

TECHNICAL PAPER: Dec 2024



SAP HANA on VMware vSphere Best Practices

By Erik Rieger

Principal SAP Global Technical Alliance Manager & Architect

Table of contents

Introduction	5
Abstract	5
Audience	5
Solution overview	6
SAP HANA on VMware vSphere	6
Solution components	7
What's new in vSphere 8.0?	7
Benefits of SAP HANA on VMware vSphere 8.0	8
Software and hardware support for SAP HANA on vSphere	8
SAP HANA production support for vSphere and VMware Cloud Foundation	8
vSphere version support for Intel CPU platforms	9
vSphere maximums	10
Release strategy	10
Supported vSphere and VMware Cloud Foundation offerings	11
Scalability and VM sizes	12
Deployment options and considerations	13
Reference architecture diagram	13
SAP HANA on VCF deployment options and sizes	19
Intel sub-NUMA clustering (SNC-2) support and deployment options	22
Support for SNC-2	22
SNC-2 supported deployment examples	24
Scale-up and scale-out deployment architecture	27
Configuration and sizing guidelines	29
Sizing compute and memory	29
Sizing process	29
Translating the SAP sizing to a VM configuration	31
Virtual memory configuration example	32
CPU calculation examples	34
VM CPU (SAPS) configuration example	34
VM vCPU configuration examples	34

Considerations for the final VM configuration.....	35
Pre-sized examples.....	36
Storage configuration and sizing	39
Tools used to verify KPIs	39
Storage connection.....	39
vSphere datastores.....	40
File system layout	40
Storage capacity calculation	41
Hardware and cloud measurement tool for SAP HANA.....	44
Network configuration and sizing	46
Workload performance and validation testing	52
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.....	66
Enhanced vMotion Compatibility, vSphere vMotion, and vSphere DRS best practices	70
Best practices for migrating SAP HANA VMs in a production environment	71
VM migration scenarios	72
Managing SAP HANA landscapes using vSphere DRS	76
SNC-2 management and operation	77
vSphere Cluster Services	80
Architecture of vCLS	80
vCLS VM details	81
Guidelines for deploying vCLS in a cluster with SAP HANA VMs.....	82
Examples of deploying vCLS in a cluster with SAP HANA VMs	83
Mixed SAP HANA and other VMs in a vSphere cluster/VCF VI WLD	83
Dedicated ESXi hosts for SAP HANA VMs on a vSphere cluster	84

SAP HANA HCI on a vSphere cluster	85
High availability best practices for virtualized SAP HANA	87
About vSphere high availability (HA)	88
vSphere HA with passthrough network adapters	92
SAP HANA system replication with vSphere (local site)	92
Disaster recovery	93
SAP HSR and vSphere Replication for a remote site.....	94
Backing up and restoring SAP HANA systems virtualized on vSphere	95
SAP HANA with persistent memory on vSphere 7	96
Prerequisites and general SAP support limitations for Intel Optane Persistent Memory (PMem)	96
What is supported?.....	96
Sizing of Optane PMem-enabled SAP HANA VMs	97
vSphere HA support for PMem-enabled SAP HANA VMs	101
SAP HANA with PMem VM configuration details	102
Monitoring and verifying an SAP HANA installation	107
SAP notes for monitoring data growth/CPU utilization and verifying configuration	107
NUMA tool: VMware NUMA Observer	107
Performance optimizations for SAP HANA VMs	109
Optimizing the configuration parameter list	109
Performance optimization for low-latency SAP HANA VMs	120
CPU affinity settings.....	120
Additional performance tuning settings for SAP HANA workloads.....	121
Examples of VMX configurations for SAP HANA VMs	123
CPU thread matrix examples	127
SAP HANA support and process	130
Open an SAP support request ticket	130
Open a VMware support request ticket	132
Conclusion	132
About the Author	133
Acknowledgments	133

Introduction

Abstract

This guide is the 2024 July edition of the best practices and recommendations for SAP HANA on [VMware Cloud Foundation](#), focusing on [VMware 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 [3rd- and 4th-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 adding more hosts/VMs to a single SAP HANA, typically for a business warehouse instance.

Most of the guidance provided here results from 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 ESXi hosts are still supported and are covered in this document as well.

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

Audience

This guide is intended for SAP and VMware hardware partners, cloud services providers, system integrators, architects, and administrators 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 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, 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, these customers can continue to seamlessly integrate their IT and SAP operations by leveraging their existing IT processes, know-how, and customer-tailored infrastructure. This allows full control of their data and the overall SAP HANA solution.

SAP HANA on VMware vSphere supports scaling database sizes up to 12TB (16TB is planned with 8-socket Sapphire Rapids systems) and scaling out with up to 16 nodes (plus high-availability nodes) with up to 48TB (depending on the host configuration).

The SAP HANA platform with VMware virtualized infrastructure provides an optimal environment for achieving a unique, secure, cost-effective solution. It also offers benefits that physical deployments of SAP HANA do not include, 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 ESXi 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 the abstraction of the hardware layer
- Higher hardware utilization rates
- Automation, standardization, and streamlining of IT operations, 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).

Note: Refer to SAP Note [2104291](#) "FAQ - SAP HANA multi-tenant database containers," page 2.

Solution components

An SAP HANA system based on VMware technologies is a fully virtualized and cloud-ready infrastructure solution running on VMware vSphere 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 virtually, abstracting 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 (out of scope for this document)
 - Optional: Aria Suite for management (out of scope for this document)
- 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 Live Recovery add-on services](#), [Unified Cloud Management](#), 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 had to manually configure 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. To achieve optimal performance before vSphere 8.0, admins had to configure the 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.

vSphere 8.0 also incorporates intelligent, adaptive NUMA scheduling and memory placement policies, eliminating the need for manual VM balancing across nodes. Admins can still use manual controls to override the default behaviors and configure the NUMA placement for performance-critical SAP HANA VMs.

To learn more about 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. The following table lists relevant SAP HANA on vSphere support notes as of March 2024.

Table 1. Relevant SAP notes

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

vSphere version support for Intel CPU platforms

The following table provides an overview, as of March 2024, of the vSphere 7/8 versions and CPU platforms supported by the SAP HANA on vSphere solution.

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

The following table summarizes the critical maximums of the different vSphere versions supported for SAP HANA.

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	< 12TB 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	< 4TB for Ice Lake and Sapphire Rapids 2-socket systems	<= 240 vCPUs	0.5-, 1- and 2-socket wide VMs. SPR 2-socket hosts require SNC for half-socket VMs
vSphere 8 U2 and later	< 8TB for Sapphire Rapids	<= 480 vCPUs	0.5-, 1-, 2-, 3-, and 4-socket wide VMs
vSphere 8 U2 and later	< 12TB 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.

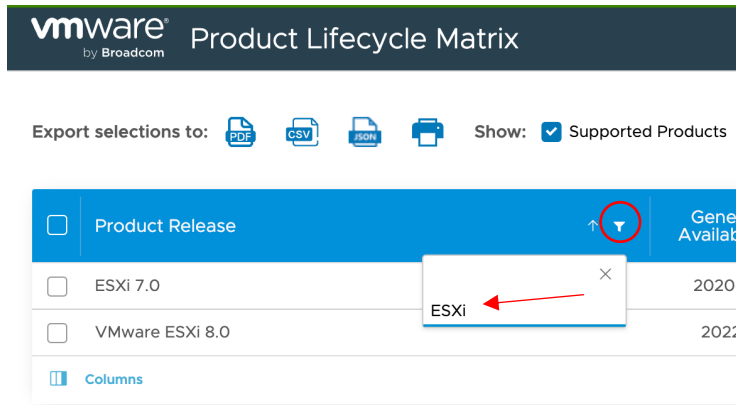
Release strategy

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 balance supporting new customers on the latest hardware and those on an older platform. You may still use vSphere 7 since this version is supported until April 2025. vSphere 8.0 Update 3 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.

To directly show the ESXi versions: Set a filter for `ESXi` at the **Product Release** table heading, as shown in the following figure.

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

vSphere Standard can also be used for smaller SAP application deployments. vSphere Essentials Plus supports Business One solutions. Make sure you are familiar with the feature, CPU, and host limitations of these products.

Notes:

- We highly recommend VMware Cloud Foundation with bundled Select Support for strategic enterprise customers.
- 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.
- This architecture reference guide and the involved test environment use only vSphere dvPortGroups to connect the SAP HANA nodes and non-SAP HANA DB / VM workloads. An upcoming version of this guide will include NSX-based vLANs or overlay segments.
- For the latest SAP-released support information, refer to the [SAP notes related to VMware](#).

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. The CPU power and RAM required for SAP workloads on HANA must be sized appropriately. VMs larger than 4 sockets 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 6TB of memory, 16 half-socket VMs with up to 750GB memory, or one large 12TB SAP HANA VM.

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

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	24TB	
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. If larger 8-socket systems are required, contact your SAP or VMware account team and 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*).

The following table shows the maximum size of a vSphere VM and some other relevant parameters, such as virtual disk size and the number of virtual NICs per VM. These VM limits are higher than the SAP HANA-supported configurations.

Table 5. vSphere guest VM maximums (extract)

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

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

Deployment options and considerations

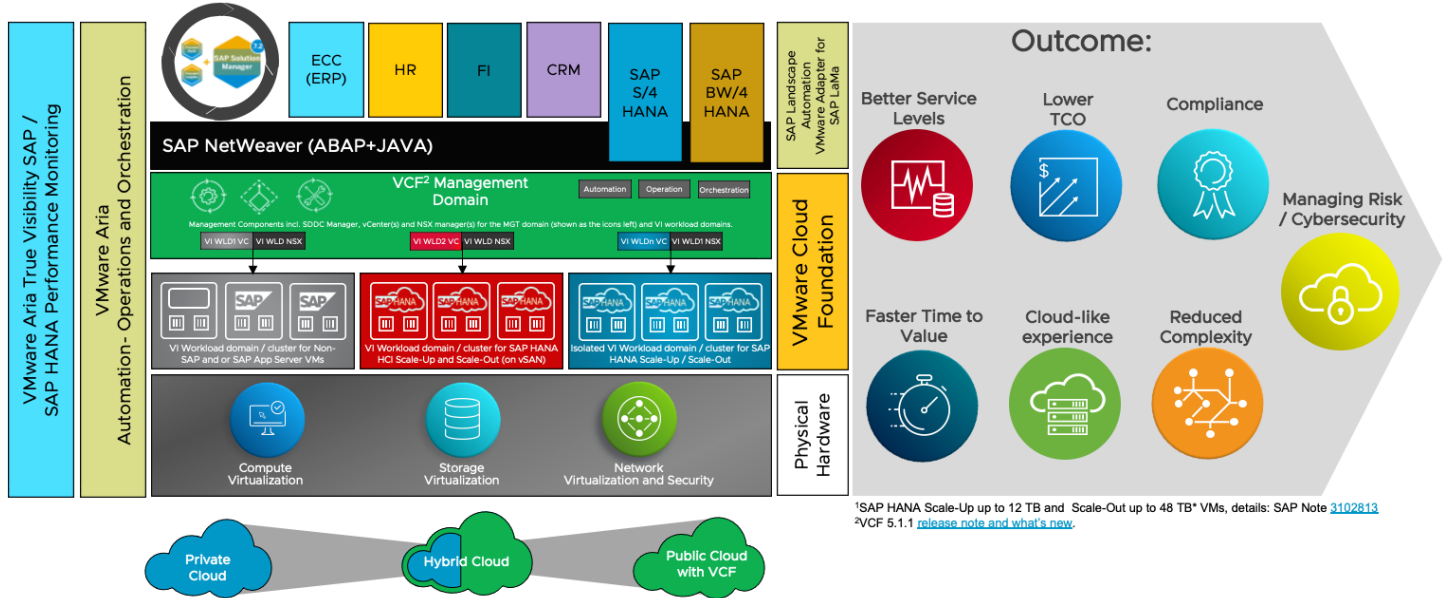
Reference architecture diagram

The following figure shows an overview of a typical VMware software-defined data center (SDDC) for SAP applications. At the center of a VMware SDDC is [VMware Cloud Foundation](#), which includes vSphere, vSAN, and NSX. VCF has VI [Workload Domains](#)—each Workload Domain is a logical unit of application-ready infrastructure that groups ESXi hosts managed by a vCenter Server instance with specific characteristics according to VMware recommended practices.

The figure shows the preferred VCF design for non-SAP HANA database VMs and SAP HANA with separate VCF VI Workload Domains. (You can include an optional, isolated VI Workload Domain.) Each VI Workload Domain has a dedicated NSX Manager cluster to provide the benefits of flexible lifecycle management, which includes the option to run different VCF software stack versions across the different VI Workload Domains. A VI Workload Domain can consist of one or more vSphere clusters provisioned by the SDDC Manager.

The red and blue VI Workload Domains are examples of SAP HANA Workload Domains with different storage types. The grey VI Workload Domain is an example of mixed SAP and non-SAP application Workload Domains with different SLAs/requirements, like ESXi hosts that are not SAP HANA-certified.

Figure 2. VMware Cloud Foundation for SAP applications



The next table provides an overview of the different Workload Domain types (WLDs) and how best to leverage these for SAP and non-SAP VMs/applications.

A VI Workload Domain (VI WLD) shares a vCenter single sign-on domain and the identity provider configuration with the Management Domain. You can manage all VI WLDs through a single pane of glass.

A VI WLD can share an NSX Manager instance with other VI WLDs. We recommend a dedicated NSX Manager per VI WLD for easier and more flexible lifecycle management.

Small and medium SAP deployments could consolidate SAP HANA and non-SAP HANA database VM workloads into a single VCF VI WLD with one or more vSphere cluster configurations to reduce the required hardware footprint, but this would lose the lifecycle management flexibility between different SAP workloads.

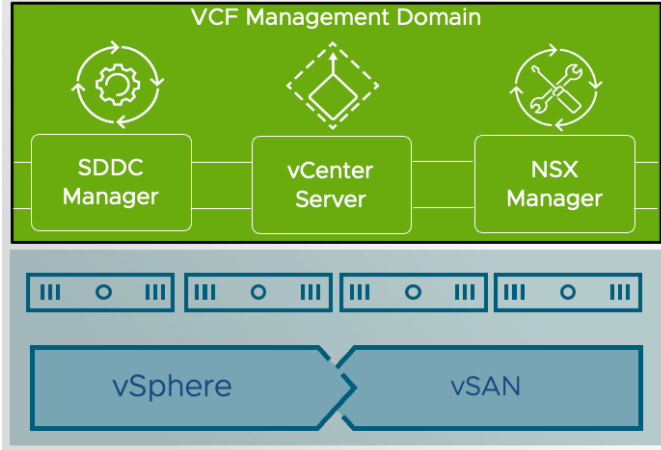
Larger deployments with SAP HANA systems with different operation, hardware, and SLA requirements should not share a VI WLD/cluster with non-SAP HANA VMs. Running these VMs in a dedicated SAP HANA VI WLD/cluster eases lifecycle management and ensures the correct cluster-wide settings for features like HA, DRS, and EVC.

An isolated VI WLD is another option for a distinct vCenter single sign-on domain; this requires an identity provider configuration. VI WLDs of this type cannot yet share an NSX Manager instance with other VI WLDs and, therefore, need a dedicated NSX Manager instance. This allows the independent lifecycle management of isolated VI WLDs and is recommended when strict isolation of user and business data is required.

Small SAP deployments could leverage the consolidated VCF architecture model, which consolidates the VCF Management components (vCenter, NSX Manager, and SDDC Manager) and all SAP workloads (non-SAP HANA database VMs and SAP HANA) with the lowest possible hardware footprint, but this loses the lifecycle management flexibility between the VCF Management Domain and the VI WLD running the different SAP workloads. If you want to deploy SAP HANA production-level VMs in a consolidated WLD, you must ensure these SAP HANA VMs run on SAP HANA TDI-supported systems and do not share a NUMA node with a non-SAP HANA VM. We advise configuring dedicated vSphere clusters to manage VCF VMs and workload VMs. For a consolidated WLD, we recommend separate cluster-wide settings for management and SAP NetWeaver/SAP HANA hosts.

Note: Every vSphere cluster requires at least 2 embedded vCLS VMs or 3 external vCLS VMs. A cluster allows different cluster-wide settings; however, we recommend you add non-SAP HANA hosts to an SAP HANA cluster to offload the vCLS VMs.

Table 6. VMware Cloud Foundation Management and VI Workload Domain types and use cases

VCF Management Domain	
<p>First domain deployed, runs on dedicated ESXi hosts (min. 4)</p> <p>Contains the following management appliances for <u>all</u> Workload Domains:</p> <ul style="list-style-type: none"> • vCenter Server • NSX Manager • SDDC Manager • Optional: VMware Aria Suite components, Management Domain NSX Edge nodes <p>The VI WLD vCenter and NSX Manager VMs are installed on the Management Domain hosts.</p> <p>Depending on the Workload Domain type, you can scale from 14 to 24 VI WLDs. A total of 1000 ESXi hosts are supported per VCF instance following the VCF standard architecture, where a single VI WLD can scale up to 800 ESXi hosts and up to 4000 registered VMs.</p> <p>A single VCF instance can have multiple availability zones (a stretched deployment).</p>	 <p>The diagram illustrates the VCF Management Domain architecture. At the top, a green box labeled 'VCF Management Domain' contains three interconnected components: SDDC Manager (with a gear icon), vCenter Server (with a diamond icon), and NSX Manager (with a wrench icon). Below this, a row of four server racks is shown, representing the hosts for the management domain. At the bottom, two blue boxes labeled 'vSphere' and 'vSAN' are connected by a double-headed arrow, indicating their integration with the management domain.</p>

Consolidated VCF SAP Domain

Represents a Management Domain that also runs customer workloads. It's limited to 100 hosts per consolidated VI WLD.

Uses resource pools to ensure sufficient resources for management components.

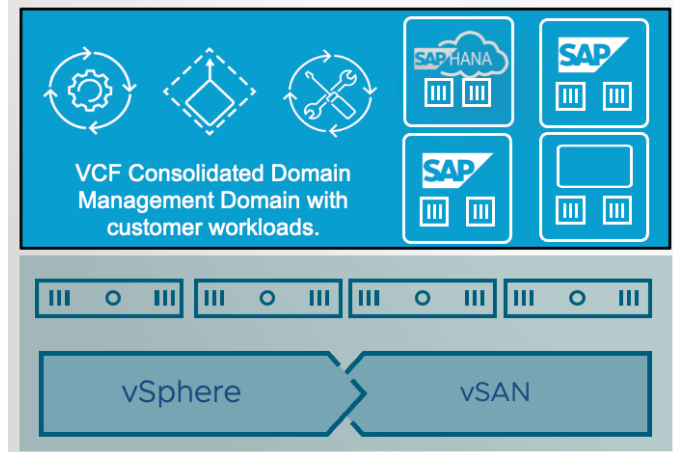
Considers the minimum possible initial hardware and management component footprint (min. 4 hosts). Only with vSAN and [SAP HANA HCI certified hosts](#) and the listed configurations for CPU, memory, and storage.

Can be scaled to a standard architecture model.

Management components and customer workloads are not isolated.

SAP HANA production-level workloads must not run on the hosts that run the management components or non-SAP HANA VMs.

This is an option for small deployments or test and dev environments. **Add HANA dedicated hosts as required if you want to deploy production-level SAP HANA VMs.**



SAP/SAP HANA VI Workload Domain

Represents an additional Workload Domain for running customer workloads. Can scale up to 14 VI Workload Domains per VCF instance.


Shares a vCenter single sign-on domain and the identity provider configuration with the Management Domain. **All Workload Domains can be managed through a single pane of glass.** Can share an NSX Manager instance with other VI Workload Domains. A dedicated NSX Manager per VI WLD is recommended for easier lifecycle management.

Allows for independent life cycle management of the different VI WLDs.

This Workload Domain type cannot provide distinct vCenter single sign-on domains for customer workloads. For this, an isolated VI WLD is required.

If SAP applications are deployed, only SAP-supported hosts can be used. SAP HANA hosts must also be SAP HANA TDI-supported systems and are limited to 2-, 4-, or 8-socket ESXi hosts with validated Intel CPUs.

This Workload Domain type cannot provide distinct vCenter single sign-on domains for customer workloads. For this, an isolated VI WLD is required.



If SAP applications are deployed, then only SAP-supported hosts can be used. SAP HANA hosts must also be [SAP HANA TDI supported systems](#) and are limited to 2-, 4- and 8-socket ESXi hosts with validated Intel CPUs. VM maximum host size: ≤ 12TB and 480 vCPUs per single SAP HANA scale-up VM. The smallest size is 0.5 socket with 128GB RAM. Scale-out deployment: VM must be ≥ 4 sockets, maximum 8 sockets, VM vRAM up to 6TB per VM depending on the use. Maximum 8 hosts + HA for a total size of 48TB. For shared ESXi or VM connected SAN or NAS storage w/o vVOLS, the storage must meet HANA KPIs. [HCI certified hosts](#) are required when vSAN is used. SAN or NAS deployments require min. 2 hosts; vSAN systems require min. 3 hosts.

SAP NetWeaver and SAP HANA hosts can run consolidated in a VI WLD if the same SLAs and lifecycle management requirements exist. If not, then use dedicated VI WLDs for SAP HANA and NetWeaver or AnyApp.

SAP/SAP HANA Isolated VI WLD

Represents an additional Workload Domain for running customer workloads. Can scale up to 24 VI Workload Domains per VCF instance

Has a distinct vCenter single sign-on domain and identity provider configuration. Workload domain vCenter Server instances are managed through different panes of glass. Can provide distinct vCenter single sign-on domains for customer workloads.

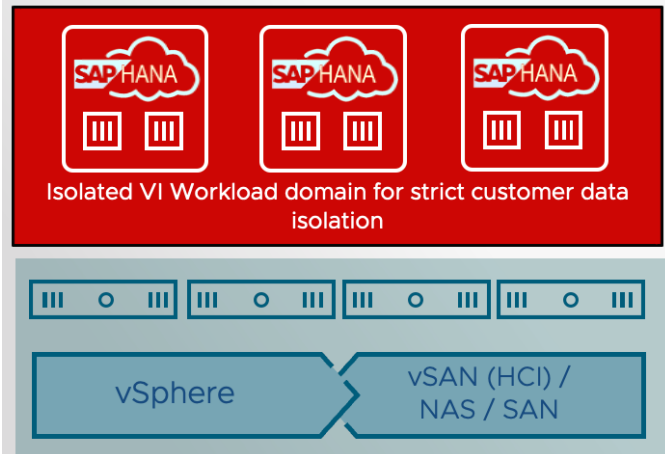
Workload domains of this type cannot share an NSX Manager instance with other VI Workload Domains.

Allows for independent life cycle management of the isolated VI WLDs and is the **recommended option when a strict isolation of user and business data is required.**

If SAP applications are deployed, then only SAP-supported hosts can be used. SAP HANA hosts must also be [SAP HANA TDI supported systems](#) and are limited to 2-, 4-, and 8-socket ESXi hosts with validated Intel CPUs.

VM maximum host size: ≤ 12TB and 480 vCPUs per single SAP HANA scale-up VM. Smallest size is 0.5 socket with 128GB RAM.

Scale-out deployment: VM must be ≥ 4 sockets, maximum 8 sockets, VM vRAM up to 6TB per VM depending on the use. Maximum 8 hosts + HA for a total size of 48TB.



For shared ESXi or VM-connected SAN or NAS storage w/o vVOLS, the storage must meet HANA KPIs. [HCI certified hosts](#) are required when vSAN is used.

SAP NetWeaver and SAP HANA hosts can run consolidated in a VI WLD if the same SLAs and lifecycle management requirements exist. If not, then use dedicated VI WLDs for SAP HANA and NetWeaver or AnyApp.

Depending on the [Workload Domain type](#), you can scale from 14 (ELM) to 24 (isolated) VI WLDs. Refer to [Maximums and Configuration Limitations](#), [VCF Deployment Options](#) (for example, deployments in multiple availability zones), and [vSphere Design for VMware Cloud Foundation](#).

In total, [1000 ESXi hosts](#) are supported per VCF instance following the VCF standard architecture, where a single VI WLD can scale up to 800 ESXi hosts and up to 4000 registered VMs.

As shown in the figure and table above, the VI WLD vCenter and the NSX Manager VMs are installed in the VCF Management Domain among the vCenter and the NSX Manager VMs to manage the VCF Management Domain.

Review the [Maximums and configuration limitations](#) and [VCF Deployment Options](#) (for example, deployments in multiple availability zones for configuration details).

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. Only the ESXi 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 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 implemented design options in VCF, like vSAN or NSX networking and security features, will 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 6TB can be chosen. 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, optimizing 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.

An SAP HANA system deployed in VMware Cloud Foundation allows you to adopt cloud-like virtual networking with NSX overlay segments, and it protects your business-critical SAP workloads with the NSX Distributed Firewall (DFW).

Note: This reference guide does not include the adoption of NSX vLAN, overlay segments, and NSX Distributed Firewall (DFW). These NSX deployment options require additional consideration and are currently out of the scope of this document.

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 ESXi host, or multiple SAP workloads/VMs can be consolidated on a single ESXi 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.

SAP HANA on VCF 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) can 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](#)).

The following table 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 guidelines for 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 with the maximum core count and memory support available for SAP HANA workloads. 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, deviating from these memory configurations is possible when an SAP HANA expert workload-based sizing is performed.

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

Legend:




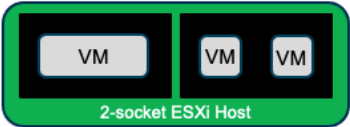
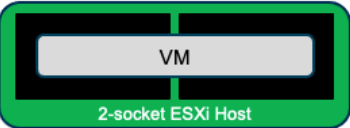
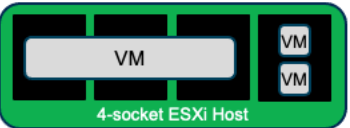


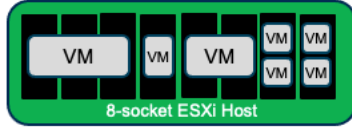
-  <=8-socket ESXi Host
-  1 physical CPU socket / NUMA Node
-  SAP HANA VM

Table 7. Overview of SAP HANA on vSphere deployment options and possible VM sizes

<p>2-socket host examples</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"> • Broadwell • Skylake, Cascade Lake, and Cooper Lake • Ice Lake <p>vSphere 8.0 supported CPU generations:</p> <ul style="list-style-type: none"> • Cascade Lake and Cooper Lake • Ice Lake • Sapphire Rapids (half-socket support only with SNC) <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 >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"> • 1GbE IPMI, 1 GbE ESXi Management network • >= 10GbE vMotion / HA • >= 10GbE HANA to App Server network • Plus optional networks like those for backup or replication
<p>4-socket host example</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"> • Broadwell • Skylake, Cascade Lake, and Cooper Lake <p>vSphere 8.0 supported CPU generations:</p> <ul style="list-style-type: none"> • Cascade Lake and Cooper Lake • Sapphire Rapids <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 >6 TB (SAP HANA standard</p>

	<p>memory size) with Cascade and Cooper Lake. Larger memory sizes based on a workload-based sizing is possible.</p> <p>Recommended network configuration per host (dual-NIC port configuration):</p> <ul style="list-style-type: none"> • 1GbE IPMI, 1GbE ESXi Management network • >= 25GbE vMotion / HA • >= 10GbE HANA to App Server network • Plus optional networks like those for backup or replication
<p>8-socket host examples</p> <p>Example 4: 1 HANA VM across 8 NUMA nodes</p>  <p>Example 5: 2 HANA VM across 4 NUMA nodes</p>  <p>Example 6: Wide VMs, single and half-socket VMs on one host</p> 	<p>vSphere 7.0 supported CPU generations:</p> <ul style="list-style-type: none"> • Broadwell • Skylake, Cascade Lake, and Cooper Lake <p>vSphere 8.0 supported CPU generations:</p> <ul style="list-style-type: none"> • Cascade Lake and Cooper Lake • Sapphire Rapids is not currently available on 8-socket hosts, but is undergoing the validation process <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 >12 TB (SAP HANA standard memory size) with Cascade and Cooper Lake.</p> <p>Note: 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"> • 1GbE IPMI, 1GbE ESXi management network • >= 25GbE vMotion / HA • >= 10GbE HANA to App Server network • Plus optional networks like those for backup or replication

SAP HANA half-socket VM support on 2-socket Sapphire Raids hosts requires Intel sub-NUMA Clustering (SNC-2) due to higher measured deviations while running two SAP VMs on 2-socket Sapphire Rapids hosts.

However, 4-socket Sapphire Rapids systems support half-socket SAP HANA VMs without requiring SNC-2. You can deploy two SAP HANA VMs, as in previous CPU generations.

SAP HANA half-socket VM support on non-SNC 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.

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

Support for SNC-2

Support for SNC-2 in SAP HANA considers the increasing density of processors, memory controllers, processor interconnects, and supporting infrastructure within a single chip as the CPU size 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 data flow 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 have equal access to the last level cache and memory. In this scenario, [a processor might access a memory controller or a portion of the last level cache located on the opposite side of the CPU chip compared to the nearest](#).

Accessing and moving data within a CPU chip takes longer, and the memory subsystem is 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. A Sapphire Rapids CPU performs much better than a Cascade Lake CPU. This includes half-socket deployments.

However, there are more fluctuations when running several SAP HANA VMs on a socket. Because of these fluctuations, SAP does not support half-socket deployments for Sapphire Rapids CPUs. The following table shows this. It also shows a Sapphire Rapids CPU's theoretical available memory channels per CPU core.

The table doesn't reflect the change in increased memory speed. It only shows that the available memory channels are more utilized because of the increased CPU core count per chip and the higher CAS latencies when first accessing data stored in memory.

Table 8. 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 [1]	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%

* [1] For more information, refer to this article: [The difference between RAM speed and CAS latency.](#)

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**.)

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

Two localization domains exist within a processor in a two-cluster SNC (SNC-2). Each domain has addresses mapped from the local memory controller and local Last Level Cache (LLC) slices. Processors in the local domain 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. This is the SAP-supported SAP HANA on vSphere configuration consisting of SNC Domains 0 and 1 and 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, limiting a VM to 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 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: SAP and VMware don't support any other configurations for SAP HANA and SAP application deployments. They only support non-SNC-2 configurations as full-socket VM configurations.

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 2TB of memory (4TB 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 they require SAP application/workload sizing.

Note: The following configuration options for running SAP HANA on vSphere are not supported on bare metal servers enabled for SNC. Sapphire Rapids 4-socket hosts support half-socket deployments without SNC-2. Therefore, the following table isn't relevant for >2-socket Sapphire Rapids hosts. It only applies to 2-socket Sapphire Rapid hosts.

Legend:





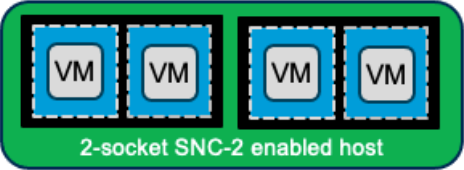
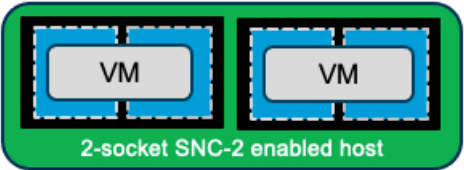
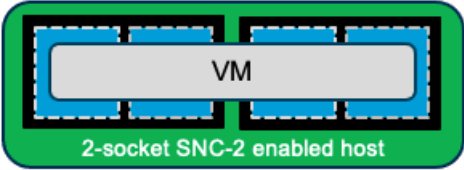
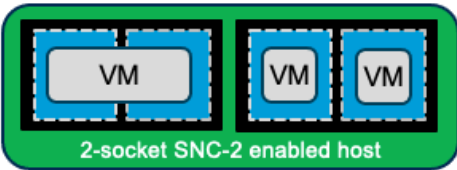
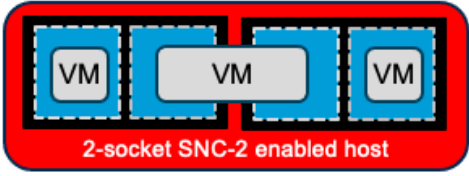
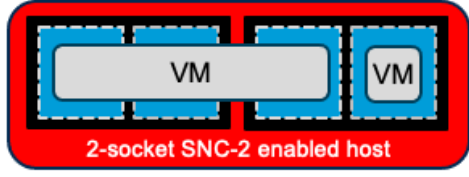
-  2-socket ESXi Sapphire Rapids host or later
-  1-CPU socket / NUMA Node
-  Sub NUMA Node (SNC-2)
-  VM

Table 9. SAP HANA-supported SNC-2 configurations

<p>Option 1: 2 VMs per NUMA node, 4 VMs in total per ESXi host</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"> 4 VMs with <=60 vCPUs, <1 TB vRAM, and 1 vSocket  <p style="text-align: center;">2-socket SNC-2 enabled host</p>	<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>Attention: 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>Option 2: 1 VM per NUMA node, 2 VMs in total per ESXi host</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"> 2 VMs with <=120 vCPUs, <2 TB vRAM and 2 vSockets  <p style="text-align: center;">2-socket SNC-2 enabled host</p>	<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>Attention: 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 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>Option 3: 1 VM across two NUMA nodes/physical CPU sockets</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"> • 1 VM with <=240 vCPUs, <4 TB vRAM and 4 vSockets 	<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 <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>Attention: 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>Option 4: 1 VM across two NUMA nodes/ physical CPU sockets</p> <p>VM size is based on the SAP HANA reference configuration for 2-socket SPR systems:</p> <ul style="list-style-type: none"> • 1 VM with <=120 vCPUs, <2 TB vRAM and 2 vSockets • 2 VM with <=60 vCPUs, <1 TB vRAM and 1 vSocket 	<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>Attention: 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 primarily using SNC-2 for half-socket VMs. If SAP HANA VMs running on SNC-2-enabled hosts experience performance issues, they should be migrated to a non-SNC-enabled host. All VMs that leverage SNC-2 require the VMX advanced parameter `sched.nodeX.affinity="Y"` to prevent unwanted NUMA node migrations.

If an SAP HANA VM running on an SNC-2-enabled host shows performance issues, migrate the VM to a whole 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 with SNC-2-enabled hosts and VMs, along with non-SNC-2 hosts and VMs.

Scale-up and scale-out deployment architecture

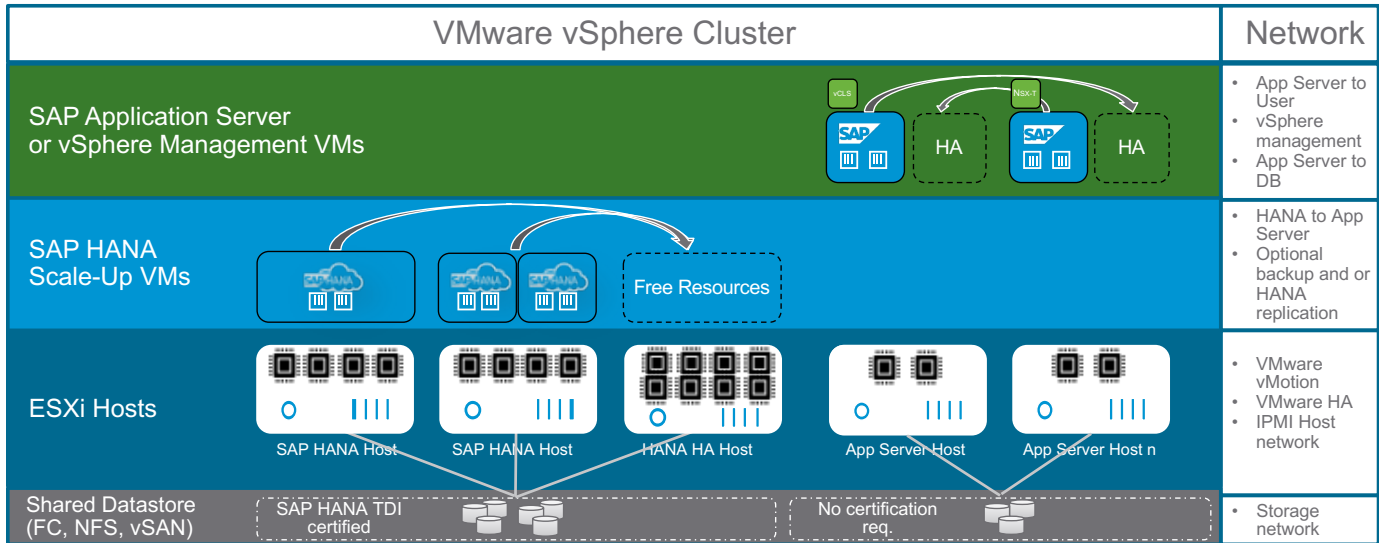
The following 2 figures describe typical scale-up and scale-out SAP HANA on vSphere architectures.

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 2TB SAP HANA VM or a 12TB 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.

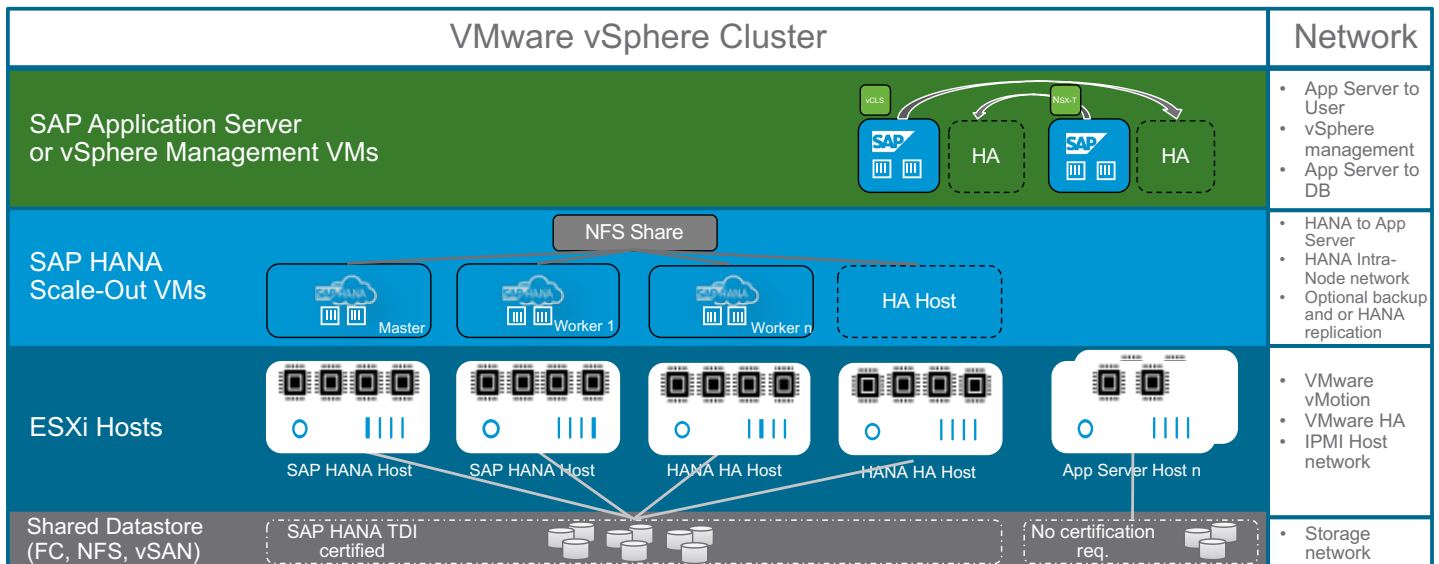
The following figure shows scale-up SAP HANA systems (single SAP HANA VMs).

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



The next figure 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. The server configuration (RAM and CPU), network configuration, and storage configuration determine which and how many SAP HANA VMs can be supported.

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 have the capacity 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 significant 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 for 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 cannot 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 might also be able to deploy larger scale-out systems.

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). The following figure 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

Note: By now, the SAP Quick Sizer allows you to calculate the needed SAPS for SAP HANA-based applications. Therefore, a manual calculation, as mentioned above in the table, is no longer required.

The SAP sizing tools do not cover network sizing. Network latency is only a general guideline—you must set your own goal. In [SAP sales and distribution \(SD\) benchmarks](#), a time below 1,000 milliseconds for the average dialog response time must be maintained.

In the “[Network configuration and sizing section](#),” 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, 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 more challenging to determine.

The following table 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. The physical server configuration influences them—for example, the number of CPUs, memory size, and installed cards like HBAs or NICs—as well as the vSphere features like NSX networking and security or vSAN in use.

Table 10. 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	64–256GB (default 128GB)	2TB 3TB	128GB 192GB
4	128–384GB (default 256GB)	4TB 6TB	256GB 384GB
8	256–768GB (default 512GB)	8TB 12TB	512GB 768GB

Note: The ESXi host memory consumption described here can range from 64GB to 768GB. However, determining the actual host memory consumption requires configuring and starting all VMs with memory reservations on the host. The last VM started might fail to run if you selected the wrong host memory reservation. In this case, you should reserve a lower amount of memory per VM to ensure 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:

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 6TB of total RAM and a 2-socket server with 2TB of RAM are available for selection. We use the default ESXi memory requirements:

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

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 6TB physical memory
- 2-CPU socket 60-core Sapphire Rapids ESXi host with 2TB physical memory

Virtual RAM calculation for the SAP HANA VM:

- **5760 GB** (6144GB – 384GB RAM) / **4 CPU sockets** = 1440GB per CPU socket > 1400GB sized HANA RAM → SAP HANA VM configuration needs 1 CPU socket
- **1920 GB** (2048GB – 128GB) / **2 CPU sockets** = 960GB per CPU socket < 1400GB 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 960GB 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, you can't determine the maximum available memory or the memory reservation for a VM until the VM has been created and started once. The amount of memory that is available depends on how the ESXi host's hardware is set up and what ESXi features are enabled and used.

After you've configured the memory for the SAP HANA VM, you verify that the available SAPS capacity is enough for the work your planned workload.

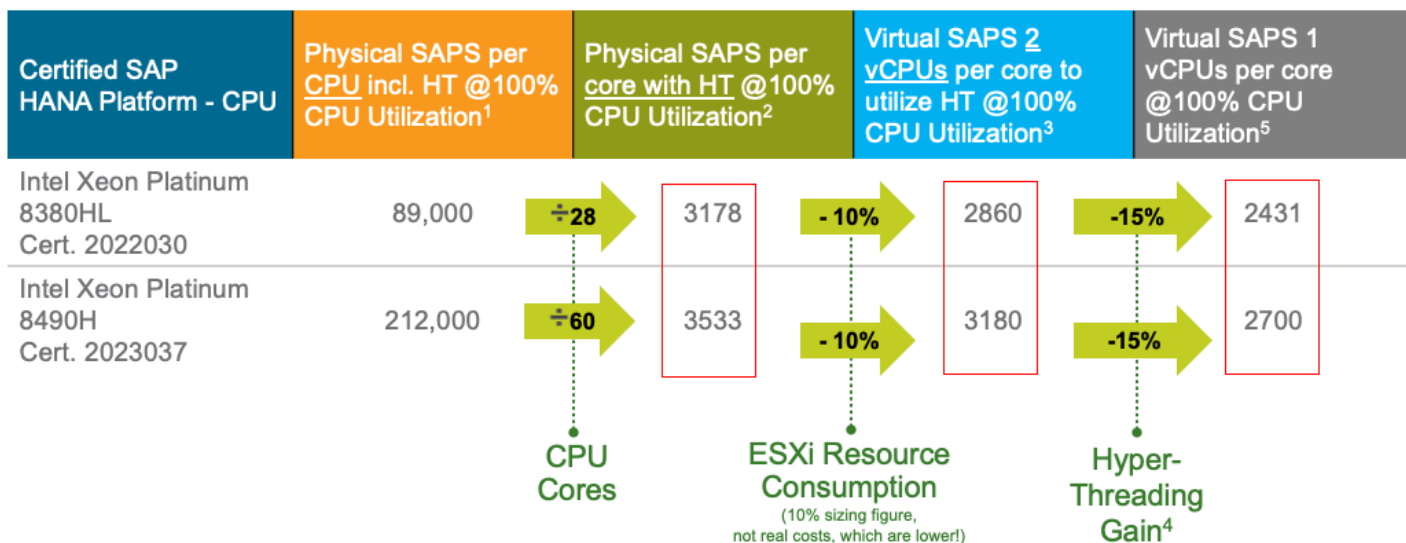
SAP and its partners measure a server's [SAPS](#) capacity and CPU by running [benchmarks and performance tests that SAP defines](#). One example is the SAP SD benchmark, which is used for all SAP NetWeaver and BW/4HANA applications. The test results of a public benchmark can be published and used in training on SAP sizing.

After you know the SAP application's SAPS resource needs, you can apply the sizing to a VM configuration.

Note: The SAP admin usually gives the CPU and SAPS requirements to the vSphere admin. If these are not available, you can use the SAP-defined standard sizes (SAP HANA appliance configurations), which have a fixed CPU type and the max memory configuration per CPU socket.

The following figure, 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 the above figure, 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 the above figure, 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 the number of vCPUs needed to power the provided SAPS resource needs of the given SAP HANA workload.

VM CPU (SAPS) configuration example

Assumed/sized SAP HANA system:

- 1400GB RAM SAP HANA system memory
- 80,000 SAPS
- Available host servers:
 - 4-socket 28-core Cooper Lake (8380HL) host with 6TB physical memory, 89,000 pSAPS (356,000 SAPS in total), and 224 logical CPUs
 - 2-socket 60-core Sapphire Rapids (8490H) host with 2TB physical memory, 212,000 pSAPS (424,000 SAPS in total), and 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: number of vCPUs = $80,000$ SAPS / 2431 SAPS = 32,91 vCPUs, rounded up 33 cores / vCPUs
- VM with HT enabled: number of vCPUs = $80,000$ SAPS / 2860 SAPS x 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 is 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: number of vCPUs = $80,000 \text{ SAPS} / 2700 \text{ SAPS} = 29,63 \text{ vCPUs}$, rounded up 30 cores / vCPUs
- VM with HT enabled: number of 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. Only half of a 60-core CPU is needed for the Sapphire Rapids CPU.

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

The sized SAP HANA database memory of 1400GB is the leading HANA sizing and VM configuration factor. On the Cooper Lake system, we have 1440GB available per CPU. This fulfills the memory and CPU requirements as calculated. In contrast, the Sapphire Rapids system has only 2TB installed, which, per CPU socket minus the ESXi host, costs only 960GB. 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 poor performance. You could solve this problem by adding memory to this host.

The final VM configuration leverages 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) 6TB host:** 1 CPU socket with 56 vCPUs (hyperthreading on) and 1440GB of virtual RAM
- **2-socket 60-core Sapphire Rapids (8490H) 2TB host:** 2 CPU sockets with 240 vCPUs (hyperthreading on) and 1920GB of 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:

- 6TB per host and the possibility to run up to 4 VMs on this host, with each VM of around 1.5TB
- 6TB in the host with 2 VMs, each around 3TB, will result in low CPU performance (high CPU utilization). Therefore, if a half-socket SNC-2 deployment cannot be used, you can use a lower-bin Sapphire Rapids CPU.

- If you want to leverage SNC-2, then an SAP HANA half-socket VM with 1.5TB 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

The following table shows possible VM sizes of specific CPU types and variants. We estimated the SAPS figures based on published SAP SD benchmarks and rounded them down. You'll need to subtract the RAM needed for the ESXi host 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 the following table 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 Sapphire Rapids 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 the following table 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 11. SAPS capacity and memory sizes for example configurations of SAP HANA on vSphere; based on published SD benchmarks and selected Intel CPUs

	Intel Xeon Platinum 8280L CPU	Intel Xeon Platinum 8380H and 8380HL CPU	Intel Xeon Platinum 8380 CPU	Intel Xeon Platinum 8490H CPU
SAP benchmark	2021009 (Jan. 26, 2021)	2020050 (Dec. 11, 2020)	2023019 (May 5, 2023)	2023037 (Aug. 15, 2023)
Max. supported RAM per CPU as of Intel datasheet	4.5TB	4.5TB (HL)	6TB	4TB
CPU cores per socket as of Intel datasheet	28	28	40	60
Max. NUMA nodes per ESXi server[1]	8	8	2	8

SAP HANA on VMware vSphere Best Practices

	Intel Xeon Platinum 8280L CPU	Intel Xeon Platinum 8380H and 8380HL CPU	Intel Xeon Platinum 8380 CPU	Intel Xeon Platinum 8490H CPU
SAP HANA supported maximum memory per CPU socket[2]	1.5TB	1.5TB	2TB	3TB
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[3]	1 to 16 x 14 physical core VM with min. 128GB RAM and max. 768GB	1 to 16 x 14 physical core VM with min. 128GB RAM and max. 768GB	1 to 4 x 20 physical core VM with min. 128GB RAM and max. 1024GB	1 to 4 x 30 physical core VM with min. 128GB RAM and max. 1536GB - with SNC-2
	vSAPS 34,000	vSAPS 34,000	vSAPS 51,255	vSAPS 95,400
	28 vCPUs	28 vCPUs	40 vCPUs	60 vCPUs
1-socket SAP HANA VM[4]	1 to 8 x 28 physical core VM with min. 128GB RAM and max. 1,536GB	1 to 8 x 28 physical core VM with min. 128GB RAM and max. 1,536GB	1 to 2 x 40 physical core VM with min. 128GB RAM and max. 2048GB	1 to 2 x 60 physical core VM with min. 128GB RAM and max. 3072GB
	vSAPS 81,000 56 vCPUs	vSAPS 80,000 56 vCPUs	vSAPS 120,600 80 vCPUs	vSAPS 190,800 120 vCPUs
2-socket SAP HANA VM[4]	1 to 4 x 56 physical core VM with min. 128GB RAM and max. 3072GB	1 to 4 x 56 physical core VM with min. 128GB RAM and max. 3072GB	1 x 80 physical core VM with min. 128GB RAM and max. 4096GB	1 x 240 physical core VM with min. 128GB RAM and max. 6144GB
	vSAPS 162,000 112 vCPUs	vSAPS 160,000 112 vCPUs	vSAPS 241,200 160 vCPUs	vSAPS 381,600 240 vCPUs
4-socket SAP HANA VM[4]	1 to 2 x 112 physical core VM with min. 128GB RAM and max. 6128GB	1 to 2 x 112 physical core VM with min. 128GB RAM and max. 6128GB	NA (ICX is a 2-socket platform only)	1 x 480 physical core VM with min. 128GB RAM and max. 8192GB
	vSAPS 324,000 224 vCPUs	vSAPS 320,000 224 vCPUs		vSAPS 763,200 480 vCPUs
8-socket SAP HANA VM[4]	1 x 224 physical core VM with min. 128GB RAM and max. 12,096GB	1 x 224 physical core VM with min. 128GB RAM and max. 12,096GB	NA (ICX is a 2-socket platform only)	Not available yet
	vSAPS 648,000 448 vCPUs	vSAPS 640,000 448 vCPUs		

- * [1] 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.
- * [2] SAP-supported appliance memory configurations are lower as specified by Intel. Cascade Lake and Cooper Lake CPUs support up to 1.5TB per CPU socket, Ice Lake up to 2TB, and Sapphire Rapids CPUs up to 3TB per CPU socket. 2TB are supported without a sizing; 3TB are supported in a mixed DIMM configuration with 128GB and 256GB DIMMs and require a workload-based SAP HANA sizing.
- * [3] 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.
- * [4] 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.

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

Sized SAP HANA system:

- 1400GB RAM SAP HANA system memory
- 80,000 SAPS (Suite on HANA SAPS)

Configuration example 1:

- 4-socket 28-core Cooper Lake (8380HL): **1.5TB per CPU** host; total 6144GB
- 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); 1400GB vRAM (optimal configuration)

Configuration example 2:

- 2-socket 60-core Sapphire Rapids (8490H) with **1TB per CPU** host; total 2048GB
- 240 vCPUs; 381,000 vSAPS available; **95,000 vSAPS per half-socket** by leveraging hyperthreading
- VM configuration: 2 sockets; 240 vCPUs; 1400GB vRAM (non-optimal configuration due to the memory configuration of the host). We recommend you upgrade the memory to 3TB per CPU.

Note: As described, the performance impact of using hyperthreading is very little, especially for systems with low CPU utilization.

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. You need a valid SAP user account to download the tools, the KPI documentation, and the user guides. 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. For more information, 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. If no VMware-specific guidelines are available, use the guidelines for physically deployed SAP HANA systems as a basis, 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 the raw device, in-guest mapped LUNs, and in-guest mounted NFS storage solutions.

However, a fully virtualized storage solution works just as well 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 of vSphere storage solutions, refer to the [vSphere Storage 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 VMs. 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](#) documentation for details.

The following table summarizes the vSphere features supported by the different storage types. All these storage types are available for virtual SAP HANA systems.

Note: For scale-out, you must install 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 logs, you can use all storage types listed in the following table, as long as the SAP HANA TDI Storage KPIs are achieved per HANA VM. vSphere doesn't support other solutions, such as the Oracle Cluster File System (OCFS) or the IBM General Parallel File System (GPFS).

Table 12. 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 [1]	Yes	Yes	vSAN 6.6 or later	No	Yes

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

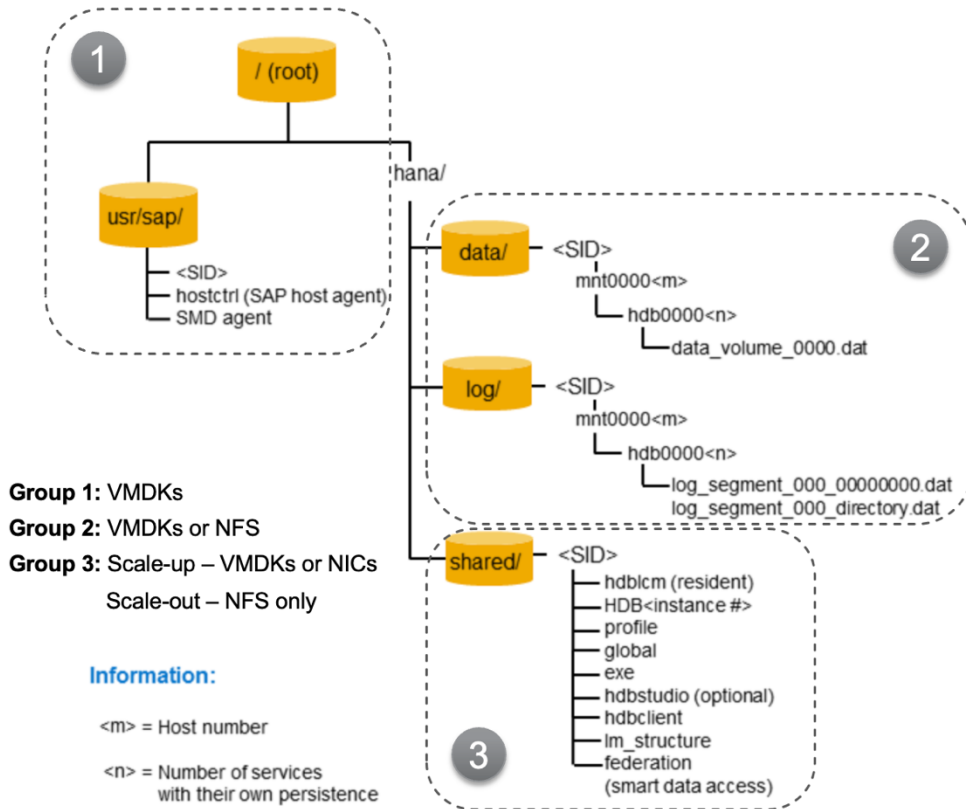
In summary, you use storage with SAP HANA on vSphere to:

- 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

The following figure shows the recommended SAP HANA Linux file system layout when running SAP HANA on vSphere. We arranged the file system layout into three groups to help you decide when to use VMDK files or an NFS mount point to store the SAP HANA files and directories.

Figure 7. SAP-recommended file system layout for SAP HANA on vSphere



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.

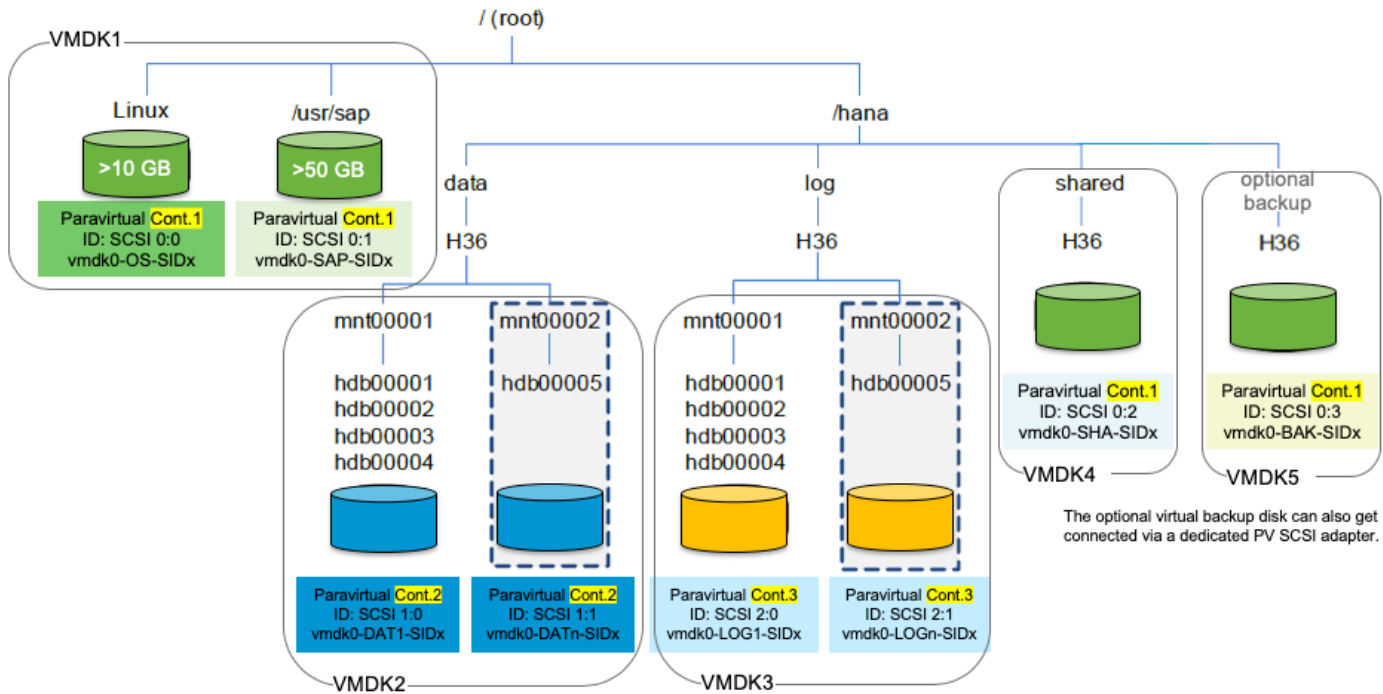
The storage subsystem configuration must meet very strict SAP-defined performance KPIs. This might result in more storage capacity than needed (even if the disk space is not needed, the number of disks might be required to provide the required I/O performance and latency).

SAP has published several architecture and sizing guidelines. The following figure and table consolidate that information—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 for your SAP HANA on vSphere system. The volumes shown in the next figure should correspond with actual VMDK files and dedicated 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 here uses three independent PVSCSI adapters or NVME controllers and at least 4 independent VMDKs. This helps parallelize I/O streams by providing the highest operational flexibility. We also use the file system layout shown here and translate it into a disk volume/VMDK disk configuration.

Figure 8. Storage layout of an SAP HANA on vSphere system [1]



* [1] Image sourced from SAP SE and modified.

The following table provides storage capacity examples of the different SAP HANA volumes. Some 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.

Table 13. 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 [1]
/ (root)	VMDK	PVSCSI (or NVMe) Contr. 1	vmdk01-OS-SIDx	SCSI 0:0	Min. 10GB for operating system We suggest 100GB, thin-provisioned.
usr/sap	VMDK	PVSCSI (or NVMe) Contr. 1	vmdk01-SAP-SIDx	SCSI 0:1	Min. 50GB for SAP binaries We suggest 100GB thin-provisioned.
shared/	VMDK or NFS	PVSCSI (or NVMe) Contr. 1	vmdk02-SHA-SIDx	SCSI 0:2	Min. 1x RAM, max. 1TB, 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 Note: 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 <= 512GB] log volume (min) = 0.5 x RAM [systems >= 512GB] log volume (min) = 512GB (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 Note: 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.

* [1] Information summarized from SAP OSS note [1900823](#), attachment: SAP HANA storage whitepaper version 2.1

Depending on the storage solution used and the 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) 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 shown in the above figure and table.

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

Note: The raw storage capacity needed 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 the following equation for your storage capacity calculation:

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 = 2TB would need the following:

- VMDK root (OS) >= 10GB (recommended 100GB, thin-provisioned)
- VMDK usr/sap >= 60GB (recommended 100GB, thin-provisioned)
- VMDK HANA shared = 1024GB (thick-provisioned)
- VMDK HANA log = 512GB (thick-provisioned)
- VMDK HANA data = 2048GB (thick-provisioned)
- VMDK HANA backup >= 2.5TB (thin-provisioned; optional)

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

To determine the cluster-wide storage capacity requirement, you must include all possible SAP HANA VMs running in the vSphere cluster in the calculation. All SAP HANA production VMs must fulfill the capacity, 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 previous 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 that the storage performance and latency of an SAP HANA VM fulfills the SAP-defined KPIs for logs 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.

The following table 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 1,000µs. 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 1,000µs, at 406µs, which shows good performance.

Table 14. 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 = 334µs
VM with FC-connected Pure X50 storage unit	16K block log overwrite latency = 406µs
Absolute difference	72µs
Percentage difference	22%
SAP HANA KPI	1,000µs (bare metal 3x and virtual 2.5x faster, as required)

The next table shows that 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 120MBps for this test.

Table 15. HCMT file system overwrite example of an 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 = 706MBps
VM with FC-connected Pure X50 storage unit	16K block log overwrite throughput = 723MBps
Absolute difference	17MBps
Percentage difference	2.4%
SAP HANA KPI	120MBps (bare metal and virtualized are both over 5 times higher than 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 [best practices configuration guidelines](#). We conducted the tests for the results shown in the tables on an 8-socket Cascade Lake system. 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.

Figure 9. Logical network connections per SAP HANA server [1]



* [1] This figure was sourced from SAP SG.

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 the ESXi host
- HANA-exposed virtual network cards in the VM

The following sections explain each.

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

vSphere offers standard and distributed switch configurations. You can use either 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 VMs to maintain a consistent network configuration as they migrate across multiple hosts.

VMs have network adapters that can be connected to port groups on the virtual switch. Every port group can use one or more physical NICs to handle their network traffic. Not having a physical NIC on a port group limits VM communication between the VMs and not the external network. You can find more information about this in the [vSphere networking guide](#).

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

Note: This reference guide doesn't include the adoption of NSX networking (VLAN or overlay segments) or NSX Distributed Firewall (DFW). These NSX deployment options require additional consideration and planning and are currently out of scope for this document.

Table 16. Recommended SAP HANA on vSphere network configuration based on 10GbE NICs (non-VCF setup)

		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)	Networklabel	Physical host management	vSphere admin	SAP app server	vMotion	Backup network (optional)	HANA replication (optional)	Scale-out (optional)	Storage network (optional)
	Typical bandwidth	1GbE	1GbE	>= 10GbE	>= 10GbE	>= 10GbE	>= 10GbE	>= 10GbE	>= 10GbE
	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 [1]	#0, 1GbE	#1, 10GbE		#2, 10GbE	#3, 10GbE	#4, 10GbE	#5, 10GbE	#6, 10GbE
	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)		

* [1] 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 25GbE 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 the next table.

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

The following table shows an example of how to group the network port failover teams. It also shows the distributed switch port groups created for different functions and the respective active and standby uplinks to balance traffic across the available ones.

The suggestions in this table will differ depending on the required and optional networks needed for vSphere, SAP HANA system replication, or scale-out internode networks. At a minimum, 3 NICs are required for a virtual HANA system leveraging vMotion and vSphere HA. Additional network adapters are required for the optional use cases.

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 10GbE NICs and group failover pairs, depending on their needed network bandwidth.

Important: Do not group 2 high-bandwidth networks, such as the internode network with the app server network.

Table 17. Minimum ESXi server uplink failover network configuration based on 10GbE NICs

Property	NIC	VLAN [1]	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)	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

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

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.

The next table shows an example with 25GbE network adapters.

Table 18. Example SAP HANA on vSphere network configuration based on 25GbE 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	1GbE	1GbE	>= 24GbE	>= 15GbE	>= 10GbE	>= 15GbE	>= 10GbE	>= 25GbE
	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 [1]	#0, 1GbE	#1, 25GbE		#2, 25GbE		#3, 25GbE		#4, 25GbE
	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)		

* [1] 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 a 4-socket large VM with a minimum of 25 GbE NICs.

Using 25GbE 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.

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

Table 19. Minimum ESXi server uplink failover network configuration based on 25GbE NICs (VCF setup and non-VCF setup)

Property	NIC	VLAN [1]	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

* [1] The network administrator usually determines the final numbers for the VLAN ID example.

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.

The next table shows an example with 100GbE network adapters.

Table 20. Example SAP HANA on vSphere Network Configuration based on 100GbE NICs (VCF setup and non-VCF setup)

		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	1GbE	1GbE	>= 24GbE	>= 15GbE	>= 10GbE	>= 15GbE	>= 10GbE	>= 25GbE
	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, 1GbE	#1, 100GbE			#2, 100GbE			
	Active physical NICport	1	1			1			

		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
	Standby physical NICport	2	2			2			
VM guest network cards	Virtual NIC (inside VM)		

Note: The network example configurations in the tables above only leverage vSphere networking. For these examples, the configured host tunnel endpoint (TEP) for NSX overlay networking via VCF is irrelevant and not shown in the tables.

The following table shows an example of how to group the network port failover teams based on 100GbE NICs.

Table 21. Minimum ESXi Server Uplink Failover Network Configuration based on 100GbE NICs

Property	NIC	VLAN [1]	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

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

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 60µs (no load) and up to 300µs of latency (high CPU load >= 80%) 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.

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 in OLTP transactional request times, which showed an overhead of up to 100ms when compared to bare metal SAP HANA systems.

We saw this rise when we used the VMXNET3 virtual NIC. This is because virtual networking adds latency to every network package sent or received, ranging from 60µs (no load) to up to 300µs (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). This is documented in Broadcom KB [83957](#).

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 applications and SAP HANA database servers. Therefore, it is difficult 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 discusses how we measured the VMXNET3 impact and how to optimize the network configuration for an SAP HANA VM for its given use case.

Note: This reference guide doesn't include the adoption of NSX networking (VLAN or overlay segments) or NSX Distributed Firewall (DFW). These NSX deployment options require additional consideration and planning and are currently out of scope for this document.

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

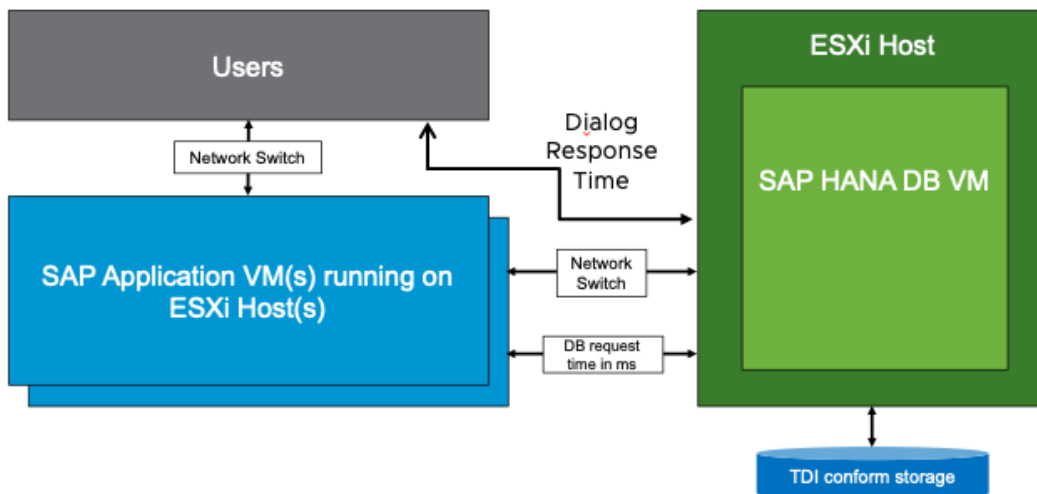
We ran performance tests to measure the impact of virtualization on SAP HANA, and ultimately on the users, to define the best possible configuration for mitigating the virtualization costs, such as increased network latencies.

Testbed configuration

A typical SAP S/4HANA environment is a three-tier landscape, with the application tier on one or more hosts separate from the database tier. Users access 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 3-tier environment (shown below) to simulate real customer systems.

Figure 10. SAP 3-tier architecture on vSphere



The table below shows our testbed configuration.

Table 22. Testbed configuration

Software	<ul style="list-style-type: none"> • SAP S/4HANA • SAP BW/4HANA • VMware vSphere
Application workloads	<ul style="list-style-type: none"> • OLTP • OLAP • OLTP/OLAP mix
Benchmark software	<ul style="list-style-type: none"> • 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 • SAP BW/4HANA – BWH benchmark
ESXi host and bare metal server hardware	<ul style="list-style-type: none"> • 4-socket Cooper Lake system • 8-socket Cooper Lake system
VM configuration	<ul style="list-style-type: none"> • Single VM. Actual size as big as the host minus memory overhead, like an 8-socket, 11.5TB large VM with up to 448 vCPUs and a VMXNET3 NIC for the SAP test-initiated network traffic.

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 hundred 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 taken 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 was reached. The test-relevant results are of users at the 35% and 65% CPU utilization points.

The result of the maxed-out measurement defines the maximum throughput number for OLTP/OLAP and the maximum possible number of SAPS for 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, allowing us to make recommendations to reduce the virtualization impact on networking performance.

Running the BWH benchmark didn't reveal any network throughput or latency issues, but the SAP S/4HANA custom mixed workload test revealed a network latency issue with the VMXNET3 driver, as documented in Broadcom KB [83957](#).

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

The following table 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%. The bare metal SAP HANA system had, on average, a netperf TCP roundtrip time of around 26 μ s (no load) and up to 95 μ s (high load). On average, the virtualized SAP HANA system showed 84 μ s and 319 μ s TCP roundtrip time while under load.

Table 23. Network latency of Netperf TCP VMXNET3 compared to bare metal, measured in μ s

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


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 database request time performance. 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 on a database request time deviation of between 27ms and approximately 82ms. See the next figure 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, the ESXi kernel was better able to 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 reserving CPU threads for the networking stack.

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



Using VMXNET3 resulted in a 25% (27ms) increase in the database request time at 35% CPU utilization but had no effect on the OLTP throughput per hour (TPH) or OLAP queries per hour (QPH). At approximately 65% CPU utilization, the database request time increased to 36% (82 ms higher) with a TPH deviation of approximately -1%. At the maximum user capacity, with 91,000 users and approximately 80% CPU utilization, the impact of TPH and QPH was approximately -8%. There was minimal impact on the OLAP request times. The following 3 figures 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 THP 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 the following figures.

Note: The measured database request time is the time between the SAP application server and the HANA database instance. We didn't 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.


The figure below shows that the VMXNET3 or passthrough network cards did not significantly impact the OLAP request times 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 have little impact on OLAP database request time (next figure).

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 following figure displays the mixed workload OLAP QPH results of a Cooper Lake system. 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



In summary, for OLAP workloads, a passthrough NIC doesn't make complex query runtime any faster. Reserving some CPU cores for the ESXi kernel by lowering the number of vCPUs in a VM could potentially improve performance during data loading, but it will result in lower QPH for SAP HANA.

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 with the user, then move on to the SAP app server tier, and then from the app server to the SAP HANA database. Low-latency switches, a flat network architecture, or newer NICs will help to 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. Additionally, using different CPU generations may yield different results. Sapphire Rapids systems show lower database request times (latencies) than older CPU generations.

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 and shorter packages. VMXNET3 latency overhead is less noticeable for OLAP workloads, which generate long-running queries but are less frequent.

The next table shows the certified results of published BWH tests with Intel Cascade Lake server systems. The results show a BWH configuration based on the L-class sizing that SAP specified for servers with 8 sockets and 6TB of memory.

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

The table below shows the test results. We ran the test on an Intel 8280L Cascade Lake server with 6TB of memory and 28-core CPUs running at 2.7GHz. The test consisted of 9 data sets (11.7 billion records).

Table 24. BWH 6TB L-class workload test: bare metal compared to VM with VMXNET3

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

As you can see in the above table, the virtualization overhead while using a VMXNET3 NIC was very low 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 be lowered even further by leveraging a passthrough (instead of a VMXNET3) NIC.

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


The benchmark that HPE ran on SAP HANA 2.0 SP6 used 6TB and 12TB large database instances to store 10,400,000,000 (BWH L-Class sizing or 5,000 QPH) and 20,800,000,000 (BWH M-Class sizing or 2,500 QPH) initial records. This is the best result ever recorded in a vSphere 7 virtualized environment, with 7,600 (cert. 2022014) and 4,676 (cert. 2022015) queries per hour (QPH).

The following table and figure show 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 12TB benchmark, which ran on the same hardware configuration with SAP HANA 2.0.

Table 25. BWH 12TB 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	2021058	448	12.28GB	20.8 billion	14,986	-	5,161	-	137	-
HPE Superdome Flex 280, Intel Xeon Platinum 8380HL, 28-core CPU - vSphere 7 VMXNET3 VM	2022015	448	11.776GB		15,275	-1.93%	4,676	-9.40%	149	8.76%
					- is better		+ is better		- is better	

Figure 15. BWH 12TB 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 6TB and an 8TB 4-socket virtualized SAP HANA Sapphire Rapids system. With the 2-socket configuration, virtualized is very close to bare metal. See the next tables and figures for details. The overall virtualization overhead for the 2-socket configuration was below 3% of that for a bare metal configuration. For your reference, the 2-socket SPR 4TB BWH benchmark was published under cert number [2023030](#).

Table 26. 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		16,158	—	5,036	—	109	—

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

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




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

Table 27. BWH 8TB 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	8,192GB	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	8,000GB		13,506	5.7%	5,880	-0.4%	156	3.3%
					- is better		+ is better		- is better	

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



Virtual SAP HANA performance evolution over 5 CPU generations

In summary, for our validation test results, we compared the different CPU generations to the same vSphere version using the SAP internal mixed workload tests and the public BHW benchmarks, which demonstrate significant performance gains.

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

Table 28. 4-Socket vSphere 8 VM vs older CPU generations running mixed SAP HANA workload test VM @ max-out CPU MP

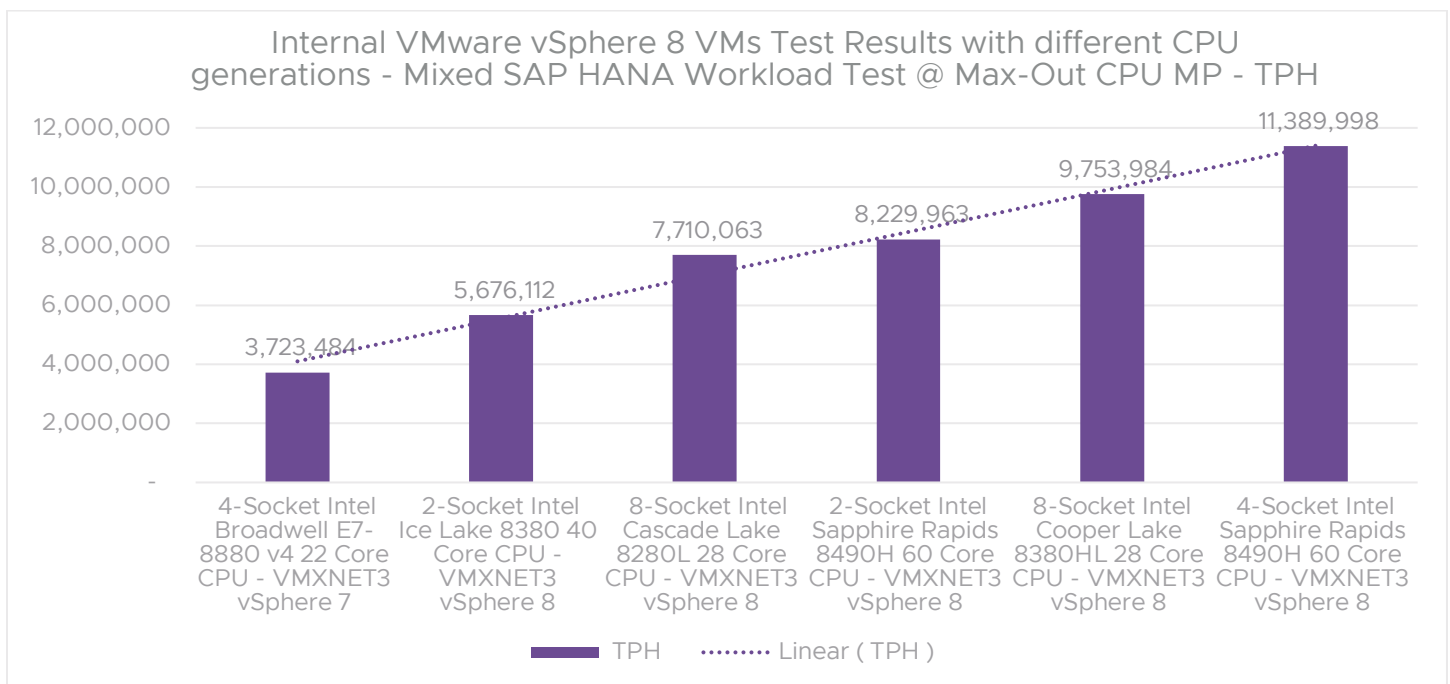
Configuration	Users @ max. CPU	Used CPU Threads	TPH	QPH
4-Socket Intel Broadwell E7-8880 v4 22 Core CPU - VMXNET3 vSphere 7 VM	35,000	168	3,723,484	6,210
2-Socket Intel Ice Lake 8380 40 Core CPU VMXNET3 vSphere 8 VM	50,000	160	5,676,112	9,487
8-Socket Intel Cascade Lake 8280L, 28 Core CPU - VMXNET3 vSphere 8 VM	85,000	448	7,710,063	12,855

Configuration	Users @ max. CPU	Used CPU Threads	TPH	QPH
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 small and for this comparison, they aren't relevant.

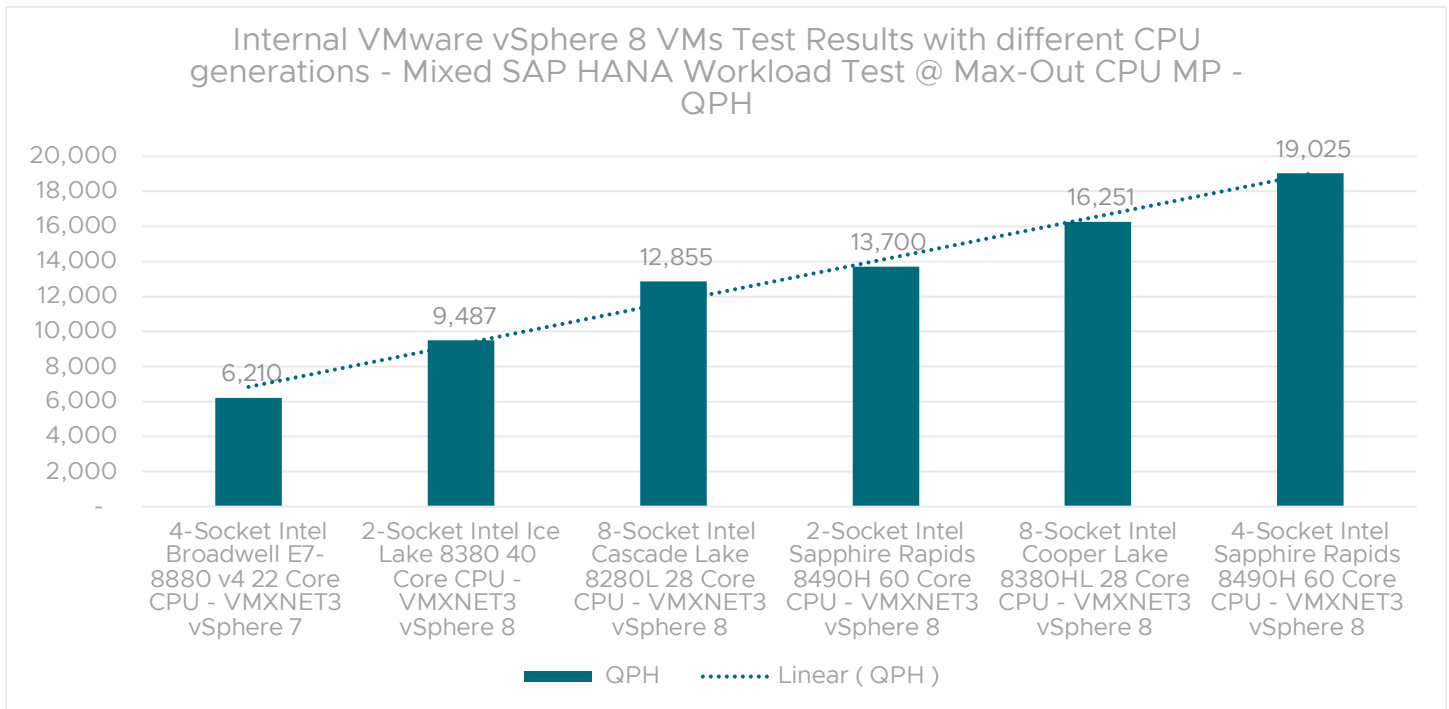
The following figure shows the throughput per hour (TPH) performance in this specific test for SAP HANA over the different CPU generations in graphical form. If you compare a 4-socket Intel Broadwell-based 4-socket server with an Intel Sapphire Rapids 4-socket system, then you see an increase in the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



The next figure 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 provide the same overview but use the publicly available SAP HANA BWH benchmark. Apart from the storage used, the benchmark environment remained identical, making the most of the test results. Phase 1 of the BHW benchmark is 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) system. Both systems run the same HANA, OS, and BHW versions with 6TB memory and 8 datasets configured. While the Sapphire Rapids system has more than 7.14% CPU threads, it shows a performance improvement of 35%. See the following figure and table for details.

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




Table 29. BWH 6TB 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	5,890GB	10.4 billion	21,423	-	5,257	-	172	-
4-Socket Intel Sapphire Rapids 8490H, 60 core CPU - VMXNET3 vSphere 8 VM	-	480	8,000GB		12,848	NC	7,111	35.27%	144	NC

SAP BWH L-Class KPIs	25,000	5,000	200
	- is better	+ is better	- is better

Using the latest Sapphire Rapids hosts with vSphere 8.0 U1 or later results in significant performance gains compared to 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. The performance boost from a Sapphire Rapids CPU mostly offsets the virtualization costs. So, you can migrate a bare metal SAP HANA installation to a virtualized Sapphire Rapids host, which saves you money because it requires less data center space.

When transitioning from an 8-socket Cooper Lake or Cascade Lake system, which can accommodate up to 12TB 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 8TB.


However, with the assistance of the Technology Development Interface (TDI), a 4-socket Sapphire Rapids host can potentially handle higher memory capacities. SAP allows up to 3TB per CPU, indicating that a 4-socket Sapphire Rapids configuration could support up to 12TB 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

A benefit 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 (shown in the following figure) 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. A high-speed network can transfer the VM's active memory and precise running state, allowing the VM to transition from running on the source ESXi host to the destination ESXi 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 ESXi hosts with different CPU generations. 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 ESXi hosts.

Recommended: To ensure high performance on a vSphere cluster, we recommend migrating SAP HANA between hosts with the same CPU type, model, and frequency. 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 ESXi 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

Using vMotion, it's possible to migrate a VM between different hardware generations of a CPU type or storage subsystems.

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 in the SAP HANA application.
- Use vMotion only during non-peak times (for example, when CPU utilization is 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 ESXi host during a vMotion migration of SAP HANA VMs. A noisy neighbor is another VM that is using up the host's resources, leaving few 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 fantastic 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

You can run a vMotion migration manually, fully, or semi-automated with vSphere DRS. To mitigate any potential impact on SAP HANA VMs during a VM migration, we recommend the following:

- 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. The network should have enough bandwidth to support a fast migration time, which depends on the active SAP HANA memory; for example, ≥ 4 GB SAP HANA VMs. A vMotion network with 25GbE 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 following figures depict all the supported scenarios. We discuss what you should do, per scenario, 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 ESXi host to a new host with a new ESXi version




Description:

- VM migration scenario to transfer VMs off a ESXi 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 because it involves different possible CPU clock speeds, and the exposure of the timestamp counter (TSC) to the VM may lead to 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: By default, the TSC inside a VM runs at the same rate as the host when it powers on. If the VM is then moved to a different host without being powered off (for example, with 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.

To resolve this issue, schedule a maintenance window to restart the VMs on non-identical hardware. This will enable the use of hardware-based TSC, instead of software rate transformation on the target host, which can be costly and lead to VM performance degradation.

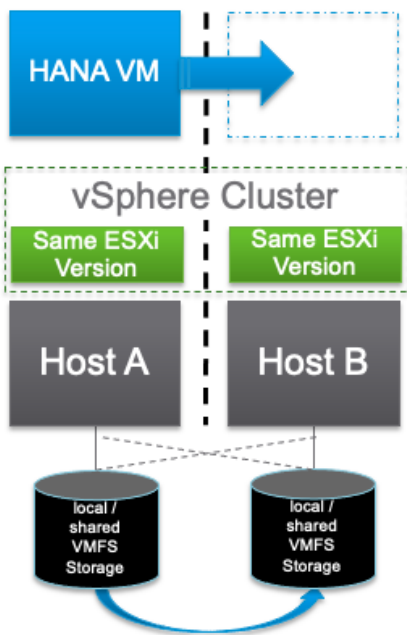
For more information, see [Synchronize the Time of a Virtual Machine Guest Operating System with the Host](#).

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 use with business priorities, 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 less utilized hosts when they need 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, you must set DRS to fully automated. When doing so, we recommend increasing the DRS host memory buffer from its default maximum of 5% of host memory to a maximum of 300GB. To do so, modify the DRS advanced option `MaxMemMBHeadroomPerHost` from the default of 128MB per host to a max of 300GB. Refer to Broadcom [KB 86013](#) for details.

SNC-2 management and operation

Building homogenous SAP HANA on vSphere clusters helps to streamline operations 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, host memory size, installed adapters like HBAs 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 are not migrated or started on non-SNC-enabled hosts. The next figure 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:

[v]

[v]

Host Group:

[v]

[CANCEL] [OK]

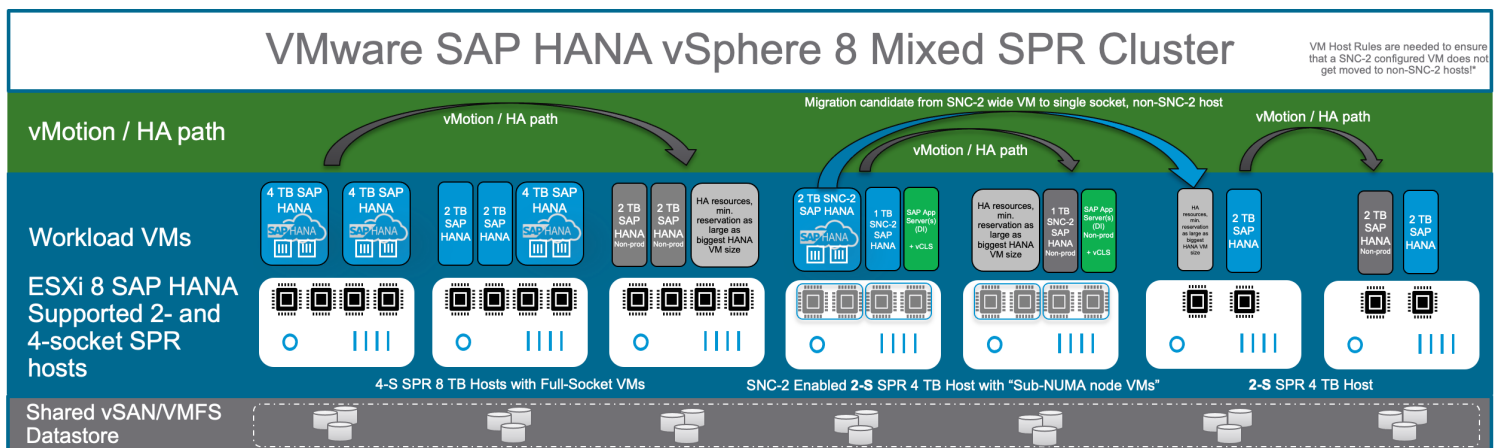
The following figure shows an example cluster with Sapphire Rapids-based 2- and 4-socket systems.

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 2TB 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 2TB per CPU can be utilized.

Attention: VM host rules are needed to ensure that half-socket (SNC-sized) VMs aren't 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 1TB, on the two SNC-2 hosts to benefit from optimized memory latencies. If an 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, 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 2-socket Sapphire Rapids ESXi host:

Note: SNC-2 is only required for 2-socket Sapphire Rapids hosts. >2-socket Sapphire Rapids hosts can deploy 2 SAP HANA VMs on a single NUMA node without SNC-2.

- 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.

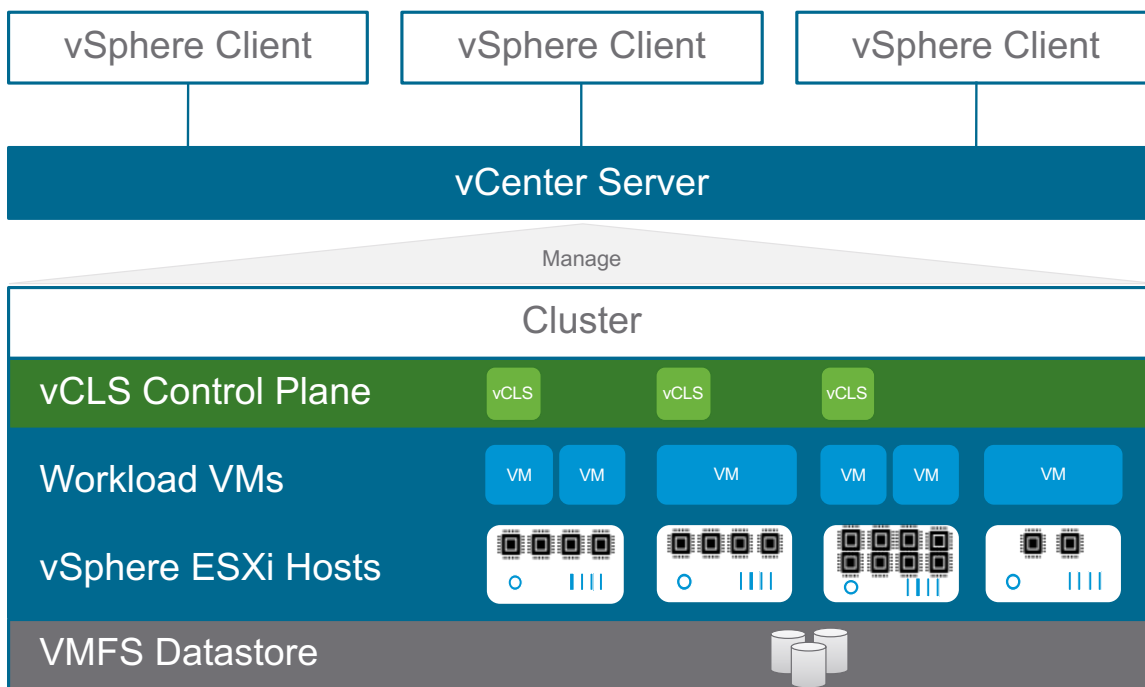
Note: vCLS is enabled when you upgrade to vSphere 7.0 Update 1 or when you have a new vSphere 7.0 Update 1 deployment. vCLS is automatically deployed as part of a vCenter Server upgrade, regardless of which ESXi version is used and is required when using VCF.

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. The next figure 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 100MHz 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 30. vCLS VM resource allocation

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

Table 31. 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/VCF VI WLD (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/VCF VI WLD.

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/VCF VI WLD

A mixed cluster/VCF VI WLD 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/VCF VI WLD 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

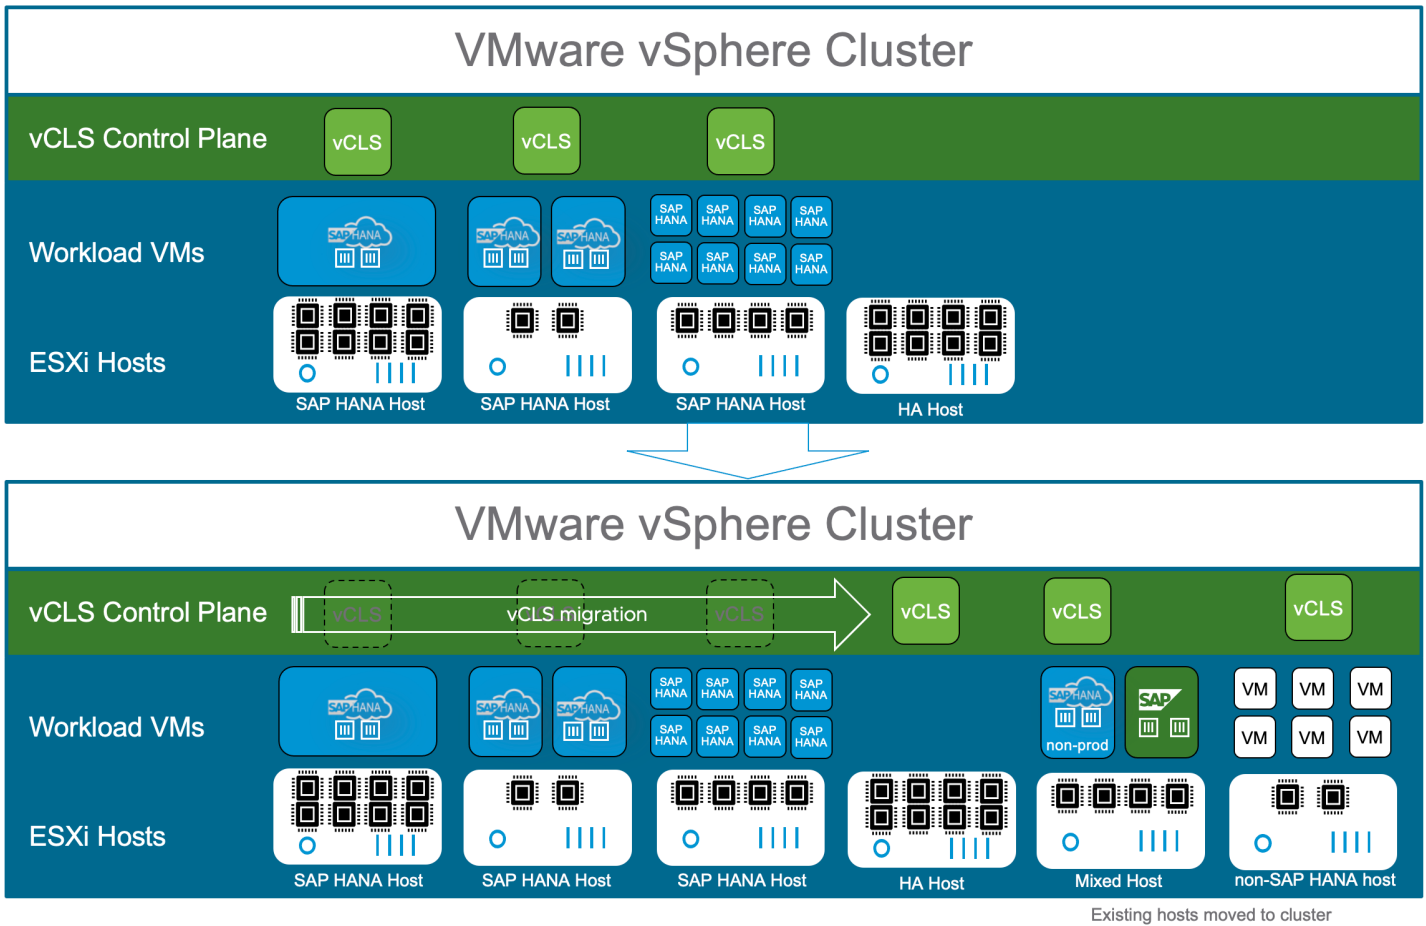


Dedicated ESXi hosts for SAP HANA VMs on a vSphere cluster

You might have deployed a vSphere cluster with ESXi hosts that run only SAP HANA VMs. In this scenario, the automatically deployed vCLS VMs reside on an SAP HANA host, making their migration to hosts without SAP HANA VMs challenging. 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.

The next figure 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 you exclusively use these nodes for SAP HANA production-level VMs, you must remove the vCLS VMs and migrate them 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, the tag name "SAP HANA") and vCLS system VMs.

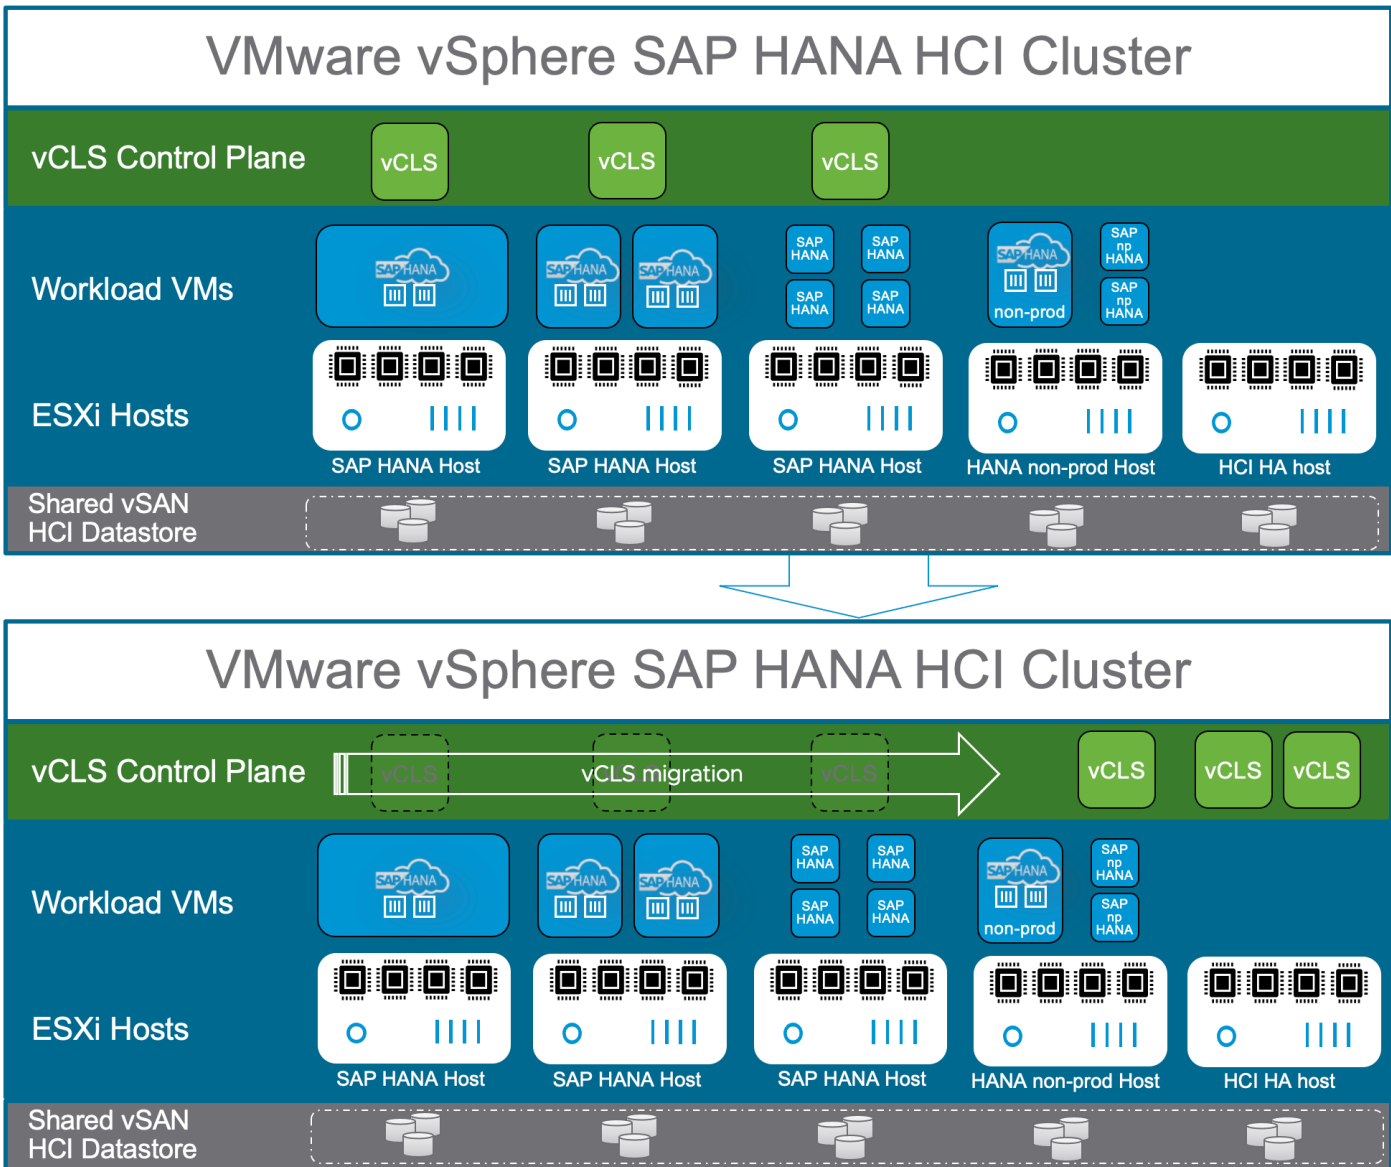
If this tag is assigned to SAP HANA VMs, the vCLS VM anti-affinity policy discourages the placement of vCLS VMs and SAP HANA VMs on the same host. Such a policy prevents the co-deployment of vCLS VMs and SAP HANA

VMs. After you create the policy and assign tags, the placement engine attempts to place vCLS VMs on hosts where tagged VMs are not running, such as the HCI vSphere HA host.

If you are unable to add an additional non-SAP HANA ESXi host to the cluster or need to validate a SAP HANA HCI partner system, you can use Retreat Mode to remove the vCLS VMs from this cluster.

Caution: Retreat Mode will impact certain cluster services, such as DRS, and is not supported when the host is part of a VCF VI WLD.

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, Broadcom is embarking on a journey to remove the vCenter dependency and possible related issues when vCenter Server is not available and provide a scalable platform for larger ESXi host deployments. When the ESXi host is part of a VCF VI WLD/cluster, the vCLS VMs cannot be set to retreat mode because the VCF lifecycle management depends on DRS to automatically put the host in maintenance mode.

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: You can use SNC-2 if you use 2-socket Sapphire Rapids hosts for vSphere or HCI (vSphere + vSAN) deployments. In this case, you could move a vCLS VM to an SNC-2 sub-NUMA node instead of adding 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

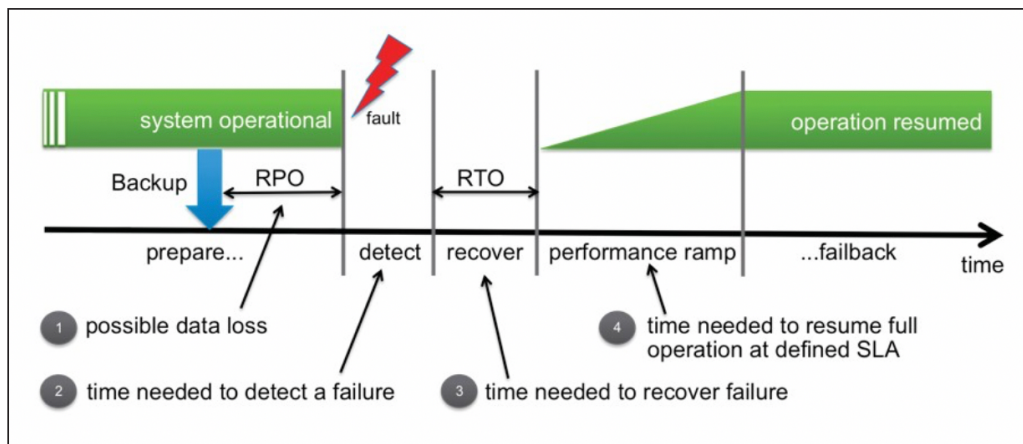
Disaster recovery includes:

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

We include SAP HANA system replication (HSR) in both recovery scenarios. When site or data recovery is required, you can use HSR as a failover or disaster recovery solution. HSR reduces startup time for large SAP HANA databases by preloading data into the replication instance's memory.

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. The following figure shows 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 during 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 start the recovery steps. This is usually done in seconds. 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 restarting the SAP HANA processes.
- **Performance ramp (4 – time needed to resume full operation at a 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, you can plan and implement the proper HA/recovery solution to meet the customer-specific RPOs and RTOs.

The goal of the IT team operating SAP HANA should be to minimize RTOs and RPOs with the available IT budget and resources. VMware virtualized SAP HANA systems allow this by highly standardizing and automating the failure detection and recovery process.

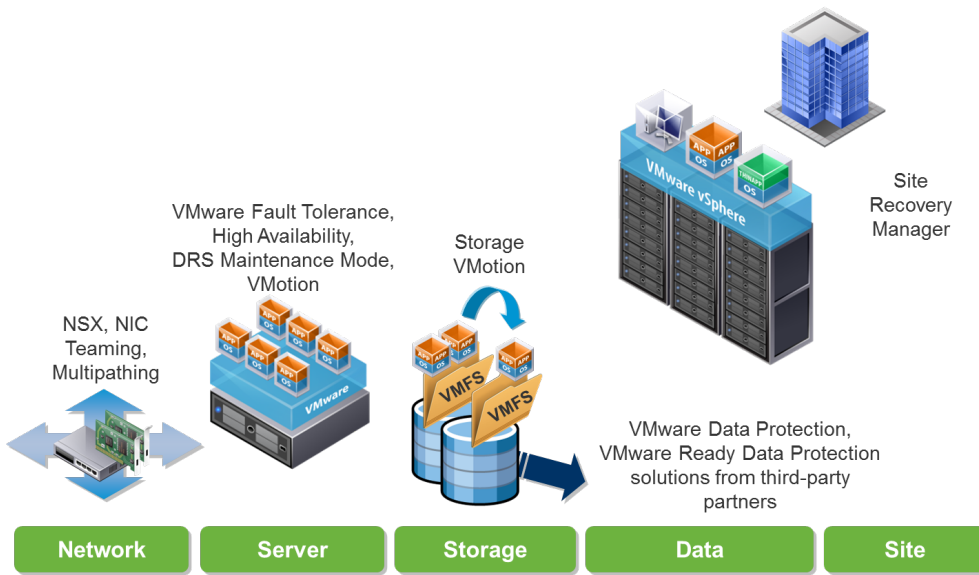
About vSphere high availability (HA)

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. Furthermore, vSphere can provide fault tolerance by incorporating redundant components like dual network and storage pathing, supporting hardware solutions like UPS, and incorporating built-in CPU features that can withstand memory model failures and ensure CPU transaction consistency.

These features are available on the ESXi host without needing to be configured on the VM or application. Organizations can meet their RPOs and RTOs with additional protections like vSphere HA.

VMware Site Recovery Manager can manage and automate various HA solutions, from component-level failures to complete site failures, as depicted in the figure below. 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 that provides uniform, cost-effective failover protection against hardware and operating system outages within a virtualized IT environment. It does this by monitoring ESXi hosts and VMs to detect hardware and guest failures. When it detects a server outage, it automatically restarts VMs on other ESXi hosts in the cluster, thereby reducing application downtime. This, combined with the SAP HANA service auto-restart feature, allows HA levels of 99.9% out of the box (EMC IT, 02/14 EMC Perspective, H12853).

The next figure shows how vSphere HA can protect against VM or host failure. Application protection solutions—such as SAP HANA service auto-restart, third-party in-guest cluster solutions, and SAP HANA system replication—can also provide disaster recovery capabilities. You can combine all these solutions with vSphere HA.

Figure 35. Virtualized SAP HANA HA solution



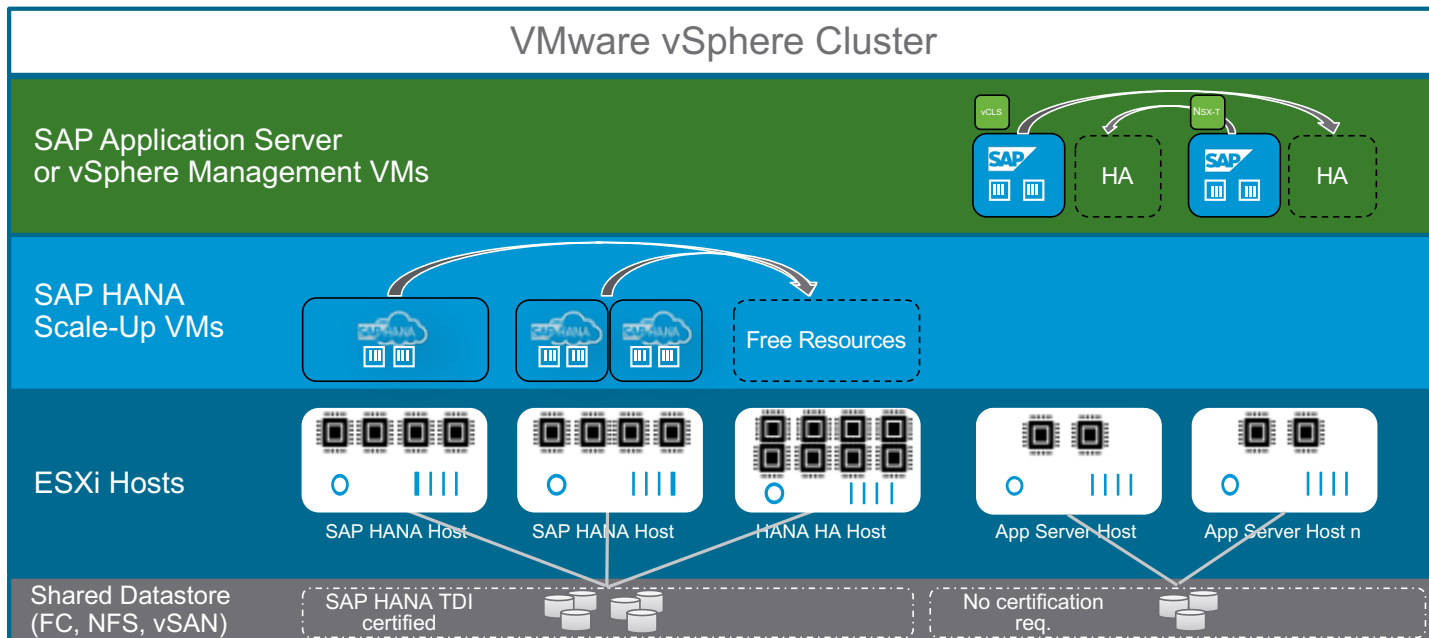
vSphere HA safeguards SAP HANA scale-up and scale-out deployments without requiring any external components, like DNS servers, or solutions like the SAP HANA Storage Connector API.

The following figure shows 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.

SAP applications and SAP HANA instances on vSphere use this vSphere HA configuration. If you need higher redundancy levels, you can use an n+2 configuration. You can also use the HA resource pool for non-critical VMs, but you must power them off before restarting an SAP HANA or SAP app server VM.

Note: The following figures don't show the vCLS control plane or the related vCLS.

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



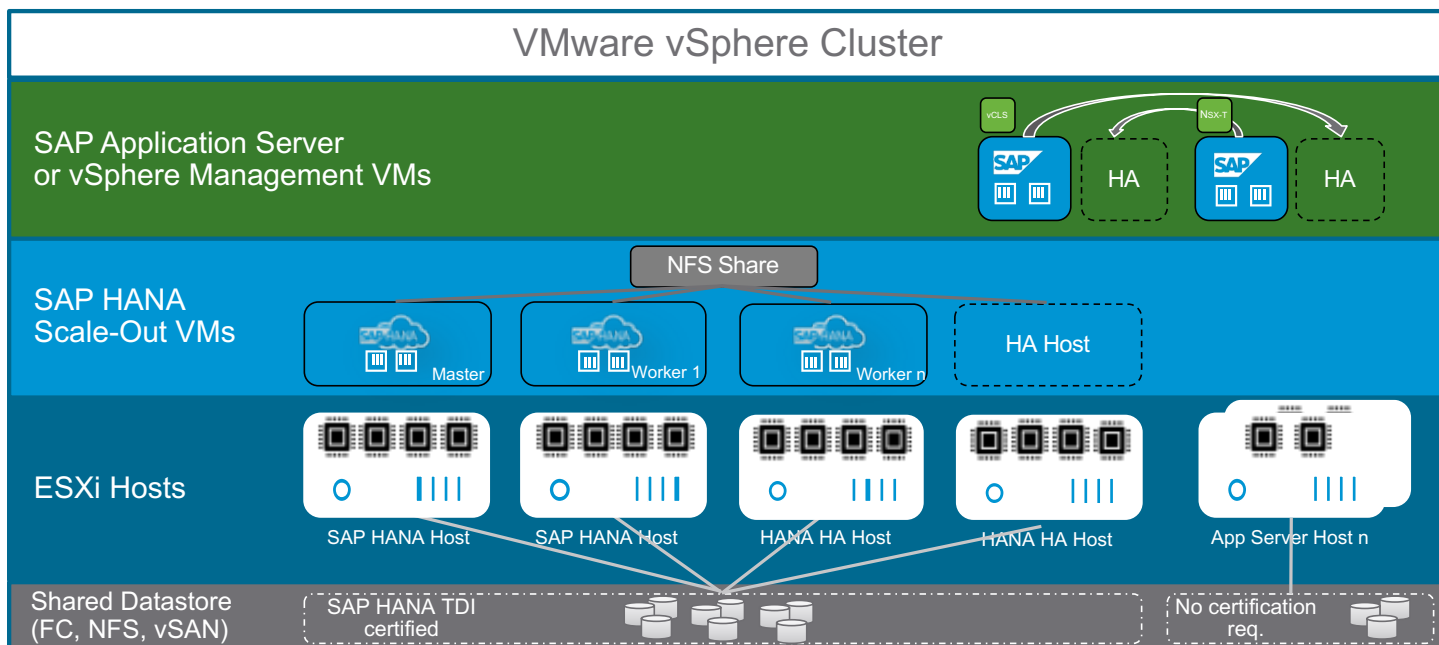
It's also possible to configure the HA cluster as active-active, where all hosts have SAP HANA VMs deployed. This guarantees the utilization of all hosts in a vSphere cluster, ensuring sufficient failover capacity for all running VMs in the event of a host failure. The arrow in the figure indicates that the VMs can failover to different hosts in the cluster. This active-active cluster configuration assumes that one host's capacity (n+1) is always available to support a complete host failure.

As previously mentioned, you can also use vSphere HA to protect an SAP HANA scale-out deployment. Unlike a physical scale-out deployment, SAP HANA on vSphere can withstand a host failure without the need for dedicated standby hosts and storage-specific implementations.

There are no dependencies on external components, such as DNS servers, the 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 VM mounts the HANA shared directory via NFS, just like physical systems do. Access to the HANA shared directory is therefore guaranteed. If the NFS server providing this share is also virtualized, you can use vSphere Fault Tolerance (FT) to protect it.

The following figure shows a configuration of 3 SAP HANA 4-socket 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's possible to use the HA node for other workloads while in normal operation. If you use the HA/standby node for other workloads, you must stop or migrate all VMs running on this host before you can restart a failed HANA scale-out VM on the host. In this case, the overall failover time could be a bit longer because, if configured correctly, vSphere HA will wait until all needed resources are available on the failover host.

You can use up to 16 scale-out nodes with up to 3TB on a 4-socket VM or 6TB on an 8-socket VM. SAP HANA on vSphere supports only 4- and 8-socket host systems with Intel Broadwell or newer CPUs. SAP doesn't support 2-socket host systems for on-premises scale-out deployments. Review the previously mentioned SAP HANA on vSphere support notes for more details.

Note: vSphere HA can only protect against OS or VM crashes or hardware failures. Logical failures or OS file system corruptions are beyond its protection.

In physical SAP HANA deployments, SAP HANA system replication is the only method to provide fault recovery. Implementing a third-party solution like SUSE HA is necessary to automate the recovery. Protecting a physical SAP HANA deployment against host failures is, therefore, relatively complex. In contrast, protecting an SAP HANA on vSphere system can be done with just a mouse click.

If you require fast failure recovery and data replication, 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, you can use it 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.

You must configure vSphere HA with passthrough NICs as a dynamic vSphere DirectPath I/O device, assigning it a unique cluster-wide hardware label. To do this, follow the instructions in [Add a PCI Device to a Virtual Machine](#). You must use the same hardware label for the passthrough NIC installed in the HA host. If you haven't configured an HA host with a passthrough NIC as a dynamic vSphere DirectPath I/O device and the same hardware label exists, then the HA failover process won't work. For more information about this topic, read [Assignable Hardware](#).

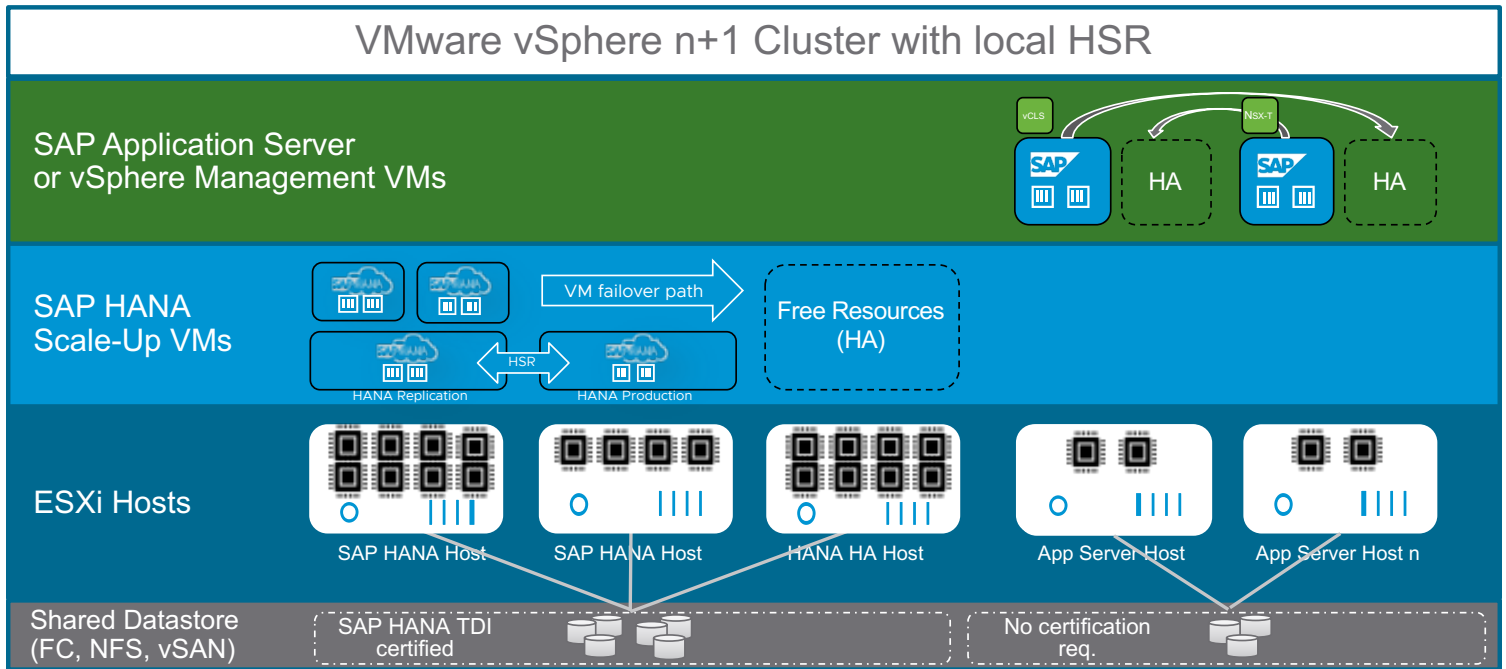
SAP HANA system replication with vSphere (local site)

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

In this case, vSphere HA would 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 of SAP HANA.

The next figure 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. If it runs in the same location, you can use HSR to recover from logical failures (if log collection is delayed) or to reduce the ramp-up time of an SAP HANA system because data can already be loaded into the replication server's memory. HSR can change direction depending on which HANA instance is the production one.

Figure 38. vSphere HA with local data and HANA system replication (HSR)



As noted, HSR doesn't provide automated failover. You must manually reconfigure the replication target VM to match the production system's identity. Alternatively, you can use 3rd-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.

Disaster recovery

You can 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.

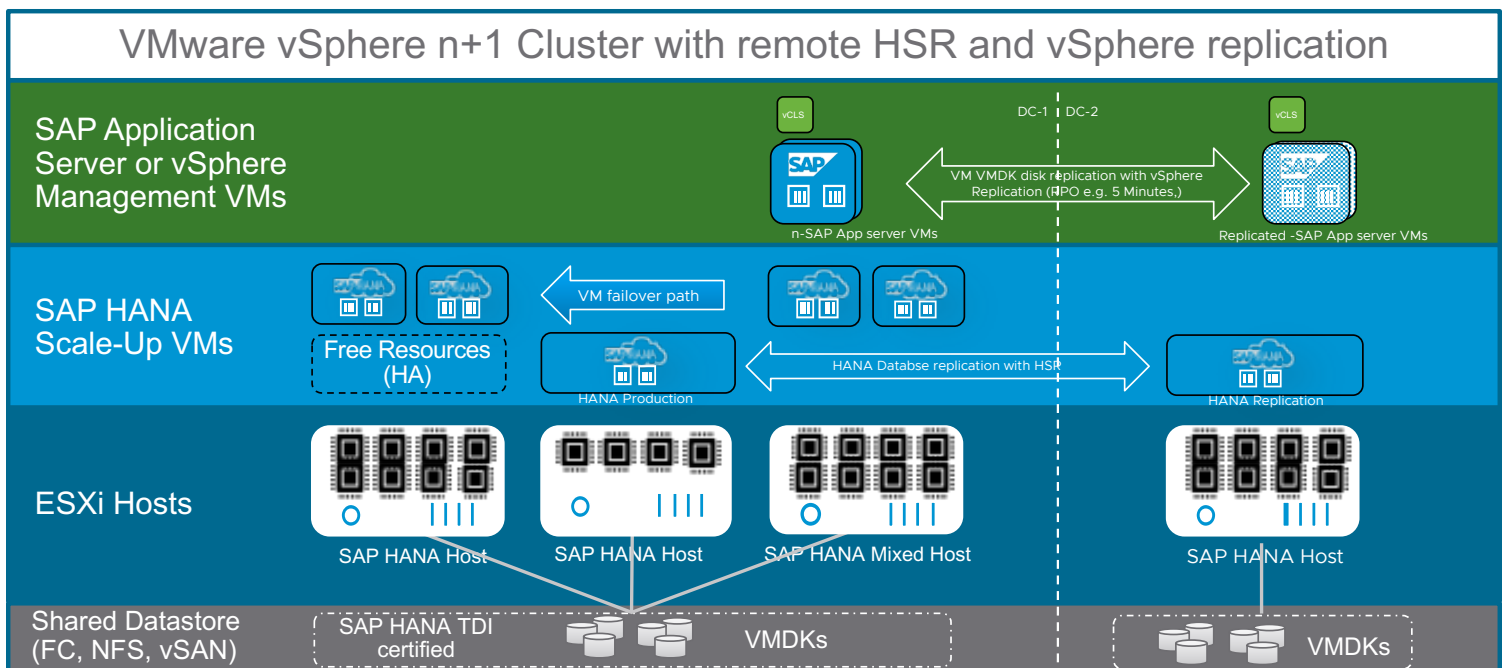
In addition to this, you also need backup and restore solutions in place for regulatory reasons and to protect the data against logical errors.

SAP HSR and vSphere Replication for a remote site

The figure below shows an HSR-protected SAP HANA instance. The concept remains unchanged from the previous figure, except for the placement of the HSR replication target in a different data center. This provides additional protection against data center failures or, if the remote data center is on another site, it protects against site failures. The ESXi host in DC-2 can be either a standalone one 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 and SAP HANA system replication (HSR)



The figure above shows HSR replication from one virtualized SAP HANA system to another. You can leverage [vSphere Replication](#) to replicate the SAP app server VMs. This will allow the operation to continue after switching 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 VMs (VMDK files). It allows recovery point objective (RPO) times ranging from 5 minutes to 24 hours. The VM replication process is nonintrusive and takes place independently of the VM's OS or applications. 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 1ms; otherwise, you should use asynchronous replication to avoid replication-related performance issues with the primary production instance. The HSR target can be a virtualized or natively installed SAP HANA replication target instance.

Caution: Avoid using vSphere replication (because it's asynchronous) if your RPO objectives are less than 5 minutes. We recommend not using vSphere replication for SAP HANA VMs for the best performance.

If they're not using HSR, admins often use vSphere Replication with clusters stretched over two separate data centers to protect their systems against local data center failures. If you're using this method, and your systems should also be protected against data center site disasters, then you need to replicate all relevant systems to a second site. As previously mentioned, you can do so using vSphere Replication, native storage replication, and SAP HANA system replication.

vSphere Replication operates at the individual VMDK level, allowing individual VMs to replicate 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

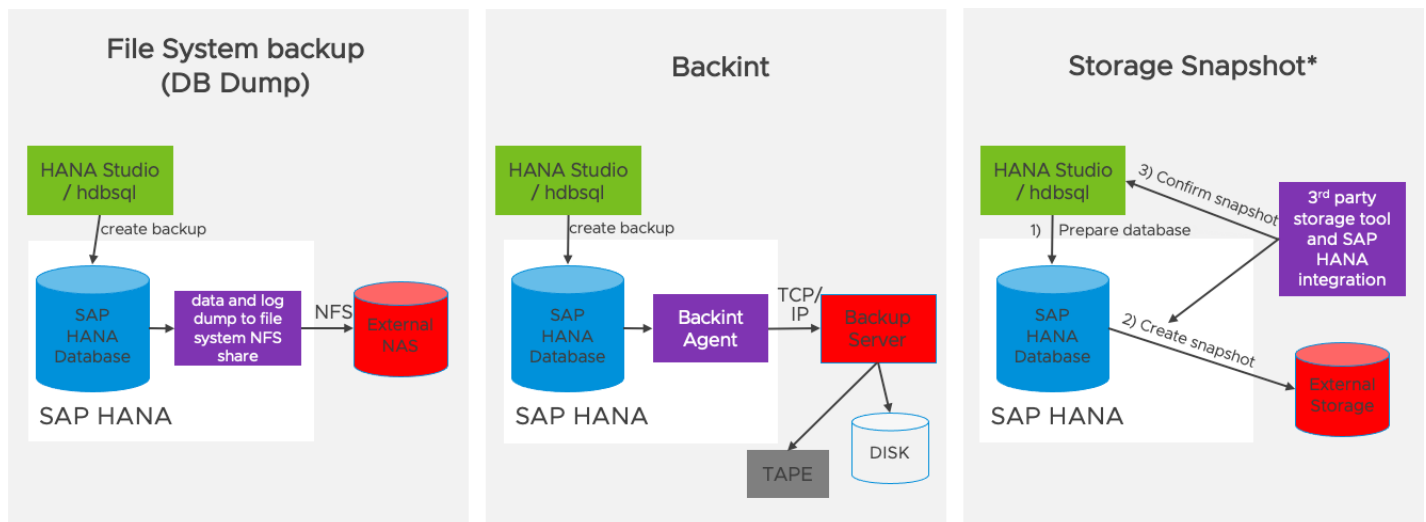
Backing up and restoring an SAP HANA database and the Linux VM supporting SAP HANA is the same as backing up bare-metal deployed SAP HANA systems.

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 the next figure). This provides a neutral storage vendor backup solution based on vSphere snapshots, resulting in shorter backup and restore times. For an example, refer to the Veeam product page: [Confidently back up your valuable SAP HANA data](#) > Resources > Documentation and [SAP HANA](#) in the Veeam Help Center.

Caution: Using vSphere snapshot backup solutions not integrated with SAP HANA can lead to complications. These issues arise because the VM must be frozen to create a snapshot while SAP HANA is running. For this reason, we advise you to 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 Broadcom KB [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). PMem doesn't support 8-socket hosts. The maximum DRAM plus PMem host memory configurations with SAP HANA support 4-socket VMs on 4-socket hosts and can be up to 15TB (the current memory limit when DRAM is combined with PMem) and must follow the hardware vendor's PMem configuration guidelines.

The maximum VM size with vSphere 7.0 is limited to 256 vCPUs and 6TB 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). The 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 PMem can be used for your specific SAP HANA workload.

Supported PMem module sizes are 128GB, 256GB, and 512GB. The following table lists the maximum supported host memory DRAM and PMem configurations. Up to two SAP HANA VMs are supported per CPU socket, and up to 4-socket ESXi hosts can be used. See the table below for the currently supported configurations.

Important: [PMem 100 series technology](#) only supports SAP HANA on vSphere 7.0 with Cascade Lake hosts. Intel has announced that [PMem is discontinued](#).

Table 32. Supported SAP HANA on vSphere with PMem ratios with Cascade Lake and vSphere 7.0

Ratio (DRAM to PMem)		DRAM (TB)	Optane PMem (TB)	Total host memory (TB)	Max. VM memory (TB) [1]	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

* [1] vSphere 7.0 U2 or later versions are required for VM sizes >6 TB.

Sizing of Optane PMem-enabled SAP HANA VMs

PMem-enabled SAP HANA VM sizings are like those for bare-metal SAP HANA systems, but are limited to a maximum size of 6TB (mix of DRAM and PMem) per VM. OLAP class-L workload sizings are limited to 3TB. Class-M sizings support up to 6TB total memory.

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

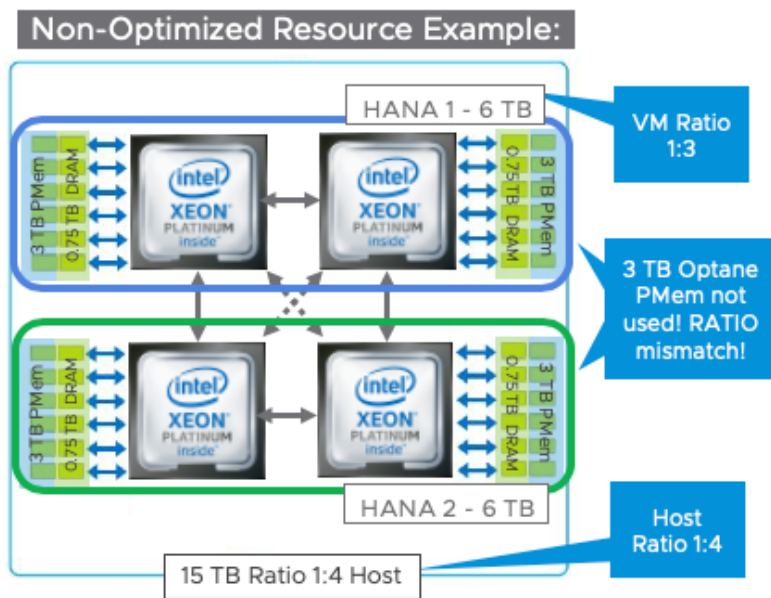
We recommend that an SAP HANA VM use the same DRAM to PMem ratio as the physical host/server DRAM to PMem ratio. If you have a growth plan, you might consider a larger physical memory configuration and upgrade the VMs and SAP HANA over your environment's lifetime.

For example, let's say you have a 1:4 PMem ratio host configured with 15TB of total RAM (3TB DRAM and 12TB PMem). To optimize resources, you'd create 4 SAP HANA VMs on this server, each with 3.75TB RAM (0.75TB DRAM and 3TB PMem). If you created 6TB VMs on the same 15TB host, you could only create 2 SAP HANA VMs.

This results in unoptimized resources because your system can only utilize 12TB of the installed 15TB memory. In this case, a 1:1 DRAM to PMem configuration with a total of 12TB (6TB DRAM and 6TB PMem) would optimize the configuration.

Important: The following examples show optimized and non-optimized memory configurations. Note that we don't currently support >6TB VMs—we've used more in the examples for convenience.

Figure 41. Non-optimized host memory configuration



4-socket host configuration:

Four 2nd Gen Intel Xeon Platinum processors

24 x 128 GB DRAM + 24 x 512GB PMem = 15TB total host memory with a 1:4 DRAM to PMem ratio

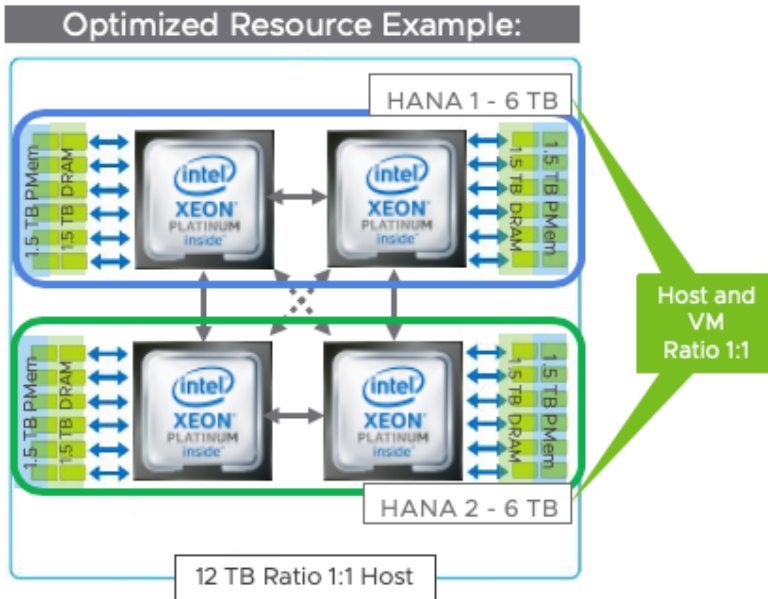
VM configuration example:

2 x 6TB SAP HANA VM with 1.5TB DRAM and 4.5TB PMem RAM, with a 1:3 DRAM to PMem ratio

Challenges:

- DRAM: The PMem ratio may not be suited for the SAP HANA workload
- The hardware configuration does not fit and will lead to unusable PMem (a ratio mismatch)

Figure 42. Optimized host memory configuration



4-socket host configuration:

Four 2nd Gen Intel Xeon Platinum processors

24 x 256 GB DRAM + 24 x 256GB PMem = 12TB total host memory with a 1:1 DRAM to PMem ratio


VM configuration example:

2 x 6 TB SAP HANA VM with 3TB DRAM and 3TB 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 PMem = 9TB total host memory

VM configuration example:

4 x VM with 0.75TB DRAM and 1.5TB PMem RAM = total RAM per SAP HANA VM 2.25TB with a 1:2 DRAM to PMem ratio

Challenges:

- Need to verify if the PMem ratio is applicable to the SAP HANA sizing and if the CPU resources are enough.

Note: Workload-based sizings are supported and allow OLAP workloads with Class-M CPU requirements to leverage up to 6TB 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 equally used on a DRAM-only VM. See the following table for details.

Using SAP HANA on vSphere allows HANA users to leverage the flexibility of vSphere capabilities, such as vMotion, which allows workloads to be migrated between ESXi 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 supported for SAP HANA VMs with 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 33. 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 4socket 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	vSphere doesn't support mixed mode	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	vSphere doesn't support mixed mode	Not supported for SAP HANA

PMem mode	App Direct (vNVDIMM)	Memory mode	Mixed mode	Storage over App Direct
vSphere vMotion	Yes, up to 6TB VM sizes (requires another PMemhost)	Yes	vSphere doesn't support mixed mode	Yes, up to 6 TB VM sizes (requires another PMem host and the usage of vSphere Storage vMotion)
vSphere DRS	Yes, up to 6TB VM sizes (requires another PMemhost)	Yes	vSphere doesn't support mixed mode	Yes, up to 6 TB VM sizes (requires another PMem host and the usage of vSphere Storage vMotion)
VMware Site Recovery Manager	No	Yes	vSphere doesn't support mixed mode	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)	vSphere doesn't support mixed mode	No
VM in guest cluster	No	Yes (can move/failover VMs to/from Memory mode and DRAM systems)	vSphere doesn't support mixed mode	No
VM snapshot	No	Yes	vSphere doesn't support mixed mode	No
SAP HSR	Yes	Yes	vSphere doesn't support mixed mode	No

vSphere HA support for PMem-enabled SAP HANA VMs

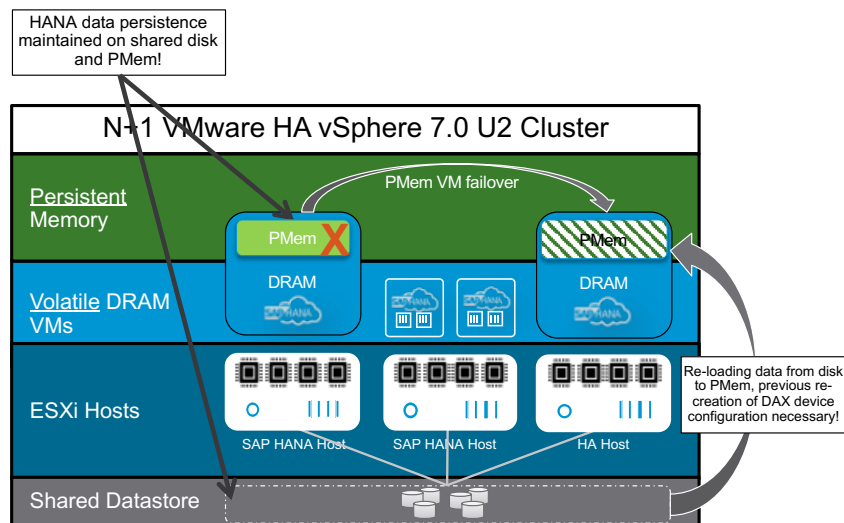
Before the vSphere 7.0 U2 release, vSphere HA did not support PMem-enabled VMs. vSphere HA now supports failing over and restarting PMem-enabled VMs. Applications using PMem must maintain data persistence on PMem and shared disks.

SAP HANA is one of the applications that provides data persistence on disk. Because of this, vSphere HA can use the data on 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 it doesn't have control over post-VM failover OS/application-specific configuration steps, such as the required re-creation of the SAP HANA DAX device configuration. This must be done manually or via a script, that is not provided by VMware or SAP.

The figure below 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, you can start SAP HANA. It will automatically load the data from the 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 lifetime of this VM (as long as it isn't migrated or deleted from the host).

The Intel SAP Solution Engineering team, as well as the Intel and VMware Center of Excellence, have developed an example script for automatically re-creating the DAX device configuration at the OS level. Run this script after the VM has failed over and restarted but **before** the SAP HANA database has been restarted. It's best to automatically run this script as part of the OS start procedure, such as a custom service. You can use the script as a template to create your own script that fits your unique environment.

Caution: Broadcom, SAP, and Intel don't maintain or support this script. Your use of it 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 the server host for PMem in its BIOS (vendor specific)
2. Create AppDirect interleaved regions and verify that they are configured for ESXi use.

VM:

1. Create a VM with hardware 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 PMem.

Details on configuration steps 2 and 3

1. Before you can add NVDIMMs to an SAP HANA VM, check if the PMem regions and namespaces were created correctly in the BIOS.
2. Also, make sure you selected all PMem as **persistent memory** and that the persistent memory type is set to **App Direct Interleaved**, as shown in the following figure.

Figure 45. Example of PMem system BIOS settings



3. After creating the PMem memory regions, you must reboot the system.
4. Install the latest ESXi version (7.0 U2 or later) and, using the web client, check 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

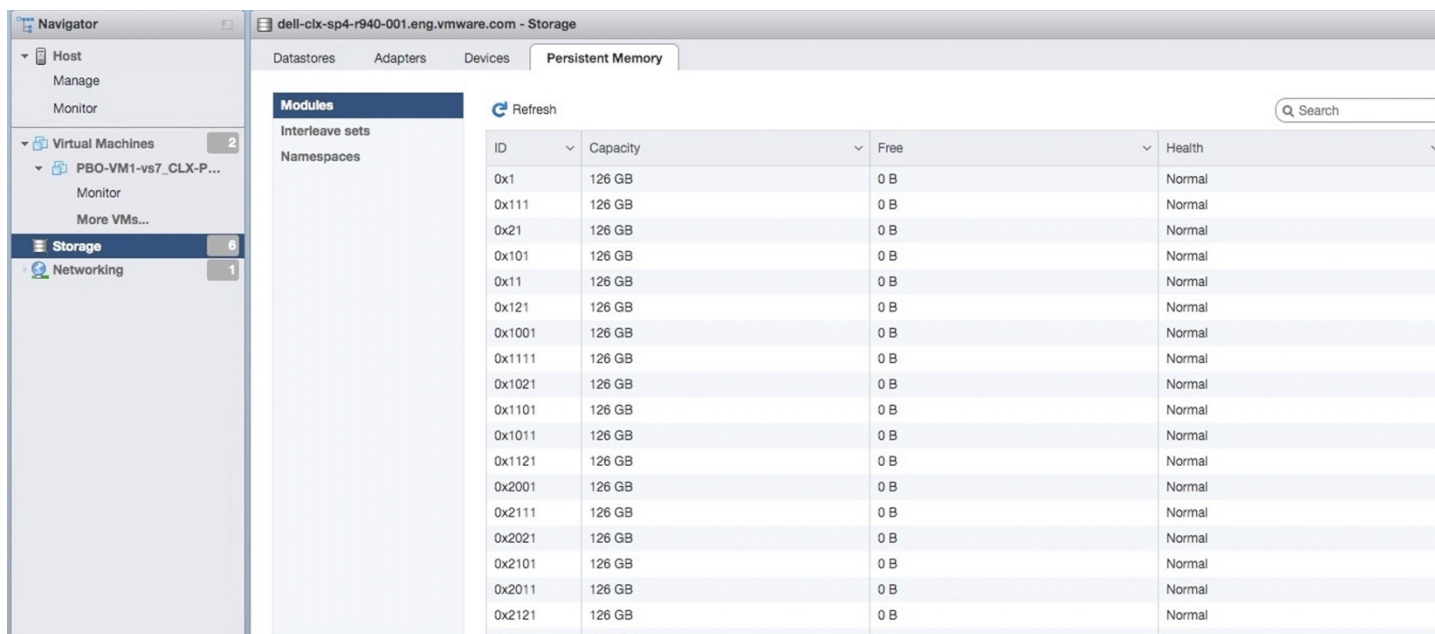


Figure 47. ESXi Persistent Memory Storage View of Interleave sets

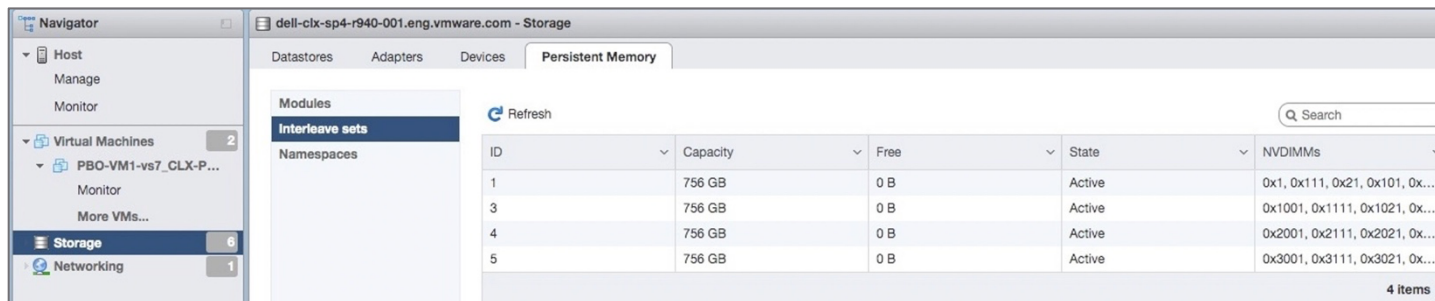
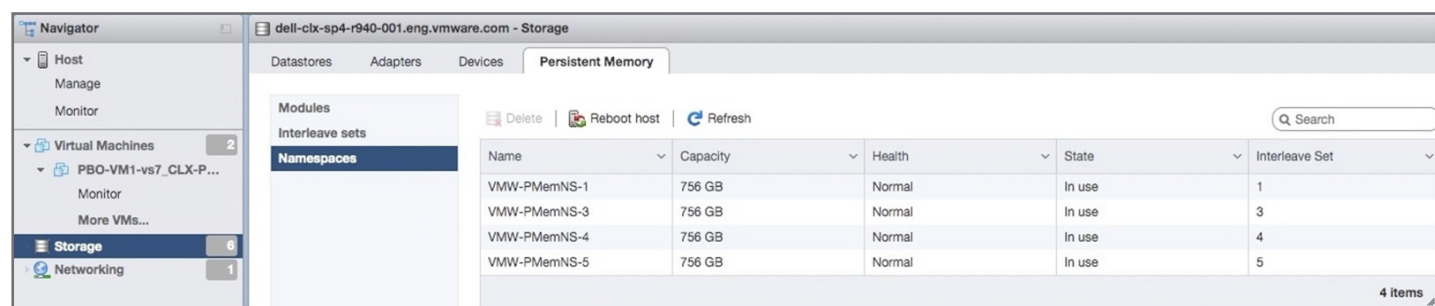


Figure 48. ESXi Persistent Memory Storage view of Namespaces



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

If the configurations is done correctly in the host's BIOS, it should look like what is shown in these figures.

5. After this, you can add NVDIMMs and NVDIMM controllers to your SAP HANA VM. Select the maximum size possible per NVDIMM; otherwise, you'll waste memory capacity.

Figure 49. The vSphere Client *Edit settings* screen, where you add NVDIMMs and NVDIMM controllers to the *SAP HANA VM*

Edit settings - PBO-VM1-vs7_CLX-PMEM-DELL (ESXi 6.7 U2 virtual machine)

▼ NVDIMM	
Size	<input type="text" value="757760"/> MB
Controller location	<input type="text" value="NVDIMM controller 0:0"/> ▼
▼ NVDIMM	
Size	<input type="text" value="757760"/> MB
Controller location	<input type="text" value="NVDIMM controller 0:1"/> ▼
▼ NVDIMM	
Size	<input type="text" value="757760"/> MB
Controller location	<input type="text" value="NVDIMM controller 0:2"/> ▼
▼ NVDIMM	
Size	<input type="text" value="757760"/> MB
Controller location	<input type="text" value="NVDIMM controller 0:3"/> ▼

To configure a PMem-enabled SAP HANA VM for optimal performance, you must align the VM configuration to the underlying hardware, especially the NUMA configuration. Broadcom KB [78094](#) provides information on how to configure the NVDIMMs (VMware's representation of Optane PMem) correctly and align the NVDIMMs to the physical NUMA architecture of the physical server.

By default, the PMem allocation in vmkernel for VM NVDIMMs does not consider NUMA. This can result in the VM running on a certain NUMA node and the PMem being allocated on a different NUMA node. This will cause the NVDIMMs to access the VM remotely instead of locally, resulting in poor performance. To solve this, you must add the following settings to a VM configuration using vCenter.

Example for a 4-socket wide VM:

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

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

Besides these parameters, you can 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 nvdimm *.vmx  
nvdimm0.present = "TRUE"  
nvdimm0:0.present = "TRUE"  
nvdimm0:0.fileName = "/vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_1.vmdk" nvdimm0:0.size = "757760"  
nvdimm0:1.present = "TRUE"  
nvdimm0:1.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_3.vmdk" nvdimm0:1.size = "757760"  
nvdimm0:2.present = "TRUE"  
nvdimm0:2.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_5.vmdk" nvdimm0:2.size = "757760"  
nvdimm0:3.present = "TRUE"  
nvdimm0:3.fileName = /vmfs/volumes/pmem:XXXXX/ PMem_SAP_HANA_VM_name_5.vmdk" nvdimm0:3.size = "757760"  
nvdimm0:0.node = "0"  
nvdimm0:1.node = "1"  
nvdimm0:2.node = "2"  
nvdimm0:3.node = "3"
```

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 changed configurations. Collect SQL statements to check and identify parameters that are configured and conflict with the SAP-recommended configuration parameters.

NUMA tool: VMware NUMA Observer

While you 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 download and use the [NUMA Observer fling](#) (downloads a .zip file).

NUMA Observer 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

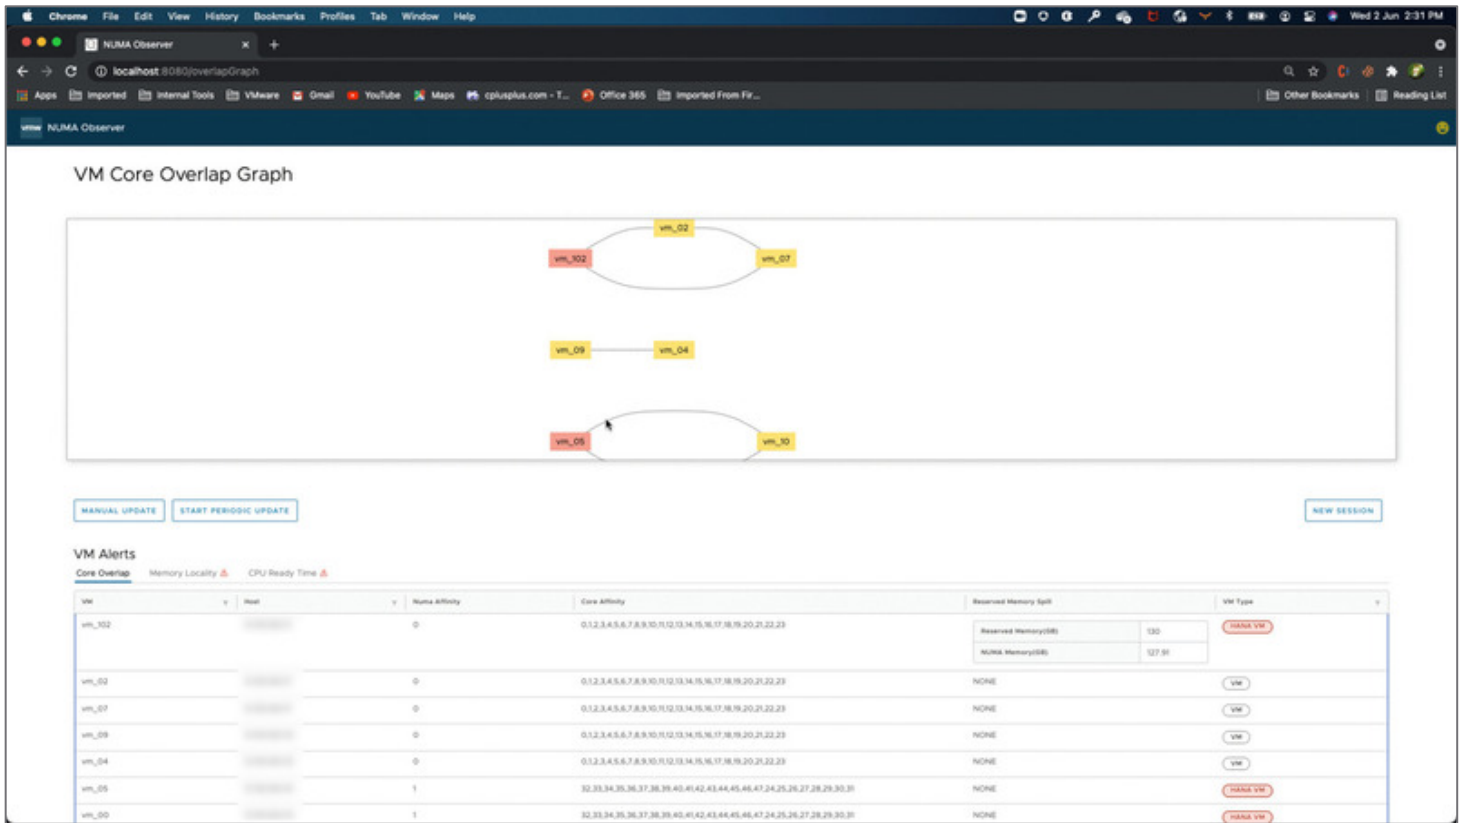
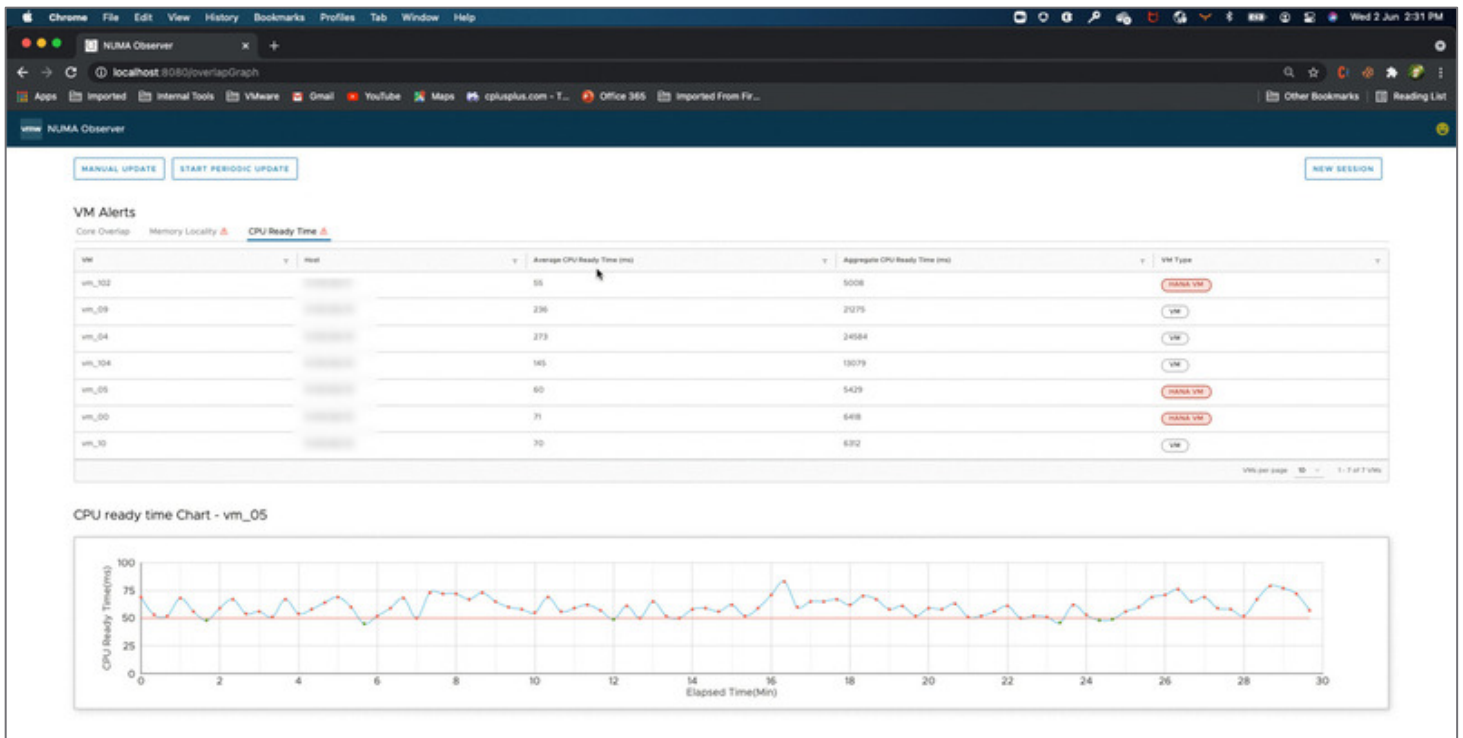


Figure 51. VMware NUMA Observer – VM Alerts



Performance optimizations for SAP HANA VMs

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

Optimizing the configuration parameter list

VMware vSphere can run a single large SAP HANA VM or multiple smaller VMs 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 best practice configurations, and, in case of an escalation, the support engineers will verify and, if not applied, they will recommend configuring these settings.

The following table shows the sub-NUMA node cluster (SNC) settings that are required when a half-socket VM runs on a 2-socket Sapphire Rapids ESXi host.

Table 34. 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 UEFI since version 5.0.
Enable Intel VT technology	Enable all BIOS virtualization technology settings.
Configure RAM hemisphere mode	Distribute DIMM or PMem modules in a way to achieve best performance (hemisphere mode) and use the fastest memory modules available for the selected memory size. Beware of the CPU-specific optimal memory configurations that depend on the available memory channels per CPU.
CPU – Populate all available CPU sockets, use a fully meshed QPI NUMA architecture	To avoid timer synchronization issues, use a multi-socket server that ensures NUMA node timer synchronization. NUMA systems that do not run synchronization will need to synchronize the timers in the hypervisor layer, which can impact performance. See the Timekeeping in VMware Virtual Machines information guide for reference. Select only SAP HANA CPUs supported by vSphere. Verify the support status with the SAP HANA on vSphere support notes. For a list of the relevant note, see the SAP Notes Related to VMware page.

Physical host BIOS parameter settings	Description
Enable CPU Intel Turbo Boost	Allow Intel automatic CPU core overclocking technology (P-states).
Disable QPI power management	Do not allow static high power for QPI links.
Set HWPE	Set to the hardware vendor default.
Enable 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 use with ESXi.</p> <p>Note: Only App Direct and Memory mode are supported with production-level VMs.</p> <p>Memory mode is only supported with a ratio of 1:4. As of today, SAP provides only non-production workload support.</p> <p>VMware vSAN does not have support for App Direct mode as cache or as a capacity tier device of vSAN. However, vSAN will work with ESXi hosts equipped with Intel Optane PMem in App Direct mode and SAP HANA VMs can leverage PMem according to SAP note 2913410.</p> <p>Important: The vSphere HA restriction (as described in SAP note 2913410) applies and must 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 35. ESXi host parameter settings

ESXi host parameter settings	Description
Networking	<p>Use virtual distributed switches to connect all hosts that work together. Define the port groups that are dedicated to SAP HANA, management, and vMotion traffic. Use at least a dedicated 10 GbE for vMotion and the SAP app server or replication networks. At least 25 GbE for vMotion for SAP HANA system; ≥ 2 TB is recommended.</p>
Settings to lower the virtual VMXNET3 network latencies	<p>Set the following settings on the ESXi host. For this to take effect, the ESXi host must be rebooted.</p> <p>Go to the ESXi console and set the following parameter (enter it on one line):</p> <pre>esxcli system settings advanced set -o /Net/NetNetqRxQueueFeatPairEnable -i 0</pre> <p>Add the following advanced VMX configuration parameters to the VMX file, and reboot the VM after adding these parameters:</p> <pre>ethernetX.pnicFeatures = "4" ethernetX.ctxPerDev = "3"</pre> <p>Change the rx-usec, lro, and rx/tx values of the VMXNET3 driver, and of the NIC used for the SAP database to app server traffic. Change them from the default value of 250 to 75 (25 is the lowest usable setting).</p> <p>Procedure: Log on to the OS running the inside the VM and use ethtool to change the following settings:</p> <pre>ethtool -C ethX rx-usec 75 ethtool -K ethX lro off ethtool -G ethX rx 512 rx-mini 0 tx 512</pre> <p>Note: Exchange x with the actual number, such as eth0. To make these ethtool settings permanent, see SLES KB 000017259 or the RHEL ethtool document.</p>
Storage configuration	<p>When creating your storage disks for SAP HANA on the VM/OS level, ensure that you can maintain the SAP-specified TDI storage KPIs for data and log files. 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 must be rebooted.</p> <p>Go to the ESXi console and set the following parameter:</p> <pre>vsish -e set /config/Disk/intOpts/VSCSIPollPeriod 100</pre> <p>If you want to use vSAN, then select one of the certified SAP HANA HCI solutions based on vSAN and follow the VMware HCI BP guide.</p>
SAP monitoring	<p>Enable SAP monitoring on the host by setting</p> <pre>Misc.GuestLibAllowHostInfo = "1"</pre> <p>For more details, see SAP note 1409604.</p> <p>Without this parameter, no host performance-relevant data will be viewable inside an SAP monitoring-enabled VM.</p>

Table 36. SAP HANA VM parameter settings

SAP HANA VM parameter settings	Description
Tips for editing a *.vmx file	Review tips for editing a *.vmx file in Broadcom KB 1714 .
UEFI BIOS guest	<p>Recommended: Use UEFI BIOS as the standard BIOS version for ESXi hosts and guests. Features such as Secure Boot are possible only with EFI.</p> <p>See VMware DOC-28494 for details. You can configure this with the vSphere Client by choosing EFI boot mode.</p> <p>If you are using vSphere 6.0 and you only see 2 TB memory in the guest, then upgrade to the latest ESXi 6.0 version.</p>
SAP monitoring	<p>Enable SAP monitoring inside the SAP HANA VM with the advancedVM configuration parameter: <code>tools.guestlib.enableHostInfo = "TRUE"</code>.</p> <p>For more details, see SAP note 1409604.</p> <p>Besides setting this parameter, the VMware guest tools must be installed. For details, see Broadcom KB 1014294.</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 Broadcom KB 2040375 .
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 socket memory for ESXi.</p> <p>Typical memory reservation for a host is between 32–64 GB for a 2-socket server, 64–128 GB for a 4-socket server, and 128–256 GB 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 or configure dedicated CPU resources per SAP HANA VM.</p> <p>You can use hyperthreads when you configure the VM 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 Broadcom KB 55636.</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 section of the vSphere Client.</p> <p>The vNUMA nodes of the VM will always be ≥ 1, depending on how many CPUs you have configured in total.</p> <p>If you need to access an additional NUMA node, use all CPU cores of this additional NUMA node, but use as few NUMA nodes as possible to optimize memory access.</p> <p>Half-Socket/sub-NUMA VMs require SNC-2 with Sapphire Rapids systems.</p>

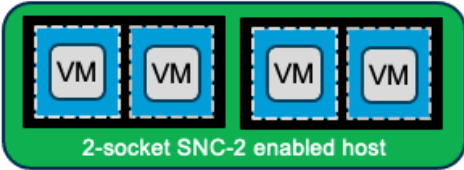
SAP HANA VM parameter settings	Description
	<p>Important: As of today, only 2-socket SPR systems are supported with SNC-2. Once SNC-2 is enabled, four sub-NUMA nodes will be available on a supported 2-socket SPR host. In the case of SNC-2, all logical CPUs of a sub-NUMA node should be assigned to such an SNC-2 VM.</p> <p>Additionally, the parameter <code>sched.nodeX.affinity="Y"</code> 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 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 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 must be utilized. For example, with a 60-core Sapphire Rapids CPU, when SNC is enabled, it becomes 30 cores per sub-NUMA node. Therefore, a half-socket VM will be configured with 60 vCPUs (including hyperthreads).</p>
Define the NUMA memory segment size	The <code>numa.memory.gransize = "32768"</code> parameter helps to align the VM memory to the NUMA memory map.
Paravirtualized SCSI (PVSCSI) controller or NVME adapter for I/O devices	Use the dedicated SCSI controllers for OS, log, and data to separate disk/I/O streams. 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 VM's file system	Use VMDK disks whenever possible to allow optimal operation via the vSphere stack. In-guest NFS mounted volumes for SAP HANA are supported as well.
Create datastores for SAP HANA data and log files	Ensure the storage configuration passes the SAP-defined storage KPIs for TDI storage. Use the SAP HANA hardware configuration check tool (HWCCT) to verify your storage configuration. For details, see SAP note 1943937 .
Use eager zero thick virtual disks for data and logs	We recommend this setting as it avoids lazy zeroing (initial write penalty).
VMXNET3	<p>Use paravirtual VMXNET 3 virtual NICs for SAP HANA VMs.</p> <p>We recommend at least 3–4 different NICs inside a VM in the SAP HANA VM (app/management server network, backup network, and, if needed, HANA system replication network). Corresponding physical NICs inside the host are required.</p>
Optimize the application server network latency if required	<p>Disable virtual interrupt coalescing for VMXNET 3 virtual NICs that communicate with the app servers or front end to optimize network latency. Do not set this parameter for throughput-oriented networks, such as vMotion or SAP HANA system replication. Use the advanced options in the vSphere Web Client or directly modify the <code>.vmx</code> file and add <code>ethernetX.coalescingScheme = "disable"</code>. X stands for your network card number.</p> <p>For details, see the Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs white paper.</p>

SAP HANA VM parameter settings	Description
Set <code>lat.Sensitivity = normal</code>	<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 is used.	<p>vSphere ESXi 8.0 introduces an enhanced virtual topology feature. This feature automatically selects optimal coresPerSocket values for VMs and optimal virtual L3 sizes.</p> <p>The intelligent, adaptive NUMA scheduling and memory placement policies in ESXi 8.0 can manage all VMs transparently, so you don't need to deal with the complexity of balancing VMs between nodes by hand.</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 the <code>sched.nodeX.affinity="Y"</code> advanced option).</p> <p>For details refer to Performance Best Practices for VMware vSphere 8.0 paper and for detailed configuration steps refer to the paper: VMware vSphere 8.0 Virtual Topology.</p> <p>If <code>sched.nodeX.affinity="Y"</code> cannot be used due to an operational reason, then NUMA action affinity must be set to avoid unwanted NUMA node co-deployments.</p>
NUMA action affinity	<p>If the <code>sched.nodeX.affinity="Y"</code> 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 Broadcom KB 2097369.</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: <code>Numa.LocalityWeightActionAffinity="0"</code></p>
<p>Associate VMs with specified NUMA nodes to optimize NUMA memory locality</p> <p>Note: While ESXi 8 provides an enhanced virtual topology feature that automatically selects optimal coresPerSocket values for VMs and optimal virtual L3 sizes to align it to the underlying NUMA topology of a host, we still recommend you associate a virtual NUMA node to a physical NUMA node. This is especially recommended for consolidation configurations, where more SAP VMs share a single host. For single large VM, these settings are not required.</p>	<p>For all SAP HANA VMs, we recommend you 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 used 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 VM is required to constrain how the ESXi scheduler can schedule a VM's CPU and memory.</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, "Associate Virtual Machines with Specified NUMA Nodes" for details.</p> <p>Procedure:</p> <ol style="list-style-type: none"> 1. Browse to the VM in the vSphere Client. 2. Right-click and select Edit Settings. 3. Select the VM Options tab and expand Advanced. 4. Under Configuration Parameters, click the Edit Configuration button.

SAP HANA VM parameter settings	Description
	<p>5. Click Add Row to add a new option.</p> <ul style="list-style-type: none"> – To specify a NUMA node for a specific virtual NUMA node on the VM: In the Name column, enter <code>sched.nodeX.affinity</code>, where X is the virtual NUMA node number. For example, <code>sched.node0.affinity</code> specifies the virtual NUMA node 0 on the VM. – In the Value column, enter the NUMA node where the VM or the virtual NUMA node can be scheduled. <p>6. For every virtual NUMA node, add a new row with a new <code>sched.nodeX.affinity</code> option.</p> <p>7. Click OK twice to close the Edit Settings dialog box.</p> <p>Example for a single NUMA node VM:</p> <ul style="list-style-type: none"> • <code>sched.node0.affinity = "2"</code> (this is physical NUMA node 2) <p>Example for a two-NUMA node VM:</p> <ul style="list-style-type: none"> • <code>sched.node0.affinity = "0"</code> • <code>sched.node1.affinity = "1"</code> <p>Not supported:</p> <ul style="list-style-type: none"> • <code>sched.node0.affinity = "1"</code> • <code>sched.node1.affinity = "2"</code> <p>Important: For SNC-enabled hosts, do not set affinities to sub-NUMA nodes that cross a physical CPU socket. See Table 37. Sub NUMA Node Cluster (SNC) settings for details.</p>
Configure VMs to use hyperthreading with NUMA	<p>For memory latency-sensitive workloads with low processor utilization, such as SAP HANA, or high interthread communication, we recommended using hyperthreading with fewer NUMA nodes instead of full physical cores spread over multiple NUMA nodes. Use hyperthreading and enforce NUMA node locality per Broadcom KB 2003582.</p> <p>This parameter is only required when hyperthreading should be leveraged for a VM. Using hyperthreading can increase the compute throughput but may increase the latency of threads.</p> <p>Note: This parameter is only important for half-socket and multi-VM configurations that do not consume the full server, such as a 3-socket VM on a 4-socket server. Do not use it when a VM leverages all installed CPU sockets (e.g., 4-socket wide VM on a 4-socket host or an 8-socket VM on an 8-socket host). If a VM has more vCPUs configured than available physical cores, this parameter is configured automatically.</p> <p>Use the vSphere Web Client and add the following advanced VM parameter:</p> <pre>numa.vcpu.preferHT="TRUE" (per VM setting) or as a global setting on the host: Numa.PreferHT="1" (host).</pre> <p>Note: For non-mitigated CPUs, such as Haswell, Broadwell, and Skylake, you may consider not using hyperthreading at all. For details, see Broadcom KB 55806.</p>
PMem-enabled VMs	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.

SAP HANA VM parameter settings	Description
	<p>Broadcom KB 78094 provides information on how to configure the NVDIMMs (VMware’s representation of Optane PMem) correctly and align the NVDIMMs to the physical NUMA architecture of the physical server.</p> <p>By default, Optane PMem allocation in the 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 NVDIMM 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> <pre>nvdimm0:0.nodeAffinity=0 nvdimm0:1.nodeAffinity=1 nvdimm0:2.nodeAffinity=2 nvdimm0:3.nodeAffinity=3</pre> <p><code>sched.pmem.prealloc=TRUE</code> is an optional parameter equivalent to eager zero thick provisioning of VMDKs and improves initial writes to 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 37. 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>The diagram shows a green rounded rectangle representing a host. Inside, there are two black-outlined rectangles representing CPU sockets. Each socket contains two blue-outlined rectangles representing VMs. Below the host, the text '2-socket SNC-2 enabled host' is written in green.</p>	<p>VM1:</p> <pre>numa.vcpu.preferHT="TRUE" sched.node0.affinity = "0"</pre> <p>VM2:</p> <pre>numa.vcpu.preferHT="TRUE" sched.node0.affinity = "1"</pre> <p>VM3:</p> <pre>numa.vcpu.preferHT="TRUE" sched.node0.affinity = "2"</pre> <p>VM4:</p> <pre>numa.vcpu.preferHT="TRUE" sched.node0.affinity = "3"</pre>
<p>1 VM per CPU socket, 2 VMs total per ESXi host</p>	<p>VM1:</p> <pre>numa.vcpu.preferHT="TRUE" sched.node0.affinity = "0" sched.node1.affinity = "1"</pre>

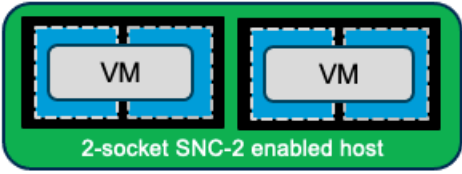
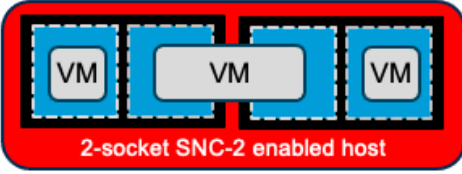
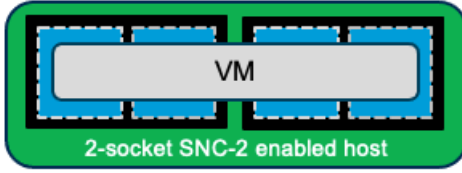
SAP SNC-2 enabled ESXi host (2-socket only)	Advanced VMX settings
 <p>2-socket SNC-2 enabled host</p>	<p>VM2:</p> <pre> numa.vcpu.preferHT="TRUE" sched.node0.affinity="2" sched.node1.affinity="3" </pre>
<p>Not supported</p>  <p>2-socket SNC-2 enabled host</p>	<p>Not supported: Do not configure a sub-NUMA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

Table 38. 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 2235581 and for settings, see SAP Note 2684254. Use SAP HANA SAPConf/SAPTune 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 or Linux VM.</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>Note: For RHEL5, edit <code>/etc/modprobe.conf</code> with the same line. Make a new <code>initrd</code> for the settings to take effect. You can do this either by using <code>mkinitrd</code>, or by re-running <code>vmware-config-tools.pl</code>.</p> <p>Starting in version 6, RHEL uses <code>modprobe.d</code>.</p> <p>Alternatively, append these to kernel boot arguments (for example, on Red Hat Enterprise Linux edit <code>/etc/grub.conf</code> or on Ubuntu edit <code>/boot/grub/grub.cfg</code>).</p> <pre>vmw_pvscsi.cmd_per_lun=254 vmw_pvscsi.ring_pages=32</pre> <p>Reboot the VM. See Broadcom KB 2053145 for details.</p> <p>Note: Review Broadcom KB 2088157 if you encounter VM freezes under heavy I/O load to recover the VM from its locked-up state.</p>
Install the latest version of VMware Tools	<p>VMware Tools is a suite of utilities that enhances the performance of the VM's guest operating system and improves the VM management. See Broadcom KB 1014294 for details.</p>
Configure NTP time server	<p>Use the same external NTP server as configured for vSphere. For details, see SAP note 989963.</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 the network latency of client/application server-facing NIC adapters. Run:</p> <pre>ethtool -K ethY lro off</pre> <p>Caution: Do not disable LRO for throughput NIC adapters such as for backup, replication, or SAP HANA internode communication networks.</p> <p>This works only with Linux kernel 2.6.24 and later and uses a VMXNET3.</p> <p>For additional details, see Broadcom KB 2055140.</p>

Table 39. 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 Reference Guide .
Disable I/O scheduling	SLES15 SP2 onwards is configured by default. The scheduler is set to <code>none</code> with <code>block_mq</code> enabled.
Disable AutoNUMA	<p>Later Linux kernels (RHEL 7 and SLES 12) support auto-migration according to NUMA statistics.</p> <p>For SLES:</p> <ol style="list-style-type: none"> # <code>yast bootloader</code> Choose Kernel Parameters tab (ALT-k) Edit the Optional Command-line Parameters section and append: <code>numa_balancing=disabled</code> <p>For RHEL:</p> <ol style="list-style-type: none"> Edit the file: <code>etc/sysctl.d/sap_hana.conf</code> Add the following: <code>kernel.numa_balancing = 0</code> Reconfigure the kernel by running: <code># sysctl -p /etc/sysctl.d/sap_hana.conf</code>
Use Block mq	<p>In the kernel parameters, add <code>scsi_mod.use_blk_mq=1</code>.</p> <p>For OS version SLES 5 SP2 and beyond, this is enabled by default.</p>
Disable transparent HugePages	<p>THP is not supported for the use with SAP HANA database because it may lead to hanging situations and performance degradations.</p> <p>To check the current configuration, run the following command:</p> <pre># cat /sys/kernel/mm/transparent_hugepage/enabled</pre> <p>Its output should read:</p> <pre>always madvise [never]</pre> <p>If this is not the case, you can disable the THP usage at runtime by issuing the following command:</p> <pre># echo never > /sys/kernel/mm/transparent_hugepage/enabled</pre> <p>For details, refer to the SAP WIKI for SAPs and the Linux OS vendor's virtualization independent recommended Linux OS settings for SAP HANA.</p>
Change the following parameters in <code>/etc/sysctl.conf</code> (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

When SAP HANA must perform as close to bare metal as possible with the lowest database access latency, you'll need to optimize the performance of SAP HANA VMs. We recommend sizing SAP HANA VMs with the fewest NUMA nodes for low latency. Configure an additional NUMA node and its resources for SAP HANA VMs that need more CPU or RAM.

To achieve optimal performance for an SAP HANA VM, use the settings as described in table 40 in addition to the previously described settings. CPU scheduling and priority settings improve performance by reducing the amount of vCPU and vNUMA migration, while increasing the priority of the SAP HANA production VM.

CPU affinity settings

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

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

Before you use a CPU affinity, consider the following:

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

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

Additional performance tuning settings for SAP HANA workloads

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

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

Table 40. Tunings for very low-latency SAP HANA VMs

SAP HANA VM parameter settings	Description
Tips about how to edit the *.vmx file	Refer to Broadcom KB 1714 .
<pre>monitor.idleLoopSpinBeforeHalt = "true" monitor.idleLoopMinSpinUS = "xx us"</pre>	<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 descheduling time.</p> <p>Background: The guest OS issues a halt instruction, which stops (or deschedules) the vCPU on the ESXi host. Keeping the VM 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 ESXi 7.0, the idle state of guest halt 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 halt will be emulated in the vmkernel, which may result in having a vCPU descheduled from the physical CPU, and this can lead to longer latencies. Therefore, we recommend you set this parameter to "TRUE" to ensure that the halt instruction is emulated inside the VM and not in the vmkernel.</p> <p>Use the vSphere Web Client to add the following advanced VM parameter:</p> <pre>monitor_control.halt_in_monitor = "TRUE"</pre>
<pre>monitor_control.disable_pause_loop_exiting = "TRUE"</pre>	<p>This parameter prevents the VM from unnecessarily exiting to the hypervisor during a pause instruction. This is specific for Intel Skylake systems.</p>

Table 41. Settings to improve NUMA alignment

SAP HANA VM parameter settings	Description
Configuring CPU affinity <code>sched.vcpuXx.affinity = "Yy-Zz"</code>	<p>Note: Remove <code>sched.nodeX.affinity</code> or <code>numa.nodeAffinity</code> settings (if they were set) and if using this parameter setting.</p> <p>By specifying a CPU affinity setting for each VM, you can restrict the assignment of VMs to a subset of the available processors (CPU cores) in multiprocessor systems. By using this feature, you can assign each VM to processors in the specified affinity set.</p> <p>See "Scheduler operation when using the CPU Affinity" (Broadcom KB 2145719) for details.</p> <p>This is required when configuring SAP HANA half-socket VMs or for very latency critical SAP HANA VMs. It is also required when the parameter <code>numa.slit.enable</code> is used.</p> <p>As with <code>sched.nodeX.affinity</code>, it is possible to decrease the CPU and memory latency by further limiting the ESXi scheduler from migrating 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 specific physical CPU thread and is necessary, for example, 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"> 1. Browse to the cluster in the vSphere Client. 2. Click the Configure tab and click Settings. 3. Under VM Options, click the Edit button. 4. Select the VM Options tab and expand Advanced. 5. Under Configuration Parameters, click the Edit Configuration button. 6. Click Add Row to add a new option. 7. In the Name column, enter <code>sched.vcpuXX.affinity</code> (XX stands for the actual vCPU want to assign to a physical CPU thread). 8. In the Value column, enter the physical CPU threads where the vCPU can be scheduled. For example, enter 0-55 to constrain the VM resource scheduling to physical CPU threads 0-55, which would be the 1st CPU of an 28 core CPU host. 9. Click OK. 10. Click OK to close the Edit VM dialog box. <p>For more information about potential performance practices, see vSphere Resource Management Guide.</p>

SAP HANA VM parameter settings	Description
>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 > 4 socket on 8-socket hosts.</p> <p>Note: <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 through the vSphere Web Client or by directly adding these parameters to the .vmx file with a text editor.

Table 42. 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 an 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>

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

Table 44. 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 45. 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 46. 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	<code>nvdimm0:0.nodeAffinity=0</code> <code>nvdimm0:1.nodeAffinity=1</code> <code>nvdimm0:2.nodeAffinity=2</code> <code>nvdimm0:3.nodeAffinity=3</code>

Table 47. 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 48. 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 49. 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 50. 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> <p>For a Skylake CPU:</p> <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.vCPUx.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	
Half-Socket	28	84	140	196	252	308	364	420	
	29	85	141	197	253	309	365	421	
	30	86	142	198	254	310	366	422	
	31	87	143	199	255	311	367	423	
	32	88	144	200	256	312	368	424	
	33	89	145	201	257	313	369	425	
	34	90	146	202	258	314	370	426	
	35	91	147	203	259	315	371	427	
	36	92	148	204	260	316	372	428	
	37	93	149	205	261	317	373	429	
	38	94	150	206	262	318	374	430	
	39	95	151	207	263	319	375	431	
	40	96	152	208	264	320	376	432	
	41	97	153	209	265	321	377	433	
	42	98	154	210	266	322	378	434	
	43	99	155	211	267	323	379	435	
	44	100	156	212	268	324	380	436	
	45	101	157	213	269	325	381	437	
	46	102	158	214	270	326	382	438	
	47	103	159	215	271	327	383	439	
	48	104	160	216	272	328	384	440	
	49	105	161	217	273	329	385	441	
	50	106	162	218	274	330	386	442	
	51	107	163	219	275	331	387	443	
	52	108	164	220	276	332	388	444	
	53	109	165	221	277	333	389	445	
	54	110	166	222	278	334	390	446	
	55	111	167	223	279	335	391	447	
Total CPU Threads	224	56	56	56	56	56	56	56	56
Total CPU Cores	112	28	28	28	28	28	28	28	28

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

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	48
Total CPU Cores	96	24	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.vCPUx.affinty = "Xx-Yy"` parameter. The list shows the starting and end range required for the "Xx-Yy" parameter. For example, for CPU 5 it would be "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	22

SAP HANA support and process

To get support for your virtualized SAP HANA systems, 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 that need escalation.

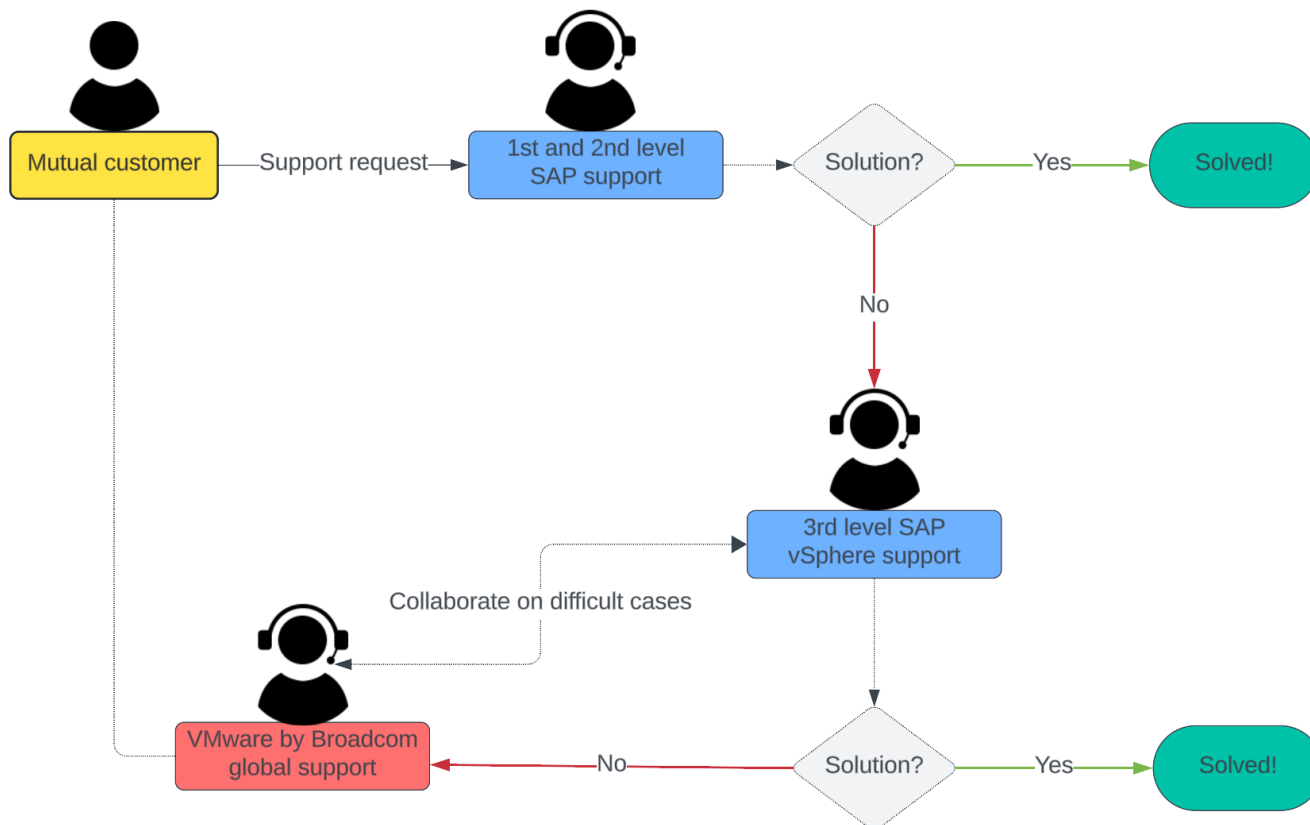
Before opening a VMware support ticket, we recommend opening a support request within the SAP support system for SAP HANA systems virtualized with vSphere. This ensures that SAP HANA and VMware specialists will work on the case and, if needed, escalate the issue to VMware product support (when it is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

You should escalate issues related to SAP HANA on vSphere to BC-OP-LNX-ESX using the SAP Solution Manager. For non-vSphere-related SAP HANA issues, the escalation will be moved to the correct support component. Figure 55 shows the support process workflow for vSphere-related SAP HANA issues.

Figure 55. SAP support workflow for vSphere-related escalations



If the problem cannot be solved by the first and second levels of SAP support, they forward it 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 VMware support because this may delay the support process. If you are uncertain that the issue is related to vSphere, open the ticket first at the general SAP HANA support component.

If the issue is related to vSphere, 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 at <https://support.broadcom.com/web/ecx/software-contact-support>.

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. Compared to a natively installed SAP HANA system, virtualizing SAP HANA with vSphere does not require any specific considerations for deployment and operation.

In addition, your SAP HANA environment gains all the virtualization benefits of easier operation, such as SAP HANA database live migration with vMotion or strict resource isolation on a virtual server, increased security, standardization, better service levels and resource utilization, an easy HA solution using vSphere HA, lower total cost of ownership, an easier way to maintain compliance, faster time to value, reduced complexity and dependencies, custom SAP HANA system sizes optimally aligned to your workload and needs, and the previously mentioned advantage of using a cloud 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 Broadcom contributed content or helped review this guide:

- Fred Abounader, Staff Performance Engineer, VCF Performance Engineering team
- Louis Barton, Staff Performance Engineer, VCF Performance Engineering team
- Julie Brodeur, Senior Technical Writer, VCF 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
- Oliver Ziltener, Senior Technical Product Manager, VCF division
- Catherine Xu, Manager of Workload, Technical Marketing team

