HITACHI

Hitachi Enterprise Storage for SAP HANA Tailored Data Center Integration

Best Practices Guide

MK-SL-469-00
May 2026

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to Docs-Feedback@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

Changes	Date
• Initial release	April 2026

Contents

Introduction

The deployment of SAP HANA under the Tailored Data Center Integration (TDI) model shifts the responsibility of hardware performance and reliability from the appliance vendor to the infrastructure architect. To provide a rock-solid foundation for SAP HANA in a TDI environment, your storage layer must be more than just "fast", it needs to be predictable, resilient, and invisible to the end-user. The storage subsystem, the persistence layer, is the most critical component for ensuring that SAP HANA meets its Key Performance Indicators (KPIs) for Data and Log volumes. SAP HANA's in-memory architecture might give the impression that storage is secondary, but in daily operations and at peak times, the storage must maintain sub-millisecond latency on the Log volume while providing high-bandwidth throughput for Data volume savepoints and system restarts.

This paper outlines the technical best practices for deploying SAP HANA TDI on the latest generation of Hitachi Vantara storage solutions.

Hitachi's Virtual Storage Platform (VSP) family is meeting the aggressive I/O demands of HANA across storage generations, ranging from massive scale-out enterprise environments to high-performance mid-range footprints.

Hitachi Storage Portfolio for SAP HANA

Hitachi Vantara's current storage lineup leverages the Storage Virtualization Operating System (SVOS) RF to provide consistent data management across different hardware tiers. These models are certified under SAP HANA TDI Phase 5 and Phase 6, supporting both traditional Fibre Channel (FC) and NFS file services via VSP One File 30 clusters.

VSP One Family
The VSP One series represents the latest evolution in Hitachi's unified data platform, streamlining management through a new unified software stack: VSP360 supports management of block, file, object, and mainframe workloads within a single, consistent management framework.

- **VSP One Block High End**: A high-throughput system engineered to handle the massive parallel I/O requirements of SAP HANA scale-out clusters and consolidated scale-up and virtualized environments. It provides the necessary I/O power for daily operations, workload peaks and large-scale data loads.
- **VSP One Block 20**: Optimized for small and mid-range HANA footprints, providing high-density NVMe storage with simplified lifecycle management. VSP One B20 models are available in two flavors: High-performance with Triple-Level Cell and cost-optimized with Quad-Level Cell NVMe media.
- **VSP One File 30**: The Hitachi VSP One File 30 provides high-performance NAS capabilities for SAP HANA TDI environments, targeting the requirements of Data and Log volume persistency and the /hana/shared and backup filesystems. This solution integrates seamlessly into the VSP One unified management framework, offering robust data protection through snapshots and high-bandwidth throughput for SAP HANA workload and rapid database backups.

VSP 5000 Series
The VSP 5000 series is a flagship scale-out enterprise array utilizing a multi-node, global-cache architecture. It is designed for large-scale consolidated HANA landscapes where extreme I/O density and 100% data availability are required. Its high-bandwidth switching fabric architecture is specifically tuned to handle the heavy writes required during HANA savepoints and log commits.

VSP E1090
The VSP E1090 is a high-performance, all-NVMe storage system designed for organizations requiring enterprise-grade performance without the larger footprint of a 5000-series array. It's the sweet spot for many TDI deployments, offering the same SVOS RF (Storage Virtualization Operating System) intelligence found in the flagship models but optimized for low-latency NVMe media.

General Best Practices for SAP HANA TDI Success

With a SAP HANA appliance deployment, the hardware vendor pre-configures all hardware components – storage and server. With Hitachi systems for SAP HANA TDI deployments, you customize each installation by using any supported servers and certified enterprise storage vendor. You might be able to leverage your existing hardware to reduce your total cost of ownership (TCO) for a high-performance SAP HANA infrastructure. If a certificate provided by SAP is for a specific operating system, you can only use that operating system for SAP HANA in that implementation.

The goal of this integration is to move beyond mere compatibility and toward an architecture where the storage scales as the HANA database grows. By leveraging the specific strengths of the VSP storage models, IT teams can guarantee that the storage remains the most reliable part of the SAP stack.

Understanding HANA IO is crucial for selecting the right storage layout and drive type. The *SAP HANA storage requirements* Whitepaper provides details about the storage requirements for your SAP HANA system. The I/O Patterns section shows a detailed table for IO happening in the different scenarios of HANA operations.

Performance Relevant HANA Volumes

The operating system and the SAP HANA share volume for binaries, trace and config files have very low storage requirements for both performance and size. Thus, we will focus on HANA data and log volumes for a more detailed look.

The Log Volume Requirement: Must satisfy extremely low write latency (often sub-millisecond) and high sequential write throughput.

The Data Volume Requirement: Large capacity, primarily read-intensive access (once loaded), with periodic large sequential writes (savepoints).

For better understanding, we will have a look at the requirements of both data and log volume, to better understand the needs and base the decision on data, business needs and SLAs rather than gut feeling.

What kind of IO can we expect on the HANA volumes?

HANA Data Volume:

- Async random write to persistent storage
- Write transactions, savepoints, snapshots and delta merges
- Contains table data, undo log information, modeling data and more
- Data is read during data backup, HANA DB or HANA node restart and table loads

HANA Log Volume:

- Redo log information
- Synchronous write to persistent storage when a transaction is committed
- Overwrite after backup operation

The following table gives information about the patterns and block sizes in the HANA volumes for different I/O scenarios.

Scenarios	Data Volume	Redo Log Volume	Backup Storage
Write transactions		WRITE OLTP – mostly 4 KB sequential I/O orders: OLAP – larger I/O order sizes (up to 1 MB)	

Scenarios	Data Volume	Redo Log Volume	Backup Storage
Savepoint, Snapshot, Delta merge	WRITE 4 KB – 64 MB asynchronous parallel I/O orders (amount of data depends on system load)		
DB Restart, Failover, Takeover	READ 4 KB – 64 MB asynchronous parallel I/O orders (amount of data depends on RowStore size)	READ 256 KB asynchronous I/O orders	
Column store table load	READ 4 KB – 16 MB asynchronous parallel I/O orders		
Data Backup	READ 4 KB – 64 MB asynchronous I/O orders copied into buffers of 512 MB		WRITE 512 MB sequential I/O orders (configurable)
Log Backup		READ 4 KB – 128 MB asynchronous I/O orders copied into buffers of 128 MB	WRITE 4 KB – 128 MB sequential I/O orders
Database Recovery	WRITE 4 KB – 64 MB asynchronous parallel I/O orders	READ 256 KB asynchronous I/O orders	READ Data backup: 512 MB I/O orders copied into buffers of 512 MB Log backup: 128 MB I/O orders copied into buffers of 128 MB

See more details around this table in the *HANA storage requirements whitepaper*.

HANA Workload

During normal operations, there are mainly write operations for the HANA data and log volume while read operations only happen during database restart or failover, column table reload (data volume only) and when performing a log- or full data-backup (both volume types respectively).

With SAP's strategy for SAP HANA as a multi-model database that can handle multiple workload types including OLAP, OLTP and mixed workload scenarios in a single database, it is very difficult to map certain application scenarios to specific IO workload patterns and requirements.

In addition to that, consolidating multiple HANA nodes onto a single VSP storage system in a TDI environment will cause an additional mix in storage requirements.

VSP Enterprise Storage SAP HANA TDI node scalability

SAP provides the SAP HANA Hardware and Cloud Measurement Tool (HCMT) for certification and scalability assessment. See *SAP Note 2493172* (SAP user credentials required) for HCMT details. This enables Enterprise Storage providers to give recommendations regarding storage performance and HANA node scalability.

The following table shows Hitachi supported maximum number of SAP HANA systems connected in parallel to a single Virtual Storage Platform to pass the SAP enterprise storage KPIs. While running storage KPI testing using HCMT, the storage was fully dedicated to the testing and not shared by other systems or applications.

Storage	Node count	Storage layout
VSP One BHE 2N	84	9 × RAID 10 (2D+2D)
VSP One BHE 4N	168	18 × RAID 10 (2D+2D)
VSP One BHE 6N	252	27 × RAID 10 (2D+2D)
VSSO for 4 × VSP One BHE 6N	1008	108 × RAID 10 (2D+2D)
VSP One B24	11	1 × DDP RAID 6 (6D+2P)
VSSO for 10 × VSP One B24	110	10 × DDP RAID 6 (6D+2P)
VSP One B26	17	1 × DDP RAID 6 (14D+2P)
VSSO for 10 × VSP One B26	170	10 × DDP RAID 6 (14D+2P)
VSP One B28	32	1 × DDP RAID 6 (14D+2P)
VSSO for 10 × VSP One B28	320	10 × DDP RAID 6 (14D+2P)
VSP 5600 2N	74	14 × RAID 10 (2D+2D) for Data and 14 × RAID 10 (2D+2D) for Log
VSP 5600 4N	148	28 × RAID 10 (2D+2D) for Data and 28 × RAID 10 (2D+2D) for Log
VSP 5600 6N	222	42 × RAID 10 (2D+2D) for Data and 42 × RAID 10 (2D+2D) for Log

Storage	Node count	Storage layout
VSP E1090	50	7 × RAID 10 (2D+2D) for Data and 7 × RAID 10 (2D+2D) for Log
VSSO for 10 × VSP E1090	500	70 × RAID 10 (2D+2D) for Data and 70 × RAID 10 (2D+2D) for Log

Hitachi Virtual Storage Scale-out (VSSO) cluster for VSP Storages was certified for SAP HANA starting with the VSP E1090 storage.
This loosely clustered VSSO provides a flexible and scalable storage solution for SAP HANA Enterprise Storage with the following:

- Achieve greater HANA node scalability by adding up to the total number of HANA nodes for multiple VSP storage systems.
- No performance impact when adding more storage devices in a loosely clustered VSSO.
- Manage all storage arrays from a single management tool – Hitachi Ops Center.
- Physical connections between SAP HANA hosts and storage are the same as non-cluster storage for SAN (using Fibre Channel switches) and DAS.

When implementing an SAP HANA TDI infrastructure on Hitachi Virtual Storage Platform you do not have to use exactly the same storage design that was used for enterprise storage KPI testing. However, the recommendation is to follow these directions and guidelines for the setup and configuration of the storage system. See *SAP HANA Tailored Data Center Integration - Frequently Asked Questions* for more details.

The scalability testing on the VSP Enterprise Storage units is done in a lab environment and might also include approximations for larger storage models. However, in a real-world environment, there are a lot of factors that will influence storage performance and scalability in one way or the other. Common factors are different disk layout and RAID level, mixed workload with HANA and non-HANA workload sharing the same storage unit and additional features enabled like encryption and compression. And last but not least, different HANA workloads will have different IO requirements and peak workloads, and high usage phases might require higher storage throughput.

VSP Enterprise Storage Solutions for HANA TDI

Hitachi Vantara offers a robust portfolio of enterprise-class storage solutions specifically engineered and validated to meet the stringent performance and reliability requirements of SAP HANA TDI. This comprehensive suite includes the Hitachi VSP 5000 series, VSP E1090 and the next-generation VSP One Block family, all of which leverage advanced architectural designs to deliver industry-leading data access and resiliency. These solutions provide the scalability and high availability required to support mission-critical production workloads across both scale-up and scale-out architectures

To ensure our storage systems perform optimally within a TDI environment, Hitachi Vantara provides several best practices and specific implementation configurations. We will examine the granular settings for each solution, including RAID group and pool setup, LUN layout and sizing strategies our engineers recommend for maximizing throughput across the VSP 5000, VSP E1090, and VSP One series.

VSP enterprise storage solutions are engineered to facilitate the consolidation of SAP HANA workloads, supporting a diverse mix of scale-up and scale-out architectures on a single storage unit. Our architecture allows for the concurrent deployment of bare-metal and virtualized instances, with the capability to run mixed applications - both SAP HANA and non-SAP HANA - in parallel, provided they meet necessary performance requirements. Consolidation of variable node sizes and memory configurations within shared, optimized provisioning pools significantly reduces total cost of ownership (TCO) by maximizing resource utilization and simplifying management across the entire SAP landscape.

VSP One Block High End

The Hitachi Virtual Storage Platform (VSP) One Block High End (BHE) represents the pinnacle of our enterprise-class all-flash storage portfolio, specifically designed to anchor the most demanding SAP HANA production landscapes. Engineered for extreme performance and “eight nines” (99.999999%) availability, the BHE series utilizes a next-generation, all-NVMe architecture that delivers the ultra-low latency and massive throughput required for high-density consolidation and real-time data processing.

Unrivaled Performance and Massive Scalability

The VSP One Block High End is optimized for massive Tailored Data Center Integration (TDI) deployments, providing a scalable foundation that grows with your enterprise requirements:

- **Extreme Consolidation:** A single VSP BHE-6N (6-node configuration) is validated to support up to 252 active SAP HANA systems. Through Hitachi Virtual Storage Scale-out (VSSO), four BHE-6N arrays can be loosely clustered into a single management image supporting an unprecedented 1,008 SAP HANA systems.

- **Industry-Leading Cache and Throughput:** To eliminate I/O bottlenecks, the BHE-6N provides up to 6,144 GB (6 TB) of cache memory. The system features high-speed 64 Gbps Fibre Channel host connectivity and a native NVMe (PCIe) dual-port backend operating at 16 Gbps per link.

- **Optimized I/O Stack:** Running SVOS RF v10.5.1 or higher, the BHE utilizes a patented flash aware I/O stack to accelerate data access, ensuring sub-millisecond response times even under the most intensive transactional and analytical pressure.

Precision Architectural Design for SAP HANA

Our validated design for the VSP BHE focuses on maximizing resource utilization through a streamlined, high-performance pool architecture:

- **Unified Provisioning Pool:** Unlike segregated pool designs, our BHE reference architecture utilizes a single SAP_HANA_Pool for all volumes (Operating System, Shared, Data, and Log). This design leverages the system's massive internal bandwidth to provide maximum optimization at a lower cost than alternative solutions.

- **High-Performance RAID Configuration:** For production workloads, we recommend RAID 10 (2D+2D) across NVMe SSDs ranging from 3.8 TB to 60 TB, ensuring the highest level of write performance and endurance for critical SAP HANA log and data volumes.

- **Adaptive Efficiency:** Integrated Adaptive Inline Data Reduction and Thin Provisioning work in concert with SAP HANA's memory-resident data structures to enhance storage efficiency without compromising the performance required for mission-critical Service Level Agreements (SLAs).

Mission-Critical Resiliency and Availability

The VSP One Block High End is the best choice for organizations requiring zero-downtime operations, offering comprehensive protection against every level of system failure:

- **"Eight Nines" Availability:** The platform is engineered with a No Single Point of Failure (NSPOF) design, ensuring 100% data availability for production nodes through redundant controllers, channel boards, and power systems.

- **Comprehensive Non-Stop Maintenance:** Every critical component – including Control PCBs, DIMMs, cache flash modules (CFMM), fans, and battery modules – supports hot-swapping and non-disruptive maintenance to ensure continuous operation during hardware life-cycling.

- **Advanced Fencing and Failover:** Scale-out configurations utilize the SAP Storage Connector API (fcClient/fcClientLVM) with SCSI-3 persistent reservations to implement robust I/O fencing, ensuring data integrity during host auto-failover events.

The Enterprise Standard for SAP HANA

The VSP One Block High End is uniquely positioned as the premier solution for SAP HANA TDI due to its ability to consolidate hundreds of production systems onto a single, highly available infrastructure. With integrated VSP 360 management and support for Ansible-based automation, organizations can deploy, manage, and scale their SAP landscapes with the precision and speed required by modern digital businesses.

SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform One Block High End provides the full reference architecture developed by our Engineering team.

VSP One Block 20

The Hitachi VSP One Block 20 series (models VSP One B24, B26, and B28) is our next-generation, enterprise-class flash storage solution, engineered to deliver high-performance and high-availability infrastructure for SAP HANA production workloads. Powered by Hitachi SVOS RF v10.4.1 or higher, the platform is optimized for NVMe SSDs, offering the agility and scalability required for both scale-up and scale-out SAP HANA TDI deployments.

Architecture and Scalable Performance

Our architecture leverages an innovative design optimized for high-density flash media, ensuring minimal latency and high throughput for mission-critical SAP applications.

- **High-Speed Interconnects:** The platform supports Fibre Channel connectivity at 16, 32, and 64 Gbps, with high-end models providing up to 1024 GB of cache memory to handle intensive transactional processing.
- **NVMe Optimization:** Built specifically for NVMe SSDs ranging from 1.9 TB to 60 TB, the system utilizes a patented flash aware I/O stack to accelerate data access and enhance application responsiveness.
- **Scalability with VSSO:** The Hitachi Virtual Storage Scale-out (VSSO) capability allows for clustering up to 10 VSP One Block storage systems, enabling extreme scalability. For example, a single VSP One Block 28 supports 32 SAP HANA systems, while a clustered VSSO configuration can scale to support up to 320 systems.

Media Choice for VSP One Block 20 models - TLC and QLC Drives

The Hitachi VSP One Block 20 series offers a strategic choice between high-performance TLC (Triple-Level Cell) and cost-optimized QLC (Quad-Level Cell) NVMe media. Because these media types cannot be mixed within a single B20 system, architects must evaluate their SAP HANA workload characteristics to determine which configuration best meets their operational SLAs.

Selecting the TLC-Based VSP One Block 20

A TLC-configured system is the recommended choice for mission-critical production environments where latency and endurance are the primary drivers.

- **Ideal Workloads:** S/4HANA ERP, CRM, and very large servers that demand the lowest possible latency, especially for write-intensive Log Volumes.
- **Key Advantage:** Provides the high endurance and consistent performance required to prevent transaction backlogs during heavy transactional pressure.

Selecting the QLC-Based VSP One Block 20

A QLC-configured system establishes a high-capacity flash tier designed to replace legacy HDD or hybrid storage without the complexity of multi-tier architectures.

- **Ideal Workloads:** Read-intensive environments, such as OLAP/Analytics with scheduled data loads, QA/DEV landscapes, and the consolidation of smaller systems.
- **Key Advantage:** Offers the most economical cost per GB and maximum density making it ideal for Data Volumes where read access is dominant and write endurance is a secondary concern.

Strategic Deployment Strategy

For organizations with a diverse SAP landscape, the decision is not whether to use TLC or QLC, but rather where to deploy each system type within the broader storage environment. For instance:

- Deploy **TLC-based VSP One Block** units for primary, high-pressure transactional systems.
- Deploy **QLC-based VSP One Block** units for reporting systems, virtualized consolidation setups, or as secondary replication targets.

By mapping the specific I/O patterns of their HANA landscapes to the appropriate system configuration, organizations can achieve a balanced infrastructure that optimizes both performance and total cost of ownership (TCO).

Resiliency and Endurance Management

Hitachi has engineered the VSP One Block platform to mitigate the inherent endurance limitations of QLC technology. Through a combination of advanced caching, wear-management, and Dynamic Drive Protection (DDP), QLC SSDs can sustain demanding production workloads. Lab testing on VSP One B26 with 24 x 30TB QLC drives demonstrated that even under heavy, constant write pressure (6.2 PB of writes over 28 days), the used-endurance indicator only reached 3%, confirming the stability of the media for enterprise deployments. In the event that write endurance reaches 99%, the system automatically triggers Dynamic Sparing to save data to available spare capacity and ensure continuous operation.

Intelligent Storage Management with SVOS RF

The VSP One Block 20 utilizes Dynamic Drive Protection (DDP) and SVOS RF to simplify storage management while maximizing resource efficiency.

- **Consolidated DDP Pools:** To achieve maximum utilization at a lower cost, our validated design utilizes a single DDP pool to provision the operating system, SAP HANA shared, data, and log volumes.

- **Dynamic Drive Protection (DDP):** DDP RAID 6 configurations (6D+2P or 14D+2P) distribute data and parity across all drives in a group, eliminating the need for dedicated spare drives and reducing RAID complexity.

- **Adaptive Data Reduction (ADR):** Integrated always-on compression and deduplication enhance storage efficiency and reduce the overall CO2 footprint through Dynamic Carbon Reduction, which optimizes CPU power consumption.

- **Unified Management:** All storage resources are managed through a single console using Hitachi VSP360.

Enterprise-Grade Resiliency and Availability

The VSP One Block 20 is designed with No Single Point of Failure (NSPOF) to ensure continuous operations for production nodes.

- **Multipath High Availability:** Utilizing native Linux Device-mapper Multipath, the solution aggregates I/O paths into a single logical path. Hitachi recommends specific tunings to optimize path performance and ensure high availability.

- **Non-Stop Maintenance**: All critical hardware components, including Control PCBs, cache memory, power supplies, fans, and flash drives, support hot-swapping and non-disruptive maintenance.

- **Advanced Fencing:** For scale-out deployments, the SAP Storage Connector API (*fcClient*/*fcClientLVM*) is utilized to manage volumes and implement fencing through SCSI-3 persistent reservations, ensuring data integrity during host failover.

The *Reference Architecture paper* developed by Hitachi Engineering provides a full overview for VSP One Block 20 deployments for SAP HANA TDI.

VSP 5000

The Hitachi Virtual Storage Platform (VSP) 5000 series is engineered as an enterprise-class flash storage solution designed to meet the rigorous performance and availability demands of SAP HANA production environments. By leveraging a high-density, scale-out architecture optimized for NVMe and storage class memory, the advanced Hitachi Storage Virtualization Operating System RF (SVOS RF), and high-availability design, the platform provides the agility and massive scaling required for large-scale SAP HANA Tailored Data Center Integration (TDI) deployments

Architecture and Performance Optimization

Our architecture is built on the Hitachi Storage Virtualization Operating System RF (SVOS RF), which incorporates a patented flash aware I/O stack to accelerate data access and minimize latency. The VSP 5000 series delivers industry-leading throughput and scalability:

- **Flash-Aware I/O Stack:** A patented I/O stack within SVOS RF that accelerates data access to optimize flash performance.

- **Massive Scalability:** The architecture allows for scaling without performance slowdowns. A single VSP 5600 system configured with three node pairs is validated to support up to 222 SAP HANA systems.

- **High Data Transfer Rates:** The system supports data transfer rates of 12 Gbps for SAS SSDs and 8 Gbps for NVMe SSDs, utilizing up to 32 Fibre Channel ports per node for high-speed connectivity.

- **Memory Efficiency:** To support intensive transactional processing, the VSP 5500 and 5600 models provide up to 2 TB of cache memory per node pair.

Advanced Storage Management with SVOS RF

Hitachi Storage Virtualization Operating System RF (SVOS RF) provides the intelligence to manage SAP HANA workloads effectively.

To optimize SAP HANA workloads, Hitachi Dynamic Provisioning (HDP) pools are created to implement a specific provisioning pool architecture that ensures maximum resource utilization. Our recommended design utilizes two distinct provisioning pools:

- **OS_SH_DT_Pool**: Consolidates the operating system, SAP HANA shared, and SAP HANA data volumes to optimize capacity.

- **LG_Pool**: A dedicated pool for SAP HANA log volumes, ensuring that critical write operations for logging are isolated for peak performance.

This ensures maximum utilization and optimization at a lower cost than alternative solutions.

From a technical and scalability standpoint, this design is engineered to support up to 74 SAP HANA production nodes on a single Virtual Storage Platform 5600 2N system with a minimal number of parity groups and storage cache. This validation confirms that the segregated pool approach meets SAP enterprise storage KPI requirements while maintaining high efficiency in high-density production environments.

Enterprise-Grade Resilience and High Availability

The VSP 5000 series is designed with No Single Point of Failure (NSPOF) to mitigate business risk and ensure continuous data availability for production nodes.

- **No Single Point of Failure (NSPOF):** The solution design uses redundant Fibre Channel ports and cables to connect node HBAs to storage target ports, ensuring both high availability and expected throughput for SAP HANA volumes.

- **Multipath High Availability:** Utilizing Device-mapper Multipath, the solution aggregates physical I/O paths into a single logical path, ensuring that LUNs remain accessible even in the event of a path failure.

- **Non-Stop Maintenance:** The platform supports the hot-swapping and non-disruptive maintenance of all critical components, including Control PCBs, cache memory, power supplies, fans, and flash drives.

- **Fencing Mechanisms:** For HANA scale-out configurations, the implementation of the SAP Storage Connector API (fcClient/fcClientLVM) utilizes SCSI-3 persistent reservations to provide robust fencing during host failover

- **Logical Volume Manager (LVM)**: We use an LVM-based storage layout to manage operating system and HANA volumes effectively and to ensure high availability and performance.

In our Reference Architecture paper you can find more details on the VSP 5000 setup and configuration for SAP HANA TDI.

VSP E1090

The Hitachi VSP E1090 is a purpose-built, all-NVMe storage solution engineered to deliver enterprise-class performance and ultra-low latency for mid-to-large-scale SAP HANA TDI workloads. Powered by Hitachi SVOS RF v9.8, the E1090 bridges the gap between mid-range flexibility and high-end enterprise power, providing a reliable foundation for mission-critical SAP applications.

Performance and Scalability

The VSP E1090 is designed to handle high-density consolidation with significant throughput and memory resources:

- **High-Performance Interconnects:** The platform features 32 Gbps Fibre Channel connectivity and supports up to 8 Channel Board (CHB) pairs, ensuring high-speed data transfer across the SAN.

- **Cache Processing Power:** With a cache memory capacity of 1024 GB, the E1090 provides the necessary buffer to maintain sub-millisecond response times during intensive SAP HANA transactional and analytical processing.

- **Massive Consolidation:** A single VSP E1090 system is validated to support up to 50 SAP HANA systems. When utilizing Hitachi VSSO to cluster ten E1090 arrays, the environment can scale to support up to 500 production systems under a single management image.

Architectural Efficiency and Management

Our HANA TDI solution for the E1090 utilizes advanced virtualization features to maximize the ROI of NVMe investments:

- **Segregated Pool Design:** To ensure peak performance for latency-sensitive operations, the E1090 implementation utilizes a dual-pool architecture. The *OS_SH_DT_Pool* manages the operating system, shared binaries, and data volumes, while the *LG_Pool* is dedicated to SAP HANA log volumes to isolate write-intensive logging traffic.

- **Flexible RAID Configurations:** The platform supports high-performance RAID 10 (2D+2D) for maximum speed or RAID 6 (6D+2P) for enhanced capacity efficiency, both of which are fully compliant with SAP performance standards.

- **SVOS RF Optimization:** The E1090 leverages the same patented flash aware I/O stack found in our high-end arrays, which accelerates data access and optimizes the endurance of NVMe SSDs ranging from 1.9 TB to 15.3 TB.

Enterprise Resiliency and High Availability

Built with a No Single Point of Failure (NSPOF) philosophy, the VSP E1090 ensures data remains available to production nodes through any hardware or path failure.

- **Non-Stop Maintenance:** All critical components – including Control PCBs, cache memory, power supplies, and fans – are hot-swappable, allowing for non-disruptive repairs and firmware updates.

- **Standardized Path Management:** The solution utilizes native Device-mapper Multipath for Linux, ensuring high availability through redundant I/O paths. Specific Hitachi tunings ensure optimal path selection and failover performance.

Find further information in our VSP E1090 Reference Architecture paper.

Virtual Storage Scale-out (VSSO) configuration

Virtual Storage Scale Out (VSSO) is a scale-out feature that allows multiple physical storage systems to be handled as a single mass storage system by defining the common virtual storage machine (VSM) for multiple physical storage systems. This doesn't change physical connections between SAP HANA host and each storage system (using either DAS or SAN).
The following provides guidance to implement VSSO:

- Make physical connections between each SAP HANA host and a storage array (via DAS or SAN), same as non-cluster storage.

- Storage provisioning:

    - Infrastructure level provisioning (creating HDP or DDP, Virtual Volumes, server registration, and volume attachment) are configured from each storage system.

- VSSO with multiple storage arrays is set up using Hitachi Ops Center (a centralized management interface for Hitachi storage environments to manage multiple storage systems from a single console).

- VSSO is configured using a new single storage serial number. When a virtual storage machine (VSM) is created, a new storage resource group is generated for a VSSO to allow you to logically group together multiple storage resources (such as pools, volumes and ports) within VSSO.

- Logical provisioning (for example, expanding VSSO with more storages, creating volumes, adding host groups, and attaching volumes) is configured using Hitachi Ops Center.

- Refer to the VSSO setup video: *Introduction to Virtual Storage Scale Out – YouTube*

VSP One File 30

The addition of a Hitachi VSP One File 30 cluster to a Hitachi VSP Enterprise Block storage provides a high-performance, Network Attached Storage (NAS)-based infrastructure for SAP HANA TDI. This solution leverages the reliability of our enterprise-grade all-flash block storages and extends their capabilities through a scalable, high-availability NAS layer specifically optimized for mission-critical file-based workloads.

NAS-Based Architecture for SAP HANA

By integrating an NAS-based setup, we offer our customers scalable file storage that seamlessly adapts to the growing data demands of SAP HANA environments. In this configuration, the SAP HANA shared, data, and log volumes are provided via the NFS protocol from the VSP One File cluster, while the underlying block storage is a certified member of the VSP family. This architecture simplifies management and enhances data availability across the network, allowing IT teams to handle complex storage tasks with significantly reduced administrative overhead.

While our *reference architecture* specifically highlights the performance of VSP One File 30 clusters in conjunction with the VSP E1090, it is important to note that these NAS clusters are certified for use with every certified Hitachi Enterprise Storage solution, providing organizations the flexibility to utilize their existing Hitachi infrastructure to reduce total cost of ownership (TCO).

VSP One File 30 Features

The VSP One File 30 series (Models 34 and 38) is engineered for operational simplification and high agility. Key features identified in our validation include:

- **Massive File Scalability:** Supports up to 130 billion files per file system and up to 20,000 shares or 10,000 exports per namespace.

- **High-Speed Networking:**

    - **VSP One File 34:** Optimized for 25GbE connectivity with 8 ports per node.
    - **VSP One File 38:** A performance model supporting up to 100GbE ports for leading-edge throughput requirements.

- **Protocol Flexibility:** Full support for NFS, SMB, iSCSI, and FTP, with specific optimizations for SAP HANA, such as NFSv3 for scale-up and NFSv4.1 for scale-out deployments.
- **Enterprise Virtual Servers (EVS):** Architecture utilizes EVS (e.g., *EVS_HANA_DATA*, *EVS_HANA_LOG*) to logically separate HANA volumes, with each EVS possessing its own IP address and metadata for enhanced isolation.
- **Advanced Data Management:** Features include always-on thin provisioning, link aggregation (LACP), and integrated analytics for rogue client handling.
- **High Performance Metrics:** The platform is validated to support up to 64,000 active connections and 1.5 million open files per node.

Integrated Storage Design and Pool Management

To ensure maximum resource utilization and performance isolation, our solution design employs **Hitachi Dynamic Provisioning (HDP)** to create a segregated block storage layer on the VSP E1090.

- **Segregated Pool Architecture:** We utilize two distinct provisioning pools: the OS_SH_Data_Pool for the operating system, shared binaries, and data volumes, and a dedicated LG_Pool for log volumes.
- **System Drive Configuration:** On both HDP pools, we create 32 virtual volumes (vVols) of equal size, which serve as the system drives for the VSP One File 30 cluster. These drives are combined into logical "spans" (e.g., *HANASHDT* and *HANALOG*) to host the file systems.
- **Cache Processing Power:** The underlying VSP E1090 provides 1024 GB of cache memory, facilitating rapid data access and low latency required for real-time SAP HANA analytics.

Operational Resiliency and High Availability

Built with a No Single Point of Failure (NSPOF) philosophy, the combined solution ensures data integrity and business continuity.

- **Network Redundancy:** Each SAP HANA node is connected to the VSP One File cluster via redundant Ethernet connections using an Ethernet switch pair. We recommend separating network traffic for SAP HANA shared volumes onto a different VLAN from the data and log volume traffic.

- **Non-Stop Maintenance:** The platform supports the hot-swapping and non-disruptive maintenance of all critical components, including Control PCBs, cache memory, power supplies, fans, and flash drives.

- **Advanced Data Protection:** Integrated features such as snapshots, replication, and automated backups ensure that mission-critical SAP HANA data remains secure and recoverable.

- **Host Auto-Failover:** For scale-out deployments, the architecture supports SAP HANA host auto-failover through NFSv4.1 file locking, ensuring that standby nodes can seamlessly take over in the event of a primary node failure.

Validated Scalability and Performance

Our extensive engineering validation, conducted with the SAP HANA Hardware and Cloud Measurement Tool (HCMT), confirms the solution's ability to meet stringent SAP KPIs.

- **System Density:** A single node of VSP One File 38 is validated to support up to 12 concurrent SAP HANA production systems.

- **Optimized Performance Parameters:** Testing utilized RAID 10 (2D+2D) configurations for both data and log pools to ensure high-speed write performance.

- **Futureproofing:** The architecture is designed to support newer SAP HANA versions and integrates with hybrid cloud environments, enabling organizations to manage data across on-premises and cloud locations effortlessly through the VSP One platform.

Validated Performance and Reliability

Our lab testing confirms that a 2-node cluster of VSP One File 38 can support up to 12 concurrent SAP HANA production systems while meeting all SAP enterprise storage KPIs. To ensure continuous operations, we utilize redundant Ethernet connections between the HANA nodes and the VSP One File cluster to achieve No Single Point of Failure (NSPOF). Furthermore, the system is designed for non-stop maintenance, allowing for the hot swapping of control PCBs, power supplies, fans, and flash drives without impacting the production landscape.

See the *Reference Architecture paper* SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform E1090 with Hitachi VSP One File 30 for full details.

NFS-Based SAP Workloads and the Role of Hitachi VSP One File 30

In many SAP landscapes particularly SAP NetWeaver and other non-HANA database deployments (Oracle, IBM Db2, Microsoft SQL Server, SAP ASE) – Network File System (NFS) storage plays a critical role in providing shared filesystem services across distributed application servers. Unlike the database layer, which is typically deployed on block storage for latency sensitive I/O, the SAP application layer relies heavily on shared

file systems that must be accessible concurrently from multiple hosts. NFS has therefore become a widely adopted mechanism for delivering scalable shared storage to SAP environments.

NFS Use Cases in SAP Non-HANA Deployments

Several core SAP directories require shared access across multiple application servers and cluster nodes. Typical examples include:

- **/sapmnt** – global SAP system directory containing kernel binaries and shared configuration files
- **/usr/sap/trans** – transport directory used for change and transport management across the SAP landscape
- **Interface directories** used for data exchange between SAP and external applications
- **SAP installation repositories and staging areas** used during patching and lifecycle management

These directories must be accessible simultaneously from all SAP application instances. NFS provides a simple and efficient method for exporting these file systems from a centralized storage system to multiple SAP hosts.

In high-availability SAP NetWeaver clusters (for example, ASCS/ERS configurations), shared file systems such as /sapmnt are often mounted from an external NFS server to ensure consistent access across cluster nodes.

Beyond core SAP directories, NFS is commonly used for:

- SAP application layer shared storage
- Transport and integration directories
- Batch interface file exchange
- Centralized log repositories
- Backup and archive staging areas

These workloads are characterized by moderate throughput, large namespace requirements, and multi-client access, making scale-out file storage platforms an ideal architectural fit.

Hitachi VSP One File 30 for SAP NFS Workloads

The Hitachi VSP One File 30 is designed as a scale-out enterprise file storage platform capable of delivering highly available NFS services for mission-critical enterprise applications such as SAP. Within SAP landscapes, the platform provides a consolidated and resilient file service layer that can support both SAP and adjacent enterprise

workloads. The unified approach of the Hitachi Virtual Storage Platform One portfolio simplifies infrastructure operations and enables consistent management across mixed workloads.

Key architectural attributes that make VSP One File 30 well suited for SAP file workloads include:

- **High Availability and Cluster Architecture**:
  VSP One File nodes can operate in clustered configurations, providing continuous file service availability. In a two-node configuration, one node can service I/O workloads while the second provides redundancy and failover protection. This architecture aligns with SAP's operational requirement for uninterrupted access to shared file systems.

- **High-Performance NFS Services**:
  The platform supports enterprise NFS features such as NFS v4.1 multi-pathing, nconnect, and link aggregation, enabling high bandwidth and resilient connectivity for Linux-based SAP hosts. These capabilities allow SAP application servers to scale throughput across multiple network connections while maintaining consistent latency characteristics.

- **Large Namespace and File System Scalability**:
  SAP landscapes frequently include multiple systems (development, quality assurance, and production), each requiring dedicated shared directories. The scale-out design of VSP One File enables organizations to consolidate these file systems into a single platform while maintaining strong performance isolation and capacity scalability.

- **Enterprise Data Management Capabilities**:
  VSP One File integrates enterprise storage services such as:
  - policy-based data tiering
  - advanced data protection
  - snapshot-based backup integration
  - compliance features such as WORM

These capabilities enable SAP administrators to protect critical application data and meet governance requirements without introducing additional storage silos.

Strategic Value for SAP Environments

From a solution architecture perspective, VSP One File 30 provides a centralized NFS storage layer optimized for SAP application workloads. By consolidating shared SAP directories and interface file systems onto a resilient file cluster, organizations gain several advantages:

- simplified SAP landscape management
- high availability for critical shared directories
- predictable performance for distributed SAP instances

- scalable capacity for growing transport and interface datasets

In large enterprise environments with multiple SAP systems, the platform enables consolidation of SAP shared storage into a single highly available file infrastructure, reducing operational complexity while maintaining enterprise-class performance and reliability.

NFS remains a key component of SAP architecture for non-HANA deployments, primarily supporting shared directories and application-layer file exchange across distributed systems. A platform such as Hitachi VSP One File 30 provides a robust enterprise-grade foundation for these workloads by combining high-performance NFS services, cluster-based availability, and scalable file system architecture. This makes it well suited to serve as the shared storage backbone for modern SAP landscapes while integrating seamlessly into broader enterprise storage strategies.

Validated Engineering and Scalability Testing

VSP One Block High End: The VSP One Block High End has undergone extensive engineering validation and scalability testing providing data that shows a VSP One BHE-6N setup can support up to 252 active SAP HANA systems.

VSP One Block 20: To ensure the VSP One Block 20 series meets or exceeds SAP enterprise storage Key Performance Indicators (KPIs). Testing confirms that our LVM-based storage layout, using specific striping (e.g., 1024 KB stripe size for Block 26 and 256 KB for Block 24/28), delivers the performance required for certified SAP HANA production environments.

VSP One File 30: Engineering validation has demonstrated that a

- VSP One File 34 node can support up to 8
- VSP One File 38 node can support up to 12

concurrent SAP HANA production systems when connected to a VSP E1090 block storage using NFSv3 for scale-up consolidation. 10 GbE network connections for NFS traffic were used for Engineering validation and scalability testing of VSP One File 30.

VSP 5000: The VSP 5000 series has undergone rigorous engineering validation to meet SAP enterprise storage KPI requirements. Testing confirmed that various models (VSP 5100, 5200, 5500, and 5600) using SSDs can handle the maximum number of validated active production nodes in an SAP HANA deployment. Specifically, the VSP 5600 provides up to 2 TB of cache memory per pair node to support intensive production processing.

VSP E1090: The VSP E1090 has undergone rigorous engineering validation using the SAP HANA Hardware and Cloud Measurement Tool (HCMT). This testing confirms that the platform meets all SAP enterprise storage Key Performance Indicators (KPIs), ensuring that consolidated SAP HANA landscapes remain stable and performant under heavy production stress.

Hitachi Vantara has conducted extensive engineering validation to ensure every certified member of the VSP family meets or exceeds SAP enterprise storage Key Performance Indicators (KPIs). Our testing protocols involve, depending on the time of testing, the SAP HANA Hardware Configuration Check Tool (HWCCT) and later the Hardware and Cloud Measurement Tool (HCMT), confirming that the storage remains compliant with SAP's stringent requirements for production-level deployments, ensuring not only Enterprise class performance but also data consistency, resiliency and high availability.

Implementation Best Practices

The successful implementation of TDI on VSP systems requires a disciplined adherence to validated deployment methodologies and engineering standards. As the complexity of SAP landscapes increases, transitioning from architectural design to operational readiness demands the integration of automation and unified management to mitigate human error and accelerate time-to-deployment.

This chapter outlines the essential implementation best practices for standardized provisioning, operating system configuration for optimized I/O performance, and the configuration of high-availability networking and storage stacks. These practices are designed to ensure that the infrastructure delivers the low latency and high throughput required for mission-critical transactional and analytical workloads. By following these guidelines, organizations can ensure that their consolidated HANA environments remain compliant with SAP performance standards while achieving the maximum levels of reliability and efficiency inherent in the Hitachi VSP portfolio.

Storage Provisioning

Based on the provided technical documentation, the following tables detail the pool and RAID group assignments, LDEV IDs, sizes, and MPU assignments for each validated Hitachi storage unit configuration.

Hitachi Virtual Storage Platform One Block High End

The following table outlines the physical parity group and LDEV configuration for the unified SAP_HANA_Pool utilizing RAID 10.

HDP Pool	PG ID	PG RAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
SAP_HANA_Pool	1	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:01	1760 GB	MPU-010
			00:00:02	1760 GB	MPU-020
			00:00:03	1760 GB	MPU-110
			00:00:04	1760 GB	MPU-120
	2	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:05	1760 GB	MPU-010
			00:00:06	1760 GB	MPU-020
			00:00:07	1760 GB	MPU-110

HDP Pool	PG ID	PG RAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
			00:00:08	1760 GB	MPU-120
	3	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:09	1760 GB	MPU-010
			00:00:0A	1760 GB	MPU-020
			00:00:0B	1760 GB	MPU-110
			00:00:0C	1760 GB	MPU-120
	4	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:0D	1760 GB	MPU-120
			00:00:0E	1760 GB	MPU-020
			00:00:0F	1760 GB	MPU-110
			00:00:10	1760 GB	MPU-120
	5	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:11	1760 GB	MPU-010
			00:00:12	1760 GB	MPU-010
			00:00:13	1760 GB	MPU-010
			00:00:14	1760 GB	MPU-010
	6	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:15	1760 GB	MPU-010
			00:00:16	1760 GB	MPU-020
			00:00:17	1760 GB	MPU-110
			00:00:18	1760 GB	MPU-120
	7	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:19	1760 GB	MPU-010
			00:00:1A	1760 GB	MPU-020
			00:00:1B	1760 GB	MPU-110
			00:00:1C	1760 GB	MPU-120
	8	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:1D	1760 GB	MPU-010
			00:00:1E	1760 GB	MPU-020
			00:00:1F	1760 GB	MPU-110
			00:00:20	1760 GB	MPU-120
	9	RAID10 (2D+2D) on 3.8 TB NVMe SSDs	00:00:21	1760 GB	MPU-010
			00:00:22	1760 GB	MPU-020
			00:00:23	1760 GB	MPU-110

HDP Pool	PG ID	PG RAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
			00:00:24	1760 GB	MPU-120

Hitachi Virtual Storage Platform One Block 20

The following table outlines the physical parity group and LDEV configuration for the unified SAP_HANA_Pool utilizing DDP RAID 6.

HDP Pool	PG ID	DDRAID Level and Disks	LDEV ID	LDEV Size	MPU Assignment
SAP_HANA_Pool	0	DDP RAID 6 (14D+2P) on 16 × 1.92 TB NVMe SSD	0	1319 GB	MPU-10
			1	1319 GB	MPU-20
			2	1319 GB	MPU-10
			3	1319 GB	MPU-20
			4	1319 GB	MPU-10
			5	1319 GB	MPU-20
			6	1319 GB	MPU-10
			7	1319 GB	MPU-20
			8	1319 GB	MPU-10
			9	1319 GB	MPU-20
			10	1319 GB	MPU-10
			11	1319 GB	MPU-20
			12	1319 GB	MPU-20
			13	1319 GB	MPU-20
			14	1319 GB	MPU-10
			15	1319 GB	MPU-20
			16	1319 GB	MPU-10
			17	1319 GB	MPU-10

Hitachi Virtual Storage Platform 5000 Series

The VSP 5000 series utilizes a segregated pool architecture for high-density production nodes.

Dynamic Provisioning Pool	PG ID	RAID Level and Disks	LDEV Name	LDEV Size	MPU Assignment
OS_SH_DT_Pool	1	RAID 6 (6D+2P) on 960 GB SSDs	OS_SH_DT_1	1320 GB	MPU-010
			OS_SH_DT_2	1320 GB	MPU-020
			OS_SH_DT_3	1320 GB	MPU-110
			OS_SH_DT_4	1320 GB	MPU-120
	2	RAID 6 (6D+2P) on 960 GB SSDs	OS_SH_DT_5	1320 GB	MPU-010
			OS_SH_DT_6	1320 GB	MPU-020
			OS_SH_DT_7	1320 GB	MPU-110
			OS_SH_DT_8	1320 GB	MPU-120
	3	RAID 6 (6D+2P) on 960 GB SSDs	OS_SH_DT_9	1320 GB	MPU-010
			OS_SH_DT_10	1320 GB	MPU-020
			OS_SH_DT_11	1320 GB	MPU-110
			OS_SH_DT_12	1320 GB	MPU-120
LG_Pool	4	RAID 6 (6D+2P) on 960 GB SSDs	LG_1	1320 GB	MPU-010
			LG_2	1320 GB	MPU-020
			LG_3	1320 GB	MPU-110
			LG_4	1320 GB	MPU-120

Hitachi VSP E1090

The standard block configuration for the E990 and E1090 utilizes **RAID 6 (6D+2P)** for capacity efficiency.

Dynamic Provisioning Pool	PG ID	RAID Level and Disks	LDEV Name	LDEV Size	MPU Assignment
OS_SH_Data_Pool	1	RAID 6 (6D+2P) on 1.9 TB SSD	OS_SH_DT_DPVOL_1	2640 GB	MPU-10
			OS_SH_DT_DPVOL_2	2640 GB	MPU-20
			OS_SH_DT_DPVOL_3	2640 GB	MPU-10
			OS_SH_DT_DPVOL_4	2640 GB	MPU-20

Dynamic Provisioning Pool	PG ID	RAID Level and Disks	LDEV Name	LDEV Size	MPU Assignment
Log_Pool	4	RAID 6 (6D+2P) on 1.9 TB SSD	LG_DPVOL_1	2640 GB	MPU-10
			LG_DPVOL_2	2640 GB	MPU-20
			LG_DPVOL_3	2640 GB	MPU-10
			LG_DPVOL_4	2640 GB	MPU-20

LDEV Setup for SAN Storage

While it is not limited to these systems, we are using the following four active SAP HANA systems as examples for a TDI deployment:

- System 1 – 384 GB
- System 2 – 768 GB
- System 3 – 1536 GB
- System 4 – 3072 GB

Provision the storage for the four SAP HANA systems listed above and follow these guidelines:

Determine the minimum sizes for operating system, data, log, and HANA shared using these formulas in the SAP white pager SAP HANA Storage Requirements, as follows:

- Every HANA node requires approximately 100 GB capacity for the operating system.
- */hana/shared* size uses formulas:

  *Single node **(**scale-up**)** — Size = MIN **(**1 × RAM**; 1 TB)***
  *Multinode **(**scale-out**)** — Size = 1 × RAM_of_worker per 4 worker nodes*

- Data size requires at least 1 × RAM on each HANA node
- Log size uses these formulas:
  - Systems with equal or less than 512 GB memory – size = ½ × RAM
  - Systems with greater than 512 GB memory – size = 512 GB

Provision the storage according to the previous sections and assign all LDEVs to the dedicated pool.

Create virtual volume LDEVs for the operating system, SAP HANA shared, log, and data volumes. The following table lists examples for HANA systems with memory of 384 GB, 768 GB, 1536 GB, and 3072 GB where *N#* stands for the Node number of the systems:

LUN ID	Volume Name	File System Mount Point	Size System 1	Size System 2	Size System 3	Size System 4
0000	HANA_OS_N#	*/(root)*	100 GB	100 GB	100 GB	100 GB
0001	HANA_SH_N#	*/hana/shared*	384 GB	768G	1536G	3072 GB
0002	HANA_LOG_N#_1	*/hana/log*	48 GB	96 GB	128 GB	128 GB
0003	HANA_LOG_N#_2		48 GB	96 GB	128 GB	128 GB
0004	HANA_LOG_N#_3		48 GB	96 GB	128 GB	128 GB
0005	HANA_LOG_N#_4		48 GB	96 GB	128 GB	128 GB
0006	HANA_DATA_N#_1	*/hana/data*	96 GB	192 GB	384 GB	768 GB
0007	HANA_DATA_N#_2		96 GB	192 GB	384 GB	768 GB
0008	HANA_DATA_N#_3		96 GB	192 GB	384 GB	768 GB
0009	HANA_DATA_N#_4		96 GB	192 GB	384 GB	768 GB

While mapping the LUN path assignment for each node, add the volumes with the LUN IDs of the above table.

Hitachi VSP One File 30 with VSP E1090

For the NAS-based setup with VSP One File 30, the underlying E1090 uses two HDP pools

HDP Pool	PG ID	RAID Level and Disks	LDEV Name	LDEV Size	MPU Assignment
OS_SH_Data_Pool	1 to 2	RAID 10 (2D+2D) on 1.9 TB NVMe SSDs	OS_SH_DT_DPVOL_1	2640 GB	MPU-10
			OS_SH_DT_DPVOL_2	2640 GB	MPU-20
			...		
			OS_SH_DT_DPVOL_7	2640 GB	MPU-10
			OS_SH_DT_DPVOL_8	2640 GB	MPU-20
Log_Pool	3	RAID 10 (2D+2D) on 1.9 TB NVMe SSDs	LG_DPVOL_1	2640 GB	MPU-10
			LG_DPVOL_2	2640 GB	MPU-20
			LG_DPVOL_3	2640 GB	MPU-10
			LG_DPVOL_4	2640 GB	MPU-20

On both HDP pools, create 32 virtual volumes (vVols) of equal size. These vVols will be the system drives (SDs) on the VSP One File 30 cluster. System drives are physical storage devices that are combined to create a storage pool. A storage pool (span) is a logical grouping of SDs that acts as a container for storing file systems which are data structures within the span that store user data, metadata, and other information.

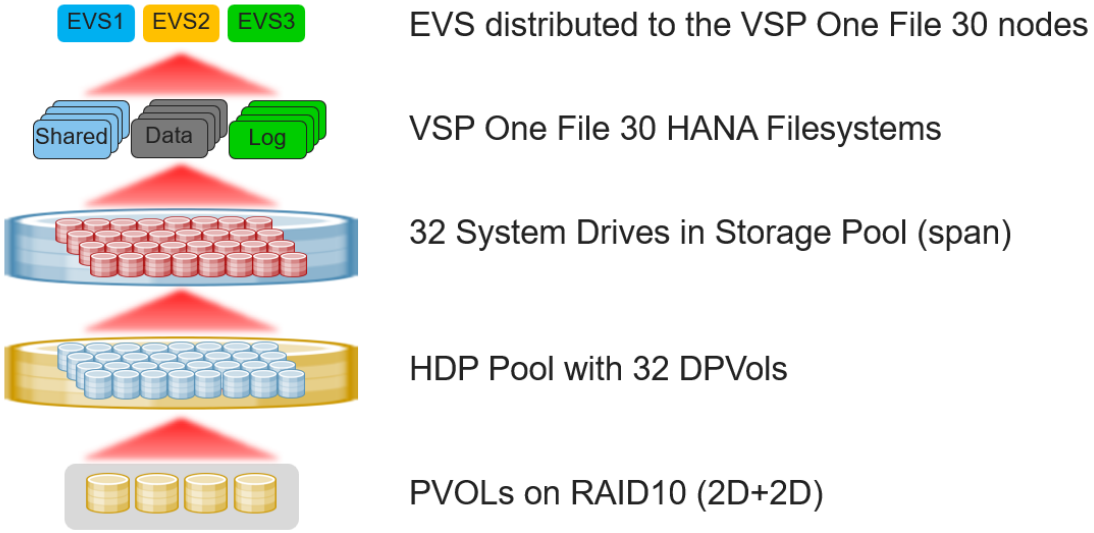

The size of the vVols on the OS_SH_Data_Pool does not have to be the same as the LG_Pool. When using different RAID or drive types for the two pools the total pool size will be different.

See the *reference architecture* guide for full details on the storage configuration for VSP One File 30 with VSP E1090. While other storage models might have a different pool layout, the VSP One File 30 part of the setup will be the same.

HANA Volumes for NAS Storage

Following these sizes, create file systems for each of your SAP HANA instances:

- *HANA_SHARED_<SID>* and *HANA_DATA_<SID>* in the *HANASHDT* span
- Select the 32KB block size option
- *HANA_LOG_<SID>* in the *HANALOG* span
- Select the 32KB block size option

Assign the file systems to the corresponding EVS created earlier.

On each EVS, create NFS exports for each file system.

Export configurations using these patterns:

- **Export Name**: */hana_shared_<SID>, /hana_data_<SID> and /hana_log_<SID>*

- **Path**: */hana/shared/<SID>, /hana/data/<SID> and /hana/log/<SID>*
- **Access configuration**: *(rw,sync,no_root_squash,no_subtree_check)*

Additional system-wide settings for the VSP One File 30 cluster are detailed in the following table.

Setting	Value
Max. NFS version	SAP HANA scale-up: NFSv3
	SAP HANA scale-out: NFSv4.1
MTU size	9000
read-ahead	disable
fs-accessed-time	off
filesystem-thin	on
security-mode	unix

Operating System Configuration for SAN Storage

This section describes the operating system setup for SAP HANA volumes storage connectivity.

Multipath configuration

This reference architecture uses Device-mapper Multipath, a native component of the Linux operating system. Using Device-mapper Multipath allows the configuration of multiple I/O paths between the server blades and storage.

Each node has two I/O paths connected with the storage. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available unless both paths fail.

Device-mapper Multipath is used for the following I/O paths:

- SAP HANA server boot volume
  Not applicable if using internal disks for the boot volume.
- SAP HANA data volume
- SAP HANA log volume
- SAP HANA shared volume

For a scale-out solution, Hitachi Vantara uses multiple paths to the disks with two options:

- Hitachi VSP One File 30 as NFS for /hana/shared.

- Global File System 2 (GFS2) for /hana/shared. This is for Hitachi-only TDI scale-out deployments where Hitachi provides server and storage for the TDI landscape.

Hitachi Vantara recommends the multipath settings listed in the following table.

Section	Option	SUSE Linux Enterprise Server for SAP Applications	Red Hat Enterprise Linux
defaults		N/A	N/A
blacklist	devnode	^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]* ^hd[a-z] ^dcssblk[0-9]*	
devices/device	vendor	HITACHI	
	product	.*	
	user_friendly_names	No	
	path_checker	direction	
	path_grouping_policy	Multibus	
	path_selector	queue-length 0	
	uid_attribute	ID_SERIAL	
	failback	Immediate	
	rr_weight	Uniform	
	rr_min_io_rq	128	1
	features	0	
	no_path_retry	5	

SAP HANA persistent storage volume configuration

For both operating systems, SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux, Hitachi Vantara uses an LVM-based storage layout. Once the operating system is installed and multipathing is correctly configured, you can see the assigned LUNs in the following directories:

- */dev/mapper*
- */dev/disk/by-id For example:*
- */dev/mapper/360060e801227fc00504027fc00000101*

- */dev/disk/by-id/scsi-360060e801227fc00504027fc00000101*

The last 6 digits of this number indicate the LDEV ID you have used during the LUN assignment. In the example above, 000101 maps to LDEV ID: 00:01:01.

For all the LUNs besides the one hosting the operating system, you need to initialize the LUNs for use by LVM, running the *pvcreate* command, which is part of the lvm2 rpm package, for example:

```
pvcreate -ff -y /dev/mapper/360060e801227fc00504027fc00000101
```

After you have prepared all the LUNs, you need to configure the volume groups using the *vgcreate* command. The names for the volume group differ between scale-up and scale-out installations.

- The volume groups for scale-up use vgdata, vglog, and vgshared.
- The volume groups for scale-out also include the SAP system ID as well as the node number. For example, *vgHITdata001*, *vgHITlog001*, *vgHITdata002*, and *vgHITlog002*.

The command to create the volume group takes no specific options. The following example creates the volume group for SAP HANA log in a scale-up scenario using four physical disks/LUNs:

```
vgcreate vglog /dev/mapper/360060e801227fc00504027fc0000010\[2,3,4,5\]
```

For creating other volume groups, use the same syntax, exchanging the volume group name as well as the physical disks or LUNs.

When creating the volume groups, you need to create a logical volume on top. The general syntax is the following:

```
lvcreate yes extents=100%VGstripes <# luns> --stripesize 1024 --name
```

Use the following table to complete the creation of logical volumes.

System Type	Number of LUNs	lv Name	vg Name
DATA	4 - following this reference architecture, or the number of assigned LUNs	lvdata	• Scale-up: vgdata • Scale-out: vg<SID>data<node number>
LOG	4 - following this reference architecture, or the number of	lvlog	• Scale-up: vglog • Scale-out: vg<SID>log<node number>

	assigned LUNs		
SHARED[1]	1 - following this reference	lvshared	• Scale-up: vgshared

Note: If you only use one LUN to create the logical volumes for data, log, or shared, the options stripes and stripesize are not needed.

Create the file system on top of the logical volume. Hitachi Vantara storage systems use the XFS file system. The following table lists the options to create and mount the file system.

System Type		Create Options	Mount Options	Mount Point
DATA	Scale-up	-F	noatime,inode64,nodiratime	/hana/data
	Scale-out		N/A	/hana/data/<SID>
LOG	Scale-up	-F	noatime,inode64,nodiratime	/hana/log
	Scale-Out		N/A	/hana/log/<SID>
SHARED	Scale-up	-F	noatime,inode64,nodiratime	/hana/shared
	Scale-out*	N/A	vers=3,proto=tcp,hard,nolock,timeo=600,retrans=2,wsize=65536,rsize=65536	/hana/shared/<SID>
* The listed options are applicable to Hitach VSP One File as NFS for /hana/shared. If you use GFS2 for /hana/shared, use mkfs.gfs2 to create the file system. The GFS2 file system will be managed by Linux Distributed Lock Manger (DLM).				

To create a file system, use the following command:

mkfs.xfs <create options> /dev/mapper/<vg id>-<lv id>

For example:

```
mkfs.xfs -F /dev/mapper/vglog-lvlog
```

SAP HANA persistent storage volume configuration for scale-up deployments

This is only for scale-up systems.

For scale-up systems, you need to persist the file systems, including the mount options, in the operating system's startup file, */etc/fstab*, to mount the file systems automatically during boot operations.

To configure the HANA persistent storage volume in a scale-up system, do the following.

Procedure

1. Add the following entry to */etc/fstab* for each file system:

   ```
   /dev/mapper/<vg id>-<lv id> <mount point> xfs <mount options> 0 0
   ```

   See the previous two tables for volume group and logical volume names as well as the mount options.

2. To create the mount points, use the following command:

   ```
   mkdir -p -m 755
   ```

   Example for scale-up:

   ```
   mkdir -p -m 755 /hana/{shared,log,data}
   ```

3. Mount the file systems.

   To mount the file systems one at a time, use the following command:

   ```
   mount
   ```

   Or, mount them all at once:

   ```
   mount -a
   ```

   You can check the mounted file system using either the *df* or *mount* command.

SAP HANA persistent storage volume configuration for scale-out deployments

This is only for scale-out systems.

To configure the SAP HANA persistent storage volume in a scale-out system, do the following.

Procedure

1. Create the mount points for SHARED, LOG, and DATA on each server. This example assumes that HIT is your SAP System ID.

   ```
   mkdir -p -m 755 /hana/{shared,log,data}/HIT
   ```

2. Configure the SAP HANA shared file system using VSP One File 30 or GFS2.

- VSP One File 30

   a. Persist the SAP HANA shared file systems in */etc/fstab* by adding the following entry:

   ```
   <ip address of nfs>:<share> <mount point> nfs 0 0
   ```

   b. Mount the file system:

   ```
   mount /hana/shared/HIT
   ```

- Global File System 2 (GFS2)

   Note: GFS2 is supported by SAP only if using Hitachi VSP storage models and Hitachi servers in TDI scale-out deployments for HANA shared filesystem.

   The logical volume manager (LVM) creates a single striped volume on which the GFS2 volume is created for the SAP HANA shared file system. The Linux Distributed Lock Manager manages shared file systems on a Linux computer cluster.

3. Verify that the file system is mounted using either the *df* or *mount* command.

SAP HANA takes care of mounting the file system used for log and data with the correct options during HANA startup. For more details, see the following:

- *SAP storage connector API Fibre Channel Client*
- *SAP HANA software installation*

SAP storage connector API Fibre Channel Client

For a scale-out configuration, SAP HANA offers a ready-to-use storage connector client for configurations with native multipaths of Fiber Channel-attached devices. This enables host auto-failover on block storage.

The Fibre Channel storage connector, *fcClient/fcClientLVM*, implements the SAP Storage Connector API that provides hooks for the following:

- Database startup
- Failing-over nodes

SAP supports this solution to enable the use of high-performance Fibre Channel devices in a scale-out installation.

The *fcClient/fcClientLVM* implementation uses standard Linux commands, such as *multipath* and *sg_persist*. Install and configure these commands.

The *fcClient/fcClientLVM* implementation is responsible for mounting the SAP HANA volumes. It also implements a proper fencing mechanism during a host failover by means of SCSI-3 persistent reservations for SAP HANA failover.

Configuration of the SAP Storage Connector API is contained within the SAP global.ini file in the */hana/shared/<SID>/global/hdb/custom/config*.

Operating System Configuration for NAS Storage

This section describes the SAP HANA volumes and operating system configuration.

SAP HANA persistent storage NFS volume configuration

For SAP HANA scale-up and scale-out systems, NFS share mount options must be persisted in one of the operating system's startup files, */etc/fstab*, to mount the file systems automatically during boot operations. Follow these step-by-step instructions to make sure your setup is prepared for SAP HANA installation. There might be differences between the commands listed here based on SLES for SAP Applications 15 and what is required for RedHat Enterprise Linux, but the general flow is the same and it will only require small adaptations.

Step-by-Step: Mounting NFS Volumes for SAP HANA on SLES 15

Install NFS client utilities:

SLES 15 includes NFS support from the *nfs-client* package.

```
zypper install nfs-client
```

This installs the required services and tools such as mount.nfs.

Create mount points:

SAP HANA uses specific directory structures. Create them as needed:

```
mkdir -p -m 755 /hana/{shared,log,data}
```

Enable and start required services:

Enable and start the NFS client locking service (nfs-client.target):

```
systemctl enable nfs-client.target
systemctl start nfs-client.target
```

Also ensure the system is using network-online.target to wait for network availability before mounting:

```
systemctl enable wicked
systemctl start wicked
```

Verify NFS server access:

Use *showmount* to check available NFS exports (/hana/shared for example):

```
showmount -e 192.168.150.26
```

Repeat for all required IP addresses.

Ensure the firewall allows NFS traffic *(rpcbind, mountd, nfs*).

Make mounts persistent in */etc/fstab*:

Edit the file:

```
vi /etc/fstab
```

Add entries either for scale-up or scale-out instances:

```
# SAP HANA NFS Mounts for HANA Scale-up
192.168.150.26:/hana_shared_HIT /hana/shared nfs rw,vers=3,rsize=65536, wsize=65
536,hard,nolock,timeo=600,retrans=2
192.168.100.27:/hana_log_HIT /hana/log nfs rw,noatime,vers=3,rsize=65536, wsize=
65536,hard,nolock,nconnect=8,timeo=600,retrans=2
192.168.100.27:/hana_data_HIT /hana/data nfs rw,noatime,vers=3,rsize=65536, wsiz
e=65536,hard,nolock,nconnect=8,timeo=600,retrans=2

# SAP HANA NFS Mounts for HANA Scale-out
192.168.150.26:/hana_shared_HIT /hana/shared nfs rw,vers=3,rsize=65536, wsize=65
536,hard,nolock,timeo=600,retrans=2
192.168.100.27:/hana_log_HIT /hana/log nfs rw,noatime,vers=4.1,lock, rsize=65536,
wsize=65536,hard,nconnect=8,timeo=600,retrans=2
192.168.100.27:/hana_data_HIT /hana/data nfs rw,noatime,vers=4.1,lock, rsize=655
36,wsize=65536,hard,nconnect=8,timeo=600,retrans=2
```

Mount the volumes:

```
mount -a
```

Set the correct permissions and owner:

Ensure correct ownership by the SAP users (*hitadm, sapsys*):

```
chown -R hitadm:sapsys /hana/data
chown -R hitadm:sapsys /hana/log
```

These users are usually created during SAP installation.

Optional: Use *systemd* mount units (advanced/production use):

For fine-grained boot-time control, use *systemd.mount* files. Systemd can handle dependency ordering (the network must be online before mounts).

Tune kernel and NFS parameters:

The following is a combination of recommended tunings for network, NFS, and kernel settings from SAP and Hitachi Vantara. Add the following to */etc/sysctl.conf*:

```
net.core.rmem_max = 262144
net.core.wmem_max = 262144
net.ipv4.tcp_rmem = 4096 87380 16777216
net.ipv4.tcp_wmem = 4096 65536 16777216
net.ipv4.tcp_rmem = 4096 87380 16777216
net.ipv4.tcp_wmem = 4096 65536 16777216
net.core.netdev_max_backlog = 50000
```

Apply the settings by running:

```
sysctl -p
```

Special requirements for SAP HANA scale-out deployments

For SAP HANA scale-out systems using stand-by nodes for host auto-failover, you must use NFSv4 or NFSv4.1 and NFS file locking. Host auto-failover for local high-availability is described in *SAP HANA – Host Auto Failover*.

With the scale-out specific mount options for the SAP HANA volumes provided previously, SAP HANA host auto-failover will work as expected. To speed up takeover time, it is possible to create a STONITH implementation which depends on the server model of the SAP HANA nodes. This is, however, not part of this document.

SAP HANA software installation

After configuring the file system for the SAP HANA data volume, log volume, and HANA shared volume, install the current version of SAP HANA supported by SAP on the SAP HANA server or servers. Refer to SAP Note *2235581 - SAP HANA: Supported Operating Systems*. The *SAP_HANA_OS_Release_Support_Matrix.pdf* file attachment outlines which combinations of OS RHEL or SLES are supported with the various SAP HANA versions.

By default, the SAP HANA database, as well as the SAP HANA client, need to be installed.

Install SAP HANA software on a scale-up environment

Follow this procedure to install SAP HANA on a scale-up environment. This procedure assumes the following:

- *HIT* is your SAP System ID
- *10* is your SAP System Number

- *saphanas.company.corp* is your fully qualified hostname

To install SAP HANA software on a scale-up environment, do the following.

Procedure

- Download a supported version of SAP HANA and follow the instructions to extract the archive.

- Install HANA by typing the following command:

```
--components=server,client \
--install_hostagent \
--number 10 \
--sapmnt=/hana/shared \
--sid=HIT \
--system_user_password=<password> -p <password> \
--saoadn_oasswird=<password> \
--datapath=/hana/data/HIT \
--logpath=/hana/log/HIT \
--hostname=saphanas.company.corp \
--certificates_hostmap=saphanas.company.corp=saphanas.company.corp
```

- During the installation you need to provide various passwords.

- Once the installation is complete, continue with *Configure the SAP HANA software*.

Install SAP HANA software on a scale-out environment

Before you can start with the installation on a scale-out environment, prepare a configuration file used by the SAP Storage Connector API Fibre Channel.

This configuration file describes the communication, persistence, and storage details for the SAP HANA installation routine with similar content to the following (assuming *HIT* is your SAP System ID).

Create this configuration file for 2+1 scale-out before following the procedure:

```
[communication]
listeninterface = .global
tcp_backlog = 2048
sslminprotocolversion = tls12

[multidb]
mode = multidb
database_isolation = low
singletenant = yes

[persistence]
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT

[storage]
```

```
ha_provider = hdb_ha.fcClientLVM
partition_*_* prtype = 5
partition_1_log lvmname = vgHITlog001-lvlog
partition_1_data lvmname = vgHITdata001-lvdata
partition_2_log lvmname = vgHITlog002-lvlog
partition_2_data lvmname = vgHITdata002-lvdata

[trace]
Ha_fcclient = info
```

See *SAP HANA Fiber Channel Storage Connector Admin Guide* for more details and configuration options used by this configuration file.

Use this configuration file during the installation. You can remove it afterwards. Save this file as */tmp/hana_install_HIT/global.ini*. Keep the file name (*global.ini*) the same, even if you choose a different directory.

Follow this configuration to complete your installation. This procedure assumes the following:

- *HIT* is your SAP System ID
- *10* is your SAP System Number
- *hana001.company.corp* is your fully qualified hostname
- *hana001ic* is the interface listening on the HANA node interconnect network
- */tmp/hana_install_HIT/global.ini* is your configuration file

To install SAP HANA software on a scale-out environment, do the following.

- Download a supported version of SAP HANA and follow the instructions to extract the archive.
- Install SAP HANA by running the following command:

```
--components=server,client \
--install_hostagent \
--number 10 \
--sapmnt=/hana/shared \
--sid=HIT \
--system_user_password=<password> -p <password> \
--sapadm_password=<password> \
--storage_cfg=/tmp/hana_install_<number> \
--hostname=hana001ic \
--certificates_hostmap=hana001ic=hana001
```

- During the installation you need to provide various passwords.

- After the installation is complete, continue with *Configure the SAP HANA software*.

To add a worker node to the existing installation, use the following command:

```
/hana/shared/HIT/hdblcm/hdblcm \
--action=add_hosts \
--addhosts=hana002ic:role=worker:group=default:storage_partition=2 \
--sid=HIT \
--install_hostagent \
--sapmnt=/hana/shared \
--sapadm_password=<password> -p <password>
```

Note: Make sure that your master installation is aware of the node you are adding. That means your configuration needs to include the partition information. The master node owns partition ID 1, the first additional node can be added as node 2, and so forth.

To add a standby node, you do not need to provide the partition information, because a standby node gets the needed information during a HANA failover. Use the following command:

```
/hana/shared/HIT/hdblcm/hdblcm \
--action=add_hosts \
--addhosts=hana003ic:role=standby:group=default \
--sid=HIT \
--install_hostagent \
--sapmnt=/hana/shared \
--sapadm_password=<password> -p <password>
```

Configure the SAP HANA software

The following is an example for a global.ini file used during Enterprise Storage validation:

SAP HANA 2.0 global.ini file:

```
[communication]
tcp_backlog = 2048
listeninterface = .global

[fileio]
max_parallel_io_requests[DATA] = 64
max_parallel_io_requests[LOG] = 64
max_submit_batch_size[DATA] = 64
max_submit_batch_size[LOG] = 64
size_kernel_io_queue[DATA] = 512
size_kernel_io_queue[LOG] = 512
async_read_submit[DATA] = on
async_read_submit[LOG] = on
async_write_submit_blocks[DATA] = all
async_write_submit_blocks[LOG] = all
min_submit_batch_size[DATA] = 16
min_submit_batch_size[LOG] = 16
async_write_submit_active[DATA] = on
async_write_submit_active[LOG] = on

[multidb]
mode = multidb
database_isolation = low
```

```
singletenant = yes

[persistence]
basepath_datavolumes = /hana/data/HIT
basepath_logvolumes = /hana/log/HIT

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_* prtype = 5
partition_1_data lvmname = vgHITdata001-lvdata
partition_1_log lvmname = vgHITlog001-lvlog
partition_2_data lvmname = vgHITdata002-lvdata
partition_2_log lvmname = vgHITlog002-lvlog
partition_3_data lvmname = vgHITdata003-lvdata
partition_3_log lvmname = vgHITlog003-lvlog

[trace]
ha_fcclientlvm = info
```

Best practices for the HANA *fileio* parameters vary between different storage setups and HANA deployments. The following table gives an overview of the general TDI best practices for the storage models described in this document.

SAP HANA *fileio* Parameter	VSP One BHE	VSP One B20	VSP One File 30	VSP E1090	VSP5000
max_parallel_io_requests[DATA]	64	64	1024	1024	1024
max_parallel_io_requests[LOG]	64	64	128	128	128
max_submit_batch_size[DATA]	64	64	1024	1024	1024
max_submit_batch_size[LOG]	64	64	128	128	128
size_kernel_io_queue[DATA]	512	512	1024	1024	1024
size_kernel_io_queue[LOG]	512	512	1024	1024	1024
async_read_submit[DATA/LOG]	on	on	on	on	on
async_write_submit_blocks[DATA/LOG]	all	all	all	all	all
min_submit_batch_size[DATA/LOG]	16	16	16	16	16
async_write_submit_active[DATA/LOG]	on	on	on	on	on

Further details can be found in the reference architecture guides of the specific storage models and deployment solutions.

Provisioning Automation

Manual provisioning of storage resources is increasingly impractical in modern enterprise environments. Traditional provisioning processes often involve multiple manual steps

across different infrastructure layers, increasing the risk of configuration errors and extending deployment timelines.

Ansible playbooks for storage provisioning

Automation frameworks such as Red Hat Ansible provide a standardized method for orchestrating infrastructure operations through declarative workflows. In combination with Hitachi storage APIs and automation modules, infrastructure teams can implement Infrastructure as Code (IaC) approaches that transform storage provisioning into a repeatable and version-controlled process.

Hitachi Vantara provides certified Ansible modules for its storage platforms, enabling administrators to automate configuration and lifecycle management tasks such as volume provisioning, host group configuration, snapshot management, and replication setup.

These integrations allow organizations to incorporate storage provisioning into broader DevOps and CI/CD pipelines while maintaining strict adherence to enterprise storage best practices.

Automation Framework for HANA TDI Storage Provisioning

Hitachi Vantara provides an official Ansible module collection for VSP One Block storage systems, enabling infrastructure teams to automate storage management tasks directly from Ansible playbooks.

This *article* enables centralized orchestration of storage operations through Ansible playbooks while maintaining direct control over storage resources via Hitachi's REST APIs.

These modules provide programmatic access to storage functionality such as:

- Storage system discovery and configuration
- Storage port configuration
- Host group creation and mapping
- Storage pool and volume creation
- LUN mapping to SAP HANA hosts

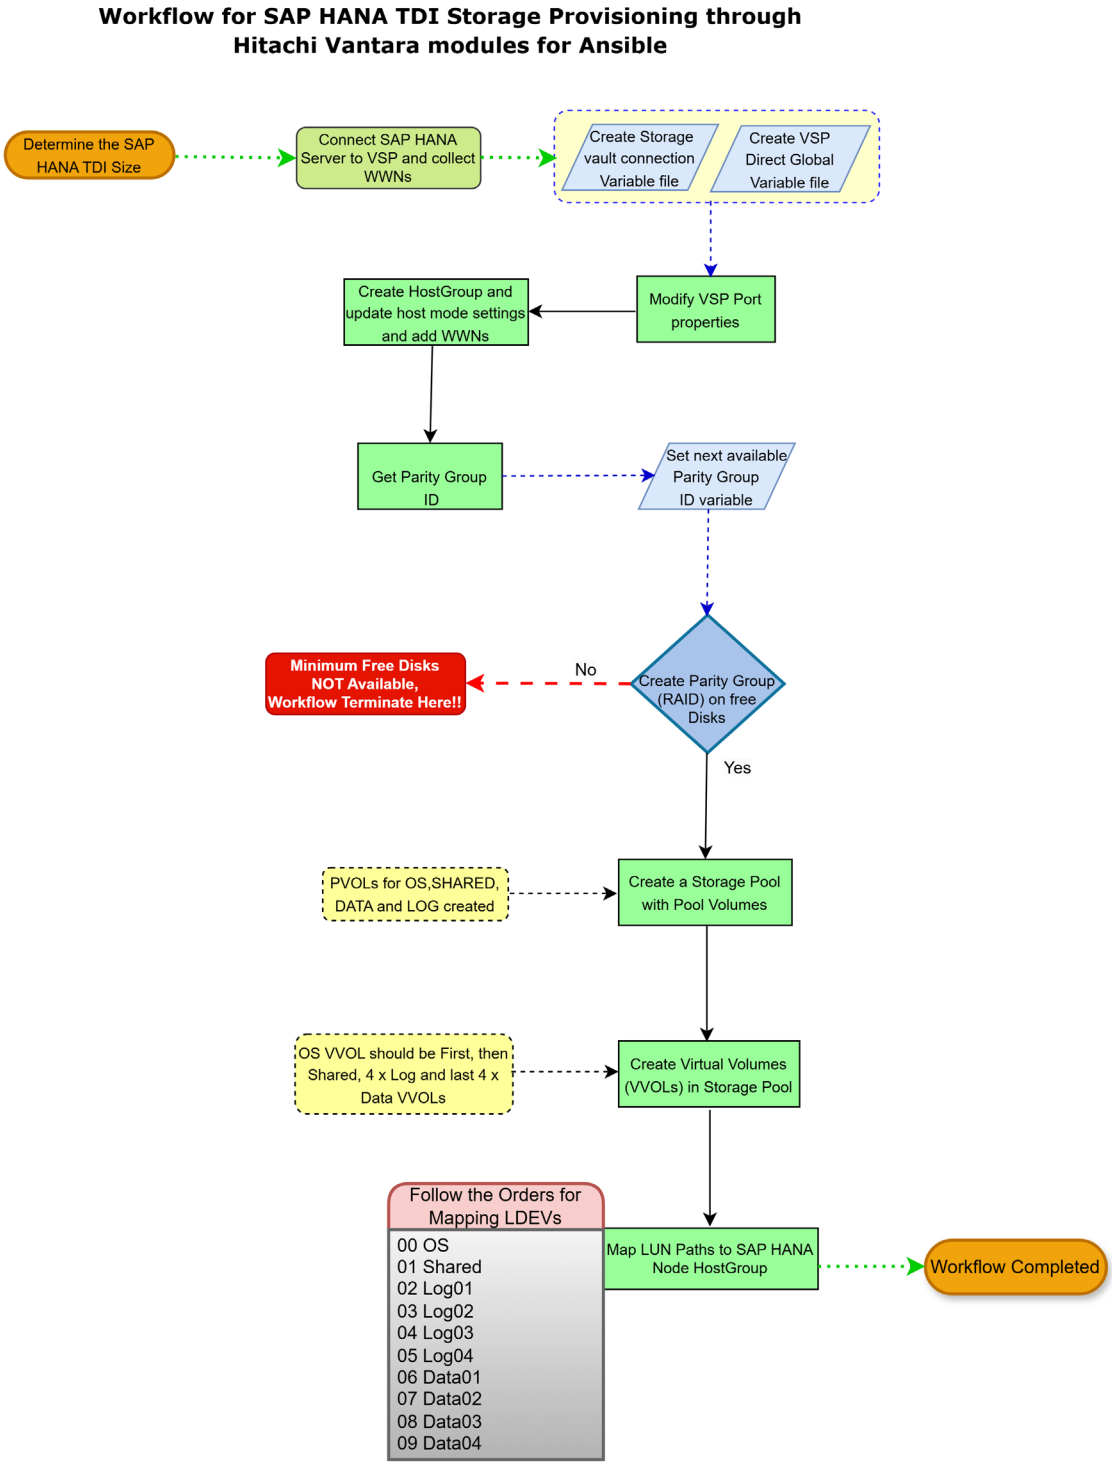

Automation pipelines can be constructed using structured Ansible workflows that execute these tasks sequentially, ensuring consistent provisioning across environments.

This approach ensures that storage configurations align with both SAP HANA best practices and Hitachi storage design guidelines.

Automation workflows are typically executed from a centralized Ansible control node, allowing organizations to manage infrastructure provisioning for multiple SAP systems from a single automation framework.

Operational Benefits

Adopting automation for SAP HANA infrastructure provisioning delivers significant operational advantages. These benefits extend beyond faster deployment and contribute directly to improved reliability and governance within enterprise SAP environments.

- **Accelerated Infrastructure Deployment**: Automated provisioning workflows allow infrastructure teams to deploy storage resources significantly faster than traditional manual processes. Tasks that previously required manual configuration across multiple management interfaces can be executed automatically in minutes.

- **Consistent and Repeatable Infrastructure**: Automation ensures that each deployment follows the same configuration process. This eliminates variability introduced by manual provisioning and ensures consistent adherence to SAP and vendor best practices.

- **Reduced Operational Risk**: Manual configuration processes are prone to human error, particularly in complex storage environments. Automation reduces this risk by implementing standardized workflows that enforce validated configuration parameters.

- **Improved Governance and Auditability**: Automation workflows can be stored in version-controlled repositories and integrated into CI/CD pipelines. This allows organizations to track infrastructure changes and maintain a clear operational audit trail.

- **Scalability Across SAP Landscapes**: Automation enables organizations to scale infrastructure operations across multiple SAP environments, including development, testing, and production landscapes, without increasing operational complexity.

Best Practices for Implementation

Based on the architecture and automation approach described in the reference article, several best practices can be derived for organizations deploying SAP HANA with Hitachi infrastructure.

- **Standardize Infrastructure Automation**: Automation frameworks such as Ansible should be used as the primary method for provisioning storage resources. This ensures consistency across environments and enables integration with modern DevOps workflows.

- **Implement Infrastructure as Code**: Infrastructure configurations should be defined as code within version-controlled repositories. This allows organizations to maintain reproducible infrastructure templates for SAP environments.

- **Use Vendor-Certified Automation Modules**: Hitachi-provided automation modules should be used whenever possible to ensure compatibility with the underlying storage platform and alignment with vendor best practices.

- **Align Automation with SAP Performance Requirements**: Automation workflows should incorporate SAP HANA storage layout guidelines and performance requirements to ensure that automatically provisioned environments meet SAP certification standards.

Ansible Playbooks

A comprehensive automation pipeline available on official *Hitachi Vantara's GitHub* and defined in the `main.yml` playbook has been developed to streamline the initial creation of SAP HANA scale-up volumes on Hitachi VSP One Block storage. The pipeline includes two dedicated workflows, each thoroughly tested to ensure reliability across all volume creation tasks.

The available pipelines efficiently orchestrates the following tasks:

- **Configure Storage Port:** Sets up the necessary storage ports on the Hitachi Vantara VSP Block Storage system using *1_storage_port_create.yml*

- **Configure Hostgroups:** Establishes host groups and adds WWNs of hosts using *2_hostgroup_create.yml*

- **Parity Group Creation:** Creates parity groups with the required RAID type for data redundancy with *4_paritygroup_create.yml,* after retrieving the parity group ids with *3_paritygroup_get_pg_id.yml*

- **Single or Two Storage Pool with Pool volumes Creation:** Create a storage pool with approx. 4 or 8 Pool Volumes (PVOLs) with *5_storagepool_create.yml*

- **Virtual Volumes (VVOLs) Creation and Mapping:** Provisions for VVOLs are made for operating system and SAP HANA components, including shared, log, and data volumes, utilizing the *6_vvol_creation_in_pool.yml* configuration file. These volumes are then mapped to the host groups that contain the host WWNs. The creation process adheres to Hitachi Vantara's best practices for SAP HANA, ensuring optimal performance and reliability.

- **Variables:** To manage configuration efficiently.

Provisioning using VSP360

In traditional environments, storage configuration for SAP HANA often requires multiple tools and manual steps across storage management interfaces. These manual workflows increase the risk of configuration inconsistencies and operational errors, particularly in scale-out SAP HANA environments where storage must be provisioned for multiple nodes.

To address these challenges, Hitachi Vantara VSP 360 provides a unified management and automation platform designed to simplify infrastructure operations across the Hitachi VSP One data platform portfolio. By introducing guided workflows and centralized management capabilities, VSP 360 enables faster, more consistent provisioning of SAP HANA storage resources while reducing operational complexity.

This section summarizes the architecture, operational workflow, and recommended practices for provisioning SAP HANA storage using Hitachi VSP 360 together with Hitachi VSP One Block storage systems provided in this *article*.

VSP 360 Platform Overview

Hitachi Virtual Storage Platform 360 (VSP 360) is a unified data infrastructure management platform designed to provide centralized administration, analytics, and automation across the VSP One storage portfolio.

Key capabilities include:

- Unified management across block, file, object, and software-defined storage
- Integrated observability and AIOps-driven insights
- Simplified provisioning workflows
- Reduced operational complexity for storage administrators
- Centralized management across heterogeneous storage environments

VSP 360 can be deployed as a virtual appliance on supported hypervisors including VMware ESXi, Microsoft Hyper-V, or KVM.

Further documentation is available through the following resources:

- Hitachi Vantara Documentation Portal
- VSP 360 Product Compatibility Guide

These resources provide detailed instructions for installation, configuration, and lifecycle management of the VSP 360 platform.

Advantages of VSP 360 for SAP HANA Deployments

Using VSP 360 for SAP HANA infrastructure management provides several operational advantages compared with traditional storage configuration tools.

Key benefits include:

- **Unified Storage Management**: VSP 360 provides a single management interface for the entire VSP One portfolio, eliminating the need to switch between multiple storage management utilities.

- **Reduced Operational Complexity**: Guided workflows simplify provisioning tasks, allowing infrastructure teams with limited storage specialization to deploy enterprise-grade configurations.

- **Integrated Observability**: Built-in analytics and monitoring provide predictive insights into infrastructure health, enabling proactive management of performance and capacity.

- **Consistency and Automation**: Automated provisioning workflows reduce manual configuration steps and ensure that storage deployments follow validated best practices.

Together, these capabilities significantly streamline infrastructure deployment and reduce the risk of operational errors in large SAP landscapes.

Storage Provisioning Workflow Using VSP 360

The storage provisioning process within VSP 360 is implemented through a **guided workflow** that automates many backend configuration steps.

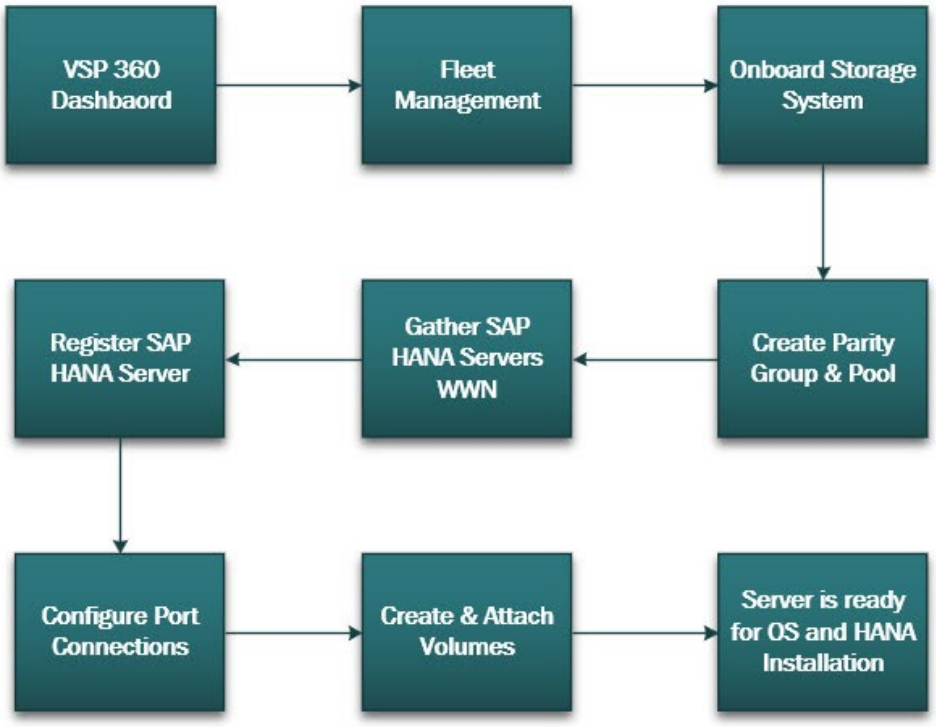

The workflow consists of the following major stages:

- Storage system onboarding
- Storage pool creation
- Volume provisioning
- Volume-to-host mapping
- Host verification and SAP HANA deployment

By combining Hitachi VSP 360 with Hitachi VSP storage, organizations can significantly simplify SAP HANA storage provisioning through centralized management and guided automation workflows.

The integrated solution delivers several operational advantages:

- Faster infrastructure deployment
- Reduced configuration errors
- Consistent implementation of best practices
- Simplified lifecycle management for enterprise storage environments

Together, VSP 360 and VSP One Block provide a robust and scalable foundation for SAP HANA TDI environments, enabling organizations to deploy and operate SAP workloads with greater efficiency, reliability, and operational visibility.

Product Descriptions

This best practice paper includes usage of the following products.

VSP One Block High End

Hitachi Virtual Storage Platform One Block High End (VSP One Block High End) all-flash NVMe block storage systems deliver ultra-fast, highly reliable, and scalable data access for the most demanding enterprise applications.

Powered by Hitachi-engineered technology, they offer eight nines of availability, next-generation connectivity, comprehensive easy-to-use management with VSP 360 and seamless workload consolidation across open systems and mainframes. With end-to-end NVMe and best-in-class data reduction, VSP One Block High End greatly enhances performance, scale and resilience in a simple, secure, sustainable way.

VSP One Block 20

Hitachi Virtual Storage Platform One Block simplifies data management while enabling enterprises to reliability and sustainability meet initiatives to fuel growth and innovation. Dynamic Carbon Reduction optimizes energy usage by switching CPUs to ECO mode during low activity. Adaptive Data Reduction (ADR) is always on, enhancing efficiency and reducing the overall CO2 footprint.

Thin Image Advanced (TIA) integrates with major snapshot ecosystems, supporting business continuity and defending against threats. CyberArk Privileged Access Manager plugins enhance block storage system security by prioritizing data confidentiality, ensuring compliance, and actively defending against security threats.

Hitachi Virtual Storage Platform One Block includes the following dedicated models:

- VSP One Block 24 – 256 GB Cache + SW Advanced Data Reduction (ADR) + 24 cores
- VSP One Block 26 – 768 GB Cache + 2 × Compression Accelerator Module (CAM) + 24 cores
- VSP One Block 28 – 1 TB Cache + 4 × CAM + 64 cores

All models support Fibre Channel, iSCSI, and NVMe TCP connectivity. The new capabilities reduce complexity: data reduction is always on, Dynamic Drive Protection removes complicated RAID setup, and Dynamic Carbon Reduction delivers real-world reduction in power consumption. In addition, the models are FIPS compliant.

In short, Hitachi Virtual Storage Platform One Block combines simplicity, sustainability, and robust security features to optimize system management, energy efficiency, and data protection.

See *https://www.hitachivantara.com/en-us/products/storage-platforms/block-storage/midrange/vsp-one-block* for more information

VSP One File 30

Hitachi Virtual Storage Platform One File (VSP One File) is a modernized Hitachi NAS that is focused on improving the customer experience, delivering on operational simplification, and enabling better agility.

The VSP One File series provides an easy and modern way to manage, monitor, and configure file systems and data across the organization.

Hitachi Virtual Storage Platform One File includes the following dedicated models:

- **VSP One File 34** – targeted for technology buyers that want the best price/performance, an all-flash deployment, and have standardized on 25 GbE network connectivity or are planning to move short/medium term to 25 GbE.
- **VSP One File 38** – built for leading technology pioneers that require the most file performance and highest rack network port density to address all-flash use cases and applications.

For high SLA workloads, all VSP One File models support SMB3 multichannel, NFS 4.1 multi-pathing, nconnect, and Link Aggregation Control Protocol (LACP) to improve availability and/or higher bandwidth for Windows and Linux clients beyond the limit of a single connection.

Enjoy consistent and secure file performance across your data infrastructure no matter where your data resides. This near-cloud deployment offers an ultra-low latency, secure connection from enterprise to public cloud.

See *https://www.hitachivantara.com/en-us/partners/vsp-one* for more information.

VSP 5000

This enterprise-class, flash array evolution storage, Hitachi Virtual Storage Platform 5000 series has an innovative, scale-out design optimized for NVMe and storage class memory. It achieves the following:

- **Agility using NVMe**: Speed, massive scaling with no performance slowdowns, intelligent tiering, and efficiency.
- **Resilience**: Superior application availability and flash resilience. Your data is always available, mitigating business risk.
- **Storage simplified**: Do more with less, integrate AI (artificial intelligence) and ML (machine learning), simplify management, and save money and time with consolidation.

See *https://www.hitachivantara.com/en-us/products/storage-platforms/block-storage/enterprise* for more information.

VSP E1090

The Hitachi Virtual Storage Platform E1090 storage system is a high-performance, large capacity data storage system. The VSP E1090 all-flash arrays (AFAs) support NVMe and SAS solid-state drives (SSDs). VSP E1090 hybrid models can be configured with both SSDs and hard disk drives (HDDs).

- The NVMe flash architecture delivers consistent, low-microsecond latency, which reduces the transaction costs of latency-critical applications and delivers predictable performance to optimize storage resources.
- The hybrid architecture allows for greater scalability and provides data-in-place migration support.

Hitachi Storage Virtualization Operating System RF

Hitachi Storage Virtualization Operating System RF powers the Hitachi Virtual Storage Platform (VSP) family. 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.

Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid arrays as storage capacity, consolidating resources for a higher ROI and providing a high-speed front end to slower, less-predictable arrays.

See *https://www.hitachivantara.com/en-us/products/storage-platforms/storage-software* for more information.

SAP HANA Tailored Data Center Integration

SAP increases flexibility and provides an alternative to SAP HANA Appliances with SAP HANA tailored data center integration (TDI) in currently 5 phases. This includes many kinds of virtualization technology. Understanding the possibilities and requirements of an SAP HANA TDI environment is crucial. One of these requirements is that all virtualized environments for SAP HANA are considered as SAP HANA TDI by SAP. SAP provides documentation around SAP HANA TDI environments that explain the 5 phases of SAP HANA TDI as well as hardware and software requirements for the whole stack:

- SAP Help Portal: SAP HANA Platform (frequently asked questions)
- SAP HANA Storage Requirements
- SAP HANA Network Requirements

Taking all this into account, Hitachi offers a solution for the whole SAP HANA TDI stack from hardware infrastructure to software deployment.

VSP 360

VSP 360 is unified data management platform from Hitachi that provides:

- centralized device management
- storage management
- data protection workflows
- AI-powered analytics
- orchestration/automation workflows

Summary

Hitachi Vantara's Virtual Storage Platform (VSP) enterprise portfolio provides a validated and resilient foundation for SAP HANA TDI environments, successfully addressing the architectural shift in responsibility from appliance vendors to infrastructure architects. By utilizing the Hitachi Storage Virtualization Operating System (SVOS) RF and its patented **flash-aware I/O stack**, these solutions deliver the predictable sub-millisecond latency and high-bandwidth throughput essential for meeting mission-critical SAP Key Performance Indicators (KPIs).

The current portfolio offers unmatched versatility and scalability to meet any enterprise requirement:

- **Extreme Scalability:** From the high-performance mid-range VSP E1090 to the VSP One Block High End, which can scale to support up to 1,008 active SAP HANA systems via Virtual Storage Scale-out (VSSO).
- **Operational Simplicity:** The integration of VSP 360 and Ansible-based automation transforms complex storage tasks into repeatable, auditable workflows, significantly reducing deployment timelines and the risk of manual configuration errors.
- **Architectural Flexibility:** Organizations can choose between specialized media types like TLC for performance or QLC for capacity and deploy either block-based SAN or high-performance NAS-based infrastructure via VSP One File 30.
- **Uncompromising Resiliency:** With a No Single Point of Failure (NSPOF) design and "eight nines" availability, Hitachi ensures that the storage layer remains the most reliable component of the SAP stack, supporting non-stop maintenance and robust data protection.

In summary, by following Hitachi's validated engineering standards and implementation best practices, IT teams can deploy a consolidated, high-density SAP landscape that optimizes Total Cost of Ownership (TCO) while ensuring maximum operational efficiency and future-proofed scalability.

Appendix

Below is the recommended list of assets and their associated paper numbers derived from our technical sources:

Consolidated Reference Architecture Portfolio

Storage Solution	Primary Reference Architecture Document Title	Paper Number (PN)
VSP One Block High End	SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform One Block High	MK-SL-423

Storage Solution	Primary Reference Architecture Document Title	Paper Number (PN)
	End with SVOS RF	
VSP One Block 20	SAP HANA Tailored Data Center Integration on Hitachi VSP One Block 20 with SVOS RF	MK-SL-300
VSP One File 30 (NAS)	SAP HANA Tailored Data Center Integration on Hitachi VSP E1090 with Hitachi VSP One File 30	MK-SL-390
VSP E1090 (Block)	SAP HANA Tailored Data Center Integration on Hitachi VSP E990 and E1090 with SVOS RF	MK-SL-197
VSP 5000 Series	SAP HANA Tailored Data Center Integration on Hitachi VSP 5000 Series with SVOS RF	MK-SL-167

Additional Strategic Assets

In addition to the core architecture guides, it is best practice to include these strategic technical publications for comprehensive lifecycle management:

- **Ansible Automation for TDI Storage:** *Automating SAP HANA TDI Storage Provisioning with Ansible and Hitachi Vantara VSP One Block Storage*
- **Automating SAP HANA Full-Stack Installation by Hitachi Vantara**: *https://community.hitachivantara.com/blogs/chandan-kumar/2026/03/30/automating-sap-hana-full-stack-installation*
- **Unified Management for VSP One Block:** *VSP360: Simplifying SAP HANA TDI Storage Provisioning on VSP One Block 20*
- **Media Selection Strategy for VSP One B20:** *QLC vs TLC Flash for SAP HANA: When to Choose Capacity Over Performance*

Hitachi Vantara

Corporate Headquarters 2535 Augustine Drive

Santa Clara, CA 95054
USA

HitachiVantara.com/contact