Enterprise Storage

Disaster Recovery Issues and Solutions

A White Paper

By Roselinda R. Schulman

December 2003

Executive Summary

Well-planned business continuity and disaster recovery solutions are critical to organizations operating in 24/7 environments. Minimal to zero disruption is the goal of today's enterprises when faced with planned or unplanned outages. This paper considers the technological advances that can serve disaster recovery issues, and the solutions, customer objectives and budgets, that contribute heavily to the architecting of the most fitting business continuity solution.

This paper does not address traditional tape backup technology or its disk declination, but focuses on replication options. It familiarizes readers with the jargon of software copy alternatives and defines currently available remote copy techniques, including network options for replication. In addition, characteristics are outlined for flexibility, interoperability, and standards compliance that are required for remote copy operations. This paper concludes with a rationale for choosing unique, enterprise-wide, simple, and elegant Point-in-Time (PiT) copies and disaster recovery solutions from Hitachi Data Systems—best-of-breed solutions that support customers as they strive to reach their business objectives.

Contents

Introduction 1

Recovery Objectives: All Data Is Not Created Equal 2

Rolling Disasters 4

Network Options for Replication 6

ESCON®-based Options 6

Fibre Channel-based Options 6

Currently Available Remote Copy Techniques 7

TrueCopy Synchronous Software 8

Semi-synchronous Option 9

Hitachi Extended Remote Copy 10

Multi-Hop and Three Data-center Copy 11

TrueCopy Asynchronous Software 12

NanoCopy Technology 13

Cascade Techniques 14

SRDF Asynchronous 14

Bandwidth Considerations 15

Summary 15

Disaster Recovery Issues and Solutions

A White Paper

By Roselinda R. Schulman

Introduction

The world has changed significantly in the past few years. Devastating terrorist acts and threats, the seemingly increased frequency of widespread power-grid disruptions, and the emergence of new government compliance regulations are placing stringent, yet necessary, data protection requirements on many organizations. Regardless of the industry, as more and more businesses operate in a 24/7 environment—especially large enterprises where global operations are the norm—an increasingly competitive edge is needed to maintain profitability and stay in business.

In the complex and challenging global environment, well-planned business continuity or proven disaster recovery practices for non-stop data availability have become critical to organizations if they are to survive any type of outage.

Most information technology-related disasters are actually logical disasters, such as data corruption, viruses, and human error, as opposed to physical disasters like fire, earthquakes, hurricanes, etc. Logical disasters occur all the time and pose a bigger threat to businesses. However, since their visibility level to the general public is low, logical disasters tend to be taken less seriously.

The real challenge lies in your organization's ability to think proactively and deploy best practices and technologies that can be leveraged to maximize business operations instead of adopting a reactive "fix-it" posture. The true test lies in the ability to prevent outages from occurring in the first place, and minimizing the effects of those incidents when they do occur. Companies today must follow the continuous business paradigm, which combines high-availability solutions with advanced disaster recovery techniques. The ultimate goal is to be able to manage both planned and unplanned situations with minimal or zero disruption.

When an unplanned event does occur, the ideal scenario is:

- Recovery will happen almost automatically with no loss of data
- Costs of the solution and resources are minimal
- Impact to the production environment is zero

While technology is moving forward at a rapid pace to reach this ideal scenario, there are many other business and technology concerns, including some significant trade-offs dictated by technology, budgets, and personnel resources that must still be considered.

Recovery Objectives: All Data Is Not Created Equal

Recovery Time Objective (RTO) and Recovery Point Objective (RPO), along with their associated costs, are important criteria when evaluating the right solution.

- **RTO** describes the time in which business functions or applications must be restored (includes time before disaster is declared and time to perform tasks).
- **RPO** describes the point in time to which data must be restored to successfully resume processing (often thought of as time between last backup and when outage occurred).

Many solutions are available depending on your organization's recovery objectives. For example, when looking at your RPO, one consideration is the cost of some data loss (typically less than five minutes, depending on the replication methodology). You may prefer having the ability to quickly perform a database restart instead of a no-data-loss option. Or, you may be more concerned with possible impacts to the production environment and an inability to recover easily (rolling disaster and database corruption) at the secondary site.

Consider the different tiers of recovery available to you (See Figure 1). While Hitachi Data Systems focuses on the higher recovery tiers in terms of its software and hardware replication offerings, we understand that not all data is created equal when it comes to disaster recovery protection (See Figure 2).

Remember that these trade-offs are business and application driven. A thorough business impact analysis is a good starting point to determine the best course of action for data protection and business continuity.

Tier #	Description of Tier	Tape Backup	Remote Disk	Real-time Disk	Remote Logging	Available System	Active System	RTO	RPO
#1	PiT	X						2-7 days	2-24 hrs
#2	Tape to Provisional Backup Site	X				X		1-3 days	2-24 hrs
#3	Disk PiT Copy, Multi-Hop			X		X		2-14 hrs	2-24 hrs
#4	Remote Logging			X	X	X		12-24 hrs	5 mins
#5	Concur ReEx (RRDF, E-Net, others)			X	X		X	1-12 hrs	5 mins
#6	Remote Copy			X			X	1-4 hrs	0-5 mins
#7	Remote Copy with Failover			X			X	5-60 mins	0-5 mins

Figure 1: Available Tiers of Recovery.

Many organizations still fall into Tier 2, although some of them will also use techniques such as remote logging of data. Typically, the higher the tier, the greater the cost. You can, however, achieve significant improvements in data currency and recovery time with the higher tiers.

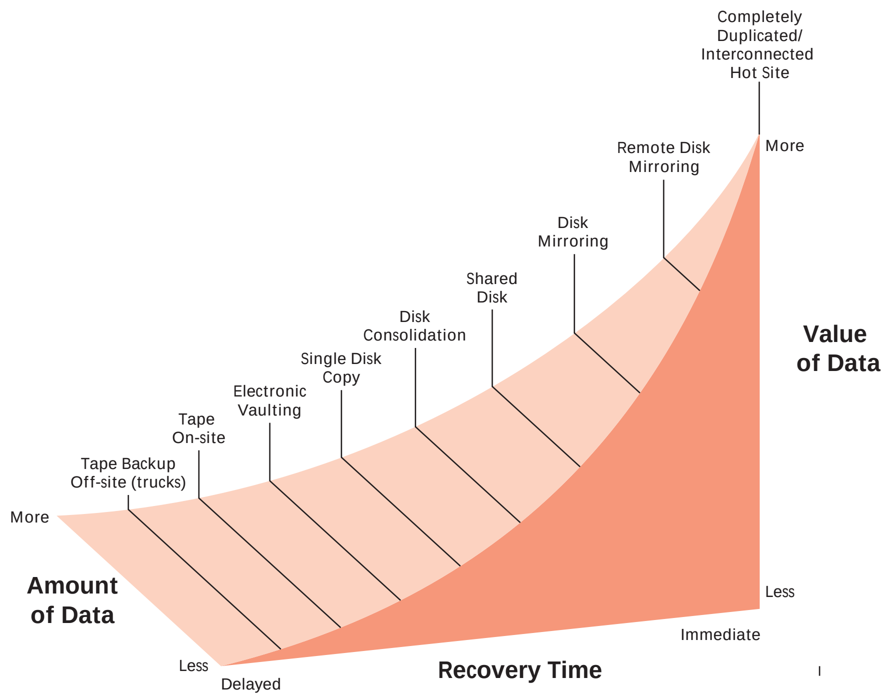

Figure 2: Data Types and Disaster Recovery. Different types of data require different levels of protection. A data audit is required to assess business criticality and cost to recover.

First, we should consider your environment. As mentioned earlier, all data is not created equal. It is likely that only a portion of a corporation's data is critical to its basic operation and that a variety of techniques could be used to secure that data, depending on the criticality of that particular business function.

Depending on how intertwined the data and applications are and the degree to which they are segregated, this can become a big undertaking. Many organizations choose not to take this approach, but rather choose to move everything.

This is a trade-off in terms of the cost of potentially having to reengineer your environment compared to the cost of using a higher tier for all data.

For companies that have a fairly local site for replication and significant bandwidth available, the option to copy everything is very attractive. However, this can change as protection from both local and regional disasters becomes necessary. While evolving network technologies help lower the cost of backing up data over significant distances, the price can still be extremely high. Therefore, it may be necessary to prioritize data or use a less current copy of data (for example, copy that is four hours old) for the second copy.

Before we start focusing on the technology, it is important to understand the differences in what has come to be known as the triumvirate of copy software product categories. The jargon of copy software alternatives becomes even more confusing when traditional backup methods are considered. Advances in technology have allowed new words and phrases, such as “real time,” “point in time” (PiT), and “snapshot” to creep into the language of enterprise-class storage. Copy products are designed to allow an enterprise to replicate, protect, and share data in dynamic new ways. Some of the terms used in copy technology are:

Remote Copy—This refers to the mirroring of data, typically in real time, to provide an I/O-consistent remote copy of that data. The purpose of remote copy is to protect the data in the event of a business interruption at a production location.

Point-in-Time (PiT) Copy—PiT Copy refers to a copy of data that is taken at a specific point in time. Ideally, this copy should be I/O-consistent. PiT copies are used in many ways, including backups and checkpoints. More recently, PiT copies have been used in architected disaster recovery solutions.

Data Duplication—This software duplicates data, as in remote copy or PiT snapshots. Data duplication differs from data migration in that with data duplication there are two copies of data at the end of the process, and with data migration there is only one.

Data Migration—This software migrates data from one storage device to another. Data migration differs from data duplication in that at the end of the process there is only one copy of data. The purpose of data migration is to reduce operational complexity and costs for storage system upgrades or equipment refurbishment. Over the last few years, many significant new technologies in both the software and hardware arenas have been brought to market. These technologies can reduce time-to-business resumption from days to hours and shorten the downtime required for backups to near zero.

When evaluating these copy technologies, there are some important points to consider. One is the consistency or integrity of the copy. While replicating data may sound simple in practice, the ability to recover from that copy of the data can be extremely complex. This is dependent not only on the technology but also on the processes that were deployed. We will discuss this in further detail as we look at available technologies with a focus on Hitachi Data Systems offerings.

Rolling Disasters

Real-time copy products are designed to maintain a duplicate image of data at a remote location so that if the primary location is lost due to a disaster, processing can continue at the second site. Although the concept of replicating updates sounds simple, surviving a disaster is actually extraordinarily difficult and challenging. To address this, three basic disaster recovery requirements should be satisfied by any disaster recovery solution:

- Surviving a rolling disaster
- Write sequencing
- Emergency restart capability following a disaster

Two discrete points in time define any catastrophic disaster: when the disaster first strikes (the beginning), and when the disaster finally completes (the end). Many seconds or even minutes may follow the beginning of a disaster. The period of time between these two events is the “rolling disaster.” Surviving a rolling disaster is the true test of any disaster recovery solution because data corruption might be occurring within this rolling disaster window (See Figure 3).

Figure 3: Rolling Disaster Window.

The real objective of any disaster recovery solution is to provide the capability to produce an image or I/O-consistent copy of data at the secondary location, as it existed at a point in time prior to the beginning of the disaster. This can be likened to the state in which data exists after a server or system crash. If update activity during the rolling disaster was also to be shadowed to the backup site, the backup copy might also be corrupted as the write order cannot always be preserved during this time. We know that the image of the shadowed data is usable at any point prior to the disaster occurring, but it cannot be guaranteed that the image is immediately usable if the potentially corrupted updates occurring during the rolling disaster are copied. In a rolling disaster, the data image may be corrupted due to “write sequencing” and “write dependency.” Write sequencing is the notion that the order or sequence of updates to the primary data structure maintains the integrity of the data. For example, the sequence in which a database and a log are updated allows the database management system (DBMS) to instantly recover the database, with data integrity, following any sudden outage (See Figure 4).

- Many applications control update sequence
- Updates applied out of sequence violate data integrity
- One transaction may involve updates to multiple volumes
- Databases have been carefully designed to protect integrity
- A DBMS has no knowledge of remote copy
- A DBMS does not control remote copy updates

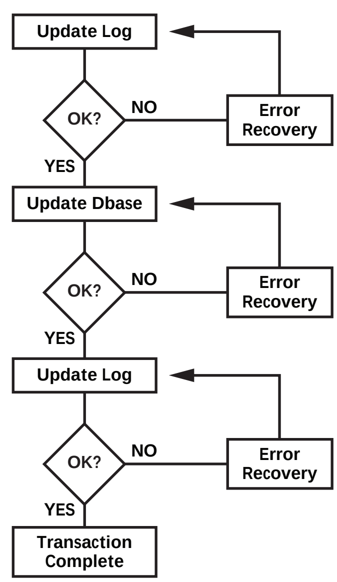

Figure 4: DBMS Crash Recovery Example.

That means a remote copy solution must be able to replicate the original sequence of updates; failure to do so will result in corrupted data at the backup site. Write dependency implies that there is a logical relationship among a series of updates, and if there is a particular update failure, the sequence of subsequent updates might change. The application controls this write sequence/dependency, but the application has no knowledge of the remote copy. There are different ways to preserve write dependency; therefore, different vendors may choose to use different methodologies. Hitachi Data Systems believes that customers should look to proven technologies, such as true synchronous remote copy products (with appropriate controls, such as freeze functions), Hitachi TrueCopy™ asynchronous remote replication software, Hitachi Extended Remote Copy (HXRC) S/390 XRC®-compatible remote replication software, IBM® XRC, IBM GDPS®, and NanoCopy™ technology, which all satisfactorily solve the write-sequencing problem.

Network Options for Replication

Before we take a look at the various remote copy technologies, we will consider some of the network options available.

ESCON®-based Options

Hardware-based replication techniques had their origins in the S/390 world and, therefore, used ESCON channels as the transport mechanism for links between the systems. Today, most customers employ Fibre Channel to transport data between systems. This is available for both open systems and mainframe environments and provides some additional benefits over ESCON-based copy.

Fibre Channel-based Options

Fibre Channel can be used in a manner similar to that employed by ESCON or direct connect using dark fibre and dense wavelength division multiplexing (DWDM) technology. In this environment,we may see a lower total cost of ownership (TCO) due to improved performance and the need for fewer links. Other benefits are gained if Fibre Channel storage area network (SAN) management expertise is already in place; the customer doesn’t have to manage yet another type of interconnection and this network option can fit into the existing SAN infrastructure environment.

When looking at Fibre Channel for extended distance copy, there are several options including replication over IP networks connecting into traditional networks, such as Point-to-Point OC3s. Newer extender technologies also allow you to connect directly through synchronous optical network technologies (SONET) networks with potentially lower overhead. When comparing the different options, you have to consider cost, ease of procurement, and the ability to control the available pipe and quality of service (QoS), so as not to compromise the real-time application by not having enough bandwidth. However, IP replication does seem to offer the ability to cover extreme distances at a lower cost.

Hitachi Data Systems and other leading IT and storage vendors were invited to participate in a transcontinental IP storage demonstration named the Promontory Project. Hitachi storage systems were connected to Fibre Channel SANs in California and New Jersey that were bridged (with gigabit IP “pipes”) with IP storage switches from Nishan Systems™ (recently acquired by McDATA®). This demonstration was undertaken to show that distance and speed are no longer barriers to mass storage architectures and solutions. Storage locations can now be implemented far enough away from primary centers so as not to share common disaster areas. TrueCopy software over Fibre Channel was tested over the gigabit IP links; basic features of the software were run without modifications over the IP backbone.

While the Promontory Project demonstration involved Nishan Systems, other vendors such as Akara™ and CNT® are currently qualified for replication over IP networks and more traditional point-to-point protocols with Fibre Channel. There are also many other vendors who are looking to get into this market as well as other emerging network replication protocols, such as iSCSI.

Currently Available Remote Copy Techniques

All of these offerings, with the exception of XRC and HXRC, are available for UNIX®, Microsoft® Windows® and S/390 operating systems, though their implementations may differ.

- Synchronous (IBM, Hitachi Data Systems, EMC®)
- SRDF™ Semi-synchronous (EMC)
- Software-based asynchronous (Hitachi Data Systems, IBM, EMC)
- XRC/HXRC (for S/390 only)
- Multi-Hop, Three Data-center Copy **(**Variations from each vendor)
- TrueCopy asynchronous (Hitachi Data Systems)
- NanoCopy technology (Hitachi Data Systems)
- Cascade Copy (typically EMC and IBM using SRDF Adaptive or Peer-to-Peer) Remote Copy (PPRC-XD)
- SRDF Asynchronous (EMC)
- Software-based and Log Replication

There are many different real-time remote copy alternatives. Implementation differs widely based on customer requirements, as previously discussed. No single alternative will satisfy every customer's objectives. In many circumstances there is little choice; for example, if a customer's backup location is hundreds of kilometers away, then only asynchronous or PiT copy technologies may be practical. As previously mentioned, it is also worth noting that not all options are created equal. The concepts of write dependency and write sequencing are not addressed in all remote copy techniques, notably semi-synchronous, PPRC-XD and adaptive copy. Other techniques, such as multi-hop, cascade, and NanoCopy, address the issue by creating consistent PiT copies on a fairly regular basis, although they all have different characteristics and cost structures.

Hitachi Data Systems has the widest range of offerings of any storage vendor, and while our competitors would argue that this makes things very confusing for the customer, that is not so. Hitachi has taken the initiative of providing both IBM-compatible and proprietary technologies to give our customers the most flexibility in choice of offerings. We use the concept of building blocks for all our software to architect the right solution, based on the customer's goals and objectives, rather than trying to fit a square peg solution into a round hole. For example, our flexibility is evident in long-distance replication. Depending on the requirements, we can offer a solution with which a customer could replicate directly to a site halfway around the world, or, if they also need a local copy, we can offer a Three Data-center solution. Other vendors can only offer cascade-style solutions, as they cannot support direct copy for write-dependent applications over distance. Multi-Hop and Three Data-center Copy are only really applicable in very specific circumstances and for a limited number of customers. EMC used to offer Multi-Hop regardless of this prior to the availability of SRDF Asynchronous. However SRDF Asynchronous is only available for the EMC Symmetrix® DMX® product line and has many limitations, so customers may still be offered Multi-Hop even if it is not the right fit for them. We will detail Multi-Hop and Three Data-center Copy further as we cover each technique.

TrueCopy Synchronous Software

TrueCopy synchronous software is one of three hardware-based, synchronous solutions from major storage vendor Hitachi; the other solutions include PPRC from IBM, and SRDF Synchronous from EMC.

In these solutions, an update to the primary data is not allowed to complete to the application until it has also been successfully secured at the secondary location. It is widely but errantly believed that with synchronous copy you are protected against a rolling disaster. In fact, if the update does not complete successfully at both locations, the action taken is dependent upon how the environment has been set up and what controls are in place. Without the correct procedures, the production environment can be impacted and/or the data at the secondary location can be corrupted (See Figure 5).

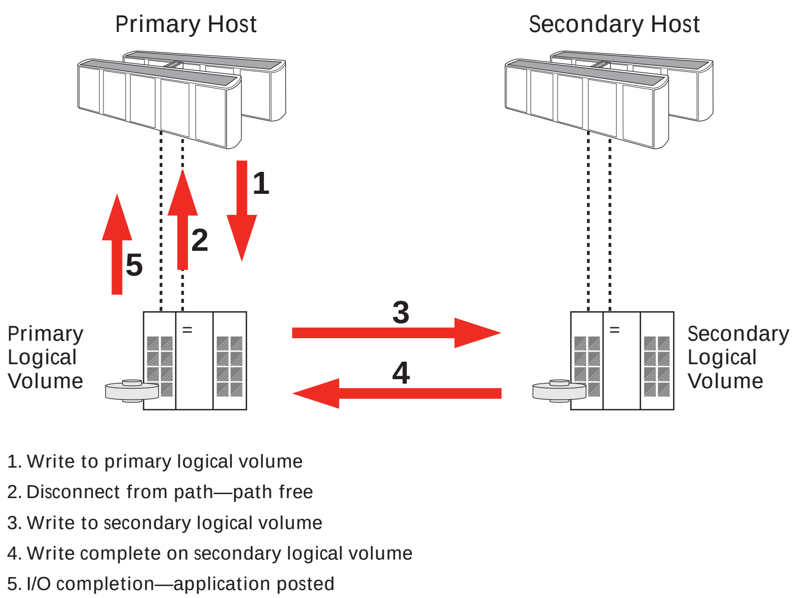

Figure 5: Synchronous Solution. The remote link is storage controller to controller. Remote copy activity is server-less and remote copy is at the LUN/volume level. Issues with this solution include performance, distance, and multiple controller coordination.

Performance is another consideration with synchronous copy, and we often get asked what the distance limitation is. The answer is that it depends on the performance sensitivity of your applications. While newer technologies, such as Fibre Channel and DWDM, provide improvements, you still cannot exceed the speed of light. It takes about 1ms for light to travel 124 miles (199.64km); both ESCON and Fibre Channel protocols support multiple round trips when used for data replication. This means that even at short distances there will be overhead when using synchronous remote copy. Having said this, there are many satisfied customers using our synchronous technology very successfully.

One other benefit of using TrueCopy software is our ability to fully support IBM GDPS® (an IBM service offering) for system failover, workload balancing, and data mirroring on systems spread across two or more sites up to 25 miles (40km) apart.

While this may not be of interest to you today, it is an evolving technology that could be an option at a future time. Other vendors have announced support for this, but it is not clear when they will deliver this capability, if ever. At the current time, almost all of the production GDPS sites around the world are using Hitachi storage systems.

Semi-synchronous Option

While Hitachi Data Systems does not offer a semi-synchronous remote copy option, it is important for the purposes of comparison to understand how it works in relation to other technologies, as it is often used by other vendors as a stand-alone solution or as part of a multi-hop type configuration.

Hitachi Data Systems does not offer this technology, as we believe it poses serious data integrity exposures, especially in a rolling disaster situation. It is not recommended for disaster recovery because the application may assume a given write is successful when, in fact, it ultimately may have failed at the secondary location. This inconsistency can very easily cause the second copy of the data to become corrupted. While improvements to this technology, such as the use of consistency groups, have reduced this risk, caution should still prevail, especially in larger, complex environments (See Figure 6).

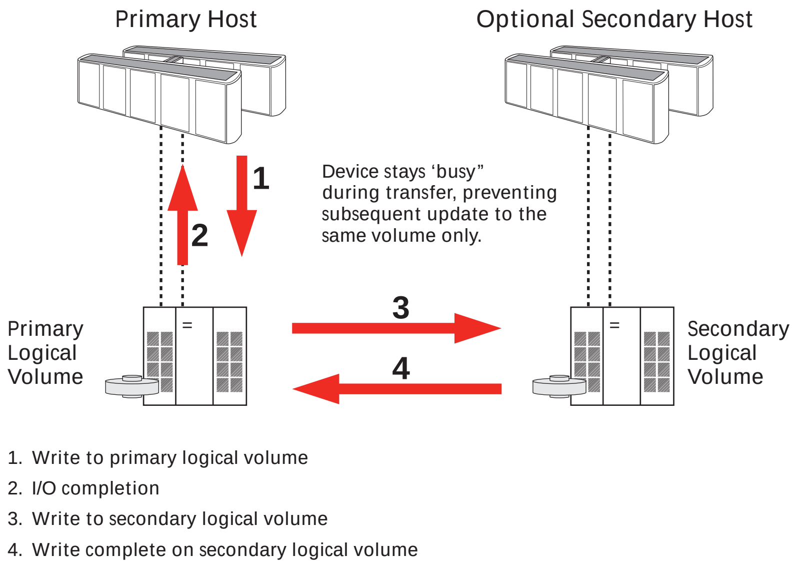

Figure 6: Semi-synchronous Solution. Issues in this solution include data integrity, multiple controller coordination, performance, and distance.

Semi-synchronous remote copy is used by some vendors in environments where synchronous performance is an issue and the customer will not consider a multi-hop or cascade-type solution. Hitachi Data Systems maintains that there are significant distance limitations as well as the data integrity issue with the semi-synchronous option. While in certain circumstances this option can improve on synchronous performance and possibly be used over greater distances, there are still typically limitations of under 100km, subject to the data patterns.

Hitachi Data Systems offers various alternatives to semi-synchronous remote copy, which can handle the replication of data directly over distance with complete integrity.

Hitachi Extended Remote Copy

This S/390 offering from Hitachi Data Systems is based upon the same architecture as IBM's XRC. Its strength is in its ability to survive a rolling disaster. This asynchronous remote copy technique is designed to provide complete data integrity for primary and remote systems. This is accomplished through the use of time stamps that allow the asynchronous updates to be formed into consistency groups by the System Data Mover (SDM) host software. Time stamping and SDM algorithms are the heart of XRC technology. While it is often considered a software offering due to the requirement of host software, the hardware has to be able to interface with the software and is, therefore, not supported by some vendors (See Figure 7).

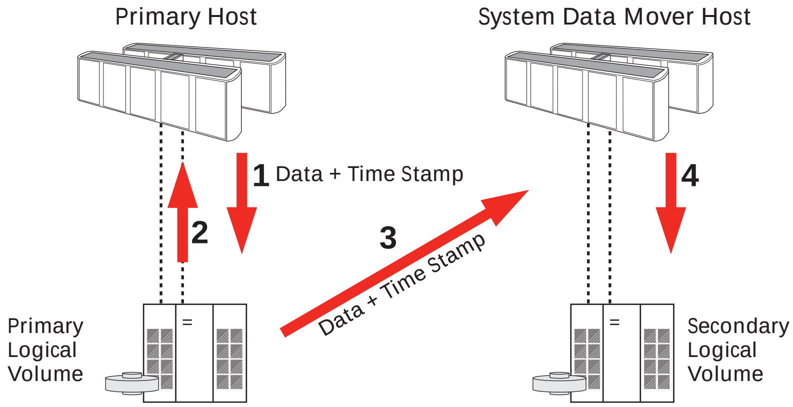

Figure 7: Hitachi Extended Remote Copy. This option provides great data integrity, long distance capability, and time stamps on all updates. Issues with this option include the requirement of host MIPs and a server at the secondary site, and cost (software license at secondary site).

Some of the issues associated with HXRC revolve around the need for a secondary site server and software; however, in order to recover quickly from any disaster, that is a requirement anyway. Recent enhancements with HXRC V3 have improved performance, reliability, and scalability, as well as adding unplanned outage support. EMC has recently added support for XRC, although it is believed the initial version is a very early version and may not have the recent enhancements.

Multi-Hop and Three Data-center Copy

These offerings revolve around creating multiple PiT copies of data and require anything from two to six copies of data and two or three storage systems. Multi-Hop from EMC was primarily devised for the express purpose of achieving consistency of data over distance, as they did not support any true data-consistent asynchronous technology. The obvious downside of these solutions is the requirement to have many additional copies of data and extra storage systems just to achieve consistency over distance, as is the case for EMC customers who do not have Symmetrix DMX storage systems or who require consistency across more than a single storage system in a mainframe environment (See SRDF-Asynchronous, page 14). This is not the case with Hitachi Data Systems replication offerings.

Other problems with Multi-Hop and Three Data-center Copy revolve around the requirement to run synchronous for the first hop, as you need to be close enough to avoid any synchronous performance issues, but far enough away to avoid the impact of a region-wide disaster as this could significantly affect the RPO of the recovery copy.

This technology does have applicability, such as for an environment where a customer wants to have two copies of data at different locations. In this scenario where there exists a requirement for two copies of data, one fairly local and one over an extreme distance, Multi-Hop and Three Data-center Copy may seem to be a good choice. The first hop gives you the local backup, with an additional copy at an alternate location. However, if a problem arises at the local backup site you have no disaster recovery at all, unless you replicate directly from your production center. If this occurred when using this type of solution from Hitachi Data Systems, you would be able to start up direct replication from your primary facility to the long-distance location with complete data integrity. This is not possible if the vendor does not have this capability.

Zero data loss over distance is another use for these solutions. By running synchronous and then an asynchronous technology, you can essentially get a synchronous-style replication over distance. It's important to remember that in order to protect your enterprise from a rolling disaster, it is unlikely that you will get zero data loss, but you may get close to it. Additionally, you need to review your RTO, as there is a catch-up phase to the recovery site that is cyclical before you can use the data. How long this takes will depend on the locality of reference of the data and the available bandwidth.

TrueCopy Asynchronous Software

This technology is unique to Hitachi Data Systems and was developed based on a requirement to achieve consistent disaster recovery solutions in both open and mainframe environments—without host software and the need for intermediate systems.

Using this technology, our customers can be confident that they can recover mission-critical business functions at facilities across the world, without the encumbrance and overhead that is inherent in other vendor's solutions.

In a mainframe environment, TrueCopy asynchronous software utilizes reliable time stamps, along with other information created by the primary system(s), to enable Hitachi storage systems to transmit updates directly to the secondary system (without any host intervention). These writes are then buffered in the secondary system's cache, queued, sorted by the time stamp, and written to the correspondent volumes in the same sequence. These rewrites are issued by the primary system over the remote link-checking-sequence numbers embedded in records to ensure no records are missing. In this way, I/O consistency is maintained. In open systems, we use sequence information to achieve the same level of consistency, though currently an application must reside on a single system. In contrast, an application in the S/390 environment can be spread across up to four primary controllers. We are working on eliminating some of these restrictions in the future. This approach allows the enterprise to achieve data integrity in a significantly simpler and less costly hardware-based solution, with no impact on server or application performance.

TrueCopy asynchronous software also uses the concept of consistency groups; this allows you to logically group applications together for consistency purposes (See Figure 8). This benefit allows you to execute operations on single applications, for example, during disaster recovery testing.

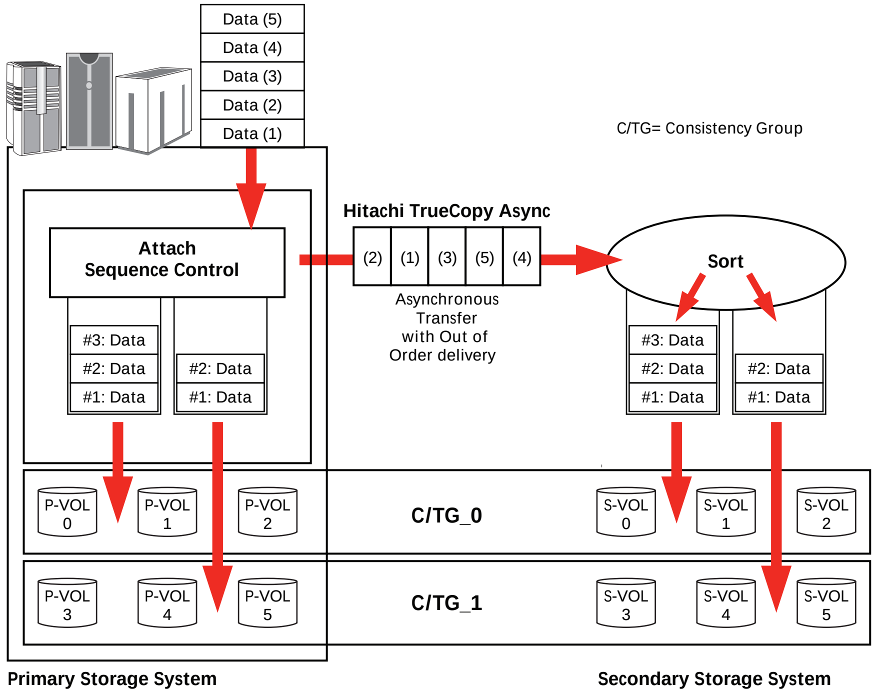

Figure 8: The Concept of Sequencing and Consistency Groups.

In areas where customers require long-distance replication to ensure business continuity or for highly performance-sensitive environments, Hitachi Data Systems has demonstrated clear leadership with TrueCopy asynchronous software:

- Data integrity is guaranteed for dependent write applications
- Excellent performance is achieved for both long- and short-distance requirements, due to its asynchronous nature
- Future enhancements will provide additional configuration flexibility

Let's compare TrueCopy asynchronous features to other asynchronous implementations:

- Other implementations are based on sending changed tracks to the secondary system instead of time-stamped I/Os
- Application of changed tracks by other implementations (and not individually time-stamped I/Os) cannot preserve the original sequence of writes, and therefore, should not be used for real-time disaster recovery unless used as part of a properly architected PiT solution.

NanoCopy Technology

NanoCopy is a PiT copy solution that allows customers to take a PiT copy (snapshot) of data without having to quiesce or otherwise interrupt the application accessing the data. It is, in fact, the only storage-based solution from any vendor that can do this. This version of NanoCopy technology is available for OS/390 environments on TrueCopy asynchronous software for any number of logical volumes and storage systems. This OS/390 capability has also been significantly enhanced through scripting to provide an alternative disaster recovery solution by taking periodic (and non-disruptive) PiT copies. These copies may be split off every 15 minutes or every few hours, according to requirements. NanoCopy technology is an architected solution that is based on your goals and objectives.

NanoCopy technology is also available on TrueCopy asynchronous software for Windows and all open systems platforms, although the source data must reside on a single Lightning 9900™ V Series system.

Cascade Techniques

Cascade copies refer to a method of creating PiT copies of data similar to NanoCopy, however there are major differences. First, most solutions in this area require you to quiesce the application in order to spin off the PiT copy. This will affect the application, and therefore, the age of the PiT copy may be affected as you may not want to do this very often. Additionally, quite often this technology is designed using four copies of data instead of three. There would be two local copies and two remote copies. This technology may seem appropriate for an environment where a customer wants to create a PIT copy every four hours, as there is some expectation of lower bandwidth requirements in these types of technology since not every update is replicated to the remote site. This is addressed in the bandwidth section later in this paper.

SRDF Asynchronous

This is a new asynchronous technology from EMC, and as mentioned previously, it is only available in the DMX series. Its technique for preserving write dependency is essentially just honoring writes without the use of other proven techniques such as sequencing or time stamping.

SRDF Asynchronous uses the concept of timed cycles known as delta sets, typically every 30 seconds, and captures all host I/O during that period in cache. If the same record is updated more than once, only the most recent update is kept. Any dependent I/O will in theory be in that delta set or in a subsequent one. Once the time is up, SRDF Asynchronous starts another delta set cycle at the primary site and begins to transmit the previous delta set to the remote side cache. Once all the data is received at the secondary site, it then promotes it to be an active cycle and the data can be destaged to the back-end disk at the secondary site. If there is a problem during transmission or a disaster occurs, then the copy of consistent data at the remote site should be the previous completely transmitted delta set. However, with multiple links and network retries, that may be of concern. This technology would also appear to use large amounts of cache, as the data is held there much longer than other asynchronous techniques. EMC also claims significant improvement in required bandwidth, but this will be highly data and network dependent. Additionally, since the time RPO of the data will be at least the 30-second cycle and the time required to transmit it, using too little bandwidth will result in elongated RPOs and the possible dropping of the SRDF-A environment due to cache overload.

SRDF Asynchronous also has many restrictions, compared to other technologies. For example, it only supports a one-to-one environment and a single consistency group, giving customers very little flexibility.

Bandwidth Considerations

One of the biggest considerations for your organization is the cost of the solution, and network bandwidth is one of the biggest contributors to that cost. In real-time remote copy, every update is sent to the remote location. If your application executes 100 writes of 10K blocks/sec you need bandwidth that accommodates the writes, plus any control information that is also sent. In PiT copy solutions for disaster recovery, data is only replicated at pre-set intervals, such as every 15 minutes, every hour, etc. During the period of time that data is not being replicated, for data that is updated at the primary site, tracks are marked as changed by the storage system. If the same record is updated 100 times, then only the last change to the track will be shipped when the data is sent to the secondary location. This in theory means that the bandwidth requirement may be lower when using PiT technologies; however, there are many factors to consider before coming to that conclusion. If the PiT copy is frequent—under six hours, for example—as is often required, then the chances are that there will not be significant reduction in bandwidth. In certain circumstances, bandwidth requirements may be the same or not significantly different. The bandwidth requirement will be very dependent on locality of reference of the data and how current you want the PiT to be at the remote location.

Careful consideration should be used when architecting Remote Copy Solutions with a view to saving network bandwidth. This includes technologies such as Cascade, NanoCopy, Multi-Hop, Three Data-center Copy, and even SRDF Asynchronous. It is important to understand your data patterns, to try and determine if there are any benefits, and to consider your RPO and RTO. Using too little bandwidth can elongate your RPO and cause other problems in the environment. Where catch-up is required, the available bandwidth will affect the time you can begin your recovery (RTO).

In environments in which RPO of the data can be looked at in hours versus minutes, a PiT solution may be a viable alternative to a real-time solution.

Summary

- Hitachi Data Systems is the only vendor to provide flexibility and choice in technologies and solutions for both open systems and S/390 environments.
- The focus of Hitachi Data Systems product offerings is on standards and interoperability.
- Lightning 9900 V Series systems provide fully compatible, high-performance S/390 (PPRC, GDPS, and XRC) solutions. This positions customers to take advantage of future software enhancements.
- TrueCopy asynchronous software and NanoCopy technology provide unique, enterprise-wide, simple, elegant PiT and disaster recovery solutions.

There are many copy technologies today that can be considered when implementing business continuity solutions, especially when traditional backup methods of copying data sets are considered. This is why it is important to choose only from best-of-breed solutions when addressing disaster recovery objectives. Hitachi Data Systems provides not only a superior range of offerings in this area, but also the expertise necessary to help our customers as they move toward their goals.

Hitachi Data Systems

www.hds.com

info@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 3468
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) 970-1001
ussalesinfo@hds.com

WHP-117-01
December 2003