

This asset was published on July 2022 and may contain out-of-date information. Please visit <https://core.vmware.com/resource/sap-hana-vmware-vsphere-best-practices-and-reference-architecture-guide> for the latest version.

SAP HANA on VMware vSphere Best Practices and Reference Architecture Guide

Table of contents

SAP HANA on VMware vSphere Best Practices and Reference Architecture Guide	4
Abstract	4
Audience	4
Solution Overview	5
SAP HANA on vSphere	5
Solution Components	5
Production Support	5
SAP HANA on vSphere Release Strategy	7
SAP HANA Supported vSphere Editions	7
SAP HANA on vSphere Scalability and VM Sizes	7
SAP HANA on vSphere Reference Architecture	9
Scope	9
Overview	9
Deployment Options and Sizes	10
High-level Architecture	11
SAP HANA on vSphere Configuration and Sizing Guidelines	13
Sizing Compute and Memory	13
SAP HANA Sizing Process	13
Storage Configuration and Sizing	18
SAP HANA Hardware and Cloud Measurement Tool	22
Network Configuration and Sizing	23
VMXNET3 SAP Workload-specific Network Latency Considerations	28
Enhanced vMotion Compatibility, vSphere vMotion, and vSphere DRS Best Practices	35
vSphere Clustering Service	38
Virtual SAP HANA Disaster Recovery	49
Backup and Restore VMware Virtualized SAP HANA Systems	50
SAP HANA with Persistent Memory on vSphere	52
Prerequisites and General SAP Support Limitations for Intel Optane PMem	52
What is Supported?	52
Sizing of Optane PMEM-enabled SAP HANA VMs	52
vSphere HA Support for PMEM-enabled SAP HANA VMs	55
SAP HANA with PMEM VM Configuration Details	55
SAP HANA Scale-Out on vSphere Installation	60
Monitoring and Verifying an SAP HANA Installation	62

VMware NUMA Observer	62
Best Practices of Virtualized SAP HANA Systems	64
Third Generation Intel Xeon 40-core Processor (Ice Lake SP) Specifics	64
Optimizing the SAP HANA on vSphere Configuration Parameter List	68
Performance Optimization for Low-Latency SAP HANA VMs	74
Additional Performance Tuning Settings for SAP HANA Workloads	75
Example VMX Configurations for SAP HANA VMs	76
CPU Thread Matrix Examples	78
SAP HANA Support and Process	83
Open an SAP Support Request Ticket	83
Open a VMware Support Request Ticket	83
Conclusion	85
VMware Acknowledgments	86

SAP HANA on VMware vSphere Best Practices and Reference Architecture Guide

Abstract

This guide is the 2021 edition of the best practices and recommendations for SAP HANA on [VMware vSphere®](#). This edition covers VMware virtualized SAP HANA systems running with vSphere 7.0 and later versions on first, [second-generation Intel Xeon Scalable processors](#), such as Broadwell, Skylake, Cascade Lake, Cooper Lake and [third-generation Intel Xeon Scalable \(Ice Lake\) processor](#) based host server systems. It also includes [Intel Optane Persistent Memory \(PMEM\) 100 series technology](#) support for vSphere 7.0 and later virtualized SAP systems.

For earlier vSphere versions or CPU generations, please refer to the [2017 edition of this guide](#).

This guide describes the best practices and recommendations for configuring, deploying, and optimizing SAP HANA scale-up and scale-out deployments. Most of the guidance provided in this guide is the result of continued joint testing conducted by VMware and SAP to characterize the performance of SAP HANA running on vSphere.

Audience

This guide is intended for SAP and VMware hardware partners, system integrators, architects and administrators who are responsible for configuring, deploying, and operating the SAP HANA platform in a VMware virtualization environment.

It assumes that the reader has a basic knowledge of vSphere concepts and features, SAP HANA, and related SAP products and technologies.

Solution Overview

SAP HANA on vSphere

By running SAP HANA on VMware Cloud Foundation™, the VMware virtualization and cloud computing platform, scale-up database sizes up to 12 TB (approximately 12,160 GB) and scale-out deployments with up to 16 nodes (plus high availability [HA] nodes) with up to 32 TB (depending on the host size) can get deployed on a VMware virtualized infrastructure.

Running the SAP HANA platform virtualized on vSphere delivers a software-defined deployment architecture to SAP HANA customers, which will allow easy transition between private, hybrid or public cloud environments.

Using the SAP HANA platform with VMware vSphere virtualization infrastructure provides an optimal environment for achieving a unique, secure, and cost-effective solution and provides benefits physical deployments of SAP HANA cannot provide, such as:

- Increased security (using VMware NSX® as a Zero Trust platform)
- Higher service-level agreements (SLAs) by leveraging vSphere vMotion® to live migrate running SAP HANA instances to other vSphere host systems before hardware maintenance or host resource constraints
- Integrated lifecycle management provided by VMware Cloud Foundation SDDC Manager™
- Standardized high availability solution based on vSphere High Availability
- Built-in multitenancy support via SAP HANA system encapsulation in a virtual machine (VM)
- Easier HW upgrades or migrations due to abstraction of the hardware layer
- Higher hardware utilization rates
- Automation, standardization and streamlining of IT operation, processes, and tasks
- Cloud readiness due to software-defined data center (SDDC) SAP HANA deployments

These and other advanced features found almost exclusively in virtualization lower the total cost of ownership and ensure the best operational performance and availability. This environment fully supports SAP HANA and related software in production[1] environments, as well as SAP HANA features such as SAP HANA multitenant database containers (MDC)[2] or SAP HANA system replication (HSR).

Solution Components

An SAP HANA virtualized solution based on VMware technologies is a fully virtualized and cloud-ready infrastructure solution running on VMware ESXi™ and supporting technologies, such as VMware vCenter®. All local server host resources, such as CPU, memory, local storage, and networking components, get presented to a VM in a virtual way, which is abstracting the underlying hardware resources.

The solution consists of the following components:

- VMware certified server systems as listed on the [VMware hardware compatibility list \(HCL\)](#)
- SAP HANA supported server systems, as listed on the [SAP HANA HCL](#)
- SAP HANA certified hyperconverged infrastructure (HCI) solutions, as listed on the [SAP HANA HCI HCL](#)
- vSphere version 7.0 and later (up to 6 TB vRAM)
- vSphere version 7.0 U2 and later (up to 12 TB vRAM)
- vCenter 7.0 and later
- A VMware specific and SAP integrated support process

Production Support

In November 2012, SAP announced initial support for single scale-up SAP HANA systems on vSphere 5.1 for non-production environments. Since then, SAP has extended its production-level support for SAP HANA scale-up and scale-out deployment options and multi-VM and half-socket support. vSphere versions 5.x and most vSphere versions 6.x are by now out of support. Therefore, any customer should by now use vSphere version 7.0 and later. Table 1 provides an overview relevant SAP HANA on vSphere support notes as of October 2022.

Table 1: Relevant SAP Notes

Key notes for virtual environments	1492000 : General support statement for virtual environments 1380654 : SAP support in cloud environments 2161991 : VMware vSphere configuration guidelines
SAP HANA on vSphere	3102813 : SAP HANA on VMware vSphere 7.0 U2 with up to 12 TB 448 vCPUs VM sizes 2937606 : SAP HANA on VMware vSphere 7.0 (incl. U1 and U2) in production 2393917 : SAP HANA on VMware vSphere 6.5 and 6.7 in production 2779240 : Workload-based sizing for virtualized environments 2718982 : SAP HANA on VMware vSphere and vSAN 2913410 : SAP HANA on VMware vSphere with Persistent Memory 2020657 : SAP Business One, version for SAP HANA on VMware vSphere in production

Table 2 provides an overview, as of October 2022, of the SAP HANA on vSphere supported and still relevant vSphere versions and CPU platforms. This guide focuses on vSphere 7.0 U2 and later.

Table 2: Supported vSphere Versions and CPU Platforms as of October 2022

vSphere version	Broadwell	Skylake	Cascade Lake	Cooper Lake	Ice Lake
vSphere 6.7 U2/U3 up to 4-socket wide VM	•	•	•	-	-
vSphere 7.0 up to 4-socket wide VM	•	•	•	-	-
vSphere 7.0 U2 and later up to 4-socket wide VM	•	•	•	•	•
vSphere 7.0 U3c and later up to 8-socket wide VM	-	-	•	•	-
PMEM Series 100	-	-	•	-	-

Table 3 summarizes the key maximums of the different vSphere versions supported for SAP HANA. For more details, see the [SAP HANA on vSphere scalability and VM sizes section](#).

Table 3: vSphere Memory and CPU SAP HANA VM Relevant Maximums per CPU Generation

vSphere version[3]	Maximum memory	Maximum CPUs[4]	CPU sockets[5]
vSphere 6.7	< 4 TB Intel Broadwell < 6 TB Intel Skylake and later	<= 128 vCPUs	0.5-, 1-, 2-, 3-, 4-socket wide VMs
vSphere 6.7 U2/U3 and 7.0/7.0 U1 and later	< 4 TB for Broadwell and Ice Lake < 6 TB Intel Skylake and later	<= 160 vCPUs for Ice Lake <= 192 vCPUs for Broadwell <= 256 vCPUs for Skylake and later, typically 224 vCPUs	0.5-, 1-, 2-, 3-, 4-socket wide VMs
vSphere 7.0 U2 and later (8S VM)	< 12 TB with Intel Cascade and Cooper Lake	<= 448 vCPUs	0.5-, 1-, 2-, 3-, 4-, 5-, 6-, 7-, and 8-socket wide VMs

SAP HANA on vSphere Release Strategy

VMware's SAP HANA certification/support strategy for vSphere is to support a single CPU generation or chipset with two versions of the hypervisor and to have a single hypervisor version span two CPU generations/chipsets.

VMware does its best to strike a balance between supporting new customers on the latest hardware and those who are remaining on an older platform. Customers may still use vSphere 6.7 as VMware has [extended the general support](#) period until October 2022. Nevertheless, vSphere 7.0 Update 2 is the most recent available version, and VMware recommends using this version for all new SAP HANA deployments on vSphere. End of general support for this version is April 2025. For more details, refer to the [VMware Product Lifecycle Matrix](#). For easier visibility, set a filter for ESXi at the product release cell.

SAP HANA Supported vSphere Editions

SAP HANA on vSphere is only supported with following editions: vSphere Standard and vSphere Enterprise Plus. Business One solutions are also supported with Acceleration or Essentials Kits. Please note the feature/CPU/host limitations that come with these kits. For SAP production environments (HANA and [Business One](#)), a support contract with a minimum [production support](#) level with a 24x7 option is highly recommended.

The [VMware vSphere Compute Virtualization white paper](#) provides an overview of the licensing, pricing and packaging for vSphere. For VMware support offerings, see the On-Premises Support section of the [VMware Support Offerings and Services page](#).

For the latest SAP released support information, please refer to the [SAP notes related to VMware](#) or the [SAP on VMware vSphere community wiki page](#).

As of SAP note [2652670: SAP HANA VM on VMware vSphere](#), usually all update and maintenance versions of vSphere hypervisors are automatically validated within the same boundary conditions.

SAP HANA on vSphere Scalability and VM Sizes

SAP HANA on vSphere is supported from the smallest SAP HANA system, which is a half-socket CPU configuration with

a minimum of 8 physical CPU cores and 128 GB of RAM, up to 8-socket large SAP HANA VMs with up to 12 TB of RAM. Actual required CPU power and RAM for a certain SAP workload on HANA must get clarified by an SAP HANA sizing exercise. VMs larger than 4-socket VMs require an 8-socket host. 8-socket servers are an optimal consolidation platform and could, for instance, host two large 4-socket VMs, each with up to 6 TB memory, or 16 half-socket VMs with up to 750 GB memory, or one single large 12 TB SAP HANA VM.

Up to 6 TB for a 4-socket large SAP HANA VM and 12 TB for an 8-socket large VM are covered by the SAP standard support for OLTP workloads. Only 50 percent of the memory defined for OLTP-type workloads is supported by this standard sizing (e.g., 3 TB for 4-socket wide OLAP VMs, or 6 TB for 8-socket wide VMs).

If more memory is required, then a workload-based/SAP expert sizing needs to be performed.

For more details, review SAP note [2779240: Workload-based sizing for virtualized environments](#). Table 4 shows the current vSphere 6.7 U2+/7.0 and 7.0 U2 or later maximums[6] per physical ESXi host.

Table 4: vSphere Physical Host Maximums

	ESXi 6.7 U3 / 7.0	ESXi 7.0 U2 / U3c
Logical CPUs per host	768	896
VMs per host	1,024	
Virtual CPUs per host	4,096	
Virtual CPUs per core	32	
RAM per host	16 TB	24 TB
NUMA nodes/CPU sockets per host	16 (SAP HANA only 8 CPU socket hosts / HW partitions)	

Note: Only SAP and VMware supported ESXi host servers with up to eight physical CPUs are supported. Contact your SAP or VMware account team if larger 8-socket systems are required and to discuss deployment alternatives, such as scale-out or memory-tiering solutions. Also note the support limitations when using 8-socket or larger hosts with node controllers (also known as glued-architecture systems/partially QPI meshed systems).

Table 5 shows the maximum size of a vSphere VM and some relevant other parameters, such as virtual disk size and the number of virtual NICs per VM. These generic VM limits are higher than the SAP HANA supported configurations, also listed in Table 5.

Table 5: vSphere Guest (VM) Maximums

	ESXi 6.7 U3 / 7.0	ESXi 7.0 U2 / U3c
Virtual CPUs per VM	256	768
RAM per VM	6,128 GB	24 TB
Virtual CPUs per SAP HANA VM with 28-core Cascade/Cooper Lake CPUs	224	448
CPU sockets per SAP HANA VM	<= 4	<= 8
RAM per SAP HANA VM[7]	<= 6,128 GB	<= 12 TB
Virtual SCSI adapters per VM	4	
Virtual NVMe adapters per VM	4	
Virtual disk size	62 TB	
Virtual NICs per VM	10	
PCI passthrough devices per VM	16	
Persistent Memory per SAP HANA VM	<= 12 TB	

SAP HANA on vSphere Reference Architecture

The reference architecture outlined in this section provides a starting point for a VMware virtualized SAP HANA project and provides best practices and guidelines. The reference architecture got tested and validated with test and workloads specified by VMware and SAP and is known to be good with the shown SAP HANA configuration. It won't provide information on specific vendors or hardware components to use.

Scope

This reference architecture:

- Provides a high-level overview of a VMware SDDC for SAP applications, such as SAP HANA
- Describes deployment options and supported sizes of SAP HANA virtual machines on vSphere
- Shows the latest supported vSphere and hardware versions for SAP HANA
- Explains configuration guidelines and best practices to help configure and size the solution
- Describes HA and [disaster recovery \(DR\) solutions](#) for VMware virtualized SAP HANA systems

Overview

Figure 1 provides an overview of a typical VMware SDDC for SAP applications. At the center of a VMware SDDC are three key products: vSphere, VMware vSAN™, and NSX. These three products are also available as a product bundle called [VMware Cloud Foundation](#) and can get deployed on premises and in a public cloud environment.

As previously explained, virtualized SAP HANA systems can get as big as 448 vCPUs with up to 12 TB RAM per VM; the vSphere 7.0 U2 or later limits are 896 vCPUs and 24 TB per VM. As of today, only 8-socket host systems with up to 448 vCPUs got SAP HANA on vSphere validated. Also, the maximum number of vCPUs available for a VM may be limited by the expected use case, such as extremely network heavy OLTP workloads with thousands of concurrent users[8].

Larger SAP HANA systems can leverage SAP HANA extension nodes or can get deployed as SAP HANA scale-out configurations. In a scale-out configuration, up to 16 nodes (GA, more upon SAP approval) work together to provide larger memory configurations. As of today, up to 16 x 4-CPU wide VMs with 2 TB RAM per node with an SAP HANA CPU sizing class L configuration, and up to 8 x 4-CPU wide VMs with up to 3 TB RAM per node and an SAP HANA CPU sizing class-M configuration are supported. Larger SAP HANA systems can get deployed upon SAP approval and an SAP HANA expert/ workload-based sizing as part of the TDI program. 8-socket wide SAP HANA VMs for scale-out deployments for OLTP and OLAP-type workloads are yet not SAP HANA supported.

An SAP HANA system deployed on a VMware SDDC based on VMware Cloud Foundation can get easily automated and operated by leveraging VMware vRealize® products. SAP HANA or hardware-specific [management packs](#) allow a top-bottom view of a virtualized SAP HANA environment where an AI-based algorithm allows the operation of SAP HANA in a nearly lights-out approach, optimizing performance and availability. A tight integration with [SAP Landscape Management](#) Automation Manager ([VMware Adapter™ for SAP Landscape Management](#)) helps to cut down operation costs even further by automating work-intensive SAP management/operation tasks.

Besides SAP HANA, most SAP applications and databases can get virtualized and are fully supported for production workloads either on dedicated vSphere hosts or running consolidated side by side.

Virtualizing all aspects of an SAP data center is the best way to build a future-ready and cloud-ready data center, which can get easily extended with cloud services and can run in the cloud. SAP applications can also run in a true hybrid mode, where the most important SAP systems still run in the local data center and less critical systems run in the cloud.

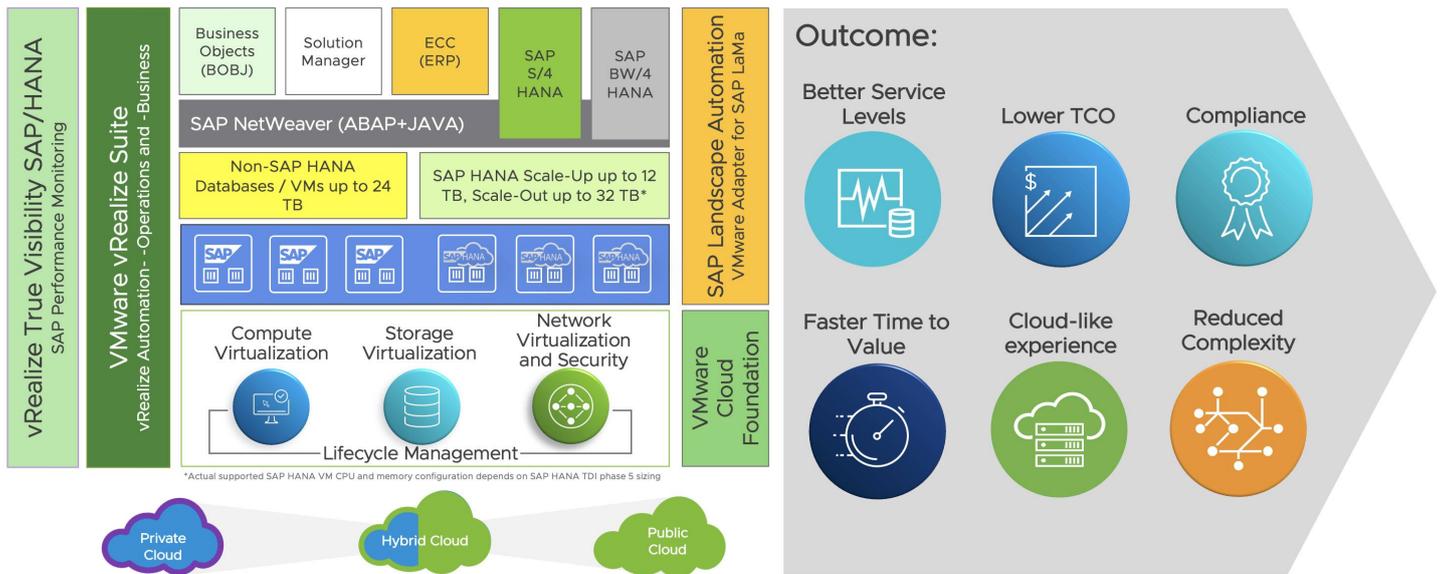


Figure 1: VMware SDDC based on VMware Cloud Foundation for SAP Applications

Deployment Options and Sizes

SAP HANA can get deployed on supported vSphere versions and validated CPU generations as scale-up and scale-out deployments as a single VM per server or multiple SAP HANA systems on a single physical server. Only 2-, 4- and 8-CPU socket VMware and SAP [supported or certified systems](#) can get used for SAP HANA production-level systems.

SAP HANA tenant databases (MDC) are fully supported to run inside a VMware VM (see SAP note [2104291](#), FAQ doc, page 2).

It is also supported to run SAP HANA VMs next to non-SAP HANA VMs, such as vSphere management VMs or SAP application servers, when these VMs run on different CPU sockets, or when an SAP HANA and SAP NetWeaver AS (ABAP or Java) run in one VM (see SAP notes [1953429](#) and [2043509](#)).

Table 6 summarizes the supported SAP HANA on vSphere VMs and host configuration and deployment options of SAP HANA (single-tenant and multitenant) instances on vSphere and provides some guidance on standard memory sizes supported by SAP and VMware based on current SAP-defined memory limitations for the top-bin Intel CPUs listed as certified appliance configurations.

Lower-bin CPUs or other CPU families may have different SAP HANA supported memory configurations. As mentioned, it is possible to deviate from these memory configurations when an SAP HANA expert/workload-based sizing gets done.

Table 6: Overview of SAP HANA on vSphere Deployment Options and Sizes

Virtual SAP HANA VM Sizes	Virtual SAP HANA VM Sizes	Virtual SAP HANA VM Sizes
<ul style="list-style-type: none"> Scale-Up: 0.5, 1 and 2-socket wide VMs* Min. 8 vCPUs and 128 GiB vRAM Max. 112 vCPUs and 3 TiB vRAM** Max. 160 vCPUs and 4 TiB vRAM*** Scale-Out: NA SAP HANA 1.0 SPS 12 Rev. 122.19 or later SAP HANA 2.0 SPS 4 or later up to 4 VMs per 2-socket host 	<ul style="list-style-type: none"> Scale-Up: 0.5, 1, 2, 3 and 4-socket wide VMs Min. 8 vCPUs and 128 GiB vRAM Max. 224 vCPUs and 6 TiB vRAM* Scale-Out: 4-socket wide VM with 224 vCPUs SAP HANA 1.0 SPS 12 Rev. 122.19 or later SAP HANA 2.0 SPS 4 or later up to 8 VMs per 4-socket host 	<ul style="list-style-type: none"> Scale-Up: 0.5, 1, 2, 3, 4, 5, 6, 7 and 8-socket wide VMs Min. 8 vCPUs and 128 GiB vRAM Max. 448 vCPUs and 12 TiB vRAM* Scale-Out: 4-socket wide VMs with 224 vCPUs SAP HANA 1.0 SPS 12 Rev. 122.19 or later SAP HANA 2.0 SPS 4 or later up to 16 VMs per 8-socket host
*Ice Lake 40 core CPUs require a 4-socket wide vSphere 7 VM **SAP HANA VM vRAM sizes depend on workload and sizing. Standard OLTP size max. 1.5 TiB and standard OLAP size max. 0.75 TiB per full SKL and later CPU socket. Broadwell maximal 192 vCPUs and 4-sockets. ***Ice Lake up to 80 vCPUs with 2 TiB per CPU socket. With a workload-based sizing larger vRAM sizes are possible. From this max. VM vRAM the memory costs of ESXi needs to get subtracted.		
Supported vSphere Versions		
<ul style="list-style-type: none"> vSphere 6.5 and 6.7 (max. vSphere support: 128 vCPUs and 4 TiB vRAM per VM) vSphere 6.7 U2/U3 and 7.0 (max. vSphere support: 256 vCPUs and 6 TiB vRAM per VM) vSphere 7.0 U2 and later (max. vSphere support: 768 vCPUs and 24 TiB vRAM per VM) Check SAP notes 2393917, 2937606 and 3102813 which CPU generation is supported with which vSphere version! PMem 100 support: SAP note 2913410 		
Physical 2-Socket Host	Physical 4-Socket Host	Physical 8-Socket Host
<ul style="list-style-type: none"> Ice Lake <= 4 TiB, max. 160 CPU threads Cooper Lake <= 3 TiB, max. 112 threads Cascade Lake <= 3 TiB, max. 112 threads <ul style="list-style-type: none"> PMem Series 100 Support Skylake <= 3 TiB, max. 112 threads Broadwell EX <= 2 TiB, max. 96 threads 	<ul style="list-style-type: none"> Cooper Lake <= 6 TiB, max. 224 threads Cascade Lake <= 6 TiB, max. 224 threads <ul style="list-style-type: none"> PMem Series 100 Support Skylake <= 6 TiB, max. 224 threads Broadwell EX <= 4 TiB, max. 192 threads 	<ul style="list-style-type: none"> Cooper Lake <= 12 TiB, max. 448 threads Cascade Lake <= 12 TiB, max. 448 threads Skylake <= 12 TiB, max. 448 threads Broadwell EX <= 8 TiB, max. 384 threads
Recommended Network Configuration: <ul style="list-style-type: none"> 1 x 1 GbE IPMI and 2 x 1 GbE ESXi mgt 2 x 10 GbE vMotion / HA 2 x 10 GbE HANA -> App Server Net + optimal network(s) like replication/backup 	Recommended Network Configuration: <ul style="list-style-type: none"> 1 x 1 GbE IPMI and 2 x 1 GbE ESXi mgt 2 x 25 GbE vMotion / HA 2 x 25 GbE HANA -> App Server Net + optimal network(s) like replication/backup 	Recommended Network Configuration: <ul style="list-style-type: none"> 1 x 1 GbE IPMI and 2 x 1 GbE ESXi mgt >= 2 x 25 GbE vMotion / HA >= 2 x 25 GbE HANA -> App Server Net + optimal network(s) like replication/backup

High-level Architecture

Figures 2 and 3 describe typical SAP HANA on vSphere architectures. Figure 2 shows scale-up SAP HANA systems (single SAP HANA VMs), and Figure 3 shows a scale-out example where several SAP HANA VMs work together to build one large SAP HANA database instance.

The illustrated storage needs to be SAP HANA certified. In the case of vSAN, the complete solution (server hardware and vSAN software) must be SAP HANA HCI certified.

The network column highlights the network needed for an SAP HANA environment. Bandwidth requirements should get defined regarding the SAP HANA size; for example, vMotion of a 2TB SAP HANA VM or a 12 TB SAP HANA VM. Latency should be as low as possible to support transaction heavy/sensitive workloads/use cases.

For more on network best practices, see the [Best practices of virtualized SAP HANA systems section](#).

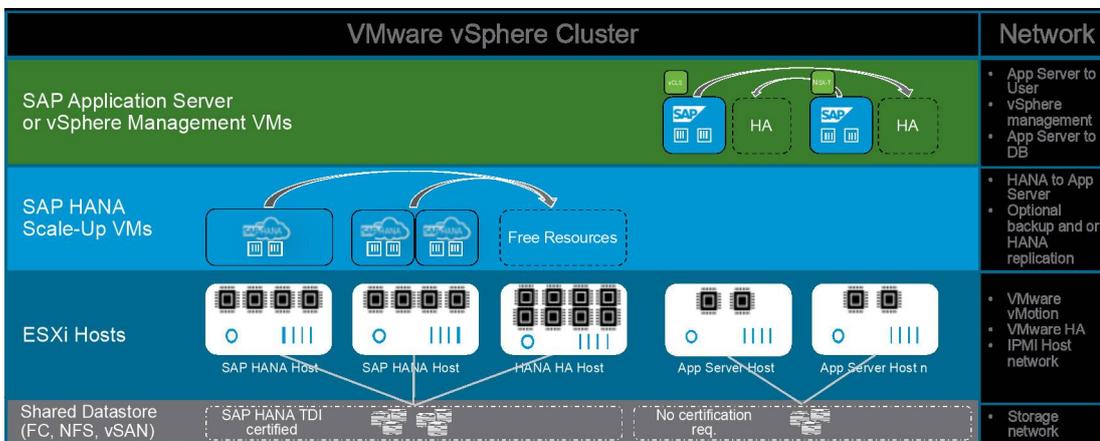


Figure 2: High-level Architecture of a Scale-up SAP HANA Deployment on vSphere

The next figure shows a Scale-Out SAP HANA deployment on vSphere ESXi 4 or 8 socket host systems.

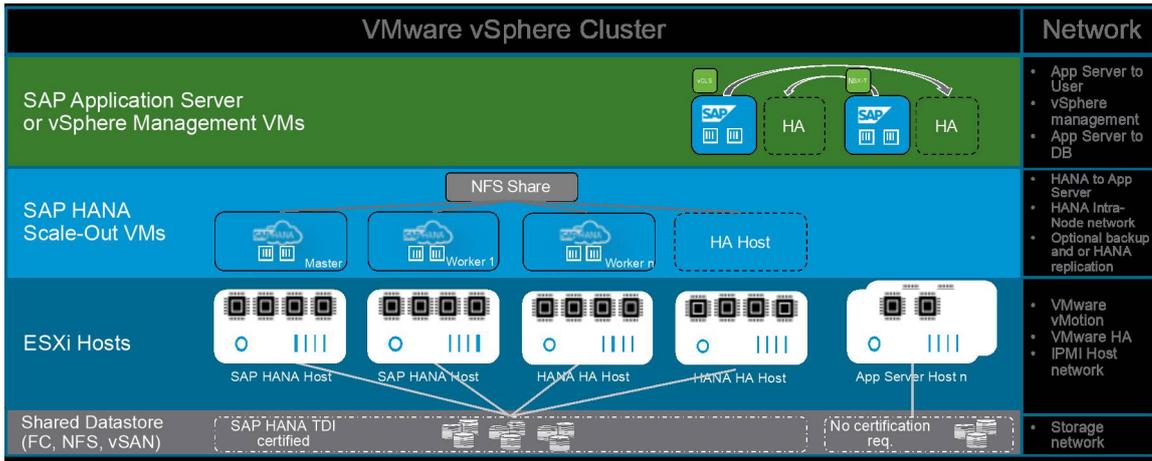


Figure 3: High-level Architecture of a Scale-out SAP HANA Deployment on ESXi 4- or 8-socket Host Systems

SAP HANA on vSphere Configuration and Sizing Guidelines

Selecting correct components and configuration is vital and the only way to achieve the performance and reliability requirements for SAP HANA. Which and how many SAP HANA VMs can get supported depends on the server configuration (RAM and CPU), the network configuration, and on the storage configuration.

Note: While it is possible to consolidate a certain number of SAP HANA VMs on a specific host with a given RAM and CPU configuration, it is important to understand that the network and storage configuration must also be able to support these SAP HANA VMs. Otherwise, a possible network or storage bottleneck will negatively impact performance of all running SAP HANA VMs on a host.

Sizing Compute and Memory

Since SAP HANA TDI Phase 5, it is now possible to perform a workload-based sizing (SAP note [2779240](#)) and not depend on appliance configurations with fixed CPU to memory ratios.

VMware virtual SAP HANA sizing gets performed just like with physically deployed SAP HANA systems. The major difference is that an SAP HANA workload needs to fit into the compute and RAM maximums of a VM and that the costs of virtualization (RAM and CPU costs of ESXi) need to get considered when planning an SAP HANA deployment.

If an SAP HANA system exceeds the available resources (virtual or physical deployments), this VM can get moved to a new host with more memory or higher performing CPUs. After this migration to this new host, the VM needs to get shut down, and the VM configuration must get changed to reflect these changes (more vCPU and/or virtual memory). If a single host is not able to satisfy the resource requirements of an SAP HANA VM, then scale-out deployments or SAP HANA extension nodes can get used.

Note: Current VMware VM maximums are 448 vCPUs and 12 TB of RAM. SAP HANA systems that fit into these maximums can get virtualized as a single scale-up system. Larger systems may be able to get deployed as scale-out systems.

SAP HANA Sizing Process

As noted, sizing a virtual SAP HANA system is just like sizing a physical SAP HANA system, plus the virtualization costs (CPU and RAM costs of ESXi). Figure 4 describes the SAP HANA sizing process.

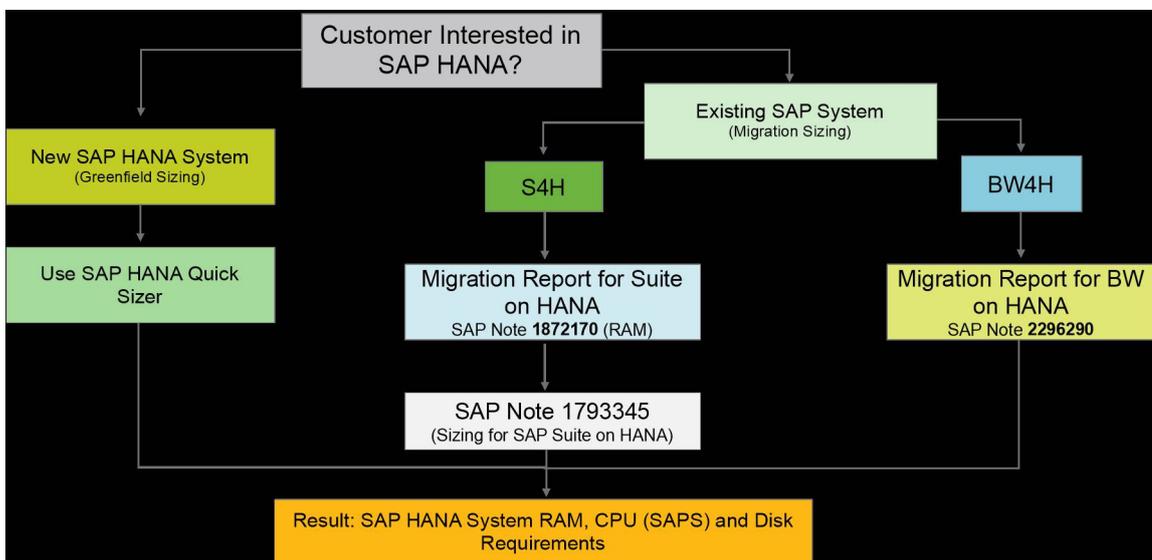


Figure 4: The SAP HANA Sizing Process

The results of an SAP HANA sizing are the needed compute (SAP Application Performance Standard [SAPS]), memory, and storage resources. Note that network sizing is not covered by the SAP sizing tools. The SAP-provided network sizing information focuses on throughput and can be found in the [SAP HANA network requirements white paper](#). Network latency is only expressed as a general guideline and is a user individual goal. In [SAP sales and distribution \(SD\) benchmarks](#), a time below 1,000ms for the average dialog response time must be maintained. See [published SD benchmarks](#) for the average response time.

In the [Network configuration and sizing section](#) of this guide, we refer to the SAP HANA network requirements white paper while we define the network infrastructure for a virtualized SAP HANA environment. SAP also provides a tool, ABAPMETER, that can be used

to measure the network performance of a selected configuration to ensure it follows the SAP defined and recommended parameters. See [SAP note 2879613: ABAPMETER in NetWeaver AS ABAP](#).

Note: The provided SAPS depend on the used SAP workload. This workload can have an OLTP, OLAP or mixed workload profile. From the VM configuration point of view, only the different memory, SAPS, network, and storage requirements are important and not the actual workload profile.

Required memory and storage resources are easy to determine. The storage capacity requirements for virtual or physical SAP HANA systems are identical. Physical memory requirements of a virtualized SAP HANA system are slightly higher and include the memory requirement of ESXi.

Table 7 shows the estimated memory costs of ESXi running SAP HANA workloads on different server configurations. Unfortunately, it is not possible to define the memory costs of ESXi upfront. The memory costs of ESXi are influenced by the physical server configuration (e.g., the number of CPUs, the number of NICs) and the used ESXi features.

Table 7: Estimated ESXi Host RAM Needs for SAP HANA Servers

Physical host CPU sockets / NUMA Nodes	Estimated ESXi memory need (rule of thumb)
2	16-128GB (default 48 GB)
4	64-256GB (default 96 GB)
8	128-512GB (default 192 GB)

Note: The memory consumption of ESXi as described in Table 7 can be from 16GB up to 512GB. The actual memory consumption for ESXi can only get determined when all VMs are configured with memory reservations and started on the used host. The last VM started may fail to start if the wrong memory reservation for ESXi was selected. In this case, the memory reservation per VM should be made lower to ensure all VMs that should run on the host fit into the host memory.

Let's use the following example to translate an SAP HANA sizing of an example 1,450GB SAP HANA system with a need of 80,000 SAPS on compute (SAP QuickSizer results).

The following formula helps to calculate the available memory for SAP HANA VMs running on an ESXi host: Total available memory for SAP HANA VMs = Total host memory - ESXi memory need

For this example, we assume that a 4-socket server with 6 TB total RAM and a 2-socket server with 1.5TB RAM is available. We use the default ESXi memory needed for a 4-socket, which is 96GB and 48GB for a 2-socket host:

- Available memory for VMs = **6048 GB** (6144 GB - 96 GB RAM) or **1488 GB** (1536 GB - 48 GB)
- Maximal HANA 4-socket VM memory = **6048 GB** or 2-socket HANA VM with **1488 GB** memory
- The maximum memory per 1 CPU socket VM is in this configuration **1512 GB** vRAM (6048 GB / 4) if ESXi memory costs get distributed equally between VMs or **744 GB** when the 2-socket server gets used.

VM Memory configuration example:

SAP HANA system memory sizing report result:

- 1500 GB RAM SAP HANA System memory

Available host server systems:

- 4-CPU socket Intel Xeon Platinum server with 6 TB physical memory.
- 2-CPU socket Intel Xeon Gold server with 1.5 TB physical memory.

VM vRAM calculation:

- $6048 / 4$ CPU sockets = 1488 GB per CPU \geq 1450 GB sized HANA RAM = 1 CPU socket
- $1512 / 2$ CPU sockets = 744 GB per CPU $<$ 1500 GB sized HANA RAM = 2 CPU socket

Temporarily VM configuration:

- In the 4-socket server case the memory of one socket is sufficient for the sized SAP HANA system. In the 2-socket server

case both CPU sockets must get used as one socket has only 744 GB available. In both cases we would select all available logical CPUs for this 1- or 2-socket wide VM.

When it comes to fine-tuning the VM memory configuration, the maximum available memory and therefore the memory reservation for a VM cannot get determined before the actual creation of a VM. As mentioned, the available memory depends on the ESXi host hardware configuration as well as ESXi enabled and used features.

After the memory configuration calculation, it is necessary to verify if the available SAPS capacity is sufficient for the planned workload.

The [SAPS](#) capacity of a server/CPU gets measured by SAP and SAP partners, which run specific [SAP defined benchmarks](#) and performance tests, such as the SAP SD benchmark, which gets used for all NetWeaver based applications and B4H. The test results of a public benchmark can be published and used for SAP sizing lectures.

Once the SAPS resource need of an SAP application is known, it is possible to translate the sized SAPS, just like the memory requirement, to a VM configuration.

The following procedure describes a way to estimate the available SAPS capacity of a virtual CPU. The SAPS capacity depends on the used physical CPU and is limited by the maximum available vCPUs per VM.

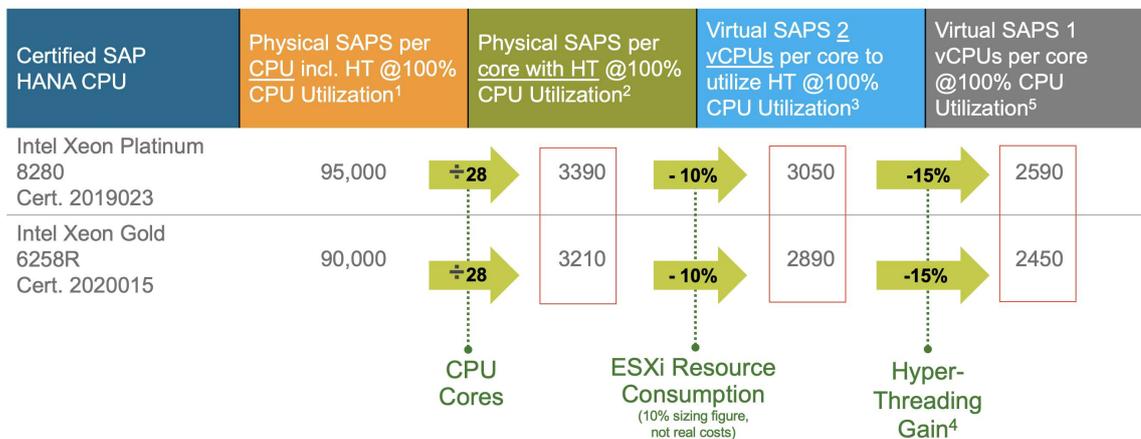


Figure 5: Physical SAPS to Virtual SAPS Example Conversion

The SAPS figures shown in Figure 5 are from published SD benchmarks. The first step is to look up the SAPS results of a benchmark of the CPU you want to use.

Next, divide this result by the number of cores of the selected CPU. In the example in Figure 5, both CPUs have 28 cores, which we use as the divisor. This provides the SAPS capacity of a hyperthreading-enabled physical CPU core (two CPU threads (PCPUs) running on one physical core).

To estimate the virtual SAPS capacity of these two CPU threads, we must subtract the ESXi CPU resource needs, which we have measured between 3-8 percent for OLTP or OLAP workloads. In Figure 5, to make the sizing easier, we use 10 percent for the virtualization costs for compute, which is subtracted from the previous result (two CPU threads running on one physical CPU core).

To define the SAPS capacity of a single vCPU running on a single CPU core, we must subtract the hyperthreading gain, which could be as little as 1 percent for very low-utilized servers or more than 30 percent for very high-loaded systems. For the sizing example in Figure 5, we assume a 15 percent hyperthreading gain.

Removing this 15 percent from the 2 vCPU result provides the SAPS capacity of a single vCPU that runs exclusively on a CPU core.

If hyperthreading should be used (default setting), ensure that `numa.vcpu.preferHT=TRUE` (per VM setting) is set to ensure [NUMA node locality](#). This is especially important for half-socket VM configurations and VM configurations that don't span all NUMA nodes of a server.

The following examples show how to calculate how many vCPUs will be needed to power the provided SAPS resource needs of the given SAP HANA workload.

VM CPU configuration example:

Assumed / sized SAP HANA system:

- 1450 GB RAM SAP HANA System memory

- 80,000 SAPS

Available host servers:

- 2-CPU socket Intel Cascade Lake Gold 6258R CPU Server with 1.5 TB memory and 56 pCPU cores / 112 threads and 180,000 pSAPS in total
- 4-CPU socket Intel Cascade Lake Platinum 8280L CPU Server with 6 TB memory and 112 pCPU cores / 224 threads and 380,000 pSAPS in total

VM vCPU configuration example:

Intel Cascade Lake Platinum 8280L CPU 4-socket Server:

- HANA CPU requirement as defined by sizing: 80,000 SAPS
- $380k / 4 = 95k \text{ SAPS} / 28 \text{ cores} = 3393 \text{ SAPS core}$
 - vSAPS per 2 vCPUs (including HT) = 3054 SAPS (3393 SAPS - 10%)
 - vSAPS per vCPU (without HT) = 2596 SAPS (3054 SAPS - 15%)
- VM without HT: #vCPUs = $80,000 \text{ SAPS} / 2596 \text{ SAPS} = 30,82 \text{ vCPUs}$, rounded up 32 cores / vCPUs
- VM with HT enabled: #vCPUs = $80,000 \text{ SAPS} / 3054 \text{ SAPS} \times 2 = 52,39 = 54 \text{ threads} / \text{vCPUs}$

VM vSocket calculation:

- $32 / 28 \text{ (CPU cores)} = 1.14$
- or $54 / 56 \text{ (threads)} = 0.96$
- The 1st result needs to get rounded up to 2 CPU sockets, since it does not leverage HT.
- The 2nd result uses HT and therefore the additional SAPS capacity of the hyperthreads are sufficient to use only one CPU socket for this system.

Intel Cascade Lake Gold 6258R CPU 2-socket Server:

- HANA CPU requirement as defined by sizing: 80,000 SAPS
- $180k / 2 = 90k \text{ SAPS} / 28 \text{ cores} = 3214 \text{ SAPS core}$
 - vSAPS per 2 vCPUs (including HT) = 2893 SAPS (3214 SAPS - 10%)
 - vSAPS per vCPU (without HT) = 2459 SAPS (2893 SAPS - 15%)
- VM without HT: #vCPUs = $80,000 \text{ SAPS} / 2459 \text{ SAPS} = 32,53 \text{ vCPUs}$, rounded up 34 cores / vCPUs
- VM with HT enabled: #vCPUs = $80,000 \text{ SAPS} / 2893 \text{ SAPS} \times 2 = 55,31 = 56 \text{ threads} / \text{vCPUs}$

VM vSocket calculation:

- $34 / 28 \text{ (CPU cores)} = 1.21$
- or $56 / 56 \text{ (threads)} = 1$

For the final VM configuration:

- The SAPS based sizing exercise showed that without HT enabled and used by the VM, 2 CPU sockets, each with 28 vCPU per socket, must get used with both server configurations.
- The sized SAP HANA database memory of 1,450GB on system RAM is below the available 1.5TB installed memory per CPU socket of the 4-socket server and will therefore lead to wasted memory (2 sockets -> 3TB of memory) when the 4-socket server gets used and if HT does not get used.
- In the case of the 2-socket server configuration, two sockets must get based on the memory calculation because one NUMA node has only around 750GB available. In the 2-socket server case, it is therefore not important to leverage hyperthreading for the given workload, but it is recommended because HT provides the additional SAPS capacity.
- As of today, CPU socket sharing with non-SAP HANA workloads is not supported by SAP and therefore, if the 2-socket server gets selected, then the two CPU sockets must be used exclusively for SAP HANA. Because of this, it is irrelevant if you leverage HT or not as all CPU resources of these two CPUs will get blocked by the SAP HANA VM. Therefore, it is recommended to use the Platinum CPU system instead of the Gold CPU because only one CPU socket would be needed instead of the two CPU sockets when the Gold CPU is used.

The final VM configuration comes out to:

- 4-socket Platinum CPU system: 1 CPU socket with 56 vCPUs and 1,450GB vRAM
- 2-socket Gold CPU system: 2 CPU sockets with 56 vCPUs and 1,450GB vRAM

If hyperthreading gets leveraged, then `numa.vcpu.preferHT=TRUE` must get set to ensure NUMA node locality of the CPU threads (PCPUs).

To simplify the SAPS calculation available in a VM, Table 8 shows possible VM sizes of specific CPU types and versions. The SAPS figures shown in Table 8 are estimated numbers based on published SAP SD benchmarks and are not measured figures. The SAPS figures got rounded down. The RAM needed for ESXi needs to get subtracted from the shown figures. The figures shown represent the virtual SAPS capacity of a CPU core with and without hyperthreading.

Note: The SAPS figures shown in Table 8 are based on published SAP SD benchmarks and can get used for Suite or BW on HANA, or S/ or BW/4HANA workloads. In the case of half-socket SAP HANA configurations, 15 percent from the SAPS capacity needs to get subtracted to consider the CPU cache misses caused by concurrent running VMs on the same NUMA node. For mixed HANA workloads, contact SAP or your hardware sizing partner. Also, if vSAN gets used (SAP HANA HCI), an additional 10 percent SAPS capacity should get reserved for vSAN.

Table 8: SAPS Capacity and Memory Sizes of Example SAP HANA on vSphere Configurations Based on Published SD Benchmarks and Selected Intel CPUs

	Intel Xeon E7-8890 v4 CPU[9]	Intel Xeon Gold 6258R CPU[10]	Intel Xeon Platinum 8280L CPU[11]	Intel Xeon Platinum 8380H CPU and HL[12]	Intel Xeon Platinum 8380 CPU[13]
SAP benchmark	2016067 (Dec. 20, 2016)	2020015 (May 5, 2020)	2019023 (April 2, 2019)	2020050 (Dec. 11, 2020)	2022019 (Sep.01.2022)
Max. supported RAM per CPU as of Intel datasheet	3 TB[14]	1,024 GB ¹⁴	4.5 TB ¹⁴	4.5 TB (HL) ¹⁴	6 TB ¹⁴
Max. CPU cores per socket as of Intel datasheet	24	28	28	28	40
Max. NUMA nodes per ESXi server	8[15]	2	8	8	2
vSAPS per CPU thread with and without HT	1,785 (core without HT), 315 (HT gain) Based on cert. 2016067	2,459 (core without HT), 434 (HT gain) Based on cert. 2020015	2,596 (core without HT), 458 (HT gain) Based on cert. 2019023	2,432 (core without HT), 429 (HT gain) Based on cert. 2020050	2,414 (core without HT), 424 (HT gain) Based on cert. 2022019
0.5-socket SAP HANA VM	1 to 16 x 12 physical core VM with min. 128 GB RAM and max. 512 GB[16]	1 to 4 x 14 physical core VM with min. 128 GB RAM and max. 512 GB ¹⁵	1 to 16 x 14 physical core VM with min. 128 GB RAM and max. 768 GB ¹⁵	1 to 16 x 14 physical core VM with min. 128 GB RAM and max. 768 GB ¹⁵	1 to 4 x 20 physical core/40 thread VM with min. 128 GB RAM and max. 1,024 GB ¹⁵
	vSAPS ¹⁶ 21,000,24 vCPUs	vSAPS ¹⁶ 34,000,28 vCPUs	vSAPS ¹⁶ 36,000,28 vCPUs	vSAPS ¹⁶ 34,000,28 vCPUs	vSAPS ¹⁶ 56,750, 40 vCPUs
1-socket SAP HANA VM	1 to 8 x 24 physical core/48 thread VM with min. 128 GB RAM and max. 1,024 GB	1 to 2 x 28 physical core/56 thread VM with min. 128 GB RAM and max. 1,024 GB	1 to 8 x 28 physical core/56 thread VM with min. 128 GB RAM and max. 1,536 GB	1 to 8 x 28 physical core/56 thread VM with min. 128 GB RAM and max. 1,536 GB	1 to 2 x 40 physical core/80 thread VM with min. 128 GB RAM and max. 2,048 GB
	vSAPS ¹⁶ 50,000,48 vCPUs	vSAPS ¹⁶ 81,000,56 vCPUs	vSAPS ¹⁶ 85,000,56 vCPUs	vSAPS ¹⁶ 80,000,56 vCPUs	vSAPS ¹⁶ 113,500, 80 vCPUs
2-socket SAP HANA VM	1 to 4 x 48 physical core/96 thread VM with min. 128 GB RAM and max. 2,048 GB	1 x 56 physical core/112 thread VM with min. 128 GB RAM and < 2,048 GB	1 to 4 x 56 physical core/112 thread VM with min. 128 GB RAM and max. 3,072 GB	1 to 4 x 56 physical core/112 thread VM with min. 128 GB RAM and max. 3,072 GB	1 x 80 physical core/80 thread VM with min. 128 GB RAM and < 4,096 GB
	vSAPS ¹⁶ 100,000,96 vCPUs	vSAPS ¹⁶ 162,000,112 vCPUs	vSAPS ¹⁶ 171,000,112 vCPUs	vSAPS ¹⁶ 160,000,112 vCPUs	vSAPS ¹⁶ 193,000, 80 vCPUs
4-socket SAP HANA VM	1 to 2 x 96 physical core/192 thread VM with min. 128 GB RAM and < 4,096 GB	-	1 to 2 x 112 physical core/224 thread VM with min. 128 GB RAM and max. 6,128 GB	1 to 2 x 112 physical core/224 thread VM with min. 128 GB RAM and max. 6,128 GB	1 x 80 physical core/160 thread VM with min. 128 GB RAM and < 4,096 GB
	vSAPS ¹⁶ 201,000,192 vCPUs	-	vSAPS ¹⁶ 342,000,224 vCPUs	vSAPS ¹⁶ 320,000,224 vCPUs	vSAPS ¹⁶ 227,000, 160 vCPUs
8-socket SAP HANA VM	-	-	1 x 224 physical core VM with min. 128GB RAM and < 12,096GB	1 x 224 physical core VM with min. 128 GB RAM and < 12,096 GB	-
	-	-	vSAPS ¹⁶ 684,000, 448 vCPUs	vSAPS ¹⁶ 640,000, 448 vCPUs	-

Using the information provided in Table 8 allows to quickly determine a VM configuration that can fulfill the SAP HANA sizing requirements on RAM and CPU performance for SAP HANA workloads.

Here’s a VM configuration example:

Sized SAP HANA system:

- 1,450 GB RAM SAP HANA system memory
- 80,000 SAPS (Suite on HANA SAPS)

Configuration Example 1:

- 2 x Intel Xeon Gold 6258R CPUs, RAM per CPU 786 GB, total of 1,536 GB
- 112 vCPUs, 162,000 vSAPS available by leveraging HT; 56 vCPUs, rounded 137,000 vSAPS without HT

Configuration Example 2:

- 1 x Intel Xeon Platinum 8280L CPU, RAM per CPU 1,536 GB
- 56 vCPUs, rounded 85,000 vSAPS with HT

Note: Review VMware [KB 55767](#) for details on 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.

Generally, it is recommended to leverage hyperthreading for SAP HANA on vSphere hosts. VMs may not utilize HT in case of security concerns or if the workload is small enough, so that HT is not required.

Storage Configuration and Sizing

Sizing a storage system for SAP HANA is very different from storage sizing for SAP classic applications. Unlike with SAP classic applications, SAP has defined strict storage key performance indicators (KPIs) for production-level SAP HANA systems for data throughput and latency. These KPIs must be achieved by any SAP HANA production-level VM running on a vSphere cluster.

The storage system connected must be able to deliver on these KPIs. The only variable in the storage sizing is the capacity, which depends on the size of the in-memory database. The SAP HANA storage performance KPIs can get verified with an SAP provided hardware configuration check tool (HWCCT) for HANA 1.0 and the Hardware and Cloud Measurement Tool (HCMT) for HANA 2.0. These tools are only available for SAP partners and customers, and the tools and the documentation with the currently valid storage KPIs can get downloaded, with a valid SAP user account, from SAP: For details, see SAP notes [1943937](#) and [2493172](#).

SAP partners provide SAP HANA ready and certified storage solutions or [certified SAP HANA HCI solutions](#) based on vSAN that met the KPIs for a specified number of SAP HANA VMs. See the [SAP HANA certified enterprise storage list](#) published on the SAP HANA Hardware Directory webpage.

Besides the storage capacity planning, the storage connection needs to get planned. Follow the available storage vendor documentation and planning guides to determine how many HBAs or NICs are needed to connect the planned storage solution. Use the guidelines for physically deployed SAP HANA systems as a basis if no VMware specific guidelines are available, and work with the storage vendor on the final configuration supporting all possible SAP HANA VMs running on a vSphere cluster.

vSphere connected storage and both raw device in-guest mapped LUNs and in-guest mounted NFS storage solutions are supported and can be used if, for instance, a seamless migration between physically and virtually deployed SAP HANA systems is required. Nevertheless, a fully virtualized storage solution works just as good as a natively connected storage solution and provides, besides other benefits, the possibility to abstract the storage layer from the operating system on which SAP HANA runs. For a detailed description on vSphere storage solutions, refer to the [vSphere documentation](#).

vSphere uses datastores to store virtual disks. Datastores provide an abstraction of the storage layer that hides the physical attributes of the storage devices from the virtual machines. VMware administrators can create datastores to be used as a single consolidated pool of storage, or many datastores can be used to isolate various application workloads.

vSphere datastores can be of different types, such as [VMFS](#), [NFS](#), [vSAN](#) or [vSphere Virtual Volumes™](#) based datastores. All these datastore types are fully supported with SAP HANA deployments. Refer to the [Working with Datastores](#) VMware documentation for details.

Table 9 summarizes the vSphere features supported by the different storage types. All these storage types are available for virtual SAP HANA systems.

Note: The VMware supported SAP HANA scale-out solution requires the installation of the SAP HANA shared file system on an NFS share. For all other SAP HANA scale-up and scale-out volumes, such as data or log, all storage types as outlined in Table 9 can be used as long the SAP HANA TDI Storage KPIs are achieved per HANA VM. Other solutions, such as the Oracle Cluster File System (OCFS) or the IBM General Parallel File System (GPFS), are not supported by VMware.

Table 9: vSphere Supported Storage Types and Features

Storage type	VM boot	vMotion	Datastore	RDM	vSphere HA and DRS
Local storage	Yes	No	VMFS versions 5 and 6	No	No
Fibre Channel	Yes	Yes	VMFS versions 5 and 6	Yes	Yes
NFS	Yes	Yes	NFS versions 3 and 4.1	No	Yes
vSAN[18]	Yes	Yes	vSAN 6.6 or later	No	Yes

In summary, with SAP HANA on vSphere, datastores can be used as follows:

- Create separate and isolated datastores for OS, SAP HANA binaries, shared folder, data, and logs.
- Enable multiple SAP HANA virtual machines to provision their virtual machine disk files on the same class of storage by maintaining the SAP HANA storage KPIs.

Figure 6 shows the SAP recommended SAP HANA Linux file system layout, which is the suggested layout when running SAP HANA virtualized. Grouping the file system layout into three groups helps you decide whether to use VMDK files or an NFS mount point to store the SAP HANA files and directories.

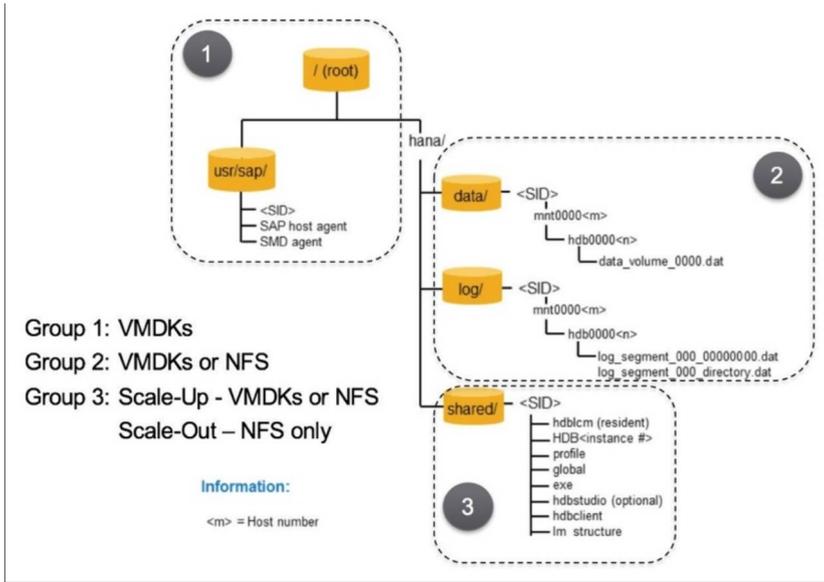


Figure 6: SAP Recommended SAP HANA File System Layout

In the next section, we use this file system layout and translate it into a disk volume/VMDK disk configuration.

Storage capacity calculation

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. This might result in more storage capacity than needed (even if the disk space is not needed, but the number of disks may be required to provide the required I/O performance and latency).

SAP has published several SAP HANA specific architecture and sizing guidelines[19], such as the [SAP HANA storage requirements](#). The essence of this guide is summarized in Figure 7 and Table 10, which can be used as a good starting point to plan the storage capacity needs of an SAP HANA system, as they show the typical disk layout of an SAP HANA system and the volumes needed. The volumes shown in Figure 7 should correspond with actual VMDK files and dedicated paravirtualized SCSI (PVSCSI) adapters, which ensures the best and most flexible configuration. The example configuration shown in Figure 7 uses three independent PVSCSI adapters and the minimum of four independent VMDKs. This helps to parallelize I/O streams by providing the highest flexibility in terms of operation.

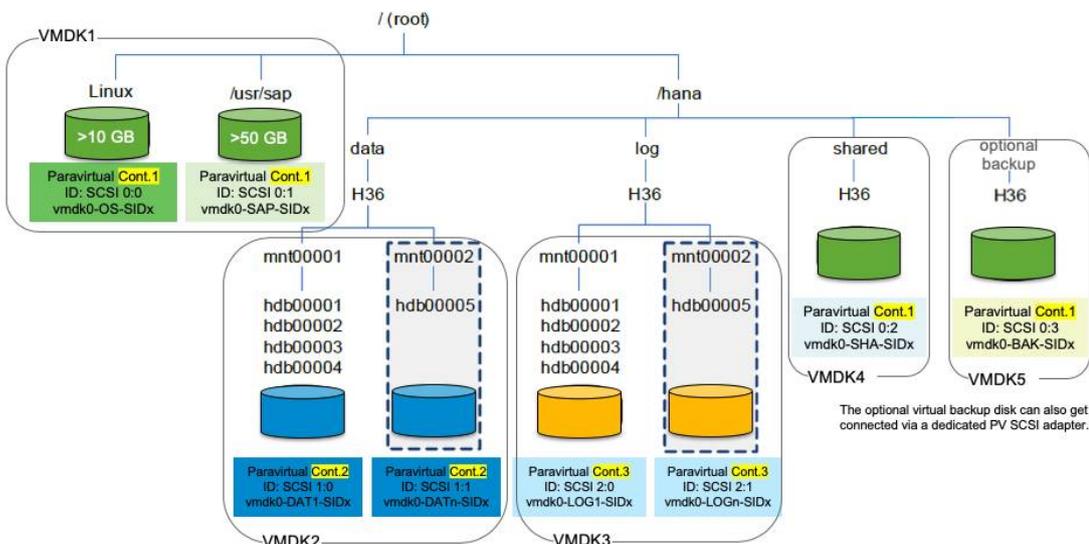


Figure 7: Storage Layout of an SAP HANA System. Figure © SAP SE, Modified by VMware

Table 10 provides storage capacity examples of the different SAP HANA volumes. Some of the volumes, such as the OS and usr/sap volumes, can be connected to and served by one PVSCSI controller. Others, such as the log and data volume, are served by

dedicated PVSCSI controllers to ensure high I/O bandwidth and low latency, which must get verified after an SAP HANA VM deployment with the hardware configuration check tools provided by SAP. Details about these tools can get found in SAP notes [1943937](#) and [2493172](#).

Beside VMDK based storage volumes especially for data, log or backup volumes ingest connected NFS volumes can get used as an alternative.

Table 10: Storage Layout of an SAP HANA System

Volume	DISK TYPE	SCSI controller if VMDK	VMDK name	SCSI ID if VMDK	Sizes as of SAP HANA storage requirements ¹⁹
/(root)	VMDK	PVSCSI Contr. 1	vmdk01-OS-SIDx	SCSI 0:0	Min. 10 GB for OS (suggested 100GB thin provisioned)
usr/sap	VMDK	PVSCSI Contr. 1	vmdk01-SAP-SIDx	SCSI 0:1	Min. 50 GB for SAP binaries (suggested 100GB thin provisioned)
shared	VMDK or NFS	PVSCSI Contr. 1	vmdk02-SHA-SIDx	SCSI 0:2	Min. 1 x RAM, max. 1TB (thick provisioned)
data/	VMDK or NFS	PVSCSI Contr. 2	vmdk03-DAT1-SIDx vmdk03-DAT2-SIDx vmdk03-DAT3-SIDx	SCSI 1:0 SCSI 1:1	Min. 1 x RAM (thick provisioned) Note: If you use multiple VMDKs, then use, for example, Linux LVM to build one large data disk.
log/	VMDK or NFS	PVSCSI Contr. 3	vmdk04-LOG1-SIDx vmdk04-LOG2-SIDx	SCSI 2:0 SCSI 2:1	[systems <= 512 GB] log volume (min) = 0.5 x RAM [systems >= 512 GB] log volume (min) = 512 GB (thick provisioned)
Backup	VMDK or NFS	PVSCSI Contr. 1 or 4	vmdk05-BAK-SIDx	SCSI 3:0	Size of backup(s) >= Size of HANA data + size of redo log Default path for backup: (/hana/shared). This path must get changed when an optional, dedicated, backup volume gets used (thin provisioned). To further optimize data throughput for backup, it is possible to use a dedicated PVSCSI adapter.

Depending on the used storage solution and future growth of the SAP HANA databases, it may be necessary to increase the storage capacity or better balance the I/O over more LUNs. The usage of Linux Logical Volume Manager (LVM) may help in this case, and it is fully supported to build LVM volumes based on VMDKs.

To determine the overall storage capacity per SAP HANA VM, sum up the sizes of all specific and unique SAP HANA volumes as outlined in Figure 7 and Table 10. To determine the minimum overall vSphere cluster datastore capacity required, multiply the SAP HANA volume requirements in Table 10 with the amount of VMs running on all hosts in a vSphere cluster.

Note: The raw storage capacity need depends on the used storage subsystem and the selected RAID level. Consult your storage provider to determine the optional physical storage configuration for running SAP HANA. NFS mounted volumes do not need PVSCSI controllers.

This calculation is simplified as:

vSphere datastore capacity = total SAP HANA VMs running in a vSphere cluster x individual VM capacity need (OS + USR/SAP + SHARED + DATA + LOG)

For example, a sized SAP HANA system with RAM = 1.5 TB would need the following:

- VMDK OS >= 10 GB (recommended 100 GB thin provisioned)
- VMDK SAP >= 60 GB (recommended 100 GB thin provisioned)
- VMDK Shared = 1,024 GB (thick provisioned)
- VMDK Log = 512 GB (thick provisioned)
- VMDK Data = 1,536 GB (thick provisioned)
- VMDK Backup >= 2,048 GB (thin provisioned) (optional)

For this example, the VM capacity requirement = 3.2 TB / with optional backup 5.2 TB

This calculation needs to get done for all possible running SAP HANA VMs in a vSphere cluster to determine the cluster-wide storage capacity requirement. All SAP HANA production VMs must fulfill the capacity as well as the throughput and latency requirements as specified by SAP note [1943937](#).

Note: SAP HANA storage KPIs must get guaranteed for all production-like SAP HANA VMs. Use the HWCCT to verify these KPIs. Otherwise, the overall performance of an SAP HANA system may be negatively impacted.

SAP HANA Hardware and Cloud Measurement Tool

Using the SAP provided hardware check tools (see SAP notes [1943937](#) and [2493172](#) and this [blog post](#)) allows you to verify, besides other aspects, if the storage performance and latency of an SAP HANA VM fulfills the SAP defined KPIs for log and data.

For its SAP HANA validation, VMware uses (besides vSAN) a modern external Fibre Channel-based flash array from [Pure Storage](#). This ensures that the conducted tests and validations use a modern external storage array and software-defined storage.

For example, what is possible with a VMware optimized Pure X50 storage system when compared to a Bare Metal (BM) running system with the same configuration following the test results for log latency, which SAP has defined as a maximum of 1,000µs.

Table 11: HCMT file System Latency Example of a BM System Compared to a Virtualized SAP HANA System

Configuration	HCMT 16K block log volume latency
BM FC connected Pure X50 storage unit	16K block log overwrite latency = 334µs
VM FC connected Pure X50 storage unit	16K block log overwrite latency = 406µs
Deviation	72µs
Deviation in percent	22%
SAP HANA KPI	1,000µs (Bare Metal (BM) 3 times and VM 2.5 times faster as required)

In Table 11, the virtualized SAP HANA system shows a 22 percent higher latency than the BM SAP HANA system running on the same configuration but with 406 µs below the SAP defined KPI of 1,000 µs.

Table 12: HCMT File System Overwrite Example of a BM System Compared to a Virtualized SAP HANA System

Configuration	HCMT 16K block log volume overwrites
BM FC connected Pure X50 storage unit	16K block log overwrite throughput= 706 MB/s
VM FC connected Pure X50 storage unit	16K block log overwrite throughput= 723 MB/s
Deviation	17 MB/s
Deviation in percent	2.4%
SAP HANA KPI	120MB/s (BM and VM over 5 times higher as required)

In Table 12, the virtualized SAP HANA system has a slightly higher throughput than the BM SAP HANA system running on the same configuration and is way over the SAP defined KPI of 120 MB/s for this test.

For these tests, we used an 8-socket wide VM running on an Intel Cascade Lake based [Fujitsu PRIMEQUEST 3800B2](#) system connected via FC on a [Pure X50 storage](#) unit configured as outlined in the Pure Storage and VMware [best practices configuration guidelines](#). Please especially apply the OS best practice configuration for SLES or RHEL, and ensure you have a log and data volume per SAP HANA VM as previously described.

Network Configuration and Sizing

To build an SAP HANA ready vSphere cluster, dedicated networks are required for SAP application, user traffic, admin, and management, as well as for NFS or software-defined storage (SDS), such as vSAN, if it is used. Follow the [SAP HANA network requirements](#) white paper to decide how many networks have to be added to support a specific SAP HANA workload in a VM and ultimately on the hosts.

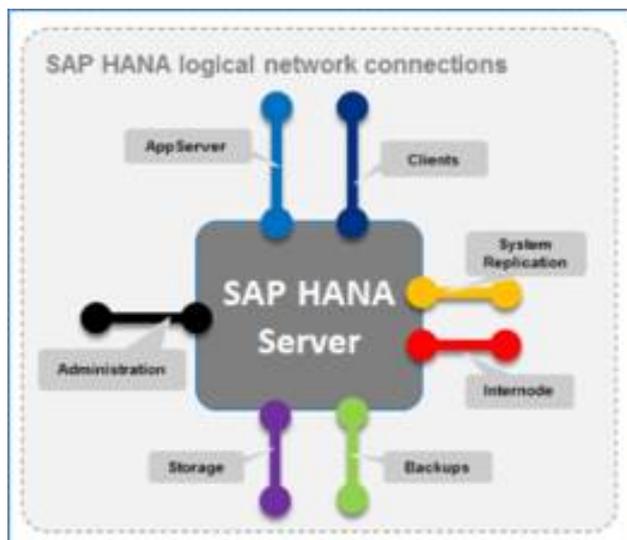


Figure 8: Logical Network Connections per SAP HANA Server, picture © SAP AG

Unlike in a physical SAP HANA environment, an administrator must plan the SAP HANA OS exposed networks as well as the ESXi host network configuration that all SAP HANA VMs that will run on the ESXi host will share. The virtual network, especially the virtual switch configuration, is done on the ESXi level, and the virtual network cards configuration that is exposed to HANA is done inside the VM. Each is described in the following sections.

Note: Planning the ESXi host network configuration must include the possibility that several SAP HANA VMs get deployed on a

single host. It may be required to have more, or higher bandwidth network cards installed when planning for a single SAP HANA instance running on a single host.

vSphere offers standard and distributed switch configurations. Both switches can be used when configuring an SAP HANA on vSphere environment. Nevertheless, it is strongly recommended to use a vSphere Distributed Switch™ for all VMware kernel-related network traffic (such as vSAN and vMotion). A vSphere 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.

Virtual machines have network adapters you connect to port groups on the virtual switch. Every port group can use one or more physical NIC(s) to handle their network traffic. If a port group does not have a physical NIC connected to it, virtual machines on the same port group can communicate only with each other but not with the external network. Detailed information can be found in the [vSphere networking guide](#).

Table 13 shows the recommended network configuration for SAP HANA running virtualized on an ESXi host with different network card configurations, based on SAP recommendations and the VMware-specific needed networks, such as dedicated storage networks for SDS or NFS and vMotion networks. In the case of SAP HANA VM consolidation on a single ESXi host, several network cards may be required to fulfill the network requirement per HANA VM.

The usage of VLANs is recommended to reduce the total amount of physical network cards needed in a server. The requirement is that enough network bandwidth is available per SAP HANA VM, and that the installed network cards of an ESXi host provide enough bandwidth to serve all SAP HANA VMs running on this ESXi host. Overbooking network cards will result in bad response times or increased vMotion or backup times.

Note: The sum of all SAP HANA VMs that run on a single ESXi host must not oversubscribe the available network card capacity.

Table 13: Recommended SAP HANA on vSphere Network Configuration based on 10 GbE NICs

		HOST iPMI / REMOTE CONTR. NETWORK	VSPHERE ADMIN + VMWARE HA NETWORK	APPLICATION SERVER NETWORK	vMotion Network	BACKUP NETWORK	SYSTEM REPLICATION (HSR) NETWORK	SCALE-OUT INTER-NODE NETWORK	NFS / vSAN / SDS Network	
Host Network Configuration (all networks are dual port NICs, configured as failover teams)	Network Label	Physical Host MGT	vSphere Admin	SAP App-Server	vMotion	Backup Network (optional)	HANA Repl. (optional)	Scale-Out (optional)	Storage Network (optional)	
	Typical Bandwidth	1 GbE	1 GbE	>=10 GbE	>=10 GbE	>=10 GbE	>=10 GbE	>=10 GbE	>=10 GbE	
	Typical MTU Size	Default (1500)	Default (1500)	Default (1500)	9000	9000	9000	9000	9000	
	VLAN ID# examples	-	200	201	202	203	204	205	206	
	Physical NIC No.#, bandwidth	#0, 1 GbE	#1, 10 GbE			#2, 10 GbE	#3, 10 GbE	#4, 10 GbE	#5, 10 GbE	#6, 10 GbE
	Active physical NIC port	1	1		1	1	1	1	1	
	Standby physical NIC port	2	2		2	2	2	2	2	
VM Guest Network Cards	virtual NIC (inside VM)	-	-	X	-	X	X	X	X	

It is recommended to create a vSphere Distributed Switch per dual port physical NIC and configure port groups for 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 for the uplink failover order as shown in Table 13. 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 14 shows an example of how to group the network port failover teams. Depending on the required optional networks needed for VMware or SAP HANA system replication or scale-out internode networks, this table/suggestion will differ. At a minimum, three NICs are required for a virtual HANA system leveraging vMotion and vSphere HA. For the optional use cases, additional network adapters are needed.

Note: Configure the network port failover teams to allow physical NIC failover. To support this, it is recommended to use NICs with the same network bandwidth, such as only 10 GbE NICs. Group failover pairs depending on the needed network bandwidth do not group, for instance, two high-bandwidth networks, such as the internode network with the app server network.

Table 14: Minimum ESXi Server Uplink Failover Network Configuration based on 10 GbE NICs

Property	NIC	VLAN[20]	Active uplink	Standby uplink
vSphere admin + vSphere HA	1	200	Nic1-Uplink1	Nic2-Uplink2
SAP application server network	1	201	Nic1-Uplink1	Nc2-Uplink2
vMotion network	2	202	Nic2-Uplink1	Nic1-Uplink2
Backup network (optional)	3	203	Nic3-Uplink1	Nic4-Uplink2
HSR network (optional)	4	204	Nic4-Uplink1	Nic3-Uplink2
Scale-out network (optional)	5	205	Nic5-Uplink1	Nic6-Uplink2
Storage network (optional)	6	206	Nic6-Uplink1	Nic5-Uplink2

Using different VLANs, as shown, to separate the VMware operational traffic (e.g., vMotion and vSAN) from the SAP and user-specific network traffic, is recommended. Using higher bandwidth network adapters reduces the number of physical network cards, cables, and switch ports. See Table 15 for an example with 25 GbE network adapters.

Table 15: Example SAP HANA on vSphere Network Configuration based on 25 GbE NICs

		HOST iPMI / REMOTE CONTR. NETWORK	VSPHERE ADMIN + VMWARE HA NETWORK	APPLICATION SERVER NETWORK	vMotion Network	BACKUP NETWORK	SYSTEM REPLICATION NETWORK	SCALE-OUT INTER-NODE NETWORK	NFS / vSAN / SDS Network	
Host Network Configuration (all networks are dual port NICs, configured as failover teams)	Network Label	Physical Host MGT	vSphere Admin	SAP App-Server	vMotion	Backup Network (optional)	HANA Repl. (optional)	Scale-Out (optional)	Storage Network (optional)	
	Typical Bandwidth	1 GbE	1 GbE	>=24 GbE	>=15 GbE	>=10 GbE	>=15 GbE	>=10 GbE	>=25 GbE	
	Typical MTU Size	Default (1500)	Default (1500)	Default (1500)	9000	9000	9000	9000	9000	
	VLAN ID# examples	-	200	201	202	203	204	205	206	
	Physical NIC No.#, bandwidth	#0, 1 GbE	#1, 25 GbE			#2, 25 GbE		#3, 25 GbE		#4, 25 GbE
	Active physical NIC port	1	1		1		1		1	
	Standby physical NIC port	2	2		2		2		2	
VM Guest Network Cards	virtual NIC (inside VM)	-	-	X	-	X	X	X	X	

Leveraging 25 GbE network adapters helps reduce the number of NICs required as well as lower the required network cables and switch ports. With higher bandwidth NICs, it is also possible to support more HANA VMs per host.

Table 16 shows an example of how to group the network port failover teams.

Table 16: Minimum ESXi Server Uplink Failover Network Configuration based on 25 GbE NICs

Property	NIC	VLAN[21]	Active uplink	Standby uplink
vSphere admin + vSphere HA	1	200	Nic1-Uplink1	Nic2-Uplink2
SAP application server network	1	201	Nic1-Uplink1	Nic2-Uplink2
vMotion network	2	202	Nic2-Uplink1	Nic1-Uplink2
Backup network (optional)	2	203	Nic2-Uplink1	Nic1-Uplink2
HSR network (optional)	3	204	Nic3-Uplink1	Nic4-Uplink2
Scale-out network (optional)	3	205	Nic3-Uplink1	Nic4-Uplink2
Storage network (optional)	4	206	Nic4-Uplink1	Nic3-Uplink2

Using different VLANs, as shown, to separate the VMware operational traffic (e.g., vMotion and vSAN) from the SAP and user-

specific network traffic is recommended. Using 100 GbE bandwidth network adapters will further help to reduce the number or physical network cards, cables, and switch ports. See Table 17 for an example with 100 GbE network adapters.

Table 17: Example SAP HANA on vSphere Network Configuration based on 100 GbE NICs

		HOST iPMI / REMOTE CONTR. NETWORK	VSPHERE ADMIN + VMWARE HA NETWORK	APPLICATION SERVER NETWORK	vMotion Network	BACKUP NETWORK	SYSTEM REPLICATION NETWORK	SCALE-OUT INTER-NODE NETWORK	NFS / vSAN / SDS Network
Host Network Configuration (all networks are dual port NICs, configured as failover teams)	Network Label	Physical Host MGT	vSphere Admin	SAP App-Server	vMotion	Backup Network (optional)	HANA Repl. (optional)	Scale-Out (optional)	Storage Network (optional)
	Typical Bandwidth	1 GbE	1 GbE	>=24 GbE	>=15 GbE	>=10 GbE	>=15 GbE	>=10 GbE	>=25 GbE
	Typical MTU Size	Default (1500)	Default (1500)	Default (1500)	9000	9000	9000	9000	9000
	VLAN ID# examples	-	200	201	202	203	204	205	206
	Physical NIC No.#, bandwidth	#0, 1 GbE	#1, 100 GbE			#2, 100 GbE			
	Active physical NIC port	1	1			1			
	Standby physical NIC port	2	2			2			
VM Guest Network Cards	Virtual NIC (inside VM)	-	-	X	-	X	X	X	X

Table 18: Minimum ESXi Server Uplink Failover Network Configuration based on 100 GbE NICs

Property	NIC	VLAN[22]	Active uplink	Standby uplink
vSphere admin + vSphere HA	1	200	Nic1-Uplink1	Nic2-Uplink2
SAP application server network	1	201	Nic1-Uplink1	Nic2-Uplink2
vMotion network	1	202	Nic1-Uplink1	Nic2-Uplink2
Backup network (optional)	2	203	Nic2-Uplink1	Nic1-Uplink2
HSR network (optional)	2	204	Nic2-Uplink1	Nic1-Uplink2
Scale-out network (optional)	2	205	Nic2-Uplink1	Nic1-Uplink2
Storage network (optional)	2	206	Nic2-Uplink1	Nic1-Uplink2

Using different VLANs, as shown, to separate the VMware operational traffic (for example, vMotion and vSAN) from the SAP and user-specific network traffic is recommended.

VMXNET3 SAP Workload-specific Network Latency Considerations

The SAP HANA on vSphere validation involves different tests, some of which focus on CPU and memory performance while others involve storage and network performance and scalability tests. These tests use SAP-specific OLTP and OLAP-type workloads and scale from single user tests up to thousands of concurrent users (up to 78,000), pushing the virtualized SAP HANA system to its limits.

A virtualized network card typically adds between 60 μ s (no load) and up to 300 μ s latency[23] (high CPU load \geq 80 percent) to every network package when compared to a bare-metal installed SAP HANA system, which impacts SAP OLTP and OLAP-type workloads issued by remote application servers/users.

This section discusses this impact and provides information on how to mitigate it by optimizing the virtual and physical network configuration.

SAP workloads and the impact on network performance

SAP workloads can get distinct by their characteristics. These can be OLTP-type workloads, which represent the classic SAP ERP type systems, and OLAP-type workloads, which represent analytical workloads that we see mainly with BW systems.

S/4 HANA combines these two workload types, and we need to consider both of their characteristics. Typical OLTP workloads are small network packages with the recommended MTU size of 1,500, whereas the recommended MTU size for OLAP-type workloads would be 9,000. It is important to understand how the S/4 HANA system gets used and select the MTU size depending on this.

Also, it is important to understand how many concurrent users will use the HANA database and how much network traffic these users will cause. In recent tests of vSphere with SAP HANA and with OLTP-type workloads, SAP and VMware observed an increase of OLTP transactional request times, which can show an overhead of up to 100 ms when compared to bare-metal installed SAP HANA systems.

This OLTP transactional request time increase can be observed when using the virtual NIC type VMXNET3. The reason for this is that virtual networking adds the mentioned latency in the range of 60 μ s (no load) and up to 300 μ s (high load, measured with 23–64K users with 4-socket and 35–78K users with 8-socket wide VMs) to each network package sent or received. These observations are documented in VMware knowledge base article [83957](#).

Unlike with storage (see the [Storage configuration and sizing section](#)), SAP did not define SAP HANA-specific network KPIs for throughput and latency, which an SAP HANA system has to maintain apart from the general recommendation to use a 10 GbE network for SAP application and SAP HANA database servers. Therefore, it is hard to define a specific network configuration, and specific tests are required to recommend a suitable network configuration for a virtualized SAP HANA environment.

The next section describes how the VMXNET3 impact got measured and how to optimize the network configuration for an SAP HANA VM for its given use case.

SAP S/4 and B/4 HANA Workload and Validation Tests

A typical SAP S/4 HANA environment is a three-tier landscape with the application tier on a separate host(s) from the database tier, with users accessing the application servers when they work within the SAP HANA database stored data. See Figure 9 for a typical SAP three-tier configuration.

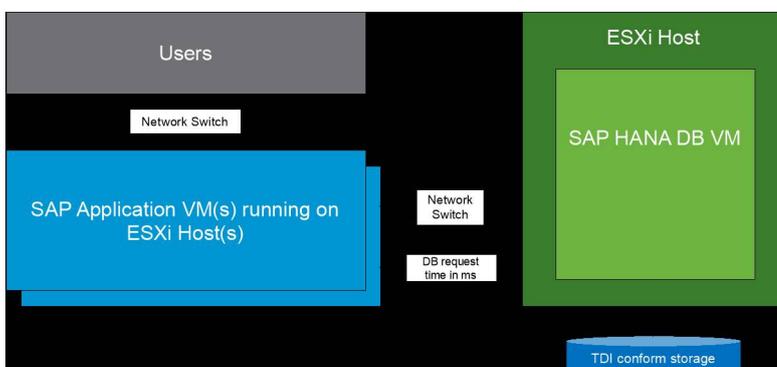


Figure 9: SAP Three-tier Architecture on vSphere

The conducted tests with S/4 SAP HANA and BW run on such a three-tier environment to simulate SAP load. The used workloads were OLTP, OLAP, and a mix of OLTP and OLAP transactions. The goals of these tests were to measure the performance of

virtualization on SAP HANA, running on such a three-tier environment.

The executed tests, simulating thousands of SAP users, executing OLTP and OLAP transactions, up to the maximal possible CPU utilization level, accessing the SAP HANA DB VM, at the same time.

These specific tests measure OLTP throughput, OLAP query per hour and database request time related values (SUM(Database Request time of all steps of one transaction (DIA+UPD+UPD2) / Number of transaction calls)).

The application server instances receive these simulated requests by a load driver and execute the SAP-specific transactions against the SAP HANA database. These transactions can cause several hundreds of database logical units of work ([LUW](#)).

Note: The measured database request time is included in the SAP customer measurable dialog response time, which should be normally under 1000 ms, and focuses on the time measured for transactions between the SAP application server and the HANA database instance.

The number of simulated SAP users depend on the SAP HANA VM size and start at approximately 20,000 concurrent users for a 4-socket wide VM and approximately 35,000 concurrent users with an 8-socket wide VM, and then increased to approximately 44,000 (4-socket) and approximately 60,000 (8-socket) concurrent users. The number of concurrent users represents a moderate to high CPU load on a modern Cascade or Cooper Lake based server of 35 percent and 65 percent on the SAP HANA database instance.

In addition to the 35 percent and 65 percent CPU utilization point, the number of users increases until the OLTP/OLAP throughput drops. In the case of the used 4-socket configuration, this happens with approximately 64,000 concurrent users at a CPU utilization of approximately 80 percent. With the 8-socket system, this point is reached with approximately 78,000 concurrent users.

To provide more details about the actual workload, one of the tests used for the validation is meant to simulate a day-in-the-life workload using common SAP online/interactive business transactions, such as the following transactions:

- VA01 - Create sales order
- VL01N - Create delivery
- VA03 - Display order
- VL02N - Post goods issue
- VA05 - List open orders
- VF01 - Create invoice

As mentioned, these transactions were executed on an 8-socket wide Cascade Lake 8280L CPU based server/VM until the maximum OLTP transactions per hour result were reached. The test-relevant results are of users at the 35 percent and 65 percent CPU utilization point. The result of the maxed-out measurement defines the maximum throughput number for OLTP/OLAP and defines the maximum possible SAPS figure for such a physical or virtualized SAP HANA system. In the B/4 HANA workload case, a public available [BWH benchmark](#) gets used and executed.

SAP S/4 and B/4 HANA Validation Tests Findings (OLTP)

Executing the described mixed workload test and BWH benchmark provides information of virtual performance and allows to provide recommendations to lower the impact of virtualization on for instance networking performance.

Whereas executing the BWH benchmark did not show any issues with either network throughput or latency, the executed S/4 HANA mixed workload test exposed a network latency increase that occurs with the VMXNET3 driver, which is documented in VMware knowledge base article [83957](#).

Following tables and figures summarize the results of these findings based on the latest Intel Cooper Lake CPUs.

Table 19 shows the minimum and maximum VMXNET3 latency deviation measured with Netperf on an SAP HANA system installed natively and virtually when running with no load and under a high user load on an 8-socket wide

test VMXNET3 VM and Bare metal server. This load represents a CPU utilization of up to 65 percent. The native installed SAP HANA system had on average a Netperf TCP roundtrip time of around 26 µs (no load) and up to 95 µs (high load). The HANA VM shows on average 84 µs and 319 µs TCP roundtrip time while under load.

Table 19: Network Latency of Netperf TCP VMXNET3 Compared to Bare Metal, Measured in µs

8-socket server/VM	Baseline latency with no load μ s	Latency at peak load with 65,000 concurrent users (μ s)
Bare metal	26	95
VMXNET3	84	319
Delta	58	224

This measured three times higher TCP roundtrip time/latency per network package sent and received while idle or while executing the test is impacting the overall SAP HANA database request time negatively and can, as of today, only be reduced by using a passthrough network card, which shows only a slightly higher latency as a native NIC, or by optimizing the underlying physical network to lower the overall latency of the network.

Looking on the observed VMXNET3 latency overhead of an 8-socket wide VM running on an Intel Cooper Lake based server with 416 vCPUs (32 CPU threads (PCPUs) were reserved to handle this massive network load) to a natively installed SAP HANA system running on the same system, we see how these microseconds accumulate to a database request time deviation between 27 ms and approximately 82 ms. See Figure 10 for this comparison.

Note: While executing this SAP-specific workload, massive network traffic gets created. Reducing the number of vCPUs from 448 to 416 helped the ESXi kernel to manage this excessive network load caused by up to 91,000 concurrent users.

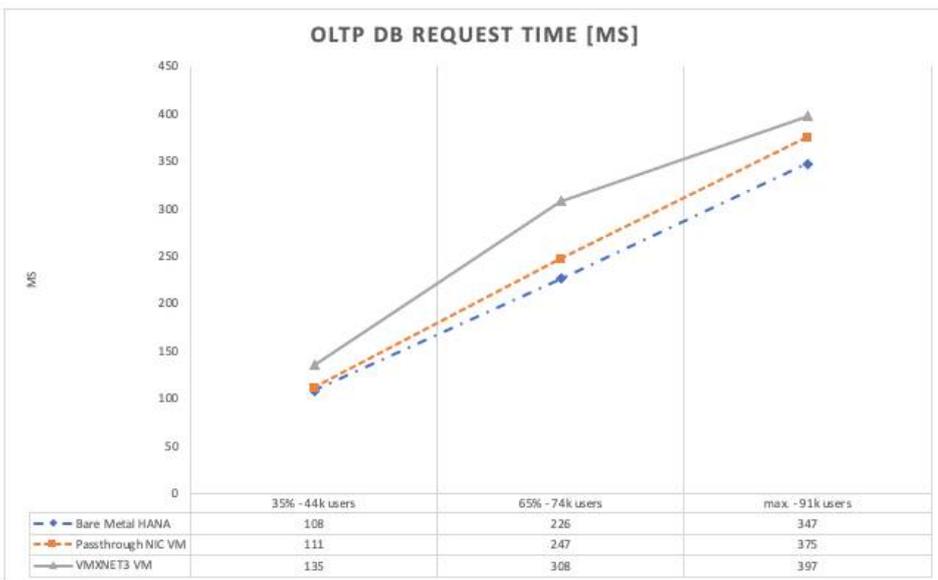


Figure 10: Mixed workload OLTP Database Request Time in Milliseconds

While the database request time gets impacted by up to 25 percent (27 ms higher) at 35 percent CPU utilization when VMXNET3 is used, the OLTP TPH and OLAP QPH results were not impacted. At approximately 65 percent CPU utilization, the database request time increased to 36 percent (82 ms higher) with a TPH deviation of approximately -1 percent. At the maximum user capacity with 91K users at approximately 80 percent CPU utilization, the TPH/QPH impact was approximately -8 percent. The OLAP request times were very little impacted. See next figures for details.

Using a passthrough network device instead of a VMXNET3 network card reduces the database request time. The measured deviation between an SAP HANA VM with a PT NIC compared to a BM SAP HANA system was approximately 3 percent (3 ms) at 35 percent CPU utilization and approximately 9 percent (21 ms) at 65 percent CPU utilization for OLTP transactional request times and at 80 percent the time deviation was still below 10% keeping the THP/QPH deviation below -3 percent at this maxed-out point.

Also, reserving CPU threads (PCPUs) to handle network traffic on the ESXi side is not necessary because the network traffic gets handled inside the VM and by the OS and not the ESXi like when VMXNET3 gets used. See Figures 10 and 11.

Note: The measured database request time is the time between the SAP application server and the HANA database instance. What did not get measured was the time a user has to wait until an application server would response on a user-initiated transaction. The user-to-application server time is significantly higher than the database request time between application servers

and a database instance.

Figure 11 shows the OLAP request times, which are very little impacted by the VMXNET3 or PT network cards when compared to a BM running SAP HANA system.

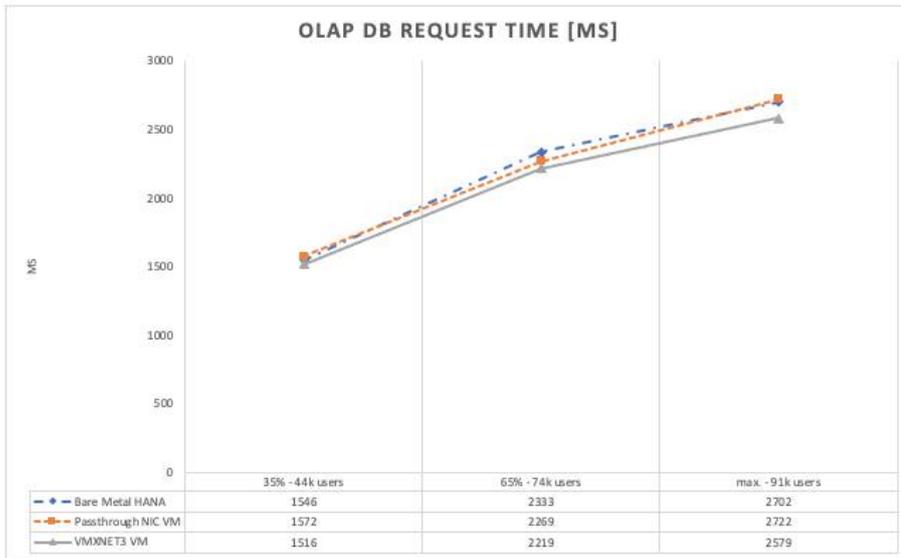


Figure 11: Mixed Workload OLAP Database Request Time in Milliseconds

While the virtualization overhead is already measurable with lower user numbers/traffic, the test results have shown that the main impact of this specific S/4 HANA mixed workload test can get measured at the higher user load numbers that generate, in a very short time, massive OLTP requests and that the impact to OLAP database request time is very little as shown in Figure 11.

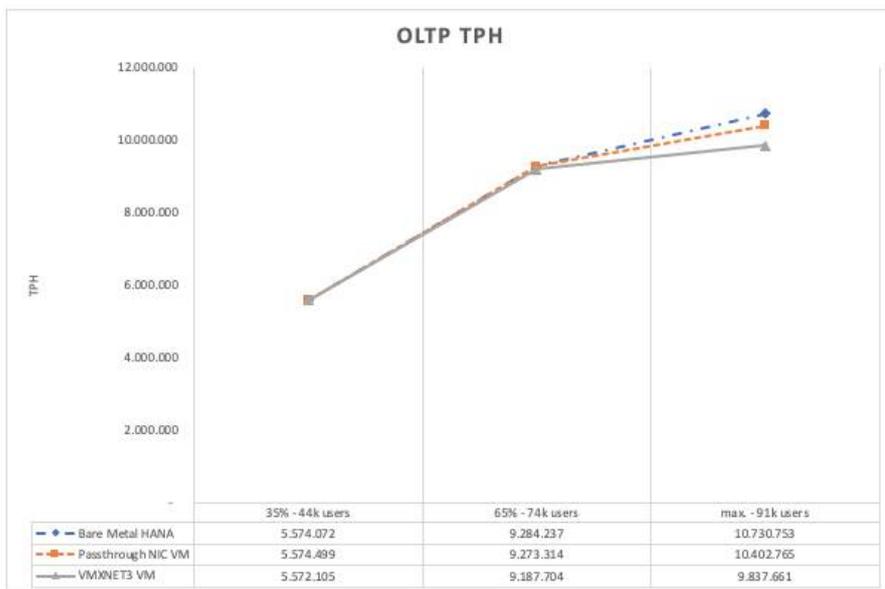


Figure 12: Mixed Workload OLTP Database TPH

At a CPU utilization between 35 percent and 65 percent, the impact on OLTP throughput is below 8% and therefore most customers shouldn't notice a high impact on database request time by using a virtualized VMXNET3 network adapter. At the so-called max-out point it was in this test -8% with VMXNET3. Note that typical SAP workload sizing's are up to the 65% CPU utilization.

Database request time sensitive customers who want to lower the network latency between the SAP app server tier and the database instance may want to consider PT NICs instead of VMXNET3 NICs at the cost of vMotion, which is not possible with PT adapters.

The next figure shows the mixed workload OLAP QPH results of a Cooper Lake CPU based system and using VMXNET3 shows very little impact on the QPH results between 35 and 65 percent (up to 1 percent) and approximately -8 percent at the max. out

measurement point at over 80 percent CPU utilization.

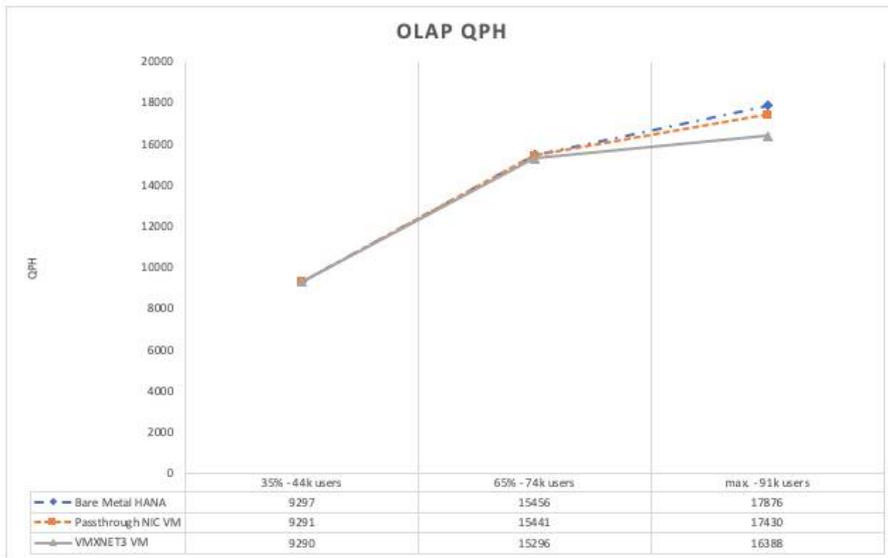


Figure 13: Mixed Workload OLAP Database QPH

SAP S/4 and B/4 HANA Validation Tests Findings (OLAP)

Business warehouse/OLAP-type workloads do not show the same impact on database request time or QPH results as OLTP- type workloads.

The explanation is simple: OLTP-type workloads tend to be more frequent and generate more and shorter network traffic/packages. OLAP-type workloads are typically less frequent and generate long running queries and are therefore less impacted by the discussed latency overhead caused by VMXNET3.

Table 20 shows public, certified results of BWH tests we have conducted with Intel Cascade Lake based server systems. The results show a BWH configuration based on an L-class sizing, which SAP has defined for 8-socket servers with 6 TB of memory.

The tests we have done are compared to a BWH system natively installed on the same hardware configuration that got later used for the virtual SAP HANA 8-socket wide 6 TB or 12 TB VM. The overhead of virtualization using a VMXNET3 NIC is very little and within 10%. Phase 3, the total runtime of complex query phase measured in seconds did not show any deviation.

Table 20 shows the results of an 8-socket BWH configuration as defined by SAP as a standard L-class configuration. The test was performed with Intel 8280L Cascade Lake 2.7GHz, 28-core CPUs with 6 TB installed and 9 data sets (11.7 billion records).

Table 20: BWH 6 TB L-class Workload Test: Bare Metal Compared to VM with VMXNET3

SAP BWH L-Class KPIs					25,000		5,000		200	
Configuration	Cert	CPUs Threads	MEM	Records	BWH Phase 1 (sec.)	Delta	BWH Phase 2 (QPH)	Delta	BWH Phase 3 (sec.)	Delta
Bare metal CLX 8S host	2020021	448	6,144GB	11.7 billion	19,551	-	5,838	-	146	-
VM with VMXNET3	2020031	448	5,890GB	11.7 billion	20,125	2.94%	5,314	-8.98%	146	0%
					- is better		+ is better		- is better	

In addition to this we did some tests with VMXNET3 and less exposed vCPUs to ensure we use the same number of vCPUs as previously used during the mixed workload test. Lowering the exposed vCPUs from 448 to 416 had a -5 percent impact on the overall QPH result

Using the same hosts with 12 TB allowed us to verify if this environment would meet the SAP-specified KPIs for an M-class BWH configuration and higher deviations. In the VMXNET3 test case, we again lowered the number of vCPUs to detect deviations in the different phases. While phase 2 shows a lower QPH result, the data load phase 1 shows a positive impact when compared to a 448 vCPU configuration. Freeing up some CPU cycles for extreme IP Load like the initial data load helped the ESXi kernel to load data faster as when 448 vCPUs get exposed for the VM. In addition to VMXNET3 we used a PT NIC to measure the impact of a PT NIC when running a BWH benchmark.

Table 21: BWH 12 TB M-Class Workload Test: VM with PT NIC vs. VM with VMXNET3

SAP BWH -Class KPIs					35,000	2,500	300			
Configuration	Cert	CPUs Threads	MEM	Records	BWH Phase 1 (sec.)	Delta	BWH Phase 2 (QPH)	Delta	BWH Phase 3 (sec.)	Delta
VM with PT	internal test	448	11,540	20.8 billion	30,342	-	3,446	-	170	-
VM with VMXNET3	internal test	416	11,540	20.8 billion	28,867	-	3,215	-6.70%	172	1.18%
		-7.14%			- is better		+ is better		- is better	

Our partner HPE recently performed an SAP BW benchmark with SAP HANA 2.0 SP6 and 6 TB and 12 TB large database instances storing 10,400,000,000 (BWH L-Class sizing) and 20,800,000,000 (BWH M-Class sizing) initial records with the highest ever measured results in a VMware virtualized environment with 7,600(cert. 2022014) and respectively 4,676 (cert. 2022015) query executions per hour (QPH). Figure 14 and table 22 show the 12 TB BM vs. VM benchmark. While the virtual benchmark is not passing the L-Class BW sizing mark, we are still within 10%, when compared to a previously published bare metal BWH 12 TB benchmark, running on the same hardware configuration, with SAP HANA 2.0.

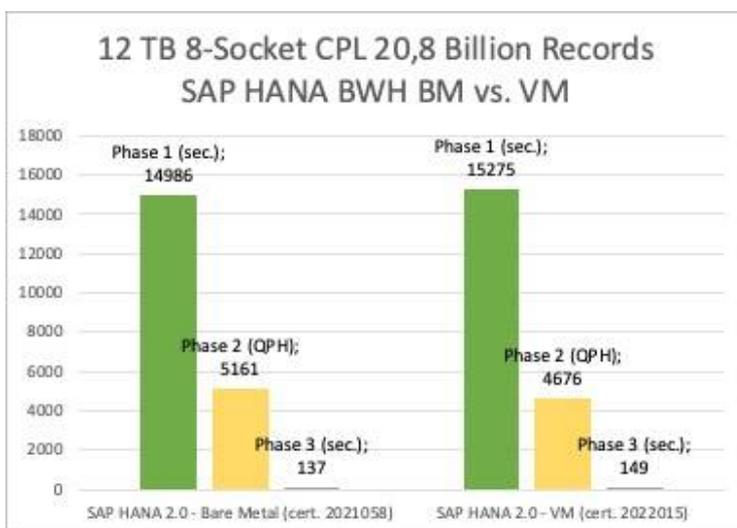


Figure 14: 12 TB 8-socket Cooper Lake BM vs. vSphere SAP HANA BWH benchmark

Table 22: BWH 12 TB M-Class Workload Test: BM vs. VM

SAP BWH M-Class KPIs					35,000		2,500		300	
Configuration	Cert	CPUs Threads	MEM	Records	BWH Phase 1 (sec.)	Delta	BWH Phase 2 (QPH)	Delta	BWH Phase 3 (sec.)	Delta
BM SAP HANA 2.0	2021058	448	12,288	20.8 billion	14,986	-	5,161	-	137	-
VM SAP HANA 2.0	2022015	448	11,776	20.8 billion	15,275	-1.93%	4,676	-9,40%	149	8,76%
		-7.14%			- is better		+ is better		- is better	

When compared previously shown old Cascade Lake CPU based 8-socket wide 6 TB SAP HANA 2.0 BWH benchmark, where we already achieved the SAP BWH sizing L-Category, we see an increase of 34% QPH in benchmark phase 2, by an 18% faster runtime in phase 3. The data load phase 1* is not comparable due to different used storage configurations. To note, also the old BM CLX benchmark is slower as our new CPL based VM benchmark. See figure 15 and table 23 for details. This increase of performance is mainly due to the newer used Intel CPU generation, but also by using a newer SAP HANA 2.0 release.

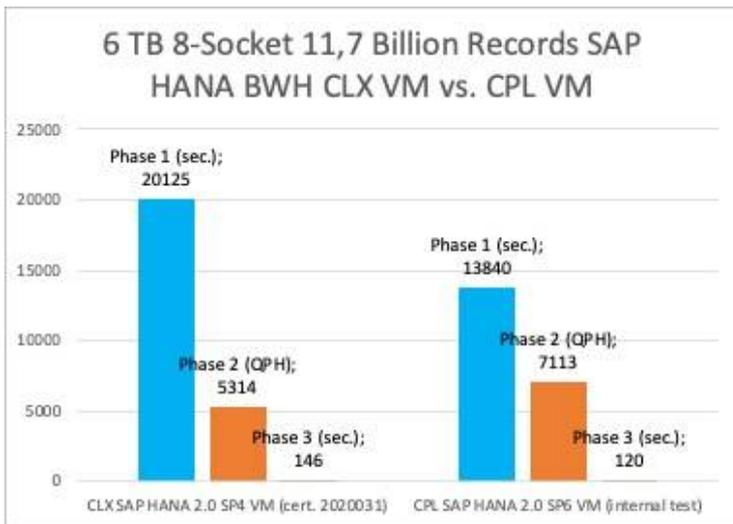


Figure 15: 6 TB 8-socket Cascade Lake VM vs. Cooper Lake VM BWH Benchmark

Table 23: BWH 6 TB L-Class Workload Test: CLX vs. CPL

SAP BWH L-Class KPIs					25,000		5,000		200	
Configuration	Cert	CPUs Threads	MEM	Records	BWH Phase 1 (sec.)	Delta	BWH Phase 2 (QPH)	Delta	BWH Phase 3 (sec.)	Delta
CLX SAP HANA 2.0 VM	2020031	448	5,890GB	11.7 billion	20,125	-	5,314	-	146	-
CPL SAP HANA 2.0 VM	internal test	448	5,720GB	11.7 billion	13,840	NC	7,113	33.85%	120	-17.81%
					- is better		+ is better		- is better	

Test summary: OLTP-type and OLAP-type Workload Tests

For OLAP-type workloads, a PT NIC has very little positive impact on the complex query runtime. Lowering the number of vCPUs inside a VM to reserve some CPU cores for the ESXi kernel helps in some situations, such as during the data load phase, but costs some performance (less QPH) available for SAP HANA.

Our recommendation for OLAP-type workloads is to use VMXNET3.

OLTP-type workloads benefit most from PT NICs or by lowering the number of vCPUs when VMXNET3 gets used. The OLTP tests as summarized in figures 10 till 13 show that, for most customers, this may not be relevant because until 44,000 users, we see very little difference between the database request time and throughput regardless of the chosen configuration.

Our recommendation is that customers start with VMXNET3 and, in the case that database request time is longer as expected, then the physical network infrastructure should get checked and, if possible, optimized for low latency before using PT NICs, which will help to achieve nearly BM NIC latencies. Again, optimizing the SAP HANA database network for low latency and throughput will have the best positive impact on the overall SAP HANA performance.

Using the latest CPU generation with vSphere 7.0 U3c and later provides significant performance gains when compared to the previous CPU generations running the same SAP HANA workload. This allows not only faster processing, but also higher memory configurations.

Note: In the case of high database request times, customers may want to consider optimizing the physical SAP network. Begin at the user to the SAP app server tier and then the app server to the SAP HANA database. The usage of low-latency switches, a flat network architecture or newer NICs will help to reduce the transaction request time experienced by users. Using PT NICs inside the HANA VM will impact only the database request time to the app servers at the cost of losing vMotion capabilities.

Enhanced vMotion Compatibility, vSphere vMotion, and vSphere DRS Best Practices

One of the key benefits of virtualization is the hardware abstraction and independence of a VM from the underlying hardware.

[Enhanced vMotion Compatibility](#), [vSphere vMotion](#), and vSphere Distributed Resource Scheduler™ (DRS) are key enabling technologies for creating a dynamic, automated, self-optimizing data center.

This allows a consistent operation and [migration](#) of applications running in a VM between different server systems without the need to change the OS and application, or to perform a lengthy restore process of a backup, which in the case of a hardware change would also need an update of the device drivers in the OS backup.

vSphere vMotion live migration (Figure 16) allows you to move an entire running virtual machine from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity. The virtual machine retains its network identity and connections, ensuring a seamless migration process. Transfer the virtual machine's active memory and precise execution state over a high-speed network, allowing the VM to switch from running on the source vSphere host to the destination vSphere host.



Figure 16: vSphere vMotion/live Migration with SAP HANA VMs

Enhanced vMotion Compatibility mode allows migration of virtual SAP HANA machines between hosts with different generations of CPUs, making it possible to aggregate older and newer server hardware generations in a single cluster. This flexibility provides scalability of the virtual infrastructure by offering the ability of adding new hardware into an existing infrastructure while extending the value of existing hosts.

Note: While Enhanced vMotion Compatibility ensures compatibility of hosts with different CPU generations, it is recommended to migrate SAP HANA between hosts with the same CPU type, model, and frequency to ensure a performant operation of SAP HANA on a vSphere cluster. The migration of an SAP HANA VM between hosts with different CPUs should be limited to situations such as

hardware upgrades or HA situations.

Moving a VM from one host to another can be done in different modes. Some modes are supported while an SAP HANA VM is powered on. Some, such as migration to another storage system, have limited support while the VM is powered on. The next section provides an overview of the different **VM migration options** and what to consider when migrating SAP HANA VMs.

Best practice for a vMotion migration of SAP HANA for production support

vMotion between different hardware generations of a CPU type or storage subsystems is possible. In the context of a performance-critical application such as SAP HANA, it is important to adhere to the following best practices when migration SAP VMs:

- You should run vMotion on SAP HANA VMs within the vSphere cluster on identical hardware (with the same CPU clock speed and synchronized TSC). This ensures that SAP HANA has the same CPU features and clock speed available.
- Do not run a live/online vMotion migration with HANA VMs while a virus scanner or a backup job is running inside the VM or while HANA gets used by users. Doing this may cause a HANA soft lock.
- VMware suggests using vMotion only during non-peak times (low CPU utilization, e.g., < 25 percent).
- You may use vMotion during a hardware refresh (non-identical source and destination host/clock speed) but should plan for VM hardware upgrade and alignment (hardware version and alignment of vCPUs to new CPU) to do a restart of the VM afterward to adapt any new CPU features of the target host.
- You may use vSphere Storage vMotion to migrate SAP HANA VMs between storage subsystems. vSphere Storage vMotion impacts the performance of an online VM. It is therefore strongly recommended to perform a storage migration while the VM is powered off or at least while the HANA database is shut down inside the VM.
- Have sufficient bandwidth allocated to the vMotion network, ideally 25 GbE or more.
- Avoid noisy neighbors during a vMotion migration.
- Check HANA patch levels (some patch levels may increase risk of OS soft lockups during migrations).
- Upgrade to vSphere 7 to leverage vMotion improvements.

Why is it critical to follow these best practices? vMotion is a great tool to ease management and operation of any SAP VMs. But if wrongly used, a VM migration activity may cause severe performance issues and may impact SAP HANA users and long running transactions, which you want to avoid.

VM migration scenarios

A vMotion migration can be done manually, fully, or semi-automated with vSphere DRS. To lower the possible impact on SAP HANA VMs during a VM migration, VMware suggests using vMotion only during non-peak times (low CPU utilization, e.g., 25 percent) and with vSphere DRS rules in place that only suggest initial placement or to allow the automated evacuation of VMs when a host gets set to maintenance mode. As previously mentioned, a dedicated vMotion network is a strict requirement, and the network should have enough bandwidth to support a fast migration time, which depends on the active SAP HANA memory (for example, for >= 4GB HANA VMs, a vMotion network with 25 GbE bandwidth is preferred).

The scenarios shown in the following figures are all supported. We will discuss per scenario what you should do to avoid possible performance issues.

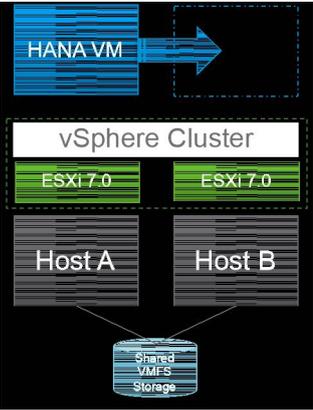
	<p>Description: Considerations:</p>
---	---

Figure 17: VMware VM Migration - Default Scenario (Live Migration)

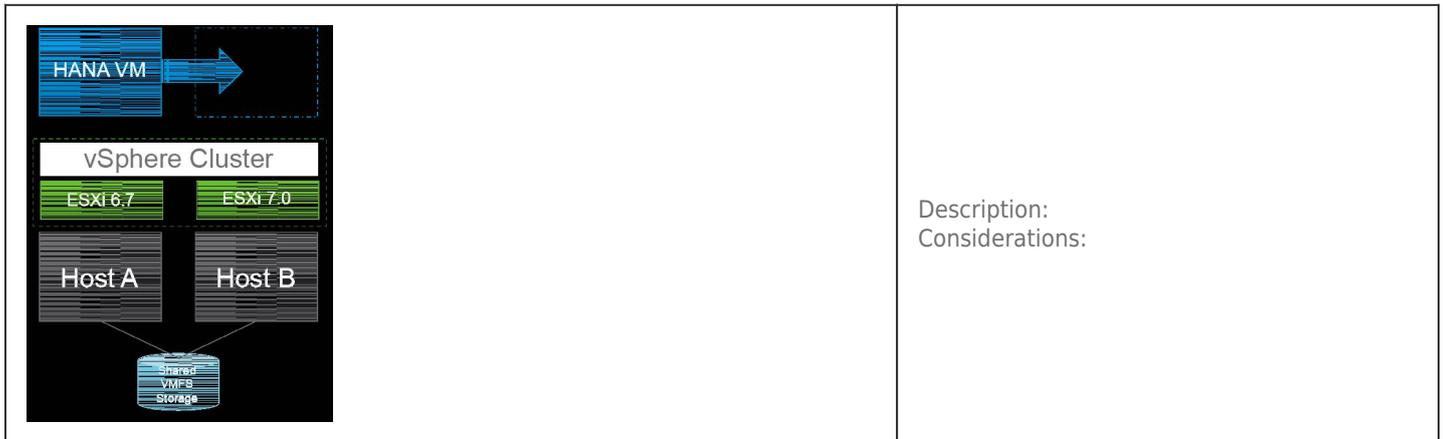


Figure 18: VMware VM Migration - ESXi Upgrade (VM Evacuation)



Figure 19: VMware VM Migration - Host HW Upgrade (VM Evacuation)

*This is a critical vMotion scenario due to the different possible CPU clock speeds, plus the exposed TSC to the VM may cause timing errors. To eliminate these possible errors and issues caused by different TSCs, vSphere will perform the necessary rate transformation. This may degrade the performance of RDTSC relative to native. Background: When a virtual machine is powered on, its TSC inside the guest, by default, runs at the same rate as the host. If the virtual machine is then moved to a different host without being powered off (for example, by using VMware vSphere vMotion), a rate transformation is performed so the virtual TSC continues to run at its original power-on rate, not at the host TSC rate on the new host machine. For details, read the document [“Timekeeping in VMware Virtual Machines.”](#)

To solve this issue, you must plan a maintenance window to be able to restart the VMs that were moved to the non-identical HW to allow the use of HW-based TSC instead of using software rate transformation on the target host, which is expensive and will degrade VM performance. Figure 20 shows the process to enable the most flexibility in terms of operation and maintenance by restarting the VM after the upgrade, ensuring the best possible performance.

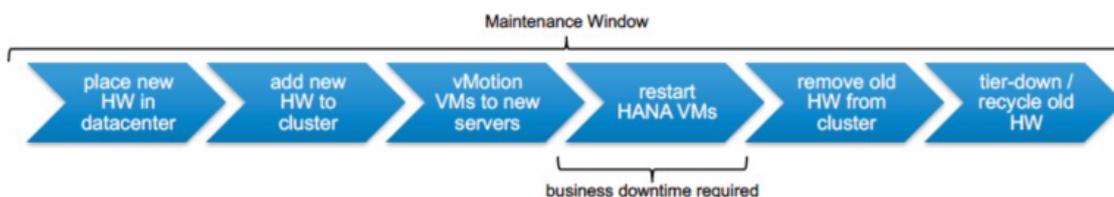


Figure 20: The Maintenance Window as a Hardware Upgrade and vMotion



Figure 21: VMware VM Migration - Storage Migration (VM datastore migration)

Managing SAP HANA Landscapes using vSphere DRS

SAP HANA landscapes can be managed using vSphere DRS, which is an automated load balancing technology that aligns resource usage with business priority. vSphere DRS dynamically aligns resources with business priorities, balances computing capacity, and reduces power consumption in the data center.

vSphere DRS takes advantage of vMotion to migrate virtual machines among a set of ESX hosts. vSphere DRS continuously monitors utilization across ESXi hosts and would be able to migrate VMs to hosts that are less utilized if a VM is not no longer in a “happy” resource state.

When deploying large SAP HANA databases (>128GB VM sizes) or production-level SAP HANA VMs, it is essential to have vSphere DRS rules in place and to set the automation mode to manual or when set to automated then set the DRS migration threshold to conservative (level 1), to avoid unwanted migrations, which may cause issues when done during peak usage times, or while a backup or virus scanners run. Migrations executed during such high usage times may negatively impact the performance of an SAP HANA system. It is possible to define which SAP HANA VMs should get excluded from automated DRS initiated migrations and, if at all, which SAP HANA VMs are targets of automated DRS migrations.

DRS can be set to these automation modes:

- **Manual** - In this mode, DRS recommends the initial placement of a virtual machine within the cluster, and then recommends the migration. The actual migration needs to be executed by the operator.
- **Semi-automated** - In this mode, DRS automates the placement of virtual machines and then recommends the migration of virtual machines.
- **Fully automated** - In this mode, DRS placements and migrations are automatic.

When DRS is configured for manual control, it makes recommendations for review and later implementation only (there is no automated activity).

Note: DRS requires the installation of a vSphere Clustering Service VM and will automatically install such a VM in a vSphere cluster.

vSphere Clustering Service

Starting with [vSphere 7.0 Update 1](#), the vSphere Clustering Service is enabled by default and runs in all vSphere clusters. VMware would like to make critical cluster services, such as vSphere HA and vSphere DRS, always available, and vSphere Clustering Service is an initiative to reach that vision.

SAP HANA as the foundation of most SAP business applications is a very critical asset of all companies using SAP solutions for their business. Due to the criticality of these applications for a business, it is important to protect and optimally operate SAP HANA.

Running SAP HANA on vSphere provides an easy way to protect and operate it by leveraging vSphere cluster services, which depend on VMware vCenter Server® availability for configuration and operation.

The dependency of these cluster services on vCenter is not ideal, and the vSphere Clustering Service 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 the future, the vSphere Clustering Service will ensure that the cluster services remain available to maintain the resources and health of the workloads that run in the clusters.

Note: vSphere Clustering Service is enabled when you upgrade to vSphere 7.0 Update 1 or when you have a new vSphere 7.0 Update 1 deployment. vSphere Clustering Service VMs are automatically deployed as part of a vCenter Server upgrade, regardless of which ESXi version gets used.

Architecture

vSphere Clustering Service uses agent virtual machines to maintain cluster services health. Up to three vSphere Clustering Service agent virtual machines are created when you add hosts to clusters. These vSphere Clustering Service VMs, which build the cluster control plane, are lightweight agent VMs.

vSphere Clustering Service VMs are required to run in each vSphere cluster, distributed within a cluster. vSphere Clustering Service is also enabled on clusters that contain only one or two hosts. In these clusters, the number of vSphere Clustering Service VMs is one and two, respectively. Figure 22 shows the high-level architecture with the new cluster control plane.

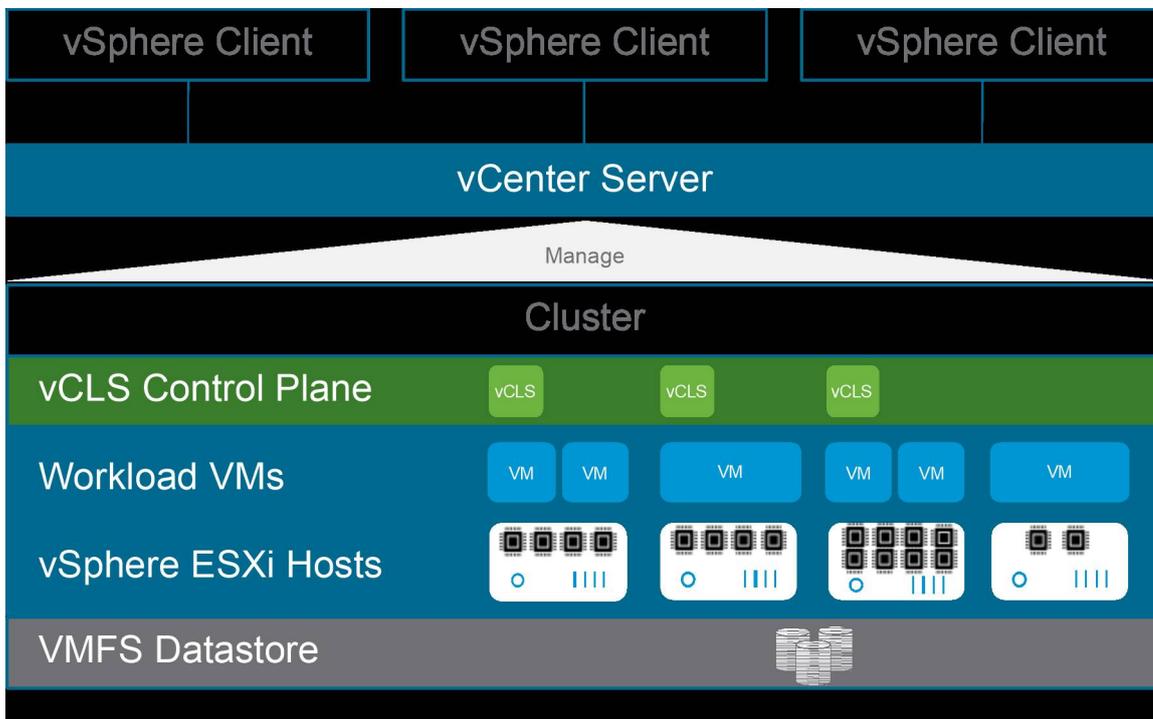


Figure 22: vSphere Clustering Service High-level Architecture

A cluster enabled with vSphere Clustering Service can contain ESXi hosts of different versions if the ESXi versions are compatible with vCenter Server 7.0 Update 1. vSphere Clustering Service works with both vSphere Lifecycle Manager™ and vSphere Update Manager™ managed clusters and runs in all vSphere license SKU clusters.

vSphere Clustering Service (vCLS) VM details

vSphere Clustering Service VMs run in every cluster, even if cluster services such as vSphere DRS or vSphere HA are not enabled on the cluster.

Each vSphere Clustering Service VM has 100 MHz and 100MB capacity reserved in the cluster. For more details, view the [Monitoring vSphere Clustering Services](#) documentation.

In the normal use case, these VMs are nearly not noticeable in terms of resource consumption. Users are not expected to maintain the lifecycle or state for the agent VMs; they should not be treated like typical workload VMs.

Table 24: vSphere Clustering Service VM Resource Allocation

Property	Size
Memory	128MB
CPU	1 vCPU
Hard disk	2 GB

Table 25: Number of vSphere Clustering Service agent VMs in Clusters

Number of hosts in a cluster	Number of vSphere Clustering Service agent VMs
1	1
2	2
3 or more	3

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

Because of these guidelines and due to the mandatory and automated installation process of vSphere Clustering Service VMs, when upgrading to vCenter 7.0 Update 1, it is necessary to check if vSphere Clustering Service VMs got co-deployed on 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 easily achieved by configuring [vSphere Clustering Service VM anti-affinity policies](#). These policies describe a relationship between VMs that have been assigned a special anti-affinity tag (e.g., a tag named SAP HANA) and vSphere Clustering Service system VMs.

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

After the policy is created and tags are assigned, the placement engine attempts to place vSphere Clustering Service VMs on the hosts where tagged VMs are not running, e.g., the HA ESXi host.

Note: Setting vSphere Clustering Service VM anti-affinity policies ensures that a vSphere Clustering Service VM does not get placed on hosts that run SAP HANA VMs. This requires hosts that do not run SAP HANA tagged VMs.

vCLS deployment examples for SAP HANA landscapes

Typically, customers deploy SAP HANA on dedicated ESXi hosts. These hosts can be part of small or large clusters (in terms of number of hosts). They can be mixed with hosts running non-SAP HANA workload VMs or can be part of a dedicated SAP HANA only cluster.

The following examples of typical SAP landscape clusters provide some guidelines on where to place up to three lightweight vSphere Clustering Service VMs.

Mixed SAP HANA and non-SAP HANA VM on vSphere cluster

A mixed cluster should be the typical scenario for most customers. In this case, check the vSphere Clustering Service VMs if they

got deployed on ESXi hosts that run production-level SAP HANA VMs. If yes, then the vSphere Clustering Service VM may run on the same CPU socket as an SAP HANA VM.

To avoid this, [configure vSphere Clustering Service anti-affinity policies](#):

Procedure:

1. Create a category and tag for each group of VMs that you want to include in a vSphere Clustering Service VM anti-affinity policy.
2. Tag the VMs that you want to include.
3. Create a vSphere Clustering Service VM anti-affinity policy.
 - a. From vSphere, click Policies and Profiles > Compute Policies.
 - b. Click Add to open the New Compute Policy wizard.
 - c. Fill in the policy name and choose vCLS VM anti affinity from the Policy type drop-down control. The policy name must be unique.
 - d. Provide a description of the policy, then use VM tag to choose the category and tag to which the policy applies. Unless you have multiple VM tags associated with a category, the wizard fills in the VM tag after you select the tag category.
 - e. Click Create to create the policy.

Figure 23 shows the initial deployed vSphere Clustering Service VMs and how these VMs get automatically migrated (green arrows) when the anti-affinity rules are activated to comply with SAP notes [2937606](#) and [3102813](#). Not shown in this figure are the HA host/HA capacity reserved for HA failover situations.

Note: In the case you have to add a new hosts to an existing SAP-only cluster to make it a mixed host cluster, ensure that you verify the prerequisites as outlined in the [Add a Host to a Cluster documentation](#).

Figure 23: vSphere Clustering Service Migration by Leveraging vSphere Clustering Service Anti-Affinity Policies (mixed host cluster)**Dedicated SAP HANA VM on vSphere cluster**

Customers may have deployed an SAP HANA cluster with dedicated hosts that run only SAP HANA workload VMs. In this case, automatically deployed vSphere Clustering Service VMs cannot get migrated easily to hosts that do not run SAP HANA VMs. The solution is to add existing hosts with non-SAP HANA workload VMs to this cluster, or to have non-tagged SAP HANA non-production VMs running on at least one host. These existing hosts may run any workload, such as SAP application server VM or infrastructure workload VMs. It is not required to buy a new host for this.

Figure 24 shows the initial deployed vSphere Clustering Service VMs and how these VMs get moved when the vSphere Clustering Service anti-affinity policy for SAP HANA gets executed to cluster-added hosts. Not shown in this figure are the HA host/HA

capacity reserved for HA failover situations.

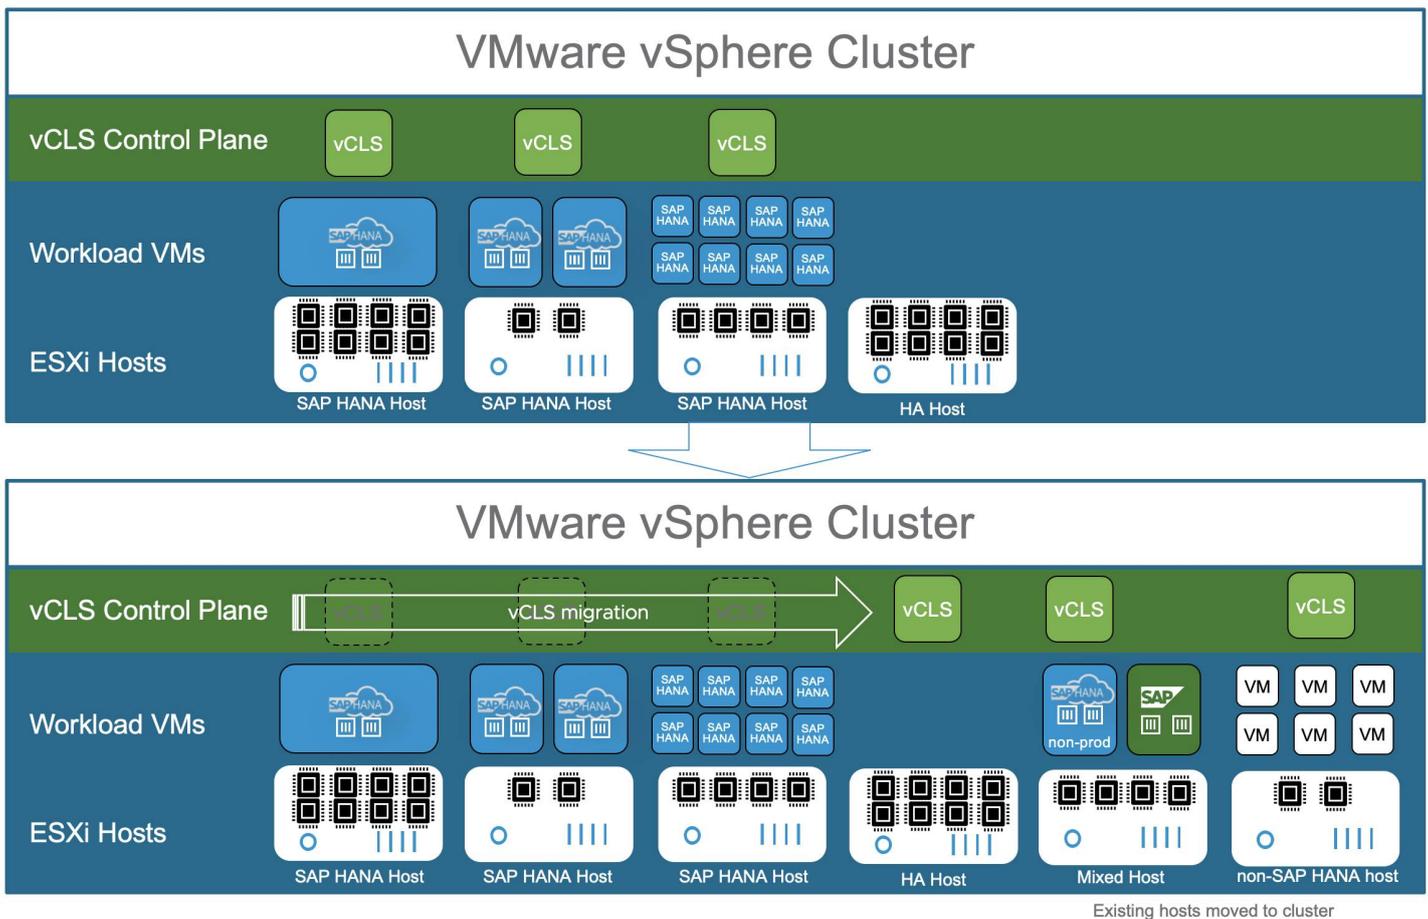


Figure 24: vSphere Clustering Service Migration by Leveraging vSphere Clustering Service Anti-Affinity Policies and Non-SAP HANA Hosts (dedicated SAP HANA host cluster)

Note: To allow the vSphere Clustering Service VM to run, as shown in Figure 21, on the same host as a non-production SAP HANA VM implies that you have not tagged the non-production SAP HANA VM with a name tag that triggers the anti-affinity policy.

SAP HANA HCI on vSphere cluster

Just as with the dedicated SAP HANA cluster, an SAP HANA HCI cluster may only run SAP HANA workload VMs. As with SAP HANA running on traditional storage, SAP HANA HCI (SAP note [2718982](#)) supports the co-deployment with non-SAP HANA VMs as outlined in SAP notes [2937606](#) (vSphere 7.0) and [2393917](#) (vSphere 6.5/6.7).

If vCenter gets upgraded to 7.0 U1, then the vSphere Clustering Service VMs will get automatically deployed on SAP HANA HCI nodes. If these nodes are exclusively used for SAP HANA production level VMs, then these vSphere Clustering Service VMs must get removed and migrated to the vSphere HCI HA host(s).

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, like the HCI vSphere HA host. 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.

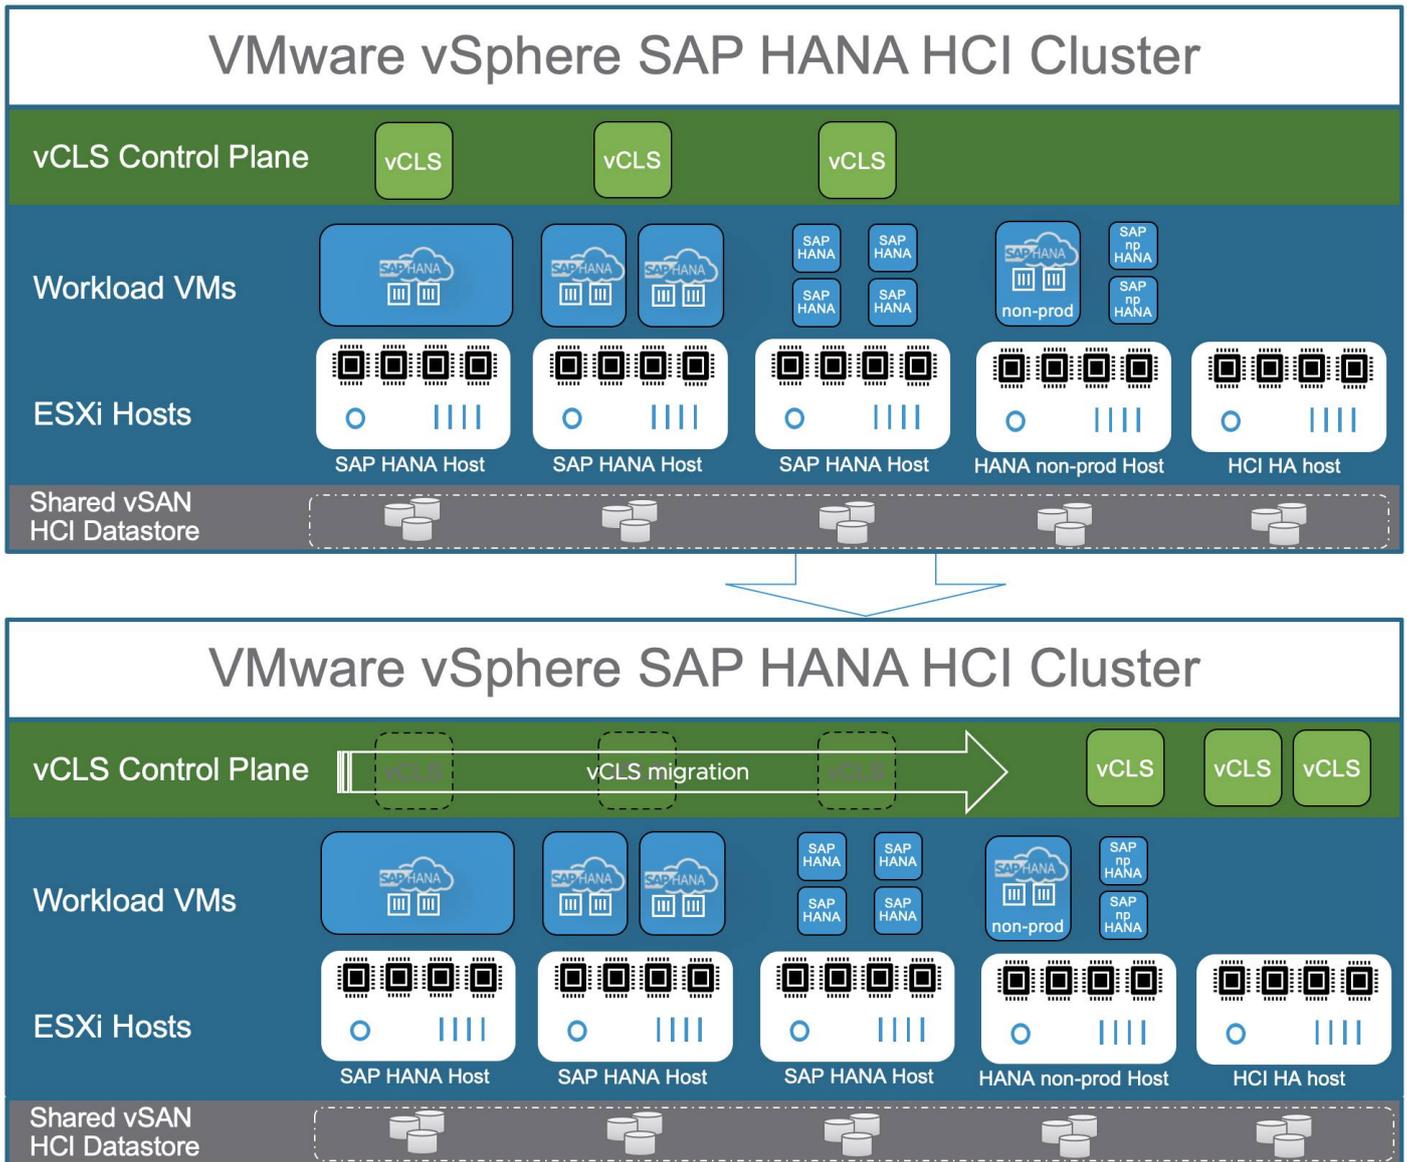


Figure 25: vSphere Clustering Service VM within an SAP HANA HCI on vSphere Cluster

Note: To allow the vSphere Clustering Service VM to run, as shown in Figures 24 and 25, on the same host as a non-production SAP HANA VM implies that you have not tagged the non-production SAP HANA VM with a name tag that triggers the anti-affinity policy.

In summary, by introducing the vSphere Clustering Service, VMware is embarking on a journey to remove the vCenter dependency and possible related issues when vCenter Server is not available and provides a scalable platform for larger vSphere host deployments.

For more information, see the following resources:

- [vSphere Clustering Services](#)
- [vSphere Clustering Service VM Anti-Affinity Policies](#)
- [Create or delete a vSphere Clustering Service VM Anti-Affinity Policy](#)

Virtualized SAP HANA high availability best practices

SAP HANA offers several methods for high availability and disaster recovery. These are auto-failover, service restart options, backups, system replication, and standby host systems. In VMware virtualized environments, all these solutions can be used. In addition, vSphere HA and vSphere Replication™ can be used to minimize unplanned downtime due to faults.

High availability support can be separated into two different areas: fault recovery and in disaster recovery.

High availability, by providing fault recovery, includes:

- SAP HANA service auto-restart
- Host auto-failover (standby host)
- vSphere HA
- SAP HANA system replication

High availability, by providing disaster recovery, includes:

- Backup and restore
- Storage replication
- vSphere Replication
- SAP HANA system replication

SAP HANA system replication is listed in both recovery scenarios. Depending on the customer requirements, SAP HANA system replication can get used as a failover solution, or as a disaster recovery solution when site or data recovery is needed. With HSR, it is also possible to lower the time needed to start a large SAP HANA database due to the possibility to preload data into the memory of the replication instance.

Different recovery point objectives (RPOs) and recovery time objectives (RTOs) can be assigned to different fault recovery and disaster recovery solutions. SAP describes the phases of high availability in their HANA HA document. Figure 26 shows a graphical view of these phases.

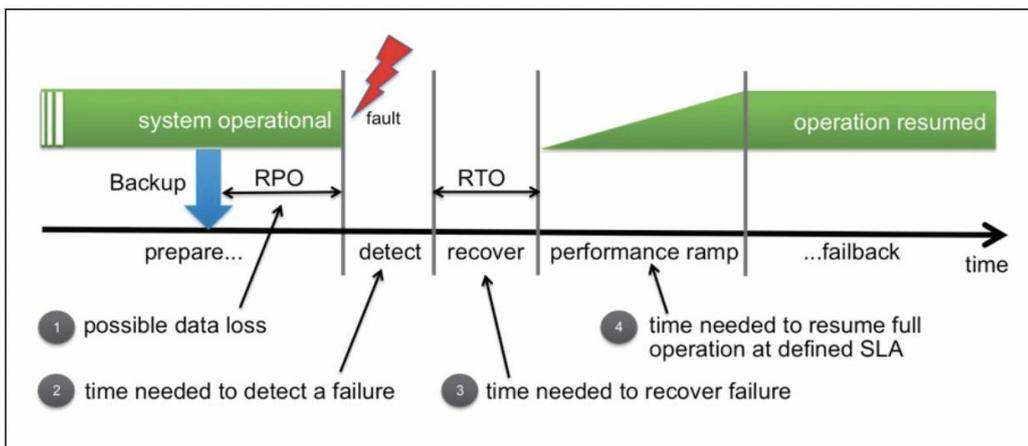


Figure 26: SAP HANA System Availability Phases

- RPO (1 in Figure 26) specifies the amount of possible data that can be lost due to a failure. It is the time between the last valid backup and/or last available SAP HANA save point, and/or the last saved transaction log file that is available for recovery and the point in time of the error situation. All changes made within this time may be lost and are not recoverable.
- Mark (2 in Figure 26) shows the time needed to detect a failure and to start the recovery steps. This is usually done in seconds for SAP HANA. vSphere HA tries to automate the detection of a wide range of error situations, thus minimizing the detection time.
- RTO (3 in Figure 26) is the time needed to recover from a fault. Depending on the failure, this may require restoring a backup or a simple restart of the SAP HANA processes.
- Ramp-up time (4 in Figure 26) shows the performance ramp, which describes the time needed for a system to run at the same service level as before the fault (data consistency and performance).

Based on this information, the proper HA/recovery solution can get planned and implemented to meet the customer specific RPOs and RTOs.

Minimizing RTOs and RPOs with the available IT budget and resources should be the goal and is the responsibility of the IT team

operating SAP HANA. VMware virtualized SAP HANA systems allow this by highly standardizing and automating the failure detection and recovery process.

vSphere HA

VMware provides vSphere built-in and optional availability and disaster recovery solutions to protect a virtualized SAP HANA system at the hardware and OS levels. Many of the key features of virtualization, such as encapsulation and hardware independence, already offer inherent protections. In addition, vSphere can provide fault tolerance by supporting redundant components, such as dual network and storage pathing, or the support of hardware solutions, such as UPS, or the support of CPU built-in features that allow to tolerate failures in memory models or that ensure CPU transaction consistency.

All these features are available on the vSphere host level with no need to configure on the VM or application level. Additional protections, such as vSphere HA, are provided to ensure organizations can meet their RPOs and RTOs.

Figure 27 shows different HA solutions to protect against component-level failures, up to a complete site failure, which can get managed and automated with VMware Site Recovery Manager™. These features protect any application running inside a VM against hardware failures, allow planned maintenance with zero downtime, and protect against unplanned downtime and disasters.

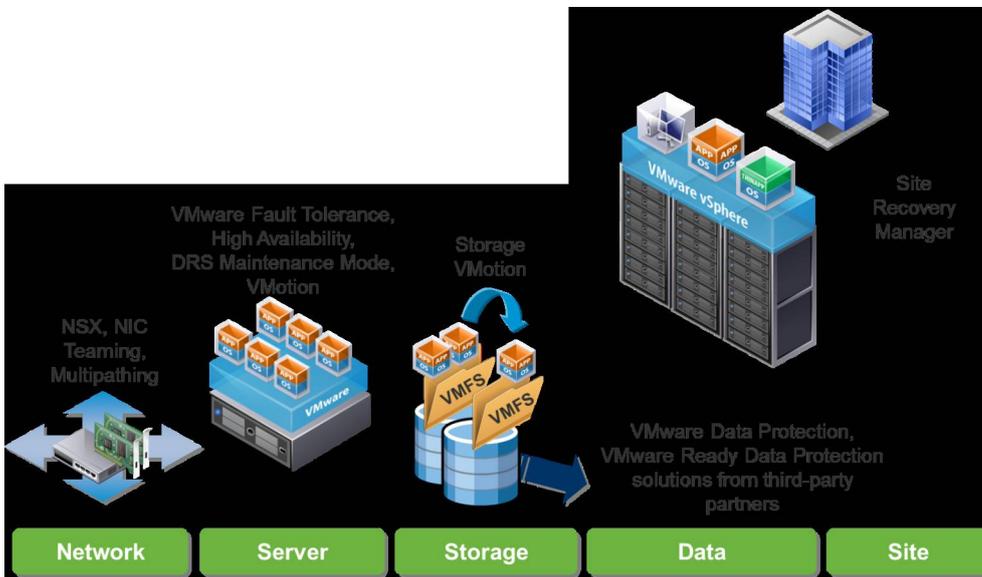


Figure 27: VMware HA and DR Solutions Provide Protection at Every Level

[vSphere HA](#) is, as already specified, a fault recovery solution and provides uniform, cost-effective failover protection against hardware and OS outages within a virtualized IT environment. It does this by monitoring vSphere hosts and virtual machines to detect hardware and guest OS failures. It restarts virtual machines on other vSphere hosts in the cluster without manual intervention when a server outage is detected, and it reduces application downtime by automatically restarting virtual machines upon detection of an OS failure. This combined with the SAP HANA service auto-restart feature allows HA levels of 99.9 percent out of the box.[25]

Figure 25 shows how vSphere HA can protect against OS or host failures and the application protection solution, such as SAP HANA service auto-restart, third-party in-guest cluster solutions, or SAP HANA system replication to even provide DR capabilities. All these solutions can get combined with vSphere HA.



Figure 28: Virtualized SAP HANA HA Solution

vSphere 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.

Figure 29 shows the how vSphere HA can get leveraged and how a typical n+1 vSphere cluster can get configured to survive a complete host failure. This vSphere HA configuration is the standard HA configuration form of most SAP applications and SAP HANA instances on vSphere. If higher redundancy levels are required, then an n+2 configuration can get used. The HA resource pool can get leveraged by non-critical VMs, which need to get powered off before an SAP HANA or SAP app server VM gets restarted.

Note: The vSphere Clustering Service control plane and the related vSphere Clustering Service VMs are not shown in the following figures.

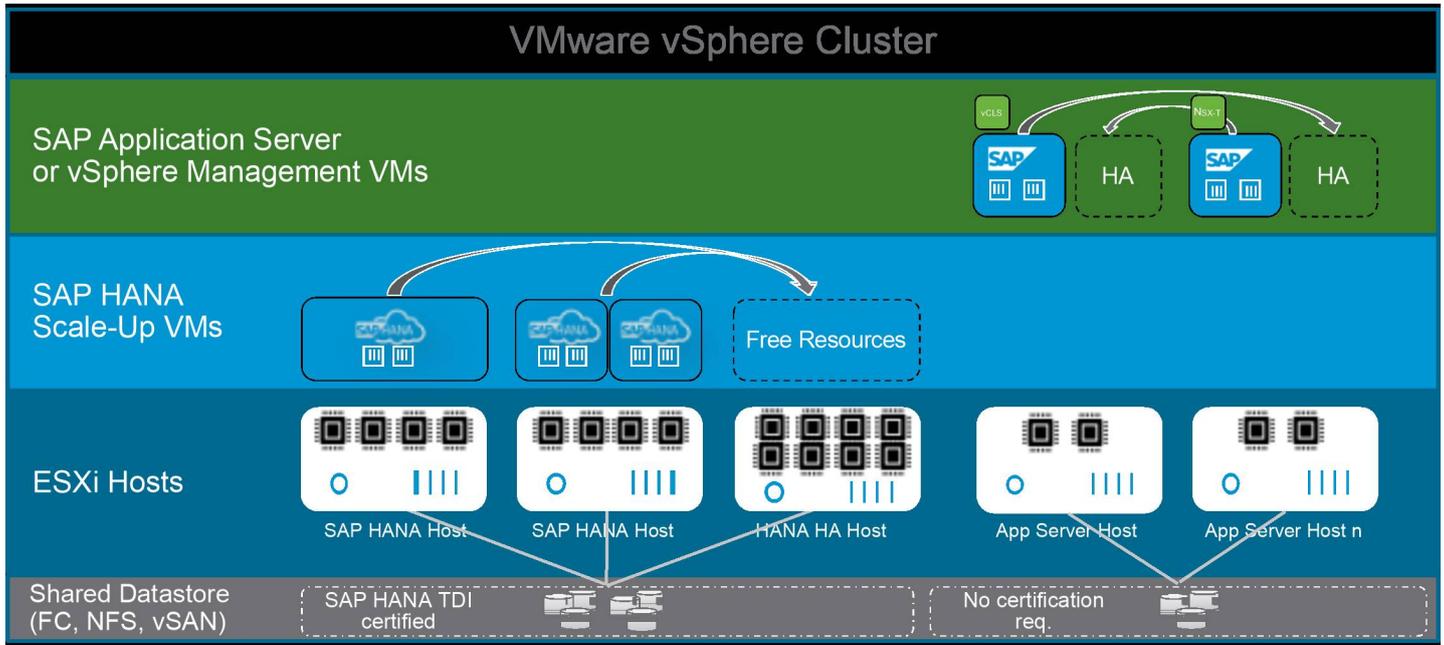


Figure 29: vSphere HA protected SAP HANA VMs in an n+1 Cluster Configuration

It is also possible to configure the HA cluster as an active-active cluster, where all hosts have SAP HANA VMs deployed. This ensures that all hosts of a vSphere cluster are used by still providing enough failover capacity for all running VMs in the case of a host failure. The arrow in the figure indicates that the VMs can failover to different hosts in the cluster. This active-active cluster configuration assumes that the capacity of one host (n+1) is always available to support a complete host failure.

As noted, vSphere HA can also get used to protect an SAP HANA scale-out deployment. Unlike with a physical scale-out deployment, no dedicated standby host and storage specific implementations are needed to protect SAP HANA against a host failure.

There are no dependencies on external components, such as DNS servers, SAP HANA Storage Connector API, or STONIT scripts. vSphere HA will simply restart the failed SAP HANA VM on the vSphere HA/standby server. The HANA shared directory is mounted via NFS inside the HANA VM, just as recommended with physical systems, and will fail over with the VM that has failed. The access to HANA shared is therefore guaranteed. If the NFS server providing this share is also virtualized, then vSphere Fault Tolerance (FT) could be used to protect this NFS server.

Figure 30 shows a configuration of three SAP HANA 4-socket wide VMs (one leader with two follower nodes) running exclusively on the host of a vSphere cluster based on 4-socket hosts. One host provides the needed failover capacity in the case of a host failure.

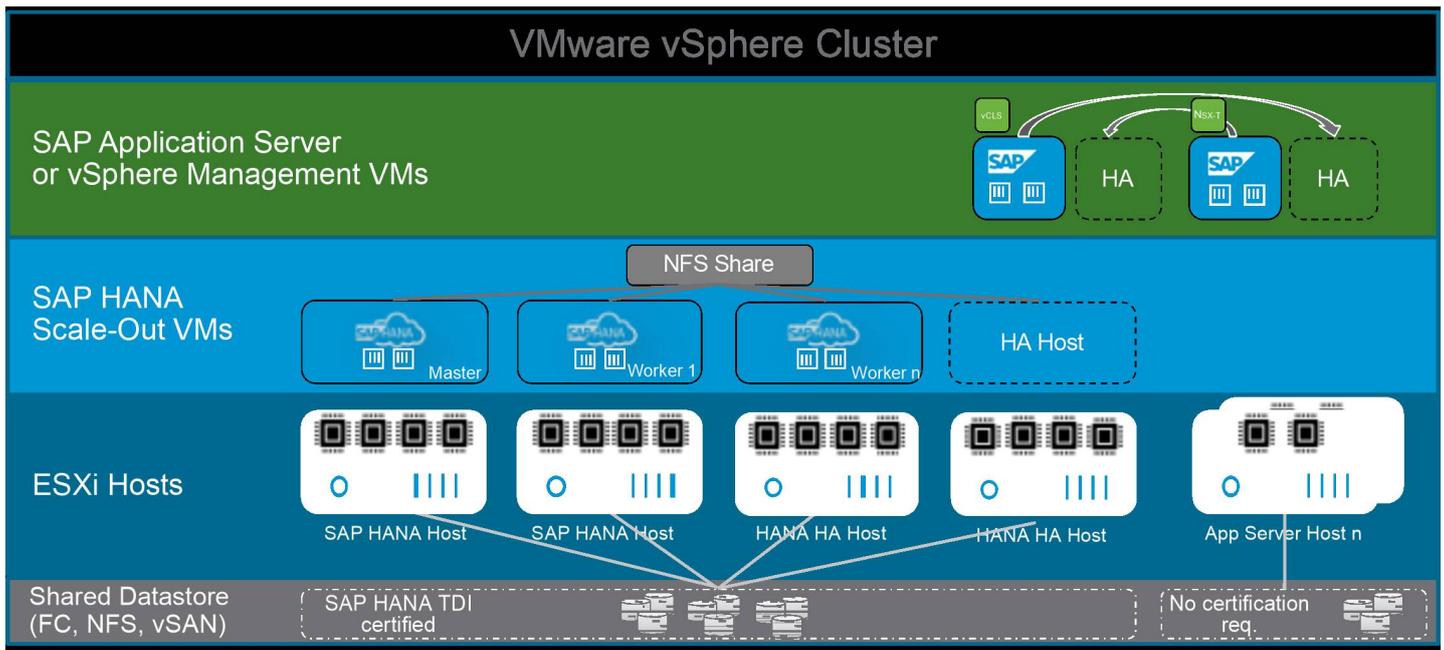


Figure 30: SAP HANA Scale-out on a vSphere HA n+1 Cluster Configuration

It is possible to use the HA node for other workloads while in normal operation. If the HA/standby node gets used for other workloads, then all potentially running VMs on this host must get terminated or migrated to another host before a failed HANA scale-out VM can get restarted on this host. In this case, the overall failover time could be a bit longer because vSphere HA will wait, if configured correctly, until all needed resources are available on the failover host.

Up to 16 scale-out nodes with up to 2TB on 4-socket VMs are GA as of today. Supported hosts would be 4- and 8-socket large host systems (Intel Broadwell CPU and newer). Review the SAP HANA on vSphere support notes for more details.

Note: vSphere HA can protect only against OS or VM crashes or hardware failures. It cannot protect against logical failures or OS file system corruptions that are not handled by the OS file system.

In physical SAP HANA deployments, SAP HANA system replication is the only method to provide fault recovery. If the recovery should get automated, then a third-party solution, such as SUSE HA, needs to get implemented. Protecting a physical SAP HANA deployment against host failures is therefore relatively complex, whereas protecting a VMware virtualized SAP HANA system is just the matter of a mouse click.

If only fast failure recovery is required, then it is recommended to use HSR. Because HSR is replicating SAP HANA data, it can be used for disaster recovery or for recovering from logical errors (depending on the log retention policy).

VMware HA with Passthrough (PT) Network Adapters

In the case that the VMXNET3 caused latency is too high for a specific use case / workload then VMware recommends to use a Passthrough (PT) NIC configuration.

To enable VMware HA with PT NICs the PT NIC must get configured as a Dynamic DirectPath I/O™ device with a unique cluster wide hardware label.

This can be done by following the instructions in the [Add a PCI Device to a Virtual Machine documentation](#). The same hardware label must get used for the PT NIC installed in the HA host. If no HA host with a PT NIC is configured as dynamic vSphere DirectPath I/O device and the same hardware label exists, then the HA failover process won't work.

Read this VMware article on [Assignable Hardware](#) for more information about this topic.

SAP HANA system replication with vSphere (local site)

To protect SAP HANA data against logical failures or disastrous failures that impact a data center, then vSphere HA can be combined with HSR.

vSphere HA would, in this case, protect against local failures, such as OS or local component failures, and HSR would protect the SAP HANA data against logical or data center impacting failures. HSR requires a running SAP HANA replication VM, which is required to receive HSR data. Alternatively, storage subsystem-based replication can get used, which would be independent from

SAP HANA.

Figure 33 shows a vSphere cluster with an SAP HANA production VM replicated to an SAP HANA replication VM. The HSR replica VM can be running on the same cluster or on another vSphere cluster/host to protect against data center impacting failures, as showed in Figure 30. If it runs in the same location, then HSR can be used to recover from logical failures (if logs get applied in a delayed manner) or to reduce the ramp-up time of an SAP HANA system because data can already be loaded into the memory of the replication server. HSR can change direction depending on which HANA instance is the production one.

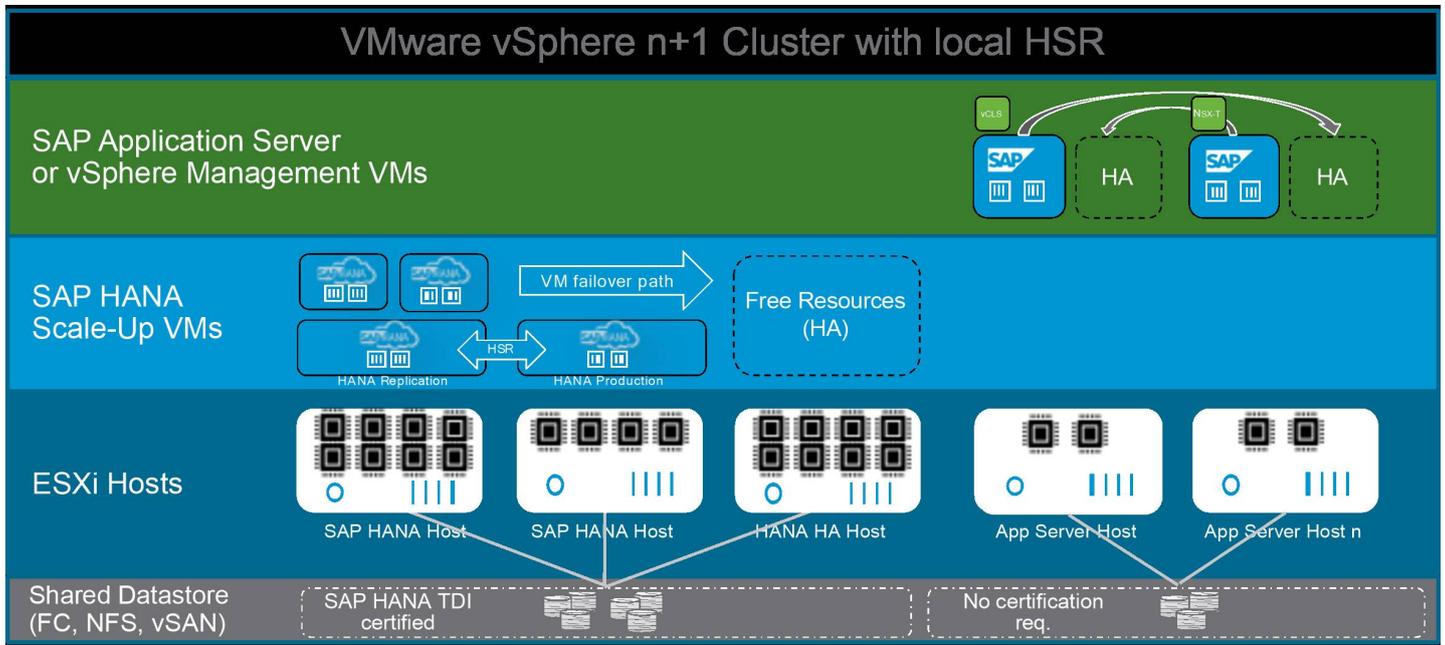


Figure 31: vSphere HA with Local Data Center HANA System Replication (HSR)

As noted, HSR does not provide an automated failover. Manual reconfiguration of the replication target VM to the production system’s identity is required. Alternatively, third-party cluster solutions, such as SLES HA or SAP Landscape Management, can be used to automate the failover to the SAP HANA replication target.

Note: SAP HSR can be combined with vSphere HA to protect the HSR source system against local failures.

To provide disaster tolerance, then it is required to place the HSR replica VM/host to another data center or even a geographically dispersed site.

Virtual SAP HANA Disaster Recovery

If data should also get protected or if in the case of a complete site failover of all IT systems, including if SAP HANA is required, then it is possible to combine the discussed HA solutions with storage/vSphere Replication and HSR to another data center/site.

Besides this, backup and restore solutions are also required to protect the data against logical errors or because of regulatory reasons.

SAP HSR to a remote site and vSphere Replication

Figure 34 shows an HSR protected SAP HANA instance. It is the same concept as discussed in Figure 33, with the only difference being that the HSR replication target is placed in another data center. This provides additional protection against data center failures or, if the remote data center is in another site, it protects against site failures. The vSphere host in DC-2 can be a standalone ESXi host or a member of a vSphere cluster. Stretched vSphere clusters are also possible and supported.

Depending on the replication requirements (synchronous or asynchronous), a roundtrip time (RTT) below 1ms may be required to be able to maintain the SAP HANA storage KPIs. If a 1ms RTT is not possible, then asynchronous replication should be used to ensure that the production SAP HANA system is not negatively impacted by the data replication.

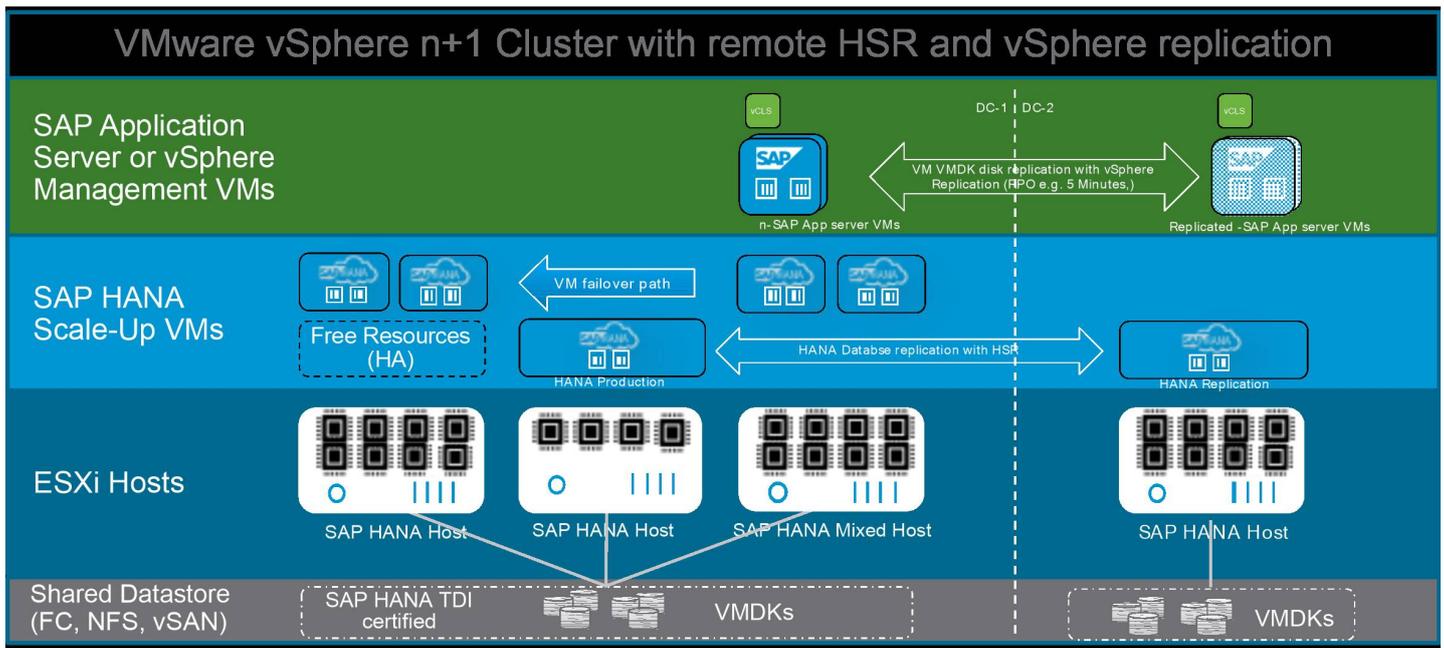


Figure 32: vSphere HA with Remote Data Center HANA System Replication

The example in Figure 34 shows the HSR replication from a virtualized SAP HANA system to another virtualized SAP HANA system. The SAP app. server VMs can get replicated to the DT side by leveraging [vSphere Replication](#). This will allow continuing the operation after a switch to this datacenter. These app. servers can also run on dedicated non-SAP HANA host systems.

vSphere Replication is a hypervisor-based, asynchronous replication solution for vSphere virtual machines (VMDK files). It allows RPO times from 5 minutes to 24 hours, and the virtual machine replication process is nonintrusive and takes place independent of the OS or applications in the virtual machine. It is transparent to protected virtual machines and requires no changes to their configuration or ongoing management.

Note: The SAP HANA performance is directly impacted by the RTT. If the RPO target is 0, then synchronous replication is required. In this case, the RTT needs to be below 1 ms. Otherwise, asynchronous replication should be used to avoid replication-related performance issues of the primary production instance. Also note that the HSR target can be a virtualized or natively installed SAP HANA replication target instance. vSphere replication is an asynchronous replication solution and should not get used if your RPO objectives are <5 minutes [26].

vSphere replication gets often used to protect non-HSR protected HANA or non-SAP HANA systems against local data center failures in combination with vSphere stretched cluster configurations over two separated data centers. If the systems should also get protected against datacenter site impacting disasters, then all IT operation relevant systems need to get replicated to a second site. This can get done as previously mentioned with vSphere Replication, native storage replication, and SAP HANA system replication.

vSphere Replication operates at the individual VMDK level, allowing replication of individual virtual machines between heterogeneous storage types supported by vSphere. Because vSphere Replication is independent of the underlying storage, it works with a variety of storage types, including vSAN, vSphere Virtual Volumes, traditional SAN, network-attached storage (NAS), and direct-attached storage (DAS).

Note: See the [vSphere Replication documentation](#) for details about supported configurations and specific requirements, such as network bandwidth.

In the case that an [SAP HANA HCI based on vSAN solution](#) gets used, data center distances of up to 5km are supported.

Backup and Restore VMware Virtualized SAP HANA Systems

Backing up and restoring an SAP HANA database and the Linux OS supporting SAP HANA is the same as when backing up bare-metal deployed SAP HANA systems.

The easiest way to perform a backup and later a recovery would be to perform a file system backup and the HANA database dump option, which can get executed within SAP HANA Studio. If a backup solution gets used, then the backint interface can be leveraged. Refer to the SAP HANA product documentation for backup and restore information and requirements.

In addition, a vSphere deployed SAP HANA system can get protected by any SAP and VMware supported backup solution that leverages vSphere snapshots (see Figure 35). This reduces the backup and restore time and provides a storage vendor-agnostic backup solution based on vSphere snapshots. Check out the [SAP HANA on vSphere specific Veeam backup and recovery solution](#) as an example.

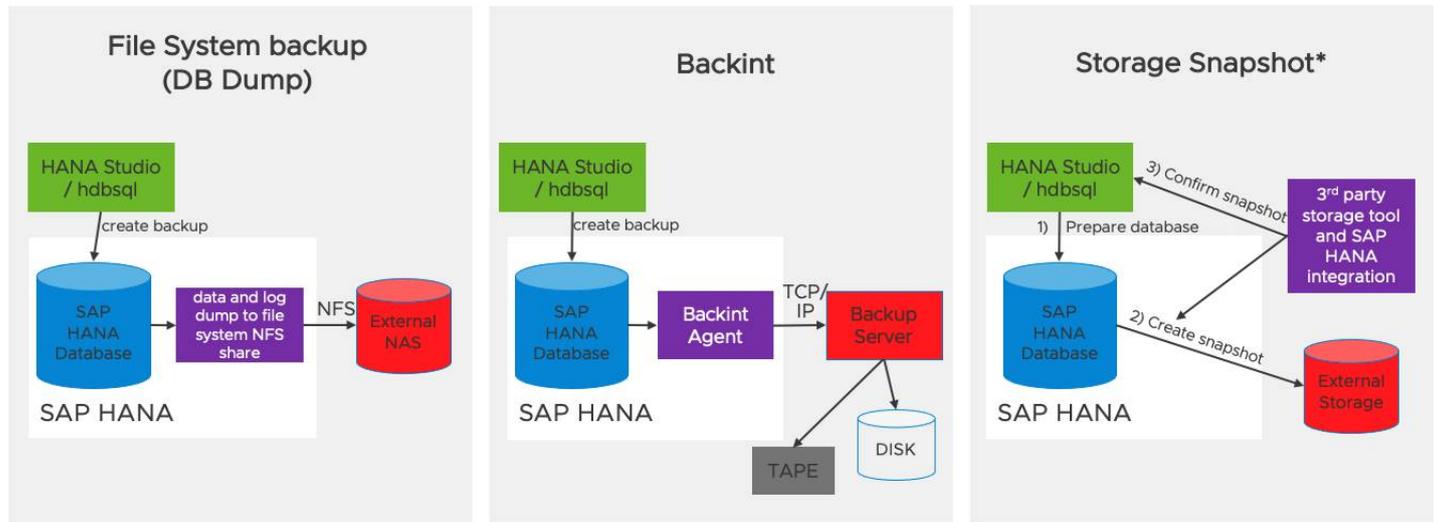


Figure 33: Virtualized SAP HANA Backup and Recovery Methods

SAP HANA with Persistent Memory on vSphere

Prerequisites and General SAP Support Limitations for Intel Optane PMem

- Review information published in SAP note [2913410](#): SAP HANA on VMware vSphere with Persistent Memory.
- Only server platforms that are [SAP HANA supported](#) and certified and listed on the [VMware Compatibility Guide](#) with the Persistent Memory feature are supported.
- Review the VMware / Intel PMEM [required](#) SAP HANA [best practices](#).
- The CPU and memory prerequisites of SAP HANA on vSphere must be met. For details, refer to SAP notes [2937606](#) and [3102813](#).
- General questions and boundary conditions about SAP HANA and Persistent Memory are addressed in SAP notes [2700084](#) and [2786237](#), and VMware knowledge base article [78094](#).
- For SAP HANA non-production support for memory mode, see SAP note [2954515: SAP HANA Persistent Memory - memory mode](#).

What is Supported?

SAP has granted support for SAP HANA 2 SPS 4 (or later) on vSphere versions[27] 6.7 (beginning with version 6.7 EP14) and vSphere 7.0 (beginning with version 7.0 P01) for 2- and 4-socket servers based on second-generation Intel Xeon Scalable processors (formerly code-named Cascade Lake). 8-socket host systems are yet not PMEM supported for SAP HANA with vSphere version 7.0 U2 and later and is in validation as of November 2021. The maximum DRAM plus Optane PMEM host memory configurations with SAP HANA supported 4-socket wide VMs on 4-socket hosts and can be up to 15 TB (current memory limit when DRAM with PMEM gets combined) and must follow the hardware vendor's Optane PMEM configuration guidelines.

The maximum VM size with vSphere 6.7 and vSphere 7.0 is limited to 256 vCPUs and 6 TB of memory. This results in SAP HANA VM sizes of 6 TB maximum for OLTP and 3 TB VM sizes for OLAP workloads (class L). vSphere 7.0 supports OLAP workloads up to 6 TB (class M). Supported DRAM to PMEM ratios are 2:1, 1:1, 1:2 and 1:4. Please refer to SAP note [2700084](#) for further details, use cases, and assistance in determining whether Optane PMEM is applicable at all for your specific SAP HANA workload.

Supported Optane PMEM module sizes are 128GB, 256GB and 512GB. Table 24 lists the supported maximum host memory DRAM and PMEM configurations. Up to two SAP HANA VMs are supported per CPU socket, **and up to a 4-socket large ESXi host system can be used**. See Table 24 for the currently supported configurations.

Note: vSphere 7.0 U2 or later versions are required for VM sizes >6 TB . These versions are, as of November 2021, not virtual SAP HANA PMem validated, and VM sizes with PMem >6 TB are not yet supported.

Table 26: Supported SAP HANA on vSphere with PMEM Ratios (as of November 2021)

Ratio [DRAM:DCPMM]	DRAM [TiB]	Optane PMem [TiB]	Total Host Memory [TiB]	Maximal VM Memory [TiB]*	Total Memory Slots (Cascade Lake)	DRAM Module Size [GiB]	Optane PMem Module Size [GiB]	SAP Supported ESXi Host Size (max. sockets)	
1	1	6	6	12	6	48	256	256	4
1	1	3	3	6	6	48	128	128	4
1	4	3	12	15	6	48	128	512	4
1	4	1,5	6	7,5	7,5	48	64	256	4
1	4	0,75	3	3,75	3,75	48	32	128	4

Sizing of Optane PMEM-enabled SAP HANA VMs

The sizing of Optane PMEM-enabled SAP HANA VMs is like bare-metal SAP HANA systems with the limitation of a maximum size of 6 TB (mix of DRAM and Optane PMEM) per VM. OLAP class-L workload sizing are limited to 3TB, class-M sizing support up to 6 TB total memory.

Please refer to SAP notes [2700084](#) and [2786237: Sizing SAP HANA with Persistent Memory](#) for details on compute and memory sizing for Optane PMEM-enabled SAP HANA systems.

We recommend that an SAP HANA VM use the same DRAM to PMEM ratio as the physical host/server DRAM to PMEM ratio. However, if you have a growth plan, you might consider a larger physical memory configuration, and upgrade the VMs and SAP

HANA over the lifetime.

For example, you have a 1:4 PMEM ratio host, configured with 15TB of total RAM (3TB DRAM and 12 TB Optane PMEM). An optimized resource scenario is to create four SAP HANA VMs on this server, each with 3.75TB RAM (0.75TB DRAM and 3TB Optane PMEM). If you are creating 6 TB VMs on this same 15TB host, then only two SAP HANA VMs can be created, which is a non-optimized resource configuration as you can only leverage 12 TB of the installed 15TB memory. In this case, a 1:1 DRAM to PMEM configuration, with a total of 12 TB (6 TB DRAM and 6 TB Optane PMEM) represents a resource-optimized configuration.

While >6 TB VMs are yet not supported, it is important to understand that, depending on the host memory configuration, not all memory can get leveraged for an SAP HANA VM. The following examples show optimized and a non-optimized memory configuration.

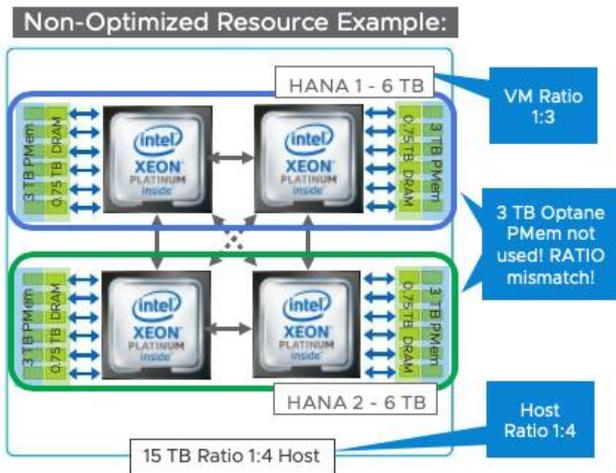


Figure 34: Non-Optimized Host Memory Configuration

4-Socket Host Configuration:

- Four 2nd Gen Intel Xeon Platinum processors, 24 x 128GB DRAM + 24 x 512GB Optane PMEM = 15 TB total host memory with a 1:4 DRAM to PMEM RATIO

VM Configuration example:

- 2 x 6 TB SAP HANA VM with 1.5 TB DRAM and 4.5 TB Optane PMEM RAM, with a 1:3 DRAM to PMEM RATIO

Challenges:

- DRAM: PMEM Ratio may not be suited for SAP HANA workload
- HW configuration does not fit and will lead to unusable PMEM (RATIO mismatch)

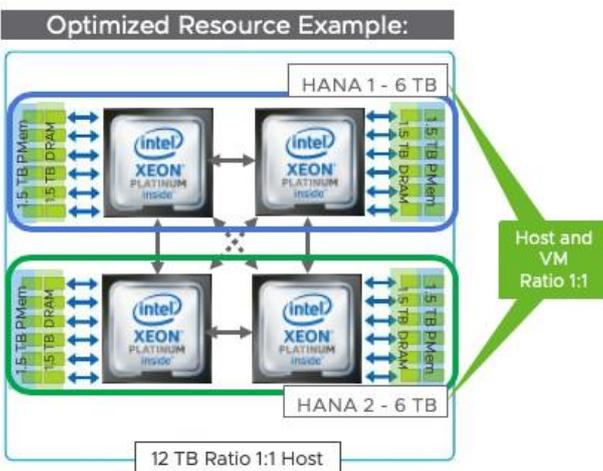


Figure 35: Optimized Host Memory Configuration

4-Socket Host Configuration:

- Four 2nd Gen Intel Xeon Platinum processors, 24 x 256GB DRAM + 24 x 256GB Optane PMEM = 12 TB total host

memory with a 1:1 DRAM to PMEM RATIO

VM Configuration example:

- 2 x 6 TB SAP HANA VM with 3 TB DRAM and 3 TB Optane PMEM RAM, with a 1:1 DRAM to PMEM RATIO

Challenges:

- Higher memory costs due to DRAM module prices

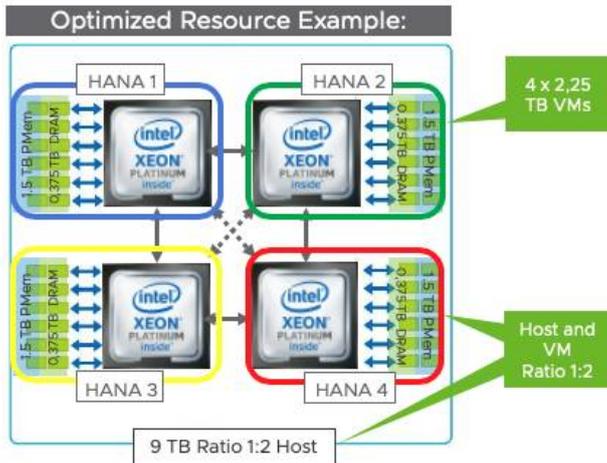


Figure 36: Optimized Host Memory Configuration

4-Socket Host Configuration:

- Four 2nd Gen Intel Xeon Platinum processors, 24 x 128GB DRAM + 24 x 256GB Optane PMEM = 9 TB total host memory

VM Configuration example:

- 4 x VM with 0.75 TB DRAM and 1.5 TB Optane PMEM RAM, total RAM per SAP HANA VM 2.25 TB with a 1:2 DRAM to PMEM RATIO

Challenges:

- SAP HANA Sizing to verify if Optane PMEM Ratio is applicable and if CPU resources are enough!

Note: WBS sizing's are supported and allow OLAP workloads with class-M CPU requirements to leverage up to 6 TB of total memory (DRAM and PMEM).

Because PMEM in App Direct mode provides data persistence in memory and is local to the host, not all vSphere features can be used equally to a DRAM-only VM. See Table 25 for details.

Using SAP HANA on vSphere allows HANA users to leverage the flexibility of vSphere capabilities, such as vMotion, which allow workloads to be migrated between vSphere hosts on Intel Xeon platforms without first having to be shut down. In addition, vSphere DRS works with a cluster of ESXi hosts to provide resource management capabilities, such as load balancing and VM placement to ensure a balanced environment for VM workloads.

vSphere HA is by now supported for SAP HANA VM with Optane PMEM use cases. For more information, read the VMware blog post, [“VMware vSphere 7.0 U2 and vSphere HA for SAP HANA with DRAM and Intel Optane PMem in App-Direct Mode.”](#)

Table 27: vSphere Features Supported with PMEM-enabled SAP HANA VMs

PMEM Modes	Usage Mode	SAP HANA Supported	Bootable Drive	Host and VM Size	VMware SMP	vSphere vMotion	VMware DRS	VMware SRM	VMware HA*	VM In-guest Cluster	VM Snapshot	SAP HSR
AppDirect (vNVDIMM)	persistence	Yes SAP Note 2913410	Not Applicable	Intel Cascade Lake 2 and 4 socket servers with up to 15 TB (up to 16 TB without PMEM). SAP HANA VM sizes up to 256 vCPUs and 6 TB of memory.	Yes, up to 4 physical CPU socket ESXi hosts and SAP HANA VMs.	Yes, up to 6 TB VM sizes (requires another PMEM host)		No	Yes, with vSphere 7.0 U2 or later, no support prior this release	No	No	Yes
Memory Mode	volatile	Only non-prod. SAP HANA support SAP Note 2954515		Only non-prod. SAP HANA support	Yes	Yes (can move/failover VMs to/from Memory Mode and DRAM systems)	Yes	Yes				
Mixed Mode	persistence/volatile	Only non-prod. SAP HANA support SAP Note 2954515		Mixed mode is not VMware vSphere supported!								
Storage over AppDirect	vPMemDisk	No	Yes	Not supported for SAP HANA		Yes, up to 6 TB VM sizes (requires another PMEM host and the usage of storage vMotion)		No (snapshot is not supported yet).	No	No	No	No

vSphere HA Support for PMEM-enabled SAP HANA VMs

vSphere HA was initially not supported for PMEM-enabled VMs before the vSphere 7.0 U2 release. Now, vSphere HA can support the failover and restart of PMEM-enabled VMs. The requirement is that the applications using PMEM maintain data persistence on PMEM as well as on shared disks.

SAP HANA is one of the applications that provides data persistence on disk. Because of this, vSphere HA can use this data on the shared disks to initiate a failover of PMEM-enabled SAP HANA VMs to another PMEM host. vSphere HA will automatically recreate the VM's NVDIMM configuration but is not in control over post VM failover OS/application-specific configuration steps, such as the required recreation of the SAP HANA DAX device configuration. This must be done manually or via a script, which is not provided by VMware nor SAP. For details on how to configure PMEM for SAP HANA, see the [Intel Optane Persistent Memory and SAP HANA Platform Configuration](#) guide.

Figure 39 illustrates the failover of a PMEM-enabled SAP HANA VM via vSphere 7.0 U2 and vSphere HA, and highlights that the PMEM NVDIMM configuration gets automatically re-created as part of the VM failover process. Once the DAX device is configured inside the OS, SAP HANA can be started and will automatically load the data from disk to the new PMEM regions assigned to this VM.

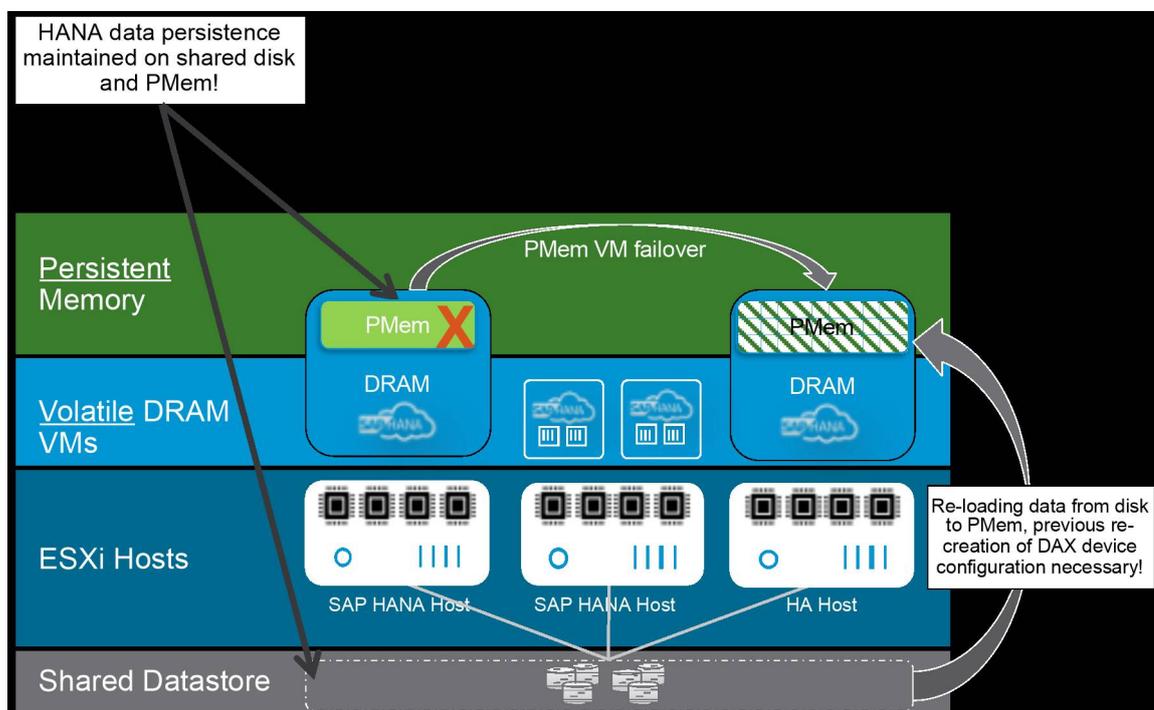


Figure 37: vSphere HA Support for PMEM-enabled VMs

After a successful failover of this PMEM-enabled VM, a garbage collector process will identify failed over VMs and free up the PMEM resources previously used by this VM on the initial host. On the host this VM now runs on, the PMEM will be blocked and reserved for the live time of this VM (as long it does not get migrated or deleted from the host).

The Intel SAP Solution Engineering team and the Intel and VMware Center of Excellence have developed an example script for the automatic [recreation of the DAX device configuration](#) on the OS level. This script must be executed after the failover and restart of the VM, prior to the restart of the SAP HANA database. It is advised to automatically run this script as part of the OS start procedure, such as a custom service. The script can be used as a template to create your own script that fits your unique environment.

Note: This script is not maintained nor supported by VMware, SAP or Intel. Any usage of this script is your own responsibility.

SAP HANA with PMEM VM Configuration Details

Utilizing PMEM in a vSphere virtualized environment requires that the physical host, ESXi and VM configuration gets configured correctly.

Follow the [Intel Optane Persistent Memory and SAP HANA Platform Configuration on VMware ESXi](#) configuration guide to prepare the needed DAX devices and see how to configure SAP HANA to enable PMEM.

The following list outlines the configuration steps. Refer to the hardware vendor-specific documentation to correctly configure PMEM for SAP HANA.

- HOST:
1. Configure Server host for PMEM using BIOS (vendor specific)
 2. Create AppDirect interleaved regions and verify that they are configured for ESXi use.
- VM:
1. Create a VM with HW version 19 (vSphere 7.0 U2 or later) with NVDIMMs and allow failover to another host while doing this.
 2. Edit the VMX VM configuration file and make the NVDIMMs NUMA aware.
- OS:
1. Create a file system on the namespace (DAX) devices in the OS.
 2. Configure SAP HANA to use the persistent memory file system.
 3. Restart SAP HANA to activate and start using Intel Optane PMEM.

Details on VM configuration steps.

Before you can add NVDIMMs to an SAP HANA VM, check if the PMEM regions and namespaces were created correctly in the BIOS. Also, ensure that you have selected all PMEM as “persistent memory” and that the persistent memory type is set to App Direct Interleaved. See the example in Figure 40.

After you have created the PMEM memory regions, a system reboot is required. Now, install the latest ESXi version (e.g., 7.0 U2 or later) and check

via the ESXi host web client if the PMEM memory modules, interleave sets, and namespaces have been set up correctly. See the examples in Figures 38–41.



Figure 38: Example of PMEM System BIOS Settings

The screenshot shows the ESXi web client interface for Persistent Memory. A table lists various memory modules, each with a 126 GB capacity and 0 B free space, all in a 'Normal' health state.

ID	Capacity	Free	Health
0x1	126 GB	0 B	Normal
0x111	126 GB	0 B	Normal
0x21	126 GB	0 B	Normal
0x101	126 GB	0 B	Normal
0x11	126 GB	0 B	Normal
0x121	126 GB	0 B	Normal
0x1001	126 GB	0 B	Normal
0x1111	126 GB	0 B	Normal
0x1021	126 GB	0 B	Normal
0x1101	126 GB	0 B	Normal
0x1011	126 GB	0 B	Normal
0x1121	126 GB	0 B	Normal
0x2001	126 GB	0 B	Normal
0x2111	126 GB	0 B	Normal
0x2021	126 GB	0 B	Normal
0x2101	126 GB	0 B	Normal
0x2011	126 GB	0 B	Normal
0x2121	126 GB	0 B	Normal

Figure 39: ESXi Persistent Memory Storage View of Modules

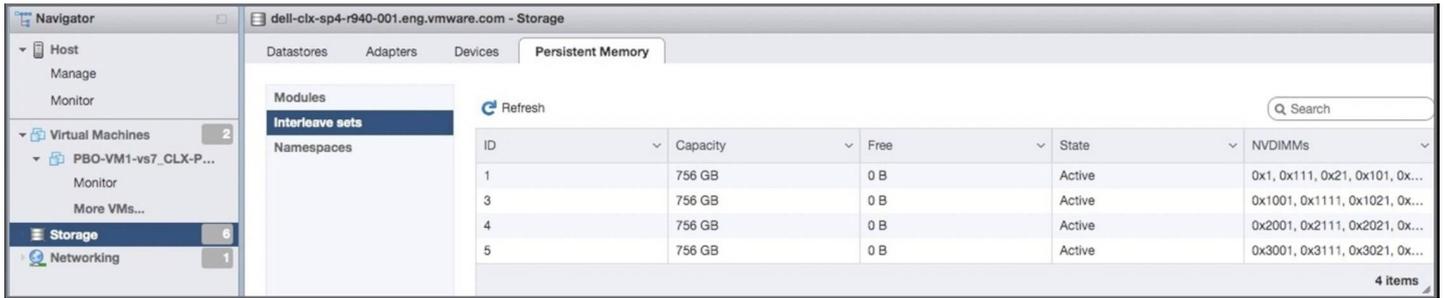


Figure 40: ESXi Persistent Memory Storage View of Interleave Sets

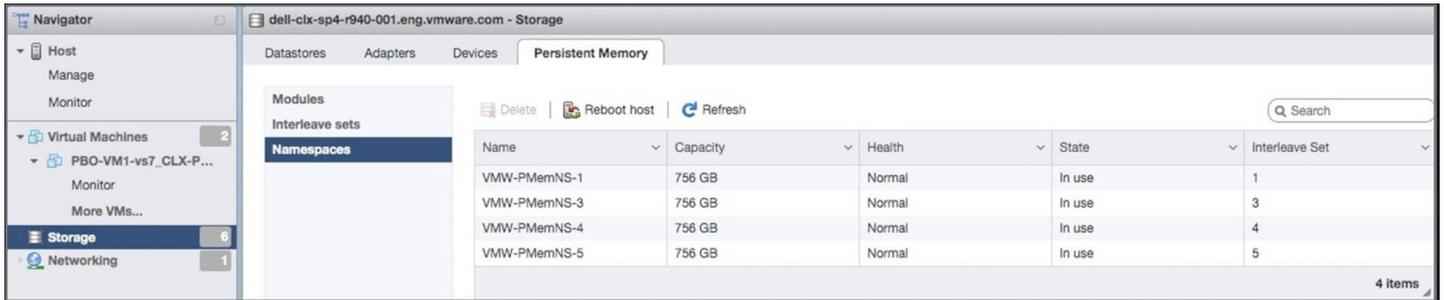


Figure 41: ESXi Persistent Memory Storage View of Namespaces

Note: The interleave set numbers shown depend on the hardware configuration and may differ in your configuration.

If the configurations were done correctly in the BIOS of the host, then the configuration should look like what is shown in Figure 42. After this, you can add NVDIMMs and NVDIMM controllers to your SAP HANA VM. Select the maximum size possible per NVDIMM, otherwise you waste memory capacity.

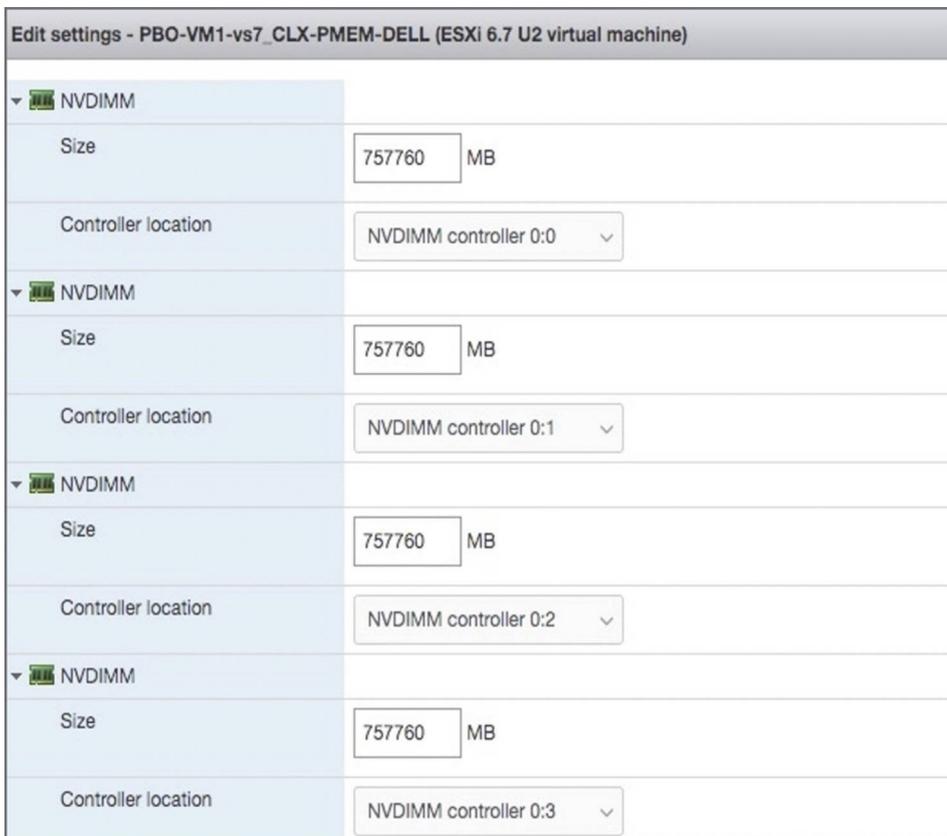


Figure 42: NVDIMM Creation via the vCenter GUI

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. VMware knowledge base article [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.

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. To solve this, you must add the following settings to a VM configuration using vCenter.

Example for a 4-socket wide VM:

- `nvdimm.mode = "independent-persistent"`
- `nvdimm0:0.nodeAffinity=0`
- `nvdimm0:1.nodeAffinity=1`
- `nvdimm0:2.nodeAffinity=2`
- `nvdimm0:3.nodeAffinity=3`
- `sched.pmem.prealloc=TRUE` (optional)

Note: `sched.pmem.prealloc=TRUE` is an optional parameter equivalent to eager zero thick provisioning of VMDKs and improves initial writes to Optane PMEM. Be aware that the first vMotion process with this parameter set will take a long time due to the pre-allocation of the PMEM in the target server.

Besides these parameters, you may also configure the CPU NUMA node affinity or CPU affinities (pinning) as described in the SAP HANA best practices parameter guidelines listed in the [Best practices of virtualized SAP HANA systems section](#).

Note: The parameters in the example above must get manually added after the PMem SAP HANA VM is created.

Verify the VMX file of the newly created VM and check if the NVDIMM configuration looks like the following example. The easiest way to do this is to use the ESXi PowerShell.

Example output of the.vmx file of a PMEM-enabled VM:

```
[root@ESXiHOSTxxx:/vmfs/volumes/XXXXXX/PMem_SAP_HANA_VM_name] grep -i nvdimm *.vmx
nvdimm0.present = "TRUE"
nvdimm0:0.present = "TRUE"
nvdimm0:0.fileName = "/vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_1.vmdk" nvdimm0:0.size = "757760"
nvdimm0:1.present = "TRUE"
nvdimm0:1.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_3.vmdk" nvdimm0:1.size = "757760"
nvdimm0:2.present = "TRUE"
nvdimm0:2.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_5.vmdk" nvdimm0:2.size = "757760"
nvdimm0:3.present = "TRUE"
nvdimm0:3.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_5.vmdk" nvdimm0:3.size = "757760"
nvdimm0:0.node = "0"
nvdimm0:1.node = "1"
nvdimm0:2.node = "2"
nvdimm0:3.node = "3"
```

manual added parameters:

```
nvdimm.mode = "independent-persistent"
nvdimm0:0.nodeAffinity=0
nvdimm0:1.nodeAffinity=1
nvdimm0:2.nodeAffinity=2
nvdimm0:3.nodeAffinity=3
sched.pmem.prealloc=TRUE (optional and will cause time delays during the first vMotion process)
```

Note: The VMDK disk numbers shown depend on the hardware configuration and may differ in your configuration.

SAP HANA Scale-Out on vSphere Installation

Unlike in physical SAP HANA environments, it is possible to operate SAP HANA as highly available without a Standby host. Host/VM failover functionality can be substituted by VMware HA.

The current version of the SAP HANA installer always expects a Standby host configuration even if none is needed. To work around this installer issue, follow these installation steps:

Install SAP HANA as a single host installation on the first VM:

- Use this command:
 - `./hdblcm`
 - `--action=install`
 - `--sid=<SID>`
 - `--number=<xx>`
 - `--sapmnt=/hana/shared/ <- Your shared folder for HANA`
 - `--datapath=/hana/data/ <- Your data folder for HANA`
 - `--logpath=/hana/log/ <- Your log folder for HANA`
 - `--restrict_max_mem=off`
 - `--nostart=off`
 - `--shell=/bin/csh`
 - `--hostname=<hostname>`
 - `--remote_execution=ssh`
 - `--install_hostagent=on`
 - `--password=<sidadm_password>`
 - `--system_user_password=<system_user_pw>`

If you want to use an extra internal network, you need to add this:

- `--internal_network=192.168.x.x/24`

If you have not extracted the HANA installables without the manifest files

- `--ignore=check_signature_file`

After a successful installation, go to „`/usr/sap/<SID>/SYS/global/hdb/custom/config`“ and add the following entries to the „`global.ini`“:

- -----
- `[communication] listeninterface = .global <- .internal will be here, if you use an internal IP adapter`
- `[persistence] basepath_shared = no`
- -----

After editing the `global.ini` file, restart this SAP HANA instance.

On the next VM you want to add as a worker node to the SAP HANA installation, use `./hdblcm` and follow the instructions.

Repeat these steps for every worker you want to add to your SAP HANA installation.

Enable Automatic Start of the SAP HANA Database:

By default, the SAP HANA database is configured so that it does not start automatically when the SAP HANA host is rebooted. If VMware HA is selected as the HA solution for SAP HANA, it is recommended to enable the automatic start feature of SAP HANA.

To do this, modify the SAP HANA database profile with the following procedure:

1. Log on to the SAP HANA host as a user with root authorization.
2. Change to the system profile directory `/usr/sap/<SID>/SYS/profile`.
3. Edit the profile `<SID>_HDB<instance number>_<host name>`.
4. Change the parameter setting `Autostart = 0` to `Autostart = 1`.
5. Save the profile `<SID>_HDB<instance number>_<host name>`.

With this modification, the SAP HANA database will start automatically when the SAP HANA VM is rebooted.

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.

VMware NUMA Observer

The next chapter discusses the best practices parameters to optimally configure an SAP HANA on VMware vSphere VM. The most critical aspect of these optimizations is that VMware administrators configure an SAP HANA VM NUMA aligned to get the best performance and lowest memory latency.

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

To check if the initial configuration is correct or to detect misalignments you can use the VMware NUMA observer, which is available to download from <https://flings.vmware.com/numa-observer>.

The NUMA Observer Fling scans your VM inventory and identifies VMs with overlapping core/NUMA affinities and generates alerts. Additionally, the Fling also collects statistics on remote memory usage and CPU starvation of critical VMs and raises alerts, see figure 46 and 47 as an example.

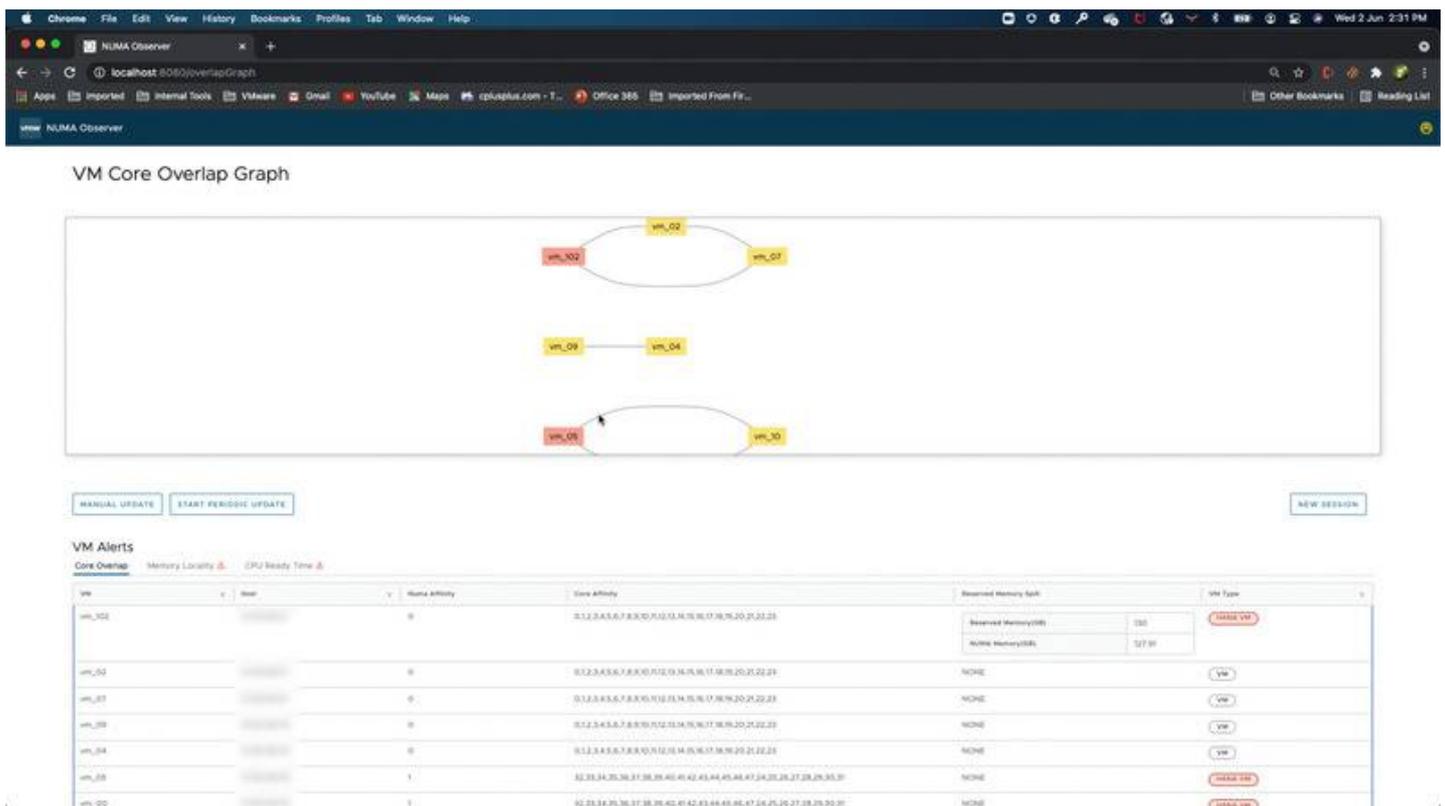


Figure 43: VMware NUMA Observer - VM Core Overlap Graph

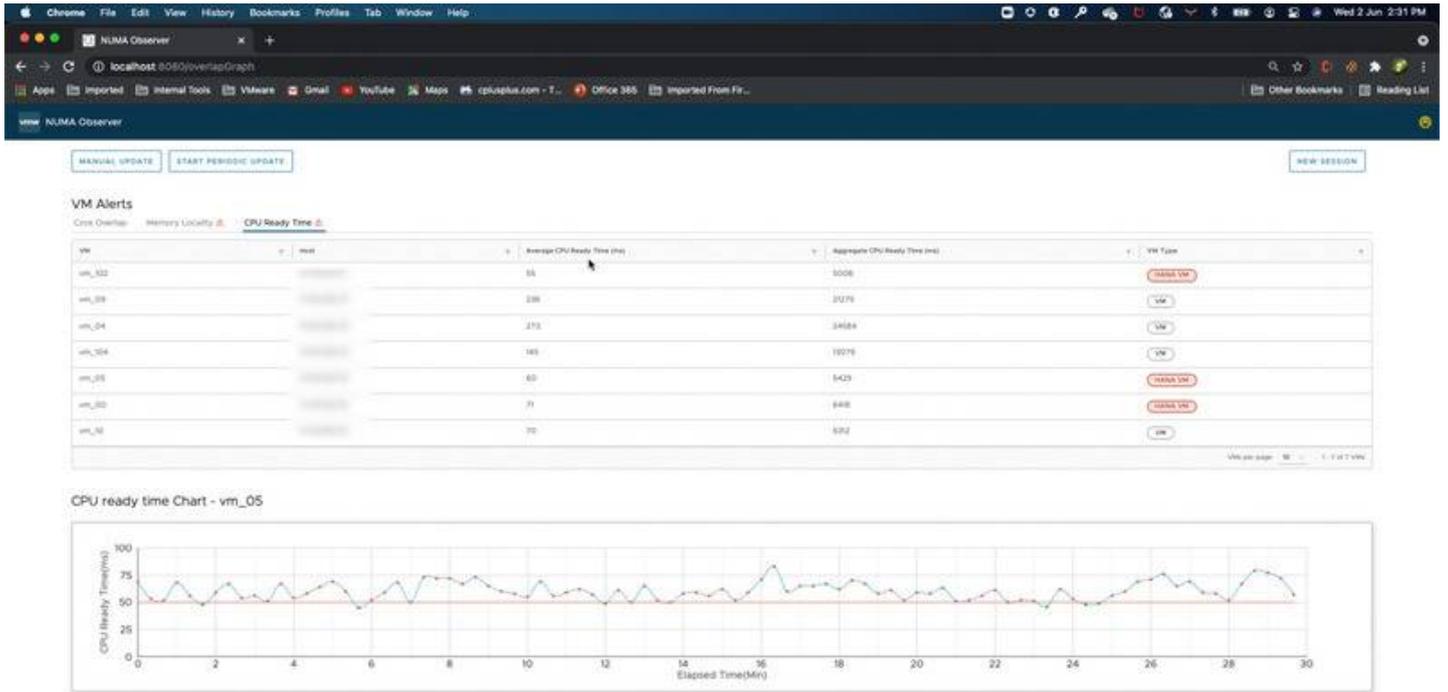


Figure 44: VMware NUMA Observer - VM Alerts

Best Practices of Virtualized SAP HANA Systems

Third Generation Intel Xeon 40-core Processor (Ice Lake SP) Specifics

While the vSphere 7 ESXi kernel supports up to 896 CPU threads (PCPUs) per physical host and up to 768 vCPUs per VM, a vSphere 7 VM configuration allows maximal 64 vCPUs per virtual socket (coresPerSocket). This is not a limitation of the hypervisor, as noted by SAP, it is a configuration maximum all vSphere ESXi versions up to 7.x have. This is documented in [Hardware Features Available with Virtual Machine Compatibility Settings](#) under: "Maximum number of cores (virtual CPUs) per socket". This configuration maximum increases in vSphere ESXi 8.0.

In the case of an Intel Ice Lake CPU with > 32-cores, like a 40-core Ice Lake CPU 2-socket based system, this configuration maximum is reached and instead of a 2-virtual socket VM a 4-virtual socket VM needs to get configured when all CPU threads (PCPUs) should get leveraged for a single large VM. This results that SAP HANA/the OS sees inside the VM a 4-virtual socket VM with, in the case of the 40-core Ice Lake CPU, 160 vCPUs, instead the expected 2-virtual socket VM. This is a normal and expected behavior and shows the correct number of virtual-CPU sockets as configured.

The VM socket information is from the vSphere ESXi perspective, a logical grouping of the CPU threads (PCPUs) presented to the OS level. All information like the CPU thread number per NUMA node gets exposed and is available for the application/SAP HANA running inside the VM.

Again, ESXi is fully aware of the underlying HW and reports the **two physical NUMA** nodes to the installed OS and SAP HANA instance. This can be checked by using LSCPU, the sapservices log file or the SAP HANA Database Server Management Console (hdbcons) output. The VM workload gets scheduled as optimally as possible per NUMA node to limit the need to access remote memory.

Since SAP HANA is NUMA aware and uses the NUMA topology information to ensure memory locality. We do not expect any virtual socket related issues when 4- instead 2-virtual CPU sockets are configured. In addition, the SAP HANA on vSphere validation test was successful and did not show any related performance or functional problems.

In the case a customer experiences issues caused by this not physical hardware aligned VM configurations (for example, wrong CPU-socket representation), follow VMware KB 89881 and open an SAP support ticket using our SAP support component BC-OP-LNX-ESX.

This configuration maximum is described in detail in VMware KB [89879](#). **No special configuration is required for Ice Lake CPUs with 32 or fewer CPUs cores.**

The upcoming vSphere 8 release improves on this configuration maximum and will allow to configure >64 vCPUs per virtual socket (coresPerSocket) VMs once it is SAP HANA supported.

Configuration examples:

Server system with 2 x **40-core Ice Lake CPU**:

- 2 physical socket (complete host) consuming configuration:
 - vCPU = 160
 - Cores per Socket = 40 -> 4 vSocket
- 2 physical socket (complete host) consuming configuration **without hyperthreading**:
 - vCPU = 80
 - Cores per Socket = 40 -> 2 vSocket
- 1 physical socket consuming configuration:
 - vCPU = 80
 - Cores per Socket = 40 -> 2 vSocket
- 1 physical socket (complete host) consuming configuration **without hyperthreading**:
 - vCPU = 40
 - Cores per Socket = 40 -> 1 vSocket
- 0.5 physical socket consuming configuration:
 - vCPU = 40
 - Cores per Socket = 40 -> 1 vSocket

Server system with 2 x **32-core Ice Lake CPU**:

- 2 physical socket (complete host) consuming configuration:
 - vCPU = 128
 - Cores per Socket = 64 -> 2 vSocket
- 1 physical socket consuming configuration:
 - vCPU = 64
 - Cores per Socket = 64 -> 1 vSocket
- 0.5 physical socket consuming configuration:
 - vCPU = 32
 - Cores per Socket = 32 -> 1 vSocket

Where to Configure?

When you build a new VM, ensure that you select the correct number of vCPUs and select the correct, for the CPU type, number of “cores per socket”, to align the VM as best as possible to the underlying HW.

Figure 45 shows a 40-core Ice Lake and because of the vSphere 7 64-cores per socket limitation, we must use 40 cores to ensure we get a reasonable virtual machine socket configuration. Figure 46 shows how to configure a 2-socket wide VM on a 32-core Ice Lake based on a ESXi vSphere 7 host, which will align with the two-socket physical 32-core Ice Lake server.

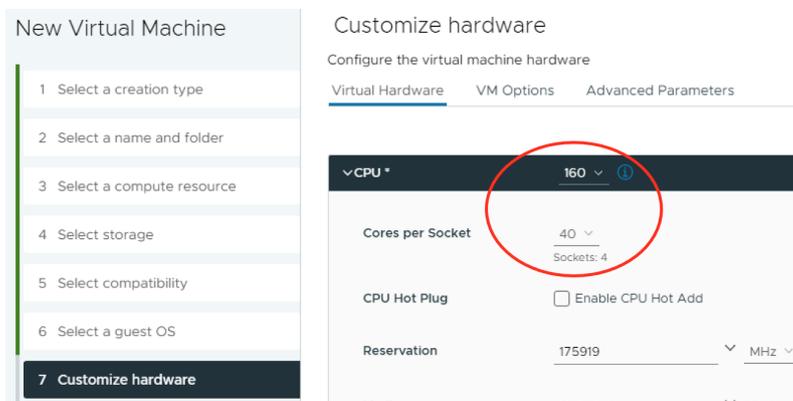


Figure 45: 4-socket Wide VM on 40-core Ice Lake vSphere 7 ESXi Host

Note: Due to the vSphere 7 Cores per socket limitation to 64, we must configure a 4-socket VM instead the recommended 2-socket wide VM.

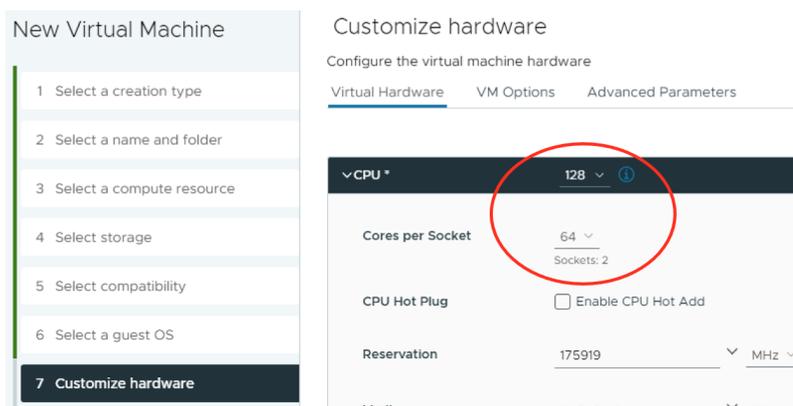


Figure 46: 2-socket wide VM on 32-core Ice Lake vSphere 7 ESXi Host

Below is a graphical representation of a VMware virtualized Ice Lake 40 Core system. The picture shows that the vCPUs are directly mapped to their corresponding CPU threads/cores regardless the vSocket configuration, which is a logical “grouping”. For SAP HANA, the NUMA topology is the important information, which is correctly exposed. See figure below and next section.

(Use zoom function to make below graphic readable.)



Figure 47: Graphical Representation of a VMware Virtualized Ice Lake 40 Core system with vSphere 7 (4-socket VM) and vSphere 8 (2-socket VM)**Where to Check the NUMA Topology?**

While all SAP HANA tools and services correctly report two NUMA Nodes of the 2-socket Ice Lake server, the socket information may differ from the underlying physical hardware when a 40-core Ice Lake CPU is configured as a 4-virtual socket /2-vNUMA VM to utilize the full 160 threads as vCPUs.

Example output of Lscpu and hdbcons of a 4-socket SAP HANA VM configured with 160 vCPUs. Lscpu shows the 4-socket that got configured to utilize all 160 CPU threads (PCPUs), but also shows correctly the two NUMA nodes important for SAP HANA:

Lscpu:

- ..> lscpu
- Architecture: x86_64
- //
- CPU(s): 160
- On-line CPU(s) list: 0-159
- Thread(s) per core: 2
- Core(s) per socket: 20
- **Socket(s): 4 (this is the number of virtual CPU sockets as defined in the VM configuration)**
- **NUMA node(s): 2 (virtual NUMA nodes)**
- //
- Model name: Intel(R) Xeon(R) Platinum 8380 CPU @ 2.30GHz
- //
- Hypervisor vendor: VMware
- Virtualization type: full
- //
- NUMA node0 CPU(s): 0-79
- NUMA node1 CPU(s): 80-159

hdbcons "mm numa -t"

- *****
- Configuration of NUMA topology
- *****
- Is NUMA-awareness enabled? (0|1) : 1
- Valid NUMA Topology? (0|1) : 1
 - Number of NUMA Nodes : 2
- Number of NUMA Nodes with logical cpus : 2
- Number of NUMA Nodes with allowed memory: 2
- Number of NUMA Nodes with both cpus and memory: 2
- Number of Logical cpus : 160
 - //
- *****

Performance of an SAP HANA 40-core Ice Lake CPU based VM

While running internal mixed workload tests with different network card configurations (VMXNET3 and PT configured NICs) with simulated S/4 HANA OLTP and OLAP workload the Ice Lake platform has shown excellent results.

Table 28 shows that the virtual BWH performance delta was in all categories way below 10% and only phase one, the data load phase showed a delta >5%. All other phases were below 3%, see table below and SAP benchmark certificates [2022023](#) and [2022024](#).

Table 28. Ice Lake Bare Metal vs. vSphere 7 VM

2-Socket 4 TB Ice Lake Server / VM	Cert	CPU Threads	MEM	Records	BWH Phase 1 (sec.)	Delta	BWH Phase 2 (QPH)	Delta	BWH Phase 3 (sec.)	Delta
Fujitsu RX2540 M6 Bare Metal	2022023	160	4096 GB	7,800,000,000	13783		3219		107	
vSphere 7 VM with VMXNET3	2022024	160	3968 GB		14701	6,7%	3145	-2,3%	109	1,9%
					"- is better"		"+ is better"		"- is better"	

Below figure shows the differences of bare metal and VM based SAP HANA on Ice Lake BWH benchmark in a graphical way.

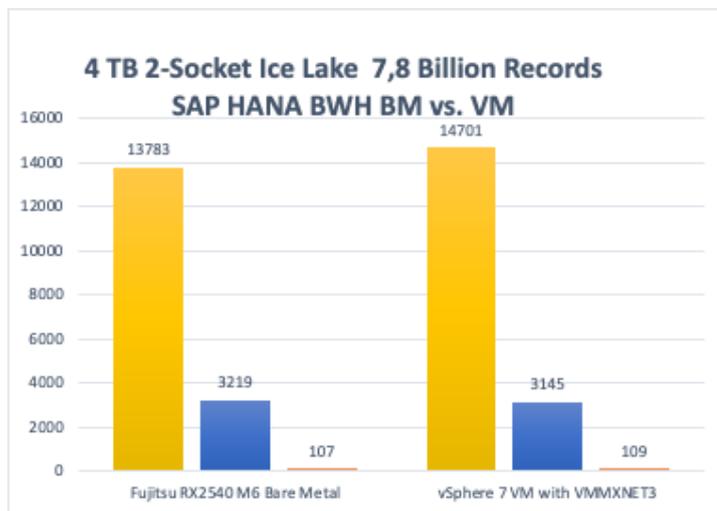


Figure 48: Ice Lake Bare Metal vs. vSphere 7 VM

The 2-socket Ice Lake platform has the same 4TB memory support from SAP as the aged Broadwell platform. This makes Ice Lake a very attractive platform for a hardware refresh since the Broadwell based VMs would be able to maintain the memory configuration.

Figure 49 and 50 show how an Ice Lake 2-socket 40c/160 thread CPU system compares against an aged 4-socket Broadwell 22c/176 thread CPU system.

All validation tests got executed in the exact same test environment we have used some time ago for our Broadwell vSphere 7 re-validation and with the same versions (test tools and SAP HANA, vSphere and SLES versions). Because of this we can directly compare the test results of the aged Intel Broadwell and new Intel Ice Lake platform.

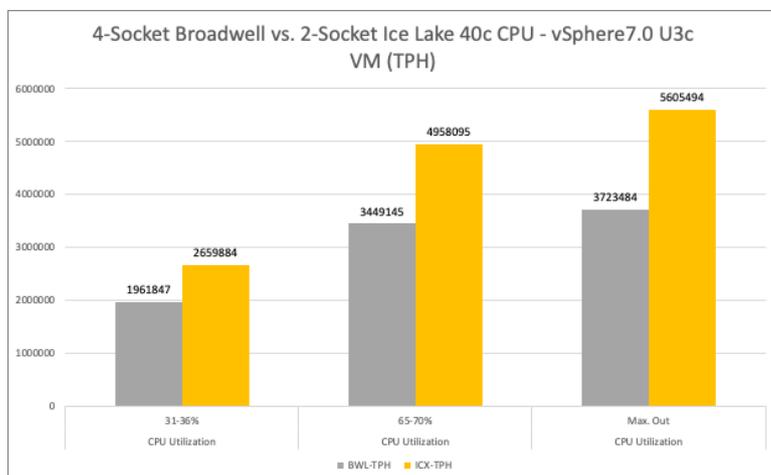


Figure 49: S/4 HANA Ice Lake Bare 160 vCPU VM vs. 176 vCPU Broadwell VM TPH Result

The performance improvements of vSphere 7 virtualized Intel Ice Lake VMs are significant and are showing up to 50% better TPH

and QPH results, while the Broadwell system had 10% more CPU threads (PCPUs).

We also have measured between -70% to -30% lower database request times (see next figures), while using up to date 10 GbE network cards and the newer CPU platform.

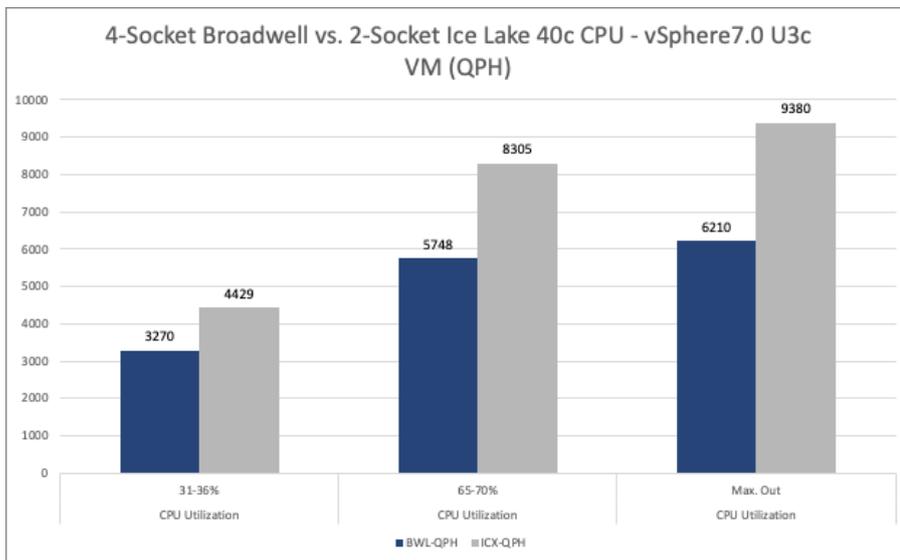


Figure 50: S/4 HANA Ice Lake Bare 160 vCPU VM vs. 176 vCPU Broadwell VM QPH result

In the case of a lower bin Ice Lake CPU used, the less CPU threads (PCPUs) must be subtracted from the results; for example, 32 c vs. 40 c Ice Lake CPU would have around 20% less TPH/QPH, with this performance it would be still more than sufficient to replace a Broadwell system.

The security and performance improvements of an Intel Ice Lake CPU (up to 50%), and the same memory support as Broadwell (up to 4 TB), makes Ice Lake a very attractive platform for a hardware and vSphere version refresh. The easy upgrade of a VM to adapt the Ice Lake platform changes is a matter of minutes and your SAP HANA installation will benefit greatly from an upgrade to vSphere 7.

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 support engineers will verify and, if not applied, will recommend configuring these settings.

Table 29. Physical Host BIOS Parameter Setting

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	To avoid timer synchronization issues, use a multi-socket server that ensures NUMA node timer synchronization. NUMA systems that do not run synchronization will need to synchronize the timers in the hypervisor layer, which can impact performance. See the Timekeeping in VMware Virtual Machines information guide for reference. Select only SAP HANA CPUs supported by vSphere. Verify the support status with the SAP HANA on vSphere support notes. For a list of the relevant note, see the SAP Notes Related to VMware page.
Enable CPU Intel Turbo Boost	Allow Intel automatic CPU core overlocking 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 (PCPUs).
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.
Uncore Frequency (scaling)	Set this BIOS parameter that it is always set to its predefined maximum. Check the HW vendor BIOS documentation how to do this. This parameter is especially important for Ice Lake based systems.
Set correct PMem mode as specified by the hardware vendor for either App Direct or Memory mode	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. Memory mode is only supported with a ratio of 1:4. As of today, SAP provides only non-production workload support. VMware vSAN does not have support for App Direct mode as cache or as a capacity tier device of vSAN. However, vSAN will work with 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.
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.

Table 30. ESXi Host Parameter Setting

ESXi host parameter settings	Description
Networking	Use virtual distributed switches to connect all hosts that work together. Define the port groups that are dedicated to SAP HANA, management and vMotion traffic. Use at least dedicated 10 GbE for vMotion and the SAP app server or replication networks. At least 25 GbE for vMotion for SAP HANA system \geq 2TB is recommended.
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> <p>Add the following advanced VMX configuration parameters to the VMX file, and reboot the VM after adding these parameters:</p> <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> <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> <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>

Table 31. SAP HANA Virtual Machine Parameter Setting

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	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. See VMware DOC-28494 for details. You can configure this with the vSphere Client™ by choosing EFI boot mode. 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.
SAP monitoring	Enable SAP monitoring inside the SAP HANA VM with the advanced VM configuration parameter tools.guestlib.enableHostInfo = "TRUE". For more details, see SAP note 1409604 . Besides setting this parameter, the VMware guest tools need to be installed. For details, see VMware KB 1014294 .
vCPU hotplug	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 .
Memory reservations	Set fixed memory reservations for SAP HANA VMs. Do not overcommit memory resources. 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. 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.
CPU	Do not overcommit CPU resources and configure dedicated CPU resources per SAP HANA VM. 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 . If you want to use hyperthreads, then you must configure 2x the cores per CPU socket of a VM (e.g., 2-socket wide VM on a 28 core Cascade Lake system will require 112 vCPUs).
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. The vNUMA nodes of the VM will always be ≥ 1 , depending on how many CPUs you have configured in total. If you need to access an additional NUMA node, use all CPU cores of this additional NUMA node. Nevertheless, use as few NUMA nodes as possible to optimize memory access.

Table 32. SAP HANA Virtual Machine Parameter Settings

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	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. Do not change this setting to "high" or "low." Change this setting under the instruction of VMware support engineers.
Associate virtual machines with specified NUMA nodes to optimize NUMA memory locality	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. Use the vSphere Web Client or directly modify the.vmx file and add NUMA. nodeAffinity=x (for example: 0,1). 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. Procedure: 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. For details, see Associate Virtual Machines with Specified NUMA Nodes .

SAP HANA virtual machine parameter settings	Description
Configure virtual machines to use hyperthreading with NUMA	<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. 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: <code>numa.vcpu.preferHT="TRUE"</code> (per VM setting) or as a global setting on the host: <code>Numa.PreferHT="1"</code> (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>
PMEM-enabled VMs	<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. 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. Example for a 4-socket wide VM: <code>sched.pmem.prealloc=TRUE</code> 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.

Table 33. Linux OS Parameter Setting

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 SAPTune to optimize the Linux OS for SAP HANA.
Check if the queue depths of the SCSI default values are aligned for optimizing large-scale workloads with intensive I/O patterns. Newer Linux version sets the parameters automatically.	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. See VMware KB 2053145 for details how to check and configure these parameters. 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.
Configure NTP time server	Use the same NTP server as configured for vSphere ESXi host. 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	This helps to lower network latency of client/application server facing NIC adapters run: "ethtool -K ethY lro off". Do not disable LRO for throughput NIC adapters such as for backup, replication, or SAP HANA internode communication networks. Works only with Linux kernel 2.6.24 and later and uses a VMXNET3. Additional details: http://kb.vmware.com/kb/2055140

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

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

Table 34. Additional Performance Tuning Settings for SAP HANA Workloads

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"	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. Edit the.vmx file and add the following two advanced parameters: monitor.idleLoopSpinBeforeHalt = "true" AND monitor.idleLoopMinSpinUS = "xx" (For example: 50) Both parameters must be configured to influence the de-scheduling time. 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.
Set monitor_control.halt_in_monitor = "TRUE"	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. 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. Use the vSphere Web Client and add the following advanced VM parameter: monitor_control.halt_in_monitor = "TRUE".
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
Configuring CPU affinity sched.vcpuXx.affinity = "Yy-Zz"	<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 (PCPUs) 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: For more information about potential performance practices, see vSphere Resource Management Guide.</p>
>4-socket VM on 8-socket hosts	<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 how to set additional VMX parameters for SAP HANA half- and full-CPU socket VMs. These parameters can get added via the vSphere Web Client or by directly adding these parameters to the.vmx file with a text editor.

Table 35. Example VMX Configurations for SAP HANA VMs

SAP HANA half-socket VM additional VMX Parameters	Settings
First half-socket VM on socket 0 on a 28-core CPU n-socket server	
First half-socket PMEM VM on socket 0 on a 28-core CPU n-socket server	
Second half-socket VM on socket 0 on a 28-core n-socket CPU server	
Second half-socket PMEM VM on socket 1 on a 28-core CPU n-socket server	

SAP HANA 1-socket VM additional VMX parameters	Settings
1-socket VM on socket 3 on a 28-core CPU 4 or 8-socket server	
1-socket PMEM VM on socket 3 on a 28-core CPU 4 or 8-socket server	

SAP HANA 2-socket VM additional VMX parameters	Settings
2-socket VM on sockets 0 and 1 on a 28-core CPU n-socket server	
2-socket PMEM VM on sockets 0 and 1 on a 28-core CPU n-socket server	

SAP HANA 3-socket VM additional VMX parameters	Settings
3-socket VM on sockets 0, 1 and 2 on a 28-core CPU 4 or 8-socket server	
3-socket PMEM VM on sockets 0, 1 and 2 on a 28-core CPU 4 or 8-socket server	

SAP HANA 4-socket VM additional VMX parameters	Settings
4-socket VM on a 28-core CPU on a 4-socket server	No additional settings are required as the VM utilizes all server resources.
4-socket PMEM VM on a 28-core CPU on a 4-socket server	

SAP HANA 4-socket VM additional VMX parameters	Settings
4-socket VM on a 28-core CPU 8-socket server running on sockets 0-3	

SAP HANA 5-socket VM additional VMX parameters	Settings
6-socket VM on a 28-core CPU 8-socket server running on sockets 0-5	

SAP HANA 8-socket VM additional VMX parameters	Settings
8-socket VM on a 28-core CPU 8-socket server	

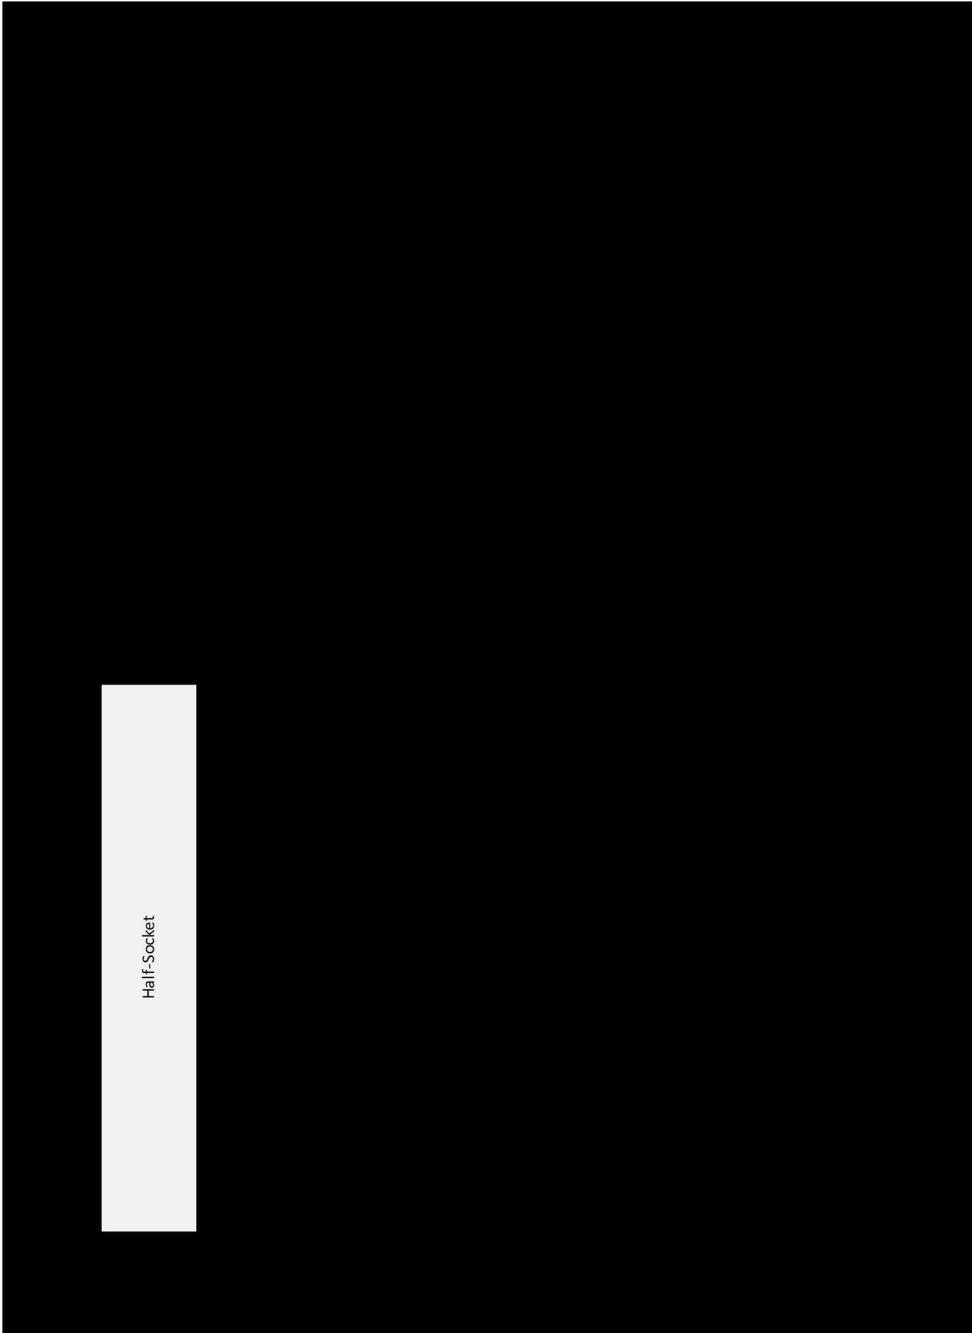
SAP HANA <u>low-latency</u> VM additional VMX parameters	Settings
n-socket <u>low-latency</u> VM on an n-socket, CPU server (Valid for all VMs when even a lower latency is required.)	

CPU Thread Matrix Examples

The following table shows the CPU thread matrix of a **28-core CPU** as a reference when configuring the `sched.vCPUx.affinity = "Xx-Yy"` parameter. The list shows the start and end ranges required for the "Xx-Yy" parameter (e.g., for CPU 5, this would be 280-335).

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

The following table shows the CPU thread matrix of a **24 core CPU** as a reference when configuring the sched.vCPUx.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".



The following table shows the CPU thread matrix of a **22 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 5 it would be "220-263".

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	

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 4 it would be "108-143"

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

SAP HANA Support and Process

In the case of supporting virtualized SAP HANA systems, customers can open a ticket directly with SAP. The ticket will be routed directly to VMware and SAP HANA support engineers, who will then troubleshoot the escalated issue.

Open an SAP Support Request Ticket

VMware is part of the SAP support organization, allowing VMware support engineers to work directly with SAP, SAP customers, and other SAP software partners, such as SUSE, as well as with hardware partners on solving issues needing escalation.

Before opening a VMware support ticket, we recommend opening a support request within the SAP support system when the SAP HANA system runs virtualized with VMware. This ensures that SAP HANA and VMware specialists will work on the case and, if needed, escalate the issue to VMware product support (when it is a VMware product issue) or to SAP support (when it is an SAP HANA issue).

The following components are available for escalating SAP on vSphere issues:

- BC-OP-NT-ESX (Windows on VMware ESX®)
- BC-OP-LNX-ESX (Linux on VMware ESX and SAP HANA)

Issues related to SAP HANA on vSphere should be escalated directly via SAP Solution Manager to BC-OP-LNX-ESX. In the case it is a non-VMware-related SAP HANA issue, the escalation will be moved to the correct support component. Figure 51 shows the support process workflow for VMware-related SAP HANA issues.

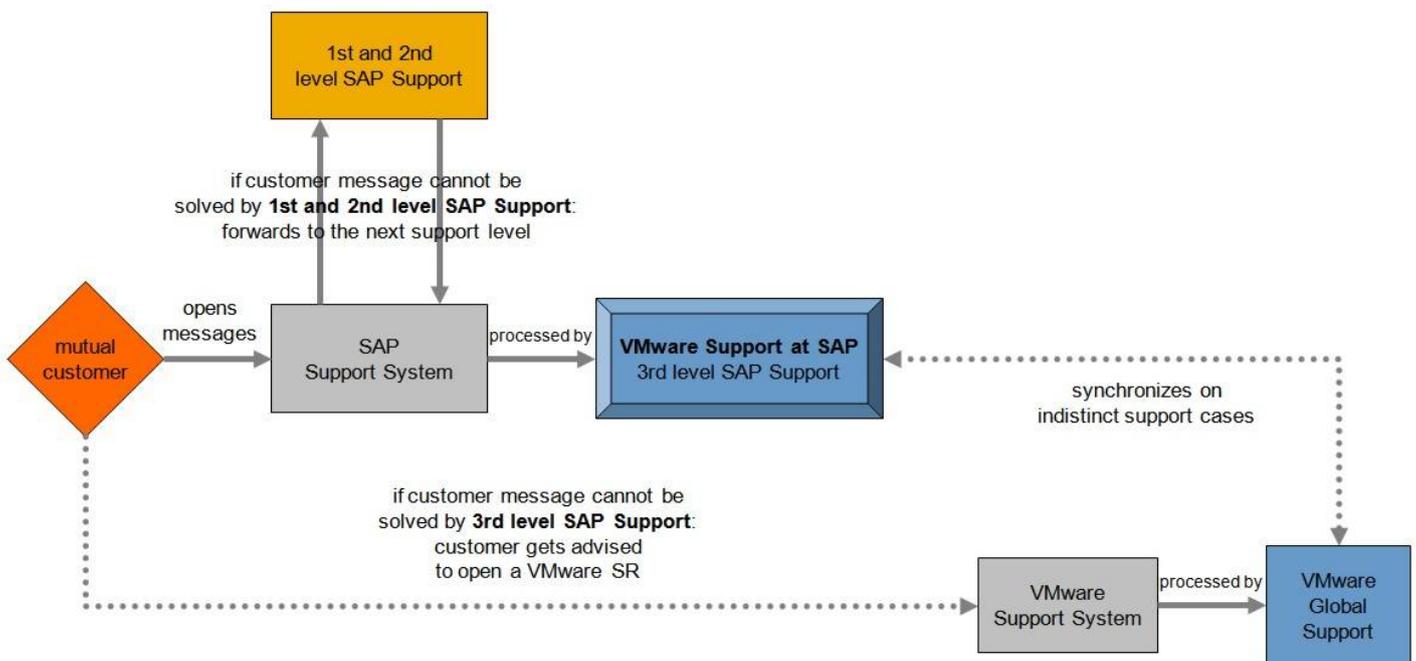


Figure 51: SAP Support Workflow for VMware-related Escalations

For example, if the issue is a Linux kernel panic or an SAP HANA product issue, we recommend that you use the correct support component instead of using the VMware support component because this may delay the support process. If you are uncertain that the issue is related to VMware, open the ticket first at the general SAP HANA support component.

If the issue is related to a VMware product, such as an ESXi driver, then you may either open the ticket via SAP Solution Manager and escalate it to BC-OP-LNX-ESX or ask the VMware customer administrator to open a support ticket directly at VMware.

Open a VMware Support Request Ticket

If there appears to be a VMware product issue or if vSphere is not configured optimally and is causing a bottleneck, file a support request on VMware Customer Connect™ at <http://www.vmware.com/support/contacts/file-sr.html>.

In addition:

- Follow the troubleshooting steps outlined in the VMware knowledge base article, [Troubleshooting ESX/ESXi virtual](#)

[machine performance issues \(2001003\)](#).

- Run the vm-support utility, and then execute the following command at the service console: vm support-s. This command collects the necessary information that VMware uses to help diagnose issues. It is best to run this command when symptoms first occur.

If you want to escalate an issue with your SAP HANA HCI solution, please 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 and selecting the HCI partners SAP support component.

Conclusion

SAP HANA on VMware vSphere/VMware Cloud Foundation provides a cloud operation model for your business-critical enterprise application and data.

For nearly 10 years, virtualizing SAP HANA with vSphere has been supported and does not require any specific considerations for deployment and operation when compared to a natively installed SAP HANA system.

In addition, your SAP HANA environment gains all the virtualization benefits in terms of easier operation, such as SAP HANA database live migration with vMotion or strict resource isolation on a virtual server level, increased security, standardization, better service levels and resource utilization, an easy HA solution via vSphere HA, lower TCO, an easier way to maintain compliance, faster time to value, reduced complexity and dependencies, custom HANA system sizes optimally aligned for your workload and needs, and the mentioned cloud-like operation model.

“I think anything ‘software-defined’ means it’s digital. It means we can automate it, and we can control it, and we can move it much faster.”

—Andrew Henderson, Former CTO, ING Bank

VMware Acknowledgments

Author:

- Erik Rieger, Principal Architect and SAP Global Technical Alliance Manager, VMware SAP Alliance

The following individuals contributed content or helped review this guide:

- Fred Abounader, Staff Performance Engineer, VMware Performance Engineering team
- Louis Barton, Staff Performance Engineer, VMware Performance Engineering team
- Pascal Hanke, Solution Consultant, VMware Professional Services team
- Sathya Krishnaswamy, Staff Performance Engineer, VMware Performance Engineering team
- Sebastian Lenz, Staff Performance Engineer, VMware Performance Engineering team
- Todd Muirhead, Staff Performance Engineer, VMware Performance Engineering team
- Catherine Xu, Senior Manager of Workload Technical Marketing Team, VMware





Copyright © 2022 VMware, Inc. All rights reserved. VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 VMware and the VMware logo are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware products are covered by one or more patents listed at vmware.com/go/patents.

Item No: 1238737aq-tech-wp-sap-hana-vsphr-best-prac-uslet-101 2/22

[1] SAP Note [2937606](#) and [3102813](#)

[2] SAP Note [2104291](#) - FAQ - SAP HANA multitenant database containers, page 2

[3] vSphere 6.5 is out of support and vSphere 6.7 Ux will be out of support October 2022. Refer to the VMware product [lifecycle matrix](#).

[4] The actual number of vCPUs depends on the actual used CPU, such as 224 vCPUs for a 4-socket Intel 28 core Cooper Lake, or 128 vCPUs for a 32-core Ice Lake based host.

[5] The maximum usable CPU sockets depend on the underlying hardware, >4-socket wide VMs require 8-socket host systems.

[6] [ESXi host and vSphere VM maximums](#)/configuration limits

[7] VM RAM > 6 TB requires vSphere 7.0 U2 and hardware version 18 or later. 7.0 supports only 6,128GB.

[8] For more details, see the [SAP HANA on vSphere configuration and sizing guidelines section](#)

[9] For more information, see the [Intel Xeon E7-8890 v4](#) processor specifications.

[10] For more information, see the [Intel Xeon Gold 6258R](#) processor specifications.

[11] For more information, see the [Intel Xeon Platinum 8280L](#) processor specifications.

[12] For more information, see the [Intel Xeon Platinum 8380H](#) or [8380HL](#) processor specifications.

[13] For more information, see the [Intel Xeon Platinum 8380](#) processor specifications.

[14] The Intel memory CPU support is larger what SAP supports in-appliance configurations, which is between 0,5 TB and 2 TB per CPU socket. The exact supported memory configuration is different between CPU models. As of Broadwell CPUs TDI phase 5/workload-based sizing can get used to determine the correct memory configuration. Example, for a 6 TB 4-socket VM or 12 TB 8-socket VM Cascade Lake configurations, an Intel L-type CPU, such as the Intel Cascade Lake 8280L CPU, is required.

[15] vSphere 6.7 U3 and 7 supports up to 64 vCPUs per vCPU socket. This vCPU maxima defines how a VM has to get configured to allow the utilization of all CPU threads (PCPUs). Up to Cooper Lake CPUs this maximum is not a limitation and it is possible to align the VM CPU socket configuration to the number of physical CPU sockets. With Ice Lake 40 core CPUs it is necessary to configure on the VM OS level 4 CPU sockets to leverage HT.

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

[17] The 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 you HW vendor.

[18] Only vSAN based SAP HANA certified HCI solutions are supported. There is no support for generic vSAN solutions for SAP HANA production workloads.

[19] For more information, see the [SAP HANA TDI storage requirements](#).

[20] For the VLAN ID example, final VLAN numbers are up to the network administrator.

[21] For the VLAN ID example, final VLAN numbers are up to the network administrator.

[22] For the VLAN ID example, final VLAN numbers are up to the network administrator.

[23] VMware. "[Increased SAP HANA response time due to network latency with VMXNET3 \(83957\)](#)."

[24] EVC mode needs to be turned on at time of cluster creation. Turning on EVC mode at a later point in time requires all hosts being in maintenance mode.

[25] EMC IT, 02/14 EMC Perspective, H12853.

[26] Due to the asynchronous replication method, it is not recommended to use vSphere Replication for SAP HANA VMs.

[27] VMware. "[Build numbers and versions of VMware ESXi/ESX \(2143832\).](#)"

