopy asynchronous, NanoCopy uses a cyclical approach to protect any amount of information at unlimited distance, regardless of applications database or data residency (even if resident on multiple controllers). Using a mainframe software routine, a user can set any time for point-in-time copies to be taken, without any outage to the system or impact on performance. Each system suspends the remote data transfers at the same time, creating an I/O-consistent copy (the same "consistency time") across any number of systems. The suspended volume pairs can be resynchronized later.

The major impact of NanoCopy on PiT copy technology is that it can create a copy, at any distance, without having to quiesce, shut down, or otherwise interrupt the application that is using (updating) the data. For the first time in the industry, a PiT copy of any amount of data can be created, at any distance, without causing an outage to the application or system and without performance impact. This unique capability vaults Hitachi Data Systems ahead of all other storage vendors for PiT copy capability.

NanoCopy offers two basic capabilities:

- PiT copy
- Disaster protection through real-time copy.

NanoCopy has the potential for providing many new solutions to the enterprise.

NanoCopy comes closest to the perfect solution by allowing customers to create a truly non-disruptive PiT copy. That is, for the first time a storage-based copy solution creates an image of any amount of data without incurring an application or system outage. Furthermore, that copy can be made to any place on the globe. For disaster protection, NanoCopy is the only product that solves the extraordinarily complex problem of ensuring that critical data can survive a rolling disaster. With outstanding data integrity, NanoCopy not only maintains a copy of critical data at any distance, but also does so using a surprisingly simple and low-cost technique that has minimal impact on critical applications.

NanoCopy has the potential for providing many new solutions to the enterprise. It clearly offers a simple and low-cost alternative to existing remote copy technologies without compromising data integrity. It also provides an industry-first capability of taking point-in-time copies without any disruption to applications. Hitachi Data Systems Professional Services consultants are available to tailor, customize, and install the NanoCopy product and associated software to meet the most rigorous customer requirements. NanoCopy implementation service is discussed further in Chapter 6.

Although NanoCopy is currently available only on the OS/390 platform, Hitachi Data Systems intends to support open systems with the same function, but limited to one primary subsystem and one remote subsystem per consistency group. The one subsystem restriction is required in an open system environment because there is no Sysplex Timer to provide a consistent time source for the entire environment. Instead, Rapid Recovery uses the primary Lightning 9900 V Series clock as a time source. The key prerequisite for NanoCopy in open systems is a dependable timestamp between primary and secondary systems. This is currently lacking in UNIX and Windows NT or 2000 environments. NanoCopy is one of many leading-edge products available from Hitachi Data Systems. NanoCopy ensures the safety, integrity, and availability of data in ways that were previously impossible.

With NanoCopy, all updates are timestamped and the secondary system uses the timestamps to recreate the original sequence of updates.

The Evaluator Group, a key industry consultant, summarizes the benefits of NanoCopy as follows:

> "The combination of Hitachi TrueCopy asynchronous and NanoCopy will provide a relatively simple and cost-effective means of implementing disaster recovery across extended distances. Evaluator Group believes that with these new products, Hitachi has assumed the leadership position in the area of disaster recovery."

NanoCopy and Rapid Recovery build on the current Hitachi TrueCopy asynchronous product, which in turn, is built upon the hardware-based concepts of Hitachi TrueCopy. All updates are timestamped and sequence-numbered, and the secondary system uses the timestamps and sequence numbers to recreate the original sequence of updates. NanoCopy has multiple components:

- Additional microcode in the Lightning 9900 V Series to support the two NanoCopy commands. These two new commands are NCFREEZE and NCQUERY.
- Host executables (currently in the form of TSO commands on OS/390 or RAID Manager for open system environments) to execute the commands to the storage subsystems.
- Hitachi TrueCopy asynchronous and ShadowImage.
- Automation scripting that allows NanoCopy to be used in a disaster recovery environment.

In addition, Hitachi Data Systems offers professional services that can adapt NanoCopy to meet specific customer needs and configurations. The enhancements that NanoCopy provide are quite simple. Since all updates are timestamped, NanoCopy uses that timestamp to create the I/O-consistent PiT copy. By using NCFREEZE, one of the two new commands available, NanoCopy extension is able to "freeze" a second disk image at an I/O-consistent point in time. This process is shown in Figure 16.

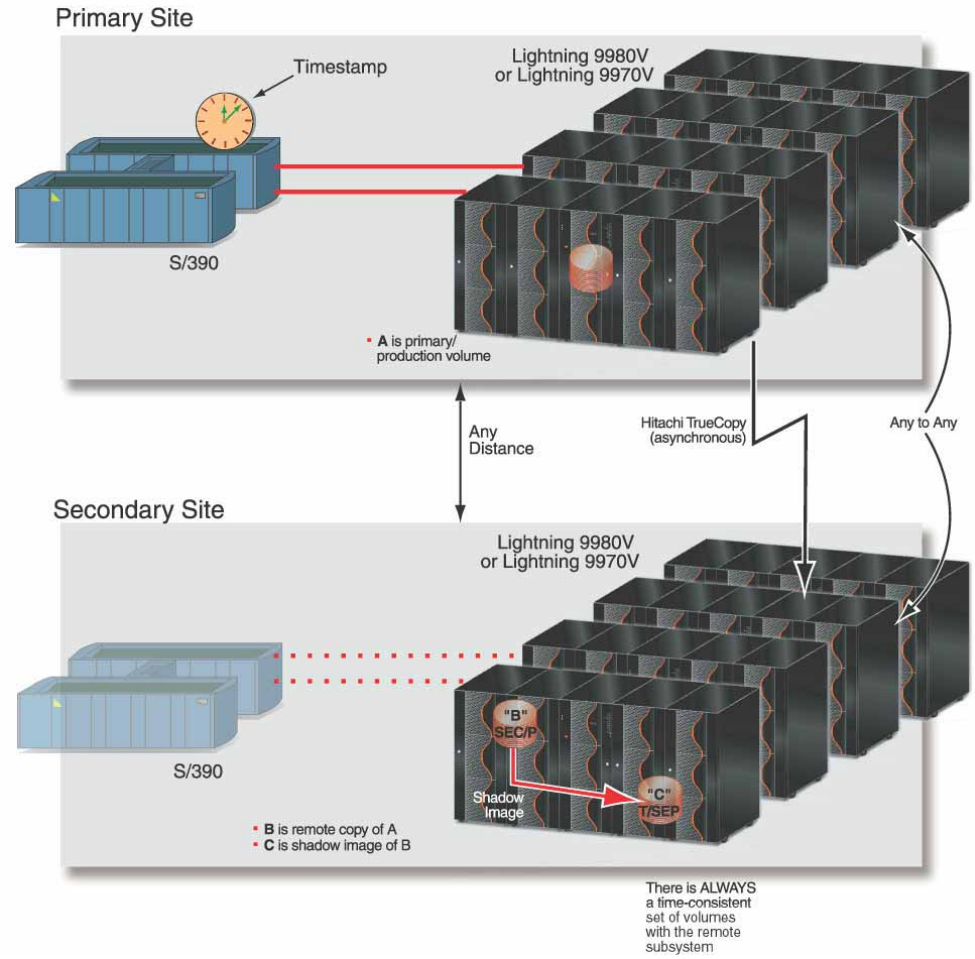

Figure 16 - For disaster recovery, NanoCopy triplets ensure there is always a time-consistent set of volumes in the remote system that is not being updated.

The PiT (Point in Time) NanoCopy NCFREEZE command allows customers to specify a future time value at which all primary systems will terminate communication with all secondary systems. For example, if customers wanted to create an I/O-consistent copy of any amount of data at 9:30, they could issue NCFREEZE TIME (9:30) commands to every storage system involved at the primary site at 9:25. By broadcasting this future time value, they tell each system at which point it is to suspend sending updates to the target. Since all write activity is timestamped, each source system will, upon detecting an update with a timestamp later than 9:30, immediately go into suspend mode. Suspend mode means the primary system would stop sending updates to the secondary system, and begin recording which data are being updated. Put another way, every update up to and including 9:30 has been sent to the target system. All updates occurring after 9:30 are recorded in a bitmap at the source storage system.

This provides the following benefits:

1) An I/O-consistent target copy at exactly 9:30 hours
2) Protection against rolling disaster
3) The ability to copy only changed data to refresh the target copy asynchronously
4) Recovery automation by removing operational decision making requests
5) Reduced recovery time at the remote site
6) Logs of all processes
7) Reports of all problem areas

Hitachi ShadowImage

Hitachi ShadowImage enables open systems users to maintain system-internal copies of Logical Units or Logical Volume Images (LUNs or LVIs) of data for open systems or mainframe environments as shown in Table 9. These copies may be used for purposes such as data backup or data duplication. The RAID-protected duplicate LUNs are created within the same Lightning 9900 Series, Lightning 9900 V Series, or Thunder 9200 systems as the primary LUN at hardware speeds. Once established, ShadowImage operations continue unattended to provide asynchronous internal data backup. ShadowImage operations are nondisruptive; the primary LUN of each ShadowImage pair remains available to all hosts for both read and write operations during normal operations. Usability is further enhanced through a resynchronization capability that reduces data duplication requirements and backup time, thereby increasing user productivity. ShadowImage also supports reverse resynchronization for maximum flexibility.

Table 9 – Summary of Hitachi ShadowImage solutions and the products they support.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390
Hitachi ShadowImage™					
Hitachi ShadowImage	✔	✔	✔	✔	✔
Hitachi ShadowImage LVD		✔	✔	✔	✔
Hitachi ShadowImage FlashCopy® Extension	✔	✔	✔	✔	✔

ShadowImage operations can be performed in conjunction with Hitachi TrueCopy to provide multiple copies of critical data at both primary and remote sites. ShadowImage also supports the Virtual LVI/LUN and the Hitachi ShadowImage LVD extension. ShadowImage FlashAccess features of the Hitachi ShadowImage FlashCopy extension ensure that all user data can be duplicated by ShadowImage operations.

In a technology exchange with IBM, ShadowImage FlashCopy extension is a software solution whereby the ShadowImage engine responds to all IBM FlashCopy commands from TSO, DFSMSdss and SIBBATCH just like it was IBM FlashCopy.

ShadowImage supports the creation of up to nine (three for S/390) RAID-protected image copies from each source volume within a Lightning 9900 V Series system at hardware speeds. This is shown in Figure 17. There are two key advantages of the Hitachi Data Systems implementation of ShadowImage:

- RAID protection
- Customer control of copies

The image copy can be of the same or different RAID type. Up to 2048 image copies are available per Lightning 9960/9910 systems and 4096 for Lightning 9980V/9970V systems. Once split from the source volume, each copy is independently available for read/write activity.

Hitachi ShadowImage further enhances usability through a resynchronization capability that reduces data duplication requirements and thus increases user productivity. When used with Hitachi TrueCopy, the duplication process allows users to maintain up to eight image copies in S/390 (Lightning 9900 Series and Lightning 9900 V Series only) and twenty in open systems, RAID-protected copies of critical data.

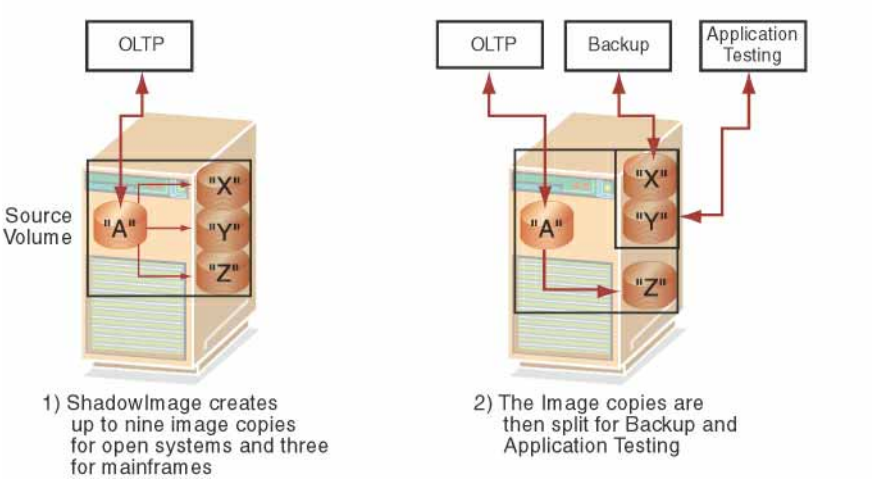

Figure 17 - Hitachi ShadowImage supports up to nine image copies in open systems and three image copies in OS/390 from a single source volume. The image copies can then be split for other uses such as backup or application testing.

There's nothing more reassuring to IT managers than the knowledge that their company's critical data is fully protected. The high-speed, problem-free data duplication provided by Hitachi ShadowImage delivers uninterrupted 24/7 availability enterprise-wide. Part of Hitachi Data Systems software solutions, ShadowImage makes it possible to duplicate LVIs non-disruptively in the Hitachi ShadowImage LVD extension. At the same time, primary or secondary remote copy pair volumes can be replicated while maintaining full remote copy capabilities. But it's not just what ShadowImage does that's important. It's what Hitachi ShadowImage enables users to do with it.

There's nothing more reassuring to IT managers than the knowledge that their company's critical information is fully protected.

Hitachi ShadowImage Provides the Only Online Backup Solution for Multi-terabyte Mission-critical Oracle® Databases

Another major benefit of the Hitachi Freedom Storage software solutions is the high-speed, non-disruptive backup-to-tape solution for users of Oracle databases, as shown in Figure 18. This new capability in data warehousing, which solves one of the most daunting problems facing IT managers today, is the result of an effort to tightly integrate Hitachi ShadowImage software with Oracle Database Management System Hot Backup capability. ShadowImage can provide up to nine copies of a database volume so that the same volume can be used simultaneously for testing, development, data-mining queries, and off-site replication, as well as backup, substantially increasing the value of the solution.

Hitachi ShadowImage software solution provides a high-speed, non-disruptive backup-to-tape solution for users of Oracle databases.

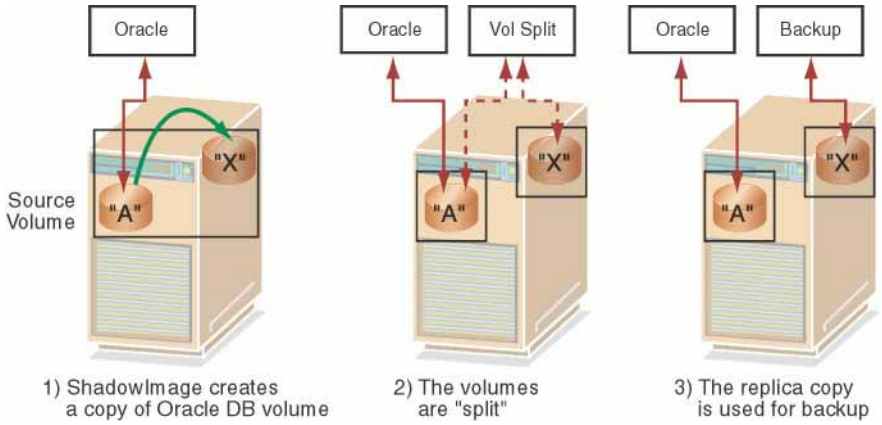

Figure 18 - Hitachi ShadowImage can provide high-speed backup of open database systems, such as Oracle.

Hitachi Extended Remote Copy (HXRC)

Hitachi Extended Remote Copy (HXRC) is compatible with IBM Extended Remote Copy (XRC) software and is based on IBM asynchronous remote copy technique that is designed to provide complete data integrity for primary and remote systems. This is accomplished through the use of timestamps that allow the asynchronous updates to be formed into consistency groups by the System Data Mover (SDM) host software. Additional support is provided through industry standard (IBM) compatibility with TSO command support and DASD ERP (error recovery programs/procedures).

XRC overcomes distance limitations imposed by synchronous remote copy at the expense of potential data loss in the event of a disaster. However, the recovery time to restore full operations is minimized through the use of the System Data Mover and in most cases, the shorter time to accurately recover from a disaster far outweighs the cost of data that might be lost. XRC is available on both Lightning 9900 and Lightning 9900 V Series.

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

Hitachi Dynamic Link Manager™ Provides Path Failover and Load Balancing

Hitachi Dynamic Link Manager is a family of Hitachi-provided middleware software utilities that are server-based as shown in Figure 19. Dynamic Link Manager enhances the availability of RAID systems by providing automatic error recovery and path failover from server-to-RAID connection failures. Dynamic Link Manager provides load balancing in addition to path failover by re-directing I/O activity to the least busy path using complex algorithms.

Just because a system is RAID-protected doesn't mean it is protected against connection bus failures, which is why Dynamic Link Manager is required for true nonstop operations. This product allows systems administrators to take advantage of the multiple paths on a Lightning 9900 V Series system by adding redundant connections between data servers and RAID systems. Hitachi Dynamic Link Manager therefore provides increased reliability and performance. Supported platforms include, IBM AIX, Sun Solaris, Windows NT, and Windows 2000. Hitachi Dynamic Link Manager is also available for Lightning 9900 and Thunder 9200 storage systems.

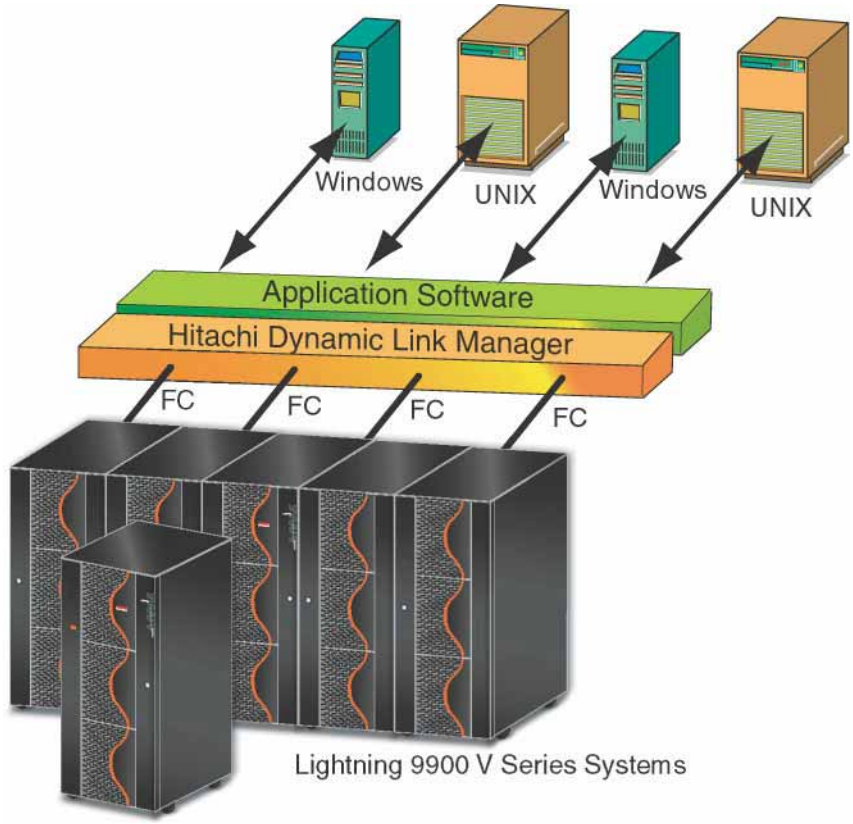

Figure 19 - Hitachi Dynamic Link Manager automatically provides path failover and load balancing for open systems.

Alternate Pathing, Host Failover, and Parallel Database Clusters

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

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

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

This type of middleware automatically switches the I/O load on a failed primary path to an alternate path on the same host system. The Lightning 9900 V Series systems support alternate pathing for AIX 4.2 and above, Compaq Tru64 UNIX, HP/UX and through PVlink, Windows NT/2000 and Sun Solaris and Sequent DYNIX/ptx. In addition, VERITAS Dynamic Path Management™ (DPM) is supported.

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

Host Failover

Host failover software supports a cluster of host processors in which one of the hosts automatically takes over the workload of any failed host in the cluster. This "takeover" includes the reassignment of networks and peripherals, as well as the restarting of

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

applications. Host clustering can also be used to create fault-tolerant workloads and scale processor capability while sharing network and disk resources.

The Lightning 9900 V Series supports all major open systems clustering schemes including: Compaq TruCluster™, HP® MC ServiceGuard®, 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 FirstWatch® for Sun Solaris.

Parallel Database Clustering

Clustering middleware supports distributed lock management.

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

The Lightning 9900 V Series systems have been certified with MC/Lock Manager for Oracle Parallel Server Version 8, and Sun PDB for Oracle Parallel Server Version 8.

The Hi-Track Remote Maintenance Tool Ensures the Accuracy of Business Continuity Solutions

The Hi-Track remote maintenance tool of the Lightning 9900 V Series systems generates service information messages (SIMs) to identify normal operations (such as Hitachi TrueCopy pair status change) as well as service requirements, errors, or failures. SIMs can be generated by the CHIP and ACP microprocessors and by the SVP. All SIMs generated by the Lightning 9900 V Series systems are stored on the SVP for use by Hitachi Data Systems personnel, logged in the SYS1.LOGREC dataset of the S/390 host system, displayed by the Web Console software, and reported via SNMP to the open-system host. The SIM display on the Lightning 9900 V Series Web Console PC enables users to remotely view the SIMs reported by the attached Lightning 9900 V Series systems. Each time a SIM is generated, the amber Message LED on the Lightning 9900 V Series control panel turns on. The Hi-Track remote maintenance tool also reports all SIMs to the Hitachi Data Systems Support Center.

Backup and Recovery Solutions

4

Overview of Hitachi Freedom Storage™ Backup and Recovery Software Solutions

Backup/Restore is the largest cost element in storage ownership and the most important to high-availability computing. In today's global IT environment, it is no longer enough to have a backup copy of one's data; businesses must be able to perform backup and restore in the shortest possible time and with minimum disruption to information availability. As shown in Table 10, Hitachi Data Systems combines its own leading-edge Business Continuity software solutions with those of industry-leading software partners to address customers' backup/restore requirements with solutions that reduce TCO.

Hitachi Data Systems provides backup/restore solutions that reduce the total cost of computing.

Table 10 - Summary of Hitachi Freedom Storage backup and recovery solutions and the products they support.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390®
Back-up and Recovery Solutions					
Hitachi Multiplatform Back-up/Restore		✔	✔		✔
Tantia Technologies® HARBOR® Backup/Recovery and HYPERtape		✔	✔		✔
Non-disruptive database backup	✔	✔	✔	✔	✔
Non-disruptive application backup	✔	✔	✔	✔	✔

Hitachi Multiplatform Backup/Restore delivers high-performance volume-level backup of open systems data. In addition to the HARBOR Backup/Recovery (HBR) and HYPERtape Backup/Restore software packages, Hitachi offers several features and services to achieve this goal. Hitachi ShadowImage™ and Hitachi TrueCopy for S/390® environments can be used for the creation of a snapshot copy for near-instantaneous backup.

All Hitachi backup/restore solutions are designed to meet the following customer requirements:

- Maximum reliability
- Centralized administration
- File-level backup/restore
- Online backup of popular databases, applications, and messaging services
- Out-of-the-box agents for popular servers such as Oracle® Server, BMC Software™
- Multiplatform backup to a mainframe or open systems server
- High-performance movement of data
- Ability to "snapshot" the backup source in UNIX®, Windows NT®/Windows® 2000 or MVS®
- Ability to simulate or test the recovery scenario.

Backup/Recovery Is Critical to High-availability Computing

What value do you place on the ability to rapidly restore data and recover applications?

Backup/Recovery is a key component in protecting information. As the number of platforms supported in an enterprise increases, it is increasingly important for an enterprise to choose "best of breed" backup/restore solutions. The GartnerGroup estimates that Backup and Restore software and operations costs account for 32 to 54 percent of total costs. These costs are expected to grow with the increasing need for 24/7/365 operations, the implementation of disaster recovery, and the trend towards new applications that

generate tremendous amounts of data. Applications such as collaborative computing (Lotus Notes® and Microsoft® Exchange), data warehousing/decision support, and business process automation (SAP®, BAAN®, PeopleSoft®) dramatically increase the need for backup/restore. Office mail systems are often backed up two or three times a day, and must be recoverable to the document or message level. Business process automation systems are built upon relational databases that provide a single view of data for the entire business from sales and order entry to shop floor control. These requirements place high value on the ability to rapidly restore data and recover applications.

Six Alternative Techniques for Backup and Restore

1. Volume Level Backup/Restore

In order to have consistent, volume-level backup of data, all applications must be stopped.

Of the six major types of backup and restore techniques, Volume Level Backup/Restore is the easiest to do from a hardware/software viewpoint. However, the enterprise must implement controls to ensure the logical and temporal consistency of the data on the volume with the application view of the data. In order to have consistent, volume-level backup of data, all applications must be stopped so that the data in memory and cache can be flushed to the disk. If this volume is part of an application database, it must be coordinated with the backup of the application's logs, indices, repository, updates, exports, and recovery scripts, which are often spread across multiple volumes or generations of volumes. A volume is an MVS concept and datasets are contained within a volume with a description of the dataset recorded in a volume table of contents (VTOC).

In UNIX or Windows NT/2000, the closest representation to a volume is a SCSI address (Logical Unit or LUN) or mount point (C disk, D disk, and so on). Unlike MVS, open systems files will often span multiple LUNs and hierarchical directories within the file system define the file locations. Unless the entire file system is backed up at the same time, the restore of volume-level (LUN) backups is not possible without file-level information from the host file system.

2. Hitachi Multiplatform Backup/Restore Provides Mainframe Benefits

With the Hitachi Multiplatform Backup/Restore feature enabled, a user can back up any open system to an MVS or VSE system.

Multiplatform volume backups between S/390 and UNIX or Windows NT/2000 are available with Hitachi Multiplatform Backup/Restore on the Hitachi Freedom Storage™ Lightning 9900™ V Series systems. With this feature enabled, a user can back up any open system to an OS/390 or VSE™ system, using any standard utility such as DFDSS or FDR. Multiplatform Backup/Restore is the fastest way to back up LUNs to a mainframe and take advantage of the strong infrastructure of mainframe backup operational procedures.

By using existing backup procedures and Automatic Tape Loading (ATL) environment, Multiplatform Backup/Restore eliminates the need for special software or procedures. The "any-to-any" connectivity of the Lightning 9900 V Series system enables S/390 to access these volumes as 3390-3 or 3390-9 devices, making them accessible via standard utilities such as DFDSS or FDR. By working with existing, mainframe-based storage management procedures, Multiplatform Backup/Restore provides a simple yet effective way for backup/restore of open systems volumes. Hitachi Multiplatform Backup/Restore may also be combined with Hitachi ShadowImage for near non-disruptive backup capability. This results in the following benefits to open systems backup:

IT managers are quickly learning that heterogeneous computing environments present significant challenges.

- The use of well-established and disciplined mainframe procedures reduces the chance of human error in backup and restore processing.
- Hitachi Multiplatform Backup/Restore capitalizes on high-speed S/390 tape systems.

- Hitachi Multiplatform Backup/Restore can even perform a restore of open system volumes without requiring any host involvement.

IT managers are quickly learning that heterogeneous computing environments present significant challenges. They must rationalize both the equipment and the processes involved in backup/restore operations for critical business data. This task is relatively easy when data resides on a single host or group of like hosts. But in multiplatform, multiserver environments, data is dispersed on a variety of platforms ranging from S/390 mainframes to all open environments.

Figure 20 demonstrates how open systems volumes can be backed up directly by the S/390 host.

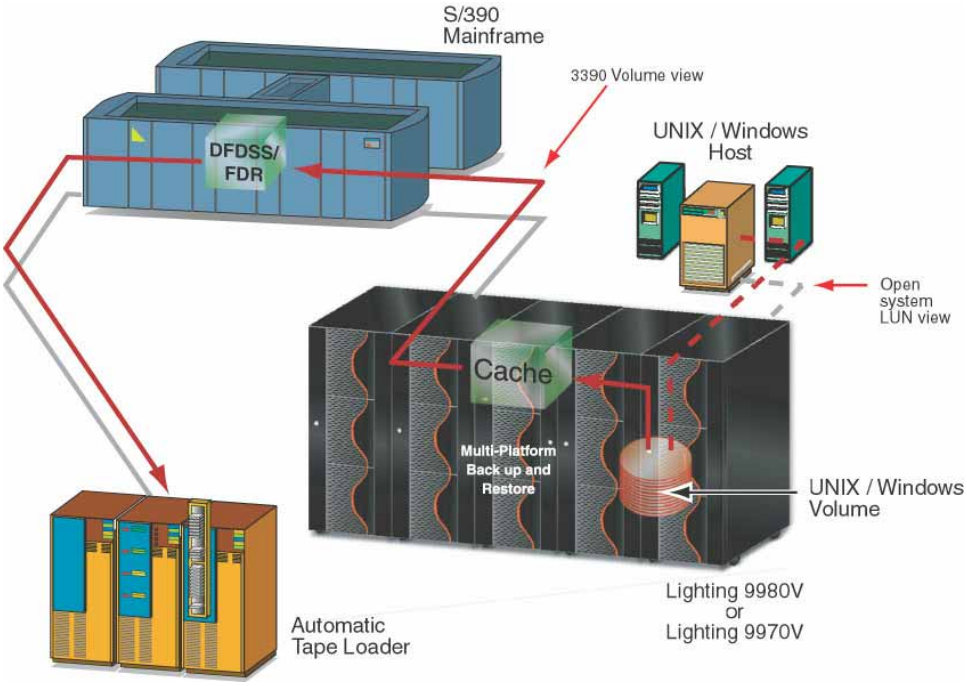

Figure 20 - Hitachi Multiplatform Backup/Restore allows UNIX or Windows NT/2000 volumes to be backed up to tape using FDR (fast dump restore) and mainframe procedures.

File-level backup between platforms requires a communication utility between a backup agent on a client system and a backup agent on the server.

3. The HARBOR File Level Backup and Restore Agent

File-level backup and restore is more difficult to do since it requires the backup software to know and communicate information on the location and extents of a dataset or file. In MVS, this information resides on the disk in catalogs and VTOCs. In UNIX or Windows NT/2000 environments, the information about files and file structures resides in the host file system in mount points and "I-nodes."

File-level backup between platforms requires a communication utility between a backup agent on a client system that can interrogate the file system and a backup agent on the server that records the file control information for retrieval. As with volume-level backup, additional software and management is required if the file is part of a database or application backup. As shown in Figure 21, HARBOR Backup/Recovery provides a solution for file-level backup between S/390 hosts and UNIX or Windows NT/2000 servers. Usually, the file-level transfer of data between UNIX or Windows NT/2000 clients is done over a network. Hitachi and HARBOR have developed a unique agent that provides file-level backup through the Lightning 9900 V Series storage systems and bypasses the network.

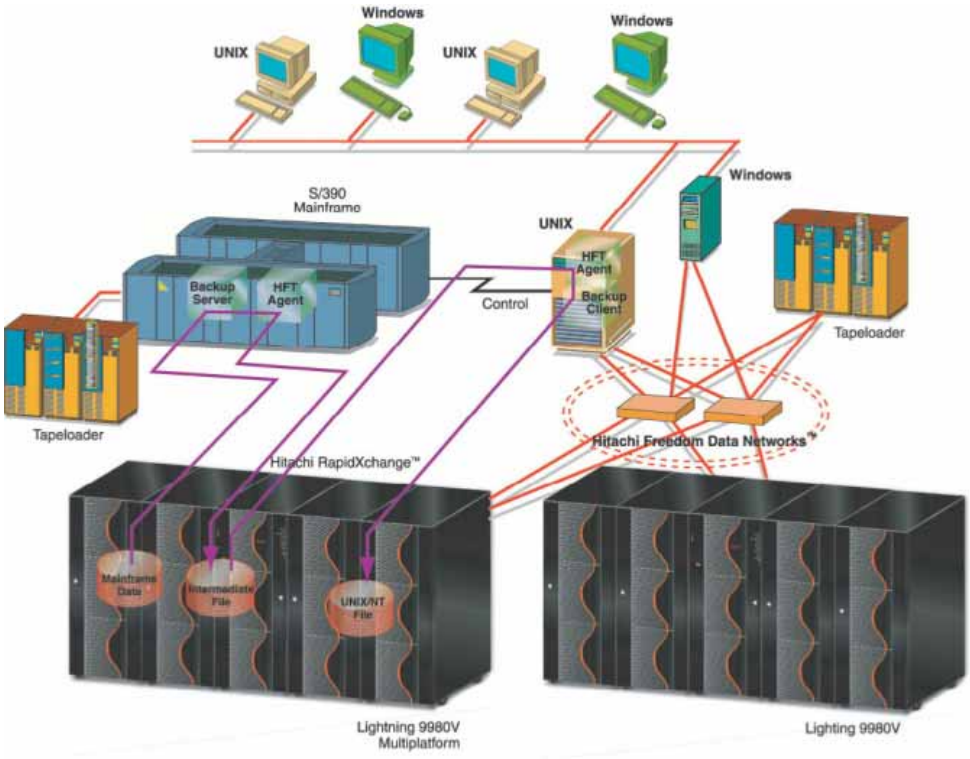

Figure 21 - A major benefit of Backup/Recovery with the HARBOR Lightning 9900 V Series Agent is the ability to restore only single files rather than restoring entire data volumes.

4. Database and Application Level Backup

Databases can be hierarchical (like IMS and VSAM) or relational (like DB2®, Oracle, Informix®, Microsoft® SQL Server™ and Sybase®). DB2 is the predominant RDBMS on mainframes, while Oracle is predominant on UNIX and Windows NT/2000 environments. Databases consist of logs, indices, data, and scripts. Relational databases can be further defined by tablespace, which may contain database tables consisting of columns and rows. Databases that are shared require a locking mechanism to maintain logical consistency.

Standalone or "cold" backup of a database requires that the applications against the database are stopped and all the in-flight data is flushed to the disk before backup is started. The applications must remain stopped until the backup is completed and the database is brought back online. In order to reduce this downtime, customers use database utilities to do online or "hot" backups. Online backups still require the applications to be stopped and data to be flushed to the disk. But once this is done to establish a synch point, the applications can resume operations in a "read only" access mode.

Some software utilities provide transaction logging that allows write updates during the online backup. Database vendors are usually the source for these utilities since they require comprehensive knowledge of the database architecture and application. There are some independent software vendors (ISVs), such as BMC, who provide products like SQL-Backtrack, which can support online backup of multiple vendor database platforms. However, the backup, and, more importantly, the restore of a database requires a database administrator who has a comprehensive understanding of the database and relationships between files. Scripts must be written, tested, and maintained for backing up all the necessary database components.

Hitachi ShadowImage can be used in conjunction with products like BMC Enterprise Snapshot and SQL products or Oracle Hot Backup Utilities to eliminate any impact or disruption caused by online backup.

Online backup of databases requires speed. The longer it takes to do the backup, the more new transactions accumulate in the logs, and the longer it takes to re-synch during the online backup. Hitachi offers several solutions to minimize this time. Hitachi ShadowImage can be used in conjunction with products like BMC Enterprise Snapshot and SQL products or Oracle Hot Backup Utilities to eliminate any impact or disruption caused by online backup and to take a snapshot copy for backup of mainframe data. The HARBOR file-level backup agent (mentioned earlier in the file-level backup section) can reduce the backup time for databases such as Oracle, Sybase, Informix, and Microsoft SQL Server.

5. Application Level Backup/Restore

Application backup and recovery involves the backup and recovery of logically related groups of database objects across multiple databases or files. Backups must be synchronized to generate a consistent recovery point across multiple files. This requires that all active transactions against target objects be placed on hold. The objects are then switched to read-only mode until the backup is completed. Application backup requires a central management server to control and coordinate backup and recovery. A repository must be maintained on information about all database instances, location of backups, and actions performed against instances. Software like BMC PATROL, PRAXIS OMNIBACK, and CA Unicenter TNG® Framework provide this type of backup.

Backups must be synchronized to generate a consistent recovery point across multiple servers.

6 . Hitachi Freedom Data Networks™

A major potential benefit of storage area networks (SANs) is LANfree and serverfree backup. Backup windows are the single major issues on the minds of storage executives as determined by a survey by ITCentrix of Fall River, Massachusetts. The Hitachi Data Systems storage area network solution is known as Freedom Data Networks. Freedom Data Networks provides an open architecture that leverages SAN technology and offers organizations freedom of choice in deploying data-access and data-sharing capabilities across the enterprise. With Freedom Data Networks, customers gain a powerful new tool that enables the consolidation of servers and storage, increased data availability, centralized storage management, and the ability to back up and migrate data without affecting the performance of enterprise LANs.

Backup windows are the single major issues on the minds of storage executives.

The Freedom Data Networks architecture goes beyond SANs.

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 Networks 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).

Industry Alliances are Key to Hitachi Freedom Data Networks Solutions

Key to the delivery of the Hitachi Freedom Data Networks solutions through its direct sales force is a growing number of industry alliances in areas such as networking, communications, and software. Each of these alliances supports Hitachi Data Systems intent to provide fully open solutions based on the use of a wide variety of interoperable components. Alliance partners for solutions developed by Hitachi Data Systems include the following:

- Brocade®, McDATA® and QLogic® for high-performance Fibre Channel switches
- JNI®, Emulex® Corporation, Troika™ Networks for Fibre Channel host bus adapters for Windows NT platforms

Tantia Technologies HARBOR Backup/Recovery (HBR) and HYPERtape

Hitachi Data Systems addresses the needs of both MVS and LAN customers with backup/restore solutions that are scalable to the enterprise level. Intended respectively for MVS and LAN environments, HARBOR Backup/Recovery (HBR) and HYPERtape support a full menu of open systems clients, mainframes, and agents for both online and offline backup of databases[6].

VERITAS® NetBackup™ Offers Backup and Recovery Solutions for Heterogeneous Environments

Scalable from the desktop to the data center, NetBackup delivers quick and reliable backup/recovery that spans terabytes to petabytes in size. The VERITAS NetBackup Array Integration option and VERITAS NetBackup ServerFree Agent are tightly integrated with ShadowImage™ and the Hitachi e-Copy feature, allows administrators to automate split-mirror backups on Hitachi Freedom Storage systems with zero impact on applications availability. The result? A significant reduction in the use of server-based resources traditionally required for backups.

[6] Hitachi Data Systems is a reseller of Harbor and HYPERtape backup/restore software packages in the U.S. and Canada.

Data Movement Solutions

5

Hitachi Data Systems developed its Data Movement Solutions to provide heterogeneous IT environments with the critical ability to transfer copy and share large amounts of data between heterogeneous platforms and applications. This capability is a key to business success in today's global information processing environments, where corporate data must be readily available to all worldwide users, regardless of their location or the nature of their computing environment. Data from heterogeneous systems needs to be exchanged regularly, either for data synchronization between discrete applications or, increasingly, for populating data warehousing systems.

Table 11 lists the Data Movement Solutions Suite that supports the Hitachi Freedom Storage™ Lightning 9900™ V Series, Lightning 9900 Series, and Thunder 9200™ storage systems.

	Product Support			O/S Support	
	Thunder	Lightning			
	9200	9900	9900 V	Open	S/390
Data Transfer Solutions					
Hitachi RapidXchange™		✔	✔		✔
Hitachi Data Migration Services *(Chapter 6)*		✔	✔		✔
Tantia Technologies® HARBOR® File Transfer		✔	✔		✔
e-Copy			✔	✔	

Table 11 - Data Movement Solutions provided by Hitachi Data Systems.

Hitachi RapidXchange Enhances Decision Support Productivity while Improving Operations Efficiency

Hitachi RapidXchange offers high-performance and high-reliability data sharing between heterogeneous host platforms such as OS/390®, UNIX or Windows.

The ability to share and move data between heterogeneous environments at very high speeds while minimizing the impact on the overall throughput of network communication links is a significant asset. It is especially useful for data warehousing and data mining, where vast amounts of data must be moved quickly from mainframe OLTP systems to UNIX® or Windows® environments. In some industries, the time it can take to refresh a data mart or data warehouse can define a competitive edge.

RapidXchange does not need additional software on an MVS® host. RapidXchange provides maximum platform flexibility and transfer speed for the quick and reliable data transfer. Any data that can be unloaded, created, or copied as a sequential dataset in MVS can be read directly by a UNIX server using RapidXchange. Multiple, disparate hosts (for example, HP®, Sun™, and Windows) can access a single copy of data. RapidXchange is especially useful for data warehousing and data mining tasks, where vast amounts of data need to be moved from mainframe OLTP systems to a UNIX environment (such as data warehouses or data marts) in a very short time. RapidXchange delivers this capability without requiring networking resources or involving intermediate media, such as tape. High-speed data transfer is achieved without putting additional data load on the network infrastructure or tape transport equipment. Also, RapidXchange has ultra-high-speed transfer rates (when used in conjunction with Hitachi FlashAccess™), which deliver the fast response demanded by today's widespread data warehousing and data mining applications.

RapidXchange provides for much higher Fibre Channel effective data transfer rates than today's networks can accommodate. Lightning 9900 V Series system users can take advantage of even faster data transfer rates by using RapidXchange in conjunction with FlashAccess, which maintains the intermediate volume in the Hitachi Freedom Storage Lightning 9900 V Series cache. Using cache as a buffer eliminates any physical disk access

during the write/read process and enables data exchange at channel speeds. Data can be transferred from a UNIX or Windows system to an MVS system by using the same procedure in the reverse direction. This process is illustrated in Figure 22.

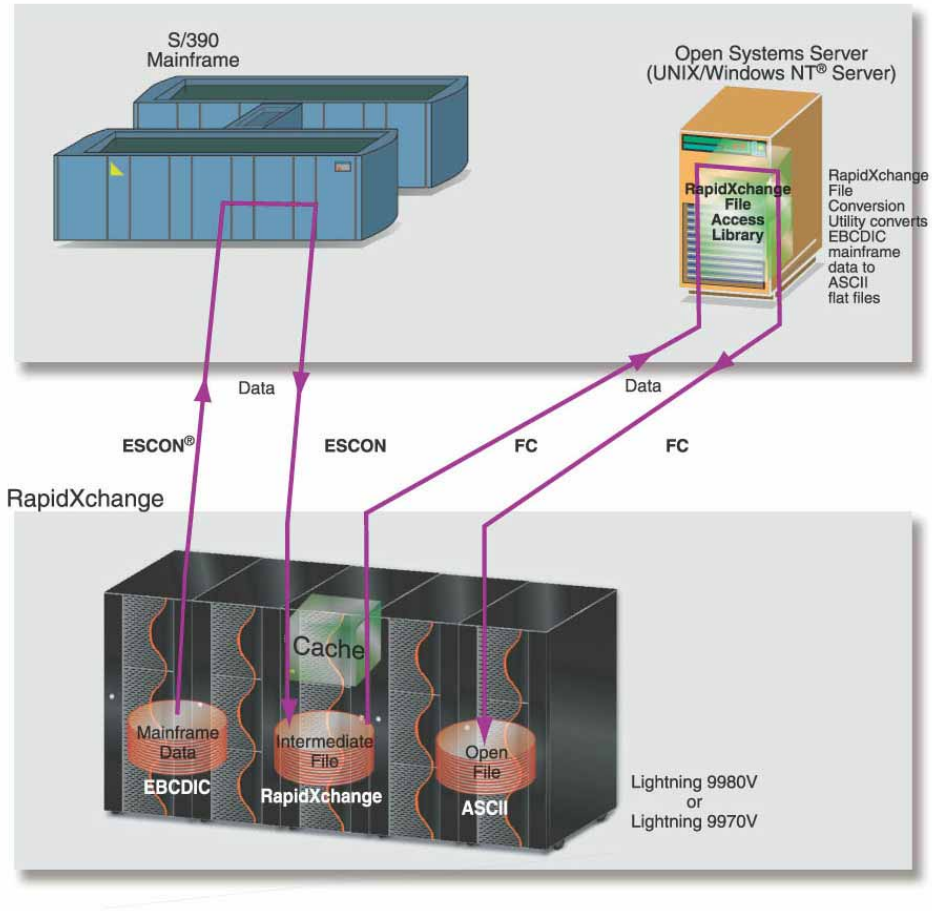

Figure 22 - Hitachi RapidXchange quickly converts mainframe EBCDIC files to open systems ACSII files for non-disruptive population of data marts and data warehouses.

RapidXchange transfers data at the speed of Fibre Channel.

Hitachi RapidXchange eliminates many of the constraints associated with sharing or moving data between heterogeneous systems. Sequential files on common storage volumes allow large amounts of data to be shared or moved. Using RapidXchange, you can "share" sequential datasets between MVS and open systems platforms.

Hitachi RapidXchange provides two options for accessing data from open systems platforms:

The UNIX flat file to receive the transformed data can be resident on any UNIX volume, including non-Hitachi storage.

File Access Library API

The Hitachi RapidXchange File Access Library provides basic API functionality in the form of open, read/write, and close commands. UNIX or Windows applications can use this API to read MVS data directly from a RapidXchange disk volume and then process that data. In many cases, existing programs that currently use tape input can be modified easily to read that data directly from the MVS system, thereby eliminating the MVS-to-tape-to-UNIX process.

File Conversion Utility

A second option is the Hitachi GUI-based File Conversion Utility (FCU). Supplied as a standard component of Hitachi RapidXchange, the FCU furnishes a GUI-based display of the intermediate MVS volume. From this display, individual

sequential datasets can be selected, converted to ASCII, and transferred as a simple UNIX or Windows flat file. Transformation options include converting data from EBCDIC to ASCII or padding out variable-length records. The UNIX or Windows flat file to receive the transformed data can be resident on any UNIX volume, including non-Hitachi storage.

Tantia Technologies HARBOR File Transfer Increases Speed and Automates Performance of Large File Transfers

HFT adds automation to the process of transferring large data files at high channel speeds.

Running with Hitachi RapidXchange, HARBOR File Transfer adds automation to the process of transferring large data files at high channel speeds in either direction between open systems and mainframe servers. It can be thought of as an ultra-fast FTP transfer except data is transferred over an ESCON link instead of Ethernet. HARBOR File Transfer enables UNIX and Windows applications such as enterprise resource planning and decision support systems to work with up-to-the-minute data extracted from S/390 environments. Multithreading capabilities and Fibre Channel speeds of 100MB/sec support HARBOR File Transfer's high-speed data movement. After automatically breaking large data files into more manageable pieces, HARBOR File Transfer offers increased transfer speed by directing data in multiple streams through the Lightning 9900 V Series system(s).

A HARBOR File Transfer agent running on the open systems client platform sets up control information (including authorization and file lists) and sends it to the MVS host over the network as shown in Figure 23. When MVS returns acknowledgment, the client sends the corresponding file data to the host through Hitachi RapidXchange. This reduces the traffic on the network, and minimizes the backup time by transferring data at ESCON, FICON, and Fibre Channel speeds rather than at network speeds.

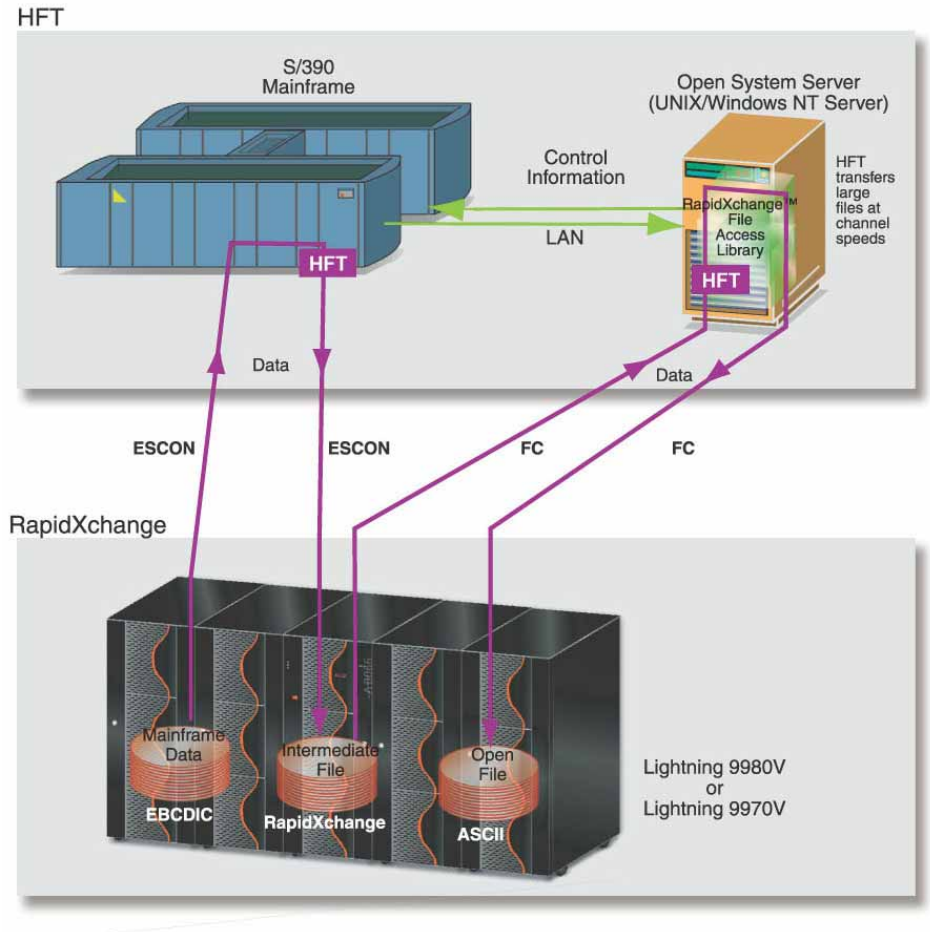

Figure 23 - HARBOR File Transfer runs on Hitachi RapidXchange and transfers large data files at channel speeds in either direction between open systems and mainframe hosts.

HARBOR File Transfer is more efficient than transfer over the much slower network lines used in conventional solutions.

This is more efficient than transfer over the much slower network lines used in conventional solutions. As shown in Figure 23, HARBOR File Transfer employs the network to transfer only "control information" between open systems and mainframes, so they remain more fully available for regular network activity. An important timesaving feature of HARBOR File Transfer is its capability in the event of a transmission failure to restart from the point of failure. This means that transfers can be continued from the point in time that a failure occurred, rather than from the very beginning of the transfer process. This saves a tremendous amount of time compared to FTP network alternatives. With the added features of alternate path routing (through the network or through the Lightning 9900 V Series system) and retry/reconnection logic, HARBOR File Transfer brings heightened reliability and fault tolerance to the file transfer process.

For "state of the art" management, HARBOR File Transfer offers an easy-to-use graphical user interface (GUI) that monitors transfer rates, active transfers, logging, and system controls, and simplifies file selection and batch file creation. HARBOR File Transfer's GUI allows both system administrators and end users to share mainframe and client system data and perform backups and restores. Efficient file transfer has become a major challenge for many organizations. Tape technology often requires manual intervention (and the risk of human error) to move data between separate systems. Today's network congestion often does not provide the performance necessary for the transfer of large volumes of data in a short time. This is due to the conflict for transmission with other network traffic and network protocol overhead.

Hitachi e-Copy Provides LANfree and Serverfree Data Movement

A common function for Hitachi Data Movement Solutions is for LANfree and Serverfree disk-to-disk or disk-to-tape data transfer for backup purposes. Hitachi e-Copy, when used with applications on the open-systems host server, such as the VERITAS® NetBackup™ series of products, and with Hitachi ShadowImage™ provides serverfree backup solutions between the Lightning 9900 V Series system and backup devices such as tape and disk devices. E-Copy enables non-disruptive backup directly from disk to tape (or disk) in SAN environments, eliminating server CPU and I/O overhead during movement of data and decreasing the time required for backup. Users can perform copy (backup) operations on the data (in units of block) stored on the multiplatform Lightning 9900 V Series system via Fibre Channel interfaces directly to the backup devices.

E-Copy supports the SCSI Extended Copy command issued from the host server to the Lightning 9900 V Series system. The Lightning 9900 V Series system receives the e-Copy commands issued by the server, and then copies the data directly to the specified backup device. E-Copy operations are configured using the LUN Manager Remote Console – Storage Navigator software. To implement e-Copy operations, you need a backup application on the host server to issue the e-Copy commands to the Lightning 9900 V Series system. The LANfree and serverfree data movement capability of the Hitachi e-Copy is shown in Figure 24.

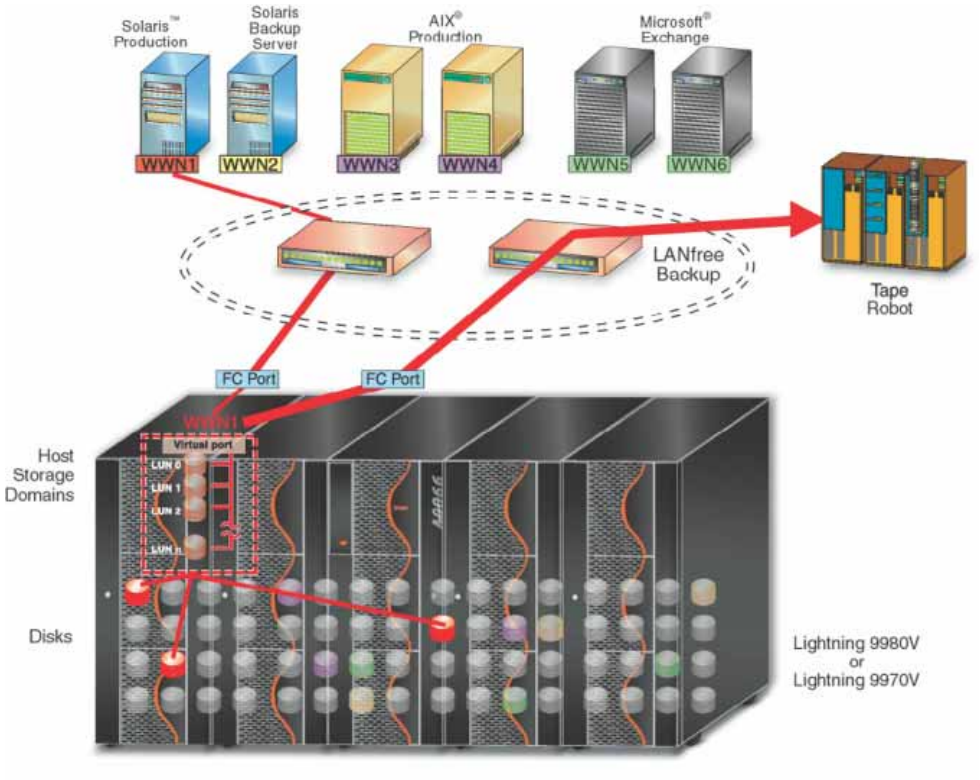

Figure 24 - Working with the VERITAS NetBackup series of products, Hitachi e-Copy provides LANfree and serverfree data movement.

Professional Services and Support

6

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 provides 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 specializes 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 Data Systems 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.

Data Protection Services

The Hitachi Data Systems 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, Data Protection Services implements essential copy solutions by moving data in order to protect it. Data Protection Services 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 is reached, 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 who may be contacted prior to beginning of the engagement. Examples of the Hitachi Data Systems 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 Freedom Storage systems and software. The Data Protection Services Lab is dedicated to resolving software and hardware issues that affect storage system consolidation, data center productivity, and business continuity.

Remote Copy Assessment and Implementation Service

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

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

NanoCopy™ Implementation Service

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

SplitSecond Custom Services

SplitSecond Custom Services extends NanoCopy functionality into the application by providing a communication layer that seamlessly integrates customer applications with ShadowImage™. This "industry first" solution tackles the complex challenges of ensuring that essential information can survive a rolling disaster. By leveraging Hitachi TrueCopy and Shadowomage capabilities, NanoCopy and SplitSecond Custom Services become an even more cost-effective and simplified technique that not only protects any amount of data with full integrity, but also enables rapid recovery following a disaster.

Continuous Availability Service

Hitachi Data Systems 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. With this service, Hitachi Data Systems consultants employ best consulting practices in both open systems and MVS® application backup and recovery (both on-site and off-site) to conduct the assessment. As a first deliverable, a baseline will be created and presented for analysis of potential strengths and weaknesses of the current processes, equipment, software, and practices. Recommendations are then discussed to improve availability and reduce the risk of data loss.

Hitachi Data Systems Data Migration Service

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

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

1) Hitachi Data Systems Professional Services methodology/procedure and skills
2) The outstanding reliability of Hitachi Freedom Storage products

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

In S/390 environments, terabytes of data can be migrated to Hitachi Freedom Storage systems from other vendor systems in a matter of hours while applications are online and processing remains completely uninterrupted. The Hitachi Data Systems 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 or 200MB-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 SAN solutions to optimize management and control data across the enterprise. The modular services include:

- The Enterprise Infrastructure Assessment (EIA) Service
- The SAN Configuration Management Service
- The SAN System Design Service
- The SAN Project Management Service
- The SAN Installation Service
- The SAN Business Continuity Service

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

Hitachi Data Systems Interoperability Laboratory Service

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

The Interoperability Laboratory Service provides clients with the knowledge of pre-tested components of a SAN in various network topologies. Customers however, will not be limited to the use of these proven interoperable SAN elements. To ensure 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 (iLab™) at its Santa Clara, California, headquarters to test new SAN elements.

Among the iLab's earliest projects was the replication and operation of the SAN developed for a major banking customer and a major telecommunications customer. The experience and methodology gained by the iLab 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 Thunder 9200 and Lightning 9900 V Series 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. This proven methodology is based on the key disciplines of software asset management practices, supports the efforts of Hitachi Data Systems "best of breed" partners, and focuses on maximizing the return on investment to acquire the greatest possible value.

Hitachi Data Systems Continuous Business Planning Questionnaire

A

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 Software®, SMS®, etc.)?

2. Support Software

2.1. Transaction Workload Balancing

2.1.1. Currently, what method of transaction routing are you using for your online 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 preventative maintenance?

2.3.2. Do you install preventative 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 online 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

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 can be sent by a variety of methods to operations personnel.

API

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

Asynchronous

Asynchronous communication (as in asynchronous remote copy) occurs when the transmission of data between two devices is 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. Asynchronous and synchronous transmissions are used extensively in the mainframe terminal environment.

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 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 as discussed in Chapters 1 and 2.

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 (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 processes 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.

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.

Common Information Model (CIM)

The Common Information Model (CIM) is an object oriented emerging standard to the central management of systems and networks across multiple organizations with heterogeneous systems and software. CIM is sponsored by the Distributed Management Task Force which receives contributions from participating technology companies. It is also supported by the SNIA.

CVS

Customized Volume Size for open systems. A feature in the Lightning 9900™ V storage systems which enables a user to define LUNs that are non-standard size.

DASD

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

Data Availability

Data availability refers to the degree to which a computer system is capable of providing data to users. See also Chapter 3 and Availability.

Data Copy

A term that refers to remote copy, data duplication and data migration. See also Chapter 3.

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

Any event that creates an inability for an organization to provide critical business functions for an undetermined period of time.

Disaster Recovery

Disaster Recovery is the profession that plans to avoid disasters and to restore services after a disaster. It also indicates the level of preparedness to respond to an interruption in services by implementing a disaster recovery plan to restore an organization's critical business functions is discussed in Chapter 3.

DR

Disaster Recovery.

Downtime

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

ECC

Error correction code.

e-Copy

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

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 taken periodically.

ESA/390®

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

ESCON®

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

ESCON Director

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

Ethernet

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

Fabric

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

Failover

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

Fast Ethernet

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

Fast write

A write operation at cache speed that does not require immediate transfer of data to a disk drive module. 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

An acronym for Fibre Channel.

FC-AL

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

FCP

An acronym for Fibre Channel Protocol – an ANSI standard covering Fibre Channel protocol for SCSI. See also Fibre Channel Protocol.

Fibre Channel

Fibre Channel is an ANSI standard designed to provide high-speed data transfers among workstations, servers, desktop computers, and peripherals. Fibre Channel makes use of a circuit/packet switched topology capable of providing multiple simultaneous point-to-point connections between devices. Fibre Channel is widely deployed in SAN implementations today. Standards for Fibre Channel SANs are being worked on by the Storage 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.

Fibre Channel Protocol (FCP)

Fibre Channel Protocol is an ANSI standard covering Fibre Channel protocol for SCSI.

FICON™

The Fibre Connector channel (FICON) is IBM's trademarked channel for zSeries or S/390 processes in SAN environments.

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 case the active file gets damaged. Backup is considered "local copy" as opposed to "remote copy." See also Remote Copy.

Gigabit Ethernet

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

GUI

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

HARBOR® File Level Backup/Recovery (HBR)

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

HARBOR File Transfer (HFT)

This software utility provides automatic, reliable, and secure data transfer between OS/390 and open systems hosts using high speed ESCON/FICON channels or network communications. HFT can be thought of a very high speed FTP that does not use a network. High performance is enabled by allowing multiple file transfers to run concurrently. Reliability is enhanced with retry logic, alternate path routing, and restart from point of failure. Ease of use is provided by a Java™-client interface, a system monitor, scheduling support, and remote control of client functions from OS/390.

HFT

See HARBOR File Transfer.

Hi-Star™ architecture

At the heart of the Lightning 9900™ V Series revolutionary design is the Hi-Star architecture, which provides multiple, redundant, non-blocking paths between the storage ports, multiple cache nodes, and multiple disk Array Control Processors (ACPs). Each path runs at 100MB/sec, permitting Lightning 9900 V Series systems to scale up to a total bandwidth of 15.9GB/sec. This is six-to-eight times the internal bandwidth of shared-bus based storage systems. For redundancy and performance, Hi-Star consists of four cache switches that are cross-connected to four cache modules and two control memory banks. These control memory banks contain addressing and control information and are also cross-connected to front-end storage ports and back-end disk array ports.

Hitachi CruiseControl™

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

Hitachi Data Systems Migration Service

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 V Series system 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®/Windows® 2000.

Hitachi FlashAccess™

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

Hitachi Freedom Data Networks™

Hitachi Freedom Data Networks provides an open architecture that leverages SAN, NAS, and IP technology and offers organizations the freedom of choice in deploying data-access and data-sharing capabilities across the enterprise. With Freedom Data Networks, customers gain a powerful new tool that enables the consolidation of servers and storage, increased data availability, centralized storage management, and the ability to back up and migrate data without affecting the performance of enterprise networks.

Hitachi Graph-Track™

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

Hitachi LUN Manager

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

Hitachi Multiplatform Backup/Restore

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

Hitachi Multiplatform Resource Sharing

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

Hitachi Priority Access

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

Hitachi RapidXchange

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

Hitachi Resource Manager™

The Hitachi Resource Manager is a comprehensive suite of management software that brings together Hitachi Graph-Track, Hitachi Virtual Logical Volume Image (VLVI) Manager, Hitachi FlashAccess Manager, and Hitachi LUN Manager into one complete package. Resource Manager 7000 manages the Hitachi Freedom Storage™ 7700 and 7700E systems. Resource Manager 9000 manages the Hitachi Freedom Storage™ Lightning 9900™ V Series systems.

Hitachi SANtinel™

Hitachi SANtinel software controls host access to Hitachi Freedom Storage 7000E or Lightning 9900 V Series LUNs in SAN environments.

Hitachi ShadowImage™

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

Hitachi ShadowImage open

Hitachi ShadowImage open is a point-in time (PIT) copy capability that allows the data replication and fast resynchronization of logical volumes for open system computers. Nine copies of a logical volume may be maintained within the same system. ShadowImage also works in concert with Hitachi TrueCopy asynchronous extension and Hitachi TrueCopy asynchronous for S/390 to provide additional copies in another system. See also: Hitachi ShadowImage.

Hitachi TrueCopy

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

Hitachi TrueCopy Asynchronous

Hitachi TrueCopy asynchronous provides asynchronous remote copy capability for data from one Lightning 9900 V Series system to another Lightning 9900 V Series system or over channel extenders available from McDATA®, CNT® and INRANGE®. Records are timestamped to ensure that data is not lost or out of order. Since the transmission is asynchronous, data can be transmitted over unlimited distances. See also Chapter 3 for a discussion of the family of Hitachi TrueCopy software utilities.

Host Storage Domains

Host Storage Domains are collections of LUNs assigned to a specific host type in the Lightning 9900™ V Series. Each Domain has its own logical (or virtual) Fibre Channel port and hosts, which are matched to their assigned Domain based on their WWNs. See also LUN, Fibre Channel, and WWNs.

Hot site

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

HXRC (Hitachi Extended Remote Copy)

This IBM XRC™-compatible host-based software is offered by Hitachi Data Systems for asynchronous remote copy in an S/390 environment. Using System Data Mover software, it guarantees data integrity for dependent write applications.

HYPERtape

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

I/O

Input/output.

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).

LAN

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

Lightning 9900™ V Series

The Hitachi Freedom Storage™ Lightning 9900™ V Series was announced in May 2002 (Lightning 9980V and Lightning 9970V). These products represent a major advance in enterprise-class storage systems. This is because the second generation Hierarchical Star Network switched internal architecture provides for many times more simultaneous transfers to and from the host compared to shared-bus architectures. (The Lightning 9960/9910 systems of the Hitachi Freedom Storage Lightning 9900 Series were announced in June and November 2000 respectively.)

Logical Fibre Channel Port

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

Logical Unit

The SCSI term for a logical disk drive.

Logical Unit Number

See LUN.

Logical Volume

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

LUN

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

LPAR

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

LUSE

The LUN Size Expansion feature of the Hitachi Freedom Storage allows standard size LUNs to be combined to create larger LUNs.

MAN

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

MIB

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

MIPS

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

Mirrored pair

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

Mirroring

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

MPLF

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

MTBF

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

MTTR

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

NanoCopy™

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

NAS

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

N+1

An N+1 power supply design provides for one redundant power supply in a power system design that provides full system power in the event of a power supply failure.

NDMP

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

Node

See Fibre Channel.

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.

Online systems

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

Open system

A system whose characteristics comply with de facto 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.

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.

Port/Port ID

See Fibre Channel.

PPRC

Peer-to-Peer Remote Copy or PPRC is an IBM feature that maintains consistent copy of a logical volume on the same controller or on another controller. Access to the remote controller occurs over an ESCON path. All writes that any attached host performs on the primary logical volume, are also performed on the secondary logical volume. The user can request that the controller perform the modifications synchronously, that is, before the controller completes the modification of the primary copy.

QoS

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

RAID

Redundant Array of Independent Disks. RAID is used to increase the reliability of disk arrays by providing redundancy either through complete duplication of the data (RAID-1, i.e., mirroring) or through construction of parity data for each data stripe in the array (RAID-3, -4, -5). RAID-5, which distributes parity information across all disks in an array, is among the most popular means of providing parity RAID since it avoids the bottlenecks of a single parity disk. The Lightning 9900 V Series system algorithms enable performance from RAID-5 that is competitive with some vendors' 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 onset of a disaster to the recovery of the critical functions.

Remote copy

Remote copy refers generically to software or hardware utilities that provide the capability to copy data from one online volume to remote volumes without disruption. Synchronous techniques are used for short distances (typically less than 25 miles) and asynchronous techniques over LAN/WAN/MAN are used at any distance. For a complete discussion of Hitachi TrueCopy (remote copy) solutions, see Chapter 3.

Remote copy links

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

Risk management

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

SAN

Storage area networks (SANs) connect storage systems to servers through Fibre Channel or Ethernet switches. The Hitachi implementation of the SAN is known as Hitachi Freedom Data Networks™. Major benefits of SANs include outboard backup, sharing of resources, pooling, and reduced cost of storage management. SANs are also defined as 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 Hitachi Freedom Data Networks and Fibre Channel.

SCSI

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

Serverless Backup

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

SIM

Service Information Messages are messages generated by Lightning 9900 Series, Lightning 9900 V Series, and Thunder 9200 systems to identify normal operations.

SLA

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

Snapshot

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

SNMP

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

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.

SWAN

Storage wide area networks (SWANs) are interconnected SANs over long distances. They are made possible by Fibre Channel and ESCON extenders.

Switch

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

Synchronous

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

TCO

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

TCP

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

TPF

Transaction Processing Facility or TPF is highest performance transaction processing software environment from IBM. TPF is used by many of the world's largest Customer Reservation Systems such as SABRE® and by the world's largest financial institutions. TPF uses a high performance record-locking scheme called MPLF. See also MPLF.

VERITAS®

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

Virtual Logical Volume Image Manager

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

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 fiber optic backbones.

Warm site

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

WWNs

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

XRC

Extended Remote Copy is an IBM implementation of a software asynchronous remote copy technique that preserves data integrity. See also Chapter 3 for a complete discussion of remote copy software and Hitachi TrueCopy product family.

Hitachi Data Systems Worldwide Offices

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

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

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

Australia/New Zealand Headquarters
Hitachi Data Systems
Level 3
82 Waterloo Road
North Ryde NSW 2113
Australia
Phone: +61-2-9325-3300
info.australia@hds.com.

Canada Headquarters
2550 Victoria Square Avenue
Suite 601
Toronto, Ontario
M2J 5A9
(416) 494-4114
info@hds.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-436-00 May 2002