

Synchronous Data Replication



Synchronous Data Replication

Business continuity has become one of the top issues facing businesses and organizations all around the world. Data growth is exploding. That's why many enterprises now require as to close to 100 percent access as possible to their data, 24/7.

At the same time, internal and external threats to data uptime are increasing every day. These can range from malicious systems attacks to natural disasters and fragile power grids. Corporations must 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. Beyond lost revenue, sporadic outages can cause adverse headlines, lower employee productivity, and negative impact on company valuation.

What Is Remote Data Replication?

Business continuity can be addressed by a wide range of technical approaches. However, the most effective insurance policy against system downtime is provided by replicating data to a remote secondary site. Remote data replication can also 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 provides a portfolio of replication solutions for open systems and mainframes over any distance, using

synchronous capabilities for shorter distances and 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: Asynchronous Long-distance Data Replication at <http://www.hds.com>.

This Solution Brief focuses on the high end of business continuity through synchronous remote data replication.

The Value of Synchronous Replication

Synchronous remote data replication is the appropriate solution for organizations seeking the fastest possible data recovery, minimal data loss, and protection against database integrity problems.

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.

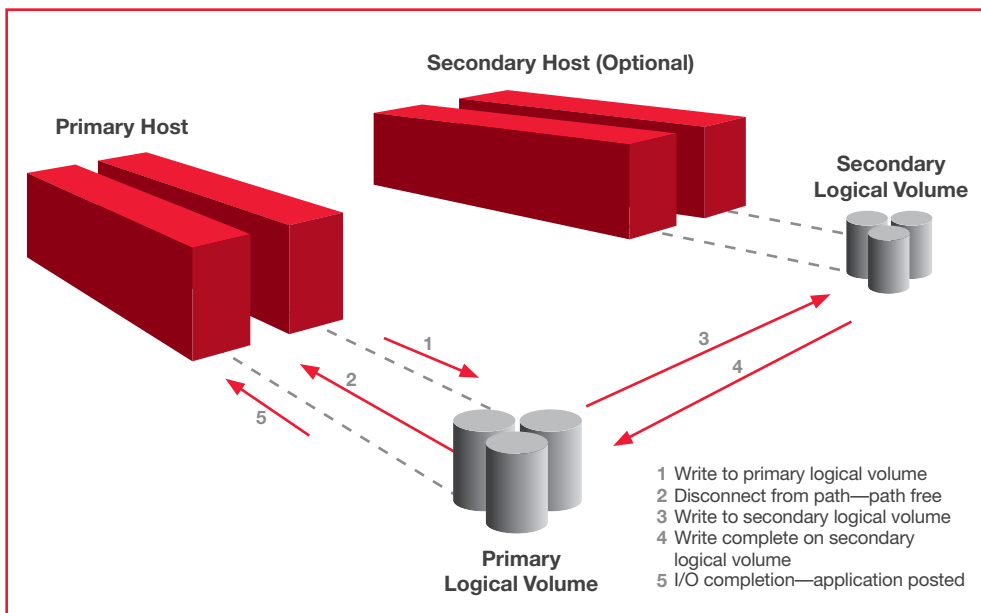


Figure 1. In a synchronous solution, serverless remote copy activity occurs at the LUN level, as a storage controller at the primary site links to one at the remote site.

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 (km) or 124 miles. However, latency quickly becomes a problem as 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 35km to 50km or 20 miles to 30 miles—possibly further—depending on the application response time tolerance and other factors; but this is not far enough to be clear of a wide-area disaster zone.

The Rolling Disaster Challenge

A rolling disaster refers to an unplanned outage causing failures to occur over seconds, minutes, or even hours. Examples of rolling disasters include the events of Sept. 11, 2001,

earthquakes, tornados, and floods. Because systems, storage, or network connections may not all fail in the same instant, the remote replication connections may fail first. This can create a situation where the system is still able to process transactions and issue updates to the primary storage devices. However, due to the first rolling disaster failures, some of the updates cannot be replicated at the recovery site. That can lead to an inconsistent set of data at the recovery site.

When the rolling disaster completes and the system is no longer able to process transactions, attempted recovery at the remote site may fail due to these inconsistencies. Thus, a basic property of data replication used to protect against rolling disasters is the ability to “freeze” replicas at points in time 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.

Clearly, the replication solution must be designed, implemented, and managed in order to avoid this situation. Hitachi TrueCopy® Synchronous software provides management tools that ensure I/O-consistent replication through a rolling disaster.

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. These include indexes, structured data tables, directories, and logs. The integrity of these internal structures is preserved 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 are completed.

Communication Facilities

Some synchronous storage-based replication implementations use IBM® ESCON® channels, a carryover from the OS/390® operating system, which has featured replication of

Data Replication Solution

Issue	Hitachi Universal Replicator software (asynchronous)	Hitachi TrueCopy® Synchronous software
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. That means you only need to 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 theoretically extends to 200km, 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 30km to 50km or 20 miles to 30 miles, and 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 significantly more performance impact than asynchronous mode, because a write from the host must wait for acknowledgment from the secondary storage system.
Data integrity	Universal Replicator software uses sequence numbers, time stamps in the data packets, and other mechanisms to ensure the remote copy is made in precisely the same write sequence as the primary copy.	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.
Recovery point objective (RPO)	High flexibility: unique capabilities afforded by journaling provide organizations with ability to ride out network outages or temporary excessive bandwidth requirements. In addition, administrators can decide to size their bandwidth below maximum requirement in order to improve their operational costs, but at the cost of a higher RPO. This is a tradeoff.	RPO = 0. In other words, data is always current, but only at a limited distance (see above).
Three data center (3DC)	Provides the distance replication leg of the 3DC configuration. In conjunction with Synchronous Remote Replication software, it provides either multitarget or cascade configurations.	Synchronous software provides the “metropolitan” or near-distance leg in either multitarget or cascade configurations.

Table 1. The pros and cons of asynchronous versus synchronous replication through Hitachi Data Systems software.

online data for more than a decade. ESCON adapters support channel distances of up to 3km, and channel extenders can increase this distance up to 43km or 25 miles for a direct connection. IBM FICON® offers higher speeds and supports distances up to 20km or more than 12 miles.


Using channel extenders over a telecommunications link is another option. With the emergence of dark fiber and dense wave division multiplexer (DWDM) technology, systems can be directly connected up to 100km to 200km or about 60 miles to 120 miles. Some enterprises use these offerings to connect two data centers in the same metro area. This 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 200km 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. Use of an efficient technique and careful design will help, but the practical distance for synchronous replication must still be considered against the application's response time requirements.

Is Synchronous Replication for You?

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 a particular enterprise. This is followed by a business impact analysis (BIA). The BIA helps determine which applications require the most protection, based on the value of the data, the business impact of downtime, and other economic factors.


The overriding goals of a business continuity plan are 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 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 your business afford to be down? RPO defines the point-in-time from which to recover. It answers the question: How much data can your business afford to lose?

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 common techniques for data replication to recover 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. The decision is

complex. That's why companies worldwide work with Hitachi Data Systems Global Solution Services (GSS). They rely on the extensive experience of the GSS team in replication implementations.

Three Data Center Solutions: The Best of Both Worlds?

It is also possible to leverage advanced configurations that combine synchronous and asynchronous replication. These are known as three data center or 3DC configurations.

For business-critical applications and data, many organizations have considered or adopted a 3DC replication model. In the 3DC model, synchronous replication—which can achieve a copy that is fully up-to-date at all times over limited distance—is employed between the primary data center and a nearby

5




Figure 2. Recovery time objective based on the value of particular data, map to a range of technical approaches, costs, and degrees of data protection.

hot site. The data is then replicated to a geographically remote site using asynchronous replication, which works over any distance. In the event of a local disaster, customers have the option of recovering to the intermediate site (if it has appropriate computing and staff resources), or to the remote site with some catch-up delay. In the event of a disaster geographically broad enough to impact the primary and intermediate sites, recovery is still possible by failing over to the remote site with some delay and some data loss. Thus, a 3DC configuration provides the best combination of protection against disaster with minimal data loss and downtime.

Table 2 summarizes the benefits of two data center (2DC) and 3DC remote-replication configurations, compared to a traditional one data center (1DC) approach that relies on physical tape movement.

Alternative replication strategies provide different levels of protection.

Table 2 shows the relative benefits of these configurations in terms of recovery speed and data currency after a primary site failure, as well as recovery speed after a regional disaster. The chart also illustrates the impact of a failure outside the primary site, which could affect the ongoing level of protection and recovery capability in case of an additional site failure.

The scope of protection increases from left to right in Table 2:

- The 1DC strategy offers the least protection from data loss during disasters, and it can require days to recover data from tape copies. The best-case scenario for recovery may be 24 to 168 hours, depending on the availability of data tapes at offsite recovery services.
- The 2DC strategies can provide good recovery speed and data currency for primary site failures, with certain trade-offs: Synchronous in-region replication is the best data currency solution, but it provides poor protection against a regional disaster. Asynchronous out-of-region replication is the best regional disaster protection, but the I/O-consistent copy at

Data Center Strategy	1DC	2DC		3DC	
Replication configuration	Onsite	Synchronous near	Asynchronous far	Cascade	Multitarget
Primary site failure/failover					
Speed of recovery (RTO)	Bad	Better	Good	Better	Better
Data currency (RPO)	Bad	Better	Good	Better	Better
Regional disaster (RTO)	Bad	Bad	Good	Good	Good
Protection after failure outside primary site	N/A	Bad	Bad	Depends	Better

Table 2. Benefits Comparison Summary

the remote site may not be 100 percent current (RPO > 0).

- The 3DC strategies can provide “zero data loss” at any distance after a primary site failure, with different levels of resilience in case of additional failures. The most familiar approach, 3DC cascade, provides superior recovery speed and data currency after primary site failures, along with good protection from regional disasters. However, a failure at the intermediate site can leave the primary site without ongoing disaster recovery protection. 3DC multitarget offers all the benefits of 3DC cascade. It also preserves ongoing replication and recovery capability if the intermediate data center or the remote data center fails while the production site survives.

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 are we the vendor of choice? Developing technology that replicates data in near real time is just the starting point. It is experience and a true culture of partnership that make the difference when the life of an organization is at stake.

This Hitachi Data Systems Solutions Brief has delved into the particulars of a Hitachi synchronous solution for your business continuity needs. Prior to choosing a solution

you may need a full understanding your business continuity risks and a business case for solving them. Hitachi Data Systems Global Solutions Services can help you choose from the variety of software and services to simplify disaster recovery, minimize downtime, speed up recovery, and protect your business’ information assets while maximizing the use of resources and personnel.

Specifically, GSS 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

Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: 1 408 970 1000 info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 info.uk@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., and/or its affiliates in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

TrueCopy is a registered trademark of Hitachi Data Systems Corporation.

IBM, ESCON, FICON, and S/390 are registered trademarks of International Business Machines Corporation.

All other trademarks, service marks, and company names are the property of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to http://www.hds.com/products/systems/licenses_warranties.html or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.