Business Continuity Solutions Brief

A Blueprint for Business Resilience and Operational Efficiency

Executive Summary

Information technology is woven into the fabric of today's always-on business operations, making fast and continuous access to your corporate data assets essential to success. Hitachi Data Systems business continuity solutions are designed to protect your organization against loss of access to data, for whatever reason, increasing business resilience and operational efficiency.

Foremost in most people's minds is site disaster, such as earthquake, flood, and power grid failure. Although these events account for only three percent of unscheduled downtime, site disasters can be devastating to your business. According to Gartner "two out of every five companies that experience a disaster go out of business within five years." Other causes—system malfunction, hacker attempts, and human error—are far more likely to interrupt operations, and can also be countered effectively with business continuity solutions. Some solutions even eliminate downtime associated with system upgrades and scheduled maintenance.

In general, advanced approaches to disaster recovery preparedness provide more than a simple insurance policy—they take your business to the next level by providing secondary access to data without impacting regular workloads.

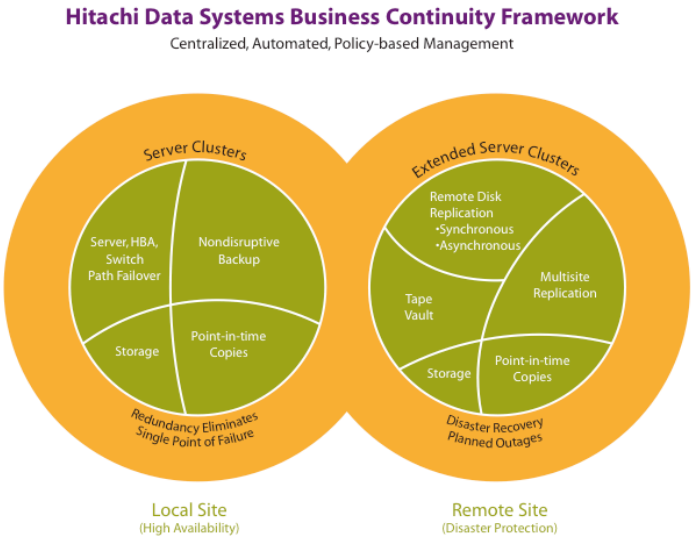

Figure 1. Hitachi Data Systems business continuity framework.

Table of Contents

Nigel Hall 2/12/07 11:21 AM
Deleted: 7

Nigel Hall 2/12/07 11:21 AM
Deleted: 7

Nigel Hall 2/12/07 11:21 AM
Deleted: 9

Securing Business Continuity

In today's highly competitive, 24/7 business environment, disasters, natural and man-made, can bring an organization to a grinding halt. In a world of just-in-time delivery, global supply chains, around-the-clock customer demand, and alternative vendors a mouse click away, delays in resuming IT operations following an unexpected disruption can have a dramatic impact on bottom-line financial performance and competitive advantage.

Meta Group estimates lost revenue from downtime at an average of US$1 million per hour and Contingency Planning Research say losses go as high as US$6.45 million per hour for retail brokerages. And the risks from unanticipated outages are not all financial. According to Binomial International, 50 percent of companies that lose critical business systems for more than 10 days never recover.

In addition to securing business data against outages, many organizations face regulatory oversight dictating mandatory levels of data protection. The United States regulations now apply to health care (HIPAA), financial services (SEC 17a-4), corporate accountability (Sarbanes-Oxley Act), life sciences (21 CFR Part 11), and government (DoD 5015.2-STD). Elsewhere in the world the story is similar, from the New Basel Capital Accord (Basel II) globally to RIPA and FAS in the U.K. and to COB in France. Like government regulators, investors and insurers are now insisting that businesses maintain feasible business continuity plans to protect critical information.

The fast restart of your business operations after a disruption is essential. Ideally, you will want to resume business at the point in time at which operations stopped functioning. To achieve this, the most up-to-date business data must be continually available at a location safely removed from the disaster area.

The perfect case scenario of a remote disaster site ready to go at a moment's notice can be costly to implement and may not be warranted for all applications. The challenge for IT is to construct a comprehensive business continuity solution from the palette of interrelated capabilities that encompass local-site replication and recovery, off-site transportation of backup data, electronic vaulting, and remote data replication.

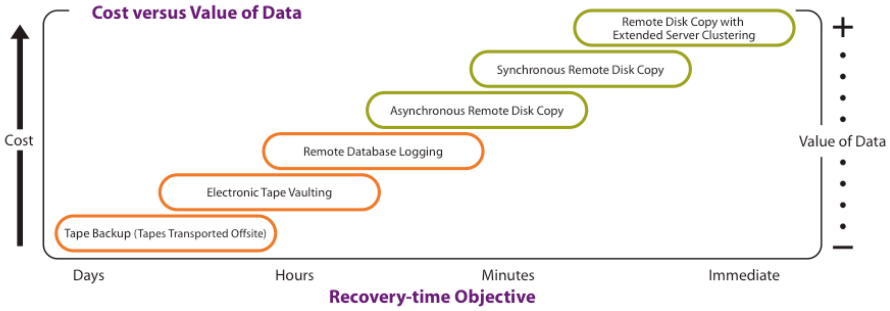

Figure 2. Application recovery-time objectives and recovery-point objectives dictate the cost of data protection.

Risk Reduction and Business Resilience

Choosing a technology to support business continuity planning starts with an assessment of the potential risks facing your organization. An operational risk management approach will clarify business requirements and reduce uncertainty, by estimating the likelihood, and loss

potential, of each type of incident. This systematic analysis of each business critical application will help to determine the extent of vulnerability, the business impact of downtime, and any regulatory requirements. With this information in hand technologies can be assessed to achieve the optimal balance of recovery speed, data value, and cost.

Two key metrics used for business continuity planning are recovery-time objective (RTO) and recovery-point objective (RPO). RTO measures how long it takes to resume essential operations—how long it takes to get back on your feet. RPO is a measure of data currency—how far behind the organization can afford to be when resuming operations after a disruption. Your RTO and RPO determine which data replication and recovery options your business needs and how much the solution will cost.

Evaluating the Options

Planning for business continuity involves matching the RTO and RPO values of each critical application to the capabilities of various data protection technologies. Figure 2 illustrates the correspondence between the RTO and RPO and the cost of a business continuity solution. For critical applications, data replication provides the most comprehensive recovery capabilities.

Remote Data Replication

Replicating data to a remote a remote secondary site represents the most effective insurance policy against system downtime. Remote data replication offers the fastest recovery time following an outage and the lowest risk of data loss. Replication eliminates the time-consuming, manual, and error prone multistep recovery process required by traditional tape-based backup. It also provides a variety of productivity benefits through secondary, or parallel, access to data, without affecting regular production workloads.

Remote data replication increases data availability by:

- Automating procedures to reduce the duration of planned events, such as system maintenance, application testing and development, and data backups
- Allowing nondisruptive backup of current production data with no impact to the production application
- Speeding failover and data restoration in the event of an outage by replacing slow and labor-intensive tape-based restores with continuously available online backups
- Allowing secondary sites to take over primary processing to eliminate scheduled downtime
- Enabling frequent, nondisruptive disaster recovery testing with an online copy of current and accurate production data

leslie steere 12/13/05 10:39 AM
Formatted: Bullets and Numbering

Two basic variations of remote data replication are available—synchronous and asynchronous.

Synchronous and Asynchronous Replication

Synchronous replication ensures that a remote copy of the data, identical to the primary copy, is created at the same time the primary copy is updated. In synchronous replication, an I/O update operation is not considered complete until confirmed at both the primary and remote sites. An incomplete operation is rolled back at both locations, guaranteeing that the remote copy is always an exact image of the primary.

Synchronous replication offers very fast recovery. After a disruption, business operations can resume immediately at the remote site, starting from the exact point the primary site stopped. Only I/Os that were in-flight at the instant of disruption will be lost. And, because neither the primary nor remote site will have a record of these transactions, rolling back a database to the last confirmed state offers full data integrity.

The drawback to synchronous replication is distance. Fibre Channel, the primary enterprise storage transport protocol, can theoretically extend a remote copy as far as 200 kilometers from the primary site. However, as distance increases, so too does latency. Delays propagating updates to the secondary site can affect application performance. Depending on response time tolerance, synchronous replication is practically limited to about 20 miles (35 kilometers). This may not be far enough to clear a wide-area disaster zone.

Asynchronous replication provides a mechanism for mirroring data across any distance, usually involving IP networks. With asynchronous replication, the primary write operation is disconnected from the remote write operation. After a successful write of data to the primary storage system, the application continues processing without waiting for an acknowledgement of a successful write at the secondary storage.

Asynchronous replication can span any distance without affecting application performance. Remote sites can be hundreds, or even thousands, of miles from the primary site, ensuring critical data is stored safely outside a disaster zone.

The downside to asynchronous replication is the potential for I/O inconsistency between the primary and remote site, and the possibility of data loss. Because of the slight time lag between data being stored at the primary and remote sites, transactions lost in-flight during an outage can mean the remote database does not pick up operations instantly at the point the primary site failed. In such a situation, asynchronous replication caching, sequence numbering, time stamps, and other techniques used to automatically preserve write-sequence fidelity at the remote site are essential.

	Hitachi TrueCopy™ Remote Replication Software	
Issue	**Asynchronous**	**Synchronous**
Data loss	Asynchronous mode may or may not lose some committed transactions in the event of an unplanned failover to the secondary site. However, a rapidly restartable data image is ensured because the remote database will be in an I/O-consistent state, resulting in the need to reapply only the most recent transaction logs to recover to the point of the outage.	Each I/O update operation waits until completion is confirmed at both the primary and mirrored sites. Any incomplete operation is rolled back at both locations; thus the remote copy is always an exact mirror image of the primary.
Distance	Asynchronous mode can span virtually any distance, because there is no propagation delay involved in confirming transactions at the remote site. Remote sites can be up to thousands of miles from the primary site, ensuring that the replicated copy of data is safely outside any likely disaster zone.	Maximum distance for synchronous theoretically extends to 200 kilometers, depending on channel extender specifications. But latency quickly becomes a problem as propagation delays (the time spent waiting for the update to travel to the remote site and confirmation to come back) lengthen with increased distance. The practical distance for synchronous replication of a busy transaction system is usually about 30 to 50 kilometers (20 to 30 miles), possibly further, depending on an application's tolerance for delayed response

		and other factors.
Performance Impact	The performance impact on the host is minimal.	Synchronous mode has a greater impact on performance than asynchronous mode, because a write from the host must wait for acknowledgment from the secondary storage system.
Data Integrity	TrueCopy software uses several mechanisms to ensure the remote copy is made in precisely the same write sequence as the primary copy, including the use of sequence numbers and time stamps in the data packets.	With synchronous replication, data integrity issues caused by out-of-sequence writes do not arise because "dependent writes" are not initiated until prior writes on which they depend complete.
Network Infrastructure	When coupled with bridging technology from Cisco or McDATA, Hitachi Data Systems replication technologies support a wide range of network infrastructures, including ESCON, FICON, IP, SONET, ATM, and others.	

Table 1: The pros and cons of asynchronous versus synchronous replication with Hitachi TrueCopy™ Remote Replication software

The Rolling Disaster Challenge

A rolling disaster occurs when an unplanned outage event takes place over a span of time. The time span can be anywhere from a few minutes to several hours.

During a rolling disaster not all systems, storage, and network connections fail at precisely the same moment. In this situation, a system may still be able to process transactions and issue updates to primary storage devices, but, due to earlier failures, updates may not replicate successfully to the secondary site. Rolling disasters pose a challenge because they often result in corrupted and unusable data at the remote site, requiring difficult and very lengthy recovery processes.

To protect against rolling disasters, a data replication technology must be able to freeze remote replicas at a point in time prior to the onset of the outage. The ability to create frozen point-in-time images of data is what differentiates remote copy technology from simple mirroring.

Because the remote and local I/O of a synchronous replication succeed or fail together, this replication approach does not introduce data inconsistencies following a disaster. Rolling disasters are primarily a challenge for remote asynchronous replication, and one of the principle problem areas is write-sequence fidelity.

Write-Sequence Fidelity

Database and file managers maintain very complex internal data structures, including indexes, structured data tables, directories, logs, and so forth. To preserve the integrity of these internal structures, each write is carefully sequenced so that, at any point in time, a correct file system or database state can be re-created.

Writes to a database often depend on other writes—for example, a transaction that applies a credit to an account, followed by a debit—and this sequence must be strictly preserved during a recovery. Careful write sequencing and strict adherence to dependencies allow file systems and databases to recover from failures, no matter what I/O activity is in progress when failure occurs.

A replication solution must address write-sequence fidelity. Writes to remote data copies must be in the same order as those at the primary storage.

Asynchronous replication introduces potential write-sequence fidelity problems because packets of data can arrive at the remote site out of order. This is especially true of

technologies that send disk track updates, which cannot preserve the original write sequence. To ensure the integrity of asynchronously replicated data in a rolling disaster, replication technology must employ techniques to automatically preserve write-sequence fidelity at the remote site.

Choosing Between Synchronous and Asynchronous Replication

Clearly, remote storage replication for recovery and business continuity requires more than just shipping data over a network. The selection process starts with an assessment of the potential risks and their probability.

If your organization cannot tolerate any data loss and operations must be resumed quickly following an outage, synchronous replication is likely to be the best choice. Of course, the decision must also factor in how far the data has to be replicated to clear any likely disaster zone, and how much degradation of application performance can be tolerated.

On the flip side, if your organization can tolerate being down while the last few transactions are reconstructed—or cannot tolerate the performance impact of synchronous propagation delays—asynchronous may prove to be a less costly option.

Another approach to be considered combines the benefits of both synchronous and asynchronous replication: the three-data-center option.

The Three-Data-Center Option

Two-data-center replication strategies are viable for most in-region recovery—for example, serving as a hot site for campus-level or metro-level server cluster—and for out-of-region recovery sites where propagation delays are not an issue. Synchronous replication provides very fast recovery time (low RTO) and good data currency (low RPO). However, asynchronous replication provides better protection against regional disasters, albeit with a less favorable RTO and RPO. A three-data-center strategy offers the best of both worlds: fast recovery and excellent data currency for local site failures, combined with advanced protection from regional disasters.

Hitachi Data Systems 12/16/05 10:40 AM
Comment: In the graphic, delete the pass through configuration

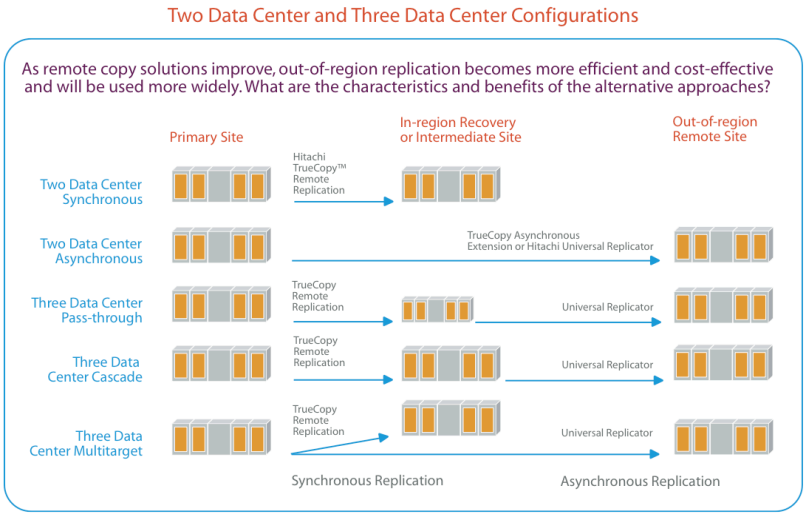

Figure 3. Comparing two-data-center and three-data-center replication strategy options

Universal Replication: A Major Breakthrough

With heterogeneous storage a feature of most midrange and enterprise data centers, the ability to replicate between systems from different vendors, locally and long distance, is critical to cost-effective business continuity planning. The downside to multivendor storage business continuity is the cost and complexity of the solution. Software, in-house expertise, and licenses must be duplicated for each vendor's replication solution.

Tiered Storage and the Universal Replicator

The Hitachi TagmaStore Universal Storage Platform and Network Storage Controller provide enterprise and midrange storage environments with a consistent interface to pooled multivendor storage. Using industry-leading controller-based virtualization, the Hitachi storage systems enable a single replication tool to operate against all heterogeneous storage resources in a tiered infrastructure. This significantly reduces the complexity and cost of replicating data, both locally and long distance.

The Hitachi Universal Replicator for TagmaStore Universal Storage Platform software provides advanced, controller-based replication among all of the storage systems certified for attachment to the Universal Storage Platform and Network Storage Controller. Universal Replicator software disk-based journaling increases replication resilience, reduces the impact of replication on production applications, and lowers replication costs. When collecting data to be replicated, the primary Universal Storage Platform writes the designated records to a special set of performance-enhanced journal volumes. The remote Universal Storage Platform then reads the records from the journal volumes, offloading the primary system by pulling them across the communication link, instead of making the primary system push them, as in most other approaches. By writing records to journal disks instead of keeping them in storage system cache, Universal Replicator software no longer consumes a large percentage of available cache, freeing resources for production transactions. With only one set of procedures, tools, and skills to be mastered, regular and comprehensive disaster recovery testing of complex heterogeneous environments becomes more feasible and affordable.

Combining Universal Replicator software and the Universal Storage Platform or Network Storage Controller enables three-data-center business continuity configurations to be achieved with less complexity, scripting, and overhead than previous solutions. This, in turn, makes three-data-center business continuity solutions more affordable for a broader range of enterprises and applications.

Local, In-System Replication

Synchronous and asynchronous replication receive much of the attention from business continuity planners, but local, in-system, copy solutions can add a valuable complementary layer to a business continuity configuration. Hitachi ShadowImage ™ In-System Replication software provides high-speed, nondisruptive replication for any Hitachi storage system or pool of storage virtualized by the Universal Storage Platform or Network Storage Controller.

ShadowImage software consistency groups allow a user-defined group of ShadowImage volume pairs to be simultaneously split, at a precise moment in time, with a single command. This method creates a consistent point-in-time (PiT) copy of an entire system, database, or any related sets of volumes—a technique that can replicate data between a primary system and secondary systems, anywhere in the world, with full data integrity.

Financial Services Firm Use Case

A financial services firm needed to enhance its business continuity infrastructure, with an out-of-region disaster recovery capability, in addition to creating daily PiT copies of the production online transaction processing (OLTP) database, for use in an online decision support system (DSS) that requires up-to-date customer data.

Hitachi Data Systems, in consultation with the IT team, decided to apply the Universal Storage Platform business continuity framework to satisfy the requirements.

The Universal Storage Platform, an evolutionary combination of the industry's highest-performing and most scalable storage hardware and software technologies, delivers unified data and storage services across different storage platforms to simplify information access and reduce management complexity.

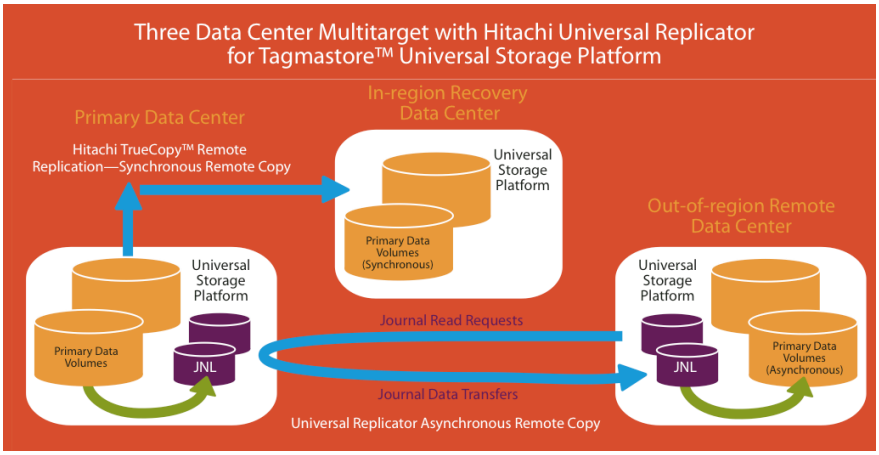

Figure x. Three-data-center, multitarget replication configuration

PiT Copies for a Large Database

The online DSS application requires up-to-date copies of customer data, and large queries often run in excess of 36 hours. Hitachi Data Systems and the company's IT group determined that maintaining multiple copies of the production data would satisfy the requirements—a new PiT copy is made every 24 hours, allowing in-process queries to complete on earlier copies of the database.

The configuration uses ShadowImage software and Hitachi Business Continuity Manager software (formerly known as CopyCentral) to split off PiT copies with ensured data consistency and integrity even through the database application reads and writes to several storage volumes. This approach creates duplicate DSS copies with no impact to production OLTP applications. As a part of the preproject assessment, the team determined that the financial benefits of running DSS queries on current data (rather than two-day-old data) would more than justify the cost of maintaining two full DSS copies of the production database.

Three-Data-Center Replication

To support its business continuity needs—including regulatory agency guidance on out-of-region disaster recovery capabilities—the project team selected an advanced three-data-center

replication solution from Hitachi Data Systems. Remote Replication software maintains an exact, synchronous copy of the production OLTP data at an in-region disaster recovery hot site 20 miles away. Universal Replicator software separately maintains an asynchronous copy at a remote disaster recovery site 1,000 miles away. Compared to a cascade or pass-through approach, this configuration eliminates the potential of having no disaster recovery should the intermediate site fail.

Backup and Testing

The solution also uses Universal Storage Platform capabilities to make copies of data on lower-cost, externally attached Hitachi Thunder 9500™ V Series systems with the SATA Intermix Option. These copies are used for backup and disaster recovery testing. The advanced three-data-center configuration also supports rapid production takeover at the remote site when required for operational recovery, full disaster recovery testing, or scheduled upgrades and configuration changes.

As this typical use case illustrates, Universal Storage Platform, ShadowImage software, Universal Replicator software, and remote replication software provide a robust and powerful foundation for meeting the requirements of large organizations with complex needs in production, development, testing, archiving, and business continuity. In addition, the solution supplies current information for more effective decision support.

Implementing Your Business Continuity Solution

This Hitachi Data Systems Business Continuity Solutions Brief has started you on the road to understanding the extent of your business continuity exposure and making a business case for solving it. With this basic understanding of what is involved in implementing the best business continuity solution for your organization, you are better equipped to choose from among the available software and services to simplify disaster recovery, minimize downtime, speed recovery, and protect your information assets while maximizing the use of resources and personnel.

Although reading is a good first step, Hitachi Data Systems strongly recommends that you engage our Global Solution Services group before you undertake a business continuity plan. Global Solution Services can help you:

- Identify and analyze your business goals, in terms of the length of time your business can afford to be down and how much data the business can afford to lose
- Identify your technical goals, based on which replication strategy makes sense, the nature of your existing infrastructure, the topologies of your storage layout and placement, and traffic patterns by application and host
- Create a migration plan to a new solution that includes test and verification components
- Document the design and create a detailed implementation and test plan
- Manage the implementation and execution phase of the plan

leslie steere 12/13/05 10:39 AM
Formatted: Bullets and Numbering

Pertinent offerings from Global Solution Services include:

- Business Continuity Readiness Evaluation Service
- Risk Analysis Workshop

leslie steere 12/13/05 10:39 AM
Formatted: Bullets and Numbering

- Business Continuity Strategic Planning Service

To learn more about how Hitachi Data Systems can help you with your business continuity plans and to read more about business continuity, please visit www.hds.com/bc or call Hitachi Data Systems at 888 234 5601, ext.950, to explore an engagement that will result in the optimal solution for your business continuity needs.