

Cisco and Hitachi Adaptive Solutions for SAP HANA

Lab Validation Report

By Hiroyuki Hayashi, Shinji Osako, and Markus Berg

July 2019

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@HitachiVantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision History

Revision	Changes	Date
MK-SL-158-00	Initial release	July 10, 2019

Table of Contents

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure	1
Cisco and Hitachi Adaptive Solutions for SAP HANA Tailored Data Center Integration (TDI)	2
Cisco and Hitachi Adaptive Solutions for SAP HANA TDI Architecture	2
Hitachi Virtual Storage Products for Cisco and Hitachi Adaptive Solutions	4
Hitachi Virtual Storage Platform F Series	4
Hitachi Virtual Storage Platform G Series	6
Hitachi Virtual Storage Product Features	7
3D Scaling of Storage Resources	9
Hitachi Dynamic Provisioning	10
Hitachi Dynamic Tiering	12
Hitachi Advanced Data Reduction by Deduplication and Compression	12
Backup Data Reduction and Storage Space Savings	13
Server Configuration	13
Storage Configuration	14
Hitachi Storage Virtualization Operating System	15
Software Components	15
Test Methodology	16
Summary of Test Results	19
Conclusions	21
Business Impact	22

Cisco and Hitachi Adaptive Solutions for SAP HANA

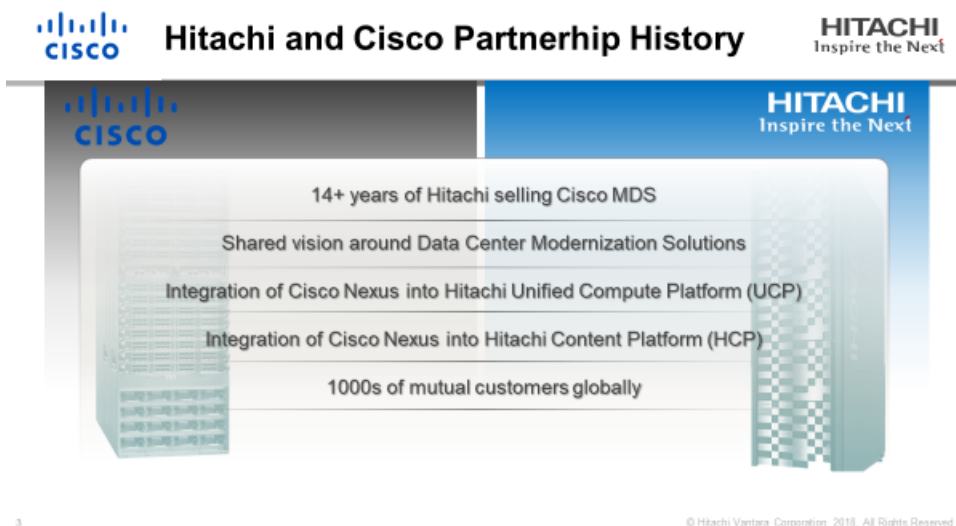
Backup Data Reduction and Storage Space Savings of SAP HANA Database with SVOS 8.1 with Hitachi Virtual Storage Platform Capacity Saving Features

Lab Validation Report

Cisco and Hitachi are working together to deliver a converged infrastructure solution that helps enterprise businesses meet the challenges of today and position themselves for the future. Leveraging decades of industry expertise and superior technology, Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a Virtual Server Infrastructure offers a resilient, agile, and flexible foundation for today's businesses. In addition, the Cisco and Hitachi partnership extends beyond a single solution, enabling businesses to benefit from their ambitious roadmap of evolving technologies such as advanced analytics, IoT, cloud, and edge capabilities. With Cisco and Hitachi, organizations can confidently take the next step in their modernization journey and prepare themselves to take advantage of new business opportunities enabled by innovative technology.

Cisco and Hitachi Adaptive Solutions for SAP HANA leverages the strengths of both Cisco and Hitachi Vantara. Figure 1 shows details of the partnership between Hitachi and Cisco.

Figure 1



Cisco and Hitachi Adaptive Solutions for Converged Infrastructure

- Adaptive Solutions for Converged Infrastructure Design Guide:
https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci_design.html
- Adaptive Solutions for Converged Infrastructure Deployment Guide:
https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci.html

Cisco and Hitachi Adaptive Solutions for SAP HANA Tailored Data Center Integration (TDI)

- Adaptive Solutions for SAP HANA TDI Design Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci_sap_design.html

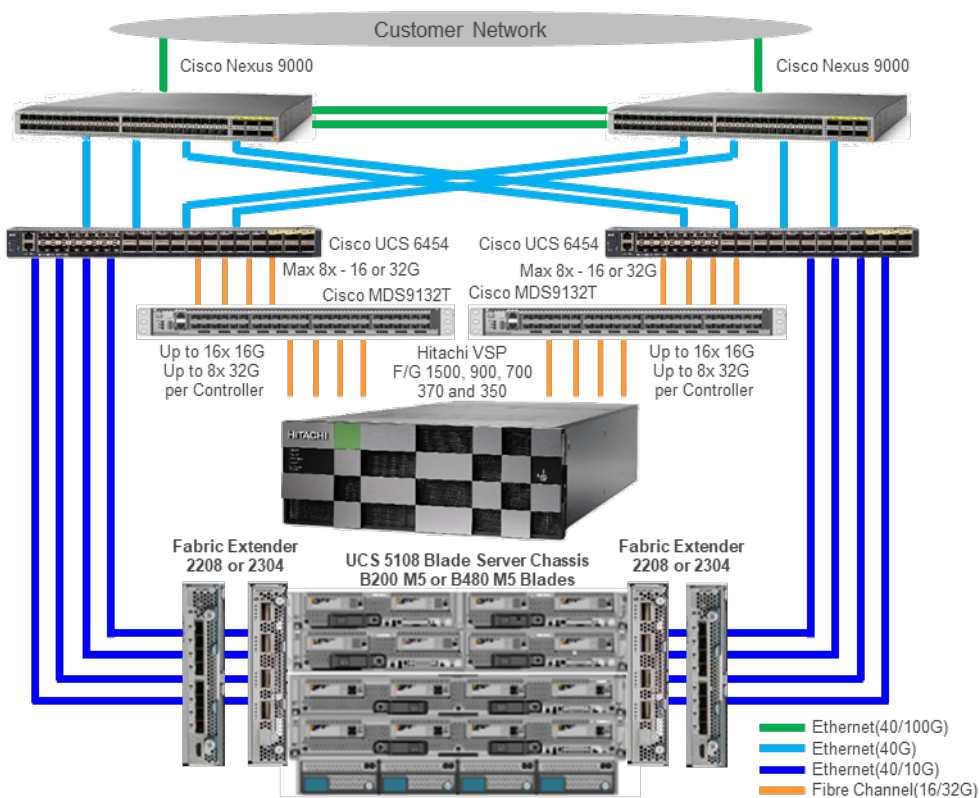
- Adaptive Solutions for SAP HANA TDI Deployment Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci_sap.html

Cisco and Hitachi Adaptive Solutions for SAP HANA TDI Architecture

Figure 2 shows the architecture of Cisco and Hitachi Adaptive Solutions for SAP HANA with Cisco Unified Computing System Server Blades and Hitachi Virtual Storage Platform family. The components are configured using the configuration and connectivity best practices from both companies to implement a powerful and scalable Infrastructure for SAP HANA in a TDI environment validated for both SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL). This design is presented as a validated reference architecture, that covers specifics of products utilized within the Cisco lab; however, equivalent supported products can be replaced which are listed within Cisco and Hitachi's published compatibility matrixes.

Figure 2



- Compute

The compute piece of the system incorporates servers based on the latest Intel x86 processors. Servers are available in blade and rack form factor, managed by Cisco UCS Manager.

- Network

The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gbps Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric. The unified fabric uses the innovative Single Connect technology to lower costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system

- Storage access

Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. Also, server administrators can pre-assign storage access policies to storage resources, for simplified storage connectivity and management leading to increased productivity

- Management

- The system uniquely integrates compute, network, and storage access subsystems, enabling it to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager increases IT staff productivity by enabling storage, network, and server administrators to collaborate on service profiles that define the desired physical configurations and infrastructure policies for applications. Service profiles increase business agility by enabling IT administrators to automate and provision resources in minutes instead of days

Hitachi Virtual Storage Products for Cisco and Hitachi Adaptive Solutions

Hitachi Virtual Storage Platform F Series



TABLE 1. HITACHI VIRTUAL STORAGE PLATFORM F SERIES

	Hitachi Virtual Storage Platform F350	Hitachi Virtual Storage Platform F370	Hitachi Virtual Storage Platform F700	Hitachi Virtual Storage Platform F900	Hitachi Virtual Storage Platform F1500
Performance (IOPS)	600K	1.2M	1.4M	2.4M	4.8M
Max. Raw Internal Capacity	2.8 PB (15 TB SSD)	4.3 PB (15 TB SSD)	6.0 PB (14 TB FMD) 13.0 PB (15 TB SSD)	8.1 PB (14 TB FMD) 17.3 PB (15 TB SSD)	8.1 PB (14 TB FMD) 34.6 PB (15 TB SSD)
Total Efficiency Guarantee Ratio*	Up to 7:1	Up to 7:1	Up to 7:1	Up to 7:1	2:1 Data Reduction

TABLE 1. HITACHI VIRTUAL STORAGE PLATFORM F SERIES (CONTINUED)

	Hitachi Virtual Storage Platform F350	Hitachi Virtual Storage Platform F370	Hitachi Virtual Storage Platform F700	Hitachi Virtual Storage Platform F900	Hitachi Virtual Storage Platform F1500
Host Port Counts (Max. without drives)	16 × Fibre Channel 8 × iSCSI	16 × Fibre Channel 8 × iSCSI	48 × Fibre Channel 24 × iSCSI	64 × Fibre Channel 32 × iSCSI	176 × Fibre Channel 176 × FICON 176 × FCoE 88 × iSCSI
Value Added Bundled Features	<ul style="list-style-type: none"> ■ 100% Data Availability Guarantee ■ Adaptive Data Reduction Storage ■ Virtualization ■ In System Replication ■ Copy Management ■ Infrastructure Analytics ■ Non-disruptive Migration ■ Total Efficiency Guarantee 				

Notes:

* The Total Efficiency Guarantee of up to 7:1 for the VSP F350, F370, F700 and F900 models covers savings from data de-duplication, compression, thin provisioning and snapshots. The VSP F1500 model is eligible for a Data Reduction Guarantee of 2:1 for capacity stored on Hitachi flash modules with data compression enabled.

1MB = 1,000,000 bytes

SFF = small form factor

LFF = large form factor, SAS = serial attached SCSI, SSD = solid state disk

FMD = flash module drive

Hitachi Virtual Storage Platform G Series



TABLE 2. HITACHI VIRTUAL STORAGE PLATFORM G SERIES

	Hitachi Virtual Storage Platform G350	Hitachi Virtual Storage Platform G370	Hitachi Virtual Storage Platform G700	Hitachi Virtual Storage Platform G900	Hitachi Virtual Storage Platform G1500
Performance (IOPS)	600K	1.2M	1.4M	2.4M	4.8M
Max. Raw Internal Capacity	2,467 TB (10 TB LFF HDD) 442 TB (2.4 TB SFF HDD) 2,889 TB (15 TB SSD)	3,642 TB (10 TB LFF HDD) 664 TB (2.4 TB SFF HDD) 4,333 TB (15 TB SSD)	11.7 PB (10 TB LFF HDD) 1.9 PB (2.4 TB SFF HDD) 6.0 PB (14 TB FMD) 13.0 PB (15 TB SSD)	14.0 PB (10 TB LFF HDD) 2.6 PB (2.4 TB SFF HDD) 8.1 PB (14 TB FMD) 17.3 PB (15 TB SSD)	6.7 PB (6 TB LFF HDD) 5.3 PB (2.4 TB SFF HDD) 8.1 PB (14 TB FMD) 34.6 PB (15 TB SSD)
Total Efficiency Guarantee Ratio*	Up to 7:1	Up to 7:1	Up to 7:1	Up to 7:1	2:1 Data Reduction

TABLE 2. HITACHI VIRTUAL STORAGE PLATFORM G SERIES (CONTINUED)

	Hitachi Virtual Storage Platform G350	Hitachi Virtual Storage Platform G370	Hitachi Virtual Storage Platform G700	Hitachi Virtual Storage Platform G900	Hitachi Virtual Storage Platform G1500
Host Port Counts (Max. without drives)	16 × Fibre Channel 8 × iSCSI	16 × Fibre Channel 8 × iSCSI	64 × Fibre Channel 32 × iSCSI	80 × Fibre Channel 40 × iSCSI	192 × Fibre Channel 176 × FICON 192 × FCoE 96 × iSCSI
Value Added Bundled Features	<ul style="list-style-type: none"> ■ 100% Data Availability Guarantee ■ Adaptive Data Reduction Storage ■ Virtualization ■ In System Replication ■ Copy Management ■ Infrastructure Analytics ■ Non-disruptive Migration ■ Total Efficiency Guarantee 				

Notes:

* The Total Efficiency Guarantee of up to 7:1 for the VSP G350, G370, G700 and G900 models, when configured entirely with flash drives, covers savings from data de-duplication, compression, thin provisioning and snapshots. The VSP G1500 model is eligible for a Data Reduction Guarantee of 2:1 for capacity stored on Hitachi flash modules with data compression enabled.

1MB = 1,000,000 bytes

SFF = small form factor, LFF = large form factor

SAS = serial attached SCSI, HDD = hard disk drive, SSD = solid state disk, FMD = flash module

Hitachi Virtual Storage Product Features

The [Hitachi Virtual Storage Platform](#) series family enables the seamless automation of the datacenter. It has a broad range of efficiency technologies that deliver maximum value while making ongoing costs more predictable. You can focus on strategic projects and consolidate more workloads while using a wide range of media choices.

The benefits start with the Hitachi Storage Virtualization Operating System (SVOS) resilience flash (RF) feature. This includes an all new enhanced software stack that offers up to three times greater performance than our previous midrange models, even as data scales to petabytes. Virtual Storage Platform G series offers support for containers to accelerate cloud-native application development. Provision storage in seconds, and provide persistent data availability, all while being orchestrated by industry leading container platforms. Move these workloads into an enterprise production environment seamlessly, saving money while reducing support and management costs.

Hitachi Virtual Storage Platform is a highly scalable, true enterprise-class storage system that can virtualize external storage and provide virtual partitioning and quality of service for diverse workload consolidation. The abilities to securely partition port, cache, and disk resources, and to mask the complexity of a multivendor storage infrastructure, make Virtual Storage Platform the ideal complement to SAP environments.

As your SAP environment grows to encompass mission-critical and Tier 1 business applications, Hitachi Virtual Storage Platform delivers the highest uptime and flexibility for your block-level storage needs, providing much-needed flexibility for Enterprise environments.

With the addition of flash acceleration technology, a single Hitachi Virtual Storage Platform is now able to service more than 1 million random read IOPS. This extreme scale allows you to increase systems consolidation, defer capital expenses and operating costs, and improve quality of service for open systems and virtualized applications.

Hitachi Virtual Storage Platform capabilities ensure that businesses can meet service level agreements (SLAs) and stay on budget.


- The industry's only 100% uptime warranty
- 3D scalable design
- 40 percent higher density
- 40 percent less power required
- Industry's leading virtualization
- Nondisruptive data migration
- Fewer administrative resources
- Resilient performance, less risk

3D Scaling of Storage Resources

Hitachi Virtual Storage Platform is the only enterprise storage platform in the industry that offers 3D scaling of resources (see Figure 3), full native VMware integration, and storage virtualization that extends features to more than 100 heterogeneous storage systems, allowing you to leverage existing investments.

- Scale up for increased performance, adding just the processing and cache you need.
- Scale out for increased capacity, with up to 2048 drives internally.
- Scale deep to leverage the industry's leading external virtualization technology and manage your external storage through common interface and management tools.


Figure 3



Hitachi Virtual Storage Platform offers a flexible choice of 2.5-inch and 3.5-inch form factors, drives, and energy efficiencies of up to 40 percent savings compared to other storage systems. With a much smaller datacenter footprint, Hitachi Virtual Storage Platform excels in performance and cost savings while addressing a multitude of enterprise business requirements. Hitachi Virtual Storage Platform also extends the life of storage assets and increases ROI by enabling the connection of legacy storage systems to Hitachi Virtual Storage Platform.

Figure 4 shows how Hitachi Virtual Storage Platform supports 3D scaling.

Figure 4



One of the unique features of Hitachi Virtual Storage Platform is its ability to use storage virtualization as a use case for Enterprise environments, which allows administrators to:

- Non-disruptively migrate to new storage
- Enable data classification for tiered storage management

The flexible, scalable, secure Hitachi Virtual Storage Platform design provides:

- Integrated data at rest encryption
- Virtualization of up to 255 PB of external storage
- Up to 1 TB of shared global cache
- Up to 192 Fibre Channel 8Gb/sec host ports
- Up to 96 Fibre Channel over Ethernet (FCoE) 10Gb/sec host ports
- Built-in hardware-based data-at-rest encryption

Hitachi Virtual Storage Platform further enables wide striping and thin provisioning with Hitachi Dynamic Provisioning (HDP) software:

- Wide striping of data across all drives in array groups assigned to the HDP pool to increase performance by striping across spindles in multiple array groups
- Thin provisioning and management of Dynamic Provisioning pools in Hitachi Command Suite management tools


Hitachi Dynamic Provisioning

Hitachi Dynamic Provisioning (HDP) is a microcode function of Virtual Storage Platform that provides thin provisioning and wide striping services for simplified provisioning operations, automatic performance optimization and storage space savings. Provisioning storage from a virtual pool dramatically simplifies capacity planning and reduces administration costs by cutting the time to provision new storage. It also improves application availability by reducing the downtime needed for storage provisioning.

Capacity is allocated to an application without it being physically mapped until it is used. In this manner it is possible to achieve overall higher rates of storage utilization with just-in-time provisioning. Dynamic provisioning separates an application's storage provisioning from the addition of physical storage capacity to the storage system. When physical storage is installed within the Hitachi Virtual Storage Platform or a virtualized external array it is mapped to one or more Dynamic Provisioning Pools. Physical storage can be non-disruptively added to or removed from the pool as needed and existing allocated capacity is automatically restriped. Additional physical storage is dynamically added to thin provisioned Host LUNs (DP-VOLS) whenever an application requires additional capacity. Physical storage that is no longer required can also be dynamically returned to the pool using Zero Page Reclaim and Hitachi Virtual Storage Platform unmap functionality.

Dynamic Provisioning also simplifies Storage Virtualization as it provides a layer of abstraction which masks the difference between internal and external storage resources. There is no difference in the way that Host LUNs are carved out of pools for internal and external storage. The only difference is the way that physical storage capacity is replenished. This is illustrated in Figure 5.

Figure 5

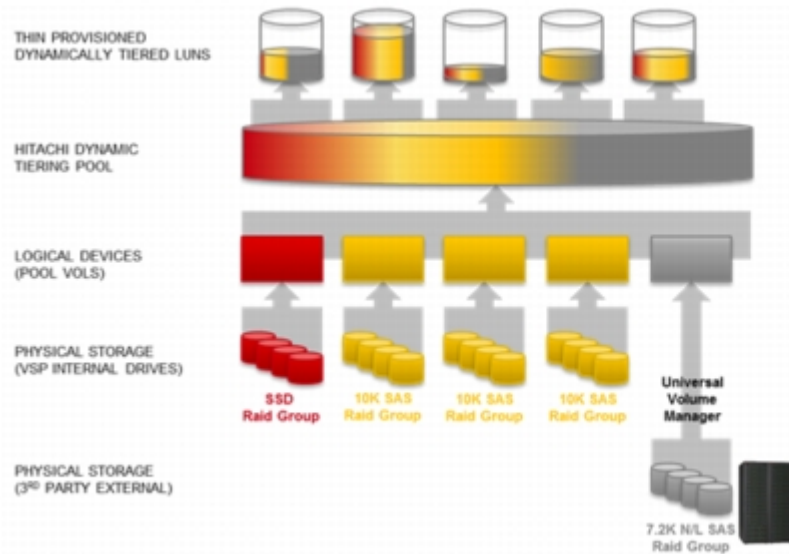


Deploying Hitachi Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. This distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Tiering

Hitachi Dynamic Tiering is an enhancement to Dynamic Provisioning that increases performance and lowers operating cost using automated data placement. Data is automatically moved according to simple rules. Two or three tiers of storage can be defined using any of the storage media types available for Hitachi Virtual Storage Platform. Tier creation is automatic based on media type and speed. Using ongoing embedded performance monitoring and periodic analysis, the data is moved at a fine grain, sub-LUN level to the most appropriate tier. The most active data moves to the highest tier. During the process the system automatically maximizes the use of storage, keeping the higher tiers fully utilized. This is illustrated in Figure 6.

Figure 6



Hitachi Virtual Storage Platform is the only 3D scaling storage platform designed for all data types. It is the only storage architecture that flexibly adapts for performance and capacity, and it extends to multivendor storage. With the unique management capabilities of Hitachi Command Suite software, it transforms the data center. You can start small and scale up, out and deep to meet your unique requirements, now and as your storage needs grow.

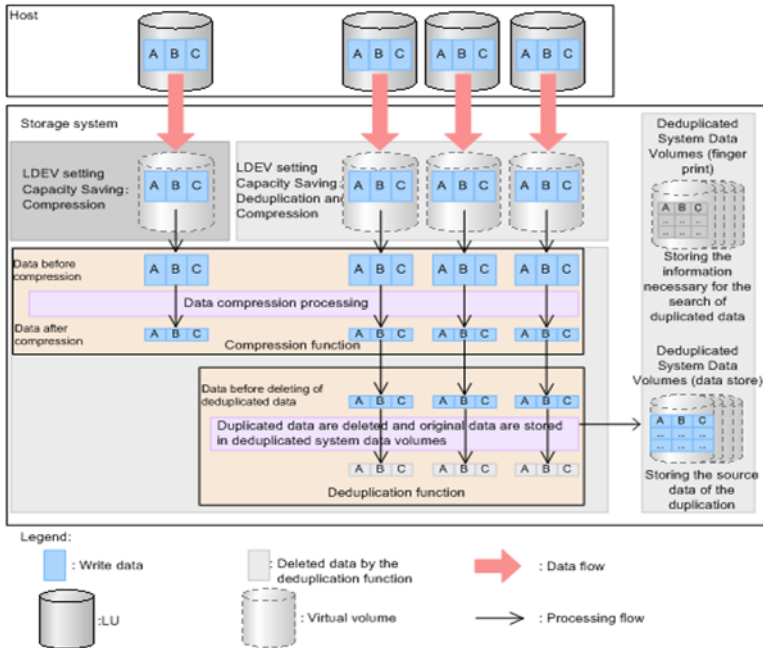
Hitachi Advanced Data Reduction by Deduplication and Compression

Hitachi Storage delivers superior adaptive data reduction (ADR) and operational efficiency, covering a broad range of efficiency services including thin provisioning, snapshots and linked clones, compression, deduplication, and cloud connect. SVOS RF adaptive data reduction intelligence is optimized for highest system throughput and response time consistency. The Virtual Storage Platform VSP F series all-flash arrays and all-flash configurations of the VSP G series storage systems deliver inline, drive-based accelerated compression to provide system-level storage efficiency savings. In addition, SVOS RF manages the ADR services used based on configuration. Figure 7 provides an illustration of adaptive data reduction.

Compression: The data compression function utilizes the LZ4 compression algorithm to compress the data. The compression function is also enabled per specific DP-VOL.

Deduplication: The data deduplication function deletes duplicate copies of data written to different addresses in the same pool and maintains only a single copy of the data at one address. The deduplication function is enabled on a Dynamic Provisioning pool and then on the desired DP-VOLs in the pool. When deduplication is enabled, data that has multiple copies between DP-VOLs assigned to that pool are removed.

Figure 7



Backup Data Reduction and Storage Space Savings

For this evaluation, Hitachi Virtual Storage Platform G700 was used in testing. These products were part of the tested configuration.

Server Configuration

The Server Blades used for this testing had these components as a SAP HANA database server:

- 4 × Intel Xeon E7-8890 v3 processors
- 2.5 GHz processor frequency
- 18 cores per processor
- 1.5 TB memory
- 1 × 2-port 10GBASE-SR LAN PCIe adapter
- 2 × 16 Gb/s 2-port Fibre Channel adapters

Storage Configuration

The hardware components used in the Hitachi Virtual Storage Platform G700 storage array used for testing were the following:

- 2 Controllers
- 16 × 16 Gb/s Fibre Channel ports
- 512 GB cache memory

One Virtual Storage Platform G700 (VSP G700, No.1 in Figure 1) was used for the following:

- SAP HANA database persistent volumes

Another Virtual Storage Platform G700 (VSP G700, No.2 in Figure 1) was used for the following:

- Backup volume

Dynamic provisioning pools created with Hitachi Dynamic Provisioning, a part of the Server Virtualization Operating System, were used in the test configuration. Its use is required for using the capacity-saving features of Virtual Storage Platform G700.

The Original Lab Testing was done by Ikki Mizumura for Hitachi LTD.

In this report, read about compression and deduplication tests performed using Hitachi Virtual Storage Platform G700 (VSP G700) storage with Hitachi Storage Virtualization Operating System 8.1 (SVOS). The tests determined the capacity saving achieved on SAP HANA database backups by enabling compression and deduplication. In Analysis, this report makes general conclusions and recommendations for all results as well as comments on specific test protocols.

Virtual Storage Platform G700 and Storage Virtualization Operating System offer data reduction features to increase usable capacity of the storage array. The capacity saving function includes data deduplication and data compression. Deduplicating and compressing the data allows you to reduce your bit cost for the stored data. The controllers of the storage system perform the data deduplication and compression.

- Deduplication — Data deduplication deletes duplicate copies of data written to different addresses in the same pool, maintaining only a single copy of the data at one address.
- Compression — Data compression utilizes the LZ4 compression algorithm to compress the data for each dynamic provisioning volume (DP-VOL).

This lab validation report describes how capacity saving features of Storage Virtualization Operating System on SAP HANA was tested. SAP HANA is an in-memory database. The column store allows efficient compression of data. This makes it less costly for a SAP HANA database to keep data in main memory.

Read details about SAP HANA compression in the [Compression Users Guide: Administration Guide](#). Since data is already efficiently compressed, capacity saving on the primary SAP HANA database storage is lower than for traditional databases. However, deduplication and compression can provide significant benefits for backup of a SAP HANA database without any direct performance impact.

This test assumes a use case of storing the full backup data of multiple generations from SAP HANA database in memory to a backup volume with deduplication and compression. This lab validation report is based on the use case documents results of tests performed for multiple backup and restore runs to identify backup time, restore time and post-process time, by changing the configuration of deduplication and compression with different drives, RAID configuration, and volume configuration to compare the performance.

Note — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Hitachi Storage Virtualization Operating System

[Hitachi Storage Virtualization Operating System](#) (SVOS) spans and integrates multiple platforms. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Storage Virtualization Operating System includes standards-based management software on a Hitachi Command Suite (HCS) base. This provides storage configuration and control capabilities for you.

Storage Virtualization Operating System uses Hitachi Dynamic Provisioning (HDP) to provide wide striping and thin provisioning. Dynamic Provisioning provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) without initially allocating any physical space. Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs).

Storage Virtualization Operating System uses Deduplication and Compression. Deduplication and Compression has two kinds of capacity saving method, Inline mode and Post-process mode.

- Inline mode: Deduplication and Compression run while writing data I/O
- Post-process mode: Deduplication and Compression run after writing data I/O

Software Components

In order to change these two modes, the RAID manager version listed in Table 1 is needed.

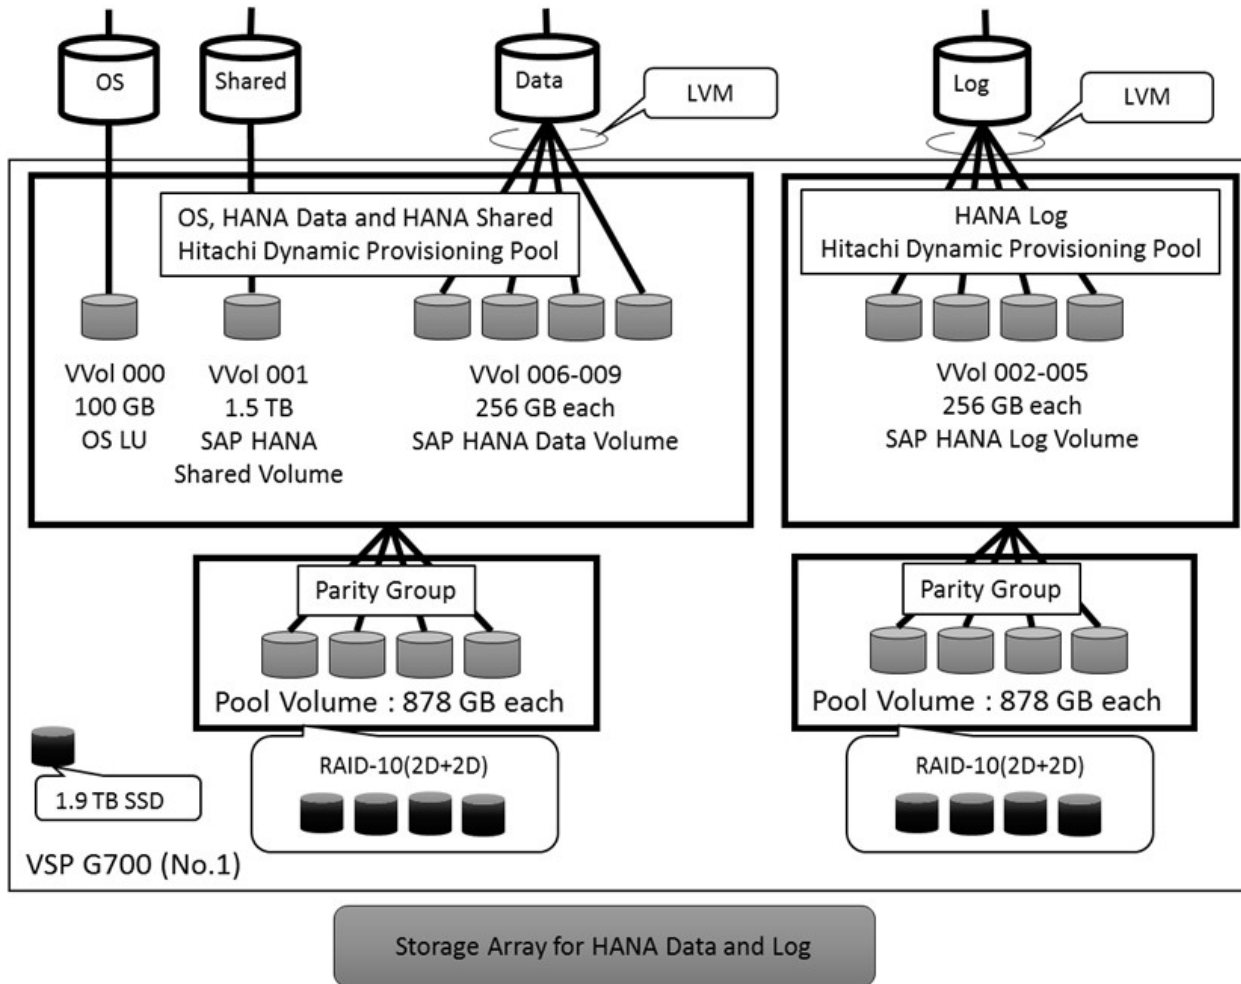
The software components listed in Table 1 were used to test compression and deduplication on Hitachi Virtual Storage PlatformG700.

TABLE 3. SOFTWARE COMPONENTS USED FOR TESTING

Software	Version
SUSE Linux Enterprise Server	SUSE Linux Enterprise Server 12 SP1
SAP HANA Platform	1.0 SPS12
Hitachi Storage Navigator	88-02-03-00/00
RAID Manager	01-46-03/02
Storage Virtualization Operating System microcode	88-01-03-60/00

Figure 8 shows the configuration of SAP HANA database persistent volumes.

Figure 8



Test Methodology

The objectives of this test were to identify possible Savings in Storage Space of multiple backup and restore runs based on the use case. In order to identify them, in each test case, the configuration of deduplication and compression changed with different drives, RAID configuration, and volume configuration as shown in Table 2, and the performance was compared on the backup volume of Hitachi Virtual Storage PlatformG700 with the SAP HANA backup tool.

Table 2 shows the summary of test cases for changing configuration of Deduplication and Compression (Dedup/Comp).

TABLE 4. SUMMARY OF TEST CASES

#	Test Case	RAID Type and VVOL Size for the Backup Volume	Backup Data Size	Dedup/ Comp	Dedup/Comp Mode
1	SSD Dedup/Comp off	SSD RAID6(6D+2P) × 1 (VVOL= 5 TB LU × 1)	237 GB	Off	-
2	SSD Dedup/Comp PostProcess Mode	SSD RAID6(6D+2P) × 1 (VVOL= 5 TB LU × 1)	237 GB	On	Post-process
3	SSD Dedup/Comp Inline Mode	SSD RAID6(6D+2P) × 1 (VVOL= 5 TB LU × 1)	237 GB	On	Inline
4	HDD Dedup/Comp PostProcess Mode	SAS HDD RAID6(14D+2P) × 2 (VVOL= 5 TB LU × 1)	237 GB	On	Post-process
5	HDD Dedup/Comp Inline Mode	SAS HDD RAID6(14D+2P) × 2 (VVOL= 5 TB LU × 1)	237 GB	On	Inline

The following test configuration of backup volumes for each test case was used:

1. Test Case #1 (SSD Dedup/Comp off): Use one large volume without deduplication and compression made from **SSD**.
2. Test Case #2 (SSD Dedup/Comp Post-Process Mode): Use one large volume with deduplication and compression (Post-Process mode) made from **SSD**.
3. Test Case #3 (SSD Dedup/Comp Inline Mode): Use one large volume with deduplication and compression (Inline mode) made from **SSD**.
4. Test Case #4 (HDD Dedup/Comp Post-Process Mode): Use one large volume with deduplication and compression (Post-Process mode) made from **HDD**.
5. Test Case #5 (HDD Dedup/Comp Inline Mode): Use one large volume with deduplication and compression (Inline mode) made from **HDD**.

The following test procedure was used:

1. Install a SAP HANA scale-up system on server blades. Follow TDI storage design using dynamic provisioning pools. More details can be found in the [SAP HANA TDI on Hitachi Virtual Storage Platform G Series and VSP F Series with Hitachi Storage Virtualization Operating System Reference Architecture Guide](#).
2. Prepare a HANA backup data and store the data to a temporary folder in the SAP HANA database server. For this test, a BW/4HANA benchmarking data (237 GB) was prepared.
3. Create a backup dynamic provisioning pool and VVols for each test case and attach them to the SAP HANA database server.
4. Load the 237 GB data prepared in step 2. into the SAP HANA database for all test cases.
5. Run backup data and restore data by HANA Studio, and check the test results as shown below.

[For Test Case #1, start from (5). For the other test cases, start from (1).]

- (1) Test Case #2, #3, #4, and #5: Run a one-time backup, and check the backup time on the OS, saving effect ratio, and post-process time from the end of backup on the OS to the time when the saving effect become stable.
- (2) Test Case #2, #3, #4, and #5: Run a one-time backup, and check the backup time on the OS, saving effect ratio and post-process times from the end of backup on the OS to the time when the saving effect become stable.
- (3) Test Case #2, #3, #4, and #5: Clean up all data in the backup volume, and check that the disk usage is zero using Hitachi Storage Navigator not the OS.

(Caution: Don't clean up the loaded HANA database in memory when the backup volume is cleaning up.)

(Caution: Check the disk usage using Hitachi Storage Navigator. Even if the OS shows that the disk usage is zero, the storage array may keep data in its volumes.)

- (4) Test Case #1, #2, #3, #4, and #5: Run backup six times, and check the backup time on the OS, saving effect ratio, and post-process time from the end of backup on the OS to the time when the saving effect becomes stable.
- (5) Test Case #1, #2, #3, #4, and #5: Run three time restore, and check the restore time on the OS.

[The test is done. Clean up all backup data for the next test case run.]

Summary of Test Results

Figure 9 shows the saving effect of some test cases. All saving effects after the first backup were roughly 1.2:1. All saving effects after the second backup were roughly 1.2:1. And all saving effects after the sixth backup were roughly 3.4:1, The saving effects were almost the same in all cases.

Figure 9




Figure 10 and Figure 11 show the capacity saving behavior of Dedup/Comp Inline mode and Post-Process Mode in this test. The graph puts the ratio of backup volume usage on the vertical axis, and the ratio of the usage without saving effect is taken as '1'. The maximum volume usage of inline mode takes roughly half of postprocess mode and the backup time of post-process mode takes roughly three times of inline mode, because deduplication and compression run while writing data I/O in inline mode. But the post-process time of postprocess mode until the capacity saving completion takes quite longer than inline mode, because deduplication and compression run after writing data I/O in post-process mode.

Figure 10





Figure 11



Note — Throughout this evaluation, the SAP HANA database is not updated. Therefore, the saving effect ratio and time for a typical database with updated data would be worse than these saving effect ratios and times. Deduplication ultimately reduces redundancy.

Conclusions

The saving effect in all test configurations was roughly 1.2:1 after the first backup, roughly 1.2:1, after the second backup, and roughly 3.4:1 after the sixth backup. The post-process Mode for Deduplication and Compression is showing significant saving in Time for Backup / Restore Operations. The volume usage of inline mode takes roughly a half of post-process mode.

The Post Process Mode shows less Impact in I/O Performance. For SAP HANA DB the recommended usage should be Post Process Mode. SAP HANA is an in-memory database, I/O still plays a critical role for the performance of the system. For a data backup the current payload of the data volumes is read and copied to the backup storage. For writing a data backup it is essential that on the I/O connection there are no collisions with other transactional operations running against the database.

Business Impact

To summarize, the benefits for an organization by using data deduplication technologies are:

- Reduced back-up costs
- Reduced costs for disaster recovery
- Reduced costs for hardware
- Increased storage efficiency

Hitachi Vantara



Corporate Headquarters
5355 Augustine Drive
Santa Clara, CA 96054 USA
HitachiVantara.com | community.HitachiVantara.com

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

© Hitachi Vantara Corporation, 2019. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. All other trademarks, service marks, and company names are properties of their respective owners

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Vantara Corporation.

MK-SL-158-00, July 2019