The Winding Road to Virtual SAP HANA® Application Workload Guidance Design for SAP S/4HANA® on VMware vSphere® 6.5

REFERENCE ARCHITECTURE LEVERAGING
APPLICATION WORKLOAD GUIDANCE DESIGN DOCUMENT
(FORMERLY *VMWARE VALIDATED DESIGN – VVD FOR SAP S4/HANA)*

JUNE 2017

Table of Contents

Business Case

SAP HANA®, the in-memory real-time platform, was initially introduced as a physical appliance and has steadily evolved to include support for virtualization with VMware vSphere® and SAP HANA tailored data center integration (TDI). Virtualized SAP HANA is now supported in scale-up and scale-out configurations in VMware® environments. Running SAP HANA on vSphere offers customers agility, resource optimization, and ease of provisioning. This solution enables SAP customers to provision instances of SAP HANA more quickly and effectively by using vSphere virtual machines (VMs). Using the SAP HANA platform with the vSphere virtualization infrastructure constitutes an optimized environment for achieving a unique, cost-effective solution. VMware capabilities such as VMware vSphere vMotion®, VMware vSphere Distributed Resource Scheduler™ (vSphere DRS), and VMware vSphere High Availability (vSphere HA) are inherent components of the virtualized SAP HANA platform.

The need exists for a comprehensive, "end-to-end" document that describes the implementation of a virtual SAP HANA deployment. VMware Solutions Labs was chosen to first develop and complete and then utilize the prescriptive approach of the *Application Workload Guidance Design* document formerly known as a VMware Validated Design (VVD) for SAP S/4HANA on vSphere 6.5. To exemplify the somewhat abstract instructions, which are generic in nature, various components were chosen to build out a working reference architecture and illustrate all facets of the virtual SAP HANA system. Although a seemingly unlimited number of viable hardware options and permutations are available for inclusion in customized architectures based on the Software-Defined Data Center (SDDC), a carefully considered set of hardware components was selected for this test. A Brocade switch and a Pure Storage FlashArray were chosen for the storage component. From a networking perspective, VMware NSX® was utilized throughout. Also, an all-flash VMware vSAN™ was configured using Western Digital SSDs to demonstrate the viability of VMware vSAN. Finally, Blue Medora's VMware vRealize® Operations Management Pack™ for SAP HANA was employed to display and generate performance statistics. Various performance tools were utilized to generate "SAP-like" mixed online transaction processing (OLTP) workloads. The full power of the SDDC was utilized and is on display throughout the system, exhibiting near-linear scalability, especially when sizing the VMs to match the NUMA topology.

Solution Overview

Virtualized SAP HANA is supported for scale-up and scale-out configurations on SAP HANA TDI–certified hardware. VMware provides best practices for deploying SAP HANA with either approach. A VVD for SAP HANA with a respective reference architecture serves as a blueprint for customers virtualizing their SAP HANA TDI environments.

Application Workload Guidance – Formerly VMware Validated Design

Application Workload Designs provide a comprehensive and extensively tested set of blueprints for building and operating an SDDC. Each design is customized for a chosen functionality and all applications running in that environment. Each holistic and standardized data center–level design spans compute, storage, networking, and management, providing a proven framework for how to deploy, configure, and operate an SDDC-based private cloud in support of a wide range of use cases. The documentation included with each VVD includes the following:

- A solution overview that details design objectives and software components
- An example design that is based on a customer scenario

Scope

This VMware Validated Design focuses on virtualized SAP HANA with tailored data center integration (TDI). Best practices for virtualizing SAP HANA for scale-up and scale-out configurations are leveraged in this design. The design is sized for various use cases, is tested on SAP-certified partner hardware, and is then validated. This reference architecture provides information specific to SAP HANA and leverages existing VVDs for general-purpose guidelines.

The physical architecture includes Dell R630 and R730 PowerEdge Servers with up to 1TB memory, a Pure Storage TDI array, and Brocade SAN and network components. The components are exemplary but are interchangeable. For example, the Pure Storage array is used for the testing, but any flash storage array can be substituted. The functional testing is not dependent on the storage type, so a VMware vSAN architecture can also be utilized. This work involves some performance testing, but it validates earlier vSphere 6.5 certification testing.

Audience

This guide is intended for administrators who are responsible for designing and deploying infrastructure for virtualized SAP S/4HANA.

Design Overview

A prototypical enterprise SAP customer with example requirements is used as a basis for all design activities in this solution. Requirements, assumptions, and constraints are carefully logged so all logical and physical design elements can be easily traced back to their source and justification.

Customer Requirements

Requirements are the key demands on the design. Sources include both business and technical representatives.

ID	REQUIREMENT
r101	SAP systems should be sized based on SAP HANA Quick Sizer output.
r102	The solution should provide appropriate sizing for the customer’s production environment.
r103	The solution should leverage standard high-performance servers.
r104	The project should size for SAP S/4HANA and SAP S/4HANA Suite on SAP HANA platform, based on module availability.
r105	The solution should use SAP HANA TDI–compliant hardware.
r106	The solution should provide end-to-end monitoring of the infrastructure and the applications.
r107	The solution should provide security that leverages VMware NSX microsegmentation.

Table 1. Customer Requirements

Design Constraints

Constraints limit the logical design decisions and physical specifications. These limitations are made independently of this engagement and might not align with stated objectives.

ID	CONSTRAINT
c101	Existing hardware from Dell, Pure Storage, Brocade, and Western Digital/SanDisk should be used for the infrastructure.
c102	The infrastructure used should be compliant with VMware Validated Designs.
c103	Solutions should leverage VMware components such as vSphere, VMware vRealize Automation™, VMware vRealize Operations™, and VMware NSX.

Table 2. Design Constraints

Design Assumptions

Assumptions are expectations about the implementation and use of a system that cannot be confirmed during the design phase and that are used to provide guidance regarding the design. If assumptions are not met, the respective design areas are at risk.

ID	ASSUMPTION
a101	Sizing information has been formulated based on typical customer requirements.
a102	Regarding SAP HANA sizing, the production workload is based on all vCPUs peaking at the same time. With one-to-one vCPU sizing, there is no overcommitment.
a103	With fewer vCPUs than cores, SAP HANA performance is linear.
a104	SAPS sizing: hyper-threading benefit for SAP OLTP = 15 percent. This benefit is used as excess capacity for the application.
a105	SAPS sizing: virtual overhead versus physical = 10 percent
a106	Minimum size for VM running an SAP database or application server (non CS) is two vCPUs.

Table 3. Design Assumptions

Use Cases

This design is targeted for the following use cases:

- Rapid provisioning to increase timelines for testing, training, and SAP upgrades; clones SAP systems quickly and easily
  - Test cases
- Enhanced application availability
  - High out-of-the-box availability, without complex configuration and setup, to protect against VMware ESXi™ server failure; no downtime due to ESXi server maintenance
  - Continuous utilization monitoring across resource pools and intelligent allocation of available resources among VMs, based on predefined rules that reflect customer business needs and changing priorities
  - VMware vSphere Fault Tolerance for SAP Central Services
- Server consolidation and data center energy cost reduction; runs SAP in VMs consolidated onto fewer physical servers that use less energy overall
- Unified performance monitoring of the virtualized SAP environment; coordinates performance reporting and analytics across the solution infrastructure stack, including the SAP and database tier, the guest operating system (OS), the hypervisor layer, and storage
- Enhanced security that leverages VMware NSX; enables restricted communication between application components

SAP HANA Deployment Options

SAP HANA Scale-Up Option

The scale-up—or single-host SAP HANA configuration, as named in the installation guide—describes in the context of virtualization a single SAP HANA instance that runs on a single VM on a vSphere host.

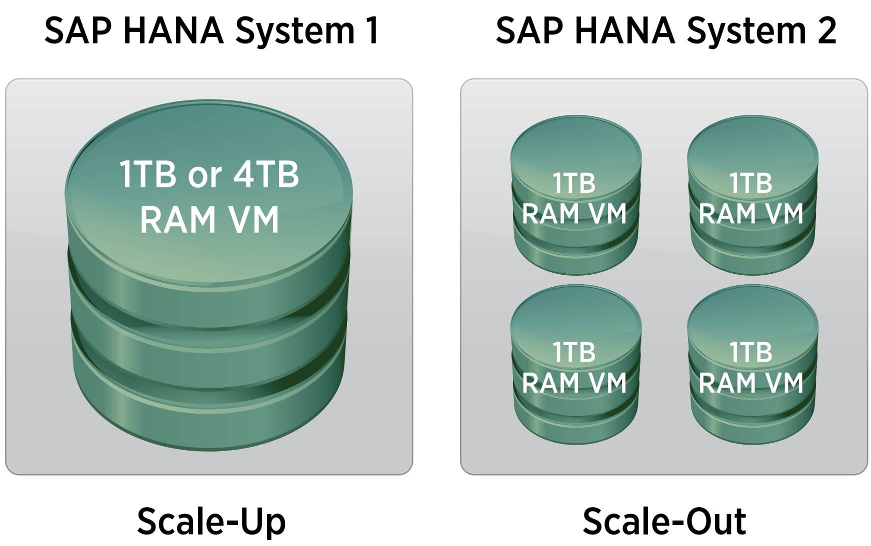

Figure 1. Scaling Options for an SAP HANA System on VMware vSphere

With a scale-up configuration, a single SAP HANA VM can be configured to consume all available host resources. After a VM reaches the host or vSphere limits, it cannot grow larger. Due to the CPU socket-to-RAM ratio limitation, the largest server systems and RAM sizes available for SAP Business Warehouse on HANA (BWoH) workload are 8-socket server systems with 2TB RAM, with up to 6TB RAM for specific SAP Business Suite workloads. Larger systems such as 16- or 32-socket server systems will become available and will enable the consolidation or codeployment of several maximum-sized SAP HANA VMs on a single server system.

Figure 2 shows such a single SAP HANA instance running on a single VM; it consumes all available server or VM resources. An SAP HANA system that already consumes the maximum server or VM limit cannot grow beyond these limitations. Adding resources or changing the VM resource configuration is required to support scale-up growth.

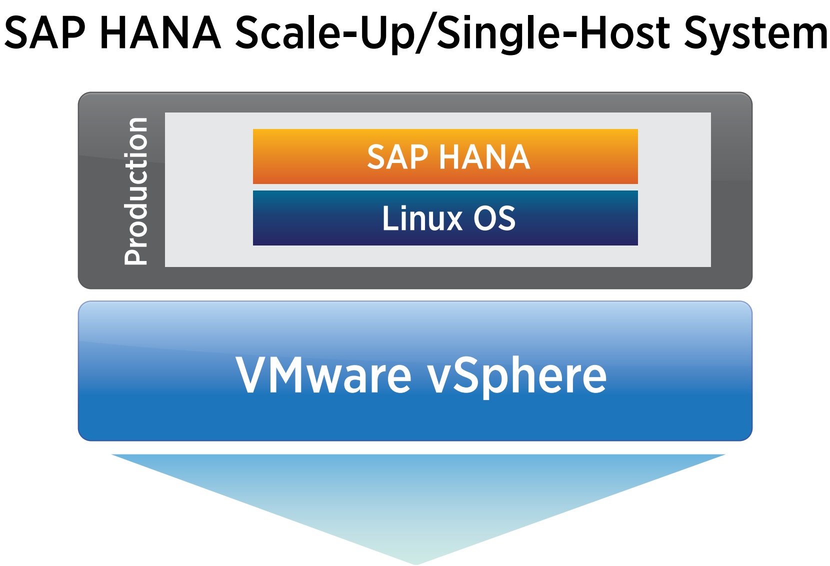

Figure 2. SAP HANA Single-Host System (Scale-Up) Configuration

SAP HANA Scale-Out Option

In an SAP HANA scale-out configuration running on vSphere, multiple SAP HANA "worker VMs" are distributed over multiple vSphere hosts. These SAP HANA worker VMs compose a single, distributed SAP HANA system. In other words, an SAP HANA scale-out configuration connects several SAP HANA instances to form one large distributed SAP HANA database system.

The benefits of this configuration are that it enables the SAP HANA database size to grow over time and it facilitates the addition of more vSphere hosts and SAP HANA worker VMs as needed.

Appliances certified for SAP HANA scale-out are available with 4 to 56 nodes, with two to eight CPU sockets and between 128GB and 3TB of RAM. Currently, scale-out configurations are limited to SAP Business Warehouse workloads and are generally supported by SAP for up to 16 SAP HANA worker systems. Business Suite workloads are not generally supported; however, for some scenarios—that is, native and virtualized—pilot phases are in process. For up-to-date details on this, refer to SAP Note 1825774. In addition, support for more than 16 SAP HANA worker nodes is available upon request from SAP.

In a physical scale-out configuration, a standby host, working as a "hot-spare," is required to provide high availability. In a virtualized environment, it is required to have only the compute resources available to support a failover and restart in case an SAP HANA VM fails. Unlike with a physical setup, spare resources can be used for other workloads so they are not wasted.

Figure 3 shows such a configuration, in which every ESXi host supports a single SAP HANA VM. The storage configuration is not shown here, and it requires a shared TDI storage subsystem. To ensure that sufficient resources are available in the event of a failure, either 1) hosts must be configured to provide spare resources or 2) resources used by noncritical VMs (shown in Figure 3 as "ANY Workload") must be available to support the failover and restart of any failed SAP HANA VMs.

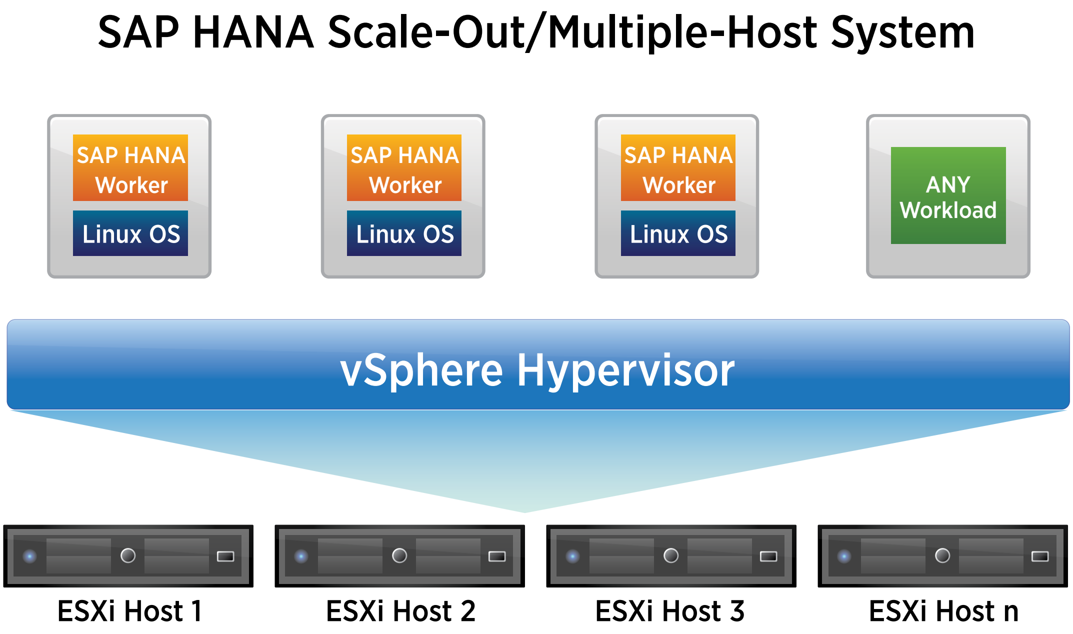

Figure 3. SAP HANA Multiple-Host System (Scale-Out) Configuration

SAP HANA Tailored Data Center Integration

The SAP HANA TDI deployment option enables customized SAP HANA solutions through which a customer can choose any supported SAP HANA server from an approved vendor, along with any supported SAP enterprise storage. TDI also enables the use of any hardware and infrastructure components that might already exist in an organization's data center and that comply with SAP HANA hardware requirements.

The following are among the customer benefits when using VMware virtualization combined with TDI:

- Faster SAP HANA deployment
- More flexibility in deploying SAP HANA
- Rightsizing instead of appliance T-shirt sizing
- No vendor lock-in
- Easier operation
- Higher SLAs due to virtualization policy–driven server management and operations

Figure 4 illustrates the appliance delivery model and the TDI model for SAP HANA systems. The appliance model shows all components in a box and indicates that these components are preconfigured, tested, and certified as an appliance. The TDI model shows the various components loosely "coupled"—such as server and storage—to highlight that a customer can choose from any supported vendor but is responsible for the installation and implementation of the overall system.

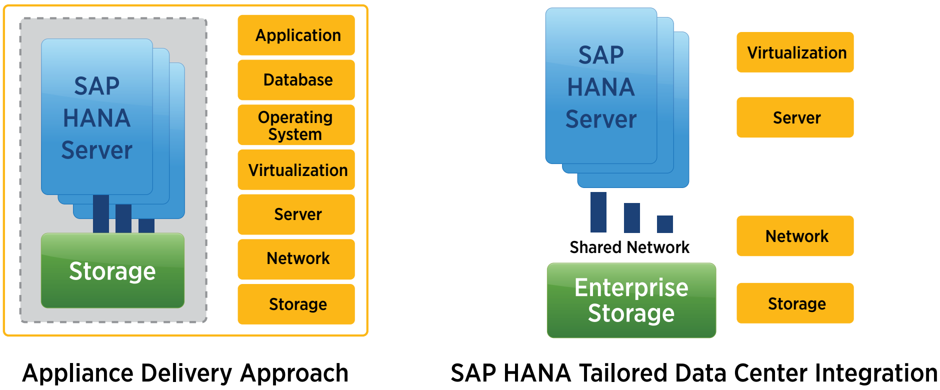

Figure 4. Appliance Delivery Model Versus SAP Tailored Data Center Integration

In scale-out configurations, use of either NAS or SAN shared storage is required for data persistence.

Currently, SAP has limited the number of active SAP HANA scale-out worker hosts to a maximum of 16 per SAP HANA TDI environment; this limit is also valid for virtualized systems. Customers that need more than 16 active worker hosts must contact the SAP HANA TDI back office to get support approval for larger SAP HANA scale-out systems. In some cases, “standby systems” do not count toward this limitation. For updated information, refer to the SAP HANA TDI FAQ document.

Best Practices Related to SAP

Recommended Configuration Settings for SAP S/4HANA on vSphere

Table 4 provides an overview of the configuration parameters and settings recommended for SAP HANA workloads on vSphere.

SERVER BIOS SETTINGS	DESCRIPTION
Enable Intel Virtualization Technology (Intel VT)	Enables all BIOS virtualization technology settings
Enable CPU Intel Turbo Boost Technology	Enables Intel automatic CPU core overclocking technology (P-states)
Disable Intel QPI Link Power Management	Disables static high power for QPI links
Enable Intel Hyper-Threading Technology	Doubles the logical CPU cores
Enable Execute Disable Feature	Enables the data execution prevention bit (NX-bit) required for vSphere vMotion
Disable Node Interleaving	
Disable C1E Halt State	Disables enhanced C-states in BIOS
Set Power Management to High Performance	Sets no power-saving modus (C-states)
Disable all unused BIOS features such as those described	Disables video BIOS shadowable, video RAM cacheable, on-board audio, on-board modem, on-board serial ports, on-board parallel ports, on-board game port, floppy drive, CD-ROM, USB

Table 4. vSphere Physical Host Server Settings

vSPHERE ESXi HOST	DESCRIPTION
PHYSICAL NETWORK ADAPTER SETTINGS	
Networking	Use VMware vSphere Distributed Switch™ instances to connect all hosts that work together. Define port groups that are dedicated to SAP HANA, management, and vSphere vMotion traffic.
STORAGE	
Log	Use dedicated HBAs or SCSI controller for DATA, LOG, and ESXi swap/scratch partition.
Data	
Swap	
MEMORY	
Configure RAM hemisphere mode	Distribute DIMM modules so as to achieve best performance (hemisphere mode); use fastest DIMM modules available for this RAM size.
CPU	
Populate all available CPU sockets; use glueless NUMA architecture	To avoid timer synchronization issues, use a multisocket server that ensures NUMA node timer synchronization. NUMA systems that do not run synchronized must synchronize the timers on the hypervisor area, which is very costly. Reference: "Timekeeping in VMware Virtual Machines"
SAP monitoring	Enable SAP monitoring on the host -> Misc.GuestLibAllowHostInfo and set the value to 1.
VM monitor	Enable vSphere to determine the best VM monitor, based on the combination of CPU and guest OS.
CPU/MMU virtualization option = automatic	Hardware-assisted memory virtualization.

Table 5. ESXi Server Settings

SAP HANA VIRTUAL MACHINE	DESCRIPTION
VMXNET3	Use paravirtual VMXNET3 virtual network adapters for SAP HANA VMs.
Set "ethernetX.coalescingScheme" to disable.	Disable virtual interrupt coalescing for VMXNET3 virtual network adapters that communicate with the application servers or front end.
Dedicated CPU and memory resources for SAP HANA instances	Do not overcommit CPU or memory resources.
Use as few NUMA nodes as possible.	Optimize VM size to as few NUMA nodes as possible.
Do not oversize the VM.	
Align virtual CPU VM configuration to actual server hardware.	Example: A VM running on a 4-socket server with 10-core CPUs should be configured with 10 cores per socket. A 20-core VM has two sockets.
numa.nodeAffinity = NUMA nodes (e.g., 0,1,2,3)	NUMA node localization is important for multi-VM configurations but is not needed for one-to-one configuration.
SAP monitoring	Enable SAP monitoring on the guest VM. tools.guestlib.enableHostInfo = true
Paravirtualized SCSI driver for I/O devices	Use multiple and dedicated SCSI controller; for details, see SAP HANA disk layout section.
Use VMware vSphere VMFS VMDKs.	
Create dedicated and isolated datastores for SAP HANA data and log files.	Ensure that the storage configuration passes the SAP HANA hardware check tool storage and file system requirements.
Use eager-zeroed thick virtual disks for data and log disk.	This prevents lazy zeroing.
Remove unused devices.	Such as floppy disks and CD-ROMs

Table 6. Virtual Machine Settings

LINUX OS	DESCRIPTION
Apply changes as documented in the following SAP and VMware notes:	
SAP HANA database: Recommended OS settings for SLES 11/SLES for SAP Applications 11 SP2	http://service.sap.com/sap/support/notes/1824819
SAP HANA Database: Recommended OS settings for SLES 11/SLES for SAP Applications 11 SP3	http://service.sap.com/sap/support/notes/1954788
Disable transparent HugePages in kernel.	echo never > /sys/kernel/mm/transparent_hugepage/ enabled verify with cat /sys/kernel/mm/ transparent_ hugepage/enabled always madvise (never) See SAP notes 2013638 (RHEL 6.5), 1954788 (SLES 11 SP3), and 1824819 (SLES 11 SP2).
Configure C-states for lower latency in Linux.	Disable in BIOS (e.g., static high performance or OS controlled). • If disabled in BIOS; optional: intel_idle.max_cstate=0 • If disabled in BIOS; optional: processor.max_cstate=0
GlibC update to avoid potential problem in the index server	Update the GlibC to at least version glibc-2.11.3-17.56.2.
Optional: Enabling LRO in Linux guest OS to lower latency for client network adapter	Optional use case to lower network latency of client traffic network adapter: "ethtool -K ethY lro on." Do not disable LRO for NFS or SAP HANA internode communication network! Works only with Linux kernel 2.6.24 and later and uses a VMXNET3.
Linux kernel	VMware strongly recommends using the latest kernel version.
Minimal supported SUSE Linux kernel 3.0.101-0.35.1se	Otherwise, customers might experience unplanned crashes or downtime when many CPUs and much memory are used.
Do not set page cache limits.	This should be done only when the system experiences a performance problem caused by heavy paging or swapping, as is described in the referenced SAP note. Linux: VMware vSphere configuration guidelines

Table 7. Virtual Machine Guest Operating System

LINUX OS	DESCRIPTION
Disable I/O scheduling.	Set kernel parameter "elevator=noop" -> Change the following parameters: Install latest VMware Tools. echo never > /sys/kernel/mm/transparent_hugepage/enabled
Disable all unnecessary Linux services such as those described.	anacron, apmd, atd, autofs, cups, cupsconfig, gpm, isdn, iptables, kudzu, netfs, and portmap Disable SUSE scheduler.
Install the following packages.	gtk2 (use version as provided with operating system distribution) java-1_6_0-ibm or later (use version as provided with operating system distribution, necessary for the SAP HANA studio on the SAP HANA system) libicu (use version as provided with operating system distribution) mozilla-xulrunner192-1.9.2.xx-x.x.x (use version as provided with operating system, but at given minimum version) ntp, sudo syslog-ng (use version as provided with operating system distribution) tcsh, libssh2-1, expect, autoyast2-installation. yast2-ncurses
Disable "Large Receive Offload."	Disabling LRO to lower latency for application that relies on TCP. Only for the application server and front-end interface. Add "ethtool -K ethX LRO off" to boot.local or edit the ethX setting through yast.
Install the latest version of VMware Tools™.	VMware Tools is a suite of utilities that enhances the performance of the VM's guest OS and improves management of the VM.
Turn off the SLES kernel dump function.	Turn off the SLES kernel dump function (kdump) if not needed for specific reasons.
Change the following parameters.	elevator=noop http://doc.opensuse.org/products/draft/SLES/SLES-tuning_sd_draft/cha.tuning.io.html vmw_pvscsi.cmd_per_lun=1024 Set them through YaST in the boot loader section. Copy them there, separated through space. vmw_pvscsi.ring_pages=32
/etc/sysctl.conf	net.core.rmem_default = 262144 net.core.wmem_max = 8388608 net.core.wmem_default = 262144 net.core.rmem_max = 8388608 net.ipv4.tcp_rmem = 4096 87380 8388608 net.ipv4.tcp_wmem = 4096 65536 8388608 net.ipv4.tcp_mem = 8388608 8388608 8388608 net.ipv4.tcp_slow_start_after_idle = 0
Adhere to the following shared memory settings.	Small, Shmmni value = 4096 if RAM > = 24 GB & < 64 GB Medium Shmmni value = 65536 if RAM > = 64 GB & < 256 GB Large Shmmni value = 53488 if RAM > 256 GB
Configure NTP time server.	Linux VMware timing problem
Align file system offset – Example: VNX 128K offset.	

Table 8. Virtual Machine Guest Operating System

Workload Design and Sizing

The SAP HANA VVD is based on a greenfield implementation using the SAP Quick Sizer tool for SAP HANA. Quick Sizer is a Web-based tool designed to facilitate a faster and easier way to size the SAP Business Suite. It has been developed by SAP in close cooperation with all platform partners and is free of cost. For more details, see SAP Quick Sizer (requires SAP login).

The Quick Sizer tool is used in the same way for virtualized deployments as for physical environments. Based on the input business requirements, the output generated by the SAP HANA Quick Sizer tool for the VVD is shown in Table 9.

SOFTWARE COMPONENT	SOFTWARE COMPONENT	CPU CAT.	SAPS (TOTAL, 2-TIER)	APP. SAPS (ABAP)	DATABASE SAPS	MEMORY CAT.	MEMORY (TOTAL, 2-TIER, MB)	APP. MEMORY (ABAP)	DATABASE MEMORY
BW/4HA_SRV	SAP BW/4HANA Server	XS	18.000	18.000	0	M	648.192	41.984	606.208
CRM SERVER	CRM Server	M	140.000	96.000	42.000	S	451.584	317.440	196.608
MDG-SERVER	Master Data Governance Server	XS	12.000	2.500	10.000	XS	235.520	9.216	226.304
S/4 SERVER	S/4 Server	S	64.000	47.000	17.000	M	904.192	132.096	773.120
SRM SERVER	SRM Server	L	150.000	130.000	21.000	S	430.080	326.656	196.608
TM SERVER	Transportation Management Server	XS	12.000	8.500	10.000	XS	199.680	24.576	196.608

Table 9. SAP HANA Quick Sizer Results

Figure 5 shows the scope of the production landscape that is sized in this chapter. The content in this section illustrates sizing and design separately for the application and the SAP HANA database.

Physical Design

This section details the ESXi hosts proposed for the vSphere infrastructure design. The hosts chosen are based on the customers' preferred hardware vendors and the standard models in use for database and application needs.

Database Host Design Specifications

The database hosts require high-memory systems with many cores to meet the requirements of the SAP S/4HANA database tier. The Dell R730 PowerEdge Server with two processors with 44 cores and 1TB RAM is a good building block for the database tier. Each node has four 10GBps network adapters for networking and two 16GBps Fibre Channel ports for high-performance SAN access.

Application Host Design Specifications

The application hosts do not require significant amounts of lot memory. Lower-priced hosts can be used to meet the requirements of the SAP S/4HANA application tier. The Dell R630 PowerEdge Server with two processors with 20 cores and 256GB RAM is the customer standard for application tier and is a good building block for this tier. Each node has four 10GBps network adapters for networking and two 16GBps Fibre Channel ports for high-performance SAN access.

The configuration and assembly process for each system is standardized, with all components installed the same on each host. Standardizing not only the model but also the physical configuration of the ESXi hosts is critical to providing a manageable and supportable infrastructure, because it eliminates variability. Consistent PCI card slot location, especially for network controllers, is essential for accurate alignment of physical-to-virtual I/O resources. The servers are described in Table 10. Both servers are on the *VMware Hardware Compatibility List*.

DELL POWEREDGE SERVER	SPECIFICATIONS	SPECint®_RATE2006 (THROUGHPUT RATING)	SAPS RATING (PHYSICAL)
R630 (Application tier)	Intel Xeon E5-2630 v4 2.2GHz, 25MB cache two processors/ 20 cores/40 threads, 1TB	835	55848 @ 100% CPU (No SAP benchmark cert. Estimated from SPECint ratings and SAP cert 2016013)
R730 (Database tier)	Intel Xeon E5-2699 v4 2.2GHz, 55MB cache two processors/ 44 cores/88 threads, 1TB	1760	117280 @ 100% CPU (SAP cert 2016013)

Table 10. Target Server SAPS Rating

The preceding SPECint ratings are available at https://www.spec.org/cpu2006/results/res2016q2/cpu2006-20160321-39711.html and https://www.spec.org/cpu2006/results/res2016q2/cpu2006-20160307-39256.html. The SAP benchmark certification is available at http://download.sap.com/download.epd?context=40E2D9D5E00EEF7CEE07089B1A1E0B4D782C09749CEA51D4E8F4C50F3C1581DB.

The Dell R730 PowerEdge Server is used to run SAP HANA VMs and is an entry-level server supported by SAP to run SAP HANA. See http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enen/entry-level-systems.html#categories=Dell&recordid=1810.

When sizing SAP on VMware with hyper-threaded Intel servers, there are two options: with one vCPU scheduled per core or with two vCPUs scheduled per core. In the latter case, the hypervisor leverages hyper-threading, which provides an extra boost in transaction throughput. This concept is explained in section 3.2.4 of http://www.vmware.com/files/pdf/business-critical-apps/sap-on-vmware-best-practices.pdf.

In this VVD design, we size based on the following requirements and assumptions:

- One vCPU per core, because it generates VM sizes with fewer vCPUs
- Hyper-threading performance benefit of 15 percent as documented in http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/whitepaper/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper.pdf.
- CPU utilization of 65 percent
- SAPS server ratings based on logical CPUs being utilized

Based on the preceding, Table 11 shows the SAPS per vCPU calculated for the two Dell servers

NOTE: The terminology used here in regard to the ESXi CPU scheduler, which impacts SAPS throughput, assumes the following:

- A *vCPU* and *a world* are interchangeably used to refer to schedulable CPU context, corresponding to a process in conventional OSs.
- On hyper-threaded systems, a physical core has two logical processors; a vCPU, or world, is scheduled on a logical processor.
- A vCPU can be scheduled on a logical processor on a core while the other logical processor of the core is idle. This is referred to as one vCPU scheduled per core.
- Two vCPUs can be scheduled on the two logical processors of the same core. This is referred to as two vCPUs scheduled per core.
- This terminology is consistent with the following guide: The CPU Scheduler in VMware vSphere 5.1.

DELL POWEREDGE SERVER	SERVER PHYSICAL SAPS (INC HT BENEFIT, @100% CPU)	SERVER VIRTUAL SAPS 2 vCPUs PER CORE (@100% CPU)	SERVER VIRTUAL SAPS 1 vCPU PER CORE (@ 100% CPU)	SAPS PER vCPU 1 vCPU PER CORE @65% CPU
		minus 10%	minus 15%	/# of cores * 0.65
R630 20 cores (Application tier)	55,848	50,264	43,707	**1,420**
R730 44 cores (Database tier)	117,280	105,552	91,784	**1,356**

Table 11. Calculate SAPS per vCPU for Each Target Server

Application Tier Sizing

This section calculates the hosts and VMs required for the application tier, based on the Quick Sizer business requirements and SAPS per vCPU rating shown in Table 11. Table 12 shows the results.

SAP COMPONENT	APPLICATION TIER SAPS REQUIREMENT	# OF vCPUs REQUIRED (@ 1,420 SAPS PER vCPU)	# OF 10-WAY VMs (NUMA NODE SIZE)	# OF 4-WAY VMs
BW/4HA_SRV	18,000	13	2	4
CRM	96,000	68	7	17
MDG	2,500	2	1	1
S/4	47,000	34	4	9
SRM	130,000	92	10	23
TM	8,500	6	1	2
TOTAL APPLICATION TIER SAPS	302,000	# of hosts required @ 65% CPU = 302,000/(43,707*.65) = 10.6 (does not include VMware HA and Central Services requirements)		
ABAP SAP Central Services	Assume standalone 2 vCPU VM for ASCS for each SAP component => 2 cores per ASCS VM Total cores required for 6 SAP systems = 6 x 2 = 12 Assume vSphere HA protection with vSphere FT = 12 cores for the shadow VMs Total cores for ASCS + vSphere FT = 24 = 24/20 = 1.2 hosts			

Table 12. Estimate Virtual Machine Size plus Total Number of Hosts for Application Tier

The application tier can scale out to multiple VMs, so there are options for choosing vCPU size and number of application server machines. Table 12 shows two VM sizes for each SAP component: 10-way (10 vCPUs) and 4-way (4 vCPUs). To determine VM size, the following guidelines are followed *for each SAP component:*

- Standardize on the same VM size.
- For high availability, there are at least two VMs.
- Minimize the total number of VMs.
- In this example, consider VM sizes with an even number of vCPUs.
- Size the vCPU count to fit within a NUMA node. vSphere automatically schedules the VM within a NUMA node and maintains best performance by minimizing remote memory access.

Given the preceding guidelines, the following application server VMs were chosen.

SAP COMPONENT	APPLICATION TIER VM vCPU	vCPU OVERSIZING (VM vCPU COUNT MINUS THE vCPUs CALCULATED BASED ON SAPS)	APPLICATION TIER MINIMUM VM MEMORY (GB) (SOURCE: QUICK SIZER OUTPUT)
BW/4HA_SRV	2 x 10-way	20-13 = 7	42/2 = 21
CRM	7 x 10-way	70-68 = 2	318/7 = 46
MDG	2 x 2-way	4-2 = 2	10/2 = 5
S/4	4 x 10-way	40-34 = 6	133/4 = 34
SRM	10 x 10-way	100-92 = 8	327/10 = 33
TM	2 x 4-way	8-6 = 2	25/2 = 13
Total Application Server vCPUs	242		
Total vCPUs for ASCS + vSphere FT Shadow VMs	24		

Table 13. Determine Final Application Server VM Sizes for Each SAP Component

The VMs listed in Table 13 are deployed on a 13-host ESXi cluster with a total of 260 cores; that is 13 x 20. The total vCPU count from Table 13 is 242 + 24 = 266. This indicates a vCPU overcommit situation; that is, the total number of vCPUs is greater than the total number of cores. This is not a concern, however, because there is an oversizing of vCPUs equal to 27 vCPUs, as is shown in the third column of the table. In the event of a single host failure, after all the impacted VMs have restarted on the remaining 12 hosts, there will be an even greater vCPU overcommitment. This will not create a serious CPU bottleneck for the following reasons:

- There is a total vCPU oversizing due to the rounding up of the VM vCPUs to a standardized size; that is, 27 vCPUs.
- Sizing has been calculated at 65-percent CPU utilization, so there is headroom.
- The hyper-threading benefit adds an additional buffer.
- In practice, not all workloads across the VMs peak at the same time.

Given the preceding specifications, after commencement of production workloads, real-time monitoring of the workload will most likely generate additional consolidation opportunities.

Application Tier Deployment on ESXi Hosts

Figure 5 shows likely scheduling and placement for the application server VMs, based on two-socket, 20-core ESXi hosts with hyper-threading enabled.

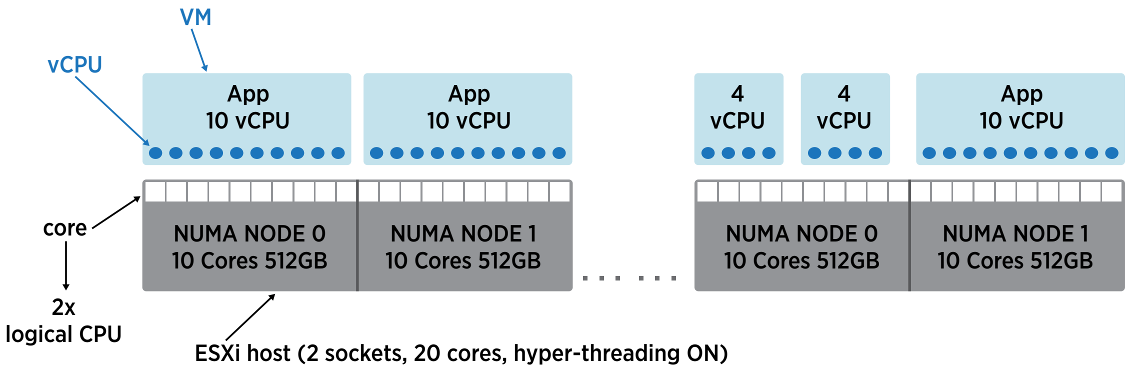

Figure 5. Example Placement of Application Tier VMs on ESXi Hosts

The diagram demonstrates the following on each ESXi host:

- Scheduling functionality of the vSphere hypervisor – Having no specific tuning and assuming all VMs are under load, the scheduler automatically performs these functions:
  - Allocates a home NUMA node for the memory of each VM
  - Schedules vCPUs of each VM on its home node, thereby maintaining local memory access
  - Schedules each vCPU on a dedicated core to enable exclusive access to core resources

As a result of the previously listed scheduling activities, the VMs are scheduled and balanced by the hypervisor across an ESXi host so a VM or groups of VMs fit within a NUMA node in the following manner:

- The memory used by all VMs in a NUMA node plus NUMA node overhead memory is less than the NUMA node memory size.
- The total number of vCPUs scheduled on a NUMA node is less than or equal to the total number of cores per NUMA node.
- A wide VM creates a different scenario: VM memory is split across NUMA nodes.

All performance benefits due to local memory access are handled by the hypervisor. No NUMA optimization inside the guest OS is required by the application.

SAP HANA Database Tier Sizing

Table 14 shows the estimated SAP HANA VM sizes based on the Quick Sizer business requirements.

SAP COMPONENT	DATABASE MEMORY REQUIREMENT (GB)	DATABASE TIER SAPS REQUIREMENT	# OF vCPUs REQUIRED (@1,356 SAPS PER vCPU)	PROPOSED VM SIZE (TARGET HOST: 2 SOCKETS, 44 CORES, 1 TB)
BW/4HA_SRV	607	n/a	2 sockets	650GB, 32 vCPU
CRM	197	42,000	31	250GB, 32 vCPU
MDG	227	10,000	8	250GB, 8 vCPU
S/4	774	17,000	13	800GB, 16 vCPU
SRM	197	21,000	16	250GB, 16 vCPU
TM	197	10,000	8	250GB, 8 vCPU
Total Memory	Total CPU 198 vCPU			

Table 14. Estimated SAP HANA Virtual Machine Size

The vCPUs for the SAP HANA VMs in Table 14 have been rounded up in the far-right column to the nearest multiple of 8. Some SAP guidelines on matching socket sizes of VMs are ignored to enable right-sizing of the environment and to reduce overprovisioning.

Data Center Design

In vSphere, a data center is the highest-level logical boundary. It can delineate separate physical sites or locations as well as vSphere infrastructures with completely independent purposes.

Within vSphere data centers, ESXi hosts typically are organized into clusters. Clusters group similar hosts into a logical unit of virtual resources, enabling the following technologies:

- vSphere vMotion
- vSphere HA
- vSphere DRS
- vSphere FT

vSphere Clusters

As part of this logical design, vSphere clusters are created to aggregate hosts. Due to the high cost of the database cluster, an N+1 cluster model is used. Having two additional hosts means 50-percent overprovisioning for high availability and maintenance. If the SLA warrants it, this should be made an N+2 cluster.

For the application cluster, an N+2 model is used to enhance redundancy for the large number of nodes in the cluster and to enable availability during maintenance.

CLUSTER SIZE CALCULATION: DATABASE TIER		
DATABASE CLUSTER	CPU	MEMORY
Total	112	2,450
Per node	44	1,024
Number of Nodes	2.55	2.39
Cluster Size (N+1)	4	4
CLUSTER SIZE CALCULATION: APPLICATION TIER		
APPLICATION CLUSTER	CPU	MEMORY
Total	266	814
Per node	20	256
Number of Nodes	13.30	3.18
Cluster Size (N+1)	15	5

Table 15. Total Number of Hosts Required

CLUSTER	DATABASE	APPLICATION
vSphere cluster size	4 hosts	15 hosts
Capacity for host failures per cluster	1 host	1 host
Capacity for maintenance	None	1 host

Table 16. Cluster Sizing for the Application and Database Tiers

Assumptions and Caveats

- The number of VMs protected by vSphere FT must be counted as doubled.
- Host failures for vSphere HA are expressed in the number of allowable host failures; that is, the expected load should be able to run on surviving hosts. vSphere HA policies should be applied on a percentage spare capacity basis to avoid slot-based admission control.
- Using dedicated hosts for maintenance assumes that a host is reserved strictly to offload running VMs from other hosts that must undergo maintenance. When not being used for maintenance, such hosts can also provide additional spare capacity to support a second host failure or for unusually high demands on resources. Having dedicated maintenance hosts can be considered somewhat conservative because spare capacity is allocated only for maintenance activities. Such spare capacity is earmarked here so there is sufficient capacity to run VMs with minimal disruption.
- Clusters can be created and organized to enforce resource allocation policies such as the following:
  - Load balancing
  - Power management
  - Affinity rules
  - Protected workloads
  - Limited number of licenses available for specific applications

Availability Design

The availability design depends on the single point of failure (SPOF) analysis of components. There are components in the SAP infrastructure that are one of a kind and are potential SPOFs; other components are capable of having multiple instances for load balancing and availability.

SAP S/4HANA Architecture

This section summarizes SAP S/4HANA architecture concepts and terminology used in this document. SAP uses the term *system landscape,* which contains all the SAP systems that have been installed. It can consist of several system groups for which SAP systems are linked by transport routes. Transport routes refer to the path of code migrations between SAP systems, for example from development (DEV) to quality assurance (QAS) to production (PRD) (https://help.sap.com/saphelp_nw74/helpdata/en/63/a30a4ac00811d2851c0000e8a57770/content.htm).

The architecture of a single SAP system is multitier and consists of the following components:

- Application servers (SAP Web application servers) – These are ABAP or Java (J2EE) based, depending on the specific SAP product or module. Two types exist:
  - Primary application server (PAS) – An application server instance that is installed with SAP Central Services in newer NetWeaver releases and is part of the base installation
  - Additional application servers (AAS) – Application servers installed as required for horizontal scalability
- SAP Message Service – The SAP Message Service is used to exchange and regulate messages between SAP instances in an SAP system. It manages functions such as determining which instance a user logs in to during client connect and scheduling batch jobs on instances configured for batch.
- SAP Enqueue Service – The SAP Enqueue Service manages the locking of business objects at the SAP transaction level. Locks are set in a lock table stored in the shared memory of the host on which the SAP Enqueue Service runs.
- Database server – SAP S/4HANA and SAP S/4HANA Suite support SAP HANA as the backend database for all applications. Each module has its own individual standalone SAP HANA database.

- The following SAP services are defined based on the Message Service and Enqueue Service:
  - SAP Central Services – In newer SAP versions, the Message Service and Enqueue Service have been grouped into a standalone service. Separate SAP Central Services exist for ABAP- and Java-based application servers. For ABAP variants, it is called ABAP SAP Central Services (ASCS); for J2EE variants, it is called SAP Central Services (SCS).
  - Replicated enqueue server – This component consists of the standalone enqueue server and an enqueue replication server. The replicated enqueue server runs on another host and contains a replica of the lock table (replication table). If the standalone enqueue server fails, it must be restarted on the host on which the enqueue replication server is running, because this host contains the replication table in a shared memory segment. The restarted enqueue server uses this shared memory segment to generate the new lock table, after which the shared memory segment is deleted. SAP Central Services and the database are both SPOFs and therefore require considerations for high availability.

Single Point of Failure Analysis

The following SPOFs exist in the SAP NetWeaver architecture:

- Database – Every application work process makes a private connection to the database at the start. If the connection is interrupted due to database instance failure, the work process attempts to set up a new connection and changes to "database reconnect" state until the database instance restarts. User sessions with database activity in process receive SQL error messages, but their logged-in sessions are preserved on the application server.
- SAP Message Service and Enqueue Service – Failure of these services has a considerable effect on the system because all transactions that contain locks must be rolled back and any SAP updates being processed fail.

The isolation of the Message Service and Enqueue Service from the central instance (CI) helps address the high-availability requirements of these SPOFs. The SAP Central Services component is "lighter" than the CI and is much quicker to start up after a failure.

SAP Application Tier Availability

Every SAP application should have a minimum of two servers. Because they are serving the same function, these servers must be separated from each other by using antiaffinity rules. An SAP application server and database server can be collocated on the same physical server by using affinity rules to optimize performance in certain situations. vSphere DRS affinity rules help protect SAP application and database components by providing appropriate separation between the primary and standby servers as well as between multiple application servers.

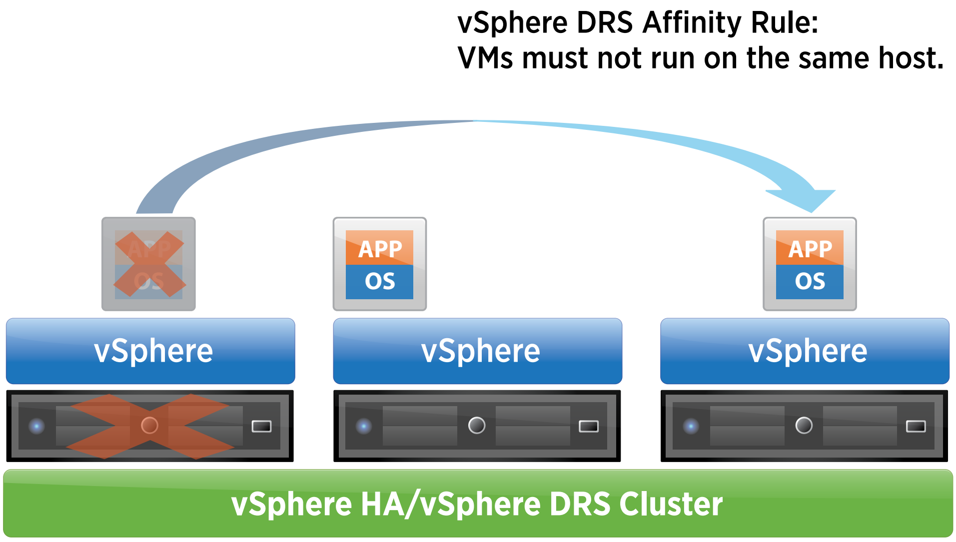

Figure 6. SAP Application Tier Availability

vSphere Fault Tolerance for SAP Central Services

vSphere FT provides continuous availability for a VM by creating and maintaining another VM that is identical and continuously available to replace the original in the event of a failover situation. The protected VM is called the primary VM. The duplicate VM—that is, the secondary VM—is created and runs on another host. The secondary VM's execution is identical to that of the primary VM. vSphere FT functionality includes the following elements:

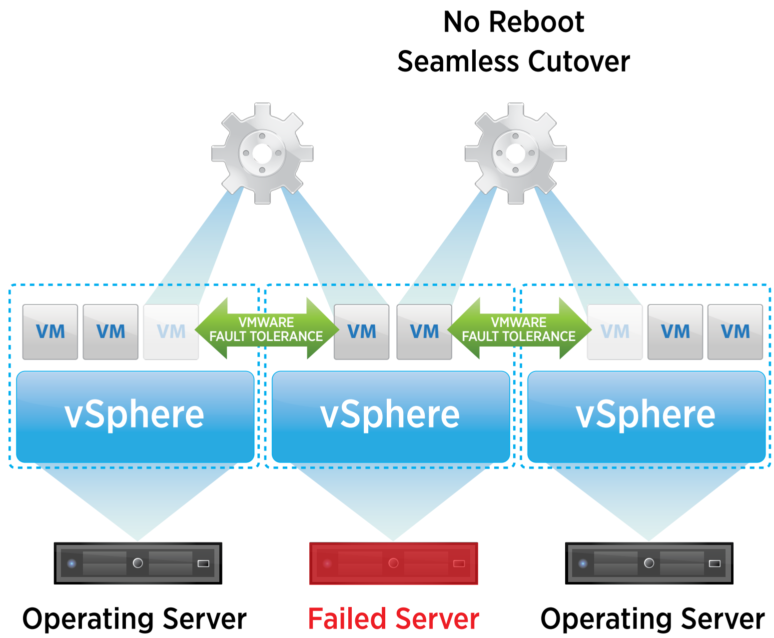

Figure 7. vSphere Fault Tolerance Functionality

- If the host running the primary VM fails, the secondary VM is immediately activated to replace it with no interruption of service to users. A new secondary VM is started, and vSphere FT redundancy is reestablished automatically.
- If the host running the secondary VM fails, it is also immediately replaced.
- A fault-tolerant VM and its secondary copy are not allowed to run on the same host.
- The primary VM can be a maximum of four vCPUs and 64GB.
- The secondary VM has its own copy of the primary VM's virtual disks.
- A dedicated 10GbE network adapter is recommended for vSphere FT logging traffic. vSphere FT logging traffic between primary and secondary VMs contains guest network and storage I/O data as well as the memory contents of the guest OS.

For more details on vSphere FT and best practices and limitations, see the *vSphere Availability Guide* (https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-availability-guide.pdf) and *Performance Best Practices for VMware vSphere 6.0* (http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf).

Based on this discussion, vSphere FT for SAP Central Services is an ideal candidate and is leveraged. The database tier is unsuitable due to the limitations of vSphere FT from a sizing perspective.

vSphere HA and vSphere DRS Antiaffinity for Application Server Availability

SAP HANA Database Tier

Table 17 shows the various options that can be leveraged for availability of the SAP HANA database layer. The database as it stands is a SPOF. The type of solution used depends on the criticality of the application and the recovery time objective (RTO) provided by the customer for local failures of the application.

HIGH-AVAILABILITY SOLUTION	vSPHERE HA	vSPHERE HA + SAP HANA AUTO-RESTART FEATURE	SAP HANA HOST AUTO-FAILOVER (STANDBY VM)
Scenario description	VMware standard vSphere HA solution. vSphere HA restarts or fails over VM to another host in case of a detected OS or hardware failure.	Standard vSphere HA combined with SAP HANA auto-restart watchdog runs inside a VM to monitor SAP HANA application status and triggers an SAP HANA process restart. OS and hardware failures are handled by vSphere HA.	SAP HANA standby VM automatically assumes the role of another SAP HANA VM in case of a detected failure.
Operating system failures	Yes	Yes	Yes
Hardware failures	Yes	Yes	Yes
Application failures	No	Yes	Yes
IP redirect/ DNS update	Not necessary	Not necessary	Not necessary
RTO	Medium (crash recovery of database)	Medium (crash recovery of database)	Short to medium (only if in-memory data loading)
RPO	0	0	0
Performance ramp	Minutes to hours (bootstrap + loading)	Minutes to hours (bootstrap + loading)	Minutes to hours (bootstrap + loading)
Cost	Included	Included	Included
Complexity	Low	Low	Medium

Table 17. SAP HANA Database Layer Availability Options

Each cluster is configured for vSphere HA to automatically recover VMs if any ESXi host fails or if there is a specific VM failure. A host is declared failed when the vSphere HA master node cannot communicate with it by means of a network heartbeat or a storage heartbeat and the host cannot be pinged. A VM is declared failed if the heartbeat inside the guest OS is no longer being received and there is no network or storage I/O.

SAP HANA Systems Protected by vSphere HA

VMware provides vSphere products with built-in and optional high-availability and disaster recovery solutions to protect a virtualized SAP HANA system at all levels.

The power behind the VMware high-availability and disaster recovery solutions is in how they are layered to protect against failures at every level of the data center—from individual components, such as network adapters and HBA card teaming, up to and including the entire site. With vSphere replication, the solutions provide protection against both planned and unplanned downtime.

Figure 8 shows the various solutions available to protect against failures, from component level to complete site.

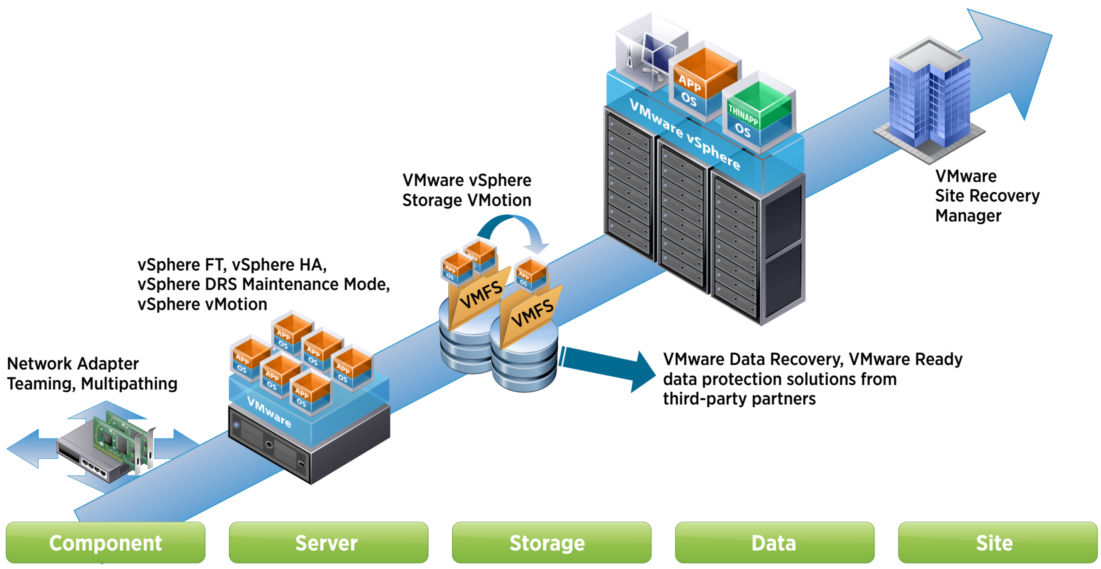

Figure 8. VMware High-Availability and Disaster Recovery Solutions – Protection at Every Level

Many of the key features of virtualization, such as encapsulation and hardware independence, already offer inherent protection. Beyond that, the vSphere platform provides these additional safeguards to ensure that organizations can meet their availability requirements:

- Protection against hardware failures
- Planned maintenance with zero downtime
- Protection against unplanned downtime and disasters

This chapter focuses on vSphere HA, the easiest way to protect an SAP HANA system against unplanned downtime caused by OS or hardware crashes.

vSphere HA for Database and vSphere FT for SAP Central Services

The ideal high-availability solution for SAP S/4HANA SPOF components is to leverage vSphere HA and other SAP-specific mechanisms for the database and vSphere FT for SAP Central Services. By having more than two application servers with antiaffinity rules, the application services can be protected.

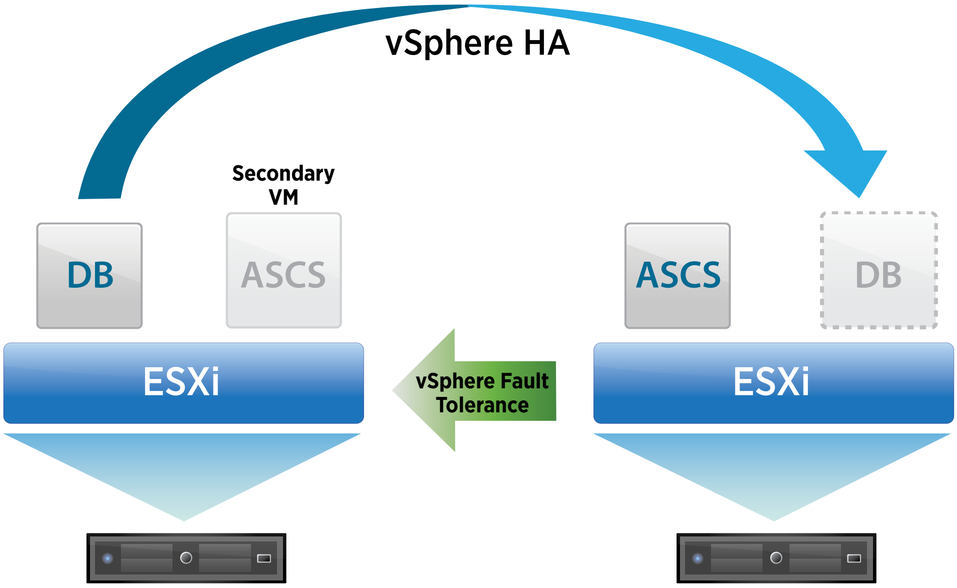

Figure 9. vSphere HA and vSphere FT – The Ideal High-Availability Solution for SAP S/4HANA SPOF Components

Storage Design

VMware storage virtualization can be categorized into three layers of storage technology. The bottom layer is the storage array, consisting of physical disks presented as logical disks—that is, storage array volumes or LUNs—to the layer above, the virtual environment occupied by vSphere. Storage array LUNs are formatted as VMware vSphere VMFS—that is, virtual machine file system—volumes in which virtual disks can be created. VMs consist of virtual disks that are presented to the guest OS as disks that can be partitioned and used in file systems. Figure 10 shows these storage layers.

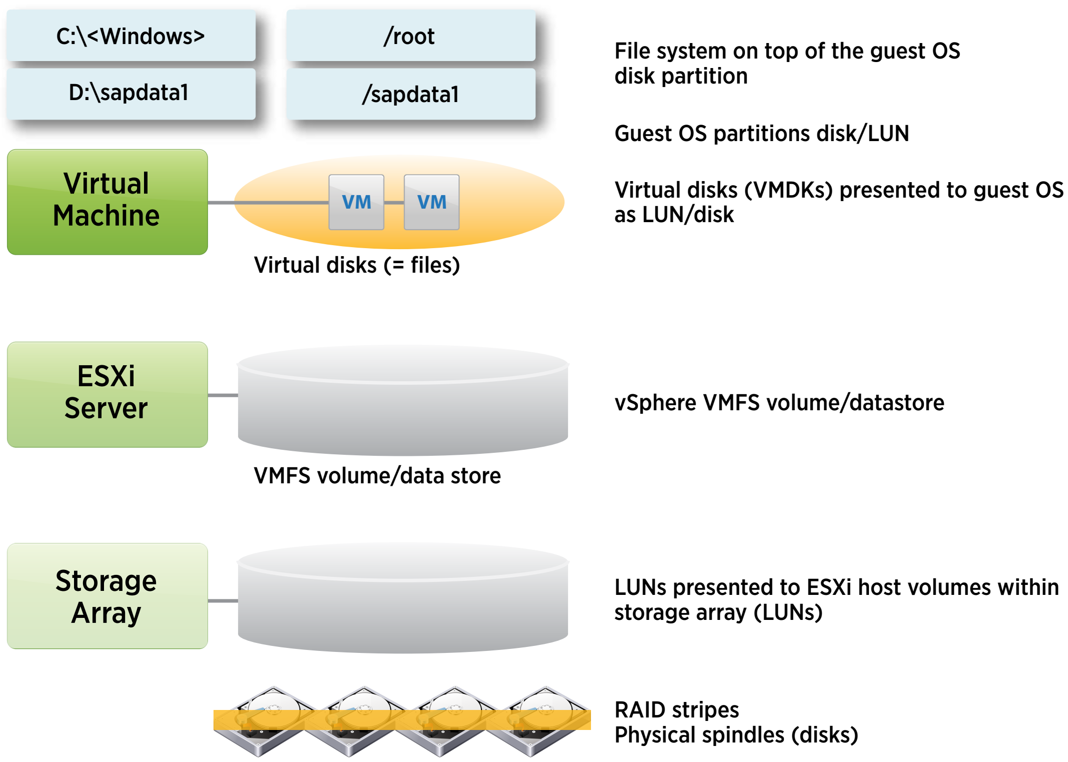

Figure 10. Storage Virtualization Layers

vSphere VMFS is a cluster file system that provides storage virtualization optimized for VMs. Each VM is encapsulated in a small set of files, and vSphere VMFS is the default storage system for these files on physical SCSI disks and partitions. VMware supports Fibre Channel and iSCSI protocols for vSphere VMFS.

VMware also supports raw device mapping (RDM). RDM enables a VM to directly access a volume on the physical storage subsystem. It can be used only with Fibre Channel or iSCSI. RDM provides a symbolic link from a vSphere VMFS volume to a raw volume. The mapping makes volumes appear as files in a vSphere VMFS volume. The mapping file, not the raw volume, is referenced in the VM configuration.

To access virtual disks, a VM uses virtual SCSI controllers. These virtual controllers include BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual. From the standpoint of the VM, each virtual disk appears as if it were a SCSI drive connected to a SCSI controller.

For more background on storage virtualization, see the *vSphere Storage Guide* at https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf.

Virtual Storage Design

This section applies the virtual storage concepts to a storage design for an SAP database on vSphere. The example here is a scale-up SAP HANA system based on the virtual SCSI controller assignments described in the *Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere* white paper (http://www.vmware.com/files/pdf/SAP_HANA_on_vmware_vSphere_best_practices_guide.pdf).

VIRTUAL SCSI DRIVER	VIRTUAL DEVICE	FILE SYSTEM
LSI Logic/Paravirtual	0:0	/usr/sapX
Paravirtual	1:0	/hana/data/<SID>
Paravirtual	2:0	/hana/log/<SID>
Paravirtual	3:0	/hana/shared/<SID>

Table 18. SAP HANA Layout for Virtual SCSI Controller

The boot drive can be configured with the paravirtual driver, but configuration depends on the guest OS version. For details, see VMware knowledge base article 1010298, *Configuring Disks to Use VMware Paravirtual SCSI (PVSCSI) Adapters* (http://kb.vmware.com/kb/1010398).

Figure 11 shows an example storage layout.

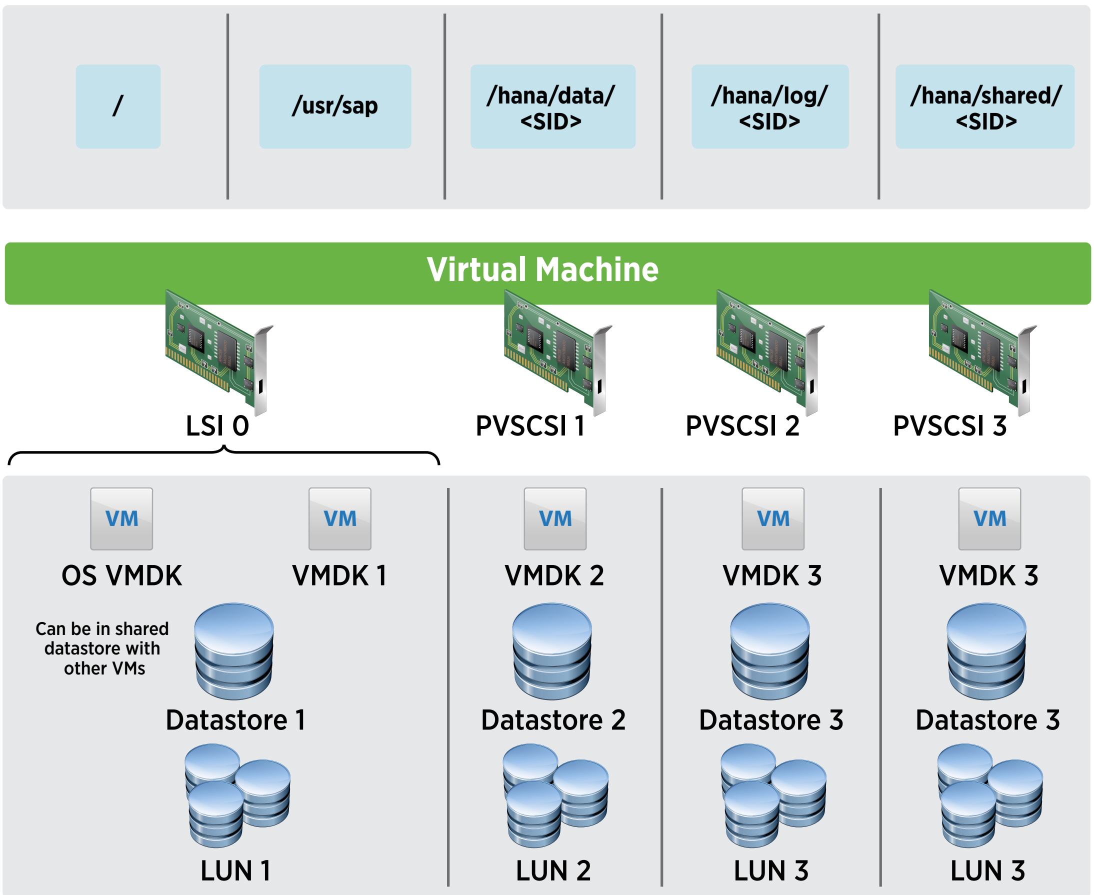

Figure 11. Example Storage Layout for Scale-Up SAP HANA

The storage design implements the following guidelines:

- To maximize performance, spread the virtual disks across all four available virtual SCSI controllers to a VM.
- To maximize I/O performance, spread the database files across multiple LUNs and LUN queues.
- To minimize I/O contention, separate database log and data files into different LUNs.
- These guidelines are similar to best practices followed in physical environments.

This storage design is an example, and some variations are possible:

- The root and `/usr/sap` file system can be the same and can be backed by one virtual disk.
- File systems can be created under the directories `/hana/data/<SID>` and `/hana/log/<SID>`, where all the database data and log files reside, because these can be backed by more virtual disks and data stores.
- Sizing and performance requirements determine the final design. VMware recommends that the sizing be jointly conducted between VMware and database and storage administrators; it should incorporate the best practices of the storage vendor. The storage vendors have guidelines for the storage design of databases, and these should be applied to the virtual environment.

Before we discuss supported storage options for SAP HANA databases running on vSphere, we provide the following important details on how VMs access different types of storage. When a VM communicates with its virtual disk stored on a datastore, it issues SCSI commands. Because datastores can exist on various types of physical storage, the SCSI commands are encapsulated into other forms, which are determined by the protocol the ESXi host uses to connect to a storage device. ESXi supports Fibre Channel, iSCSI, Fibre Channel over Ethernet (FCoE), and NFS protocols. Regardless of the type of storage device a physical host uses, the virtual disk always appears to the VM as a mounted SCSI device. The virtual disk hides the physical storage layer from the VM's OS. This enables OSs to run in a very standardized way without any specific storage requirements, such as SAN HBA drivers or other settings.

ATTRIBUTE	SPECIFICATION
Storage type	Fibre Channel SAN
Number of storage processors	2 (redundant) Pure Storage FlashArray//M50 instances
Number of switches	
Number of ports per host per switch	2 (redundant) 2
LUN size	Variable
Total LUNs	Depends on cluster and database requirements
vSphere VMFS datastores per LUN	1
vSphere VMFS version	6

Table 19. Shared Storage Design Specifications

Shared Storage Physical Design Specifications

This section details the physical design specifications of the shared storage corresponding to the previous section that describes the logical design specifications.

ATTRIBUTE	SPECIFICATION
Vendor and model	Pure Storage FlashArray//M50
Type	Active/Active
ESXi host multipathing policy	Round Robin
Minimum/maximum speed rating of switch ports	16GBps

Table 20. Shared Storage Physical Design Specifications

VMware vSphere Storage I/O Control

VMware vSphere Storage I/O Control is enabled, despite the fact that production SAP databases are designed with dedicated LUNs. This provides cluster-wide storage I/O prioritization and the ability to control the amount of storage I/O that is allocated to VMs during periods of I/O congestion. The shares are set per VM and can be adjusted, based on need, for each VM.

Shared I/O Control Settings Explanation

The following terms are pertinent to this discussion:

Storage I/O Enabled – Storage I/O is enabled per datastore. Navigate to **Configuration** and then to **Properties** to verify that the feature is enabled.

Storage I/O Shares – Storage I/O shares are similar to VMware CPU and memory shares. Shares define the hierarchy of the VMs for distribution of storage I/O resources during periods of I/O congestion. VMs with higher shares have higher throughput and lower latency.

Limit IOPS – By default, the IOPS allowed for a VM are unlimited. By allocating storage I/O resources, the IOPS allowed to a VM are limited. If a machine has multiple disks, the same IOPS value must be set for all the disks that access that VM.

Storage Sizing for SAP S/4HANA

All SAP HANA worker VMs have a database log, data, root, local SAP, and shared SAP volume. The storage capacity sizing calculation of these volumes is based on the overall amount of memory needed by the SAP HANA in-memory database. SAP has defined very strict performance KPIs that must be met when configuring a storage subsystem. This might result in more storage capacity than is needed—even if the disk space is not needed—but that particular number of spindles might be required to provide the needed I/O performance and latency.

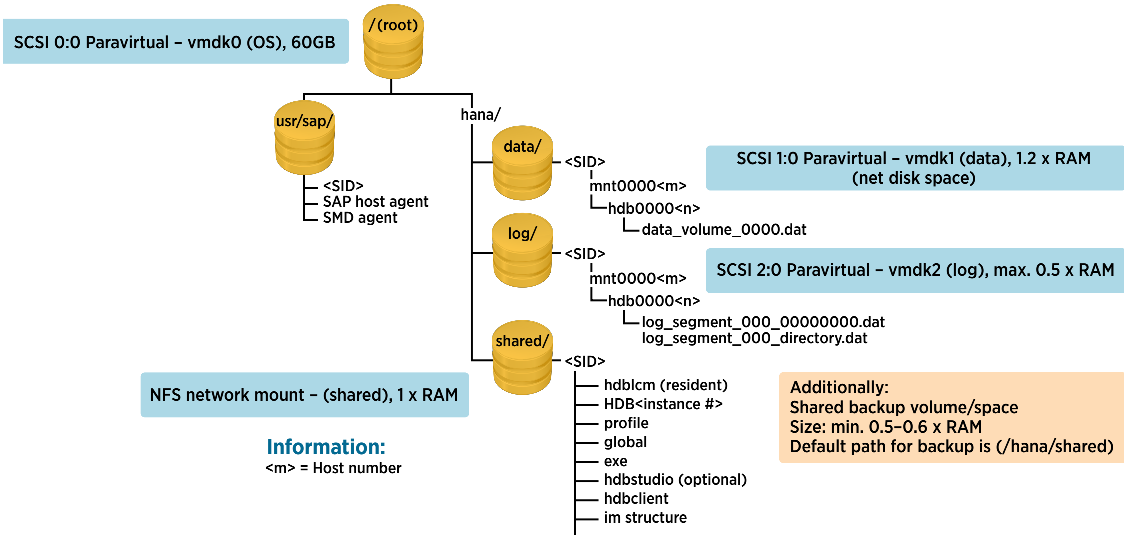

Figure 12. Virtualized SAP HANA Scale-Out Fibre Channel SAN–Based Storage Disk Layout

SAP has published guidelines for calculating and sizing the correct storage requirements for SAP HANA TDI configurations. These guidelines change from time to time, so always download the latest white paper on SAP HANA storage requirements. In addition, always ensure that the storage vendor that offers an SAP HANA TDI storage solution for SAP HANA scale-out configurations has followed the correct space and IP sizing guidelines.

For the most part, the figures for storage volume size that are used in this guide follow the SAP HANA storage requirements published in January 2015. These figures should be used only as an initial sizing guideline that must be aligned to the available SAP sizing reports, Quick Sizer input, or storage vendor calculation. VMware cannot guarantee that the figures and formulas are 100-percent correct.

ROOT volume:	60GB
Data volume:	1.2 x RAM (net disk space)
Redo log:	Systems less than or equal to 512GB = 0.5 x RAM Systems greater than 512GB = 512GB
SAP HANA shared:	1 x RAM
Backup volume:	Minimum 0.5–0.6 x RAM

Storage Sizing by Module Based on SAP Guidelines

MODULE	MEMORY GB	DATA GB	LOG GB	SHARED GB	BACKUP GB
BW/4HA_SRV	650	780	325	650	390
CRM	250	300	125	250	150
MDG	250	300	125	250	150
S/4	800	960	400	800	480
SRM	250	300	125	250	150
TM	250	300	125	250	150

Table 21. Database Component Storage Sizing for SAP S/4HANA

As a default, the backup volume is stored under the `/hana/shared` volume. To store backup data, creation of a dedicated VMDK volume or NFS mount point is recommended.

NOTE: *Disk capacity sizing must be shown as the minimum requirement of the storage configuration. Actual storage sizing might require more disks to fulfill the I/O requirements of SAP HANA. If the I/O needs are fulfilled with a specific SAP HANA TDI configuration, the system must be verified with the SAP HANA HWCCT utility.*

Storage Area Network

The SAP S4/HANA solution storage area network (SAN) is based on the Brocade G620 switch with Generation 6 Fibre Channel and Brocade Fabric Vision technology. With enterprises migrating to high-performance flash storage for their business-critical applications, the G620 switch is better suited for building underlying Fibre Channel SAN to support applications such as SAP S4/HANA.

The G620 switch is backward compatible with previous-generation 16Gbps and 8Gbps connectivity. Although the solution is tested with the Pure Storage FlashArray//M50 with 16Gbps speed, the following features are unique to Generation 6 Fibre Channel and benefit this deployment:

- Provides high scalability in an ultradense, 1U, 64-port switch to support high-density server virtualization, cloud architectures, and flash-based storage environments
- Increases performance for demanding workloads across 32Gbps links and helps application performance reach up to 100 million IOPS
- Provides proactive, nonintrusive, real-time monitoring and alerting of storage I/O health and performance with Brocade IO Insight, the industry's first integrated network sensors
- Enables VM visibility in a storage fabric to monitor and optimize VM performance and identify VM anomalies
- Increases resiliency by automatically discovering and recovering from device or network errors
- Leverages Brocade Fabric Vision technology to simplify administration, quickly resolve problems, increase uptime, and reduce costs

See Appendix A for more details on features used and SAN best practices.

VMware vSAN

VMware vSAN is the VMware software-defined storage solution for hyperconverged infrastructure, a software-driven architecture that delivers tightly integrated compute, networking, and shared storage from a single virtualized x86 server.

VMware vSAN provides enterprise-class scale and performance as well as new capabilities that broaden the applicability of the proven VMware vSAN architecture to business-critical environments. VMware vSAN includes these important features:

- Deduplication and compression – Software-based deduplication and compression optimize all-flash storage capacity, providing as much as 7 times data reduction with minimal CPU and memory overhead.
- Erasure coding – Erasure coding increases usable storage capacity by up to 100 percent while keeping data resiliency unchanged. It is capable of tolerating one or two failures with single-parity or double-parity protection.
- Quality of service (QoS) with IOPS limit – Policy-driven QoS limits and monitors IOPS consumed by specific VMs, eliminating noisy neighbor issues and managing performance SLAs.
- Software checksum – End-to-end data checksum detects and resolves silent errors to ensure data integrity; this feature is policy driven.
- Client cache – Client cache leverages local dynamic random access memory (DRAM) to VMs to accelerate read performance. The amount of memory allocated is 0.4 percent of total host memory up to 1GB per host to local VMs.

VMware vSAN provides the following advantages:

- VMware Hyper-Converged Software–powered all-flash solutions available at up to 50 percent less than the costs of competing hybrid solutions in the market.
- Increased storage utilization by as much as 10 times through new data efficiency features including deduplication and compression as well as erasure coding.
- Future-proof IT environments with a single platform supporting business-critical applications, OpenStack, and containers with up to 100K IOPS per node at submillisecond latencies.

All-Flash VMware vSAN Architecture

All-flash VMware vSAN delivers extremely high IOPS with predictable low latencies. Two different grades of flash devices are commonly used in an all-flash VMware vSAN configuration: lower-capacity and higher-endurance devices for the cache layer; more cost-effective, higher-capacity, and lower-endurance devices for the capacity layer. Writes are performed at the cache layer and then destaged to the capacity layer as needed. This helps extend the usable life of lower endurance flash devices in the capacity layer.

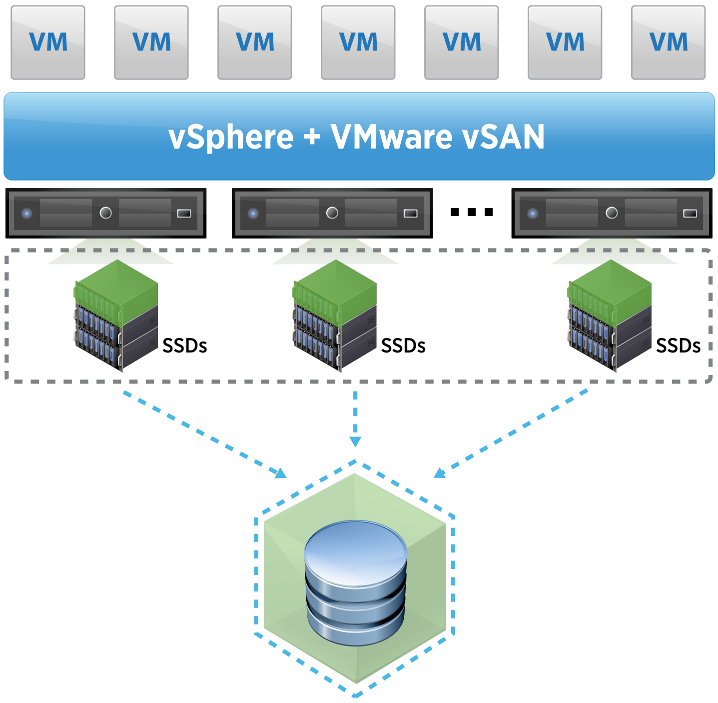

Figure 13. VMware vSAN All-Flash Datastore

All-Flash VMware vSAN Configuration for SAP S/4HANA Validation

VMware vSAN was used for the management cluster and as an alternate for the TDI storage for SAP HANA. All storage components used were certified for VMware vSAN and provided by Western Digital.

VMware vSAN features the following components:

- Two disk groups per server
- Each disk group contains
  - One NVMe 1.5TB drive for caching
  - Two 3.5TB SSD drives for capacity

The total capacity of all-flash VMware vSAN is 52TB.

SAP HANA Network Design Considerations

The network components were designed based on the foundational VVD concepts. VMware NSX is available as part of this design, to provide networking and security for application workloads. Details on the VVD network design and components can be found in *Architecture and Design* in VMware Validated Designs Documentation. VMware NSX is the foundation for networking in the Software-Defined Data Center (SDDC), which is leveraged for SAP S/4HANA.

Network Sizing for SAP S/4HANA Virtualized Systems

Standard VMware network sizing guidelines apply. The total number of supported 1Gb and 10Gb network adapter ports is limited by the vSphere version being used. vSphere 5.5 supports eight 10Gb and four 1Gb network ports; vSphere 6.0 supports sixteen 10Gb and four 1Gb network ports.

Table 22 summarizes network requirements and provides an overview of the dedicated networks required by a virtualized SAP HANA scale-out configuration.

NOTE: *A dedicated 10GB network port is needed per SAP HANA VM; otherwise, the SAP HANA network KPI might not be achieved.*

	MANAGEMENT, HIGH-AVAILABILITY, AND CLIENT NETWORK	SAP HANA SHARED (NFS MOUNT) NETWORK	vSPHERE vMOTION NETWORK	SAP HANA INTERNODE NETWORK	SAP HANA NFS STORAGE NODES 2–8*
Network label	External	HANA-NFS-shared	vMotion	HANA-internal	SAP HANA-NFS-storage 2–8
Bandwidth	1GbE	1GbE	10GbE	10GbE	10GbE
Network adapter cards	1		1	1	1
Network adapter ports**	2		2	2	2
NAS storage	ü				
FC storage	ü				N/A

* Minimum one dedicated 10GbE network adapter per SAP HANA VM; possible maximum of an additional 6 network adapters with vSphere 5.5 and 14 with vSphere 6.0 SAP HANA NFS storage networks.

** With network port redundancy.

Table 22. SAP HANA Scale-Out vSphere Network Sizing

Mission-Critical Architectures: Completing the SDDC Picture with VMware NSX

Figure 14 shows the *SAP HANA Network Requirements Guide*, which expresses the basis of why networks should be virtualized. Currently, the components of an SAP HANA system communicate via various network channels. SAP recommends that a well-defined network topology be instituted to control and limit access into only the required channels in order to apply the appropriate security measures when necessary.

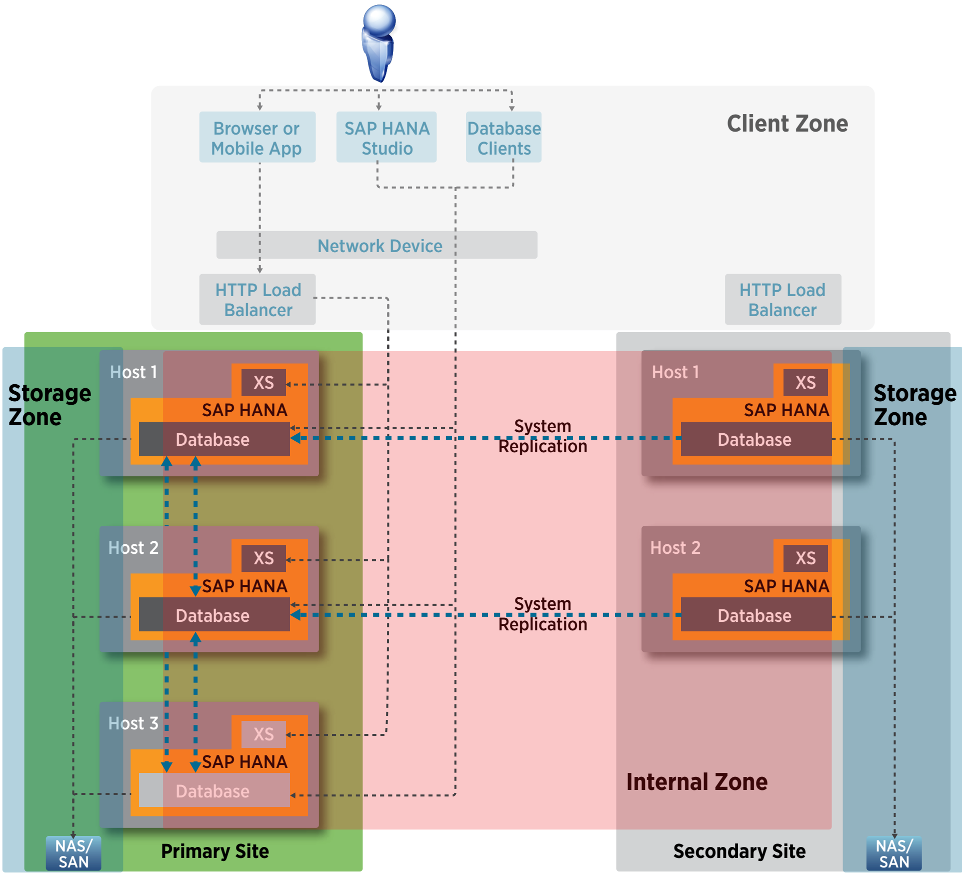

Figure 14. SAP HANA Network Zones

The following describes the various SAP HANA network zones:

Client zone – Access is granted to various clients, such as SQL clients on SAP application servers. There are also browser applications using HTTPS to access the SAP HANA server, as well as other data sources such as BI, that require a network communication channel to the SAP HANA database.

Internal zone – This handles communication between hosts in a distributed SAP HANA system as well as SAP HANA system replication communication between two SAP HANA sites.

Storage zone – Although SAP HANA holds its data in memory, it is a fully ACID-compliant relational database, so data also persists in storage. SAP HANA data must be accessed via a network to provide protection against power failures or if a host becomes unavailable.

The dynamic tiering zone, which is not shown here, also has its own set of unique networking requirements.

In the physical world, this can become quite complex because each communication zone requires special attention with regard to setup, security, and performance requirements. VMware NSX highly enables the creation of a virtual network architecture based on the *SAP HANA Network Requirements Guide*.

SAP HANA Client Zone VMware NSX Design

VMware NSX greatly facilitates the creation and deployment of network services. The actual abstraction of these services, however, is a collaborative effort. To create an optimized virtual architecture, at a minimum this involves the network operations team, the storage team, virtual infrastructure administrators, application owners, and database administrators. This should not be an isolated task.

From a desktop environment using VMware NSX, the *SAP HANA Network Requirements Guide* for tailored data center integration deployments can be decomposed and translated into a virtual network design. Regarding the client zone, the application server network and the network for other SAP HANA clients such as the SAP HANA studio, business intelligence clients, and so on, can be either on the same network or on separate networks. For scalability, management, and security purposes, using separate networks via microsegmentation is recommended. Microsegmentation enables customers to secure, isolate, and characterize network traffic from workload to workload. We also recommend a distributed routing scheme rather than a centralized routing scheme, to optimize for both north–south traffic and east–west traffic.

The east–west dynamic routing capabilities of the DLR are key here because of how SAP HANA load-balances its client connections. SAP uses techniques such as statement routing, connection selection, and command splitting to direct queries to the proper nodes. These routing techniques are based on the type of query and on which node or nodes the data lives on.

As the database grows, the data distribution and types of queries can change. DLRs enable the central management and optimization of east–west traffic by proactively reacting to database growth and traffic patterns.

For external communication with SAP HANA servers that are initiated by a Web browser or a mobile application, a VMware NSX edge services perimeter gateway manages and optimizes north–south network traffic and leverages its multifunction capabilities as a firewall, load balancer, and virtual private network (VPN) device.

And because these external connections can have vastly different security requirements, customers can use VMware NSX to associate firewall rules at the router or at the VM level to achieve greater granularity.

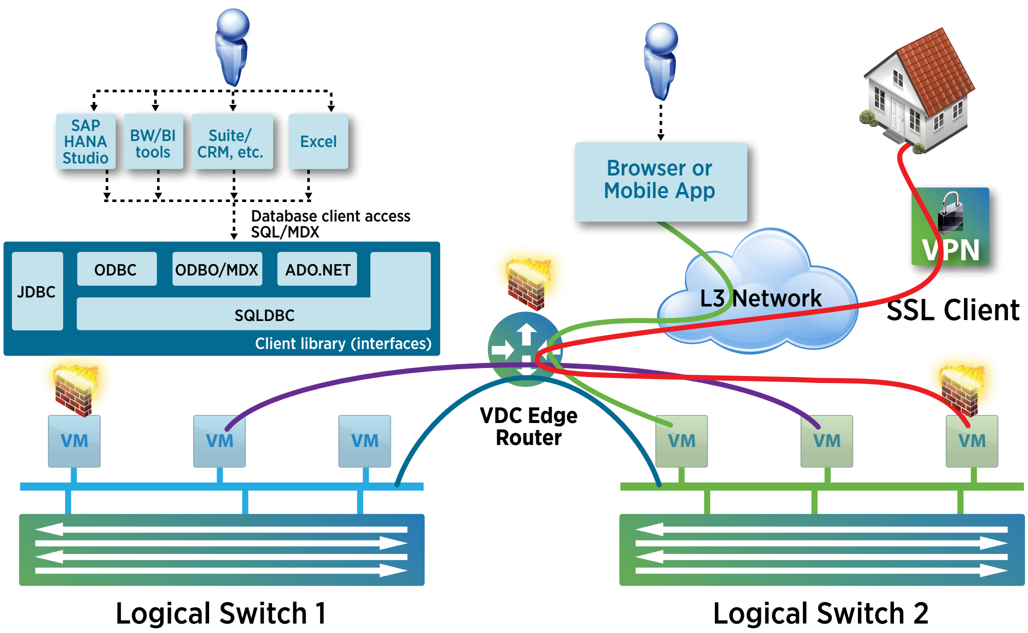

Figure 15. SAP HANA Client Zone VMware NSX Network Design

SAP HANA Internal Zone VMware NSX Design

This design also includes the transport zone and transit switch. Microsegmentation creates various tiers that are fully isolated and secured zones. The transport zone enables central management of and communication with these tiers, as well as communication with one or more ESXi clusters, spanning multiple logical switches through the transport zone. The transport zone is also connected to VMware NSX and other VMware components such as the control cluster, VMware NSX Manager™, and VMware vCenter Server®.

Regarding the internal zone, SAP requires for performance and security reasons that the internode network be configured with its own network adapter. The internal network requires separate network interface cards that are exclusively configured as part of a private network, using separate IP addresses and ports.

Depending on the deployment needs of the customer, SAP HANA tables can be distributed or replicated among multiple nodes or assigned to a single node. In addition, it might be necessary to join and aggregate tables, which can produce different network traffic patterns. As with the client zone, the VMware NSX DLR can be used to dynamically manage, load-balance, and optimize network traffic, based on growth, data distribution, and queries in unique SAP HANA clusters.

SAP S4/HANA Security

Standard vSphere security best practices and hardening practices must be applied to secure the infrastructure. Follow guidelines for access control and security policies provided in the VMware Validated Designs Documentation. We address only the application-specific security for SAP S/4HANA in this section.

Microsegmentation for Securing SAP S/4HANA Workloads

SAP microsegmentation enables flexible security policies that align to the multitier architecture of an individual SAP system—presentation, application and database tiers—and also to the landscape of the SAP environment, separating production from nonproduction. Figure 16 shows an SAP microsegmentation example based on the NetWeaver ABAP stack with a backend database. The following are the various tiers and components of the SAP architecture:

- Client tier – In this example, the SAP client “SAPGUI” accesses the application tier. Customer environments can include browser-based access, load balancers, and a Web tier.
- Application or Central Services tier – Application servers based on the NetWeaver ABAP stack. SAP Central Services handles SAP locking services, messaging between the application servers, and an NFS share required by all the application servers.
- Database tier – Services are database dependent.
- The components are isolated into their own respective VMware NSX security groups. Although other classifications are possible, a VMware NSX security group in this example is a definition in VMware NSX and corresponds to a logical grouping of VMs within which there is free communication flow. Communication flow in and out of a security group, and from and to another group, depends on the firewall rules.

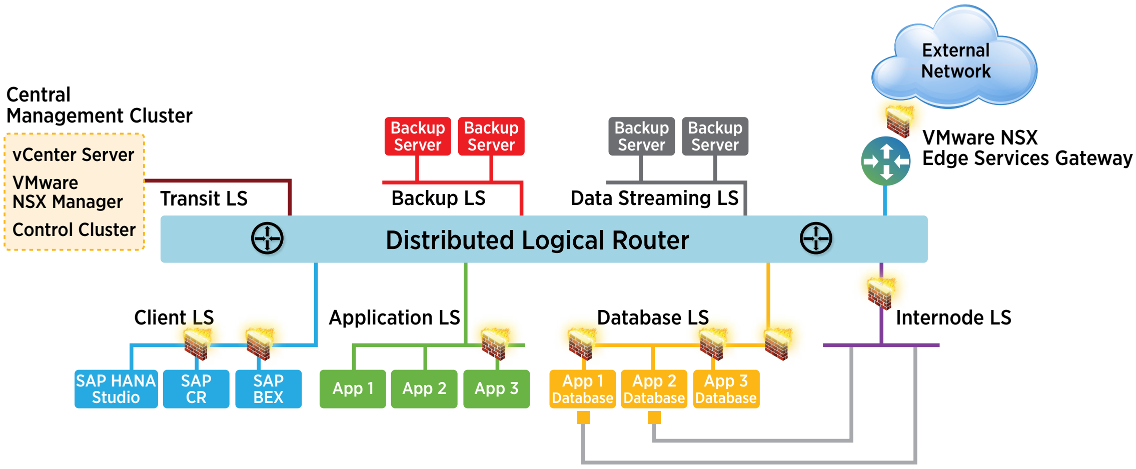

Figure 16. Segmentation of SAP S4/HANA with VMware NSX

Security policies shown in Figure 16 provide the following controls:

- Controlled communication path limited to specific services and protocols between tiers
- External access permitted to the application tier via the SAP presentation service
- Access between application and database VMs only via specific database services
- SAP services ports that vary depending on the "instance number" assigned to the application servers and SAP Central Services. Some values are shown here.

SAP S/4HANA Monitoring and Management

Nearly every component of the IT stack contributes to application performance, which can make it challenging to identify the cause of issues when they arise. For many organizations, a lack of visibility can lead to mean-time-to-innocence hunts that waste time and create alert storms that drain the productivity of business teams.

With a complex application such as SAP S/4HANA, performance issues can be even more difficult to specify because the application requires resources from the virtual environment, the network, and databases. However, integrating monitoring into a single console—such as VMware vRealize Operations Manager™—can provide visibility into SAP workloads and other IT relationships to impact performance.

The Management Pack for SAP S/4HANA enhances vRealize Operations Manager by adding three dashboards that include the following features:

- SAP overview dashboard – See heat maps depicting the overall health of SAP landscapes, host systems, and popular instance types such as Java, ABAP, and dual stack.
- SAP relationships – Access relationships, badges, health trees, and metrics for a particular SAP resource.
- SAP host overview – View top alerts, heat maps, and relationships to VMware VMs, parent VM KPIs, and CPU and memory metrics for SAP hosts.
- SAP landscape overview – See top alerts, heat maps, CPU and memory usage metrics, and services utilization for an SAP landscape.
- SAP HANA environment – See details about SAP HANA resources, including alerts and key metrics.
- SAP HANA host details – Access detailed information about SAP HANA hosts, including workload, capacity remaining, statements summary, and connections.
- SAP HANA overview – Select a system to see properties, workload, capacity remaining, request summary, and topology view of relationships.

Reports and Dashboards

Performance issues, particularly across a wide-sprawling application such as SAP, can be challenging to ascertain, especially with the growing complexity of the IT stack in today's environment. Having the ability to clearly see where issues develop can be a game changer in ensuring availability and a better experience for end users.

Reporting and dashboards can extend visibility into key areas of the SAP environment and can identify issues as soon as they arise rather than when they are wreaking havoc on the system.

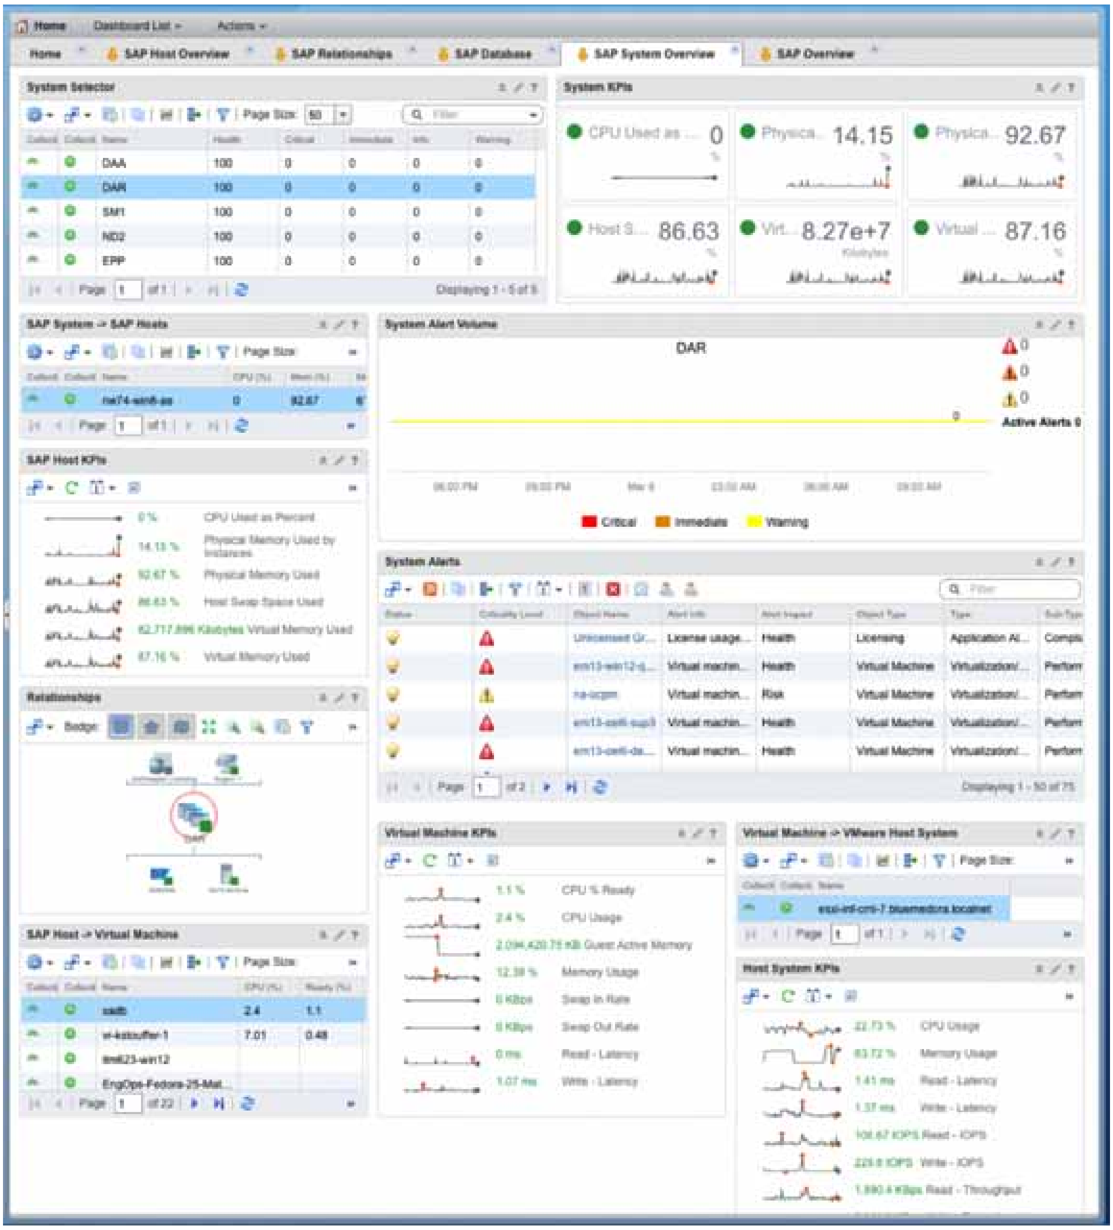

Figure 17. Example of the SAP System Overview Dashboard in vRealize Operations

Capacity Planning

Predictive analytics in vRealize Operations can provide the insight required to optimize the capacity and health of an SAP environment.

Analysis badges offer a visual indication of the current condition of the virtual environment. Updates in real time quickly help determine whether capacity issues are being caused by various indicators such as workload, capacity, or stress on the application.

Capacity definitions help extend that visibility into specific areas of an SAP application and enable reporting on key elements that help determine trends and how to improve application performance.

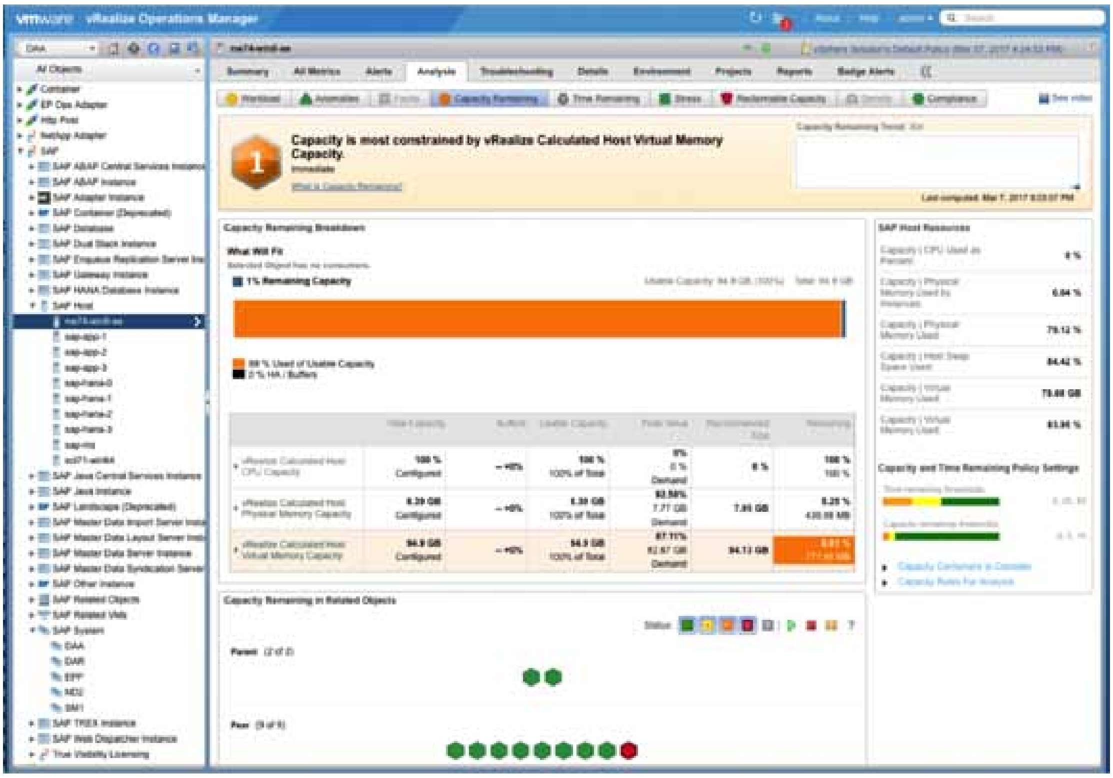

Figure 18. Example of Analysis Badge for SAP S/4HANA in vRealize Operations

Alerts and Notifications

Given the complexity of the SDDC, it can be challenging to track issues as they occur. This often means that businesses are unaware of an issue until there is a dip in performance—or worse, an impact on end users.

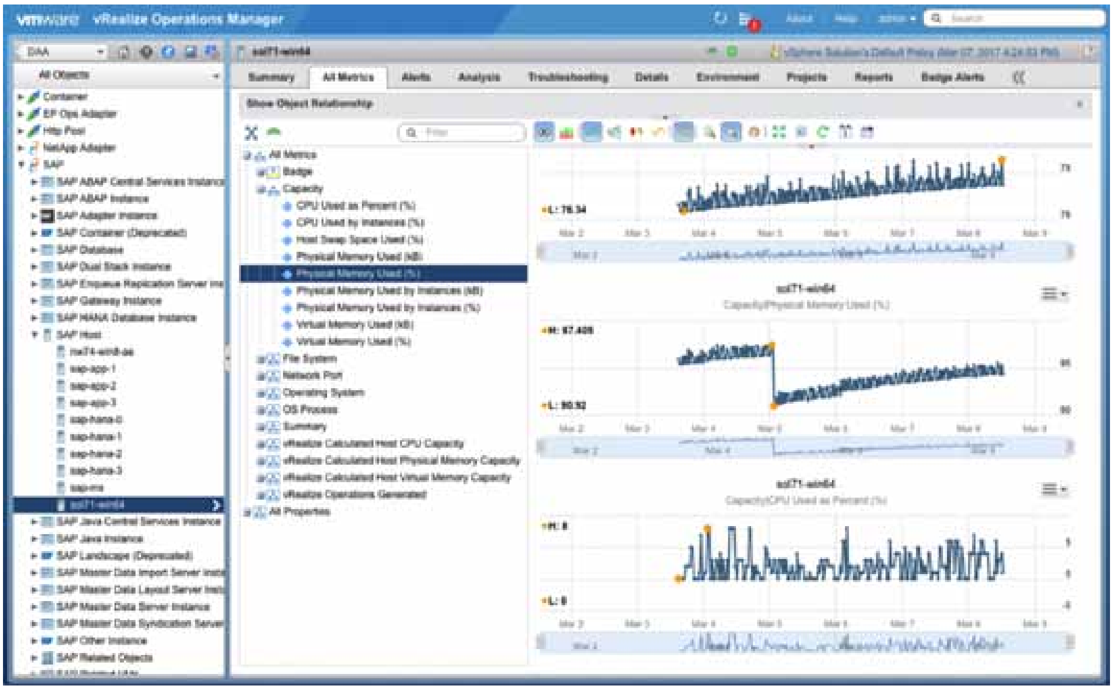

Figure 19. Simplified Troubleshooting in an Easy-to-Read Dashboard for SAP Resources and Objects

vRealize Operations helps define thresholds for providing detailed alerts and notifications when issues arise. This simplifies troubleshooting by offering insight into the root cause of issues and enables concentrated focus on the actual problem, all within a single console in vRealize Operations.

The Blue Medora SAP Management Pack for vRealize Operations extends the capabilities of vRealize Operations, providing increased access to reports and dashboards, analysis and capacity badges, and key alerts and notifications. This drives better performance across SAP workloads.

SAP S/4HANA Automation

VMware Adapter for SAP Landscape Management Solution Overview

VMware Adapter for SAP Landscape Management, part of the VMware private cloud solution for SAP is a virtual appliance that integrates SAP Landscape Management with VMware management software—that is, vCenter Server and vRealize Automation. This delivers unique automation capabilities that radically simplify how SAP basis administrators and end users provision and manage SAP landscapes. The appliance accepts application calls from SAP Landscape Management, then uses vRealize Automation or VMware vRealize Orchestrator™ workflows to execute commands to vCenter Server for operations related to VMware products, such as starting and stopping a VM. Furthermore, IT administrators can now leverage SA-API to automatically provision SAP systems from templates with vRealize Automation in conjunction with SAP Landscape Management.

Key Benefits

The deployment of new SAP systems can take days or even weeks before systems are ready for use. Customers have long used various cloning methods to speed up the deployment process. However, these processes are complex and labor intensive. The VMware Adapter for SAP Landscape Management – Connector for vRealize Automation greatly simplifies the deployment process by utilizing vSphere cloning and SAP Landscape Management to create new SAP workloads in an automated and repeatable form and from proven templates.

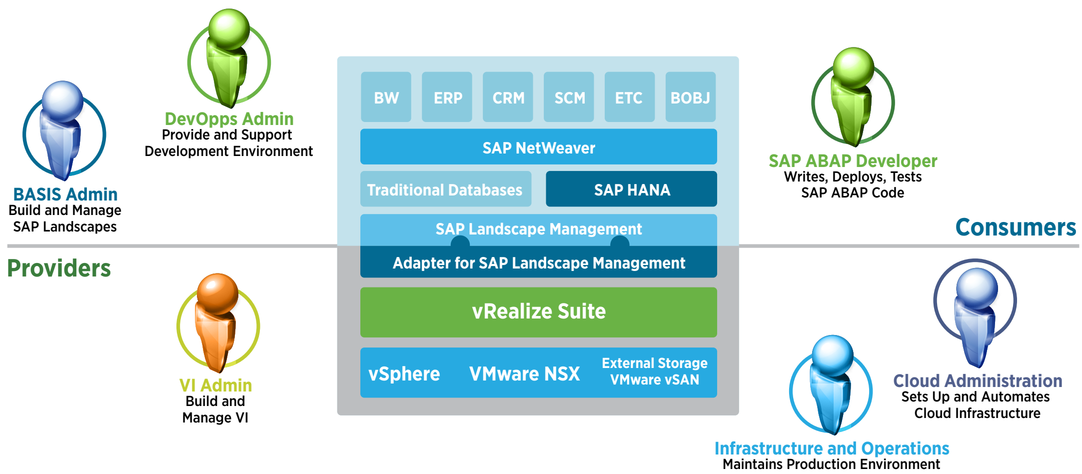

Figure 20. VMware Adapter for SAP Landscape Management

NOTE: *Refer to Appendix C for additional details on this solution.*

SAP S/4HANA Backup and Recovery

SAP HANA on vSphere Backup and Restore

After reviewing how SAP HANA data persistence works, ensure that the savepoints and logs SAP HANA uses to persist its data are backed up and stored securely to have them available to recover from data loss via restoring this data.

An SAP HANA database backup consists of data backup—that is, snapshots—and transaction log backups. The data backup can be scheduled or started manually within SAP HANA Studio, DBA Cockpit, or via SQL commands. Logs are saved automatically in an asynchronous way whenever a log segment is full or the timeout for log backup has elapsed.

Transaction redo logs are used to record any changes made to the database. In the case of failure, the most recent consistent state of the database can be restored by replaying the changes recorded in the log, redoing completed transactions, and rolling back incomplete ones. Savepoints are created and described as periodic, representing the data stored in the SAP HANA database. They are coordinated across all processes—called SAP HANA services—and instances of the database to ensure transaction consistency. New savepoints normally overwrite older savepoints, but it is possible to freeze a savepoint for future use; this is called a snapshot. Snapshots can be replicated in the form of full data backups, which can be used to restore a database to a specific point in time. Snapshots can also be used to create a database copy for SAP HANA test-and-development systems. Periodic backup of the snapshots and logs ensures the ability to recover from fatal storage faults with minimal loss of data.

Backup and recovery of virtualized SAP HANA systems is similar to that of any physically deployed SAP HANA system. The backup of the necessary files can be performed as a normal file system backup to an external NFS server. When a backint-compatible backup solution is used, the backup can be performed directly via the backint interface to a backup server and then to the final backup device. Using storage built-in snapshot functionality to create backups is another option. This method is the fastest way to create a backup. Some vendors work on backups that are vSphere snapshot compatible. Storage systems that today support the new VVOL standard already enable snapshotting VMs with the full awareness of the virtual disks belonging to a VM. Figure 21 provides an overview of the SAP HANA backup and recovery methods.

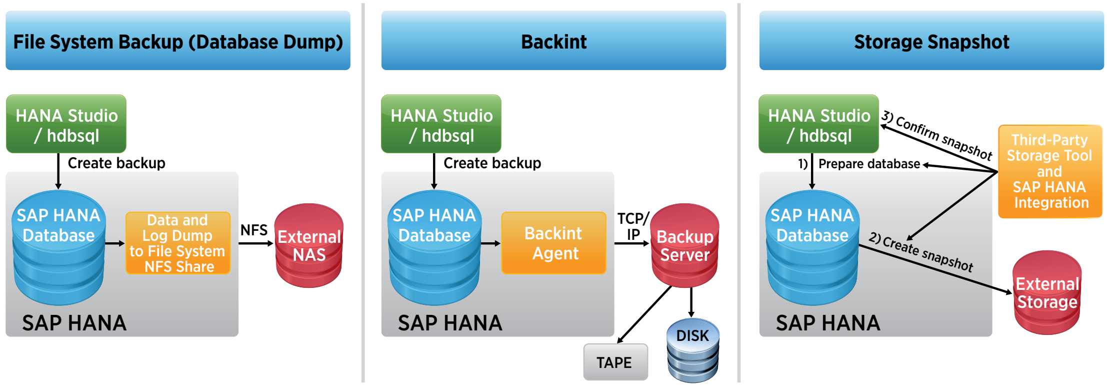

Figure 21. Overview of SAP HANA Backup and Recovery Methods

Recovery works the same way and is differentiated only by how the backup files must be handled. Recovery via backing, for example, can be done in an automated manner by selecting the right recovery set inside SAP HANA Studio. If a storage vendor has implemented snapshots-based backup into SAP HANA, the recovery can be similarly easy.

NOTE: *Backup and recovery of virtual SAP HANA systems is done similarly to physical SAP HANA systems. Storage snapshot solutions might not be available for VMDK-based systems and might require raw device mappings. Some backup vendors work on leveraging vSphere snapshot and cloning or VVOL implementations.*

Disaster Recovery Solutions with vSphere for SAP S/4HANA

We have already discussed recovery solutions for local failures—component or OS failures, for example. In addition to these solutions for local failures, SAP HANA offers disaster recovery solutions supported by vSphere that replicate the data from the primary data center to VMs in a secondary data center. SAP HANA system replication provides a very robust solution to replicate the SAP HANA database content to a secondary disaster recovery site; this storage-based system replication can be used as well.

When using SAP HANA System Replication, the same number of SAP HANA VMs must exist at the disaster recovery site. These VMs must be configured and installed similarly to a natively running SAP HANA system with System Replication enabled.

SAP HANA System Replication provides various modes for system replication:

- Synchronous
- Synchronous in-memory
- Asynchronous

Depending on requirements, the disaster recovery VMs can consume higher or lower amounts of resources on the disaster recovery vSphere cluster. For instance, the synchronous in-memory mode consumes the same amount of RAM as with the primary systems. This mode is required only if the customer requests the shortest recovery time. In most customer scenarios, using synchronous data replication should be sufficient. SAP states that by replicating only the data, about 10 percent of the system resources are required enabling up to 90 percent of the resources to continue to be used by other systems such as test or QA systems.

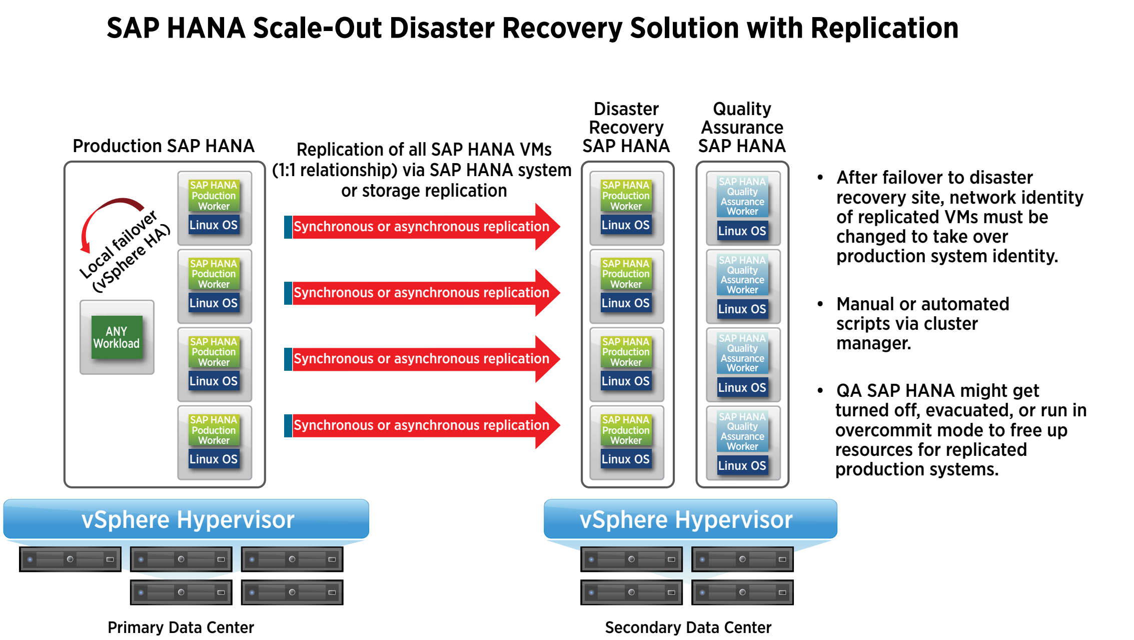

Figure 22. SAP HANA Scale-Out Solution with Replication

In this scenario, resource overcommitment is used to enable the deployment of such an environment. By using resource pools and resource shares, required resources can be provided to the disaster recovery SAP HANA scale-out VMs. The deployed system, with fewer resource shares, experiences performance degradation after the disaster recovery systems are used following a site failover. Evacuate these VMs to other available vSphere systems to free up all resources for the disaster recovery SAP HANA VMs. This is another option, as opposed to running the two systems in parallel—with resource limitations—on the same platform.

System replication via storage—or the SAP HANA replication solution—requires additional steps after a site failover has taken place, to switch the network identity (IP redirect) of the replicated systems from the disaster recovery configuration to the production configuration. This can be done manually or via automated tools such as HP ServiceGuard, SUSE cluster connector, SAP Landscape Virtualization Management (LVM), or other cluster managers. The configuration of such a solution in a virtualized environment is similar to that of natively running systems. Contact your storage vendor to discuss a cluster manager solution supported by its storage solution.

VMware Site Recovery Manager

VMware Site Recovery Manager™ can help reduce the complexity of a system replication disaster recovery solution by automating the complex disaster recovery steps on any level.

Site Recovery Manager is designed for disaster recovery of a complete site or a data center failure. It supports both unidirectional and bidirectional failover. It also supports “shared recovery site,” enabling organizations to fail over multiple protected sites into a single, shared recovery site. This site can, for instance, also be a cloud data center provided by VMware vCloud® Air™, the VMware cloud service offering.

The following key elements compose a Site Recovery Manager deployment for SAP:

- Site Recovery Manager – Designed for virtual-to-virtual disaster recovery. Site Recovery Manager requires a vCenter Server management server at each site. These two vCenter Server instances are independent, each managing its own site. Site Recovery Manager informs them of the VMs they must recover if a disaster occurs.
- Site Recovery Manager manages, updates, and executes disaster recovery plans. It is managed via a vCenter Server plug-in.
- Site Recovery Manager relies on storage vendors’ array-based replication Fibre Channel or NFS storage that supports block-level replication of SAP HANA data and log files to the disaster recovery site. Site Recovery Manager communicates with the replication process via storage replication adapters offered by the storage vendor and that have been certified for Site Recovery Manager.
- vSphere Replication has no such restrictions on use of storage type or adapters. It can be used for VMs that are either static or are not performance critical, such as infrastructure services or SAP application servers with a recovery point objective (RPO) of 15 minutes or longer.

Figure 23 shows an example SAP landscape protected by Site Recovery Manager and storage. The VMs running on the primary site contain all needed infrastructure and SAP components such as LDAP, SAP HANA database, and SAP application servers, as in an SAP Business Suite implementation. The VMs can be replicated, depending on the RPO, via vSphere, SAP HANA, or storage replication. vSphere replication can be used with VMs that tolerate an RPO time of 15 minutes or longer.

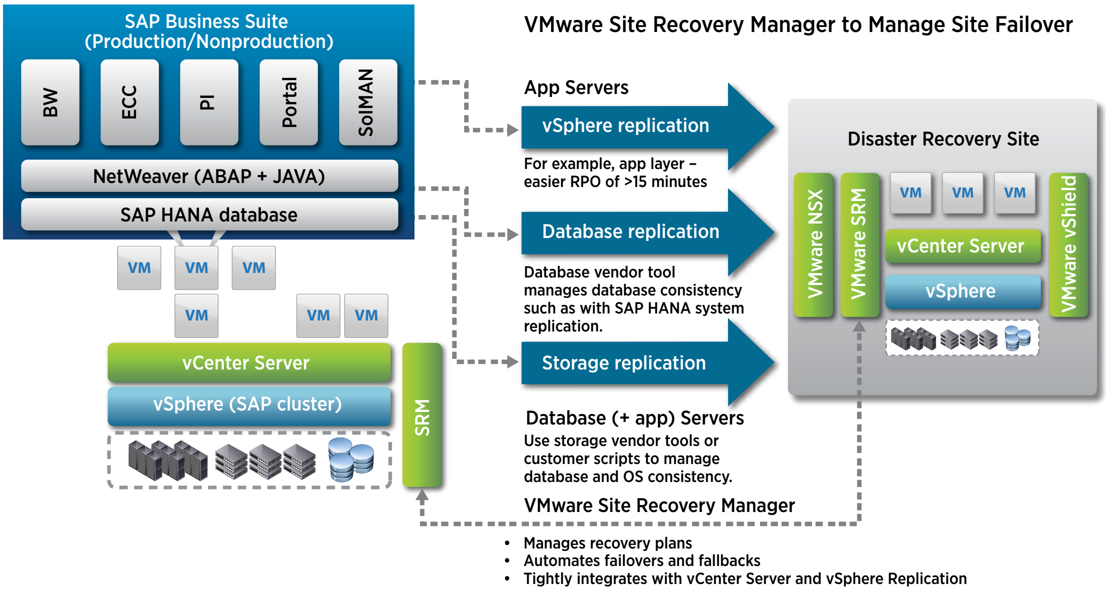

Figure 23. SAP Landscape Protected by VMware Site Recovery Manager

The Site Recovery Manager datasheet details the benefits of using Site Recovery Manager for managing the disaster recovery process for SAP landscapes. Here is a summary

- Reduce the cost of disaster recovery by up to 50 percent.
- Application-agnostic protection eliminates the need for application-specific point solutions.
- Support for vSphere Replication and array-based replication offers choices and options for synchronous replication with zero data loss.
- Centralized management of recovery plans directly from VMware vSphere Web Client replaces manual runbooks.
- Self-service, policy-based provisioning via vRealize Automation automates protection.
- Frequent, nondisruptive testing of recovery plans ensures highly predictable recovery objectives.
- Automated orchestration of site failover and failback with a single click reliably reduces RTO.
- Planned migration workflows enable disaster avoidance and data center mobility.

SAP S/4HANA Design Validation

Validation of an SAP design is often difficult because of the absence of publicly available validation and performance tools. This design utilizes best practices derived from vendor testing conducted in SAP labs. The SAP HANA database tier is critical to the infrastructure and must be validated. So as part of this SAP S/4HANA VVD solution, some SAP standard validation tools were used to exercise the designed infrastructure.

As part of the application workload guidance process, the design was created to depict an example customer environment based on common business requirements. Even though the design is based on best practices developed by VMware, it is a good practice to conduct performance validation for the solution.

Validation was performed for the SAP HANA database tier for the SAP S/4HANA modules that were designed earlier. The database tier was deployed in its own dedicated cluster on large memory hosts with 44 cores and 1TB RAM each. The VMs were deployed on SLES 12 for SAP with memory and CPU sizes from the infrastructure design. All VMware best practices were adhered to in the deployment of the various SAP S/4HANA deployments.

Validation Tests

Standardized OLTP and OLAP tests for SAP S/4 HANA were used for the validation. The SAP HANA systems in the design had 8, 16, or 32 vCPUs each. The validation tests were run across these three sizes of VMs and were compared for performance and scalability. The storage used for the SAP HANA deployments was initially SAP HANA TDI storage hosted on the Pure Storage FlashArray//M50. The VMs then were migrated via vSphere vMotion to the all-flash VMware vSAN configuration. The tests were then repeated.

OLTP Tests

The OLTP tests included inserts, updates, deletes, joins, and join–select operations. These tests were run for the various VM sizes.

Results for Tests with Storage on Pure Storage FlashArray//M50

OLTP tests were run multiple times, and the average runtimes were calculated and tabulated as shown in Table 23:

	8 vCPU	16 vCPU	32 vCPU
Insert	13.7115	8.318	4.912
Delete	13.131	7.802	4.4335
Update	3.2555	1.845	1.091
Select	222.2555	110.454	63.018
Join–Select	6.1375	3.152	1.983

Table 23. OLTP Time Comparison for Various Operations on Different SAP HANA VM Sizes

The time in seconds shows the duration required to run these tests. Lower times are better because they imply that the test ran more quickly.

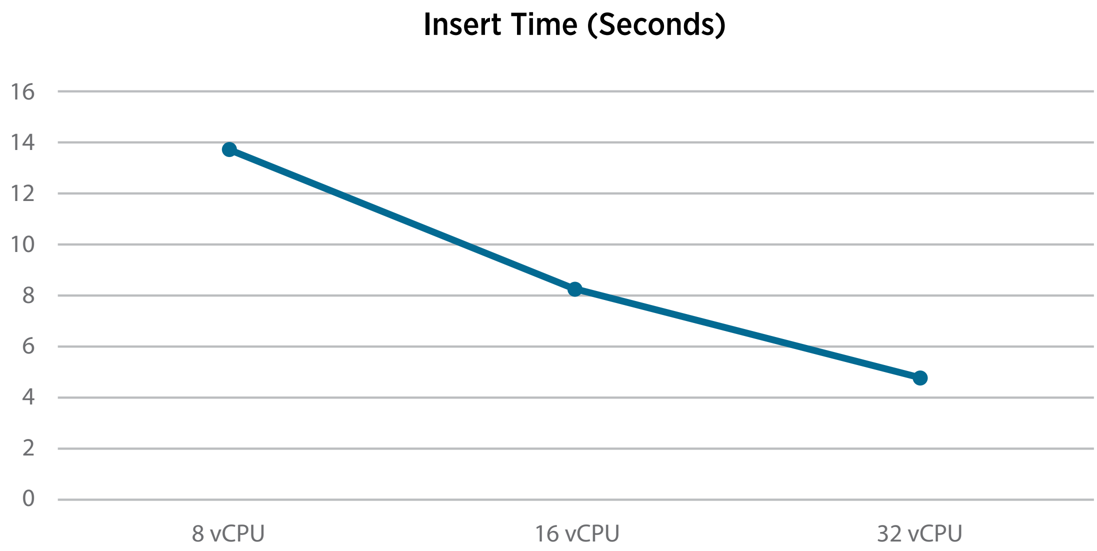

Figure 24. Insert Time Comparison for Different SAP HANA VM Sizes

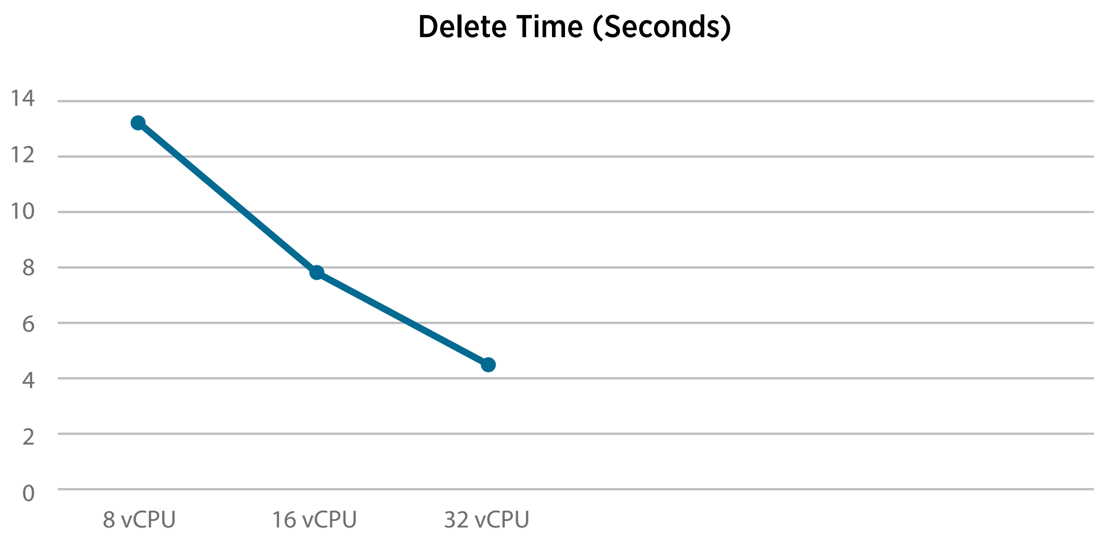

Figure 25. Delete Time Comparison for Different SAP HANA VM Sizes

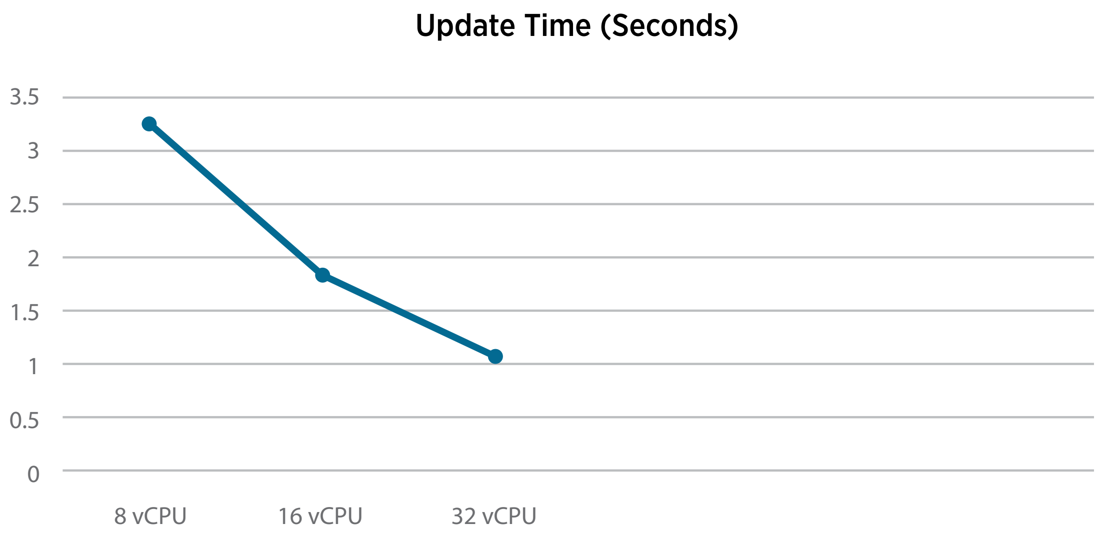

Figure 26. Update Time Comparison for Different SAP HANA VM Sizes

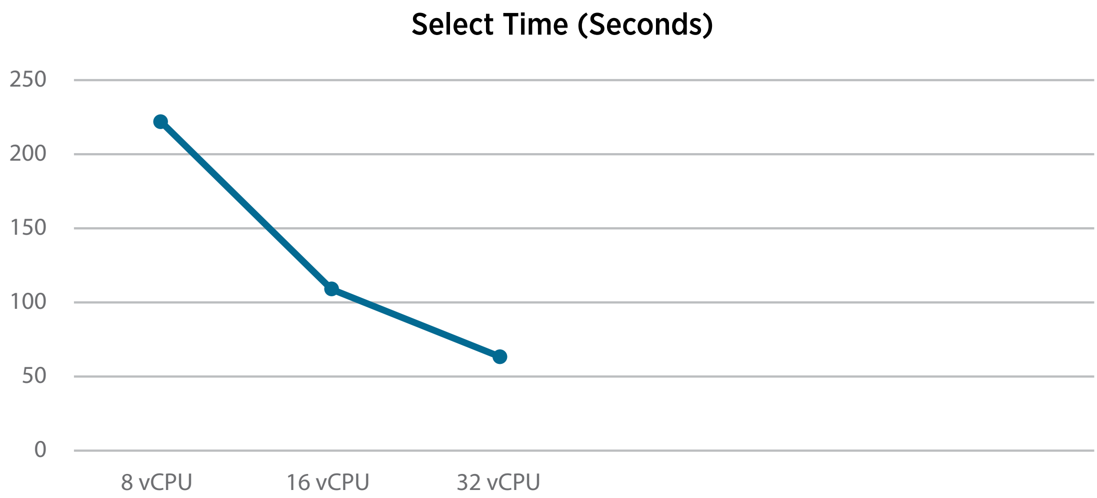

Figure 27. Select Time Comparison for Different SAP HANA VM Sizes

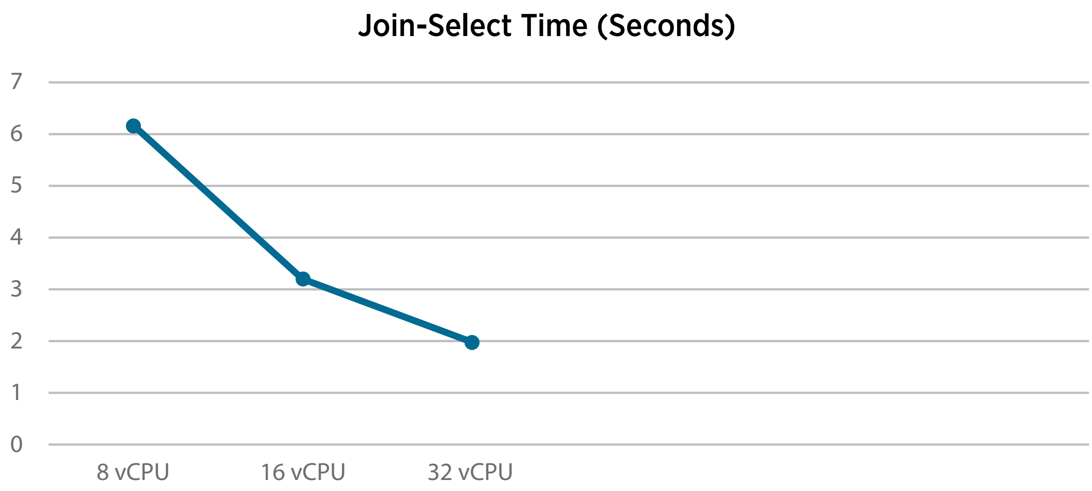

Figure 28. Join-Select Time Comparison for Different SAP HANA VM Sizes

Results from OLAP Type Tests on Pure Storage FlashArray//M50

Two different types of OLAP tests were performed on the SAP HANA database servers: VDM and TPC-DS. The results of these tests are shown in Table 24.

VDM

The performance for this test is measured in queries per second (QPS). There are two different types of queries, labeled QPS1 and QPS2. The QPS for the two types of queries are compared in Table 24 for the three sizes of SAP HANA database VMs.

vCPU	QPS1	QPS2
8	3.93	45.17
16	7.69	80.18
32	11.41	135.52

Table 24. QPS Time Comparison Across Different SAP HANA VM Sizes

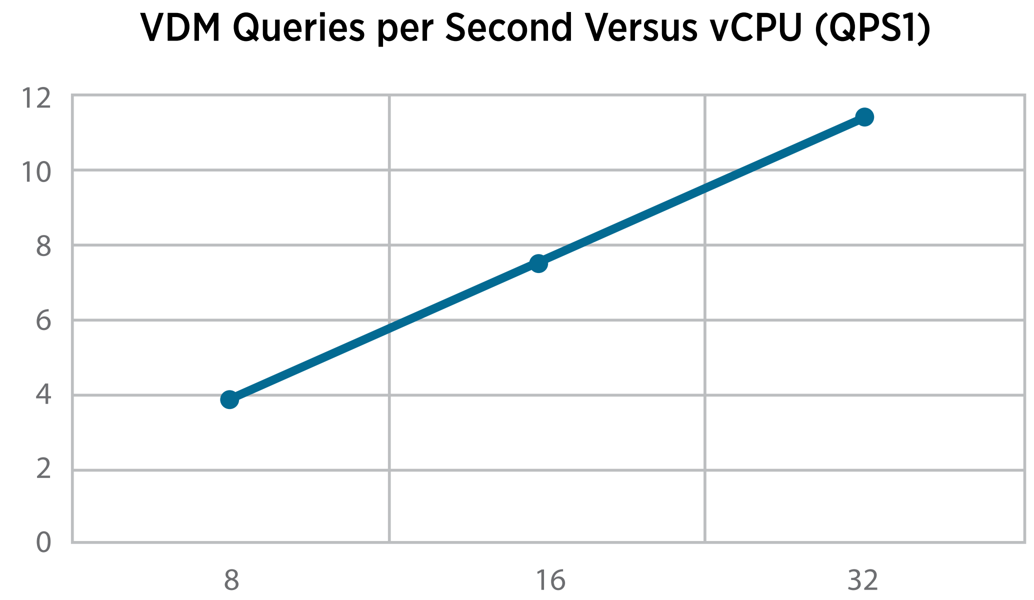

Figure 29. Query Time Comparison for Different SAP HANA VM Sizes for Query Type 1

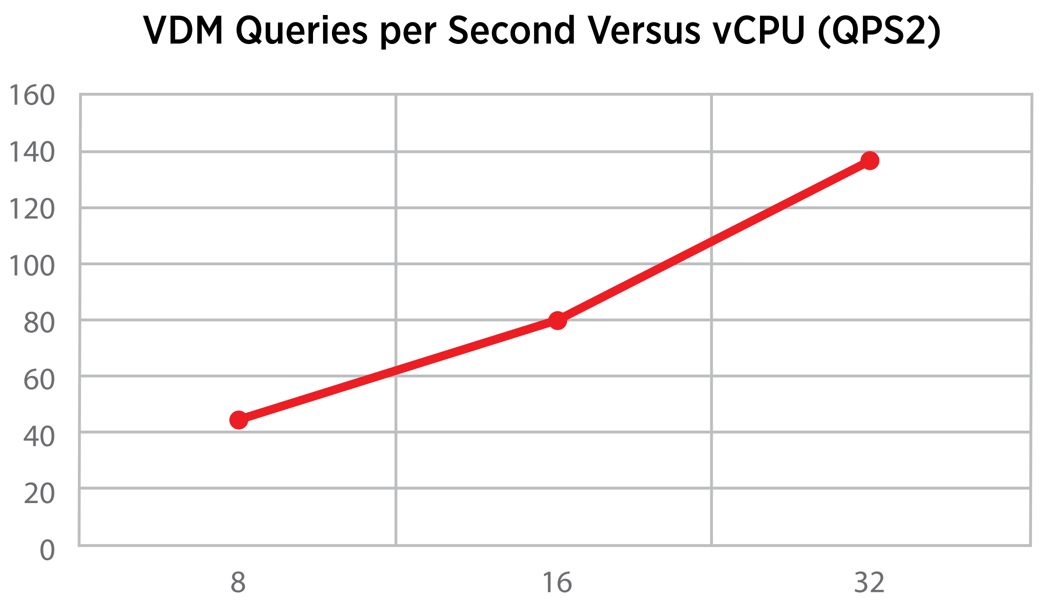

Figure 30. Query Time Comparison for Different SAP HANA VM Sizes for Query Type 2

TPC-DS is a standard benchmark for measuring the performance of decision support solutions. The benchmark models several generally applicable aspects of a decision support system, including queries and data maintenance. It includes a set of 99 queries that test various aspects of OLAP-based systems. The entire suite of queries was run against the SAP HANA databases, and the top 25 queries were compared for the three sizes of database servers. The results are shown in Figure 31. The 8-vCPU database system is used as the baseline and represents 100 percent from a query time perspective. The 16- and 32-vCPU systems' query times are compared with the baseline as shown in Figure 31.

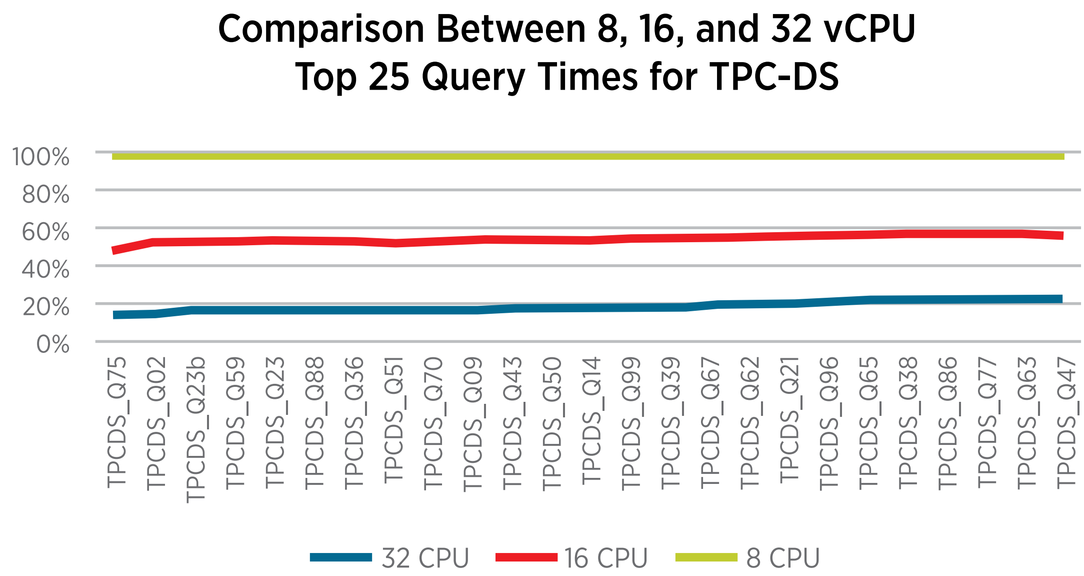

Figure 31. TPC-DS Query Time Comparison for Different SAP HANA VM Sizes

Results for Tests with Storage on All-Flash VMware vSAN

OLTP tests were run multiple times. The average runtimes were calculated and tabulated as shown in Table 25.

	8 vCPU	16 vCPU	32 vCPU
Insert	16.09	7.61	5.15
Delete	13.6	7.30	4.44
Update	3.26	1.84	1.04
Select	239.96	110.50	67.87

Table 25. OLTP Time Comparison for Different Operations on Different SAP HANA VM Sizes

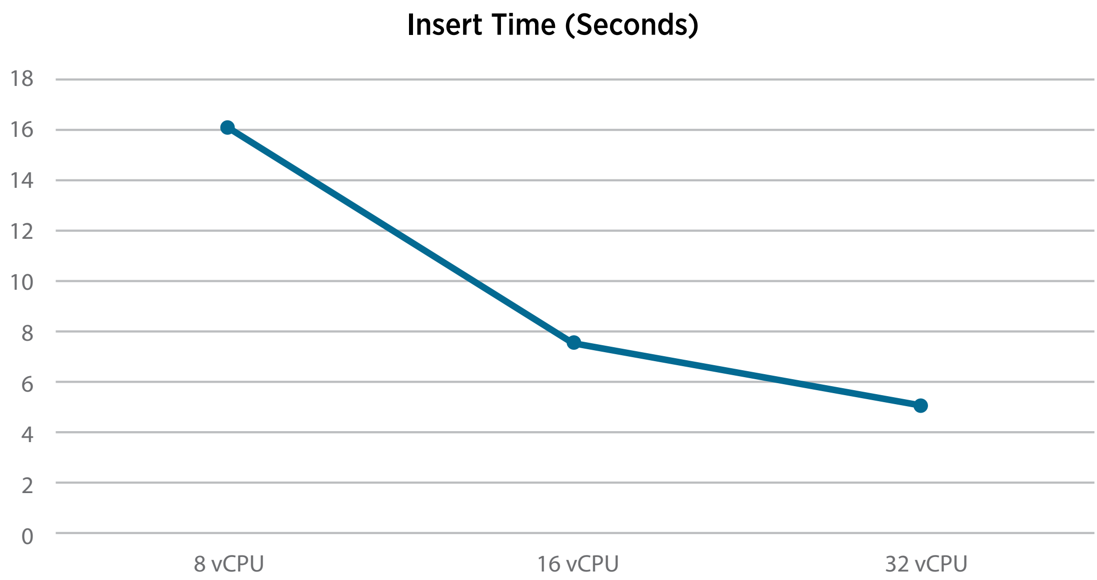

Figure 32. Insert Time Comparison for Different SAP HANA VM Sizes

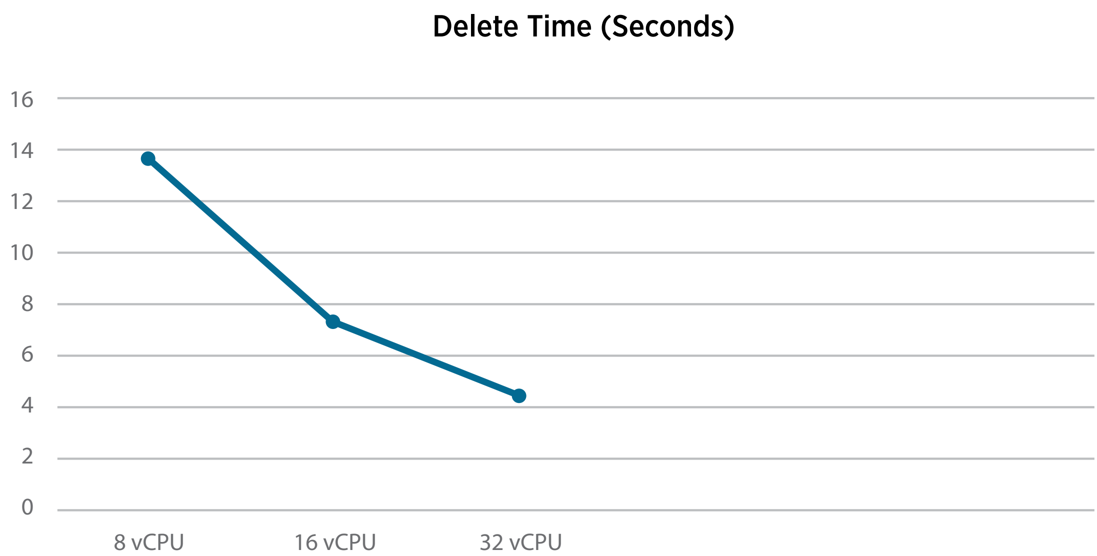

Figure 33. Delete Time Comparison for Different SAP HANA VM Sizes

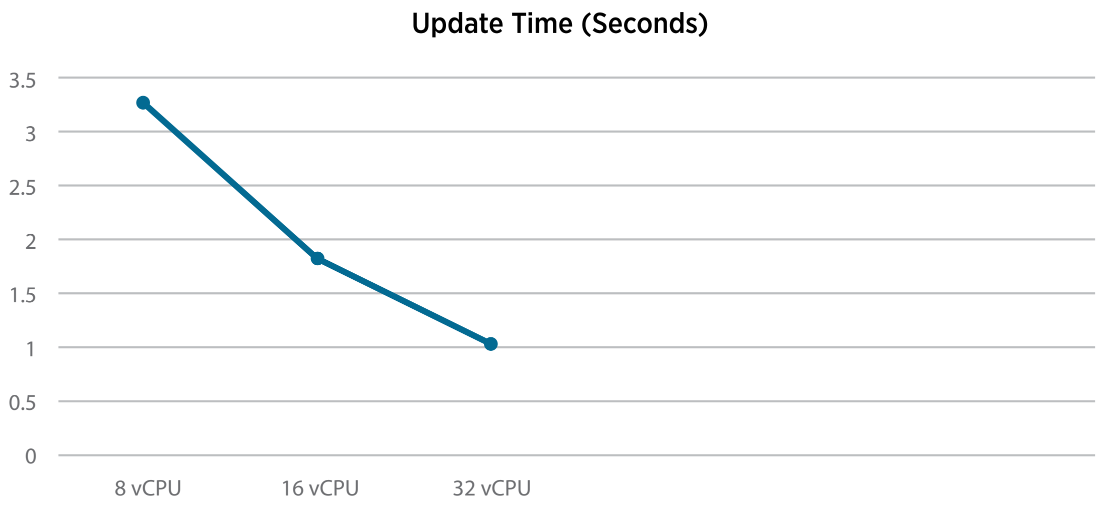

Figure 34. Update Time Comparison for Different SAP HANA VM Sizes

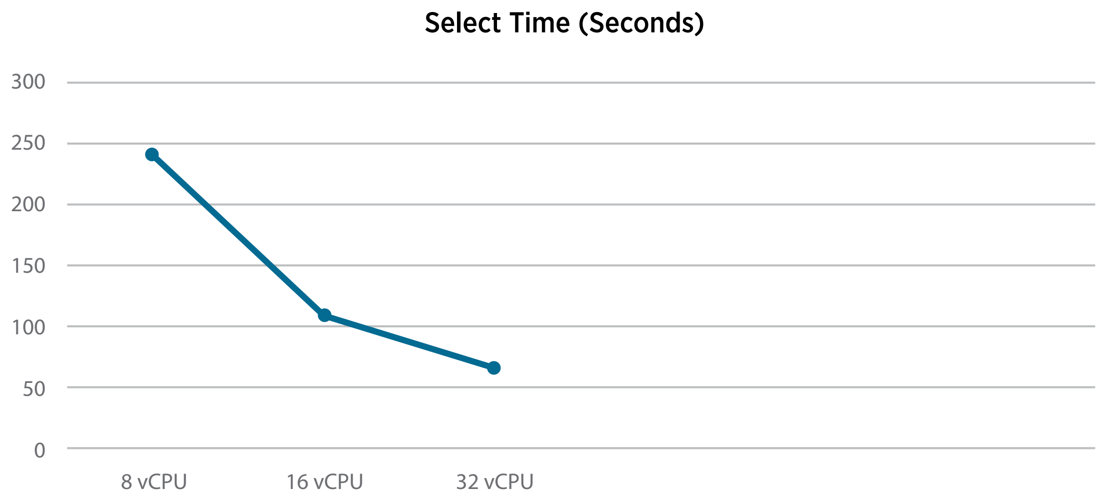

Figure 35. Select Time Comparison for Different SAP HANA VM Sizes

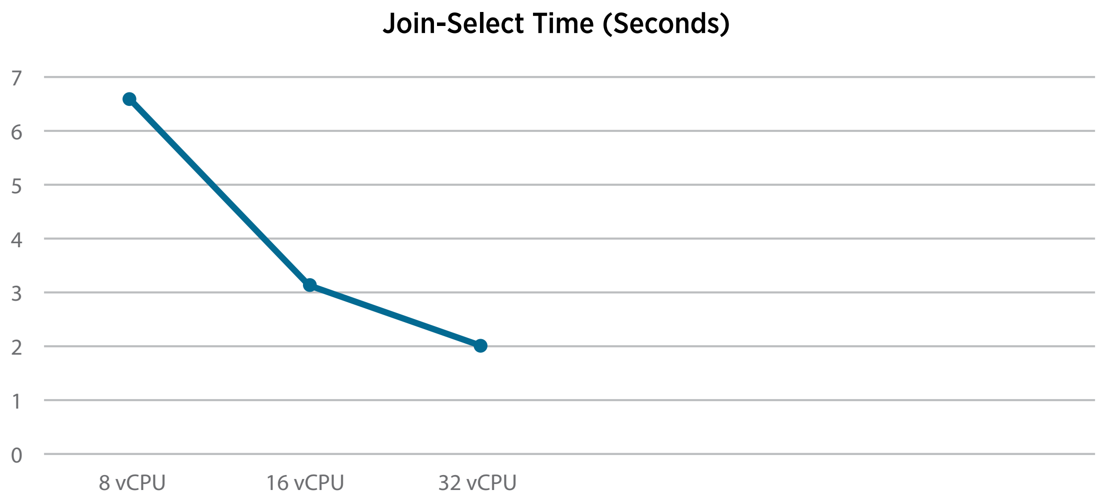

Figure 36. Join-Select Time Comparison for Different SAP HANA VM Sizes

Results from OLAP-Type Tests on All-Flash VMware vSAN

Two different types of OLAP tests were performed on the SAP HANA database servers: VDM and TPC-DS. The results of these tests are shown in Table 26.

VDM

The performance for this test is measured in queries per second (QPS). There are two different types of queries, labeled QPS1 and QPS2. The QPS for the two types of queries are compared in Table 26 for the three sizes of SAP HANA database VMs.

vCPU	QPS1	QPS2
8	4.08	45.33
16	7.08	79.45
32	11.75	136.35

Table 26. QPS Time Comparison Across Different SAP HANA VM Sizes

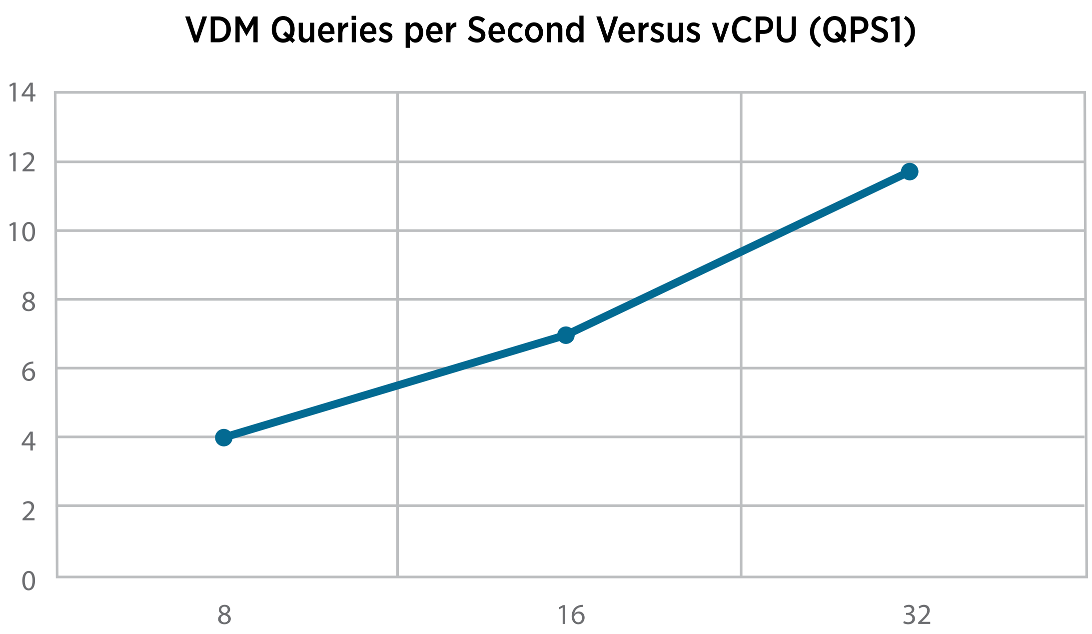

Figure 37. Query Time Comparison for Different SAP HANA VM Sizes for Query Type 1

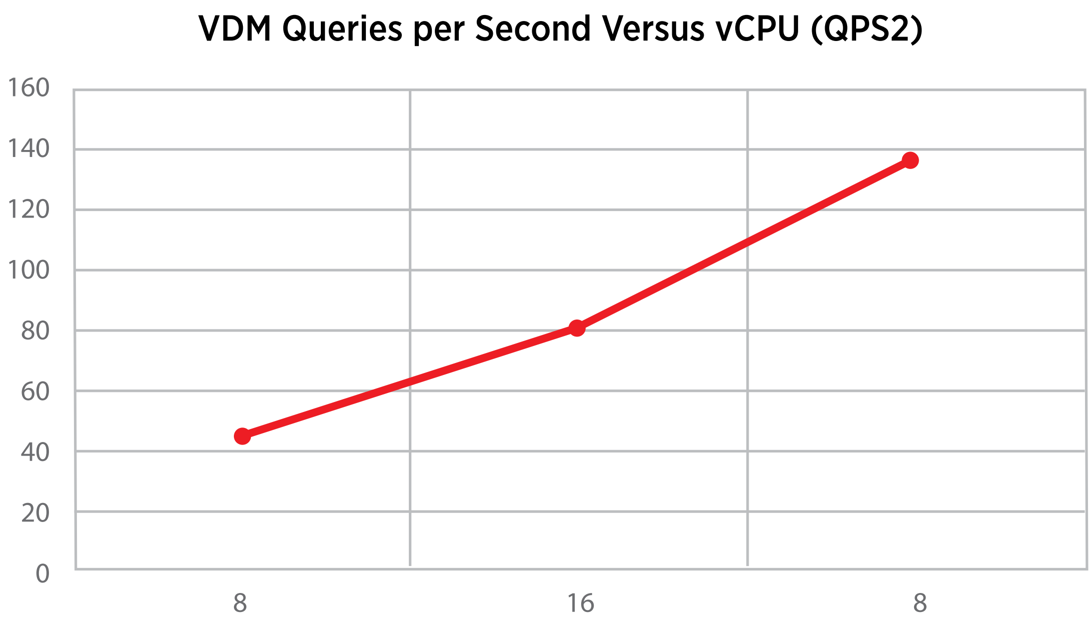

Figure 38. Query Time Comparison for Different SAP HANA VM Sizes for Query Type 2

Analysis

The results of the tests clearly show that the vSphere platform provides linear scalability and consistency in performance across different types of SAP S/4HANA database workloads. Linear CPU, memory, and I/O performance and scalability are reflected in these results.

TPC-DS on All-Flash VMware vSAN

TPC-DS is a standard benchmark for measuring the performance of decision support solutions. The benchmark models several generally applicable aspects of a decision support system, including queries and data maintenance. It includes a set of 99 queries that test various aspects of OLAP-based systems. The entire suite of queries was run against the SAP HANA databases, and the top 25 queries were compared for the three sizes of database servers. The results are shown in Figure 39. The 8-vCPU database system was used as the baseline and represents 100 percent from a query time perspective. The 16- and 32-vCPU systems query times are compared with the baseline as shown in Figure 39.

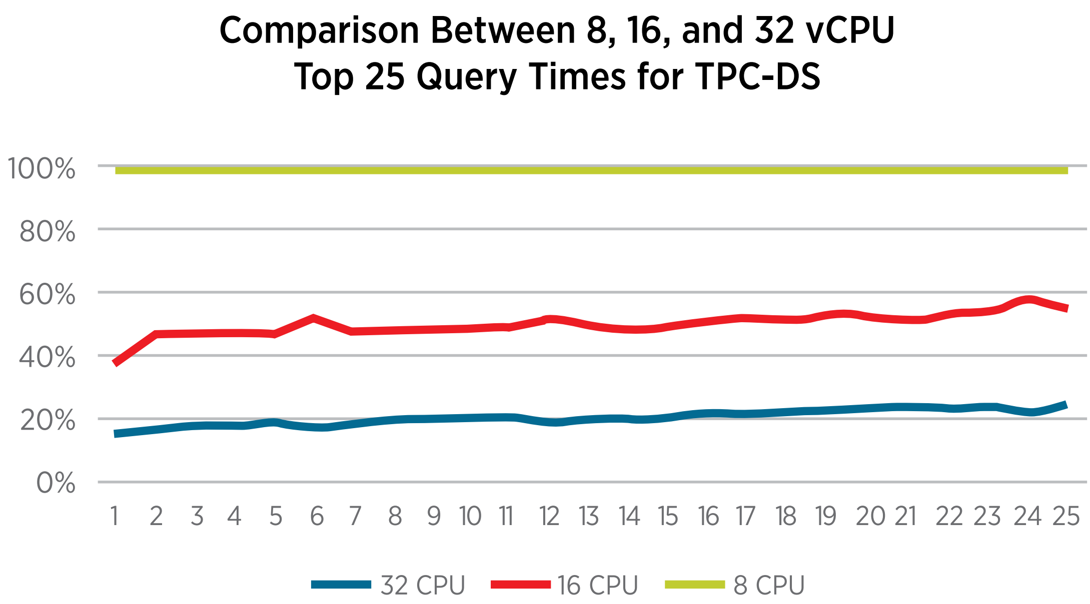

Figure 39. TPC-DS Query Time Comparison for Different SAP HANA VM Sizes

Analysis

The results of the tests clearly show that the vSphere platform provides linear scalability and consistency in performance across different types of SAP S/4HANA database workloads. Linear CPU, memory, and I/O performance and scalability are reflected in these results.

Conclusion

Virtualizing SAP HANA scale-out systems enables an organization to benefit from all supported VMware virtualization solutions and options, such as live migration via vSphere vMotion, to increase SLAs and reduce TCO. The recent joint SAP and partner testing and the subsequent release of the VMware virtualized BW-EML benchmark of an SAP HANA scale-out system have shown reliable operation with a very small impact on overall system performance. SAP HANA scale-out support in controlled availability provides additional benefits for the customer by offering additional consultation from SAP, VMware, and the hardware partner. This ensures that virtualized SAP HANA configurations work well and run optimally in a customer's virtualized environment.

A comprehensive implementation of the Software-Defined Data Center (SDDC) using the prescriptive advice found in this *Application Workload Guidance Design* document will lead to the successful "end-to-end" and linearly scalable virtual SAP HANA system. The components used within this set of tests constitute a reference architecture. Each component, although ideal for this type of implementation, is interchangeable with a variety of products widely available through the industry. The overhead imposed by the SDDC is extremely minimal. Even when using VMware vSAN, the tests in this report reveal that any overhead introduced by VMware vSAN processing is negligible.

This *Application Workload Guidance Design* document based on a VMware Validated Design, in addition to providing an end-to-end infrastructure for running business-critical applications, includes a VVD infrastructure that is designed, deployed, and validated for the entire SAP S/4HANA landscape. The SAP Quick-Sizer approach was used to convert business requirements to allocated compute resources. The sizing and designing of the SAP S4/HANA infrastructure followed SAP and VMware best practices. All VVD components were leveraged to build out the end-to-end solution for SAP S/4HANA. The following are some of the major components of the VVD that were used in this infrastructure:

- vSphere for virtualized compute and infrastructure
  - High-performance Dell R630 and R730 PowerEdge Servers
  - Pure Storage–based SAP TDI and Western Digital all-flash VMware vSAN storage
  - Brocade Generation 6 SAN fabric
- Separate vSphere clusters for management, SAP applications, and SAP HANA databases
- vSphere HA and vSphere FT for simplified high availability
- VMware NSX for networking with microsegmentation for security on Brocade VDX switches
- vRealize Automation with Adapter for SAP Landscape Management
- vRealize Operations with Blue Medora SAP dashboards for operations and capacity planning

The solution was then successfully validated to meet and exceed the requirements. The results of the validation tests clearly show that vSphere offers a robust platform for running SAP S/4HANA production workloads.

This set of tests was conducted in a truly comprehensive manner, both broad in scope and striking in depth. In addition to what has been highlighted in this voluminous work, there are components of VMware technology that were used as "common practice" during the testing procedure that are well accepted in "everyday" virtualized environments and therefore don't necessarily warrant detailed or complex analysis and explanation in the text. For example, vSphere vMotion was used on the virtual SAP HANA VMs to transition them from servers using the Pure Storage array to those configured to use VMware vSAN. This is a highly impressive capability, but it is one that in 2017 is so commonly employed in virtualized environments that it is unnecessary to highlight.

The time has come to definitively recognize that the well-known VMware motto "No Application Left Behind," which was adopted in the vSphere 6.x time frame and signifies that every implemented application and database is a candidate for virtualized infrastructure, is an observable fact. The benefits and value of the VMware SDDC will enhance any SAP HANA implementation. The long and winding road to virtual SAP HANA has been completed.

Acknowledgments

Authors

VMware

- Mohan Potheri
- Vas Mitra
- Erik Rieger
- Don Sullivan

Partners

- Krishna Satyavarupu (Pure Storage)
- Nanjunda Somayaji (Brocade)
- Craig Lee (Blue Medora)

The following individuals contributed to the creation of this guide:

- Arne Arnold
- Dirk Karjack
- David Galant
- Bob Goldsand

Resources

For more information on the SAP HANA and VMware products and technologies discussed in this guide, see the following links and references.

References

1. *Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere*
2. *SAP HANA Master Guide*
3. *SAP HANA Server Installation and Update Guide*
4. *Certified and Supported SAP HANA Hardware Directory*
5. *SAP HANA Tailored Data Center Integration Frequently Asked Questions*
6. *Security of the VMware vSphere Hypervisor*
7. *SAP HANA Tailored Data Center Integration (TDI) Overview*

8 http://service.sap.com/sap/support/notes/1943937 [requires SAP login]

9. *SAP HANA on VMware vSphere for Production Environments*
10. *vSphere Storage*
11. *Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on EMC VNX Series Unified Storage Systems*
12. EMC IT, 02/14 EMC Perspective, H12853

Appendix A: Brocade Storage Fabric and Capabilities for VMware Validated Design

Fibre Channel SAN Design Considerations

When designing the SAN for an all-flash array, understanding the application workloads and the intended scalability, redundancy, and resiliency requirements is the main factor to consider.

When deploying a Fibre Channel storage array, the main design consideration is the adequate sizing of the ISLs between the edge switches where the servers are connected and the core switches where the storage arrays are connected. In implementations where the Fibre Channel storage arrays are deployed to serve specific latency- and I/O-intensive applications, such as OLTP database servers, connecting both servers and Fibre Channel storage array to the core backbone switches can be advantageous.

When deploying dedicated SANs for Fibre Channel storage array services, the SAN design can be tailored directly, based on the application workloads and scalability requirements.

Storage Performance and Scaling

A best practice for SAN infrastructure design is to plan for a 3- to 5-year life span for the solution to be deployed. Considerations for the length of the life of a SAN infrastructure include a combination of equipment depreciation, the limited predictability of business transformation and development, and technology refresh and improvements. Having clear indications or understanding of application needs over the same period determines the range of scalability and flexibility necessary in the design.

SAN Topology Redundancy and Resiliency

Redundancy refers to duplicated components in a system—in this case, the SAN—whereas resiliency refers to the capability to continue when failures occur. So although redundancy supports resiliency, the degree of redundancy determines the level of resiliency. And although SANs are deployed with redundant fabrics, each fabric can also be designed without any SPOF and therefore be resilient by design. Designing fabrics to be resilient requires dual-core or full-mesh topologies. For the VMware Validated Design (VVD), a redundant SAN fabric leveraging two Brocade switches should be included.

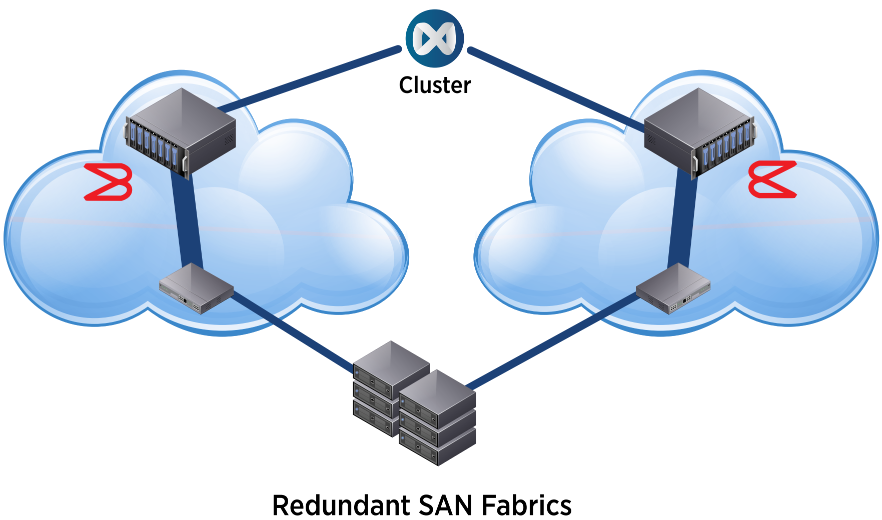

Figure A1. SAN Topology Redundancy

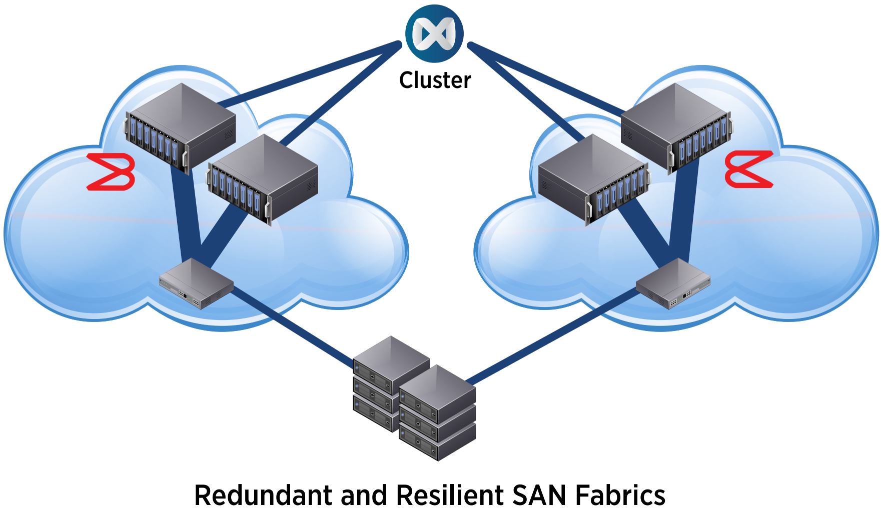

Figure A2. SAN Topology Redundancy and Resiliency

SAN Design Guidelines

Designing SANs fully dedicated with Fibre Channel storage is often a straightforward sizing exercise. Calculate the ISL bandwidth necessary for the server density and number of arrays in each Fibre Channel storage array cluster combined with the anticipated server and storage growth during the life cycle of the computing infrastructure.

For core-edge designs, a good rule of thumb is to design with ISL bandwidth equal to total capacity of all array ports that are accessible—that is, zoned—to hosts attached to each top-of-rack (ToR) switch.

For collapsed SAN designs, calculate the number of host and storage ports for the lifetime of the infrastructure.

Brocade Best Practice for Zoning

Zoning is a fabric-based service in SANs that groups host and storage nodes that must communicate. It enables nodes to communicate with each other only if they are members of the same zone. Nodes can be members of multiple zones, enabling a great deal of flexibility when a SAN is implemented using zoning.

Zoning not only prevents a host from unauthorized access of storage assets but also stops undesired host-to-host communication and fabric-wide Registered State Change Notification (RSCN) disruptions.

Brocade recommends that users always implement zoning, even if they are using LUN masking. Also, PWWN identification for zoning is recommended for both security and operational consistency. For details, refer to the *Brocade SAN Fabric Administration Best Practices Guide*.

Fibre Channel SAN Availability and Reliability

This solution uses Fibre Channel SAN based on Brocade Generation 6 technology with Brocade Fabric Vision technology and industry-standard 32Gb connectivity. The key features available in Brocade Fabric Vision and its integration with VMware vRealize Suite help build a highly available, performant, scalable, and reliable solution architecture. The following sections discuss the key features of Brocade Fabric Vision.

Automation

Monitoring and Alerts Policy Suite (MAPS) – MAPS is a policy-based monitoring and alerting tool that proactively monitors the health and performance of the SAN infrastructure based on predefined policies that cover more than 170 customizable rules, ensuring application uptime and availability. Administrators who want a pristine network can set an "aggressive" policy level that has rules and actions with strict thresholds to minimize the possibility of data errors. With the ability to tailor the MAPS policies, all-flash ports can be monitored more closely for faster identification of any performance degradation.

SAN Availability

Configuration and Operational Monitoring Policy Automation Services Suite (Compass) – Compass is an automated configuration and operational monitoring policy tool that enforces consistency of configuration across the fabric and monitors changes, simplifying SAN configuration and alerting users to changes. In medium- to large-sized environments, this can prevent inadvertent changes to switch configurations that can impact the preferred parameters set across the fabric to optimize performance.

Buffer Credit Recovery – Buffer Credit Recovery automatically detects and recovers buffer credit loss at the virtual channel level, providing protection against performance degradation and enhancing application availability.

Forward Error Correction (FEC) – FEC automatically detects and recovers from bit errors, enhancing transmission reliability and performance. It can reduce latency time significantly by preventing the need to retransmit frames with bit errors.

Monitoring and Troubleshooting

Dashboard – Dashboard includes customizable health and performance dashboard views, providing all critical information on one screen. Viewable dashboard "widgets" that should be monitored include errors on all-flash array–facing ports, top 10 flows, memory usage, and port health.

I/O Insight – I/O Insight monitors individual host and storage devices to gain deeper insight into the performance of the network to maintain SLA compliance and obtain total I/O, first response time, I/O latency—that is, exchange completion time (ECT)—and outstanding I/O performance metrics for a specific host or storage device, to diagnose I/O operational issues, enable tuning of device configurations with integrated I/O metrics, and optimize storage performance.

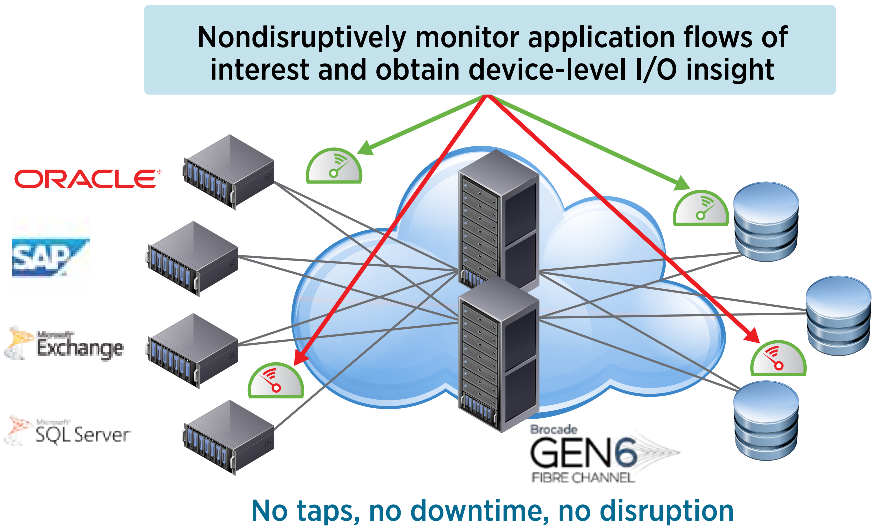

Figure A3. I/O Insight

VM Insight – VM Insight uses standards-based, end-to-end VM tagging to gain VM visibility in a storage fabric. It enables storage administrators to monitor VM-level application performance and set baseline workload behavior. Using this information, they can quickly determine whether a storage fabric is the source of performance anomalies for applications running on VMs. VM Insight also enables fast correlation with other Fabric Vision metrics to identify the root cause of problems before operations are affected. It provides the visibility for administrators to provision and plan storage networks based on application requirements and to fine-tune the infrastructure to meet service-level objectives.

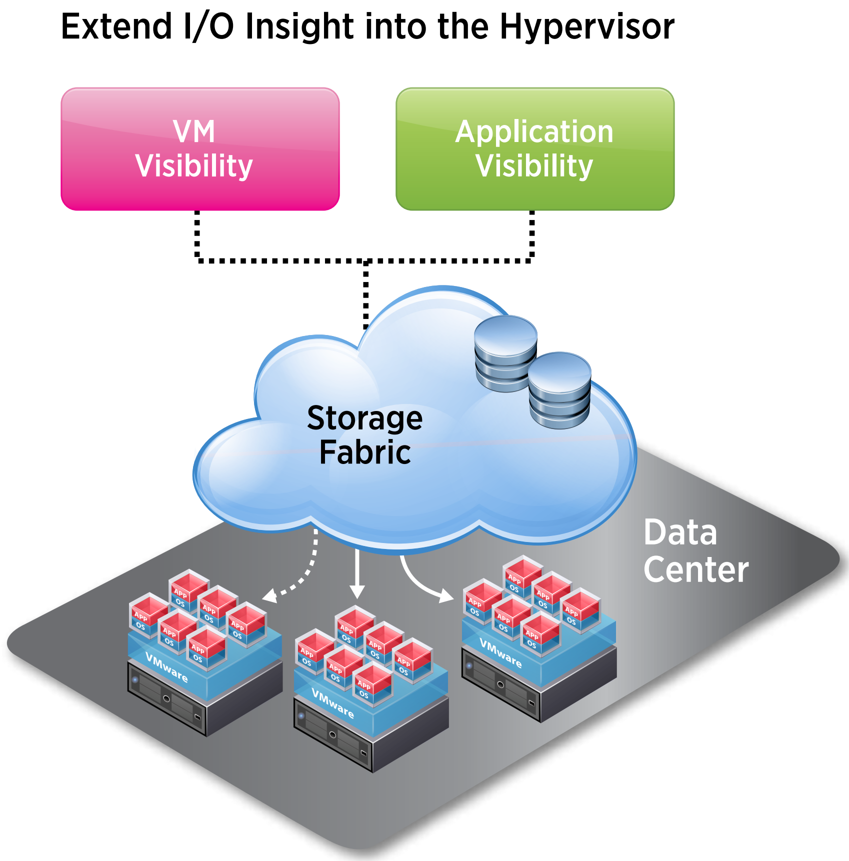

Figure A4. VM Insight

Flow Vision – Flow Vision is a comprehensive tool that enables administrators to identify, monitor, and analyze specific application and data flows to maximize performance, avoid congestion, and optimize resources. Flow Vision consists of the following elements:

- Flow Monitoring—Monitors specified traffic flows from source to destination through the SAN
- Flow Generator—Generates traffic between any two ports in a Generation 5 fabric
- Flow Mirroring—Captures packet data as it flows through the SAN, then displays and analyzes the captured packet's data

Flow Vision is best suited for temporary use while troubleshooting high-latency conditions. Continual use results in the collection of substantial amounts of diagnostic data. Flow Vision can be used as needed to verify optimal performance for the most demanding applications requiring optimal performance.

Fabric Performance Impact (FPI) Monitoring – FPI Monitoring identifies and alerts administrators to device or ISL congestion and high levels of latency in the fabric, which can have a severe impact on all-flash array performance. It provides visualization of bottlenecks and identifies slow drain devices and impacted hosts and storage.

ClearLink Diagnostics – Diagnostic Port (D_Port) provides loopback test capabilities for link latency and distance measurement at the optical and electrical level to validate the integrity and performance of optics and cabling, ensuring signal and optical quality and optimal performance across SAN and WAN connections. Prevalidating the integrity of cables and optics with ClearLink Diagnostics prior to deployment identifies potential support issues before they occur and enhances the resiliency of high-performance fabrics. For all-flash–array performance, any impurity in the physical infrastructure can impact performance.

VMware vRealize and VMware vRealize Log Insight™ – Brocade has worked closely with VMware to offer deeper investigation capabilities for root-cause analysis and remediation of SANs in virtualized environments. The Brocade SAN Content Pack eliminates noise from millions of events and amplifies critical SAN alerts to accelerate troubleshooting with actionable analytics. Faster troubleshooting provides time to proactively drive the value of IT resources. Leveraging Brocade Fabric Vision technology provides thorough knowledge of the behavior and health of Brocade SAN fabrics over time to identify and remediate patterns impacting VM performance and application responsiveness.

Appendix B: Pure Storage FlashArray//M50

The Pure Storage FlashArray//M50 was used as the TDI storage for SAP HANA.

Figure B1. Pure Storage FlashArray//M50

The following are some of the advantages of running SAP HANA on Pure Storage:

1. Data reduction – SAP HANA performs data compression in column stores on two levels: dictionary and advanced. Further data reduction results of around 1.9–2.3 have been experienced. SAP HANA does not compress row store tables, so the data reduction is even higher when there are row store tables. This use case is analytics, so there are no row store tables.
2. Encryption off-loading – SAP HANA encrypts only data volumes; it does not encrypt log volumes. Switch off the encryption on SAP HANA and use the Pure Storage FlashArray encryption, which can preserve valuable SAP HANA CPU cycles. That in turn can be used more efficiently to process analytical queries.
3. SAP homogeneous system copy and backup and recovery using storage snapshots – System copy, backup, and recovery using Pure Storage snapshots are extremely fast and efficient. The entire process can be automated very easily via a script and can be scheduled on a regular basis. There is no need to invest in backint third-party SAP-certified tools. This process enables customers to achieve very low RTO and RPO.

Additional Key Benefits of Pure Storage Arrays

The following are other key benefits of Pure Storage arrays:

- Consistent performance and scalability
  - Consistent submillisecond latency with 100-percent flash storage
  - Scalability through a design for hundreds of discrete servers and thousands of VMs; the capability to scale I/O bandwidth to match demand without disruption
- Operational simplicity
  - Reduced management complexity
  - No storage tuning or tiers necessary
  - Autoaligned 512Mb architecture that eliminates storage alignment problems
- Improved TCO
  - Dramatic savings in power, cooling, and space with 100-percent flash
- Enterprise-grade resiliency
  - Highly available architecture and redundant components
  - Nondisruptive operations
  - Ability to upgrade and expand without downtime or performance loss
  - Native data protection: snapshots and replication

Appendix C: VMware Adapter for SAP Landscape Management – Connector for vRealize Automation

VMware Adapter for SAP Landscape Management – Connector for vRealize Automation is a software suite that enables VMware vRealize Automation administrators to memorialize an SAP system—typically a verified one—into a *Gold Master* (GM). The connector instantiates blueprints from the Gold Masters and the necessary workflows for deployment and imports them into vRealize Automation and VMware vRealize Orchestrator, respectively. The blueprints are then entitled and made available to the end users in the vRealize Automation service catalog. vRealize Automation end users with requisite permission can then at any time select the blueprint for deployment from the service catalog, furnish the requested values, and be able to create a new and verified SAP system based on the Gold Master. This process is called *instantiating the Gold Master*. The connector is built on top of the VMware Adapter for SAP Landscape Management, which is based on the SAP Landscape Management (LaMa) offering. For an understanding of this VMware product, see http://www.vmware.com/products/adapter-sap-lvm.html. Figure C1 illustrates the components of a connector environment and their relationship to one another.

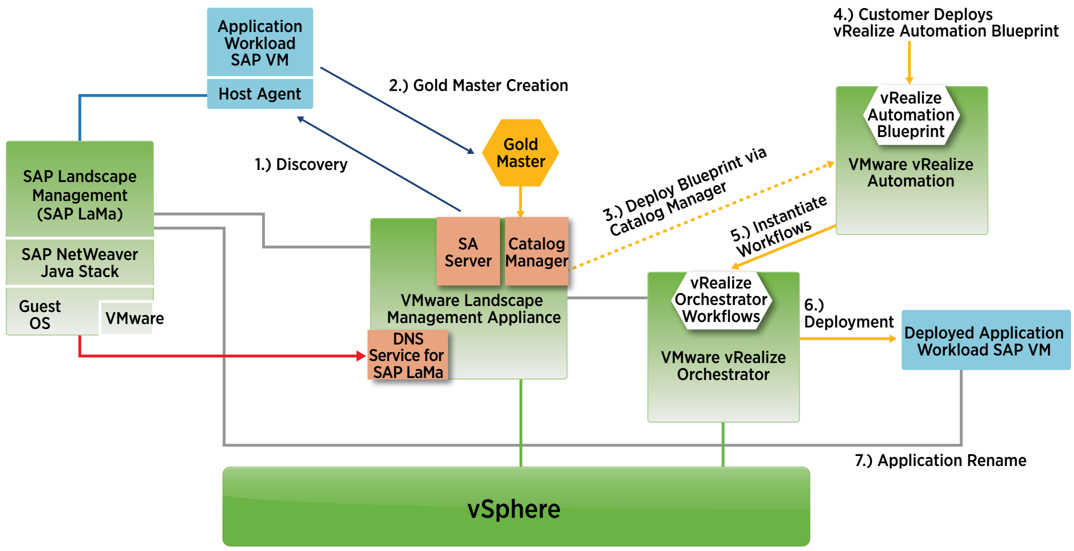

Figure C1. SAP Automation Solution Overview

Figure C1 shows the following key components:

- Verified SAP Systems – Each of these systems consists of software running on one or more machines—bare metal or in the case of VMware Landscape Management Appliance environments, VMs hosted on VMware vSphere products; that is, VMware ESXi systems managed by VMware vCenter Server. The machines perform a business function, such as order processing, accounts payable, general ledger, inventory management, and so on. Each SAP System consists of one or more components. When all the components are up and running, the SAP System is running. However, when all of the components are stopped, the SAP System is stopped. If some systems are running and some are not, the SAP System is in an intermediate state.
- SAP Landscape Management (SAP LaMa) VM – The SAP Landscape Management application runs on an SAP NetWeaver Application Server Java Edition. It provides a Web-based user interface for SAP BASIS administrators to centralize SAP Landscape Management and operations of SAP Systems and landscapes running on physical or virtual infrastructures. SAP Landscape Management has an extensible architecture that enables SAP and third-party vendors—VMware, for example—to create plug-ins to extend certain features.
- VMware Adapter for SAP Landscape Management – This is a plug-in to SAP Landscape Management that extends the way SAP LaMa integrates with the underlying systems virtualized with vSphere, optimizing

and extending the functionality for certain operations, such as activating (powering on) and deactivating (powering off), copying and cloning systems, and automation of these copying and cloning operations.

- VMware vSphere – ESXi is the premier hypervisor product from VMware. Virtual infrastructure administrators typically install it on bare-metal server-class computers, with VMs running guest operating systems (OSs) with SAP Systems as applications within the guests. vCenter Server is the premier VMware product for managing environments virtualized with ESXi. Collectively called vSphere, these products provide an enterprise-class environment with features for creating clusters, load-balancing VMs between host systems (ESXi instances), fault tolerance, virtual networking, virtual storage, and more. In VMware Landscape Management Appliance environments, the appliance runs in a VM on this infrastructure.
- VMware vRealize Orchestrator – This VMware product helps VI administrators automate their environments by creating workflows—essentially scripts—that perform VI administrative actions, including complex actions that can require multiple steps and involve loops, conditions, and so on. vRealize Orchestrator workflows can handle exceptions automatically or can pause to wait for a VI administrator to mitigate an issue.
- VMware Landscape Management Appliance – This part of the product is a virtual appliance. Collectively, it consists of one or more Web services that accept commands from VMware Adapter for SAP Landscape Management and take appropriate actions to implement the commands, typically with the help of vRealize Orchestrator. The following are examples:
  - When an SAP BASIS administrator activates—that is, powers on—an SAP System via SAP Landscape Management, VMware Adapter for SAP Landscape Management sends commands to VMware Landscape Management Appliance to power on the underlying VM. VMware Landscape Management Appliance in turn invokes a workflow specific to it on vRealize Orchestrator to turn on the VM in the underlying vSphere infrastructure. An analogous action occurs when an SAP BASIS administrator deactivates—that is, powers off—an SAP System.
  - When an SAP BASIS administrator copies an SAP System, VMware Adapter for SAP Landscape Management sends commands to VMware Landscape Management Appliance, which in turn invokes a vRealize Orchestrator workflow specific to it to create vSphere copies of the VMs on which the SAP Systems reside, configuring the VMs according to the parameters provided by the SAP BASIS administrator in the SAP Landscape Management Web user interface.
- Gold Master – SA-API provides a model and mechanism that enable organizations to memorialize verified SAP Systems in a *Gold Master*: images of the SAP host system VMs plus certain metadata. A copy or clone of the SAP application can quickly be deployed from the Gold Master, always leaving the verified Gold Master as is. This is called instantiation of the Gold Master.
- VMware vRealize Automation – vRealize Automation enables users to create and manage multivendor cloud infrastructures. With it, end users can self-provision VMs, cloud machines, physical machines, applications, and IT services as per the policies denied by the administrators. vRealize Automation provides a secure portal through which authorized administrators, developers, or business users can request new IT services. They can manage specific cloud and IT resources while ensuring compliance with business policies. There is a common service catalog through which users can request IT service, including infrastructure, applications, and desktops. This ensures a consistent user experience.
- vRealize Automation blueprint – A machine blueprint is a complete specification of a virtual, cloud, or physical machine that defines resources, attributes, policies, and methods of provisioning for a new system. A blueprint is required to provision an SAP System from a Gold Master.

Service catalog items on vRealize Automation are unique SAP System blueprints that the end user can browse and is entitled to request. The end user selects the blueprint for provisioning and provides the required values. The values for different fields can be typed in, or the user might also be prompted to select values from a drop-down list for any custom properties defined in the blueprint. The user can then monitor the current status of the provisioning request in the portal.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com