



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

BEST PRACTICES AND REFERENCE ARCHITECTURE GUIDE

Table of Contents

Abstract	3
Executive Summary	3
Audience	3
Solution Overview	4
SAP HANA on vSAN	4
Solution Components	5
SAP HANA HCI on vSAN—Support Status as of Note 2718982.....	6
Reference Architecture	7
Scope.....	7
SAP HANA HCI on vSAN—Architecture Overview	7
SAP HANA HCI on vSAN—Scalability and Sizes	8
VMware HA to Protect Scale-Up and Scale-Out SAP HANA Systems	12
Reference Test System Configuration	14
Reference Environment	14
Hardware Resources	15
Software Resources.....	18
SAP HANA HCI on vSAN Configuration Guidelines	19
SAP HANA HCI Network Configuration	19
VMware ESXi Server: Storage Controller Mode	21
vSAN Disk Group Configuration	21
vSAN Storage Configuration	23
vSAN Storage Policy	23
SAP HANA HCI Storage Planning and Configuration.....	25
SAP HANA Storage Capacity Calculation and Configuration	25
Usage of Linux Logical Volume Manager (LVM)	27
vSAN Storage Capacity Calculation	28
SAP HANA HCI VMware vSAN Detection Script	29
SAP HANA HCI on vSAN—Support and Process	32
Support Status	32
Support Process	32
vSAN Support Insight.....	32
Conclusion	32
Appendix	33

Best Practices of Virtualized SAP HANA for Skylake Based Server Host Systems	33
Optimizing SAP HANA on vSphere	33
Performance Optimization for Extreme Performance/Latency Critical SAP HANA VMs	40
CPU Affinities	40
Additional Performance Tunings for Maximum Performance	41
Monitoring and Verifying an SAP HANA Installation	43
References	44
Author and Contributors	44

Abstract

This document provides best practices and architecture guidelines to design and build an SAP HANA environment on Hyperconverged Infrastructure (HCI) solutions based on 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” in the [SAP Integration and Certification Center \(ICC\)](#).

This paper also includes information about the support processes and requirements of SAP HANA on vSAN.

Executive Summary

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](#). An SAP valid s-user is required to view the note.

Audience

This paper is intended for SAP and VMware hardware partners that 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.

Following figure summarizes the benefits of VMware based SPA 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 6.5 update 2 or vSphere 6.7, see SAP note [2393917](#) for details.
- VMware vSAN Enterprise 6.6 or 6.7, see SAP note [2718982](#) for details.
- VMware vCenter 6.5 and later.
- 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 the [Certified HCI Solutions](#) section on the SAP HANA Supported and Certified Hardware Directory section.

Figure 1 provides an overview of the solution components.

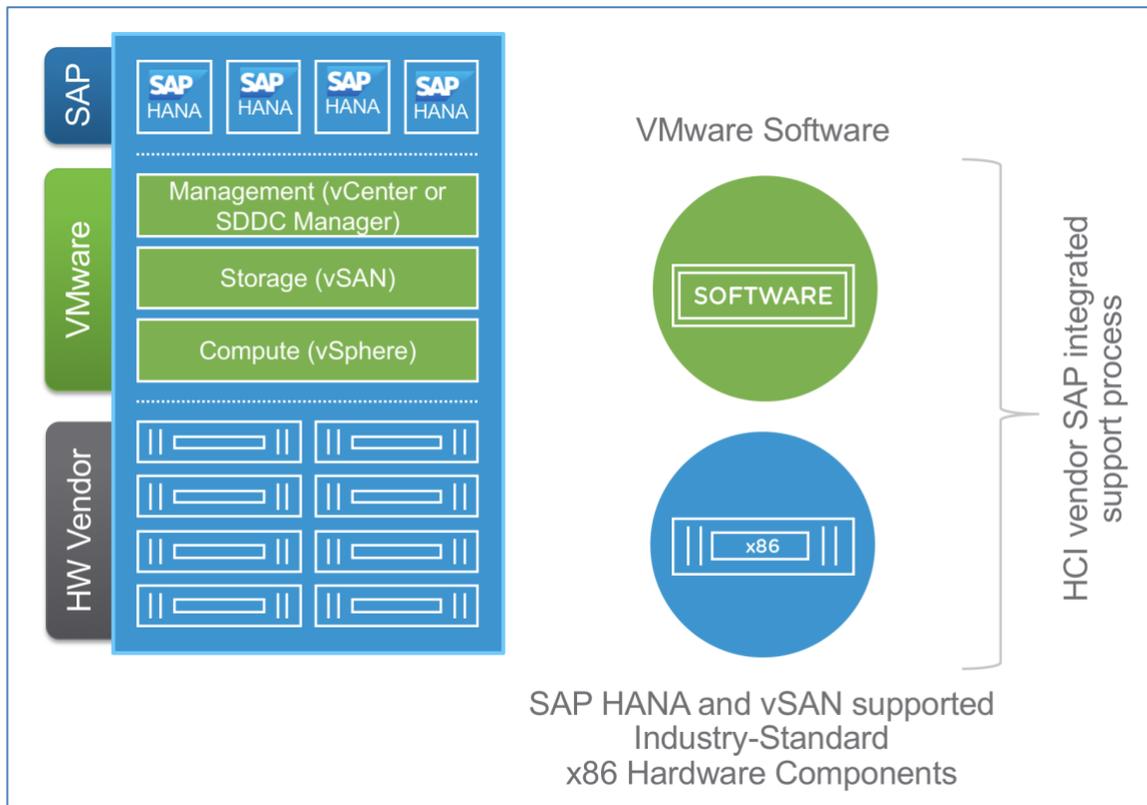


Figure 1. SAP HANA HCI Solution Overview

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 and 6.7**
 - SAP HANA SPS 12 PL 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.
- **VMware vSAN 6.6 and 6.7**
 - The HCI vendor can select either the 6.6 or 6.7 version of vSAN for their HCI solutions.
 - The vSAN versions 6.6 and 6.7 are included in vSphere version 6.5 U2 and vSphere version 6.7 respectively.
 - As of today, a maximum of 4 SAP HANA VMs are supported per 2, 4 and 8 socket vSAN cluster nodes.
 - No CPU sockets need to be reserved for vSAN
 - Scale-out is currently not supported with VMware vSAN.

- **CPU architecture:**
 - Intel Skylake processors 2,4, 8-socket servers as of TDIv5 Model with the existing VMware vSphere virtualization boundaries.
 - The certified Hardware for HCI solutions can be found on the [Certified Hyperconverged Infrastructure Solutions](#).

Note: This SAP note will be continuously updated with the latest support information without further notice. Above support status is from November 2018.

Reference Architecture

The purpose of the introduced vSAN for SAP HANA reference architectures is to provide a starting point for HCI solution vendors that 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 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.

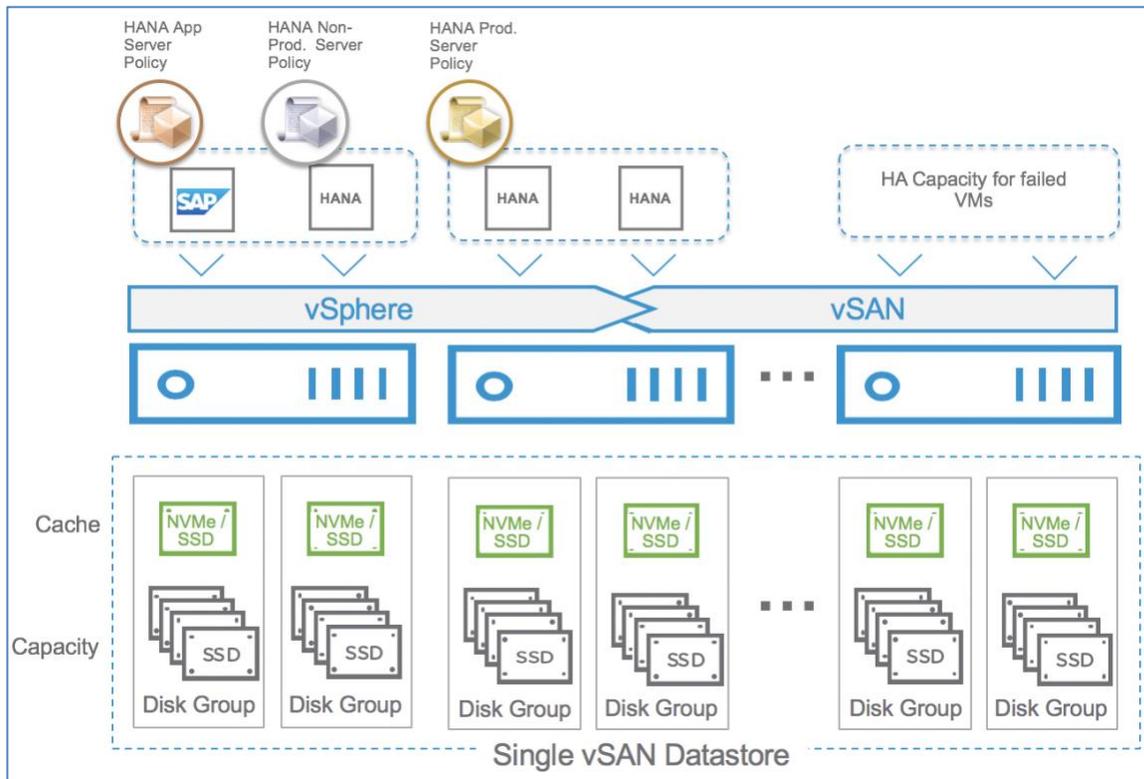


Figure 2. SAP HANA Applications on vSAN High-Level Architecture

[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 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 sizes of SAP HANA systems. 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 is 4 socket SAP HANA VMs with 6 TB of RAM and up to 128 vCPUs. Larger than 4 socket VMs require an 8-socket server and are a subject of an expert sizing and special SAP approval.

vSphere 6.5 and 6.7 provide the following maximums per physical host. For SAP HANA, only servers with up to 8 physical CPU sockets are currently supported.

Table 1. vSphere Host Maximums

	ESXi 6.7	ESXi 6.5
<i>Logical CPUs per host</i>	768	576
<i>Virtual machines per host</i>	1,024	
<i>Virtual CPUs per host</i>	4,096	
<i>Virtual CPUs per core</i>	32	
<i>RAM per host</i>	16 TB	12 TB
<i>NUMA Nodes 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 like virtual disk size or number of virtual NICs per VM.

Table 2. vSphere Guest (VM) Maximums

	ESXi 6.5 and 6.7
<i>Virtual CPUs per virtual machine</i>	128
<i>RAM per virtual machine</i>	6,128 GB
<i>Virtual SCSI adapters per virtual machine</i>	4
<i>Virtual NVMe adapters per virtual machine</i>	4
<i>Virtual disk size</i>	62 TB
<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. Up to 64 cluster nodes are supported by vSAN.

Table 3. vSAN Maximums

<i>VMware vSAN 6.6 and 6.7</i>	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.

As of SAP note [2393917](#), the maximum size of a single SAP HANA VM is 6 TB of RAM and 128 vCPUs. SAP HANA VM can be configured as half-socket, full-socket, 2, 3 or 4 sockets. 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 on our vSAN configuration. For detailed information, see the Reference Test System Configuration section.

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.

Table 4. SAP HANA on VMware Powered HCI Solution Maximums Based on the Selected Intel Skylake SP CPUs without Hyperthreading

	Intel Xeon Gold 6140 CPU	Intel Xeon Platinum 8176 CPU	Intel Xeon Platinum 8180M CPU
Max supported RAM per CPU as of Intel datasheet	768 GB	768 GB	1536 GB
CPU Cores per socket as of Intel datasheet	18	28	28
Max. NUMA Nodes per ESXi Server	2	8 ¹	8 ¹
0.5 Socket SAP HANA VM (Half-Socket)	1 to 4 x 9 physical core VM with min. 128 GB RAM and max. 384GB ²	1 to 16 x 14 physical core VM with min. 128 GB RAM and max. 384 GB ²	1 to 16 x 14 physical core VM with min. 128 GB RAM and max. 768 GB ²
	vSAPS ⁴ 14,400, 9 vCPUs	vSAPS ⁴ 28,000, 14 vCPUs	vSAPS ⁴ 32,000, 14 vCPUs
1 Socket SAP HANA VM	1 to 2 x 18 physical core VM with min. 128 GB RAM and max. 768 GB	1 to 8 x 28 physical core VM with min. 128 GB RAM and max. 768 GB	1 to 8 x 28 physical core VM with min. 128 GB RAM and max. 1536 GB
	vSAPS ³ 28,500, 18 vCPUs	vSAPS ³ 56,000, 28 vCPUs	vSAPS ³ 64,000, 28 vCPUs
2 Socket SAP HANA VM	1 x 36 physical core VM with min. 128 GB RAM and max. 1536 GB	1 to 4 x 56 physical core VM with min. 128 GB RAM and max. 1536 GB	1 to 4 x 56 physical core VM with min. 128 GB RAM and max. 3072 GB
	vSAPS ³ 57,000, 36 vCPUs	vSAPS ³ 112,000, 56 vCPUs	vSAPS ³ 128,000, 56 vCPUs
3 Socket SAP HANA VM	–	1 to 2 x 84 physical core VM with min. 128 GB RAM and max. 2304 GB	1 to 2 x 84 physical core VM with min. 128 GB RAM and max. 4608 GB
	–	vSAPS ³ 168,000 84 vCPUs	vSAPS ³ 193,000, 84 vCPUs
4 Socket SAP HANA VM	–	1 to 2 x 112 physical core VM with min. 128 GB RAM and max. 3072 GB	1 to 2 x 112 physical core VM with min. 128 GB RAM and max. 6128 GB
	–	vSAPS ³ 224,000, 112 vCPUs	vSAPS ³ 257,000, 112 vCPUs

Notes:

¹A maximum of 128 vCPUs can be configured per VM with vSphere 6.5 and 6.7. A Skylake SP Platinum 8180 processor provides up to 112 physical CPU cores. An 8-socket SKL server has up to 224 physical CPU cores. For 6 TB VM configurations, an Intel Skylake SP 8180M CPU, which supports up to 1.5 TB RAM per CPU is required. Intel Xeon Gold or the Intel Xeon Platinum 8180 model supports only up to 768 GB per CPU.

²The half-socket RAM figures listed in the table show even configured half-socket VM RAM sizes. The maximum RAM of a half-socket configuration can be as large as maximum RAM installed per NUMA node minus 128 GB. For example, 768 GB per NUMA node of an SKL CPU, the first half socket VM has 128 GB while the second half-socket VM can have up to 640 GB RAM.

³vSAPS figures without 20% hyperthreading gain and minus 10% virtualization costs. Hyperthreading gain for real live SAP HANA workloads is normally much lower as the assumed 20% gain and virtualization costs are normally lower than 10%. The shown figures are rounded figures and based on the average and rounded SAPS performance figures per CPU from published [SAP SD benchmarks](#).

The SAP'S figures in table 4 are calculated values without hyperthreading. To leverage hyperthreading, it is required to configure two vCPUs per physical CPU core. Due to the Intel L1 terminal fault, VMware does not recommend using hyperthreading in VM configurations.

Note: Review VMware KB 55767 for details about the performance impact leveraging or not leveraging hyperthreading. As described, especially with low CPU utilized systems, the performance impact of leveraging hyperthreading is very little.

As discussed in KB 55767, hyperthreading gain can be as low as 1% for low utilized systems and as high as 32% for highly utilized servers.

Table 5. KB55767 Test Results to Investigate the OLTP Database Workload Gain of Hyperthreading

<i>Application Workload / Guest OS</i>	% Host CPU Utilization before Enabling the ESXi SCA Scheduler	Remaining Usable Host CPU Capacity before enabling ESXi SCA Scheduler	Performance Degradation after Enabling the ESXi SCA Scheduler
<i>OLTP Database / Linux</i>	90%	0%	32%
	77%	~15%	20%
	62%	~30%	1%

Note: In this example, the peak application throughput is reached at 90% host CPU utilization. Enabling the SCA scheduler as show in the table above means that the virtual machine does not use CPU hyperthreading.

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

VMware vSphere HA delivers a vSphere built-in high availability solution, which can be used to protect SAP HANA. VMware HA provides uniform, 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.

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.

Note: See the [SAP HANA on VMware vSphere best practice guide](#) 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 virtualized 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 3 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.

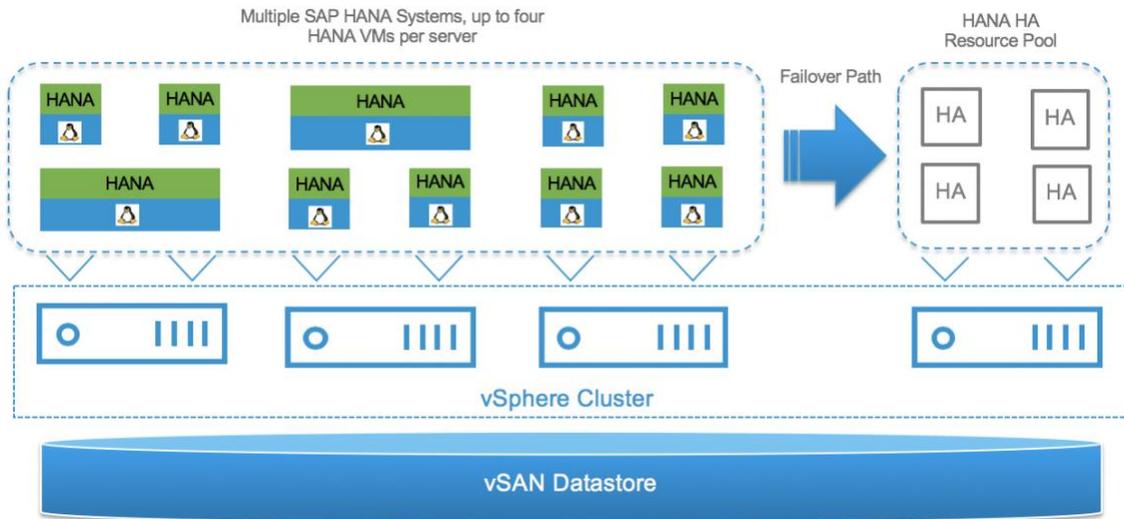


Figure 3. VMware HA Protected SAP HANA VMs in an n+1 vSAN Cluster Configuration

Figure 4 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.

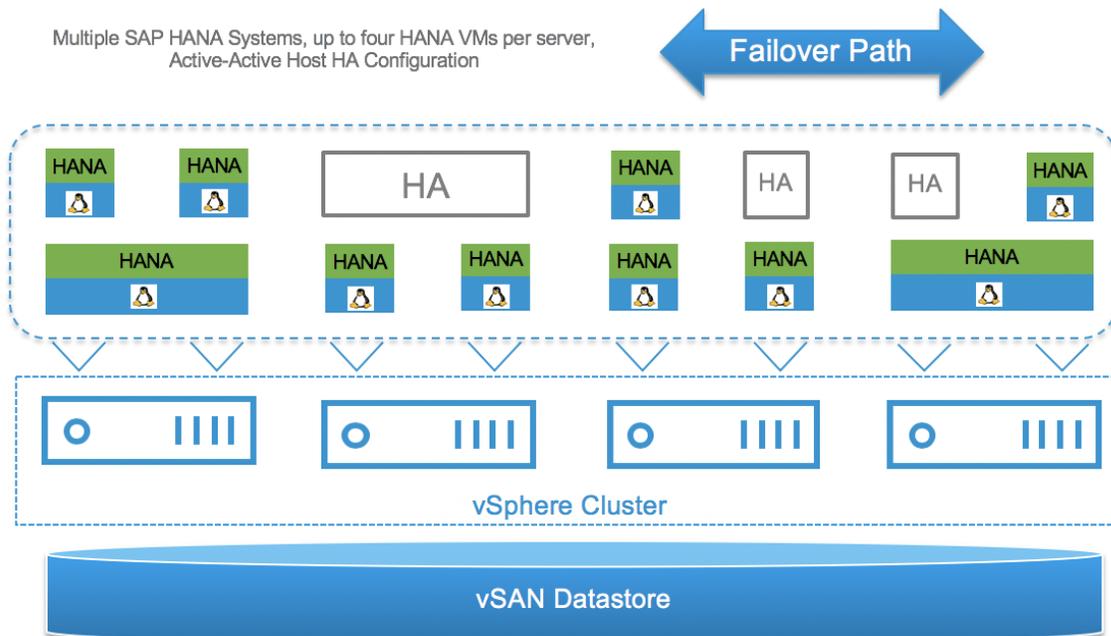


Figure 4. VMware HA Active-Active n+1 vSAN Cluster Configuration

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 5 shows a configuration of 3 SAP HANA VMs (1 x master with 2 worker nodes) running exclusively on the virtual SAP HANA HCI platform.

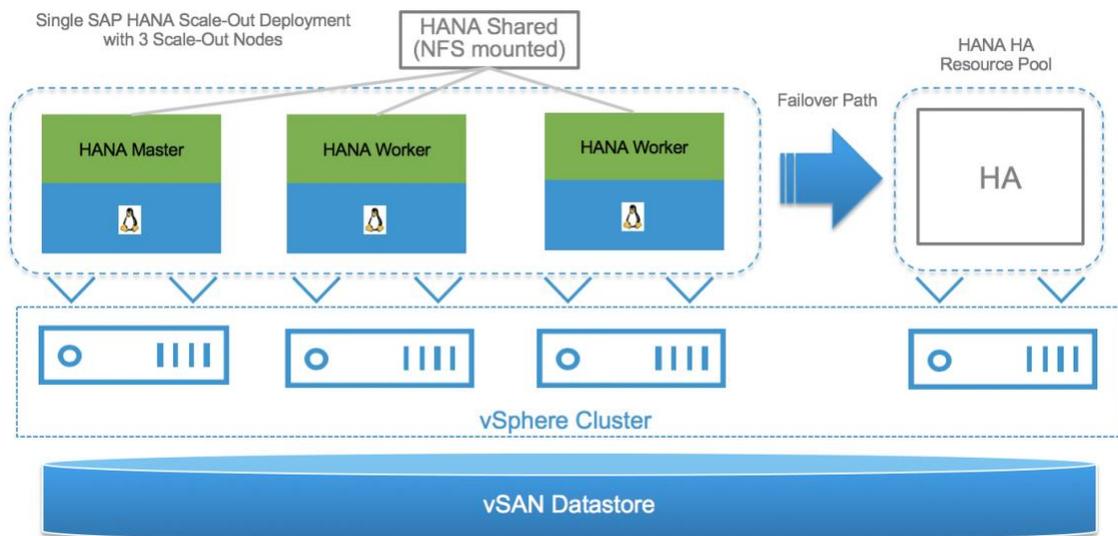


Figure 5. SAP HANA Scale-out VMware HA n+1 vSAN Cluster Configuration

Unlike physical deployments, there are no dependencies on external components, such as DNS servers, SAP HANA Storage Connector API, or STONIT 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 fail over 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.

Reference Test System Configuration

This section introduces the resources and configurations for the solution including:

- Reference environment
- 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 test scenario). These test tools validate performance, scalability, and consistency.

We validated two different testing environments, which were based on 2- and 4-socket SAP HANA and vSAN supported Skylake based server systems, configured as a 3-host and a 4-host cluster with two or four vSAN disk groups per server. We also tested with NVMe and SAS SSD devices for caching and with both devices, it is possible to meet the SAP HANA KPIs.

Hardware Resources

Two vSAN clusters in the 3-host and 4-host configuration respectively:

- One vSAN SAP HANA HCI environment with four 2-socket based Skylake SP servers configured with two and four vSAN disk groups.
- One vSAN SAP HANA HCI environment with three 4-socket based Skylake SP servers configured with four vSAN disk groups.

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

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 for data around 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 Table 6.

Table 6. ESXi Server Configuration of the 2-socket/2-disk Server SAP HANA 4-node HCI vSAN Cluster

PROPERTY	SPECIFICATION
ESXi server model	HW Vendor SAP HANA Supported Skylake Server Model
ESXi host CPU	2 x Intel(R) Xeon(R) Gold 6130 CPU , 2 Sockets, 18 Physical Cores/Socket.
ESXi host RAM	768 GB (384 GB per CPU)
Network adapters	2 x 25 Gbit as a shared network for vSAN, SAP application and admin network traffic. See the network section of this guide for the recommended network configurations. SAP HANA HCI Network Configuration
Storage adapter	1 x 8 channel PCI-EXPRESS 3.0 SAS 12GBS/SATA 6GBS SCSI controller
Storage configuration	Disk Groups: 2 Cache Tier: 2 x 1.46 TB NVMe Capacity Tier: 8 x 3.49 TB SAS SSDs

Linux LVM configuration	Used Linux LVM for log and data. Used 8 VMDK disks in LVM to create the needed log and a data volume for SAP HANA.
--------------------------------	---

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

The 2-socket/4-disk group test environment had four vSphere ESXi servers configured and used direct-attached SAS and SATA SSDs inside the hosts to provide a vSAN datastore that meets the SAP HANA HCI test requirements and KPIs.

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

Each vSphere ESXi server had four disk groups each consisting of one 800 GB **SAS SSD** for caching and four capacity tier 1.8 TB **SATA SSDs**. The raw capacity of the vSAN datastore of this 4-node vSAN cluster was for data around 115 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 10 Gbit network adapters were used as a dedicated vSAN network in active-standby failover manner. Our tests have shown, that a dedicated 10 GbE network for vSAN is required for 2 SAP HANA VMs per server. Using higher bandwidth network adapters, like 25 Gbit adapters for vSAN instead, will allow you to support more SAP HANA VMs per host/vSAN datastore and other network traffic like SAP application server network.

Each ESXi Server in the vSAN Cluster had the following configuration as shown in Table 7.

Table 7. ESXi Server Configuration of the 2-socket/4-disk Server SAP HANA 4-node HCI vSAN Cluster

PROPERTY	SPECIFICATION
ESXi server model	HW Vendor SAP HANA Supported Skylake Server Model
ESXi host CPU	2 x Intel(R) Xeon(R) Gold 6154 CPU , 2 Sockets, 18 Physical Cores/Socket
ESXi host RAM	768 GB (384 GB per CPU)
Network adapters	2 x 10 Gbit dedicated network ports for vSAN. Plus, optional network cards for SAP application and admin network traffic. See the network section of this guide for the recommended network configurations.
Storage adapter	1 x 8 channel PCI-EXPRESS 3.0 SAS 12GBS/SATA 6GBS SCSI controller
Storage configuration	Disk Groups: 4 Cache Tier: 4 x 800 GB SAS SSD Capacity Tier: 16 x 1.92 TB SATA SSD

Linux LVM configuration	Used Linux LVM for log and data. Used 8 VMDK disks in LVM to create the needed log and a data volume for SAP HANA.
--------------------------------	---

4-Socket SAP HANA HCI vSAN Cluster Example

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 up to the vendor.

Each vSphere ESXi server had **four disk groups** each consisting of one 1.6 TB NVMe devices for caching and six capacity tier 1.92 TB SAS SSDs. The **raw capacity** of the vSAN datastore of this **3-node** vSAN cluster was for data around **138.24 TB**. The vSAN cluster limit is up to 64 cluster nodes.

Two 25 Gbit network adapters were used for vSAN and another 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 8.

Table 8. ESXi Server Configuration of the 4-socket Server SAP HANA 3-node HCI vSAN Cluster

PROPERTY	SPECIFICATION
ESXi server model	HW Vendor SAP HANA Supported Skylake Server Model
ESXi host CPU	4 x Intel(R) Xeon(R) Platinum 8180M CPU, 4 Sockets, 28 Physical Cores/Socket
ESXi host RAM	6,144 GB
Network adapters	2 x 25 Gbit as a shared network for vSAN, SAP application and admin network traffic. See the network section of this guide for recommended network configurations.
Storage adapter	1 x 8 channel PCI-EXPRESS 3.0 SAS 12GBS/SATA 6GBS SCSI controller
Storage configuration	Disk Groups: 4 Cache Tier: 4 x 1.6 TB NVMe Capacity Tier: 24 x 1.92 TB SAS SSD
Linux LVM configuration	Linux LVM was not used.

--	--

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 GA build#: 8217866 6.5 U2b build#: 8815520	VMware vCenter Server provides a centralized platform for managing VMware vSphere environments.
<i>vSphere ESXi</i>	6.7 GA build#: 8169922 6.5 U2b build#: 8935087	vSphere ESXi cluster to host virtual machines and provide vSAN Cluster.
<i>VMware vSAN</i>	6.7 GA build#: 8169922 6.6.1 U2 build#: 8294253	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 like the 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	SYSTEM REPLICATION NETWORK
Host Network Configuration (all networks are dual port NICs, configured as failover teams)	Network Label	Admin	App-Server	vMotion	vSAN	Backup	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	
VM Guest Network Cards	Virtual NIC (inside VM)	0	1	-	-	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 larger than 10 Gbit NIC port is needed if more than two SAP VMs should be deployed per host.

Beside using a high bandwidth, 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 10GbE** 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 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, **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 4-disk group configuration. As noted, the HW vendors are free to decide which configuration they want to certify.

In our 2-socket/4-disk group and 4-host server vSAN cluster configuration, we used a write optimized 800 **GB SAS SSD for caching** and four 1.92 TB SATA SSDs for the capacity tier per disk group. In total we used 20 SSDs per server, all connected on one vSAN certified SCSI adapter. With this vSAN configuration, we were able to validate 2 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 a 1.6 TB **NVMe device for caching** and six SAS 1.92 TB large SSDs for the capacity tier per disk group. With this vSAN configuration, we were able to validate successfully 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 6 shows a single server of our 4-disk group test configuration in a graphical way.

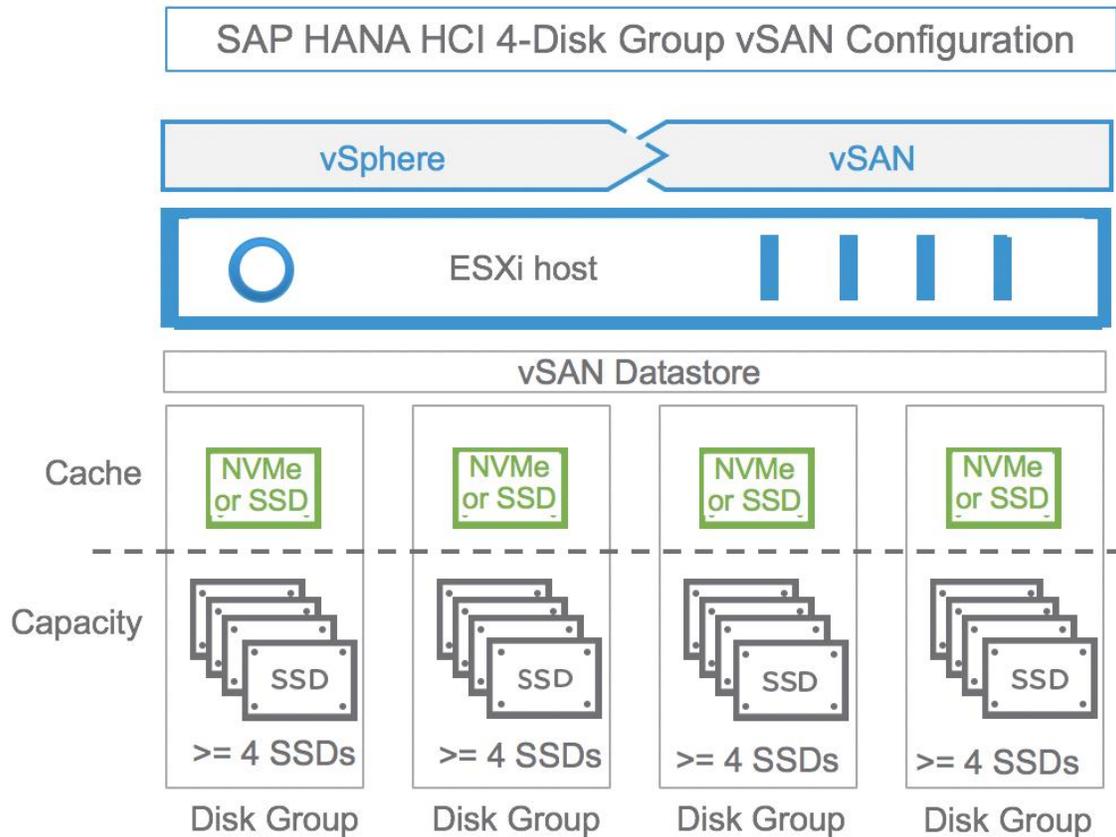


Figure 6. 4-disk Group Configuration Example of Our SAP HANA HCI Test Configuration

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.

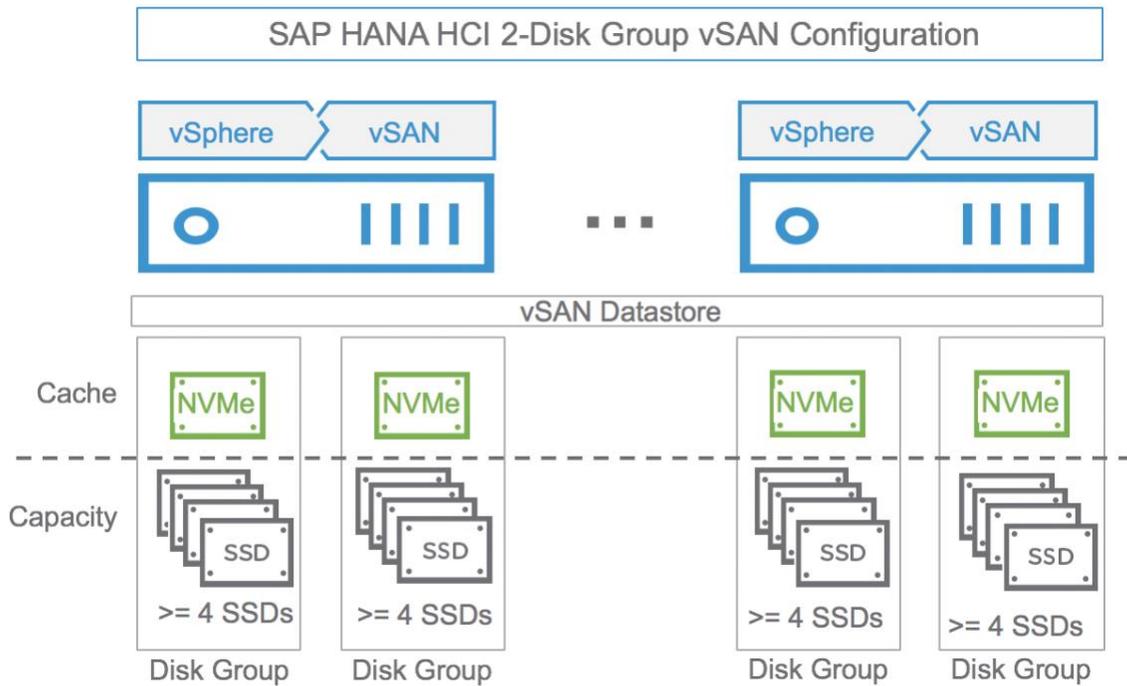


Figure 7. 2-disk Group Configuration of Our SAP HANA HCI Test Configuration

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 an into 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.

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

STORAGE CAPABILITY	SETTING
Failures 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 5/6.

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 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.

HWCCT is a tool available only for SAP partners and customers that test a system and measure it against the currently valid storage KPIs. The tool and the documentation can be downloaded from SAP with a valid SAP user account. See [SAP note 1943937](#).

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 spindles (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 has to 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 8 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).

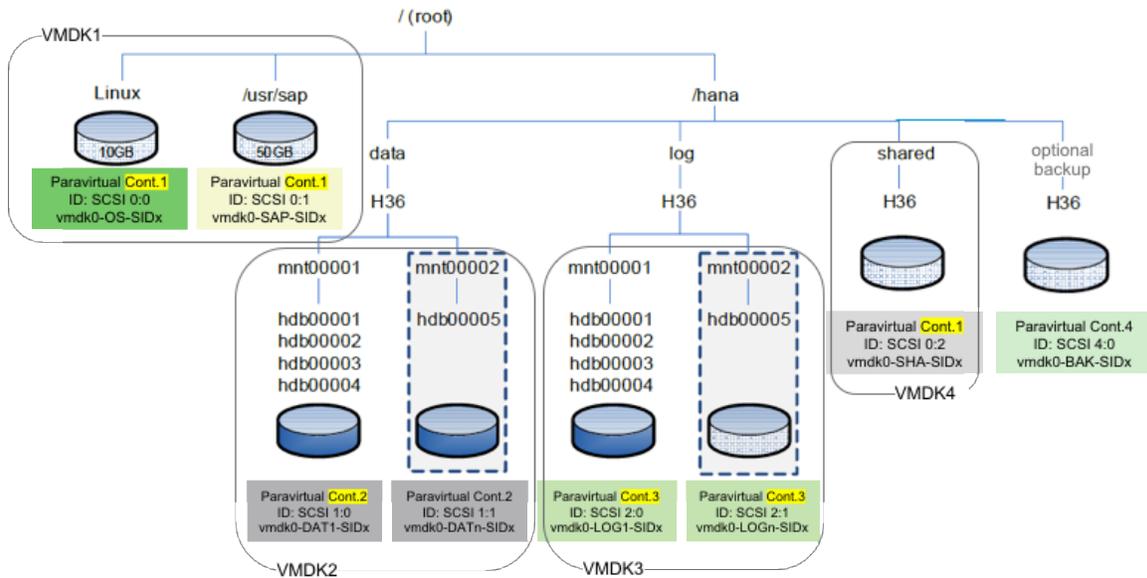


Figure 8. Storage Layout of an SAP HANA System. Figure © SAP SE, Modified by VMware

Table 16 is used to determine 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 10 GiB for OS (suggested 100 GB Thin)
usr/sap	✓	PVSCSI Contr. 1	vmdk01-SAP-SIDx	SCSI 0:1	Minimum 50 GiB for SAP binaries (suggested 100 GB thin)

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 vmdk03-DAT3-SIDx	SCSI 1:0 SCSI 1:1	Minimum 1 x RAM Note: If you use multiple VMDKs, use for example 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 9 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.

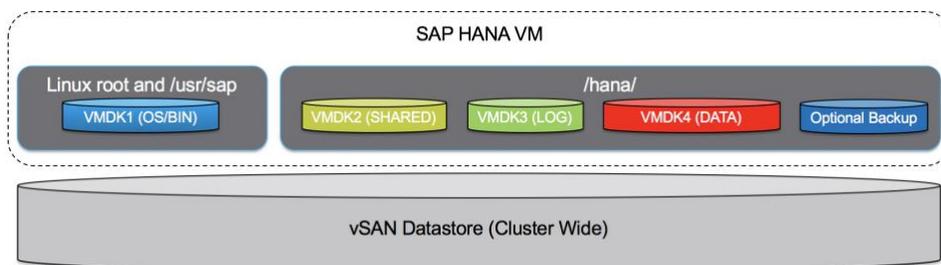


Figure 9. Dedicated VMDK Files per SAP HANA Volume

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 10 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.

Note that the number of VMDK disks used for LVM should not exceed the overall number physical capacity disks used in a host. For example, if only 6 capacity disks are used in a server, 8 VMDK disks for LVM would require because two disks will have to store the objects that represent two VMDK files. This would result in an unbalanced configuration. Whereas 8 VMDK disks would be used if 8 or more capacity disks are installed. For example, if only 6 capacity disks are used in a server, up to 6 VMDK disks for LVM would be ideal.

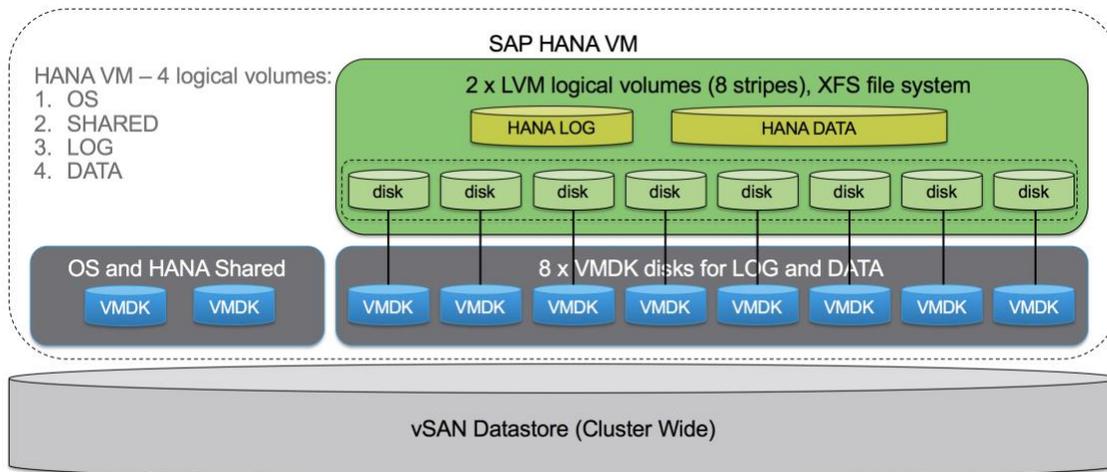


Figure 10. Usage of LVM Managed Volumes

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 8 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 + USR/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.

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

As [shown earlier](#), we have over 138 GB 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:

- Enable SSH for the time of installation. Deactivate it after the installation of the script.
- 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 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

```

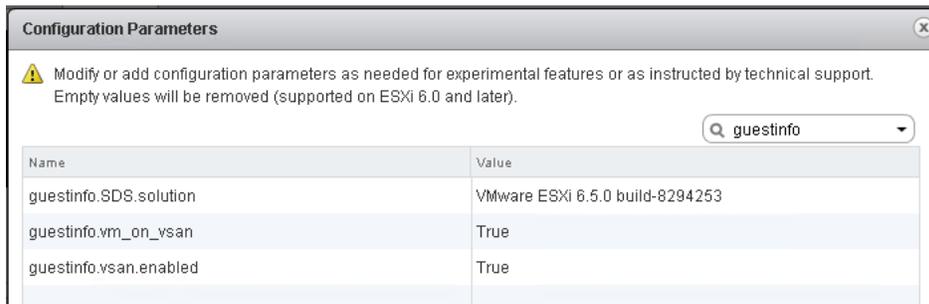
Compare those values above with the values extracted from the ESXi, then update the values on the VM if necessary.

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 vm disk on vSAN), or moved out from vSAN to other datastores, the guestinfo will be updated within 5 mins.

If the VM gets storage vMotioned to another host that does not have the script running, you need to manually delete the 3 variables from the VM configuration file to ensure consistency. You can check inside vCenter if these variables are set for a VM. See Figure 11 for details:



The screenshot shows a window titled "Configuration Parameters" with a warning icon and text: "Modify or add configuration parameters as needed for experimental features or as instructed by technical support. Empty values will be removed (supported on ESXi 6.0 and later)." Below this is a search bar containing "guestinfo". A table lists the following parameters:

Name	Value
guestinfo.SDS.solution	VMware ESXi 6.5.0 build-8294253
guestinfo.vm_on_vsan	True
guestinfo.vsan.enabled	True

Figure 11. 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 get 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

Best Practices of Virtualized SAP HANA for Skylake Based Server Host Systems

Optimizing SAP HANA on vSphere

Note: This section is independent from vSAN and for all SAP HANA systems running on vSphere relevant.

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 in a standard vSphere environment or SAP HANA on vSAN based SAP HANA HCI configurations.

The listed settings show the parameter settings recommended for the physical server, the vSphere ESXi host, the VM, and the Linux OS to achieve optimal operational readiness and stable performance for SAP HANA on vSphere. These parameter settings are the default settings that should be configured for SAP HANA. The settings described in the section “extreme performance settings” should only be applied in rare situations where SAP HANA has to 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 these settings.

Physical HOST BIOS SETTINGS	DESCRIPTION
UEFI BIOS HOST	It is recommended to use (U)EFI BIOS as the standard BIOS version for physical vSphere hosts. All SAP HANA appliance server configurations leverage (U)EFI 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 modules in a way to achieve best performance (hemisphere mode), use the fastest DIMM modules available for the selected RAM size. Beware of the CPU specific optimal RAM 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. See Timekeeping in VMware Virtual Machines for reference.</p> <p>Select only SAP HANA vSphere supported CPUs. Verify the support status with SAP note 2652670—SAP HANA VM on VMware vSphere.</p>

Enable CPU Intel Turbo Boost	Allow Intel automatic CPU core overclocking technology (P-States)
Disable QPI Power Management	Static high power for QPI links
Enable hyperthreading	Double the logical CPU cores to allow ESXi to take advantage for the more available CPU threads. Enable HT regardless the Intel L1 Terminal Fault issue, but make sure to install the latest VMware vSphere patches for this issue
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 like C-States. Configure in the BIOS static high performance
Disable all unused BIOS features, such as:	Video BIOS Shadowable, 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, USB

VSPHERE ESXi HOST SETTINGS	DESCRIPTION
Memory	
Networking	Use: Virtual Distributed Switches (vDS) 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 10GbE for vMotion and the SAP App server or replication networks. At least 25 GbE for vMotion for SAP HANA system >= 2 TB is recommended.
Storage Configuration	Use only VMware vSAN certified storage components. Use the existing BOM of the vSAN ready nodes as a template for an SAP HANA HCI solution.
SAP Monitoring	Enable SAP monitoring on the host -> Misc.GuestLibAllowHostInfo = "1" Details: SAP note 1409604 . Without this parameter, no host performance relevant data will be viewable inside an SAP monitoring enabled VM.

6.5 and later scheduler changes	<p>Set the advanced host parameter Numa.FollowCoresPerSocket="1". If you want to change this globally for all VMs on this host.</p> <p>If you want to change this per VM, then check the per VM settings (see the "6.5 and later scheduler" section in the next table). Setting this parameter per VM is recommended, especially if you do not have dedicated vSphere hosts/clusters and if non-SAP HANA VMs run on the same host.</p> <p>This parameter will revert the vNUMA layout optimization to the old vSphere 6.0 method and allows to influence the vNUMA layout of a VM via vSphere client configuration parameters "CPU" and "cores per socket".</p>
--	--

SAP HANA VIRTUAL MACHINE	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 (U)EFI BIOS as the standard BIOS version for vSphere hosts and guests. Only with EFI features like Secure Boot are possible. See VMware DOC-28494 for details. You can configure this with the vSphere client by choosing EFI boot mode.</p> <p>If you are using vSphere 6.0 and you only see 2 TB 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 tools.guestlib.enableHostInfo = "TRUE".</p> <p>Details: SAP note 1409604.</p> <p>Beside setting this parameter, the VMware guest tools need to be installed. See VMware KB 1014294 for details.</p>
vCPU hotplug	Ensure that vCPU hotplug is deactivated , otherwise vNUMA is disabled and SAP HANA will have a negative performance impact. Details: VMware KB 2040375 .
Memory reservations	Set fix memory reservations for SAP HANA VMs. Do not overcommit memory resources.
CPU	Do not overcommit CPU resources.
vNUMA nodes	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>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>
Align virtual CPU VM configuration to actual server HW	<p>Example: A VM running on a 4-socket server with 28-core CPUs, should be configured with 14 virtual CPUs and 14 cores per CPU socket. Per half socket VM, 28 CPUs and cores per socket; per full socket VM, 56 CPUs and 28 cores per socket for a 2-CPU socket VM, 84 CPUs and 28 cores for a 3-socket VM and finally 112 CPUs and 28 cores per CPU socket to get a 4-socket VM.</p>
Paravirtualized SCSI driver for I/O devices	<p>Use the dedicated SCSI controllers for OS, log, and data, to separate disk I/O streams. For details see SAP HANA disk layout section.</p>
Use Virtual Machine File System	<p>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.</p>
Create datastores for SAP HANA data and log files	<p>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. See SAP note 1943937 for details.</p>
We recommend eagerzeroed thick virtual disks for data and log disk	<p>This avoids lazy zeroing (initial write penalty).</p>
6.5 and later scheduler changes	<p>If not set globally on the host, set the per VM advanced parameter numa.vcpu.followcorespersocket = "1"</p> <p>Setting this parameter per VM is recommended, especially if you do not have the dedicated vSphere hosts/clusters and if non-SAP HANA VMs run on the same host.</p>
VMXNET3	<p>Use paravirtual VMXNET 3 virtual NICs for SAP HANA virtual machines.</p> <p>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.</p>
Optimize the application server network latency if required	<p>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 like vMotion or SAP HANA system replication. Use the advanced options in the vSphere Web Client or directly</p>

	<p>modify the .vmx file and add ethernetX.coalescingScheme = "disable". X stands for your network card number. See Best Practice for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs for detail.</p>
<p>Set lat.Sensitivity = normal</p>	<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 have to 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>
<p>Associate Virtual Machines with Specified NUMA Nodes</p>	<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>This is only needed when more than one SAP HANA VM run on a host, especially with the half-socket configuration.</p> <p>Use the vSphere Web Client or directly modify the .vmx file and add numa.nodeAffinity=x (for example: 0,1)</p> <p>Attention:</p> <p>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>See Associate Virtual Machines with Specified NUMA Nodes for details.</p>
<p>Configure virtual machines to use hyperthreading with NUMA</p> <p>Currently using hyperthreading is not recommended due to the Intel L1 Terminal Fault vulnerabilities. See VMware KB 55636 for details.</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 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. It is recommended for SAP HANA VM configurations that fill a complete server, like a 4-socket VM on a 4-socket server.</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>

Remove unused devices	Such as floppy disks or CD-ROM to release resources and to mitigate possible errors.
------------------------------	--

LINUX OS	DESCRIPTION
VMware Specific Configuration Recommendations:	
Linux kernel	VMware strongly recommends using the latest SAP HANA supported kernel version
Configure C-States and BIOS EPB for lower latency in Linux:	<p>Disable in BIOS (for example, static high performance or OS controlled) If disabled in BIOS, optional: intel_idle.max_cstate=0</p> <p>If disabled in BIOS, optional: processor.max_cstate=0</p> <p>Set Energy Performance in BIOS to "maximum performance", optional: cpupower set -b 0</p>
Optional: Disable Large Receive Offload (LRO) in the Linux guest operating system 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: "ethtool -K ethY lro off".</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>
SLES Linux page cache limits	Follow SAP Note 1557506 and VMware vSphere configuration guidelines, SAP note 2161991 to use the optimal configuration settings.
Disable I/O scheduling	Set kernel parameter " elevator=noop " for virtual machines. This disables I/O scheduling. The VMware hypervisor has its own I/O scheduling mechanisms; therefore, scheduling I/O inside the Guest Operating System can cause unnecessary overhead. For details, see Red Hat DOC-5428 and SUSE KB 7009616 (applies to other SLES versions as well)
To optimize large-scale workloads with intensive I/O patterns, change the queue depths of the SCSI default values	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>Create a file of any name in the <code>/etc/modprobe.d/</code> directory with this line:</p> <pre>options vmw_pvscsi cmd_per_lun=254 ring_pages=32</pre> <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>. 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> <pre>vmw_pvscsi.cmd_per_lun=254 vmw_pvscsi.ring_pages=32</pre> <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	<p>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.</p>
Configure NTP time server	<p>Use the same external NTP server as configured for vSphere. For details: SAP note 989963.</p>

General SAP HANA Linux Configuration Recommendations:

Linux Operating System with SAP HANA Reference Guide	<p>Refer to the SAP WIKI for SAPs and the Linux OS vendors virtualization independent recommended Linux OS settings for SAP HANA.</p>
Disable AutoNUMA	<p>Later Linux kernel (RHEL 7 and SLES 12) supporting auto-migration according to NUMA statistics.</p> <p>For SLES: # yast bootloader, choose "Kernel Parameters" tab (ALT-k) and edit the "Optional Commandline Parameters" section by appending numa_balancing=disabled</p> <p>For RHEL add "kernel.numa_balancing = 0" to <code>/etc/sysctl.d/sap_hana.conf</code> and reconfigure the kernel by running: # sysctl -p <code>/etc/sysctl.d/sap_hana.conf</code></p>
Disable transparent HugePages	<p>THP is not supported for the use with SAP HANA DB, as it may lead to hanging situations and performance degradations.</p> <p>To check the current configuration run the following command: # <code>cat /sys/kernel/mm/transparent_hugepage/enabled</code></p>

	<p>Its output should read: always madvise [never]</p> <p>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</p> <p>For details, refer to the above SAP Wiki page.</p>
<p>Change the following parameters in /etc/sysctl.conf</p> <p>(important for SAP HANA scale-out deployments)</p>	<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>

Performance Optimization for Extreme Performance/Latency Critical 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 best/extreme performance, 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.

Before you use 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.
- 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 affinity, 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, 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 virtual machine that is already assigned to the certain processors using affinity.
- Affinity 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 potential performance practices, see the [vSphere Resource Management Guide](#).

Additional Performance Tunings for Maximum Performance

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

SAP HANA VIRTUAL MACHINE	DESCRIPTION
Tips about how to edit the .vmx file	Review 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 to improve performance of a VM at the cost of CPU time on the ESXi host and should only get configured for SAP HANA workload that runs anyhow as the only workload on a NUMA node/complete server.</p> <p>Edit the vmx file and add the following two advanced parameters: monitor.idleLoopSpinBeforeHalt = "true" AND monitor.idleLoopMinSpinUS = "xx" (For example: 150)</p> <p>Both parameters have to be configured to influence the de-scheduling time.</p> <p>Background: The guest operating system 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-processors wake up requests.</p>
Set control.halt_in_monitor = "TRUE"	In the default configuration of ESX 6.0, idle state of guest HLT instruction will be emulated w/o leaving VMM if a vCPU has exclusive affinity. If affinity is non-exclusive, guest HLT will be emulated in vmkernel, which may result in having 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

	<p>instruction gets emulated inside the VMM and not in the vmkernel.</p> <p>Use the vSphere Web Client and add the following advanced VM parameter: <code>monitor_control.halt_in_monitor = "TRUE"</code>.</p>
<p>Set monitor_control.disable_pause_loop_exiting = "TRUE"</p>	<p>This parameter prevents the VM from exiting to the hypervisor unnecessarily during a pause instruction. This is specific to Intel Skylake CPU based systems.</p>
<p>Configuring CPU Affinity</p> <p>Note: Remove numa.nodeAffinity settings if set and if the CPU affinities are configured.</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>Setting CPU affinities can decrease the CPU and memory latency by not allowing the ESXi scheduler to migrate VM threads to other logical processors.</p> <p>See Scheduler operation when using CPU Affinity (2145719) for details.</p> <p>Procedure:</p> <ol style="list-style-type: none"> 1. Find the virtual machine in the vSphere Web Client inventory. <ol style="list-style-type: none"> a. To find a virtual machine, select a data center, folder, cluster, resource pool, or host. b. Click the Related Objects tab and click Virtual Machines. 2. Right-click the virtual machine and click Edit Settings. 3. Under Virtual Hardware, expand CPU. 4. Under Scheduling Affinity, select physical processor affinity for the virtual machine. Use '-' for ranges and ',' to separate values. For example, "0, 2, 4-7" would indicate processors 0, 2, 4, 5, 6 and 7. 5. Select the processors where you want the virtual machine to run and click OK. <p>For more information about potential performance practices, see vSphere Resource Management Guide.</p>

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 vSAN demos, labs, and reference architectures](#)
- [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 Support Insight](#)
- [vSAN SAP support note 2718982 - SAP HANA on VMware vSphere and vSAN](#)
- [SAP HANA Certified Hyperconverged Infrastructure Solutions](#)

Author and Contributors

Erik Rieger, Global Solution Architect SAP, Global Partner Organization, Solution Engineering team wrote the original version of this paper.

Chen Wei, Sr. Solution Architect in the Product Enablement team of the Storage and Availability Business Unit contributed to the original version of this paper.

Todd Muirhead, VMware Staff Performance of the VMware Performance Unit contributed to the original version of this paper.

Sebastian Lenz, VMware SAP Performance and Certification Engineer of the VMware Performance Unit contributed to the original version of this paper.

Catherine Xu, Senior Technical Writer in the Product Enablement team of the Storage and Availability Business Unit edited this paper to ensure that the contents conform to the VMware writing style.