HITACHI
Inspire the Next

WHITE PAPER

Hitachi Vantara's Approach to 3-Data-Center Business Continuity and Disaster Recovery

Achieve Zero Data Loss and Geo-Dispersion

By Hitachi Vantara

October 2017

Contents

Executive Summary 2

Introduction 3

Understand the Risks 3

Why Do You Need a 3DC Architecture? 4

Types of 3DC Architectures 4

Cascade 3DC Disaster Recovery Architecture 5

Multitarget 3DC Disaster Recovery Architecture 6

Storage Clustering With Asynchronous Replication 6

Simplify Operations: Consolidate and Automate End-to-End Recovery Processes 7

What About Using Cloud for the Third Site? 7

Recommendations 8

Next Steps 8

Executive Summary

The risk of permanent data loss is a growing problem for organizations as the amount of data under management is growing at an annual compound rate of 40% or more. In addition, the complexity and interdependency of the IT environment is increasing, and the number and size of risks to the organization continue to rise. Previously, organizations had manual procedures and paper trails they could utilize to recover data, but those are largely a thing of the past. Today's interconnected world can create an illusion of business continuity; however, the risk to revenue and reputation from an extended outage has never been greater.

The standard for disaster tolerance in most organizations is a two-site disaster recovery architecture, with one production site and a secondary disaster recovery location. The distance between these sites helps to determine the amount of data at risk of permanent loss. The further the distance, the greater the data transfer latency, and the greater the amount of data that will not be in the secondary site at the time of a production site outage. However, having the two sites within close proximity results in the risk of both sites being impacted by a major regional event, such as a large earthquake or violent storm.

For many organizations, these risks are becoming unacceptable, and reducing the probability of permanent data loss is a business imperative.

Recent technological advances in storage-based data services, including active-active storage clustering and highly efficient long-distance replication, have enabled the cost-effective implementation of 3-data-center (3DC) disaster recovery architectures. These architectures dramatically improve protection against permanent loss and corruption of data. Moreover, this approach provides the foundation for much faster system recovery, or can even mitigate the need for recovery.

Introduction

When a disaster strikes, an organization may lose data as well as access to data, and thus the ability to function. Recovering from such a catastrophe is a business imperative. Leveraging techniques that focus on risk quantification and mitigation is key: It will help organizations to choose which technology to use and balance that decision against the determination of how much to spend.

Specifically, this decision revolves around three fundamental service level requirements for business-critical functions and their associated applications:

- Recovery point objective: How much data loss can be tolerated during recovery? For critical data, the ideal is usually zero data loss.
- Recovery time objective: What is the acceptable time within which to recover systems and operations? The goal for critical operations should be measured in seconds or minutes.
- Return on investment: What is the right balance between risk and mitigation? The total cost of disaster protection should be less than the anticipated impact of a potential major disruption.

Understand the Risks

What could possibly happen to disrupt your operations? Depending on your location, a lot (see Table 1).

Table 1. Levels and Types of Disruptions

Categories	Threats
System Events	Hardware or software failures, network problems, corrupted data, viruses, glitches, bugs
Internal Events	Human error, fire, plumbing leaks, electrical spikes, construction defects, angry employee
External Events	Utility interruptions, sabotage or terrorism, hacking, accidents
Acts of Nature	Floods, hurricane or typhoon, tornados, earthquakes
Interdependence Threats	Supply chain disruptions, partner failures, labor strikes

These decisions are not independent of one another. Data loss and loss of data integrity can significantly increase recovery time.

Historically, the emphasis has been on reducing the time required to recover systems. All recovered systems lose some data. Today, there is significant pressure on organizations to implement recovery solutions that give a very high probability of zero data loss; this has been a business imperative for the financial industry for some time. More and more, organizations are finding it inordinately costly to recover if data is lost, and the corresponding recovery time is significantly increased. The business processes to recover lost data manually after a disaster become more difficult as processes are increasingly computerized. Automatic remote recovery for all major IT systems can simplify business processes and reduce costs.

Recovery decisions can no longer be made in isolation. Within an organization, the failure of one system can quickly have a domino effect and bring down other systems. Government agencies and business groups are mandating increasingly stringent recovery objectives to ensure that industries and society can recover quickly from manmade and natural disasters. As the world of digital business evolves, the attributes of business continuity and disaster recovery are transitioning from being highly desirable targets to being absolute “must haves.”

A SERIOUS OUTAGE CAN RUIN YOUR BUSINESS

According to the U.S. National Archives and Records Administration, 93% of companies that lost their data for 10 days or more filed for bankruptcy within one year of the disaster, and 50% filed for bankruptcy immediately.

Why Do You Need a 3DC Architecture?

Recovery solutions have an important technology constraint: Zero data loss cannot be achieved over long distances. The practical limitation is usually less than 50 miles, or whatever latency the application will tolerate when its data is synchronously replicated.

If your critical applications and data require always-on access, you will need an active-active high-availability topology between sites. However, if you can't get two sites within acceptable synchronous distance that meet the requirements for geographic dispersion, you will need a third site to protect against a major disaster that impacts the first two.

By introducing a three-site architecture with two data recovery nodes (one at close distance and one at long distance), a very high probability of zero data loss and fast recovery times can be achieved.

Types of 3DC Architectures

The architectures of 3DC disaster recovery solutions can consist of a combination of technologies to enable very high probabilities of zero data loss at local and long distances. They combine synchronous replication (local recovery node) with asynchronous replication (remote recovery node).

The local recovery node can accommodate very rapid recovery with a high probability of zero permanent data loss. Active-active storage clustering adds even greater resilience to the design. Testing of this environment is simplified, and IT personnel can be shared between the primary node and the local backup nodes.

The remote recovery node provides for recovery with low permanent data loss "in the unlikely event" that both the primary and local recovery nodes are impacted.

As is the case in many IT decisions, choices available for 3DC protection can be deployed and even combined to meet an organization's specific needs for various business continuity and disaster recovery scenarios. Hitachi Vantara offers the following options, among other combinations, with Hitachi Virtual Storage Platform (VSP) G series and VSP F series storage systems. These models use the remote replication technologies in Hitachi Storage Virtualization Operating System (SVOS), as shown in Table 2.

Table 2. Types of 3DC Architectures

Architecture	Technologies
Cascade: Synchronous + Asynchronous	Hitachi TrueCopy + Hitachi Universal Replicator
Multitarget: Synchronous + Asynchronous	TrueCopy + Universal Replicator + Universal Replicator
Storage Clustering + Asynchronous	Global-Active Device Feature + Universal Replicator

Cascade 3DC Disaster Recovery Architecture

This approach (see Figure 1) is sometimes known as “multihop,” and it combines technologies to provide a high probability of zero permanent data loss for most disaster scenarios over a long distance.

Figure 1. Multihop Topology

Typically, Hitachi TrueCopy synchronous remote replication is used between the production and local recovery sites to minimize data loss, with Hitachi Universal Replicator used to copy the data asynchronously from the local recovery site to the remote recovery site. Alternatively, Universal Replicator can be used for both links.

DON’T GIVE YOUR CUSTOMERS A REASON TO SWITCH

Business continuity and disaster recovery, using a 3-data-center topology, will be essential contributors to sustaining brand loyalty and brand value, especially in a world in which changing vendors can be as simple as going to a different website.

Recovery time depends on various factors: the speed of the long-distance link between the local and remote recovery nodes, what time of day or year the primary node goes down, and the complexity of the recovery process. Recovery can be made at the remote node in under an hour or within a few hours.

There are two main options within this topology:

- The local recovery node can be a minimal disk-only “bunker.” Its primary function is to ensure that data can continue flowing, to bring the remote recovery node completely up to date should the primary node go down. The local recovery node is often an unmanned storage site. This configuration is the most cost-effective way of providing a high probability of zero data loss at a remote recovery node, with very good recovery time characteristics.
- Less frequently, the local recovery node can be a full data center (often with failover and failback systems). This approach provides zero data loss and very rapid recovery for disasters at the primary node. Going forward, this configuration is less likely, as the multitarget topology discussed in the next section is a cost-effective option that gives better protection.

One trade-off with cascade topology is seen in the following example. In the event that the local recovery node goes down, the remote recovery node is frozen with the data it has received at that point in time. The organization then must decide whether to continue to run the IT production systems without ongoing protection.

If it does, the remote recovery node gets further behind, and if a rolling disaster takes out the primary node as well, significant permanent data loss can occur. Alternatively, the organization can stop the systems at the primary node until the secondary node is recovered or a communications link can be established between the primary node and

the remote recovery node. In this case, the recovery time is elongated, but the probability of permanent data loss is minimized.

For organizations within a small geographical area, the cascade three-node topology makes good business sense. A disaster that takes down both the primary and local recovery sites is likely to affect most local customers. For interstate and international business, and especially for organizations that provide critical infrastructure services, this topology may not meet more exacting requirements.

Multitarget 3DC Disaster Recovery Architecture

The difference between the cascade topology and the multitarget is that in the multitarget topology, the primary data node backs up data to both nodes simultaneously. See Figure 2.

Figure 2. Multitarget Topology

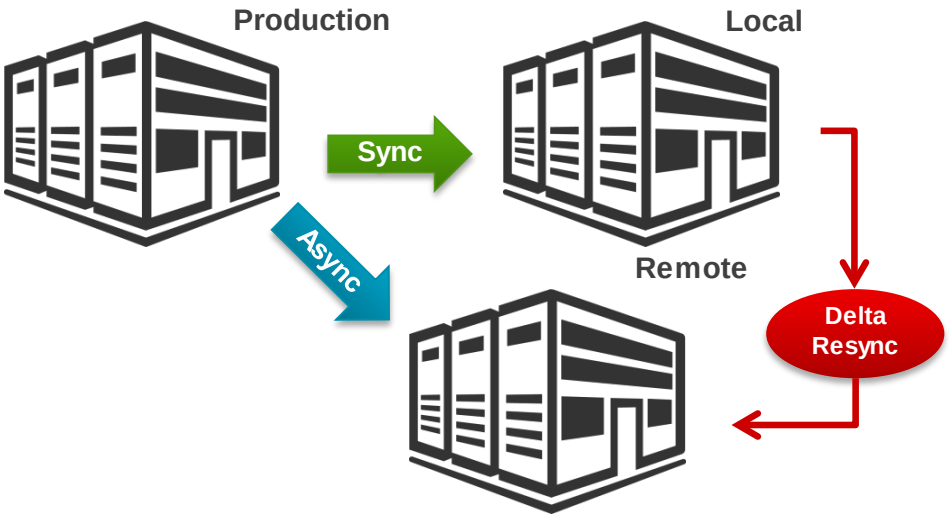

This is a recent technological capability, and very-high-performance storage controllers are required to manage this process. This approach ensures that there is no permanent data loss if either the primary or local recovery node is lost.

Either node can communicate data to the remote recovery node to ensure zero data loss. To ensure rapid recovery, the storage controller technology must be able to resynchronize the controllers at the remote recovery node with either the primary or local node, and pass just the changed data (delta resynchronization). In the cascade topology, if the local recovery node is down, no data can be transferred to the remote recovery node, as discussed above.

The major disadvantage of the multitarget architecture is the higher cost of telecommunication lines. A major advantage is that if there are backup servers in the local recovery node, there can be failover and failback between the primary and local nodes. This significantly enhances recovery times.

Additionally, remote snapshots or clones can be created and mounted in either of the backup sites to enable secondary processes. These include operations such as backing up to tape, refreshing development systems, or enabling recovery testing without impacting the performance or availability of the production systems.

Storage Clustering With Asynchronous Replication

The most recent advances in storage resiliency and data availability are found in global-active device storage clustering technology. This Hitachi feature is part of the Storage Virtualization Operating System of VSP G series

and VSP F series systems. With global-active device, there are two production sites, each with an active copy of all data. If a failure occurs at either site, its data is transparently available at the other site, with no need to fail over or fail back.

Universal Replicator asynchronous replication is used to copy the data from either production site to the remote recovery site. All additional benefits of the multitarget configuration described above apply to this architecture. See Figure 3.

Figure 3. 3DC With Active-Active Storage Clustering

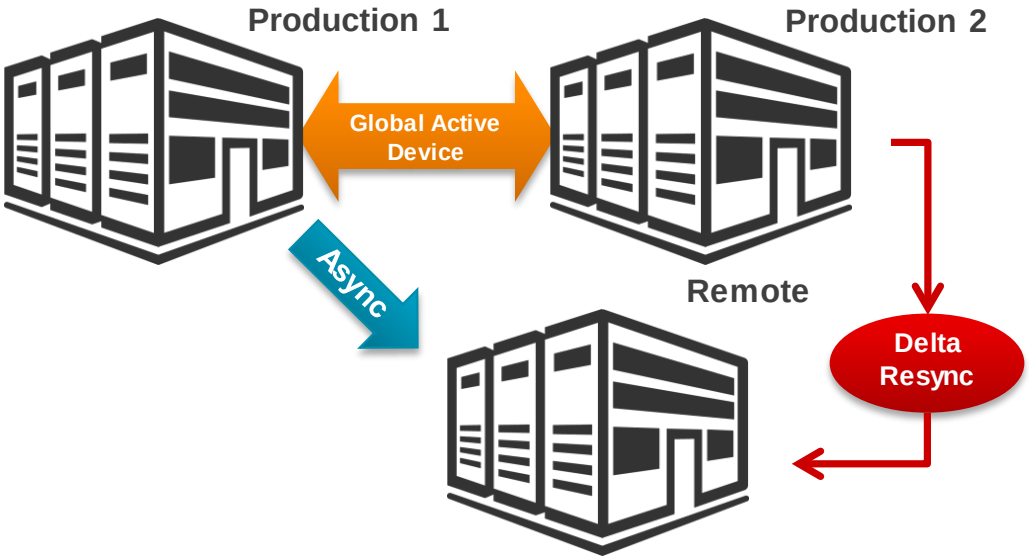

The storage clustering 3DC model provides the greatest levels of data availability and resiliency with zero data loss.

Simplify Operations: Consolidate and Automate End-to-End Recovery Processes

Many disaster recovery strategies consist of a collection of application-specific point solutions that must be individually maintained and managed. These must then be individually executed at the time of a disaster. This situation results in “hidden” costs and risks to the organization. These hidden costs and risks are a function of the recovery complexity, which only gets worse as the organization scales up, provisioning new applications and/or data.

Hitachi storage can consolidate all of your business continuity and data recovery processes using a flexible combination of in-system and remote replication capabilities. Hitachi Data Instance Director software also dramatically simplifies the configuration, management, automation and orchestration of these modern data protection technologies. Easily combine local snapshots with remote replication and then remote snapshots. Enable local application-consistent operational recovery, remote business continuity, disaster recovery and data repurposing, all in a single policy-based workflow.

What About Using Cloud for the Third Site?

Many organizations are considering the use of a third-party “disaster-recovery-as-a-service” (DRaaS) solution as an option for storing a copy of their data, either in a vault or in the cloud. The primary benefits of these services are cost reduction and replacing large equipment capital costs with smaller monthly operating expenses. However, trusting your data with a third party exposes the organization to several potential risks.

The most obvious risk is data security. The service providers may employ more advanced security protection than a typical business can. However, they are potentially a much bigger target for hackers, and you have no direct control of their employees and contractors. You must also put your faith in their ability to recover from a disaster of their own to ensure continuous availability of your data, if that is even part of the service.

Their low “pay-as-you-go” pricing models may be attractive today as the service providers seek to gain market share at the expense of profit margins, but as the service industry consolidates those prices will likely rise. And, as has been demonstrated several times, with such companies as Nirvanix and Symantec with Backup Exec.cloud, the business failure of a service provider can make it very challenging to retrieve or move the data they have been storing for their customers.

Hitachi Vantara continues to develop and provide cost-effective data storage, business resiliency and disaster recovery solutions for organizations that prefer to maintain total control of their data assets. Equipment leasing options are available that provide an attractive alternative to pay-as-you-go services from third parties. Other solutions, such as managed or hosted services, can be tailored to meet specific requirements.

Recommendations

- Understand your risks, the likelihood of them occurring, and what their financial impact may be.
- Prioritize your systems, applications and data, based on service level agreement (SLA) requirements, including operational resilience, operational recovery, disaster recovery and regulatory compliance.
- Evaluate the technology choices available that mitigate the risks, and balance the costs and the benefits to arrive at the right solution for each requirement. The investment should also be balanced across the business process, including servers, storage, networking and people.
- Take advantage of automation and orchestration tools to simplify configuration and management.
- Public or shared cloud services may be an attractive disaster recovery option, but understand the operational and security risks associated with entrusting your data to a third party.

Next Steps

You do not need to take this journey alone. Hitachi Vantara and its partners offer professional services that include assessment, planning, design, installation, implementation, transition, migration, management, optimization and support.

To learn more, please visit hitachivantara.com/go/protect, or contact DP-Sales@hitachivantara.com or your local Hitachi Vantara sales office to schedule an initial conversation with our business continuity and disaster recovery experts.

Hitachi Vantara

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hitachivantara.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com
Asia Pacific: +852 3189 7900 or info.marketing.apac@hitachivantara.com

WP-564-A R. Vining October 2017.