TECHNICAL PAPER: November 2025

SAP HANA on Hyperconverged Infrastructure (HCI) Solutions Powered by VMware vSphere and vSAN 8

By Erik Rieger

Principal SAP Global Technical Alliance Manager & Architect

Table of contents

Introduction

Abstract

This document provides best practices and architectural guidelines for designing and building an SAP HANA environment on **Hyperconverged Infrastructure (HCI)** solutions using **Intel Xeon Scalable** processor-based systems running **VMware vSphere** and **VMware vSAN (OSA and ESA)**.

It serves as a reference for hardware vendors developing their own SAP HANA HCI solutions based on vSphere and vSAN. Customers can also use this guide to evaluate whether a vSAN-based SAP HANA HCI solution meets their requirements for running SAP HANA without additional external storage or dedicated storage network infrastructure, helping to **lower TCO of an SAP HANA environment**.

Vendors are responsible for selecting and configuring vSAN-ready hardware components, VMware vSphere, and vSAN-specific settings for their unique HCI solutions. Additionally, HCI solution vendors must certify their solutions under the **SAP HANA Hardware Certification Hyperconverged Infrastructure** certification scenario (**HANA-HyCI 1.0 v1.2**). For details and requirements, contact the [SAP Integration and Certification Center (ICC)].

This paper also outlines the support processes and requirements for running SAP HANA on vSAN, including vendor-specific solutions.

This document focuses on **vSphere and vSAN version 8.0** as part of VMware Cloud Foundation (VCF 5.x) and **Intel Xeon Sapphire Rapids and Emerald Rapids CPUs**. It provides guidance on leveraging **vSAN Express Storage Architecture (vSAN ESA)**, VMware's latest **vSAN storage architecture.** For guidance on older vSAN versions refer to previous revisions of this document.

Executive summary

For more than 13 years, VMware has successfully virtualized SAP applications on Intel x86 server hardware. In 2019, VMware received SAP support to virtualize the SAP HANA storage layer with vSAN, further advancing its capabilities. With VMware Cloud Foundation (VCF), VMware takes the next step in data center virtualization, providing the foundation for enterprises to build their own private cloud. VCF offers a simplified, standardized, and cost-efficient approach to private cloud infrastructure, ensuring full control, ownership, and sovereignty over critical assets like their data and unique business processes.

SAP HANA-based applications demand a high-performance, reliable, and dynamic compute environment, supported by predictable and resilient storage and networking subsystems. These systems must deliver optimal workload throughput, security, and adaptability to evolving business needs.

Reaching this goal is not an easy task. But VMware Cloud Foundation (VCF) includes VMware vSAN, a software-defined storage solution, as a key component. This allows enterprises to build and manage highly flexible, application-driven storage configurations that fully meet SAP HANA storage requirements.

From a storage perspective, this translates to delivering the right balance of data capacity, IOPS performance, and protection levels for SAP HANA VMs at an optimized cost. Traditional storage solutions, such as large Fibre Channel SANs or standalone SAP HANA appliances with internal storage, struggle to keep pace with today's fast-changing, resource-intensive environments. Their static configurations and reliance on dedicated storage network infrastructure make them less agile and more expensive than Ethernet-based storage solutions like VMware vSAN.

For detailed SAP support information on VMware vSAN, refer to: SAP Note 2718982 - SAP HANA on VMware vSphere and vSAN 7 and SAP note 3406060 - SAP HANA on VMware vSphere 8 and vSAN 8 A valid SAP S-user ID is required to view these notes.

Audience

This paper is intended for SAP and VMware hardware partners who want to build and offer SAP HANA certified HCI solutions based on VMware technologies and customers who want to evaluate whether a vSAN-based SAP HANA HCI solution meets their requirements for running an SAP HANA environment on a completely software-defined datacenter on premise or on VMware cloud provider clouds.

Solution overview

SAP HANA on VMware Cloud Foundation (VCF)

SAP HANA is an in-memory data platform that converges database and application platforms, eliminates data redundancy, and improves performance with lower latency. SAP HANA architecture enables converged, online transaction processing (OLTP) and online analytical processing (OLAP) within a single in-memory, column-based data store that adheres to the Atomicity, Consistency, Isolation, and Durability (ACID) compliance model.

SAP environments with SAP HANA can use VMware Cloud Foundation (VCF) with VMware vSAN included running on SAP and VMware certified x86 servers to eliminate traditional IT silos of compute, storage, and networking with a hyperconverged solution.

All intelligence and management are moved into a single software stack, allowing for SAP HANA VM- and SAP application-centric policy-based control and automation.

This brings better security, predictable performance, operational simplicity, and cost-effectiveness into SAP HANA environments. vSAN also simplifies the complexity of traditional IT silos for compute, storage, and networking, as it allows managing the SAP HANA storage just like compute resources.

HCI lowers costs by combining the economics and simplicity of local storage with the features and benefits of shared storage and can be used for both SAP HANA Scale-Up and Scale-Out deployments.

The following figure summarizes the benefits of VMware-based SAP HCI solutions.

Table 1. Benefits of VMware-based SAP HCI solutions

Radically simple storage and SAP HANA infrastructure	Elastic scalability, high performance	Lower TCO
	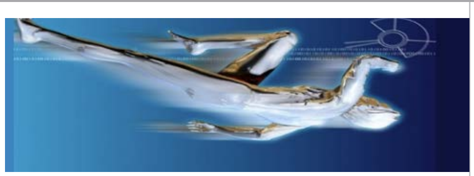	
• Two clicks to enable • Single pane of glass management	• Built on high-performance standard x86 servers	• Server-side economics

• Policy driven • Self-tuning • Integrated with VMware stack • Same one-stop SAP/VMware support process as with vSphere/VCF • Ready for any SAP workloads • Fully supported by SAP HANA	• Embedded in vSphere kernel with direct I/O path • Fast and flexible SAP deployments • Up to 4 SAP HANA VMs per vSAN-based SAP HCI cluster node • Scale from 2 to 64 nodes • Granular and linear scaling for SAP applications and SAP HANA	• No large up-front investments—grow as you go • Easy to operate with powerful automation • No specialized skillset required • Reliable and standard x86 hardware • Built-in high availability • SAP HANA virtual infrastructure turnkey solution

Note: Only SAP and VMware HCI solution vendor offered solutions that get **SAP HANA HCI solution certified** (SAP HANA Hardware Certification Hyperconverged Infrastructure (HANA-HyCI) test scenario) are supported for SAP HANA workloads. SAP applications other than HANA are fully supported on a VMware vSAN cluster. No additional support notes or specific support processes are required from SAP when running other applications on vSAN, such as an SAP application server.

vSAN overview

vSAN is a software-defined storage solution that leverages local storage resources from multiple ESXi hosts to create a shared, centrally managed storage pool. This eliminates the need for dedicated storage hardware, simplifying management and reducing costs.

vSAN employs advanced data resilience techniques, such as space-efficient erasure coding, to ensure data durability and high availability. Additionally, it optimizes storage utilization through Storage Policy-Based Management (SPBM) and features like compression.

As a hypervisor-integrated software-based solution, vSAN can be leveraged for SAP HANA databases and applications without the need to reserve dedicated resources, unlike other software-based storage solutions that are not integrated into the hypervisor.

This tight integration allows for unified management of compute and storage resources through the familiar vSphere Client, simplifying operations and reducing the learning curve for administrators. Furthermore, vSAN utilizes core vSphere features like vMotion, DRS, and HA, ensuring seamless workload mobility, resource optimization, and high availability.

This streamlined approach significantly reduces administrative overhead and enables more efficient resource allocation within the VCF stack. Figure 1 provides an overview of the available storage solutions for VCF and highlights why vSAN is an excellent choice for a fully virtualized SAP HANA HCI environment.

Figure 1. VMware Cloud Foundation storage options

	vSAN	vVol 2.0	VMFS on FC	NFS
Provides a Full SDDC Solution	✓			
Storage Policy Based Management	✓	✓		
Automated Deployment and Scale	✓	✓		✓
Automated LCM, Patching and Upgrades	✓			
Stretched Clusters	✓			
Remote Clusters	✓	✓	✓	✓
Compute Only Clusters	✓			

vSAN's ability to rapidly adopt the latest technology is a significant advantage, driven by the fast-paced development of server hardware. Unlike traditional storage arrays, which have longer development cycles, vSAN is software-defined and can immediately leverage the latest server hardware as soon as it becomes available.

This enables vSAN to quickly integrate advancements in CPU, memory, and storage technologies, delivering enhanced performance and efficiency to your virtualized environment. With vSAN, you can stay at the cutting edge of technology without being locked into a specific hardware platform or waiting for lengthy hardware refresh cycles. Figure 2 illustrates this capability in a comprehensive manner.

Figure 2. vSAN technology adoption and progress from its release to today

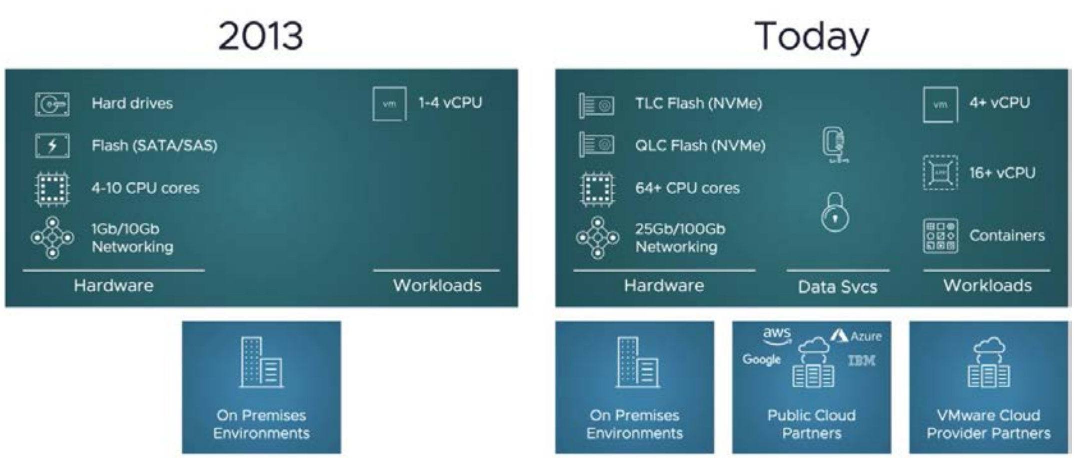

vSAN's scale-out architecture provides exceptional flexibility for SAP HANA deployments. By simply adding more hosts to the VCF/SAP HANA HCI environment, you can seamlessly expand both storage capacity and performance.

vSAN supports both scale-up and scale-out approaches to increasing storage:

- Scale-up: Adding disks to existing hosts
- Scale-out: Adding more hosts to the cluster

This flexibility enables SAP HANA HCI hardware vendors to provide granular scaling based on SAP workload demands, optimizing resource allocation and cost efficiency for customers.

For example, a vSAN cluster can start with as few as three hosts and scale up to 64 hosts, with each host contributing both storage capacity and performance to the SAP HANA HCI cluster. vSAN also supports dynamic resource allocation and workload mobility across VCF workload domains, ensuring efficient resource utilization and adaptability to changing business needs.

Additionally, vMotion enables seamless VM migration between hosts and even across workload domains without downtime, allowing for workload balancing and optimization.

vSAN is built for resilience. Its distributed architecture replicates data across multiple hosts, ensuring redundancy and high availability. Even in the event of hardware failure, applications stay online. Seamless integration with vSphere HA and DRS enhances resilience through automated recovery and optimized resource utilization, ensuring business continuity.

vSAN streamlines disaster recovery. Its integration with VMware Live Site Recovery enables automated failover and recovery, ensuring business continuity. vSAN stretched clusters replicate data across sites for high availability, while SAP HANA system replication, which is fully supported, offers a storage-independent alternative.

vSAN also ensures continuous data resilience. It automatically "self-heals" non-compliant data and uses Adaptive Resynchronization to balance VM performance and data recovery. By prioritizing VM I/O during high demand and accelerating resync during low demand, vSAN minimizes disruption and swiftly restores full protection, maintaining a highly resilient infrastructure.

vSAN's lifecycle management is tightly integrated with vSphere; it simplifies deployment, maintenance, and updates as summarized here:

- **Automated Deployment**: vSAN is enabled within a few clicks in the vSphere Client, with guided wizards minimizing manual setup and errors.
- **Effortless Maintenance**: vSphere Lifecycle Manager (vLCM) provides unified updates for vSAN, ESXi, and hardware firmware (if implemented by the HCI HW vendor), ensuring compatibility and stability.
- **Proactive Monitoring**: Health checks and alerts help administrators detect and resolve issues before they escalate, reducing downtime.
- **Centralized Firmware & BIOS Management**: vLCM can streamline hardware maintenance (if implemented from the HCI HW vendor), maintaining a validated configuration and preventing drift-related performance or security issues.

By automating updates and enforcing a consistent state, vLCM reduces complexity, enhances stability, and improves overall system resilience.

vSAN 8 offers a new storage architecture, vSAN Express Storage Architecture (vSAN ESA). This new storage architecture, which is an alternative to vSAN Original Storage Architecture (vSAN OSA), is now available for approved vSAN ReadyNodes™ and SAP HCI-certified systems based on Intel Xeon Scalable Generation 4 (Sapphire Rapids) and Generation 5 (Emerald Rapids) processors, see [SAP note 3406060] for details.

Why vSAN Express Storage Architecture is Superior to the Original

At a high level, the original vSAN architecture (vSAN OSA) was a two-tier setup, with dedicated cache and data disks designed to accommodate a wide range of storage devices. In contrast, the vSAN ESA in vSAN 8 is a single-tier design, eliminating the need for dedicated cache devices and optimized for high-performance NVMe-based TLC flash. The new log-structured file system and I/O path optimizations to the write path reduce write amplification and lower write latency. Figure 3 shows a high-level overview of the two storage architectures in vSAN 8.

Figure 3. High-level overview of the two storage architectures in vSAN 8

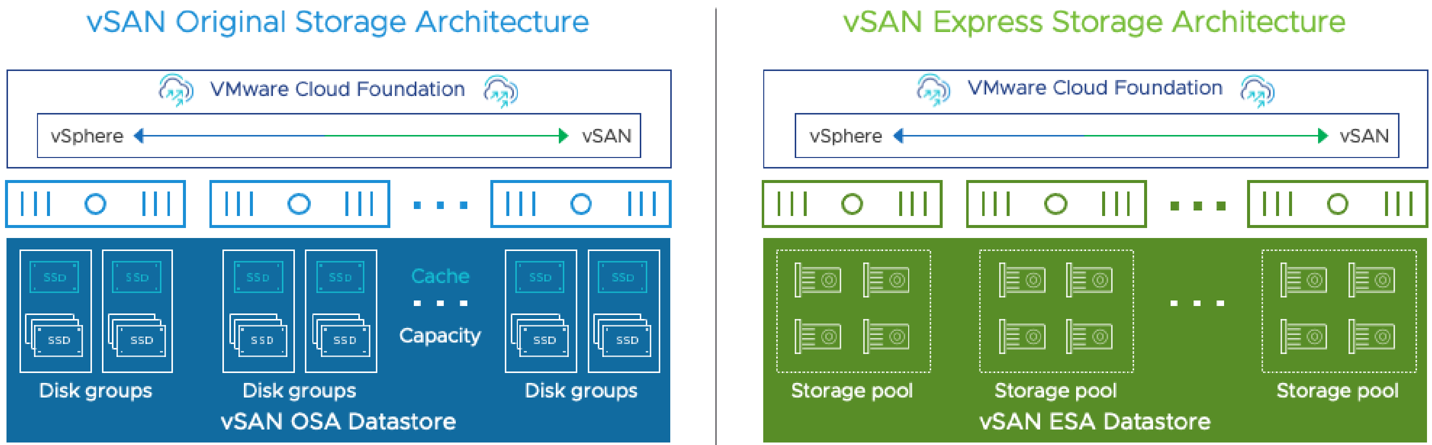

- Dedicated cache and capacity devices
- Accommodates older, wide ranging storage devices
- Uses disk groups to provide performance through caching and capacity

- Flexible architecture using single tier in vSAN 8
- Optimized for high performance NVMe based flash devices
- No disk groups. All devices contribute to performance and capacity

These optimizations not only allow for higher storage capacities, leveraging RAID 5/6 at the performance level of RAID 1 for SAP HANA VMs, but also reduce the number of components (no cache devices needed anymore). This significantly reduces the cost per GB for VMware vSAN SAP HANA HCI clusters.

When running on qualified hardware, ESA consistently delivers superior performance, scalability, resilience, and easier operation for SAP HANA VMs compared to OSA. During our SAP HANA vSAN 8 validation we used SAP HCMT to measure the performance of eight SAP HANA VMs running in parallel on both vSAN OSA- and ESA-configured datastores. Using identical hardware and software configurations (eight NVMe devices), we observed significant performance improvements with ESA, while managing and operating the vSAN configuration was much easier and less error prone. In addition, we used compression with ESA, and it further improved performance.

Note: vSAN ESA is the recommended option for any SAP HANA HCI deployments, starting with Intel Sapphire Rapids (SPR) CPU-based systems. OSA is the storage architecture of choice for older, but still vSAN 8-supported SAP HANA HCI systems.

Solution components

An SAP HANA HCI solution based on VMware technologies is a fully virtualized and cloud-ready infrastructure turnkey solution running on VMware Cloud Foundation (VCF), which include VMware vSphere and vSAN. All components, including compute, storage, and networking, are delivered together as an HCI vendor-specific SAP HANA solution. It's based on a tightly integrated VMware software stack, where vSAN is part of the VMware ESXi kernel to efficiently provide great performing storage.

There are two ways to build a vSAN SAP HANA HCI cluster for an SAP HCI OEM vendor:

- Select certified VCF or vSAN ReadyNodes as the basis for an SAP HANA ready HCI solution meeting the SAP HANA HCI KPIs.
- Use the certified vSAN components and build an SAP HANA ready HCI solution.

Broadcom strongly recommends using certified vSAN ReadyNodes that are validated to provide predictable performance and scalability for your vSAN deployment. If you would still like to build your own vSAN SAP HANA HCI solution, with certified components, select components from the relevant vSAN-certified component lists.

The solution consists of the following components:

- VCF and vSAN certified solutions, as listed on the vSAN ESA HCL or vSAN OSA HCL lists.
- SAP HANA supported server components, as listed here: vSAN I/O Controller, vSAN over RDMA, IO Devices, or vSAN SSD links.
- VMware Cloud Foundation (VCF) that include vSphere and vSAN 8 or later. For Emerald Rapids-based servers vSphere and vSAN 8.0U3 or later is required.
- Review SAP note 3406060 for supported vSAN 8.0 and SAP note 3372365 for supported vSphere 8 versions and possible deployment rules and restrictions, such as with SAP HANA half-socket deployments.
- VMware vCenter version 8.0 or later. Note that vCenter versions need to be aligned with the vSphere version being used.
- HCI vendor-implemented vSAN HCI detection script provided by VMware (adds vSAN HCI identifier to the VM configuration file, see the SAP HANA HCI VMware vSAN Detection Script section for details). Download location published in SAP note 3406060.
- HCI vendor SAP-integrated support process.

After the solutions are defined and built, SAP ICC validates and lists them under **Certified HCI Solutions** in the SAP HANA Supported and Certified Hardware Directory. Figure 4 provides an overview of the solution components and illustrates an example SAP HCI workload domain managed by a VCF management domain. If only vSphere and vSAN are used then the environment is managed by a vCenter Server instance.

> **Note:** It is the responsibility of the SAP HANA HCI hardware vendor, not Broadcom, to design, validate and certify their unique configuration with the SAP HANA Hardware Certification Hyperconverged Infrastructure test scenario (HANA-HyCI).
>
> Broadcom can only support SAP HANA HCI configurations that demonstrably maintain compliance with the SAP-defined HANA-HyCI test scenario and HANA TDI KPI requirements.

Figure 4. Overview of solution components

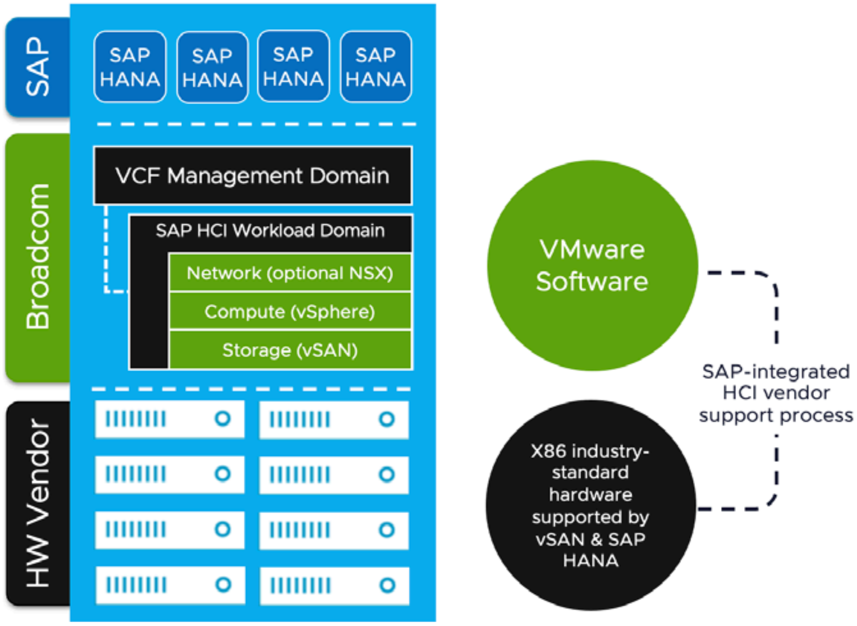

Supported software versions, hardware, and vSAN features

SAP HANA HCI solutions are currently supported with the following VMware vSphere and vSAN versions:

- VMware vSphere 8.0 and later, SAP note 3372365.
  - For older, supported versions see SAP Note 2393917, 2937606, and 3102813.
- VMware vSAN 8.0 and later, SAP note 3406060.
  - For older, supported versions see SAP Note 2718982.
- SAP HANA SPS 12 Rev. 122.19 (or later releases) for production single-VM and multi-VM use cases. See SAP Note 2393917 for detailed information and constraints.

Supported CPUs and SAP HANA VM sizes per SAP HANA deployment option as of this document version and released SAP support notes. For older CPUs refer to older versions of this document.

The **minimum** size of a virtual SAP HANA instance is a half socket (represented by at least **8 physical cores**) and **128 GiB** of **RAM**.

The **maximum size** (as of SAP note 3372365) of a virtual SAP HANA instance for **standard sizing** depends on the CPU architecture:

- **Intel Cascade/Cooper Lake: 56 vCPUs per socket**, the maximum memory is limited to **1.5 TiB per socket** (standard sizing, refer to the workload-based sizing note 2779240 for different configurations). For the maximum supported 8-socket VM size, this would translate to 448 vCPUs and 12 TiB.

  - **OLAP** workload is supported with up to **6 TiB** (0.75 TiB per socket) of memory (Class L) or **12 TiB** (1.5 TiB per socket, Class M - customers must size accordingly) on 8-socket wide VMs (half of these memory sizes if the host or VM is only a 4-socket configuration).
- **Intel Ice Lake: 80 vCPUs per socket**, the maximum memory is limited to **2 TiB per socket**. For the largest supported 2-socket VM size, this would translate to 160 vCPUs and 4 TiB. The **OLAP** workload is supported with up to **2 TiB** (1 TiB per socket) of memory (Class L) or **4 TiB** (2 TiB per socket, Class M); customers must size accordingly.
- **Intel Sapphire Rapids**: 120 vCPUs per socket, the maximum memory is limited to 2 TiB per socket (standard sizing, refer to the workload-based sizing note 2779240 for different configurations). For the largest supported 4-socket **standard** VM size, this would translate to 480 vCPUs and 8 TiB.
  - **2-socket hosts: OLAP workload** is supported with up to **3 TiB** (1.5 TiB per socket) of memory (Class L) or **4 TiB** (2 TiB per socket, Class M) on 2-socket wide VMs; customers must size accordingly.
  - **4-socket hosts: OLAP workload** is supported with up to **8 TiB** (2 TiB per socket) of memory (Class L and M) with no sizing restrictions (depending on the underlying hardware) on 4-socket wide VMs.
- **Intel Emerald Rapids**: 128 vCPUs per socket, the maximum memory is limited to 2 TiB per socket (standard sizing, refer to the workload-based sizing note 2779240 for different configurations). For the largest supported 2-socket **standard** VM size, this would translate to 256 vCPUs and 4 TiB.
  - **2-socket hosts: OLAP workload** is supported with up to **4 TiB** (2 TiB per socket) of memory (Class L) on 2-socket wide VMs.

SAP HANA HCI VM socket specifics

VM size - full socket, maximum of four VMs per vSAN host:

- Intel Cascade/Cooper Lake
  - Minimum 1 socket / maximum 8 sockets
  - 1-, 2-, 3-, 4-, 5-, 6-, 7-, and 8-socket VMs on up to 8-socket servers (fully & partly QPI meshed)
  - Maximum of 4 VMs per 4- and 8-socket hosts, 2 VMs per 2-socket host
- Intel Ice Lake
  - Minimum 1 socket / maximum 2 sockets
  - 1- and 2-socket VMs on up to 2-socket server (fully & partly QPI meshed)
  - Maximum of 2 VMs per 2-socket host
- Intel Sapphire Rapids
  - Minimum 1 socket / maximum 4 sockets
  - 1-, 2-, 3-, and 4-socket VMs on up to 4-socket server (fully & partly QPI meshed)
  - Maximum of four VMs per 4-socket hosts, 2 VMs per 2-socket host
- Intel Emerald Rapids

  - Minimum 1 socket / maximum 2 sockets
  - 1- and 2-socket VMs on up to 2-socket server (fully & partly QPI meshed)
  - Maximum of two VMs per 2-socket host

VM size – “half socket” (NUMA Node sharing VM), maximum of 4 VMs per vSAN host:

- Intel Cascade/Cooper Lake
  - NUMA Node sharing VMs on 2-, 4-, 6-, and 8-socket servers: Supported
  - No odd multiples of 0.5 (half) sockets like 1.5-socket VMs, 2.5-socket VMs, etc.
  - Maximum of four VMs per 2-, 4- and 8-socket host
- Intel Ice Lake
  - NUMA Node sharing VMs on 2-socket server: Supported
  - No odd multiples of 0.5 (half) sockets like 1.5-socket VMs, 2.5-socket VMs, etc.
  - Maximum of four VMs per 2-socket host
- Intel Sapphire Rapids
  - NUMA Node sharing VMs on 2-socket server: Not supported, 4-socket hosts: Supported.
  - No odd multiples of 0.5 (half) sockets like 1.5-socket VMs, 2.5-socket VMs, etc.
  - Maximum of four VMs per 4- and 8-socket hosts, SNC-2 is not supported with the current vSAN version and therefore 2-socket SPR hosts are limited to one VM per CPU socket.
- Intel Emerald Rapids
  - NUMA Node sharing VMs on 2-socket server: Supported
  - No odd multiples of 0.5 (half) sockets like 1.5-socket VMs, 2.5-socket VMs, etc.
  - Maximum of four VMs per 2-socket host

SAP HANA Scale-Out Deployment (OLAP)

CPU architecture:

- Intel Cascade Lake processors of TDIv5 Model: 8 sockets
- Intel Cooper Lake processors of TDIv5 Model: 8 sockets
- Intel Sapphire Rapids processors of TDIv5 Model: 4 sockets

> **Note:** SAP HANA scale-out deployments are only supported by SAP on 4- or 8-socket hosts with at least a 4-socket (full-socket) large VM configuration.

VM size – full socket:

- Node size 6 TiB (Cascade Lake and Cooper Lake)
  - 8-socket VM on **8-socket hosts** (fully and partly UPI meshed)
  - Up to one master plus seven worker nodes are possible. HA nodes should be added according to the cluster-size. FAQ: SAP HANA High Availability [2]
  - Up to 6 **TiB** of **RAM** per scale-out node (VM).
  - Sizing up to **class M** workloads
- Node size 6 TiB (**Sapphire Rapids**)
  - 4-socket VM on **4-socket hosts** (fully and partly UPI meshed)
  - Up to one master plus seven worker nodes are possible. HA nodes should be added according to the cluster-size. FAQ: SAP HANA High Availability [2]
  - Up to 6 **TiB** of **RAM** per scale-out node (VM).
  - Sizing up to **class L** workloads

The SAP HANA BW needed CPU Sizing Class can be found in the BW sizing report under the heading "RSDDSTAT ANALYSIS DETAILS." See SAP note 2610534 for details.

Supported vSAN features and SAP HANA VM on HCI specifics

- No SNC CPU support with vSAN 8.0 / vSAN 8.x.
- vSAN 8 OSA and ESA. ESA supports compression; encryption is not yet supported. ESA is only supported with Ice Lake and later CPUs.
- Minimum of 3, maximum of 64 nodes per SAP HANA HCI vSAN cluster, 4 or more vSAN nodes are recommended for optimal failure resilience and maintenance operations.
- No CPU sockets need to be reserved for vSAN. Note that just like ESXi, vSAN will need system resources. Refer to Broadcom KB 2113954 for information about vSAN resource needs.

> **Note:** SAP Note 3406060 - SAP HANA on VMware vSphere 8 and vSAN 8 will be continuously updated with the latest support information without further notice. The above support status is from July 2025. Before planning an SAP HANA HCI vSAN-based solution, check this SAP note for any updates or newly released vSAN versions.

Reference architecture

The vSAN for SAP HANA reference architectures presented on the following pages serve as a starting point for HCI solution vendors looking to build a VCF/vSAN-based solution that meets SAP HCI certification requirements. Additionally, they provide best practices and guidelines for effectively designing and implementing such solutions.

The reference architecture examples have been tested and validated using SAP-specified HCI tests and have proven to perform well with the provided SAP HANA configuration.

Scope

This reference architecture:

- Provides a technical overview of SAP HANA on a VMware HCI solution.
- Shows the components used to test and validate vSAN with the SAP HCI tests and to meet the storage performance requirements of SAP HANA Tailored Datacenter Integration (TDI).
- Explains configuration guidelines and best practices.
- Provides information about vSAN/HCI solution sizing for SAP HANA VMs.
- Explains the deployment options from a single data center (DC) to stretched cluster configurations.

Architecture overview

vSAN SAP HANA HCI Clusters consist of three or more physical hosts that contain either a combination of SSDs and NVMe devices (for OSA) or only NVMe devices (for ESA). All vSAN nodes contribute cache and capacity to the vSAN distributed/clustered datastore and are interconnected through a dedicated 25 GbE (minimum) or better path-redundant network.

> **Note:** In SAP HANA environments, a four-host setup is recommended for easier maintenance and Stretched Cluster readiness. However, a three-host vSAN cluster meets the minimum requirement for SAP HANA HCI. Additionally, a dedicated 25 GbE or better redundant network card, segmented via VLANs for SAP, vSAN, vMotion, and management traffic, is recommended to optimize SAP-related network performance.

Figure 5 illustrates the high-level architecture of an SAP landscape running on a VMware-based SAP HANA HCI solution, built on multiple vSphere hosts with vSAN OSA and ESA in a single data center configuration. Unlike other SAP HANA HCI solutions, a VMware vSAN-powered SAP HANA HCI solution can utilize all available ESXi host CPU and RAM resources.

As shown in Figure 5, one host serves as a dedicated high availability host, ensuring sufficient compute resources for full host failover. This is crucial for restarting all VMs running on a failed host or during maintenance activities, such as BIOS or vSphere ESXi upgrades.

Different storage policies for SAP HANA and non-SAP HANA database and application servers, based on the failure tolerance, capacity, and I/O requirements, as shown in **Figure 5** (bronze, silver, and gold), ensure that the VMs receive the necessary and configured performance, failure tolerance, and quality of service settings.

OSA and ESA differ in storage policies: **ESA** supports compression and RAID 5/6, whereas **OSA** typically requires RAID 1 with no compression to meet SAP HANA TDI storage KPIs.

Figure 5. **SAP HANA applications on vSAN OSA and ESA high-level architecture**

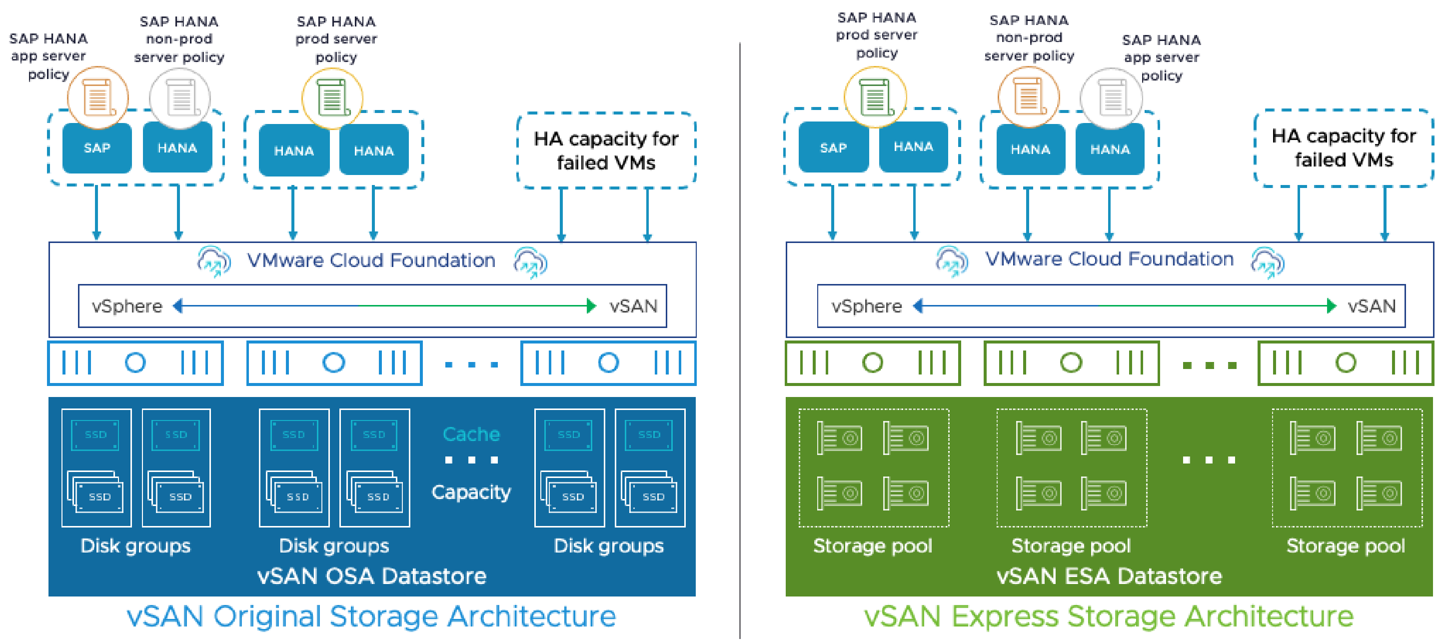

In this environment, storage policies can be dynamically adjusted to meet changing SAP application needs. These policies enable protection levels or I/O limits to be modified without disrupting SAP VM operations, ensuring availability while maintaining defined SLAs and performance requirements.

Scalability and SAP HANA VM Sizes

SAP HANA on vSphere supports a variety of system and VM sizes. The smallest configuration is a half-socket CPU with at least 8 physical cores and 128 GiB of RAM. The largest configuration currently available is an 8-socket wide VM with 12 TiB of RAM and 448 vCPUs. Scale-out SAP HANA deployments can include up to eight nodes, each with 6 TiB, allowing for a total of 48 TiB of system memory in SAP HANA vSAN HCI solutions for OLAP workloads.

vSphere 8.0 U3 and later update versions are supported for vSAN deployments and provide the following maximums per physical host (Table 2). For SAP HANA, only servers with up to eight physical CPU sockets are currently supported.

Table 2. Current vSphere HCI supported host maximums*

	ESX 7.0 U3+	ESX 8.0 U3+
Logical CPUs per host	896	960
Virtual machines per host	1024	
Virtual CPUs per host	4096	
Virtual CPUs per core	32	
RAM per host	24 TiB	
Physical NICs (25 GbE ports)		16
Physical NICs (100 GbE ports)		8
NUMA Nodes / CPU sockets per host	16 (SAP HANA max 8 CPU sockets)	

Table 3 shows the maximum size of an SAP HANA VM and some relevant parameters, such as virtual disk size and number of virtual NICs per VM. You can find additional information in vSphere guest (VM) maximums.

Table 3. **vSphere** guest (VM) maximums

	ESX 7.0 U3+	ESX 8.0 U3+
Virtual CPUs per virtual machine	768	
RAM per non-SAP HANA virtual machine	24 TiB	
Virtual CPUs per HANA virtual machine	448	768
RAM per HANA virtual machine	12 TiB	16 TiB
Virtual SCSI adapters per virtual machine	4	
Virtual NVMe adapters per virtual machine	4	
Virtual disk size	62 TiB	
Virtual NICs per virtual machine	10	

Table 4 shows the vSAN maximum values. The component number or the size of a vSAN SAP HANA HCI cluster depends on the SAP HANA HCI vendor. The minimum SAP HANA supported cluster configuration is three vSAN cluster nodes, and up to 64 cluster nodes are supported by vSAN.

Table 4. vSAN maximums

vSAN 8 Configuration limits	vSAN OSA ESXi host	vSAN ESA ESXi host	vSAN Cluster	vSAN Virtual Machine
vSAN disk groups per host	5	N/A		
Cache disks per disk group	1	N/A		
Components per vSAN host	9,000	27,000		
Capacity tier maximum devices per disk group	7	N/A		
Cache tier maximum devices per host	5	N/A		
Capacity tier maximum devices	35	24		
Number of vSAN datastores per cluster			1	
Stretched Cluster maximum number of data hosts			40 (20 per site)	
Number of vSAN hosts per cluster			64	
Virtual machines per host **(SAP HANA: up to 4 VMs)**				200
Virtual machines per cluster (SAP HANA: **maximum** 4 x (#hosts-1)				6400
Virtual machine virtual disk size				62 TiB
Disk stripes per object				12
Percentage of flash read cache reservation				100
Percentage of object space reservation				100
vSAN networks/physical network fabrics				2

SAP HANA VM can be configured as half-socket, full-socket, 2-, 3-, 4-, 5-, 6-, 7- or 8-socket wide VMs. No CPU sockets need to be reserved for vSAN. If an SAP HANA system fits into these maximums, it can be virtualized with vSphere in a single VM. If it is larger, in terms of CPU or RAM needs, scale-out or SAP HANA data-tiering solutions can instead be used to meet the requirements.

We have validated that four SAP HANA VMs can work consolidated on a vSAN configuration running on 2- or more-socket systems. For detailed information, see the Reference Test System Configuration section.

Table 5 shows the example VM sizes that can be deployed on a vSAN SAP HANA HCI node for various CPUs.

> **Note:** RAM must be reserved for ESXi. The RAM figures listed in Table 4 show the theoretical RAM maximums of a VM. The RAM for ESXi needs to get subtracted from this RAM size. No CPU sockets or other system resources need to be reserved for vSAN. Broadcom KB 314302 about the vSAN memory consumption in ESXi. For sizing purposes an additional 10% sizing buffer should be subtracted from the virtual SAP HANA SAPS capacity to factor in the vSAN CPU resource costs. This buffer is a sizing value; the actual consumption can be much lower in typical customer environments or, in rare situations, such as during vSAN storage re-organization tasks, can be higher.

Table 5. Example of SAP HANA VM configurations and capacity of VMware virtualized SAP HANA instances based on selected 2nd Generation INTEL XEON Scalable Processors[5]

	Intel Xeon Platinum 8280L CPU (Cascade Lake)	Intel Xeon Platinum 8380H CPU and HL (Cooper Lake)	Intel Xeon Platinum 8490H CPU Sapphire Rapids	Intel Xeon Platinum 8592+ CPU Emerald Rapids
SAP Benchmark	2022021 (23. Feb 2022)	2023009 (07. Feb. 2023)	2024018 (20. Mar. 2024)	2024062 (31. May 2024)
Max supported RAM per CPU as of Intel datasheet.	4.5 TiB	4.5 TiB (HL)	4.0 TiB	4.0 TiB
SAP standard supported RAM per CPU[1]	0.75 TiB OLAP 1.5 TiB OLTP	0.75 TiB OLAP 1.5 TiB OLTP	2 TiB OLAP 1.5 TiB OLTP	2 TiB OLAP 1.5 TiB OLTP
CPU Cores per socket (from Intel datasheet)	28	28	60	64
Max. CPUs per ESXi Host CPU threads per platform	8 448	8 448	8 960	2 256
vSAPS per CPU thread w/ and w/o HT	2295 (core w/o HT) 405 (HT gain) Based on SD benchmark cert. 2019023	2459 (core w/o HT) 434 (HT gain) Based on SD benchmark cert. 2023009	2754 (core w/o HT) 486 (HT gain) Based on SD benchmark cert. 2024018	3084 (core w/o HT) 544 (HT gain) Based on SD benchmark cert. 2024062
0.5-Socket SAP HANA VM (Half-Socket)	1 to 16 x max. 14 physical core VMs with min. 128 GiB RAM and max. 768 GiB[3]	1 to 16 x max. 14 physical core VMs with min. 128 GiB RAM and max. 768 GiB[3]	1 to 16 x max. 30 physical core VMs with min. 128 GiB RAM and max. 1024 GiB[3]	1 to 4 x max. 32 physical core VMs with min. 128 GiB RAM and max. 1024 GiB[3]
	vSAPS[4,5] 32,130, 28 vCPUs	vSAPS[4,5] 34,425, 28 vCPUs	vSAPS[4,5] 82,620, 60 vCPUs	vSAPS[4,5] 98,685, 64 vCPUs
1 Socket SAP HANA VM	1 to 8 x max. 28 physical core VM with min. 128 GiB RAM and max. 1536 GiB	1 to 8 x max. 28 physical core VM with min. 128 GiB RAM and max. 1536 GiB	1 to 8 x max. 60 physical core VM with min. 128 GiB RAM and max. 2048 GiB	1 to 2 x max. 64 physical core VM with min. 128 GiB RAM and max. 2048 GiB
	vSAPS[4,5] 75,600, 56 vCPUs	vSAPS[4,5] 81,000 56 vCPUs	vSAPS[4,5] 194,400, 120 vCPUs	vSAPS[4,5] 232,200, 128 vCPUs
2 Socket SAP HANA VM	1 to 4 x max. 56 physical core VM with min. 128 GiB RAM and max. 3072 GiB	1 to 4 x max. 56 physical core VM with min. 128 GiB RAM and max. 3072 GiB	1 to 4 x max. 120 physical core VM with min. 128 GiB RAM and max. 4096 GiB	1 max. 256 physical core VM with min. 128 GiB RAM and max. 4096 GiB

	vSAPS[4,5] 151,200, 112 vCPUs	vSAPS[4,5] 162,000, 112 vCPUs	vSAPS[4,5] 388,800, 240 vCPUs	vSAPS[4,5] 464,400, 256 vCPUs
3 Socket SAP HANA VM	1 to 2 x max. 84 physical core VM with min. 128 GiB RAM and max. 4608 GiB	1 to 2 x max. 84 physical core VM with min. 128 GiB RAM and max. 4608 GiB	1 to 2 x max. 180 physical core VM with min. 128 GiB RAM and max. 6144 GiB	NA
	vSAPS[4,5] 226,800, 168 vCPUs	vSAPS[4,5] 243,000, 168 vCPUs	vSAPS[4,5] 583,200, 360 vCPUs	NA
4 Socket SAP HANA VM	1 to 2 x max. 112 physical core VM with min. 128 GiB RAM and max. 6128 GiB	1 to 2 x max. 112 physical core VM with min. 128 GiB RAM and max. 6128 GiB	1 to 2 x max. 240 physical core VM with min. 128 GiB RAM and max. 8192 GiB	NA
	vSAPS[4,5] 302,400, 224 vCPUs	vSAPS[4,5] 324,000, 224 vCPUs	vSAPS[4,5] 777,600, 480 vCPUs	NA
8 Socket SAP HANA VM	1 x max. 224 physical core VM with min. 128 GiB RAM and max. 12096 GiB	1 x max. 224 physical core VM with min. 128 GiB RAM and max. 12096 GiB	1 x max. 384 physical core VM with min. 128 GiB RAM and max. 16384 GiB	NA
	vSAPS[4,5] 604,800, 448 vCPUs	vSAPS[4,5] 648,000, 448 vCPUs	vSAPS[4,5] 1,244,160, 768 vCPUs[2]	NA

[1] In appliance configuration SAP supports up to 1.5 TiB for OLTP and 0.75 TiB for OLAP based SAP HANA applications with CLX and CPL CPUs. With SPR and EMR CPUs 2 TiB for OLTP and 1.5 TiB for OLAP are supported. More RAM is possible with a TDI phase 5 / workload-based sizing, details see SAP note 2779240.

[2] A maximum of 768 vCPUs can be configured per VM with vSphere 8.0. Actual vCPU configuration depends on the CPU model.

[3] The half-socket RAM figures listed in the table show evenly configured half-socket VM CPU configurations and RAM sizes.

[4] The listed vSAPS figures are based on published SAP SD benchmark results using hyper-threading (2-vCPU configuration), with a 10% assumed virtualization overhead. Actual virtualization costs depend on the CPU generation and utilization, and are typically below 3%, rising up to 10% when CPU utilization is near maximum (e.g., 90%). For simplicity in this sizing guide, we assume a 10% virtualization cost. In the case of a half-socket configuration, an additional 15% reduction must be applied to the SD capacity, on top of the 10% virtualization cost. The figures shown are rounded estimates derived from SAPS values published in SAP SD benchmarks.

[5] In the SAP HANA HCI (vSAN) use case an additional 10% sizing buffer should be removed from the shown vSAPS capacity.

Scale-out deployments are supported with up to **eight** 6 TiB 4-socket wide VMs (one master plus seven worker VMs) on Sapphire Rapids-based systems, up to 3 TiB 4-socket VMs are supported on CLX and CPL-based systems, see SAP note 3406060 for details. This provides for an OLAP-type SAP HANA instance up to 48 TiB with SPR and 24 TiB on CLX/CPL systems in total system RAM by utilizing the scale-out deployment model.

Figure 9 shows an example 4+1 SAP HANA HCI BW cluster. The required SAP HANA “Shared Directory” gets mounted directly to SAP HANA VMs via NFS from an SAP HANA HCI external NFS server. A minimum of one vSAN node is needed to support HA.

VMware vSphere High Availability deployment

VMware vSphere High Availability (HA) delivers a vSphere built-in high availability solution, which can be used to protect SAP HANA. VMware HA provides unified and cost-effective failover protection against hardware and operating system outages and works with a vSAN data store just like it does with any other supported shared storage solution with VMFS. These options are for single site.

VMware HA protects SAP HANA scale-up and scale-out deployments without any dependencies on external components, such as DNS servers, or solutions, such as the SAP HANA Storage Connector API or STONITH scripts.

Note: See the Architecture and Best Practices Guidelines document of SAP HANA on VMware vSphere for detailed information about how VMware HA is used to protect any SAP HANA deployment on vSphere and how to install an SAP HANA scale-out system on VMware vSphere.

In the event of a hardware or OS failure, the entire SAP HANA VM is restarted either on the same host or on another host in the vSphere cluster. Because all virtual disks, such as OS, data, and log VMDKs are stored on the shared vSAN datastore, no specific storage tasks, STONITH scripts, or extra cluster solutions are required.

Figure 6 shows the concept of a typical n+1 configuration for multiple independent SAP HANA scale-up VMs, which is the standard deployment of an SAP HANA HCI solution.

Recommended: Deploy an n+2 configuration for easy maintenance and to protect your systems against a failure while you perform maintenance.

Figure 6. VMware HA protected SAP HANA scale-up VMs in an n+1 vSAN example cluster configuration

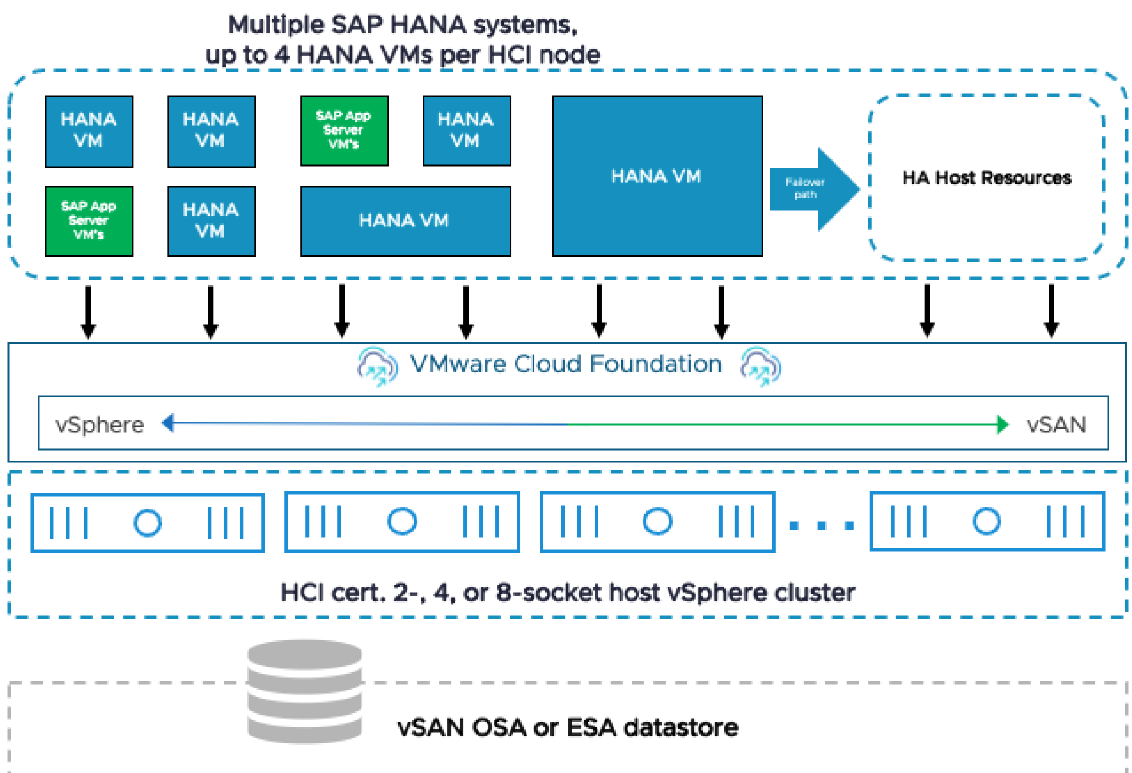

Figure 7 shows the same configuration with an active-active cluster. This ensures that all hosts of a vSAN cluster are used by still providing enough failover capacity for all running VMs in the case of a host failure.

Figure 7. VMware HA active-active n+1 vSAN example cluster configuration

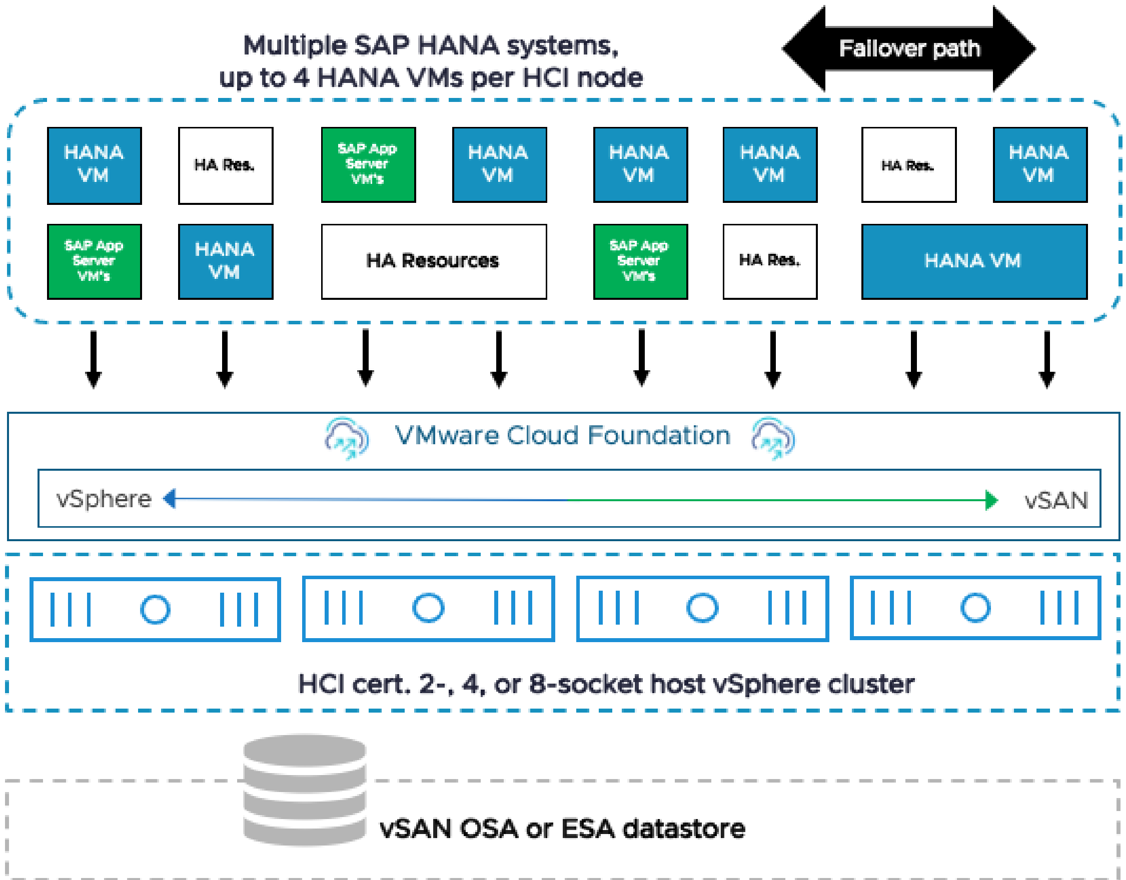

The same concept is also used for SAP HANA scale-out configurations. Note that for SAP HANA scale-out, only 4-socket or 8-socket servers are supported. The minimum SAP HANA VM must be configured with four virtual CPU sockets. Figure 8 shows a configuration of three SAP HANA VMs (one master with two worker nodes) running on the virtual SAP HANA HCI platform. The NFS mounted SAP HANA shared directory is not part of the vSAN SAP HANA scale-out cluster and needs to be provided as a vSAN external NFS service.

Figure 8. SAP HANA scale-out VMware HA n+1 vSAN example cluster configuration based on HCI nodes

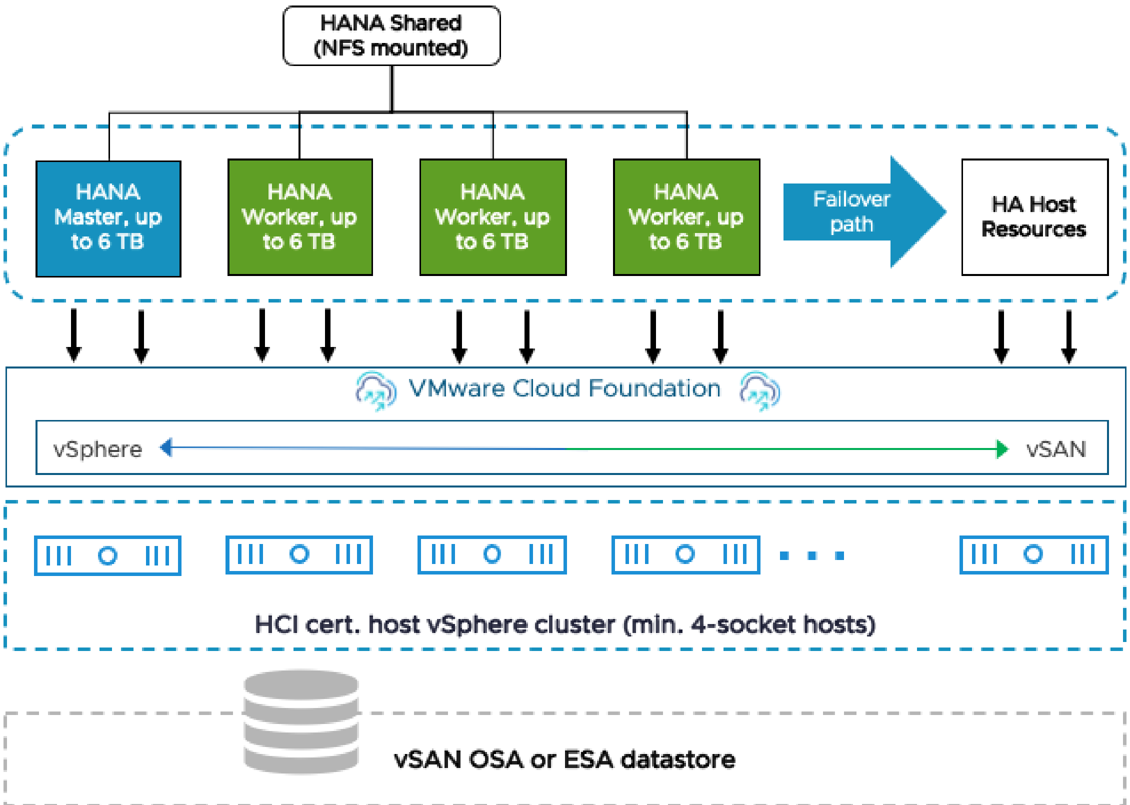

Unlike physical deployments, there are no dependencies on external components, such as DNS servers, SAP HANA Storage Connector API, or STONITH scripts. VMware HA will simply restart the failed SAP HANA VM on the VMware HA/standby server. The /hana/shared directory is mounted via NFS, just as recommended with physical systems, and will move with the VM that has failed. The access to /hana/shared is therefore guaranteed.

Note: VMware HA can protect only against OS or VM crashes or HW failures. It cannot protect against logical failures or OS file system corruptions that are not handled by the OS file system. If necessary, third-party HA solutions like Pacemaker can be used instead of VMware HA. vSAN File Service (NFS) can be used when the File Services VM is deployed on a dedicated CPU socket, which may be combined with another SAP VM like for SAP application server VMs. Co-deployments with an SAP HANA VM are not supported and the usage for SAP HANA scale-out deployments is therefore limited due to this requirement.

vSphere Cluster Services (vCLS) deployment

Starting with vSphere 7.0 Update 1, upgrading vCenter enables the vSphere Cluster Service (vCLS) and runs it in all vSphere clusters. VMware would like to make critical cluster services like vSphere High Availability (HA) and vSphere Distributed Resource Scheduler (DRS) always available and vCLS is an initiative to reach that vision.

The dependency of these cluster services on vCenter is not ideal and in the vSphere 7 Update 1 release VMware introduced the vSphere Clustering Service (vCLS), which is the first step to decouple and distribute the control plane for clustering services in vSphere and to remove the vCenter dependency. In the future, if vCenter Server becomes unavailable, vCLS will ensure that the cluster services remain available to maintain the resources and health of the workloads that run in the clusters.

As detailed in the vSphere specific SAP HANA support notes, it is not supported to run a non-SAP HANA VM on the same NUMA node where an SAP HANA VM already runs:

"SAP HANA VMs can get co-deployed with SAP non-production HANA or any other workload VMs on the same vSphere ESXi host, if the production SAP HANA VMs are not negatively impacted by the co-deployed VMs. In case of negative impact on SAP HANA, SAP may ask to remove any other workload." Also, "no NUMA node sharing between SAP HANA and non-HANA allowed."

Due to the mandatory and automated installation process of vCLS VMs, when upgrading to vCenter 7.0 Update 1 or any later vCenter version, it's necessary, because of the above guidelines, to check if vCLS VMs got co-deployed on vSphere ESXi hosts that run SAP HANA production level VMs. If this is the case, then these VMs must be migrated to hosts that do not run SAP HANA production level VMs.

This can be achieved by configuring vCLS VM Anti-Affinity Policies. A vCLS VM anti-affinity policy describes a relationship between VMs that have been assigned a special anti-affinity tag (e.g., tag name SAP HANA) and vCLS system VMs.

If this tag is assigned to SAP HANA VMs, the vCLS VM anti-affinity policy discourages placement of vCLS VMs and SAP HANA VMs on the same host. With such a policy it can be assured that vCLS VMs and SAP HANA VMs do not get co-deployed.

After the policy is created and tags are assigned, the placement engine attempts to place vCLS VMs on the hosts where tagged VMs are not running.

In the case of an SAP HANA HCI partner system validation, or if an additional non-SAP HANA ESXi host cannot be added to the cluster, then Retreat Mode can be used to remove the vCLS VMs from this cluster. Note the impacted cluster services (DRS) due to the enablement of Retreat Mode on a cluster.

> **Note:** Setting vCLS VM Anti-Affinity Policies ensures that a vCLS VM does not get placed on hosts that run SAP HANA VMs. This requires that some hosts that do not run SAP HANA tagged VMs. **If this is not possible then you can use Retreat Mode to remove the vCLS VMs from the cluster.** In the case of VCF an additional node (e.g., the HA node) is required per cluster in order to offload the vCLS VMs.

For more details and vCLS deployment examples in the context of SAP HANA, refer to this blog.

Embedded vSphere Cluster Services (Embedded vCLS)

With vSphere 8.0 Update 3, **Embedded vSphere Cluster Services (Embedded vCLS)** were introduced. Embedded vCLS is automatically activated on vSphere clusters that are managed by vCenter 8.0 Update 3 or later and include at least one ESXi 8.0 Update 3 or later host.

Embedded vCLS uses the **CRX runtime**, originally developed for the vSphere IaaS control plane. These VMs have no storage footprint and run entirely in host memory. The ESXi host directly creates the embedded vCLS VMs, there is no OVF deployment initiated by vCenter, and ESX Agent Manager (EAM) is no longer involved.

When using Embedded vCLS the number of vCLS VMs per cluster has also been reduced from up to three VMs to up to two VMs. A single-node cluster will use a single Embedded vCLS VM and clusters of two or more hosts will use two Embedded vCLS VMs.

In our SAP HANA tests, we observed **no measurable performance impact**. The required CPU host resources and memory consumption are minimal, typically falling within the standard ESXi host CPU and memory overhead. As a result, Embedded vCLS does **not affect SAP HANA sizing or performance**.

While **anti-affinity rules** using Compute Policies and **Retreat Mode** are still available for troubleshooting or specific deployments, they are **not required** to avoid co-locating Embedded vCLS VMs with SAP HANA VMs. This marks an improvement over vSphere 7, where anti-affinity was often needed due to traditional vCLS VMs using shared storage and VM-usable resources.

Figure 9 shows an SAP HANA HCI vSAN-based cluster with two embedded vCLS VMs and their corresponding resource usage. Since the vCLS VMs are embedded, only the utilized host resource memory is shown.

Note: Setting vCLS VM Anti-Affinity Policies Retreat Mode **is not required** for vSphere 8.

Figure 9. SAP HANA HCI vSAN based cluster with two running vCLS VMs

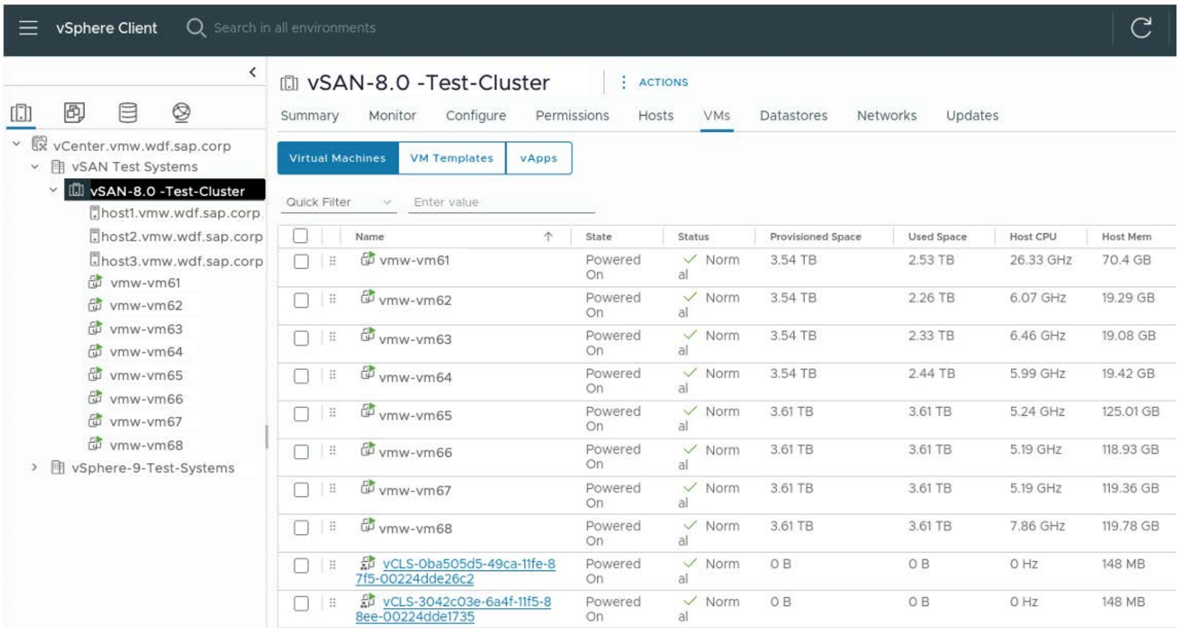

VMware vSAN stretched clusters

Stretched clusters extend the vSAN cluster from a single data site to two sites for a higher level of availability and cross-site load balancing. Stretched clusters are typically deployed in environments where the distance between data centers is limited, such as metropolitan or campus environments. The two sites typically have a similar number of hosts, and are connected using a high bandwidth/low latency connection known as an inter-site link, or ISL. A third site, or tertiary site, hosts a vSAN Witness Host that is connected to both data sites.

You can use stretched clusters to manage planned maintenance and avoid disaster scenarios because maintenance or loss of one site does not affect the overall operation of the cluster. In a stretched cluster configuration, both data sites are active. If either site fails, vSAN uses the storage on the other site. vSphere HA restarts any VMs that must be restarted on the remaining active site.

Support for SAP HANA VMs in vSAN stretched clusters

SAP HANA has been tested and is supported on vSAN 7 Update 1 and later and vSAN 8 Update 3 and later stretched cluster environments. A key requirement for SAP HANA is that log file writes must have a latency of less than 1 millisecond (ms), and this latency must be maintained consistently across both primary and secondary sites.

The underlying network infrastructure plays a critical role in meeting this requirement, as the maximum supported distance between the two sites is determined solely by the ability to maintain the 1 ms SAP HANA log file write latency, which is typically achieved when the round-trip time (RTT) remains below 1 ms.

Testing with one of our HCI hardware partners confirmed that all SAP HANA performance throughput and latency KPIs were met when the sites were located up to 30 km (18.6 miles) apart. However, actual supported distances could be shorter depending on the specific network components and architecture in use. SAP HANA vSAN stretched clusters are normally used campus wide, therefore distances up to 5 km are supported with the RTT requirement below 1 ms for standard support. Longer distances may be evaluated upon request, provided that the round-trip time (RTT) remains below 1 ms.

It is the responsibility of the SAP HANA HCI hardware vendor and system integrator, not Broadcom, to design and validate the configuration against SAP HANA TDI network and storage KPIs, like the 1 ms log file latency KPI. SAP note 2493172 (SAP HANA Hardware and Cloud Measurement Tools) provides the necessary details about how to measure and analyze the configuration for SAP HANA readiness.

> **Note:** Broadcom can only support stretched cluster configurations that demonstrably maintain compliance with the SAP HANA TDI storage KPI requirements.

Designing a vSAN stretched cluster

vSAN uses a construct called a **fault domain** to distribute data in a resilient and fault-tolerant manner. By default, each host is treated as a separate fault domain, ensuring data remains available in the event of a host failure. vSAN also offers an optional feature called **vSAN Fault Domains**. When one defines a group of hosts as a “fault domain” it identifies this as an additional logical boundary of failure. It can use this information to write the data in such a way that if that defined fault domain is offline, the data will remain available. The Fault Domains feature is commonly used to ensure rack-level resilience in a data center and is compatible with both vSAN OSA and ESA.

In **vSAN stretched cluster** configurations, the fault domain concept is extended across physical sites rather than within a single data center. Each fault domain represents an entire site, helping ensure VM availability even in the event of a complete site outage. The roles of the sites in a vSAN stretched cluster are:

- **Preferred Site** – This is the primary data site and the default owner of the objects. In the event of site isolation, the Preferred site always takes precedence.
- **Non-Preferred Site** – This serves as the secondary or alternate data site and is the second fault domain.
- **Witness Site** – A third location that hosts the Witness appliance, which contains only witness components and no data.

If the connection between the two data sites is lost, vSAN continues operation with the Preferred site, provided it is operational and not in a resynchronization or degraded state. vSAN always prioritizes the site that ensures the highest data availability.

Figure 10 illustrates a simplified stretched cluster setup:

- **Prod Site X** is the Preferred site, which can range from a full data center to a single rack.
- **Non-Prod Site Y** serves as the secondary / failover site.
- To handle localized failures, high availability (HA) capacity can be provisioned within each site.

For a successful failover, **Site Y must have sufficient capacity** to host the failed-over SAP HANA VMs. This can be achieved by pre-provisioning extra capacity, reserving failover resources, or restricting failover to critical VMs only.

The **Witness Appliance** must be deployed at a third, independent location (**Site Z**). It plays a critical role in failover decision-making and must run on an ESXi host that supports the SAP application server VMs.

Figure 10. SAP HANA stretched cluster example configuration

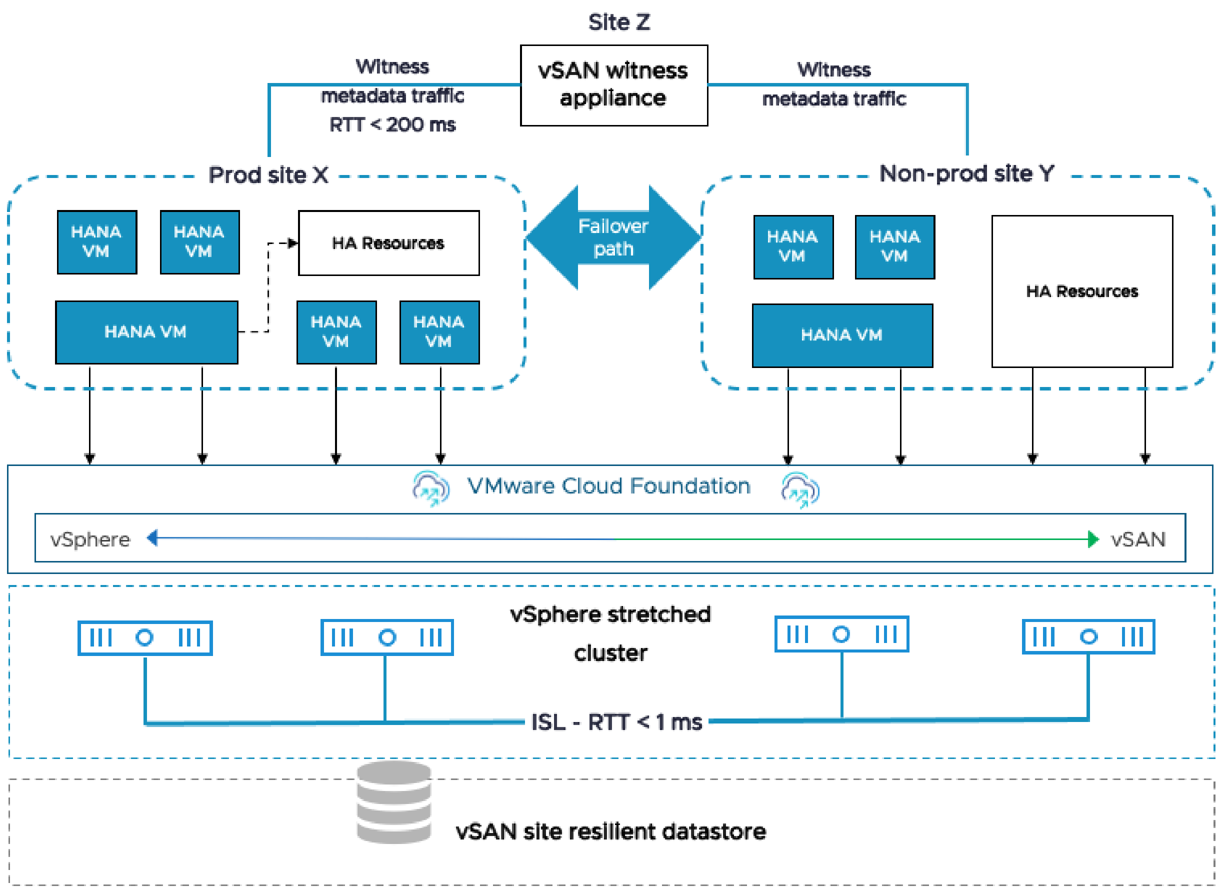

While stretched clusters are typically deployed across geographically separated sites for disaster recovery, the fault domains can also represent locations within a large building or campus. Although this setup is not suitable for natural disaster protection, it might still fulfill specific customer requirements by maintaining SAP HANA storage KPIs.

For scenarios requiring protection against natural disasters, **asynchronous replication** is necessary to preserve storage KPIs, albeit with a trade-off of potential data loss due to the nature of asynchronous replication.

Review the [vSAN Stretched Cluster Guide](), Recommendations for vSAN 8 U3, as a part of VMware Cloud Foundation 5.2 document for detailed information and recommendations to design a vSAN stretched cluster.

Summary:

- Tested distances up to 30 km sustain the less than 1 ms KPI for SAP HANA log writing.

- SAP HANA vSAN stretched clusters are normally used campus wide, therefore distances up to 5 km are supported with the RTT requirement below 1 ms. Alternatively, longer distances that can maintain a below 1 ms RTT (or < 0.5 ms one way) time might be supported upon request.
- The witness VM can run on any vSphere host at a third site Z; the resource requirements are very light.
- Failover capacity is n+1, where 1 stands for one host failure. N+2 would support two host failures.
- Failover can be bi-directional.

Setting the storage policy for an SAP HANA VM on a vSAN stretched cluster

The first stage in setting up a stretched cluster is to configure the fault domains according to the physical location of the vSAN hosts. Detailed information on how to set up a stretched cluster can be found in the vSAN Stretched Cluster Guide.

Figure 11. Fault domain – configure stretched cluster

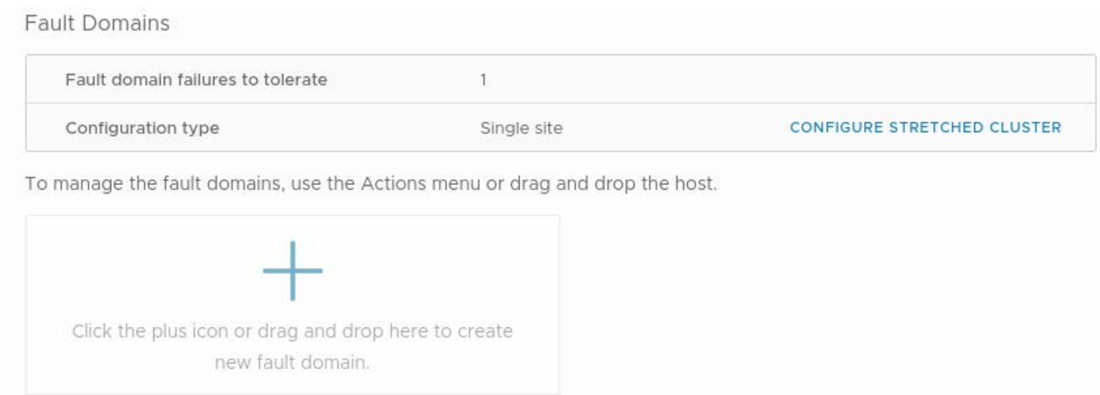

For an SAP HANA VM in a vSAN stretched cluster, you must change the VM storage policy to **Site mirroring – stretched cluster**. If you use vSAN OSA, set **Failures to tolerate** to **1 failure – RAID-1 (Mirroring)**, as shown in Figure 12. If you use vSAN ESA you have the additional options of **1 failure - RAID-5 (Erasure Coding)**, **2 failures – RAID-1 (Mirroring)**, or **2 failures – RAID-6 (Erasure Coding)**. The 2-failure options will require more storage capacity.

Figure 12. SAP HANA VM Storage Policy in vSAN stretched cluster

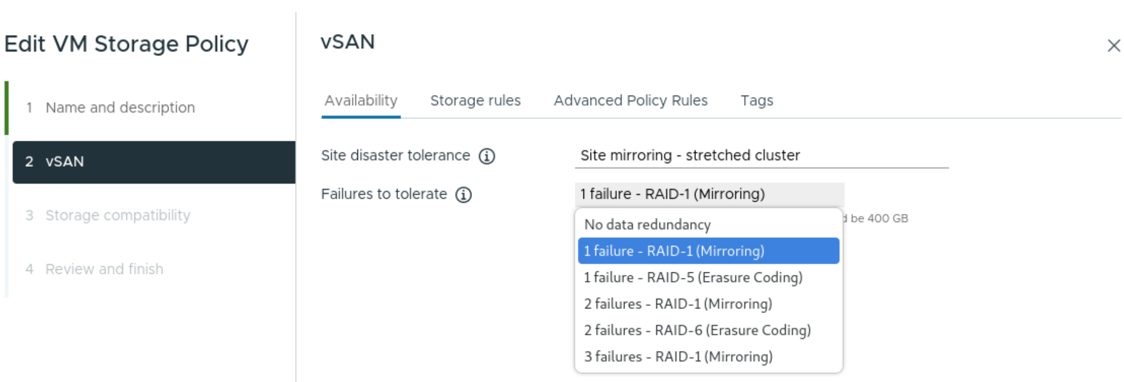

When **Site disaster tolerance** is set to **Dual site mirroring (stretched cluster)**, a copy of the data goes to both sites. When the **Failures to tolerate** rule is equal to **1**, writes will be mirrored across sites; therefore, make sure sufficient disk capacity is available at each site.

Reference test system configuration

This section introduces the resources and configurations including:

- Reference environment
- Quality of service policy
- Hardware resources
- Software resources

> **Note:** The reference configuration shown is an example, not a requirement. The **HCI solution vendor** may or may not follow the example configurations and might select different components. The only requirements are that only components supported/certified by the HCI vendor and Broadcom are used, and that the system meets the SAP HANA TDI storage KPIs and SAP HCI defined certification test requirements.

Reference environment

To ensure that a vSAN-based HCI solution is SAP HANA ready, use SAP test tools and test cases (HANA-HyCI 1.0 v1.2 test scenario). These test tools validate the performance, scalability, and consistency of the system. On the "Certified Hyper-Converged Infrastructure Solutions" web page, vSAN-supported configurations from different vendors are listed.

VMware has validated the reference configurations listed on the next pages, and they provide examples for scale-up and scale-out configurations. The scale-up example configurations are based on 2-,4-, and 8-socket systems configured as 3-host and 4-host vSAN OSA clusters with two disk groups and four disk groups per server, or with ESA, where no disk groups are defined. The configurations had NVMe devices from different vendors as the caching layer and capacity disks. With either configuration (OSA or ESA), it's possible to meet the SAP HANA KPIs defined for each configuration.

The scale-out test configuration ran on a 4-host vSAN cluster. Each host had four CPUs, and four disk groups installed and configured. In the 8-socket host case, two individual scale-out systems could be consolidated on an 8-socket host configuration.

> **Note:** Only 4-socket or larger host systems are supported by SAP for SAP HANA scale-out deployments.

IOPS limit per partition policy: quality of service option to storage throughput

vSAN storage policies provide the option to set IOPS on the partition (VMDK) that this policy is applied to. You can set this IOPS value at the time you create a policy or you can edit it later. A value of "0" means that there's no IOPS limit set for the partition. Setting an IOPS limit helps to control the environment and can ensure that the VMDK disks / partitions provide the IOPS needed for the use case. It can also help to ensure fair treatment and avoid situations where one VM gets more IOPS than another.

Besides the OS, an application, and /hana/shared, an SAP HANA instance has log and data partitions. We have done internal testing where the IOPS limit is set on the data partition. To do this, configure a policy with limits and allocate this policy to a VMDK that you mount as a data partition. A minimum setting of around 13,700 IOPS on the test setup's data and log partition has shown it can handle the KPI requirements of the data and log partition. It's not necessary or recommended to set IOPS limits for OS and SAP binary / share partition. This is the minimum value and anything above this value is also valid. However, for your unique configuration, we advise you to test the values you set to make sure they allow you to maintain the HANA HCMT KPIs over all running SAP HANA VMs in your test setup, e.g., four VMs per host.

Use cases:

- Depending on application requirements, a policy with IOPS limits set will ensure that there is fair I/O distribution among the different VMs for that specific partition. This is an option you can use if there are multiple SAP HANA VMs running that require control to ensure a minimum I/O distribution among them.
- You want to control or limit the IOPS of other VMs, like non-HANA VMs or non-prod HANA VMs, to ensure that the production SAP HANA VM is getting the most I/O resources.

Hardware resources

The selected hardware configuration, such as CPU type (e.g., frequency, number of cores), and particularly the choice of storage devices (NVMe-only or a combination of NVMe and SSDs), directly impact the number of SAP HANA VMs that can be deployed on a given vSAN setup. Additionally, vSAN OSA and vSAN ESA have distinct hardware requirements, that could further influence and potentially limit the number of SAP HANA VMs that can be supported in your vSAN cluster.

The configurations on the next pages are example configurations we've used successfully for our vSAN validation efforts; these should allow you to pick a suitable vSAN HCI cluster configuration for your needs.

We configured and tested different vSAN 8 OSA and ESA clusters in the 3- and 4-host environments, respectively:

- A vSAN 8 OSA SAP HANA HCI environment with three 2-socket Ice Lake processor hosts with 2 SAP HANA VMs per CPU socket (example 1).
- A vSAN 8 OSA and ESA SAP HANA HCI environment with three 4-socket Sapphire Rapids processor hosts with 1 SAP HANA VM per CPU socket (example 2).
- A vSAN 8 ESA SAP HANA HCI environment with three 2-socket Emerald Rapids processor hosts with 2 SAP HANA VMs per CPU socket (example 3).

Note: For older CPU types and system configurations, like CLX- or CPL-based systems, refer to older revisions of this document.

Configurations are not limited to the ones listed. The example configurations are only a starting point to determine an SAP HANA HCI hardware vendor's specific and unique solution, which might use lower core count CPUs and different memory, network, and storage capacity configurations.

For our tests, we minimally used **two 25 GbE** links as a dedicated vSAN network link across all configurations, alongside a minimal **10 GbE** SAP client and admin network, and **two 25 GbE** for the vSphere (e.g., vMotion) network. Additionally, we used **two 100 GbE** network cards that supported **VLAN-based separation** of different traffic types (vSAN, vMotion, SAP application, etc.).

Example 1: Ice Lake 2-socket / 2-HANA VMs per CPU – vSAN 8 OSA

The Ice Lake test environment was set up using three 2-socket Ice Lake–based hosts with vSAN 8 OSA and 16 Performance Class F NVMe devices per host. This configuration created a vSAN datastore that meets the SAP HANA HCI test requirements and KPIs for running four SAP HANA VMs per host. The third node also provided high availability (HA) for all eight SAP HANA VMs running in the vSAN cluster in a 2+1 configuration.

Each vSphere host was configured with four disk groups, each consisting of one 1.6 TiB NVMe device for caching and three 1.6 TiB NVMe capacity devices of the same type. The raw storage capacity of the 3-node vSAN cluster totaled 57.6 TiB (3 hosts × 4 disk groups × 3 capacity devices × 1.6 TiB). Storage capacity can be scaled by adding more capacity devices, using larger devices, or adding additional vSAN nodes. While vSAN clusters support up to 64 nodes, the actual number of nodes supported in an SAP HANA HCI solution depends on the HCI vendor.

Each host was equipped with two 100 GbE network adapters configured as an active-standby failover pair for ESXi management, vMotion, vSAN, and SAP application-to-database traffic. All networks were logically segmented using VLANs.

Each vSphere host in the vSAN cluster had the following configuration and could support up to four SAP HANA VMs.

Table 6. Configuration for Example 1: Ice Lake 2-socket / 2-HANA VMs per CPU – vSAN 8 OSA

PROPERTY	SPECIFICATION
ESXi host model	Hardware vendor and SAP HANA-supported Ice Lake 2-socket server model
ESXi host CPU	2 x Intel Xeon Platinum 8362, 32 physical cores per socket
ESXi host RAM	1.0 TiB
Number of hosts	Minimum of three hosts (2 active, one HA host), recommended four hosts (in a 3+1 or 2+2 configuration)
Network adapters	2 x 1 GbE BMC / IPIMI for bare metal server management 2 x 100 GbE as a shared, VLAN separated network for vSAN, vMotion, ESXi admin, and SAP network traffic See the "Network configuration" section of this guide for the recommended network configurations
vSAN Storage adapter	N/A, we direct-attached the NVMe devices
Storage configuration	vSAN 8 OSA Disk groups: 4; 16 NVMe devices in total Cache tier: 4 x Performance Class F NVMe or better, like 1.6 TiB Kioxia Type: KCM61VUL1T60 Capacity tier: 4 x 3 x Performance Class F NVMe or better, like 1.6 TiB Kioxia Type: KCM61VUL1T60
Linux LVM configuration	Optional usage of Linux LVM for log and data In our setup we used 4 VMDK disks in LVM each, to create the needed log and data volumes for SAP HANA
SAP HANA VMs	Two per CPU (half-socket), four per Ice Lake HCI host

Example 2: Sapphire Rapids 4-Socket / 1-HANA VM per CPU vSAN 8 OSA and ESA

The Sapphire Rapids 4-socket test environment was set up using three 4-socket Sapphire Rapids-based hosts with vSAN 8 OSA and for a second test with ESA. For both tests we used 16 Performance Class F NVMe devices per host. This configuration created a vSAN datastore based on vSAN OSA and ESA that meets the SAP HANA HCI test requirements and KPIs for running four SAP HANA VMs per host. The third node also provided high availability (HA) for all eight SAP HANA VMs running in the vSAN cluster in a 2+1 configuration.

In the OSA test case each vSphere host was configured with four disk groups, each consisting of one 1.6 TiB NVMe device for caching and three 1.6 TiB NVMe capacity devices of the same type. The raw storage capacity of the 3-node vSAN cluster totaled 57.6 TiB (3 hosts × 4 disk groups × 3 capacity devices × 1.6 TiB). Storage capacity can be scaled by adding more capacity devices, using larger devices, or adding additional vSAN nodes. While vSAN clusters support up to 64 nodes, the actual number of nodes supported in an SAP HANA HCI solution depends on the HCI vendor. In the ESA test case, we had all 16 Performance Class F 1.6 TiB NVMe devices assigned to the ESA datastore, totaling to a raw storage capacity of 76.8 TiB (3 hosts × 16 capacity devices × 1.6 TiB), which is 33% more storage capacity than with ESA when the same devices are used.

Each host was equipped with two 100 GbE network adapters configured as an active-standby failover pair for ESXi management, vMotion, vSAN, and SAP application-to-database traffic. All networks were logically segmented using VLANs.

Each vSphere host in the vSAN cluster had the following configuration and could support up to four SAP HANA VMs.

Table 7. Configuration for Example 2: Sapphire Rapids 4-socket / 1-HANA VMs per CPU – vSAN 8 OSA and ESA

PROPERTY	SPECIFICATION
ESXi host model	Hardware vendor and SAP HANA-supported Sapphire Rapids 4-socket server model
ESXi host CPU	4 x Intel Xeon Platinum 8490H, 60 physical cores per socket
ESXi host RAM	2.0 TiB
Number of hosts	Minimum of three hosts (2 active, one HA host), recommended four hosts (in a 3+1 or 2+2 configuration)
Network adapters	2 x 1 GbE BMC / IPIMI for bare metal server management 4 x 25 GbE as a shared, VLAN separated network for vSAN, vMotion, ESXi admin, and SAP network traffic See the "Network configuration" section of this guide for the recommended network configurations
vSAN Storage adapter	N/A, we direct-attached the NVMe devices
Storage configuration	vSAN 8 OSA Disk groups: 4; 16 NVMe devices in total Cache tier: 4 x Performance Class F NVMe or better, like 1.6 TiB Kioxia Type: KCM61VUL1T60 Capacity tier: 4 x 3 x Performance Class F NVMe or better, like 1.6 TiB Kioxia Type: KCM61VUL1T60
Storage configuration	vSAN 8 ESA 16 x Performance Class F NVMe or better, like 1.6 TiB Kioxia Type: KCM61VUL1T60
Linux LVM configuration	Not used in this test
SAP HANA VMs	1 per CPU (full-socket), four per 4-socket Sapphire Rapids HCI host

Example 3: Emerald Rapids 2-Socket / 2-HANA VM per CPU vSAN 8 ESA

The Emerald Rapids test environment was set up using three 2-socket Emerald Rapids-based hosts with vSAN 8 ESA and 16 Performance Class F NVMe devices per host. This configuration created a vSAN datastore that meets the SAP HANA HCI test requirements and KPIs for running four SAP HANA VMs per host. The third node also provided high availability (HA) for all eight SAP HANA VMs running in the vSAN cluster in a 2+1 configuration.

Each vSphere host was configured with four disk groups, each consisting of one 1.6 TiB NVMe device for caching and three 1.6 TiB NVMe capacity devices of the same type. The raw storage capacity of the 3-node vSAN cluster totaled 57.6 TiB (3 hosts × 4 disk groups × 3 capacity devices × 1.6 TiB). Storage capacity can be scaled by adding more capacity devices, using larger devices, or adding additional vSAN nodes. While vSAN clusters support up to 64 nodes, the actual number of nodes supported in an SAP HANA HCI solution depends on the HCI vendor.

Each host was equipped with two 100 GbE network adapters configured as an active-standby failover pair for ESXi management, vMotion, vSAN, and SAP application-to-database traffic. All networks were logically segmented using VLANs.

Each vSphere host in the vSAN cluster had the following configuration and could support up to four SAP HANA VMs.

Table 8. Configuration for Example 3: Emerald Rapids 2-socket / 2-HANA VMs per CPU – vSAN 8 ESA

PROPERTY	SPECIFICATION
ESXi host model	Hardware vendor and SAP HANA-supported Emerald Rapids 2-socket server model
ESXi host CPU	2 x Intel Xeon Platinum 8592+, 64 physical cores per socket
ESXi host RAM	2.0 TiB
Number of hosts	Minimum of three hosts (2 active, one HA host), recommended four hosts (in a 3+1 or 2+2 configuration)
Network adapters	2 x 1 GbE BMC / IPIMI for bare metal server management 2 x 100 GbE as a shared, VLAN separated network for vSAN, vMotion, ESXi admin, and SAP network traffic See the "Network configuration" section of this guide for the recommended network configurations
vSAN Storage adapter	N/A, we direct-attached the NVMe devices
Storage configuration	vSAN 8 ESA 16 x Performance Class F NVMe or better, like 1.6 TiB Kioxia Type: KCM61VUL1T60
Linux LVM configuration	Not used with this setup

Software resources

Table 9 shows the VMware software resources used in this solution. From the SAP side, both SAP HANA version 1.0 and 2.0 are supported.

Table 9. VMware software resources

SOFTWARE	VERSION	PURPOSE
vCenter Server	8.0 U3	Provides a centralized platform for managing vSphere environments
ESXi	8.0 U3	Hosts VMs and provides the vSAN cluster
vSAN	8.0 U3	Provides a software-defined storage solution for hyperconverged infrastructure (HCI)

SAP HANA HCI on vSAN configuration guidelines

Selecting the correct components and a proper configuration is vital to achieving the performance and reliability requirements for SAP HANA. The server and vSAN configurations (RAM and CPU) determine how many and which SAP HANA VMs are supported.

The vSAN configuration specifics—such as the tested storage policies, number of disk groups, and the selected flash devices for the caching and capacity storage tier or the vSAN network configuration—define how many SAP HANA VMs a vSAN datastore can support.

Network configuration

To build an SAP HANA HCI system you need dedicated or VLAN-separated vSAN networks for the SAP application, user, admin, and ESXi management traffic. Follow the SAP HANA Network Requirements document to decide how many additional network cards must be added to support a specific SAP HANA workload on the servers.

Table 10 provides an example configuration for an SAP HANA HCI configuration based on dual-port network cards.

Table 10. Recommended SAP HCI on vSphere network configuration

		Admin and VMware HA network	Application server network	vMotion network	vSAN network	Backup network	Scale-out internode network	System replication network
Host network configuration[1]	Network label	Admin	App-Server	vMotion	vSAN	Backup	Scale-Out	HANA-Replication
	Recommended MTU size	Default (1500)	Default (1500)	9000	9000	9000	9000	9000
	Bandwidth	>=10 GbE	>=10 GbE	>=25 GbE[2]	>=25 GbE[2]	>=10 GbE[2]	>=10 GbE[2]	>=10 GbE[2]
	Physical NIC	0	0	1	1	Optional	Optional	Optional
	Physical NIC port	1	2	1	2	Optional	Optional	Optional
	VLAN ID	100	101	200	201	202	203	204
VM guest network cards	Virtual NIC (inside VM)	0	1	—	—	2	2	3

[1] All networks use dual-port NICs, configured as failover teams.

[2] The selected network card bandwidth for vSAN directly influences the number of SAP HANA VMs that can be deployed on a host of the SAP HANA HCI vSAN cluster. For instance, a network >=25 GbE is needed if vSAN ESA is used. In the stretched cluster case, the witness appliance VM can leverage the App server network. A dedicated witness network is not required.

In addition to using a high-bandwidth, low-latency network for vSAN, we recommend using a vSphere Distributed Switch for all vSphere-related network traffic (such as vSAN and vMotion).

A Distributed Switch acts as a single virtual switch across all associated hosts in the data cluster. This setup allows VMs to maintain a consistent network configuration as they migrate across multiple hosts.

The vSphere Distributed Switch in the SAP HANA HCI solution uses at least two 25 GbE dual-port network adapters for teaming and failover. A port group defines properties regarding security, traffic shaping, and NIC teaming. We recommend using the default port group setting and changing the uplink failover order as shown in Table 11. The table also shows the distributed switch port groups created for different functions and the respective active and standby uplinks to balance traffic across the available uplinks.

Tables 10 and 11 show examples of how to group the network port failover teams. They do not show the optional networks needed for vSphere, SAP HANA system replication, or scale-out internode networks, which require additional network adapters.

Table 11. Minimum vSphere uplink failover network configuration (if all NICs have the same bandwidth to allow physical NIC failure recovery)

Network component	VLAN[1]	Active uplink	Standby uplink
ESXi Admin + VMware HA network	100	Nic0-Uplink1	Nic1-Uplink1
SAP Application Server network	101	Nic0-Uplink2	Nic1-Uplink2
vMotion network	200	Nic1-Uplink1	Nic0-Uplink1
vSAN network	201	Nic1-Uplink2	Nic0-Uplink2

Table 12. Minimum vSphere uplink failover network configuration (if all NICs have a different bandwidth with no physical NIC failover)

Network component	VLAN[1]	Active uplink	Standby uplink
ESXi Admin + VMware HA network	100	Nic0-Uplink1	Nic0-Uplink2
SAP Application Server Network	101	Nic0-Uplink2	Nic0-Uplink1
vMotion network	200	Nic1-Uplink1	Nic1-Uplink2
vSAN network	201	Nic1-Uplink2	Nic1-Uplink1
HANA backup Network (>=10 GbE)	301	Nic2 port 0-uplink3	Nic2 port 1-uplink4

[1] VLAN id example, final VLAN numbers are determined by the network administrator.

We used different VLANs to separate the vSphere operational traffic (for example, vMotion and vSAN) from the SAP and user-specific network traffic.

vSphere and vSAN requirements for passthrough and RAID

The server storage and RAID controller normally support both pass-through and RAID. However, in vSAN environments you should only use controllers that support pass-through mode because you must deactivate RAID and use only pass-through mode.

vSAN OSA disk group configuration

Depending on the hardware configuration—such as the number of CPUs and flash devices—you should configure **a minimum of two disk groups per host**. Using more disk groups improves storage performance. The current maximum is **five disk groups per vSAN cluster node**.

For the **cache tier**, use **low-latency, high-bandwidth SAS SSDs** or **Performance Class F or better NVMe devices**. The recommended minimum capacity is **600 GiB**. Refer to the vSAN Flash Read Cache Compatibility Guide for supported devices.

In our **2- and 4-socket vSAN OSA reference configurations**, we used **four disk groups per host** to fully leverage the server hardware and support **up to four SAP HANA VMs per node**. The number of disk groups should be selected based on the number of SAP HANA VMs you plan to run on a single node.

In some test scenarios, a **2-disk group configuration**—with **NVMe for caching** and **four capacity disks** per group—may be sufficient to meet SAP HANA KPIs for **one VM per CPU socket** (two VMs per host). However, **SSDs used as cache devices** were not sufficient to meet SAP HANA performance requirements in these cases and a 4-disk group configuration was required.

Figures 13 and 14 show example configurations that include SSDs as device options. However, since NVMe devices offer significantly better performance and are now available at comparable price points, it is strongly recommended to use NVMe for both cache and capacity tiers whenever possible.

Recommended: If you need to deploy 4 SAP HANA VMs on a server using OSA, we recommend using 4 disk groups. The better, easier to manage, and more performant solution is to use ESA.

In addition, the PCIe lanes of the PCIe slots of the NVMe should be distributed evenly across all the CPUs. If multiple NVMe devices are connected by a single CPU link, make sure you use separate ports. As noted, the hardware vendors decide which configurations they certify.

Figure 14 shows the 4-disk group configuration we used for our 3-node Sapphire Rapids 4-Socket vSAN cluster. We first used SSD devices and then replaced the SSDs with Performance Class F NVMe devices for cache and data. We validated 4 SAP HANA VMs per server with this the SSD and NVME vSAN configuration.

Figure 13. 4-disk group configuration example of the SAP HANA HCI test configuration

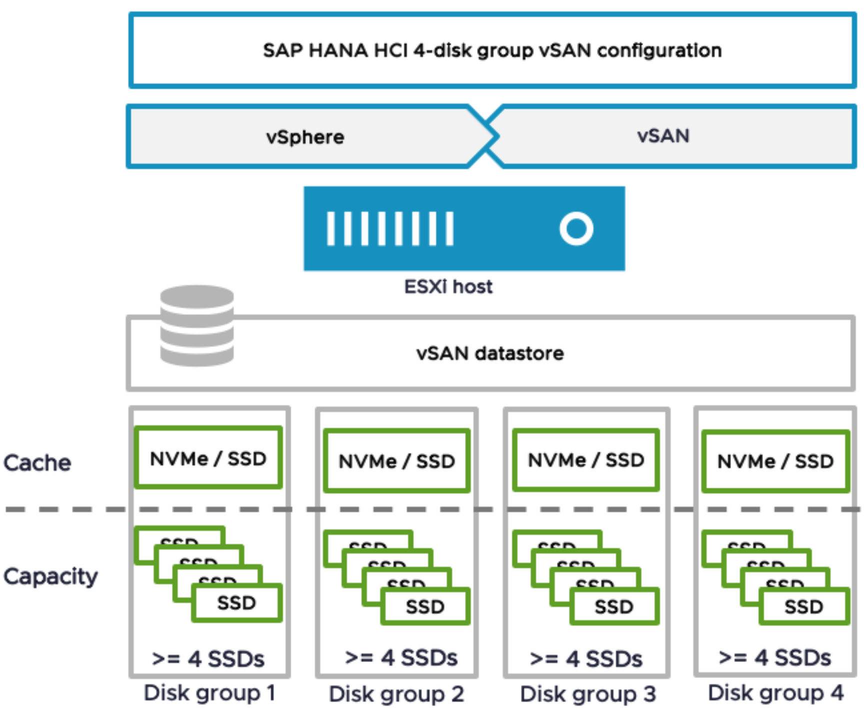

Replacing the SAS SSDs for caching with high-bandwidth, low-latency NVMe devices helped us reduce the number of vSAN disk groups from four to two while still maintaining the SAP HANA storage KPIs for two SAP HANA VMs per server. See figure 14 for the example.

Note: If 4 or more SAP HANA VMs run on a vSAN node, we recommend using at least a 4-disk group configuration with NVME devices instead of SSDs.

Figure 14. 2-disk group configuration for two SAP HANA VMs

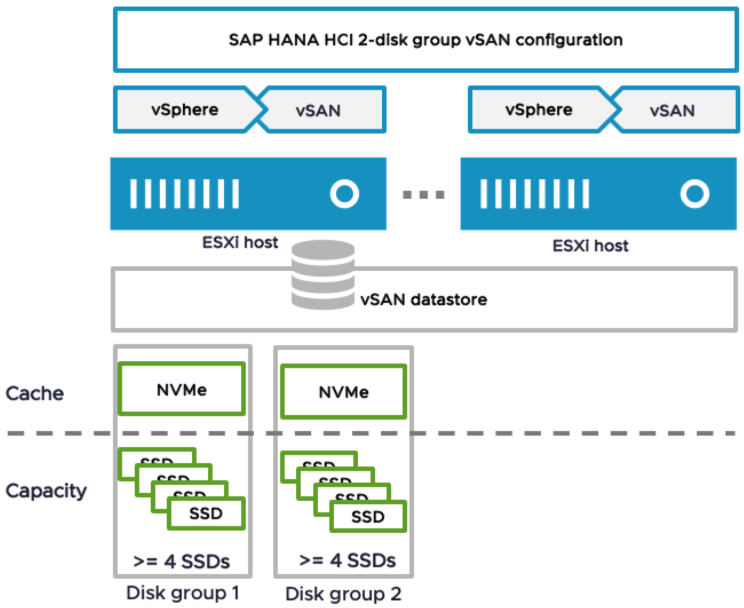

Note: The number of capacity drives per disk group and the number of disk groups per host impact the maximum number of SAP HANA VMs per server.

vSAN ESA storage pool configuration

Unlike with OSA, an ESA-based vSAN configuration does not have disk groups with dedicated cache and data tiers and instead pools the server internal NVMe devices into a single tier. Each host that contributes storage contains a single storage pool of one or more flash devices. Each flash device provides caching and capacity to the cluster.

Figure 15. vSAN ESA NVME device storage pool configuration

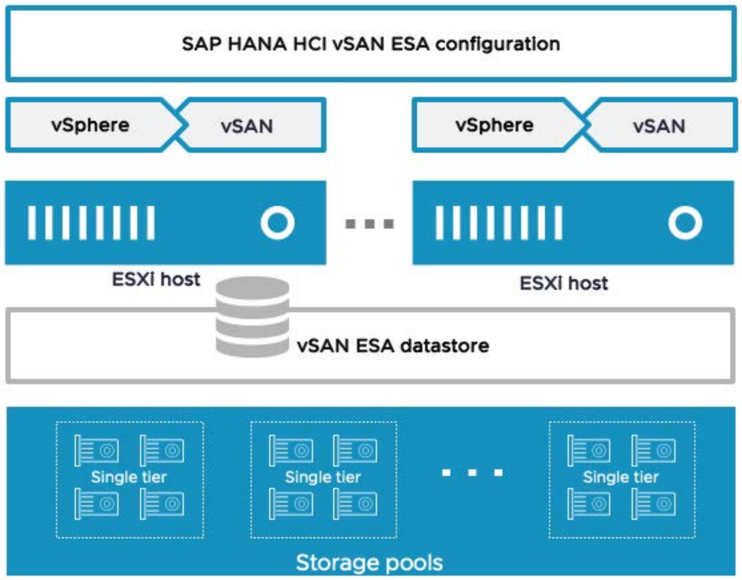

vSAN storage configuration

We recommend configuring all vSAN disk groups in OSA, or the storage pool in ESA, as a single, cluster-wide shared datastore. Specific vSAN storage policies should be applied based on the use case (e.g., OS, logs, or data) defining the storage characteristics for each VMDK and therefore VM disk.

In our tests, we used tailored storage policies for the operating system, SAP HANA bin, shared, log, and data volumes.

Consider modifying the storage policies and adapt them to your unique vSAN ReadyNode / planned SAP HANA HCI system hardware components.

vSAN storage policy

The SAP HANA HCI solutions are very specific configurations and are optimized for the high-performance demands of an SAP HANA database. Just like an SAP HANA appliance, the storage subsystem is optimized to meet the SAP HANA KPIs for reliability, consistency, and performance. SAP validated if our example test configurations could maintain the SAP HANA TDI storage KPIs.

Review the "What are vSAN Policies" document for more details and available options.

You should select and apply a vSAN storage policy to the SAP HANA OS, bin, shared, log, and data disks according to the vSAN configuration. The vSAN policy is one of the parameters that controls how many SAP HANA VMs a server can support.

Tables 13 and 14 show the vSAN storage policy we used for our testing. You should apply the storage policy to both the data and log disks of each SAP HANA database VM. We selected RAID-1 (mirroring), which provides the highest performance, but consumes double the capacity.

For volumes that don't require high performance, such as SAP HANA bin and shared, you can use the vSAN default storage policy or use a custom RAID 5 or 6 storage policy to lower their storage capacity needs.

For example, a VM with an FTT of 1 protected with RAID-1 requires twice the virtual disk size, but with RAID 5 it requires just 1.33 times the virtual disk size. With vSAN OSA, RAID-5 requires a minimum of four nodes per vSAN cluster. With vSAN ESA, RAID-5 requires a minimum of just three nodes per vSAN cluster. To use RAID 5, set Failure tolerance method to RAID-5/6 (Erasure Coding) and Failures to tolerate to 1. To use RAID 6, set Failure tolerance method to RAID-5/6 (Erasure Coding) and Failures to tolerate to 2. RAID 5 or RAID 6 erasure coding does not support a Failures to tolerate value of 3.

To use RAID 1, set Failure tolerance method to RAID-1 (Mirroring). In vSAN OSA RAID 1 provides slightly better performance than RAID 5 or 6. For example, a cluster resynchronization takes less time to complete with RAID 1. In vSAN ESA RAID-5 is actually faster than RAID-1, and RAID-6 (tolerating up to two failures; FTT=2) is faster than FTT=2 RAID-1.

The number of stripes configured should equal the number of capacity devices contained in a host. Note that the maximum number of stripes you can configure in a storage policy is 12.

> **Note:** VMware tested and used only the listed vSAN storage policies during the vSAN SAP HANA validation. Storage compression is only supported with ESA. Encryption is yet not supported.

Table 13. vSAN storage policy settings for data and logs for a configuration with 12 or more capacity disks

Storage Capability	Setting (OSA)	Setting (ESA)
Failure Tolerance Method (FTM)	RAID 1 (mirroring)	RAID 5
Failures to tolerate (FTT)	1	1
Minimum hosts required	3	3
Disk Groups	4 (for 2S /4S)	NA
Number of disk stripes per object	12	1
Flash read cache reservation	0%	0%
Object space reservation	100%	100%
Disable object checksum	No	No
Compression and Deduplication	No	No
Compression only	No (OSA)	Yes

Table 14 shows the vSAN policy for a 2-disk group vSAN OSA cluster configuration.

Table 14. vSAN storage policy settings for data and logs for a configuration with 8 capacity disks

Storage Capability	Setting (OSA)	Setting (ESA)
Failure Tolerance Method (FTM)	RAID 1 (mirroring)	RAID 5
Failures to tolerate (FTT)	1	1
Minimum hosts required	3	3
Number of disk stripes per object	8	1
Flash read cache reservation	0%	0%
Object space reservation	100%	100%
Disable object checksum	No	No
Compression and Deduplication	No	No
Compression only	No (OSA)	Yes (ESA)

Table 15 shows the vSAN policy for less performance-critical disks like OS or SAP HANA bin and shared.

> **Note:** With vSAN OSA, RAID-5 requires a minimum of four nodes per vSAN cluster. With vSAN ESA, RAID-5 requires a minimum of just three nodes per vSAN cluster. Both OSA and ESA require a minimum of six vSAN nodes for RAID 6 (FTT=2).

Table 15. vSAN storage setting for OS or other less performance-critical volumes like SAP HANA bin or shared

Storage Capability	Setting
Failure tolerance method (FTM)	RAID 5/6 (Erasure Coding)
Failures to tolerate (FTT)	1 (2 when RAID 6)
Minimum hosts required	4 (OSA), 3 (ESA)
Number of disk stripes per object	1
Flash read cache reservation	0%
Object space reservation	0%
Disable object checksum	No
Compression and Deduplication	No
Compression only	No (OSA) / Yes (ESA)

SAP HANA HCI storage planning and configuration

When planning a virtualized SAP HANA environment running on vSAN, you need to understand the storage requirements of such an environment and the applications running on it. You must define the storage capacity and IOPS needs of the planned workload.

Sizing a storage system for SAP HANA is very different from sizing a storage system for SAP classic applications. Unlike SAP classic applications, where you size the storage subsystem according to the IOPS requirements of a specific SAP workload, you define the SAP HANA storage according to the specified SAP HANA TDI storage KPIs. These data throughput and latency KPIs are fixed values per SAP HANA instance.

Depending on how many SAP HANA VMs a vendor supports on an SAP HANA HCI cluster, the vSAN datastore must maintain the performance KPIs and must be able to provide the storage capacity of all installed SAP HANA VMs on the vSAN cluster.

VMware has successfully tested 2- and 4-socket, 3-node and 4-node vSAN cluster, vSAN-based SAP HANA HCI configurations with eight SAP HANA VMs (four SAP HANA VMs per active node). These configurations represent up to two SAP HANA VMs per CPU socket, leaving the third or fourth node in the cluster used as an HA node. This leaves enough capacity available to support either a single node or two nodes' failures at any time.

Use the SAP HANA Hardware and Cloud Measurement Tools (HCMT) to check if the vSAN storage performance meets the SAP HANA requirements. You can download the tool and the documentation from SAP with a valid SAP S-User account. See SAP note 2493172.

SAP HANA storage capacity calculation and configuration

All SAP HANA instances have database log, data, root, local SAP binaries, and shared SAP volumes. The storage capacity sizing calculation of these volumes is based on the overall amount of memory needed by SAP HANA's in-memory database.

SAP has set very strict performance KPIs for configuring storage subsystems. These KPIs will determine the sizing for your storage subsystem. You'll need to keep in mind what hardware components are available from your HCI vendors and size your storage subsystem to meet the necessary I/O performance and latency defined by the SAP KPIs.

The capacity of the overall vSAN datastore must be large enough to store the SAP HANA data, logs, and binaries by providing the minimum FTT=1 storage policy. SAP has published several SAP HANA specific architecture and sizing guidelines, such as "[SAP HANA Storage Requirements."](#) Figure 16 and Table 16 summarize the essence of this guide, which you can use as a good starting point for planning the storage capacity needs of an SAP HANA system. Figure 16 shows the typical disk layout of an SAP HANA system, and the volumes needed. The volumes shown in the document should correspond with VMDK files to ensure flexibility for changing storage needs (which will typically be a need for increased storage capacity).

Figure 16. Suggested storage Layout of an SAP HANA System (figure from SAP SE, modified by VMware)

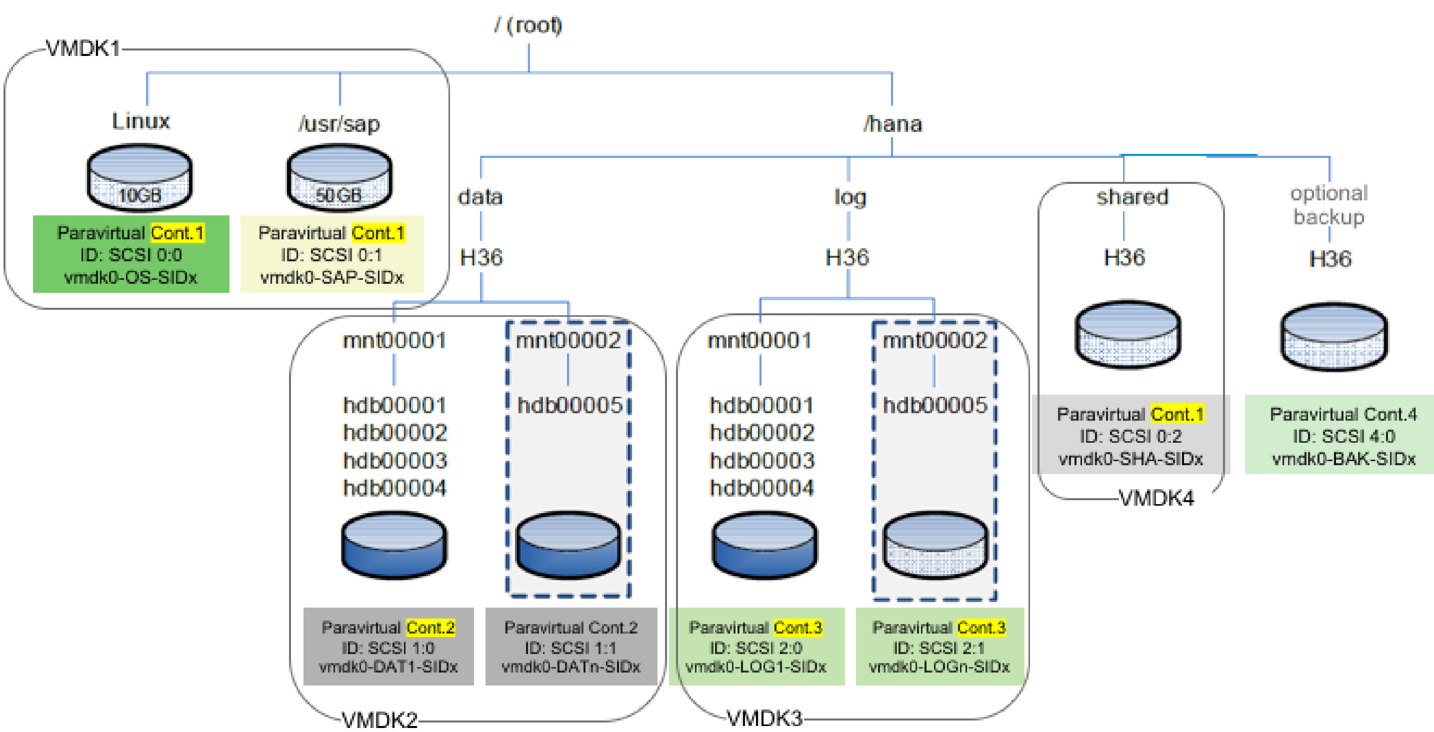

Table 16 determines the size of the different SAP HANA volumes. It's based on the [SAP HANA Storage Requirements](#) document. A single PVSCSI controller can connect and serve some volumes, such as those for OS and usr/sap. Others, such as the log and data volumes, require dedicated PVSCSI or NVME controllers to ensure high I/O bandwidth and low latency.

Table 16. Storage layout of an SAP HANA system

Volume	VMDK	Controller	VMDK name	SCSI ID	Sizes (SAP HANA storage requirements)
/(root)	✓	PVSCSI Contr. 1	vmdk01-OS-SIDx	SCSI 0:0	Minimum 40 GiB for OS
usr/sap	✓	PVSCSI Contr. 1	vmdk01-SAP-SIDx	SCSI 0:1	Minimum 50 GiB for SAP binaries
hana/shared (Scale-Up)	✓	PVSCSI / NVME Contr. 1	vmdk02-SHA-SIDx	SCSI 0:2	Minimum 1 x RAM max. 1 TiB
hana/shared (Scale-Out)	✓	PVSCSI / NVME Contr. 1	vmdk02-SHA-SIDx	SCSI 0:2	Minimum 1 x RAM per 4 worker nodes (e.g. 5+1 nodes, 1 TiB per node = 2x1 TiB = 2 TiB)
hana/data	✓	PVSCSI / NVME Contr. 2	vmdk03-DAT1-SIDx (vmdk03-DAT2-SIDx, vmdk03-DAT3-SIDx)	SCSI 1:0 (SCSI 1:1, SCSI 1:2)	Minimum 1.5 x RAM Note: If you use multiple VMDKs, you can use Linux LVM to build one large data disk.
hana/log	✓	PVSCSI / NVME Contr. 3	vmdk04-LOG1-SIDx (vmdk04-LOG2-SIDx)	SCSI 2:0 (SCSI 2:1)	[systems <= 512 GiB] log volume (min) = 0.5 x RAM [systems >= 512 GiB] log volume (min) = 512 GiB
hana/backup (optional)	✓	PVSCSI / NVME Contr. 4	vmdk05-BAK-SIDx	SCSI 3:1	Size backups >= Size of HANA data + size of redo log. The default path for backup is /hana/shared. Change the default value when using a dedicated backup volume.

Figure 17 shows the default SAP HANA storage layout, as documented in the SAP HANA storage white paper. It depicts the dedicated VMDKs per OS, HANA shared, log, and data volume.

Figure 17. Dedicated VMDK files per SAP HANA system

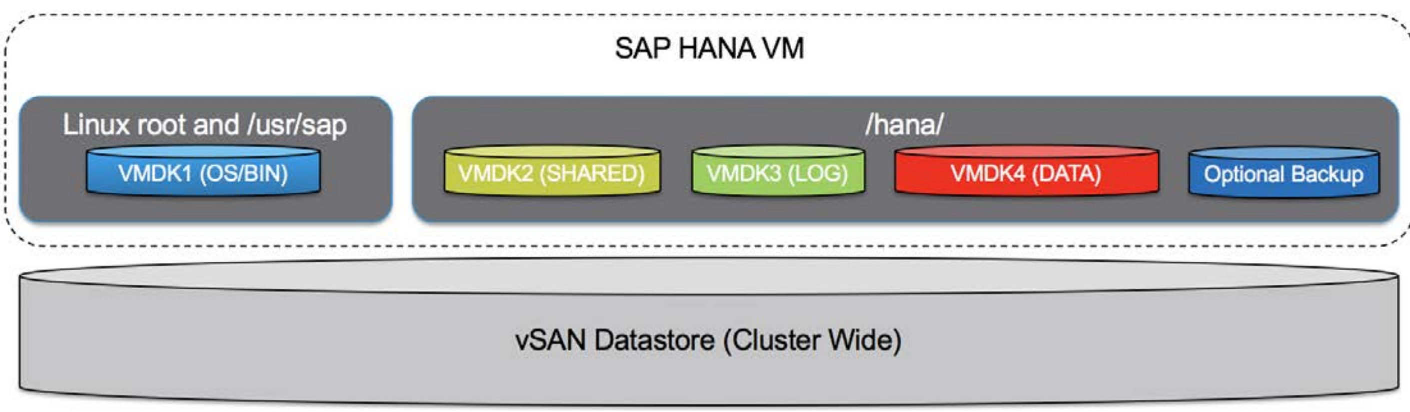

A "pure" VMDK solution provides the easiest way to get started and is enough to achieve the SAP HANA HCI performance and scalability KPIs.

Use of Linux Logical Volume Manager (LVM)

We didn't use Linux LVM in our tests, but you can use it if you'd like. Using Linux LVM can help leverage more vSAN resources, balance the VMDK component distribution on vSAN, and align it to storage layouts you may have from older bare metal installations. On the other hand, using LVM adds configuration complexity. In the end, the HCI solution vendor determines whether or not LVM is supported in their unique HCI solution.

Figure 18 shows an example of how you can use LVM for the SAP HANA log and data volumes and eight VMDKs for the data partition and co-deployed log partition.

Note that the number of VMDK disks used for LVM should not exceed the number of capacity disks used for vSAN in a host. For example, if there are six capacity disks in a server, up to six VMDK disks for LVM would be ideal.

Figure 18. Usage of LVM managed volumes

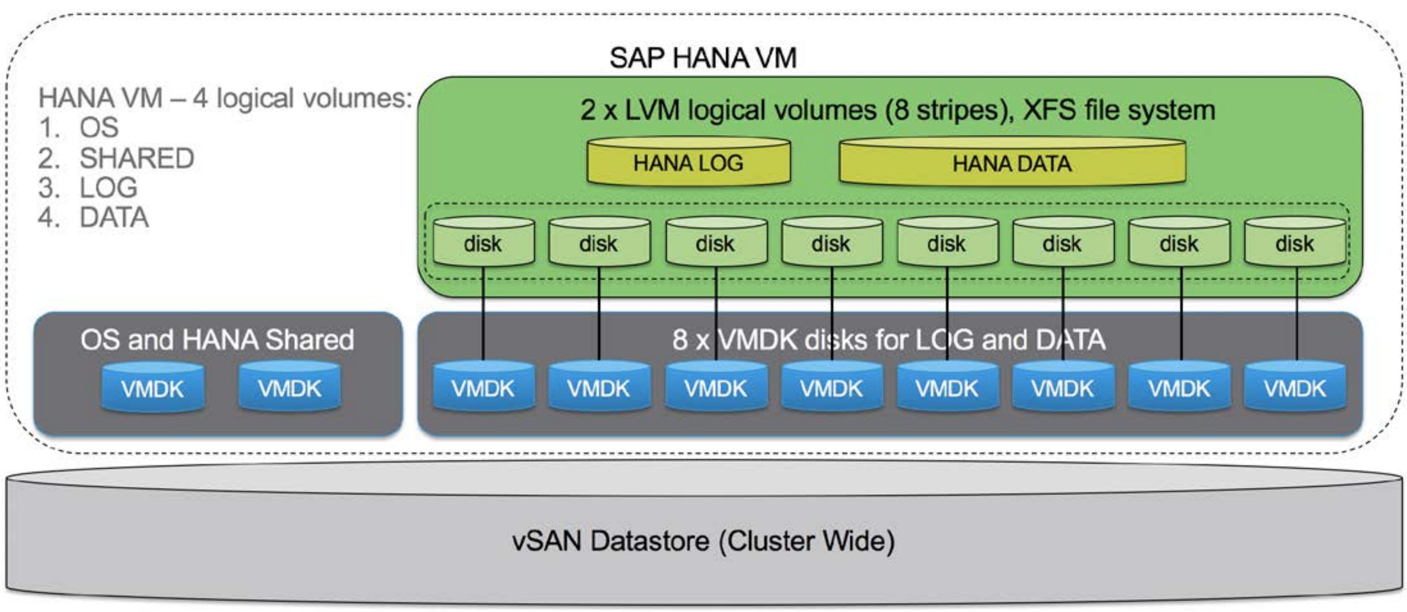

Calculating vSAN storage capacity

To calculate the storage capacity per SAP HANA VM, summarize the sizes of all specific and unique SAP HANA volumes as outlined in Figure 16 and Table 16.

To determine the overall vSAN datastore capacity required (the data tier in OSA; the sum of all disks in ESA), multiply the above SAP HANA volume requirements by the number of VMs that can run on the SAP HANA HCI cluster. Consider the impact of the storage policies selected, such as FTT value and RAID level.

Simplified calculation:

```
vSAN datastore capacity = total HANA VMs × (OS + USR/SAP + SHARED + DATA + LOG) × 2 (FTT=1, RAID-1) + 30% slack space
```

Example:

Suppose we have a cluster made up of three servers with a total of 18 TiB RAM and eight SAP HANA VMs. We use one server for HA, and we need to store 12 TiB of SAP HANA data.

```
8 × (10 GiB + 60 GiB + 1 TiB + 1.5 x 12 TiB / 8 VMs + 512 GiB) × 1 × 2 × 1.3 = 8 × 3910 GiB × 2 × 1.3 = ~78,213 TiB
```

Based on our previous example with 16 x 1.6 TiB NVMe devices in a 4-host cluster, we have over 100 TiB of vSAN raw datastore capacity available when we select ESA as the storage architecture. OSA would require a cache, reducing the number of usable data disks to 12. The OSA based vSAN datastore capacity is around 77 TiB, which is barely enough for the up to eight SAP HANA VMs with a total of up to 12 TiB of SAP HANA system RAM capacity.

vSAN detection script

SAP requested a way to detect whether an SAP HANA VM runs on an SAP HANA HCI solution. VMware and SAP worked together to develop a script that allows SAP to detect at runtime if an SAP HANA VM is running on a vSAN datastore. The download location for this script is documented in the SAP HANA vSAN support note 3406060. The SAP HANA HCI solution vendor must install the script on their SAP HANA/VMware HCI systems and is responsible that this script remains running over the lifetime of the HCI solution. The vSAN servers hosting the SAP HANA HCI solution will run this Python script every 5 minutes. It will add the following information to the VM configuration file:

- guestinfo.vsan.enabled (true or false)
- guestinfo.SDS.solution (the ESXi version, which is also the vSAN release version)
- guestinfo.vm_on_vsan (true or false)

The SAP CIM provider and SAP support tools (like SAPsysinfo) will parse these parameters in the script, thus detecting whether an SAP HANA HCI on VMware vSAN solution is used for this specific VM. You can find the monitored SAP/VMware key metrics in the OSS Note 2266266—Key Monitoring Metrics for SAP on VMware vSphere.

Installing the script

Prerequisites

1. Enable SSH for the duration of the installation. (Deactivate it after the script is installed.)
2. To simplify the installation process, we recommend you configure all the ESXi hosts with an identical superuser password.
3. Download the script from Github.com (link).

Installation and execution of the health check installation script

1. Enable SSH on the vSphere ESXi hosts.
2. Place the installation script and setInfo.vib into the same directory on any Linux or MacOS system.
3. Run ./install_vib_script_remotely.sh and type the following information as prompted:

```
root@Centos [ ~/hana_health_check ]# ./install_vib_script_remotely.sh -h
Please specify one of the following parameters:
-i or --install: install the vib on specified hosts
-r or --remove: remove the vib on specified hosts
-u or --update: update the vib on specified hosts
-h or --help: display this message
```

4. To install the vib on the specified host:

```
[w3-dbc302:hana_vib]$./install_vib_script_remotely.sh -i
Please specify the hosts (separate by whitespace):
10.xx.xx.xx
login as: root
```

```
Password:

Validating connectivity to 10.xx.xx.xx
10.xx.xx.xx connectivity validated

Uploading setInfo_v1.3.vib to 10.xx.xx.xx
Successfully uploaded setInfo_v1.3.vib to 10.xx.xx.xx:/tmp

Installing VIB in 10.xx.xx.xx
Host acceptance level changed to 'CommunitySupported'.
Installation Result:
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed: VMware_bootbank_setInfo_1.3.0-6.7.0
   VIBs Removed:
   VIBs Skipped:
Successfully installed VIB in 10.xx.xx.xx
```

5. Disable SSH on the vSphere ESXi hosts.

Functions and scenarios of the script

You can increase how often the script is run from its default of every 5 minutes by changing the value of RUN_PYC_INTERVAL at the beginning of the bash script.

For all the VMs on the host, the script gets the following information from each VM:

- vsan.enabled = true|false
- sds.solution = ESXi/vSAN_Version
- vm_on_vsan = true|false

Testing the script

To test the script:

1. Add a Linux VM to one of your ESXi hosts. SAP HANA must support the Linux version.
2. Ensure that the latest Open VM guest tools are installed on this test VM.
3. Enable host and VM monitor data exchange, as documented in SAP note 1606643.
4. Download and copy the latest sapsysinfo script (>= revision 114), see SAP note 618104 for details and run it on your test VM.
5. Example output vSAN section (search for vSAN in the document):

```
-------------------------------------------------------------------------------
VMware vSAN info:

guestinfo.vsan.enabled =
guestinfo.SDS.solution =
guestinfo.vm_on_vsan =

-------------------------------------------------------------------------------
Linux system information:

-------------------------------------------------------------------------------
```

If your output isn't the same as the above, check the prerequisites and ensure that the hypervisor used is vSphere and the SDS solution is vSAN.

If a VM is moved to vSAN storage, a newly created vmdk is on the vSAN, or a VM is moved from vSAN to another datastore, guestinfo will be updated within 5 mins.

If the VM is Storage vMotioned to another host that doesn't have the script running, manually delete the three variables from the VM configuration file to ensure consistency. You can check inside vCenter if these variables are set for a VM. See Figure 19 for details.

Figure 19. SAP HCI vSAN specific configuration parameters

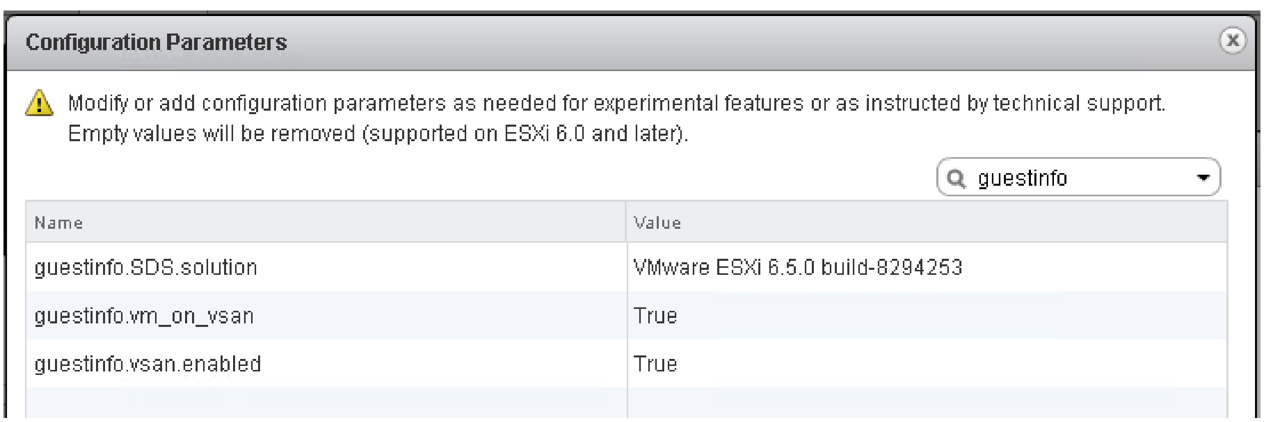

Updating the script

To update the vib to a higher version, replace the vib file within your Linux/MacOS directory and run the following command:

```
[w3-dbc302:hana_vib]$./install_vib_script_remotely.sh -u
Please specify the hosts(Separate by whitespace):
10.xx.xx.xx
login as: root
Password:
```

```
Validating connectivity to 10.xx.xx.xx
10.xx.xx.xx connectivity validated

Uploading setInfo_v1.3.5.vib to 10.xx.xx.xx
Successfully uploaded setInfo_v1.3.5.vib to 10.xx.xx.xx:/tmp

Updating VIB in 10.xx.xx.xx
Installation Result:
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed: VMware_bootbank_setInfo_1.3.5-6.7.0
   VIBs Removed: VMware_bootbank_setInfo_1.3.0-6.7.0
   VIBs Skipped:
Successfully update VIB in 10.xx.xx.xx
```

Uninstalling the script

To uninstall the script from ESXi, run:

```
[w3-dbc302:hana_vib]$./install_vib_script_remotely.sh -r
Please specify the hosts(Separate by whitespace):
10.xx.xx.xx
login as: root
Password:

Validating connectivity to 10.xx.xx.xx
10.xx.xx.xx connectivity validated

Removing VIB in 10.xx.xx.xx
Removal Result:
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed:
   VIBs Removed: VMware_bootbank_setInfo_1.3.5-6.7.0
   VIBs Skipped:
Successfully cleared up VMs info, removed vib from 10.xx.xx.xx
```

Monitoring and verifying an SAP HANA installation

SAP notes for monitoring data growth/CPU utilization and verifying configuration

- SAP note 1698281 provides information about how you can monitor data growth and utilization of actual memory. With this, it's also possible to detect and diagnose memory leaks during operation.
- SAP note 1969700 covers all the major HANA configuration checks and presents a tabular output with changed configurations. Collect SQL statements to check and identify parameters that are configured and conflict with the SAP-recommended configuration parameters.

NUMA tool: VMware NUMA Observer

While you might configure large critical VMs with affinities to unique logical cores or NUMA nodes, maintenance, and HA events can change this unique mapping. An HA event would migrate VMs to other hosts with spare capacity and those hosts may already be running VMs affined to the same cores or sockets. This results in multiple VMs constrained/scheduled to the same set of logical cores. These overlapping affinities may result in a CPU contention and/or non-local allocation of memory.

To check if the initial configuration is correct or detect misalignments, you can download and use the NUMA Observer fling (downloads a .zip file).

NUMA Observer scans your VM inventory, identifies VMs with overlapping core/NUMA affinities, and generates alerts. Additionally, the fling also collects statistics on remote memory usage and CPU starvation of critical VMs and raises alerts; see Figures 19 and 20 for examples.

Figure 19. VMware NUMA Observer – VM Core Overlap Graph

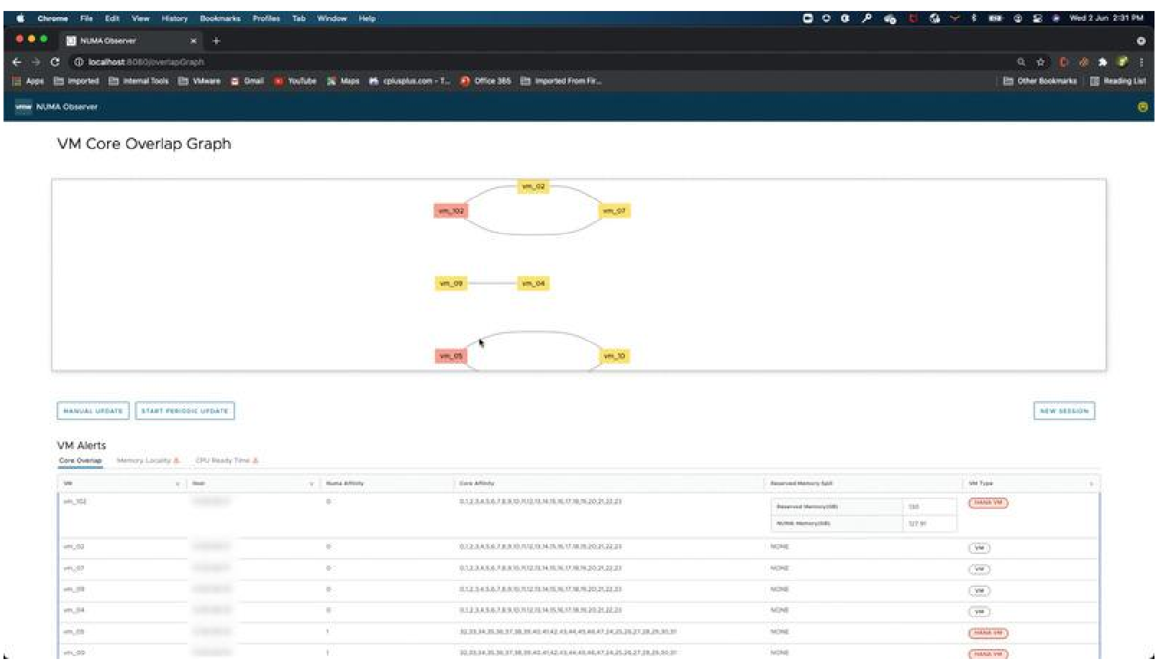

Figure 20. VMware NUMA Observer – VM Alerts

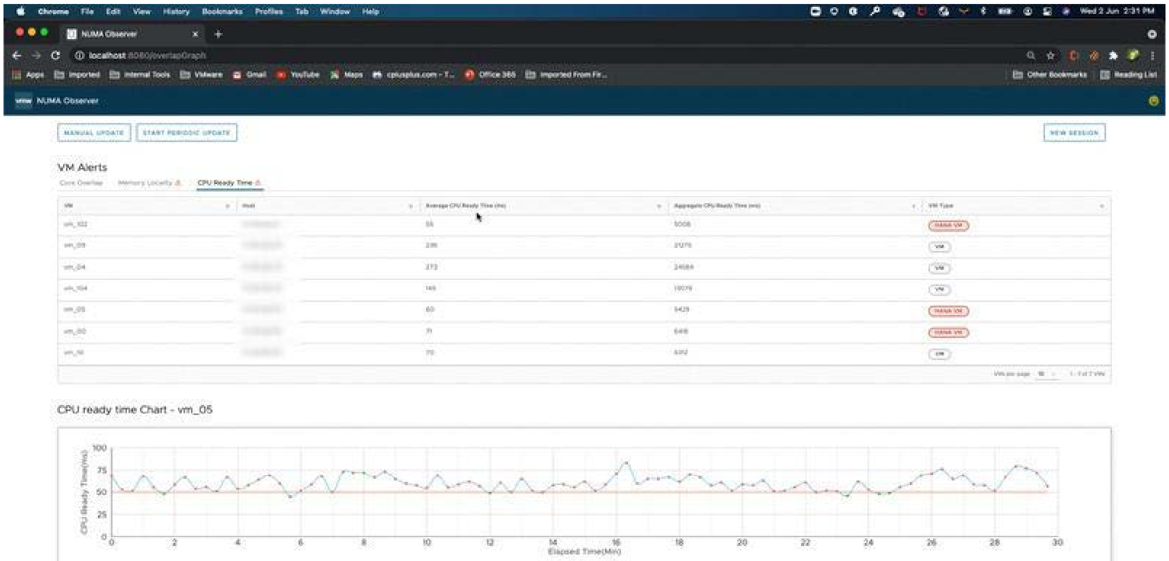

SAP HANA support and SAP HANA HCI vendor support process

Support status

SAP HANA on vSAN is supported only as part of a certified SAP HANA HCI solution, as outlined in this document and in SAP Notes 2718982 and 3406060. The SAP HANA validation for vSAN and SAP ICC certification goes beyond traditional TDI storage certification, as it includes the full solution stack.

The SAP HANA HCI solution vendor serves as the primary support contact for customers, with escalation paths to VMware and SAP. This support model is comparable to the SAP HANA Appliance model, with the key difference being that the HCI vendor, not SAP, leads support escalations.

The HCI solution vendor is responsible for defining the usable storage capacity and performance available for SAP HANA VMs, as well as the supported number of SAP HANA VMs per host and the maximum number of vSAN cluster nodes.

For a list of supported SAP HANA HCI solutions based on VMware products, refer to the SAP HANA Certified HCI Directory.

vSAN Support Insight

vSAN Support Insight is a next-generation platform for analytics of health, configuration, and performance telemetry. Its purpose is to help vSAN users maintain reliable and consistent computing, storage, and network environment; it can be used in addition to the HCI vendor specific support contract.

For more details, see Broadcom KB 327053 - vSAN Health Service - Online Health - vSAN Support Insight.

Support process

If you need to escalate an issue with your SAP HANA HCI solution, work directly with your HCI vendor and follow the defined and agreed-upon support process, which typically begins by opening a support ticket through the SAP support system.

The SAP HANA HCI vendors and Broadcom are integrated into the SAP support organization, enabling non-SAP support engineers to collaborate directly with SAP, SAP customers, and other SAP software partners (e.g., SUSE) to resolve escalated issues.

Before opening a support ticket with Broadcom, we recommend first submitting a support request to your SAP HANA HCI vendor or through the SAP support portal. This ensures the HCI vendor's specialists can address the issue initially and escalate it to Broadcom (for Broadcom product issues) or to SAP (for SAP HANA-related issues) as appropriate.

Conclusion

With VMware Cloud Foundation (VCF) and vSAN in specific, VMware empowers hardware and OEM partners to deliver fully integrated, SAP HANA-ready hyperconverged infrastructure (HCI) solutions. VCF combines compute, memory, networking, and software-defined storage into a unified, flexible, and scalable platform built entirely on industry-standard x86 servers. This drastically transforms the economics and deployment models for SAP HANA infrastructure.

VCF and vSAN enable enterprises to move beyond static, hardware-bound SAP HANA deployments by offering a streamlined, automated, and resilient private cloud foundation. SAP hardware partners can now offer turnkey solutions that simplify procurement and accelerate time-to-value—allowing customers to deploy SAP HANA in days rather than weeks or months.

By leveraging VCF's software-defined infrastructure approach and tight integration with vSphere and vSAN, businesses gain not only performance and scalability, but also full control and sovereignty over their data and infrastructure. This positions VCF-based SAP HANA HCI solutions as a strategic enabler for modern SAP landscapes, supporting dynamic, mission-critical workloads while reducing complexity and cost.

Appendix

Optimizing SAP HANA on vSAN - Configuration parameter List

VMware vSphere supports running either a single large SAP HANA virtual machine (VM) or multiple smaller VMs on a single host. This section outlines how to optimally configure vSphere and vSAN within a VMware VCF environment for SAP HANA. These recommendations apply to SAP HANA VMs running on both vSAN and external storage; vSAN-specific settings are provided in a separate table.

The listed parameters include recommended settings for the physical server (BIOS), ESXi host, vSAN, VM, and Linux guest OS. These configurations are required to ensure optimal operational readiness and stable performance for SAP HANA on vSphere.

The parameters described here are baseline best practices that **should always be applied** for SAP HANA on vSphere. Additional settings provided in the section Performance optimization for low-latency SAP HANA VMs should be used only in exceptional cases where the lowest possible latency is required.

In case of performance issues or support escalations, VMware or the HCI vendor's SAP support engineers will verify these configuration settings and may request their application if not already in place.

Table 17. Physical host BIOS parameter setting

Physical host BIOS parameter settings / best practices	Description
UEFI BIOS host	Use only UEFI BIOS as the standard BIOS version for the physical ESXi hosts. All SAP HANA appliance server configurations leverage UEFI as the standard BIOS. vSphere has fully supported UEFI since version 5.0.
Enable Intel VT technology	Enable all BIOS virtualization technology settings.
Configure RAM hemisphere mode	Distribute DIMM or PMem modules in to achieve the best performance (hemisphere mode) and use the fastest memory modules available for the selected memory size. Beware of the CPU-specific optimal memory configurations that depend on the available memory channels per CPU.
CPU – Populate all available CPU sockets, use a fully meshed QPI NUMA architecture	To avoid timer synchronization issues, use a multi-socket server that ensures NUMA node timer synchronization. NUMA systems that do not run synchronization will need to synchronize the timers in the hypervisor layer, which can impact performance. See the Timekeeping in VMware Virtual Machines information guide for reference. Select only SAP HANA CPUs supported by vSphere. Verify the support status with the SAP HANA on vSphere support notes. For a list of the relevant note, see the SAP Notes Related to VMware page.
Enable CPU Intel Turbo Boost	Allow Intel automatic CPU core overclocking technology (P-states).
Disable QPI power management	Do not allow static high power for QPI links.
Set HWPE	Set to the hardware vendor default.
Enable hyper-threading	Always enable hyper-threading on the ESXi host. This will double the logical CPU cores to allow ESXi to take advantage of more available CPU threads.

Physical host BIOS parameter settings / best practices	Description
Enable execute disable feature	Enable the Data Execution Prevention bit (NX-bit), required for vMotion.
Disable node interleaving	Disable node interleaving in BIOS.
Disable C1E Halt state	Disable enhanced C-states in BIOS.
Set power management to high performance	Do not use any power management features on the server, such as C-states. Configure in the BIOS static high performance.
Set correct PMem mode as specified by the hardware vendor for either App Direct or Memory mode	In deployments containing Pmem, follow the vendor documentation and enable PMem for use with ESXi. Note: Only App Direct and Memory mode are supported with production-level VMs. Memory mode is only supported with a ratio of 1:4. As of today, SAP provides only non-production workload support. VMware vSAN does not have support for App Direct mode as cache or as a capacity tier device of vSAN. However, vSAN will work with ESXi hosts equipped with Intel Optane PMem in App Direct mode and SAP HANA VMs can leverage PMem according to SAP note 2913410. Important: The vSphere HA restriction (as described in SAP note 2913410) applies and must be considered.
Enable SNC-2 for sub-NUMA VMs	If Sapphire Rapids 2-socket systems and "half-socket" / sub-NUMA VMs are required, enable SNC-2. Do not enable SNC-2 if this is not required or on ESXi hosts other than 2-socket Sapphire Rapids systems.
Disable all unused BIOS features	This includes video BIOS, video RAM cacheable, on-board audio, on-board modem, on-board serial ports, on-board parallel ports, on-board game ports, floppy drives, CD-ROMs, and USB ports.

Table 18. ESXi host parameter settings

ESXi host parameter settings / best practices	Description
Networking	Use virtual distributed switches to connect all hosts that work together. Define the port groups that are dedicated to SAP HANA, management, and vMotion traffic. Use at least a dedicated 10 GbE for vMotion and the SAP app server or replication networks. At least 25 GbE for vMotion for SAP HANA system with >= 2 TiB of RAM is recommended.
Storage configuration	When creating your storage disks for SAP HANA on the VM/OS level, ensure that you can maintain the SAP-specified TDI storage KPIs for data and log files. Use the storage layout as a template, as explained in this guide. Configure the following setting on the ESXi hosts. Note: For this change to take effect, the ESXi host must be rebooted. From the ESXi console, set the following parameter: `vsish -e set /config/Disk/intOpts/VSCSIPollPeriod 100`

ESXi host parameter settings / best practices	Description
SAP monitoring	Enable SAP monitoring on the host by setting `Misc.GuestLibAllowHostInfo = "1"` For more details, see SAP note 1409604. Without this parameter, no host performance-relevant data will be viewable inside an SAP monitoring-enabled VM.

Table 19. vSAN parameter settings / requirements

vSAN settings / best practices	Description
HW requirements	• A SAP HANA HCI node should be based upon a vSAN-certified ReadyNode. In addition, the selected HCI node must be an SAP HANA-supported server. • Only vSAN certified components can be used. • For new SAP HANA HCI systems only all-flash configurations are supported, NVMe devices are prefered. • SNC-2 is only supported with Intel Granite Rapids or newer CPUs.
vSAN Detection Script	Script to detect if an SAP HANA VM is deployed on a vSAN based datastore. The SAP HANA HCI solution vendor must ensure that the script is deployed and running on their SAP HANA/VMware HCI certified systems. The download location for this script is documented on the SAP HANA vSAN support note 3406060. The script will add the following information to the VM configuration file: • guestinfo.vsan.enabled (true or false) • guestinfo.SDS.solution (the ESXi version, which is also the vSAN release version) • guestinfo.vm_on_vsan (true or false)
Supported vSAN storage policies OSA	**Failure Tolerance Method (FTM):** RAID-1 for log and data, RAID-1 and RAID-5 for OS, HANA bin, shared, and backup **Minimum number of failures to tolerate (FTT):** 1 **Minimum hosts / fault domains required**: 3 (2 x FTT + 1) **Encryption services:** No encryption **Space efficiency:** Compression not supported with OSA **Number of disk stripes per object:** Between 8 and 12 for log and data with RAID 1. OS, HANA bin, shared, and backup use the default. **IOPS limit for object:** Set if needed, normally 0 for production SAP HANA. The minimum setting for log and data is 13500. **Object space reservation:** 100% (Thick provisioning) **Flash read cache reservation (%):** 0% (all-flash cluster) **Disable Object Checksum:** No **Force provisioning:** No Details: https://techdocs.broadcom.com/us/en/vmware-cis/vsan/vsan/8-0/vsan-administration/using-vsan-policies/about-vsan-policies.html

vSAN settings / best practices	Description
Supported vSAN storage policies ESA	**Failure Tolerance Method (FTM):** RAID-1 for log and data, RAID-1, RAID-5, and RAID-6 for OS, HANA bin, shared, or backup **The minimum number of failures to tolerate (FTT):** 1 **Minimum hosts / fault domains required**: 3 (2 x FTT + 1) **Encryption services:** No encryption **Space efficiency:** Compression supported with ESA **Number of disk stripes per object:** Between 8 and 12 for log and data when RAID 1, if RAID-5 is used then disk stripe = 1. OS, HANA bin, shared, and backup use the default. **IOPS limit for object:** Set if needed, normally 0 for production SAP HANA. The minimum setting for log and data is 13500. **Object space reservation:** 100% (Thick provisioning) **Flash read cache reservation (%):** 0% **Disable Object Checksum:** No **Force provisioning:** No Details: https://techdocs.broadcom.com/us/en/vmware-cis/vsan/vsan/8-0/vsan-administration/using-vsan-policies/about-vsan-policies.html

Table 20. SAP HANA VM parameter settings

SAP HANA VM parameter settings / best practices	Description
Tips for editing a *.vmx file	Review tips for editing a *.vmx file in Broadcom KB 1714.
UEFI BIOS guest	Required: Use UEFI BIOS as the standard BIOS version for ESXi hosts and VMs. Features such as Secure Boot are possible only with UEFI. You can configure this with the vSphere Client by choosing UEFI boot mode.
SAP monitoring	Enable SAP monitoring in the SAP HANA VM with the advanced VM configuration parameter: `tools.guestlib.enableHostInfo = "TRUE"`. For more details, see SAP note 1409604. Besides setting this parameter, the VMware guest tools must be installed. For details, see Broadcom KB 1014294.
vCPU hotplug	Ensure that vCPU hotplug is deactivated, otherwise vNUMA is disabled, which will reduce SAP HANA performance. For details, see Broadcom KB 2040375.
Memory reservations	Set fixed memory reservations for SAP HANA VMs. Do not overcommit memory resources. You must also reserve memory for the ESXi host. The amount you reserve depends on the amount of CPU socket memory for ESXi. Typical memory reservation for a host is between 32–64 GiB for a 2-socket server, 64–128 GiB for a 4-socket server, and 128–256 GiB for an 8-socket server. These are not absolute figures, as the memory needs of ESXi depend strongly on the actual hardware, ESXi and VM configuration, and enabled ESXi features, such as vSAN.
CPU	Do not overcommit CPU resources or configure dedicated CPU resources per SAP HANA VM. You can use hyper-threading when you configure the VM to gain additional performance. For CPU generations older than Cascade Lake, review KB 302535 and KB 317621 for details about the effects of hyper-threading in relation to the security vulnerability known as "L1 Terminal Fault - VMM" (aka Foreshadow). If you want to use hyper-threading, you must configure your VM with twice the cores per CPU socket (e.g., a 2-socket wide VM on a 60-core Sapphire Rapids system will require 240 vCPUs).
vNUMA nodes	SAP HANA on vSphere can be configured to leverage half-socket and full-socket VMs. A half-socket VM is configured with only half of the available physical cores of a CPU socket in the VM configuration section of the vSphere Client. The vNUMA nodes of the VM will always be >=1, depending on how many CPUs you have configured in total. If you need to access an additional NUMA node, use all CPU cores of this additional NUMA node, but use as few NUMA nodes as possible to optimize memory access.

SAP HANA VM parameter settings / best practices	Description
	In Sapphire Rapids systems half-Socket/sub-NUMA VMs require SNC-2. **Important:** As of today, only 2-socket SPR systems are supported with SNC-2. Once SNC-2 is enabled on a supported 2-socket SPR host, four sub-NUMA nodes will be available. In the case of SNC-2, all logical CPU cores of a sub-NUMA node should be assigned to such an SNC-2 VM. Additionally, the parameter `sched.nodeX.affinity="Y"` must be set to ensure that the VM runs on the correct sub-NUMA node and does not get migrated by the scheduler to another sub-NUMA node.
Align virtual CPU VM configuration to actual server hardware	**Example:** A half-socket VM running on a server with 28-core CPUs should be configured with 28 vCPUs to use 14 cores and 14 hyper-threads per half CPU socket. Similarly, a full-socket VM should be configured with 56 vCPUs to use all 28 cores and all 28 available hyper-threads per socket. In the case of a Sapphire Rapids CPU, SNC-2 is required for half-socket VMs, and an SNC-enabled host will provide two sub-NUMA nodes per CPU socket. In this scenario, all logical CPU cores of each sub-NUMA node must be utilized. For example, with a 60-core Sapphire Rapids CPU, when SNC is enabled, it becomes 30 cores per sub-NUMA node. Therefore, a half-socket VM will be configured with 60 vCPUs (including hyper-threads).
Define the NUMA memory segment size	The `numa.memory.gransize = "32768"` parameter helps to align the VM memory to the NUMA memory map.
Paravirtualized SCSI (PVSCSI) controller or NVME adapter for I/O devices	Use the dedicated SCSI controllers for OS, log, and data to separate disk I/O streams. Based on the used storage devices, select a PVSCSI controller (SAS/SATA disks and FC) or an NVMe adapter (NVMe devices). For details, see the SAP HANA storage and disk layout section.
VMDK disks	Use VMDK disks whenever possible to allow optimal operation via the vSphere stack. In-guest NFS mounted volumes for SAP HANA are supported as well.
Create dedicated datastores for SAP HANA data and log files	Ensure the storage configuration passes the SAP-defined storage KPIs for TDI storage. Use the SAP HANA hardware configuration check tool (HWCCT) to verify your storage configuration. For details, see SAP note 1943937.
Use eager zero thick virtual disks for data and logs	We recommend this setting as it avoids lazy zeroing (initial write penalty).
VMXNET3	Use paravirtual VMXNET3 virtual NICs for SAP HANA VMs. We recommend at least 3–4 different NICs inside a VM in the SAP HANA VM (app/management server network, backup network, and, if needed, HANA system replication network). Corresponding physical NICs inside the host are required.
Set `lat.Sensitivity = normal`	Use the vSphere Client to check the VM configuration and ensure that the value of Latency Sensitivity Settings is set to "normal." If you must change this setting, restart the VM. Do not change this setting to **high** or **low**. Change this setting only under the instruction of VMware support engineers.

<table>
<tr><th>SAP HANA VM parameter settings / best practices</th><th>Description</th></tr>
<tr><td>Configuring virtual topology when vSphere 8 and later is used.</td><td>vSphere ESXi 8.0 introduces an enhanced virtual topology feature that automatically selects optimal coresPerSocket values for VMs and optimal virtual L3 sizes.

The intelligent, adaptive NUMA scheduling and memory placement policies in ESXi 8.0 can manage all VMs transparently, so you don't need to deal with the complexity of balancing VMs between nodes by hand.

Manual controls are still available to override this default behavior, however, and SAP HANA production or performance-critical VMs should manually configure NUMA placement (through the <code>sched.nodeX.affinity="Y"</code> advanced option).

For detailed configuration steps refer to VMware vSphere 8.0 Virtual Topology.</td></tr>
<tr><td>NUMA action affinity</td><td>If the <code>sched.nodeX.affinity="Y"</code> setting cannot be used due to specific operational constrains, it's possible to avoid VM co-deployment on a CPU socket/NUMA node by deactivating the NUMA locality affinity, as described in Broadcom KB 314059.

Use the vSphere Client and add the following advanced VM parameter per VM setting:
<code>Numa.LocalityWeightActionAffinity="0"</code></td></tr>
<tr><td>Associate VMs with specified NUMA nodes to optimize NUMA memory locality

Note: While ESXi 8 provides an enhanced virtual topology feature that automatically selects optimal coresPerSocket values for VMs and optimal virtual L3 sizes to align it to the underlying NUMA topology of a host, we still recommend you associate a virtual NUMA node to a physical NUMA node. This is especially recommended for consolidation configurations, where more SAP VMs share a single host. When a host has only a single large VM these settings are not required.</td><td>For all SAP HANA VMs, we recommend you use affinity to associate the virtual NUMA node with a physical NUMA node. This provides direct control over how a VM is placed on a specific NUMA node. It also ensures that a virtual NUMA node does not get migrated, which could occur if there are idle NUMA nodes. While these migrations help to load balance how a server is used by VMs, for SAP HANA VMs this is mostly negative and could impact memory latency and therefore performance. Due to this, associating the physical NUMA nodes with a VM is required to constrain how the ESXi scheduler can schedule a VM's CPU and memory.

Use the vSphere Client or directly modify the .vmx file and add
<code>sched.nodeX.affinity="Y"</code>

Refer to the VMware article, "Associate Virtual Machines with Specified NUMA Nodes" for details.

Procedure:
– Browse to the VM in the vSphere Client.
– Right-click and select Edit Settings.
– Select the VM Options tab and expand Advanced.
– Under Configuration Parameters, click the Edit Configuration button.
– Click Add Row to add a new option.
– To specify a physical NUMA node for a specific virtual NUMA node on the VM: In the Name column, enter <code>sched.nodeX.affinity</code>, where X is the virtual NUMA node number. For example, <code>sched.node0.affinity</code> specifies the virtual NUMA node 0 on the VM.
– In the Value column, enter the NUMA node where the VM or the virtual NUMA node can be scheduled.
– For every virtual NUMA node, add a new row with a new <code>sched.nodeX.affinity</code> option.
– Click OK twice to close the Edit Settings dialog box.

Example for a VM with a single virtual NUMA node:
• <code>sched.node0.affinity ="2"</code> (this is physical NUMA node 2)</td></tr>
</table>

SAP HANA VM parameter settings / best practices	Description
	Example for a VM with two virtual NUMA nodes: ● `sched.node0.affinity ="0"` ● `sched.node1.affinity ="1"` Not supported: ● `sched.node0.affinity =`"**1**" ● `sched.node1.affinity =`"**2**"
Configure VMs to use hyper-threading with NUMA	For memory latency-sensitive workloads with low processor utilization, such as SAP HANA, or high interthread communication, we recommended using hyper-threading with fewer NUMA nodes instead of full physical cores spread over multiple NUMA nodes. Use hyper-threading and enforce NUMA node locality as described in Broadcom KB 326180. This parameter is only required when hyper-threading should be leveraged for a VM. Using hyper-threading can increase the compute throughput but may increase the latency of threads. **Note:** This parameter is only important for half-socket and multi-VM configurations that do not consume the full server, such as a 3-socket VM on a 4-socket server. Do not use it when a VM leverages all installed CPU sockets (e.g., 4-socket wide VM on a 4-socket host or an 8-socket VM on an 8-socket host). If a VM has more vCPUs configured than available physical cores, this parameter is configured automatically. Use the vSphere Client and add the following advanced VM parameter: `numa.vcpu.preferHT="TRUE"` (per VM setting) or as a global setting on the host: `Numa.PreferHT="1"` (host). **Note:** For CPU generations older than Cascade Lake, review KB 302535 and KB 317621 for details about the effects of hyper-threading in relation to the security vulnerability known as "L1 Terminal Fault - VMM" (aka Foreshadow).
Remove unused devices	Remove unused devices, such as floppy disks or CD-ROMs, to release resources and to mitigate possible errors.

Table 21. Linux operating system parameter settings

Linux OS parameter settings / best practices	Description
Linux version	VMware strongly recommends using only the SAP HANA-supported Linux and kernel versions. See SAP Note 2235581; for settings see SAP Note 2684254. In particular, generic SLES or RHEL versions are not supported. Only **SAP HANA SAPConf/SAPTune** optimized Linux OS configurations can be supported for SAP HANA.
Install the latest version of VMware Tools	VMware Tools is a suite of utilities that enhances the performance of the VM’s guest operating system and improves the VM management. See Broadcom KB 1014294 for details.

Linux OS parameter settings / best practices	Description
Configure NTP time server	Use the same external NTP server as configured for vSphere. For details, see SAP note 989963.

Performance optimization for low-latency SAP HANA VMs

Should the network latencies prove too high, we recommend further reducing them by applying the following parameters (alongside those described in the previous section).

Table 22. Additional network parameter settings to further reduce network latency

Additional Network Parameter settings	Description

Settings to lower the virtual VMXNET3 network latencies	Configure the following settings on the ESXi host. Note: For these changes to take effect, the ESXi host must be rebooted.
	
	Go to the ESXi console and set the following parameter:
	`vsish -e set /config/Net/intOpts/NetNetqRxQueueFeatPairEnable 0`
	
	Add the following advanced VMX configuration parameters to the VMX file, and reboot the VM after adding these parameters:
	
	`ethernetX.pnicFeatures = "4"`
	`ethernetX.ctxPerDev = "3"`
	
	Change the rx-usec, lro, and rx/tx values of the VMXNET3 driver, and of the NIC used for the SAP database to app server traffic. Change them from the default value of 250 to 75 (25 is the lowest usable setting).
	
	Procedure: Log on to the guest OS and use ethtool to change the following settings:
	
	`ethtool -C ethX rx-usec 75`
	`ethtool -K ethX lro off`
	`ethtool -G ethX rx 512 rx-mini 0 tx 512`
	
	Note: Replace ethX with the actual network interface name, such as eth0. To make these ethtool settings permanent, see SLES KB 000017259 or the RHEL ethtool document.

Optimize the application server network latency if required	Disable virtual interrupt coalescing for VMXNET3 virtual NICs that communicate with the app servers or front end to optimize network latency. Do not set this parameter for throughput-oriented networks, such as those used for vMotion or SAP HANA system replication. Use the advanced options in the vSphere Client or directly modify the .vmx file and add `ethernetX.coalescingScheme = "disable"`. X stands for your network card number. For details, see Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs.
Optional: Disable large receive offload (LRO) in the Linux guest OS to lower latency for client/application server-facing NIC adapter	This helps lower the network latency of client/application server-facing NIC adapters. Run: `ethtool -K ethX lro off` **Note:** Replace ethX with the actual network interface name, such as eth0. To make these ethtool settings permanent, see SLES KB 000017259 or the RHEL ethtool document. **Caution:** Do not disable LRO for throughput-oriented NIC adapters such as those used for backup, replication, or SAP HANA internode communication networks. This works only with VMs running Linux kernel 2.6.24 or later and using a VMXNET3 virtual network adapter. For additional details, see Broadcom KB 2055140.

When SAP HANA must perform as close to bare metal as possible with the lowest database access latency, you'll need to optimize the performance of SAP HANA VMs. We recommend sizing SAP HANA VMs with the fewest

NUMA nodes for low latency. Configure an additional NUMA node and its resources for SAP HANA VMs that need more CPU or RAM.

To achieve optimal performance for an SAP HANA VM, use the settings described in Table 22 in addition to the previously described settings. CPU scheduling and priority settings improve performance by reducing the amount of vCPU and vNUMA migration, while increasing the priority of the SAP HANA production VM.

CPU affinity settings

By specifying a CPU affinity setting for each VM, you can restrict the assignment of VMs to a subset of the available CPU cores (the specified affinity set).

Setting CPU affinities can decrease the CPU and memory latency by preventing the ESXi scheduler from migrating VM threads to other logical CPU cores. Setting CPU affinities is required when configuring SAP HANA half-socket VMs.

Before you use a CPU affinity, consider the following:

- ESXi systems perform automatic load balancing. Avoid manually setting VM affinity to allow the scheduler to optimally balance load across processors.
- An affinity setting can interfere with the ESXi host's ability to meet the reservation and shares specified for a VM.
- Because CPU admission control does not consider affinities, a VM with manual affinity settings might not always receive its full reservation. VMs that do not have manual affinity settings are not adversely affected by VMs with manual affinity settings.
- When you move a VM from one host to another, a manually set affinity might no longer apply because the new host might have a different number of processors.
- The NUMA scheduler might not be able to manage a VM that is already assigned to the processor using a manually set affinity.
- An affinity setting can affect the host's ability to schedule VMs on multicore or hyper-threaded processors, preventing the host from taking full advantage of the resources shared on those processors.

For more information about performance practices, see the vSphere Resource Management documentation and the VMware documentation on specifying NUMA controls.

Additional performance tuning settings for SAP HANA workloads

Caution: The following parameters are optional; they are needed only for systems that require the lowest CPU latency. Make sure you read the previous section, "CPU affinity settings" before changing them.

Important: Only set these advanced VM parameters for an SAP HANA workload that is the only workload running on a NUMA node/compute server. These settings may help a VM perform better, but at the expense of CPU time on the ESXi host.

Table 23. Tunings for very low-latency SAP HANA VMs

SAP HANA VM parameter settings	Description
Tips about how to edit the *.vmx file	Refer to Broadcom KB 1714.

`monitor.idleLoopSpinBeforeHalt = "true"`
`monitor.idleLoopMinSpinUS = "xx us"`
	Edit the .vmx file and add the following two advanced parameters:
	`monitor.idleLoopSpinBeforeHalt = "true"`
	`monitor.idleLoopMinSpinUS = "xx"`
	(Where "xx" could be "50", for example.)
	Both parameters must be configured to influence the descheduling time.
	Background: The guest OS issues a halt instruction, which stops (“deschedules”) the vCPU on the ESXi host. Keeping the VM spinning longer before halt reduces the number of inter-processor wake-up requests.

`monitor_control.halt_in_monitor = "TRUE"`	In the default configuration the idle state of guest halt instruction will be emulated without leaving the VM if a vCPU has an exclusive affinity. If the affinity is non-exclusive, the guest halt will be emulated in the vmkernel, which might result in having a vCPU descheduled from the physical CPU, and this can lead to longer latencies. We therefore recommend you set this parameter to "TRUE" to ensure that the halt instruction is emulated inside the VM and not in the vmkernel. Use the vSphere Client to add the following advanced VM parameter: `monitor_control.halt_in_monitor = "TRUE"`
`monitor_control.disable_pause_loop_exiting = "TRUE"`	This parameter prevents the VM from unnecessarily exiting to the hypervisor during a pause instruction. This is specific for Intel Skylake systems.

Table 24. Settings to improve NUMA alignment

SAP HANA VM parameter settings	Description
Configuring CPU affinity `sched.vcpuXX.affinity = "Yy-Zz"`	**Note:** Remove `sched.nodeX.affinity` or `numa.nodeAffinity` settings (if they were set) and if using this parameter setting. By specifying a CPU affinity setting for each VM, you can restrict the assignment of VMs to a subset of the available CPU cores (the specified affinity set). See "Scheduler operation when using the CPU Affinity" (Broadcom KB [334171](334171)) for details. This is required when configuring SAP HANA half-socket VMs or for very latency critical SAP HANA VMs. It's also required when the parameter `numa.slit.enable` is used. As with `sched.nodeX.affinity`, it's possible to decrease the CPU and memory latency by leveraging the `sched.vCPUxxx.affinity` VMX parameter to further limit the ESXi scheduler from migrating VM threads to other logical processors/CPU threads. In contrast to parameter `sched.nodeX.affinity`, it is possible to assign a vCPU to a specific physical CPU thread and is necessary, for example, when configuring half-socket SAP HANA VMs. Use the vSphere Client or directly modify the .vmx file (recommended way) and add `sched.vcpuXx.affinity = "Yy-Zz"` (for example: `sched.vcpu0.affinity = "0-55"`) for each virtual CPU you want to use. 1. Browse to the cluster in the vSphere Client. 2. Click the **Configure** tab and click **Settings**. 3. Under **VM Options**, click the **Edit** button. 4. Select the **VM Options** tab and expand **Advanced**. 5. Under **Configuration Parameters**, click the **Edit Configuration** button. 6. Click **Add Row** to add a new option. 7. In the **Name** column, enter `sched.vcpuXX.affinity` (XX stands for the actual vCPU to assign to a physical CPU thread). 8. In the **Value** column, enter the physical CPU threads where the vCPU can be scheduled. For example, enter 0-55 to constrain the VM resource scheduling to physical CPU threads 0-55, which would be the entire first CPU of a host with 28-core CPUs with hyper-threading enabled. 9. Click **OK**. 10. Click **OK** to close the **Edit VM** dialog box. For more information about potential performance practices, see vSphere Resource Management documentation.
> 4-socket VM on 8-socket hosts (in the case the HCI vendor supports 8-socket SAP HANA HCI hosts)	Add the advanced parameter: `numa.slit.enable = "TRUE"` to ensure the correct NUMA map for VMs > 4 socket on 8-socket hosts. **Note:** `sched.vcpuXx.affinity = "Yy-Zz"` must be configured when `numa.slit.enable` is set to `"TRUE"`.

Examples of VMX configurations for SAP HANA VMs

The following examples provide an overview of how to set additional VMX parameters for SAP HANA half-socket and full-socket VMs. These parameters can be added through the vSphere Client or by directly adding these parameters to the .vmx file with a text editor.

Table 25. Additional VMX parameters to set for half-socket SAP HANA VMs on 28-core/n-socket CPU host; socket 0

Half-socket SAP HANA VM additional VMX parameters	Settings

First half-socket VM on socket 0 on an example 28-core CPU n-socket host	`numa.vcpu.preferHT="TRUE"`
	`sched.vcpu0.affinity = "0-27"`
	`sched.vcpu1.affinity = "0-27"`
	`sched.vcpu2.affinity = "0-27"`
	…
	`sched.vcpu26.affinity = "0-27"`
	`sched.vcpu27.affinity = "0-27"`

Second half-socket VM on socket 0 on an example 28-core n-socket CPU host	`numa.vcpu.preferHT="TRUE"`
	`sched.vcpu0.affinity = "28-55"`
	`sched.vcpu1.affinity = "28-55"`
	`sched.vcpu2.affinity = "28-55"`
	…
	`sched.vcpu26.affinity = "28-55"`
	`sched.vcpu27.affinity = "28-55"`

Table 26. Additional VMX parameters to set for single-socket SAP HANA VMs on a 28-core/8-socket CPU host; socket 3

1-socket SAP HANA VM additional VMX parameters	Settings
1-socket VM on socket 3 on an example 28-core CPU 4-socket host	`numa.vcpu.preferHT="TRUE"` `sched.node0.affinity="3"`

1-socket PMem VM on socket 3 on an example 28-core CPU 4 or 8-socket host	`nvdimm0:0.nodeAffinity=3`
	`numa.vcpu.preferHT="TRUE"`
	`sched.node0.affinity="3"`

Table 27. Additional VMX parameters to set for dual-socket SAP HANA VMs on a 28-core/n-socket CPU host; sockets 0-1

2-socket SAP HANA VM additional VMX parameters	Settings

2-socket VM on sockets 0 and 1 on an n-socket host	`numa.vcpu.preferHT="TRUE"`
	`sched.node0.affinity="0"`
	`sched.node1.affinity="1"`

Table 28. Additional VMX parameters to set for 3-socket SAP HANA VMs on a 28-core 4-socket CPU host; sockets 0-2

3-socket SAP HANA VM additional VMX parameters	Settings

3-socket VM on sockets 0, 1, and 2 on a 4-socket host	`numa.vcpu.preferHT="TRUE"`
	`sched.node0.affinity="0"`
	`sched.node1.affinity="1"`
	`sched.node2.affinity="2"`

Table 29. Additional VMX parameters to set for 4-socket SAP HANA VMs on a 28-core 4-socket CPU host

4-socket SAP HANA VM additional VMX parameters	Settings
4-socket VM on a 4-socket server (maximum wide VM)	No additional settings are required as the VM utilizes all server resources.

Table 30. Additional VMX parameters to set for an SAP HANA VM running a latency-sensitive application

Low latency SAP HANA VM additional VMX parameters	Settings
n-socket low-latency VM on an n-socket, CPU server (Valid for all VMs when an even lower latency is required.)	`monitor.idleLoopSpinBeforeHalt = "TRUE"` `monitor.idleLoopMinSpinUS = "50"`

CPU thread matrix

The following table shows the CPU thread matrix of a 28-core CPU as a reference when configuring the `sched.vCPUXx.affinity = "Xx-Yy"` parameter. The list shows the start and end range required for the `"Xx-Yy"` parameter. For example, for the fourth CPU (CPU 3), it would be 168–223.

		28 Core, 56 Thread CPU							
		CPU0 Threads	CPU1 Threads	CPU2 Threads	CPU3 Threads	CPU4 Threads	CPU5 Threads	CPU6 Threads	CPU7 Threads
	Half-Socket	0	56	112	168	224	280	336	392
		1	57	113	169	225	281	337	393
		2	58	114	170	226	282	338	394
		3	59	115	171	227	283	339	395
		4	60	116	172	228	284	340	396
		5	61	117	173	229	285	341	397
		6	62	118	174	230	286	342	398
		7	63	119	175	231	287	343	399
		8	64	120	176	232	288	344	400
		9	65	121	177	233	289	345	401
		10	66	122	178	234	290	346	402
		11	67	123	179	235	291	347	403
		12	68	124	180	236	292	348	404
		13	69	125	181	237	293	349	405
		14	70	126	182	238	294	350	406
		15	71	127	183	239	295	351	407
		16	72	128	184	240	296	352	408
		17	73	129	185	241	297	353	409
		18	74	130	186	242	298	354	410
		19	75	131	187	243	299	355	411
		20	76	132	188	244	300	356	412
		21	77	133	189	245	301	357	413
		22	78	134	190	246	302	358	414
		23	79	135	191	247	303	359	415
		24	80	136	192	248	304	360	416
		25	81	137	193	249	305	361	417
		26	82	138	194	250	306	362	418
		27	83	139	195	251	307	363	419
	Half-Socket	28	84	140	196	252	308	364	420
		29	85	141	197	253	309	365	421
		30	86	142	198	254	310	366	422
		31	87	143	199	255	311	367	423
		32	88	144	200	256	312	368	424
		33	89	145	201	257	313	369	425
		34	90	146	202	258	314	370	426
		35	91	147	203	259	315	371	427
		36	92	148	204	260	316	372	428
		37	93	149	205	261	317	373	429
		38	94	150	206	262	318	374	430
		39	95	151	207	263	319	375	431
		40	96	152	208	264	320	376	432
		41	97	153	209	265	321	377	433
		42	98	154	210	266	322	378	434
		43	99	155	211	267	323	379	435
		44	100	156	212	268	324	380	436
		45	101	157	213	269	325	381	437
		46	102	158	214	270	326	382	438
		47	103	159	215	271	327	383	439
		48	104	160	216	272	328	384	440
		49	105	161	217	273	329	385	441
		50	106	162	218	274	330	386	442
		51	107	163	219	275	331	387	443
		52	108	164	220	276	332	388	444
		53	109	165	221	277	333	389	445
		54	110	166	222	278	334	390	446
		55	111	167	223	279	335	391	447
Total CPU Threads	**224**	**56**	**56**	**56**	**56**	**56**	**56**	**56**	**56**
Total CPU Cores	**112**	**28**	**28**	**28**	**28**	**28**	**28**	**28**	**28**

The following table shows the CPU thread matrix of a 24 core CPU as a reference when configuring the sched.vCPUXx.affinity =”Xx-Yy” parameter. The list shows the start and end range required for the “Xx-Yy” parameter. For example, for the fourth CPU (CPU 3) it would be 144-191.

	24 Core, 48 Thread CPU								
		CPU0 Threads	CPU1 Threads	CPU2 Threads	CPU3 Threads	CPU4 Threads	CPU5 Threads	CPU6 Threads	CPU7 Threads
	Half-Socket	0	48	96	144	192	240	288	336
		1	49	97	145	193	241	289	337
		2	50	98	146	194	242	290	338
		3	51	99	147	195	243	291	339
		4	52	100	148	196	244	292	340
		5	53	101	149	197	245	293	341
		6	54	102	150	198	246	294	342
		7	55	103	151	199	247	295	343
		8	56	104	152	200	248	296	344
		9	57	105	153	201	249	297	345
		10	58	106	154	202	250	298	346
		11	59	107	155	203	251	299	347
		12	60	108	156	204	252	300	348
		13	61	109	157	205	253	301	349
		14	62	110	158	206	254	302	350
		15	63	111	159	207	255	303	351
		16	64	112	160	208	256	304	352
		17	65	113	161	209	257	305	353
		18	66	114	162	210	258	306	354
		19	67	115	163	211	259	307	355
		20	68	116	164	212	260	308	356
		21	69	117	165	213	261	309	357
		22	70	118	166	214	262	310	358
		23	71	119	167	215	263	311	359
		24	72	120	168	216	264	312	360
	Half-Socket	25	73	121	169	217	265	313	361
		26	74	122	170	218	266	314	362
		27	75	123	171	219	267	315	363
		28	76	124	172	220	268	316	364
		29	77	125	173	221	269	317	365
		30	78	126	174	222	270	318	366
		31	79	127	175	223	271	319	367
		32	80	128	176	224	272	320	368
		33	81	129	177	225	273	321	369
		34	82	130	178	226	274	322	370
		35	83	131	179	227	275	323	371
		36	84	132	180	228	276	324	372
		37	85	133	181	229	277	325	373
		38	86	134	182	230	278	326	374
		39	87	135	183	231	279	327	375
		40	88	136	184	232	280	328	376
		41	89	137	185	233	281	329	377
		42	90	138	186	234	282	330	378
		43	91	139	187	235	283	331	379
		44	92	140	188	236	284	332	380
		45	93	141	189	237	285	333	381
		46	94	142	190	238	286	334	382
		47	95	143	191	239	287	335	383
Total CPU Threads	**192**	**48**	**48**	**48**	**48**	**48**	**48**	**48**	**48**
Total CPU Cores	**96**	**24**	**24**	**24**	**24**	**24**	**24**	**24**	**24**

The following table shows the CPU thread matrix of a 32 core CPU as a reference when configuring the sched.vCPUXx.affinity =”Xx-Yy” parameter. The list shows the start and end range required for the “Xx-Yy” parameter. For example, for the fourth CPU (CPU 3) it would be 192-255.

32 Core, 64 Thread CPU					
	CPU0 Threads	CPU1 Threads	CPU2 Threads	CPU3 Threads	
Half-Socket	0	64	128	192	
	1	65	129	193	
	2	66	130	194	
	3	67	131	195	
	4	68	132	196	
	5	69	133	197	
	6	70	134	198	
	7	71	135	199	
	8	72	136	200	
	9	73	137	201	
	10	74	138	202	
	11	75	139	203	
	12	76	140	204	
	13	77	141	205	
	14	78	142	206	
	15	79	143	207	
	16	80	144	208	
	17	81	145	209	
	18	82	146	210	
	19	83	147	211	
	20	84	148	212	
	21	85	149	213	
	22	86	150	214	
	23	87	151	215	
	24	88	152	216	
	25	89	153	217	
	26	90	154	218	
	27	91	155	219	
	28	92	156	220	
	29	93	157	221	
	30	94	158	222	
	31	95	159	223	
Half-Socket	32	96	160	224	
	33	97	161	225	
	34	98	162	226	
	35	99	163	227	
	36	100	164	228	
	37	101	165	229	
	38	102	166	230	
	39	103	167	231	
	40	104	168	232	
	41	105	169	233	
	42	106	170	234	
	43	107	171	235	
	44	108	172	236	
	45	109	173	237	
	46	110	174	238	
	47	111	175	239	
	48	112	176	240	
	49	113	177	241	
	50	114	178	242	
	51	115	179	243	
	52	116	180	244	
	53	117	181	245	
	54	118	182	246	
	55	119	183	247	
	56	120	184	248	
	57	121	185	249	
	58	122	186	250	
	59	123	187	251	
	60	124	188	252	
	61	125	189	253	
	62	126	190	254	
	63	127	191	255	
Total CPU Threads	**64**	**64**	**64**	**64**	**256**
Total CPU Cores	**32**	**32**	**32**	**32**	**128**

About the author

Erik Rieger is a Principal SAP Global Technical Alliance Manager and Architect, working with Broadcom's VMware Cloud Foundation Division. He is responsible for defining and managing the SAP solution and validation roadmap for the Broadcom VCF division. Erik works closely with SAP, partners, and customers to define the architectures of VMware-based SAP solutions so that they have the best mix of functionality to help transform businesses into real-time enterprises. Erik has more than 25 years of experience in the IT sector and has a Bachelor of Science degree in electronics and a Master of Science degree in Information Systems and Management.

Acknowledgments

The following individuals from VMware by Broadcom contributed content or helped review this guide:

- **Sathya Krishnaswamy**, Staff Performance Engineer, Performance Engineering team
- **Sebastian Lenz**, Staff Performance Engineer, Performance Engineering team
- **Louis Barton**, Staff Performance Engineer, Performance Engineering team
- **Julie Brodeur**, Senior Technical Writer, Performance Engineering team