



SAP HANA on Hyperconverged Infrastructure (HCI) Solutions Powered by VMware vSAN™

BEST PRACTICES AND REFERENCE ARCHITECTURE GUIDE FOR 2ND GENERATION INTEL® XEON® SCALABLE PROCESSORS BASED SYSTEMS RUNNING VSPHERE AND VSAN 6.7 U3 AND 7.0 U1/U2/U3C

DOCUMENT VERSION 2.3

Table of Contents

Abstract	3
Executive Summary	3
<i>Overview</i>	3
<i>Audience</i>	4
Solution Overview	4
<i>SAP HANA on vSAN</i>	4
<i>Solution Components</i>	5
<i>SAP HANA HCI on vSAN—Support Status as of Note 2718982</i>	7
Reference Architecture	9
<i>Scope</i>	9
<i>SAP HANA HCI on vSAN—Architecture Overview</i>	9
<i>SAP HANA HCI on vSAN—Scalability and Sizes</i>	10
<i>VMware HA</i>	15
VMware HA to Protect Scale-Up and Scale-Out SAP HANA Systems	15
vSphere Clustering Service (vCLS).....	17
VMware vSAN Stretched Cluster Support for SAP HANA Systems	18
<i>SAP HANA Stretched Clusters VM Storage Policy</i>	20
<i>Reference Test System Configuration</i>	20
Reference Environment.....	21
IOPS Limit per Partition Policy—Quality of Service Option to Storage Throughput	21
Hardware Resources	22
Software Resources	28
SAP HANA HCI on vSAN Configuration Guidelines	29
<i>SAP HANA HCI Network Configuration</i>	29
<i>VMware ESXi Server: Storage Controller Mode</i>	31
<i>vSAN Disk Group Configuration</i>	31
<i>vSAN Storage Configuration</i>	33
<i>vSAN Storage Policy</i>	33
<i>SAP HANA HCI Storage Planning and Configuration</i>	35
SAP HANA Storage Capacity Calculation and Configuration	36
Usage of Linux Logical Volume Manager (LVM)	38
vSAN Storage Capacity Calculation	39
SAP HANA HCI VMware vSAN Detection Script	40
SAP HANA HCI on vSAN—Support and Process	44

<i>Support Status</i>	44
<i>Support Process</i>	44
<i>vSAN Support Insight</i>	44
Conclusion	44
Appendix	45
<i>Optimizing the SAP HANA on vSphere Configuration Parameter List</i>	45
<i>Performance Optimization for Low Latency SAP HANA VMs</i>	54
<i>Additional Performance Tuning Settings for SAP HANA Workloads</i>	55
<i>Example VMX Configurations for SAP HANA VMs</i>	57
<i>CPU Thread Matrix</i>	60
<i>Monitoring and Verifying an SAP HANA Installation</i>	64
References	65
Revision History	65
Authors	65
Contributors	65

Abstract

This document provides best practices and architecture guidelines to design and build an SAP HANA environment on Hyperconverged Infrastructure (HCI) solutions based on Skylake, Cascade Lake, and Cooper Lake (second generation of INTEL® XEON® SCALABLE processor) based systems running with VMware vSphere® and VMware vSAN™.

The document provides information any hardware vendor can use as a starting point to build their own unique SAP HANA HCI solutions.

The vendor is responsible to define and select the best vSAN ready hardware components, VMware vSphere, and vSAN configuration for its unique solution. The HCI solution vendor is also responsible to certify its solution accordingly 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 includes information about the support processes and requirements of SAP HANA on vSAN.

The version 2.3 of this document includes support for 8-socket wide VM configurations and vSphere and vSAN 7.0 U3c.

Executive Summary

Overview

For more than 10 years, VMware has successfully virtualized SAP applications on x86 server hardware. It is a logical next step to virtualize and use software to define the storage for these applications and their data. SAP has provided production support for their applications on VMware vSphere for the last 10 years, and for HANA since 2014.

SAP HANA based applications require dynamic, reliable, and high-performing server systems coupled with a predictable and reliable storage subsystem. The servers and storage must deliver the needed storage throughput and IOPS capacity at the right protection level and costs for all SAP HANA VMs in the environment.

Storage solutions like large Fibre Channel SAN or standalone SAP HANA appliance configurations, are not able to quickly react to today's changing and dynamic environment. Their static configurations have difficulty keeping pace with modern and resource-demanding applications.

VMware vSAN, a software-defined storage solution, allows enterprises to build and operate highly flexible and application-driven storage configurations, which meets SAP HANA application requirements. Detailed SAP support information for VMware vSAN can be found in SAP Note [2718982 - SAP HANA on VMware vSphere and vSAN](#). A valid SAP s-user is required to view the note.

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.

Solution Overview

SAP HANA on vSAN

SAP HANA is an in-memory data platform that converges database and application platforms, which 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, Durability (ACID) compliance model.

SAP environments with SAP HANA can use vSAN running on certified x86 servers to eliminate traditional IT silos of compute, storage, and networking with this hyperconverged solution.

All intelligence and management are moved into a single software stack, allowing for a VM- and 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 get used for SAP HANA Scale-Up and Scale-Out deployments.

Following figure summarizes the benefits of VMware based SAP HCI solutions.

Radically Simple Storage and SAP HANA Infrastructure	Elastic Scalability, High Performance	Lower TCO
 <ul style="list-style-type: none"> • Two-clicks to enable • Single pane of glass management • Policy-driven • Self-tuning • Integrated with VMware stack • Same one-stop SAP and VMware support process as with vSphere • Ready for any SAP workloads • Fully SAP HANA supported! 	 <ul style="list-style-type: none"> • Built on high-performance standard x86 servers • Embedded in vSphere kernel with direct I/P path • Fast and flexible SAP deployments • Up to four 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 	 <ul style="list-style-type: none"> • Server-side economics • No large upfront investments, grow-as-you-go • Easy to operate with powerful automation • No specialized skillset required • Reliable and standard x86 hardware • Built-in HA • SAP HANA virtual infrastructure turnkey solution.

Only SAP and VMware certified and tested components, like server and storage systems, are supported. These components must comply with the SAP HANA Hardware Certification Hyperconverged Infrastructure (HANA-HyCI) test scenario.

Note: Only SAP and VMware HCI solution vendor offered solutions, that get [SAP HANA HCI solution certified](#), 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, like an SAP application server.

Solution Components

An SAP HANA HCI solution based on VMware technologies is a fully virtualized and cloud ready infrastructure turnkey solution running on VMware vSphere and vSAN. All components like compute, storage, and networking are delivered together as an HCI vendor specific SAP HANA solution. It is 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 OEM:

- Turnkey deployment using the existing vSAN/VCF appliances and modify these systems to be SAP HANA HCI ready.
- Use the certified vSAN Ready Nodes and components and build an SAP HANA ready HCI solution.

The solution consists of the following components:

- VMware certified vSAN HW partner solutions, as listed on the [VMware vSAN HCL](#) (vSAN Appliances or Ready Nodes with vSAN Ready Components).
- SAP HANA supported server systems, as listed on the [SAP HANA HCL](#).

- VMware vSphere Enterprise Plus starting with version 6.5 update 2. For Cascade Lake based servers vSphere 6.7u3 or later and for Cooper Lake based servers vSphere 7.0 U2 or later is required. Review SAP note [2393917 for an updated list of supported vSphere versions](#).
- VMware vSAN Enterprise starting with version 6.6. For Cascade Lake based servers vSAN 6.7u3 or later is required. Review SAP note [2718982 for an updated list of supported vSAN versions](#).
- VMware vCenter starting with version 6.5. Note that vCenter versions need to be aligned with the used vSphere version.
- HCI vendor implemented the 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).
- HCI vendor SAP-integrated support process.

After the solutions are defined and built, SAP ICC validates and lists these solutions in [Certified HCI Solutions](#) on the SAP HANA Supported and Certified Hardware Directory section.

Figure 1 provides an overview of the solution components.

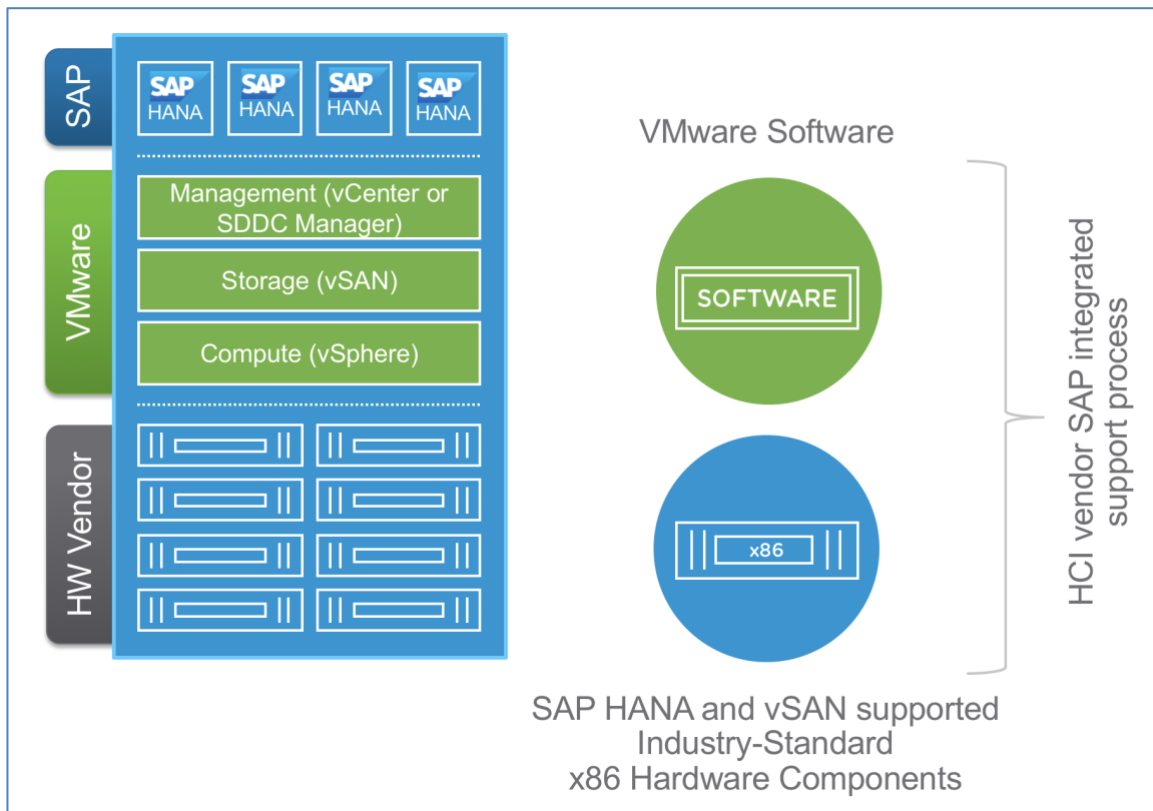


Figure 1. SAP HANA HCI Solution Overview

SAP HANA HCI on vSAN—Support Status as of Note [2718982](#)

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

- **VMware vSphere Hypervisor 6.5, 6.7, 6.7 U2/U3, 7.0 U1/U2/U3c**
 - 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. SAP HANA on vSAN has additional constraints, see next section.
 - vSphere 7.0 U1/U2/U3c are supported as of today for SAP HANA HCI solutions.
- **VMware vSAN 6.6, 6.7, 6.7 U1/U2/U3), 7.0 U1/U2/7.0 U3c¹**
 - vSAN 6.6, 6.7, 6.7 U1/U2/U3, 7.0 U1/U2/U3c: Maximum 4 SAP HANA VMs on 2, 4, and 8 socket servers

¹ vSAN 7.0 is not supported.

- Starting from above vSAN versions, the HCI vendor can select the appropriate vSAN version for their HCI solution.
 - The vSAN versions 6.6, 6.7, 7.0 U1/U2/U3c are included in vSphere version 6.5 U2 and vSphere version 6.7, 6.7 U2/U3, 7.0 U1/U2/U3c respectively.
 - Minimal 3, maximal 64 nodes per SAP HANA HCI vSAN cluster.
 - 4 or more vSAN nodes are recommended for optimal failure resilience.
 - No CPU sockets need to get reserved for vSAN. Note that vSAN will need, just like ESXi, system resources. Refer to [VMware KB 2113954](#) for information about vSAN resource needs.
- **CPU architecture and VM maximums:**
 - Any supported Intel Skylake, Cascade Lake or Cooper Lake processors with 2, 4, and 8 socket servers with existing VMware vSphere virtualization boundaries, for example:
 - max. **12 TB RAM and 448 vCPUs per SAP HANA Scale-Up VM**, see SAP Note [3102813](#) for details. Also note that the actual usable vCPUs depend on the used CPU type: the current SAP HANA vSphere maximum vCPU number is 448 with 8-socket host systems.
 - The certified hardware for HCI solutions can be found on the [Certified Hyper-Converged Infrastructure Solutions](#)
- **SAP HANA Deployment Options (SAP note [3102813](#)):**
 - **SAP HANA Scale-Up:**
 - OLTP or mixed type workload up to 12 TB RAM and up to 448 vCPUs per VM
 - OLAP type workload:
 - SAP sizing Class L: up to 6 TB RAM and up to 448 vCPUs per VM
 - SAP sizing Class M: up to 12 TB RAM and up to 448 vCPUs per VM
 - **SAP HANA Scale-Out:**
 - Up to 1 master and 7 worker nodes are possible, plus HA node(s). HA nodes should be added according to the cluster size. [FAQ: SAP HANA High Availability](#)
 - SAP HANA Scale-Out system max. 16 TB memory with SAP HANA CPU L-Class* sizing:
 - Up to 2 TB of RAM and up to 256 vCPUs per Scale-Out node (VM), sizing up to SAP HANA CPU Sizing Class L type workloads
 - SAP HANA Scale-Out system max. 24 TB memory with SAP HANA M-Class* CPU sizing:
 - Up to 3 TB of RAM and up to 256 vCPUs per Scale-Out node (VM), sizing up to SAP HANA CPU Sizing Class M** type workloads
 - Additional requirements:
 - Scale-Out is only support by SAP with **4 or 8** socket servers and minimal 4-socket wide VMs that use up to 256 logical CPU threads, in the case of Cascade Lake 28 core CPUs this represents a maximal VM size with 224 vCPUs.

Note: SAP Note 2718982 - SAP HANA on VMware vSphere and vSAN will be continuously updated with the latest support information without further notice. Above support status is from February 2022. Before planning an SAP HANA HCI vSAN based solution, check note 2718982 for any updates.

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.

Reference Architecture

The purpose of the introduced vSAN for SAP HANA reference architectures is to provide a starting point for HCI solution vendors who want to build a vSAN based solution and to provide best practices and guidelines.

The reference architecture is tested and validated with the SAP specified HCI tests and is known to be good with the shown SAP HANA configuration.

Scope

This reference architecture:

- Provides a technical overview of an 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 single DC to multiple suite configurations

SAP HANA HCI on vSAN—Architecture Overview

vSAN Clusters consist of two or more physical hosts that contain either a combination of magnetic disks and flash devices (hybrid configuration) or all-flash devices (all-flash configuration). All vSAN nodes contribute cache and capacity to the vSAN distributed/clustered datastore and are interconnected through a dedicated 10 GbE (minimum) or better path redundant network.

Note: In SAP HANA environments, four nodes are recommended for easier maintenance and Stretched Cluster readiness, but the minimum requirement of a three-host vSAN cluster for SAP HANA HCI can be used.

Figure 2 shows the high-level architecture of an SAP landscape running on a VMware-based SAP HANA HCI solution built upon several vSphere hosts and vSAN for a single data center configuration. Unlike with other SAP HANA HCI solutions, a VMware vSAN powered SAP HANA HCI solution can leverage all available ESXi host CPU and RAM resources. As shown in the figure, one host has enough compute resources available to support a complete host failover.

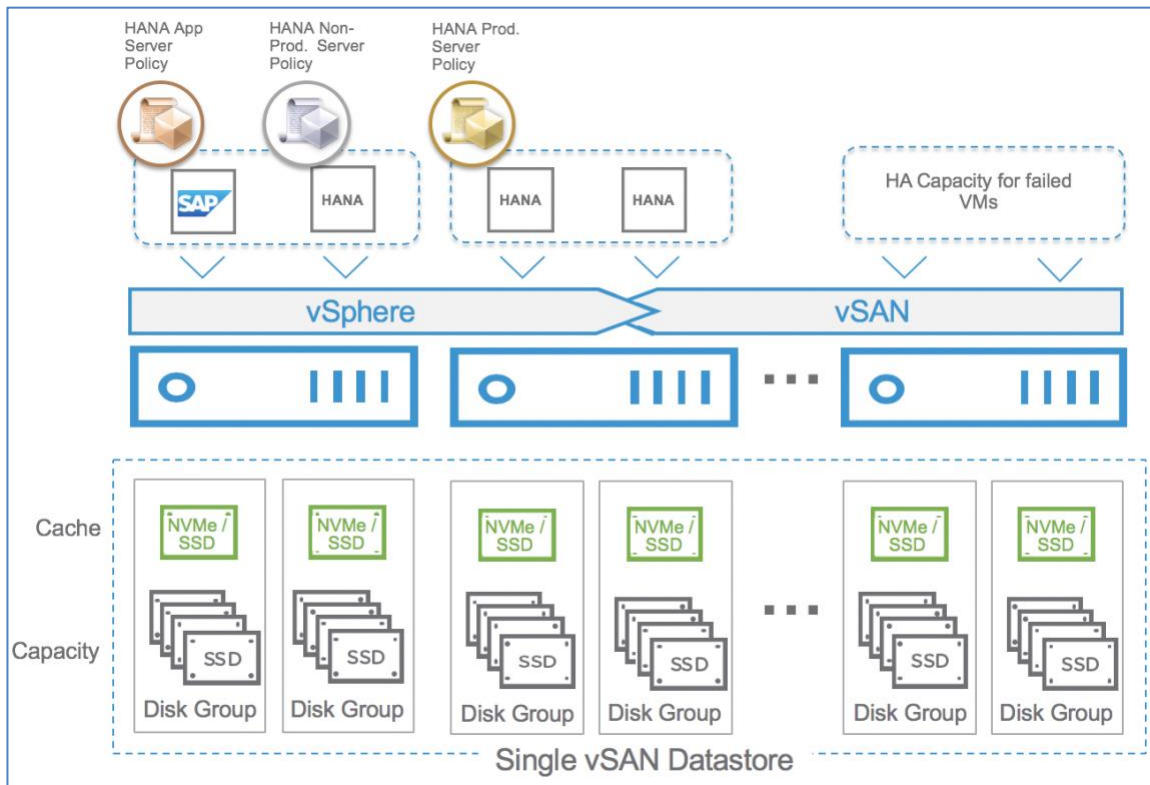


Figure 2. SAP HANA Applications on vSAN High-Level Architecture

[Storage policies](#) for less demanding SAP HANA application servers and for IO demanding HANA VMs, as shown in Figure 2 (bronze, silver, and gold), will ensure that all VMs get the performance, resource commitments, failure tolerance, and quality of service as a required policy-driven dynamic environment. With such an environment, changing the assigned storage policy in response to the changing needs of a virtualized application can directly influence the underlying infrastructure. The different storage policies allow changing the level of protection or possible IO limits without disturbing operation of a VM, allowing it to remain available and maintain the defined SLAs and performance KPIs.

SAP HANA HCI on vSAN—Scalability and Sizes

SAP HANA on vSphere is supported across a range of SAP HANA system size. The smallest size is a half socket CPU configuration with a minimum of 8 physical CPU cores and 128 GB of RAM. The largest size, as of today, is a 8 socket wide VM with 12 TB of RAM and 448* vCPUs. Scale-Out SAP HANA deployments are also supported with up to eight 3 TB nodes. In total, up to 24 TB system memory can be deployed on SAP HANA vSAN HCI solutions for OLAP type workloads. Starting with vSphere 6.7 U2, there is also an alternative sizing possibility based on the SAP HANA workload. More information in [SAP note 2779240](#).

vSphere 6.7 U3, 7.0 U1/U2 and U3c* as supported for vSAN provides the following [maximums per physical host](#) (table 1). For SAP HANA, only servers with up to 8 physical CPU sockets are currently supported.

*Actual usable number of vCPU depends on the selected CPU type and vSphere version. vSphere 7.0 U2 or later supports VMs with up to 24 TB and 768 vCPUs.

Table 1. Current vSphere HCI Supported Host Maximums

	ESX 6.7 U3 / 7.0	ESX 7.0 U2/U3c
<i>Logical CPUs per host</i>	768	896
<i>Virtual machines per host</i>	1024	
<i>Virtual CPUs per host</i>	4096	
<i>Virtual CPUs per core</i>	32	
<i>RAM per host</i>	16 TB	24 TB
<i>NUMA Nodes / CPU sockets per host</i>	16 (SAP HANA only 8 CPU sockets)	

Table 2 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. The current VM size is limited for SAP HANA to 448 vCPUs and 12 TB of RAM (8-socket wide VM).

Table 2. vSphere Guest (VM) Maximums

	ESX 6.7 U3 / 7.0	ESX 7.0 U2/U3c
<i>Virtual CPUs per virtual machine</i>	256	768
<i>RAM per virtual machine</i>	6128 GB	24 TB
<i>Virtual CPUs per HANA virtual machine</i>	224	448
<i>RAM per HANA virtual machine</i>	6128 GB	<=12 TB
<i>Virtual SCSI adapters per virtual machine</i>	4	
<i>Virtual NVMe adapters per virtual machine</i>	4	
<i>Virtual disk size</i>	62TB	
<i>Virtual NICs per virtual machine</i>	10	

Table 3 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 3 vSAN cluster nodes, and up to 64 cluster nodes are supported by vSAN.

Table 3. vSAN Maximums

VMware vSAN 6.7 (incl. U3) and 7.0 U1/U2	vSAN ESXi host	vSAN Cluster	vSAN Virtual Machines
<i>vSAN disk groups per host</i>	5		
<i>Cache disks per disk group</i>	1		
<i>Capacity tier maximum devices per disk group</i>	7		
<i>Cache tier maximum devices per host</i>	5		
<i>Capacity tier maximum devices</i>	35		
<i>Components per vSAN host</i>	9,000		
<i>Number of vSAN hosts in an All-Flash cluster</i>		64	
<i>Number of vSAN hosts in a hybrid cluster</i>		64	
<i>Number of vSAN datastores per cluster</i>		1	
<i>Virtual machines per host</i>			200
<i>Virtual machine virtual disk size</i>			62 TB
<i>Disk stripes per object</i>			12
<i>Percentage of flash read cache reservation</i>			100
<i>Percentage of object space reservation</i>			100
<i>vSAN networks/physical network fabrics</i>			2

Table 4 shows the size of SAP HANA VMs possible on a supported SAP HANA HCI solution for production use case with vSphere and vSAN 6.7 U3, 7.0 U1/U2 and U3c.

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 be used alternatively to meet the requirements.

We have validated that **4 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 4 shows the example CPU's VM sizes that can be deployed on a vSAN SAP HANA HCI node.

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 get reserved for vSAN. VMware KB 2113954 provides details about the vSAN and memory consumption in ESXi. For sizing purposes an additional 10% sizing buffer should get subtracted from the virtual SAP HANA SAPS capacity to factor the vSAN CPU resource

costs in. This sizing buffer is a sizing value, which can be much lower in typical customer environments, or in rare situations like during vSAN storage re-organization tasks, can be higher.

Table 4. Example SAP HANA VM Configurations and Capacity of VMware virtualized SAP HANA instances based on selected 2nd Generation INTEL® XEON® Scalable Processors⁵

	Intel Xeon EX E7-8890 v4 CPU	Intel Xeon Gold 6258R CPU	Intel Xeon Platinum 8280L CPU	Intel Xeon Platinum 8380H CPU and HL
SAP Benchmark	2016067 (20. Dec. 2016)	2020015 (05. May 2020)	2019023 (02. April 2019)	2020050 (11. Dec. 2020)
Max supported RAM per CPU as of Intel datasheet ₁	3 TB, (maximal supported RAM by SAP is 1 TB ¹)	1024 GB	4.5 TB	4.5 TB (HL)
CPU Cores per socket as of Intel datasheet.	24	28	28	28
Max. NUMA Nodes per ESXi Server	8 ²	2	8 ²	8 ²
vSAPS per CPU thread w/ and w/o HT	1785 (core w/o HT) 315 (HT gain) Based on SD benchmark cert.2016067	2459 (core w/o HT) 434 (HT gain) Based on SD benchmark cert. 2020015	2596 (core w/o HT) 458 (HT gain) Based on SD benchmark cert. 2019023	2432 (core w/o HT) 429 (HT gain) Based on SD benchmark cert.2020050
0.5 Socket SAP HANA VM (Half-Socket)	1 to 16 x 12 physical core VM with min. 128 GB and max. 512 GB ³	1 to 4 x 14 physical core VMs with min. 128 GB RAM and max. 512GB ³	1 to 16 x 14 physical core VMs with min. 128 GB RAM and max. 768 GB ³	1 to 16 x 14 physical core VMs with min. 128 GB RAM and max. 768 GB ³
	vSAPS ^{4,5} 21,000, 24 vCPUs	vSAPS ^{4,5} 34,000, 28 vCPUs	vSAPS ^{4,5} 36,000, 28 vCPUs	vSAPS ^{4,5} 34,000, 28 vCPUs
1 Socket SAP HANA VM	1 to 8 x 24 physical core VM with min. 128 GB RAM and max. 1024 GB	1 to 2 x 28 physical core VM with min. 128 GB RAM and max. 1024 GB	1 to 8 x 28 physical core VM with min. 128 GB RAM and max. 1536 GB	1 to 8 x 28 physical core VM with min. 128 GB RAM and max. 1536 GB
	vSAPS ^{4,5} 50,000, 48 vCPUs	vSAPS ^{4,5} 81,000, 56 vCPUs	vSAPS ^{4,5} 85,000, 56 vCPUs	vSAPS ^{4,5} 80,000 56 vCPUs
2 Socket SAP HANA VM	1 to 4 x 48 physical core VM with min. 128 GB RAM and max. 2048 GB	1 x 56 physical core VM with min. 128 GB RAM and max. 2048 GB	1 to 4 x 56 physical core VM with min. 128 GB RAM and max. 3072 GB	1 to 4 x 56 physical core VM with min. 128 GB RAM and max. 3072 GB
	vSAPS ^{4,5} 100,000, 96 vCPUs	vSAPS ^{4,5} 162,000, 112 vCPUs	vSAPS ^{4,5} 171,000, 112 vCPUs	vSAPS ^{4,5} 160,000, 112 vCPUs
3 Socket SAP HANA VM	1 to 2 x 72 physical core VM with min. 128 GB RAM and max. 3072 GB	-	1 to 2 x 84 physical core VM with min. 128 GB RAM and max. 4608 GB	1 to 2 x 84 physical core VM with min. 128 GB RAM and max. 4608 GB
	vSAPS ^{4,5} 151,000, 144 vCPUs	-	vSAPS ^{4,5} 256,000, 168 vCPUs	vSAPS ^{4,5} 240,000, 168 vCPUs
4 Socket SAP HANA VM	1 to 2 x 96 physical core VM with min. 128 GB RAM and max. 4096 GB	-	1 to 2 x 112 physical core VM with min. 128 GB RAM and max. 6128 GB	1 to 2 x 112 physical core VM with min. 128 GB RAM and max. 6128 GB
	vSAPS ^{4,5} 201,000, 192 vCPUs	-	vSAPS ^{4,5} 342,000, 224 vCPUs	vSAPS ^{4,5} 320,000, 224 vCPUs
8 Socket SAP HANA VM	-	-	1 x 224 physical core VM with min. 128 GB RAM and max. 12096 GB	1 x 224 physical core VM with min. 128 GB RAM and max. 12096 GB

	-	-	vSAPS ^{4,5} 684,000, 448 vCPUs	vSAPS ^{4,5} 640,000, 448 vCPUs
--	---	---	---	---

¹SAP supports in appliance configuration up to 1.5 TB for Suite and 1 TB for BW SAP HANA applications with SKL and later CPU's. More RAM may be possible with a TDI phase 5 / workload-based sizing.

²Maximal 224 vCPUs can get configured per VM with vSphere 6.7 U2 and later with a Skylake, Cascade Lake or Cooper Lake CPU.

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

⁴The listed vSAPS figures are based on published SD benchmark results with Hyperthreading (2-vCPU configuration) and minus 10% virtualization costs. In the case of a half-socket configuration in addition to the 10% virt. costs, 15% from the SD capacity must get subtracted. The shown figures are rounded figures and based on rounded SAPS performance figures published [SAP SD benchmarks](#), and can get only used for Suite or BW on HANA or BW/4HANA workloads. For mixed HANA workloads sizing parameters contact SAP or your HW vendor.

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

Scale-Out deployments are supported with up to eight² 3 TB 4-socket wide VMs. This provides for an OLAP type SAP HANA instance up to 24 TB in total system RAM by utilizing the Scale-Out deployment model.

Figure 3 shows a 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. Minimal one vSAN node is needed to support HA.

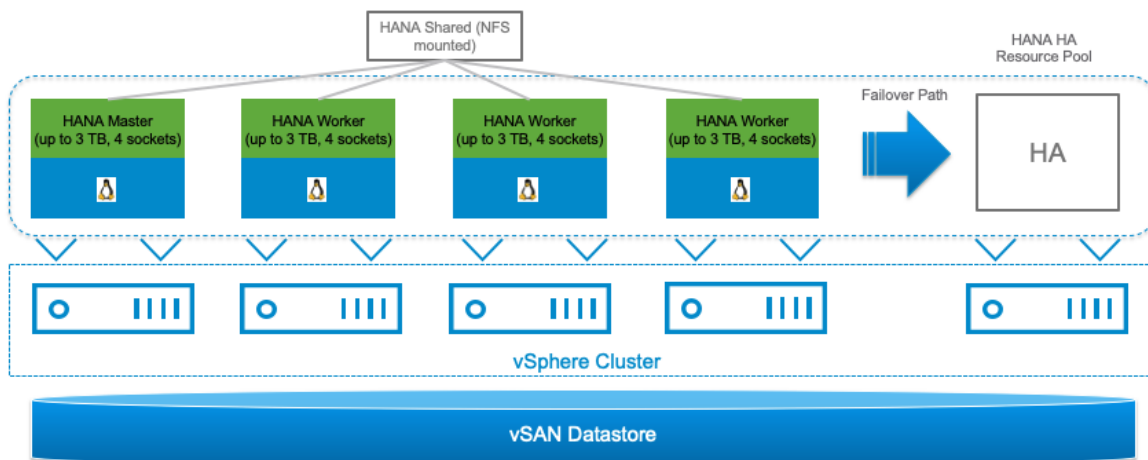


Figure 3. SAP HANA Scale-Out Example Cluster on vSAN (HCI servers)

² This is a support rather than technical limitation.

VMware HA

VMware vSphere 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 to Protect Scale-Up and Scale-Out SAP HANA Systems

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 Guidelines and Best Practices for Deployments 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 or other hosts 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 4 shows the concept of a typical n+1 configuration, which is the standard deployment of an SAP HANA HCI solution. It is recommended deploying an n+2 configuration for easy maintenance and to protect your systems against a failure while you perform maintenance.

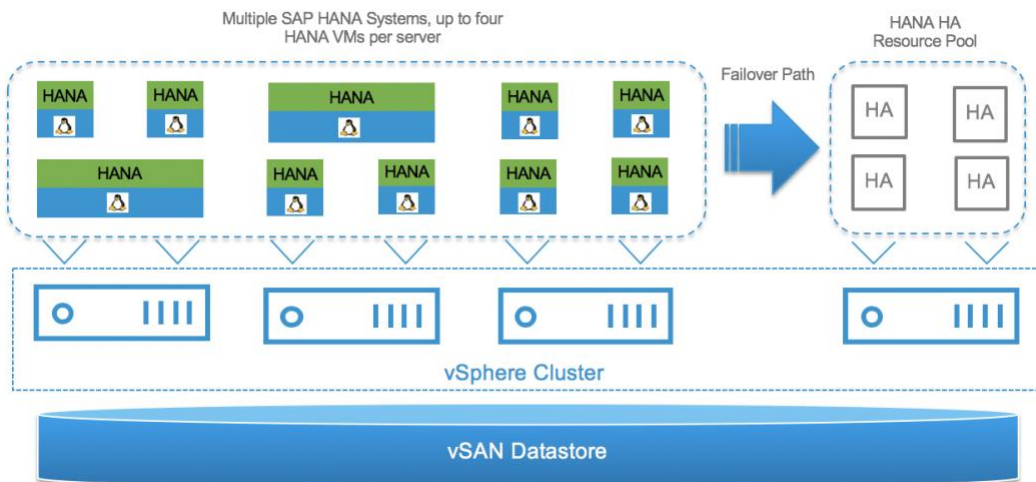


Figure 4. VMware HA Protected SAP HANA VMs in an n+1 vSAN Example Cluster Configuration

Figure 5 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.

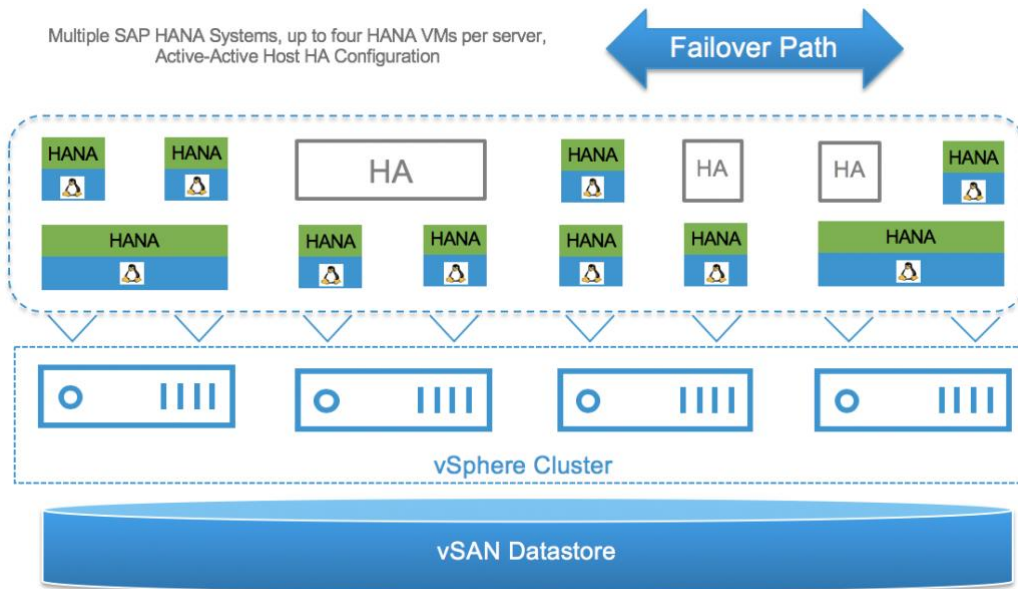


Figure 5. VMware HA Active-Active n+1 vSAN Example Cluster Configuration

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 4 virtual CPU sockets respectively. Figure 6 shows a configuration of 3 SAP HANA VMs (1 x master with 2 worker nodes) running on the virtual SAP HANA HCI platform. The NFS mounted SAP HANA shared directory is not part of the vSAN cluster and needs to be provided as an vSAN external NFS service*.

* vSAN 7.0 U1/U2 File Service (NFS) is not supported. **[Stretched Cluster Documentation](#)

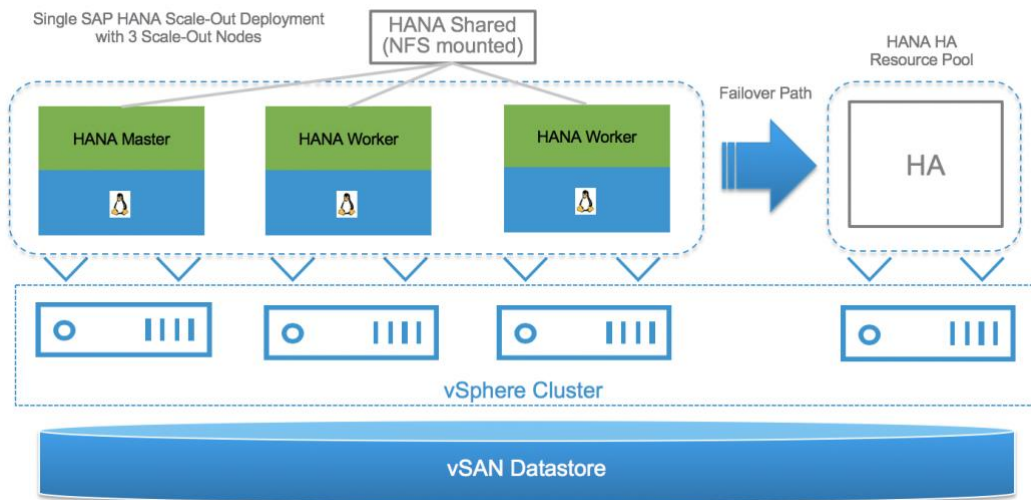


Figure 6. SAP HANA Scale-Out VMware HA n+1 vSAN Example Cluster Configuration

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. HANA shared 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. vSAN 7.0 U1 File Service (NFS) are yet not supported.

vSphere Clustering Service (vCLS)

Starting with vSphere 7.0 Update 1, vSphere Cluster Service (vCLS) is enabled by upgrading the vCenter to 7.0U1 and runs 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 the vision.

The dependency of these cluster services on vCenter is not ideal and VMware introduced in the [vSphere 7 Update 1](#) release 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. If vCenter Server becomes unavailable, in future, vCLS will ensure that the cluster services remain available to maintain the resources and health of the workloads that run in the clusters.

vCLS Deployment Guidelines for SAP HANA Landscapes

As of SAP notes [2937606](#) and [3102813](#), it is not supported to run a non-SAP HANA VM on the same NUMA node where a 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, as long as 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, it is necessary, because of 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 get migrated to hosts that do not run SAP HANA production level VMs.

This can get 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 get assured that vCLS VMs and SAP HANA VMs do not get co-deployed.

After the policy is created and tags were 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 get added to the cluster then the [Retreat Mode](#) can get used to remove the vCLS VMs from this cluster. Please 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 hosts that do not run SAP HANA tagged VMs. **If this is not possible then you can use the [Retreat Mode](#) to remove the vCLS VMs from the cluster.** In the case of VCF an additional node to offload the vCLS VMs is required per cluster (e.g. the HA node).

For more details and vCLS deployment examples in the context of SAP HANA please refer to this [blog](#).

VMware vSAN Stretched Cluster Support for SAP HANA Systems

Stretched clusters extend the vSAN cluster from a single data site to two sites for a faster 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. SAP HANA is now tested and is supported on vSAN stretched cluster 7.0 U1/U2/U3c environments. SAP HANA requires that log file writes have a less than 1ms latency, and this log file write latency should not exceed 1ms across primary and secondary sites. The underlying network infrastructure plays a key role, as the distances supported between the two sites solely depends on ability to achieve latency of 1ms across these sites for SAP HANA.

Tests conducted with one of our HCI hardware partners showed that distances up to 30 km between the two sites, all SAP HANA performance throughput and latency KPIs were passed across the tested SAP HANA VMs. Actual distances could be lower since the network latency is depending on the used network components and architecture. Therefore, distances up to 5 km are supported. SAP HANA HCI support of stretched cluster configurations is upon the SAP HANA HCI hardware vendor and not VMware.

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 sites. If either site fails, vSAN uses the storage on the other site. vSphere HA restarts any VM that must be restarted on the remaining active site.

You must designate one site as the preferred site. The other site becomes a secondary or nonpreferred site. If the network connection between the two active sites is lost, vSAN continues operation with the preferred site. The site designated as preferred typically is the one that remains in operation unless it is resyncing or has another issue. The site that leads to maximum data availability is the one that remains in operation.

Below example shows a simplified stretched cluster configuration. Site X is the preferred (prod. site), which can be a complete datacenter or a single rack. Site Y is the secondary site and in this case the failover site in the case of site failure. Local failures can get compensated by providing HA capacity inside the site and a site failover can get initiated by providing enough failover capacity at site Y, or by freeing up failover capacity or by only failing over production level SAP HANA VMs. The witness appliance must run on a host placed in an independent site Z. The witness decides ultimately on a site failover and can run for instance on an ESXi host which supports the SAP application server VMs.

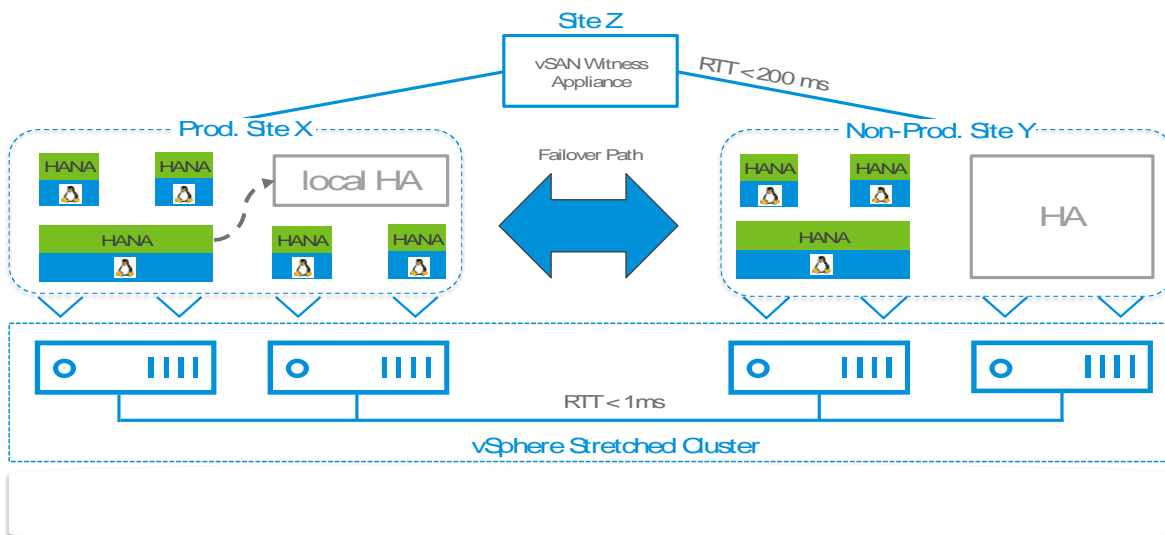


Figure 7. SAP HANA Stretched Cluster Example Configuration

Summary:

- Tested distances maintain ≤ 1 ms KPI for SAP HANA log writing up to 30 km
- Stretched clusters are normally used campus wide, therefore distances up to 5 km supportable.
- Witness VM can run on any vSphere host in 3rd size “z”.
- Failover capacity is $n+1$, where 1 stands for one host failure. $N+2$ would support two host failures.
- Failover can be bi-directional.

SAP HANA Stretched Clusters VM Storage Policy

For an SAP HANA VM in vSAN stretched cluster, the VM storage policy should be enhanced for Dual Site Mirroring and with Failures to tolerate equal to 1 failure – RAID1 (Mirroring) as shown in figure 8.

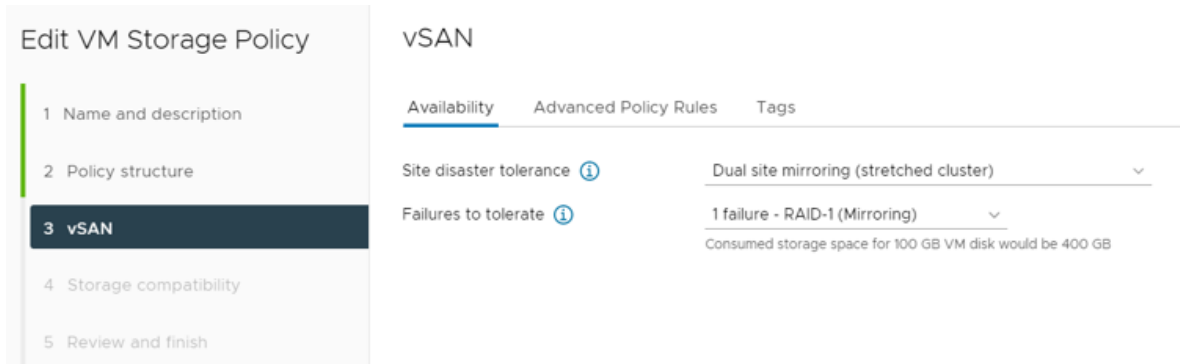


Figure 8. SAP HANA VM Storage Policy in vSAN Stretched Cluster

When Site Disaster Tolerance is set to “Dual Site Mirroring”, a copy of the data goes to both sites, When the “Primary Failures to Tolerate” rule is equal to 1, writes will continue to be written in a mirrored fashion across sites, therefore check and ensure that the required disk capacity is available on 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 shown reference configuration should be treated as an example but not as a requirement. The HCI solution vendor may or may not follow the shown example configurations and select different components. The only requirements are that the SAP HANA HCI KPIs are met and that only supported/certified components are used.

Reference Environment

To ensure that a vSAN based HCI solution is SAP HANA ready, SAP provided test tools and test cases (HANA-HyCI 1.0 v1.2 test scenario). These test tools validate performance, scalability, and consistency. Under the “[Certified Hyper-Converged Infrastructure Solutions](#)”, a complete set of vSAN supported configurations from different vendors are listed. For example, this [link](#) provides a list for Cascade Lake and vSAN supported configurations.

The reference configurations listed on the next pages are validated by VMware and provide an example for Scale-Up and Scale-Out configurations. The Scale-Up example configurations are based on 2-,4 and 8-socket systems, configured as a 3-host and 4-host vSAN cluster with 2 disk groups and 4 disk groups per server. The configurations had Intel Optane NVMe devices as caching layer, and SSD based capacity disks. With both configurations, it is 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 3 TB memory, 4 CPU sockets, and 4 disk groups installed and configured. In the 8-socket host case two individual Scale-Out systems could get consolidated on an 8-socket host configuration.

IOPS Limit per Partition Policy—Quality of Service Option to Storage Throughput

vSAN storage policies provide multiple options, many features are built into storage policy that has capability to limit the IOPS on the partition that this policy is enforced upon. This IOPS value can be set when a policy is created. Typically, a value of “0” ensures that there is no IOPS limit on partition that this policy is enforced upon. After the policy is created, it can be applied to any partition or partitions across VMs or single VM depending on requirements.

A typical HANA instance has log and data partition. We have done internal testing scenarios where the IOPS limit is set on data partition only, across a set of 8 HANA VMs RAID 1, running on a 3-host cluster with one HA host. To do this, configure a policy with limits and allocate this policy on VMDK which will be mounted as data partition. A minimum setting of around 13,700 IOPS on the test setup's data partition, has shown capability to handle the KPI requirements of data partition. This is the minimum value and anything above this value is also valid. However, for your configuration, it is advisable to do a sanity check with around above ballpark value. For log partition, it is not recommended to set any limits at this time.

Use cases :

1. Depending on application requirements, such a policy will ensure that there is fair IO distribution among different VMs for that specific partition. This is an option that customers can use, if there are multiple HANA VMs running, which requires control to ensure a minimum IO distribution among HANA VMs.
2. Other use cases are that you would want to control the IOPS done by other VMs but want to have one main database VM (unlimited IOPS) as the one getting most of IO resources.

Hardware Resources

Different vSAN clusters were configured in the 4-host and 3-host environments respectively:

- A vSAN SAP HANA HCI environment with four 2-socket based Cascade Lake processors with 2 disk groups per each vSAN host (example 1).
- A vSAN SAP HANA HCI environment with 4 hosts, having a combination of three 4 socket Cascade Lake and one 2 socket Cascade Lake processors configured with 4 disk groups per each vSAN host (example 2).
- A vSAN SAP HANA HCI environment with 3 hosts, having a combination of two 8 socket Cascade Lake and one 4 socket Cascade Lake processors configured with 3 disk groups per each vSAN host (example 3).
- A vSAN SAP HANA HCI environment with 4 hosts, with 4 x 4 socket Cascade Lake processors configured with 4 disk groups per each vSAN host as a scale-out SAP HANA system (example 4).
- A vSAN SAP HANA HCI environment with 4 hosts, with 3 x 4 socket Cooper Lake processors configured with 4 disk groups per each vSAN host (example 5).

The idea using different configurations is to test a variant of system configurations and to provide to the SAP HANA HCI hardware vendors a wide range of proven test configurations they can choose from when building their own solutions. We used mainly 2 x 25 GbE as a shared network link for all tested configurations. For easier administration, an optional 1/10 GbE admin network is recommended.

2-Socket/2-Disk Group SAP HANA HCI vSAN Based—Example 1

The 2-socket/2-disk group test environment had four vSphere ESXi servers configured and used direct-attached SAS and NVMe devices inside the hosts to provide a vSAN datastore that meets the SAP HANA HCI test requirements and KPIs for two SAP HANA VMs per server.

The fourth node was part of the vSAN datastore and provided HA capabilities for the total 6 running SAP HANA VMs in the vSAN cluster.

Each vSphere ESXi server had two disk groups each consisting of one 1.46 TB NVMe device for caching and four capacity tier 3.49 TB SAS SSDs. The raw capacity of the vSAN datastore of this 4-node vSAN cluster was 112 TB. The vSAN cluster limit is up to 64 cluster nodes. The vSAN cluster node number that an HCI vendor supports is up to the vendor.

Two 25 Gbit network adapters were used for vSAN and another network traffic in active-standby failover manner.

Each ESXi Server in the vSAN Cluster had the following configurations as shown in the following table, based on this configuration, up to two running SAP HANA VMs per HCI server.

PROPERTY

SPECIFICATION

ESXi server model	HW Vendor SAP HANA Supported Cascade Lake Server Model
ESXi host CPU	2 x Intel® X®(R) Gold 6248 CPU, 2 Sockets, 20 Physical Cores/Socket.
ESXi host RAM	1.5 TB
Network adapters	2 x 25 Gbit or better as a shared network for vSAN, SAP application and admin network traffic. Optional a dedicated NIC for the admin network. See the network section of this guide for the recommended network configurations. SAP HANA HCI Network Configuration
Storage adapter	4 x 8 channel PCI-EXPRESS 3.0 SAS 12GBS/SATA 6GBS SCSI controller
Storage configuration	Disk Groups: 2 Cache Tier: 2 x 750 GB Intel Optane DC 4800X Series Capacity Tier: 12 X 3.8 TB 3840 Micron 5200 TCG-E
Linux LVM configuration	Used Linux LVM for log and data. Used 4 VMDK disks in LVM each (total 8) to create the needed log and a data volume for SAP HANA.

4-Socket/4-Disk Group SAP HANA HCI vSAN Cluster—Example 2

The **4-socket** test environment had three vSphere ESXi servers configured and used direct-attached SSDs and NVMe inside the hosts to provide a vSAN datastore that meets the SAP HANA HCI test requirements and KPIs.

The third node was part of the vSAN datastore and provided HA capabilities for the total 8 running SAP HANA VMs. Adding the fourth node would allow you to run more SAP HANA VMs on the vSAN datastore. The vSAN cluster limit is up to 64 cluster nodes. The vSAN cluster node number that an HCI vendor supports is determined by the vendor.

Each vSphere ESXi server had **four disk groups** with each consisting of 800 GB NVMe devices for caching and four capacity tier 890GB SAS SSDs. The **raw capacity** of the vSAN datastore of this **3-node** vSAN cluster was for data around **43 TB**. The vSAN cluster limit is up to 64 cluster nodes.

Two 25 Gbit network adapters were used for vSAN and other network traffic in active-standby failover manner.

Review the network section for details about how the network cards impact the number of SAP HANA VMs per host and the sizing section for how many SAP HANA VMs can be deployed on a Skylake based server system.

Each ESXi Server in the vSAN Cluster had the following configuration as shown in Table 5. This configuration supports up to four SAP HANA VMs per HCI server.

Table 5. ESXi Server Configuration of the 4-socket Server SAP HANA on 3-node HCI vSAN Cluster

PROPERTY	SPECIFICATION
ESXi server model	HW Vendor SAP HANA Supported Cascade Lake supported Model
ESXi host CPU	4 x Intel® Xeon(R) Platinum 8280M CPU 2.70GHz , 4 Sockets, 28 Physical Cores/Socket
ESXi host RAM	1.5 TB
Network adapters	2 x 25 Gbit or better as a shared network for vSAN, SAP application and admin network traffic. Optional a dedicated NIC for the admin network. See the network section of this guide for recommended network configurations.
Storage adapter	SAS 12GBS/SATA 6GBS SCSI controller
Storage configuration	Disk Groups: 4 Cache Tier: 4 x Express Flash PM1725a 800GB SFF Capacity Tier: 16 x SAS 894 GB SSD

8-Socket/4-socket 3 Disk Group SAP HANA–HCI vSAN Cluster—Example 3

The **8-socket** and 4-socket test environment had a total of three vSphere ESXi servers configured and used direct-attached SAS SSDs and NVMe cache drives inside the hosts to provide a vSAN datastore that meets the SAP HANA HCI test requirements and KPIs.

The third node was part of the vSAN datastore and provided HA capabilities for the total 8 running SAP HANA VMs. Adding the fourth node would allow you to run more SAP HANA VMs on the vSAN datastore. The vSAN cluster limit is up to 64 cluster nodes. The vSAN cluster node number that an HCI vendor supports is determined by the vendor.

Each vSphere ESXi server had **three disk groups** with each consisting of 3TB NVMe devices for caching and four capacity tier 1.6TB SAS SSDs. The **raw capacity** of the vSAN datastore of this **3-node** vSAN cluster was for data around **58 TB**. The vSAN cluster limit is up to 64 cluster nodes.

Two 25 Gbit network adapters were used for vSAN and other network traffic in active-standby failover manner.

Review the network section for details about how the network cards impact the number of SAP HANA VMs per host and the sizing section for how many SAP HANA VMs can be deployed on a Skylake based server system.

Each ESXi Server in the vSAN Cluster had the following configuration as shown in Table 6. This configuration supports up to four SAP HANA VMs per HCI server.

Table 6. ESXi Server Configuration of the 4-socket Server SAP HANA on 3-node HCI vSAN Cluster

PROPERTY	SPECIFICATION
ESXi server model	HW Vendor SAP HANA Supported Cascade Lake supported Model
ESXi host CPU	Intel® Xeon®Platinum 8276L CPU @ 2.2GHz, 28 cores, 2 hosts of 8 socket 1 host of 4 Sockets
ESXi host RAM	3 TB
Network adapters	2 x 25 Gbit or better as a shared network for vSAN, SAP application and admin network traffic. Optional a dedicated NIC for the admin network. See the network section of this guide for recommended network configurations.
Storage adapter	SAS 12GBS SCSI controller
Storage configuration	Disk Groups: 3 Cache Tier: 3 x 3.2 TB P4610 NVMe cache disks Capacity Tier: 12 x 1.6 TB PM1643a SSD capacity disks

4-Socket/4-Disk Group SAP HANA Scale-Out HCI vSAN Cluster — Example 4

The **4-socket** test environment had four vSphere ESXi host servers in the vSAN cluster. Each host had 4 Intel Optane NVMe 750G for caching layer, 3TB memory and class M based CPU. The SSD based capacity drives provided the vSAN datastore for SAP HANA scale out configuration. Up to 8 hosts (7+1 HA) can be configured for the scale out environment.

The shared partition for NFS can be provided by vSAN data store with a VM running on the HA node or a separate NFS based storage. In this config, a separate NFS based storage was used for shared partition. BWH benchmark was run and a separate ESXi host running NetWeaver VM is used as a driver to do user throughput testing.

Each vSphere ESXi server had **four disk groups** each consisting of 1.6TB devices for caching and four capacity tier 4TB SAS SSDs.

Two 25 Gbit network adapters were used for vSAN and other network traffic in active-standby failover manner.

Review the network section for details on how the network cards impact the number of SAP HANA VMs per host and the sizing section for how many SAP HANA VMs can be deployed on a Skylake based server system.

Each ESXi Server in the vSAN Cluster had the following configuration as shown in Table 7. This configuration supported running up to four SAP HANA VMs per HCI server.

Table 7. SAP HANA Scale-Out Example Configuration

PROPERTY	SPECIFICATION
ESXi server model	HW Vendor SAP HANA Supported Cascade Lake Server Model
ESXi host CPU	4 x Intel(R) Xeon(R) Gold 6230 20C CPU 2.10GHz , 4 Sockets, 20 Physical Cores/Socket
ESXi host RAM	3 TB
Network adapters	2 x 25 Gbit or better as a shared network for vSAN, SAP application, and admin network traffic. Optional a dedicated NIC for the admin network. See the network section of this guide for the recommended network configurations.
Storage adapter	Dell HBA330 Adapter SAS 12GBS/SATA 6GBS SCSI controller
Storage configuration	Disk Groups: 4 Cache Tier: 4 NVMe p4610 1.6TB Capacity Tier: 16 x SAS p4510 4TB SFF
Linux LVM configuration	Linux LVM used with 8 VMDK for Data partition and 8 VMDK for Log partition

4-Socket/4-Disk Group on Cooper Lake SAP HANA HCI vSAN Cluster — Example 5

This 4-socket and 4-disk group test environment had a total of three vSphere ESXi servers configured and used an all-NVMe cache and capacity drives inside the hosts to provide a vSAN datastore that meets the SAP HANA HCI test requirements and KPIs.

The third node was part of the vSAN datastore and provided HA capabilities for the total 8 running SAP HANA VMs. Adding the fourth node would allow you to run more SAP HANA VMs on the vSAN datastore. The vSAN cluster limit is up to 64 cluster nodes. The vSAN cluster node number that an HCI vendor supports is determined by the vendor.

Each vSphere ESXi server had **four disk groups** with each consisting of one 750G Optane NVMe devices for caching and four capacity tier 1.6TB NVMe. The **raw capacity** of the vSAN datastore of this **3-node** vSAN cluster was for data around **58 TB**. The vSAN cluster limit is up to 64 cluster nodes.

Two 100 Gbit network adapters were used for vSAN and other network traffic in active-standby failover manner.

Review the network section for details about how the network cards impact the number of SAP HANA VMs per host and the sizing section for how many SAP HANA VMs can be deployed on a Skylake based server system.

Table 8. SAP HANA Cooper Lake Example Configuration

PROPERTY	SPECIFICATION
ESXi server model	HW Vendor SAP HANA Supported Cooper Lake Server Model
ESXi host CPU	4 x Intel(R) Xeon(R) Platinum 8376HL 28Cores CPU @ 2.6GHz, 4 Sockets, 28 Physical Cores/Socket
ESXi host RAM	1.5 TB
Network adapters	2 x 100 Gbit as a shared network for vSAN, SAP application, and admin network traffic. Optional a dedicated NIC for the admin network. See the network section of this guide for the recommended network configurations.
Storage configuration	Disk Groups: 4 Cache Tier: 4 x 750G Intel Optane SSD DC P4800x based PCIe adapters Capacity Tier: 16 x 1.6TB Intel SSD D7-P5600 based NVMe PCIe adapters
Linux LVM configuration	Linux LVM with 8 VMDK for Log partition

Software Resources

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

Table 9. Software Resources

SOFTWARE	VERSION	PURPOSE
<i>VMware vCenter Server</i>	6.7 U3 , 7.0 U1/U2/U3c	VMware vCenter Server provides a centralized platform for managing VMware vSphere environments.
<i>vSphere ESXi</i>	6.7 U3, 7.0 U1/U2/U3c	vSphere ESXi cluster to host virtual machines and provide vSAN Cluster.
<i>VMware vSAN</i>	6.7 U3, 7.0 U1/U2/U3c	Software-defined storage solution for hyperconverged infrastructure

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 SAP HANA VMs and which VMs are supported.

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

SAP HANA HCI Network Configuration

To build an SAP HANA ready HCI system, dedicated vSAN networks for SAP application, user traffic, admin and management are required. Follow the [SAP HANA Network Requirements](#) document to decide how many additional network cards have to 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 (all networks are dual port NICs, configured as failover teams)	Network Label	Admin	App-Server	vMotion	vSAN	Backup	Scale-Out	HANA-Repl	
	Recommended MTU Size	Default (1500)	Default (1500)	9000	9000	9000			
	Bandwidth	1 or 10 GbE	1 or 10 GbE	≥ 10 GbE*		≥ 10 GbE*			
	Physical NIC	0			1		OPTIONAL		
	Physical NIC port	1	2	1	2				
	VLAN ID#	100	101	200	201	202	203	204	

VM Guest Network Cards	Virtual NIC (inside VM)	0	1	-	-	2	2	3
------------------------	-------------------------	---	---	---	---	---	---	---

*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 exceeding 10 GE is needed if more than two SAP VMs should be deployed per host. 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 and low latency network for vSAN, it is also strongly recommended using a VMware vSphere Distributed Switch™ for all VMware 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 virtual machines 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 10 GbE** dual port network adapters for the teaming and failover purposes. A port group defines properties regarding security, traffic shaping, and NIC teaming. It is recommended to use the default port group setting except the uplink failover order that should be changed as shown in Table 11. It also shows the distributed switch port groups created for different functions and the respective active and standby uplink to balance traffic across the available uplinks.

Table 11 and Table 12 show an example of how to group the network port failover teams. It does not show the optional networks needed for VMware or SAP HANA system replication or Scale-Out internode networks. Additional network adapters are needed for these networks.

Table 11. Minimum ESXi Server Uplink Failover Network Configuration (if all NICs have the same bandwidth to allow physical NIC failure recovery)

<i>PROPERTY</i>	<i>VLAN*</i>	<i>ACTIVE UPLINK</i>	<i>STANDBY UPLINK</i>
<i>Admin plus VMware HA Network</i>	100	Nic0-Uplink1	Nic1-Uplink1
<i>SAP Application Server Network</i>	101	Nic0-Uplink2	Nic1-Uplink2
<i>vMotion Network</i>	200	Nic1-Uplink1	Nic0-Uplink1
<i>vSAN Network</i>	201	Nic1-Uplink2	Nic0-Uplink2

Table 12. Minimum ESXi Server Uplink Failover Network Configuration (if NICs have different bandwidth, with no physical NIC failover possibility)

<i>PROPERTY</i>	<i>VLAN*</i>	<i>ACTIVE UPLINK</i>	<i>STANDBY UPLINK</i>
<i>Admin plus VMware HA Network</i>	100	Nic0-Uplink1	Nic0-Uplink2
<i>SAP Application Server Network</i>	101	Nic0-Uplink2	Nic0-Uplink1
<i>vMotion Network</i>	200	Nic1-Uplink1	Nic1-Uplink2
<i>vSAN Network</i>	201	Nic1-Uplink2	Nic1-Uplink1

*VLAN id example, final VLAN numbers are determined by the network administrator.

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

VMware ESXi Server: Storage Controller Mode

Server storage or RAID controller normally support both pass-through and RAID mode.

In vSAN environments, only use controllers that support pass-through mode, as it is a requirement to deactivate RAID mode and use only pass-through mode.

vSAN Disk Group Configuration

Depending on the hardware configuration, like the number of CPUs and flash devices, minimal **two or more disk groups should be configured per host**. More disk groups improve the storage performance. The current disk group maximum is five disk groups per vSAN cluster node.

The cache disk type should be a low-latency high-bandwidth SAS SSD or an NVMe device. The capacity should be 600 GB or larger. Review VMware vSAN design considerations for [flash caching devices](#) and the [All-Flash Cache Ratio Update](#) document for details.

In our 2 and 4-socket server reference configuration, we used two and four disk group configurations to optimally use the server hardware components installed. The number of disk groups should be selected depending on how many SAP HANA VMs will be deployed on a single vSAN cluster node. In some test configurations, we noticed that a 2-disk group configuration with NVMe for caching and four capacity disks was enough to meet the SAP HANA KPIs for two SAP HANA VMs per server. **If four SAP HANA VMs should be deployed on a server, we recommend using a minimum of 3 or up to 4-disk group configuration. In addition, the PCIe lanes of the PCIe slots of the NVMe should be distributed evenly across all the CPUs. If multiple NVMe are connected on single CPU link, ensure that separate ports are used.** As noted, the HW vendors are free to decide which configuration they want to certify.

In our Cascade Lake 2-socket/2-disk group and 4-host server vSAN cluster configuration, we used a write optimized **2x 750 GB Intel Optane DC 4800X Series for caching** and 6 X 3.8TB 3840 Micron 5200 TCG-E SAS SSDs for the capacity tier per disk group. In total, we used 12 SSDs per server, all connected on one vSAN

certified SCSI adapter. With this vSAN configuration, we were able to validate two SAP HANA VMs per server successfully.

In the next 4-disk group configuration, we used 4-socket in a three node vSAN cluster configuration. We used an 800GB **NVMe device for caching** and four SAS 900GB SSDs for the capacity tier per disk group. With this vSAN configuration, we were able to validate 4 SAP HANA VMs per server. Additional tests with other configurations also found that fewer capacity SSDs are needed to pass the tests by using NVMe in the caching tier. The number of capacity SSDs depends mainly on the size of the SAP HANA VMs that will run on the HCI solution.

Figure 9 shows a single server of our 4-disk group test configuration in a graphical way.

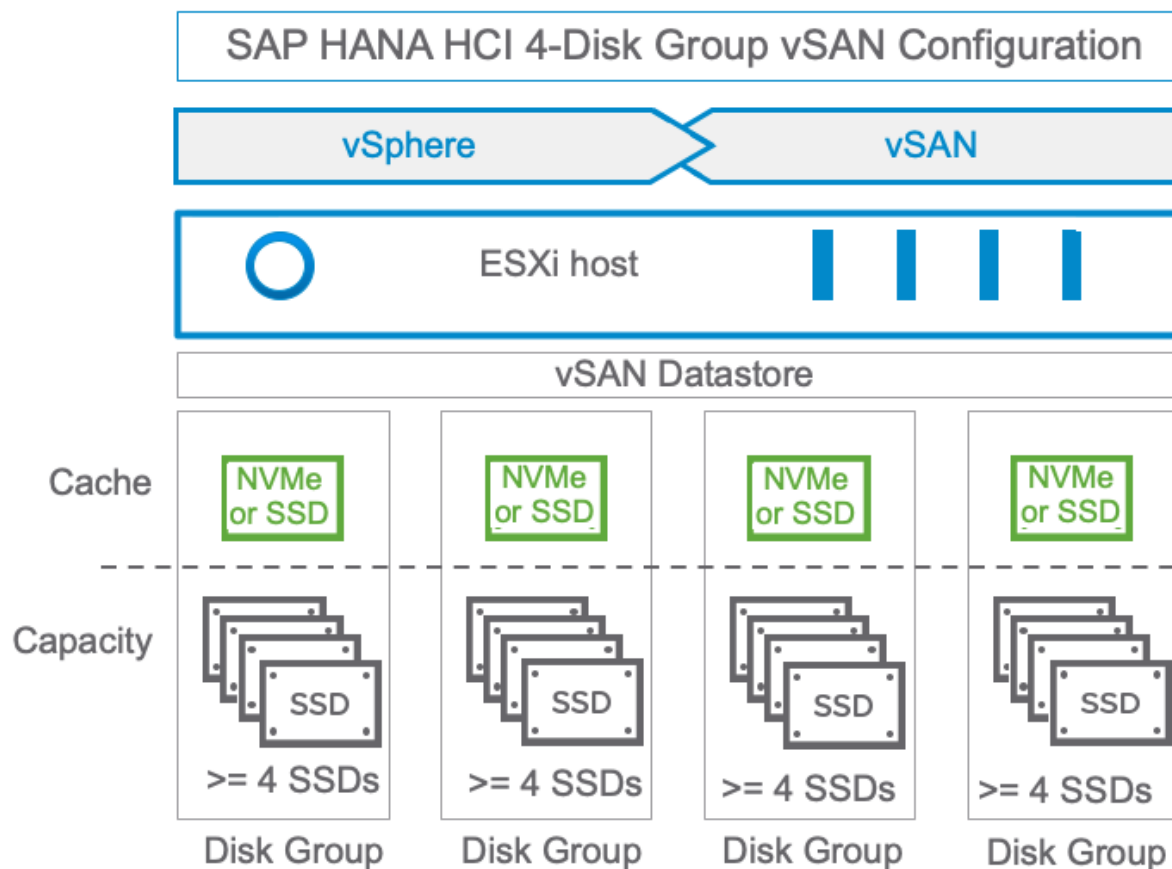


Figure 9. 4-disk Group Configuration Example of Our SAP HANA HCI Test Configuration

Replacing the SAS SSDs for caching with a higher bandwidth, low latency NVMe devices helped us to 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.

Note: If 4 or more SAP HANA VMs are supported per vSAN node, it is recommended to use at least a 4-disk group configuration.

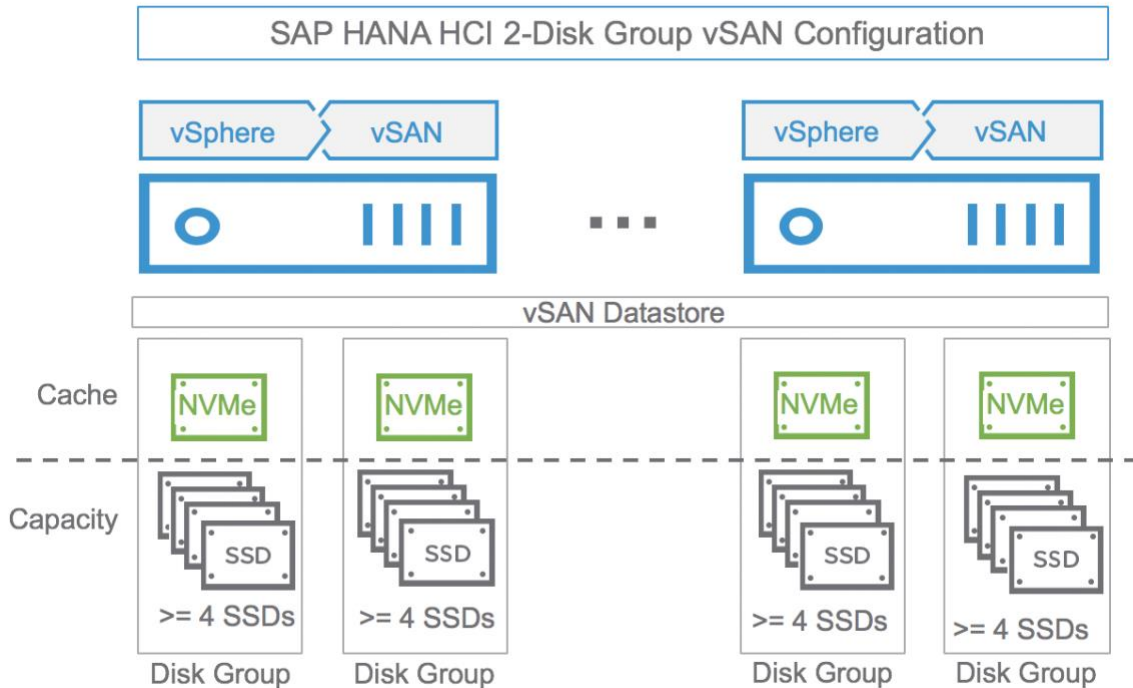


Figure 10. 2-disk Group Configuration of Our SAP HANA HCI Test Configuration

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 Storage Configuration

All vSAN disk groups are configured as a single cluster-wide shared datastore. The actual storage configuration for a VMDK file will be defined by a specific vSAN storage policy that fits its usage (OS, log, or data). For SAP HANA, we use the same storage policy for log and data. For the VMDK that contains the OS or SAP HANA binaries, the storage policy can be modified to use erasure coding to optimize capacity instead of performance.

vSAN Storage Policy

The SAP HANA HCI solution is a very specific configuration and is 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. These KPIs are validated, tested, and finally certified by SAP.

Depending on the actual vSAN hardware configuration, a vSAN storage policy is selected and applied to the SAP HANA OS, log and data disks. The vSAN policy is one of the parameters that determine the number of HANA VMs that can be supported on a server.

Table 13 and Table 14 show the vSAN storage policy we used for our testing. This storage policy should be applied on both the data disks and the log disk of each SAP HANA Database VM.

For non-performance critical volumes, the vSAN default storage policy or a custom [RAID 5 storage policy](#) can be used to lower the storage capacity need of these volumes.

For example, a VM protected by a primary level of failures to tolerate value of 1 with RAID 1 requires twice the virtual disk size, but with RAID 5, it requires 1.33 times the virtual disk size. For RAID 5, you need a minimum of 4 vSAN nodes per vSAN cluster.

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

Note: Only the listed vSAN storage policies got tested by VMware and got used during the vSAN SAP HANA validation.

Table 13. vSAN Storage Policy Settings for Data and Log for a Configuration with 12 or More Capacity Disks

STORAGE CAPABILITY	SETTING
Failure Tolerance Method (FTM)	RAID 1 (Mirroring)
Number of failures to tolerate (FTT)	1
Minimum hosts required	3
Number of disk stripes per object	12
Flash read cache reservation	0%
Object space reservation	100%
Disable object checksum	No

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

Table 14. vSAN Storage Policy Settings for Data and Log for a Configuration with 8 Capacity Disks

STORAGE CAPABILITY	SETTING
Failure Tolerance Method (FTM)	RAID 1 (Mirroring)

Number of FTT	1
Minimum hosts required	3
Number of disk stripes per object	8
Flash read cache reservation	0%
Object space reservation	100%
Disable object checksum	No

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

Note: You need a minimum of 4 vSAN nodes to use RAID 5, and a minimum of 6 vSAN nodes for RAID 6.

Table 15. vSAN Storage Setting for OS or Other Less Performance Critical Volumes

STORAGE CAPABILITY	SETTING
Failure tolerance method (FTM)	RAID 5/6 (Erasure Coding)
Number of FTT	1
Minimum hosts required	4
Number of disk stripes per object	1
Flash read cache reservation	0%
Object space reservation	0%
Disable object checksum	No

SAP HANA HCI Storage Planning and Configuration

When planning a virtualized SAP HANA environment running on vSAN, it is important to understand the storage requirements of such an environment and the applications running on it. It is necessary to define the storage capacity and IOPS needs of the planned workload.

Sizing a storage system for SAP HANA is very different from storage sizing for SAP classic applications. Unlike with SAP classic applications, where the storage subsystem is sized regarding the IOPS requirements of a specific SAP workload, the SAP HANA storage is defined by 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 is supporting on an SAP HANA HCI cluster, the vSAN datastore must be able to maintain the performance KPIs and must be able to provide the storage capacity of all installed SAP HANA VM on the vSAN cluster.

VMware has successfully tested a 4-socket 3-node vSAN based SAP HANA HCI configuration with an 8 SAP HANA VM configuration (4 SAP HANA VMs per one active node) and a 2-socket 4-node cluster with 2 SAP HANA VM per server configuration. These configurations represent one SAP HANA VM per socket, leaving the third or fourth node in the cluster used as an HA node. This leaves enough capacity available to support a full node failure 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. The tool and the documentation can be downloaded from SAP with a valid SAP user account. See [SAP note 2493172](#).

SAP HANA Storage Capacity Calculation and Configuration

All SAP HANA instances have a database log, data, root, local SAP, and shared SAP volume. 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 defined very strict performance KPIs that have to be met when configuring a storage subsystem. These storage performance KPIs are the leading sizing factor. How many capacity devices are used to provide the required I/O performance and latency depends on the HCI vendor hardware components. The capacity of the overall vSAN datastore must be large enough to store the SAP HANA data, log, and binaries by providing the minimal FTT=1 storage policy. SAP has published several SAP HANA specific architecture and sizing guidelines, like the [SAP HANA Storage Requirements](#). The essence of this guide is summarized in Figure 11 and Table 16, which can be used as a good start for planning the storage capacity needs of an SAP HANA system. It 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 the flexibility for the changing storage needs (the increased storage capacity most likely).

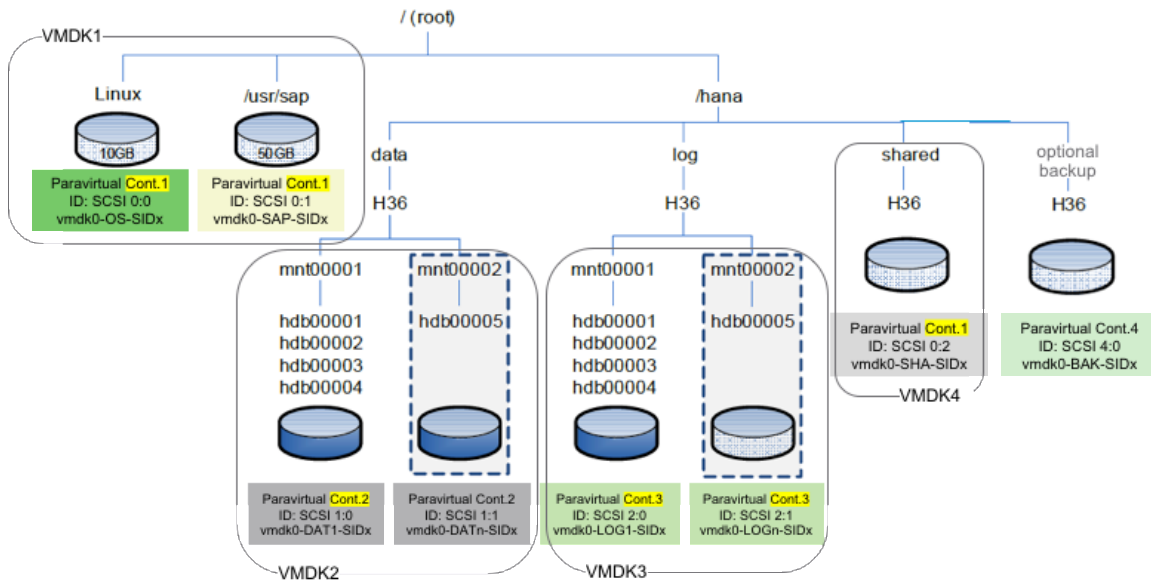


Figure 11. Storage Layout of an SAP HANA System. Figure © SAP SE, Modified by VMware

Table 16 determines the size of the different SAP HANA volumes. It is based on the [SAP HANA Storage Requirements](#) document. Some of the volumes like the OS and usr/sap volumes can be connected and served by one PVSCSI controller, others like the LOG and data volume are served by dedicated PVSCSI controllers to ensure high IO bandwidth and low latency.

Table 16. Storage Layout of an SAP HANA System

Volume	VMDK	SCSI Controller	VMDK Name	SCSI ID	Sizes as of 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	✓	PVSCSI Contr. 1	vmdk02-SHA-SIDx	SCSI 0:2	Minimum 1x RAM max. 1 TB
hana/data	✓	PVSCSI Contr. 2	vmdk03-DAT1-SIDx (vmdk03-DAT2-SIDx)	SCSI 1:0 (SCSI 1:1)	Minimum 1 x RAM Note: If you use multiple VMDKs, you can use

			vmdk03-DAT3-SIDx)	SCSI 1:2)	Linux LVM to build one large data disk.
hana/log	✓	PVSCSI Contr. 3	vmdk04-LOG1-SIDx (vmdk04-LOG2-SIDx)	SCSI 2:0 (SCSI 2:1)	[systems <= 512GB] log volume (min) = 0.5 x RAM [systems >= 512GB] log volume (min) = 512GB
hana/backup (optional)	✓	PVSCSI 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 a dedicated backup volume is used.

Figure 12 shows the default SAP HANA storage layout as documented in SAP’s HANA storage white paper, in a graphical view with the dedicated VMDKs per OS, HANA shared, log, and data volumes.

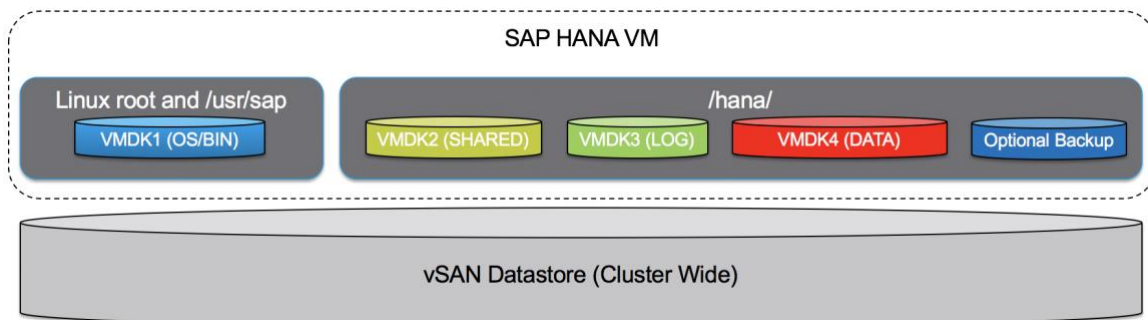


Figure 12. Dedicated VMDK Files per SAP HANA Volume

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

Usage of Linux Logical Volume Manager (LVM)

Although we did not use Linux LVM with our 4-socket, 3 node cluster tests, the usage of LVM is supported and can be used if desired. Using Linux LVM can help to leverage more vSAN resources and balance the VMDK components distribution on vSAN. On the other hand, using LVM is adding configuration complexity. It is determined by the HCI solution vendor to use LVM or not.

Figure 13 shows an example how LVM can be used for the SAP HANA LOG and DATA volumes. We used this configuration with our 4-disk group SAS cache and 2-disk group configurations. Linux LVM used with 8 VMDK for data partition and a min 4 to 8 VMDK for log partition.

Note that the number of VMDK disks used for LVM should not exceed the capacity disks used for vSAN in a host. For example, if 6 capacity disks are used in a server, up to 6 VMDKs disks for LVM would be ideal.

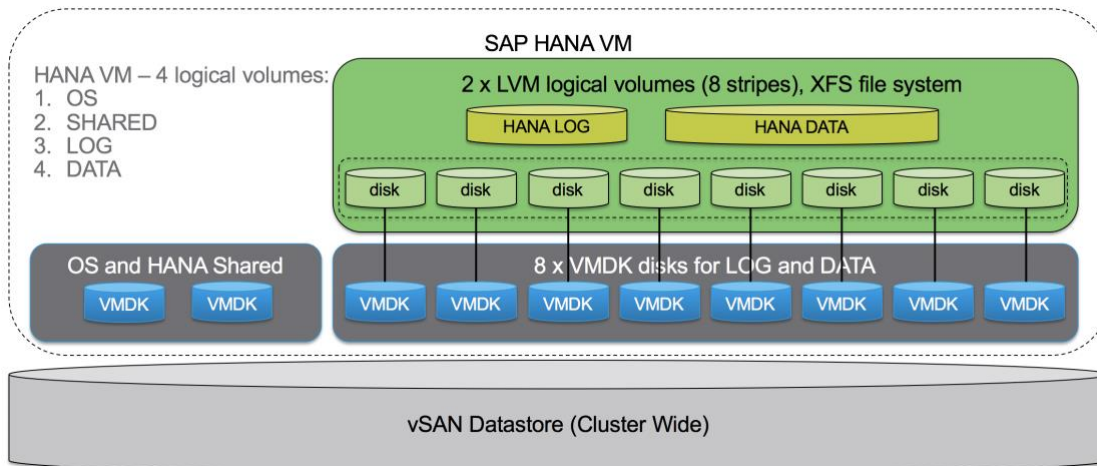


Figure 13. Usage of LVM Managed Volumes

vSAN Storage Capacity Calculation

To determine the overall storage capacity per SAP HANA VM, summarize the sizes of all specific and unique SAP HANA volumes as outlined in Figure 11 and Table 16.

To determine the overall vSAN datastore capacity required, without cache tier, it is required to multiply the above SAP HANA volume requirements with the amount of VMs that can run on the SAP HANA HCI cluster.

Simplified calculation:

vSAN datastore capacity = total HANA VM x (OS + OSR/SAP + SHARED + DATA + LOG) x 2 (FTT=1, RAID 1) + 30% Slack space.

Example:

A cluster with 3 servers with **18 TB RAM and 8 SAP HANA VMs in total**, one server is HA -> 12 TB of SAP HANA data needs to be stored.

$$8 \times (10 \text{ GB} + 60 \text{ GB} + 1 \text{ TB} + 12 \text{ TB} / 8 \text{ VMs} + 512 \text{ GB}) \times 2 \times 1.3 = 8 \times 3,142 \text{ GB} \times 2 \times 1.3 = \sim 65.4 \text{ TB}$$

We have over 138 TB of vSAN raw datastore capacity available with our example 3-node 4-socket SAP HANA HCI configuration, which is much more than the amount required for up to 8 SAP HANA VMs with a total up to 12 TB of SAP HANA system RAM capacity. The additional space can be used to store optional backups. This lowers the TCO since no additional nearline backup storage is needed. Nevertheless, moving the backup files to a permanent external backup location is recommended and required.

SAP HANA HCI VMware vSAN Detection Script

SAP requested to implement a method to detect if 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 the runtime if an SAP HANA VM is running on a vSAN datastore.

The VMware vSAN detection script is available on request from VMware. Contact the author or contributors of this document for the script. The script needs to be installed by the SAP HANA HCI solution vendor on their VMware based SAP HANA HCI solution systems.

The script is a Python script that will be executed every 5 minutes on the vSAN hosts that build the SAP HANA HCI solution and will add the following information to the VM configuration file:

- "guestinfo.vsan.enabled": True/False,
- "guestinfo.SDS.solution": esxi_version, which is also the vSAN release version
- "guestinfo.vm_on_vsan": True/False

This data will be parsed by the SAP CIM provider and SAP support tools like SAPsysinfo, and it detects if an SAP HANA HCI on VMware vSAN solution is used for this specific VM. Details about the VMware SAP key monitoring metrics can be found in OSS Note [2266266](#)—Key Monitoring Metrics for SAP on VMware vSphere (version 5.5 u3 and higher).

Prerequisites:

1. Enable SSH for the time of installation. Deactivate it after the installation of the script.
2. To simplify the installation process, it is recommended to have identical password of the super user on all the ESXi servers.

Installation and execution of the health check installation script

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

```
root@Centos [ ~/hana_health_check ]# ./install_vib_script_remotely.sh -h
Please specify the 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:

```
[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
```

Functions and scenarios of the script

The SAP HANA HCI health check Python script is executed every 5 minutes, if you want to shorten the time to 1 min, modify 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 info from each individual VM:

```
vsan.enabled = True/False
sds.solution = ESXI/vSAN_Version
vm_on_vsan = True/False
```

Testing of the script

1. Install an SAP HANA supported Linux version based VM on one of your host servers.
2. Ensure that the latest Open VMware guest tools are installed on your 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 execute it on your test VM.
5. Example output (search for vSAN in the document):
 - a. -----
 - b. VMware vSAN info:
 - c.
 - d. guestinfo.vsan.enabled =
 - e. guestinfo.SDS.solution =
 - f. guestinfo.vm_on_vsan =
 - g.
 - h. -----
 - i. Linux system information:
 - j.
 - k. -----
6. If you do not get the above outputs, check the prerequisites and ensure that the used hypervisor is a VMware hypervisor and the used SDS solution is vSAN.

If a VM is moved to vSAN (or have a newly created vmdk on vSAN), or moved out from vSAN to other datastores, the guestinfo will be updated within 5 mins.

If the VM is storage vMotioned to another host that does not have the script running, you need to 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 14 for details:

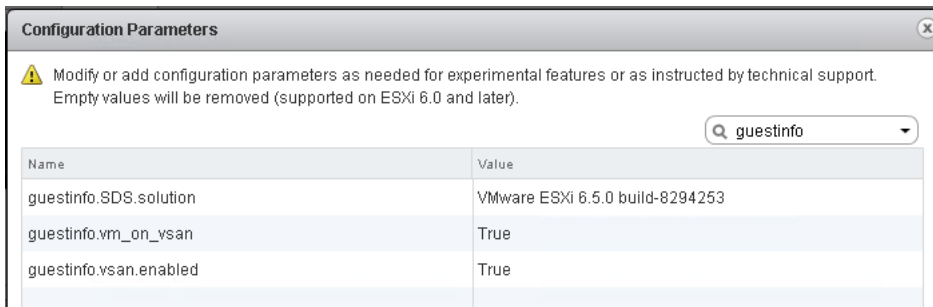


Figure 14. SAP HCI vSAN Specific Configuration Parameters

Updating the script

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

```
[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
```

SAP HANA HCI on vSAN—Support and Process

Support Status

vSAN is supported as part of an SAP HANA HCI vendor-specific solution, as described in this document. See SAP Note [2718982](#).

The SAP HANA HCI solution vendor is the primary support contact for the customer with escalation possibilities to VMware and SAP. The support of an SAP HANA HCI solution is comparable and very similar to the SAP HANA Appliance support model.

The SAP HANA HCI solution vendor must define the overall usable storage capacity and performance available for SAP HANA VMs. The vendor also defines how many SAP HANA VMs per host are supported and how many vSAN cluster nodes are supported by the vendor.

For supported SAP HANA HCI solutions based on VMware products, check out [SAP HANA Certified HCI Solution](#).

Support Process

If you want to escalate an issue with your SAP HANA HCI solution, work directly with your HCI vendor and follow the defined and agreed support process, which normally starts by opening a support ticket within the SAP support tools.

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 and can be used, in addition to the HCI vendor specific support contract, with an active support contract running vSphere 6.5 U1d or later.

How does it work?

See [How your Experiences Improve Products and Services](#) for details.

Conclusion

With vSAN, VMware equips its hardware and OEM partners to build SAP HANA ready HCI solutions. Not only compute, memory, and network, but also internal storage components can be leveraged from a standard x86 server, which dramatically changes the SAP HANA infrastructure economics and the deployment

methodology. With SAP HANA HCI powered by VMware solutions, an SAP hardware partner vendor can build true turnkey SAP HANA infrastructure solutions offering customers to deploy SAP HANA in days instead of weeks or months.

Appendix

Optimizing the SAP HANA on vSphere Configuration Parameter List

VMware vSphere can run a single large or multiple smaller SAP HANA virtual machines on a single physical host. This section describes how to optimally configure a VMware virtualized SAP HANA environment. These parameters are valid for SAP HANA VMs running vSphere and vSAN based SAP HANA HCI configurations.

The listed parameter settings are the recommended BIOS settings for the physical server, the ESXi host, the VM, and the Linux OS to achieve optimal operational readiness and stable performance for SAP HANA on vSphere.

The parameter settings described in this section are the default settings that should be always configured for SAP HANA on vSphere. The settings described in the [Performance optimization for low-latency SAP HANA VMs section](#) should only be applied in rare situations where SAP HANA must perform with the lowest latency possible.

The shown parameters are the best practice configuration parameters and in case of an escalation, the VMware or HCI vendor SAP support engineers will verify and, if not applied, will recommend configuring these settings.

Physical host BIOS parameter settings	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 fully supports EFI since version 5.0.
Enable Intel VT technology	Enable all BIOS virtualization technology settings.
Configure RAM hemisphere mode	Distribute DIMM or PMem modules in a way to achieve 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	<p>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.</p> <p>See the Timekeeping in VMware Virtual Machines information guide for reference.</p> <p>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.</p>
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 support to	Set to HW vendor default
Enable hyperthreading	Always enable hyperthreading on the ESXi host. This will double the logical CPU cores to allow ESXi to take advantage of more available CPU threads.
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	<p>Follow the vendor documentation and enable PMem for the usage with ESXi. Note: Only App Direct and Memory mode are supported with production-level VMs.</p> <p>Memory mode is only supported with a ratio of 1:4. As of today, SAP provides only non-production workload support.</p> <p>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 vSphere hosts equipped with Intel Optane PMem in App Direct mode and SAP HANA VMs can leverage PMem according to SAP note 2913410. Please especially note that the vSphere HA restriction (as described in SAP note 2913410) applies and needs to be considered.</p>
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 port, floppy drive, CD-ROM, and USB.

ESXi host parameter settings	Description
Networking	<p>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 dedicated 10 GbE for vMotion and the SAP app server or replication networks. At least 25 GbE for vMotion for SAP HANA system >= 2TB is recommended.</p>
Settings to lower the virtual VMXNET3 network latencies	<p>Set the following settings on the ESXi host. For this to take effect, the ESXi host needs to be rebooted.</p> <p>Procedure: Go to the ESXi console and set the following parameter(s):</p> <ul style="list-style-type: none"> vsish -e set /config/Net/intOpts/NetNetqRxQueueFeatPairEnable 0 <p>Add the following advanced VMX configuration parameters to the VMX file, and reboot the VM after adding these parameters:</p> <ul style="list-style-type: none"> ethernetX.pnicFeatures = "4" ethernetX.ctxPerDev = "3" <p>Change the rx-usec, lro and rx / tx values of the VMXNET3 OS driver, and of the NIC used for SAP database to app server traffic, from the default value of 250 to 75 (25 is the lowest usable setting). Procedure: Log on to OS running the inside the VM and use ethtool to change the following settings, then execute:</p> <ul style="list-style-type: none"> ethtool -C ethX rx-usec 75 ethtool -K ethX lro off ethtool -G ethX rx 512 rx-mini 0 tx 512 <p>Note: Exchange X with the actual number, such as eth0. To make these ethtool settings permanent, see SLES KB 000017259 or the RHEL ethtool document.</p>
Storage configuration	<p>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. Use the storage layout as a template as explained in this guide.</p> <p>Set the following settings on the ESXi host. For this to take effect, the ESXi host needs to be rebooted.</p> <p>Procedure: Go to the ESXi console and set the following parameter(s):</p> <ul style="list-style-type: none"> vsish -e set /config/Disk/intOpts/VSCSIPollPeriod 100 <p>If you want to use vSAN, then select one of the certified SAP HANA HCI solutions based on vSAN and follow the VMware HCI BP guide.</p>
SAP monitoring	<p>Enable SAP monitoring on the host -> Misc.GuestLibAllowHostInfo = "1"</p> <p>For more details, see SAP note 1409604.</p> <p>Without this parameter, no host performance relevant data will be viewable inside an SAP monitoring enabled VM.</p>

SAP HANA virtual machine parameter settings	Description
Tips how to Edit the *.vmx file	Review tips for editing a *.vmx file in VMware KB 1714 .
UEFI BIOS guest	<p>It is recommended to use UEFI BIOS as the standard BIOS version for vSphere hosts and guests. Features such as Secure Boot are possible only with EFI.</p> <p>See VMware DOC-28494 for details. You can configure this with the vSphereClient™ by choosing EFI boot mode.</p> <p>If you are using vSphere 6.0 and you only see 2TB memory in the guest, then upgrade to the latest ESXi 6.0 version.</p>
SAP monitoring	<p>Enable SAP monitoring inside the SAP HANA VM with the advanced VM configuration parameter <code>tools.guestlib.enableHostInfo = "TRUE"</code>.</p> <p>For more details, see SAP note 1409604.</p> <p>Besides setting this parameter, the VMware guest tools need to be installed. For details, see VMware KB 1014294.</p>
vCPU hotplug	<p>Ensure that vCPU hotplug is deactivated, otherwise vNUMA is disabled and SAP HANA will have a negative performance impact. For details, see VMware KB 2040375.</p>
Memory reservations	<p>Set fixed memory reservations for SAP HANA VMs. Do not overcommit memory resources.</p> <p>It is required to reserve memory for the ESXi host. It is recommended to reserve, depending on the amount of CPU sockets memory for ESXi.</p> <p>Typical memory reservation for a host is between 32–64GB for a 2-socket server, 64–128GB for a 4-socket server, and 128–256GB for an 8-socket server. These are not absolute figures as the memory need of ESXi depends strongly on the actual hardware, ESXi and VM configuration, and enabled ESXi features, such as vSAN.</p>
CPU	<p>Do not overcommit CPU resources and configure dedicated CPU resources per SAP HANA VM.</p> <p>You can use hyperthreads when you configure the virtual machine to gain additional performance. For CPU generations older than Cascade Lake, you should consider disabling hyperthreading due to the Intel Vulnerability Foreshadow L1 Terminal Fault. For details, read VMware KB 55636.</p> <p>If you want to use hyperthreads, then you must configure 2x the cores per CPU socket of a VM (for example, 2-socket wide VM on a 28 core Cascade Lake system will require 112 vCPUs).</p>

vNUMA nodes	<p>SAP HANA on vSphere can be configured to leverage half-CPU and full-CPU sockets. A half-CPU socket is configured by only half of the available physical cores of a CPU socket in the VM configuration GUI.</p> <p>The vNUMA nodes of the VM will always be ≥ 1, depending on how many CPUs you have configured in total.</p> <p>If you need to access an additional NUMA node, use all CPU cores of this additional NUMA node. Nevertheless, use as few NUMA nodes as possible to optimize memory access.</p>
-------------	--

SAP HANA virtual machine parameter settings	Description
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 virtual CPUs to leverage 14 cores and 14 hyperthreads per CPU socket. A full-socket VM should also be configured to use 56 vCPUs to leverage all 28 physical CPU cores and available hyperthreads per socket.
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 driver for I/O devices	Use the dedicated SCSI controllers for OS, log and data to separate disk/I/O streams. For details, see the SAP HANA disk layout section.
Use the virtual machine's file system	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 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 .
Eager zero thick virtual disks for data and log disk	We recommend this setting as it avoids lazy zeroing (initial write penalty).
VMXNET3	Use paravirtual VMXNET 3 virtual NICs for SAP HANA virtual machines. We recommend at least 3–4 different NICs inside a VM in the HANA VM (app/ management server network, backup network, and, if needed, HANA system replication network). Corresponding physical NICs inside the host are required.
Optimize the application server network latency if required	Disable virtual interrupt coalescing for VMXNET 3 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 vMotion or SAP HANA system replication. Use the advanced options in the vSphere Web Client or directly modify the .vmx file and add ethernetX.coalescingScheme = "disable". X stands for your network card number. For details, see the Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs white paper.

Set lat.Sensitivity = normal	<p>Check with the vSphere Client and ensure that, in the VM configuration, the value of Latency Sensitivity Settings is set to “normal.” If you must change this setting, restart the VM.</p> <p>Do not change this setting to “high” or “low.” Change this setting under the instruction of VMware support engineers.</p>
------------------------------	--

SAP HANA virtual machine parameter settings	Description
Associate virtual machines with specified NUMA nodes to optimize NUMA memory locality	<p>Associating a NUMA node with a virtual machine to specify the NUMA node affinity is constraining the set of NUMA nodes on which NUMA can schedule a virtual machine’s virtual CPU and memory.</p> <p>Use the vSphere Web Client or directly modify the .vmx file and add numa. nodeAffinity=x (for example: 0,1).</p> <p>Note: This is only needed if the VM is < the available CPU sockets of the ESXi host. For half-socket SAP HANA VM configurations, use the VMX parameter sched.vCPUx.affinity as documented in the next section instead.</p> <p>Procedure:</p> <ol style="list-style-type: none"> 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 numa.nodeAffinity. 8. In the Value column, enter the NUMA nodes where the virtual machine can be scheduled. Use a comma-separated list for multiple nodes. For example, enter 0,1 to constrain the virtual machine resource scheduling to NUMA nodes 0 and 1. 9. Click OK. 10. Click OK to close the Edit VM dialog box. <p>Attention: When you constrain NUMA node affinities, you might interfere with the ability of the ESXi NUMA scheduler to rebalance virtual machines across NUMA nodes for fairness. Specify the NUMA node affinity only after you consider the rebalancing issues.</p> <p>For details, see Associate Virtual Machines with Specified NUMA Nodes.</p>

SAP HANA virtual machine parameter settings	Description
<p>Configure virtual machines to use hyperthreading with NUMA</p>	<p>For memory latency-sensitive workloads with low processor utilization, such as SAP HANA, or high interthread communication, we recommended using hyperthreading with fewer NUMA nodes instead of full physical cores spread over multiple NUMA nodes. Use hyperthreading and enforce NUMA node locality per VMware KB 2003582.</p> <p>This parameter is only required when hyperthreading should be leveraged for a VM. Using hyperthreading can increase the compute throughput but may increase the latency of threads.</p> <p>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 gets configured automatically.</p> <p>Use the vSphere Web Client and add the following advanced VM parameter:</p> <p>numa.vcpu.preferHT="TRUE" (per VM setting) or as a global setting on the host: Numa.PreferHT="1" (host).</p> <p>Note: For non-mitigated CPUs, such as Haswell, Broadwell and Skylake, you may consider not to use hyperthreading at all. For details, see VMware KB 55806.</p>
<p>PMem-enabled VMs</p>	<p>To configure an Optane PMem-enabled SAP HANA VM for optimal performance, it is necessary to align the VM configuration to the underlying hardware, especially the NUMA configuration.</p> <p>VMware KB 78094 provides information on how to configure the NVDIMMs (VMware's representation of Optane PMem) correctly and align the NVDIMMs to the physical NUMA architecture of the physical server.</p> <p>By default, Optane PMem allocation in vmkernel for VM NVDIMMs does not consider NUMA. This can result in the VM running on a certain NUMA node and Optane PMem allocated from a different NUMA node. This will cause NVDIMMs access in the VM to be remote, resulting in poor performance.</p> <p>To solve this, you must add the following settings to a VM configuration using vCenter.</p> <p>Example for a 4-socket wide VM:</p> <ul style="list-style-type: none"> • nvdim0:0.nodeAffinity=0 • nvdim0:1.nodeAffinity=1 • nvdim0:2.nodeAffinity=2 • nvdim0:3.nodeAffinity=3 <p>sched.pmem.prealloc=TRUE is an optional parameter equivalent to eager zero thick provisioning of VMDKs and improves initial writes to Optane PMem.</p> <p>Besides these parameters, the CPU NUMA node affinity or CPU affinities must also be configured.</p>

Remove unused devices	Remove unused devices, such as floppy disks or CD-ROM, to release resources and to mitigate possible errors.
Linux OS parameter settings	Description
Linux version	VMware strongly recommends using only the SAP HANA supported Linux and kernel versions. See SAP note 2235581 and settings, see note 2684254. Use SAP HANA SAPConf/SAPTune to optimize the Linux OS for SAP HANA.
To optimize large-scale workloads with intensive I/O patterns, change the queue depths of the SCSI default values	<p>The large-scale workloads with intensive I/O patterns require adapter queue depths greater than the Paravirtual SCSI (PVSCSI) default values. The default values of PVSCSI queue depth are 64 (for device) and 254 (for adapter). You can increase PVSCSI queue depths to 254 (for device) and 1024 (for adapter) inside a Windows virtual machine or Linux virtual machine.</p> <p>Create a file of any name in the <code>/etc/modprobe.d/</code> directory with this line:</p> <ul style="list-style-type: none"> options vmw_pvscsi cmd_per_lun=254 ring_pages=32 <p>Note: For RHEL5, edit <code>/etc/modprobe.conf</code> with the same line. Make a new <code>initrd</code> for the settings to take effect. You can do this either by using <code>mkinitrd</code>, or by re-running <code>vmware-config-tools.pl</code>.</p> <p>Starting in version 6, RHEL uses <code>modprobe.d</code>.</p> <p>Alternatively, append these to kernel boot arguments (for example, on Red Hat Enterprise Linux edit <code>/etc/grub.conf</code> or on Ubuntu edit <code>/boot/grub/grub.cfg</code>).</p> <ul style="list-style-type: none"> vmw_pvscsi.cmd_per_lun=254 vmw_pvscsi.ring_pages=32 <p>Reboot the virtual machine. See VMware KB 2053145 for details.</p> <p>Note: Review the VMware KB article 2088157 to ensure that the minimum VMware patch level is used to avoid possible virtual machine freezes under heavy I/O load.</p>
Install the latest version of VMware Tools	VMware Tools is a suite of utilities, which enhances the performance of the virtual machine's guest operating system and improves the vm management. See http://kb .vmware com/kb/1014294 for details.
Configure NTP time server	Use the same external NTP server as configured for vSphere. For details, see SAP note 989963 .
Optional: Disable large receive offload (LRO) in the Linux guest OS to lower latency for client/application server- facing NIC adapter	<p>This helps to lower network latency of client/application server facing NIC adapters run: <code>"ethtool -K ethY lro off"</code>.</p> <p>Do not disable LRO for throughput NIC adapters such as for backup, replication, or SAP HANA internode communication networks.</p> <p>Works only with Linux kernel 2.6.24 and later and uses a VMXNET3.</p> <p>Additional details: http://kb.vmware.com/kb/2055140</p>

Linux OS parameter settings	Description
General SAP HANA Linux configuration recommendations (These settings will get automatically get configured when you use SAP Tune/SAPConf)	
Linux Operating System with SAP HANA Reference Guide	See the recommended operating system configuration settings for running SAP HANA on Linux in the Reference Guide .
Disable I/O scheduling	SLES15SP2 onwards this is configured by default, scheduler is set to none with block_mq enabled.
Disable AutoNUMA	Later Linux kernel (RHEL 7 and SLES 12) supporting auto- migration according to NUMA statistics. For SLES: # yast bootloader, choose "Kernel Parameters" tab (ALT-k) and edit the "Optional Commandline Parameters" section by appending numa_balancing=disabled For RHEL add "kernel.numa_balancing = 0" to /etc/sysctl.d/sap_hana.conf and reconfigure the kernel by running: # sysctl -p /etc/sysctl.d/sap_hana.conf
Use Block mq	In the kernel parameters add scsi_mod.use_blk_mq=1. For OS version Sles15 SP2 and beyond this is enabled by default.
Disable transparent HugePages	THP is not supported for the use with SAP HANA DB, as it may lead to hanging situations and performance degradations. To check the current configuration, run the following command: # cat /sys/kernel/mm/transparent_hugepage/enabled Its output should read: always madvise [never] If this is not the case, you can disable the THP usage at runtime by issuing the following command: # echo never > /sys/kernel/mm/ transparent_hugepage/enabled For details, refer to the SAP WIKI for SAPs and the Linux OS vendors virtualization independent recommended Linux OS settings for SAP HANA.
Change the following parameters in /etc/sysctl.conf (important for SAP HANA Scale-Out deployments)	<pre> net.core.rmem_default = 262144 net.core.wmem_max = 8388608 net.core.wmem_default = 262144 net.core.rmem_max = 8388608 net.ipv4.tcp_rmem = 4096 87380 8388608 net.ipv4.tcp_wmem = 4096 65536 8388608 net.ipv4.tcp_mem = 8388608 8388608 8388608 net.ipv4.tcp_slow_start_after_idle = 0 </pre>

Example Linux kernel boot loader parameters	intel_idle.max_cstate=0 processor.max_cstate=0 numa_balancing=disabled transparent_hugepage=never elevator=noop vmw_pvscsi.cmd_per_lun=254 vmw_pvscsi.ring_pages=32
---	---

Performance Optimization for Low Latency SAP HANA VMs

Further optimization of virtual SAP HANA performance can be required when SAP HANA must perform as close to bare metal as possible and with the shortest latency in terms of database access times. When optimizing SAP HANA for low latency, we recommend sizing an SAP HANA VM with the least number of NUMA nodes. When an SAP HANA VM needs more CPU or RAM than a single NUMA node provides, configure an additional NUMA node and its resources.

To achieve the optimal performance for an SAP HANA virtual machine, use the settings as described in the next table in addition to the previously described settings. In terms of CPU scheduling and priority, these settings improve performance by reducing the amount of vCPU and vNUMA migration, while increasing the priority of the SAP HANA production virtual machine.

CPU affinities

By specifying a CPU affinity setting for each virtual machine, you can restrict the assignment of virtual machines to a subset of the available processors (CPU cores) in multiprocessor systems. By using this feature, you can assign each virtual machine to processors in the specified affinity set.

Setting CPU affinities can decrease the CPU and memory latency by not allowing the ESXi scheduler to migrate VM threads to other logical processors. Setting CPU affinities is required when configuring SAP HANA half-socket VMs.

Before you use a CPU affinity, you need to take the following items into consideration:

- For multiprocessor systems, ESXi systems perform automatic load balancing. Avoid the manual specification of virtual machine affinity to improve the scheduler's ability to balance load across processors.
- An affinity can interfere with the ESXi host's ability to meet the reservation and shares specified for a virtual machine.
- Because CPU admission control does not consider affinities, a virtual machine with manual affinity settings might not always receive its full reservation. Virtual machines that do not have manual affinity settings are not adversely affected by virtual machines with manual affinity settings.
- When you move a virtual machine from one host to another, an 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 virtual machine that is already assigned to the certain processors using an affinity.

- An affinity setting can affect the host's ability to schedule virtual machines on multicore or hyperthreaded processors to take full advantage of resources shared on such processors.

For more information about performance practices, see the [vSphere Resource Management Guide](#) as well as the VMware documentation around [specifying NUMA controls](#).

Additional Performance Tuning Settings for SAP HANA Workloads

Additional Performance Tuning Settings for SAP HANA Workloads

Note: The following are optional parameters that are only needed for the lowest CPU latency. Set these parameters with caution.

SAP HANA virtual machine parameter settings	Description
Tips about how to edit the *.vmx file	Review the tips for editing a *.vmx file in VMware KB 1714 .
monitor.idleLoopSpinBeforeHalt = "true" and monitor.idleLoopMinSpinUS = "xx us"	<p>Setting these advanced VM parameters can help improve performance of a VM at the cost of CPU time on the ESXi host and should only be configured for an SAP HANA workload that runs as the only workload on a NUMA node/compute server.</p> <p>Edit the .vmx file and add the following two advanced parameters: monitor.idleLoopSpinBeforeHalt = "true" AND monitor.idleLoopMinSpinUS = "xx" (For example: 50)</p> <p>Both parameters must be configured to influence the de-scheduling time.</p> <p>Background: The guest OS issues a Halt instruction, which stops (or de-schedules) the vCPU on the ESXi host. Keeping the virtual machine spinning longer before Halt negates the number of inter-processor wake-up requests.</p>
Set monitor_control.halt_in_monitor = "TRUE"	<p>In the default configuration of ESX 7.0, the idle state of guest HLT instruction will be emulated without leaving the VM if a vCPU has an exclusive affinity.</p> <p>If the affinity is non-exclusive, the guest HLT will be emulated in vmkernel, which may result in having a vCPU de-scheduled from the physical CPU, and can lead to longer latencies. Therefore, it is recommended to set this parameter to "TRUE" to ensure that the HLT instruction gets emulated inside the VM and not in the vmkernel.</p> <p>Use the vSphere Web Client and add the following advanced VM parameter: monitor_control.halt_in_monitor = "TRUE".</p>
Set monitor_control.disable_pause_loop_exiting = "TRUE"	This parameter prevents the VM from exiting to the hypervisor unnecessarily during a pause instruction. This is specific for Intel Skylake CPU-based systems.

SAP HANA virtual machine parameter settings	Description
<p>Configuring CPU affinity</p> <p>sched.vcpuXx.affinity = "Yy-Zz"</p> <p>For more information about potential performance practices, see vSphere Resource Management Guide.</p>	<p>Note: Remove the numa.nodeAffinity settings if set and if the CPU affinities with sched.vCPUxxx.affiity are used.</p> <p>By specifying a CPU affinity setting for each virtual machine, you can restrict the assignment of virtual machines to a subset of the available processors (CPU cores) in multiprocessor systems. By using this feature, you can assign each virtual machine to processors in the specified affinity set.</p> <p>See Scheduler operation when using the CPU Affinity (2145719) for details.</p> <p>This is especially required when configuring so-called SAP HANA "half-socket" VM's or for very latency critical SAP HANA VMs. It is also required when parameter numa.slit.enable gets used.</p> <p>Just like with numa.NodeAffinity it is possible to decrease the CPU and memory latency by further limit the ESXi scheduler to migrate VM threads to other logical processors / CPU threads by leveraging the sched.vCPUxxx.affinity VMX parameter in contrast to parameter numa.NodeAffinity it is possible assign a vCPU to a specify physical CPU thread and is for instance necessary for instance when configuring half-socket SAP HANA VMs.</p> <p>Use the vSphere Web 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.</p> <p>Procedure:</p> <ol style="list-style-type: none"> 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 want to assign to a physical CPU thread). 8. In the Value column, enter the physical CPU threads where the vCPU can be scheduled. E.g. enter 0-55 to constrain the virtual machine resource scheduling to physical CPU threads 0-55, which would be the 1st CPU of an 28 core CPU host. 9. Click OK. 10. Click OK to close the Edit VM dialog box.
<p>>4-socket VM on 8-socket hosts</p>	<p>Add the advanced parameter: numa.slit.enable = "TRUE" to ensure the correct NUMA map for VMs > 4 socket on 8-socket hosts.</p> <p>Note: sched.vcpuXx.affinity = "Yy-Zz" must get configured when numa.slit.enable is set to "TRUE".</p>

Example VMX Configurations for SAP HANA VMs

The following examples provide an overview about how to set the **additional** VMX parameters for SAP HANA half- and full-CPU socket VMs. **These parameters should be added in addition to the standard parameters** via the vSphere Web Client or by directly adding these parameters to the .vmx file with a text editor.

SAP HANA VM additional VMX Parameters	Settings
First half-socket VM on socket 0 on a 28-core CPU n-socket server	<ul style="list-style-type: none"> • 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"
First half-socket PMem VM on socket 0 on a 28-core CPU n-socket server	<ul style="list-style-type: none"> • nvdimm0:0.nodeAffinity=0 • 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 a 28-core n-socket CPU server	<ul style="list-style-type: none"> • 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"
SAP HANA VM additional VMX Parameters	Settings

<p>Second half-socket PMem VM on socket 1 on a 28-core CPU n-socket server</p>	<ul style="list-style-type: none"> • nvdim0:0.nodeAffinity=1 • numa.vcpu.preferHT="TRUE" • sched.vcpu0.affinity = "84-111" • sched.vcpu1.affinity = "84-111" • sched.vcpu2.affinity = "84-111" • ... • sched.vcpu26.affinity = "84-111" • sched.vcpu27.affinity = "84-111"
<p>1-socket VM on socket 3 on a 28-core CPU 4 or 8-socket server</p>	<ul style="list-style-type: none"> • numa.vcpu.preferHT="TRUE" • numa.nodeAffinity=3
<p>1-socket PMem VM on socket 3 on a 28-core CPU 4 or 8-socket server</p>	<ul style="list-style-type: none"> • nvdim0:0.nodeAffinity=3 • numa.vcpu.preferHT="TRUE" • numa.nodeAffinity=3
<p>2-socket VM on sockets 0 and 1 on a 28-core CPU n-socket server</p>	<ul style="list-style-type: none"> • numa.vcpu.preferHT="TRUE" • numa.nodeAffinity=0,1
<p>2-socket PMem VM on sockets 0 and 1 on a 28-core CPU n-socket server</p>	<ul style="list-style-type: none"> • nvdim0:0.nodeAffinity=0 • nvdim0:1.nodeAffinity=1 • numa.vcpu.preferHT="TRUE" • numa.nodeAffinity=0,1
<p>3-socket VM on sockets 0, 1 and 2 on a 28-core CPU 4 or 8-socket server</p>	<ul style="list-style-type: none"> • numa.vcpu.preferHT="TRUE" • numa.nodeAffinity=0,1,2
<p>3-socket PMem VM on sockets 0, 1 and 2 on a 28-core CPU 4 or 8-socket server</p>	<ul style="list-style-type: none"> • nvdim0:0.nodeAffinity=0 • nvdim0:1.nodeAffinity=1 • nvdim0:2.nodeAffinity=2 • numa.vcpu.preferHT="TRUE" • numa.nodeAffinity=0,1,2
<p>SAP HANA VM additional VMX Parameters</p>	<p>Settings</p>

<p>4-socket VM on a 28-core CPU 8-socket server running on sockets 0–3</p>	<ul style="list-style-type: none"> • numa.slit.enable = "TRUE" • sched.vcpu0.affinity = "0-55" • sched.vcpu1.affinity = "0-55" • sched.vcpu2.affinity = "0-55" • ... • sched.vcpu222.affinity = "168-223" • sched.vcpu223.affinity = "168-223"
<p>6-socket VM on a 28-core CPU 8-socket server running on sockets 0–5</p>	<ul style="list-style-type: none"> • numa.slit.enable = "TRUE" • sched.vcpu0.affinity = "0-55" • sched.vcpu1.affinity = "0-55" • sched.vcpu2.affinity = "0-55" • ... • sched.vcpu334.affinity = "280-335" • sched.vcpu335.affinity = "280-335"
<p>8-socket VM on a 28-core CPU 8-socket server</p>	<ul style="list-style-type: none"> • numa.slit.enable = "TRUE" • sched.vcpu0.affinity = "0-55" • sched.vcpu1.affinity = "0-55" • sched.vcpu2.affinity = "0-55" • ... • sched.vcpu446.affinity = "392-447" • sched.vcpu447.affinity = "1392-447"

<p>SAP HANA low-latency VM additional VMX Parameters</p>	<p>Settings</p>
<p>n-socket <u>low-latency</u> VM on an n-socket, CPU server (Valid for all VMs when even a lower latency is required.)</p>	<ul style="list-style-type: none"> • monitor.idleLoopSpinBeforeHalt = "TRUE" • monitor.idleLoopMinSpinUS = "50" • monitor_control.disable_pause_loop_exiting = "TRUE" (when Skylake)

CPU Thread Matrix

The following table shows the CPU thread matrix of a 28 core CPU as a reference when configuring the sched.vCPUXx.affinty="Xx-Yy" parameter. The list shows the starting and end range required for the "Xx-Yy" parameter. For example, for CPU 5 it would be "280-335".

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.affinty ="Xx-Yy" parameter. The list shows the starting and end range required for the "Xx-Yy" parameter. For example, for CPU 5 it would be "240-287".

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
Total CPU Cores	96	24	24	24	24	24	24	24	

The following table shows the CPU thread matrix of a 22 core CPU as a reference when configuring the sched.vCPUXx.affinty = "Xx-Yy" parameter. The list shows the starting and end range required for the "Xx-Yy" parameter. For example, for CPU 5 it would be "220-263".

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

The following table shows the CPU thread matrix of a 18 core CPU as a reference when configuring the sched.vCPUx.affinity="Xx-Yy" parameter. The list shows the starting and end range required for the "Xx-Yy" parameter. For example, for CPU 3 it would be "108-143".

18 Core, 36 Thread CPU					
		CPU0 Threads	CPU1 Threads	CPU2 Threads	CPU3 Threads
Half-Socket	0	36	72	108	
	1	37	73	109	
	2	38	74	110	
	3	39	75	111	
	4	40	76	112	
	5	41	77	113	
	6	42	78	114	
	7	43	79	115	
	8	44	80	116	
	9	45	81	117	
	10	46	82	118	
	11	47	83	119	
	12	48	84	120	
	13	49	85	121	
	14	50	86	122	
	15	51	87	123	
	16	52	88	124	
	17	53	89	125	
Half-Socket	18	54	90	126	
	19	55	91	127	
	20	56	92	128	
	21	57	93	129	
	22	58	94	130	
	23	59	95	131	
	24	60	96	132	
	25	61	97	133	
	26	62	98	134	
	27	63	99	135	
	28	64	100	136	
	29	65	101	137	
	30	66	102	138	
	31	67	103	139	
	32	68	104	140	
	33	69	105	141	
	34	70	106	142	
	35	71	107	143	
Total CPU Threads	144	36	36	36	36
Total CPU Cores	72	18	18	18	18

Monitoring and Verifying an SAP HANA Installation

- [SAP Note 1698281](#) provides information about how you can monitor the data growth and the utilization of actual memory. With this, it is also possible to detect and diagnose the memory leaks during operation.
- [SAP Note 1969700](#) covers all the major HANA configuration checks and presents a tabular output with configurations that are changed. The collection of SQL statements is very helpful in checking and identifying parameters that are configured and conflict with the SAP recommended configuration parameters.

References

- [VMware Docs webpage](#)
- [VMware Tech Zone](#)
- [vSAN Support Insight](#)
- [SAP HANA on VMware vSphere](#)
- [SAP on VMware vSAN](#)
- [SAP HANA on VMware vSphere Best Practices Guide](#)
- [The Winding Road to Virtual SAP HANA Application Workload Guidance Design for SAP S/4HANA on VMware vSphere 6.5](#)
- [SAP Notes related to VMware](#)
- [vSAN SAP support note 2718982 - SAP HANA on VMware vSphere and vSAN](#)
- [SAP HANA Certified Hyperconverged Infrastructure Solutions \(vSAN\)](#)

Revision History

We included the following updates for vSphere 7.0 U2/U3c:

- Cooper Lake support, review SAP note [2393917](#) for an updated list of supported vSphere versions.
- Added vCLS information.
- Stretched cluster deployment, see [vSphere Clustering Service \(vCLS\)](#)
- **Starting with vSphere 7.0 Update 1**, vSphere Cluster Service (vCLS) is enabled by upgrading the vCenter to 7.0U1 and runs 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 the vision.

The dependency of these cluster services on vCenter is not ideal and VMware introduced in the vSphere 7 Update 1 release 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. If vCenter Server becomes unavailable, in future, vCLS will ensure that the cluster services remain available to maintain the resources and health of the workloads that run in the clusters.

vCLS Deployment Guidelines for SAP HANA Landscapes

As of SAP notes 2937606 and 3102813, it is not supported to run a non-SAP HANA VM on the same NUMA node where a 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, as long as 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, it is necessary, because of 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 get migrated to hosts that do not run SAP HANA production level VMs.

This can get 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 get assured that vCLS VMs and SAP HANA VMs do not get co-deployed.

After the policy is created and tags were 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 get added to the cluster then the **Retreat Mode** can get used to remove the vCLS VMs from this cluster. Please 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 hosts that do not run SAP HANA tagged VMs. **If this is not possible then you can use the Retreat Mode to remove the vCLS VMs from the cluster.** In the case of VCF an additional node to offload the vCLS VMs is required per cluster (e.g. the HA node).

For more details and vCLS deployment examples in the context of SAP HANA please refer to this [blog](#).

- VMware vSAN Stretched Cluster Support for SAP HANA Systems.
- 8-Socket support, see [8-Socket/4-socket 3 Disk Group SAP HANA–HCI vSAN Cluster—Example 3](#).
- New sizing samples and CPU Thread Matrix examples.
- Quality of Service configuration guidance, see [IOPS Limit per Partition Policy—Quality of Service Option to Storage Throughput](#).

Authors

- Erik Rieger, Principal SAP Global Technical Alliance Manager and Architect, VMware Global SAP Alliance
- Sathya Krishnaswamy, VMware Staff SAP Performance and Certification Engineer, Performance Engineering

Contributors

- Sebastian Lenz, VMware Staff SAP Performance and Certification Engineer, Performance Engineering
- Louis Barton, VMware Staff SAP Performance and Certification Engineer, Performance Engineering
- Chen Wei, Senior Manager of Workload Technical Marketing, VMware

- Catherine Xu, Manager of Workload Technical Marketing, VMware