IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Business Continuity: More than Just Waiting for Disaster

Sponsored by: Hitachi Data Systems

Richard L. Villars Bill North

September 2003

When most people think of business continuity, they focus on unplanned downtime and its costs. Less frequently mentioned is planned downtime because many businesses operate nearly continuously. More recently, business volatility emerged as a concern. When business lags, IT investments are underutilized and the total cost of ownership (TCO) rises proportionally. When business booms, IT becomes a limiting factor, opportunity costs are incurred due to lost business, and TCO rises as IT scrambles to catch up.

These formidable challenges and the growing complexity of IT systems that support business operations make business continuity a high priority for IT professionals.

In this IDC white paper, we examine the technical and economic challenges that make effective business continuity planning (BCP) for storage systems so difficult. We also discuss strategies and technologies that allow companies to quickly and effectively improve their data protection and rapid recovery capabilities, including two case studies of companies that are using these solutions. Finally, we examine the business continuity offering of Hitachi Data Systems (HDS) and highlight areas where Hitachi must add capabilities to address emerging customer requirements.

KEEPING BUSINESSES IN BUSINESS

Companies need better, faster, and more cost-effective ways to collect, store, analyze, and access information as they struggle to remain competitive during difficult economic times. Furthermore, to meet customer, employee, and regulatory requirements, they must store a growing variety of information types (e.g., medical records, purchasing histories, customer correspondence, images, and videos) for longer periods of time.

When information systems that store and process this information are unavailable, the enterprise is "out of business." Productivity declines as staff go idle, revenue stops flowing, profits erode, customers go unserved, and the company's reputation is damaged. Today, with so much at stake, enterprises must be especially mindful of their plans for ensuring business continuity. Examples of the challenges that enterprises face are described below:

- Dependence on email, customer relationship management (CRM), and other applications that force an expansion in the scope of data protection efforts
- Regulatory mandates in specific industries (e.g., finance and healthcare) that call for the retention of all internal and external correspondence while guaranteeing data integrity and easy (but secure) accessibility for extended periods
- Scrutiny of overall business continuity plans by senior executives, investors, regulators, and customers following the tragic events of September 11, 2001

These formidable challenges and the expanding scope and complexity of the IT systems that support business operations make business continuity planning a daunting but high priority for IT professionals.

BUSINESS CONTINUITY: MORE THAN DISASTER RECOVERY

When most people think of business continuity, they focus on unplanned downtime. A multitude of factors can interrupt IT services, including those listed below:

- Storage, server, or application failures
- Loss of electrical power or network access due to factors outside the control of IT managers (e.g., cable cuts, regional brownouts/blackouts, or weather-related outages)
- Software and operator error that can often wreak havoc unexpectedly by corrupting data or preventing the reliable storage and movement of data

Planned downtime is less frequently mentioned but of growing concern for IT managers because many businesses operate continuously or nearly continuously. While evenings and weekends have traditionally been a time when systems could be taken down without consequence, today's CRM and supply chain management applications extend the concept of "normal business hours" to 24 x 7.

Planned downtime is needed for routine maintenance activities. For example, a server needs additional attached-storage capacity or upgraded processors. Operating systems require upgrades, and application software must be enhanced and modified to keep it in alignment with business processes. Finally, tape backup and restore of data is most easily accomplished when production systems are quiescent — though the increased bandwidth and connectivity of storage area networks (SANs) have helped minimize the impact of tape backup on application availability. Business volatility is a more recent and increasingly important business continuity concern. Enterprises continually struggle to gauge marketplace behavior and predict the demands employees, customers, and business partners will place on IT systems. Miscalculations can be very costly in terms of wasted effort and wasted resources. Overly optimistic business forecasts encourage IT planners to increase capacity. When business lags, IT investments are underutilized and the TCO rises proportionally. Conversely, pessimistic business forecasts lead to smaller investments in IT systems. When business booms, IT becomes a limiting factor, opportunity costs are incurred due to lost business, and TCO rises as the need to satisfy demand outstrips the normal focus on efficiency.

In summary, business continuity is no longer just about planning for the next system failure, network outage, or local catastrophe. Enterprises must develop business continuity plans that allow them to operate more efficiently and react quickly to changing business conditions.

HURDLES TO BUSINESS CONTINUITY PLANNING

Business continuity planning spans all elements of the enterprise, including IT, and, in today's tightly linked economy, frequently extends beyond the enterprise to include business and IT operations of customers or suppliers or both. With respect to IT, BCP requires the development and application of procedures and technologies to ensure that critical information systems remain available or can be brought back into service quickly when unplanned and planned outages, such as power blackouts or system upgrades, occur.

Any BCP effort will be dominated by the most extreme challenge: the possible loss of a datacenter. Disaster recovery plans call for alternative facilities as well as staff at often-distant locations and depend upon the availability of recent copies of relevant data. How recent these copies of data need to be depends on a number of factors. The most significant are listed below:

- Cost (in terms of revenue, reputation, or legal liability) of downtime for individual applications (recovery time)
- Acceptable size of the gap between the last copy of data and the onset of the outage (data loss)
- Ability of transaction processing applications to tolerate the performance impact associated with frequent copying of critical data

BALANCING RECOVERY, INTEGRITY, AND COST

Best practices in BCP dictate that the degree of resiliency, including recovery time and levels of redundancy, should be provided in proportion to the cost or consequences of downtime; the value of continued, timely access to information; and the value of the data itself.

A real-time settlement system in a financial services business has no tolerance for unavailability or data loss because the value of the data is extremely high for an extended period. Even minor problems can have an instantaneous and severe impact on the central mission of the enterprise. Duplicate or triplicate investments to ensure instantaneous recovery for critical systems are, therefore, a basic cost of doing business for financial services companies and are often prescribed by regulation.

Downtime for other applications may not have the same severe consequences, but it may still hamper operations in the same enterprise. For example, monthly payroll and time-reporting applications are less critical to the daily operations of most enterprises. The higher expense associated with closely coupled, highly resilient datacenters would be wasted in such a case. For systems that can be out of service for hours or days without a negative impact on the enterprise, routine backup and offsite storage of data and applications have been the preferred solutions.

Today, however, ensuring data integrity is an increasingly critical concern that is changing the dynamics of BCP. Although instantaneous recovery is not always necessary, ensuring backup data integrity is a growing concern. For example, financial institutions are now required to retain all internal and external electronic correspondence (e.g., email and instant messaging) for a set period (typically two years). They must also be able to prove that the correspondence has been unaltered. BCP for these systems must take a different approach. Recovery times are less important (hours to recovery are acceptable), but data loss or corruption is unacceptable, subjecting the enterprise to financial and even criminal liability. In this environment, reducing backup increments and staging data to a remote archive facility are now the most important requirements.

The range of information that companies need to protect is expanding rapidly, and the pressure to ensure adequate protection for all corporate data continues to rise. Concurrently, ongoing economic and business uncertainties place a growing strain on budgets for IT systems and on the people who install and maintain them. These often-conflicting influences are making business continuity planning extremely difficult. They also preclude the possibility of finding a "one size fits all" data protection and recovery solution.

HISTORICAL DATA RECOVERY AND PROTECTION OPTIONS

In the past, concerns about business continuity centered on data replication for critical information systems including online transaction processing (OLTP) and other large-scale business applications deployed by large enterprises and hosted in one or more datacenters. IT managers backed up databases, files, or data sets after business hours and moved copies to remote storage archives via shipment of removable media (e.g., magnetic tape or optical disks). For the most critical information in transaction-intensive applications, IT managers also had the option of replicating (mirroring) data to disaster recovery sites via very expensive high-speed network links.

Tape backup is the oldest and most widely deployed data recovery and data protection solution. As a removable medium, tape is easy to transport, making it easy to protect data by moving it out of harm's way. The limitations of tape backup include the following:

- Extended periods where important data (e.g., changes made in the past 24 hours) is not protected until the next backup cycle
- Slow recovery times (tapes must be physically transported to and from remote sites) that can range from half a day to more than two days
- Risk of damage to tape media that prevents recovering data
- Inconsistent ability to recover due to uncertainties in the quality of specific backups (the severity of this uncertainty varies widely from high-end mainframe systems that are very reliable to more dispersed solutions for open systems, where failure rates can approach 50%)

In today's increasingly complex IT environment, where multiple applications and data sets are interdependent, data synchronization creates a new challenge for tape-based disaster recovery strategies. IT managers must be able to create simultaneous point-in-time copies for multiple data sets, thereby creating I/O-consistent copies.

The advent of disk-to-disk data mirroring to a remote disaster recovery site overcame many of the shortcomings of tape through improved protection of recently created data as well as faster and more reliable application recovery. The deployment of such solutions, however, was often very expensive due to the replication of facilities and extremely high telecommunications costs (often 50–70% of ongoing costs). As a consequence, companies protected only the most mission-critical applications with this type of solution, leaving the vast majority of data and applications dependent on tape.

IMPROVING BUSINESS CONTINUITY OPTIONS

These limited data replication options don't mesh well with IT managers' current business continuity requirements (i.e., accommodating a wide variety of application and data types with different application recovery and data protection thresholds). Today's BCP must focus on three challenges:

- Decrease the likelihood that IT systems will fail
- Accelerate recovery when IT systems do fail
- Accommodate the on-demand provisioning of processing and storage resources to compensate for greater volatility in IT requirements

ENHANCING SYSTEM RESILIENCY

Many improvements in system resiliency are the result of enhancements to existing products, including the following:

- Advances in the use of redundant, hot-swappable components
- Development of predictive system-monitoring capabilities to eliminate potential failures before they occur
- Support for nondisruptive microcode updates to systems
- Software tools that automate processes to avoid human/operator error (one of the leading causes of unplanned downtime)

As enterprises review their business continuity plans, IT managers will be pleased to see that improvements in overall IT architectures and nondisruptive data movement are helping to make business applications more resilient at a lower cost.

SERVER AND STORAGE CONSOLIDATION

Current efforts to consolidate servers and storage that were often widely dispersed and managed in uncoordinated ways set the stage for simpler and better data protection schemes. IDC research shows that server and data consolidation is a high priority for many enterprises.

Rather than backing up server-attached storage from a multitude of departmental email servers, an approach that requires widely dispersed tape backup systems and administrative support, IT storage managers can leverage a SAN. Inside the dedicated storage network, large amounts of data, including departmental email databases, can be:

- Replicated to protect from device failure
- Centrally provisioned with respect to capacity as well as speed
- Migrated to alternate media and offsite locations for safekeeping

The consolidated storage and SAN approach makes applications more resilient at a lower cost, but it does drive the need for more heterogeneous and scalable data protection solutions.

NONDISRUPTIVE TECHNOLOGIES

Using nondisruptive techniques can also reduce scheduled downtime. Storage systems, for example, have point-in-time copy capabilities. While business applications continue to read and update data, a replica of that data can be split off and used for other purposes. Unrelated to SANs, point-in-time copy can completely eliminate concern for backup windows as static copies of data can be used as source files for tape backups.

These point-in-time copies can also play a critical role in application development. Developers can nondisruptively use actual system data (a point-in-time copy) to test new versions or patches of operating systems and business applications for stability and to assess performance characteristics before production deployment.

Case Study: Marshfield Clinic

Improving data availability and application development performance was a critical issue for Carl Christensen, CIO at Marshfield Clinic in Marshfield, Wisconsin. From a practice of six doctors in 1916, Marshfield has grown to encompass 700 doctors at 41 centers in Wisconsin. The Clinic also operates its own HMO and a statewide immunization registry, manages a decade-old regional electronic medical records service, performs laboratory testing work for healthcare institutions throughout the Midwest, and operates both research and education foundations.

Marshfield moved exclusively (with the exception of certain specialized medical systems) to Windows and clustered servers in 1999 and develops many of its own applications on an internally created database. To maintain all the systems for Marshfield's different groups, Christensen has an IT staff of about 200, split evenly between administration and development. They are located at the primary datacenter in Marshfield.

Marshfield has slowly expanded the use of electronic documents since 1985 and is now ready to deploy tablet PCs throughout the Clinic system in the next six months. "We wanted the mobility, security, and privacy advantages of going electronic with the tablets," Christensen said, but now availability becomes a critical concern.

"The paper medical chart has traditionally served as a great downtime system, and we're going to lose that," he said. "We had to improve our data recovery systems and are looking for 99.99% availability for our customers, so our storage system had to be 99.999% available."

Achieving these types of availability targets required significant changes in the company's storage infrastructure:

- Marshfield decided to consolidate much of its storage on a SAN connected to the clustered Microsoft 2000 servers. This would allow Christensen and his team to deploy a consistent backup and disaster recovery program for all of their diverse databases and applications.
- The company needed to take its disaster recovery system to the next level. Previously, it had been copying data to its backup datacenter over its corporate WAN. The team decided that they needed to mirror all record updates simultaneously at both the primary and secondary sites.

Marshfield had significant time constraints and selected Hitachi Lightning 9980 V systems and Hitachi TrueCopy software for this consolidation effort. "We were most concerned about getting a system that performed well and had sufficient disaster recovery features. Hitachi's availability guarantee really got our attention," said Christensen. "From the time we said 'ready to go' to the time they had a system onsite for benchmarking was less than a week. Full production took less than two months."

Following this deployment, the Marshfield development team was able to use snapshots of the mirrored production data located at the backup site to speed application development and testing. Given Marshfield's heavy use of self-developed applications, any new capability that accelerates the development and testing cycle is a great value.

Going forward, Christensen would like to see Hitachi Data Systems help him with two issues. First, he still maintains several specialized medical systems (e.g., for X-ray images and MRI scans) that aren't part of the SAN. Christensen would like Hitachi to work with companies such as Siemens and GE Medical to change this situation. Second, Marshfield uses network-attached storage (NAS) for many applications and would like to move to a NAS solution that complements its existing SAN-based data protection strategy.

Disaster recovery plans begin with the assumption that an entire datacenter is no longer available. How the recovery proceeds depends on the distance to the backup datacenter and the replication method used, either synchronous or asynchronous.

Distant data replication is an important BCP concern. Enterprises must mitigate the risk of a failure affecting an entire datacenter or (increasingly) all datacenters in a metropolitan area or region (e.g., caused by an earthquake or flood).

SYNCHRONOUS DATA REPLICATION

Within a campus, metropolitan area, or region (less than 100 miles), synchronous data replication (mirroring) techniques provide the most reliable method for instantaneous duplication of data and rapid recovery.

Storage and database systems that support OLTP applications use synchronous replication to store data locally and at a second datacenter, while the OLTP monitor holds a transaction in suspense. The transaction is completed when both the local and remote sites acknowledge the operation. The advantages of synchronous data replication are the speed and completeness of the recovery. Should the primary datacenter suffer an outage, then all completed transactions are stored at the backup facility. Assuming that servers and application software are available, the failover from one datacenter to another occurs in seconds or minutes.

Synchronous systems require access to dedicated high-speed fiber-optic circuits, not so much for bandwidth as for latency. For high-performance transaction processing systems (the ones most likely to require mirroring), even a few milliseconds (ms) of additional delay can cause unacceptable performance degradation.

This latency sensitivity becomes critical as the fiber-optic links between datacenters pass critical distances. Bits of data travel at a constant speed — the speed of light in optical fiber. Although the speed of light may seem infinitely fast, over continental and global distances, latencies become a noticeable factor. Latency accrues in optical fiber at the rate of 2ms per 125mi round trip. The minimum latency for a U.S. coast-to-coast round trip (from datacenter A to datacenter B and back again) is 50ms; traveling halfway around the world and back in fiber takes approximately 200ms, minimum.

Even at limited distances (greater than 100 miles), performance will quickly become unacceptable.

ASYNCHRONOUS DATA REPLICATION

Beyond 100 miles, speed-of-light limitations often make synchronous replication impractical for data protection and recovery. IT managers must turn to asynchronous replication solutions.

Asynchronous solutions replicate data to a distant site. In contrast to synchronous replication, however, transactions commit based on acknowledgment by the local storage system without waiting for acknowledgment from the secondary site.

Asynchronous techniques don't have the same latency requirements as synchronous techniques and can leverage existing corporate networks (e.g., an IP network) to forward data at a much lower cost on an ongoing basis.

Even at a shorter distance, the impact on OLTP application response time from synchronous replication can be a problem. Therefore, a number of enterprises are deploying asynchronous replication over shorter distances to ensure an acceptable response time for OLTP applications. For database applications, they use asynchronous replication for the database data space and synchronous replication for the database log file. This method minimizes the application response time impact while ensuring a very rapid recovery with no lost data.

Asynchronous techniques are also playing an increasingly important role in BCP efforts because they are key elements in solutions for improving data protection and rapid recovery for nontransactional data. As noted previously, synchronous replication requires very expensive high-speed fiber-optic links to minimize latency for transactions. Non-transaction-intensive applications and data sources do not have these same latency requirements even over short distances.

For example, a weekly market analysis of transactional data doesn't require backup on a real-time, continuous basis. IT managers can make point-in-time copies of historical transaction data, asynchronously replicate the data to a remote site, and allow developers and business analysts to use the information for application testing and data mining.

ON-DEMAND STORAGE

The flexibility of SANs and networked storage has led to server and storage consolidation within companies and has made it possible for IT managers to consider a new element in their BCP strategies: on-demand storage provisioning.

The unexpected exhausting of available storage capacity is a common threat to business continuity, particularly when storage is server attached. Traditionally, IT managers sought to mitigate this risk by overprovisioning — deploying excess capacity "just in case."

In today's cost-sensitive IT environment, just-in-case provisioning leads to unacceptable levels of underutilization and unjustifiable costs. In a datacenter built around SANs, IT managers can consolidate storage through virtualization and logical volume management, thereby reducing overprovisioning. In addition, networked storage services, such as those provided by a SAN, can be rerouted when necessary. In the case of a server failure, a backup server can be switched online without the need to move data to a new location.

Consolidated storage will also make it much easier for companies to roll out consistent disaster recovery and remote datacenter solutions. IT managers no longer have to set up and manage many inconsistent disaster recovery solutions for different applications and services. They can look to a single solution that leverages many different data replication and recovery functions as well as provides a foundation for automating repetitive storage management tasks that consume valuable administrative resources.

The technologies to enable robust provisioning of network storage across an enterprise are still in the early stages of development and deployment; therefore, this approach is not yet suitable for all applications or businesses. Despite these limitations, however, any company that is currently evaluating its BCP strategy for the next two years should be considering the impact of these technologies.

Case Study: Financial Institution

Finding a reliable and cost-effective solution for extending data protection to a new set of applications was a critical challenge for the IT vice president at a division of a major U.S. financial institution. In the spring of 2002, the company decided to consolidate existing business applications running on Sun and Windows 2000 servers. As part of the project, the company also decided to consolidate its storage assets, over 20TB in total.

"We wanted to replace all existing JBOD storage with a SAN, and tied to that change were new disaster recovery requirements," noted the vice president. "It's not considered a hot site, but there was a two-hour requirement to get the applications up and running at the backup site." This requirement included the condition that the data be protected in case of regionwide disruptions, so any remote datacenter would have to be located more than 100 miles away, essentially mandating an asynchronous remote copy approach.

Because the IT team was guaranteeing availability as part of the new implementation, it was critical that they find a solution that effectively met the recovery requirements while controlling ongoing costs.

The company selected Hitachi Freedom Storage Lightning 9900 V Series storage system and Hitachi TrueCopy asynchronous software as the foundation for this project. It made the selection in late spring of 2002 and was up and running in the early fall. The company set up a remote datacenter with a SAN and 20TB of storage and then connected it to the primary site with two T3 lines (45Mbps each). According to the vice president, "The data migration was flawless. If I was to pat Hitachi on the back for anything, the data migration was very well thought out. I only had a 36-hour window to get everything done on a scheduled weekend. "

Since the solution came online, the IT team has also seen significant improvements in development efforts. The ability to take a snapshot of production data for testing saves hours of development time. The development staff reports that it led to a 50% improvement in the times required to develop and test applications.

Going forward, the vice president and his team would like Hitachi Data Systems to continue adding functions to improve both system and management scalability. They want to replicate between more systems and leverage different classes of storage (e.g., low-cost systems employing ATA drives) at different locations without losing or complicating existing data protection and disaster recovery processes. They expect Hitachi Data Systems to be an important partner as the company explores emerging storage options.

THREE GOALS FOR BUSINESS CONTINUITY SUCCESS

Business executive and IT managers must develop business continuity plans for application recovery and data protection that meet three specific goals:

- Accommodate the backup and recovery needs of different classes of information (e.g., structured transactions, unstructured fixed content, and long-term archival data)
- Allow IT managers to deploy a range of recovery and data protection solutions across heterogeneous applications and systems so they can balance costs, data integrity, performance, and recovery times
- Provide an overall business continuity architecture and management framework that ensures enterprisewide consistency in policies and procedures while reducing administrative overhead

One of the key steps to meeting these BCP goals is the adoption of a storage infrastructure that allows IT managers to take advantage of critical data protection and data replication capabilities. The remainder of this white paper examines the portfolio of BCP solutions offered by Hitachi Data Systems.

HITACHI DATA SYSTEMS' REPLICATION SOLUTIONS

Hitachi Data Systems is a supplier of storage systems, software, and services to enterprises around the globe. HDS conducts business through direct sales and resellers in the public, government, and private sectors in more than 170 countries. Both Sun Microsystems and Hewlett-Packard resell Hitachi systems as part of their high-end storage offerings.

HDS offers two families of storage systems:

- **Hitachi Freedom Storage Lightning 9900 V Series,** a family of high-end storage systems featuring a switched architecture and optimized for enterprisewide storage consolidation and business continuity
- **Hitachi Freedom Storage Thunder 9500 V Series,** a modular storage system family that extends storage consolidation and business continuity capabilities into environments that require lower costs and smaller footprints

HDS also offers a family of complementary storage software and services solutions that help customers get the most out of their storage systems. These software products include HDS-developed software as well as products from key partners such as CommVault, InterSAN, Sun, and VERITAS. The Freedom Storage Software products are grouped in five suites that target specific storage management issues:

- **Business continuity:** Data replication and data validation
- **Backup and recovery:** Tape backup integration
- **Performance enhancement:** Performance optimization
- **Data movement:** Data migration
- **Storage area management:** Device and SAN management, LUN/volume management

Hitachi Data Systems also provides a number of professional services programs through its Global Solution Services organization that help companies design, implement, and maintain effective storage management and business continuity strategies. These service programs provide assistance to customers dealing with a number of business issues, including storage consolidation, data migration, disaster recovery, storage integration following mergers and acquisitions, fixed content management, and regulatory compliance.

In the business continuity area, HDS offers a number of service programs to help customers effectively implement Freedom Storage Software products.

In this white paper, IDC focuses on Hitachi Data Systems' business continuity and data replication software solutions: ShadowImage, TrueCopy, and NanoCopy.

SHADOWIMAGE: POINT-IN-TIME COPY

Hitachi ShadowImage is a point-in-time copy solution that lets IT managers perform routine data management tasks that would traditionally be impossible without affecting information availability. ShadowImage replicates large volumes of information into secondary volumes that can then be split from the host application. IT and business applications can use these secondary copies for other purposes while the primary business application continues running at full capacity. Some of the most effective ways to use secondary copies are the following:

- Enhanced backup through elimination of "backup windows" that limit application availability and the enablement of LAN-free backup
- Rapid restoration of applications in case of server or software failure
- Accelerated application development through improved application testing
- Comprehensive data warehousing/data mining analysis without disrupting production applications

ShadowImage includes one other important capability: consistency groups.

MANAGING COMPLEX APPLICATION RECOVERY

Today's business applications are often highly interdependent (e.g., manufacturers need to link ecommerce, inventory, and supply chain management systems). In this increasingly complex IT environment of interdependent applications, backup, recovery, and application testing can no longer be done on an application-by-application basis. IT managers must be able to carry out simultaneous point-in-time copies for multiple data sets, thereby creating I/O-consistent copies. ShadowImage enables such group consistency functions across multiple applications without the need for host-based software.

TRUECOPY: FLEXIBLE REMOTE DATA REPLICATION

Hitachi TrueCopy is a set of data replication solutions that help IT managers safeguard critical data up to the point of failure at a primary datacenter. Hitachi TrueCopy uses automated copy processes to help enterprises reduce recovery times and resume business operations from a backup datacenter.

Hitachi Data Systems offers both synchronous and asynchronous versions of Hitachi TrueCopy for open systems and OS/390 environments.

Hitachi TrueCopy synchronous software provides a high degree of data integrity by creating mirrored, real-time copies of data at a secondary site connected via high-speed fiber-optic circuits over limited distances (less than 100 miles). In the case of failures at the primary site, servers at the secondary site can quickly take over processing of critical transactions. Once the primary site comes back online, synchronous TrueCopy also enables the resynchronization of data needed to complete recovery back to the primary servers.

Hitachi TrueCopy asynchronous software provides improved data consistency when protecting data over long distances (more than 100 miles) and for expanding protection to cost-sensitive, non-transaction-intensive data. It is a hardware-based solution that connects two storage systems (one at the primary site and one at the remote site). Connection can be via ESCON and Fibre Channel over a dedicated fiber-optic circuit as well as via channel extenders supporting WAN services such as IP, T1/T3, SONET, and ATM.

FAILING OVER WITH STYLE

One key challenge addressed by TrueCopy asynchronous software is "rolling disasters." Many datacenter failures do not happen instantaneously. Different systems and applications fail during time intervals that range from milliseconds to minutes. The window between the start of a failure and full impact can lead to significant data inconsistency and corruption if not managed properly, especially in today's world of interdependent applications. This reduced data integrity translates into significantly long recovery times.

As is the case with ShadowImage, TrueCopy asynchronous software enables consistency groups for different application volumes, thereby ensuring that I/O copy to the remote site occurs simultaneously for all critical data. Administrators at the remote site can then leverage this "crash consistent" data to more quickly restart/restore critical database applications.

NANOCOPY: PROTECT ANY AMOUNT OF DATA OVER ANY DISTANCE

NanoCopy is a technology that builds on the TrueCopy asynchronous software product to provide a completely nondisruptive point-in-time copy of any amount of data at any distance. Using this technique, IT managers can make shadow copies of critical data every few minutes, thereby ensuring the availability of critical data at a point in time *prior* to any disaster that may occur.

NanoCopy is a technology for OS/390 environments that protects against rolling disasters. NanoCopy builds on the TrueCopy asynchronous software product to create a continuous stream of nondisruptive point-in-time copies across any number of primary and secondary storage systems. Using this technique, IT managers can specify exact times at which all primary systems suspend sending updates to all secondary systems. As a result, a known good copy with full data integrity — made at

a known time that is prior to the start of a rolling disaster — is always available for recovery procedures. At the remote site, a ShadowImage copy is made of the I/O-consistent data. Copying between primary and secondary systems can then resume.

The major benefit of NanoCopy is that it can create a copy of any amount of data across any number of storage systems at any distance at specified times, without having to shut down or otherwise interrupt the application that is using (updating) the data. This capability allows the I/O-consistent copies to be made as frequently as desired, resulting in optimum data currency and shorter and smoother recovery in the event of an outage.

CHALLENGES FOR HITACHI

Companies' growing desire to consolidate storage assets on SANs as well as their continued need to more cost-effectively deploy and manage storage assets pose a number of challenges for Hitachi Data Systems and its competitors. IDC believes that HDS must focus on three specific issues.

Hitachi Data Systems must deliver storage systems that allow customers to deploy different classes of storage under a consistent storage management and data replication infrastructure. For example, companies want to deploy high-performance (but often expensive) disks and storage systems at the primary site while leveraging lower-performance but also lower-cost (e.g., low-cost ATA base drives) systems at the secondary site.

Such an approach as well as emerging disk-to-disk-to-tape strategies for rapid recovery require that HDS extend the capabilities of its TrueCopy and ShadowImage software products to a broad range of systems.

As a corollary to the above issue, IDC also sees companies working to establish more consistent storage management across storage assets from multiple vendors. These efforts are driven by companies' need to improve the utilization of existing storage assets during storage consolidation efforts and their need to implement more consistent business continuity standards without adding IT staff.

Hitachi Data Systems and its competitors must all develop solutions that allow companies to leverage their storage management and business continuity solutions across heterogeneous storage environments.

To fully take advantage of multiple classes of storage as well as support storage management across heterogeneous storage environments, storage vendors such as HDS must also develop more integrated and flexible storage management solutions. Existing storage management solutions, while powerful, are typically sold as isolated functions with minimal internal consistency or an ability to share information across applications.

Hitachi Data Systems must provide a single, centralized, and automated enterprisewide replication management tool to manage all HDS replication products as well as the ability to integrate management of other vendors' replication products. This solution must also be more tightly integrated with other applications that support data migration and storage provisioning.

THREE-TIER BUSINESS CONTINUITY: PUTTING THE PIECES TOGETHER

The tragic events of September 11, 2001, as well as studies of the success and failure of disaster recovery plans in the aftermath of September 11 are driving changes in what is minimally acceptable in a sound business continuity plan. Prior to this event, secondary sites located even a few miles away were considered an acceptable option because they provided the best guarantee of rapid recovery from most common datacenter problems.

Today, IT managers and regulators in specific industries are much more sensitive to the risks posed by regional disasters that can quickly knock out both the primary and secondary sites. The August 14, 2003, blackout, which affected cities and states throughout the northeast United States and parts of Canada, will further reinforce concerns about exposure to a regional crisis.

As a result, enterprises are now evaluating a three-tier strategy: primary datacenter, secondary datacenter (within 100 miles), and tertiary datacenter (several hundred to thousands of miles away). This strategy calls for a multisite cascade that combines synchronous and asynchronous replication to achieve the best of both worlds: a solution with no data loss over virtually any distance.

With this option, data from the primary site is replicated synchronously to a second intermediate site nearby, creating an exact mirrored copy with complete data consistency between the two sites. Periodically, the mirrored copy in the intermediate system is suspended from its pairing to replicate it asynchronously to a recovery site that can be located virtually any distance away.

The intermediate pair is then resynchronized again to make them identical, and the copy cycle to the recovery site is repeated periodically. Following a disaster at the primary site, the intermediate site contains the most current data, which can then be copied to the recovery site. The recovery site will then have the same data and consistency as if it had been in a direct synchronous remote copy session with the primary site. The ultimate goal of this approach is a current, I/O-consistent copy at the remote site that can be used for very rapid application restart.

Effectively managing such a business continuity architecture will require coordinated use of many different data protection services, including local point-in-time copy, synchronous replication, and asynchronous replication. Hitachi TrueCopy and ShadowImage products are a solid foundation upon which enterprises can move to the next level of actually deploying robust business continuity solutions.