

TECHNICAL PAPER: April 2024



# SAP HANA on VMware vSphere Best Practices

By Erik Rieger

Principal SAP Global Technical Alliance Manager & Architect

## Table of contents

<b>Introduction</b> .....	<b>6</b>
Abstract .....	6
Audience .....	6
<b>Solution overview</b> .....	<b>7</b>
SAP HANA on VMware vSphere .....	7
Solution components .....	8
What's new in vSphere 8.0? .....	8
Benefits of SAP HANA on VMware vSphere 8.0.....	9
<b>Software and hardware support for SAP HANA on vSphere</b> .....	<b>9</b>
SAP HANA production support for vSphere and VMware Cloud Foundation .....	9
vSphere version support for Intel CPU platforms .....	10
vSphere maximums.....	11
Release strategy .....	11
Supported vSphere and VMware Cloud Foundation offerings .....	12
Scalability and VM sizes .....	13
<b>Deployment options and considerations</b> .....	<b>14</b>
Reference architecture diagram .....	14
Deployment options and sizes .....	16
Intel sub-NUMA clustering (SNC-2) support and deployment options .....	19
Support for SNC-2 .....	19
SNC-2 supported deployment examples.....	21
Scale-up and scale-out deployment architecture .....	24
<b>Configuration and sizing guidelines</b> .....	<b>26</b>
Sizing compute and memory .....	26
Sizing process .....	26
Translating the SAP sizing to a VM configuration.....	28
Virtual memory configuration example .....	29
CPU calculation examples .....	31

VM CPU (SAPS) configuration example .....	31
VM vCPU configuration examples .....	31
Considerations for the final VM configuration.....	32
Pre-sized examples.....	33
Storage configuration and sizing .....	36
Tools used to verify KPIs .....	37
Storage connection.....	37
vSphere datastores.....	37
File system layout .....	38
Storage capacity calculation .....	39
Hardware and cloud measurement tool for SAP HANA .....	43
Network configuration and sizing .....	44
<b>Workload performance and validation testing .....</b>	<b>51</b>
SAP workload characterization and the impact on network performance.....	52
Performance tests run with SAP S/4HANA and BW/4HANA: workload and validation testing ....	53
Testbed configuration .....	53
Workload information .....	54
Test methodology.....	54
Test results for OLTP workload .....	55
SAP S/4HANA and B/4HANA validation test findings (OLAP) .....	61
BWH L-Class vs. M-class sizings for SAP HANA VMs.....	62
Virtual SAP HANA performance evolution over 5 CPU generations .....	67
<b>Enhanced vMotion Compatibility, vSphere vMotion, and vSphere DRS best practices .....</b>	<b>72</b>
Best practices for migrating SAP HANA VMs in a production environment.....	73
VM migration scenarios.....	74
Managing SAP HANA landscapes using vSphere DRS.....	78
<b>SNC-2 management and operation .....</b>	<b>78</b>
<b>vSphere Cluster Services.....</b>	<b>81</b>
Architecture of vCLS .....	81
vCLS VM details .....	82

- Guidelines for deploying vCLS in a cluster with SAP HANA VMs ..... 83
- Examples of deploying vCLS in a cluster with SAP HANA VMs ..... 84
  - Mixed SAP HANA and other VMs in a vSphere cluster ..... 84
  - Dedicated ESXi hosts for SAP HANA VMs on a vSphere cluster ..... 85
- SAP HANA HCI on a vSphere cluster ..... 86
- High availability best practices for virtualized SAP HANA ..... 88**
  - About vSphere high availability (HA) ..... 90
  - vSphere HA with passthrough network adapters ..... 94
  - SAP HANA system replication with vSphere (local site) ..... 94
  - Disaster recovery ..... 95
    - SAP HSR and vSphere Replication for a remote site ..... 96
  - Backing up and restoring SAP HANA systems virtualized on vSphere ..... 97
- SAP HANA with persistent memory on vSphere 7 ..... 98**
  - Prerequisites and general SAP support limitations for Intel Optane Persistent Memory (PMem) ... 98
  - What is supported? ..... 99
  - Sizing of Optane PMem-enabled SAP HANA VMs ..... 100
  - vSphere HA support for PMem-enabled SAP HANA VMs ..... 105
  - SAP HANA with PMem VM configuration details ..... 106
- Monitoring and verifying an SAP HANA installation ..... 111**
  - SAP notes for monitoring data growth/CPU utilization and verifying configuration ..... 111
  - NUMA tool: VMware NUMA Observer ..... 111
- Performance optimizations for SAP HANA VMs ..... 113**
  - Optimizing the configuration parameter list ..... 113
  - Performance optimization for low-latency SAP HANA VMs ..... 124
    - CPU affinity settings ..... 124
    - Additional performance tuning settings for SAP HANA workloads ..... 125
    - Examples of VMX configurations for SAP HANA VMs ..... 128
    - CPU thread matrix examples ..... 132
- SAP HANA support and process ..... 135**
  - Open an SAP support request ticket ..... 135

Open a VMware support request ticket .....	137
<b>Conclusion .....</b>	<b>137</b>
<b>About the Author .....</b>	<b>138</b>
<b>Acknowledgments .....</b>	<b>138</b>

# Introduction

## Abstract

This guide is the 2024 edition of the best practices and recommendations for SAP HANA on [VMware Cloud Foundation](#), with a focus on vSphere. It describes the best practices and recommendations for configuring, deploying, and optimizing SAP HANA scaled-up and scaled-out deployments, with a focus on vSphere 8.0 running on [third and fourth generation Intel Xeon Scalable processors](#), such as Cascade Lake, Cooper Lake, Ice Lake, and Sapphire Rapids systems.

**Note:** Scale-up SAP HANA vSphere deployments use a larger host and therefore larger VMs to scale from 2-socket up to 8-socket large VMs on a single 8-socket ESXi host. Scale-out means to add more hosts/VMs to a single SAP HANA, typically for a business warehouse instance.

Most of the guidance provided here is the result of ongoing joint testing by VMware and SAP to characterize the performance of SAP HANA running on vSphere.

vSphere 7.0 and vSphere 7.0 update versions on second and third-generation Intel Xeon Scalable processors, such as Broadwell, Skylake, Cascade Lake, Cooper Lake, and Ice Lake vSphere hosts are still supported and are covered in this document as well.

Information on Intel Optane Persistent Memory (PMem) 100 series technology, which is only supported with Cascade Lake and vSphere 7.0 virtualized SAP systems, is provided as well. SAP HANA does not support later CPU generations with PMem, and Intel has announced it is [discontinuing this technology](#).

For earlier vSphere versions or CPU generations, please refer to the [2017](#) or the [2022](#) edition of this guide.

## Audience

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

It assumes you have a basic knowledge of VMware Cloud Foundation concepts and features, SAP HANA, and related SAP products and technologies.

## Solution overview

### SAP HANA on VMware vSphere

[Per SAP](#), 80% of their German and over 51% of their US customers run SAP on-premises or in a private cloud. According to an IDC study mentioned in the same SAP article, up to 68% of SAP workloads will stay on-premises in the United States, especially for large customers. SAP has supported vSphere for production use cases for 10 years now, and most of the on-premises customers (over 70%) use VMware solutions as their private cloud (SDDC) solution for SAP applications like SAP HANA.

By continuing to validate SAP HANA on the latest Intel CPU generations and latest vSphere and Cloud Foundation (VCF) versions, all these customers can continue to seamlessly integrate their IT and SAP operations by leveraging their existing IT processes, know-how, and customer tailored infrastructure that allows full sovereignty over their data and the overall SAP HANA solution, which supports scaling database sizes up to 12 TB (16 TB is planned with 8-socket Sapphire Rapids systems) and scaling out with up to 16 nodes (plus high availability nodes) with up to 48 TB (depending on the host configuration) running on a vSphere virtualized infrastructure.

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

- On-premises security and control
- Locality for consistent and predictable performance
- Regulatory demand for cloud neutrality
- Sovereignty over data and business transactions
- Increased security (using VMware NSX as a zero-trust platform)
- Higher service-level agreements (SLAs) by leveraging vSphere vMotion to migrate live 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 HA
- Built-in multitenancy support via SAP HANA system encapsulation in a VM
- Easier hardware upgrades or migrations due to abstraction of the hardware layer
- Higher hardware utilization rates
- Automation, standardization, and streamlining of IT operation, processes, and tasks
- Public cloud operating model and 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. As mentioned in SAP Notes [2937606](#), [3102813](#), and [3372365](#), and SAP KB [2101244](#), this environment fully supports SAP HANA and related software in production environments, as well as SAP HANA features such as multi-tenant database containers (MDC) and 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 are presented to a VM in a virtual way, which abstracts the underlying hardware resources.

The solution consists of the following components:

- VMware certified server systems as listed in the VMware hardware compatibility list (HCL)
- SAP HANA supported server systems, as listed in the SAP HANA HCL
- SAP HANA certified hyperconverged infrastructure (HCI) solutions, as listed in the SAP HANA HCI HCL
- VMware Cloud Foundation or vSphere Foundation with VMware products like:
  - vSphere and vSAN version 7.0 U2 and later, Sapphire Rapids based systems, vSphere 8.0 and vSAN U2 and later
  - vCenter 7.0 and later, vCenter 8.0 with vSphere 8.0
  - Optional: NSX Networking and Security
  - Optional: Aria Suite for management
- A VMware-specific and SAP-integrated support process

## What's new in vSphere 8.0?

From “[Introducing vSphere 8: The Enterprise Workload Platform](#),” VMware vSphere 8, the enterprise workload platform, brings the benefits of the cloud to on-premises workloads, supercharges performance through DPUs and GPUs, and accelerates innovation with an enterprise-ready integrated Kubernetes runtime. Additionally, there are significant operational benefits, like [VMware Cloud Disaster Recovery add-on services](#), [VMware Aria](#), a unified multi-cloud management solution that provides capacity planning and optimization for your infrastructure with the right size to fit the current and future needs of your SAP workloads, [host different GPU workloads on a single GPU](#), and perform pre-staged ESXi upgrades, to name some. Visit the VMware by Broadcom specific product webpages for details and up-to-date information.



## Benefits of SAP HANA on VMware vSphere 8.0

The latest advancements in vSphere 8.0 bring significant benefits to SAP customers, facilitating the creation of robust, cost-efficient, manageable, and high-performing SAP HANA environments. A notable enhancement for SAP users in vSphere 8.0 is the revamped approach to presenting the physical system/processor topology to virtual machines (VMs).

Previously, vSphere administrators were tasked with manually configuring SAP HANA VMs to align with the underlying host hardware, including non-uniform memory access (NUMA) alignment. This alignment is crucial for optimizing SAP HANA performance. However, prior to vSphere 8.0, achieving optimal performance required the manual configuration of advanced parameters for each VM.

The introduction of the enhanced virtual topology feature in vSphere 8.0 marks a significant improvement. This feature automatically determines optimal coresPerSocket values and virtual L3 cache sizes for VMs, simplifying configuration and enhancing performance.

Furthermore, vSphere 8.0 incorporates intelligent, adaptive NUMA scheduling and memory placement policies, eliminating the need for manual VM balancing across nodes. While manual controls remain available to override default behavior, advanced administrators may still opt for manual NUMA placement for performance-critical SAP HANA VMs.

For comprehensive guidance on optimizing performance, refer to the "[Performance Best Practices for VMware vSphere 8.0](#)" paper. For detailed configuration steps, consult the "[VMware vSphere 8.0 Virtual Topology](#)" paper.

In addition, vSphere 8 Update 1 marks the first release supporting up to 960 logical CPUs per physical host, further expanding scalability.

vSphere 8 Enterprise is available as part of VMware Cloud Foundation and VMware vSphere Foundation, as well as standalone editions such as vSphere 8 Standard and Essentials.

## Software and hardware support for SAP HANA on vSphere

### SAP HANA production support for vSphere and VMware Cloud Foundation

In November 2012, SAP announced initial support for scaling up SAP HANA systems on vSphere 5.1 for non-production environments. Since then, SAP has extended its production-level support for scale-up and scale-out SAP HANA deployment options and multi-VM and half-socket support. vSphere versions 5.x and 6.x are no longer supported and vSphere 7.0 will be unsupported in April 2025. Therefore, you should plan to upgrade to vSphere versions 7.0 or 8.0. Table 1 provides an overview of relevant SAP HANA on vSphere support notes as of March 2024.

Table 1: Relevant SAP notes

Key notes for virtual environments	<a href="#">1492000</a> : General support statement for virtual environments <a href="#">1380654</a> : SAP support in cloud environments <a href="#">2161991</a> : VMware vSphere configuration guidelines
SAP HANA on vSphere	<a href="#">3102813</a> : SAP HANA on VMware vSphere 8 <a href="#">3102813</a> : SAP HANA on VMware vSphere 7.0 U2 with up to 12 TB 448 vCPUs VM sizes <a href="#">2937606</a> : SAP HANA on VMware vSphere 7.0 (incl. U1 and U2) in production <a href="#">2393917</a> : SAP HANA on VMware vSphere 6.5 and 6.7 in production <a href="#">2779240</a> : Workload-based sizing for virtualized environments <a href="#">2718982</a> : SAP HANA on VMware vSphere and vSAN 6.x / 7.x <a href="#">2718982</a> : SAP HANA on VMware vSphere and vSAN 8.x <a href="#">2913410</a> : SAP HANA on VMware vSphere with Persistent Memory <a href="#">2020657</a> : SAP Business One, version for SAP HANA on VMware vSphere in production

## vSphere version support for Intel CPU platforms

Table 2 provides an overview, as of March 2024, of the SAP HANA on vSphere supported and still relevant vSphere versions and CPU platforms. This guide focuses on vSphere 7.0 and 8.0 updates.

Table 2: Supported vSphere versions and CPU platforms as of March 2024

vSphere version	Broadwell	Skylake	Cascade Lake	Cooper Lake	Ice Lake	Sapphire Rapids
vSphere 7.0 U2 and later up to 8-socket wide VM	✓	✓	✓	✓	✓	
PMem Series 100 (vSphere 7)			✓			
vSphere 8.0 U2 and later up to 2-socket wide VM			✓	✓	✓	✓
vSphere 8.0 U2 and later up to 4-socket wide VM			✓	✓		✓
vSphere 8.0 U2 and later up to 8-socket wide VM			✓	✓		

## vSphere maximums

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 relevant maximums per CPU generation as defined by SAP**

vSphere version	SAP HANA maximum virtual memory	Maximum CPUs for SAP HANA deployments	CPU sockets for SAP HANA
vSphere 7 U2 and later	< 12 TB with Cascade and Cooper Lake	<= 448 vCPUs	0.5-, 1-, 2-, 3-, 4-, 5-, 6-, 7-, and 8-socket wide VMs
vSphere 8 U2 and later	< 4 TB for Ice Lake and Sapphire Rapids 2-socket systems	<= 240 vCPUs	0.5-, 1- and 2-socket wide VMs. SPR requires SNC for half-socket VMs
vSphere 8 U2 and later	< 8 TB for Sapphire Rapids	<= 480 vCPUs	1-, 2-, 3-, and 4-socket wide VMs
vSphere 8 U2 and later	< 12 TB with Cascade and Cooper Lake	<= 448 vCPUs	0.5-, 1-, 2-, 3-, 4-, 5-, 6-, 7-, and 8-socket wide VMs

**Note:** These configurations may vary if smaller core count CPUs or different memory configurations are used and require an SAP HANA TDI/workload-based sizing. >8 socket Sapphire Rapids systems are in SAP HANA vSphere validation.

## Release strategy

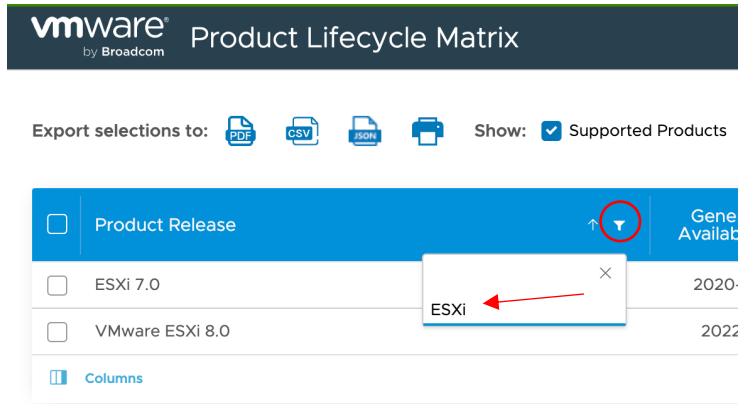
VMware's SAP HANA certification and support strategies for vSphere are to support a single CPU generation or chipset with two versions of the hypervisor and to have a single hypervisor version span two CPU generations/chipsets.

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. You may still use vSphere 7, since this version is in general supported until April 2025. vSphere 8.0 Update 2 is the most recent available version, and we recommend using this version for all new SAP HANA deployments on vSphere. The end of general support for this version is April 2027. For more details, refer to the [VMware Product Lifecycle Matrix](#).

## To directly show the ESXi versions:

- Set a filter for `ESXi` at the **Product Release** table heading, as shown in figure 1.

Figure 1: Select the filter icon next to Product Release and type `ESXi`



## Supported vSphere and VMware Cloud Foundation offerings

SAP HANA on vSphere is supported only with the following editions:

- VMware Cloud Foundation
- vSphere Foundation

For smaller SAP application deployments, vSphere Standard can be used as well. Business One solutions are also supported with vSphere Essentials Plus. Make sure you are familiar with the feature, CPU, and host limitations that come with these products.

### Notes:

- For strategic enterprise customers, we highly recommend VMware Cloud Foundation with its bundled in Select Support offering.
- Per 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.

### References:

- 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, refer to the [SAP notes related to VMware](#) or the [SAP on VMware vSphere community wiki page](#).

## Scalability and VM sizes

SAP HANA on vSphere is supported on 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 be properly sized. 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.

SAP supports VMware virtualized OLTP workloads up to 8 TB for a 4-socket large SAP HANA VM and 12 TB for an 8-socket large VM as standard sizes when selecting the SAP HANA supported top-bin Intel CPUs. Only 50 % of the memory defined for OLTP-type workloads is supported by this standard sizing (for example, 3 TB for 4-socket wide OLAP VMs, or 6 TB for 8-socket wide VMs). You will need to do a workload-based SAP expert sizing if additional memory is needed.

For more details, review SAP Note [2779240: Workload-based sizing for virtualized environments](#). Table 4 shows the current [vSphere maximums](#) per physical ESXi host.

**Table 4: vSphere physical host maximums (extract)**

Relevant Maximums	ESXi 7.0 U2 and later	ESXi 8.0 U1 and later
Logical CPUs per host	896	960
VMs per host	1,024	
Virtual CPUs per host	4,096	
Virtual CPUs per core	32	
RAM per host	24 TB	
NUMA nodes/CPU sockets per host	16 (SAP HANA only 8 CPU socket hosts / HW partitions)	

**Note:** ESXi hosts with up to 8 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-tier solutions. Also note the support limitations when using 8-socket or larger hosts with node controllers (also known as *glued-architecture systems* or *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 (extract)

Maximums	ESXi 7.0 U2 and later	ESXi 8.0 U1 and later
Virtual VM Hardware Version <sup>1</sup>	19	21
Virtual CPUs per VM	Up to 768	
RAM per VM	Up to 24 TB	
CPU sockets per SAP HANA VM	<= 8	
RAM per SAP HANA VM	<= 12 TB	
Virtual SCSI adapters per VM	4	
Virtual NVMe adapters per VM	4	
Virtual disk size	62 TB	
Virtual NICs per VM	10	
Persistent Memory per SAP HANA VM	<= 12 TB	Not supported

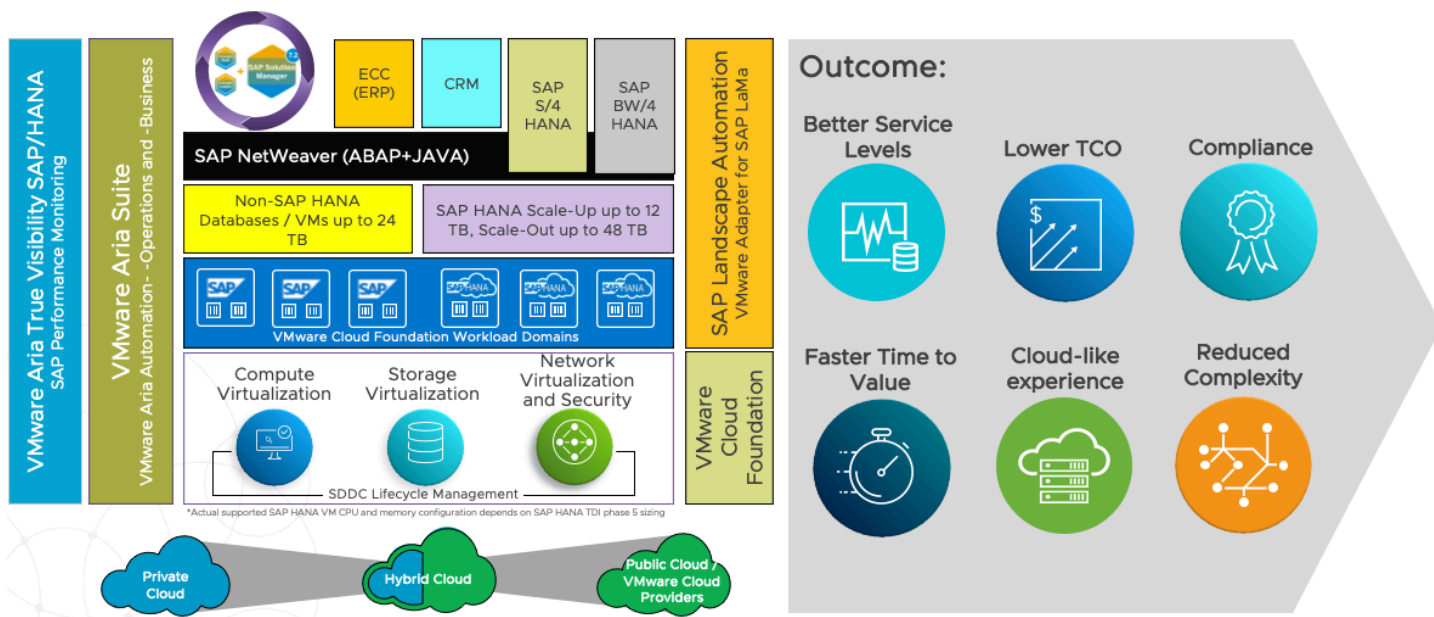
## Deployment options and considerations

### Reference architecture diagram

Figure 2 provides an overview of a typical VMware SDDC for SAP applications. At the center of a VMware SDDC is the [VMware Cloud Foundation](#) with its key products: vSphere, vSAN, and NSX, where SAP applications run in a dedicated VCF workload domain, managed by a VCF management domain.

<sup>1</sup> Review the [Hardware Features Available with Virtual Machine Compatibility Settings](#) web page for a detailed list of the guest hardware capacities. You must use (or upgrade to) hardware version 21 for VMs on (or migrated to) Sapphire Rapids hosts.

Figure 2: VMware SDDC based on VMware Cloud Foundation for SAP applications



As previously explained, virtualized SAP HANA systems are currently supported on VMware vSphere with up to 448 vCPUs and 12 TB RAM per VM on Intel Cascade / Cooper Lake and with up to 480 vCPUs and 8 TB on Intel Sapphire Rapids systems; the vSphere 7.0 U2 and vSphere 8.0 U1 VM guest limits are 768 vCPUs and 24 TB per VM. As of today, only the vSphere host systems shown in table 2 are validated on SAP HANA.

**Note:** The following may limit the maximum number of vCPUs and vRAM available for a VM:

- The selected CPU type
- The virtualized SAP HANA workload type (OLTP or OLAP)
- An SAP HANA use case of very network heavy OLTP workloads with thousands of concurrent users, which may be required to reserve CPU threads to handle such an intensive network load.
- The used VCF features and options, like vSAN or NSX, may reduce the available memory per SAP HANA VM.

Larger SAP HANA systems can leverage SAP HANA extension nodes or be deployed as SAP HANA scale-out configurations. In a scale-out configuration, up to 16 nodes (more upon SAP approval) work together to provide larger memory configurations. A scale-out SAP HANA node's memory size depends on the selected CPU generation and 4- or 8-socket systems with memory sizes per host and up to 6 TB can be selected. Refer to the relevant SAP Notes ([3102813](#) and [3372365](#)) for detailed information on supported host configurations. In addition, the [SAP HANA on vSphere SAP Help Portal](#) page provides an overview of supported configurations and is a good starting point.

An SAP HANA system deployed on a VMware SDDC based on VMware Cloud Foundation can be easily automated and operated by leveraging VMware Aria products. SAP HANA or hardware-specific [management packs](#) allow a top-to-bottom view of a virtualized SAP HANA environment where an AI-

based algorithm allows the operation of SAP HANA in a nearly automated approach, which optimizes performance and availability. A tight integration with [SAP Landscape Management Automation Manager](#) via the [VMware Adapter for SAP Landscape Management](#) helps to cut down operation costs even further by automating work-intensive SAP management and operation tasks.

In addition to SAP HANA, most SAP applications and databases can be virtualized and are fully supported for production workloads. Each SAP workload can run on its own vSphere host or multiple SAP workloads / VMs can run consolidated on a single vSphere host.

Virtualizing all aspects of an SAP data center is the best way to plan for quick growth and easily move to a cloud-based infrastructure. 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.

## Deployment options and sizes

You can install SAP HANA on SAP-supported vSphere versions and validated CPU generations as scale-up and scale-out deployments on a single large VM or on multiple SAP HANA VMs on a single ESXi host. You may use only 2-, 4-, and 8-CPU socket VMware and SAP-[supported or certified systems](#) 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). Running SAP HANA VMs next to non-SAP HANA VMs, such as vSphere management VMs or SAP application servers, is also supported 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 shows the supported host configurations for SAP HANA on vSphere and the deployment options for both single-tenant and multi-tenant SAP HANA instances on vSphere. It also gives some guidance about the standard memory sizes that SAP and VMware support based on the current SAP-defined memory limits for the top-bin Intel CPUs listed as certified appliance configurations.

The examples show the SAP-selected, top-bin CPUs for SAP HANA workloads with the maximum core count and memory support available. Refer to the [SAP HANA Certified and Supported SAP HANA Hardware directory](#) for details. 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 is done.

**Note:** The maximum available memory for a virtualized SAP HANA system is limited to the maximum memory tested with vSphere, which is currently 12 TB per single scale-up VM, or 48 TB with scale-out deployments.



**Legend:**




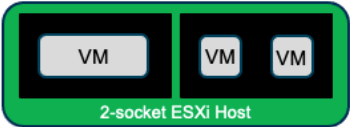
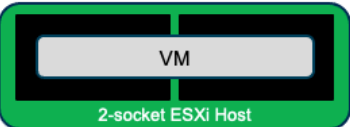
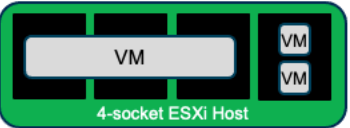


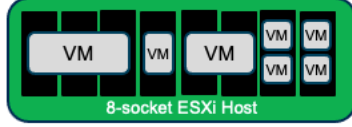
-  <=8-socket ESXi Host
-  1 physical CPU socket / NUMA Node
-  SAP HANA VM

Table 6: Overview of SAP HANA on vSphere deployment options and possible VM sizes

<p><b>2-socket host examples</b></p> <p>Example 1: Single and Half-Socket HANA VMs</p>  <p>Example 2: 2-Socket Wide HANA VM</p> 	<p>vSphere 7.0 supported CPU generations:</p> <ul style="list-style-type: none"> <li>• Broadwell</li> <li>• Skylake, Cascade Lake, and Cooper Lake</li> <li>• Ice Lake</li> </ul> <p>vSphere 8.0 supported CPU generations:</p> <ul style="list-style-type: none"> <li>• Cascade Lake, and Cooper Lake</li> <li>• Ice Lake</li> <li>• Sapphire Rapids (half-socket support only with SNC)</li> </ul> <p>VM sizes:</p> <p>0.5, 1, 2-socket wide VMs on 2-socket hosts with a minimum of 8 vCPUs and 128 GB vRAM and up to 240 vCPUs and &gt;4 TB (SAP HANA standard memory size) with Ice Lake and Sapphire Rapids, 3 TB (standard) or larger memory sized based on a workload-based sizing possible with Cascade or Cooper Lake.</p> <p>Recommended network configuration per host (dual-NIC port configuration):</p> <ul style="list-style-type: none"> <li>• 1 GbE IPMI, 1 GbE ESXi management network</li> <li>• &gt;= 10 GbE vMotion / HA</li> <li>• &gt;= 10 GbE HANA to App Server network</li> <li>• Plus optional networks like those for backup or replication</li> </ul>
<p><b>4-socket host example</b></p> <p>Example 3: Wide and Half-Socket HANA VM on one host</p> 	<p>vSphere 7.0 supported CPU generations:</p> <ul style="list-style-type: none"> <li>• Broadwell</li> <li>• Skylake, Cascade Lake, and Cooper Lake</li> </ul> <p>vSphere 8.0 supported CPU generations:</p> <ul style="list-style-type: none"> <li>• Cascade Lake and Cooper Lake</li> <li>• Sapphire Rapids (no-half-socket support)</li> </ul> <p>VM sizes:</p> <p>0.5, 1, 2, 3 and 4-socket wide VMs on 4-socket hosts with minimal 8 vCPUs and 128 GB vRAM and up to 224 vCPUs and &gt;6 TB (SAP HANA standard memory size) with Cascade and Cooper Lake. Larger memory sizes based on a workload-based sizing possible.</p>

	<p>Recommended network configuration per host (dual-NIC port configuration):</p> <ul style="list-style-type: none"> <li>• 1 GbE IPMI, 1 GbE ESXi management network</li> <li>• &gt;= 25 GbE vMotion / HA</li> <li>• &gt;= 10 GbE HANA to App Server network</li> <li>• Plus optional networks like those for backup or replication</li> </ul>
<p><b>8-socket host examples</b></p> <p><b>Example 4: 1 HANA VM across 8 NUMA nodes</b></p>  <p><b>Example 5: 2 HANA VMs across 4 NUMA nodes</b></p>  <p><b>Example 6: Wide VMs, single and half-socket VMs on one host</b></p> 	<p>vSphere 7.0 supported CPU generations:</p> <ul style="list-style-type: none"> <li>• Broadwell</li> <li>• Skylake, Cascade Lake, and Cooper Lake</li> </ul> <p>vSphere 8.0 supported CPU generations:</p> <ul style="list-style-type: none"> <li>• Cascade Lake and Cooper Lake</li> <li>• Sapphire Rapids is not currently available on 8-socket hosts, but is undergoing the validation process</li> </ul> <p>VM sizes:</p> <p>0.5, 1, 2, 3, 4, 5, 6, 7 and 8-socket wide VMs on 8-socket hosts with minimal 8 vCPUs and 128 GB vRAM and up to 448 vCPUs and &gt;12 TB (SAP HANA standard memory size) with Cascade and Cooper Lake.</p> <p><b>Note:</b> 6-socket ESXi host configurations are supported as well.</p> <p>Recommended network configuration per host (dual-NIC port configuration):</p> <ul style="list-style-type: none"> <li>• 1 GbE IPMI, 1 GbE ESXi management network</li> <li>• &gt;= 25 GbE vMotion / HA</li> <li>• &gt;= 10 GbE HANA to App Server network</li> <li>• Plus optional networks like those for backup or replication</li> </ul>

**Note:** Starting with the [Intel Sapphire Rapids](#) platform, SAP HANA "half-socket VM" support is only available with Intel Sub-NUMA Clustering (SNC-2) enabled on vSphere 8 or later ESXi **2-socket** hosts.

SAP HANA "half-socket VM" support on non-SCN enabled 2-socket Sapphire Rapids hosts or >2-socket Sapphire Rapids systems are not yet supported.

NUMA node/CPU sharing half-socket VMs are supported up to Ice Lake CPUs with vSphere 7 and where applicable with vSphere 8, as described in table 6.

# Intel sub-NUMA clustering (SNC-2) support and deployment options

## Support for SNC-2

Support for SNC-2 in SAP HANA takes into account the increasing density of processors, memory controllers, processor interconnects, and supporting infrastructure within a single chip, as the size of the CPU decreases. Increasing the number of CPU cores and, consequently, CPU performance is advantageous for users. However, this increased density [is associated with a longer data transfer time between different parts of the CPU chip](#).

The process of accessing the memory within a CPU is facilitated by a uniform memory access (UMA) domain, which offers a unified and contiguous address space that is interleaved among all the memory controllers. UMA lacks a mechanism for optimizing the flow of data from the closest available resources. Processor affinity is typically employed to specify the processor(s) that a particular software thread utilizes. Consequently, in a Uniform Memory Access (UMA) system, all cores possess equal access to both the last level cache and memory. In this scenario, there is a [possibility of a processor accessing a memory controller or a portion of the last level cache located on the opposite side of the CPU chip compared to the nearest](#).

It takes longer to access and move data within a CPU chip, and the memory subsystem is being used more than before. This is clear when you compare the latest Sapphire Rapids CPUs to Cascade Lake CPUs. In this case, the number of CPU cores and memory controllers doubled (from 2 to 4 memory controllers) or slightly more than doubled (from 28 to 60 cores), but the number of memory channels only went up by 33.33%. The Sapphire Rapids platform has a faster memory bandwidth (MT/s), but this means that more CPU threads need to share a memory channel. DDR5 memory modules also have a longer CAS latency than DDR4 modules. [CAS latency](#) is a timing parameter that measures the delay between a memory controller sending a request for data and the memory module responding with the requested data. A longer CAS latency leads to bigger differences when memory-sensitive applications are used; for example, when two SAP HANA VMs share a NUMA node or CPU socket. When compared to a Cascade Lake CPU, a Sapphire Rapids CPU has much better overall performance. This includes half-socket deployments.

However, when running several SAP HANA VMs on a socket, there are more fluctuations. Because of these fluctuations, SAP does not provide support for half-socket deployments for Sapphire Rapids CPUs. Table 7 illustrates this and shows that the [theoretical available memory channels per CPU core of a Sapphire Rapids CPU is lower as with older CPU generations](#).

Table 7 does not reflect the change in increased memory speed and shows only that the available memory channels are higher utilized because of the increased CPU core count per chip and the higher CAS latencies when first accessing data stored in memory.

Table 7: Overview of CPU memory channels and CAS latency of selected SAP HANA relevant CPUs

CPU	Cores	# of Memory Controllers	Max # of Memory Channels	Core/Channel Ratio	DDR Type	Percentage Change	Typical CAS Latency	Percentage Change
Intel Xeon Platinum 8280 Processor "Cascade Lake"	28	2	6	4,667	4	0%	21	0%
Intel Xeon Platinum 8380 Processor "Ice Lake"	40	4	8	5,000	4	7%	22	5%
Intel Xeon Platinum 8490H Processor "Sapphire Rapids"	60	4	8	7,500	5	61%	40	90%

The [number of memory channels](#) refers to the bandwidth operation for the real-world application. (After following the hyperlink, click on (?) at **Max # of Memory Channels**.)

To decrease the latency of data movements across the CPU, Intel introduced sub-NUMA clusters (SNC). Sapphire Rapids CPUs support SNC-2 and SNC-4, **but SAP HANA vSphere VMs support only SNC-2 on 2-socket host configurations.**

In a two-cluster SNC (SNC-2), two localization domains exist within a processor. Each domain has addresses mapped from the local memory controller and local Last Level Cache (LLC) slices. When processors are in the local domain, they will use the local memory controller and local LLC slices. This means that LLC and memory accesses in the local domain will have lower latency than accesses to locations outside the same SNC domain.

SNC has a unique location for every address in the LLC, and it is never duplicated within the LLC banks. Localization of addresses within the LLC for each SNC domain applies only to addresses mapped to the memory controllers in the same socket. All addresses mapped to memory on remote sockets are uniformly distributed across all LLC banks independent of the SNC mode. Therefore, even in the SNC mode, the entire LLC capacity on the socket is available to each core, and the LLC capacity reported through the CPUID is not affected by the SNC mode.

Refer to figure 2 – "Block Diagram Representing Domains Of sub-NUMA With Two Clusters" and figure 3 – "Block Diagram Representing Domains Of sub-NUMA With Four Clusters" on the [Intel product webpage](#), which shows the block diagrams of a sub-NUMA node enabled CPU. Figure 2 represents a two-cluster configuration, which is the SAP-supported SAP HANA on vSphere configuration that consists of SNC Domains 0 and 1, in addition to their associated cores, LLCs, and memory controllers. Each SNC-2 domain contains half of the processors on the socket, half of the LLC banks, and half of the memory controllers with its associated DDR channels, which is limiting a VM to exactly these resources.

According to Intel, the affinity of cores, LLC, and memory within a domain are expressed via software using the NUMA affinity parameters in the operating system.

**Note:** SNC is enabled at the BIOS level and requires that the memory be symmetrically populated.

When SNC-2 is used for SAP vSphere VMs on 2-socket Sapphire Rapids, then we advise that either all hosts in a Sapphire Rapids vSphere cluster are SNC-2 enabled or that a VM host rule is applied to ensure that the SNC-2 enabled VMs do not get migrated or started on non-SNC enabled hosts.

### SNC-2 supported deployment examples

SAP HANA and SAP Application servers, as of today, are supported by SAP and VMware with SNC-2 only on 2-socket Sapphire Rapids servers with vSphere 8 and can be leveraged as described in the table, below.

**Important:** Any other configurations are not SAP nor VMware supported for SAP HANA and SAP Application deployments. Non-SNC-2 configurations are supported only as full-socket VM configurations as described in table 8.

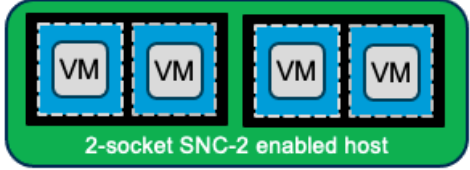
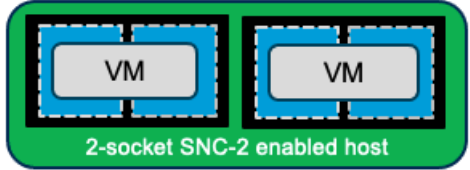
The SAP HANA reference architecture for 2-socket Intel Xeon Platinum [8490H Processor](#) Sapphire Rapids-based systems defines a 60-core Sapphire Rapids CPU to support up to 2 TB of memory (4 TB in total per 2-socket Sapphire Rapids host) according to SAP HANA appliance sizing. Different memory sizes and lower core count CPUs, like a 32 or 48 core Sapphire Rapids CPU can also be used but require SAP application/workload sizing.

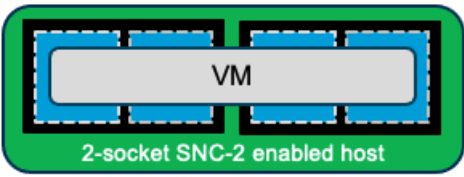
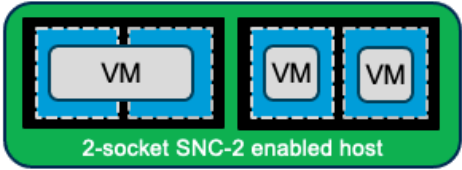
**Note:** The following configuration options (table 8) for running SAP HANA on vSphere are not supported on bare metal servers enabled for SNC.

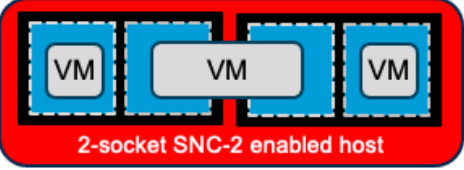
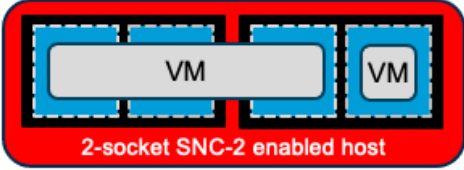
Legend:

- 2-socket ESXi Sapphire Rapids host or later
- 1-CPU socket / NUMA Node
- Sub NUMA Node (SNC-2)
- VM

Table 8: SAP HANA supported SNC-2 configurations

<p><b>Option 1: 2 VMs per NUMA node, 4 VMs in total per ESXi host</b></p> <p>VM size is based on the SAP HANA reference configuration for 2-socket Sapphire Rapids systems:</p> <ul style="list-style-type: none"> <li>4 VMs with &lt;=60 vCPUs, &lt;1 TB vRAM, and 1 vSocket</li> </ul> 	<p>Option 1 shows 4 SAP HANA VMs running on an SNC-2 enabled host. Each VM operates on a sub-NUMA node that is fully isolated, providing exclusive access to all available CPU resources.</p> <p>This configuration replaces the NUMA node sharing half-socket configuration seen in older CPU generations running SAP HANA VMs shared on a CPU. The benefit lies in the dedicated CPU resources separated by SNC, eliminating resource sharing between VMs as seen without SNC. This full isolation and shorter distance from a CPU thread to memory offer the best possible performance per VM when 4 VMs run on a host.</p> <p>Additionally, SAP supports, with SNC, the deployment of non-SAP HANA VMs on another sub-NUMA node on the same physical CPU socket. This was not supported in the past, allowing VMs for infrastructure management (such as vCLS VMs) or SAP application servers to be co-deployed with SAP HANA VMs on the same CPU socket.</p> <p><b>Attention:</b> The drawback of an SNC-divided CPU is that VMs sharing a CPU socket, as in the past, cannot utilize idle resources not used by a co-deployed VM. This may result in lower peak performance due to the split of CPU resources compared to a non-SNC configuration. However, in an SAP environment with its strict sizing rules, this drawback is not significant, and the benefits of lower latencies and more predictable performance outweigh it.</p>
<p><b>Option 2: 1 VM per NUMA node, 2 VMs in total per ESXi host</b></p> <p>VM size is based on the SAP HANA reference configuration for 2-socket Sapphire Rapids systems:</p> <ul style="list-style-type: none"> <li>2 VMs with &lt;=120 vCPUs, &lt;2 TB vRAM and 2 vSockets</li> </ul> 	<p>Option 2 shows 2 SAP HANA VMs running on an SNC-2 enabled host. SAP supports expanding an SAP HANA VM across two sub-NUMA nodes to utilize a full CPU socket. This is analogous to when a single CPU socket VM spans two CPU sockets on an older CPU generation. SAP HANA is NUMA-aware and optimizes memory access based on memory latencies.</p> <p>This allows a user to scale up an SAP HANA "half-socket" VM to occupy a full physical CPU socket, without the need to migrate a VM to a non-SNC host, which remains an option and may be the preferred solution due to the SNC-based CPU limitations.</p> <p><b>Attention:</b> Be aware that leveraging two SNC-2 NUMA nodes to allocate more memory or CPU resources to this VM may result in lower performance compared to a single CPU socket VM running on a non-SNC-2 enabled host. If you observe this performance, then this VM must be migrated to a non-SNC configured host, because SAP (not VMware) can address this issue in their software.</p> <p>In approximately 3% of all tested cases, we observed a negative impact related to SNC-2 on a 2-sub-NUMA Node-wide VM compared to a 1- NUMA Node SAP HANA VM (without SNC-2) when transactions had to leverage more CPU threads or memory that crossed the sub-NUMA</p>

	<p>node boundary. The impact we measured ranged between 5% and 23% for runtime or throughput. The positive impact, when a task stayed inside the sub-NUMA node, was not measured.</p>
<p><b>Option 3: 1 VM across two NUMA nodes / physical CPU sockets</b></p> <p>VM size is based on the SAP HANA reference configuration for 2-socket Sapphire Rapids systems:</p> <ul style="list-style-type: none"> <li>1 VM with &lt;=240 vCPUs, &lt;4 TB vRAM and 4 vSocket</li> </ul> 	<p>Option 3 shows a 4-SNC 2-node wide single SAP HANA VM running on an SNC-2 enabled host.</p> <p>On an SNC-2 enabled 2-socket SAP HANA host, SAP supports spanning an SAP HANA VM across 4-SNC 2 nodes. In this configuration, SAP HANA detects a '4 NUMA node' server and attempts to optimize memory latency based on NUMA locality.</p> <p>This capability allows a user to scale up an SAP HANA VM to the maximum size of a 2-socket Sapphire Rapids server, as specified in the SAP HANA reference architecture, which can accommodate up to 240 logical CPUs with &lt;4 TB of memory per host. When an SAP HANA VM requires all CPU resources of a host, we recommend you migrate this VM to a non-SNC host.</p> <p><b>Attention:</b> Be aware that leveraging four SNC-2 sub-NUMA nodes may result in lower performance compared to a two NUMA node wide VM running on a non-SNC-2 enabled host. If you observe this performance degradation, then this VM must be migrated to a non-SNC configured host, because SAP (not VMware) can address this issue in their software.</p> <p>In approximately 4% of all tested cases, we observed a negative impact related to SNC-2 on a 4 sub-NUMA node wide VM compared to a 2 NUMA node wide SAP HANA VM (without SNC-2) when transactions had to leverage more CPU threads or memory that crossed the sub-NUMA node boundary. The impact we measured ranged between 5% and 14% for runtime or throughput. The positive impact, when a task stayed inside the sub-NUMA node, was not measured.</p>
<p><b>Option 4: 1 VM across two NUMA nodes / physical CPU sockets</b></p> <p>VM size is based on the SAP HANA reference configuration for 2-socket SPR systems:</p> <ul style="list-style-type: none"> <li>1 VM with &lt;=120 vCPUs, &lt;2 TB vRAM and 2 vSockets</li> <li>2 VM with &lt;=60 vCPUs, &lt;1 TB vRAM and 1 vSocket</li> </ul> 	<p>Option 4 shows a supported configuration with 3 SAP HANA VMs, two single SNC-2 wide VMs and one VM spanning two SNC nodes.</p> <p>This option is the last supported configuration.</p> <p><b>Attention:</b> Be aware that leveraging four SNC-2 sub-NUMA nodes may result in lower performance compared to a two NUMA node wide VM running on a non-SNC-2 enabled host. If you observe this performance degradation, then this VM must be migrated to a non-SNC configured host, because SAP (not VMware) can address this issue in their software.</p>

<p>Option 5: No support for NUMA node / sockets crossing SNC-2 VMs</p>  <p>2-socket SNC-2 enabled host</p>	<p>Not supported for SAP HANA VM deployments.</p>
<p>Option 6: No support for NUMA node / sockets crossing SNC-2 VMs</p>  <p>2-socket SNC-2 enabled host</p>	<p>Not supported for SAP HANA VM deployments.</p>

**Important:** Due to the lower memory bandwidth associated with an SNC-NUMA node, we recommend to primarily use SNC-2 for half-socket VMs. In the event of performance issues with SAP HANA VMs running on SNC-2 enabled hosts, these SAP VMs should be migrated to a non-SNC enabled host to resolve SNC-related performance issues. All VMs that leverage SNC-2 require the VMX advanced parameter `sched.nodeX.affinity="Y"` to prevent unwanted NUMA node migrations.

In case of performance issues with SAP HANA VMs running on an SNC-2 enabled host, an SAP HANA VM shall be migrated on a full socket, non-SNC enabled host as part of the troubleshooting.

**Recommended:** Build a vSphere cluster based on the same CPU model and generation. This allows for easy VM migration.

When you deploy SNC-2 enabled Sapphire Rapids systems, you must build a dedicated 2-socket SNC-2 enabled Sapphire Rapids cluster to ensure seamless VM migration. We don't recommend adding non-SNC enabled hosts to this cluster.

Review the "[SNC-2 management and operation](#)" section to learn how to best build and operate a vSphere SAP HANA cluster that has SNC-2 enabled hosts and VMs, along with non-SNC-2 hosts and VMs.

## Scale-up and scale-out deployment architecture

Figures 3 and 4 describe typical scale-up and scale-out SAP HANA on vSphere architectures.

The storage must 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 be defined regarding the SAP HANA size; for example, vMotioning a 2 TB SAP HANA VM or a 12 TB SAP HANA VM will take significantly more time to be migrated or will need a higher bandwidth network to minimize the migration time. Latency should be as low as possible to support the transactions of heavy or sensitive workloads and use cases. For SAP applications, the [average response time](#) for dialog transactions (online transactions) should be below one second.

Figure 3 shows scale-up SAP HANA systems (single SAP HANA VMs).

Figure 3: High-level architecture of a scale-up SAP HANA deployment on vSphere

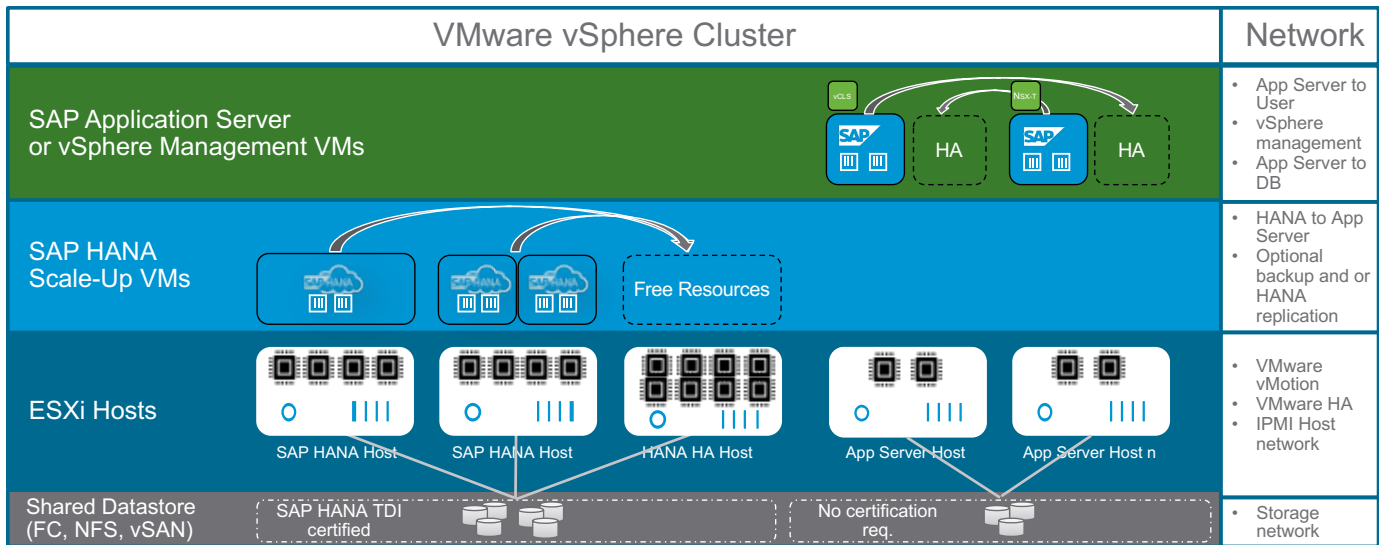
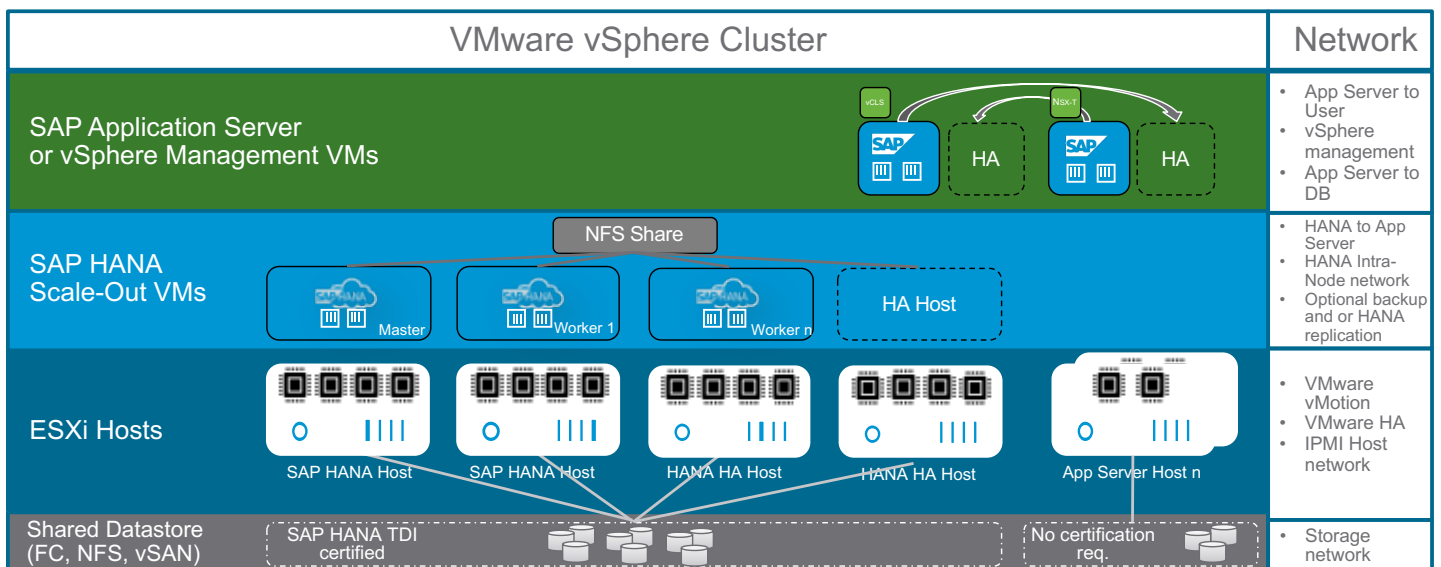


Figure 4 shows a scale-out example where several SAP HANA VMs work together to build one large SAP HANA database instance.

Figure 4: High-level architecture of a scale-out SAP HANA deployment on ESXi 4- or 8-socket host systems



## Configuration and sizing guidelines

You must select the correct components and configuration to achieve the performance and reliability requirements for SAP HANA. Which and how many SAP HANA VMs can be supported depends on the server configuration (RAM and CPU), the network configuration, and the storage configuration.

**Note:** It is possible to consolidate a certain number of SAP HANA VMs on an ESXi host with a given RAM and CPU configuration; however, the network and storage configurations must be able to support these SAP HANA VMs. Otherwise, a network or storage bottleneck could impact the performance of all running SAP HANA VMs on the host.

### Sizing compute and memory

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

You perform VMware virtual SAP HANA sizing in much the same way as you do for physically deployed SAP HANA systems. The major difference is that an SAP HANA workload must fit into the compute and RAM maximums of a VM. You also need to consider the costs of virtualization (RAM and CPU of the ESXi host) when planning an SAP HANA deployment.

If an SAP HANA system exceeds the available resources on virtual deployments, the VM can be moved to a new host with more memory or higher performing CPUs. After this migration to the new host, the VM must be shut down and the VM configuration must be 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 you can use scale-out deployments or SAP HANA extension nodes.


**Note:** The current VMware SAP HANA VM maximums are 448 vCPUs and 12 TB of RAM with Intel Cascade and Cooper Lake 8-socket systems and 480 vCPUs and 8 TB of RAM with Intel Sapphire Rapids 4-socket systems. SAP HANA systems that fit into these maximums can be virtualized as a single scale-up system. You may be able to deploy larger scale-out systems as well.

### Sizing process

As noted, sizing a virtual SAP HANA system is just like sizing a physical SAP HANA system plus the virtualization costs of CPU and RAM on the ESXi host). Figure 5 depicts the SAP HANA sizing process.



Figure 5: The SAP HANA sizing process



An SAP HANA sizing results in the following required resources:

- Compute - per the SAP Application Performance Standard (SAPS)
- Memory
- Storage

The SAP sizing tools do not cover network sizing. You can find the SAP-provided network sizing information, which only focuses on throughput, in the [SAP HANA network requirements white paper](#). Network latency is only expressed as a general guideline—you must set your own goal. In [SAP sales and distribution \(SD\) benchmarks](#), a time below 1,000 ms 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 when we define the network infrastructure for a virtualized SAP HANA environment. SAP also provides a tool, ABAPMETER, you can use 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 SAP workload you use. 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.

The storage capacity requirements for virtual or physical SAP HANA systems are identical. However, the CPU resource and physical memory requirements of a virtualized SAP HANA system are slightly higher than when deployed on bare metal servers; the virtualized requirements also include the virtual

CPU and memory of the VM. We can express the CPU requirements with a fixed factor like 3-10%, but the memory needs are harder to determine.

Table 9 shows the estimated memory requirements of an ESXi host running SAP HANA workloads on different server configurations. Unfortunately, you cannot determine the memory cost of ESXi upfront. They are influenced by the physical server configuration—for example: the number of CPUs, size of the memory, and installed cards like HBAs or NICs—and the vSphere features, like NSX or vSAN in use.

**Table 9: Estimated ESXi host RAM needs for SAP HANA VMs based on the number of sockets and memory installed on the ESXi host**

Physical host CPU sockets	Estimated ESXi host memory requirement (guidelines)	Example memory requirement for ESXi host based on installed memory	
2	64–256 GB (default 128 GB)	2 TB	128 GB
		3 TB	192 GB
4	128–384 GB (default 256 GB)	4 TB	256 GB
		6 TB	384 GB
8	256–768 GB (default 512 GB)	8 TB	512 GB
		12 TB	768 GB

**Note:** The ESXi host memory consumption as described here can be from 64 up to 768 GB. But you can't determine the actual host memory consumption until all the VMs are configured with memory reservations and started on the host. The last VM started might fail to start if you selected the wrong host memory reservation. In this case, you should reserve a lower amount of memory per VM to ensure there is enough host memory for all the VMs running on that host.

## Translating the SAP sizing to a VM configuration

Let's use the following example to translate an SAP HANA sizing of a 1,400 GB SAP HANA database system that needs 80,000 SAPS on compute, according to the [SAP Quick Sizer](#) results.

The following formula calculates the available memory for SAP HANA VMs running on an ESXi host:

**Equation 1** Total available memory for SAP HANA VMs = Total host memory – ESXi memory need

For this example, let's assume that a 4-socket server with 6 TB total RAM and a 2-socket server with 2 TB RAM are available for selection. We use the default ESXi memory requirements:

- For the 4-socket ESXi host with a total of 6 TB RAM, that's around 384 GB of ESXi memory needed, see table 9 for details.
- For the 2-socket ESXi host with a total of 2 TB RAM, that's around 128 of ESXi memory needed, see table 9 for details.

## Virtual memory configuration example

SAP HANA example system memory sizing report result:

- 1400 GB RAM SAP HANA example system memory requirement

Available host server systems:

- 4-CPU socket 28-core Cooper Lake ESXi host with 6 TB physical memory
- 2-CPU socket 60-core Sapphire Rapids ESXi host with 2 TB physical memory

Virtual RAM calculation for the SAP HANA VM:

- **5760 GB** (6144 GB – 384 GB RAM) / **4 CPU sockets** = 1440 GB per CPU socket > 1400 GB sized HANA RAM → SAP HANA VM configuration needs 1 CPU socket
- **1920 GB** (2048 GB – 128 GB) / **2 CPU sockets** = 960 GB per CPU socket < 1400 GB sized HANA RAM → SAP HANA VM configuration needs 2 CPU sockets

Temporary 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 be used because one socket has only 960 GB available.
- In both cases, you should select all available logical CPUs per CPU socket, like in the Cooper Lake case 56 vCPUs (28 cores plus hyperthreading).

**Note:** When fine-tuning the VM memory configuration, determining the maximum available memory and consequently the memory reservation for a VM cannot occur until the VM is actually created and started once. As previously noted, the available memory is contingent upon the hardware configuration of the ESXi host, as well as the enabled and utilized features of ESXi.

After the memory configuration for the SAP HANA VM, you must verify if the available SAPS capacity is sufficient for your planned workload.

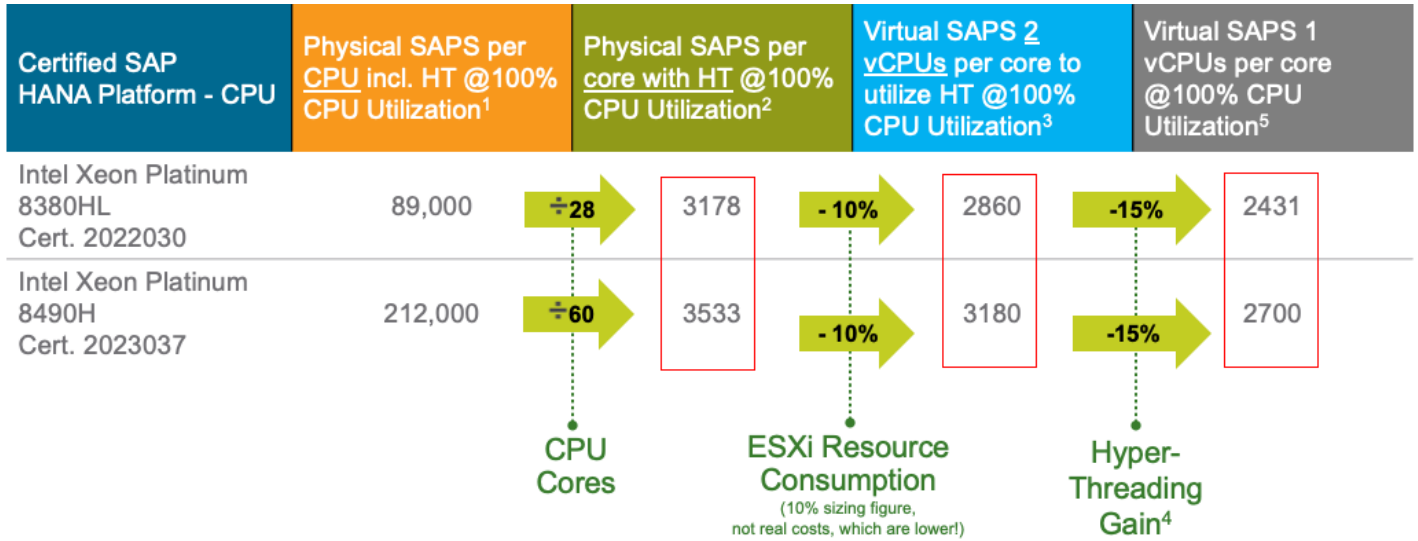
SAP and SAP partners measure the [SAPS](#) capacity of a server and its CPU gets measured by running specific [SAP defined benchmarks](#) and performance tests, such as the SAP SD benchmark, which is used for all SAP NetWeaver and BW/4HANA applications. The test results of a public benchmark can be published and used for SAP sizing lectures.

After you know the SAPS resource need of an SAP application, you can translate the sized SAPS, just like the memory requirement, to a VM configuration.

**Note:** Typically, a VMware administrator gets the CPU / SAPS requirement from the SAP administrator. If not available, the SAP-defined standard sizes (SAP HANA appliance configurations) with its fixed CPU type and maximum memory configuration per CPU socket can be used.

Figure 6, which is from a published SD benchmark, shows 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 6: Physical SAPS to virtual SAPS example conversion



First, look up the SAPS results of a benchmark of the CPU you want to use. Then, divide this result by the number of cores of the selected CPU. In the figure 6 example, the Cooper Lake (8380HL) CPU has 28 cores, and the Sapphire Rapids (8490H) CPU has 60 cores. Use the number of cores as the divisor. This provides the SAPS capacity of a hyperthreaded physical CPU core (you can use two CPU threads or two logical CPUs).

To estimate the virtual SAPS capacity of these two CPU threads, subtract the ESXi CPU resource needs, which are between 3%–10% for OLTP or OLAP workloads. In figure 6, to make the sizing easier, we use 10% 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, subtract the hyperthreading gain, which could be as little as 1% for very low-utilized servers or more than 30% for very high-loaded systems. For the sizing example in figure 6, we assume a 15% hyperthreading gain. Removing this 15% from the 2-vCPU result provides the SAPS capacity of a single vCPU (virtual SAPS 1 vCPU) that runs exclusively on a CPU core.

**Recommended:** Use hyperthreading (this is the default). Make sure that `numa.vcpu.preferHT=TRUE` (per VM setting) is set for the SAP HANA VM 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.

## CPU calculation examples

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 (SAPS) configuration example*

Assumed/sized SAP HANA system:

- 1400 GB RAM SAP HANA system memory
- 80,000 SAPS

Available host servers:

- 4-socket 28-core Cooper Lake (8380HL) host with 6 TB physical memory, 89,000 pSAPS (356,000 SAPS in total) with 224 logical CPUs
- 2-socket 60-core Sapphire Rapids (8490H) host with 2 TB physical memory, 212,000 pSAPS (424,000 SAPS in total) with 240 logical CPUs

### *VM vCPU configuration examples*

4-socket 28-core Cooper Lake (8380HL) host:

- HANA CPU requirement as defined by sizing: 80,000 SAPS
- $356,000 / 4 = 89,000$  SAPS / 28 cores = 3178 pSAPS per core  
vSAPS per 2 vCPUs (including HT) = 2860 SAPS (3,178 SAPS – 10%)  
vSAPS per vCPU (without HT) = 2,431 SAPS (2,860 SAPS – 15%)
- VM without HT: #vCPUs =  $80,000 \text{ SAPS} / 2431 \text{ SAPS} = 32,91$  vCPUs, rounded up 33 cores / vCPUs
- VM with HT enabled: #vCPUs =  $80,000 \text{ SAPS} / 2860 \text{ SAPS} \times 2$  (threads per core) = 55,94 = 56 threads / vCPUs

VM vSocket calculation:

- $33 / 28$  (CPU cores) = 1.18 or  $56$  (threads) /  $56$  (CPU threads) = 1
- 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. This example shows why it is important to leverage hyperthreads.

2-socket 60-core Sapphire Rapids (8490H) host:

- HANA CPU requirement as defined by sizing: 80,000 SAPS
- $424k / 2 = 212k \text{ SAPS} / 60 \text{ cores} = 3533 \text{ SAPS core}$   
vSAPS per 2 vCPUs (including HT) = 3180 SAPS (3533 SAPS – 10%)  
vSAPS per vCPU (without HT) = 2700 SAPS (3180 SAPS – 15%)
- VM without HT: #vCPUs =  $80,000 \text{ SAPS} / 2700 \text{ SAPS} = 29,63 \text{ vCPUs}$ , rounded up 30 cores / vCPUs
- VM with HT enabled: #vCPUs =  $80,000 \text{ SAPS} / 3180 \text{ SAPS} \times 2 \text{ (threads per core)} = 50,31 = 51 \text{ threads} / \text{vCPUs}$

VM vSocket calculation:

- $30 / 60 \text{ (CPU cores)} = 0.5$  or  $51 \text{ (threads)} / 120 \text{ (CPU threads)} = 0.43$

## Considerations for the final VM configuration

The SAPS-based sizing exercise showed that without HT enabled and used by the VM, 2 CPU sockets of a Cooper Lake system are needed to fulfill the SAPS requirements. With the Sapphire Rapids CPU, only half of a 60-core CPU is needed.

**Recommended:** Use hyperthreading in the VM configuration of the Cooper Lake system, because only one CPU socket of the Cooper Lake system is sufficient per the CPU requirements.

The sized SAP HANA database memory of 1,400 GB is the leading HANA sizing and VM configuration factor. On the Cooper Lake system, we have 1440 GB available per CPU. This fulfills the memory and CPU requirements as calculated. In contrast, the Sapphire Rapids system has only 2 TB installed, which, per CPU socket minus the ESXi host, costs only 960 GB. To fulfill the sized memory requirements of this example, you need to use 2 CPU sockets on the Sapphire Rapids host, which will result in very low CPU performance utilized system. You could solve this problem by adding memory to this host.

The final VM configuration will leverage all CPU and memory resources related to the CPU sockets needed. The resulting required configurations for the ESXi servers are:

- **4-socket 28-core Cooper Lake (8380HL) 6 TB host:** 1 CPU socket with 56 vCPUs (hyperthreading on) and 1,440 GB virtual RAM
- **2-socket 60-core Sapphire Rapids (8490H) 2 TB host:** 2 CPU sockets with 240 vCPUs (hyperthreading on) and 1,920 GB virtual RAM



Optimizing the Sapphire Rapids host configuration to have more memory per CPU socket will let you use a half-socket (SNC-2) configuration with a total of:

- 6 TB per host and the possibility to run on this host up to 4 VMs, each of around 1.5 TB
- 6 TB in the host with 2 VMs, each around 3 TB, which will result in low CPU performance (high CPU utilization). Therefore, you can use a lower bin Sapphire Rapids CPU if a half-socket SNC-2 deployment cannot be used.
- If you want to leverage SNC-2, then an SAP HANA “half-socket” VM with 1.5 TB and enough SAPS performance would be the best case for this example SAP HANA VM.

**Note:** If hyperthreading is used on the host, then you must set `numa.vcpu.preferHT=TRUE` per SAP HANA VM to ensure NUMA node locality of the vCPU threads.

### Pre-sized examples

Table 10 shows possible VM sizes of specific CPU types and variants. The SAPS figures are estimated based on published SAP SD benchmarks and rounded down. You will need to subtract the RAM needed for the ESXi from the shown figures. The figures represent the virtual SAPS capacity of a CPU core with and without hyperthreading.

**Note:** We based the SAPS figures shown in table 10 on published SAP SD benchmarks; you can use these sizings for Suite on HANA, BW on HANA, S/4HANA, or BW/4HANA workloads.

- For half-socket SAP HANA configurations, you must subtract 15% from the SAPS capacity to consider the CPU cache misses caused by VMs running concurrently on the same NUMA node.
- For SNC-2 configurations that use half-socket VMs on SPR or other supported systems, this reduction is not necessary.
- For mixed HANA workloads, contact SAP or your hardware sizing partner.
- For SAP HANA on Hyperconverged Infrastructure (HCI), reserve an additional 10% SAPS capacity to take vSAN into account.

Use table 10 to quickly determine a VM configuration that will fulfill the SAP HANA sizing requirements for RAM and CPU performance for SAP HANA on vSphere workloads.

Table 10: SAPS capacity and memory sizes for example configurations of SAP HANA on vSphere; based on published SD benchmarks and selected Intel CPUs

	<a href="#">Intel Xeon Platinum 8280L CPU</a>	Intel Xeon Platinum <a href="#">8380H</a> and <a href="#">8380HL</a> CPU	<a href="#">Intel Xeon Platinum 8380 CPU</a>	<a href="#">Intel Xeon Platinum 8490H CPU</a>
SAP benchmark	<a href="#">2021009</a> (Jan. 26, 2021)	<a href="#">2020050</a> (Dec. 11, 2020)	<a href="#">2023019</a> (May 5, 2023)	<a href="#">2023037</a> (Aug. 15, 2023)
Max. supported RAM per CPU as of Intel datasheet	4.5 TB	4.5 TB (HL)	6 TB	4 TB
CPU cores per socket as of Intel datasheet	28	28	40	60
Max. NUMA nodes per ESXi server <sup>2</sup>	8	8	2	8
SAP HANA supported maximum memory per CPU socket <sup>3</sup>	1.5 TB	1.5 TB	2 TB	3 TB
vSAPS per CPU thread with and without HT	2,459 (core without HT) 434 (HT gain) Based on cert. 2021009	2,432 (core without HT) 429 (HT gain) Based on cert. 2020050	2,563 (core without HT) 452 (HT gain) Based on cert. 2023019	2,703 (core without HT) 477 (HT gain) Based on cert. 2023037
0.5-socket SAP HANA VM <sup>4</sup>	1 to 16 x 14 physical core VM with min. 128 GB RAM and max. 768 GB <sup>15</sup>	1 to 16 x 14 physical core VM with min. 128 GB RAM and max. 768 GB <sup>15</sup>	1 to 4 x 20 physical core VM with min. 128 GB RAM and max. 1024 GB	1 to 4 x 30 physical core VM with min. 128 GB RAM and max. 1536 GB - with SNC-2

<sup>2</sup> Maximum NUMA nodes per CPU sockets depend on the Intel CPU architecture. Cascade Lake, Cooper Lake, and Sapphire Rapids CPUs support up to 8 socket CPU platforms. Ice Lake is a 2-socket platform only. Sapphire Rapids is only validated on up to 2-CPU sockets for SAP HANA on vSphere systems.

<sup>3</sup> SAP-supported appliance memory configurations are lower as specified by Intel. Cascade Lake and Cooper Lake CPUs support up to 1.5 TB per CPU socket, Ice Lake up to 2 TB, and Sapphire Rapids CPUs up to 3 TB per CPU socket. 2 TB are supported without a sizing; 3 TB are supported in a mixed DIMM configuration with 128 GB and 256 GB DIMMs and require a workload-based SAP HANA sizing.

<sup>4</sup> The SAPS (SAP Application Performance Standard) figures provided are derived from published SD benchmark results, incorporating hyperthreading with a 2-vCPU configuration. To account for virtualization costs, a standard deduction of 10% is applied. For half-socket configurations, an additional deduction of 15% from the SD capacity is required, along with the standard 10% virtualization cost deduction. SNC-2 half-socket VMs are exclusive to 2-socket SPR systems and are calculated with a standard 10% deduction only. There's no need to subtract an additional 15%. All figures presented are rounded and based on rounded SAPS performance figures from published SAP SD benchmarks. These figures should be used exclusively for Suite on HANA, BW on HANA, or BW/4HANA workloads. For sizing parameters related to mixed HANA workloads, please consult SAP or your hardware vendor directly.

	<a href="#">Intel Xeon Platinum 8280L CPU</a>	<a href="#">Intel Xeon Platinum 8380H and 8380HL CPU</a>	<a href="#">Intel Xeon Platinum 8380 CPU</a>	<a href="#">Intel Xeon Platinum 8490H CPU</a>
	vSAPS 34,000 28 vCPUs	vSAPS 34,000, 28 vCPUs	vSAPS 51,255 40 vCPUs	vSAPS 95,400 60 vCPUs
<b>1-socket SAP HANA VM<sup>5</sup></b>	1 to 8 x 28 physical core VM with min. 128 GB RAM and max. 1,536 GB	1 to 8 x 28 physical core VM with min. 128 GB RAM and max. 1,536 GB	1 to 2 x 40 physical core VM with min. 128 GB RAM and max. 2048 GB	1 to 2 x 60 physical core VM with min. 128 GB RAM and max. 3072 GB
	vSAPS17 81,000 56 vCPUs	vSAPS17 80,000 56 vCPUs	vSAPS 120,600 80 vCPUs	vSAPS 190,800 120 vCPUs
<b>2-socket SAP HANA VM<sup>5</sup></b>	1 to 4 x 56 physical core VM with min. 128 GB RAM and max. 3,072 GB	1 to 4 x 56 physical core VM with min. 128 GB RAM and max. 3,072 GB	1 x 80 physical core VM with min. 128 GB RAM and max. 4096 GB	1 x 240 physical core VM with min. 128 GB RAM and max. 6144 GB
	vSAPS 162,000 112 vCPUs	vSAPS 160,000 112 vCPUs	vSAPS 241,200 160 vCPUs	vSAPS 381,600 240 vCPUs
<b>4-socket SAP HANA VM<sup>5</sup></b>	1 to 2 x 112 physical core VM with min. 128 GB RAM and max. 6,128 GB	1 to 2 x 112 physical core VM with min. 128 GB RAM and max. 6,128 GB	NA (ICX is a 2-socket platform only)	1 x 480 physical core VM with min. 128 GB RAM and max. 8,192 GB
	vSAPS 324,000 224 vCPUs	vSAPS 320,000 224 vCPUs		vSAPS 763,200 480 vCPUs
<b>8-socket SAP HANA VM<sup>5</sup></b>	1 x 224 physical core VM with min. 128GB RAM and max. 12,096GB	1 x 224 physical core VM with min. 128 GB RAM and max. 12,096 GB	NA (ICX is a 2-socket platform only)	Not available yet
	vSAPS 648,000 448 vCPUs	vSAPS 640,000 448 vCPUs		

Here's the VM configuration example again. This time, we use the information provided in table 10 to quickly determine the VM configuration.

<sup>5</sup> The listed SAPS figures are based on published SD benchmark results with hyperthreading (2-vCPU configuration) minus a 10% virtualization cost. In the case of a half-socket configuration, in addition to the 10% virtualization cost, 15% from the SD capacity must be subtracted. SNC-2 half-socket VMs are calculated with only -10%. The shown figures are rounded and based on rounded SAPS performance figures from published SAP SD benchmarks, and can be used only for Suite on HANA, BW on HANA, or BW/4HANA workloads. For mixed HANA workload sizing parameters, contact SAP or your hardware vendor.

Sized SAP HANA system:

- 1,400 GB RAM SAP HANA system memory
- 80,000 SAPS (Suite on HANA SAPS)

Configuration example 1:

- 4-socket 28-core Cooper Lake (8380HL): **1.5 TB per CPU** host; total 6,144 GB
- 224 total CPU threads, 320,000 vSAPS available, **80,000 vSAPS per CPU** by leveraging hyperthreading
- VM configuration: 1 socket; 56 vCPUs (80,000 vSAPS); 1,400 GB vRAM (optimal configuration)

Configuration example 2:

- 2-socket 60-core Sapphire Rapids (8490H) with **1 TB per CPU** host; total 2,048 GB
- 240 vCPUs; 381,000 vSAPS available; **95,000 vSAPS per half-socket** by leveraging hyperthreading
- VM configuration: 2 sockets; 240 vCPUs; 1400 GB vRAM (non-optimal configuration due to the memory configuration of the host). We recommend you upgrade the memory to 3 TB per CPU.

**Note:** Review VMware [KB 55767](#) for details on the performance impact of using vs. not using hyperthreading. As described, especially with low CPU utilized systems, the performance impact of using hyperthreading is very little.

**Recommended:** In most cases, use hyperthreading for SAP HANA on vSphere hosts. You might make an exception for:

- Small workloads, which don't require hyperthreading to increase performance.
- Workloads with very low latency requirements.
- Security risks exposed by hyperthreading that exist on some processors.

## Storage configuration and sizing

Sizing a storage system for SAP HANA is different from sizing storage for classic SAP applications: SAP defines strict storage key performance indicators (KPIs) for production-level SAP HANA systems regarding data throughput and latency.

**Important:** Your production storage systems for SAP HANA VMs must meet these Key Performance Indicators (KPIs). If multiple SAP HANA VMs are running on one host, the TDI storage must be capable of supporting the specified number of SAP HANA instances/VMs.

The only variable in storage sizing is the capacity, which depends on the size of the in-memory database.

## Tools used to verify KPIs

You can verify your storage system meets the KPIs by using the following SAP tools:

- Hardware configuration check tool (HWCCT) for HANA 1.0
- Hardware and Cloud Measurement Tool (HCMT) for HANA 2.0

These tools are available only to SAP partners and customers. To download the tools, the KPI documentation, and the user guides, you need a valid SAP user account. 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 meet the KPIs for a specified number of SAP HANA VMs. Refer to the [Certified and Supported SAP HANA Hardware](#) web page.

## Storage connection

Along with the storage capacity, you must plan the storage connection. 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.

If, for example, you require a seamless migration between physically and virtually deployed SAP HANA systems, you can use vSphere-connected storage with raw device, in-guest mapped LUNs and in-guest mounted NFS storage solutions. Both are supported.

However, a fully virtualized storage solution works just as good as a natively connected storage solution and provides all the benefits of vSphere virtualization, including the possibility to abstract the storage layer from the operating system on which SAP HANA runs. For a detailed description about vSphere storage solutions, refer to the [vSphere documentation](#).

## vSphere datastores

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. For example, you can create datastores to be used as a single consolidated pool of storage, or you can use many datastores to isolate various application workloads.

vSphere datastores can be of different types: [VMFS](#), [NFS](#), [vSAN](#), or [vSphere Virtual Volumes](#). SAP HANA fully supports all of these. Refer to the [Working with Datastores](#) VMware documentation for details.

Table 11 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 11: vSphere supported storage types and features

Storage type	VM boot	Migration with vMotion	Datastore	Raw device mapping (RDM)	vSphere high availability (HA) and distributed resource scheduling (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 <sup>6</sup>	Yes	Yes	vSAN 6.6 or later	No	Yes

In summary, you use storage with SAP HANA on vSphere to:

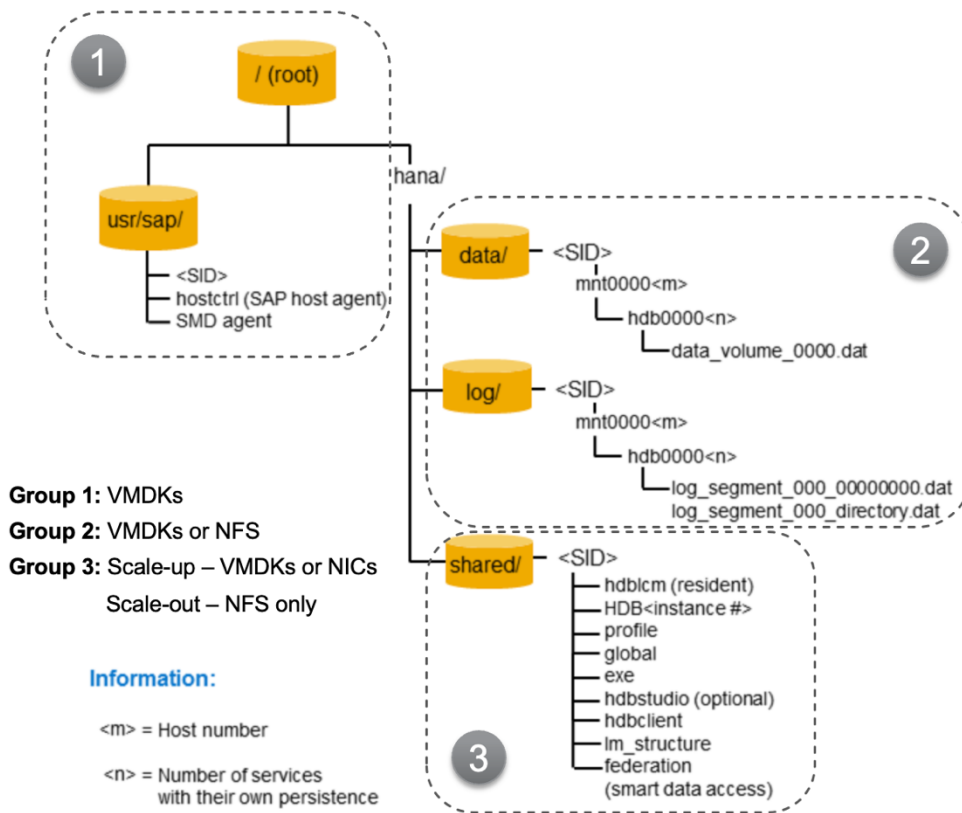
- Create separate and isolated datastores for operating systems, SAP HANA binaries, shared folders, data, and logs.
- Enable multiple SAP HANA VMs to provision their VMDK files on the same class of storage to meet SAP HANA storage KPIs.

## File system layout

Figure 7 shows the recommended SAP HANA Linux file system layout, which is the suggested layout when running SAP HANA virtualized on vSphere. 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.

<sup>6</sup> Only vSAN-based SAP HANA certified HCI solutions are supported. There is no support for generic vSAN solutions for SAP HANA production workloads.

Figure 7: SAP-recommended file system layout for SAP



## 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 the SAP HANA in-memory database.

SAP has defined very strict performance KPIs that must 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 architecture and sizing guidelines<sup>7</sup>, such as the [SAP HANA storage requirements](#). Figure 8 and table 12 consolidate that info—they show the typical disk layout of an SAP HANA system and the volumes needed. Use the information here as a starting point to plan the storage capacity needs of an SAP HANA on vSphere system. The volumes shown in figure 8 should correspond

<sup>7</sup> For more information, see the [SAP HANA TDI storage requirements](#).

with actual VMDK files and dedicated paravirtualized SCSI (PVSCSI) adapters or NVME controllers, which ensures the best and most flexible configuration.

**Recommended:** Use NVME controllers instead of PVSCSI because most modern systems have NVME devices instead of disks or SSDs installed.

The example configuration shown in figure 8 uses three independent PVSCSI adapters or NVME controllers and at least four independent VMDKs. This helps to parallelize I/O streams by providing the highest flexibility in terms of operation. We also use the file system layout as described in figure 8 and translate it into a disk volume/VMDK disk configuration.

Figure 8: Storage layout of an SAP HANA on vSphere system<sup>8</sup>

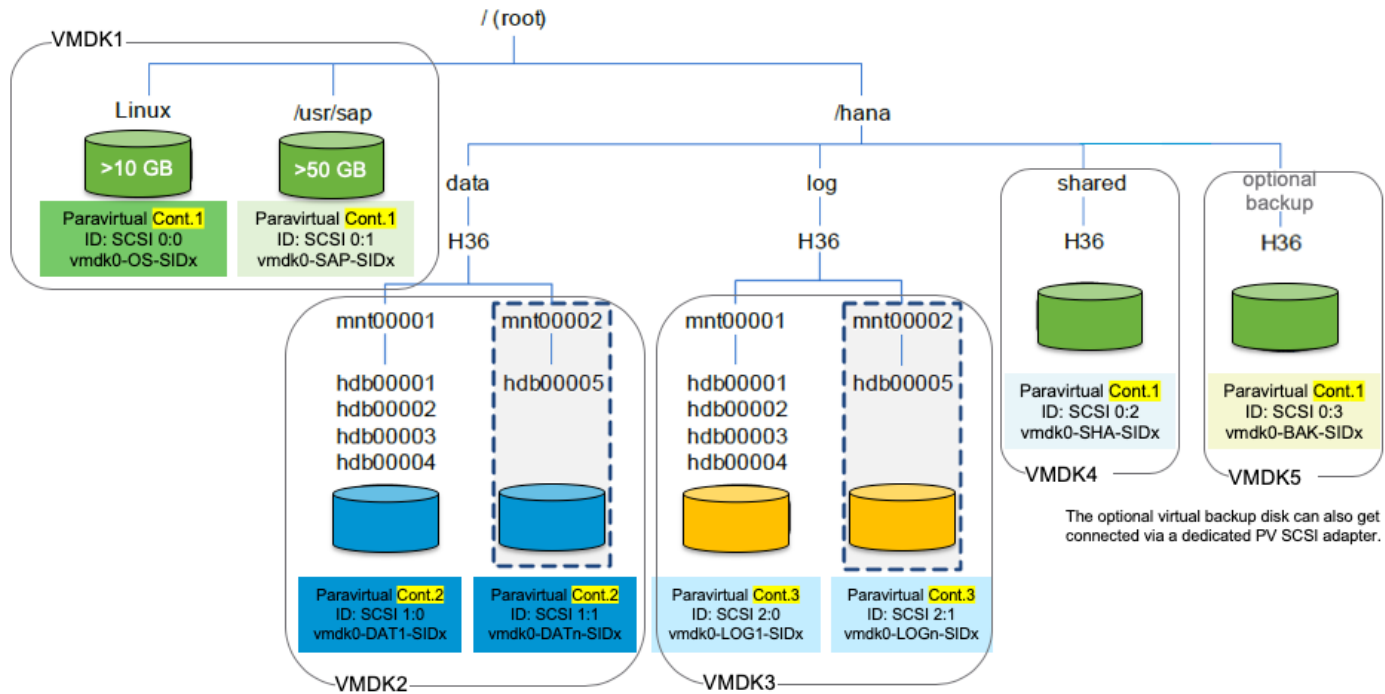


Table 12 provides storage capacity examples of the different SAP HANA volumes. Some of the volumes, such as the operating system and `usr/` SAP volumes, can be connected to and served by one PVSCSI adapter. Others, such as the log and data volumes, are served by dedicated PVSCSI controllers to ensure high I/O bandwidth and low latency.

**Important:** You must verify the performance of the I/O bandwidth and latency after an SAP HANA VM deployment with SAP hardware configuration checking tools. Refer to SAP Notes [1943937](#) and [2493172](#).

Instead of VMDK-based storage volumes, especially for data, you can use log or backup volumes that ingest connected NFS volumes.

<sup>8</sup> Image sourced from SAP SE and modified.



Table 12: Storage layout of an SAP HANA on vSphere system

Volume	Disk type	SCSI controller if VMDK	VMDK name	SCSI ID if VMDK	Sizes per SAP HANA storage requirements <sup>9</sup>
/ (root)	VMDK	PVSCSI (or NVME) Contr. 1	vmdk01-OS-SIDx	SCSI 0:0	Min. 10 GB for operating system We suggest 100 GB, thin-provisioned.
usr/sap	VMDK	PVSCSI (or NVME) Contr. 1	vmdk01-SAP-SIDx	SCSI 0:1	Min. 50 GB for SAP binaries We suggest 100 GB thin-provisioned.
shared/	VMDK or NFS	PVSCSI (or NVME) Contr. 1	vmdk02-SHA-SIDx	SCSI 0:2	Min. 1x RAM, max. 1 TB, thick-provisioned
data/	VMDK or NFS	PVSCSI (or NVME) Contr. 2	vmdk03-DAT1-SIDx vmdk03-DAT2-SIDx vmdk03-DAT3-SIDx	SCSI 1:0 SCSI 1:1	Min. 1x RAM, thick-provisioned <b>Note:</b> If you use multiple VMDKs, then use a tool to build one large data disk; for example, Linux LVM.
log/	VMDK or NFS	PVSCSI (or NVME) 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 - default path is /hana/shared	VMDK or NFS	PVSCSI (or NVME) Contr. 1 or 4	vmdk05-BAK-SIDx	SCSI 3:0	size of backups >= size of SAP HANA data + size of redo log <b>Note:</b> You must change the default backup path of /hana/shared when an optional, dedicated backup volume is used (thin-provisioned).  To further optimize data throughput for backup, you can use a dedicated PVSCSI or NVME adapter.

Depending on the storage solution used and future growth of the SAP HANA databases, you might need to increase the storage capacity or better balance the I/O over more LUNs. In this case, you can use the Linux Logical Volume Manager (LVM), which is fully supported, to build LVM volumes based on VMDKs.

To determine the overall storage capacity per SAP HANA VM, sum the sizes of all specific and unique SAP HANA volumes as outlined in figure 8 and table 12.

<sup>9</sup> Information summarized from SAP OSS note [1900823](#), attachment: SAP HANA storage whitepaper version 2.1

To determine the minimum overall vSphere cluster datastore capacity required, multiply the SAP HANA volume requirements in table 12 by the number of VMs running on all hosts in the vSphere cluster.

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

Use equation 2 for your storage capacity calculation:

**Equation 2** vSphere datastore capacity = total SAP HANA VMs running in a vSphere cluster × individual VM capacity needed (OS + usr/sap + /shared + /data + /log)

For example, a sized SAP HANA system with RAM = 2 TB would need the following:

- VMDK root (OS) >= 10 GB (→ recommended 100 GB, thin-provisioned)
- VMDK usr/sap >= 60 GB (→ recommended 100 GB, thin-provisioned)
- VMDK HANA shared = 1,024 GB (thick-provisioned)
- VMDK HANA log = 512 GB (thick-provisioned)
- VMDK HANA data = 2,048 GB (thick-provisioned)
- VMDK HANA backup >= 2,5 TB (thin-provisioned; optional)

In this example, the VM capacity requirement = 3.7 TB with optional backup of the total is 6.2 TB.

To determine the cluster-wide storage capacity requirement, you must do this calculation for all possible running SAP HANA VMs in the vSphere cluster. 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 be guaranteed for all production SAP HANA VMs. Use the HWCCT to verify these KPIs. Otherwise, the overall performance of the SAP HANA system might be lower than required. Refer to the above section, "

Tools used to verify KPIs."

## Hardware and cloud measurement tool for SAP HANA

Using the SAP hardware check tools (refer to SAP Notes [1943937](#) and [2493172](#) and this [blog post](#)) allows you to verify, in addition to 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) an external Fibre Channel flash array from [Pure Storage](#). This ensures that the conducted tests and validations use a modern external storage array and software-defined storage.

Table 13 shows what is possible with a VMware-optimized Pure X50 storage system compared to a bare metal system running the same configuration and following the test results for log latency, which SAP has defined as a maximum of 1,000  $\mu\text{s}$ . SAP HANA running on a VM shows a 22% higher latency than SAP HANA running the same configuration on bare metal. However, the virtualized system still runs quite a bit below the SAP-defined KPI of 1,000  $\mu\text{s}$ , at 406  $\mu\text{s}$ , which shows good performance.

**Table 13: HCMT file system latency example of SAP HANA running on a bare metal system vs on a vSphere VM**

Configuration	HCMT 16K block log volume latency
Bare metal system with FC-connected Pure X50 storage unit	16K block log overwrite latency = 334 $\mu\text{s}$
VM with FC-connected Pure X50 storage unit	16K block log overwrite latency = 406 $\mu\text{s}$
Absolute difference	72 $\mu\text{s}$
Percentage difference	22%
SAP HANA KPI	1,000 $\mu\text{s}$ (bare metal 3x and virtual 2.5x faster, as required)

In table 14, the virtualized SAP HANA system has only slightly higher throughput than the bare metal SAP HANA system running on the same configuration. However, both results are over the SAP-defined KPI of 120 MB/s for this test.

Table 14: HCMT file system overwrite example of a SAP HANA bare metal system compared to a virtualized system


Configuration	HCMT 16K block log volume overwrites
Bare metal system with FC-connected Pure X50 storage unit	16K block log overwrite throughput = 706 MB/s
VM with FC-connected Pure X50 storage unit	16K block log overwrite throughput = 723 MB/s
Absolute difference	17 MB/s
Percentage difference	2.4%
SAP HANA KPI	120MB/s (BM and VM over 5 times higher as required)

For these tests, we used an 8-socket wide VM running on an Intel Cascade Lake [Fujitsu PRIMEQUEST 3800B2](#) and lately 2- and 4-socket Intel Sapphire Rapids based [ThinkAgile VX650](#) and [VX850 V3 systems from Lenovo](#) connected via Fibre Channel on a [Pure X50 storage](#) unit configured as outlined in the Pure Storage and VMware [best practices configuration guidelines](#). The results shown in the tables were conducted on the 8-socket Cascade Lake systems. Make sure you apply the operating system 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

An SAP HANA VM on a vSphere cluster requires dedicated networks for the 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 you need to support a specific SAP HANA workload in a VM and ultimately on the hosts.

Figure 9: Logical network connections per SAP HANA server<sup>10</sup>



In contrast to a physical SAP HANA environment, you must plan the SAP HANA operating system-exposed networks and the ESXi host network configuration that all SAP HANA VMs will share. You need to configure:

- The virtual network, especially the virtual switch, on ESXi
- HANA-exposed virtual network cards in the VM

The following sections explain each.

**Note:** You should consider the need for multiple SAP HANA VMs on one ESXi host when configuring the host network. A single SAP HANA instance on a single host might require more or higher bandwidth network cards.

vSphere offers standard and distributed switch configurations. Both switches can be used when configuring an SAP HANA on vSphere environment.

**Recommended:** 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.

<sup>10</sup> Figure sourced from SAP SG.

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

Table 15 describes the recommended network configuration for SAP HANA running virtualized on an ESXi host using different network card configurations. The information there is based on SAP recommendations and includes vSphere-specific networks like vMotion and dedicated storage networks for software defined storage (SDS) or NFS. You might need to use multiple network cards in order to meet the network requirements for each SAP HANA VM when consolidating all of them onto a single ESXi host.

Table 15: Recommended SAP HANA on vSphere network configuration based on 10 GbE NICs

		Host IPMI/ remote control network	vSphere admin + vSphere HA network	Application server network	vMotion network	Backup network	System replication (HSR) network	Scale-out internode network	NFS/ vSAN/SDS network
Host network configuration  (all networks are dual-port NICs configured as failover teams)	Network label	Physical host management	vSphere admin	SAP app server	vMotion	Backup network (optional)	HANA replication (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(1,500)	Default (1,500)	Default (1,500)	9,000	9,000	9,000	9,000	9,000
	VLAN ID# examples		200	201	202	203	204	205	206
	Physical NIC#, bandwidth <sup>11</sup>	#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)			.		.	.	.	.

<sup>11</sup>The network card bandwidth influences how many SAP HANA VMs are supported on the vSAN datastore and how long a vMotion migration process will take to finish. Depending on the SAP HANA memory sizes, we recommend using 4- and 8-socket host systems and 4-socket large VMs with a minimum of 25 GbE NICs.

**Recommendations:**

- Use VLANs to reduce the total number of physical network cards needed in a server.
- Make sure there's enough network bandwidth available for each SAP HANA VM.
- Make sure the installed network cards of an ESXi host provide enough bandwidth to serve all SAP HANA VMs running on it. Oversubscribing network cards will result in poor 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.

**Recommended:** Create a vSphere Distributed Switch per dual-port physical NIC and configure port groups for teaming and failover purposes. Use the default port group setting, except for the uplink failover order, as shown in table 16.

A port group defines properties regarding security, traffic shaping, and NIC teaming.

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

Depending on the required optional networks needed for vSphere or SAP HANA system replication or scale-out internode networks, the suggestions from this table 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.

You should configure the network port failover teams to allow physical NIC failover.

**Recommended:** For physical NIC failover, you should use NICs with the same network bandwidth, such as only 10 GbE NICs, and group failover pairs depending on their needed network bandwidth.

**Important:** Do not group, for instance, two high-bandwidth networks, such as the internode network with the app server network.

**Table 16: Minimum ESXi server uplink failover network configuration based on 10 GbE NICs**

Property	NIC	VLAN <sup>12</sup>	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

<sup>12</sup> For the VLAN ID example, final VLAN numbers are up to the network administrator.

Property	NIC	VLAN <sup>12</sup>	Active uplink	Standby uplink
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

**Recommended:** Use different VLANs, as shown, to separate the VMware operational traffic (for example, vMotion and vSAN) from the SAP and user-specific network traffic. Also use higher bandwidth network adapters to reduce the number of physical network cards, cables, and switch ports.

Table 17 shows an example with 25 GbE network adapters.

Table 17: Example SAP HANA on vSphere network configuration based on 25 GbE NICs

		Host IPMI/ remote control network	vSphere admin + vSphere HA network	Application server network	vMotion network	Backup network	System replication network	Scale-out internode network	NFS/ vSAN/SDS network
Host network configuration (all networks are dual-port NICs configured as failover teams)	Network label	Physical host mgmt	vSphere admin	SAP app server	vMotion	Backup network (optional)	HANA replication (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 (1,500)	Default (1,500)	Default (1,500)	9,000	9,000	9,000	9,000	9,000
	VLAN ID# examples		200	201	202	203	204	205	206
	Physical NIC#, bandwidth <sup>13</sup>	#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	Virtual NIC (inside VM)			.		.	.	.	.

<sup>13</sup> The selected network card bandwidth influences how many SAP HANA VMs are supported on the vSAN datastore or how long a vMotion migration process will take to complete. Depending on the HANA memory sizes, we recommend you use 4- and 8-socket host systems and 4-socket large VM with minimum of 25 GbE NICs.



Using 25 GbE network adapters helps reduce the number of NICs, network cables, and switch ports. With higher bandwidth NICs, the vSphere system can support more SAP HANA VMs per host.

Table 18 shows an example of how to group the network port failover teams.

**Table 18: Minimum ESXi server uplink failover network configuration based on 25 GbE NICs**

Property	NIC	VLAN <sup>14</sup>	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

**Recommended:** Use different VLANs, as shown, to separate the vSphere operational traffic (for example, vMotion and vSAN) from the SAP and user-specific network traffic. Use 100 GbE bandwidth network adapters to further reduce the number of physical network cards, cables, and switch ports.

<sup>14</sup> The network administrator usually determines the final numbers for the VLAN ID example.

Table 19 shows an example with 100 GbE network adapters.

Table 19: Example SAP HANA on vSphere Network Configuration based on 100 GbE NICs

		Host IPMI/ remote control network	vSphere admin + vSphere HA network	Application server network	vMotion network	Backup network	System replication network	Scale-out internode network	NFS/ vSAN/SDS network
Host network configuration (all networks are dual-port NICs configured as failover teams)	Network label	Physical host management	vSphere admin	SAP app server	vMotion	Backup network (optional)	HANA replication (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 MTUsize	Default (1,500)	Default (1,500)	Default (1,500)	9,000	9,000	9,000	9,000	9,000
	VLAN ID # examples		200	201	202	203	204	205	206
	Physical NIC #, bandwidth	#0, 1 GbE	#1, 100 GbE			#2, 100 GbE			
	Active physical NICport	1	1			1			
	Standby physical NICport	2	2			2			
VM guest network cards	Virtual NIC (inside VM)			.		.	.	.	.

Table 20 shows an example of how to group the network port failover teams based on 100 GbE NICs.

**Table 20: Minimum ESXi Server Uplink Failover Network Configuration based on 100 GbE NICs**

Property	NIC	VLAN <sup>15</sup>	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

**Recommended:** Use different VLANs, as shown, to separate the VMware operational traffic (for example, vMotion and vSAN) from the SAP and user-specific network traffic.

## Workload performance and validation testing

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 workloads and scale from single user tests up to thousands of concurrent users (for example, up to 78,000 with Cooper Lake processors), pushing the virtualized SAP HANA system to its limits.

**VMXNET3 network latency considerations:** Our performance testing found that the virtualized network card (VMXNET3) typically adds between 60  $\mu$ s (no load) and up to 300  $\mu$ s of latency (high CPU load  $\geq$  80 percent) to every network package compared to a bare metal installed SAP HANA system, which impacts SAP OLTP and OLAP workloads issued by remote application servers or users. The following sub-sections include information about this impact and describe how to mitigate it by optimizing the virtual and physical network configuration.

<sup>15</sup> For the VLAN ID example, final VLAN numbers are up to the network administrator.

# SAP workload characterization and the impact on network performance

We can characterize SAP workloads in three ways:

- OLTP workloads, which represent classic SAP ERP systems
- OLAP workloads, which represent the analytical workloads of Business Warehouse (BW) systems
- Both OLTP and OLAP, which represent SAP S/4HANA systems

Because SAP S/4HANA combines these two workload types, we need to consider both of their characteristics. Typical OLTP workloads are small network packages with a recommended MTU size of 1,500, whereas the recommended MTU size for OLAP workloads is 9,000. You need to understand how the SAP S/4HANA system is used to choose the correct MTU size.

You also need to understand how many concurrent users will interact with the SAP HANA database and how much network traffic they will create. In recent tests of vSphere with SAP HANA with OLTP workloads, SAP and VMware observed an increase of OLTP transactional request times, which showed an overhead of up to 100 ms when compared to bare metal SAP HANA systems.

We observed this increase when using the VMXNET3 virtual NIC; it's because virtual networking adds the mentioned latency in the range of 60  $\mu$ s (no load) and up to 300  $\mu$ s (high load, wide VM, measured with 23,000–64,000 users running on a 4-socket Cooper Lake processor host and 35,000–78,000 users running on an 8-socket Cooper Lake processor host) to each network package sent or received. VMware knowledge base article [83957](#) documents this.

Unlike for storage (see the "Storage configuration and sizing" section), SAP did not define SAP HANA-specific network KPIs for throughput and latency that an SAP HANA system must 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 we measured the VMXNET3 impact and how to optimize the network configuration for an SAP HANA VM for its given use case.

## Performance tests run with SAP S/4HANA and BW/4HANA: workload and validation testing

We ran performance tests to measure the impact of virtualization on SAP HANA, and ultimately on the users, to define the best possible configuration for mitigating the virtualization costs, such as increased network latencies.

### Testbed configuration

A typical SAP S/4HANA environment is a three-tier landscape with the application tier on one or more separate hosts from the database tier, with users accessing the application servers when they work within the stored data of the SAP HANA database.

We configured tests with SAP S/4HANA and BW/4HANA to run on the three-tier environment shown in figure 10 to simulate real customer systems.

Figure 10: SAP three-tier architecture on vSphere

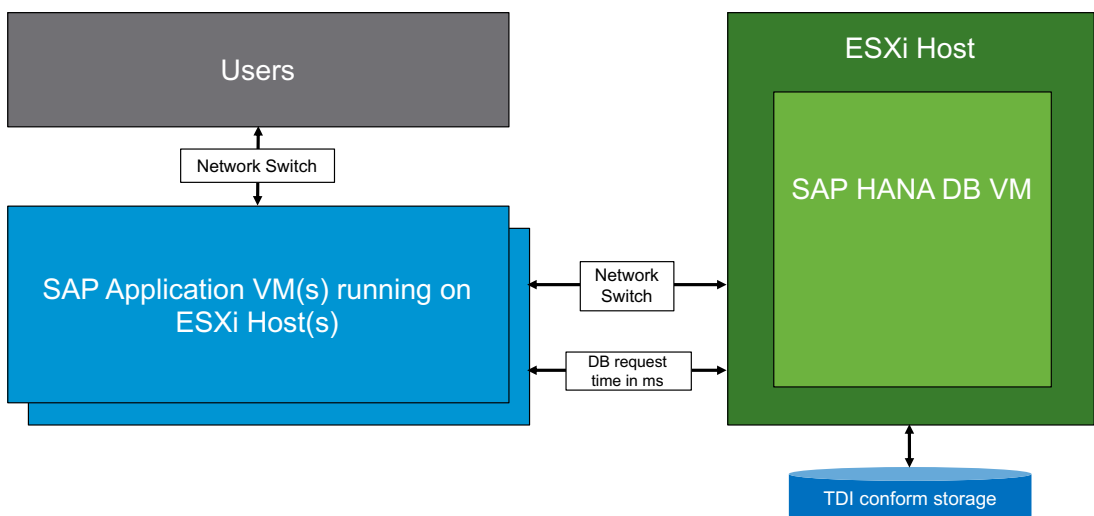


Table 21 shows our testbed configuration.

**Table 21: Testbed configuration**

<b>Software</b>	<ul style="list-style-type: none"> <li>• SAP S/4HANA</li> <li>• SAP BW/4HANA</li> <li>• VMware vSphere</li> </ul>
<b>Application workloads</b>	<ul style="list-style-type: none"> <li>• OLTP</li> <li>• OLAP</li> <li>• OLTP/OLAP mix</li> </ul>
<b>Benchmark software</b>	<ul style="list-style-type: none"> <li>• SAP S/4HANA – custom mixed workload that exercised the CPU utilization at 35% for the 4-socket server and 65% for the 8-socket server</li> <li>• SAP BW/4HANA – <a href="#">BWH benchmark</a></li> </ul>
<b>vSphere host and bare metal server hardware</b>	<ul style="list-style-type: none"> <li>• 4-socket Cooper Lake system</li> <li>• 8-socket Cooper Lake system</li> </ul>
<b>VM configuration</b>	<ul style="list-style-type: none"> <li>• Info to come</li> </ul>

## Workload information

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:

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

## Test methodology

The tests simulated OLTP and OLAP transactions up to the maximum possible CPU utilization level. A load driver initiated the application workloads, which simulated thousands of SAP users accessing the SAP HANA VM at the same time.

The application server instances received these requests and ran the SAP-specific transactions against the SAP HANA database. These transactions created several hundreds of database logical units of work ([LUW](#)) and were measured as a database request time in milliseconds.

**Note:** The measured database request time is the time measured for transactions between the SAP application server and the SAP HANA database instance. We did not measure the time a user needed to wait until an application server responded to a user-initiated transaction. The user-to-application server time is normally significantly higher than the database request time between application servers and a database instance.

The number of simulated SAP users per test run depended on the SAP HANA VM size and started at approximately 23,000 concurrent users for a 4-socket wide VM and approximately 35,000 concurrent users for an 8-socket wide VM (low load), and then increased to approximately 44,000 (4-socket) and approximately 60,000 (8-socket) concurrent users (moderate load) and then increased again to approximately 64,000 (4-socket) and approximately 78,000 (8-socket) concurrent users until the throughput dropped, which represents the maximum number of users possible (high load).

The number of concurrent users on the SAP HANA database instance represents a moderate to high CPU load:

- 35% CPU utilization → low CPU load
- 65% CPU utilization → moderate CPU load
- >65% CPU utilization → high CPU load (for example, 85%)

The number of users increased until the OLTP/OLAP throughput dropped. With this requirement:

- The 4-socket server achieved approximately 64,000 concurrent users at a CPU utilization of approximately 80%.
- The 8-socket system achieved approximately 78,000 concurrent users, also at a CPU utilization of about 80%.

As mentioned, these transactions were run 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 number of SAPS for such a physical or virtualized SAP HANA system.

## Test results for OLTP workload

The test results of running the custom mixed workload and the BWH benchmark provide information on virtual network performance and let us make recommendations to lower the impact of virtualization on networking performance.

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

The following tables and figures summarize the results of these findings based on the latest Cooper Lake processors.

Table 22 shows the minimum and maximum VMXNET3 latency deviation measured with netperf on 8-socket wide virtualized (VMXNET3 VM) and non-virtualized SAP HANA systems when running first with no load, and next under a high user load. The high user load represents a CPU utilization of up to 65 percent. The bare metal SAP HANA system had, on average, a netperf TCP roundtrip time of around 26  $\mu$ s (no load) and up to 95  $\mu$ s (high load). The virtualized SAP HANA system showed, on average, 84  $\mu$ s and 319  $\mu$ s TCP roundtrip time while under load.

Table 22: Network latency of Netperf TCP VMXNET3 compared to bare metal, measured in  $\mu$ s

8-socket server/VM	Baseline latency with no load $\mu$ s	Latency at peak load with 65,000 concurrent users ( $\mu$ s)
Bare metal	26	95
VMXNET3	84	319
<b>Overall change in value (delta)</b>	<b>58</b>	<b>224</b>

The virtualization overhead here is 3x that of bare metal. The overhead specifically refers to the higher TCP round trip time (latency) per network package sent and received (while idle: no load; running: high load). This reduces the overall performance of the database request time. We found two ways to increase performance:


- Use a network card configured as passthrough. This shows only a slight latency compared to a physical NIC.
- Optimize the underlying physical network to lower the overall latency of the network.

When we observed the VMXNET3 virtualization overhead (latency) of an 8-socket wide VM running on an Intel Cooper Lake server with 416 vCPUs (32 CPU threads were reserved to handle this massive network load) to a natively installed SAP HANA system running on the same server, we saw how these microseconds accumulate to a database request time deviation of between 27 ms and approximately 82 ms. See Figure 11 for this comparison.

**Note:** Running the application workload with 91,000 concurrent users results in a significant volume of network traffic. By reducing the number of vCPUs from 448 to 416, helped the ESXi kernel to better manage this intense network load when using VMXNET3. Subsequent tests using vSphere 8 with the same Cooper Lake configuration demonstrated that reducing the vCPUs to 432 was adequate. Lower user loads typically do not require to reserve CPU threads for the networking stack.



Figure 11: Mixed workload OLTP database request time in milliseconds (ms)



The database request time was impacted by up to 25% (27 ms higher) at 35% CPU utilization when VMXNET3 was used; however, the OLTP throughput per hour (TPH) and OLAP queries per hour (QPH) results were not impacted. At approximately 65% CPU utilization, the database request time increased to 36% (82 ms higher) with a TPH deviation of approximately -1 percent. At the maximum user capacity with 91,000 users at approximately 80% CPU utilization, the impact of TPH and QPH was approximately -8%. The OLAP request times were not impacted much. Figures 12, 13, and 14 show these results.

We found that using a network device configured as passthrough (instead of using a VMXNET3 network card) reduced the database request time for the OLTP application workload:

- At 35% CPU utilization, the VM with a passthrough NIC showed only a 3% difference (3 ms) compared to bare metal.
- At 65% CPU utilization, the VM with a passthrough NIC showed approximately a 9% (21 ms) difference.


80% of the time, the difference was still below 10%, keeping the TPH and QPH deviations below -3% while maxed out.

Reserving CPU threads to handle network traffic on the ESXi side was not necessary because the network traffic was handled inside the VM and by the OS—not by ESXi as when VMXNET3 is used. Refer to figures 11 and 12.

**Note:** The measured database request time is the time between the SAP application server and the HANA database instance. We did not measure the time a user had to wait until an application server responded to the user-initiated transaction. The user-to-application server time was significantly higher than the database request time between application servers and a database instance.


Figure 12 shows the OLAP request times were not impacted much by the VMXNET3 or passthrough network cards when compared to a bare metal server.

Figure 12: Mixed Workload OLAP database request time in milliseconds (ms)



While the virtualization overhead can already be measured with lower user numbers (network traffic), this SAP S/4HANA mixed workload test has shown that the main impact can be measured at higher user load numbers that quickly generate massive OLTP requests and with little impact to OLAP database request time (figure 13).

Figure 13: Mixed Workload OLTP Database TPH




Using a virtualized VMXNET3 network adapter typically has no significant effect on database request time at CPU utilizations between 35% and 65%, which means that the impact on OLTP throughput is less than 8%. In this test, the so-called max-out point was -8% using VMXNET3. Keep in mind that standard SAP workloads can reach 65% CPU utilization.

If your application is sensitive to database request times and you want to lower the network latency between the SAP app server tier and the database instance, consider using passthrough NICs instead of VMXNET3 NICs. However, passthrough doesn't support vMotion.

The mixed workload OLAP QPH results of a Cooper Lake system are displayed in figure 14. Using VMXNET3 has very little effect on the QPH results, which are between 35% and 65% (up to 1%), and at the max-out point at over 80% CPU utilization, it has about -8% of an impact.

Figure 14: Mixed workload OLAP database QPH



As a summary, for OLAP workloads, a passthrough NIC doesn't make complex query runtime any faster. Lowering the number of vCPUs in a VM to reserve some CPU cores for the ESXi kernel might help in some situations, like when the data is being loaded, but it means SAP HANA has less performance (lower QPH).

For OLAP workloads, we recommend using the VMXNET3 NIC. OLTP workloads benefit most from passthrough NICs. If you don't use a passthrough NIC and need to use VMXNET3 for some reason, you could instead reduce the number of vCPUs.

We recommend you begin with VMXNET3, and if the database request time is longer than expected, then check the physical network infrastructure and, if possible, optimize it for low latency before considering passthrough NICs. This can help achieve nearly bare-metal NIC latencies. Once again, optimizing the SAP HANA database network for low latency and throughput will have the most beneficial impact on overall SAP HANA performance.

**Note:** When database request times are high, consider optimizing the physical SAP network. Start from the user to the SAP app server tier, then from the app server to the SAP HANA database. The use of low-latency switches, a flat network architecture, or newer NICs will help reduce the transaction request time experienced by users. The use of passthrough NICs inside the SAP HANA VM will only impact the database request time to the app servers at the expense of losing vMotion capabilities.

## SAP S/4HANA and B/4HANA validation test findings (OLAP)

Business warehouse (OLAP) workloads affect database request time and QPH results less than OLTP workloads—OLTP workloads generate more frequent network traffic shorter packages. VMXNET3 latency overhead is less noticeable for OLAP workloads, which generate long running queries but are less frequent.

Table 23 shows the results of published BWH tests with Intel Cascade Lake server systems. These results are public and have been certified. The results show a BWH configuration based on an L-class sizing that SAP specified for servers with 8 sockets and 6 TB of memory.

We compared the virtual test results to a BWH system natively installed on the same hardware configuration.

Table 23 shows the test results of an 8-socket BWH set up according to the standard SAP L-class configuration. We ran the test on an Intel 8280L Cascade Lake server with 6 TB of memory and 28-core CPUs running at 2.7GHz. The test consisted of 9 data sets (11.7 billion records).

**Table 23: BWH 6 TB L-class workload test: bare metal compared to VM with VMXNET3**

SAP BWH -Class KPIs					25,000		5,000		200	
Configuration	Cert	CPUs Threads	Memory	Records	BWH Phase 1 (sec.)	Delta	BWH Phase 2 (QPH)	Delta	BWH Phase 3 (sec.)	Delta
Bare metal CLX 8S host	<a href="#">2020021</a>	448	6.144 GB	11.7 billion	19,551	-	5,838	-	146	-
VM with VMXNET3	<a href="#">2020031</a>	448	5.890 GB		20,125	2.94%	5,314	-8.98%	146	0%
					- is better		+ is better		- is better	

As you can see in table 23, the virtualization overhead while using a VMXNET3 NIC was very little and within 10%. In Phase 3 (measured in seconds), there was no change in the total runtime of the complex query phase.

In very critical environments this overhead can get lowered even further by leveraging a PT NIC instead of a VMXNET3 NIC.

## BWH L-Class vs. M-class sizings for SAP HANA VMs


HPE performed an SAP BWH benchmark with SAP HANA 2.0 SP6, including 6 TB and 12 TB large database instances that stored 10,400,000,000 (BWH L-Class sizing (5,000 QPH)) and 20,800,000,000 (BWH M-Class sizing (2,500 QPH)) initial records—that’s the highest ever measured results in a VMware vSphere 7 virtualized environment with 7,600 (cert. 2022014) and 4,676 (cert. 2022015) queries per hour (QPH).

Table 24 and figure 15 shows the benchmark results of a bare metal vs. virtual environment. While the virtual results don’t pass the BWH L-Class sizing mark, they are still within 10% of a previously published bare metal BWH 12 TB benchmark, which ran on the same hardware configuration with SAP HANA 2.0.

**Table 24: BWH 12 TB Cooper Lake, 8-socket M-Class workload test: bare metal vs. virtual**

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
HPE Superdome Flex 280, Intel Xeon Platinum 8380HL,28 core CPU - Bare Metal	<a href="#">2021058</a>	448	12.288 GB	20.8 billion	14,986	-	5,161	-	137	-
HPE Superdome Flex 280, Intel Xeon Platinum 8380HL,28 core CPU – vSphere 7 VMXNET3 VM	<a href="#">2022015</a>	448	11.776 GB		15,275	-1.93%	4,676	-9.40%	149	8.76%
					- is better		+ is better		- is better	

Figure 15: BWH 12 TB Cooper Lake, 8-socket M-Class workload test: bare metal vs. virtual




With the new Sapphire Rapids based systems with vSphere 8 we achieved the BWH L-Class sizing mark with both a 6 TB and an 8 TB 4-socket virtualized SAP HANA Sapphire Rapids system and with the 2-socket configuration we are very close to it, see the next tables and figures for details. The overall virtualization overhead for the 2-socket configuration, was below 3% to a bare metal configuration. For your reference, the 2-socket SPR 4 TB BWH benchmark was published under cert number [2023030](#).

Table 25: BWH 4 TB Sapphire Rapids, 2-socket M-Class workload test: bare metal vs. virtual

SAP BWH M-Class KPIs					35,000		2,500		300	
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
Lenovo ThinkAgile VX650 V3 CN, Intel® Xeon® Platinum 8490H, 60 core CPU - Bare Metal	-	240	4,096 GB	7.8 billion	16,158	-	5,036	-	109	-
Lenovo ThinkAgile VX650 V3 CN, Intel® Xeon® Platinum 8490H, 60 core CPU- vSphere 8 VMXNET3 VM	2023030	240	3,937 GB		15,738	-2.6%	4,933	-2%	110	-1.3%
					- is better		+ is better		- is better	



Figure 16: BWH 4 TB Sapphire Rapids, 2-socket M-Class workload test: bare metal vs. virtual




The same benchmark running with mor datasets on a 4-socket Sapphire Rapids system with 6 TB and 8 TB and 480 logical CPUs shows also very little virtualization overhead (below 6%) and with both memory configurations we achieved the BHW L-Class CPU sizing category.

Table 26: BWH 8 TB Sapphire Rapids, 4-socket M-Class workload test: bare metal vs. virtual, achieving L-Class CPU sizing mark

SAP BWH M-Class KPIs					35,000		2,500		300	
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
Lenovo ThinkAgile VX850 V3 CN, Intel® Xeon® Platinum 8490H, 60 core CPU - Bare Metal	-	480	8,192 GB	14,3 billion	12,778	-	5,902	-	151	-
Lenovo ThinkAgile VX850 V3 CN, Intel® Xeon® Platinum 8490H, 60 core CPU- vSphere 8 VMXNET3 VM	-	480	8,000 GB		13,506	5.7%	5,880	-0.4%	156	3.3%
					- is better		+ is better		- is better	

Figure 17: BWH 8 TB Sapphire Rapids, 4-socket M-Class workload test: bare metal vs. virtual



## Virtual SAP HANA performance evolution over 5 CPU generations

Summarizing our validation test results, we have gone over the different CPU generations with the SAP internal mixed workload tests and the public BHW benchmarks, which show significant performance gains when the different CPU generations are compared to the same vSphere version.

The next tables provide a performance snapshot comparing different virtual SAP HANA systems running on Broadwell and now Sapphire Rapids systems.

**Table 27: 4-Socket vSphere 8 VM vs older CPU generations running mixed SAP HANA workload test VM @ max-out CPU MP**

Configuration	Users @ max. CPU	Used CPU Threads	TPH	QPH
4-Socket Intel Broadwell E7-8880 v4 22 Core CPU - VMXNET3 vSphere 7 VM	35,000	168	3,723,484	6,210
2-Socket Intel Ice Lake 8380 40 Core CPU VMXNET3 vSphere 8 VM	50,000	160	5,676,112	9,487
8-Socket Intel Cascade Lake 8280L, 28 Core CPU - VMXNET3 vSphere 8 VM	85,000	448	7,710,063	12,855
2-Socket Intel Sapphire Rapids 8490H 60 Core CPU - VMXNET3 vSphere 8 VM	77,000	240	8,229,963	13,700
8-Socket Intel Cooper Lake 8380HL 28 Core CPU - VMXNET3 vSphere 8 VM	91,000	448	9,753,984	16,251
4-Socket Intel Sapphire Rapids 8490H 60 Core CPU - VMXNET3 vSphere 8 VM	110,000	480	11,389,998	19,025

**Note:** The Broadwell system used vSphere 7 instead of vSphere 8. The performance deviations between vSphere 7 and 8 are very little and for this comparison are not relevant.

Figure 18 shows the throughput per hour (TPH) performance in this specific test for SAP HANA over the different CPU generations in a graphical form. If you compare a 4-socket Intel Broadwell-based 4-socket server with an Intel Sapphire Rapids 4-socket system, then you see an increase of the TPH test result of factor 3 or 300%.

Figure 18: Performance evolution of CPU generations running mixed SAP HANA workload test VM @ max-out CPU MP measured in throughput per hour (TPH)





Figure 19 shows the queries per hour (QPH) performance gain over the different CPU platforms.

Figure 19: Performance evolution of CPU generations running mixed SAP HANA workload test VM @ max-out CPU MP measured in queries per hour (QPH)



The next table and figures will provide the same overview but using the public available SAP HANA BWH benchmark. Apart from the storage used, the benchmark environment remained identical, making most of the test results, but Phase 1 of the BWH benchmark comparable. Also, here we can see a significant improvement from an 8-socket Cascade Lake (448 logical CPU) system to a 4-socket Sapphire Rapids (480 logical CPU) based system. Both systems run the same HANA, OS and BWH versions with 6 TB memory and 8 datasets configured. While the SPR system has more 7.14% CPU threads, it shows a performance improvement of 35%. See figure 20 and table 28 for details.

Figure 20: BWH 6 TB Sapphire Rapids, 4-Socket vSphere 8 VM vs. Cascade Lake, 8-Socket vSphere 8 VM - 35% Performance Gain




Table 28: BWH 6 TB Sapphire Rapids, 4-socket vSphere 8 VM vs. Cascade Lake, 8-socket vSphere 8 VM

SAP BWH L-Class KPIs					25,000		5,000		200	
Config	Cert	CPU Threads	Mem	Records	BWH Phase 1 (sec.)	Delta	BWH Phase 2 (QPH)	Delta	BWH Phase 3 (sec.)	Delta
8-Socket Intel Xeon Platinum 8280L, 28 core CPU - VMXNET3 vSphere 8 VM	-	448	5,890 GB	10.4 billion	21,423	-	5,257	-	172	-
4-Socket Intel Sapphire Rapids 8490H, 60 core CPU - VMXNET3 vSphere 8 VM	-	480	8,000 GB		12,848	NC	7,111	35.27%	144	NC
					- is better		+ is better		- is better	

Using the latest Sapphire Rapids hosts with vSphere 8.0 U1 or later results in significant performance gains over previous CPU generations running the same SAP HANA workload. This not only allows for faster processing, but also larger memory configurations and a smooth transition from older CPU generations to Sapphire Rapids hosts. Virtualization costs are mostly offset by the performance boost provided by a Sapphire Rapids CPU. As a result, even bare-metal SAP HANA installations can be upgraded to a virtualized Sapphire Rapids host, requiring less data center space.

When transitioning from an 8-socket Cooper Lake or Cascade Lake system, which can accommodate up to 12 TB of memory, to a 4-socket Sapphire Rapids system, it's essential to note that the standard memory support for the latter is limited to 8 TB.

However, with the support of TDI (Technology Development Interface), a 4-socket Sapphire Rapids system can potentially handle higher memory capacities. SAP allows up to 3 TB per CPU, indicating that a 4-socket Sapphire Rapids configuration could support up to 12 TB of memory. Moreover, given the enhanced performance capabilities of the CPU, you could seamlessly perform such an upgrade.


## 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 its 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 operating system and application, or to perform a lengthy backup. A change in hardware would also add time because you'd need to update the device drivers in the operating system too.

vSphere vMotion live migration (figure 21) allows you to move an entire running VM from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity. The VM retains its network identity and connections, ensuring a seamless migration process. The VM's active memory and precise running state can be transferred over a high-speed network, allowing the VM to switch from running on the source vSphere host to the destination vSphere host.

Figure 21: vSphere vMotion allows the live migration of SAP HANA VMs



With enhanced vMotion compatibility mode, you can migrate SAP HANA VMs between vSphere hosts with different generations of CPUs. This lets you combine older and newer server hardware generations into a single cluster, thereby saving money because you don't need to replace the older vSphere hosts.

**Recommended:** We recommend you migrate SAP HANA between hosts with the same CPU type, model, and frequency to ensure high performance of SAP HANA on a vSphere cluster. You should limit the migration of SAP HANA VMs between hosts with different CPUs to situations such as hardware upgrades or HA.



You can move a VM from one vSphere host to another in different modes. vSphere supports some modes for a powered-on SAP HANA VM. Other modes, 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 practices for migrating SAP HANA VMs in a production environment

vMotion between different hardware generations of a CPU type or storage subsystems is possible.

**Important:** When migrating VMs for a performance-critical application like SAP HANA, you should follow these best practices:

- Run 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.
- Don't live migrate SAP HANA VMs while a virus scanner or a backup job is running inside the VM or while people are using the SAP HANA VMs because this can cause a soft lock of the SAP HANA application.
- Don't live migrate SAP HANA VMs while a VM snapshot-based backup job is running outside the VM this can cause a soft lock or data inconsistencies of the SAP HANA application.
- Use vMotion only during non-peak times (for example, during CPU utilization of less than 25%).
- If the source host and destination host clock speeds are different, you can still use vMotion. However, you should plan to upgrade the VM's hardware version and align the vCPUs to the new CPU so that you can restart the VM and use any new CPU features of the target host.
- You may use vSphere Storage vMotion to migrate SAP HANA VMs between storage subsystems. Because vSphere Storage vMotion impacts the performance of a running VM, we strongly recommend you perform a storage migration while the VM is powered off, or at least while the SAP HANA database is shut down inside the VM.
- Allocate sufficient bandwidth to the vMotion network, ideally 25 GbE or more.
- Avoid having "noisy neighbors" active on the vSphere host during a vMotion migration of SAP HANA VMs. A noisy neighbor is another VM that is using up all of the host's resources, leaving little left for other activities.
- Check SAP HANA patch levels (some patch levels may increase the risk of soft lockups to the operating system during migrations).
- Upgrade to vSphere 7 or 8 to leverage vMotion improvements.

**Caution:** While vMotion is a great tool that helps you manage and operate production SAP HANA VMs, be very careful migrating SAP HANA VMs because doing so may cause severe performance issues that impact SAP HANA users and long-running transactions.


## VM migration scenarios

With vSphere DRS, a vMotion migration can be run manually, fully, or semi-automated. To lower the possible impact on SAP HANA VMs during a VM migration, we suggest that you:

- Use vMotion only during non-peak times (low CPU utilization: less than 25%).
- Have vSphere DRS rules in place that:
  - Only suggest initial placement or
  - Allow VMs to be automatically moved when a host is set to maintenance mode.
- 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,  $\geq 4$  GB SAP HANA VMs. A vMotion network with 25 GbE or higher bandwidth is preferred. Multiple vMotion network cards will help parallelize the vMotion process and lower the impact on the VM performance and time.

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

Figure 22: Live migration of an SAP HANA VM (default scenario)




### Description:

- Typical VM migration scenario mainly used for load balancing or to transfer VMs off the host for server maintenance.
- All hosts have the same vSphere version and identical HW configuration.
- Hosts are connected to a shared VMFS datastore and are in the same network.

### Considerations:

- Perform a manual vMotion migration during non-peak hours.
- Enhanced vMotion Compatibility is not required because all hosts use the same CPU.

Figure 23: Transferring VMs off the vSphere host to a new host with a new ESXi version



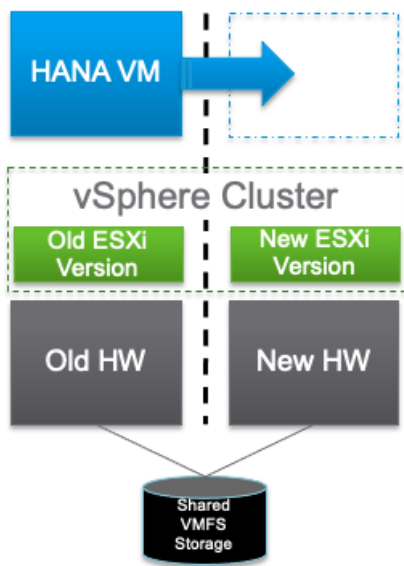
#### Description:

- VM migration scenario to transfer VMs off a vSphere host (set to maintenance mode) before a vSphere upgrade.
- All hosts have an identical HW configuration.
- Hosts are connected to a shared VMFS datastore and are in the same network.

#### Considerations:

- Perform a manual vMotion migration during non-peak hours.
- Enhanced vMotion Compatibility is not required because all hosts use the same CPU.
- Live migration is possible, but double-check the virtual hardware version first:
  - Before you upgrade the virtual hardware, you might need to upgrade VMware Tools first. Double-check the dependencies in the release notes.
  - You might need to upgrade the virtual hardware of the VM; for example, hardware version 16 to 21.
  - You might need to align the virtual hardware to the new vSphere maximums; for example, 6 TB VM now to 12 TB.
  - You must restart the VM after you change its virtual hardware.

Figure 24: Migrating a SAP HANA VM during a host hardware upgrade

**Description:**

- VM migration scenario to migrate a VM to a new host with a different CPU (CPU type or new CPU generation). This is a critical vMotion scenario (see below).
- Hosts are connected to a shared VMFS datastore and are in the same network.

**Considerations:**

- Perform a manual vMotion migration during non-peak hours.
- EVC is required to allow the use of hosts with different CPUs in one vSphere cluster.
- Online migration of VMs is possible, but if the hardware has changed, then the VMs need to be aligned to the changed HW and CPU configuration.
  - Plan offline time for VM maintenance to perform a virtual hardware upgrade of the VM (for example, hardware version 20 to 21).
  - Align the number of vCPUs per socket and VM memory to the physical CPU capabilities and NUMA node memory configuration.
  - Power on the VMs on the new host.
- If the memory and CPU core configurations remained the same, then a simple restart of the VMs is enough to ensure the timestamp counter is synchronized; otherwise, the VM performance could be negatively impacted (more on this below).


If the hardware configuration of the host (CPU cores and/or memory or number of sockets) has changed, you need to power off the VM and change the VM configuration to align it to the new host configuration.

**Critical vMotion scenario:** If the hardware configuration stays the same, but the CPU frequency was changed, you need to reboot the VM. This is a critical vMotion scenario due to the different possible CPU clock speeds, plus the exposed timestamp counter (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 the RDTSC instruction relative to a bare metal system.

**Background:** When a VM 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 vMotion), a rate transformation is performed so the virtual TSC continues to run at its original power-on rate, not at the new host TSC rate. For details, read the document [Timekeeping in VMware Virtual Machines](#).


To solve this issue, you must plan a maintenance window, so you can restart the VMs that were moved to the non-identical hardware to allow the use of hardware-based TSC instead of using software rate transformation on the target host, which is expensive and will degrade VM performance. Figure 25 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 25: The maintenance window as a hardware upgrade and vMotion



\*Business downtime required

Figure 26: VMware VM migration – storage migration (VM datastore migration)



**Description:**

- VM migration scenario to migrate a VM to a new host with a different VMFS datastore.
- Migration between local and shared or different shared storage subsystems possible.
- Hosts are connected to the same network.

**Considerations:**

- Requires manual vMotion during off-peak hours.
- EVC might be required if the hosts have different CPUs installed.
- Online migration is possible, but in the SAP HANA case, we don't recommend a migration with Storage vMotion.
- Plan offline time for VM maintenance to perform the storage migration. Storage vMotion takes more time and has a higher impact on VM performance.
- Ensure the new storage meets the SAP HANA TDI storage KPIs.

## Managing SAP HANA landscapes using vSphere DRS

vSphere DRS, an automated load balancing technology that aligns resource usage with business priority, can be used to manage SAP HANA landscapes. 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 VMs among a set of ESXi hosts. vSphere DRS continuously monitors utilization across ESXi hosts and can migrate VMs to hosts that are less utilized if a VM needs more resources.

When deploying large SAP HANA databases or production-level SAP HANA VMs, it is essential to have vSphere DRS rules in place and to set the automation mode to manual. If you don't set the automation mode to manual and use the automated mode, then you must set the DRS migration threshold to conservative (level 1), to avoid unwanted migrations, which may negatively impact the performance of an SAP HANA system. It is possible to define which SAP HANA VMs should be 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** – DRS recommends the initial placement of a VM within the cluster, and then recommends the migration. You must manually place and migrate the VM.
- **Partially automated** – DRS automatically places a VM when it is being powered on, but during load balancing, DRS displays a list of vMotion recommendations for you to select and apply.
- **Fully automated** – DRS placements and migrations are automatic.

**Note:** DRS requires a vSphere Clustering Service VM and automatically installs this VM in a vSphere cluster. If you are using VMware SDDC manager to manage and update your environment, then DRS must be set to fully automated.

## SNC-2 management and operation

Building homogenous vSphere SAP HANA clusters helps to streamline operation and eliminate possible issues when an SAP HANA VM is started on another host in the cluster. You must use the same CPU model and type, size of host memory, installed adapters like HBA or NICs, and storage configuration. Adding SNC-2-enabled 2-socket Sapphire Rapids hosts to a cluster provides some benefits, like the co-deployment of other types of VMs (not SAP HANA VMs) on the same physical CPU socket. But it makes the operation of such an environment more challenging. To make your job easier when adding SNC-2 enabled hosts, use VM host rules.

This VM-specific SNC-2 VM/host rule must be applied to all VMs that leverage SNC-2. This ensures the SNC-2-enabled VMs do not get migrated or started on non-SNC-enabled hosts. Figure 27 shows an example VM/host rule that ensures VMs created on SNC nodes run only on SNC-enabled hosts.

Figure 27: VM SNC-2 host rule

**Edit VM/Host Rule | DZ10** [X]

Name:   Enable rule.

Type:

Virtual machines that are members of the Cluster VM Group SNC\_VM\_test must run on host group SNC\_test.

VM Group:

Host Group:

[CANCEL] [OK]

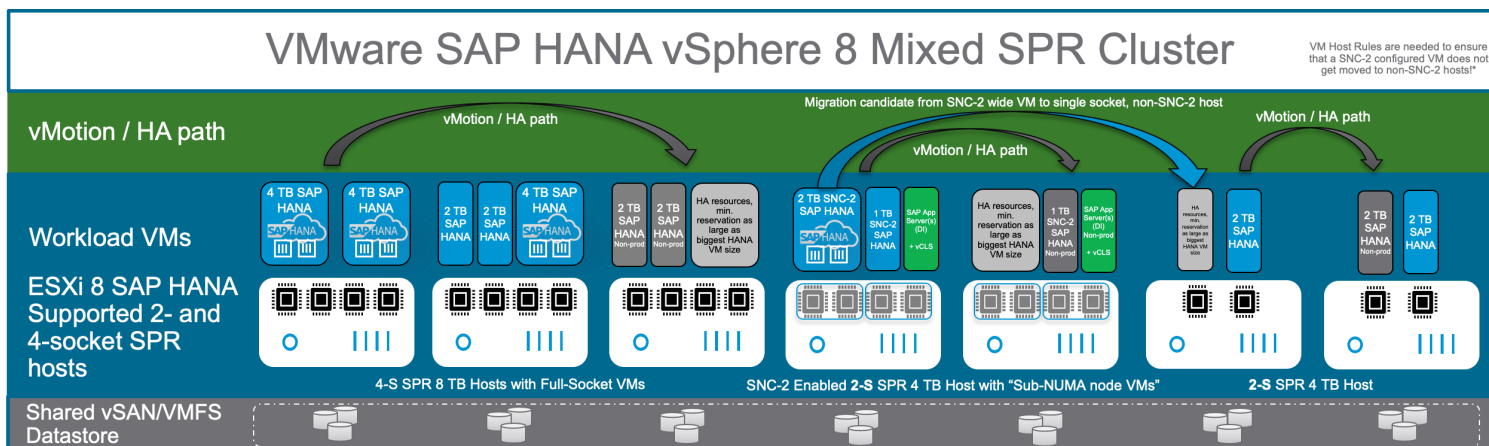
Figure 28 shows an example cluster with Sapphire Rapids-based 2- and 4-socket systems.

The 4-socket systems are used for the larger SAP HANA VMs that require more memory. All hosts adhere to the SAP HANA memory standard configuration of 2 TB per CPU. OLAP workloads running on these systems, alongside the OLTP workload VMs, are verified with an SAP HANA OLAP expert sizing to ensure that 2 TB per CPU can be utilized.

**Attention:** VM host rules are needed to ensure that half socket (SNC-sized) VMs do not get started on non-SNC-2 hosts. Figure 28 depicts the typical vMotion or HA path for a VM if it needs to be migrated due to maintenance or an HA event.

You should consolidate smaller SAP HANA VMs, in this case up to 1 TB, on the two SNC-2 hosts to benefit from optimized memory latencies. If a SNC-2 sub-NUMA node VM needs to be extended, you should migrate it to one of the 2- or 4-socket non-SNC-2 hosts as shown in the figure to get the best possible improvements of a single NUMA node configuration with the same size.

Figure 28: VMware SAP HANA vSphere 8 mixed Sapphire Rapids (SPR) cluster



**Important:** Due to the lower memory bandwidth associated with an SNC NUMA node, we recommend you primarily use SNC-2 for half-socket VMs. In the event of performance issues with SAP HANA VMs running on SNC-2 enabled hosts, you should migrate these VMs to a non-SNC enabled host to resolve SNC-related performance issues. All VMs that leverage SNC-2 require the VMX advanced parameter `sched.nodeX.affinity="Y"` to prevent unwanted NUMA node migrations.

#### Guidelines for deploying an SAP HANA VM on an SNC-2 enabled Sapphire Rapids ESXi host:

- SNC requires that the ESXi host memory be symmetrically populated.
- Enable SNC-2 and hyperthreading in the BIOS of the ESXi host.
- Use SNC-2 for SAP HANA VMs exclusively on 2-socket SPR hosts or later.
- Size the SAP HANA VM according to available logical threads and memory per CPU socket.
- Apply the `sched.nodeX.affinity="Y"` VMX advanced parameter to all VMs that leverage SNC-2 to prevent unintended NUMA node migrations.
- Use SNC-2 primarily as a consolidation platform for half-socket VMs or SAP HANA VMs requiring low memory latency due to lower memory bandwidth associated with an SNC NUMA node.
- SAP HANA VMs within a 2-socket SNC-2 host can be extended up to 4 sub-NUMA nodes (refer to figure 28, above, for details).
- Address SNC-2 performance issues by offline migrating SAP HANA VMs from SNC-2-enabled hosts to non-SNC-enabled hosts and adjusting the VM configurations accordingly to the non-SNC configuration.
- The vMotion of VMs that leverage SNC-2 is only supported between SNC-2 enabled hosts. Avoid migrating SNC-2-configured VMs to non-SNC-2 hosts to prevent performance issues. Migration in the reverse direction is also unsupported for the same reason.
- Use vMotion host rules to prevent the migration of SNC-2 VMs to non-SNC-2 ESXi hosts.



## vSphere Cluster Services

Starting with [vSphere 7.0 Update 1](#), vSphere Cluster Service (vCLS) is enabled by default and runs in all vSphere clusters. vSphere makes critical cluster services, such as vSphere HA and vSphere DRS, highly available, and vCLS makes cluster services even more available.

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, 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 vCLS, which depends on vCenter Server availability for configuration and operation.

The dependency of these cluster services on vCenter Server is not ideal, and the use of vCLS is the first step to decouple and distribute the control plane for clustering services in vSphere and to remove the vCenter Server dependency. If vCenter Server should ever become unavailable, vCLS will ensure that the cluster services remain available to maintain the resources and health of critical workloads.

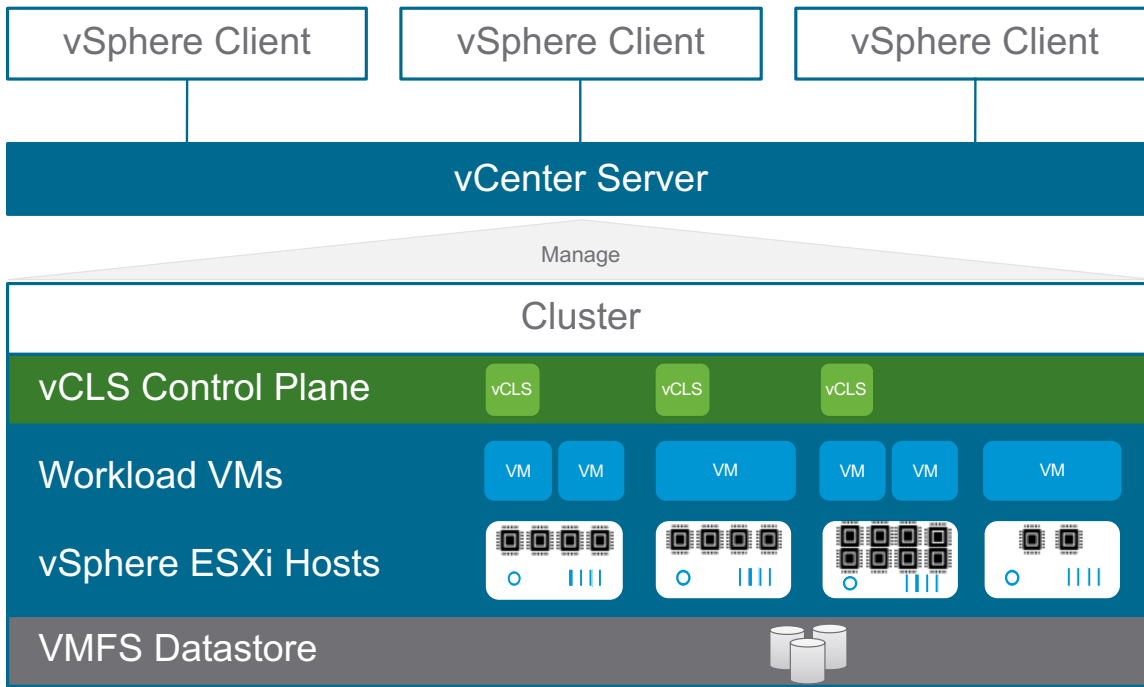
**Note:** vCLS is enabled when you upgrade to vSphere 7.0 Update 1 or when you have a new vSphere 7.0 Update 1 deployment. vCLS is automatically deployed as part of a vCenter Server upgrade, regardless of which ESXi version is used.

### Architecture of vCLS

vCLS uses agent VMs to maintain cluster service health. Up to three vCLS agent VMs, which are lightweight and build the cluster control plane, are created when you add hosts to clusters.

vCLS VMs must run and are distributed in each vSphere cluster. vCLS is also enabled on clusters that contain only one or two hosts. In these clusters, there are only two vCLS VMs. Figure 26 shows the high-level architecture with the new cluster control plane.

Figure 29: vCLS high-level architecture



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

## vCLS VM details

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

Each vCLS VM has 100 MHz and 100MB capacity reserved in the cluster. For more details, refer to the [Monitoring vSphere Cluster Services](#) documentation.

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

Table 29: vCLS VM resource allocation

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

Table 30: Number of vCLS agent VMs in clusters

Number of hosts in a cluster	Number of vCLS agent VMs
1	1
2	2
3 or more	3

## Guidelines for deploying vCLS in a cluster with SAP HANA VMs

As of SAP notes [3102813](#) and [3372365](#), you can't run a non-SAP HANA VM on the same NUMA node where an SAP HANA VM already runs: "SAP HANA VMs can be 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 is allowed."

**Note:** These statements do not apply to a 2-socket Sapphire Rapids platform with SNC-2. Here, it is possible to co-deploy a non-SAP HANA VM like an SAP application server or a VMware vCLS VM. Co-deployment of other VMs, like a web server, is still not supported.

All other CPU generations and >2-socket Sapphire Rapids systems have restrictions because of SAP guidelines and the mandatory and automated installation process of vCLS VMs. When upgrading to vCenter 7.0 Update 1, you must check if vCLS VMs were co-deployed on ESXi hosts that run SAP HANA production-level VMs. If this is the case, then you must migrate these VMs to hosts that do not run SAP HANA production-level VMs or to another SNC-2 CPU socket host when SNC-2 is used.

You'll need to configure [vCLS VM anti-affinity policies](#). These policies describe a relationship between VMs that have been assigned a special anti-affinity tag (for example, a tag named "SAP HANA") and vCLS VMs.

When you assign this tag to SAP HANA VMs, the policy discourages placement of vCLS VMs and SAP HANA VMs on the same host. After you create the policy and assign tags, the placement engine will attempt to place vCLS VMs on the hosts where tagged VMs are not running—for example, the HA ESXi host. For this to work, you need to have some VMs on the cluster that are not tagged with "SAP HANA".

## Examples of deploying vCLS in a cluster with SAP HANA VMs

Typically, you deploy SAP HANA VMs 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 cluster.

The following examples of typical SAP HANA clusters provide some guidelines on where to place up to three lightweight vCLS VMs.

### *Mixed SAP HANA and other VMs in a vSphere cluster*

A mixed cluster with SAP HANA VMs and other VMs is a typical scenario. In this case, check the vCLS VMs to see if they were deployed on ESXi hosts that run production-level SAP HANA VMs. If yes, then the vCLS VM can run on the same CPU socket as an SAP HANA VM.

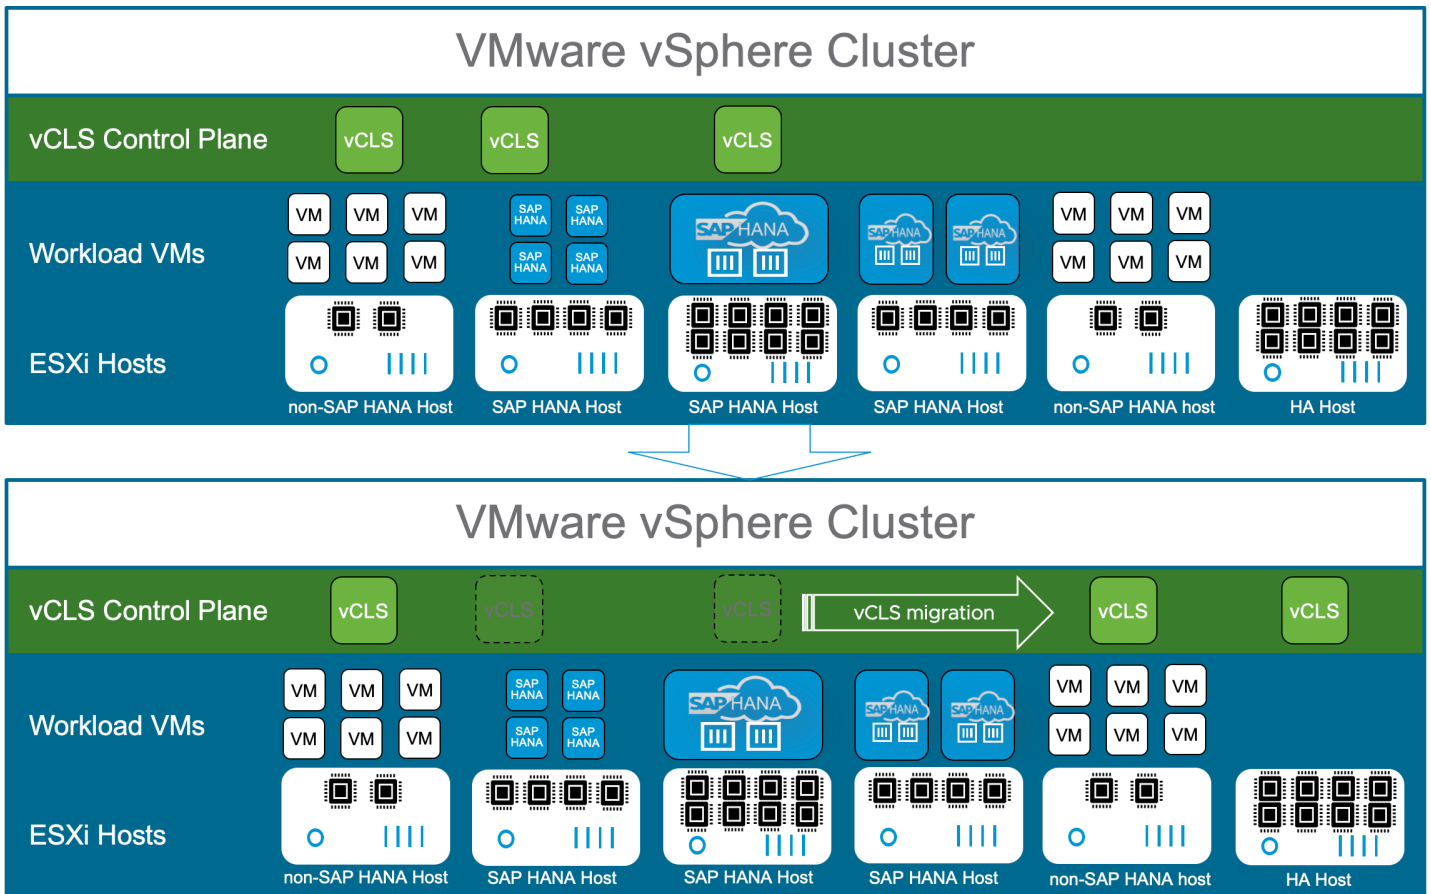
To avoid this, [configure vCLS anti-affinity policies](#):

1. Create a category and tag for each group of VMs that you want to include in a vCLS VM anti-affinity policy.
2. Tag the VMs that you want to include.
3. Create a vCLS 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 30 shows the initially deployed vCLS VMs and how these VMs get automatically migrated (green arrows) when the anti-affinity rules are activated to comply with SAP notes [3102813](#) and [3372365](#). Not shown in this figure are the HA host/HA capacity reserved for HA failover situations.

**Note:** If you need to add 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 30: vCLS VM migration for a mixed host cluster: migrate the vCLS VMs to an ESXi host that runs workloads other than SAP HANA VMs




### *Dedicated ESXi hosts for SAP HANA VMs on a vSphere cluster*

You might have deployed a vSphere cluster with ESXi hosts that run only SAP HANA VMs. In this case, the automatically deployed vCLS VMs will be on a SAP HANA host and cannot be easily migrated 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 a SAP application server VM or infrastructure workload VMs. You don't need to buy a new host for this.

Figure 31 shows the initially deployed vCLS VMs and how these VMs are moved when the vCLS anti-affinity policy for SAP HANA VMs is added to new hosts.

Figure 31: Migrating vCLS VMs in a dedicated SAP HANA host cluster



**Important:** To allow the vCLS VM to run on the same host as a non-production SAP HANA VM, make sure you have not tagged the non-production SAP HANA VMs with a name tag that triggers the anti-affinity policy.

## SAP HANA HCI on a vSphere cluster

Just as with the dedicated SAP HANA cluster, an SAP HANA HCI cluster may only run SAP HANA workload VMs in a VCF workload domain. 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 is upgraded to 7.0 U1, then the vCLS VMs will be automatically deployed on SAP HANA HCI nodes. If these nodes are exclusively used for SAP HANA production-level VMs, then the vCLS VMs must be removed and migrated to the vSphere HCI HA host (or hosts).

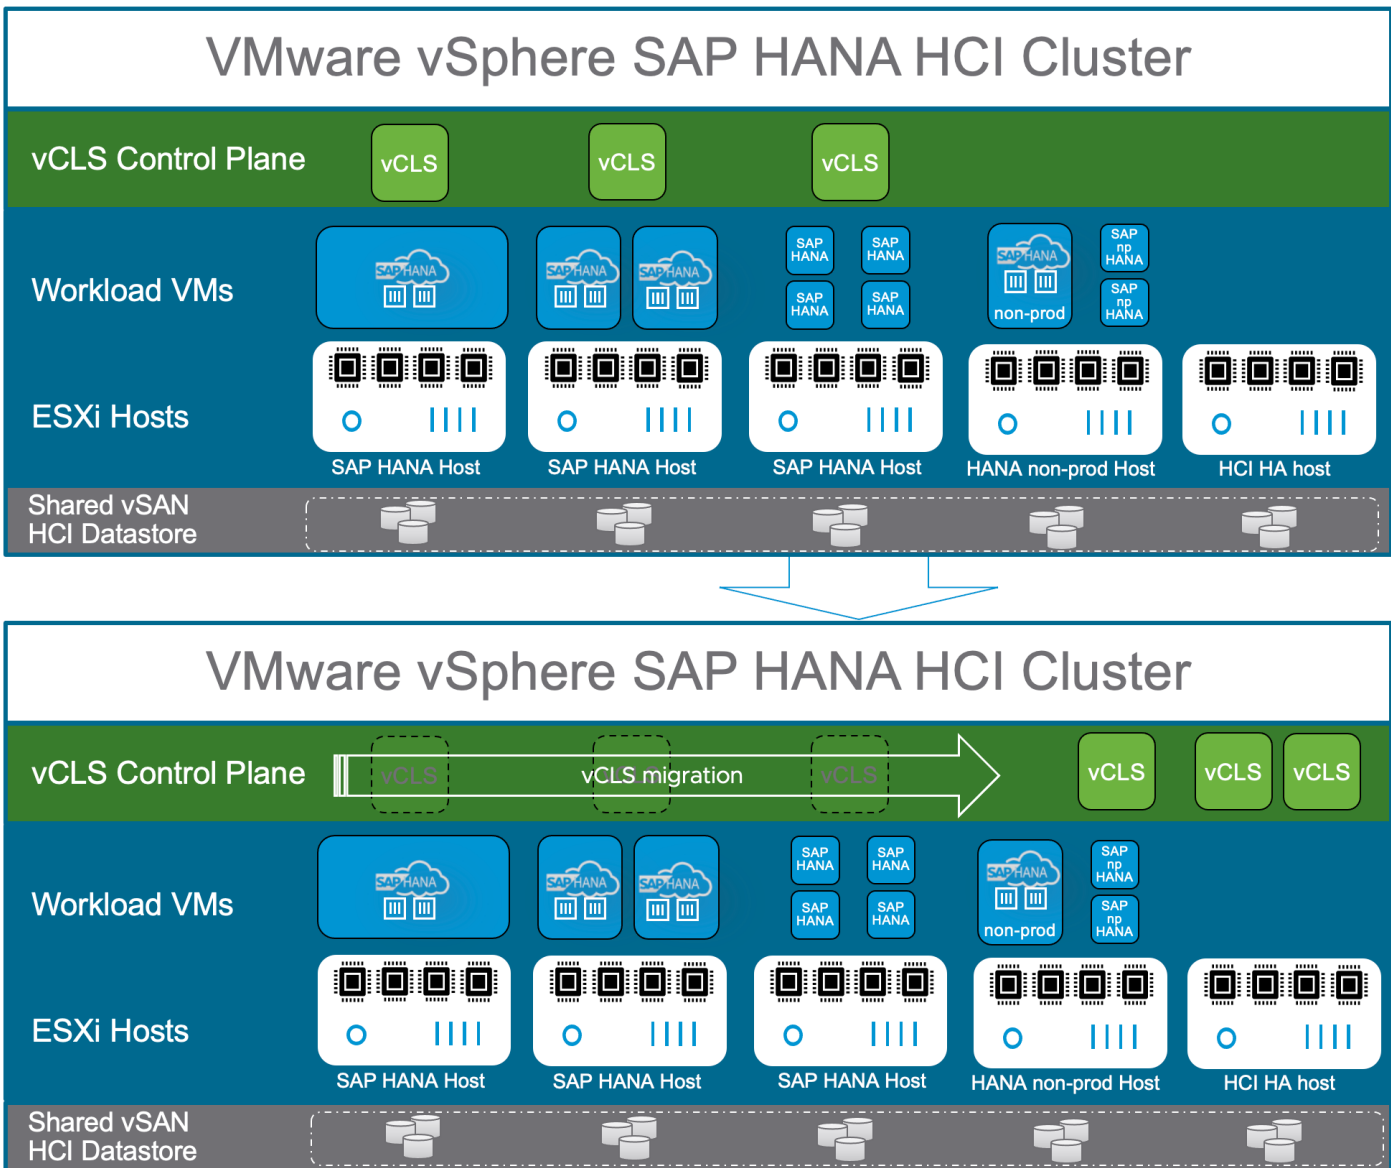
You can do so by configuring [vCLS VM Anti-Affinity Policies](#). A vCLS VM anti-affinity policy describes a relationship between VMs that have been assigned a special anti-affinity tag (for example, tag name "SAP HANA") and vCLS system VMs.

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, vCLS VMs and SAP HANA VMs are not co-deployed. After you create the policy and assign tags, 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 a SAP HANA HCI partner system validation or if an additional non-SAP HANA ESXi host cannot get added to the cluster, then you can use [Retreat Mode](#) to remove the vCLS VMs from this cluster.

**Caution:** Retreat Mode will impact certain cluster services, such as DRS.

Figure 32: vCLS VMs in an SAP HANA HCI vSphere cluster should be migrated off servers hosting SAP HANA VMs that are in a production environment



**Important:** To allow the vCLS VMs to run, as shown in figures 31 and 32, on the same host as a non-production SAP HANA VM, you must not tag the non-production SAP HANA VM with a name tag that triggers the anti-affinity policy.

In summary, by introducing vCLS, 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, please 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](#)

**Note:** In the case of 2-socket Sapphire Rapids hosts get used for vSphere or HCI (vSphere + vSAN) deployments then SNC-2 can get used and a vCLS VM could get moved to an SNC-2 sub-NUMA Node, instead to add an additional host for vCLS offloading.

## High availability best practices for virtualized SAP HANA

SAP HANA offers several methods for high availability and disaster recovery: auto-failover, service restart options, backups, system replication, and standby host systems. In VMware virtualized environments, you can use any of these solutions. In addition, you can use vSphere HA and vSphere Replication to minimize unplanned downtime due to faults.

There are two areas of high availability support: fault recovery and disaster recovery.

Fault recovery includes:

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

Disaster recovery includes:

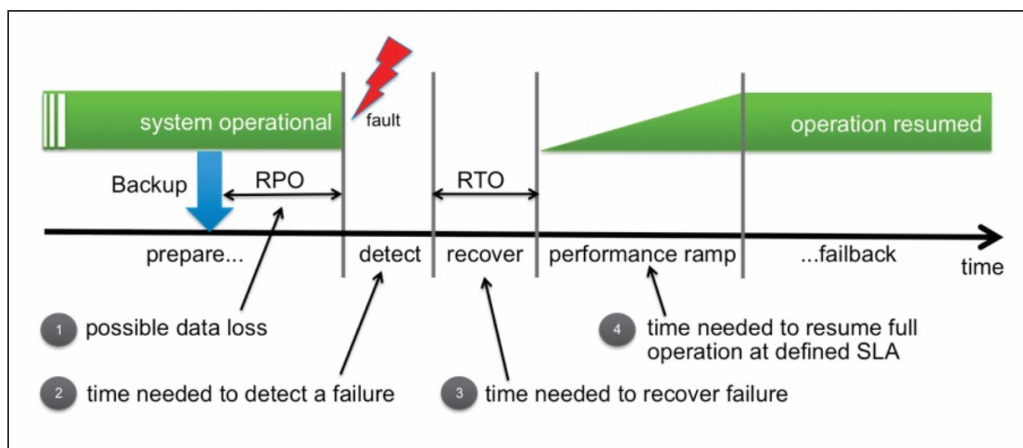
- Backup and restore
- Storage replication
- vSphere Replication
- SAP HANA system replication



We include SAP HANA system replication (HSR) in both recovery scenarios. Depending on your requirements, SAP HANA system replication can be used as a failover or disaster recovery solution when site or data recovery is needed. Because HSR allows data to be preloaded into the replication instance's memory, it can also reduce the startup time for large SAP HANA databases.

You can assign different recovery point objectives (RPOs) and recovery time objectives (RTOs) to different fault recovery and disaster recovery solutions. SAP describes the phases of high availability in their HANA HA document. Figure 33 shows a graphical view of these phases.

Figure 33: SAP HANA system availability phases



- **RPO (1 - possible data loss)** 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.
- **Detect (2 – time needed to detect a failure)** 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, recover (3 - time needed to recover failure)** 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.
- **Performance ramp (4 – time needed to resume full operation at defined SLA)** 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 be 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.

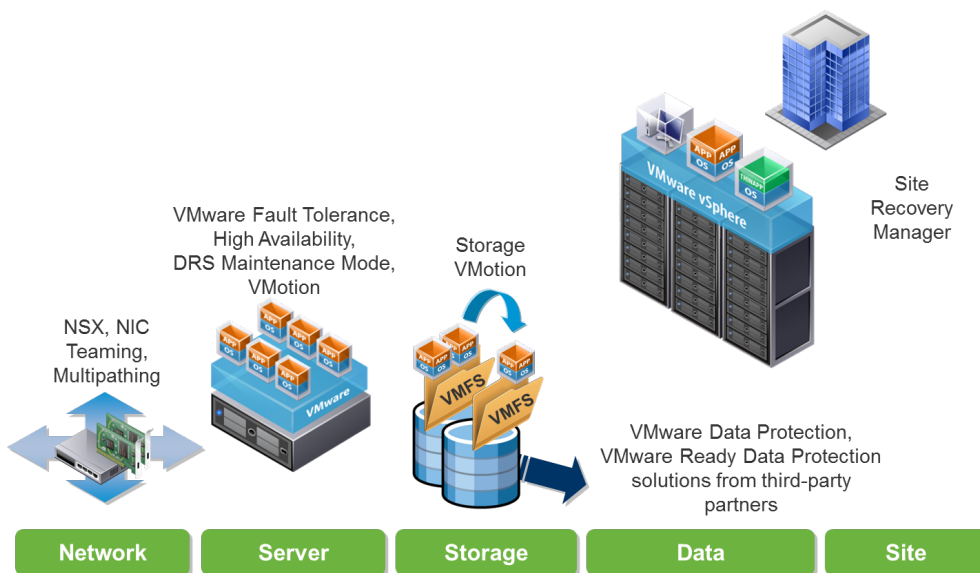
## About vSphere high availability (HA)

VMware provides vSphere built-in and optional availability and disaster recovery solutions to protect a virtualized SAP HANA system at the hardware and operating system 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 tolerate failures in memory models or that ensure CPU transaction consistency.

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

Figure 34 shows different HA solutions to protect against component-level failures, up to a complete site failure, which can be 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 34: VMware HA and DRS solutions provide protection at every level




[vSphere HA](#) is a fault recovery solution and provides uniform, cost-effective failover protection against hardware and operating system outages within a virtualized IT environment. It does this by monitoring vSphere hosts and VMs to detect hardware and guest failures. It restarts VMs on other vSphere hosts in the cluster without manual intervention when a server outage is detected, and it reduces application downtime by automatically restarting VMs upon detection of an operating system failure. This, combined with the SAP HANA service auto-restart feature, allows HA levels of 99.9% out of the box.<sup>16</sup>

<sup>16</sup> EMC IT, 02/14 EMC Perspective, H12853.

Figure 35 shows how vSphere HA can protect against VM or host failure and the application protection solution, such as SAP HANA service auto-restart, third-party in-guest cluster solutions, or SAP HANA system replication can also provide disaster recovery capabilities. All these solutions can be combined with vSphere HA.

Figure 35: Virtualized SAP HANA HA solution



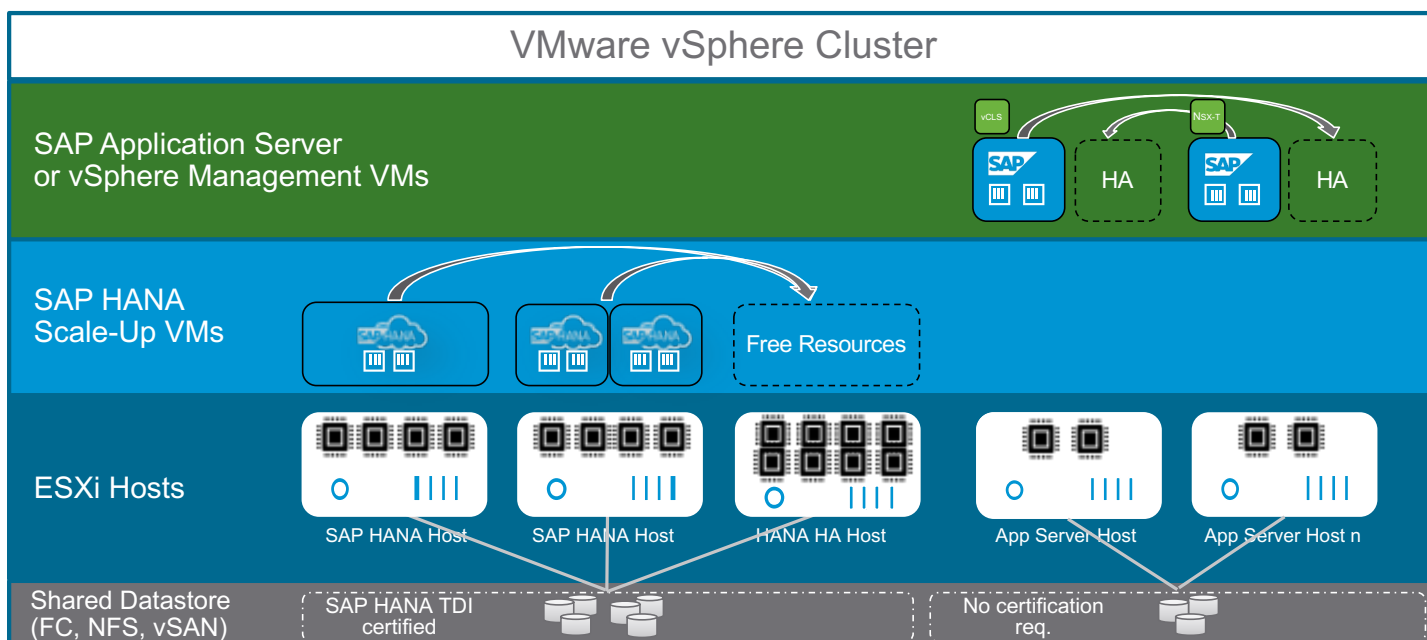
vSphere HA protects SAP HANA scale-up and scale-out deployments without any dependencies on external components, such as DNS servers; or solutions, such as the SAP HANA Storage Connector API.

Figure 36 shows the how critical SAP HANA VMs and non-SAP HANA VMs that run on different, lower cost server systems can leverage vSphere HA, and how a typical n+1 vSphere cluster can be configured to survive a complete host failure.

This vSphere HA configuration is used by SAP applications and SAP HANA instances on vSphere. If higher redundancy levels are required, then an n+2 configuration can be used. Non-critical VMs can also leverage the HA resource pool; they need to be powered off before an SAP HANA or SAP app server VM is restarted.

The vCLS control plane and the related vCLS VMs are not shown in the following figures.

Figure 36: vSphere HA protected SAP HANA VMs in an n+1 cluster configuration



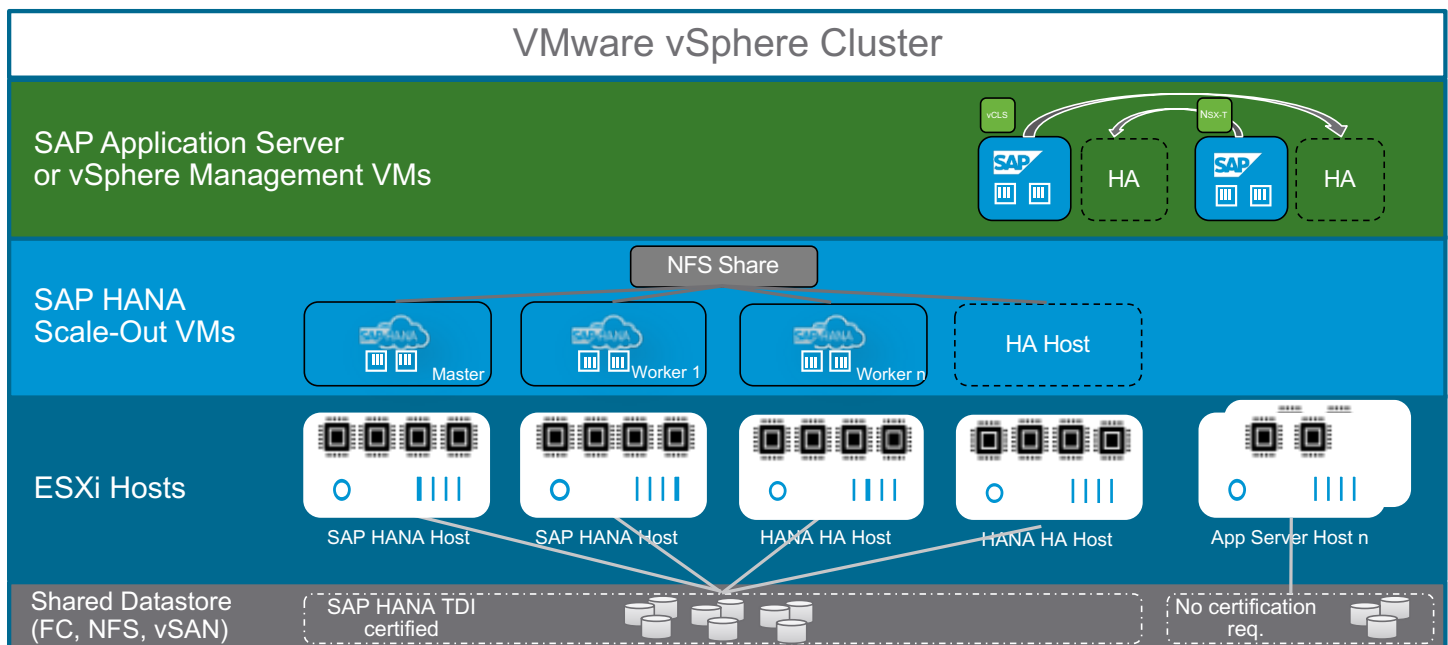
It is also possible to configure the HA cluster as active-active, where all hosts have SAP HANA VMs deployed. This ensures that all hosts of a vSphere cluster are used by still providing enough failover capacity for all running VMs if there is a host failure. The arrow in the figure indicates that the 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 be 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 failover with the VM that has failed. The access to the HANA shared directory 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 37 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 for a host failure.

Figure 37: 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 is used for other workloads, then all potentially running VMs on this host must be terminated or migrated to another host before a failed HANA scale-out VM can be 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 3 TB on 4-socket and 6 TB on 8-socket VMs are generally available (GA) as of today. Supported hosts are only 4- and 8-socket large host systems (Intel Broadwell CPU and newer). SAP does not support 2-socket host systems for on-premises scale-out deployments. Review the mentioned SAP HANA on vSphere support notes for more details.

**Note:** vSphere HA can only protect 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 be automated, then a third-party solution, such as SUSE HA, needs to be implemented. Protecting a physical SAP HANA deployment against host failures is therefore relatively complex, whereas protecting a VMware virtualized SAP HANA system is just a mouse click.

If fast failure recovery and data replication are required, then we recommend you use HANA System Replication (HSR) in combination with a SAP-supported Linux cluster solution like Pacemaker. Because HSR replicates SAP HANA data, it can be used for disaster recovery or for recovering from logical errors (depending on the log retention policy).

## vSphere HA with passthrough network adapters

If the VMXNET3-caused latency is too high for a specific use case/workload, we recommend you use a passthrough NIC configuration.

To enable VMware HA with passthrough NICs, it must be 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 be used for the passthrough NIC installed in the HA host. If no HA host with a passthrough NIC is configured as a dynamic vSphere DirectPath I/O device and the same hardware label exists, then the HA failover process won't work.

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

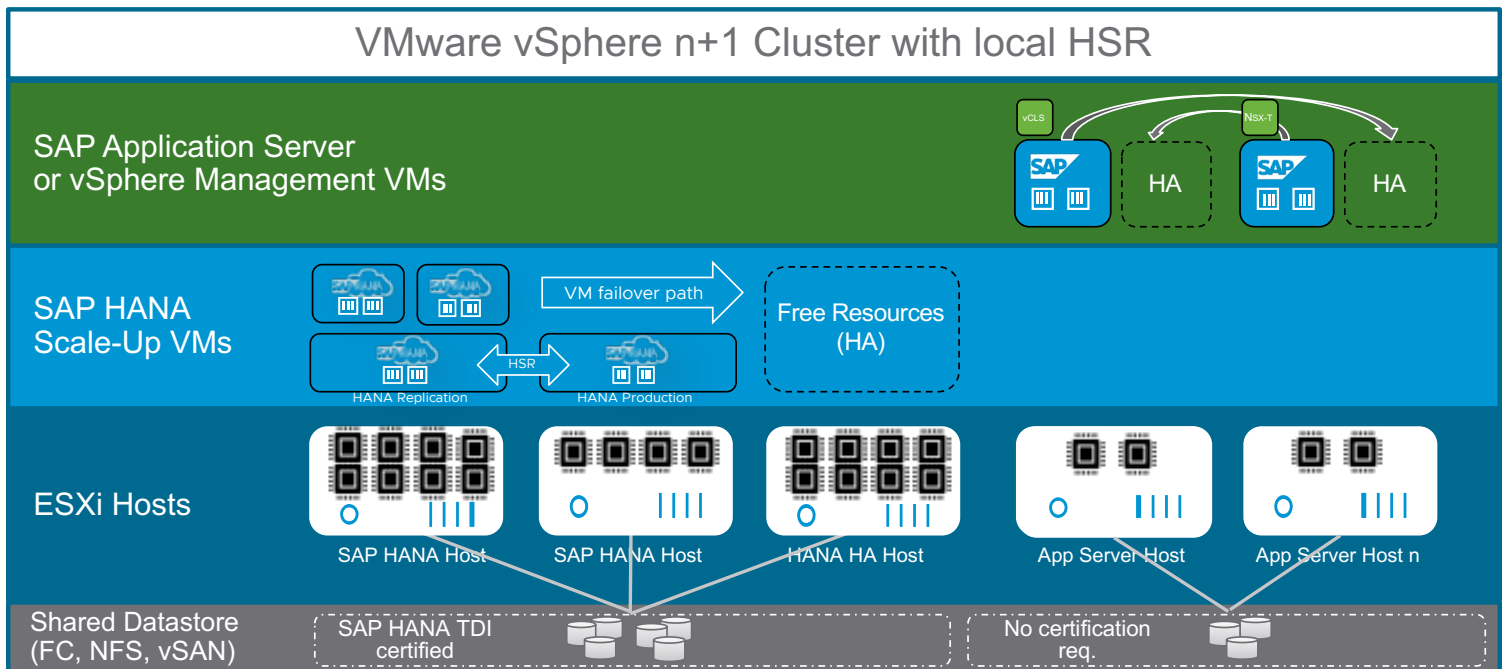
## SAP HANA system replication with vSphere (local site)

vSphere HA can be combined with [HANA System Replication \(HSR\)](#) to protect SAP HANA data against logical or disastrous failures that impact a data center.

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 failures. HSR requires a running SAP HANA replication VM, which must receive HSR data. Alternatively, you can use storage subsystem-based replication, which would be independent from SAP HANA.

Figure 38 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 failures, as showed in Figure 38. If it runs in the same location, then HSR can be used to recover from logical failures (if logs are 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 38: vSphere HA with local data venter HANA system replication (HSR)



As noted, HSR does not provide automated failover. You must manually reconfigure the replication target VM to the production system’s identity. Alternatively, you can use third-party cluster solutions such as SAP HANA HA Linux solutions from [SUSE](#) or [Red Hat](#), or SAP Landscape Management to automate the failover to the SAP HANA replication target.

**Note:** You can combine SAP HSR with vSphere HA to protect the HSR source system against local failures.

To provide disaster tolerance, you must place the HSR replica VM/host on another data center or even a geographically dispersed site.

## Distaster recovery

It is possible to combine the discussed HA solutions with storage/vSphere Replication and HSR to another data center/site if data protection is also a priority or if a complete site failover of all IT systems—including SAP HANA—is necessary.

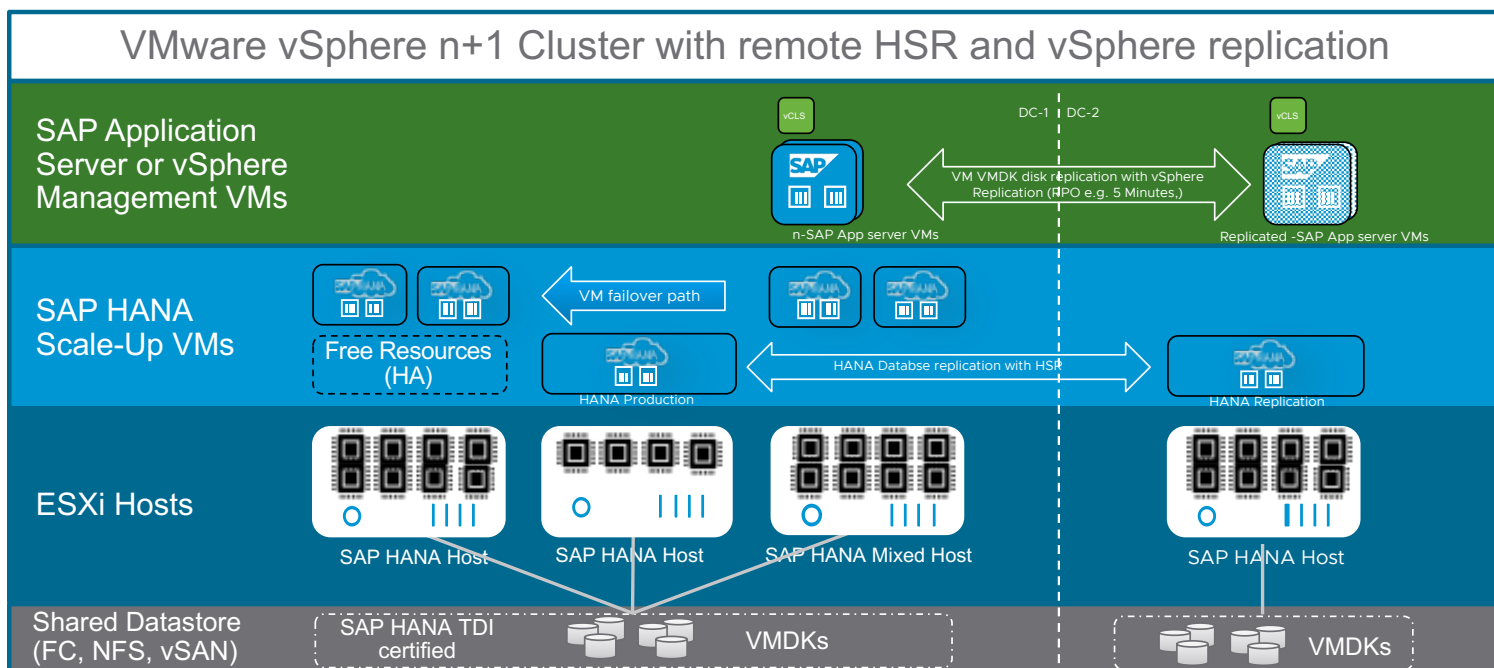
In addition to this, you also need backup and restore solutions in place for regulatory reasons and to protect the data against logical errors.

## SAP HSR and vSphere Replication for a remote site

Figure 39 shows an HSR-protected SAP HANA instance. It is the same concept as discussed in figure 38, except 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.

Some replication requirements (synchronous or asynchronous) might call for a roundtrip time (RTT) below 1 millisecond to maintain the SAP HANA storage KPIs. If a 1 millisecond RTT is not possible, then you should use asynchronous replication to ensure that the data replication doesn't negatively impact the production SAP HANA system.

Figure 39: vSphere HA with remote data center SAP HANA system replication (HSR)



The example in figure 39 shows the HSR replication from a virtualized SAP HANA system to another virtualized SAP HANA system. You can leverage [vSphere Replication](#) to replicate the SAP app server VMs to the DT side. This will allow operation to continue after a switch to this data center. 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 recovery point objective (RPO) times from 5 minutes to 24 hours, and the VM replication process is nonintrusive and takes place independent of the OS or applications in the VM. It is transparent to protected VMs and requires no changes to their configuration or ongoing management.



**Note:** SAP HANA performance is directly impacted by round trip time (RTT). If the RPO target is 0, then synchronous replication is required. In this case, the RTT needs to be below 1 ms; otherwise, you should use asynchronous replication 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 be used if your RPO objectives are <5 minutes<sup>17</sup>.

vSphere replication is 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 be protected against data center site disasters, then all relevant systems need to be replicated to a second site. This can be 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 VMs between heterogeneous storage types that vSphere supports. Because vSphere Replication is independent of the underlying storage, it works with a variety of storage types, including vSAN, vSphere Virtual Volumes (vVols), traditional SAN, network-attached storage (NAS), and direct-attached storage (DAS).

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

If you use an [SAP HANA HCI based on vSAN solution](#), data center distances of up to 5 kilometers/3 miles are supported.

## Backing up and restoring SAP HANA systems virtualized on vSphere

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

The simplest approach involves performing a file system backup along with a HANA database dump, which you can conveniently run within SAP HANA Studio. If you are using a backup solution, you can leverage the backint interface. Refer to the SAP HANA product documentation for backup and restore information and requirements.

In addition, any SAP- and VMware-supported backup solution that leverages SAP HANA studio/backint with vSphere snapshots can protect a vSphere-deployed SAP HANA system (refer to figure 40). This offers a neutral storage vendor backup solution based on vSphere snapshots and

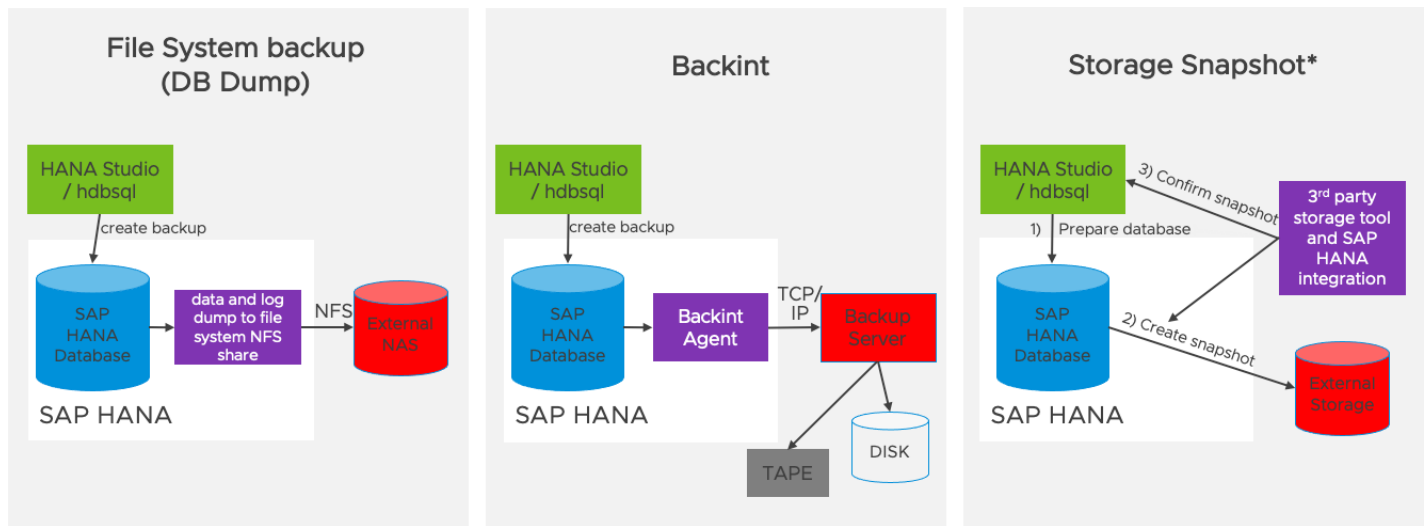
---

<sup>17</sup> Due to the asynchronous replication method, we do not recommend you use vSphere Replication for SAP HANA VMs.

reduces backup and restore times. For an example, refer to the [Veeam backup and recovery solution designed for vSphere and SAP HANA](#).

**Caution:** Using vSphere snapshot backup solutions not integrated with SAP HANA can lead to complications. These issues arise because of the need to freeze the VM to create a snapshot while SAP HANA is running. For this reason, we advise you use VMware snapshots only when SAP HANA has been stopped or when all IO activities have been paused to ensure data consistency.

Figure 40: Virtualized SAP HANA backup and recovery methods



## SAP HANA with persistent memory on vSphere 7

### Prerequisites and general SAP support limitations for Intel Optane Persistent Memory (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 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 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 31 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 get used. See table 32 for the currently supported configurations.

**Important:** [Intel Optane Persistent Memory \(PMem\) 100 series technology](#) is only supported with Cascade Lake and vSphere 7.0 virtualized SAP systems. Later CPU generations are not SAP HANA Intel PMem supported and Intel has announced the [discontinuance of Intel Optane Persistent Memory](#).

Table 31: Supported SAP HANA on vSphere with PMem ratios with Cascade Lake and vSphere 7

Ratio (DRAM to PMem)		DRAM (TB)	Optane PMem (TB)	Total host memory (TB)	Max. VM memory (TB) <sup>18</sup>	Total memory slots (Cascade Lake)	DRAM module size (GB)	Optane PMem module size (GB)	SAP supported ESXi host size (max. sockets)
2	1	6	3	9	6	48	256	128	4
1	1	6	6	12	6	48	256	256	4
1	1	3	3	6	6	48	128	128	4
1	2	3	6	9	6	48	128	256	4
1	2	1.5	3	4.5	4.5	48	64	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's are limited to 3TB, class-M sizing's 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.

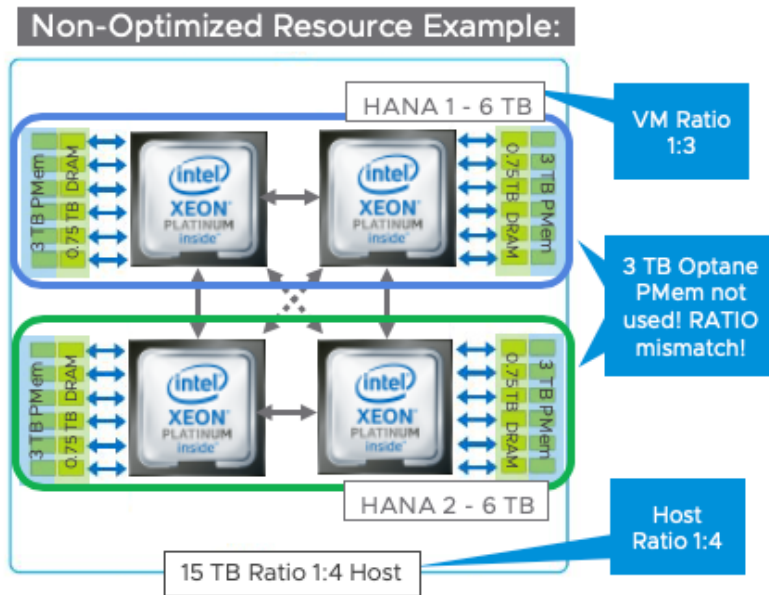
For example, you have a 1:4 PMem ratio host, configured with 15TB of total RAM (3TB DRAM and 12 TB Optane PMem). An optimized resource scenario is to create four SAP HANA VMs on this server, each with 3.75TB RAM (0.75TB DRAM and 3TB Optane PMem). If you are creating 6 TB VMs on this same 15TB host, then 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

<sup>18</sup>vSphere 7.0 U2 or later versions are required for VM sizes >6 TB.

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 configurations.

Figure 41: 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

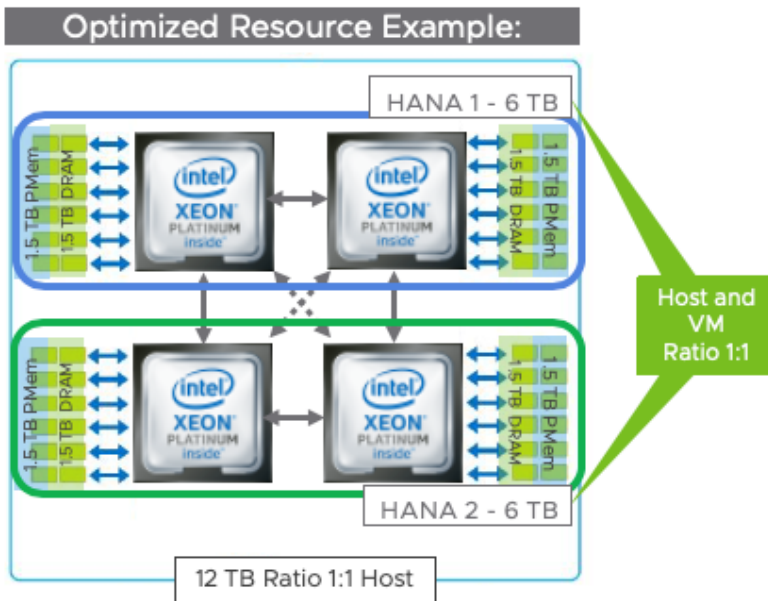
#### VM configuration example:

- 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 42: 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 43: 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 sizings 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 32 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 32: vSphere features supported with PMem-enabled SAP HANA VMs

PMem mode	App Direct (vNVDIMM)	Memory mode	Mixed mode	Storage over App Direct
Usage mode	Persistence	Volatile	Persistence/volatile	vPMemDisk
SAP HANA supported	Yes (SAP note 2913410)	Only non-prod SAP HANA support (SAP note 2954515)	Only non-prod SAP HANA support (SAP note 2954515)	No
Bootable drive	N/A	N/A	N/A	Yes
Host and VM size	Intel Cascade Lake 2 and 4 socket servers with up to 15TB (up to 16 TB without PMem); SAP HANA VM sizes up to 256 vCPUs and 6 TB of memory	Only non-prod SAP HANA support	Mixed mode is not supported by vSphere	Not supported for SAP HANA
VMware SMP	Yes, up to 4 physical CPU socket ESXi hosts and SAP HANA VMs	Only non-prod SAP HANA support	Mixed mode is not supported by vSphere	Not supported for SAP HANA
vSphere vMotion	Yes, up to 6 TB VM sizes (requires another PMem host)	Yes	Mixed mode is not supported by vSphere	Yes, up to 6 TB VM sizes (requires another PMem host and the usage of vSphere Storage vMotion)
vSphere DRS	Yes, up to 6 TB VM sizes (requires another PMem host)	Yes	Mixed mode is not supported by vSphere	Yes, up to 6 TB VM sizes (requires another PMem host and the usage of vSphere Storage vMotion)
VMware Site Recovery Manager	No	Yes	Mixed mode is not supported by vSphere	No (snapshot is not yet supported)
vSphere HA	Yes, with vSphere 7.0 U2 or later (no support prior to this release)	Yes (can move/failover VMs to/from Memory mode and DRAM systems)	Mixed mode is not supported by vSphere	No
VM in guest cluster	No	Yes (can move/failover VMs to/from Memory mode and DRAM systems)	Mixed mode is not supported by vSphere	No
VM snapshot	No	Yes	Mixed mode is not supported by vSphere	No



PMem mode	App Direct (vNVDIMM)	Memory mode	Mixed mode	Storage over App Direct
SAP HSR	Yes	Yes	Mixed mode is not supported by vSphere	No

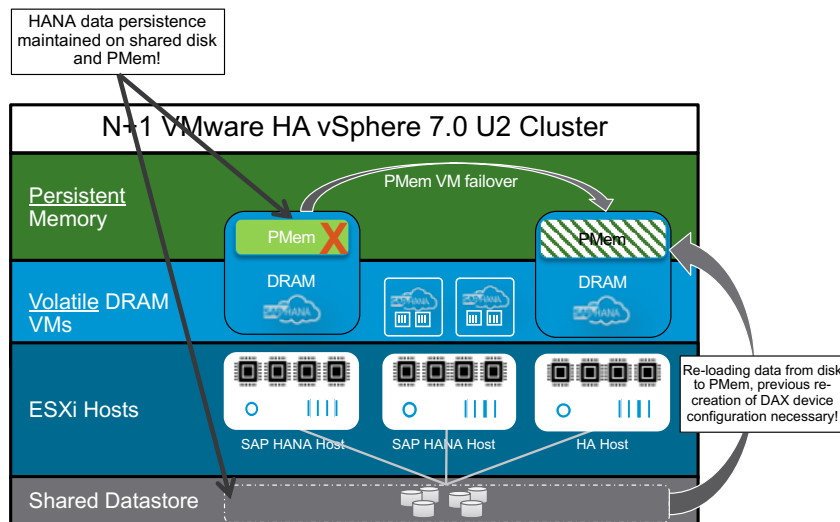
## vSphere HA support for PMem-enabled SAP HANA VMs

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 44 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 is 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 44: 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 run 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 or supported by VMware, SAP, or Intel. Any use of this script is your own responsibility.

## SAP HANA with PMem VM configuration details

Using PMem in a vSphere virtualized environment requires that the physical host, ESXi, and VM configurations are correctly configured.

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.


### Operating system:

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 configuration steps 2 and 3

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 45.

Figure 45: Example of PMem system BIOS settings



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 46–49.

Figure 46: ESXi Persistent Memory Storage view of Modules

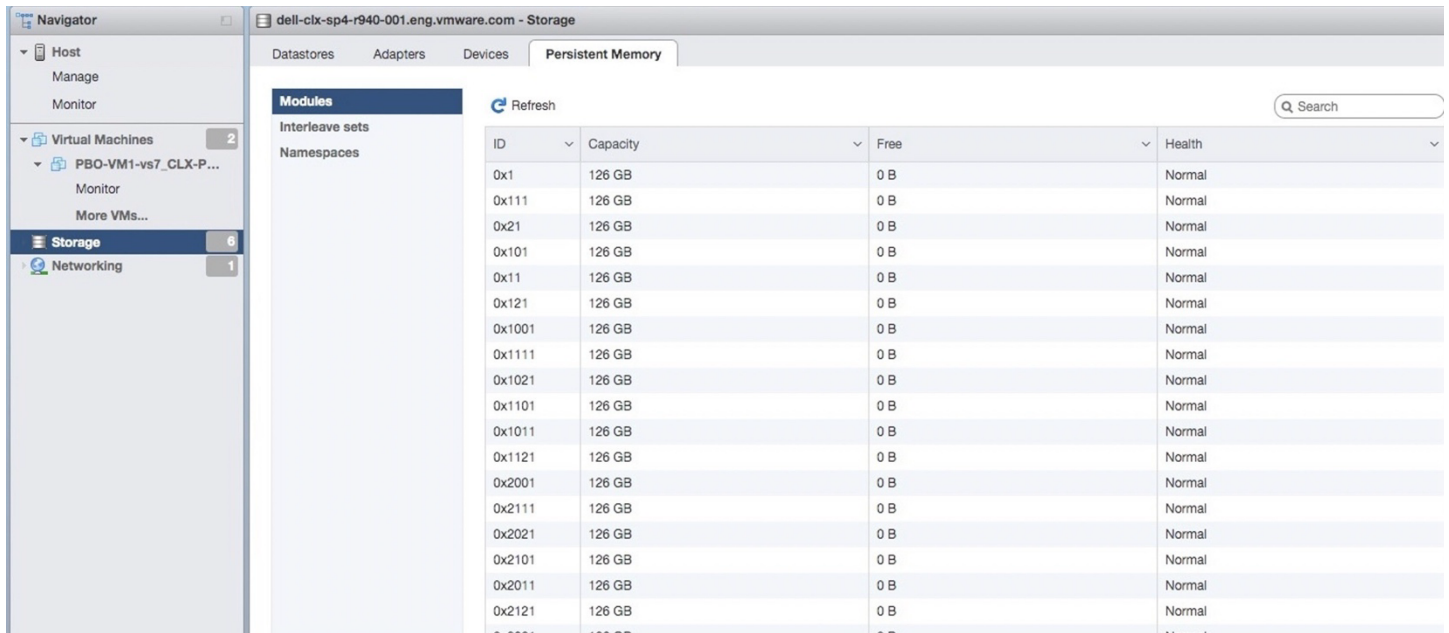


Figure 47: ESXi Persistent Memory Storage View of interleave sets

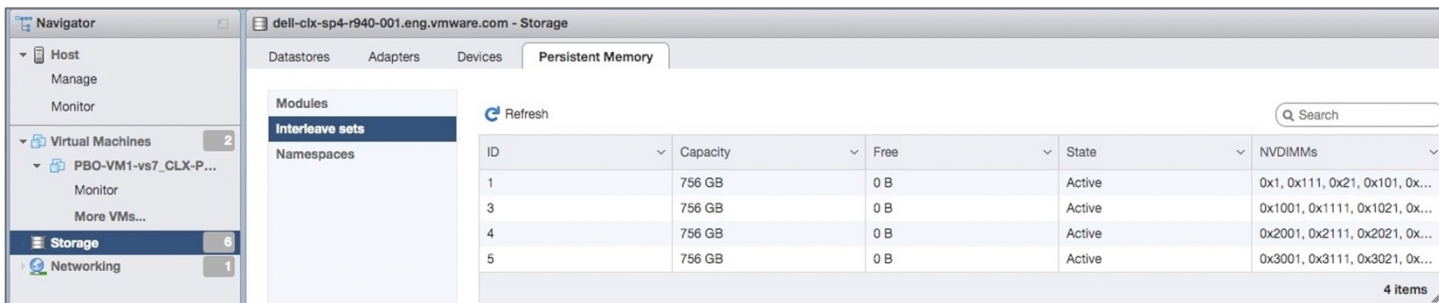
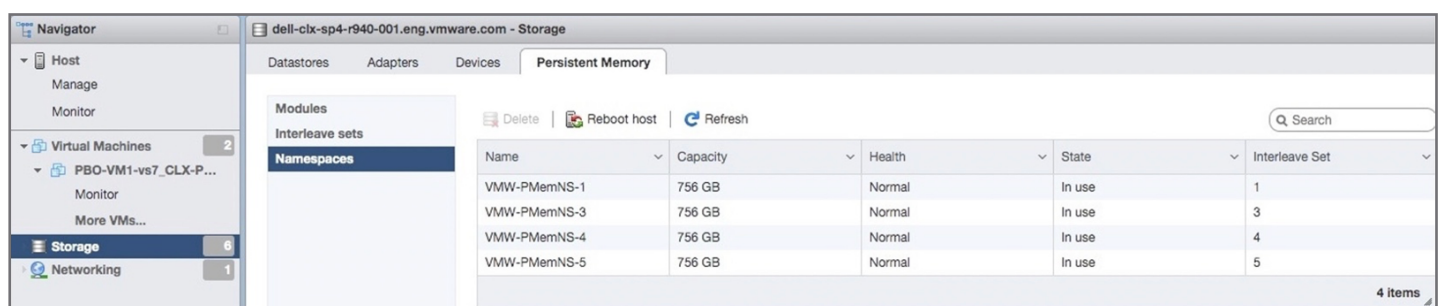


Figure 48: ESXi Persistent Memory Storage view of Namespaces



**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, the configuration should look like what is shown in Figures 47–49. 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 49: NVDIMM creation via the vSphere Client

Edit settings - PBO-VM1-vs7_CLX-PMEM-DELL (ESXi 6.7 U2 virtual machine)	
▼ NVDIMM	
Size	757760 MB
Controller location	NVDIMM controller 0:0
▼ NVDIMM	
Size	757760 MB
Controller location	NVDIMM controller 0:1
▼ NVDIMM	
Size	757760 MB
Controller location	NVDIMM controller 0:2
▼ NVDIMM	
Size	757760 MB
Controller location	NVDIMM controller 0:3

To configure an Optane PMem-enabled SAP HANA VM for optimal performance, you must 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:

- `nvdimmmode = "independent-persistent"`
- `nvdimmm0:0.nodeAffinity=0`
- `nvdimmm0:1.nodeAffinity=1`
- `nvdimmm0:2.nodeAffinity=2`
- `nvdimmm0:3.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 preallocation of the PMem in the target server.

Besides these parameters, you can 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 be 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 nvdimmm *.vmx nvdimmm0.present = "TRUE"
nvdimmm0:0.present = "TRUE"
nvdimmm0:0.fileName = "/vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_1.vmdk" nvdimmm0:0.size = "757760"
nvdimmm0:1.present = "TRUE"
nvdimmm0:1.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_3.vmdk" nvdimmm0:1.size = "757760"
nvdimmm0:2.present = "TRUE"
nvdimmm0:2.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_5.vmdk" nvdimmm0:2.size = "757760"
nvdimmm0:3.present = "TRUE"
nvdimmm0:3.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_5.vmdk" nvdimmm0:3.size = "757760"
nvdimmm0:0.node = "0"
nvdimmm0:1.node = "1"
nvdimmm0:2.node = "2"
nvdimmm0:3.node = "3"
```

Manually added parameters:

- `nvdimmmode = "independent-persistent"`
- `nvdimmm0:0.nodeAffinity=0`
- `nvdimmm0:1.nodeAffinity=1`
- `nvdimmm0:2.nodeAffinity=2`
- `nvdimmm0:3.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.

## Monitoring and verifying an SAP HANA installation

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

- SAP note [1698281](#) provides information about how you can monitor 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.

### NUMA tool: VMware NUMA Observer

While you 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 figures 50 and 51 for examples.

Figure 50: VMware NUMA Observer – VM Core Overlap Graph


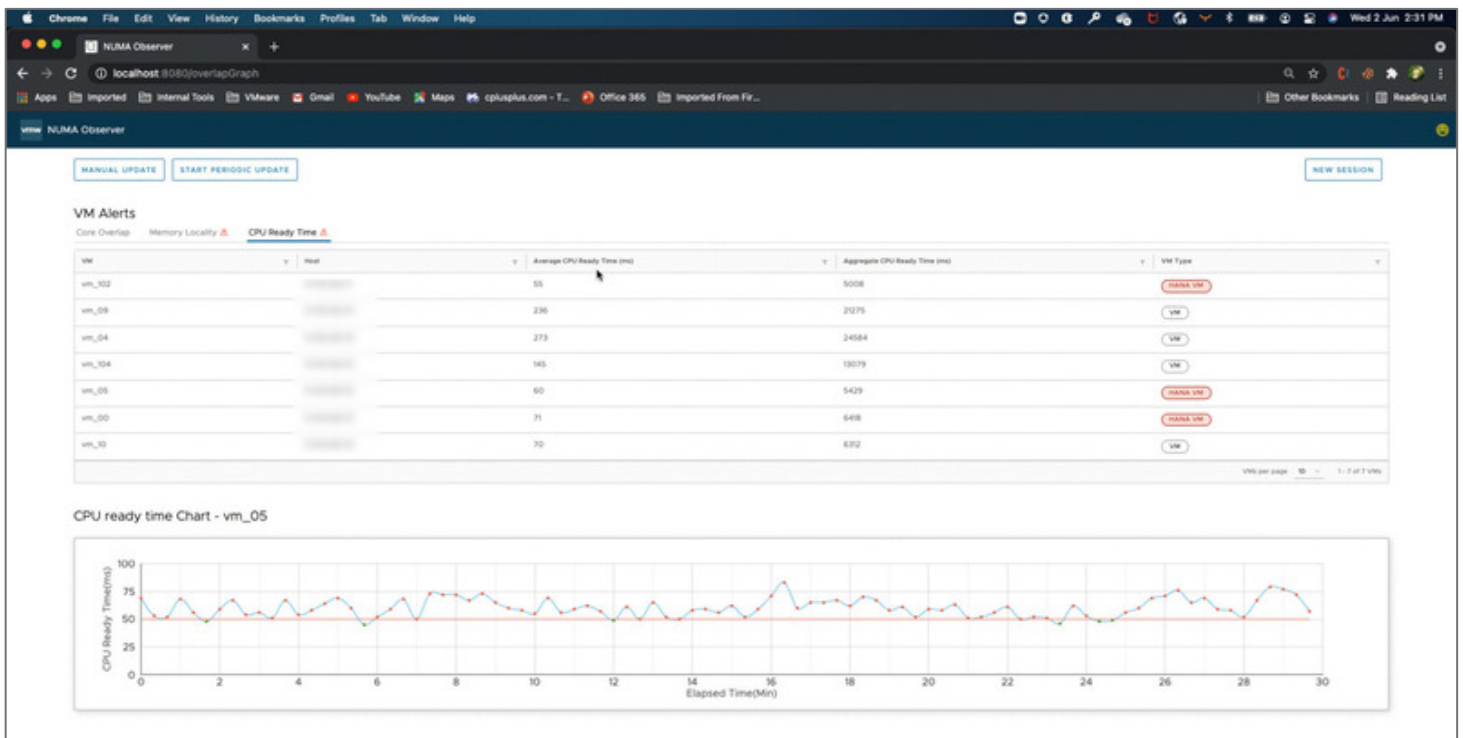


Figure 51: VMware NUMA Observer – VM Alerts





## Performance optimizations for SAP HANA VMs

This section discusses the best practices parameters to optimally configure an SAP HANA VM. The most critical aspect of these optimizations is that you properly align NUMA nodes to get the best performance and lowest memory latency.

### Optimizing the 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 always be configured for SAP HANA on vSphere. The settings described in the [Performance optimization for low-latency SAP HANA VMs section](#) should be applied only 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 36 shows the Sub NUMA Node Cluster (SNC) specific settings and are required when a half-socket VM should get run on a 2-socket Sapphire Rapids ESXi host.

**Table 33: 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.

Physical host BIOS parameter settings	Description
CPU – Populate all available CPU sockets, use a fully meshed QPI NUMA architecture	<p>To avoid timer synchronization issues, use a multi-socket server that ensures NUMA node timer synchronization. NUMA systems that do not run synchronization will need to synchronize the timers in the hypervisor layer, which can impact performance.</p> <p>See the Timekeeping in VMware Virtual Machines information guide for reference.</p> <p>Select only SAP HANA CPUs supported by vSphere. Verify the support status with the SAP HANA on vSphere support notes. For a list of the relevant note, see the SAP Notes Related to VMware page.</p>
Enable CPU Intel Turbo Boost	Allow Intel automatic CPU core overclocking technology (P-states).
Disable QPI power management	Do not allow static high power for QPI links.
Set HWPE support to	Set to HW vendor default
Enable hyperthreading	Always enable hyperthreading on the ESXi host. This will double the logical CPU cores to allow ESXi to take advantage of more available CPU threads.
Enable execute disable feature	Enable the Data Execution Prevention bit (NX-bit), required for vMotion.
Disable node interleaving	Disable node interleaving in BIOS.
Disable C1E Halt state	Disable enhanced C-states in BIOS.
Set power management to high performance	Do not use any power management features on the server, such as C-states. Configure in the BIOS static high performance.
Set correct PMem mode as specified by the hardware vendor for either App Direct or Memory mode	<p>Follow the vendor documentation and enable PMem for the usage with ESXi.</p> <p><b>Note:</b> Only App Direct and Memory mode are supported with production-level VMs.</p> <p>Memory mode is only supported with a ratio of 1:4. As of today, SAP provides only non-production workload support.</p> <p>VMware vSAN does not have support for App Direct mode as cache or as a capacity tier device of vSAN. However, vSAN will work with vSphere hosts equipped with Intel Optane PMem in App Direct mode and SAP HANA VMs can leverage PMem according to SAP note <a href="#">2913410</a>.</p> <p><b>Important:</b> The vSphere HA restriction (as described in SAP note 2913410) applies and needs to be considered.</p>
Enable SNC-2 for sub-NUMA VMs	If Sapphire Rapids 2-socket systems and "half-socket" / sub-NUMA VMs are required, enable SNC-2. Do not enable SNC-2 if this is not required or on another ESXi host that is not a 2-socket Sapphire Rapids system.
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 34: ESXi host parameter settings

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 >= 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>Go to the ESXi console and set the following parameter:</p> <ul style="list-style-type: none"> <li>vsish -e set /config/Net/intOpts/NetNetqRxQueueFeatPairEnable 0</li> </ul> <p>Add the following advanced VMX configuration parameters to the VMX file, and reboot the VM after adding these parameters:</p> <ul style="list-style-type: none"> <li>ethernetX.pnicFeatures = "4"</li> <li>ethernetX.ctxPerDev = "3"</li> </ul> <p>Change the rx-usec, lro and rx/tx values of the VMXNET3 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 operating system running the inside the VM and use ethtool to change the following settings, then execute:</p> <ul style="list-style-type: none"> <li>ethtool -C ethX rx-usec 75</li> <li>ethtool -K ethX lro off</li> <li>ethtool -G ethX rx 512 rx-mini 0 tx 512</li> </ul> <p><b>Note:</b> Exchange X with the actual number, such as eth0. To make these ethtool settings permanent, see <a href="#">SLES KB 000017259</a> or the <a href="#">RHEL ethtool document</a>.</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>Go to the ESXi console and set the following parameter:</p> <ul style="list-style-type: none"> <li>vsish -e set /config/Disk/intOpts/VSCSIPollPeriod 100</li> </ul> <p>If you want to use vSAN, then select one of the certified <a href="#">SAP HANA HCI solutions based on vSAN</a> and follow the <a href="#">VMware HCI BP guide</a>.</p>
SAP monitoring	<p>Enable SAP monitoring on the host by setting Misc.GuestLibAllowHostInfo = "1"</p> <p>For more details, see SAP note <a href="#">1409604</a>.</p> <p>Without this parameter, no host performance relevant data will be viewable inside an SAP monitoring enabled VM.</p>

Table 35: SAP HANA virtual machine parameter settings

SAP HANA virtual machine parameter settings	Description
Tips for editing a *.vmx file	Review tips for editing a *.vmx file in VMware KB <a href="#">1714</a> .
UEFI BIOS guest	<p><b>Recommended:</b> Use UEFI BIOS as the standard BIOS version for vSphere hosts and guests. Features such as Secure Boot are possible only with EFI.</p> <p>See VMware <a href="#">DOC-28494</a> for details. You can configure this with the vSphereClient by choosing EFI boot mode.</p> <p>If you are using vSphere 6.0 and you only see 2TB memory in the guest, then upgrade to the latest ESXi 6.0 version.</p>
SAP monitoring	<p>Enable SAP monitoring inside the SAP HANA VM with the advanced VM configuration parameter <code>tools.guestlib.enableHostInfo = "TRUE"</code>.</p> <p>For more details, see SAP note <a href="#">1409604</a>.</p> <p>Besides setting this parameter, the VMware guest tools need to be installed. For details, see VMware KB <a href="#">1014294</a>.</p>
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 <a href="#">2040375</a> .
Memory reservations	<p>Set fixed memory reservations for SAP HANA VMs. Do not overcommit memory resources.</p> <p>You must reserve memory for the ESXi host. The amount you reserve depends on the amount of CPU sockets memory for ESXi.</p> <p>Typical memory reservation for a host is between 32–64GB for a 2-socket server, 64–128GB for a 4-socket server, and 128–256GB for an 8-socket server. These are not absolute figures as the memory need of ESXi depends strongly on the actual hardware, ESXi and VM configuration, and enabled ESXi features, such as vSAN.</p>
CPU	<p>Do not overcommit CPU resources and configure dedicated CPU resources per SAP HANA VM.</p> <p>You can use hyperthreads when you configure the virtual machine to gain additional performance. For CPU generations older than Cascade Lake, you should consider disabling hyperthreading due to the Intel Vulnerability Foreshadow L1 Terminal Fault. For details, read VMware KB <a href="#">55636</a>.</p> <p>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 60 core Sapphire Rapids system will require 240 vCPUs).</p>
vNUMA nodes	<p>SAP HANA on vSphere can be configured to leverage half-CPU and full-CPU sockets. A half-CPU socket is configured by only half of the available physical cores of a CPU socket in the VM configuration GUI.</p> <p>The vNUMA nodes of the VM will always be <math>\geq 1</math>, depending on how many CPUs you have configured in total.</p> <p>If you need to access an additional NUMA node, use all CPU cores of this additional NUMA node. Nevertheless, use as few NUMA nodes as possible to optimize memory access.</p>

SAP HANA virtual machine parameter settings	Description
	<p>Half-Socket/sub-NUMA VMs require SNC-2 with Sapphire Rapids systems.</p> <p><b>Important:</b> As of today, only 2-socket SPR systems are supported with SNC-2. Once SNC-2 is enabled, four sub-NUMA nodes will be available on a supported 2-socket SPR host. In the case of SNC-2, all logical CPUs of a sub-NUMA node should be assigned to such an SNC-2 VM.</p> <p>Additionally, the parameter 'sched.nodeX.affinity="Y"' must be set to ensure that the VM runs on the correct sub-NUMA node and does not get migrated by the scheduler to another sub-NUMA node.</p>
Align virtual CPU VM configuration to actual server hardware	<p>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. Similarly, a full-socket VM should also be configured to use 56 vCPUs to utilize all 28 physical CPU cores and available hyperthreads per socket.</p> <p>In the case of a Sapphire Rapids (SPR) CPU, SNC-2 is required for half-socket VMs, and an SNC-enabled host will provide two sub-NUMA nodes per CPU socket. In this scenario, all logical CPUs of each sub-NUMA node need to be utilized. For example, with a 60-core SPR CPU, when SNC is enabled, it becomes 30 cores per sub-NUMA node. Therefore, a "half-socket" VM will be configured with 60 vCPUs (including hyperthreads).</p>
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 controller or NVME adapter for I/O devices	Use the dedicated SCSI controllers for OS, log and data to separate disk I/O streams. Select, based on the used storage devices, a PVSCSI controller (SAS/SATA disks and FC) or NVME adapter (NVME devices). For details, see the SAP HANA storage and disk layout section.
Use the virtual machine's file system	Use VMDK disks whenever possible to allow optimal operation via the vSpherestack. 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 <a href="#">1943937</a> .
Eager zero thick virtual disks for data and log disk	We recommend this setting as it avoids lazy zeroing (initial write penalty).
VMXNET3	<p>Use paravirtual VMXNET 3 virtual NICs for SAP HANA virtual machines.</p> <p>We recommend at least 3–4 different NICs inside a VM in the HANA VM (app/ management server network, backup network, and, if needed, HANA system replication network). Corresponding physical NICs inside the host are required.</p>

SAP HANA virtual machine parameter settings	Description
Optimize the application server network latency if required	<p>Disable virtual interrupt coalescing for VMXNET 3 virtual NICs that communicate with the app servers or front end to optimize network latency. Do not set this parameter for throughput-oriented networks, 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.</p> <p>For details, see the <a href="#">Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs</a> white paper.</p>
Set lat.Sensitivity = normal	<p>Check with the vSphere Client and ensure that, in the VM configuration, the value of Latency Sensitivity Settings is set to "normal." If you must change this setting, restart the VM.</p> <p>Do not change this setting to "high" or "low." Change this setting under the instruction of VMware support engineers.</p>
Configuring virtual topology when vSphere 8 and later gets used.	<p>vSphere ESXi 8.0 introduces an enhanced virtual topology feature. This feature automatically selects optimal coresPerSocket values for virtual machines and optimal virtual L3 sizes.</p> <p>The intelligent, adaptive NUMA scheduling and memory placement policies in ESXi 8.0 can manage all virtual machines transparently, so that administrators don't need to deal with the complexity of balancing virtual machines between nodes by hand.</p> <p>Manual controls are still available to override this default behavior, however, and SAP HANA production or performance critical VMs should manually configure NUMA placement (through sched.nodeX.affinity="Y advanced option).</p> <p>For details refer to <a href="#">Performance Best Practices for VMware vSphere 8.0</a> paper and for detailed configuration steps refer to paper: <a href="#">VMware vSphere 8.0 Virtual Topology</a>.</p> <p>In the case sched.nodeX.affinity="Y cannot get used due to operational reason, then NUMA action affinity must get set to avoid unwanted NUMA node cd-deployments.</p>
NUMA action affinity	<p>If the sched.nodeX.affinity="Y" setting is not usable due to specific operational constraints, it is possible to avoid VM co-deployment on a CPU socket/NUMA node by deactivating the NUMA locality affinity. Follow VMware KB <a href="#">2097369</a>.</p> <p>This setting can be used for SNC and non-SNC configurations.</p> <p>Use the vSphere Web Client and add the following advanced VM parameter per VM setting: Numa.LocalityWeightActionAffinity="O"</p>
Associate virtual machines with specified NUMA nodes to optimize NUMA memory locality  <b>Note:</b> While ESXi 8 provides an enhanced virtual topology feature, that automatically selects optimal coresPerSocket values for virtual machines and optimal virtual L3 sizes to align it to the underlying NUMA topology of a host, it is still recommended to associate a virtual NUMA node to a physical NUMA node. This is especially recommended for consolidation configurations, where more SAP	<p>For all SAP HANA VMs, it is recommended to affinity the virtual NUMA node with a physical NUMA node. This provides direct control over how a VM is placed on a specific NUMA node. It also ensures that a virtual NUMA node does not get migrated, which may occur if there are idle NUMA nodes. While these migrations help to load balance how a server is utilized by VMs, for SAP HANA VMs, this is mostly negative and may impact memory latency and therefore performance. Due to this, associating the physical NUMA nodes with a virtual machine is required to constrain how the ESXi scheduler can schedule a VM's CPU and memory.</p>

SAP HANA virtual machine parameter settings	Description
<p>VMs share a single host. In the case of single large VM these settings are not required.</p>	<p>Use the vSphere Web Client or directly modify the .vmx file and add <code>sched.nodeX.affinity="Y"</code></p> <p>Refer to the VMware article, "<a href="#">Associate Virtual Machines with Specified NUMA Nodes</a>" for details.</p> <p>Procedure:</p> <ol style="list-style-type: none"> <li>1. Browse to the VM in the vSphere Client.</li> <li>2. Right-click and select Edit Settings.</li> <li>3. Select the VM Options tab and expand Advanced.</li> <li>4. Under Configuration Parameters, click the Edit Configuration button.</li> <li>5. Click Add Row to add a new option. <ul style="list-style-type: none"> <li>– To specify a NUMA node for a specific virtual NUMA node on the VM: In the <b>Name</b> column, enter <code>sched.nodeX.affinity</code>, where X is the virtual NUMA node number. For example, <code>sched.node0.affinity</code> specifies the virtual NUMA node 0 on the virtual machine.</li> <li>– In the <b>Value</b> column, enter the NUMA node where the virtual machine or the virtual NUMA node can be scheduled.</li> </ul> </li> <li>6. Add for every virtual NUMA node the VM should use a new row with a new <code>sched.nodeX.affinity</code> option.</li> <li>7. Click <b>OK</b> twice to close the <b>Edit Settings</b> dialog box.</li> </ol> <p>Example for a single NUMA node VM:</p> <ul style="list-style-type: none"> <li>• <code>sched.node0.affinity = "2"</code> (this is physical NUMA node 2)</li> </ul> <p>Example for a two-NUMA node VM:</p> <ul style="list-style-type: none"> <li>• <code>sched.node0.affinity = "0"</code></li> <li>• <code>sched.node1.affinity = "1"</code></li> </ul> <p>Not supported:</p> <ul style="list-style-type: none"> <li>• <code>sched.node0.affinity = "1"</code></li> <li>• <code>sched.node1.affinity = "2"</code></li> </ul> <p><b>Important:</b> For SNC-enabled hosts, do not set affinities to sub-NUMA nodes that cross a physical CPU socket. See table 31, "Sub-NUMA Node Cluster (SNC) Specific Settings" for details.</p>
<p>Configure virtual machines to use hyperthreading with NUMA</p>	<p>For memory latency-sensitive workloads with low processor utilization, such as SAP HANA, or high interthread communication, we recommended using hyperthreading with fewer NUMA nodes instead of full physical cores spread over multiple NUMA nodes. Use hyperthreading and enforce NUMA node locality per VMware KB <a href="#">2003582</a>.</p> <p>This parameter is only required when hyperthreading should be leveraged for a VM. Using hyperthreading can increase the compute throughput but may increase the latency of threads.</p> <p><b>Note:</b> 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</p>

SAP HANA virtual machine parameter settings	Description
	<p>an 8-socket VM on an 8-socket host). If a VM has more vCPUs configured than available physical cores, this parameter gets configured automatically.</p> <p>Use the vSphere Web Client and add the following advanced VM parameter:</p> <p>numa.vcpu.preferHT="TRUE" (per VM setting) or as a global setting on the host: Numa.PreferHT="1" (host).</p> <p><b>Note:</b> For non-mitigated CPUs, such as Haswell, Broadwell and Skylake, you may consider not to use hyperthreading at all. For details, see VMware KB <a href="#">55806</a>.</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.</p> <p>VMware KB <a href="#">78094</a> provides information on how to configure the NVDIMMs (VMware's representation of Optane PMem) correctly and align the NVDIMMs to the physical NUMA architecture of the physical server.</p> <p>By default, Optane PMem allocation in vmkernel for VM NVDIMMs does not consider NUMA. This can result in the VM running on a certain NUMA node and Optane PMem allocated from a different NUMA node. This will cause NVDIMMs access in the VM to be remote, resulting in poor performance.</p> <p>To solve this, you must add the following settings to a VM configuration using vCenter.</p> <p>Example for a 4-socket wide VM:</p> <ul style="list-style-type: none"> <li>• nvdim0:0.nodeAffinity=0</li> <li>• nvdim0:1.nodeAffinity=1</li> <li>• nvdim0:2.nodeAffinity=2</li> <li>• nvdim0:3.nodeAffinity=3</li> </ul> <p>sched.pmem.prealloc=TRUE is an optional parameter equivalent to eager zerothick 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 36: Sub NUMA Node Cluster (SNC) settings

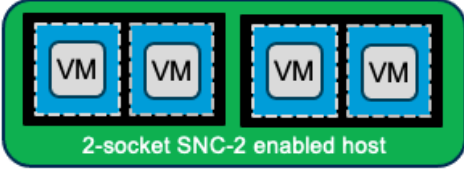
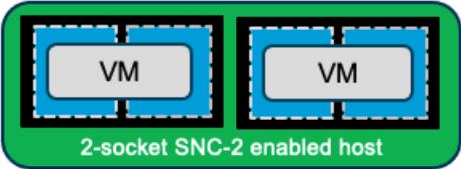
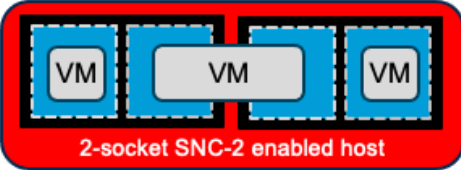
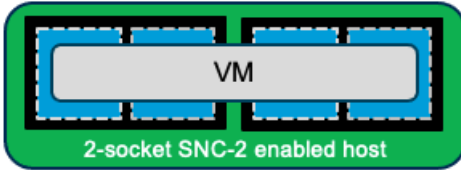
SAP SNC-2 enabled ESXi host (2-socket only)	Advanced VMX settings
<p>2 VMs per CPU socket, 4 VMs total per ESXi host</p> 	<p>VM1:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity = "0"</code></p> <p>VM2:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity = "1"</code></p> <p>VM3:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity = "2"</code></p> <p>VM4:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity = "3"</code></p>
<p>1 VM per CPU socket, 2 VMs total per ESXi host</p> 	<p>VM1:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity = "0"</code>  <code>sched.node1.affinity = "1"</code></p> <p>VM2:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity="2"</code>  <code>sched.node1.affinity="3"</code></p>
<p>Not supported</p> 	<p><b>Not supported:</b> Do not configure a sub-NUMA node, wide VM that crosses the QPI link.</p> <p>VM:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity="1"</code>  <code>sched.node1.affinity="2"</code></p>
<p>1 VM across two CPU sockets</p> 	<p>VM:  <code>numa.vcpu.preferHT="TRUE"</code>  <code>sched.node0.affinity="0"</code>  <code>sched.node1.affinity="1"</code>  <code>sched.node2.affinity="2"</code>  <code>sched.node3.affinity="3"</code></p>

Table 37: Linux operating system parameter settings

Linux OS parameter settings	Description
Linux version	<p>VMware strongly recommends using only the SAP HANA supported Linux and kernel versions. See SAP Note <a href="#">2235581</a> and for settings, see SAP Note <a href="#">2684254</a>.</p> <p>Use <b>SAP HANA SAPConf/SAPTune</b> to optimize the Linux OS for SAP HANA.</p>
<p>To optimize large-scale workloads with intensive I/O patterns, change the queue depths of the SCSI default values.</p>	<p>The large-scale workloads with intensive I/O patterns require adapter queue depths greater than the Paravirtual SCSI (PVSCSI) default values. The default values of PVSCSI queue depth are 64 (for device) and 254 (for adapter). You can increase PVSCSI queue depths to 254 (for device) and 1024 (for adapter) inside a Windows virtual machine or Linux virtual machine.</p> <p>Create a file of any name in the <code>/etc/modprobe.d/</code> directory with this line:</p> <pre>options vmw_pvscsi cmd_per_lun=254 ring_pages=32</pre> <p><b>Note:</b> For RHEL5, edit <code>/etc/modprobe.conf</code> with the same line. Make a new <code>initrd</code> for the settings to take effect. You can do this either by using <code>mkinitrd</code>, or by re-running <code>vmware-config-tools.pl</code>.</p> <p>Starting in version 6, RHEL uses <code>modprobe.d</code>.</p> <p>Alternatively, append these to kernel boot arguments (for example, on Red Hat Enterprise Linux edit <code>/etc/grub.conf</code> or on Ubuntu edit <code>/boot/grub/grub.cfg</code>).</p> <pre>vmw_pvscsi.cmd_per_lun=254 vmw_pvscsi.ring_pages=32</pre> <p>Reboot the virtual machine. See VMware <a href="#">KB 2053145</a> for details.</p> <p><b>Note:</b> Review the <a href="#">VMware KB article 2088157</a> to ensure that the minimum VMware patch level is used to avoid possible virtual machine freezes under heavy I/O load.</p>
Install the latest version of VMware Tools	<p>VMware Tools is a suite of utilities, which enhances the performance of the VM's guest operating system and improves the VM management. See <a href="http://kb.vmware.com/kb/1014294">http://kb .vmware com/kb/1014294</a> for details.</p>
Configure NTP time server	<p>Use the same external NTP server as configured for vSphere. For details, see SAP note <a href="#">989963</a>.</p>
<p>Optional: Disable large receive offload (LRO) in the Linux guest OS to lower latency for client/application server-facing NIC adapter</p>	<p>This helps to lower network latency of client/application server facing NIC adapters run:</p> <pre>ethtool -K ethY lro off</pre> <p>Do not disable LRO for throughput NIC adapters such as for backup, replication, or SAP HANA internode communication networks.</p> <p>Works only with Linux kernel 2.6.24 and later and uses a VMXNET3.</p> <p>Additional details: <a href="http://kb.vmware.com/kb/2055140">http://kb.vmware.com/kb/2055140</a></p>

**Table 38: General SAP HANA Linux configuration recommendations (These settings will be automatically configured when you use SAPtune/SAPConf.)**

Linux parameter settings	Description
Linux with SAP HANA Reference Guide	See the recommended operating system configuration settings for running SAP HANA on Linux in the <a href="#">Reference Guide</a> .
Disable I/O scheduling	SLES15 SP2 onwards is configured by default, scheduler is set to none with block_mq enabled.
Disable AutoNUMA	<p>Later Linux kernel (RHEL 7 and SLES 12) supporting auto- migration according to NUMA statistics.</p> <p><b>For SLES:</b></p> <ol style="list-style-type: none"> <li>1. # yast bootloader</li> <li>2. Choose <b>Kernel Parameters</b> tab (ALT-k)</li> <li>3. Edit the Optional Commandline Parameters section and append: numa_balancing=disabled</li> </ol> <p><b>For RHEL:</b></p> <ol style="list-style-type: none"> <li>1. Edit the file: etc/sysctl.d/sap_hana.conf</li> <li>2. Add the following: kernel.numa_balancing = 0</li> <li>3. Reconfigure the kernel by running: # sysctl -p /etc/sysctl.d/sap_hana.conf</li> </ol>
Use Block mq	In the kernel parameters, add scsi_mod.use_blk_mq=1. For OS version Sles15 SP2 and beyond this is enabled by default.
Disable transparent HugePages	<p>THP is not supported for the use with SAP HANA DB, as it may lead to hanging situations and performance degradations.</p> <p>To check the current configuration, run the following command: # cat /sys/kernel/mm/transparent_hugepage/enabled</p> <p>Its output should read: always madvise [never]</p> <p>If this is not the case, you can disable the THP usage at runtime by issuing the following command: # echo never &gt; /sys/kernel/mm/ transparent_hugepage/enabled</p> <p>For details, refer to the SAP WIKI for SAPs and the Linux OS vendors virtualization independent <a href="#">recommended Linux OS settings</a> for SAP HANA.</p>
Change the following parameters in /etc/sysctl.conf (important for SAP HANA Scale-Out deployments)	<pre>net.core.rmem_default = 262144 net.core.wmem_max = 8388608 net.core.wmem_default = 262144 net.core.rmem_max = 8388608 net.ipv4.tcp_rmem = 4096 87380 8388608 net.ipv4.tcp_wmem = 4096 65536 8388608</pre>

Linux parameter settings	Description
	<pre>net.ipv4.tcp_mem = 8388608 8388608 8388608 net.ipv4.tcp_slow_start_after_idle = 0</pre>
Example Linux kernel boot loader parameters	<pre>intel_idle.max_cstate=0 processor.max_cstate=0 numa_balancing=disabled transparent_hugepage=never elevator=noop vmw_pvscsi.cmd_per_lun=254 vmw_pvscsi.ring_pages=32</pre>

## 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 table 39 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 affinity settings

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

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

Table 39: Tunings for very low-latency SAP HANA VMs

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

SAP HANA VM parameter settings	Description
<code>monitor_control.disable_pause_loop_exiting = "TRUE"</code>	This parameter prevents the VM from exiting to the hypervisor unnecessarily during a pause instruction. This is specific for Intel Skylake systems.

Table 40: Settings to improve NUMA alignment

SAP HANA virtual machine parameter settings	Description
Configuring CPU affinity <code>sched.vcpuXx.affinity = "Yy-Zz"</code>	<p><b>Note:</b> Remove <code>sched.nodeX.affinity</code> or <code>numa.nodeAffinity</code> settings if set and if the CPU affinities with <code>sched.vCPUxxx.affiity</code> 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 (VMware KB <a href="#">2145719</a>) 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 <code>numa.slit.enable</code> gets used.</p> <p>Just like with <code>sched.nodeX.affinity</code> it is possible to decrease the CPU and memory latency by further limit the ESXi scheduler to migrate VM threads to other logical processors/CPU threads by leveraging the <code>sched.vCPUxxx.affinity</code> VMX parameter in contrast to parameter <code>sched.nodeX.affinity</code> 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 <code>.vmx</code> file (recommended way) and add <code>sched.vcpuXx.affinity = "Yy-Zz"</code> (for example: <code>sched.vcpu0.affinity = "0-55"</code>) for each virtual CPU you want to use.</p> <ol style="list-style-type: none"> <li>1. Browse to the cluster in the vSphere Client.</li> <li>2. Click the <b>Configure</b> tab and click <b>Settings</b>.</li> <li>3. Under VM Options, click the <b>Edit</b> button.</li> <li>4. Select the <b>VM Options</b> tab and expand <b>Advanced</b>.</li> <li>5. Under <b>Configuration Parameters</b>, click the <b>Edit Configuration</b> button.</li> <li>6. Click <b>Add Row</b> to add a new option.</li> <li>7. In the <b>Name</b> column, enter <code>sched.vcpuXX.affinity</code> (XX stands for the actual vCPU want to assign to a physical CPU thread).</li> <li>8. In the <b>Value</b> column, enter the physical CPU threads where the vCPU can be scheduled. For example, enter 0-55 to constrain the virtual machine resource scheduling to physical CPU threads 0-55, which would be the 1st CPU of an 28 core CPU host.</li> <li>9. Click <b>OK</b>.</li> </ol>

SAP HANA virtual machine parameter settings	Description
	<p>10. Click <b>OK</b> to close the <b>Edit VM</b> dialog box.</p> <p>For more information about potential performance practices, see <a href="#">vSphere Resource Management Guide</a>.</p>
>4-socket VM on 8-socket hosts	<p>Add the advanced parameter:</p> <pre>numa.slit.enable = "TRUE"</pre> <p>to ensure the correct NUMA map for VMs &gt; 4 socket on 8-socket hosts.</p> <p><b>Note:</b> <code>sched.vcpuXx.affinity = "Yy-Zz"</code> must be configured when <code>numa.slit.enable</code> is set to "TRUE".</p>

## Examples of VMX configurations for SAP HANA VMs

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

**Table 41: Additional VMX parameters to set for half-socket SAP HANA VMs on 28-core/n-socket CPU server; socket 0**

Half-socket SAP HANA VM additional VMX parameters	Settings
First half-socket VM on socket 0 on a example 28-core CPU n-socket server	<pre> numa.vcpu.preferHT="TRUE" sched.vcpu0.affinity = "0-27" sched.vcpu1.affinity = "0-27" sched.vcpu2.affinity = "0-27" ... sched.vcpu26.affinity = "0-27" sched.vcpu27.affinity = "0-27" </pre>
First half-socket PMem VM on socket 0 on an example 28-core CPU n-socket server	<pre> nvdimm0:0.nodeAffinity=0 numa.vcpu.preferHT="TRUE" sched.vcpu0.affinity = "0-27" sched.vcpu1.affinity = "0-27" sched.vcpu2.affinity = "0-27" ... sched.vcpu26.affinity = "0-27" sched.vcpu27.affinity = "0-27" </pre>
Second half-socket VM on socket 0 on an example 28-core n-socket CPU server	<pre> numa.vcpu.preferHT="TRUE" sched.vcpu0.affinity = "28-55" sched.vcpu1.affinity = "28-55" sched.vcpu2.affinity = "28-55" ... sched.vcpu26.affinity = "28-55" sched.vcpu27.affinity = "28-55" </pre>
Second half-socket PMem VM on socket 1 on an example 28-core CPU n-socket server	<pre> nvdimm0:0.nodeAffinity=1 numa.vcpu.preferHT="TRUE" sched.vcpu0.affinity = "84-111" sched.vcpu1.affinity = "84-111" sched.vcpu2.affinity = "84-111" ... sched.vcpu26.affinity = "84-111" sched.vcpu27.affinity = "84-111" </pre>



Table 42: Additional VMX parameters to set for single-socket SAP HANA VMs on a 28-core/8-socket CPU server; socket 3

1-socket SAP HANA VM additional VMX parameters	Settings
1-socket VM on socket 3 on an example 28-core CPU 4 or 8-socket server	<code>numa.vcpu.preferHT="TRUE"</code> <code>sched.node0.affinity="3"</code>
1-socket PMem VM on socket 3 on an example 28-core CPU 4 or 8-socket server	<code>nvdimm0:0.nodeAffinity=3</code> <code>numa.vcpu.preferHT="TRUE"</code> <code>sched.node0.affinity="3"</code>

Table 43: Additional VMX parameters to set for dual-socket SAP HANA VMs on a 28-core/n-socket CPU server; sockets 0-1

2-socket SAP HANA VM additional VMX parameters	Settings
2-socket VM on sockets 0 and 1 on a n-socket server	<code>numa.vcpu.preferHT="TRUE"</code> <code>sched.node0.affinity="0"</code> <code>sched.node1.affinity="1"</code>
2-socket PMem VM on sockets 0 and 1 a n-socket server	<code>nvdimm0:0.nodeAffinity=0</code> <code>nvdimm0:1.nodeAffinity=1</code>  <code>numa.vcpu.preferHT="TRUE"</code> <code>sched.node0.affinity="0"</code> <code>sched.node1.affinity="1"</code>

Table 44: Additional VMX parameters to set for 3-socket SAP HANA VMs on a 28-core/4- or 8-socket CPU server; sockets 0-2

3-socket SAP HANA VM additional VMX parameters	Settings
3-socket VM on sockets 0, 1 and 2 on a 4 or 8-socket server	<code>numa.vcpu.preferHT="TRUE"</code> <code>sched.node0.affinity="0"</code> <code>sched.node1.affinity="1"</code> <code>sched.node2.affinity="2"</code>
3-socket PMem VM on sockets 0, 1, and 2 on a 4 or 8-socket server	<code>nvdimm0:0.nodeAffinity=0</code> <code>nvdimm0:1.nodeAffinity=1</code> <code>nvdimm0:2.nodeAffinity=2</code>  <code>numa.vcpu.preferHT="TRUE"</code> <code>sched.node0.affinity="0"</code> <code>sched.node1.affinity="1"</code> <code>sched.node2.affinity="2"</code>

Table 45: Additional VMX parameters to set for 4-socket SAP HANA VMs on a 28-core/4-socket CPU

4-socket SAP HANA VM additional VMX parameters	Settings
4-socket VM on a 4-socket server	No additional settings are required as the VM utilizes all server resources.
4-socket PMem VM on a 4-socket server	<pre> nvdimm0:0.nodeAffinity=0 nvdimm0:1.nodeAffinity=1 nvdimm0:2.nodeAffinity=2 nvdimm0:3.nodeAffinity=3 </pre>

Table 46: Additional VMX parameters to set for 4-socket SAP HANA VM on a 28-core/8-socket CPU; sockets 0-3

4-socket SAP HANA VM additional VMX parameters	Settings
4-socket VM on an example 28-core CPU 8-socket server running on sockets 0-3	<pre> numa.slit.enable = "TRUE" sched.vcpu0.affinity = "0-55" sched.vcpu1.affinity = "0-55" sched.vcpu2.affinity = "0-55" ... sched.vcpu222.affinity = "168-223" sched.vcpu223.affinity = "168-223" </pre>

Table 47: Additional VMX parameters to set for 6-socket SAP HANA VM on a 28-core/8-socket CPU server; sockets 0-5

6-socket SAP HANA VM additional VMX parameters	Settings
6-socket VM on an example 28-core CPU 8-socket server running on sockets 0-5	<pre> numa.slit.enable = "TRUE" sched.vcpu0.affinity = "0-55" sched.vcpu1.affinity = "0-55" sched.vcpu2.affinity = "0-55" ... sched.vcpu334.affinity = "280-335" sched.vcpu335.affinity = "280-335" </pre>

**Table 48: Additional VMX parameters to set for 8-socket SAP HANA VM on a 28-core/8-socket CPU server**

SAP HANA 8-socket VM additional VMX parameters	Settings
8-socket VM on an example 28-core CPU 8-socket server	<pre> numa.slit.enable = "TRUE" sched.vcpu0.affinity = "0-55" sched.vcpu1.affinity = "0-55" sched.vcpu2.affinity = "0-55" ... sched.vcpu446.affinity = "392-447" sched.vcpu447.affinity = "1392-447"                     </pre>

**Table 49: Additional VMX parameters to set for an SAP HANA VM running a latency-sensitive application**

Low latency SAP HANA VM additional VMX parameters	Settings
n-socket low-latency VM on an n-socket, CPU server (Valid for all VMs when even a lower latency is required.)	<pre> monitor.idleLoopSpinBeforeHalt = "TRUE" monitor.idleLoopMinSpinUS = "50"                     </pre> For a Skylake CPU: <pre> monitor_control.disable_pause_loop_exiting = "TRUE"                     </pre>

## CPU thread matrix examples

Figure 52 shows the CPU thread matrix of a 28-core CPU as a reference when configuring the `sched.vCPUXx.affinty = "Xx-Yy"` parameter. The list shows the start and end ranges required for the "Xx-Yy" parameter (for example, for CPU 5, this would be 280–335).

Figure 52: CPU thread examples for a 28-core/56-thread CPU server

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

Figure 53 shows the CPU thread matrix of a 24 core CPU as a reference when configuring the `sched.vCPUXx.affinty = "Xx-Yy"` parameter. The list shows the starting and end range required for the "Xx-Yy" parameter. For example, for CPU 5 it would be "240-287".

Figure 53: CPU thread examples for a 24-core/48-thread CPU server

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

Figure 54 shows the CPU thread matrix of a **22 core CPU** as a reference when configuring the `sched.vCPUXx.affinty = "Xx-Yy"` parameter. The list shows the starting and end range required for the "Xx-Yy" parameter. For example, for CPU 5 it would be "220-263".

Figure 54: CPU thread examples for a 22-core/44-thread CPU server

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

## 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.

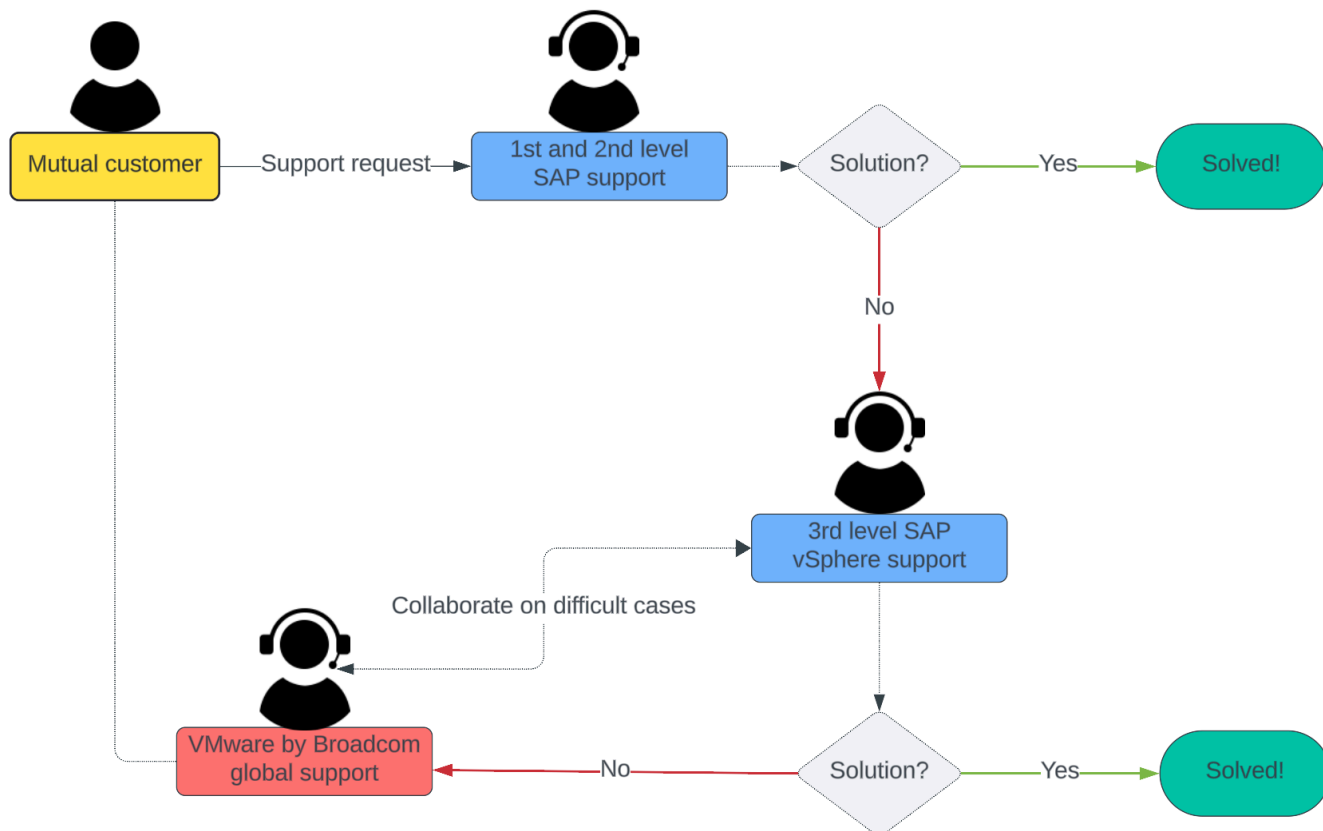
Before opening a VMware support ticket, we recommend opening a support request within the SAP support system when the SAP HANA system runs virtualized with VMware. This ensures that SAP HANA and VMware specialists will work on the case and, if needed, escalate the issue to VMware product support (when 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 ESXi)
- **BC-OP-LNX-ESX** (Linux on VMware ESXi 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 55 shows the support process workflow for VMware-related SAP HANA issues.

Figure 55: SAP support workflow for VMware-related escalations



If the customer message cannot be solved by first- and second-level SAP support, forwards to the next support level.

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 run 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 over 12 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*

## About the Author

**Erik Rieger** is a Principal SAP Global Technical Alliance Manager & Architect, working with VMware by Broadcom's Global SAP Alliance. He defines and manages the VMware SAP solution and validation roadmap. Erik works closely with SAP, partners, and customers to define the architectures of VMware-based SAP solutions so that they have the best mix of functionality to help transform businesses into real-time enterprises. Erik has more than 20 years of experience in the IT sector, has a technical degree in electronics, and a Master of Science degree in Information Systems and Management.

## Acknowledgments

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

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

