HITACHI
Inspire the Next

Hitachi Universal Replicator Software Architecture Overview

A White Paper

By Christophe Bertrand, Bill Martin, Roselinda Schulman, and Claus Mikkelsen

March 2007

Hitachi Data Systems

Executive Summary

Business continuity requirements are changing. When establishing and updating their business continuity plans, enterprises must respond to new business drivers, such as around-the-clock operations, higher service-level expectations, closer regulatory scrutiny with the emergence of stringent out-of-region data protection requirements, and increased sensitivity to loss of data and information assets.

The challenge is to reduce risk and increase business resilience, while also reducing costs and increasing efficiency.

Hitachi Universal Replicator for TagmaStore® Universal Storage Platform or Network Storage Controller is an advanced technology for asynchronously replicating data hosted on Universal Storage Platform or Network Storage Controller—or on externally attached storage systems from Hitachi Data Systems and other vendors. Universal Replicator software evolves and improves industry-leading Hitachi asynchronous remote copy solutions.

Universal Replicator software uses disk-based journaling and a pull-style replication engine to reduce resource consumption and costs, while increasing performance and operational resilience. These features make it more flexible, efficient, and cost-effective than traditional replication methods.

Universal Replicator software also makes out-of-region replication more cost-effective, and it enables wider adoption of advanced three data center (3DC) business continuity strategies that provide "no data loss" replication and business continuity protection over extended distances.

With the recent enhancement of 4 by 4 architecture for Universal Replicator software, 2DC multiple controller consistency groups are now supported and, in the near future, 3DC multicontroller consistency groups will also be supported. This means that Hitachi Data Systems now offers a very compelling alternative to IBM® XRC® for large IBM z/OS® customers.

With these advanced technology capabilities, Universal Replicator software with the Universal Storage Platform or Network Storage Controller enables enterprises to improve business continuity, reduce costs and complexity, and increase business resilience.

Contents

Hitachi Universal Replicator Software Architecture Overview

A White Paper

By Christophe Bertrand, Bill Martin, Roselinda Schulman, and Claus Mikkelsen

Introduction

Business continuity requirements are changing. When establishing and updating their business continuity plans, enterprises must respond to new business drivers, such as around-the-clock operations, higher service-level expectations, closer regulatory scrutiny with the emergence of stringent out-of region data protection requirements, and increased sensitivity to loss of data and information assets.

The challenge is to reduce risk and increase business resilience, while also reducing costs and increasing efficiency.

Risk Reduction and Business Resilience

Many organizations—along with stakeholders and regulatory authorities—are focusing on operational risk assessment, and one key area of focus is information systems and data protection. Companies typically need to enhance system availability, and develop or update a robust disaster recovery strategy.

Compliance Mandates. In many industries and geographies, government regulations require companies to have effective business continuity plans that will enable them to protect information assets and maintain their service capabilities in spite of a local or regional disaster. The most commonly regulated industries likely to adopt out-of-region strategies worldwide include telecom, transportation, banking and other financial services, government, utilities, health care, and e-commerce.

Business Resilience. In many cases, IT is still falling short in providing organizations with a fully resilient environment. One key requirement is data replication for business continuity—with guaranteed integrity and consistency.

Wider Scope of Protection. While some organizations can define a very narrow scope for data protection and replication, many demanding environments will need to replicate essentially all of their actively changing data sets. And the volume of data is growing exponentially.

Better and More Flexible RTO and RPO. For business continuity planning, two of the key metrics are recovery time objective (RTO) and recovery point objective (RPO). The RTO measures how long it takes to resume essential operations: how long it takes the organization to get back on its feet. The most demanding environments may have a hot site that is fully redundant—with the needed equipment, applications, people, and processes—using technologies such as server clustering to minimize RTO in case of a primary site failure. The RPO is a measure of data currency: how far behind the organization can afford to be when resuming operations after an interruption. Depending on the data location and the nature of the interruption, the copy maintained at the recovery site could be out of date by minutes or even hours.

Many organizations are seeking to improve RTO and RPO, while also protecting their operations and data from a wider range of threats and disaster scenarios. This leads into the area of operational risk, which is increasingly being looked at versus the more traditional RTO and RPO metrics. The idea is that the organization needs to assess the impact of risk to their total environment and look at mitigation strategies based on what will have the biggest impact and return on investment in the environment.

Cost Reduction and Efficiency

Cost pressures continue to constrain IT organizations in their business continuity initiatives, so organizations must find ways to meet requirements while reducing costs and increasing efficiency.

Migration and Consolidation. Organizations may need to replicate data to and from multiple data centers for testing, upgrades, or load balancing to improve IT operations. Consolidation of computing assets may improve efficiency and reduce costs, but it is also akin to putting more eggs in one basket. The consolidated platform or data center requires a higher level of data protection and disaster resilience.

Resource Utilization. IT organizations are expected to support higher levels of application performance and responsiveness, while also improving the level of protection and resilience. Business continuity solutions can use remote replication to improve resilience and rapid recovery capabilities, but the replication solutions themselves may consume scarce resources and introduce management complexity.

Management Complexity. IT organizations must manage complex and heterogeneous server and storage infrastructures, and this makes business continuity planning and testing more difficult to manage. In the absence of centralized and universal management tools, organizations must support local and remote replication using a combination of disjointed replication solutions. This makes business continuity testing a challenge, and it increases the risks and costs of failure in a real recovery situation.

This white paper takes a closer look at the strengths and issues of existing approaches to remote replication and business continuance, and introduces a new technology that helps IT organizations meet their data movement and protection needs more simply and efficiently.

Remote Replication: Strengths and Issues

Data replication is widely used for business continuance, to protect businesses from interruption of critical applications and business processes. Organizations have tried a variety of replication approaches based on server, storage, and network capabilities. Each has its strengths and issues, but storage-based replication is generally considered the best approach for replication of critical data, particularly in heterogeneous server and application environments.

Historically, storage systems were not designed for replication. However, replication capabilities have been added to storage systems so successfully that storage-based replication is widely used today. It provides the best performance for demanding applications, and it has a proven track record for data protection and disaster recovery.

Strengths

For business continuity, disk-based remote replication has greatly improved RTO and RPO, as well as risk reduction, compared with earlier tape-based methods. Customers and analysts agree that storage-based replication is the platform of choice for replication of critical data; it is generally superior to server-based or fabric-based replication. As one Hitachi Data Systems customer remarked, “I use (a server-based replication product) for some of the data, but for my key applications I use Hitachi replication products.” Hitachi Data Systems customers are quick to note that they have never lost data with their Hitachi storage and Hitachi replication software

Local Disasters. Synchronous replication to an in-region hot site is now widely used for business continuity, to protect businesses from interruption of critical applications and data access at a production data center. However, synchronous replication is limited to relatively short distances—typically less than 50 miles—and is suitable for replication to in-region recovery sites. This approach does not protect from regional disasters that may affect both the production site and the in-region recovery site.

Regional Disasters. Use of asynchronous replication is growing as companies respond to operational risks and compliance requirements. Asynchronous remote replication enables organizations to maintain very current data copies at out-of-region recovery sites. Hitachi is the leader in asynchronous remote replication, based on the success of Hitachi TrueCopy® Asynchronous software.

Figure 1. RPO Flexibility from Near Synchronous to Hours

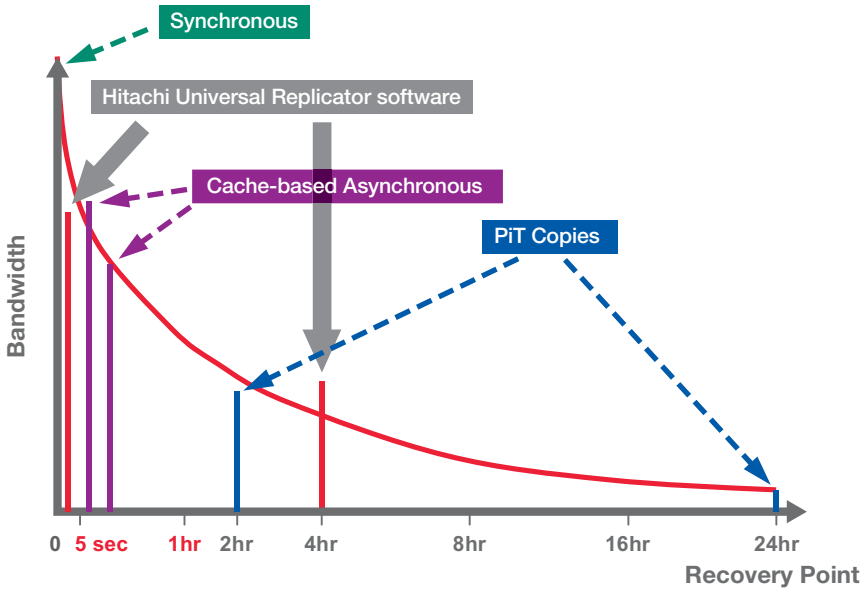

Universal Replicator software allows organizations not just to achieve near synchronous RPO, but also to have the flexibility for elongated RPOs of minutes to hours if that meets their requirements.

Hitachi Universal Replicator bridges an RPO gap that typical cache-based asynchronous technologies can't fill. In typical cache-based solutions, while RPOs of seconds can be achieved, given enough bandwidth and resources, not enough bandwidth and resources causes the replication to suspend very quickly. Point-in-time copies serve organizations that require RPOs of hours, as they work best with long cycle times to be able to see a reduction in bandwidth requirements. Universal Replicator software allows organizations not only to achieve near synchronous RPO, but also to have the flexibility for elongated RPOs of minutes to hours if that meets their requirements. It also avoids the problem of the replication suspending due to lack of network resources.

However, as requirements have evolved over time, and with soaring amounts of replicated data, organizations have pushed the limits of current asynchronous replication technology. Furthermore, regulations worldwide are endorsing (if not requiring) out-of-region replication for critical industries such as banking and securities trading. A number of issues have arisen that indicate the need for an improved solution. Hitachi Data Systems now delivers that improved solution with Universal Replicator software on the Hitachi TagmaStore® Universal Storage Platform or Network Storage Controller.

Issues

For all their benefits, current remote replication technologies are not perfect. Enterprise IT managers have identified common issues with remote replication solutions from all vendors. Here are a few typical comments:

- "This stuff uses up all my cache."
- "I know replication is crucial, but I'd like to focus system resources on my applications' transactions."
- "Bandwidth requirements are killing me."
- "I wish I had only one tool for replicating all my data."
- "How can I implement out-of-region disaster recovery without losing data?"

The issues generally involve resource utilization, network bandwidth, management complexity, and the trade-offs between alternative replication strategies.

Resource Utilization. Current remote copy technologies can consume tremendous amounts of resources. In server-based approaches, replication consumes server memory to buffer the changed data during replication. In storage-based solutions, replication uses part of the storage system cache to capture changes and transmit them to the other side. It also uses processing cycles on the storage systems—primarily at the originating (production) data center. These resources are, in effect, taken away from production applications. The result is lower application performance, or the increased cost of adding resources to maintain required performance and throughput.

Data growth makes these problems worse as resource consumption increases at the production site. Organizations are looking for a better solution. Universal Replicator software, with its replication engine, now returns these resources to where they belong: the application.

Network Problems. Replication network bandwidth is another key resource consumed by remote replication technologies. Bandwidth is expensive, and replication requires more of it. Also, since momentary link failures or network congestion can cause cache-based replication to fail, most organizations plan and provision bandwidth to meet the maximum peak-load requirement. If they don't, the replication can fail, forcing organizations to perform difficult processes to recover and reestablish replication. This can be painful and disruptive.

Replication Management Complexity. Simplification of management complexity is another requirement from businesses. Because of the number of storage vendors and products in a heterogeneous storage environment, many organizations must learn and use different replication tools for the various systems. Many are asking for a single tool for managing multiple storage-based replication platforms.

Replication Strategy Choices. In-region replication using synchronous remote copy technology can provide recovery with little or no data loss, but it is vulnerable to regional disasters. Out-of-region replication using asynchronous replication protects against regional disasters, but it may lose some recent transactions. Multi–data center configurations can combine both approaches. However, existing solutions require storage for multiple copies of the data, as well as complex management and scripting.

Universal Replicator software builds on existing Hitachi strengths in remote replication technology, and introduces advanced technology to address these emerging business issues with a more efficient and cost-effective solution. The next section describes Universal Replicator features and benefits, and later sections outline improved business continuity approaches that Universal Replicator software makes technically possible and economically feasible.

Universal Replicator Software: Overview

Hitachi Universal Replicator software for TagmaStore Universal Storage Platform and Network Storage Controller is an advanced technology for asynchronously replicating data hosted on the Universal Storage Platform, Network Storage Controller, or externally attached storage systems from Hitachi Data Systems and other vendors.[1]

Universal Replicator software uses disk-based journaling and an optimized replication engine to reduce resource consumption and costs, while increasing performance and operational resilience. These features make it more efficient and cost-effective than traditional replication methods.

Universal Replicator software also enables more resilient and cost-effective business continuity solutions, especially when combined with Hitachi TrueCopy Synchronous software in three data center (3DC) replication and recovery strategies.

The following sections describe two key technical innovations that underlie the strengths and benefits of Universal Replicator software:

- Performance-optimized, disk-based journaling
- Pull-style replication engine

Journal-based Replication

Figure 2 illustrates Universal Replicator software's application of performance-optimized, disk-based journaling. In this illustration, online transaction processing (OLTP) data volumes at the primary data center are being replicated to a second Universal Storage Platform at the remote data center. (This example might also be configured with the Network Storage Controller.)

Figure 2. Hitachi Universal Replicator Software: Asynchronous Remote Copy

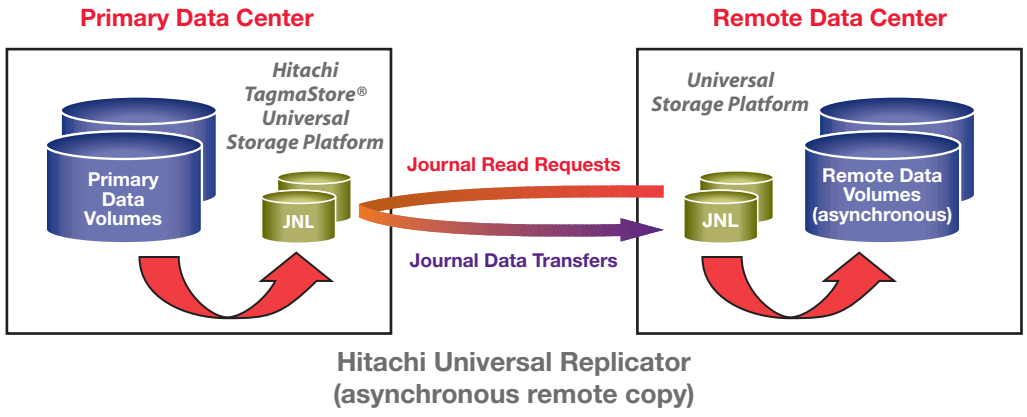

Universal Replicator software uses disk-based journal volumes and a pull-style replication engine to move data from the primary site to the remote site.

[1] For a discussion of the Universal Storage Platform/Network Storage Controller and business continuity applications of its capabilities, including Universal Replicator software, see *Business Continuity and the Hitachi TagmaStore® Universal Storage Platform and Network Storage Controller*, Hitachi Data Systems white paper WHP-163. Please check with Hitachi Data Systems on supported platforms and configurations.

When collecting the data to be replicated, the primary Universal Storage Platform writes the designated records to cache and then destages them to a special set of journal volumes. The remote storage system can then read the records from cache or the journal volumes, pulling them across the communication link as described in the next section.

By writing the records to journal disks instead of keeping them in cache, Universal Replicator software overcomes the limitations of earlier asynchronous replication methods. Writes to the journal are cached for application performance reasons, and if there is available bandwidth a read journal request can be satisfied from cache; however, the data is quickly de-staged to disk to minimize cache usage. The journal disks are specially architected and optimized for maximum performance.

Guaranteed Data Integrity. In addition to the records being replicated, the journal contains metadata for each record to ensure the integrity and consistency of the replication process. Each transmitted record set includes time stamp (mainframe only) and sequence number information, enabling the replication engine to verify that all the records are received at the remote site, and to arrange them in the correct write order for storage. These processes build on the proven algorithms of TrueCopy software. The journaling and replication processes also support consistency across multiple volumes. Thus, Universal Replicator software builds on the proven integrity of existing Hitachi replication solutions, while introducing innovative journaling technology to improve performance and efficiency.

Pull-based Replication

In addition to disk-based journaling, Universal Replicator software introduces pull-style replication. The primary storage system does not push data across the replication link unless requested by the remote site. The replication engine on the remote system receives the data from the primary system's journal volume and writes it to the journal volume at the receiving site. The replication engine then applies the journaled writes to the remote data volumes, using metadata and consistency algorithms to ensure data integrity.

In the default configuration, Universal Replicator software considers replication complete when the data is received in mirrored system cache at the remote system, written to the journal disk, and applied to the remote data volumes. Since the processes that control asynchronous replication are located on the remote system, fewer resources are consumed on the primary storage system, thus improving production application performance. In effect, Universal Replicator software restores primary site storage to its intended role as a transaction processing resource, not a replication engine.

Figure 3. Single Consistency over Multiple Controllers

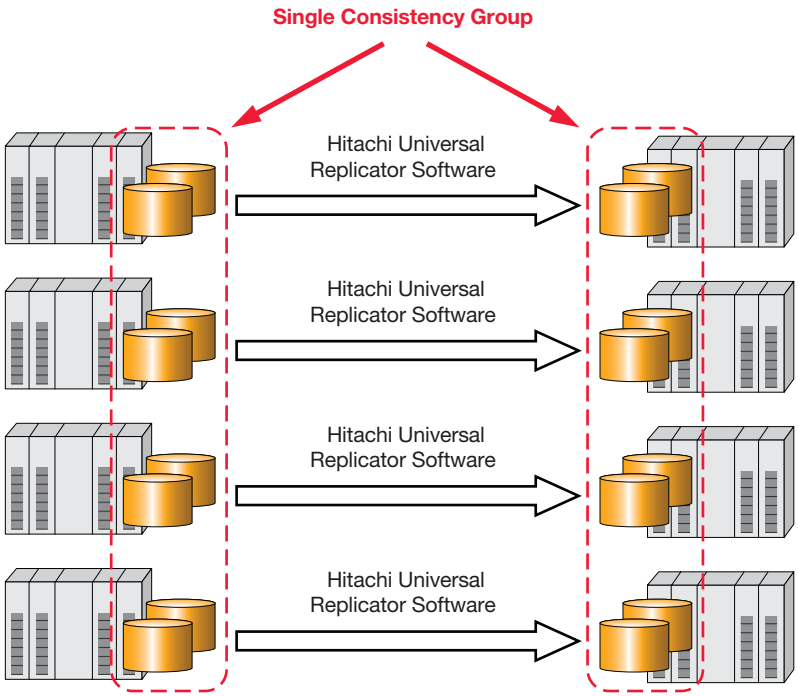

Universal Replicator software provides the capability to have one consistency group spread across multiple controllers in IBM® z/OS® environments.

Because Universal Replicator software provides the capability to have one consistency group spread across multiple controllers in IBM® z/OS® environments, large users gain the ability to have up to 64K volumes be consistent to a single point in time, while maintaining a near synchronous RPO. This capability provides a compelling alternative for current IBM XRC® customers: it lowers their TCO, as this solution does not require any host software at the remote site. In addition, managing many thousands of replicated volumes is accomplished easily using Hitachi Business Continuity Manager management software, when compared to issuing thousands of z/OS commands.

Universal Replicator Software: Benefits

Universal Replicator software significantly reduces resource consumption. It sets a new standard for data protection and Application Optimized Storage™ solutions from Hitachi Data Systems by enabling replication over any distance for business continuity or improved IT operations, without the need for redundant servers or replication appliances.

Universal Replicator software, running on the Universal Storage Platform, delivers a number of important benefits:

- Resource optimization
- Mitigation of network problems
- Simplified replication management
- Advanced business continuity configurations

Resource Optimization

Universal Replicator software significantly reduces consumption of production-site resources, such as storage system cache and processing power. Essentially, it liberates resources and improves performance.

By using local disk-based journaling and a pull-style remote replication engine, Universal Replicator software releases critical resources that are consumed by other asynchronous replication approaches at the primary site, such as disk array cache in storage-based solutions, or server memory in host-based software approaches. Universal Replicator software improves cache utilization, lowering costs and improving performance of production transaction applications. It also maximizes the use of bandwidth by better handling the variations of the replication network resources, enabling enterprises to manage bandwidth cost and data currency more flexibly and intelligently.

The pull-style replication engine also contributes to resource optimization. It controls the replication process from the secondary system and frees up valuable production resources on the primary system. It also provides the ability to read and transfer larger blocks of data, therefore maximizing the bandwidth utilization.

Mitigation of Network Problems

In traditional asynchronous replication, typical issues include temporary communication problems, such as communication link failure or insufficient bandwidth for peak-load requirements. These conditions can cause cache-based “push” replication methods to fail. When this happens, traditional replication solutions suspend the replication process and go into bitmap mode, noting changed tracks in a bitmap for future resynchronization. Recovery typically involves a destructive process such as rewriting all the changed tracks, with possible loss of data consistency for ordered writes.

Increased Resilience. In contrast, Universal Replicator software logs every change to the journal disk at the primary site, including the metadata needed to apply the changes consistently. Should the replication link between sites fail, Universal Replicator software keeps logging changes in the local journal so that they can be transmitted later, without interruption to the protection process or the application. The journal data is simply transferred after the network link failure or bandwidth limitation is corrected, with no loss of consistency. The RPO may be extended a bit during temporary network failures or congestion, but the asynchronous replication process does not fail, and the catch-up process is simple and automatic. Data consistency is preserved, and the user can maintain a more current RPO during outage situations, compared with existing replication methodologies. This capability is also critical in providing trans-oceanic replication services, where, in some cases, the network bandwidth may not optimum, but RPO requirements are still stringent.

Improved Bandwidth Costs and RPO. With Universal Replicator software, the remote storage system pulls data from the primary journal volumes—over the data replication network—as fast as the bandwidth allows, adjusting to available network conditions. If available bandwidth does not support optimal replication (for example, during peak-load spikes in transaction volume), the primary journal volumes simply buffer the data on disk until more bandwidth becomes available.

This technology approach gives enterprises a more flexible way to manage network bandwidth. They can improve bandwidth utilization and lower their communication costs by provisioning bandwidth for the average traffic level—not for peak usage—and sizing the journal volumes to handle the peak load. Enterprises can manage the currency of their data more intelligently, in relation to infrastructure and communication costs. The overall result is better replication with improved RPO, lower TCO over a typical three-year planning period, and greater resilience (please see “Hitachi's Storage Management Vision Expands Reach and Range,” an ITCentrix, Inc., analyst report at www.hds.com/tools/analyst_reports/).

Simplified Replication Management

Universal Replicator software will simplify replication management by leveraging Universal Storage Platform or Network Storage Controller support for externally attached storage from multiple vendors. Universal Replicator features will ultimately support real-time replication for all internal and externally connected disk storage volumes, making it a very powerful solution for remote copy operations.[2] Thus, Universal Storage Platform/Network Storage Controller with Universal Replicator software will support the data replication and movement requirements of disparate applications and business processes, managed from a single pane of glass. It will enable a consistent replication approach and a common set of procedures, tools, and skills to manage diverse, heterogeneous storage and application environments.

Advanced Business Continuity Configurations

Universal Replicator software, in combination with TrueCopy Synchronous software, supports improved performance and cost reduction in multisite business continuity configurations. Advanced 3DC configurations intelligently combine synchronous and asynchronous replication to provide a "no data loss" replication capability at any distance.

Advanced configurations can also use fewer intermediate copies than alternative approaches, while placing the copies on less expensive storage and reducing resource consumption at the primary site. Thus, Universal Replicator software makes the business continuity benefits of 3DC configurations practical and affordable for a wider range of organizations, applications, and business processes.

Remote Replication Configurations for Business Continuity

Universal Replicator software improves remote copy solutions, making out-of-region replication more efficient and cost-effective, and thus enabling its wider use.[3] More and more organizations will find that out-of-region replication is both necessary and feasible, so this section provides an overview of the characteristics and benefits of several alternative approaches. It includes a description of several two data center (2DC) and 3DC configurations, followed by a technical comparison and a summary of the relative benefits of each approach. The next section then shows how Universal Replicator software improves these alternatives—particularly the 3DC configurations—and makes them more efficient, affordable, and cost-effective for a wider range of organizations.

Figure 4 illustrates several replication configurations involving two or three data centers.

[2] Customers should check with their Hitachi Data Systems representatives regarding the scheduled availability of Universal Replicator software support for externally connected disk systems from Hitachi Data Systems and other vendors.

[3] To support business continuity objectives, many enterprises send copies of production data to remote sites for use in disaster recovery processes. Early disaster recovery approaches used local replication or backup software to create copies on removable media, and then shipped the media to an offsite recovery provider. Such approaches support RPOs and RTOs of several days, and may be appropriate for less critical applications and some fixed content. However, many enterprises would suffer substantial economic damage if critical business applications were shut down for several days after a data center disaster. Electronic replication strategies involving interconnected data centers can provide faster recovery, less risk of data loss, and protection from a wider range of disruptive events.

Figure 4. Remote Replication Configurations: Two Data Center and Three Data Center

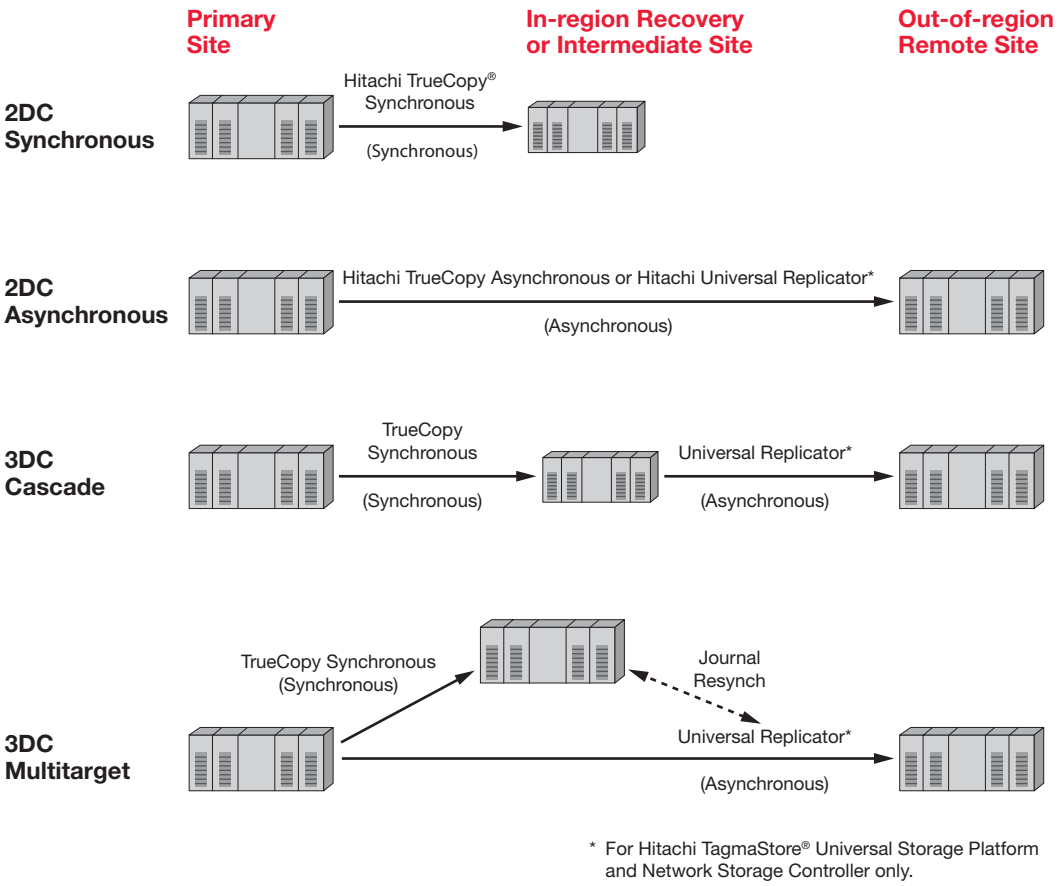

Remote replication strategies may involve two data centers or three data centers and use in-region or out-of-region replication—or both.

2DC Strategies: These involve replication from the primary data center to a second data center. The second data center might be an in-region recovery site, such as a hot site serving a campus-level or metro-level server cluster. Or, the second site might be an out-of-region recovery site, located hundreds of miles away.

2DC Synchronous: Synchronous replication can provide a very fast recovery time (low RTO) and good data currency (low RPO) when recovering at an in-region hot site. But in-region replication does not protect from regional disasters (such as power grid failures, earthquakes, and floods) that can impact critical business operations and recovery capabilities at both primary and secondary data centers.

2DC Asynchronous: To reduce business risks from regional disasters, many organizations have established remote data replication to out-of-region data centers. Because synchronous replication over long distances can impact application performance, out-of-region replication typically uses asynchronous replication.[4] As we have already shown, Universal Replicator software makes 2DC asynchronous replication more resilient in response to network issues, and more cost-effective in terms of resource utilization.

[4] *Synchronous replication* requires the application to wait for the remote site to confirm each write operation before sending the next write operation. As the replication distance increases, the time lag between synchronous write operations gets longer and longer, and these delays become intolerable for high-volume, write-intensive transaction processing applications. *Asynchronous replication* allows the production application to continue, and it keeps track of the pending writes until they are complete. A well-designed asynchronous replication solution maintains an I/O-consistent copy at the remote site. It may deliver a somewhat lengthened RPO and RTO, but it does protect the enterprise from major data loss in case of a regional disaster.

3DC Strategies: These provide the best of both worlds: fast recovery and data currency for local site failures, combined with good protection from regional disasters. Figure 5 illustrates several configurations that we will introduce in general terms and then describe in more detail.

3DC Cascade: Synchronous replication supports a current copy of production data at an in-region site that does provide processing capabilities for recovery. The in-region site also forwards the data to a remote recovery site, using asynchronous replication. 3DC cascade is the traditional 3DC replication configuration, and, of course, it benefits from Universal Replicator software's more resilient and cost-effective approach to asynchronous replication.

3DC Multitarget: Synchronous replication supports a current copy of production data at an in-region site that provides processing capabilities for recovery. A separate asynchronous replication process copies data from the primary site to the out-of-region recovery site, over a separate replication network. Multitarget will also support the ability to have a journal at the intermediate site, which will allow replication and disaster recovery to continue asynchronously from the metropolitan site to the remote site after a primary site failure or switchover, without having to undertake a full copy. However, it only requires the difference data, between the synchronous and asynchronous copies, to be moved. Until now, 3DC multitarget solutions have been deployed only in the most critical business continuity environments, However, Universal Replicator software makes 3DC multitarget configurations feasible and practical for a wider range of organizations and applications (see Figure 5).

Figure 5. 3DC 4 by 4 by 4

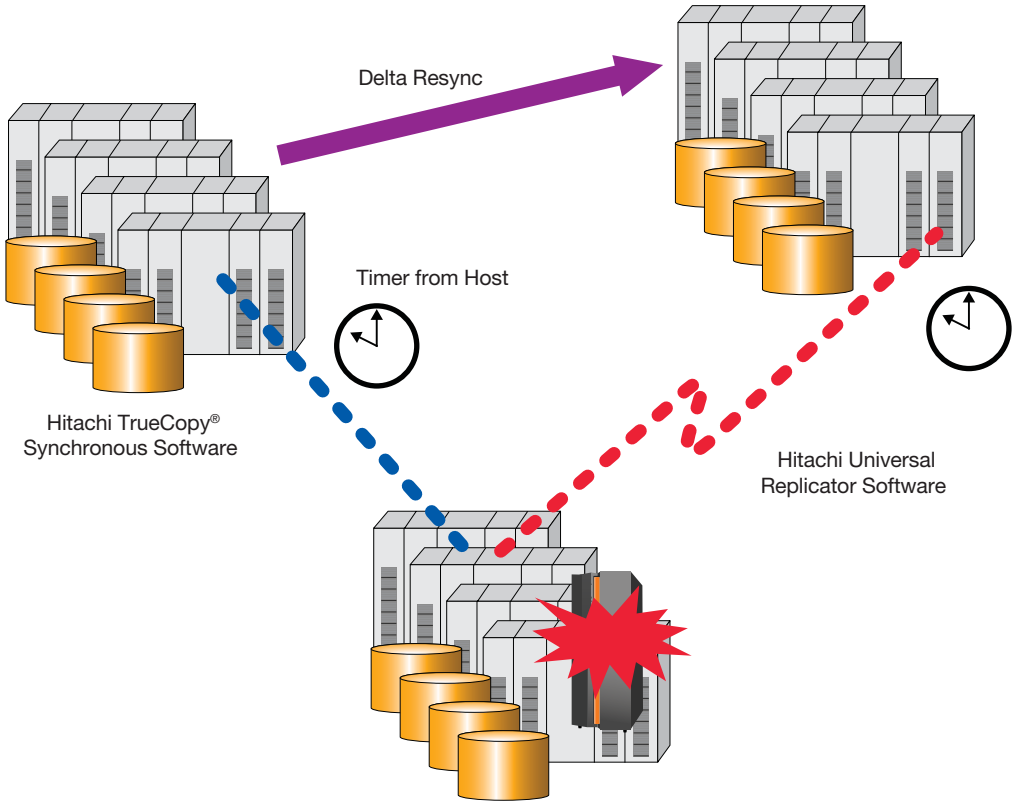

Universal Replicator software makes 3DC multitarget configurations feasible and practical for a wider range of organizations and applications.

This technology extends the replication capability for very large z/OS environments that have a requirement for greater protection and flexibility, by providing for 64K volumes in a consistency group for both the synchronous and asynchronous legs of a multitarget solution. This provides superior capabilities for organizations that may be running XRC today but cannot currently take advantage of a 3DC environment. All the features of 3DC multitarget are supported including the ability to do a catch-up between the two remote sites in the event of a primary site failure.

Technical Comparison of 2DC and 3DC Strategies

This section presents a more detailed comparison of the 2DC and 3DC replication configurations introduced above. Table 1 shows a number of quantitative parameters of each 2DC and 3DC configuration. For comparison, the traditional tape-based disaster recovery approach as the "one data center" (1DC) configuration is included.[5] The RPO and RTO ranges in Table 1 are Hitachi Data Systems estimates based on customer experience in real-world enterprise application environments.

Alternative replication strategies provide different levels of protection in terms of data currency and recovery time, for various failure or disaster recovery scenarios.

Table 1. Technical Comparison: 2DC and 3DC Replication Configurations

Data Center Strategy	*1DC**	*2DC*		*3DC*	
Remote Replication Configuration	Local	Synchronous with Failover	Asynchronous Out of Region	Cascade	Multitarget
Technology					
Synchronous Replication	N/A	TrueCopy Synchronous software	N/A	TrueCopy Synchronous software	TrueCopy Synchronous software
Asynchronous Replication	N/A	N/A	TrueCopy Asynchronous software, Universal Replicator software	Universal Replicator software	Universal Replicator software
Characteristics					
Recovery Locations Supported	0	1	1	2	2
In-region Site Failover	No	Yes	No	Yes	Yes
Out-of-region Site Failover	No	No	Yes	Yes	Yes
"No Data Loss" Available	No	Yes	No	Yes	Yes
Minimum # of Logical Full-volume Copies **	1	2	2	3	3
Primary Site Failure					
RPO	24–168 Hours	0–2	0–5	0–2	0–2

[5] Note on 1DC strategy: Organizations with one data center need to make other arrangements for disaster recovery, such as a warm site with capability to recover data from tapes and restart applications. Estimates in this table assume a traditional disaster recovery strategy with an offsite recovery site provider, and that current tape backups exist (including journal tapes if needed). This is the best-case scenario.

Data Center Strategy	1DC*	2DC		3DC	
		Minutes	Minutes	Minutes	Minutes

Table 1. Technical Comparison: 2DC and 3DC Replication Configurations *(Continued)*

Data Center Strategy	1DC*	2DC		3DC	
RTO***		Minutes	Minutes	Minutes	Minutes
	48–168	5–60	1-8	5–60	5–60
	Hours	Minutes	Hours	Minutes	Minutes
Regional Disaster					
RPO	24–168	24–168	0–5	0–5	0–5
RTO	Hours	Hours	Minutes	Minutes	Minutes
	48–168	48–168	1–8	1–8	1–8
	Hours	Hours	Hours	Hours	Hours
Impact of Single Failure Outside Primary Site					
Recovery/Intermediate Site Failure (primary still up)	N/A	No Disaster Recovery	N/A	No Disaster Recovery	OK
Remote Out-of-region Site Failure (primary still up)	N/A	N/A	No Disaster Recovery	OK	OK
Link Failures Only—Primary to In-region	N/A	No Disaster Recovery	N/A	No Disaster Recovery	OK
Link Failures Only—Primary to Out-of-region Remote Link Failures Only—Intermediate to Remote	N/A	N/A	No Disaster Recovery	N/A	OK
	N/A	N/A	N/A	OK	N/A

* 1DC strategy represents a traditional tape-based disaster recovery approach.

** Best Practices recommend at least an additional full volume copy of the data be available at the recovery site.

*** RTO with automated failover. RTO assumes people, processes, and systems are in place.

Technology. The chart shows that synchronous remote replication is supported by TrueCopy Synchronous software, and asynchronous replication is supported by TrueCopy Asynchronous software (on all Hitachi Lightning 9900™ Series and Lightning 9900 V Series storage systems) and by Universal Replicator software (on the TagmaStore Universal Storage Platform and Network Storage Controller only).

Characteristics. The 2DC configurations each support one recovery location: either an in-region recovery site or an out-of-region remote site, as illustrated in Figure 2. The 3DC cascade and multitarget configurations both support two recovery locations.

The “no data loss” capability provided by synchronous replication (in 2DC synchronous, 3DC cascade, and 3DC multitarget configurations) is best thought of as excellent data currency—within zero to two minutes—rather than absolutely no transactions lost.[6]

Primary Site Failure. The data currency is excellent, with target RPO levels of zero to two minutes, for 2DC synchronous and for all of the 3DC configurations, because they can all recover from a synchronous copy at the recovery site. In contrast, 2DC asynchronous replication can lag farther behind the primary site data, with target RPO values up to five minutes or more, depending on available bandwidth and other factors. 2DC synchronous and the other two 3DC configurations provide the best RTO performance.

Regional Disaster. In the case of a regional disaster, the 2DC synchronous copy at the in-region site can be lost along with the primary-site copy, forcing the organization to recover from backup tapes—with RPO and RTO comparable to traditional 1DC approaches. In contrast, the 2DC asynchronous and the 3DC approaches provide similar RPO and RTO during a regional disaster: they all require recovery from an asynchronous copy, at the out-of-region data center.

Impact of Single Failure outside Primary Site. Even if the production site remains up and running, a failure or disruption that impacts one of the remote sites or communication links can become a problem. Such a failure can eliminate the disaster recovery capability that the primary site depends on, leading to longer outages and more data loss in the event of a subsequent failure of the primary site.

For example, the 3DC cascade configuration provides ongoing protection if the out-of-region site fails, so the chart shows this scenario as “OK.” In contrast, failure of the intermediate site (or either communication link) prevents ongoing replication to the remote site as well, leaving no active remote replication for primary-site data. Thus, the chart shows “No Disaster Recovery” for this failure scenario.

In this respect, the 3DC multitarget configuration provides a significant benefit: either of the remote data centers can fail, and the other one will continue providing disaster recovery protection for the production site. If the out-of-region site fails, the in-region site continues to provide synchronous remote copy and excellent RPO and RTO in the event of a subsequent failure at the primary site. If the in-region recovery site fails, the out-of-region site provides ongoing recovery capabilities with asynchronous replication. Multitarget will also support the ability to have a journal at the intermediate site, which will allow replication and disaster recovery to continue asynchronously from the metropolitan site to the remote site after a primary site failure or switchover.

Considering all these different disaster scenarios and impacts, 3DC multitarget provides the broadest protection and the best combination of capabilities. The price can be higher as well, and each company must assess its risks and requirements when evaluating these advanced business continuity configurations. However, as shown in the next section, Universal Replicator software makes 3DC configurations more cost-effective, thus enabling more enterprises and applications to implement improved protection.

[6] Real-world synchronous replication configurations—especially for transaction-processing databases—involve settings of “criticality” parameters that impact application performance as well as data currency. In practice, they are rarely set for “critical” write ordering. The technical details are beyond the scope of this paper. See resources in Appendix B.

Benefits Comparison

Table 2 summarizes the benefits of these 2DC and 3DC remote-replication configurations, compared to a traditional 1DC approach that relies on physical tape movement.

Table 2. Benefits Comparison Summary

Data Center Strategy	*1DC*	*2DC*		*3DC*	
Replication Configuration	On-site	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

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 can it 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 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 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, and it also preserves ongoing replication and recovery capability if the intermediate data center or the remote data center fails while the production site survives.

The final section of this paper discusses the benefits of Universal Replicator advanced technology in each of the 3DC strategies.

Universal Replicator Software: Improving Out-of-region Replication

Out-of-region replication is increasingly important for business continuity, compliance, and business resiliency. Three data center strategies combine in-region and out-of-region replication to provide the strongest protection: fast recovery and data currency for local site failures, combined with good protection from regional disasters. However, multiple data centers and data copies increase costs, so robust 3DC strategies have typically been limited to large organizations with extremely critical business continuity needs.

Universal Replicator software makes remote replication more efficient, enabling more organizations and applications to benefit from 3DC replication strategies. This section summarizes the advantages and benefits of Universal Replicator software in each type of 3DC configuration.

3DC Cascade

The most familiar 3DC configuration is 3DC cascade. Figure 6 illustrates a 3DC cascade configuration that uses TrueCopy Synchronous software to maintain a current copy of the production data at an in-region data center. As noted earlier, 3DC cascade configurations make sense when the in-region hot site provides processing capabilities for recovery.

Figure 6. 3DC Cascade with Universal Replicator Software and the Universal Storage Platform

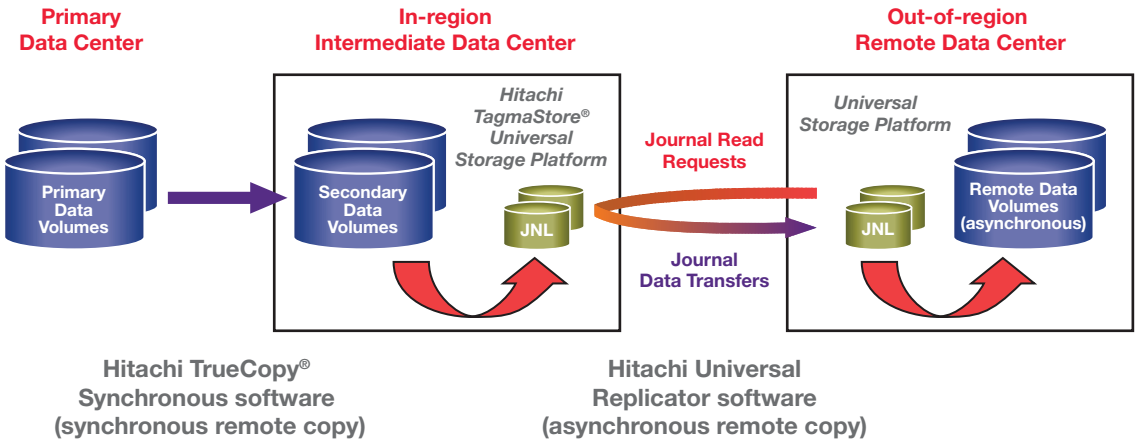

Three data center cascade operation provides full volume copies at the in-region data center.

The Universal Storage Platform at the in-region site also cascades the data to an out-of-region recovery site, using Universal Replicator software's asynchronous replication capabilities.

In comparison with other asynchronous replication technologies, Universal Replicator software does not require an additional point-in-time copy of the data volume at the intermediate site. Universal Replicator software stages the data to the journal disk, which is relatively small compared with a complete data copy. This feature saves physical disk space and reduces the cost of the 3DC configuration.

Universal Replicator software also provides a more robust asynchronous replication technology that enables organizations to reduce bandwidth costs for the remote replication network. It can ride through temporary network outages and peak-load congestion without suspending replication or requiring human intervention or

complex recovery procedures. In case of production-site failure, processing can resume at the intermediate site, using a current TrueCopy software replica of production data. The in-region hot site can also support planned failover when needed for maintenance, upgrades, or business continuity testing. Meanwhile, Universal Replicator software provides ongoing replication to the out-of-region site, maintaining robust business continuity protection. In case of a regional disaster, the out-of-region site enables rapid recovery with a consistent copy of production data.

3DC Multitarget

The multitarget strategy is similar to the 3DC cascade approach in its support of full data copies at the in-region recovery data center. The multitarget configuration is appropriate when the in-region site provides processing capabilities for recovery, and when the organization needs to maintain business continuity protection despite failure of the in-region recovery site or one part of the replication network. (See Figure 7.)

Figure 7. 3DC Multitarget with Universal Replicator Software and the Universal Storage Platform

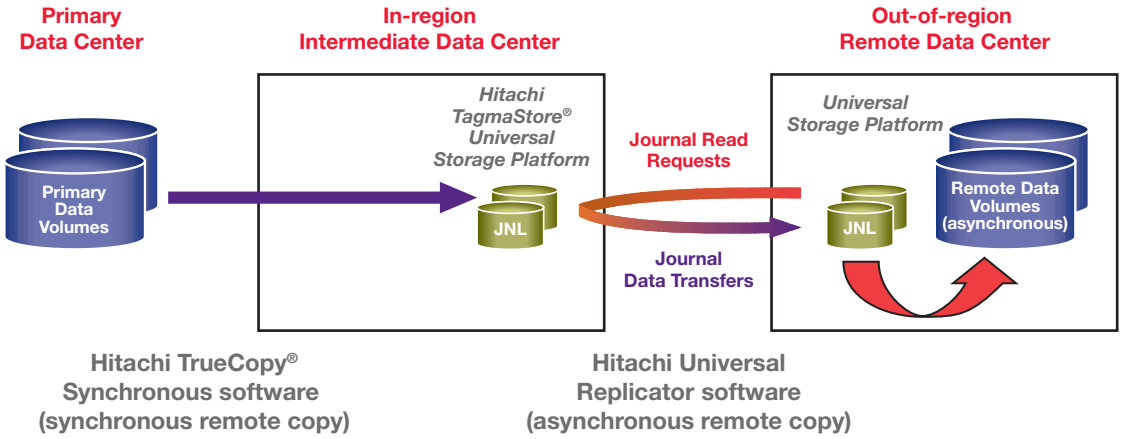

3DC multitarget provides independent replication paths to two recovery sites, and delivers the best protection against a wide range of potential outages at any site.

Figure 7 illustrates a 3DC multitarget configuration, in which data is replicated to two remote sites in parallel. TrueCopy Synchronous software maintains a current copy of the production data at an in-region recovery data center. At the same time, the Universal Storage Platform at the primary site replicates the data to an out-of-region recovery site, using Universal Replicator software's asynchronous replication across a separate replication network.

In case of production site failure, processing can resume at the in-region recovery site, using a current TrueCopy software replica of production data. In case of a regional disaster, the out-of-region data center can recover rapidly with a slightly older but fully consistent copy of production data. It is possible to establish an asynchronous replication process with Universal Replicator software from the in-region site, thus ensuring ongoing disaster recovery protection.

This is accomplished using a delta resynchronization capability, which allows only the difference data between the in-region (Intermediate) site and the remote to be moved to the remote site to achieve ongoing disaster recovery protection. This is achieved by collecting journal data at the synchronous site and comparing to the journals at the remote site after a primary site failure. The benefit of this capability is to get a fully functional disaster recovery site

up and running as quickly as possible after a primary site failure to ensure ongoing protection and compliance with regulations. This capability could also be utilized instead of cascade to achieve a synchronous copy of data to an out-of-region site, by not having servers at the intermediate site, and by executing the resync capability immediately upon failure of the primary site. This would be similar to running a cascade solution, but without the risk of having no disaster recovery if the in-region (intermediate) site fails.

Universal Replicator software makes 3DC multitarget configurations feasible and practical for a wider range of organizations and applications by reducing resource consumption and bandwidth costs and by providing a more resilient technology for out-of-region replication. Because the primary data center must support synchronous and asynchronous replication simultaneously, Universal Replicator software is particularly beneficial for a multitarget strategy because it reduces primary-site resource consumption for the asynchronous copy function.

Compared to a cascade approach, the multitarget configuration eliminates the potential of having no disaster recovery should the in-region recovery site fail.

In all of these 3DC configurations, Universal Storage Platform with Universal Replicator software enables enterprises to maintain “no data loss“ replication and out-of-region recovery capability with less cost, complexity, scripting, and overhead than previous solutions.

Conclusion

Enterprises are faced with the challenges of meeting business applications’ service levels, including increased resilience and protection from local and regional disruptions, while dealing with complex infrastructures and tight budgets. Application Optimized Storage solutions from Hitachi Data Systems enable organizations to precisely match business application requirements to Hitachi storage system attributes such as performance, availability, data value, and cost.

Within this robust framework, Hitachi Universal Replicator for TagmaStore Universal Storage Platform and Network Storage Controller delivers an advanced technology for asynchronously replicating data hosted on the Universal Storage Platform, the Network Storage Controller, or externally attached storage systems from Hitachi Data Systems and other vendors.

Universal Replicator software evolves and improves industry-leading Hitachi asynchronous remote copy solutions. It improves data currency and recovery time for minor outages. It also enables improved utilization of storage system cache and replication network bandwidth, reducing the cost of asynchronous replication between Universal Storage Platform or Network Storage Controller models.

Universal Replicator software also makes out-of-region replication more cost-effective. And it enables wider adoption of advanced 3DC business continuity strategies that provide “no data loss” replication and business continuity protection over extended distances.

With these advanced technology capabilities, Universal Replicator software and the Universal Storage Platform or Network Storage Controller enable enterprises to improve business continuity, reduce costs and complexity, and increase business resilience.

Please visit www.hds.com for customer success stories and case studies.

Appendix A: Complementary Solutions

Hitachi ShadowImage™ Heterogeneous In-System Replication software
ShadowImage Heterogeneous In-System Replication software is a natural partner for Universal Replicator software, as these products enhance each other for advanced data protection solutions.

Hitachi TrueCopy® Synchronous software
TrueCopy Synchronous software is a perfect combination with Universal Replicator software for 3DC configurations. This combination enables a business continuity solution that provides "virtually no data loss over any distance."

Hitachi TrueCopy Asynchronous software
TrueCopy Asynchronous software is the choice for platform combinations that do not include two Universal Storage Platforms (such as Hitachi Lightning 9900™ Series or Lightning 9900 V Series systems).

Hitachi Business Continuity Manager software
Business Continuity Manager software helps organizations limit or eliminate complex scripting for business continuity, and it is supported on mainframe systems.

Hitachi HiCommand® Replication Monitor software
With a single, easy-to-use display for monitoring, managing, and visualizing volume replication configurations and status, HiCommand Replication Monitor software simplifies administration of the entire suite of Hitachi replication products for open systems environments. This software interfaces with HiCommand Device Manager software and the replication management functions for Hitachi storage systems and replication software.

Storage Cluster for Microsoft Environments Service
Delivered as a turnkey solution incorporating hardware, software, and professional services, this service helps improve mission-critical availability and reliability from Microsoft Windows Server 2003 and Windows 2000 environments by leveraging Microsoft clustering over Hitachi TrueCopy Synchronous or Hitachi TrueCopy Asynchronous software.

Hitachi Data Systems Corporation

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

WHP-165-04 LWD March 2007