


White Paper

Hitachi VSP 5000 Storage – Innovative and Compatible

With Hitachi Virtual Storage Platform 5000 Series



White Paper

Table of Contents

03	Executive Summary
04	The Hitachi Vantara Two-Pillar Approach to Mainframe Storage
06	Enhanced Storage Features for Today's Mainframe Environments Information Availability and Data Resilience
18	Summary

Executive Summary

With over 40 years of experience supporting IBM® mainframe environments, Hitachi Vantara is committed to continued support of these environments. We provide new innovative storage solutions designed to improve storage processing, performance, availability, recoverability and management in mainframe environments.

Hitachi Vantara introduced a 100% data accessibility warranty when the first generation of Hitachi RAID was rolled out to the market in 1995 (7700 Classic).

Hitachi Vantara has announced a new solution to allow enterprises to store immutable images of the production data in a fortress should the company have a logical corruption or, worse, be the subject of a cyberattack. This solution also allows the organization to mount stored images to check the data periodically.



Hitachi Vantara is a leading provider of mainframe storage. Hitachi, Ltd. and Hitachi Vantara (Hitachi Vantara is a subsidiary of Hitachi, Ltd.) use a two-pillar approach in the design and development of IBM mainframe-compatible storage systems.

The first pillar is compatibility: the focus on support and qualification of essential IBM mainframe storage features. The second pillar is innovation: development and testing of unique Hitachi storage and storage management features to provide mainframe customers with additional capability and value.

IBM successfully completed qualification tests of Hitachi Vantara Virtual Storage Platform (VSP) with IBM GDPS using Hitachi support for Metro Mirror® (MM or PPRC), with IBM FlashCopy® and with Hitachi Universal Replicator (HUR) 3 data center (3DC) delta resync. (Hitachi 4DC was also qualified with GDPS). GDPS Multitarget Metro Mirror (MTMM) has also been qualified by IBM.

Hitachi VSP 5000 series IBM qualification was completed in 2025 with GDPS 4.7. Previously, an IBM basic connectivity and features qualification with IBM z15® was carried out. Hitachi has participated in this IBM qualification process since 2005.

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

Hitachi Vantara has a installed base of storage systems connected to IBM Z® family mainframes via IBM FICON® networks as well as mainframe-based Linux environments via Fibre Channel networks. Hitachi Vantara provides high-speed flash storage for mainframe environments with NVMe connectivity. 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 storage innovation, IBM mainframe compatibility, performance and data resilience, your Hitachi storage system will meet the demands of your mainframe and open systems storage environments.

The Hitachi Two-Pillar Approach to Mainframe Storage

Hitachi and Hitachi Vantara (subsidiary of Hitachi, Ltd.) design and develop 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 40 years, Hitachi, Ltd., has provided mainframe technology to business and industry, including solutions for IBM z/OS®, z/VM®, z/VSE®, and z/TPF, and Linux for 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 IBM's PPRC, XRC, FlashCopy®, FlashCopy Space Efficient, Remote Pair FlashCopy (Preserve Mirror), Parallel Access Volumes (PAVs), HyperPAV, SuperPAV, Multiple Allegiance (MA), High-Performance FICON® (zHPF), zHyperWrite, Copy Services Manager (CSM) Basic HyperSwap and IFCES/EDiF. 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, Ltd., and 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 Platform 5000 series, VSP G1500 and VSP F1500 are qualified and supported on the IBM Z series processors, with the latest qualification being z16 processors. They are also supported under the latest z/OS version (3.1 at the time of this writing).

Hitachi has implemented technologies that are compatible with IBM's Metro Mirror (PPRC), Multi Target Metro Mirror (MTMM), FlashCopy V2 and FlashCopy Space Efficient on Hitachi storage systems. We are the first storage vendor to go through this broad qualification. IBM also successfully completed the qualification testing for GDPS/HyperSwap, FlashCopy, and 3 and 4 Data Center (3DC and 4DC) multitarget replication configuration using VSP 5000 series GDPS/HyperSwap support and Hitachi Universal Replicator with delta resync along with Hitachi Business Continuity Manager. IBM also qualified MTMM 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 will be capable of supporting these new features. In addition, Hitachi has its own testing facilities to ensure that Hitachi-delivered products and features are compatible with the IBM solutions.

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

- FICON 32Gbit
- FICON 32Gbit with IFCES (authentication & EDiF)
- Forward Error Correction
- Read Diagnostic Parameters (RDP)
- FICON Dynamic Routing (FIDR)
- High Performance FICON (zHPF)
 - zHPF single/multi-track
 - 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
- Enhanced Dynamic cache management
- Cache Fast Write (CFW) and DASD Fast Write (DFW)
- Sequential Data Striping
- PDS Search Assist
- MIDAW
- Priority I/O Queuing
- Multiple Allegiance
- Parallel Access Volumes (PAV)
- HyperPAV
- SuperPAV
- Thin provisioning
- Extent Space Efficient volumes
- Extended Address Volume (EAV, up to 1TB disks)
- Erase on scratch (multitracks)
- User Directed Space Reclaim
- Dynamic Volume Expansion (DVE)
- Basic HyperSwap with Copy Services Manager
- Geographically Dispersed Parallel Sysplex™ (GDPS) HyperSwap
- Query Host Access
- Soft Fence
- Hardware Reserve Transfer
- Dynamic Path Management (GDPS)
- PPRCSUM (GDPS)
- zHyperWrite
- IBM® Metro Mirror (PPRC)
- IBM Multi-Target Metro Mirror (MTMM)
- 3DC with Delta Resync
- 4DC with Delta Resync
- MTMM with URz
- FlashCopy V2 (including multiple incremental FC)
- FlashCopy Space Efficient (FlashCopy SE)
- Remote Pair FlashCopy— Preserved Mirror
- Concurrent Copy
- z/OS® Discovery and Auto Configuration (zDAC) support
- DB2® Cast Out Accelerator
- FCP support
- LinuxONE support

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

Hitachi Vantara and IBM have a long-term agreement to provide assistance to each other to quickly resolve any joint customer issues related to Hitachi storage and IBM mainframes and 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 products are identified as the source of the problem.

Innovation

Hitachi innovation, the second pillar of Hitachi Vantara 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 of the Fortune 100 organizations today. These represent some of the most critical mainframe environments in the world.

Data in Place Non-Disruptive Migration

Hitachi Vantara has a proven history of availability. The VSP 5x00 series offers eight 9's of availability (without GDPS). One of the major risks for customer data is the migration to a new generation of storage. This time-consuming task and complex planning can be eliminated by the capability of the Hitachi storage to do a data-in-place nondisruptive upgrade from one generation to the latest generation. For example, VSP 5500 can be upgraded nondisruptively to a VSP 5600. This will ensure an upgrade to our next-generation storage architecture, protecting customer investment.

Cyber-Resiliency Solution

Hitachi has engineered a solution to protect enterprises from logical corruption, or worse, a cyberattack. The unique features of the VSP 5000 series allow organizations to make an image of production without the need for an I/O freeze. Virtualization, another unique feature of the VSP 5000 series, allows this delta copy triggered on mainframe storage to reside on open systems storage, creating a true air gap, and provides another layer of protection from attacks as it is unknown by the 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. 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 on different storage from the primary storage. For security purposes, any critical change does require a two-user acknowledgement.

Hitachi Dynamic Provisioning for Mainframe

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 the existing mainframe provisioning process while improving and simplifying performance and capacity optimization. HDPM automatically spreads the content of application datasets across a storage pool, containing many physical media (from SAS physical disks to SSD-based disks with NVMe interfaces as well as SCM), distributing I/O activity across available physical resources and back-end paths. This action eliminates the challenge of manually spreading application data over many back-end physical media to optimize performance and throughput. At the same time it reduces storage management concerns.

Hitachi Dynamic Tiering for Mainframe

Hitachi Dynamic Tiering for Mainframe (HDTM) is an extension of dynamic provisioning. HDTM offers an additional level of automated, optimized storage management by managing data across a full range of storage tiers from high-performance storage to low-cost storage. Data is automatically moved between tiers according to simple policies. If data is frequently being moved from a low-performance tier to cache, then the data is automatically moved to a high-performing tier to optimize performance.

HDTM complements mainframe storage provisioning processes, such as DFSMS. At the same time it offers the full benefits of Hitachi Dynamic Provisioning for Mainframe to improve and simplify performance and capacity optimization. Existing DFSMS storage groups and ACS routines can be aligned to different tiered storage pools.



Policy-Based Management With Hitachi Tiered Storage Manager for Mainframe

Hitachi Tiered Storage Manager (HTSM) for Mainframe is an IBM z/OS software management product that enables control of service levels. The controls, based on performance and/or time to facilitate meeting mainframe service level agreements (SLAs) in a multitiered configuration, place data on the best tiers to meet requirements. HTSM for Mainframe allows DFSMS storage groups to match with the storage tier/s.

Mainframe Analytics Recorder

Ever since the mainframe analytics recorder capability we included in Hitachi Storage Virtualization Operating System (SVOS), we have delivered storage-system-specific detailed performance integrated with other SMF/RMF reporters. This capability is key to providing system programmers with an integrated view of the internal performance of their mainframe storage. 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.

The mainframe analytics reporter capability 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 other virtualized storage. Hitachi Mainframe Analytics Recorder (HMAI) can be used to visualize the mainframe analytics recorder data using Grafana.

Hitachi Universal Replicator

With Hitachi Universal Replicator, administrators have an innovative remote replication solution. HUR offloads most of the processing for asynchronous replication to the remote storage system using journaling to keep track of all updates. It also uses unique Hitachi pull technology. This frees up processing power at the primary site for critical production activities. It provides the best performance available, as well as the capability to continue to replicate with degraded intersite bandwidth and offers the lowest available RPO on the market. Hitachi has also developed support for extended consistency groups to provide consistency of data cross storage systems during replication.

Enhanced Storage Features for Today's Mainframe Environments

Information Availability and Data Resilience

Data is at the core of your data center, and any data center must provide for the movement, provisioning, access, and protection of data that is provided by storage systems. An important element of your information availability strategy is ensuring your data is resilient (see Figure 1).

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. 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 (SLAs) and requirements for high availability, with specific recovery time objectives (RTO) and recovery point objectives (RPO). Solutions to provide data resilience, including data recoverability, should be in place to counter any logical corruption that can put the company at risk.

Hitachi Vantara Global Solutions Support

- IBM GDPS HyperSwap and Basic HyperSwap support
- IBM MTMM/MTIR support
- Hitachi TrueCopy (synchronous replication)
- 3 data center multitarget with delta resync support
- 3 data center multitarget and cascade support
- 4 data center support
- Mixed mainframe and open systems consistency group support
- Asynchronous replication extended consistency groups support
- Hitachi TrueCopy consistency groups support
- Hitachi Universal Replicator (Asynchronous replication)
- Compatible FlashCopy and FlashCopy Space Efficient
- Hitachi ShadowImage Replication (clones)
- Hitachi Mainframe Cyber Resiliency

Figure 1. Hitachi Vantara's Data Resilience Solutions

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.

With industry-leading Hitachi Vantara replication and replication management solutions, which support both IBM GDPS HyperSwap and Basic HyperSwap, you can be confident of your ability to restart operations in case of a disaster or other outage.

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

- Enabling application testing against a full copy of data to simulate full production.
- Supporting parallel processing for reports, queries, analytics, artificial intelligence and data warehouse applications.
- Allowing consistency checks and 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 RPO.
- Minimizing 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. By coupling this replication or copy capability with virtualization with static and dynamic tiering capabilities you can best align cost and risk.

Due to the increase of cyberattacks, vulnerability from tight coupling between mainframe and open systems, and the need to quickly recover from a logical corruption, we offer the Hitachi Mainframe Cyber Resiliency solution. This solution supports multiple consistent images of production, rapid recovery from malicious attacks, production images validation, examination of 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, Universal Replicator, ShadowImage for Mainframe and Hitachi Business Continuity Manager (BCM). Using BCM, all three of these solutions can be seamlessly integrated into enterprise mainframes environments. Hitachi Vantara works with IBM to ensure compatibility of Hitachi replication solutions with IBM key technologies. To this end, the solutions employ unique Hitachi replication software and Hitachi enterprise-class storage support for GDPS HyperSwap. Also, as a lower cost, high-availability solution for storage, we support Basic HyperSwap with IBM Copy Services Manager (CSM).



In-System Replication

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

Compatible Mirroring for IBM FlashCopy SE (HCFC SE) provides an in-system solution for data protection that is IBM mainframe compatible. This solution combines snapshots, virtual data management and dynamic provisioning to create logical point-in-time copies. When an HCFC SE source-target volume relationship is established, HCFC SE uses HDPM 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 HCFC 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. 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 write-time stamps allows the system to perform an at-time split, providing a consistent image of the data without the need for freezing I/Os.

Consistency groups are a standard feature of Hitachi In-System Replication bundle offerings. They 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.

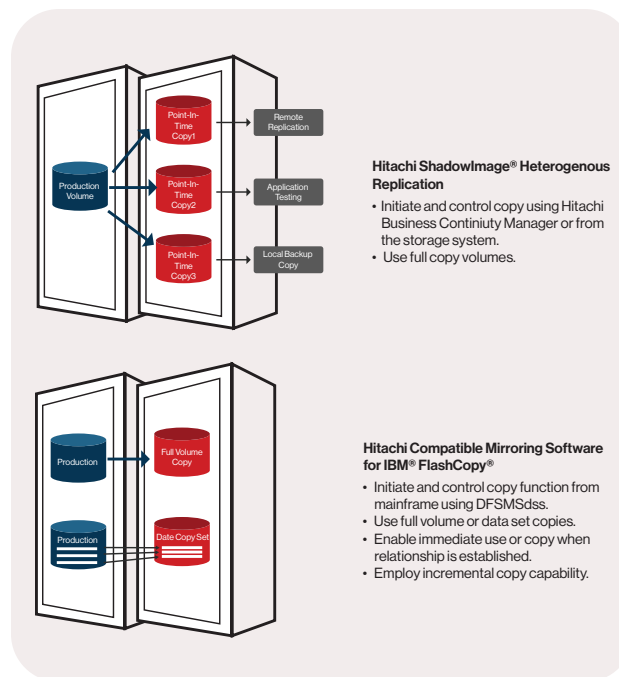


Figure 2. Local In-System Replication With Hitachi ShadowImage Heterogeneous Replication and IBM FlashCopy




Figure 3.
Remote Synchronous Replication
Using Hitachi TrueCopy
(synchronous capabilities)

Remote Replication

Hitachi Vantara replication solutions for mainframe environments support intermix environments, which can include all the latest generations of Hitachi enterprise-class storage systems, including the Hitachi Virtual Storage Platform 5000 series, VSP G1x00 and VSP F1500. Using Hitachi Vantara data replication solutions for mainframe, companies can implement GDPS, GDPS/HyperSwap, or basic HyperSwap. This includes the ability to implement 3DC and 4DC 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.

The synchronous capabilities of Hitachi TrueCopy Remote Replication bundle software provide a continuous, non-disruptive, 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 supports host application-dependent, write-order consistency across multiple controllers.

TrueCopy:

- 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.
- Works in conjunction with Hitachi Universal Replicator to provide advanced 3DC and 4DC.
- Up to 300km for environments that can sustain the distance impact on response time.

Hitachi Vantara Universal Replicator delivers simplified asynchronous data replication for internal and external storage attached to Hitachi Vantara enterprise-class storage platforms (see Figure 4). This software delivers the enterprise-class performance associated with storage system-based replication, while providing resilient business continuity without the need for redundant servers or replication appliances. HUR 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. HUR supports consistency groups across volumes on Hitachi enterprise storage systems.

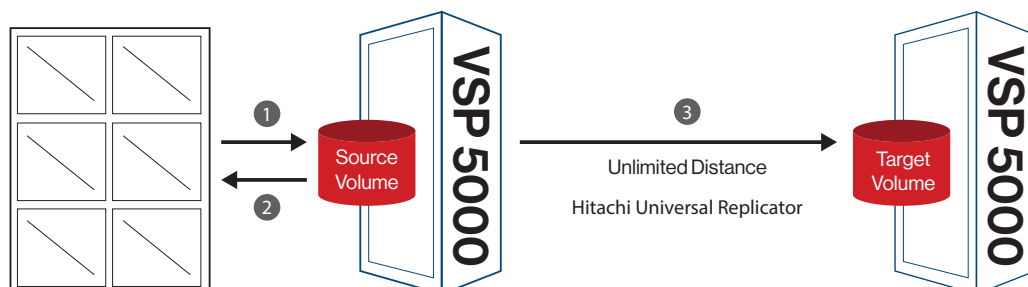


Figure 4.
Remote Asynchronous
Replication Using Hitachi
Universal Replicator

HUR provides administrators with an innovative remote replication solution. HUR offloads the majority of processing for 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. The journaling approach provides resiliency for a shorter 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:

- Supports near-zero RPO while providing resiliency to minimize RPO elongation during unexpected network or workload anomalies.
- Creates no impact on application response time as a result 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's synchronous capabilities.

For mainframe environments, these groups ensure consistency of volumes across storage system boundaries when replicating or making point-in-time copies. For 2-data-center environments, these extended consistency groups can include up to 16 storage systems in each location (64x64). For 3DC and 4DC data center environments, they can include up

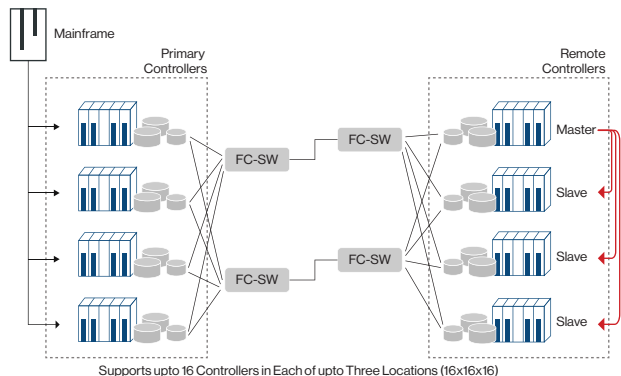


Figure 5. Extended Consistency Groups for Mainframe

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 (HUR target) site.

3 Data Center Replication

The 3DC 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 replication 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 HUR 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 Vantara enterprise-class storage system (such as VSP 5000 series, VSP G1x00 and VSP F1500) 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 enterprise storage system using Hitachi Universal Replicator propagating the TrueCopy (no data loss) copy to the remote (HUR) 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.
- Ensures no data loss when losing primary site.

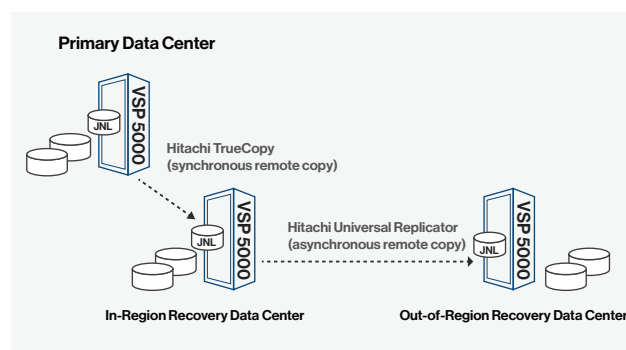


Figure 6. 3DC Cascade Solution With Hitachi TrueCopy

Multitarget (concurrent) replication synchronously replicates data from the production site via TrueCopy's synchronous capabilities 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. If the production site experiences a failure, the hot-standby site takes over and uses the remote site for disaster recovery. It only has to copy 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

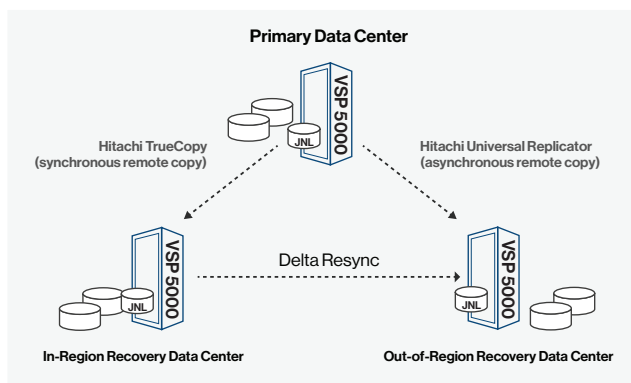


Figure 7: 3 DC Multitarget Remote Replication With Hitachi TrueCopy (synchronous) and Hitachi Universal Replicator

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.

IBM's® Multi Target Metro Mirror synchronously replicates data from the production site using TrueCopy synchronous capabilities to two local sites (within metro distance). With Multi Target Metro Mirror, protection can be resumed with no data loss between the two remaining sites (hot standby site and remote site) using IBM's Multi Target Incremental Resync (see Figure 8).

The multitarget replication solution:

- Provides two-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.

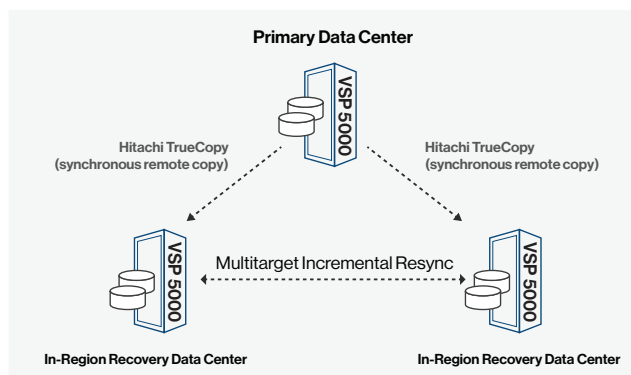


Figure 8: 3 DC Multitarget Metro Mirror With Hitachi TrueCopy (synchronous capabilities)

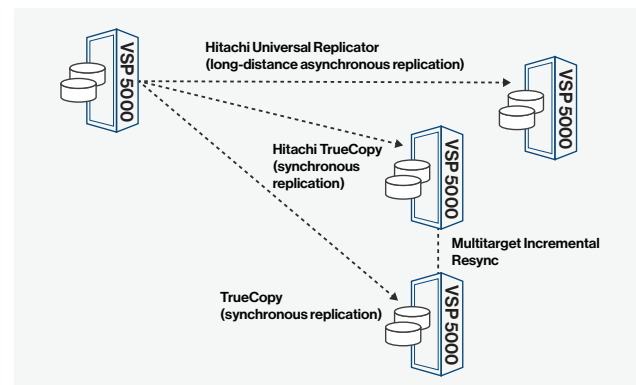


Figure 9: 1: Hitachi Universal Replicator From Primary Data Center

At-time split is a feature that eliminates the need to freeze I/Os at the primary site or HUR remote site to create a consistent copy of data. In conjunction with HUR, it eliminates the need to actually split the HUR 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 HUR copy group in a duplex state. At-time split uses mainframe 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. 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. The storage at the remote site is also primary for synchronous replication. When ShadowImage replication 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 HUR copy sessions or elongating disaster recovery RPO.




Figure 10: 4 DC Multitarget Metro Mirror with Hitachi TrueCopy (synchronous capabilities)

Automated Disaster Restart — IBM HyperSwap Support

HyperSwap Process

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

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

The microcode of Hitachi Virtual Storage Platform 5000 series supports native PPRC commands.

In normal circumstances, production runs on an IBM Z series server 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 series server is written in primary VSP 5000 series cache, transmitted and written to the secondary VSP 5000 series cache. The secondary cache then acknowledges the write to the primary server, which triggers acknowledgement of the write to the Z series server.

Using this technique allows a recovery point objective of zero.

The latency (round trip time) in glass is 1ms for 100km, which impacts all writes.

The latency (round-trip time) in glass being 1ms for 100 km, it does impact all writes.

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

On the VSP 5000 series, synchronous replication is possible up to 300km. However, most customers are in below 10km range, with asynchronous replication being preferred for long distance.




Figure 11: Synchronous Replication Process

To perform a HyperSwap, CSM or GDPS sets up the possibility to swap the unit control blocks (UCBs). The approach is to use system services, which set up UCBs by preparing them for HyperSwap.

To prepare the UCB, the synchronous replication must be active (duplex) and the Z series must be capable of seeing the primary volume and its target.

During HyperSwap process (planned or unplanned), the I/Os are frozen for a few seconds and device addresses are swapped, so the I/Os are redirected.

The operations are more complex, as the following sequence of operations is done:

- 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 must be reassigned.
- I/Os are resumed to the previous secondary storage, which has been promoted to primary.

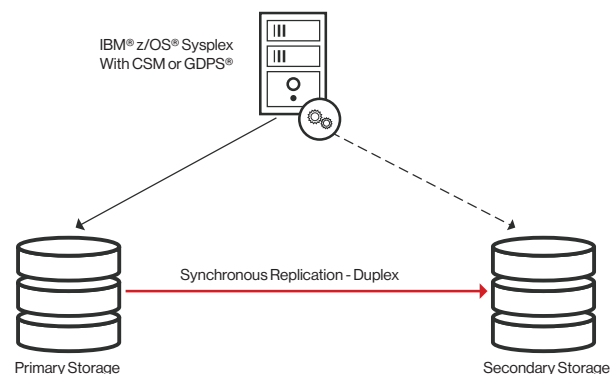


Figure 12: Usual Production Environment

After the HyperSwap process, the configuration is usually as follows:




Figure 13: Usual Configuration After HyperSwap

After the HyperSwap Process

The secondary has been promoted to the new primary, and the replication pairs are generally suspended after the HyperSwap process (at least unplanned HyperSwap). Note that the IBM command set allows the simulation of an unplanned HyperSwap process.

With GDPS it is possible to do a HyperSwap resync for a planned HyperSwap process: That is, after the HyperSwap process is completed, the PPRC pairs are reversed and duplex. From the application standpoint, this is transparent. Note that a planned HyperSwap process allows the failback to the initial configuration.

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

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, among others.

The Different HyperSwap Technologies

- CSM: Copy Services Manager Basic Edition for IBM z Systems® and Copy Services Manager for z Systems.
- GDPS: several versions, with the most common being HyperSwap Manager, MM, MTMM.

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

Note that the maintenance is not included in the maintenance fees; it must be purchased separately. It does allow you to manage HyperSwap — that is protection against disk access failure.

The Copy Services Manager for Systems is a licensed product that brings 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, allowing you to manage CSM when the primary CSM address space or server is not available.

Although the processes involved are the same ones, there is no upgrade path between CSM and GDPS. The GDPS family is broader, allowing storage outage protection (HyperSwap) to dynamic site switching. In the case of GDPS family, the HyperSwap is completed by a specific system: the controlling system.

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

- GDPS HyperSwap Manager manages HyperSwap like CSM would do, but also allows the upgrade to a broader version of the GDPS family.
- GDPS MM allows the user to HyperSwap z/OS, but also, using a proxy, z/VM (usually Linux for z is running under z/VM)
- GDPS MTMM allows the user to manage and HyperSwap a PPRC multitarget configuration: that is, two synchronous copies out of the same source volume.

Hitachi Storage Technologies and HyperSwap

Hitachi for many years has participated in the IBM® qualification program from GDPS, and the latest qualification was done for VSP 5000 series. At the end of the qualification, IBM publishes a qualification letter on this website <https://www.ibm.com/it-infrastructure/z/technologies/gdps> (see qualification section).

The qualification included : GDPS MM (formerly GDPS/PPRC) for HyperSwap, GDPS 3DC [i.e., GDPS MM combined with Universal Replicator delta resync managed with business continuity management (BCM)], GDPS MTMM, GDPS MTMM combined with Universal Replicator, and combination of Hitachi Vantara 4DC with GDPS MM on both regions.

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

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

The VSP 5000 series family appears in the IBM compatibility list at:

<https://www.ibm.com/support/pages/node/656273>

The VSP 5000 series GDPS qualifications for GDPS MM (2DC), GDPS MM with 3DC delta resync, GDPS MTMM, GDPS MTMM with Universal Replicator, and GDPS synergy with Hitachi Vantara 4DC can be found at: <https://www.ibm.com/products/gdps>

Replication Management

Hitachi Vantara provides flexibility in storage management for mainframe environments with both host-based and storage system-based storage management solutions. Hitachi Business Continuity Manager (BCM) provides a host-based solution that uses familiar TSO/ISPF screens to view and manage storage replication. BCM has a REST API like interface, allowing it to manage and query the replication with responses in a JSON format. Hitachi has developed a GUI based on this API. 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, 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 or WEB API calls to manage these replication offerings. In-band BCM commands can be used to manage replication at the remote site from the primary site. BCM includes a monitor, which allows it to detect and observe pairs and take action should the status being incorrect.

Virtual Tape Library Support

As disk technology has advanced, use of lower-cost, nearline SAS disk storage for data backup and archive has become a popular alternative to tapes. Both the Secure Agent Secure Data Solution (SDS) and the Luminex Channel Gateway provide virtual tape library solutions for all IBM Z series and compatible mainframes.

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.

Hitachi Vantara works with and provides solutions from both SecureAgent Software and Luminex for virtual tape library (VTL) solutions. SecureAgent Software's SDS is an "all in one" solution that supports open systems, IBM iSeries® and all IBM Z series operating systems, including z/TPF. Luminex Channel Gateway supports z/OS, z/VM and z/VSE. Both 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 help customers lower their operating costs. More and more customers have embraced cloud solutions. We are working with BMC AMI Cloud (formerly

Model9) to provide a solution that can lower the millions of service units (MSU) usage for the data movement for tape processing and copy the data over IP to an Amazon S3 compatible unit. This approach allows data to be hosted in the cloud, locally — for example on HCP.

- There is no need to change current backup JCLs to write on HCP. The link is made through flexible automatic code synchronizer (ACS) routines.
- For this specific backup to HCP, the DFSMSdss is replaced by java code running on a zIIP without any MSU cost, reducing the software bill for z/OS.
- HCP allows versioning and protecting data from any changes.
- A recent customer experience reduced the backup time, compared to his VTL, by 17-fold.

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

Hitachi Mainframe Cyber Resiliency

Taking advantage of the various Hitachi technologies, Hitachi Vantara has implemented a solution to enable recovery from logical data corruption. The solution uses Business Continuity Manager to trigger a consistent copy of production data using the ShadowImage at-time split capability.

The modifications since last cycle are copied and stored into a virtual volume, which is external to the mainframe storage (implying that it does not take extra space in the current production VSP 5000 series) and physically stored in a virtual volume located in an open systems storage external to the mainframe storage. This open system storage prevents data in the fortress from being modified or deleted before the expiration date is over. The fortress keeps the number of cycles that the administrator is planning available.

When the last cycle modifications are stored on the open system external volume, system external volume, a script is triggered to index/store to index/store the modifications in the fortress. When this last process is finished, the mainframe is warned that the last cycle is stored in the fortress.

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

As this fortress does not reside on the production storage, 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 impact the integrity of the current image (changing cycles, changing retention period, for example) would require a two-user acknowledgment for security purposes.

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

Data Mobility and Migration

Hitachi Vantara is committed to ensuring that your data is available at the right place and at the right time. We provide solutions for moving and/or migrating your information without impacting the performance of your system. These solutions ensure continuous access, business continuity and data resilience. They enable you to make the most efficient use of your storage assets and optimize the deployment of your critical data without impacting your business operations. Hitachi Vantara offers solutions for movement of data in mainframe environments for migration or just for data mobility.

Hitachi Cross-OS File Exchange

Hitachi 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 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.

Hitachi FICON Data Migration

FICON data migration is a unique feature by which z/OS data can be copied from legacy storage systems (Hitachi, DELL/EMC and IBM) to the Hitachi Virtual Storage Platform 5000 series without using host software. In addition, no conversion of data is necessary. The legacy system is connected to the Hitachi enterprise storage platform using FICON connections. The data is then moved while the mainframe applications continue to execute. It is intended purely as a means of migration. No other vendor offers this unique capability.

Integration With Databases and Applications

Hitachi 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 and Hitachi Vantara have 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.

Mainframe Productivity

Enterprise storage solutions from Hitachi Vantara are designed to improve productivity of storage administrators. With virtualization provided by Hitachi Universal Volume Manager, external storage from Hitachi Vantara or other vendors can be easily virtualized behind Hitachi Vantara enterprise storage controllers, such as the VSP 5000 series, VSP G1x00 and VSP F1500.

With the introduction of Hitachi Dynamic Provisioning for Mainframe (HDPM), storage administrators do not have to manually provision storage. Rather, the storage controller will be able to provision storage, as it is needed. Hitachi Dynamic Tiering for Mainframe (HDTM) will enable the automatic movement of data between tiers. HDTM moves highly accessed blocks of data to the highest tier storage and migrates less frequently accessed data to the lowest tiers. This significantly reduces the time storage administrators must spend analyzing storage usage and managing the movement of data to optimize performance.

Hitachi Tiered Storage Manager for Mainframe (HTSM for Mainframe) online storage service level control HDTM policies. With HTSM for Mainframe control of HDT for Mainframe policies, storage service levels can be controlled to improve performance and avoid problems before they happen. HTSM for Mainframe works with SMS storage group constructs, giving users the capability of managing HDTM environments from the z/OS point of view. HTSM for Mainframe also provides reporting and control for ongoing management and monitoring of HDTM operations.

Data-at-Rest Encryption (DARE)

Our data-at-rest encryption feature 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 (HDD/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 function is implemented using cryptographic chips included as part of the encryption hardware. The VSP 5000 series is compliant to FIPS 140-2 Level 2.

Hitachi Technology Partners

21st Century Software

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

BMC Compuware

INNOVATION Data Processing's 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 (formerly 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.

Broadcom (Vantage)

Broadcom's 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 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.

IBM's IntelliMagic (IntelliMagic Vision)

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 Vision ingests Hitachi storage SMF records (recordings of the internals of the storage).

Luminex

Luminex Channel Gateway virtual tape solutions deliver support for z/OS, z/VM and z/VSE. 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's 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+ quickly detects 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 (Iron Stream)

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). These records are used to trigger alerts and can be merged with mainframe records for deep performance analysis and sent to Splunk.

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 enables access to the data recordings of the storage internals).

SecureAgent Software

SecureAgent Software's SDS solution delivers a fully integrated capability for IBM Z series 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 Vantara has a strong relationship with IBM to ensure our 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 their data.

Technologies such as Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering 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 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 Hitachi Mainframe Cyber Resiliency solution to offer customers an approach that offers a real air gap: 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 take multiple consistent images from production, and permits them to perform checking, forensic analysis, surgical or full restoration of the production data from selected image.

Hitachi, Ltd., and its subsidiary, Hitachi Vantara, are leaders in providing leading-edge technology and storage solutions for both large and not so large enterprise customers with IBM mainframe infrastructures.

Hitachi Virtual Storage Platform 5000 series.
Fastest data platform in the world.

Learn more →

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.



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

© Hitachi Vantara LLC 2025. All Rights Reserved. HITACHI and Pentaho are trademarks or registered trademarks of Hitachi, Ltd. All other trademarks, service marks and company names are properties of their respective owners.

HV-BTD-WP-VSP5000-Storage-Series-10Dec24-E