White Paper

Mainframe Storage Compatibility and Innovation With Hitachi Virtual Storage Platform One Block High End

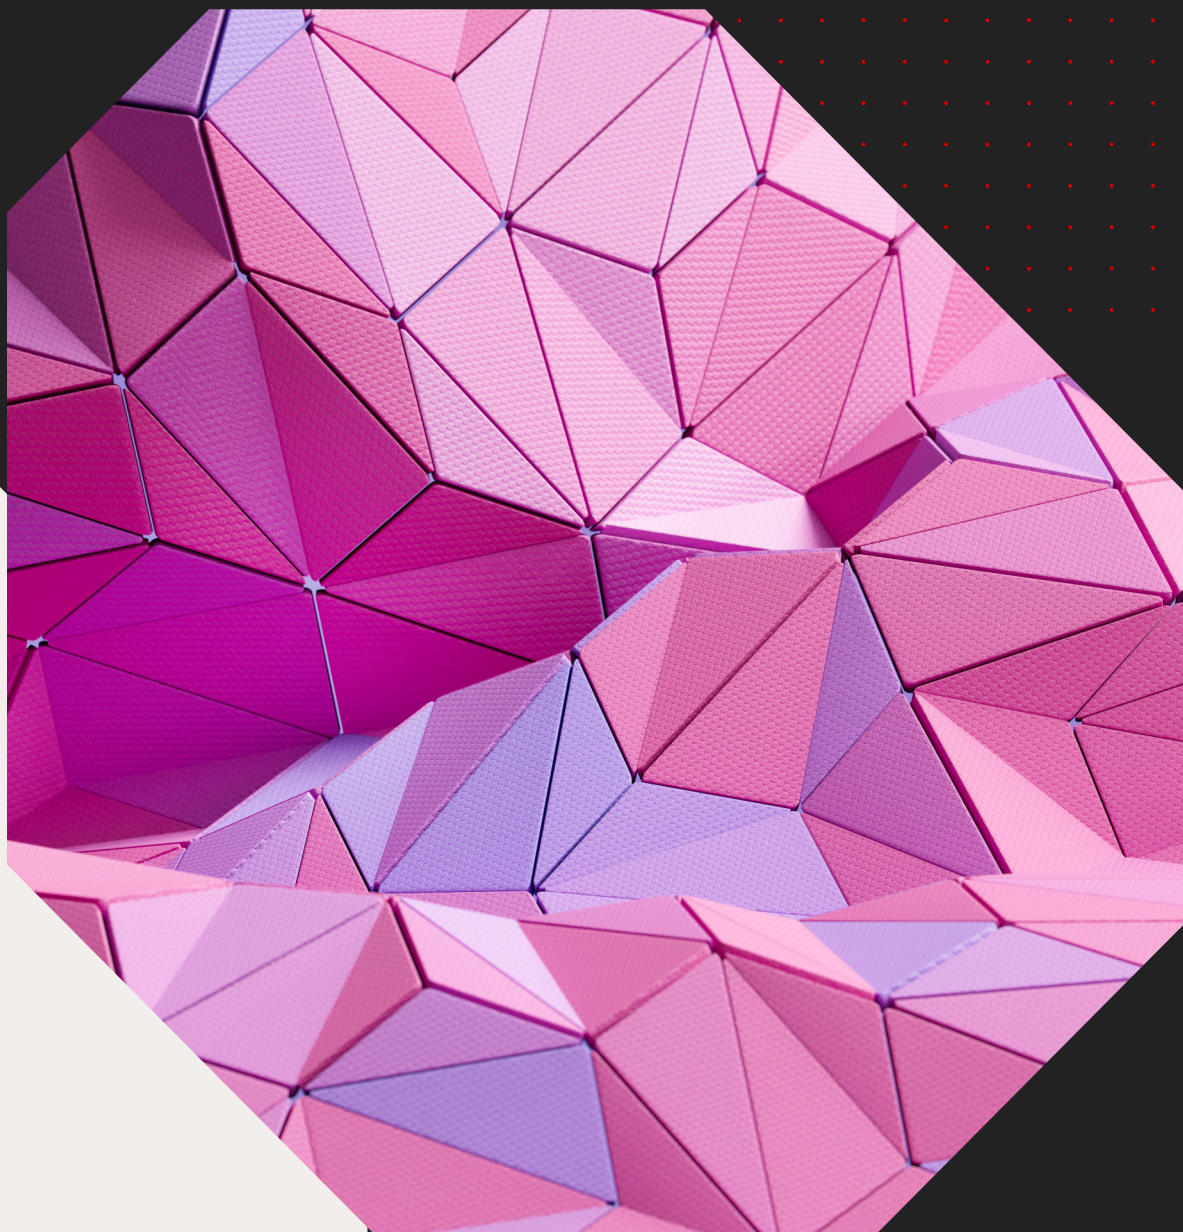

Hitachi Vantara

With over 45 years of experience supporting IBM® mainframe environments, Hitachi Vantara is committed to continued support of these environments. We provide innovative new storage solutions designed to improve storage processing, performance, availability, recoverability and management in mainframe environments.

Hitachi Vantara Delivers for Mainframe Environments

- Introduced a 100% data accessibility warranty with first generation of Hitachi RAID in 1995 (7700 Classic).
- Was the first — and is still the only — company to offer eight 9s of availability of a standalone storage subsystem, beginning with the introduction of Hitachi Virtual Storage Platform (VSP) 5500 in 2019.
- Offered the lowest asynchronous replication RPO of the industry with Hitachi Universal Replicator, which was introduced in 2004 with Hitachi Universal Storage Platform (USP) series.
- Hitachi was the first vendor to introduce Thin Provisioning for Mainframe on Hitachi Universal Storage Platform V (USP V) in 2009.

Today, we're delivering the solution that allows enterprises to store immutable images of the production data in a fortress — essential protection, should the company suffer from a logical corruption or worse, be subject to a cyber attack. This solution also allows organizations to check the data periodically by mounting stored images.

Qualifications and Support for IBM

Hitachi Vantara, a subsidiary of Hitachi, Ltd., is a leading provider of mainframe storage. We use a two-pillar approach in the design and development of IBM mainframe compatible storage systems. The first is compatibility, with a focus on support and qualification of essential IBM mainframe storage features. The second is innovation, development and testing of Hitachi unique storage and storage management features to provide mainframe customers with additional capability and value.

Hitachi Vantara has a track record of earning IBM qualifications for Hitachi storage, including Basic qualification and IBM GDPS® qualification. Basic qualification ensures compatibility between IBM servers and Hitachi storage. GDPS qualification ensures that Hitachi storage will meet the requirements for business continuity. Hitachi achieves GDPS qualification with IBM GDPS [including Metro Mirror (MM), Peer-to-Peer Remote Copy (PPRC)], with IBM FlashCopy®, with Hitachi Universal Replicator 3DC Delta Resync, with Universal Replicator 4DC, and Multi-Target Metro Mirror (MTMM) with Multi-Target Incremental Resynchronization (MTIR).

Hitachi has participated in this IBM qualification process since 2005. This qualification testing is one more way organizations can be confident in the Hitachi commitment to providing compatibility with essential mainframe storage features. And they gain full support of unique, innovative value-added replication solutions, including Hitachi TrueCopy, Hitachi Universal Replicator and Hitachi ShadowImage for Mainframe.

Hitachi Vantara has a large installed base of Hitachi storage systems connected to IBM Z® mainframes (formally known as z Systems®) via IBM FICON® networks as well as mainframe-based Linux® environments via Fibre Channel networks. Hitachi provides high-speed flash storage for mainframe environments with full NVMe connectivity (since 2019). Hitachi is the first storage manufacturer to offer virtualized storage in the mainframe environment, including the ability to virtualize external multivendor storage.

With industry-leading Hitachi storage innovation, IBM mainframe compatibility, performance and data resilience, you can trust that your Hitachi storage system will meet the demands of your mainframe and open systems storage environments.

Hitachi Vantara Virtual Storage Platform One Block High End (VSP One Block High End) inherits Hitachi storage history and customers will benefit from improvements and new features.

The Hitachi Two-Pillar Approach to Mainframe Storage

Hitachi designs and develops IBM mainframe compatible storage systems with a two-pillar approach: compatibility and innovation.

Compatibility

Hitachi and IBM share APIs to ensure compatibility between Hitachi storage and IBM mainframes. For over 45 years, Hitachi Vantara has provided mainframe technology to business and industry, including solutions for z/OS®, z/VM®, z/VSE®, z/TPF® and Linux for IBM Z. During this time, license agreements have provided Hitachi access to IBM APIs and interface specifications for mainframe storage platforms as well as IBM access to selected Hitachi APIs.

Access to these APIs has enabled Hitachi storage systems to support PPRC, FlashCopy, FlashCopy Space Efficient, Remote Pair FlashCopy (Preserved Mirror), Parallel Access Volumes

(PAV), HyperPAV, SuperPAV, Multiple Allegiance (MA), High Performance FICON (zHPF), CSM and Basic HyperSwap® support (as well as Metro Mirror Failover/Failback with HyperSwap) and IBM Fibre Channel Endpoint Security.

The license agreements have also given IBM access to Hitachi APIs for NanoCopy disaster recovery manager technology from Hitachi Vantara and Hitachi ShadowImage for Mainframe. This licensing agreement also enables Hitachi Vantara to develop and deliver storage management software for mainframe environments. Hitachi Vantara participates in the Early Shipment Program for z/OS, which means that Hitachi has early access to z/OS versions for testing. This access enables Hitachi to ensure compatibility with the new versions of z/OS when they become generally available.

Hitachi storage systems and features are tested and qualified with IBM mainframes, operating systems and critical features. This testing and qualification includes features such as support for GDPS MM and MTMM, FlashCopy and zHPF. Hitachi Virtual Storage Platforms (VSPs) are qualified and supported by the IBM Z series processors. They are also supported with the latest z/OS version.

Hitachi has implemented on the VSP series technologies that are compatible with IBM's Metro Mirror (PPRC), Multi-Target Metro Mirror, FlashCopy V2 and FlashCopy Space Efficient. We are the first storage vendor to go through this broad qualification.

IBM also successfully completed the qualification testing for GDPS/HyperSwap, FlashCopy, and three- and four-site data center architectures (3DCz and 4DCz) with multitarget, multisite replication configuration using VSP, GDPS/HyperSwap support and Universal Replicator along with Hitachi Business Continuity Manager (BCM). IBM also qualified MTMM/MTIR support (https://www.ibm.com/it-infrastructure/z/technologies/gdps qualification section).

The license agreements ensure that as IBM announces and delivers new features and functions for mainframe environments, Hitachi Vantara will have the capability to support these new features. In addition, Hitachi has its own testing facilities to ensure that Hitachi Vantara delivered products and features are compatible with the IBM solutions.

The Hitachi Vantara and IBM license agreement provides Hitachi Vantara access to the design specifications for key technologies from IBM, which enables Hitachi Vantara to deliver support for these key technologies:

- FICON 32Gbit with IFCES
- Enhanced Dynamic Cache Management
- Cache Fast Write (CFW) and DASD Fast Write (DFW)
- Sequential Data Striping
- PDS Search Assist
- Parallel Access Volumes (PAV)
- HyperPAV
- SuperPAV
- Multiple Allegiance
- MIDAW
- Priority I/O Queuing
- Extended Address Volume (EAV, up to 1TB disks)
- Dynamic Volume Expansion (DVE)
- Forward Error Correction
- Read Diagnostic Parameters (RDP)
- FICON Dynamic Routing (FIDR)
- Query Host Access
- Soft Fence
- zHyperWrite
- High Performance FICON (zHPF)
  - zHPF Single/Multitrack
  - zHPF DB2® List Prefetch
  - zHPF Bi-directional Channel Program
  - zHPF List Prefetch Optimizer
  - zHPF QSAM/BSAM Support
  - zHPF Format Writes
  - zHPF Extended Distance
  - zHPF Extended Distance II
  - zHPF z/TPF
- IBM Metro Mirror (PPRC)
- IBM Multi-Target Metro Mirror (MTMM)
- FlashCopy V2 (including multiple incremental FC)
- FlashCopy Space Efficient (FlashCopy SE)
- Remote Pair FlashCopy – Preserved Mirror
- Geographically Dispersed Parallel Sysplex™ (GDPS) HyperSwap
- Basic HyperSwap with Copy Services Manager
- Metro Mirror Failover/Failback with HyperSwap With Copy Services Manager
- PPRC Summary (PPRCSUM)
- Hardware Reserve Transfer
- Concurrent Copy
- z/OS Discovery and Auto Configuration (zDAC) support
- DB2 Cast Out Accelerator
- LinuxONE™ Support
- FCP Support

Hitachi Vantara not only provides support for IBM compatible replication solutions such as Metro Mirror (PPRC), but also provides ongoing innovation, delivering alternative high performance replications solutions, such as Universal Replicator.

Hitachi Vantara and IBM have a long-term agreement to help each other to quickly resolve any joint customer issues related to Hitachi storage and IBM mainframes as well as mainframe operating systems. The focus is on determining the source of the problem with the root cause and solution being provided by the vendor whose product(s) is identified as the source of the problem.

Innovation

Innovation, the second pillar of Hitachi two-pillar approach, delivers unique Hitachi storage and storage management features, providing mainframe customers with additional capability and value.

Hitachi enterprise storage platforms are used in most Fortune 100 organizations today. These represent some of the most critical mainframe environments in the world.

As mentioned throughout this document, Hitachi has made several innovations in the past and keeps on innovating.

Cyber Resiliency Solution

Hitachi has engineered a solution to protect enterprises from logical corruption or worse — a cyber attack. The unique features of the VSP series allow us to make an image of production without the need of an I/O freeze. Virtualization, another unique feature of the VSP series, allows this delta copy triggered on mainframe storage to reside on open systems storage, creating a true air gap and another layer of protection from attacks as it is unknown from mainframe. The copy and up to 1,022 point-in-time, space-efficient snapshots are stored in a fortress, making the copy and delta stored on that open system storage immutable, associated with the predefined retention period. The images stored in the fortress are immutable. Users can mount any of the images stored in the fortress, with read and write capability for verification, forensic analysis, surgical restore or full restore. Please also note that restore can be done using a different storage than the primary storage.

Dynamic Provisioning

Hitachi was the first to introduce this technology in the mainframe world, and it has been adopted by the market. Hitachi Dynamic Provisioning for Mainframe (HDPM) complements existing mainframe provisioning process while improving performance and simplifying performance and capacity optimization. HDPM automatically spreads the content of application data sets across a storage pool, containing many physical media (SSD-based disks with NVMe interface), distributing I/O activity across available physical resources and backend paths. This action eliminates the challenge of manually spreading application data over many backend physical media to optimize performance and throughput. At the same time, it reduces storage management concerns.

Mainframe Analytics Recorder

With the mainframe analytics recorder capability included in Hitachi Storage Virtualization Operating System (SVOS) we deliver storage-system-specific detailed performance integrated with other SMF® (System Management Facility)/ RMF™ (Resource Management Facility) reporters. This capability is key to providing system programmers with an integrated view of their mainframe storage internal performance. It enables time- and load-synchronized reporting, an essential capability that is very hard (and laborious) to achieve with a mix of different open-based array reporting and mainframe tools.

Mainframe analytics reporter feature shortens the time needed to analyze and optimize the system, and to solve time-critical performance problems. It can also help you to avoid them in the first place. And it enables third-party partners to better extend their analysis and management capabilities for Hitachi and virtualized storage.

Universal Replicator

With Hitachi Universal Replicator, administrators have an innovative remote replication solution. Universal Replicator offloads most of the processing for asynchronous replication to the remote storage system using journaling to keep track of all updates as well as unique Hitachi pull technology. This frees up processing power at the primary site for critical production activities and provides the best performance available. Hitachi Vantara has also developed support for Extended Consistency Groups to provide consistency of data cross storage systems during replication.

Enhanced Storage Features for Mainframe Environments in the Age of AI

Data is at the core of your data center, and any data center must provide for the movement, provisioning, access and protection of data, which is provided by storage systems. An important element of your information availability strategy is ensuring your data is resilient.

Information Availability and Data Resilience

VPS One Block High End represents the third generation of disk subsystems from Hitachi Vantara to deliver 99.999999% availability (eight 9s) — a level of reliability that remains unmatched in the industry for standalone storage systems. Unlike other solutions that rely on complex HyperSwap technologies to achieve similar availability, our architecture accomplishes this independently.

Given the mission-critical nature of our customers' operations, our system supports nondisruptive microcode upgrades, maintaining full availability and potentially avoiding any performance impact. This is made possible by our ability to upgrade one core at a time, ensuring continuous service.

In contrast, cluster-based architectures often require node shutdowns during upgrades, introducing significant risks to availability due to the absence of failover mechanisms. Hitachi architecture has evolved specifically to overcome these limitations, enabling consistent high availability. Competitive systems have not undergone similar advancements and, as a result, have not achieved eight 9s availability in their latest models.

Data resilience does not mean just disaster recovery; it means that in the case of a system or site failure your data and information can be available in an alternate storage system at the same or a different location (see Figure 1). It means you have availability that allows your business to continue business operations in the event of a planned or unplanned outage.

Solutions to provide data resilience include service level agreements and requirements for high availability, with specific recovery time objectives (RTO) and recovery point objectives (RPO). Solutions to provide data resilience also include data recoverability should logical corruption happen and put the company at risk.

An increasing corporate understanding of the value of data and importance of the IT infrastructure has resulted in IT organizations being tasked to ensure the IT infrastructure will protect these critical assets across the corporation. This means that data must be protected and guaranteed to be consistent across complex applications that share data and require the best local and remote replication and restart available.

Hitachi advanced replication and replication management — supporting IBM GDPS HyperSwap, Basic HyperSwap, and Metro Mirror Failover/Failback with HyperSwap — enable enterprises to maintain uninterrupted operations and ensure rapid recovery in the event of a disruption. These capabilities provide a robust foundation for high availability and business continuity, allowing executives to operate with greater confidence in their organization's ability to restart critical systems seamlessly and without impact to service delivery.

Hitachi Vantara data resilience offerings are designed to improve the availability of business data. The data resilience offerings from Hitachi Vantara can be used to improve the availability of a business's applications by:

- Enabling application testing against full copy of data to simulate full production.
- Supporting parallel processing for reports, queries, analytics, artificial intelligence and data warehouse applications.

Hitachi Vantara Global Solutions Support

IBM® GDPS® Hyperswap and Basic Hyperswap Support

Hitachi TrueCopy (synchronous replication – MM) Support

Hitachi TrueCopy Consistency Groups Support

IBM MTMM (with MTIR) Support

Hitachi Universal Replicator (asynchronous replication) Support

Extended Consistency Group Support

3DC Multisite Replication with Delta Resync Support

3DC Multisite Multitarget and Cascaded Replication Support

4DC Multisite Replication Support

Compatible IBM FlashCopy® and FlashCopy Space Efficient Support

Hitachi ShadowImage Support

Hitachi Mainframe Cyber Resiliency Support

Figure 1. Hitachi Vantara Data Resilience Solutions.

- Allowing consistency checks, forensic analysis of databases.
- Performing upgrades and maintenance to hardware and software nondisruptively.
- Migrating or moving data easily and nondisruptively.
- Providing consistency across data volumes and arrays for replicated data.
- Improving the time for backup and restore of data.
- Providing protection of data in both local and remote locations.
- Surviving a major outage or disaster with the ability to continue or restart business operations.
- Reducing your RPO.
- Minimizing your RTO.

Recovery and Protection Offerings

The ability to protect and recover your data requires creation of and access to a totally separate copy of the data. Depending on the situation, the copies may be local and/or remote and provide the ability to select the copy based on location RTO or RPO.

These separate copies not only provide protection and data availability in the event of an outage or disaster but also enable testing against real data. They also enable the performance of support operations, such as backup to virtual tapes or consistency checking.

Due to the increase of cyber attacks, vulnerability from tight coupling between mainframe and open systems, the need to quickly recover from a logical corruption (or worse), we offer Hitachi Mainframe Cyber Resiliency. This solution allows you to keep multiple immutable consistent images of production data, rapidly recover from malicious attacks, validate images, check copies for possible infection or logical corruption, and much more.

Hitachi Vantara provides best-in-industry proprietary solutions for replication and replication management, such as TrueCopy synchronous, Universal Replicator, ShadowImage for Mainframe and Business Continuity Manager. Using BCM, all three of these solutions can be seamlessly integrated into enterprise mainframe environments.

Hitachi Vantara works with IBM to ensure compatibility of Hitachi Vantara replication solutions with IBM key technologies. To this end, the solutions employ Hitachi Vantara unique replication software and Hitachi Vantara enterprise-class storage support for GDPS HyperSwap. Also, as a lower-cost, high-availability solution for storage, we support Basic HyperSwap (as well as Metro Mirror Failover/Failback with HyperSwap) with IBM Copy Services Manager (CSM).

In-System Replication

Hitachi Compatible Mirroring for IBM FlashCopy provides support compatible with IBM FlashCopy V2 and supports application-aware, near-instant snapshot backups, with minimal performance impact on applications. It also provides near-instant restore capability to improve service levels and application availability (see Figure 2).

Hitachi Compatible Mirroring for IBM FlashCopy SE (FC-SE) provides an in-system solution for data protection that is IBM mainframe compatible. This solution combines snapshots,

Hitachi ShadowImage Heterogeneous Replication

- Initiate and control copy from JCL, Hitachi Business Continuity Manager, or storage.
- It is a clone of the primary volume.
- It can be triggered on time schedule.

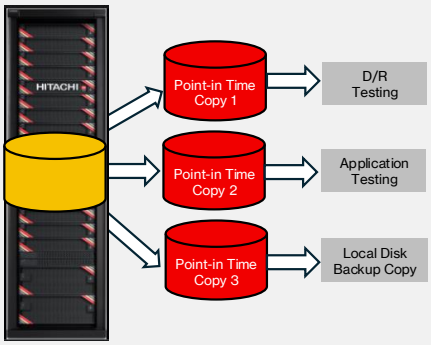

Hitachi Compatible Mirroring Software for IBM® FlashCopy® (replication)

- Initiate and control copy from mainframe using JCL (DFSMSdss, ICKDSF, etc.).
- Full volume or dataset copy.
- Enable immediate use of copy when relationship is established.
- 1 to 16 or 16 to 1 relationships.
- Incremental capability.
- Preserved mirror capability.

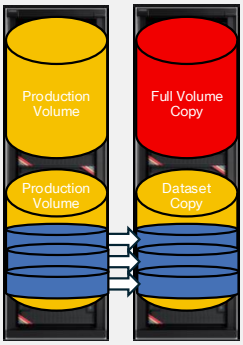

Figure 2. Local in-system replication with Hitachi ShadowImage Heterogeneous Replication and Hitachi Compatible Mirroring Software for IBM FlashCopy.

virtual data management and dynamic provisioning to create logical point in time (PiT) copies. When a FlashCopy SE source-target volume relationship is established, FC-SE uses HDP storage pools to allocate storage capacity, which is only utilized when a change has occurred within the source volume. As changed data is copied to the target volume, storage resources are consumed on an as-needed basis, depending on the amount of data that is changed and copied. With FC-SE, storage resource utilization is minimized, potentially lowering overall storage costs compared to full physical mirroring.

ShadowImage for Mainframe software provides nondisruptive, host-independent data replication to create copies of any customer accessible data within a single Hitachi storage system. These copies can provide immediate, nondisruptive access and sharing of information for decision support, test and development. They can also optimize tape backup operations. Additionally, ShadowImage copies provide nearly instant recovery from logical data corruption (see Figure 2).

ShadowImage can use HDPM storage pools to allocate target storage capacity, which is only used when data is copied from the source volume. Storage resources are consumed on an as needed basis, depending on the amount of data that is copied. Using the write time stamps allows us to perform an at-time split, providing a consistent image of the data without the need for freezing the I/Os.

Consistency groups, a standard feature of Hitachi in-system replication offerings, ensure consistency of cross-volume data when replicating or making point-in-time copies.

Multiple tiers provide the ability with ShadowImage to make first- and second-generation copies of volumes that are being copied by ShadowImage. This supports disk-based backup as well as additional copies for testing or backing up to virtual tape.

The interface to Hitachi replication solutions for enterprise storage platforms has remained consistent across the last four generations of devices. It provides investment protection and eliminates the need to build new scripts and procedures when implementing the newest generation of Hitachi enterprise storage systems.

Remote Replication

Hitachi Vantara replication solutions for mainframe environments support intermixed environments, which can include all the latest generations of Hitachi enterprise-class storage systems, including VSP One Block High End. Using Hitachi data replication solutions for mainframe, companies can implement GDPS, GDPS/HyperSwap, or Basic HyperSwap (as well as Metro Mirror Failover/Failback with HyperSwap). This includes the ability to implement 3 and 4 data center solutions with the ability to synchronize the data at alternate sites using Delta Resync if the primary site should become unavailable. This can be managed using innovative Business Continuity Manager software.

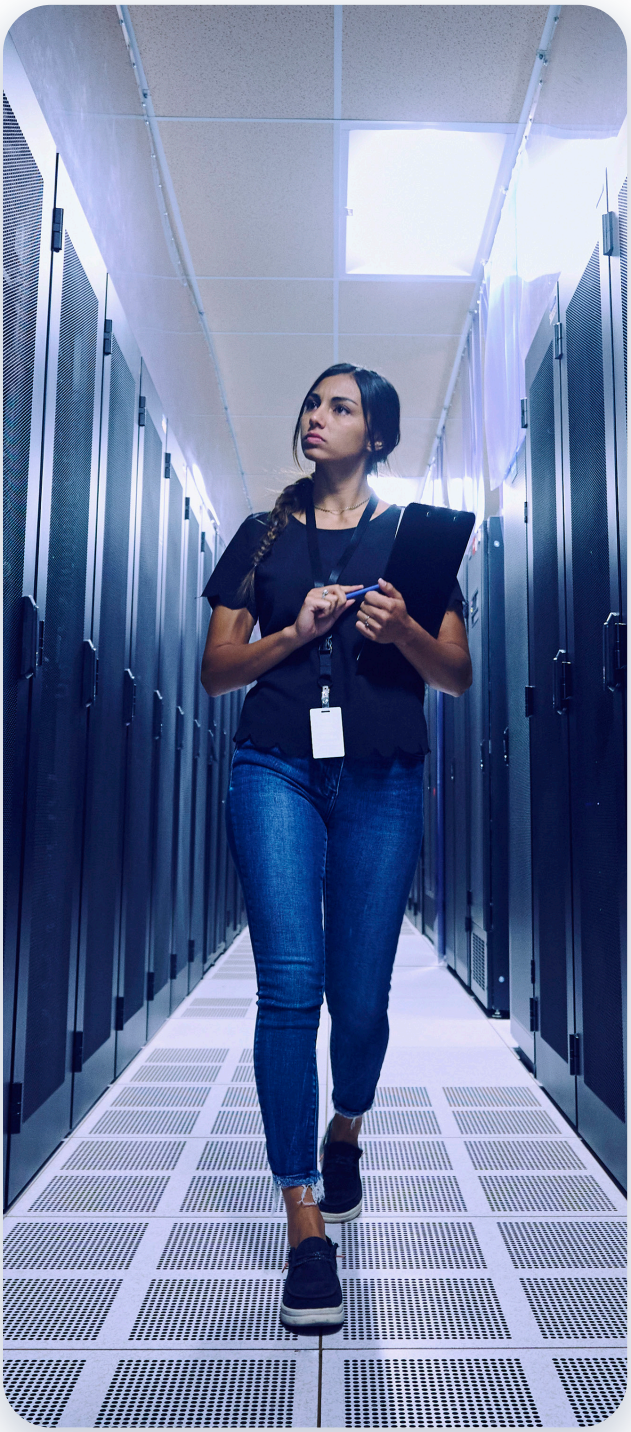

Hitachi Vantara also provides support for extended consistency groups, providing consistency of data during replication across multiple storage systems to multiple storage systems at the remote location.

TrueCopy synchronous remote replication software provides a continuous, nondisruptive, host independent, zero-data-loss solution. This rapid-restart remote data replication solution addresses disaster recovery or data migration for distances within the same metropolitan area (see Figure 3).

For mainframe environments, TrueCopy synchronous software ensures application consistency across multiple controllers by respecting write order fidelity and dependent writes.

1. Write from the host.
2. Write is sent to the remote site.
3. Write is acknowledged from remote site.
4. I/O is acknowledged to the host.

TrueCopy synchronous software:

- Offers a zero data loss solution (RPO = 0).
- Ensures remote copy is always a mirror image of source volume.
- Supports consistency across multiple storage systems (12x12).
- Works in conjunction with Universal Replicator to provide advanced 3 data center (3DC) or 4 data center (4DC) configurations.
- Up to 300km for environments that can sustain the distance impact on write response time.

Universal Replicator delivers simplified asynchronous data replication for internal and external storage attached to Hitachi enterprise-class storage platforms (see Figure 4). This software delivers enterprise-class performance associated with storage system-based replication, while providing resilient business continuity without the need for redundant servers or replication appliances. Universal Replicator provides unique RPO control using storage system journal-based replication and leverages asynchronous replication driven by the remote site to minimize impact on primary production systems. Universal Replicator supports consistency groups across volumes on Hitachi enterprise storage systems.

1. Write from the host.
2. The write I/O is acknowledged to the host.
3. The write is pulled from the remote site asynchronously.

Universal Replicator provides administrators with an innovative remote replication solution. Universal Replicator offloads most of the processing for maintaining consistency with asynchronous replication to the remote storage system using unique Hitachi pull technology. It also uses journaling to keep track of all updates at the primary site until the updates are applied on the remote storage system with write order fidelity. The journaling approach provides resiliency for a shorted RPO while protecting production performance during network anomalies. Pull technology frees up processing power at the primary site for critical production activities and, as you grow, limits the need to upgrade your primary site storage infrastructure, avoiding costs.

Universal Replicator:

- Provides the lowest RPO of the industry.
- Supports near-zero RPO while providing resiliency to minimize RPO elongation during unexpected network or workload anomalies.

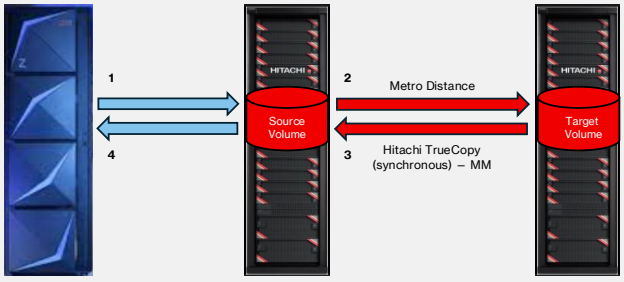

Figure 3. Remote synchronous replication using TrueCopy synchronous.

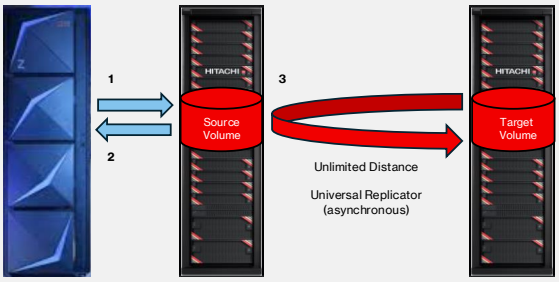

Figure 4. Remote asynchronous replication using Universal Replicator.

- Creates no impact on application response time because of distance or latency.
- Supports journal-based replication.
- Manages replication with target controllers; frees source controller resources.
- Supports extended consistency groups across multiple storage systems and between two storage complexes (up to 64x64).
- Provides advanced 3DC and 4DC configurations with extended consistency between the storage complexes in each data center (up to 16x16x16) when used with TrueCopy synchronous software.

For mainframe environments these groups ensure consistency of volumes across storage system boundaries when replicating or making point-in-time copies (see Figure 5). For two-data-center environments, these extended consistency groups can include up to 64 storage systems in each location (64x64). For 3DC and 4DC environments, they can include up to 16 storage systems in each location (16x16x16). Consistency groups and extended consistency groups use a mainframe time stamp to ensure consistency across volumes and storage controllers in the remote (Universal Replicator target) site.

3 Data Center Replication

The 3 data center replication solution uses both TrueCopy synchronous replication and Universal Replicator asynchronous replication. It provides a synchronous copy to a separate Hitachi storage system within metro-distance using TrueCopy. It provides an asynchronous copy to a separate Hitachi storage system that can be located at any distance from the primary site. When ShadowImage for Mainframe is used with Universal Replicator and its at-time-split feature, it provides consistent copies between Universal Replicator and ShadowImage (clone) copies without stopping or resynchronizing the Universal Replicator copy sessions or elongating disaster recovery RPO. There are two implementations of 3DC replication: cascaded replication and multitarget replication.

Cascaded replication provides synchronous replication from a production Hitachi enterprise-class storage system such as VSP One Block High End to a secondary storage system using TrueCopy synchronous replication software. It then asynchronously replicates the data from the secondary storage system to an out-of-region Hitachi Vantara enterprise

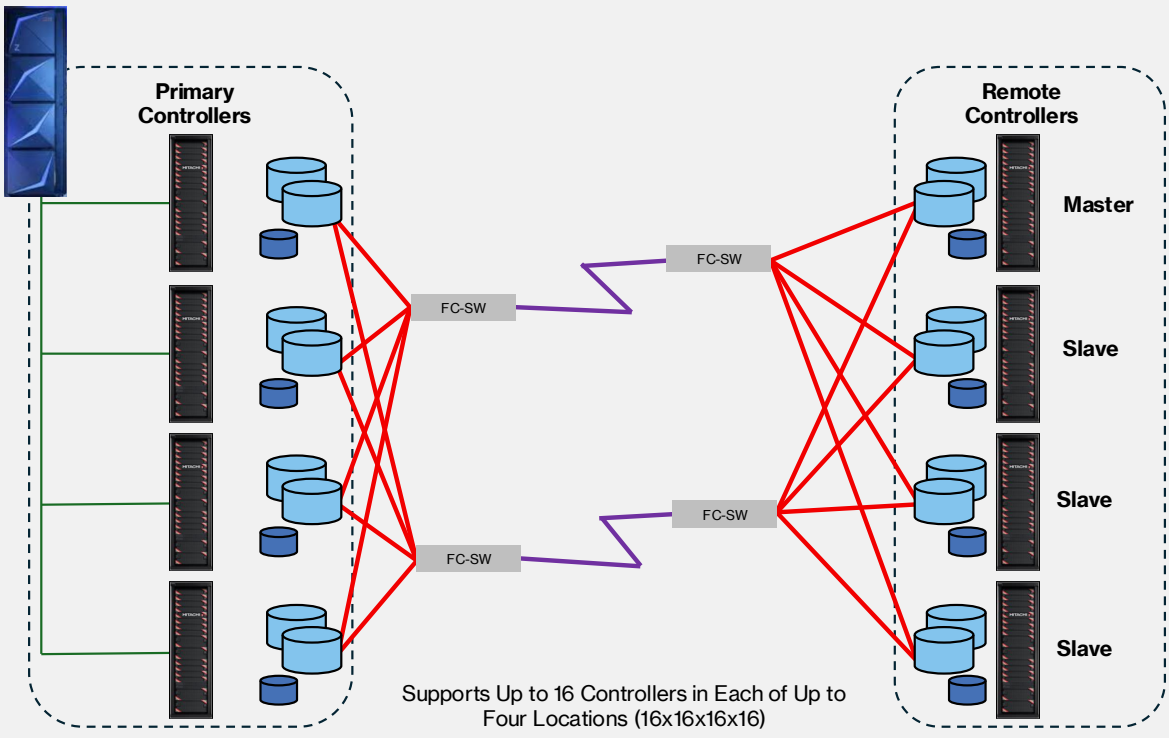

Figure 5. Extended consistency groups for mainframe.

storage system using Universal Replicator (asynchronous replication) propagating the TrueCopy (synchronous, no data loss) copy to the remote (Universal Replicator) site in the event of a production site outage or disaster (see Figure 6).

The 3DC cascade solution:

- Provides a synchronous mirror image of primary data and an asynchronous image at an out-of-region site.
- Enables fast recovery using the in-region data center.
- Provides out-of-region copy to allow recovery from a regional disaster.
- Supports a no-data-loss solution when losing primary site.

Multitarget (Concurrent) Replication synchronously replicates data from the production site using TrueCopy synchronous software to a local site (within metro distance). It asynchronously replicates to an out-of-region remote site using Universal Replicator for optimal data protection. With multitarget replication, protection can be resumed with no data loss between the two remaining sites (hot standby site and remote site) using Delta Resync (see Figure 7).

The multitarget replication solution:

- Provides campus or metro copy plus out-of-region copy.
- Enables fast recovery and excellent data currency for local site failures combined with advanced protection from regional disasters.
- Resyncs, using Delta Resync, the out-of-region asynchronous copy with the metro synchronous copy in event of primary site outage.

Delta Resync provides additional disaster recovery protection. With a multitarget 3DC configuration as described above, the production site mirrors data synchronously to a hot standby site using TrueCopy synchronous replication. It mirrors the data asynchronously to a third remote site using Universal Replicator. Should the production site experience failure, the hot standby site takes over and uses the remote site for disaster recovery. It only copies differential data, the updated data that has not yet been propagated to the remote site. A full copy of the data is normally not required. With Delta Resync, only the changes not recognized between the hot standby and remote site, as received at the remote site when the outage occurred, are transferred from the hot standby site (new production site) to the remote site (new disaster recovery site). As a result, disaster recovery protection for the hot standby site resumes within minutes instead of hours.

Multitarget Metro Mirror synchronously replicates data from the production site using TrueCopy synchronous software to two local sites (within metro distance). With multitarget Metro Mirror, protection can be resumed with no data loss between the two remaining sites (hot standby site and remote site) using the IBM MultiTarget Incremental Resynchronization (MTIR) capability (see Figure 8).

The multitarget replication solution:

- Provides 2 campus or metro copy and can be coupled with an out-of-region copy.
- Enables fast recovery and excellent data currency for local site failures, combined with advanced protection from regional disasters.

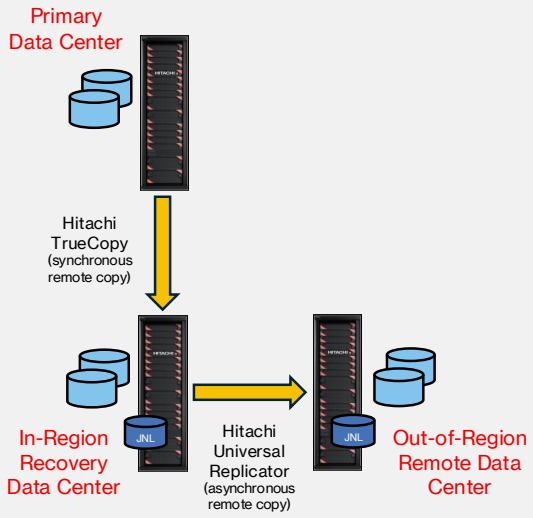

Figure 6. 3DC cascade solution with TrueCopy synchronous and Universal Replicator.

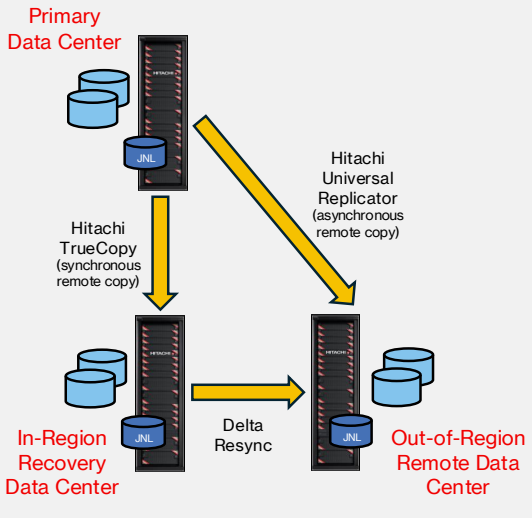

Figure 7. 3DC multitarget remote replication with TrueCopy synchronous and Universal Replicator.

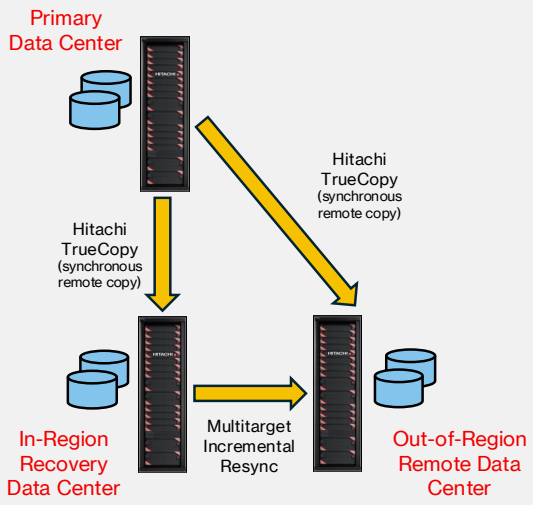

Figure 8. 3DC multitarget Metro Mirror with TrueCopy synchronous.

At-Time Split is a feature that eliminates the need to freeze I/Os at primary site or Universal Replicator remote site to create a consistent copy of data. In conjunction with Universal Replicator, it eliminates the need to split the Universal Replicator pairs when creating a tertiary copy with ShadowImage. Rather, the at-time split operation is performed on the specified remote ShadowImage copy group and creates a data-consistent snapshot at a specified time while sustaining the Universal Replicator copy group in a duplex state. At-time split uses mainframe write timestamps to ensure consistency when creating the tertiary copy.

4 Data Center Replication

The 4DC replication solution uses both TrueCopy synchronous replication and Universal Replicator asynchronous replication (see Figure 9). It provides a synchronous copy to a separate Hitachi storage system within metro-distance using TrueCopy. It provides an asynchronous copy (with Delta Resync to protect from failures on one primary site) to a separate Hitachi storage system that can be located at any distance from the primary site. The storage at the remote site is also primary for synchronous replication. When ShadowImage for Mainframe is used with Universal Replicator and its at-time split feature, it provides consistent copies between Universal Replicator and ShadowImage (clone) copies without stopping or resynchronizing the Universal Replicator copy sessions or delaying disaster recovery RPO.

IBM HyperSwap Process

The synchronous replication technology has been implemented at most customers for a very long time. Compatibility with IBM standard Metro Mirror (formerly PPRC) has been the foundation for HyperSwap.

Running a set of programs on the host enables HyperSwap to happen: Copy Services Manager (CSM) or GDPS.

The microcode of the VSP One Block High End supports native PPRC commands.

In normal circumstances, production runs on an IBM Z family server (formerly zServer®) on primary site and data is synchronously replicated to a secondary site, usually below 10km distance to minimize performance impact of the replication.

In synchronous replication, the write for the Z family server is written in primary VSP One Block High End cache, transmitted and written to secondary VSP One Block High End cache which acknowledges the write to the primary, itself acknowledging the write to the Z server.

Using this technique allows a for a recovery point objective of zero.

The latency (round trip time) in glass is 1 ms for 100km, and all writes are impacted as per synchronous replication (see Figure 10).

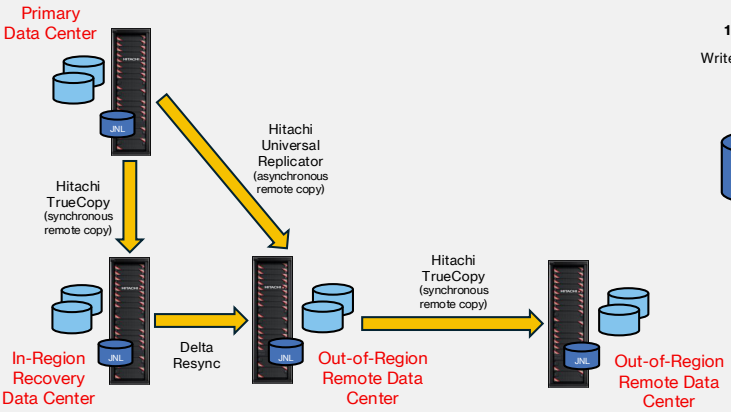

Figure 9. 4DC multitarget Metro Mirror with Hitachi TrueCopy (synchronous capabilities).

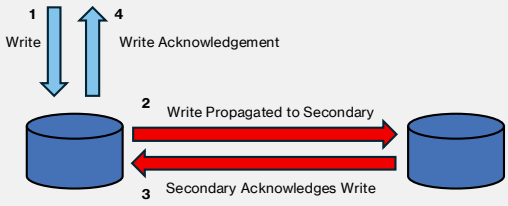

Figure 10. Synchronous replication process.

In the modern disk subsystems, writes are always done in cache (write cache hit 100% of the time), with data being destaged asynchronously to the physical media, so it improves the response time.

With VSP One Block High End, synchronous replication is possible up to 300km; however, most customers are within or below the 10km range, as asynchronous replication is preferred for long distance.

Automated Disaster Restart — HyperSwap support

To perform HyperSwap, CSM or GDPS uses the possibility to swap the UCBs (uncontrolled capacity buffers) by using system services.

In order to prepare the UCBs (to be swappable), the synchronous replication has first to be active, i.e. duplex, and the IBM Z systems (desendents of the IBM zSystem and zSeries® products) can see the primary volume and its target.

During HyperSwap process (planned or unplanned), the I/Os are frozen for a few seconds, and devices addresses are swapped; so the I/Os are redirected, and the replication is also reversed.

But the reality is that operations are more complex, as the following sequence of operations must be completed:

- Freeze happens upon a trigger.
- Replicated pairs are suspended.
- Devices go to long busy state.
- PPRC paths are reversed.
- UCBs are swapped in memory.
- HyperPAV aliases are reassigned.
- I/Os are resumed to the previous secondary storage, which has been promoted to new primary.

The secondary has been promoted to the new primary. The replication pairs are generally suspended after a HyperSwap (at least after an unplanned HyperSwap).

With GDPS it is possible to do a HyperSwap resync for a planned HyperSwap; i.e., after the HyperSwap is completed, the PPRC pairs are reversed and duplex.

This process is transparent from the application standpoint.

When possible, the replication can be resumed between the new primary and the new secondary.

A planned HyperSwap allows the failback to the initial configuration.

There are many HyperSwap triggers, including an I/O error on a disk, power loss of primary disk subsystem and very high response time of primary storage ....

IBM command set allows the simulation of unplanned HyperSwap.

The Different HyperSwap Technologies

Depending upon the level of protection required (and associated costs to implement), there are different ways to meet customer needs:

CSM:

- Copy Services Manager Basic Edition for zSystems
- Copy Services Manager for zSystems

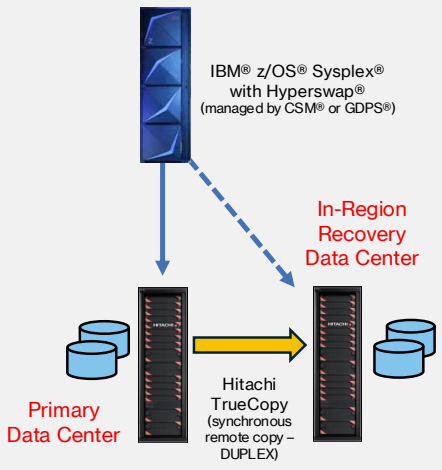

Figure 11. Usual production environment.

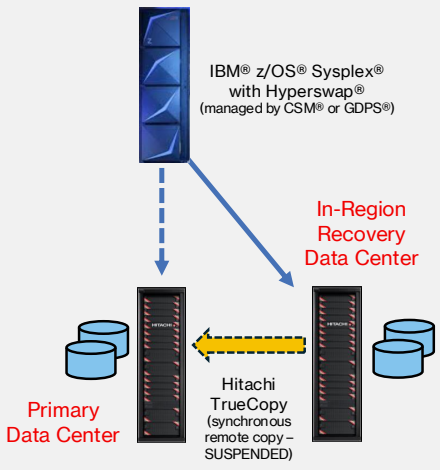

Figure 12. Usual configuration after HyperSwap.

GDPS – several versions, with the most common being:

- HyperSwap Manager
- MM
- MTMM

The Copy Services Manager Basic Edition for z Systems comes as part of z/OS.

For information, the maintenance is not included in the maintenance fees, it must be taken separately.

It allows you to manage HyperSwap, and thus, to protect against disk access failure.

The Copy Services Manager for z Systems is a licensed product that delivers two advantages over the basic edition:

- A freeze policy, protecting against remote path failure with suspending the pairs (this can be done with the basic edition by capturing the IOS message when the last remote path is lost).
- A standby CSM that allows you to manage CSM when the primary CSM address space or server is not available.

Despite the process involved being the same, there is no upgrade path between CSM and GDPS.

GDPS family is broader, allowing storage outage protection (HyperSwap) to dynamic site switching.

In the case of GDPS family, the HyperSwap is done by a specific system called controlling system.

The difference between CSM and GDPS is the large customization allowed in the freeze policy, which lets users have many choices on how to react on a trigger.

- GDPS HyperSwap Manager allows organizations to not just manage HyperSwap like CSM would do but also upgrade to a broader version of the GDPS family.
- GDPS MM allows administrators to HyperSwap z/OS, and, using a proxy, z/VM (usually Linux for Z is running under z/VM).
- GDPS MTMM allows administrators to manage and HyperSwap a PPRC multitarget configuration, i.e. two synchronous copies out of the same source volume.

Hitachi Vantara disks can integrate with GDPS MM, MTMM, 3DC and 4DC.

Hitachi Vantara Storage Technologies and HyperSwap

Hitachi Vantara for many years has participated in the IBM qualification program from GDPS.

At the end of the qualification, IBM publishes a qualification letter on the IBM website (https://www.ibm.com/it-infrastructure/z/technologies/gdps qualification section).

The GDPS MM (formerly GDPS/PPRC) is not only qualification for HyperSwap, but also for GDPS 3DC with Universal Replicator Delta Resync managed with BCM as well as 4DC.

IBM also qualified GDPS MTMM with MTIR on Hitachi storage.

Please also note that Hitachi Vantara is also going through an IBM basic connectivity qualification so customers can be fully confident with Hitachi Vantara compatibility with IBM Z series.

Hitachi Vantara has many large customer references who are running HyperSwap on Hitachi storage, so please contact your local representative to learn more.

The IBM compatibility list is available at: https://www.IBM.com/support/pages/node/656273.

The GDPS qualification for 2DC, 3DC Delta Resync, 4DC and MTMM can be found at: https://www.IBM.com/downloads/cas/JR5J4LWY.

It is the intention of Hitachi Vantara to continue participating in the qualification process, and we're currently proceeding toward basic and GDPS qualifications for VSP One Block High End.

Replication Management

Hitachi Vantara provides flexibility in storage management for mainframe environments with both host-based and storage-system-based storage management solutions. Business Continuity Manager provides a host-based solution that uses familiar TSO/ISPF screens to view and manage storage replication. With BCM, the storage administrator can set up and manage Universal Replicator, ShadowImage for Mainframe and TrueCopy synchronous remote replication software operations.

Business Continuity Manager, a z/OS-based management software solution from Hitachi Vantara, automates complex disaster recovery and planned outage functions, dramatically reducing recovery times. BCM provides auto discovery of enterprise-wide storage configuration and replication objects, thereby eliminating tedious, error-prone data entry that can cause outages, as well as centralized management of Universal Replicator, TrueCopy and ShadowImage software solutions, wherever they reside. BCM uses ISPF, or optionally, REXX scripts, WEB API calls (REST API type call with JSON format response) or a GUI to manage these replication offerings. In-band BCM commands can be used to manage replication at the remote site from the primary site. The BCM monitor allows administrators to monitor pairs and act should the status be incorrect.

Hitachi Mainframe Replication Extensions, a z/OS-based software solution, allows administrators to monitor and export the asynchronous replication RPO and issue an alert should the SLA be violated.

Virtual Tape Library Support

Hitachi Vantara has several partners in that area, with each one having his own advantages and specificities. SecureAgent®

Secure Data Solution® (SDS), Luminex™ Channel Gateway™ and Rocket® Software provide virtual tape library solutions for IBM Z (formerly zSeries), enabling writes to midrange or Amazon S3® storage. Log-On Software™ VTFM-NewGen runs on IBM Z Integrated Information Processor (zIIP) and can use midrange or S3 storage as well and offers the capability to share virtual tapes. BMC AMI® Cloud runs on zIIP; the tape management software must be replaced by the embedded one and writes on S3 storage.

These solutions greatly reduce or eliminate a data center's dependence on physical tapes by transparently replacing high-cost tape subsystems, including automatic tape libraries (ATL) and virtual tape systems (VTS) with new, low-cost storage. They also can provide a very competitive solution to Transparent Cloud Tiering without being tight to the vendor. Using S3 (object) storage also brings immutability and performance at lower cost, as well as immutability and expiration management.

Hitachi Vantara works with and provides solutions from SecureAgent Software, Luminex for virtual tape library (VTL) solutions, Rocket Software, Log-On Software and BMC AMI Cloud. SecureAgent SDS is an "all in one" solution that supports open systems, IBM iSeries® and all IBM zSeries operating systems, including z/TPF. Luminex Channel Gateway supports z/OS, z/VM and z/VSE. Rocket Software, Log-On Software and BMC support z/OS.

All these solutions provide lower tape processing costs, improved end-user service levels and professional services to assist enterprise mainframe customers with their tape migration. Data can be stored to the cloud or to an object storage solution like Hitachi Content Platform (HCP).

Backup to the Cloud Support

Hitachi Vantara is committed to helping organizations lower their operating costs. More and more, customers have embraced cloud solutions. We work with BMC (BMC AMI Cloud — ex Model9) and Log-On Software to provide administrators with a solution that lowers million service units (MSU), CPU consumption and thus software licensing cost for the data movement for tape processing. It copies the data over IP to an Amazon S3 compatible unit, so backup data can be hosted in the cloud or on-premises — for example, on HCP.

There is no need to change current backup JCLs to write on the HCP. The link is made through flexible automatic class selection (ACS) routines.

For this specific backup to HCP, the IBM DFSMSdss™ is replaced by Oracle Java® code running on a zIIP without any MSU cost, reducing the software bill for z/OS.

HCP is an object storage, and as such, allows versioning and protection of data from any changes.

A recent customer reduced backup time with HCP, compared to the company's VTL, by 17-fold.

We believe that this is a better solution for customers than transparent cloud tiering as it is vendor agnostic: It does not create a vendor lock-in situation, and it gives much better throughput.

Cyber Resiliency Solution

Taking advantage of various Hitachi technologies, Hitachi Vantara has implemented a solution to help organizations recover from logical data corruption. The solution, based on Business Continuity Manager, triggers a consistent copy of production data using the ShadowImage at-time split capability.

Modifications since the last cycle are stored into a virtual volume, which is external to the mainframe storage and physically stored in a virtual volume located in an open systems storage internal or external to the mainframe storage. This open system storage format prevents data in the fortress from being modified or deleted before the expiration date is over. The fortress keeps the number of cycles that organizations are planning to keep available.

When the last cycle modifications are stored in open system volumes, Ansible templates are triggered to index/store the modifications in the fortress using Hitachi Thin Image Advanced technology. When this last process is finished, the mainframe is warned that the last cycle is stored in the fortress.

Hitachi Mainframe Cyber Resiliency allows customers to choose an image and to expose it to mainframe to check the data and do forensic analysis. Should it be needed, partial data can be restored to production, or restoration of all data can be completed.

Data can be restored to another storage or eventually a different location, assuming Fibre Channel connectivity is provided. The fortress can be up to 100km away from its storage connection.

From security standpoint, the Mainframe Cyber Resiliency Solution has been engineered as zero trust architecture. Any modification that would involve a change that would have an impact on the integrity of the current image (changing cycles, changing retention period, for example) would require a two-user acknowledgement for security purposes.

For cyber security protection, Hitachi Vantara is partnering with MainTegrity®.

Data Mobility and Migration

Hitachi Vantara is committed to ensuring that data is available at the right place and at the right time. Hitachi Vantara provides solutions for moving and/or migrating information without impacting the performance of the system. These solutions ensure continuous access, business continuity and data

resilience. They enable organizations to make the most efficient use of storage assets and optimize the deployment of critical data without impacting business operations. Hitachi Vantara offers solutions for movement of data in mainframe environments for migration or just for data mobility.

Cross-OS File Exchange

Cross-OS File Exchange software delivers high-performance, high-reliability data sharing between heterogeneous host platforms and moves vast amounts of data quickly between mainframe and UNIX® (IBM AIX®, HP-UX® and Oracle Solaris®), Linux (RedHat® and SUSE®) and Microsoft® Windows® environments without tying up networking resources or intermediate tape media. File transfer speeds are amplified 5 to 10 times. It enables data stored on the Hitachi storage systems to be converted and transferred between mainframe and open system platforms and between different open system platforms.

Integration With Databases and Applications

Hitachi Vantara provides tight integration between storage-system-based software and utilities with application and database functions to optimize resources, minimize complexity and provide efficient and effective management of storage resources.

Hitachi Vantara has strong partnerships with companies such as BMC, 21st Century Software®, and other independent software vendors, to strengthen Hitachi storage management and effectiveness in mainframe environments.

Security: Data-at-Rest Encryption (DARE), Root of Trust

Data-at-rest encryption offers the possibility to encrypt the data stored on Hitachi storage. There is the possibility to use an external key management server to generate and store the keys or to back up the keys generated into the storage. The keys can also be self-generated in the storage system. There is one unique key per encrypted internal drive (SSD). The keys are AES 256 bit. Cryptographic erasure (media sanitization) of data is performed when an internal encrypted drive is removed from the storage system. The data-at-rest encryption or DARE functionality is implemented using cryptographic chips included as part of the encryption hardware. VSP One Block High End is compliant to FIPS 140-3 Level 2. Hardware root of trust (HRoT) and secure boot has been implemented for a more secure environment and control.

Sustainability

Hitachi Vantara has a track record of publishing the green gas emission of the storage through industry recognized organizations. The figures are not from Hitachi Vantara but from the measurements from these organizations and, as such, trusted figures and not marketing figures.

We also have a track record of shrinking green gas emissions by 30% to 40% from one generation to the next.

To help with energy savings, Hitachi Vantara has introduced Titanium plus power supply.

VSP One Block High End is also reusing up to 60% of recycled parts, which is also unique in the industry.

Hitachi Vantara Technology Partners

21st Century Software

Essential Mainframe Migration for z/OS (EMMz) offers the possibility to migrate mainframe storage nondisruptively to Hitachi storage.

BMC and CompuWare®, a BMC Company

The Compuware INNOVATION Data Processing FDR® (Fast Dump Restore) Suite offers business resiliency for z/OS customers with solutions that help to eliminate business downtime for Hitachi Vantara customers. The suite applies IBM FlashCopy data replication technology to the basic storage management operations of backup, restore and disaster recovery.

- FDRPAS® allows nondisruptive migration from old storage to Hitachi storage.
- The BMC AMI Cloud (ex Model9) solution eliminates the dependency on virtual tapes for mainframe data backup and archive and reduces the CPU and I/O overhead associated with large tape environments, such as virtual tape management, recycling and expiration.
- BMC AMI Security is an umbrella for many different tools and products to ensure and enforce security.
- BMC AMI Data is a toolbox for managing files and data bases.
- BMC AMI Ops allows to manage performance, availability, automation, storage resources and capacity management.

Broadcom® (Vantage™)

Broadcom Mainframe Software Division empowers customers to amplify the value of their mainframe investments. Vantage allows centralized monitoring and policy-based automation for your mainframe storage infrastructure.

Hitachi Vantara storage can be managed using Vantage.

Brocade®, a Broadcom company

With Brocade Fibre Channel technology-based directors and switches from Broadcom, you've got the firepower to deliver high-performance connectivity across the data center and across the globe. Scale your network on demand — move more data more places — as you keep costs of ownership reined in.

Cisco®

Data must be secure, accessible and always available. Build and operate a storage network that can do it all with Cisco MDS 9000 Series.

Data Kinetics®

Tablespace allows DB2 and VSAM ultra-fast access in memory rather than fetching data from disk. This brings many benefits, including shrinking transaction time, reducing MSU consumption, improving the batch and so forth.

EPS®

Pivotor is aware of the Hitachi Vantara MAR records. Hitachi Vantara is teaming with EPS for independent performance studies.

Infotel

Infotel provides smart DB2 tools as well as archiving tools, which respects the regulation constraints.

IntelliMagic™ (IntelliMagic Vision™)

IntelliMagic delivers its solutions to Fortune 500 companies and some of the world's largest IT data centers. IntelliMagic Vision for z/OS unlocks the full potential of mainframe infrastructure performance and configuration data by automatically applying z/OS-aware expert knowledge. Its modernized, intelligent interpretation detects risks before they impact production, uncovers true root causes and identifies optimization opportunities. Ultimately, this enables a higher level of application service reliability at optimal cost. Intellimagic has been acquired by IBM.

Log-On Software

Log-On Software offers a smart VTL and TCT alternative that runs on zIIP, does not require replacement of tape management software and offers PAT, which allows multiple concurrent accesses to the same virtual tape. The target can be traditional Fibre Channel attached storage or S3 compatible storage (HCP, cloud).

Luminex

Luminex Channel Gateway virtual tape solutions deliver support for z/OS, z/VM and VSEn. All major tape application and tape management systems are supported. Luminex offers several options, including compression, Luminex replication, replication monitoring, tape monitoring and allocation control systems, plus encryption key management. Professional services offer a smooth migration from customer current virtual tape system.

MainTegrity (FIM+)

MainTegrity is an innovator, delivering file integrity monitoring software (FIM+) to improve IBM mainframe cybersecurity. FIM+ strengthens customer defenses in ways never before possible, while interoperating with existing security tools. It provides whitelists, backup verification and automated forensics to combat ransomware and other malicious attacks. FIM+ allows users to quickly detect unwanted encryption. FIM+ delivers improved compliance with NIST, PCI, GDPR and bank resiliency requirements. FIM+ also provides deployment audit and integration with the DevOps toolchain.

Precisely® (Ironstream®)

Precisely is a global software company specializing in big data, high-speed sorting products and data integration software and services for Hadoop®, Microsoft Windows, UNIX, Linux and mainframe systems.

IronStream® ingests Hitachi storage SMF records (recordings of the internals of the storage) and these are used to trigger alerts and can merge with mainframe records for deep performance analysis. The ingested records can be sent to Splunk® or equivalent tools in real time.

SAS-MXG®

With SAS, you can deliver analytic results to all users to share insights and drive fact-based decisions. SAS-MXG ingests Hitachi storage SMF records (Mainframe Analytics Recorder, which brings access to the data recordings of the storage internals).

Rocket Software

Rocket Software has a large collection of mainframe software that can bring solutions for customers.

SecureAgent

SecureAgent SDS solution delivers a fully integrated capability for IBM zSeries and commonly installed tape management systems, providing a mainframe virtual tape solution. The SDS solution addresses organizational requirements for improved tape efficiencies, regulatory compliance and governance. It also improves compression, encryption, automation, file authenticity, secure long-term preservation and retention of mainframe content.

Summary

Hitachi Vantara is committed to providing industry-leading enterprise-class storage for IBM mainframe computing environments. For over 40 years Hitachi has produced and delivered storage with leading edge function and capabilities. Hitachi has a strong relationship with IBM to ensure their ability to support the latest features and functions of z/OS when they are delivered. Additionally, Hitachi has developed and continues to develop new technology and capabilities to enable customers to maximize the efficiency and utilization of their storage and the availability of the data.

Technologies such as Hitachi Dynamic Provisioning, along with virtualization of both internal and external storage, enable customers to get the best application performance and automated management of their storage tiers. In addition, Hitachi Vantara business continuity solutions are qualified by IBM to work with the latest releases of z/OS and GDPS, and provide organizations with solid, tested recovery capabilities.

Recently, Hitachi Vantara has introduced a Mainframe Cyber Resiliency Solution to offer a solution that offers a real air gap as the mainframe data is stored in a fortress residing on open systems storage, which is virtualized on the mainframe storage. This solution allows customers to have multiple consistent images from production, and perform checking, forensic analysis, surgery or full restoration of the production data from selected images.

Hitachi Vantara (subsidiary of Hitachi Ltd) provides leading-edge technology and storage solutions for both large and not so large enterprise customers with IBM mainframe infrastructures.

About Hitachi Vantara

Hitachi Vantara is transforming the way data fuels innovation. A wholly owned subsidiary of Hitachi, Ltd., we're the data foundation the world's leading innovators rely on. Through data storage, infrastructure systems, cloud management and digital expertise, we build the foundation for sustainable business growth.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
hitachivantara.com | community.hitachivantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
hitachivantara.com/contact

HV-BTD-WP-Mainframe-Storage-Compatibility-and-Innovation-27Aug25-A