HITACHI
DATA SYSTEMS

Business Continuity Solution Blueprint:

Synchronous Data Replication

Business continuity has become one of the top issues facing enterprises globally. Data growth is exploding, and more and more enterprises must have 100 percent access to data 24/7.

At the same time, internal and external threats to data uptime are increasing every day. The caprices of Mother Nature and fragile power grids make headlines worldwide and malicious system attacks require corporations to proactively think about—and plan for—successful disaster recovery and business continuity. Simply put, your business cannot afford to just wait and react to the havoc that internal or external forces can wreak on your data.

Most companies cannot tolerate more than a few hours or even minutes of downtime without serious impact to the bottom line. Meta Group estimates lost revenue from downtime at an average of US$1 million/hour. Contingency Planning Research says losses go as high as US$6.45 million/hour for retail brokerages. Beyond the loss of revenue, there are adverse headlines and the potential impact on company valuation to consider, not to mention lower employee productivity caused by sporadic outages.

It adds up to this: Your data really *is* your business.

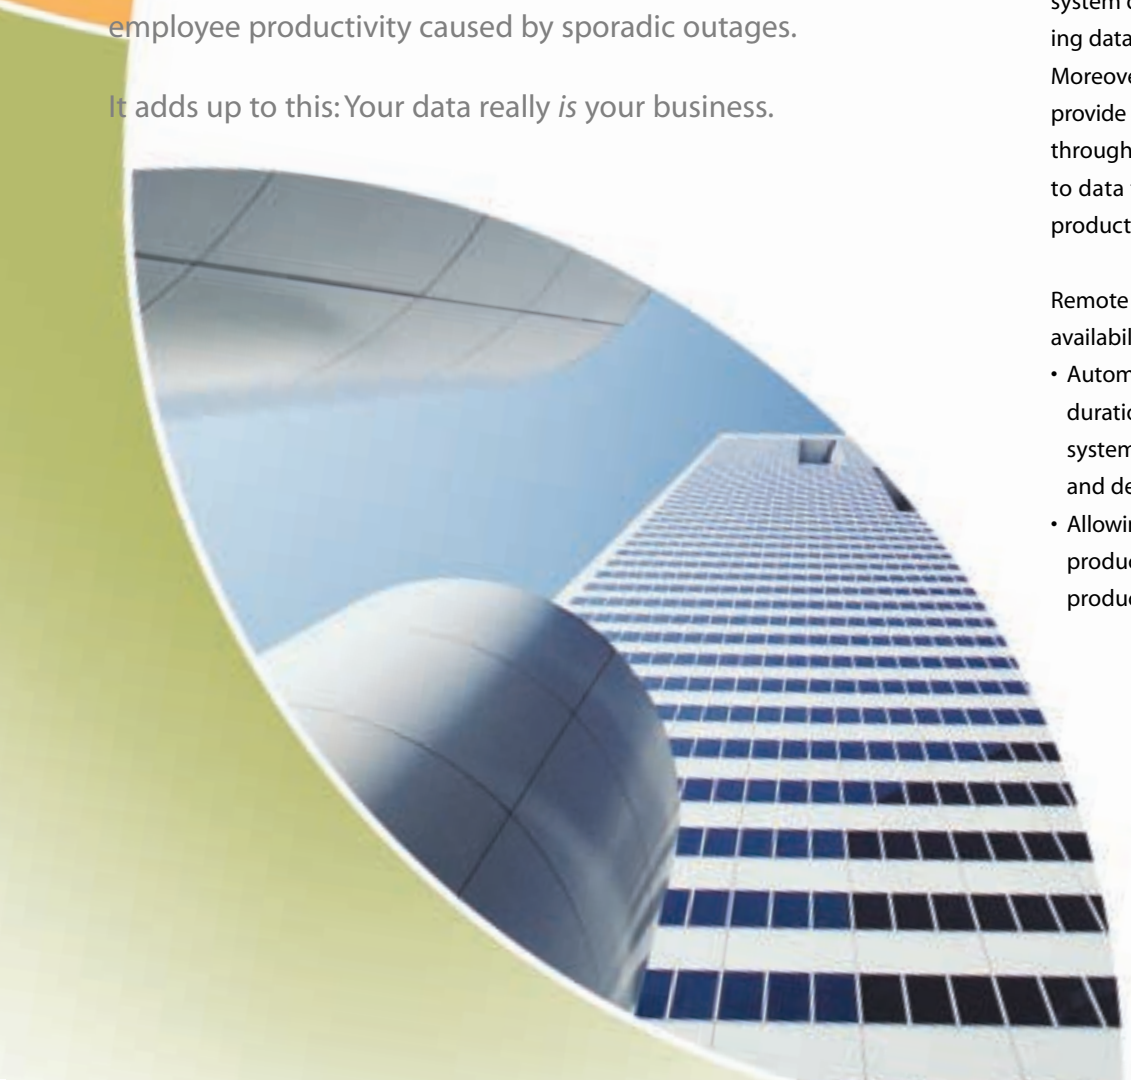

Governments are also getting into the act with mandates to protect business data. For example, United States regulations on data protection now apply to health care (HIPAA), financial services (SEC 17a4), corporate accountability (Sarbanes-Oxley Act), life sciences (21 CFR Part 11), and government (DoD 5015.2-STD). Elsewhere in the world the story is similar, from the New Basel Capital Accord (Basel II), which will exclude businesses from global finance markets in 2007 if they haven't complied to RIPA and FAS in the United Kingdom and to COB in France. Beyond government regulators, investors and even insurers are insisting that businesses put in place feasible disaster recovery and business continuity plans to protect critical information.

What Is Remote Data Replication?

Of the wide range of technical approaches to the business continuity problem, the most effective insurance policy against system downtime is provided by replicating data to a remote secondary site. Moreover, remote data replication can provide a variety of productivity benefits through secondary or parallel access to data that does not impact 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

Hitachi Data Systems is a leading provider of high-availability storage solutions, and our systems tend to be deployed in support of an organization's most critical operations. Accordingly, this Solution Blueprint focuses on the high end of business continuity through synchronous remote data replication.

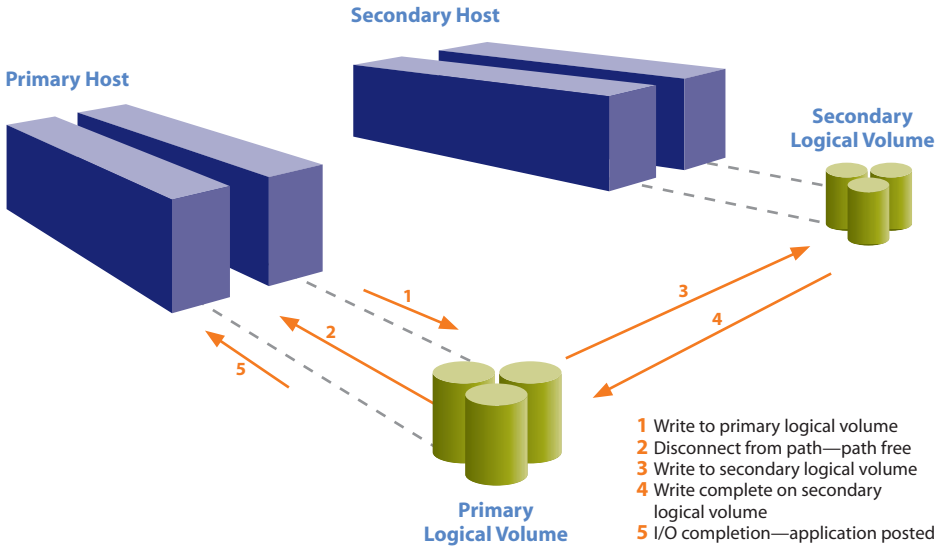

Figure 1: *In a synchronous solution, serverless remote copy activity occurs at the LUN level, as a storage controller at the primary site links to one at the remote site.*

Synchronous remote data replication is ideal for shorter distances and near-real-time data copy; as such, it is the appropriate solution for many organizations seeking the fastest possible data recovery, minimal data loss, and protection against database integrity problems.

Hitachi Data Systems provides a portfolio of replication solutions for open systems and mainframes over any distance, using synchronous for shorter distances or asynchronous solutions for longer distances, depending on the business recovery objectives. See Table 1 for a pro-and-con comparison of the various solutions, or read our companion publication: *Business Continuity Solution Blueprint: Asynchronous Long-Distance Data Replication* at http://www.hds.com/bc.

The Value of Synchronous Replication

Synchronous replication ensures that a remote copy of the data, which is identical to the primary copy, is created at the time the primary copy is updated. In synchronous replication, an I/O-update operation is not considered done until completion is confirmed at both the primary and mirrored sites. An incomplete operation is rolled back at both locations, ensuring that the remote copy is always an exact mirror image of the primary (see Figure 1).

The main advantage of synchronous replication is that data can be recovered quickly. Operations at the remote, mirrored site can begin immediately at the point in time where the primary site stopped should operations at the primary site be disrupted. Only the few operations in process at the instant of disruption may be lost. Because neither the primary nor remote site will have a record of those transactions, the database rolls back to the last confirmed state.

The chief drawback to synchronous replication is its distance limitation. Fibre Channel, the primary enterprise storage transport protocol, can theoretically extend as far as 200 kilometers. However, latency quickly becomes a problem because propagation delays lengthen with increased distance. Propagation delays can significantly slow down a system by forcing it to wait for confirmation of each storage operation at the local and the remote sites. This means that the practical distance for synchronous replication for a busy system is about 35 to 50 kilometers or 20 to 30 miles, depending on the application response time tolerance and other factors—not far enough to be clear of a wide area disaster zone.

The Rolling Disaster Challenge

The rolling disaster challenge refers to unplanned outages that occur over a time span of anywhere from minutes to hours. These events are called rolling disasters because not all systems, storage, or network connections fail at precisely the same moment. During such a disaster, components fail independently, resulting in corrupted and unusable data that often requires difficult and very lengthy recovery. Examples of rolling disasters include the events of September 11, 2001, earthquakes, tornados, and floods.

One common consequence of a rolling disaster is undetected corruption of online data. Information services may continue to function, but because of faulty input or other undetected failures, critical data gradually becomes corrupted. If this scenario occurs, any data-protection solution that replicates data updates in real time would faithfully reproduce the corruption at the recovery site. Thus, a basic property of data replication used to protect against rolling disasters is the ability to "freeze" replicas at points in time so that recovery procedures can at least restore data to its state at some point prior to the onset of a rolling disaster. The ability to create point-in-time frozen images of data is what differentiates remote copy technology from simple mirroring.

Special point-in-time techniques are required for asynchronous replication over long distances to handle this need to freeze the data image when one portion of the system fails. It is widely but erroneously assumed that synchronous replication does this as a matter of course. After all, if the application does not complete its write until all replicas have been successfully written, doesn't that effectively freeze the image in a self-consistent pre-disaster state? In fact, the action taken depends on how the environment has been set up and what controls are in place. Without the correct procedures and freeze commands, the production environment can be affected and the data at the secondary site corrupted, or both.

	Data Replication Solution	
Issue	TrueCopy Asynchronous	TrueCopy 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 only reapply 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 since 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 mode 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) lengthens with increased distance. The practical limit for synchronous replication of a busy transaction system is usually 20 to 30 miles (about 35 to 50 kilometers), depending on the application's tolerance for delayed response and other factors.
Performance Impact	The performance impact on the host is minimal.	Synchronous mode has significantly more performance impact than asynchronous mode, because a write from the host must wait for acknowledgment from the secondary storage system.
Data Integrity	TrueCopy 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.

Table 1: *The pros and cons of asynchronous versus synchronous replication*

Write-Sequence Fidelity

The other major challenge with asynchronous replication technology is the "write sequencing" problem. Database and file managers maintain some very complex internal data structures, including indexes, structured data tables, directories, logs, and so forth. They preserve the integrity of these internal structures by carefully sequencing every write that affects them so that at every stage, a correct file system or database state can be restored if necessary. With synchronous replication, sequencing is not a problem because "dependent writes" are not initiated until prior writes on which they depend complete.

Communication Facilities

Some synchronous storage-based replication implementations use ESCON® channels, a carryover from the OS/390® operating system, which has featured replication of online data for more than a decade. ESCON adapters support channel distances of up to 3 kilometers, and channel extenders can increase this distance up to 43 kilometers or 25 miles for a direct connection. ESCON has given way to FICON™ and Fibre Channel. FICON offers higher speed and supports distances up to 20 kilometers.

Using channel extenders over a telecommunications link is another option. With the emergence of dark fiber and dense wavelength division multiplexing (DWDM) technology, systems can be directly connected up to 100 to 200 kilometers or about 60 to 120 miles. Some enterprises use these offerings to connect two data centers that, while in the same metro area, may provide more protection against a metro-wide disaster in addition to higher bandwidths.

Regardless of connectivity type, synchronous replication has performance limitations that may appear before distance limitations are reached. For example, Fibre Channel, the primary enterprise storage transport protocol, theoretically extends up to 200 kilometers over DWDM. Yet regardless of the connectivity type, latency rapidly becomes a problem. Propagation delays can slow down a busy system dramatically by forcing it to wait for confirmation of each storage operation. As a result, the practical distance for synchronous replication of a large, busy online transaction processing (OLTP) system application is about 20 to 30 miles, depending on the application's response time tolerance and other factors—not necessarily far enough to clear a disaster's impact zone.

Is Synchronous Replication for You?

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 for your particular enterprise. Next comes a business impact analysis (BIA). The BIA helps determine which applications require the most protection based on the value of the data and the business impact of downtime as well as other economic factors.

The overriding goals of a business continuity plan must be to survive a disaster and resume operations as quickly as possible. The best recovery plan to achieve those goals depends on how your organization chooses to balance three factors—recovery speed, value of data, and cost. Determining the recovery-time objective (RTO) and the recovery-point objective (RPO) will define how quickly an enterprise needs to recover to survive and how much data loss can be tolerated.

RTO defines the time frame in which specific business operations must be restored. It answers the question: How long can a business afford to be down? RPO defines the point in time from which to recover. It answers the question: How much data can the business afford to lose?

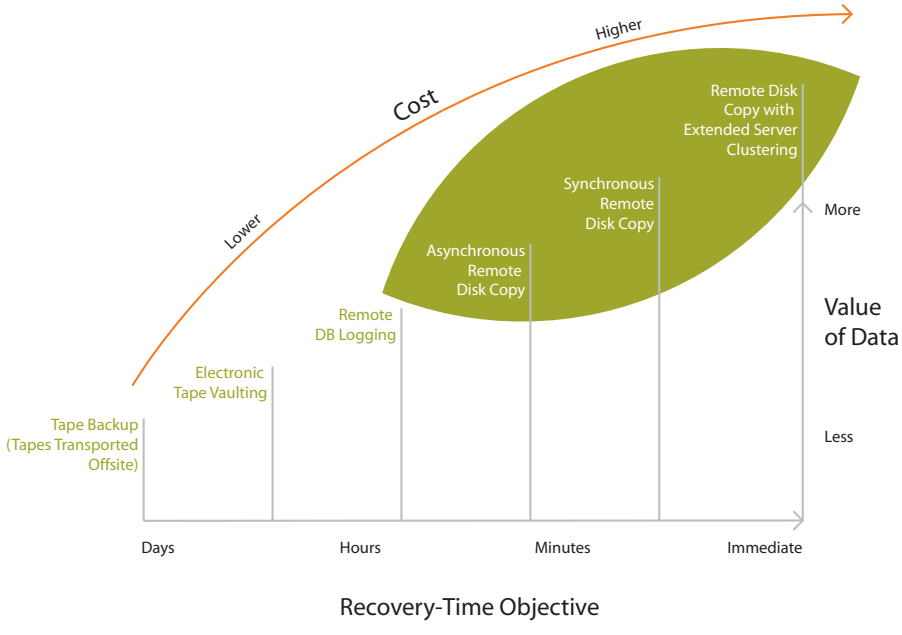

Figure 2: *Recovery-time objective combined with a recovery-point objective, based on the value of particular data, map to a range of technical approaches, costs, and degrees of data protection.*

The specific RTO and RPO determine which data replication and recovery option your business needs and how much it will cost. The next step is to find the method that satisfies your RTO and RPO values, which in effect matches the value of data to the cost of the solution.

Figure 2 shows the spectrum of the common techniques for data replication to protect from disasters. No one method fits every application.

If your business cannot tolerate any data loss and operations must be resumed quickly, then synchronous replication may be the answer. Of course, that depends on how far the data has to be replicated to clear the likely disaster zone and how much degradation of the specific application's performance can be tolerated.

On the flip side, the organization that can tolerate being down until it can reconstruct the last few transactions—or that cannot tolerate the performance impact of synchronous propagation delays—might opt for a less costly asynchronous option. Clearly, the decision is a complex one, which is why companies worldwide are working with Hitachi Data Systems Global Solution Services to benefit from our extensive experience in replication implementations.

Implementing Your Business Continuity Solution

Hitachi Data Systems can help you implement the appropriate solution derived from the full range of possible approaches. We have uniquely replicated thousands of terabytes of critical data synchronously for our customers. Why were we the vendor of choice? Developing technology that replicates data in near-real time is just a starting point. It is experience and a true culture of partnership that makes the difference when the life of an organization is at stake.

This Hitachi Data Systems Solution Blueprint has delved into the particulars of a Hitachi TrueCopy™ Synchronous solution to your business continuity needs, but first you may need a full understanding of your business continuity risks and a business case for solving it. Our Global Solution Services group is available to help you choose from among the available software and services to simplify disaster recovery, minimize downtime, speed up recovery, and protect your business's information assets while maximizing the use of resources and personnel.

Specifically, Hitachi Data Systems Global Solution Services can help you:

- Identify and analyze your business goals in terms of how long 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 system 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

Achieving Solid System Uptime with TrueCopy Synchronous

The customer—the IT arm of a government body—needed to ensure reliable data uptime for the myriad mainframe and distributed open systems it supported. IT knew what was needed. Storage connected through a Fibre Channel SAN was to replace direct attach storage from a number of vendors. Synchronous remote copy, based on either Fibre Channel or FICON, transmitted over dark fiber to a secondary site approximately 10 kilometers away would provide parallel processing, backup, and disaster recovery.

In addition, IT wanted a vendor that also offered asynchronous replication capability to handle all data integrity considerations for possible future support of a distant secondary site.

The first hurdle was clearing the numerous requirements set by IT. For mainframes, they used an IBM-based platform, which included the z/OS operating system and FICON interfaces. The infrastructure also included open systems connected by Fibre Channel—and both required support by the same storage system. IT specified compatibility with many IBM mainframe performance-enhancing features including adaptive cache management, sequential-read pre-staging, and sequential data striping. RAID-5 would be acceptable for the striping capability, but IT specifically excluded use of RAID-S because of overhead issues. They noted their admiration for the enhanced RAID-5 implementation in the Hitachi Freedom Storage™ Lightning 9900V™ storage system, which minimizes the write penalty incurred by standard RAID-5 implementations by keeping write data in cache until an entire stripe can be built, and then writing the entire data stripe to the disk drives.

More importantly from a business continuity standpoint, IT demanded the following features to provide system resiliency:

- Redundant components including power supplies, cooling, cache, controllers, and internal data path.
- Dynamic sparing
- Remote diagnostics
- Nondisruptive updates to microcode

IT was also concerned about the response-time impact of the replication technology. They specified an average 5 millisecond response or less, even taking into account the remote copy response-time delay over a 10 kilometer fiber medium.

Finally, there was the practical matter of proven experience. The organization wanted to know if Hitachi Data Systems could provide nondisruptive migration of data from the existing storage systems to the new local and remote pair. IT also wanted to hear from other customers who had successfully implemented synchronous replication of critical data.

Hitachi TrueCopy Synchronous provided the answer

Hitachi Data Systems was able to provide local references who had been very satisfied by the migration and replication capabilities provided to them by Hitachi Data Systems. They vouched for the ability of Hitachi TrueCopy operations to be nondisruptive, allowing the primary (main) volume of each TrueCopy volume pair to remain online to all hosts for both read and write I/O operations while TrueCopy provided continuous, real-time data backup.

Once implemented, TrueCopy's time delay over 10 kilometers was less than one millisecond, and thus, response time from Lightning 9900V storage systems was more than adequate. This, combined with Hitachi Data Systems experience in replicating thousands of terabytes of critical data synchronously for customers made Hitachi the right choice for this customer.

Throughout the data migration and replication implementation process, Hitachi Data Systems Global Solution Services worked with the customer to maintain fidelity to the comprehensive synchronous data replication plan, ensuring the overall solution met the customer's business and technical objectives, staff requirements, migration issues, and training needs.

To learn more about how Hitachi Data Systems can help you with your business continuity plans, please visit us online www.hds.com/bc. Or, call Hitachi Data Systems at (888) 234-5601, ext. 950, to start on the path toward the optimal solution for your business continuity needs.

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

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

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

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

HITACHI
DATA SYSTEMS

DISK-464-00

BC Solution Blueprint: Synchronous Data Replication