JANUARY 2017

ARCHITECTURE GUIDELINES AND BEST PRACTICES FOR DEPLOYMENTS OF SAP HANA ON VMWARE VSPHERE

ARCHITECTURE AND TECHNICAL CONSIDERATIONS GUIDE

Table of Contents

Overview 9

Document change log 9

Audience 9

SAP Acknowledgments 9

Introduction 9

Production Support 10

Enabling Products and Technologies 13

SAP HANA 13

Real-time Analytics 14

Real-time Applications 15

OLTP Workloads 15

SAP HANA Deployment Options 15

SAP HANA Scale-Up Option 17

SAP HANA Scale-Out Option 21

SAP HANA Appliance Model 23

SAP HANA Tailored Datacenter Integration (TDI) 24

SAP HANA TDI Support Process 25

VMware vSphere 26

VMware vSphere 6.x–New Features 28

vSphere 5.5 vs 6.0 30

VMware Resource Isolation and Security–Basis for SAP HANA Multi-VM 32

VMware vMotion (Live Migration) 35

Migrating SAP HANA Databases with VMware vMotion 37

Why is it critical to follow these best practices? 40

Managing SAP HANA Landscapes Using VMware DRS 42

VMware High Availability 44

SAP HANA VM Clones and Template 46

VMware Host Profiles 47

Migration from vSphere 5.5 to vSphere 6.x 47

Virtualized SAP HANA Reference Architecture 48
General Concepts and Overview. 49
Architecture Overview 49
Storage Layer 53
Virtual Machine File System. 54
Datastores. 54
NFS Datastores 55
SCSI Adapters 56
File System Considerations and Alignment. 57
Server and Network Layer. 61
SAP HANA on vSphere Server Considerations 62
Virtual CPUs (vCPUs) and Hyperthreading 62
Memory Considerations 64
Non-Uniform Memory Access. 65
SAP HANA on vSphere Network Considerations 68
Virtualization Layer 74
Application and OS Layer 76
Management Layer. 76
Availability Considerations for vCenter Server. 77
SAP HANA on vSphere High Availability (HA) 77
SAP HANA High Availability Solutions Overview 79
Fault Recovery 80
SAP HANA Service Auto-Restart. 80
Host Auto-Failover (standby host) 80
VMware vSphere HA for SAP HANA. 80
VMware vSphere HA with External Cluster or Monitoring Solution for SAP HANA. 85
Disaster Recovery (DR) 87
SAP HANA Data Persistence. 87
SAP HANA on vSphere Backup and Restore 88
SAP HANA on vSphere Storage Replication 89
Storage-based Replication. 90
VMware vSphere Replication 91

SAP HANA system replication 93

VMware vSphere Site Recovery Manager (SRM). 95

Virtual SAP HANA Sizing 96

Virtual SAP HANA Sizing Considerations 96

SAP HANA Sizing Procedure 100

Using SAP Quick Sizer HANA version (new systems) 103

Using SAP HANA sizing reports to determine the memory and disk requirements of SAP HANA (existing systems) 111

HANA database server selection process based on memory requirements 112

Virtual HANA Configuration Examples 120

Deployment Options and Support Status. 122

Best Practices–Configuration. 124

Best Practices–Optimizing SAP HANA on vSphere. 126

Monitor and verify an SAP HANA installation. 132

Performance Optimization for Production-Level SAP HANA VMs. 132

Support Process for SAP HANA on VMware vSphere. 138

Summary. 140

VMware Acknowledgments 140

List of Figures

Figure 1: SAP HANA: The Platform for All Applications 13

Figure 2: SAP HANA In-memory Platform for a New Class of Real-time Analytics and Mission-Critical Applications[10] 14

Figure 3: Scaling Options of an SAP HANA System on VMware vSphere 17

Figure 4: SAP HANA Single-Host System (Scale-Up) Configuration 18

Figure 5: SAP HANA Multi-VM Single-Host System (Scale-Up) Configuration . . . 19

Figure 6: SAP HANA Landscape Consolidation Example 20

Figure 7: SAP HANA Multiple-host System (Scale-Out) on vSphere Configuration. 22

Figure 8: SAP HANA Scale-Out/Multiple VM Landscape Consolidation. 23

Figure 9: SAP Tailored Datacenter Integration (TDI) 25

Figure 10: VMware vSphere Virtual Infrastructure 28

Figure 11: HPE vSphere 5.5 vs 6.0 SAP HANA BW-EML Benchmark Results 31

Figure 12: Comparison between vSphere 5.5, 6.0 and 6.5, *1, with specific OEM partners, *2 as of writing this guide only 4 TiB are currently supported. 6 TiB tests are on the roadmap. 31

Figure 13: Co-deployment and Consolidation of Independent and Different SAP Workloads on a Server 33

Figure 14: Transparent Page Sharing–Page-Content Hashing 33

Figure 15: Network Isolation of a Virtualized Server Environment. 34

Figure 16: VMware vMotion. 35

Figure 17: Non-critical vMotion Scenarios 41

Figure 18: Critical vMotion Scenario 41

Figure 19: Maintenance Window After a HW Upgrade and vMotion 42

Figure 20: VMware Distributed Resource Scheduler 43

Figure 21: VMware HA/DR Solutions–Protection at Every Level. 44

Figure 22: VMware High Availability. 45

Figure 23: Virtualized SAP HANA High-level Architecture 50

Figure 24: VMware vCenter Server. 51

Figure 25: Converting a Virtualized SAP HANA Scale-Up to a Scale-Out System 52

Figure 26: SAP HANA TDI Scale-Up and Scale-Out Systems Sharing a Single TDI Storage 52

Figure 27: Virtual Machines Accessing Different Types of Storage 53
Figure 28: SAP Recommended SAP HANA File System Layout 56
Figure 29: Virtualized SAP HANA Disk Layout of a Scale-Up System 59
Figure 30: SAP HANA on vSphere 6.0 Certification Architecture 61
Figure 31: CPU Utilization Across All Virtual Machines 64
Figure 32: Node Boundaries and VM Placement 66
Figure 33: SAP HANA on vSphere – NUMA Best Practices 66
Figure 34: NUMA Node Sharing Example with Haswell 18-Core System 67
Figure 35: SAP HANA VM with 36 vCPUs and 500 GiB RAM, Sized to Fit into a Single NUMA Node (CPU Socket) of an 18-Core CPU with 512 GiB RAM Attached 67
Figure 36: vSphere Standard Switch Architecture. 69
Figure 37: Logical Network Connections per SAP HANA Server 70
Figure 38: vSphere SAP HANA Scale-Out Network Configuration of a Virtualized 3+1 SAP HANA Scale-Out System 73
Figure 39: vSphere Versions–Compared 75
Figure 40: SAP HANA Hardware Directory–OS Filter Criteria 76
Figure 41: SAP HANA System High Availability Phases 78
Figure 42: VMware vSphere HA Combined with SAP HANA Service Auto-Restart 81
Figure 43: SAP HANA Protected Systems via vSphere HA. 83
Figure 44: VMware HA–VM Restart After OS Failure 83
Figure 45: VMware HA - Protection Against Host Failures 84
Figure 46: SAP HANA Service Auto-Restart Feature Combined with VMware HA 85
Figure 47: SAP HANA System Replication as Fast RTO Cluster Solution 87
Figure 48: SAP HANA Data Persistence. 88
Figure 49: SAP HANA Backup and Recovery Methods 89
Figure 50: SRM–Configure Replication per VM 92
Figure 51: Configure RPO per VM. 92
Figure 52: SAP HANA on vSphere Replication. 94
Figure 53: SAP HANA Site Failover Managed by VMware SRM. 96
Figure 54: General SAP HANA Sizing Procedure with Building Blocks, Sizing Guide. 101

Figure 55: Sizing Result of an SAP Simple Finance on HANA Application, Quick Sizer Version 244 105
Figure 56: SoH Sizing Example Based on Quick Sizer Result 106
Figure 57: Sizing Result of an SAP Business Warehouse Powered by SAP HANA Application, Quick Sizer Version 244 106
Figure 58: BWoH Sizing Example Based on Quick Sizer Result 107
Figure 59: Excerpt from a BW Sizing Report 111
Figure 60: Certified SAP HANA Hardware Directory 114
Figure 61: VMware Compatibility Guide 115
Figure 62: VMware HCL Results View 115
Figure 63: Detailed View of a Server in the VMware HCL 116
Figure 64: HANA Certified Directory Server Detailed View 116
Figure 65: Server Configuration Example with Three Independent SAP HANA Instances and 1.5% RAM Overhead 118
Figure 66: Example Server Non-production Configuration with Five Independent SAP HANA Instances 120
Figure 67: Example Server Configuration for SoH Instances on 4-Socket 4 TB E7-v4 Systems 121
Figure 68: Example Server Configuration for BWoH Instances on 4-Socket 2 TB E7-v4 Systems 121
Figure 69: SAP Support Workflow for VMware Related Escalations 139

List of Tables

Table 1: SAP HANA vSphere Version Support Status–as of December 2016 11

Table 2: SAP HANA vSphere 5.5 Support Status per Use Case and Workloads–as of December 2016 12

Table 3: SAP HANA vSphere 6.0 Support Status per Use Case and Workloads–as of December 2016 12

Table 4: SAP HANA on vSphere Supported Hardware Configurations–as of December 2016 12

Table 5: vSphere 5.5 vMotion Migration Policies 38

Table 6: ECV Levels (source: VMware KB100321) 40

Table 7: vSphere Features Supported by Storage Type 55

Table 8: Virtualized SAP HANA Storage Configuration Based on SAP HANA Storage Requirements Document 58

Table 9: Example TDI Configuration, Number of RAID Groups–All of Type RAID6(6D+2P), Based on a HDS Storage System 60

Table 10: SAP HANA vSphere Minimum Recommended Host and VM Network Configuration 71

Table 11: SAP HANA vSphere Bandwidth Optimized Host and VM Network Configuration 72

Table 12: SAP HANA on vSphere High Availability Solutions 79

Table 13: SAP HANA Platform Sizing Figures for Selected CPU Types 97

Table 14: Maximum SAP HANA VM Configuration 100

Table 15: SAP HANA Platform Sizing Figures for E7-v3 and E7-v4 CPUs 108

Table 16: Virtual HANA Configuration Examples 121

Table 17: SAP HANA on vSphere Supported Capabilities and Options 124

Table 18: vSphere Physical Host Server Settings 126

Table 19: vSphere ESXi Server 127

Table 20: SAP HANA VM Settings 129

Table 21: Virtual Machine Guest Operating System 132

Overview

This document provides best practices and configuration setting guidelines to ensure the best possible performance for production-level SAP HANA® virtual machines (VMs). The settings, such as CPU scheduler affinity, may not apply to other applications or databases, and therefore should be treated as specific to SAP HANA.

This document is the second edition of the best practices and recommendations guidelines for SAP HANA on VMware vSphere®. It describes the best practices and recommendations for configuring, deploying, and optimizing SAP HANA Scale-Up and Scale-Out deployments that run on the VMware virtualization infrastructure. Therefore, it combines the earlier Scale-Up and Scale-Out guides into a single source of truth. Most of the guidance provided in this document is the result of continued joint testing conducted by VMware and SAP to characterize the performance of SAP HANA, powered by VMware vSphere. The support status of the specific vSphere versions is summarized in several tables, and you must consider them when you plan an SAP HANA on vSphere environment.

Document change log

This document replaces the *Best Practices and Recommendations for Scale-Up Deployments of SAP® HANA® on VMware vSphere® Deployment and Considerations* Guide, published in 2014, and the *Best Practices and Recommendations for Scale-Out Deployments of SAP HANA on VMware vSphere Deployment and Technical Considerations Guide*, published in 2015, which includes information for vSphere 5.5 and vSphere 6.0.

Audience

This guide is intended for architects, engineers, and administrators who are responsible for configuring and deploying the SAP HANA platform in a VMware virtualization environment. It assumes that the reader has a basic knowledge of VMware vSphere concepts and features, SAP HANA, and related SAP products and technologies.

SAP Acknowledgments

The author wants to thank the SAP HANA product management, development, QA, and sizing teams for their constructive feedback, especially: Zora Caklovic, Jan Teichmann, Oliver Rebholz, Sebastian Schmitt, Thomas Becker, and Maik Gaertner from SAP.

The sections on product support, support process, virtualized SAP HANA architecture, high availability (HA) options, sizing, and best practices were reviewed in detail and approved by SAP. The remaining sections provide detailed information on how to plan, integrate, and operate a virtualized SAP HANA environment.

Introduction

SAP® and partner solutions that use the market-leading SAP HANA in-memory data platform enable businesses to run in real time. Optimized for both transactional and analytical processing, SAP HANA solutions can dramatically accelerate analytics, business processes, sentiment data processing, and predictive capabilities. SAP HANA solutions also allow businesses to analyze Big Data quickly, which drives rapid innovation and provides a platform for you to develop your own business applications that support real-time analytics, or lets the customer choose ready to run applications from SAP that are SAP HANA optimized, such as SAP S/4HANA®.

By running SAP HANA on the VMware vSphere virtualization and cloud computing platform, Scale-Up database sizes up to 4 TiB[1] can now be deployed on top of a VMware® virtualized infrastructure. Larger database sizes can be deployed when leveraging the SAP HANA Scale-Out deployment option.

Running the SAP HANA platform virtualized on VMware vSphere delivers a new deployment architecture to SAP HANA customers. The SAP HANA with VMware solution provides SAP customers with a data center platform that provides agility, high availability, cost savings, and easy provisioning. This solution also gives SAP customers the ability to provision instances of SAP HANA in virtual machines faster.

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

- Increased security and SLAs (e.g., using VMware vSphere® vMotion® or VMware vSphere® Distributed Resource Scheduler™ [DRS])
- Live migration of running SAP HANA instances to other vSphere host systems with vSphere vMotion
- Standardized high availability solution based on VMware vSphere® High Availability (HA)
- Built-in multi-tenancy support via system encapsulation in a virtual machine (VM)
- Abstraction of the hardware (HW) layer
- Higher hardware utilization rates
- Streamlining IT operation, processes, and standards

These and other advanced features found almost exclusively in virtualization lower the total cost of ownership and ensure the best operational performance and availability.

Production Support

In November 2012, SAP announced initial support for SAP HANA on vSphere 5.1 for non-production environments. Since then, SAP has extended its support to SAP HANA Scale-Up production environments for vSphere 5.5, including SAP HANA Tailored Datacenter Integration deployments, as well as to multiple production SAP HANA VMs running on a single vSphere host, and SAP HANA Scale-Out deployments on vSphere 5.5.

SAP now supports Scale-Up VMware vSphere 6.0 and 6.5 deployments in production. Table 1 through Table 4 provide the SAP HANA 1.0 and 2.0 on vSphere support status– you need to consider the information in these tables for all deployment options described in this paper.

[1] TiB = Tebibyte = 240 Bytes

Note: HANA on VMware vSphere is supported on certified HANA hardware configurations only. The same scalability strategy and limits apply to HANA on VMware as well as bare-metal HANA.

With vSphere 6.x, the guest and host resource limits have increased significantly, so customers can deploy even more complex and larger landscapes. An example of some relevant mission–critical physical vSphere host resource limits are as follows: RAM is 12 TiB[2] , 16 non-uniform memory access (NUMA) nodes, 480 logical CPUs, 1024 virtual machines, 2048 virtual disks, and 64 TiB LUN (logical unit) sizes. Leveraging vSphere 6.0 with its larger supported VM sizes or SAP HANA Scale-Out configurations, customers can add new physical server resources incrementally or add complete host server systems with SAP HANA VMs running them. Each virtualized SAP HANA system can be installed with up to 1 TiB of memory (vSphere 5.5) or up to 4 TiB of memory (vSphere 6.0 and 6.5) and once supported up to 6 TiB per VM with vSphere 6.5, depending on the available physical hardware and individual VM sizing and SAP limitations[3], to adapt easily and quickly to data growth.

Table 1 through Table 4 summarize the support status of SAP HANA deployments on VMware vSphere, where:

- GA = generally available
- TBA = to be announced
- NA = not available
- Limited = As of December 2016, limited to one SAP HANA production instance co-deployed with non-production or non-SAP HANA VMs.

Check out SAP Note *1788665*–SAP HANA Support for virtualized/partitioned (multi-tenant) environments for the latest SAP HANA on VMware vSphere support status before planning a virtualized SAP HANA environment.

WORKLOAD	vSPHERE 5.5[4]	vSPHERE 6.0[5]	vSPHERE 6.5[6]
Support for production workloads	GA	GA	GA
Support for non-production workload	GA	GA	GA

Table 1: SAP HANA vSphere Version Support Status–as of December 2016[4, 5, 6]

[2] 12 TiB of RAM is supported on a specific OEM certified platform. Refer to VMware Hardware Compatibility Limits for guidance on the platforms that support vSphere 6.0 with 12 TiB of physical memory.

[3] SAP HANA supported RAM sizes depend on the used CPU and are limited (as of December 2016) with SPS11 to 512 GiB per CPU socket for BW on HANA and 768 GiB per CPU socket for Suite on HANA. For more details, see SAP Note 2227464 _ SAP HANA Platform SPS 11 Release Notes.

[4] SAP Note 1788665 – SAP HANA Support for virtualized/partitioned (multi-tenant) environments

[5] SAP Note 2315348 – Single SAP HANA VM on VMware vSphere 6.0 in production

[6] SAP note 2393917 – Single SAP HANA VM on VMware vSphere 6.5 in production

USE CASE–VSPHERE 5.5	BWOH	SOH/S4H
Support for single SAP HANA VM	GA	GA
Support for multiple SAP HANA VMs	GA	GA
Support for Scale-Out SAP HANA VMs	GA	NA[7]

Table 2: SAP HANA vSphere 5.5 Support Status per Use Case and Workloads–as of December 2016[7]

USE CASE-VSPHERE 6.0/6.5	BWOH	SOH/S4H
Support for single SAP HANA VM	GA	GA
Support for multiple SAP HANA VMs	TBA	Limited
Support for Scale-Out SAP HANA VMs	TBA	NA

Table 3: SAP HANA vSphere 6.0/6.5 Support Status per Use Case and Workloads–as of December 2016

SUPPORTED HW/CPU CONFIGURATION FOR SAP HANA	vSPHERE 5.5	vSPHERE 6.0	vSPHERE 6.5
Intel Xeon E7 CPU Support	Ivy Bridge (E7 –v2), Haswell (E7 –v3)	Ivy Bridge, Haswell, Broadwell (E7 –v4)	Broadwell (E7 –v4)
Intel Xeon E5 CPU Support	Ivy Bridge, Haswell, (E5 –v3)	Ivy Bridge, Haswell, Broadwell	Broadwell (E7 –v4)
CPU socket support	2 sockets with E5, with E7 up to 4 socket systems	2 sockets with E5, with E7 –v3 up to 8 and E7 –v4 up to 4 socket systems, 8 socket support in certification.	2 sockets with E5, with E7 –v4 up to 4 socket systems, 8 socket support in certification.

Table 4: SAP HANA on vSphere Supported Hardware Configurations–as of December 2016

Note: This document does not differentiate between vSphere 5.5 and 6.0/6.5 and different support situations. When a VMware virtualized SAP HANA gets deployed in production environments, verify with the above tables if the proposed vSphere Version has GA production support. More recent support information may be found in SAP Note 1788665 – *SAP HANA Support for VMware vSphere*.

[7] Scale-out SoH support: http://service.sap.com/sap/support/notes/1781986

Enabling Products and Technologies

The key enabling products and technologies for SAP and VMware include the SAP HANA platform and the VMware vSphere platform. In particular, this solution uses VMware vSphere advanced features including vMotion, VMware Distributed Resource Scheduler, and VMware High Availability. Each of these products and technologies is described in the following sections.

SAP HANA

SAP and partner solutions that use the market-leading SAP HANA in-memory data platform enable businesses to run in real-time. SAP HANA[8] is an in-memory data platform that is deployable as an on-premises appliance or in public and hybrid cloud deployments.

SAP HANA platform converges database and application platform capabilities in-memory to transform transactions, analytics, text analysis, predictive processing, and spatial processing to enable businesses to operate in real time. Figure 1 presents SAP HANA as "The Platform for All Applications" and shows the broad application, processing and integration services support available with the SAP HANA platform.

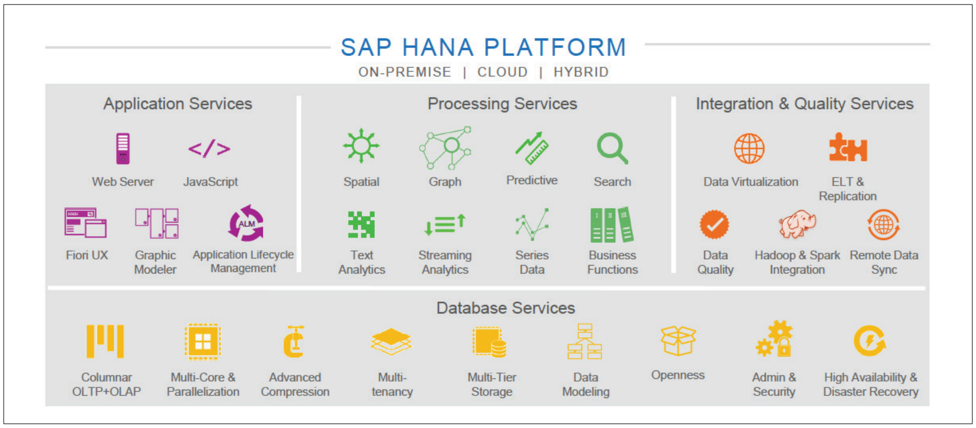

Figure 1: SAP HANA: The Platform for All Applications[9]

At the core of this real-time data platform is the SAP HANA database. The SAP HANA database is fundamentally different from any other database engine in the market today, as shown in Figure 2.

[8] SAP HANA Master Guide, page 4.

[9] SAP HANA Product Roadmap, Road Map Revision: 2016.02.17, pages 8 and 13.

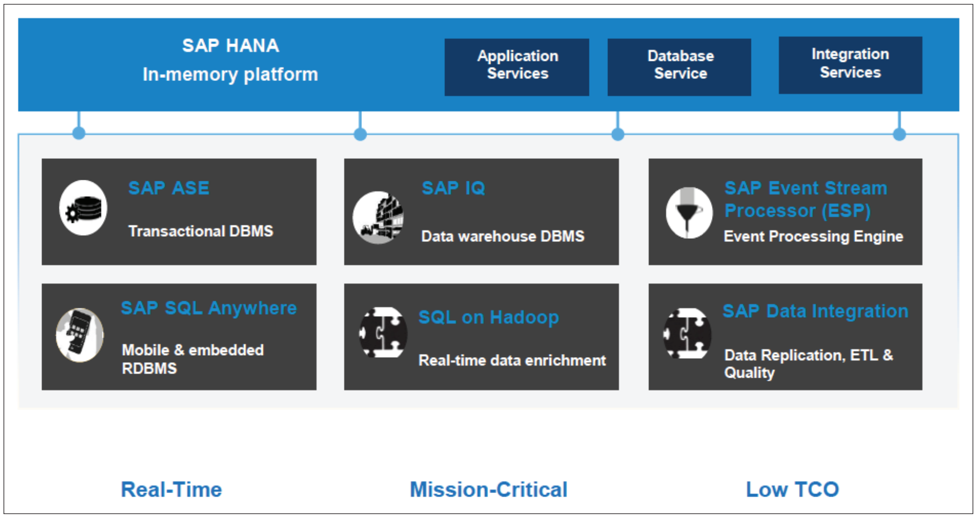

Figure 2: SAP HANA In-memory Platform for a New Class of Real-time Analytics and Mission-Critical Applications[10]

The SAP HANA architecture enables converged, online transaction processing (OLTP) and online analytical processing (OLAP) within a single in-memory, column-based data store using the ACID (Atomicity, Consistency, Isolation, Durability) compliance model, and it eliminates data redundancy and latency.

SAP HANA is well-suited for scenarios where businesses need to:

- Go deep within their data sets to ask complex and interactive questions
- Go broad, working with enormous data sets that are different types and from different sources at the same time

Increasingly, there is a need for this data to be recent and preferably in real-time.

In addition, businesses require this scenario to be performed at high speed (very fast response time and true interactivity), all without any pre-fabrication (no data preparation, no pre-aggregates, and no tuning).

In this type of scenario, only SAP HANA is well suited to address this unique set of requirements effectively. In fact, SAP HANA is at its best when fulfilling these requirements, or any subset these needs, in any combination.

Real-time Analytics

Examples of the real-time analytics SAP HANA can perform include:

- Operational reporting. Provides real-time insights from transaction systems such as custom or SAP ERP.
- Data warehousing (SAP NetWeaver® BW, powered by SAP HANA). BW customers can run their entire BW application on the SAP HANA platform, leading to unprecedented BW performance. For example, queries run 10 to 100 times faster, data loads 5 to 10 times faster, and calculations 5 to 10 times faster.

- Predictive and text analysis on Big Data. SAP HANA provides the ability to perform predictive and text analysis on large volumes of data in real-time using its in-database predictive algorithms and its R integration capability. SAP HANA text search/analysis capabilities provide a robust way to take advantage of unstructured data.

Real-time Applications

Examples of the real-time applications SAP HANA can perform include:

- Core process accelerators. Accelerate business reporting by using ERP Accelerators. They provide non-disruptive ways to take advantage of in-memory technology.
- Planning, optimization applications. SAP HANA excels at applications that require complex scheduling with fast results, and SAP is delivering solutions that no other vendor can match.
- Sense and response applications. These applications offer real-time insights on Big Data such as smart meter data, point-of-sale data, social media data, and more. They involve complexities such as personalized insight and recommendations, text search and mining, and predictive analytics. Typically, these processes are data-intensive, and many of these processes could not be deployed in the past due to cost and performance constraints.

OLTP Workloads

In addition to real-time analytic workloads (OLAP), SAP HANA is well suited to run traditional online transaction processing (OLTP) workloads. With Suite on HANA (SoH), SAP supports SAP HANA to run the SoH workload, in addition to regular databases in an integrated scenario. With SoH, scenarios that used to be performed on a relational database are now performed exclusively on an SAP HANA database. This enables the existing and optimized functions on an SAP HANA database to be used, instead of separating the analytics from the online processing.

For detailed SAP HANA information, see the SAP HANA Platform at: https://help.sap.com/hana_platform

SAP HANA Deployment Options

SAP HANA can be deployed in two different deployment options: Scale-Up or Scale-Out. Which deployment option an organization selects depends strongly on the planned use case, the application, and the size and data growth.

Business Warehouse workloads can be used through both deployment options, Scale-Up and Scale-Out, whereas currently the only deployment option for Suite on HANA workloads is the Scale-Up deployment option. When to Scale-Up or Scale-Out depends on two key components:

- Size of the source database and expected data growth
- Ability of the application to semantically divide and parallelize the processing of data

SAP recommends deploying SAP Business Suite applications in a Scale-Up configuration. Scale-Out configurations for SoH are not widely available yet for production workloads, and support for Suite on HANA workloads using Scale-Out must be verified by SAP. For the latest support updates for SoH on Scale-Out, refer to

SAP Note 2075461–*Suite on SAP HANA Scale-Up for large environments* and *SAP Note 1781986*[10].

The SAP HANA installation guide refers to these deployment options as SAP HANA single- or multiple-host systems (distributed).[11]

Scaling an SAP HANA system depends on the deployed system type:

- In a Scale-Up environment, you may add more compute resources:
  - RAM and CPU power to the physical host or
  - Extend the VM size to a maximum of 1 TiB RAM and 64 vCPUs with vSphere 5.5, and up to 4 TiB and 128 vCPUs with vSphere 6.0 and in the future with vSphere 6.5, once certified, up to 6 TiB and 128 vCPus per VM.
- In a Scale-Out environment, you can add SAP HANA nodes by:
  - Adding physical servers to a multiple-host configuration
  - Adding worker VMs

Figure 3 shows a simplified view of these two deployment options (Scale-Up and Scale-Out) on a VMware vSphere virtualized environment and Scale-Up or Scale-Out in a larger SAP HANA system.

The figure shows a physical 4 TiB system and four smaller 1024 GiB[12] systems, which provide the same amount of total memory; e.g., in the context of an SAP BW on HANA scenario. Adding more of these nodes is an easy way to extend the available SAP HANA BW memory.

The current RAM size limitations of a VM are 1 TiB for vSphere 5.5 and 4 TiB for vSphere 6.0. These maximum RAM sizes (minus approximately 0.5-3% RAM reserved for the usage of the hypervisor) and the current cluster limitations of 32 with vSphere 5.5 and 64 with vSphere 6.0 define the maximum RAM of an SAP HANA Scale-Out system configured when running virtualized.

[10] http://service.sap.com/sap/support/notes/2075461 and http://service.sap.com/sap/support/notes/1781986

[11] SAP HANA Master Guide. SAP HANA system types, page 26.

[12] GiB = Gibibytes = 2^{30} Byte

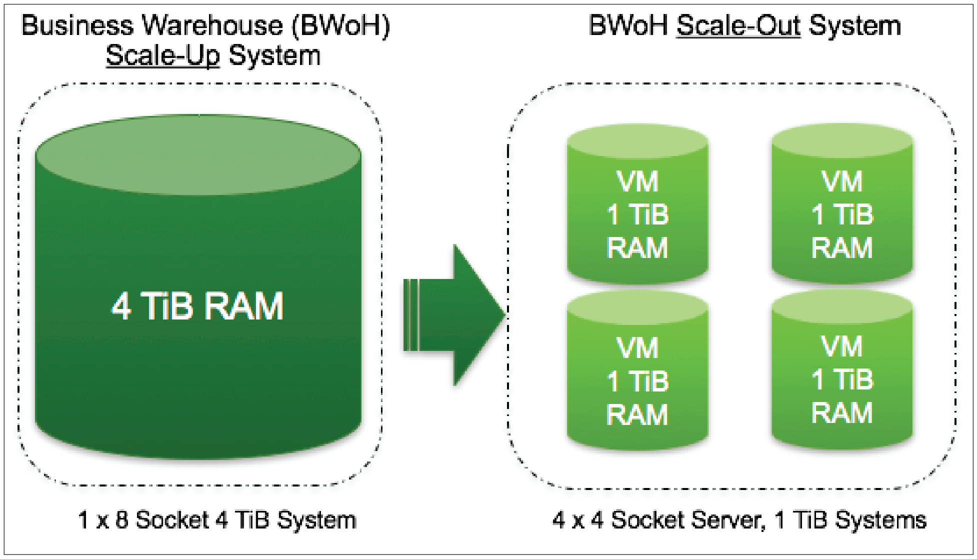

Figure 3: Scaling Options of an SAP HANA System on VMware vSphere

Both solutions provide their unique benefits, such as easier management or faster data access for a Scale-Up configuration, or cost-efficient expansion for greater storage requirements. For instance, it may be more expensive to expand the memory of a single server than it is to buy smaller servers with smaller, less expensive memory module sizes for use in a Scale-Out configuration.

Beside the purchasing costs of a physical server system, there also may be additional dependencies on hardware purchasing decisions driven by the organization's IT hardware standards and strategy. When deciding on scaling an SAP HANA system up or out, the overall costs of a server landscape–not just the initial hardware price–need to be considered.

When scaling an SAP HANA instance up or out, remember that RAM sizing must meet SAP sizing requirements and follow the currently existing fixed-CPU socket-to-RAM ratios for BW and Business Suite workloads, which will change over the time with support for newer CPU generations.

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 inside a single (but not necessary unique) VM on a vSphere host.

With a Scale-Up configuration, a single SAP HANA VM may consume all available host resources, or multiple SAP HANA VMs can share the host resources. For both options, single SAP HANA VM and multiple SAP HANA VMs, deployments are supported.

Once a VM reaches the virtual or physical host or vSphere limits, it cannot grow larger. Due to the SAP HANA CPU socket-to-RAM ratio, the biggest server systems and RAM sizes available for the BW on HANA (BWoH) workload are currently 8-socket server systems, with 4 TiB RAM and up to 8 TiB RAM for specific SAP Business Suite workloads[13]. Larger systems such as 16-socket server systems are also available, and allow the consolidation/co-deployment of several maximum-sized SAP HANA VMs on a single server system, but won't allow the deployment of a VM larger than the vSphere RAM or CPU maximums in a 8 socket partition. For a complete list of supported server configurations, visit the SAP HANA certified appliance web page.

Note: The current vSphere 5.5u2 virtual machine (VM) RAM limitations are 1 TiB and 6 TiB per physical host. For vSphere 6.0, the resource limits are 4 TiB per VM and 12 TiB per selected physical host and for vSphere 6.5 it would be 6 TiB per VM, whereas only 4 TiB got tested and certified as of writing of this document.

Figure 4 shows such a single SAP HANA instance running on a single VM that consumes all available server/VM resources. An SAP HANA system that already consumes the maximum server or VM limit won't be able to grow beyond these limitations. Adding resources or changing the VM resource configuration is required to support further growth to up to 1 TiB (vSphere 5.5) or 4 TiB (vSphere 6.0/6.5) RAM per VM.

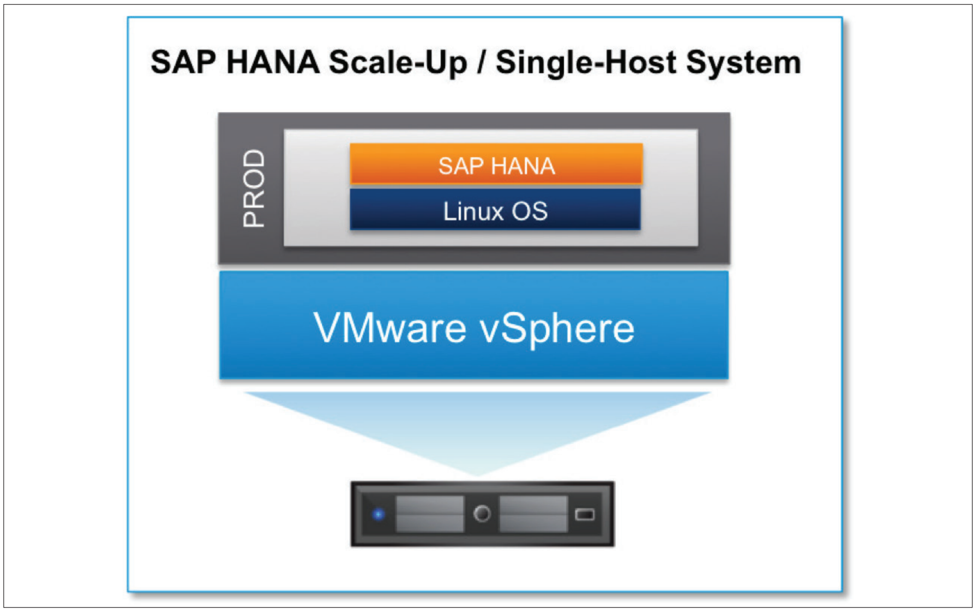

Figure 4: SAP HANA Single-Host System (Scale-Up) Configuration

Figure 5 and Figure 6 show Multi-VM configurations: single, independent SAP HANA systems running co-deployed on a single host system. SAP HANA multi-VM configurations that run on a single host are supported in production with general availability with vSphere 5.5. With vSphere 6.0/6.5, it is currently supported to run one production instance, dedicated or co-deployed with several non-production instances on a single host as shown in Figure 6. Joint VMware SAP Multi-VM configuration

[13] Review the SAP HANA Certified SAP HANA Hardware Directory, for the most current information on supported memory configurations. The shown memory sizes require an Intel E7-v3 CPU and SAP HANA SPS11.

certification tests with several production SAP HANA instances are ongoing and depending on the findings will get supported once testing is finished.

The host system is required to have enough resources (CPU, RAM, network, and storage) available to support all co-deployed VMs, and that CPU and memory resources of production VMs get reserved and not overcommitted. Non-production, non-performance critical workload can share resources to better utilize available hardware resources. If several production-level SAP HANA VMs are co-deployed on a single host, required resources must be reserved for all VMs.

Leveraging SAP HANA multi-VM support, virtualizing SAP HANA gets even more attractive. In addition to easier and more cost-efficient operation of a virtualized SAP HANA platform, now all the benefits of consolidation and co-deployment of SAP HANA instances help to lower the total cost of ownership (TCO) of an SAP HANA environment, allowing now IT organizations to deploy SAP HANA as just any other business-critical application by maintaining the required high performance and SLAs.

Figure 5 shows the consolidation of flexible-sized production level SAP HANA VMs that run co-deployed, fully isolated with the required resources committed on a single SAP HANA supported server. The figure shows a vSphere 5.5u2 host, which supports host RAM sizes up to 6 TiB and VM sizes of up to 1024 GB.

Note: As of May 016, vSphere 6.0/6.5 Multi-VM support is limited to one production workload SAP HANA VM and several non-production level VMs. Figure 5with four production-level VMs is currently supported only with vSphere 5.5.

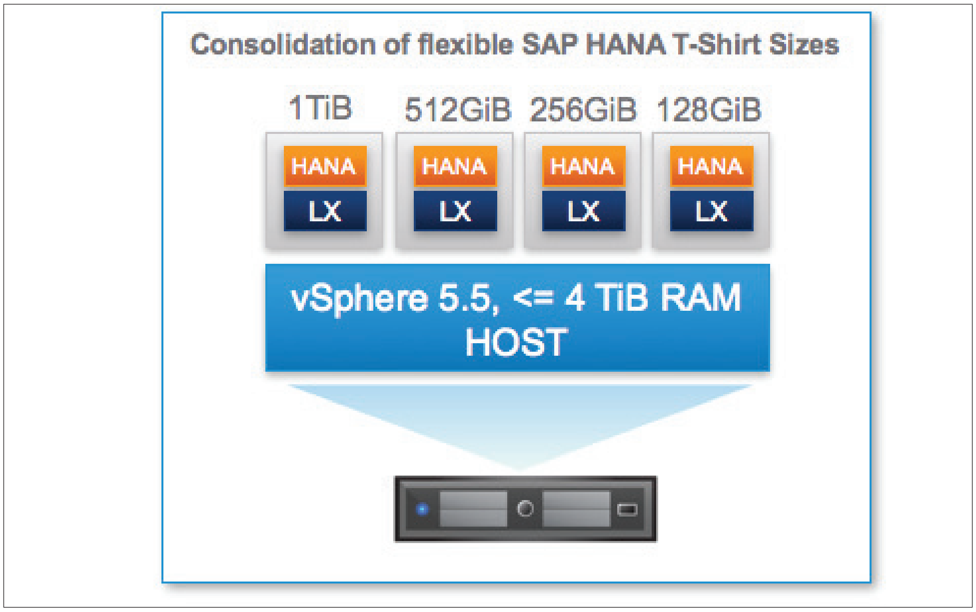

Figure 5: SAP HANA Multi-VM Single-Host System (Scale-Up) Configuration

Figure 6 shows a consolidated landscape scenario, where dedicated parts of a server (NUMA nodes) are used for an SoH production instance, and the rest of the server is shared by non-production instances such as development and test systems. This scenario illustrates how a larger server such as an 8-socket, 6–8 TiB RAM server can be used very efficiently to provide resources for a complete SAP HANA prod/test/dev landscape. The HA server needed to protect the production SAP HANA instance could be a 4-socket, 3 TiB server with VMware HA affinity rules configured because only this instance will require HA. (There is further information regarding HA in the HA section of this document.) This again helps to lower the TCO of an SAP HANA environment by not only consolidating workloads, but also by using server systems better aligned to the workload instead of buying fixed standard T-shirt sized configurations.

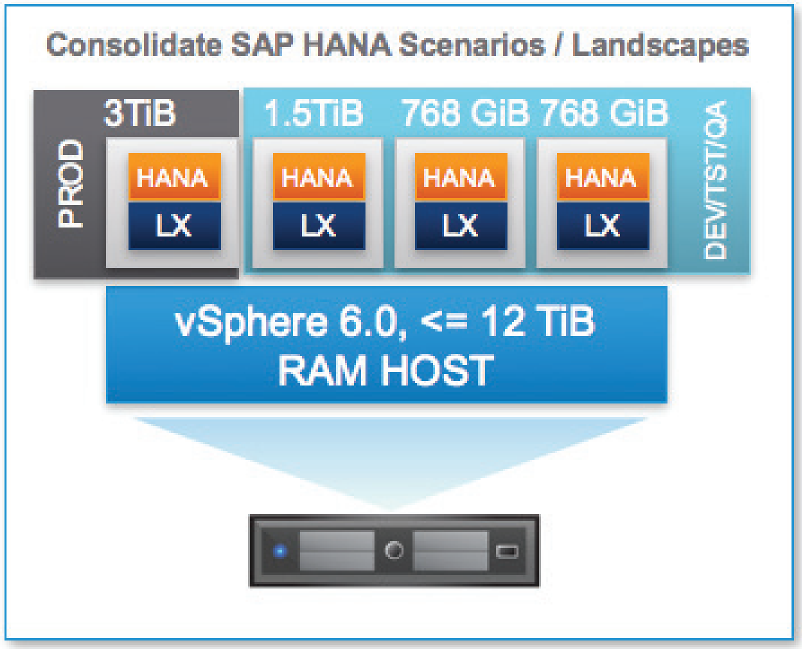

Figure 6: SAP HANA Landscape Consolidation Example

Note: The scenario in Figure 6 is fully supported, and GA and allows to consolidate production-level SAP HANA VM with several non-production level VMs with vSphere 6. The host limitation is currently 4 sockets and the maximum 4 TiB physical RAM for SoH/S4H on Broadwell CPU based server systems. 8 socket server systems with larger physical RAM sizes are in certification.

SAP HANA Scale-Out Option

In an SAP HANA Scale-Out configuration, several SAP HANA instances are connected together into one large, distributed SAP HANA database system. Running on vSphere, these SAP HANA "worker VMs" are distributed over multiple vSphere hosts, while the number of physical hosts may be smaller than the number of actual VMs or even just one. All SAP HANA worker VMs together to comprise a single, distributed SAP HANA system.

The benefit of this configuration is that it is possible to grow the SAP HANA database size over time by adding more SAP HANA VMs and, if needed, additional physical servers. Once a new SAP HANA VM is available, it is possible to add it to an existing SAP HANA Scale-Out configuration via the SAP HANA installer program.

SAP HANA certified Scale-Out configurations are available from 4 to 94 nodes with 2- to 8-CPU sockets, and between 128 GiB and 4 TiB of RAM (with Intel Xeon E7-v4 CPUs). Currently, Scale-Out configurations for production workloads are limited to SAP BW workloads only. Business Suite workloads are not yet supported. For some specific scenarios, SAP supports pilot phases and CA programs (native and virtualized). For up-to-date details and status, refer to SAP Note 1825774. Also, support for more than 16 SAP HANA worker nodes in BWoH cases are available upon request from SAP.

In a physical Scale-Out configuration, a standby host (working as a "hot-spare") is required to implement a failover. In a virtualized environment, if you use VMware HA, it is required to have only the compute resources available to support a failover and restart process of a VM or a complete host. Unlike with a physical setup, a dedicated standby host is not required, and standby resources can be used for other workloads.

Figure 7 shows such a configuration, where every VMware ESXi™ host supports a single SAP HANA VM. The storage configuration is not shown here, and requires a shared TDI storage subsystem. To ensure enough resources are available in the event of a failure, hosts must be configured to provide free resources or resources used by non-critical VMs (shown in the figure as "Free or non-prod workload"), which have to be migrated or shut down before the actual failover and restarted of a failed SAP HANA Scale-Out VM. (See more information on how to handle Scale-Out VM failover in the HA section of this guide.)

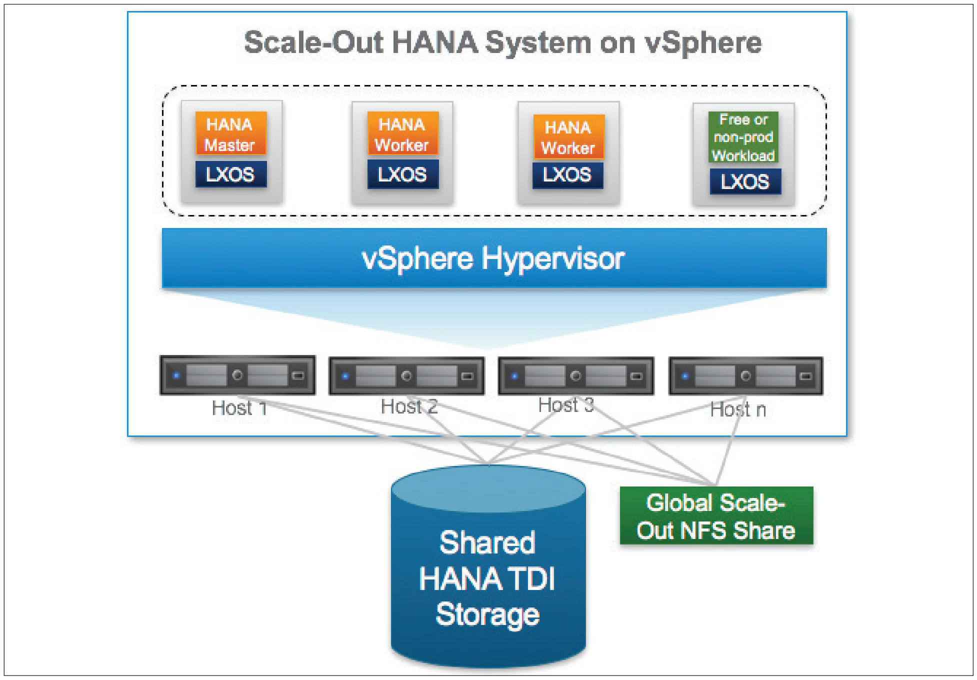

Figure 7: SAP HANA Multiple-host System (Scale-Out) on vSphere Configuration

Unlike in a Scale-Up configuration, shared storage is required to provide all workers access to the central SAP HANA directory and configuration files.

When the configured ESXi hosts are able to provide enough resources, you can also co-deploy other SAP HANA Scale-Out systems on the same physical vSphere ESXi host cluster (or any other workload that an organization may want to run in virtualized mode). Figure 8 shows such a configuration, where two independent SAP HANA systems (1 and 2) are running on several vSphere ESXi hosts that work together as two independent SAP HANA systems (SAP HANA System 1 and 2). When selecting the SAP HANA system—or any other virtualized workload that has overlapping resource peaks and valleys—it is possible to optimally utilize all available hardware resources of the vSphere cluster.

The co-deployment of SAP HANA Scale-Out systems allows the consolidation of multiple production SAP HANA Scale-Out systems or the consolidation of a complete SAP HANA landscape including test and development systems. In this example, the failover of a complete host can be compensated because enough spare resources are available in the vSphere cluster. The spare resources could also be used for uncritical business VMs, which may be shut down at any time to free up the resources when a HANA VM needs to be failed over. The used storage must be able to provide enough resources (capacity and I/O performance) to allow the co-deployment of multiple HANA systems. If only one SAP HANA instance is used for production workloads, then the storage must be able to provide the SAP HANA storage-related KPIs to the SAP HANA VM at any time. Additionally to tis, the storage must have enough resources available to support any other possible co-deployed VM.

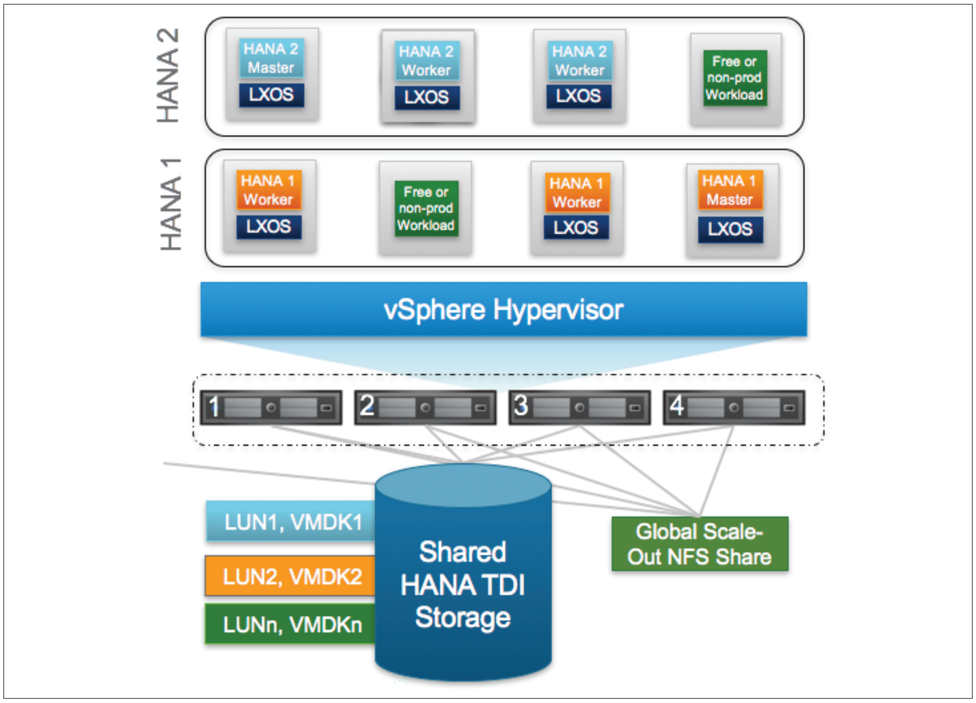

Figure 8: SAP HANA Scale-Out/Multiple VM Landscape Consolidation

For detailed SAP HANA deployment information, refer to SAP HANA documentation at: https://help.sap.com/ hana_platform or review the SAP HANA Master Guide[14] and SAP HANA Server Installation and Upgrade Guide.[15]

Follow the installation steps in the Best Practices and Configuration Settings section how to install a Scale-Out configuration on vSphere.

SAP HANA Appliance Model

The SAP HANA appliance is a flexible, multi-purpose, data-source-agonistic in-memory appliance that combines SAP software components with Intel based hardware, optimized and delivered by SAP's leading partners such as Hitachi Data Systems (HDS), VCE, Dell, Cisco, Lenovo, Hewlett Packard Enterprise (HPE), NEC, Fujitsu, and others. For a list of certified SAP HANA appliances and supported hardware components, visit the "SAP HANA Hardware Directory".[16]

The SAP HANA appliance delivers the following capabilities:

- Single SAP HANA database with native support for row and columnar data stores, providing full ACID (atomicity, consistency, isolation, durability) transactional capabilities
- Powerful and flexible data calculation engine

[14] http://scn.sap.com/docs/DOC-60318

[15] http://help.sap.com/hana/SAP_HANA_Server_Instal lation_Guide_en.pdf

[16] http://global12.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/appliances.html

- SQL and MDX interfaces
- Unified information modeling design environment
- Data repository for persistent views of business information (you can unplug the power and your information is safely stored)
- Data integration capabilities for accessing SAP (SAP NetWeaver BW, ERP, etc.) and non-SAP data sources
- Integrated lifecycle management capabilities
- Low TCO, as it is optimized for commodity hardware from a large number of partners

A VMware virtualized SAP HANA system can—but doesn't have to—come preconfigured as an SAP HANA appliance, such as those offered by HPE, Dell, or HDS and other vendors. The minimum requirements are:

- The OEM hardware must be supported by SAP for SAP HANA workloads.
- The server hardware and operating system version must correspond to the bill of material (BOM) of a certified SAP HANA appliance or is listed as a supported SAP HANA server system.
- The selected server must be VMware certified.

In addition to the core server system, Scale-Out SAP HANA appliances must also include the required components to interconnect multiple SAP HANA hosts. Unlike Scale-Up configurations, a shared storage area network (SAN) or a network-attached storage (NAS) storage system is used instead of local attached disks, and non-blocking 10 GB Ethernet network switches must be used.

Note: Today, most virtualized SAP HANA systems get deployed as TDI configurations and support Intel Xeon E7 and E5 CPU- based server systems, which are SAP HANA supported.

SAP HANA Tailored Datacenter Integration (TDI)

The SAP HANA Tailored Datacenter Integration option, also called TDI, allows customers to use certain parts of their existing hardware and infrastructure components for the SAP HANA environment. Typically, an SAP HANA appliance comes with all of the necessary components preconfigured, as provided by certified SAP HANA hardware partners. TDI targets the usage of certain hardware and infrastructure components that might already exist in a customer's landscape, instead of using the corresponding components that are delivered with a HANA appliance.

Depending on the existing data center layout, SAP HANA TDI can:

- Reduce hardware and operational costs by reusing existing hardware components and operation processes.
- Mitigate risks and optimize time-to-value by enabling existing IT management processes for the SAP HANA implementation.
- Deliver more flexibility in hardware vendor selection by making the best use of the existing ecosystem.
- Enable an easier and more cost-effective migration to SAP HANA.
- Allow the leveraging of VMware vMotion, DRS, and HA features.

SAP HANA TDI is the VMware preferred deployment model because many VMware features require shared storage and external network devices. As of today, no vendor offers a SAP certified VMware virtualized appliance, therefore all VMware SAP HANA configurations are TDI configurations.

Figure 9 describes the appliance delivery model and the TDI model for SAP HANA systems. The graphic for the appliance model shows all components in a box, and shows how these components are all preconfigured, tested, and certified as an appliance. If the virtualization layer is part of the appliance certification, then it would be treated like an appliance. The TDI model shows the different components, such as the server and storage, loosely "coupled" to highlight that a customer can choose from any supported vendor, but is responsible for the installation and implementation of the overall system.

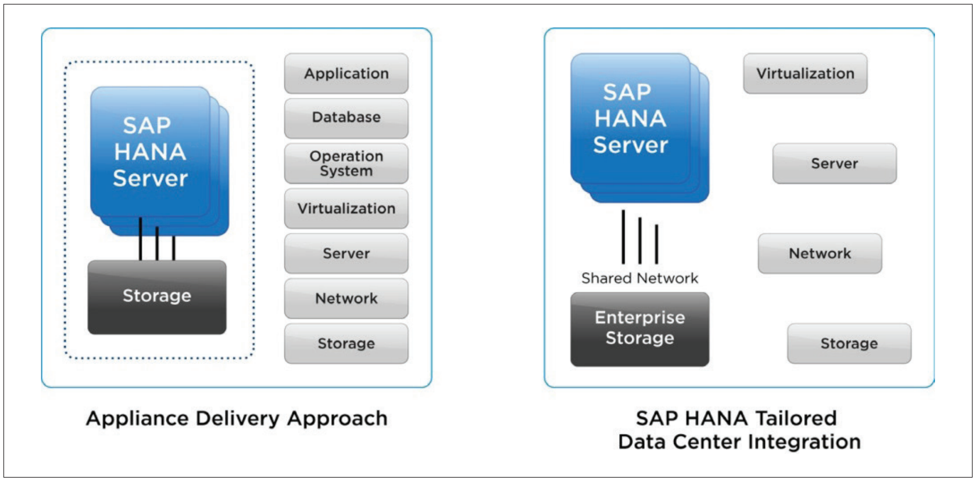

Figure 9: SAP Tailored Datacenter Integration (TDI)[17]

In Scale-Out configurations, 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[18] document.

SAP HANA TDI Support Process

In the SAP HANA appliance delivery model, the hardware vendor validates and ensures the HANA readiness of the solution. In the SAP HANA TDI approach, it is the responsibility of the customer to install and implement HANA in its environment. SAP HANA implementation partners are available to perform this task for a customer.

[17] SAP HANA Tailored Data Center Integration Frequently Asked Questions, page 4

[18] http://scn.sap.com/docs/DOC-62942

If support is required for virtualized SAP HANA systems, customers can open a ticket directly at SAP. The ticket will, if needed, get routed to VMware specialized support engineers for SAP HANA, who will then troubleshoot the escalated issue. This support process ensures fast problem resolution and is different from the general TDI support process for non-virtualized SAP HANA TDI installations. In those instances, hardware/appliance vendors take the lead role in resolving SAP support tickets for problems related to SAP HANA infrastructure and installation. Once VMware is involved then VMware takes the lead instead of the hardware vendor.

SAP provides the "HW Configuration Check Tool for SAP HANA,"[19] a method for checking the minimum number of KPIs for HANA readiness in the installed environment.

Note: If case support is required for virtualized SAP HANA systems, customers can open a ticket directly at SAP. The ticket will, if needed, get routed to VMware specialized support engineers for SAP HANA, who will then troubleshoot the escalated issue. This ensures that the customer has to open only a ticket at SAP and gets the needed support.

- The "HW Configuration Check Tool for SAP HANA" provides tests for customers to determine if the hardware configuration meets the minimum performance criteria required to run SAP HANA in production. This tool also allows ongoing checks and measurements to ensure required performance levels when additional SAP HANA VMs are deployed on the same hardware. In the past, it was mandatory to run this tool; now, it is the customer's responsibility to use this tool to verify the SAP HANA landscape.
- SAP provides go-live-checks and other support services for SAP HANA customers as part of their SAP support license. For details, review the SAP AGS IT Planning webpage for more details.
- In addition, VMware provides SAP HANA health check and architecture reviews for virtual deployments on top of vSphere. For details, visit the VMware pages on SAP SCN at http://scn.sap.com/docs/DOC-60470.
- Freedom of choice of support contracts allows organizations to harmonize SAP HANA support with other support processes and contracts an organization may already have in place for other business-critical applications, such as existing SAP ERP or database applications.

VMware vSphere

VMware vSphere is the market-leading virtualization platform and enabling technology for cloud computing architectures. vSphere enables IT to meet service level agreements (SLAs) for the most demanding business-critical applications at the lowest total cost of ownership (TCO). VMware vSphere delivers control over all IT resources with the highest efficiency and choice in the industry.

[19] http://service.sap.com/sap/support/notes/1943937

VMware vSphere virtualization solutions provide:

- **Consolidation.** VMware virtualization allows multiple application servers to be consolidated onto one physical server, with little or no decrease in overall performance. This helps to minimize or eliminate underutilized server hardware, software, and infrastructure.
- **Manageability.** The live migration of virtual machines from server to server and the associated storage is performed with no downtime using VMware vSphere vMotion, which simplifies common operations such as hardware maintenance, and VMware vSphere® Storage vMotion®.
- **Availability.** High availability can be enabled to reduce unplanned downtime and enable higher service levels for applications. VMware vSphere High Availability (HA) ensures that, in the event of an unplanned hardware failure, the affected virtual machines are automatically restarted on another host in a VMware cluster.
- **Automation.** VMware automated load balancing takes advantage of vMotion and Storage vMotion to migrate virtual machines among a set of VMware ESX® hosts. VMware vSphere Distributed Resource Scheduler (DRS) and VMware vSphere® Storage DRS™ enable automatic resource relocation and optimization decisions for virtual machines and storage.
- **Provisioning.** VMware virtualization encapsulates an application into an image that can be duplicated or moved, greatly reducing the cost of application provisioning and deployment.
- **Control.** Single pane of glass control, management, and operation of the complete virtualized infrastructure via the VMware vCenter® server.

Figure 10 shows a VMware vSphere virtual infrastructure, and the different layers and components of such an environment. SAP or any other applications VMs, which run virtualized, on top of a virtualized infrastructure resource pools of compute, storage and network on VMware vSphere. These VMs can be operated with very little effort with vCenter. The vCenter Server is the single pane of glass management system that allows VMware administrators to configure, manage, and operate the complete virtualized infrastructure. For SAP HANA administrators, this is transparent because they would use SAP HANA Studio to manage the actual SAP HANA installation inside a VM.

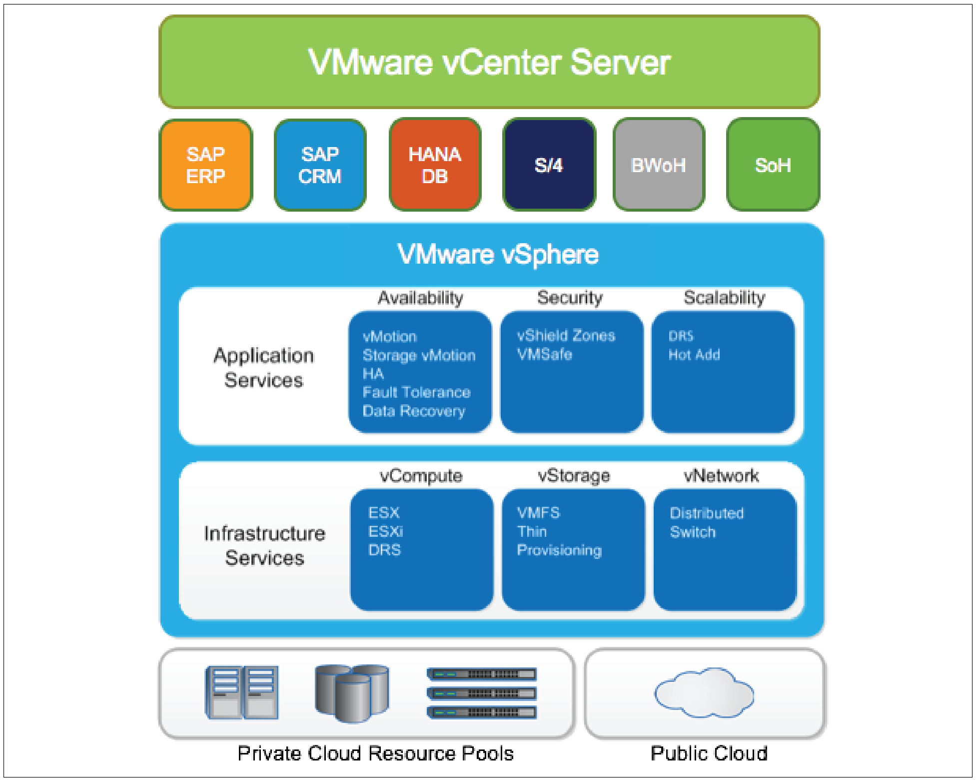

Figure 10: VMware vSphere Virtual Infrastructure

VMware vSphere creates a layer of abstraction between the resources required by an application, the operating system, and the underlying hardware that provides those resources. vSphere enables multiple, isolated execution environments to share a single hardware platform. It implements each environment with its own set of hardware resources.

For more information about vSphere, see the "Resources" section in this document.

With the release of vSphere 6.0 and 6.5, larger systems are now supported, and new features and enhancements for storage, networking, availability or manageability were introduced.

VMware vSphere 6.x–New Features

Note: Please review the "What's New in VMware vSphere 6.5–Technical Whitepaper" see the VMware Compatibility Guide[20] for details. As of writing this guide vSphere 6.5 was not released and therefore this section focuses on vSphere 6.0.

[20] http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/whitepaper/vsphere/vmw-white-paper-vsphr-whats-new-6-5.pdf

As a summary, VMware vSphere® 6.5 offers simplified customer experience for automation and management at scale, comprehensive built-in security for protecting data, infrastructure, and access and an universal application platform for running any application anywhere. With vSphere 6.5, customers can now run, manage, connect, and secure their applications in a common operating environment, across clouds and devices.

Relevant new features or enhancements[21] of vSphere 6.0 for SAP HANA include:

- Compute:
  - Increased Scalability–Increased configuration maximums, where virtual machines support up to 128 virtual CPUs (vCPUs) and 4 TiB virtual RAM (vRAM). Hosts support up to 480 CPUs and 12 TiB of RAM, 1024 virtual machines per host, and 64 nodes per cluster.
  - Expanded Support–Expanded support for the latest x86 chip sets, devices, drivers, and guest operating systems. For a complete list of guest operating systems, hardware components, and CPUs supported, see the VMware Compatibility Guide.[22]
  - Instant Clone–Technology built into vSphere 6.0 that lays that foundation to rapidly clone and deploy virtual machines, as much as 10X faster than what is possible today.
- Storage:
  - Transform Storage for your Virtual Machines–VMware vSphere® Virtual Volumes™ (VVOLs) enables your external storage arrays to become VM-aware. Storage Policy-Based Management (SPBM) allows common management across storage tiers and dynamic storage class of service automation. Together, they enable exact combinations of data services (such as clones and snapshots) to be instantiated more efficiently on a per VM basis.
- Network:
  - Network I/O Control–New support for per-VM Distributed vSwitch bandwidth reservations to guarantee isolation and enforce limits on bandwidth.
  - Multiple TCP/IP Stack for vMotion-Allows vMotion traffic a dedicated networking stack. Simplifies IP address management with a dedicated default gateway for vMotion traffic.
- Availability:
  - vMotion Enhancements–Perform nondisruptive live migration of workloads across virtual switches and vCenter Servers and over distances of up to 150 ms RTT. The astonishing 10X increase in RTT offered in long-distance vMotion[23] now makes it possible for data centers physically located in New York and London to migrate live workloads between each another.
- Management:
  - Cross-vCenter Clone and Migration–Copy and move virtual machines between hosts on different vCenter Servers in a single action.

[21] http://www.vmware.com/files/pdf/vsphere/VMW-WP-vSPHR-Whats-New-6-0-PLTFRM.pdf

[22] http://www.vmware.com/resources/compaTiBility/search.php

[23] http://blogs.vmware.com/performance/2015/02/vmware-pushes-envelope-vsphere-6-0-vmotion.html

For detailed information of the new features of vSphere 6.0 and 6.5, download review the vSphere product page.[24]

vSphere 5.5 vs 6.0

The most relevant vSphere 6.0 enhancement is the increased scalability that allows SAP HANA VMs to be as big as 4 TiB with vSphere 6.0 and 6.5 and once certified, up to 6 TiB with vSphere 6.5. Besides the features listed above and the increased scalability, vSphere 6.x works just as fine as vSphere 5.5. When migrating a vSphere 5.5 based SAP HANA VM to vSphere 6.x, the VM configuration settings must be updated and where needed, changed to run optimally on vSphere 6.x. Review the best practices and configuration parameter section at the end of this document for guidelines. Some changes to the VM configuration, such as upgrading the VM's hardware level, will require a restart of the VM.

VMware internal performance tests have shown that an identically configured SAP HANA vSphere 6.x VM, running on the same server system, has a similar performance to a vSphere 5.5 VM. The tests have shown little performance benefits for vSphere 6.x in the area of memory management. vSphere 6.x VMs get access to more compute resources and provide more resources to an application, which is then able to consume double the amount of CPU threads (128) instead of up to 64 CPU threads with vSphere 5.5.

HPE has performed recent testing that shows a 6 percent performance gain when using vSphere 6.0 instead of vSphere 5.5 on the same hardware configuration (see Figure 11). The figure shows the results of an HPE ConvergedSystem 500 based on the HPE ProLiant DL580 Gen9 Server running the Intel Xeon E7 v3 processor configured with VMware vSphere 5.5 and 64 vCPUs, compared to an identical HPE ConvergedSystem 500 configured with VMware vSphere 6.0 and 72 vCPUs. The increase of vCPUs was possible due to vSphere 6.0. For details about this test, review the HPE document "Virtualized SAP HANA performance evolution with the HPE Converged System 500 and VMware vSphere 6.0."[25]

[24] http://www.vmware.com/products/vsphere/

[25] "Virtualized SAP HANA performance evolution.." http://h20195.www2.hp.com/v2/getpdf.aspx/4AA6-6194ENW.pdf

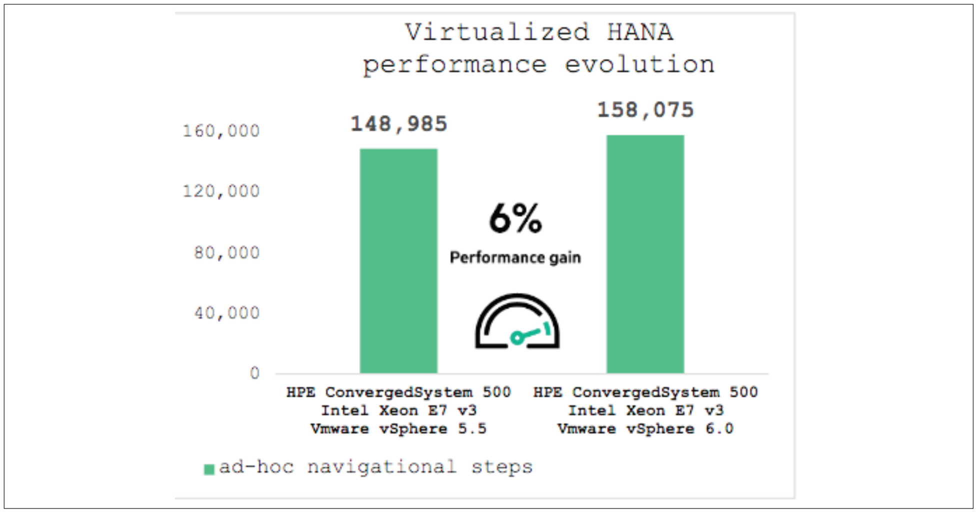

Figure 11: HPE vSphere 5.5 vs 6.0 SAP HANA BW-EML Benchmark Results

This demonstrates that the vSphere 6.0 SAP HANA VMs can be more powerful as vSphere 5.5 SAP HANA VMs.

Figure 12 compares vSphere 5.5 and 6.0, with a focus on scalability.

	vSphere 5.5u2	vSphere 6.0	vSphere 6.5
Hosts per Cluster	32	64	64
VMs per Cluster	4,000	8,000	8,000
CPUs per Host	320	480	576
RAM per Host	6 TB	12 TB[1]	12 TB[1]
VMs per Host	512	1,024	1,024
Virtual CPUs per VM	64	128	128
Virtual RAM per VM	1 TB	4 TB	6 TB[2]

Figure 12: Comparison between vSphere 5.5[26], 6.0[27] and 6.5[28], *1, with specific OEM partners, *2 as of writing this guide only 4 TiB are currently supported. 6 TiB tests are on the roadmap.

26 https://www.vmware.com/pdf/vsphere5/r55/vsphere-55-configuration-maximums.pdf, up to 6 TB is supported for ESXi 5.5 update 2 and later

27 https://www.vmware.com/pdf/vsphere6/r60/vsphere-60-configuration-maximums.pdf

28 https://www.vmware.com/pdf/vsphere6/r65/vsphere-65-configuration-maximums.pdf

On the following pages, some of the SAP HANA relevant vSphere features will be introduced. The features that help operate an SAP HANA system on VMware vSphere are highlighted.

VMware Resource Isolation and Security–Basis for SAP HANA Multi-VM

From its core up to its management interfaces, the vSphere hypervisor is designed to run multiple, independent co-deployment VMs. These VMs run side by side on the same physical infrastructure (e.g., compute and storage) by maintaining high security standards and by enforcing strict resource commitment policies as defined by the customer.

The SAP HANA TDI deployment option allows an organization to deploy several SAP HANA systems in production and non-production use cases, depending on available compute resources; this can include non-SAP HANA mixed workloads (OLTP and OLAP) running on the same physical server when virtualized with VMware vSphere.

Co-deployment and consolidation of independent workloads running on different VMs is one of the primary use cases of virtualization, but must follow strict isolation and encapsulation of the memory and instructions a VM is using. The VMware document "Security of the VMware vSphere Hypervisor"[29] describes it this way:

> "Isolation of CPU, memory, and I/O now is done at a hardware level, with the vSphere ESXi hypervisor managing how much of the hardware resources a virtual machine can use, similar to a choreographer or traffic officer. This part of the hypervisor is called the virtual machine monitor (VMM). With the ability to leverage these CPU extensions, the attack surface of the hypervisor shrinks considerably.
>
> Memory pages that are identical in two or more virtual SAP HANA machines are stored once in the host system's RAM, and each of the virtual machines has read-only access. Such shared pages are common, for example, if many virtual machines on the same host run the same OS. As soon as any one virtual machine attempts to modify a shared page, it gets its own private copy. Because shared memory pages are marked copy-on-write, it is impossible for one virtual machine to leak private information to another through this mechanism. Transparent page sharing is controlled by the VM kernel and VMM and cannot be compromised by virtual machines. Transparent page sharing can be disabled on a per-host or per-virtual machine basis."

Figure 13 illustrates the shared and nonshared/strictly isolated components of a virtualized environment running on vSphere. Shared components are the underlying physical hardware/infrastructure components, such as network devices. Memory segments or instructions are not shared and are strictly isolated between VMs.

[29] http://www.vmware.com/files/pdf/techpaper/vmw-wp-secrty-vsphr-hyprvsr-uslet-101.pdf, page 5

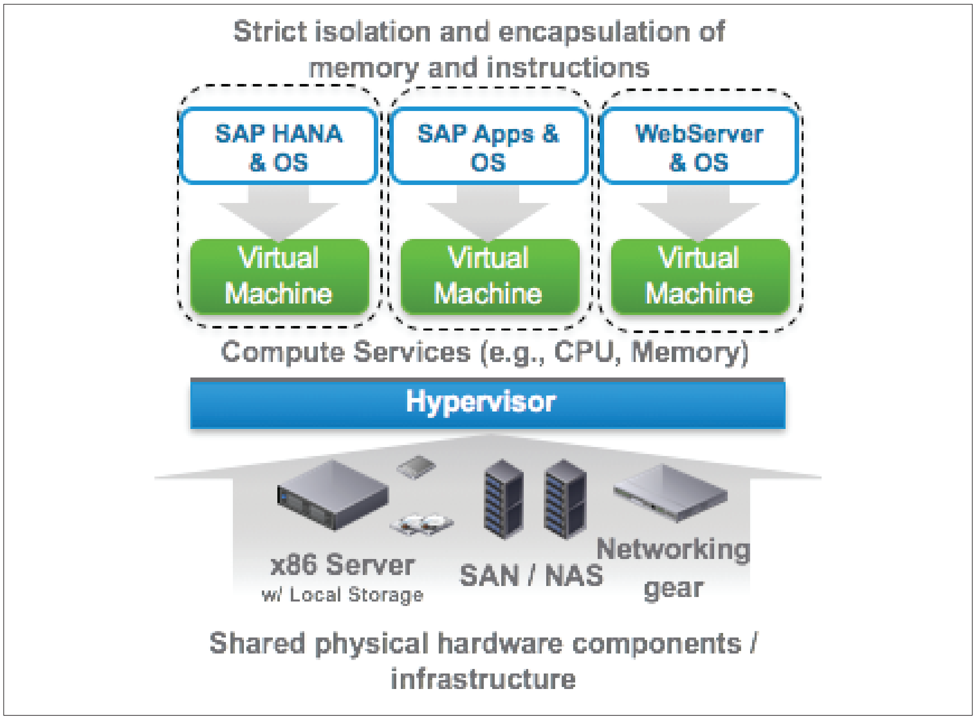

Figure 13: Co-deployment and Consolidation of Independent and Different SAP Workloads on a Server

Using Transparent Page Sharing will reduce the amount of RAM of co-deployed, similar VMs. For instance, the sharing of RAM pages used by the operating system (OS) and SAP kernel can help reduce the overall consumed RAM on a vSphere host, and would allow the deployment of larger systems on a single vSphere host. Figure 14 shows this concept, as discussed in the VMware document, "Security of the VMware vSphere Hypervisor."

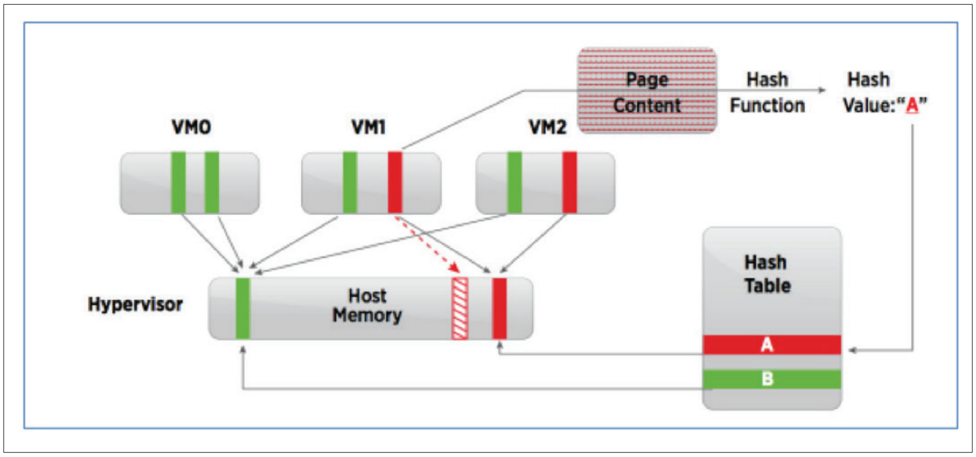

Figure 14: Transparent Page Sharing–Page-Content Hashing

For SAP HANA VMs to run in a secure and protected mode, instructions and memory must be isolated; this is fully provided and supported by VMware vSphere. Therefore, you can safely run as many concurrent SAP HANA VMs as your underlying hardware is able to support.

Besides CPU instruction and memory isolation, virtualized device drivers are also limited to instruction and memory elements of their VM, and are unable to access information from other VMs.

In a virtualized environment, compute resources usually are shared among all virtual machines, but also can be configured to be dedicated to one particular virtual machine. To avoid resource overcommitments, SAP HANA systems that run in production must be configured with strict resource reservations, or run the risk of being affected by resource constraints caused by the co-deployment of other VMs.

Note: By default, co-deployments of SAP HANA VMs–and the operation of transactions within the VMs–is secure, reliable, and protected by vSphere.

Virtual Networking

Virtual network communication and security works the same as it does with any other physical OS or application that is natively installed on a server and connected to a network. The primary difference is that it is much easier to isolate and protect network traffic of SAP HANA VMs running on vSphere by leveraging the VMware provided software for virtual network and solutions.

Figure 15 shows how network communications can become limited and secured by leveraging the built-in hypervisor network capabilities in a way that is similar to virtual switches.

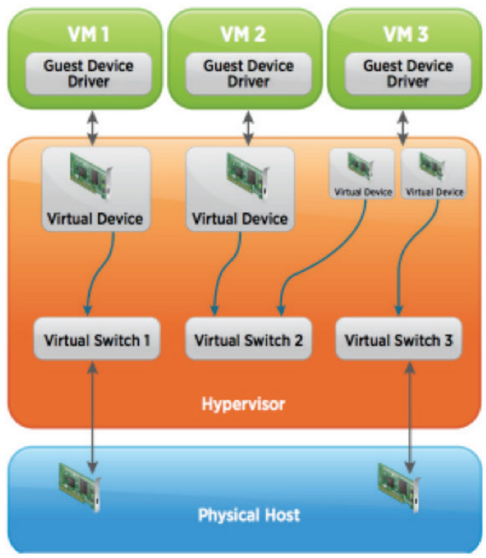

Figure 15: Network Isolation of a Virtualized Server Environment

Virtual machine network isolation, as shown in Figure 15 works similarly to this:

- Virtual Machine 1 (VM 1) does not share a virtual switch with any other virtual machines within the host. VM 1 is therefore completely isolated from other virtual networks within the host, but is able to communicate via Virtual Switch 1 to any network devices connected to the same network.
- If no physical network adapter is configured for Virtual Machine 2 (VM 2), then the virtual machine is completely isolated from any physical networks. In this example, the only access to a physical network is if Virtual Machine 3 acts as a router between Virtual Switch 2 and Virtual Switch 3 (VM 3). The communication between VM 2 and VM 3 gets done internally in the host server with network bandwidth outperforming physical networks.
- A virtual machine can span two or more virtual switches only if configured by the administrator. Virtual Machine 3 shows such a configuration.

Through the use of a virtualized network interface card (vNIC)–level firewall, a virtual machine can be isolated from other virtual machines, even on the same switch (using "Layer 2 isolation").

Note: Consolidation, co-deployment and multi-tenant SAP HANA environments require high security and resource control. A VMware virtualized SAP HANA environment provides these attributes, and allows secure, reliable, and controllable co-deployments of multiple, independent SAP HANA systems.

VMware vMotion (Live Migration)[30]

Perform Live Migrations

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

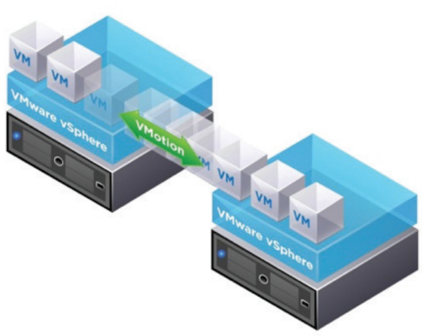

Figure 16: VMware vMotion

[30] https://www.vmware.com/products/vsphere/features/vmotion

Note: During a vMotion process, which can take over one hour, the response time of an SAP HANA database that runs inside a VM may get delayed and, at worst case, time-critical transactions that depend on predefined and timely responses may fail because of this delay. Live migration tasks are therefore recommended to be performed during off-peak hours! Tests during the SAP HANA vSphere validations have shown that live migration of SAP HANA VM, with around 500 GiB, takes 5–6 minutes to complete on a dedicated 10 gigabit Ethernet network.

vMotion is a key enabling technology for creating the dynamic, automated, and self-optimizing data center. This capability makes hardware maintenance possible at any time, and vMotion does not require clustering or redundant servers, as described below. VMware vMotion provides following features and benefits:

VMware vMotion:

- Moves entire running virtual machines and performs live migrations with zero downtime, undetectable to the user
- Manages and schedules live migrations with ease at predefined times, without an administrator's presence, and with the reliability and manageability that is derived from a production-proven product
- Performs multiple concurrent migrations of a virtual machine running any operating system, across any type of hardware and storage that is supported by vSphere, complete with an audit trail
- Moves online workloads from one ESXi server host machine to another in order to maintain service levels and performance goals
- Continuously and automatically optimizes, in combination with DRS, virtual machine placement within resource pools. Proactively moves virtual machines away from failing or underperforming servers
- Performs hardware maintenance without the need to schedule downtime and disrupt business operations

Note: VMware vMotion is not a High Availability solution preventing downtime because of failures. Instead, VMotion helps to reduce or completely avoid planned downtime by allowing for live migration of running applications.

Technical Details[31]

Using the Virtual Machine File System (VMFS)

The entire state of a virtual machine is encapsulated by a set of shared storage files such as Fibre Channel (FC), iSCSI storage area network (SAN), or network attached storage (NAS). VMware vSphere VMFS[32] allows multiple installations of VMware ESX® to access the same virtual machine files concurrently.

vMotion Across Boundaries and Over Long Distances

VMware revolutionized the concept of nondisruptive, live migration with the introduction of vMotion.

[31] https://www.vmware.com/products/vsphere/features/vmotion

[32] http://www.vmware.com/files/pdf/VMware-vStorage-VMFS-DS-EN.pdf

vSphere vMotion capabilities have been enhanced in vSphere 6.0, enabling users to perform live migration of virtual machines across virtual switches, vCenter Server systems, and long distances of up to 150 ms[33] RTT. These new vSphere vMotion enhancements enable greater flexibility when designing vSphere architectures that were previously restricted to a single vCenter Server system due to scalability limits and multisite or metro design constraints.

Storage vMotion Integration

VMware vSphere 5.1 and later versions combine standard vMotion with VMware vSphere Storage vMotion in a single migration. This means you can live-migrate an entire virtual machine between hosts, between clusters, or between data centers–without disruption or shared storage between the involved hosts. Storage vMotion can be used as a tool to help migrating offline SAP HANA systems from one storage to another storage; e.g., from internal server storage to external shared TDI storage solutions.

Preserving Network State

The underlying vSphere host virtualizes the networks used by the virtual machine. This ensures that even after the migration, the virtual machine network identity and network connections are preserved. Since a virtual machine migration with vMotion preserves the precise execution state, network identity, and active network connections, the result is no downtime or disruption to users.

Transaction Integrity

vMotion can transfer the active memory and precise execution state of the virtual machine over a high-speed network, allowing it to switch from running on the source vSphere host to the destination vSphere host. vMotion keeps the transfer period imperceptible to users by tracking ongoing memory transactions in a bitmap. Once the entire memory and system state have been copied to the target vSphere host, vMotion suspends the source virtual machine, copies the bitmap to the target vSphere host, and resumes the virtual machine on the target vSphere host, thus ensuring transaction integrity.

Migrating SAP HANA Databases with VMware vMotion

vMotion can migrate running SAP HANA databases from one host to another with no downtime. During the migration of the SAP HANA database VM, there may by an impact on response time. This needs to be taken into account when using vMotion.

In the event that an SAP HANA host experiences increased hardware alerts, an administrator can proactively migrate SAP HANA databases that reside on the host to another vSphere host in order to avoid downtime or costs to the business. This gets done by initiating a vMotion process in the vCenter front end.

The migration time of a VM depends on the size of the VM, the size of the attached virtual machine disks (VMDKs), and the quality of the underlying hardware such as network bandwidth and latency, and the change rate of data stored in the SAP HANA database. For instance, individual customer tests migrating SAP HANA VMs with sizes of approximately 1 TiB took in their unique 10GbE network environment and used

[33] Long Distance vMotion requirements in VMware vSphere 6.0 (2106949)

storage hardware up to 30 minutes, whereas tests in lab environments have shown that live migration of an SAP HANA VM, with approximately 500 GiB, takes 5–6 minutes to complete on a dedicated 10 gigabit Ethernet network with optimized hardware configuration.

Theoretically, the time to migrate a 4 TiB SAP HANA database over once 10 GbE network link, under optimal conditions will be just over 1 hour. In shared 10GbE environments it may be significantly longer! By adding additional physical network ports and vmkernel ports to the dedicated vMotion network, this time can get lowered significantly.

SAP HANA runs in-memory and it has a large memory footprint. When executing a live migration, vMotion preserves the entire state of the SAP HANA memory while query or transaction processing continues, with a possible performance hit, until the migration is completed.

When a standard vMotion process gets triggered, then by default it will get started with the “high priority” vMotion migration policy, which is also the recommendation for SAP HANA.

When a vMotion process starts during a CPU high load phase, vMotion may have issues finishing the migration due to the high change rate. In this case, the vMotion process will slow down the CPU change rate to ensure that all changes during a vMotion process can be replicated to the other system.

Starting a vMotion process with “standard priority” will have less impact on the performance of the VM and may be an option when a vMotion process is not allowed to impact too much of the SAP HANA performance, but may lead to longer vMotion times. The vMotion times stated on the previous page require high priority.

Table 5 describes the impact of these policies and the requirements for a successful migration process.

OPTION	DESCRIPTION
High priority	vCenter Server attempts to reserve resources on both the source and destination hosts to be shared among all concurrent migrations with vMotion. vCenter Server grants a larger share of host CPU resources. If sufficient CPU resources are not immediately available, vMotion is not initiated.
Standard priority	vCenter Server reserves resources on both the source and destination hosts to be shared among all concurrent migrations with vMotion. vCenter Server grants a smaller share of host CPU resources. If there is a lack of CPU resources, the duration of vMotion can be extended.

Table 5: vSphere 5.5 vMotion Migration Policies

Note: vMotion processes with a high migration policy will try to slow down a VM when the change rate inside this VM is too high. In extreme cases, this may lead to disconnects between the application server and the SAP HANA database. Scheduling vMotion during non-peak hours helps to avoid this issue.

In contrast, if SAP HANA has to shut down and restart, it will have following effect:

- Query or transactions processing aborted
- Temporary tables and computations lost
- SAP HANA performs a lazy restart by loading the system tables
- Columns to be preloaded–or upon query, into memory from disk

Customers that have HANA system replication implemented and configured can alternatively use it for planned maintenance situations as well. The requirement is that all services need to be available on the target site and that a failback procedure of SAP HANA is implemented as well. If System Replication is implemented in this way, then a restart time of near-zero would result.

Migrating SAP HANA via vMotion from an older hardware platform to a newer platform, such as from a Haswell EX based server system to a Broadwell EX based server, is possible. The enabling feature to do this is VMware Enhanced vMotion Compatibility (EVC)[34], which simplifies vMotion compatibility issues across CPU generations and has to be used to enforce CPU compatibility. A long-term operation of SAP HANA on a vSphere cluster requires the use of hosts in the cluster that use the same CPU type, model, and frequency.

Note: EVC has to be used and the EVC baseline has to be at least level L4, to ensure that only servers with Intel Sandy Bridge or later based server systems can be selected for vMotion. To avoid possible time-stamp counter (TSC) issues when migrating virtual machines, it is also recommended that the source and target ESXi host servers are required to use the same CPU type, model, and frequency.

EVC automatically configures server CPUs with Intel FlexMigration technologies to be compatible with older servers. After EVC is enabled for a cluster in the vCenter Server inventory, all hosts in that cluster are configured to **present identical CPU features and ensure CPU compatibility for vMotion**. The features presented by each host are determined by selecting a predefined EVC baseline. This ensures CPU compatibility for vMotion even though the underlying hardware might be different from host to host. Identical CPU features are exposed to virtual machines regardless of which host they are running on, so the virtual machines can migrate between any hosts in a cluster.

vCenter Server does not permit the addition of hosts that cannot be configured to be compatible automatically with the EVC baseline. For details about the EVC level and baseline, refer to Table 6.

[34] https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003212

EVC LEVEL	EVC BASELINE	DESCRIPTION
L4	Intel "Sandy Bridge" Generation	Applies baseline feature set of Intel "Sandy Bridge" Generation processors to all hosts in the cluster. Compared to the Intel "Westmere" Generation mode, this EVC mode exposes additional CPU features, including: AVX and XSAVE. **Note:** Intel "Sandy Bridge" processors that do not support AESNI and PCLMULQDQ cannot be admitted to EVC modes higher than the Intel "Nehalem" Generation mode.
L5	Intel "Ivy Bridge" Generation	Applies baseline feature set of Intel "Ivy Bridge" Generation processors to all hosts in the cluster. Compared to the Intel "Sandy Bridge" Generation EVC mode, this EVC mode exposes additional CPU features, including: RDRAND, ENFSTRG, FSGSBASE, SMEP, and F16C. **Note:** Some Intel "Ivy Bridge" processors do not provide the full "Ivy Bridge" feature set. Such processors cannot be admitted to EVC modes higher than the Intel "Nehalem" Generation mode.
L6	Intel "Haswell" Generation	Applies baseline feature set of Intel "Haswell" Generation processors to all hosts in the cluster. Compared to the Intel "Ivy Bridge" Generation EVC mode, this EVC mode exposes additional CPU features, including: ABMX2, MOVBE, FMA, PERMD, RORX/MULX, INVPCID, VMFUNC.

Table 6: ECV Levels (source: VMware KB100321)

Best Practice for SAP HANA vMotionon VMware for Production Support

vMotion between different hardware generations of a CPU type is possible, but in the context of a performance-critical application such as SAP HANA, it is important to follow these two best practices when using this feature:

- The customer **should** run SAP HANA VMs within the vSphere cluster on identical hardware only (with the same CPU clock speed and synchronized TSC).
- The customer may use vMotion during a hardware refresh (non-identical source and destination host/clock speed), but **should** plan for a restart of the VM afterwards.

Why is it critical to follow these best practices?

In **non-critical vMotion scenarios**, where a VM is moved from one host to another host in a cluster, and where all hosts have the same hardware version and level with the same CPU clock speed and synchronized time-stamp counter (TSC), there are no impact or performance issues. The scenarios shown in Figure 17 work without any problems and there is **no degradation in performance of the RDTSC instruction**.

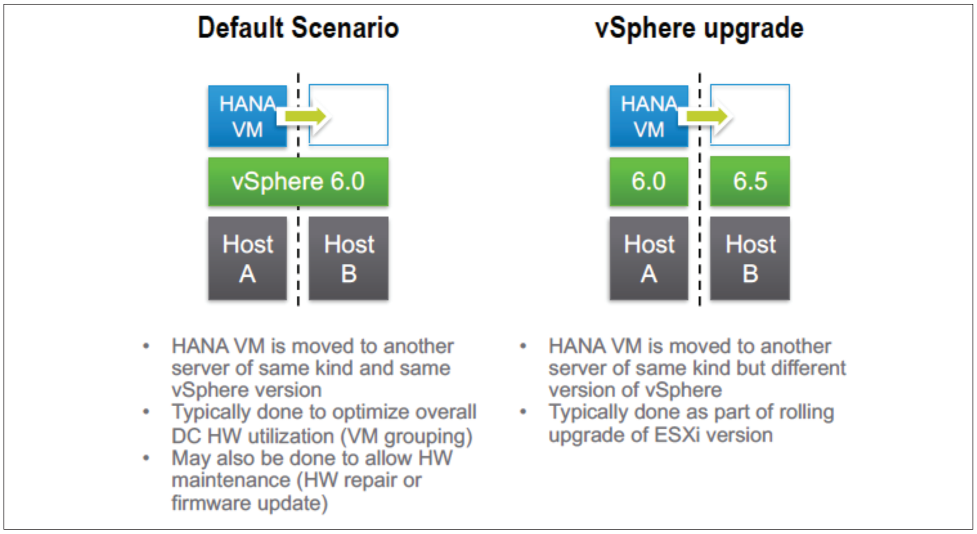

Figure 17: Non-critical vMotion Scenarios

A critical vMotion operation is when an in-homogenous vSphere cluster is used. This can happen when new servers with a different CPU version are added to an existing vSphere cluster (for example, when an existing Intel Xeon E7-v3 [Haswell] server-based cluster is upgraded or extended with Intel Xeon E7-v4 [Broadwell] based server systems, as shown in Figure 18).

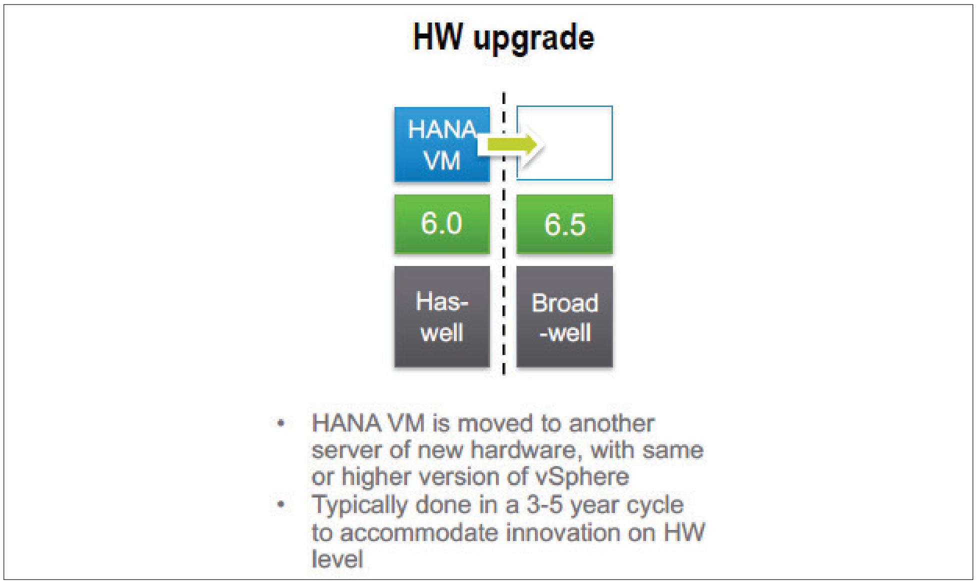

Figure 18: Critical vMotion Scenario

This a critical vMotion scenario due the different possible CPU clock speeds, plus the exposed TSC to the VM may cause timing errors. To eliminate these possible errors and issues caused by different TSCs, vSphere will perform the necessary rate transformation. This may degrade the performance of RDTSC relative to native.

Background: *When a virtual machine is powered on, its TSC inside the quest, by default, runs at the same rate as the host. If the virtual machine is then moved to a different host without being powered off (for example, by using VMware vSphere vMotion), a rate transformation is performed so the virtual TSC continues to run at its original power-on rate, not at the host TSC rate on the new host machine. For details, read the document "Timekeeping in VMware Virtual Machines."*[35]

To solve this issue, you must plan a maintenance window to be able to restart the VMs that were moved to the non-identical HW to allow the use of HW-based TSC instead of using software rate transformation on the target host, which is expensive and will degrade VM performance. Figure 19 shows the process to enable the most flexibility in terms of operation and maintenance by restarting the VM after the upgrade, ensuring the best possible performance.

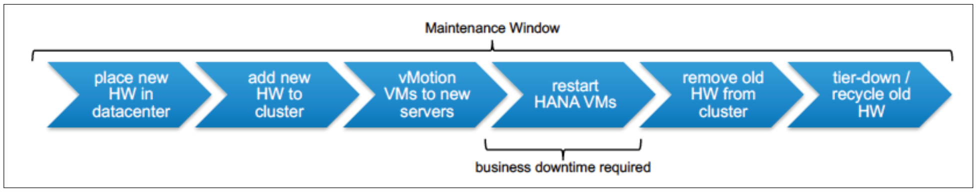

Figure 19: Maintenance Window After a HW Upgrade and vMotion

Summary:

- During normal operation, a vMotion does not cause any issues regarding the use of TSC.
- However, if used to for online replacement/refresh of the HW, the HANA VM may face TSC performance degradation due to the required software rate transformation on the target host.
- The use of rate transformation is temporary until VM is restarted and the guest OS can re-establish the TSC rating during boot time.

Managing SAP HANA Landscapes Using VMware DRS[36]

SAP HANA landscapes can be managed using VMware Distributed Resource Scheduler (DRS), which is an automated load balancing technology that aligns resource usage with business priority. DRS dynamically aligns resources with business priorities, balances computing capacity, and reduces power consumption in the data center, as shown in Figure 20.

[35] http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/Timekeeping-In-VirtualMachines.pdf

[36] https://www.vmware.com/products/vsphere/features/drs-dpm.html

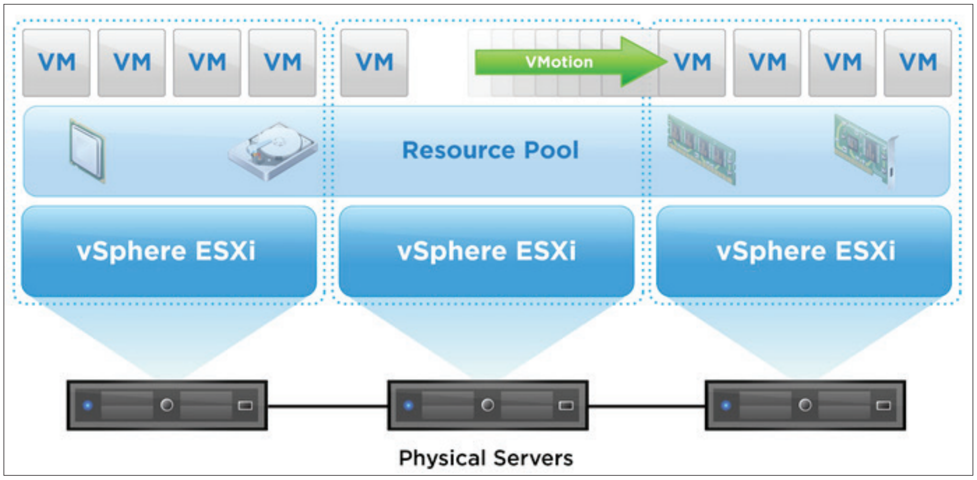

Figure 20: VMware Distributed Resource Scheduler

DRS takes advantage of vMotion to migrate virtual machines among a set of ESX hosts. DRS continuously monitors utilization across vSphere servers and would be able to migrate VMs to hosts that are less utilized.

DRS enables automatic initial virtual machine placement on any of the hosts in the cluster. It also makes automatic resource relocation and optimization decisions as hosts and virtual machines are added or removed from the cluster. When DRS is configured for manual control, it makes recommendations for review and later implementation only (there is no automated activity).

DRS can be set to these automation modes:

- Manual. In this mode, DRS recommends the initial placement of a virtual machine within the cluster, and then recommends the migration.
- Semi-automated. In this mode, DRS automates the placement of virtual machines and then recommends the migration of virtual machines.
- Fully automated. In this mode, DRS placements and migrations are automatic.

When deploying SAP HANA databases, it is essential to have DRS rules in place and to set the automation mode to manual in order to avoid unwanted migrations:

- As an example, a vSphere 6.0 and 6.5 cluster can include up to 64 hosts. If the cluster contains only three hosts that are SAP HANA certified, DRS potentially could move SAP HANA to a non-certified host. In this case, DRS should be set with an Affinity rule that allows DRS to migrate SAP HANA only to an SAP HANA certified server.
- DRS also can be used to set an Anti-Affinity rule to enforce that; for example, during quarter end close, SAP HANA is the only virtual machine that can run on a particular vSphere host. This rule can also include a condition that only one instance of SAP HANA can run on a host.

Note: For SAP HANA environments, it is recommended that you set DRS automation mode to "manual" to avoid unwanted migrations.

VMware High Availability

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

The power behind VMware's High Availability/Disaster Recovery (HA/DR) solutions is how they are layered to protect against failures at every level of the data center, from individual components such as NIC or HBA card teaming, all the way up to the entire site.

Figure 21 shows the different solutions to protect against component-level failures to complete site failures.

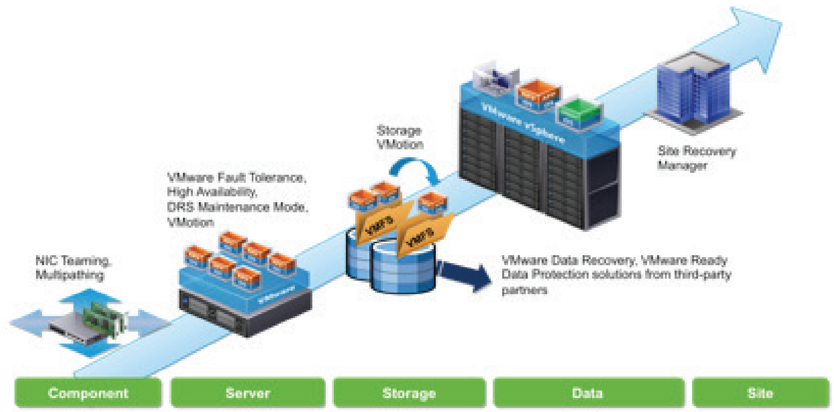

Figure 21: VMware HA/DR Solutions–Protection at Every Level

Many of the key features of virtualization, such as encapsulation and hardware independence, already offer inherent protections. From there, additional protections throughout the vSphere platform are provided to ensure organizations can meet their availability requirements and provide the following features:

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

VMware vSphere HA delivers[37] the availability required by most applications running in virtual machines, independent of the OS and the application running in it. It provides uniform, cost-effective failover protection against hardware and OS outages within a virtualized IT environment. It does this by monitoring vSphere hosts and virtual machines to detect hardware and guest OS failures. It restarts virtual machines on other vSphere hosts in the cluster without manual intervention when a server outage is detected, and it reduces application downtime by automatically restarting virtual machines upon detection of an OS failure.

[37] http://www.vmware.com/files/pdf/vsphere/VMW-WP-vSPHR-Whats-New-6-0-PLTFRM.pdf, page 13

With the growth in size and complexity of vSphere environments, the ability to prevent and recover from storage issues is more important than ever. VMware vSphere HA now includes Virtual Machine Component Protection (VMCP), which provides enhanced protection from All Paths Down (APD) and Permanent Device Loss (PDL) conditions for block (FC, iSCSI, FCoE) and file storage (NFS).

Prior to vSphere 6.0, vSphere HA could not detect APD conditions and had limited ability to detect and remediate PDL conditions. When those conditions occurred, applications were impacted or unavailable for longer periods of time and administrators had to help resolve these issues. vSphere VMCP detects APD and PDL conditions on connected storage, generates vCenter alarms, and automatically restarts impacted virtual machines on fully functional hosts. By doing this, it greatly improves the availability of virtual machines and applications without requiring more effort from administrators.

Note: To ensure that all VMware HA vSphere 5.5 and 6.x features can be used, ensure that the SAP HANA TDI storage to be used is also VMware vSphere supported and certified. Check out the VMware HCL for details.[38]

---

vSphere 6.0 also is fully compatible with VMware Virtual Volumes, VMware vSphere Network I/O Control, IPv6, VMware NSX®, and cross vCenter Server vSphere vMotion. Now, vSphere HA can be used in more and in larger environments and with less concern for feature compatibility.

Figure 22 shows a vSphere resource cluster with some of the VMs protected by VMware HA. If a server fails, the VMs (and all the VMs running SAP HANA instances) will be moved and restarted automatically on another vSphere host in the cluster.

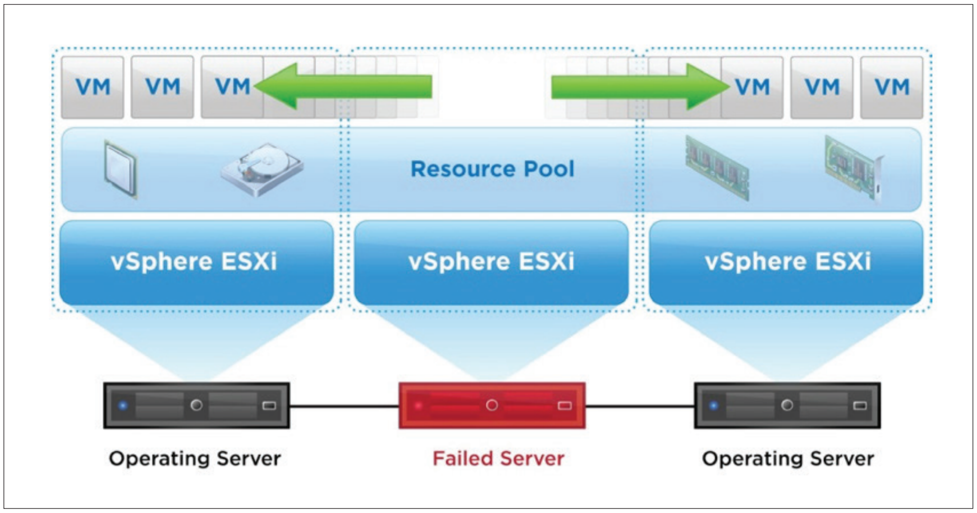

Figure 22: VMware High Availability

[38] http://www.vmware.com/resources/compatibility/search.php?deviceCategory=san

An SAP HANA specific HA and DR solution is described in the HA section of this document.

VMware HA allows IT organizations to:

- Minimize unplanned downtime and IT service disruptions, while eliminating the need for dedicated standby hardware and the installation of additional software.
- Provide affordable uniform high availability across the entire virtualized IT environment, without the cost and complexity of failover solutions that are tied to either operating systems or specific applications.

As mentioned, VMware HA is transparent for any OS or application running inside a VM. With VMware HA, it is easy to build an HA standard for any SAP or non-SAP application or database that runs virtualized on VMware vSphere. Details on how to protect SAP HANA on vSphere are in following section, "SAP HANA on vSphere High Availability (HA)."

SAP HANA VM Clones and Template

After a virtual SAP HANA database is installed, a VMware template or clone can be created with the Linux® operating system already configured and tuned, preloaded with all of the necessary drivers, and with any additional software that is optimized to run on vSphere.

VMware vSphere can enable copies of an SAP HANA virtual machine using either a VMware clone or a VMware template. A VMware clone is an exact copy of an SAP HANA virtual machine. When a virtual machine is cloned, a copy of the entire virtual machine is created, including its settings, any configured virtual devices, installed software, and other contents of the virtual machine's disks. Optionally, guest operating system customization can be used to change some of the properties of the clone, such as the vSphere host name and networking settings. In combination with SAP Landscape Virtualization Management (SAP LVM), it is possible to customize SAP applications such as SAP HANA after initial deployment to create, for example, an LVM-initiated VMware clone from a production system, deploying it new and changed as a QA system. For details, check out the SAP LVM webpages (LVM main page[39] and SAP HANA LVM supported operation features[40]).

A note on VMware snapshots: in contrast to VMware clones, which are exact copies of a virtual machine, a VMware snapshot represents the state of a virtual machine at the time it was taken, and it might negatively affect the performance of the virtual machine. This is based on how long it has been in place, the number of snapshots taken, and how much the virtual machine and its guest operating system have changed since the time it was taken. When a snapshot should be taken; for instance, before installing a new SAP HANA patch, take the snapshot and do not select the option "Snapshot the virtual machine's memory". The general recommendation is that you shut down the SAP HANA VM prior a snapshot being created.

[39] http://help.sap.com/nwlvm

[40] http://help.sap.com/saphelp_lvment21/helpdata/en/a3/305f149e0346bf94c9dcfdd4e6b900/content.htm

Note: Taking VMware snapshots and keeping them for backup reasons is, as with other performance-critical applications, not recommended or supported. Taking a snapshot for a short time period is possible; for instance, to create a recovery point prior to installing an SAP HANA or OS patch or to make a snap/clone based backup, but using snapshots and keeping them as a backup and restore solution is not recommended!

A VMware template is recommended once the SAP HANA database is fully configured and optimized for vSphere. A VMware template is a master copy of an SAP HANA virtual machine that can be used to create many SAP HANA clones. This template can be used to create and provision multiple SAP HANA virtual machines.

Templates cannot be powered on or edited. A template offers a more secure way of preserving repeatable SAP HANA virtual machine configurations for deployment.

Note: In combination with SAP Landscape Virtualization Management or other deployment and lifecycle tools, clones and templates can help to lower the burden of managing an SAP HANA environment. Please refer to the VMware SAP LVM adapter webpage for more information.[41]

VMware Host Profiles

A VMware Host Profile encapsulates a reference host configuration and turns it into a profile or template that can then be applied across new or existing vSphere hosts. Administrators can use a Host Profile to ensure that all vSphere host configurations are consistent and that configuration drifts can be prevented easily through automatic compliance checks.

Checking compliance of the vSphere host cluster on a regular basis ensures that the host or cluster continues to be correctly configured. Once the profile to a host or cluster is attached, it can be used to check the compliance status from various places in the vSphere Client.

In addition, when firmware upgrades or other events occur that require storage, network, or security configuration changes on multiple hosts in a cluster, administrators can edit the Host Profile and apply it across the cluster for consistent configuration updates.

When creating SAP HANA landscapes, it is critical for all virtualized hosts and the SAP HANA virtual machines to be configured consistently. Using VMware clones, templates, and Host Profiles to create the SAP HANA environment ensures that deployments are rapid and consistent, and achieve consistent performance levels.

Migration from vSphere 5.5 to vSphere 6.x

Migrating from an existing vSphere 5.5 environment to a vSphere 6.x environment is relatively easy, and can be done without any downtime of the SAP HANA VMs that are running.

The first component that needs to be upgraded is the vCenter management server to allow management of vSphere versions 5.5, 6.0 and 6.5. Before upgrading any vSphere host systems, verify the certification status of the currently used server

[41] http://www.vmware.com/de/products/adapter-sap-lvm.html

systems to ensure that they are supported with vSphere 6.0/6.5 and determine if there are SAP HANA VMs deployed that limit the usage of vSphere 6.0 because of currently non-supported deployment options, such as a Scale-Out configuration.

Upgrading the SAP HANA VM's virtual hardware version to version 11 or 12 (vSphere 6.5) and to correspondingly add more RAM and CPU resources requires a restart of the VM. The actual vSphere host upgrade can be performed as mentioned, however, without any downtime.

Perform the Perform the following steps:

1. Install a new vCenter 6.5 server, which includes in this version now the Update Manager from an pre.configured appliance fOVA file.
2. Install vSphere 5.5 and 6.0/6.5 hosts as needed for SAP HANA and respect the SAP HANA support status and deploy virtualized SAP HANA systems on the hosts and select the correct VM hardware version of the VMs. Version 10 for vSphere 5.5 (1 TiB limit) and version 11 for vSphere 6.0 (4 TiB limit) and HW version 12 for vSphere 6.5..
3. Select the correct hardware version to ensure that that hardware version 11/12 VM won't run on vSphere 5.5 hosts and that the SAP support limitations get enforced.
4. Set the vSphere 5.5 host to maintenance mode and evacuate all running VMs. When the host is empty, perform a manual or automated vSphere upgrade by using for the vSphere Update Manager Server to upgrade the host from 5.5 to 6.x.
5. When upgraded and up and running again, migrate the running VMs back to the version 6.x host server.
6. Do this with all vSphere 5.5 host systems.
7. Plan for a downtime window to update VM hardware version 10 to 11/12, and upgrade the VM vCPU and vRAM sizes as needed
8. YOU ARE DONE!

Note: Before upgrading a vSphere environment, ensure that the used server hardware is supported with the new vSphere version. Prior to the upgrade, you are required to perform a backup of the host and vCenter configuration.

The document "vSphere Upgrade[42]" provides all necessary information to perform a successful upgrade to the new vSphere version.

Virtualized SAP HANA Reference Architecture

Note: When planning a production-level virtualized SAP HANA environment, validate the support status of the planned vSphere version because the different versions of vSphere may have a different support status. This document provides the support status as of December 2016. More recent support information may be found in SAP Note 1788665–SAP HANA Support for VMware vSphere Environment.

---

[42] https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-upgrade-guide.pdf

General Concepts and Overview

SAP HANA can be delivered as an appliance, preinstalled by one of the hardware vendors, or be installed on-site by a customer or service contractor in the Tailored Datacenter Integration (TDI) framework. Currently, no SAP HANA appliance has been certified virtualized, yet some hardware vendors still deliver SAP HANA preinstalled on VMware. The main usage of SAP HANA on VMware is within the TDI framework.

When SAP HANA is installed on-site at the customer as a TDI configuration, then an SAP Certified Technology Specialist (Edition 2013)–SAP HANA Installation (E_HANAINS131) trained and certified person should do the SAP HANA installation.

Virtualized SAP HANA configurations can be deployed Scale-Up and Scale–out configurations. An SAP HANA Scale-Out configuration depends on external infrastructure components for storage and network, which interconnects the server systems hosting the SAP HANA instances that build an SAP HANA Scale-Out system. As mentioned, SAP HANA Scale-Up configurations also benefit from leveraging external infrastructure components and, in fact, the most virtualized SAP HANA Scale-Up configuration is installed on an external shared storage array.

Regardless if deployed as a virtual SAP HANA appliance with predefined local storage or by leveraging external infrastructure components, all these components must meet the minimum requirements SAP has defined for SAP HANA configurations, and should pass the tests as defined in the hardware Configuration Check Tool for SAP HANA.[43] The selected server systems and configuration (per a BOM) need to be listed on the SAP HANA Certified Hardware Directory[44] and the VMware vSphere Hardware Compatibility List.[45] A server can be used for virtualized SAP HANA deployments only when it is on both lists.

This chapter describes a VMware vSphere virtualized SAP HANA architecture in detail and covers the current system maximums, the network configuration, and storage options such as Fibre Channel-SAN or NFS attached shared storage systems. It also provides information on high availability, disaster tolerance, and vRealize Operations.

The following documents describe SAP HANA TDI prerequisites and recommendations, and will be referenced later in this guide.

- "SAP HANA Tailored Data Center Integration Frequently Asked Questions"
- "SAP HANA Tailored Data Center Integration Overview Presentation"
- "SAP HANA TDI Storage Requirements"
- "SAP HANA Server Installation and Update Guide"

Architecture Overview

Figure 23 shows the different layers of the architecture from the storage layer up the management layer. The architecture highlights the shared nothing architecture of the separated hosts, where only resources can be shared between VMs inside a server or when external components such as shared storage is used. It is possible to operate and configure vSphere 5.5 and vSphere 6.0/6.5 based server systems with one single

[43] http://service.sap.com/sap/support/notes/1943937

[44] https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html

[45] https://www.vmware.com/resources/compatibility/search.php

vCenter 6.5 instance and to configure TDI and appliance-based servers in one management zone. When SAP HANA certified Intel Xeon E7 server systems and E5 supported systems are used, then it is required to run them in different resource clusters, which again can be configured and operated with one single vCenter. The shared TDI storage should be sized to support all running SAP HANA VMs and to deliver the needed HANA KPIs. If the systems should be protected against failures, then sufficient spare resources have to be available to support a failover and restart process. The needed external network components are not shown in the figure.

Every layer of the architecture and the SAP HANA specific requirements are explained on the following pages.

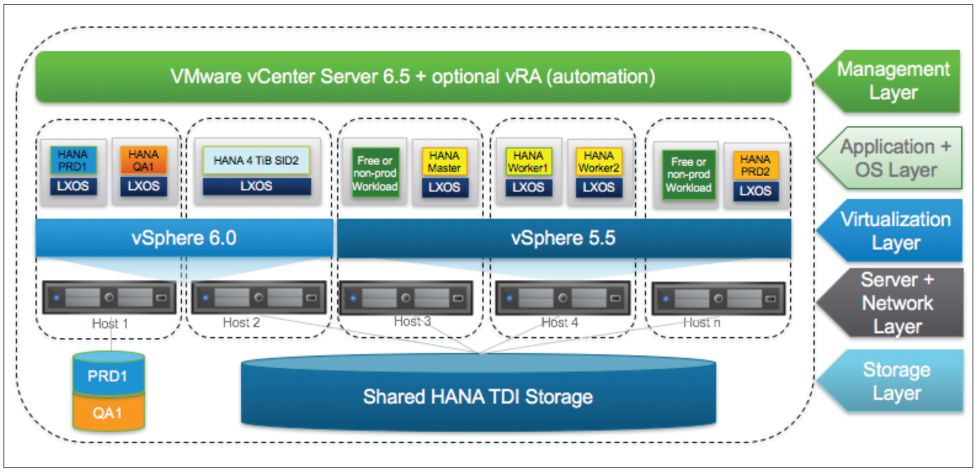

Figure 23: Virtualized SAP HANA High-level Architecture

All components and layers are configurable and manageable via one single management server named **VMware vCenter Server** (in a single management domain), and allows a vSphere or SAP administrator to manage all systems and resources using the same tools, processes, and methods. The virtualized SAP appliance servers provide local storage capacity for the SAP HANA VMs running on them.

With vSphere 6.0 and 6.5, resource clusters with up to 64 hosts are possible; with vSphere 5.5, up to 32 hosts had been possible. A single vCenter Server can manage multiple vSphere clusters with up to 1000 hosts in total. The VM maximum that can be managed via a single vCenter Server instance is 10,000 powered-on VMs[46] or 15,000 powered-off VMs . Up to 10 vCenter Servers can be linked together to manage even larger configurations.

Figure 24 shows the main functions of the vCenter Server and how SAP HANA deployments benefit from running on VMware vSphere by providing automation, scalability, and visibility of the whole infrastructure landscape supporting SAP HANA.

[46] https://www.vmware.com/pdf/vsphere6/r60/vsphere-60-configuration-maximums.pdf, page 17

One or multiple SAP HANA VMs can be operated on the vSphere hosts. The number of SAP HANA VMs are limited only by the underlying hardware resources and SAP HANA sizing rules. Additional compute, storage, network capacity, and components can be added easily when needed.

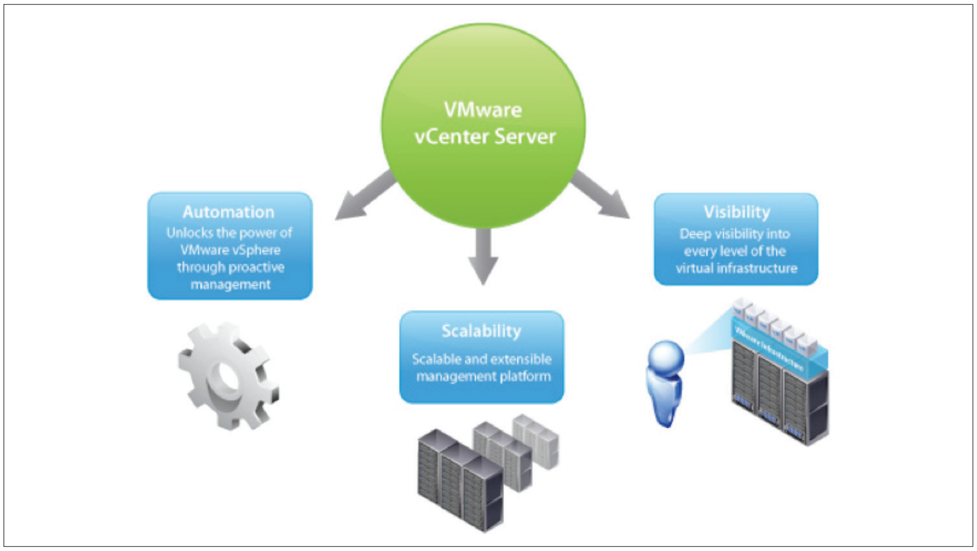

Figure 24: VMware vCenter Server

By installing vSphere on an SAP HANA appliance system, it is possible to gain the management operation and SLA advantages a VMware vSphere environment is able to provide. The migration of complete, installed HANA systems between physical server systems using local or shared storage, and vice versa, is possible via enhanced vMotion, which migrates not only the VM, but also the assigned virtual disks.

Deploying an SAP HANA BWoH Scale-Out system on such an environment is relatively easy because it is possible to extend the configuration with the needed NFS share and additional compute, network, and storage capacity.

Figure 25 shows how to convert a Scale-Up BW system to a Scale-Out system by adding resources.

Note: The current version of the SAP HANA installer always expects a standby host when a Scale-Out system is installed, even if none is needed (e.g., due to VMware HA). Follow the steps in the Best Practice section to work around this issue.

Converting a virtualized SAP HANA Scale-Up to a Scale-Out system may be required when the system runs out of compute resources, and additional resources such as RAM cannot get extended. It may be also required when SAP HANA VM reaches its 1 TiB (vSphere 5.5) or 4 TiB (vSphere 6.0) maximum, and converting to a physical SAP HANA system is not an option. SAP currently supports up to 16 Scale-Out BWoH nodes in GA[47], which would allow 4-socket Intel Xeon E7-v3 (Haswell) systems (2 TiB

[47] Generally Available

per host) close to 32 TiB of total VM RAM or when one host failure should get supported close to 30 TiB of total vRAM (n+1 configuration). The SAP HANA BW RAM maximum for 8-socket Haswell server systems is currently 4 TiB per host.

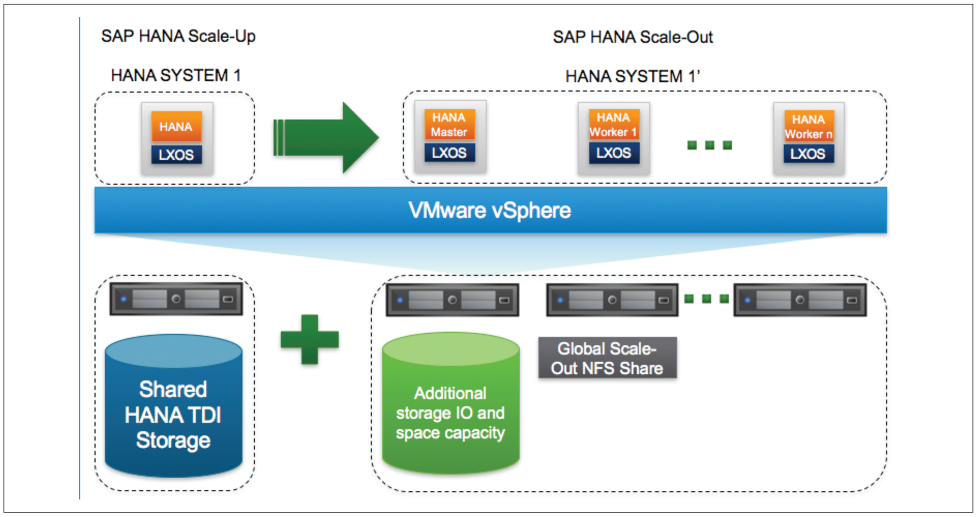

Figure 25: Converting a Virtualized SAP HANA Scale-Up to a Scale-Out System

Figure 26 shows the previously described environment, where the one SAP HANA system running on one of the appliance server systems was converted to a Scale-Out configuration, leveraging the existing shared SAP HANA TDI storage system. Before converting the system, an NFS share had to be created for the global SAP HANA Scale-Out share, and the storage might have to be upgraded to provide enough storage capacity and IOPS performance for the additional connected systems.

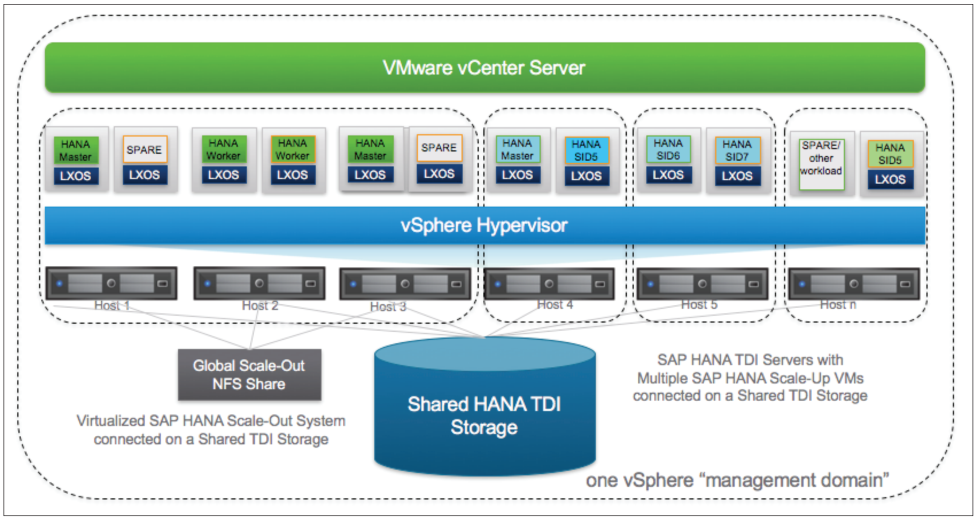

Figure 26: SAP HANA TDI Scale-Up and Scale-Out Systems Sharing a Single TDI Storage

The next sections describe the different layers in more detail, and highlight the virtual SAP HANA requirements on the specific layers.

Storage Layer

Before we discuss which storage options are supported for SAP HANA databases running on VMware vSphere, here are some important details on how virtual machines access different types of storage.

When a virtual machine 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 (FC), Internet SCSI (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 virtual machine as a mounted SCSI device.

The virtual disk hides the physical storage layer from the virtual machine's operating system. This allows running operating systems in a very standardized way without any specific storage requirements, such as SAN host bus adapter (HBA) drivers or other settings.

Figure 27 shows the storage types that vSphere uses, how the storage is connected to the vSphere host, and how a VM accesses it. More details about the different storage types can be found in the "vSphere storage guide"[48]. For SAP HANA systems, currently only NAS over NFS and fiber-based storage arrays formatted with Virtual Machine File System (VMFS) are supported.

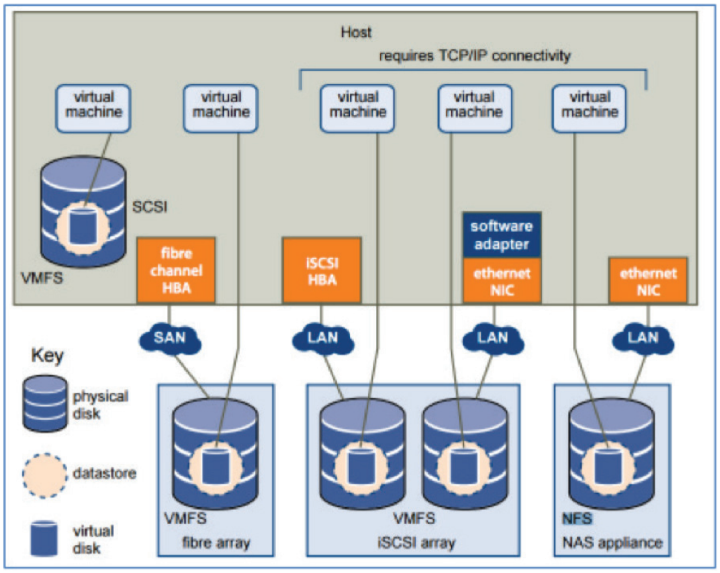

Figure 27: Virtual Machines Accessing Different Types of Storage

[48] https://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-storage-guide.pdf

The figure shows how the physical disks/logical unit numbers (LUNs) are connected via a network card (NFS) or a Fibre Channel HBA (fiber-based SAN storage) on the vSphere host, and how the actual VMDK file of a VM is placed inside a datastore. Datastores on fiber arrays are formatted with VMware's own cluster file system, VMFS; datastores on NAS storage use NFS as the cluster file system.

A VM with an NFS client installed can directly access any NFS-based NAS in the network. This feature is required when installing the SAP HANA Scale-Out deployment option on vSphere because:

- SAP HANA server software in a Scale-Out configuration must be installed partly on a shared NFS based storage system.
- The file system has to be mounted by all VMs in the SAP HANA system.

For example, Figure 26 shows this file system share as "Global Scale-Out NFS Share." The individual SAP HANA host/VM files (such as SAP HANA database or log files) can be installed either on NFS or SAN storage.

The following section provides more details on VMFS and VMware vSphere datastores.

Virtual Machine File System

VMware vSphere VMFS provides high-performance, clustered storage virtualization that is optimized for virtual machines. With VMFS, each virtual machine is encapsulated into a small set of files. VMFS is the default storage management interface that is used to access those files on physical SCSI disks and partitions. VMFS allows multiple ESX instances to access shared virtual machine storage concurrently. It also enables virtualization-distributed infrastructure services, such as VMware vMotion, DRS, and High Availability, to operate across a cluster of ESX hosts.

To balance performance and manageability in a virtual environment, an accepted best practice is to deploy databases using VMFS. Raw device mapping (RDM) is sometimes erroneously selected to provide increased performance. The two dominant workloads associated with databases, random read/write and sequential writes, have nearly identical performance throughput characteristics when deployed on VMFS or using RDM.

Datastores

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

In traditional storage area network (SAN) deployments, an accepted best practice is to create a dedicated datastore if the application has a demanding I/O profile. Databases fall into this category. The creation of dedicated datastores with vSphere allows database administrators to define individual service level agreements (SLAs) for different applications. This is analogous to provisioning dedicated logical units (LUNs) in the physical world.

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

- Create separate and isolated datastores for SAP HANA data and logs.
- Enable multiple SAP HANA virtual machines to provision their data and log virtual machine disk files on the same class of storage.

NFS Datastores

vSphere allows customers to leverage enterprise-class NFS arrays to provide datastores with concurrent access by all of the nodes in an ESX cluster. For TDI deployments, data for SAP HANA system volumes are connected through the Network File System (NFS) protocol, and these volumes are mounted directly into the guest operating systems. Because all data for SAP HANA instances, including performance-critical data and log volumes for the database, is provided through NFS in this solution, proper network design and configuration are crucial[49].

Table 7 summarizes the vSphere features supported by the different storage types. All these storage types are available for virtual SAP HANA systems.

Note: The VMware supported Scale-Out solution requires the installation of the SAP HANA shared file system on an NFS share. For all other SAP HANA Scale-Up and Scale–out volumes such as data or log, all storage types as outlined in Table 5 could be used, as long the SAP HANA TDI Storage KPIs are achieved per HANA VM. Other solutions–such as Oracle Cluster File System (OCFS) or IBM General Parallel File System (GPFS)–are not supported by VMware.

STORAGE TYPE	BOOT VM	VMOTION	DATASTORE	RDM	VM IN-GUEST CLUSTER	VMWARE HA AND DRS
Local storage	Yes	No	VMFS	No	Yes*	No
Fibre Channel	Yes	Yes	VMFS	Yes	Yes*	Yes
iSCSI	Yes	Yes	VMFS	Yes	No*	Yes
NAS over NFS	Yes	Yes	NFS	No	No*	Yes

*In guest cluster software support, e.g., SUSE Linux Enterprise High Availability Extension[50]

Table 7: vSphere Features Supported by Storage Type

Figure 28 shows the SAP recommended SAP HANA file system layout, which is the suggested layout when running SAP HANA virtualized. Grouping the file system layout into three groups helps you decide whether to use VMDK files or an NFS mount point to store the SAP HANA files and directories.

[49] "SAP HANA on VMware vSphere and NetApp FAS Systems Reference Architecture," http://www.netapp.com/us/media/tr-4338.pdf

[50] https://www.suse.com/products/highavailability/

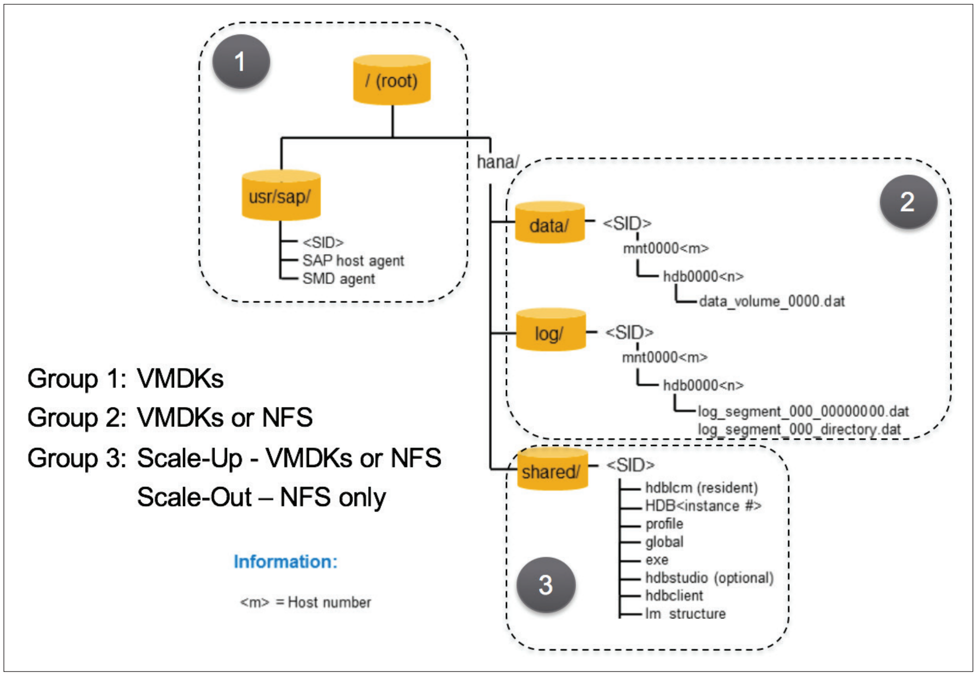

Figure 28: SAP Recommended SAP HANA File System Layout

All SAP HANA instances have a database log, data, root, local SAP, and shared SAP volume. The storage capacity sizing calculation of these volumes is based on the overall amount of memory needed by the SAP HANA in-memory database. SAP has defined very strict performance KPIs that have to 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 the amount of spindles may be required to provide the required I/O performance and latency. Therefore, storage sizing is very important and should be done by the storage vendor or predefined SAP HANA TDI storage configurations should be used.

SCSI Adapters

Paravirtual SCSI (PVSCSI) adapters are high-performance storage adapters that can result in greater throughput and lower CPU utilization, which is good for VMs running high demanding database systems such as SAP HANA.

It is a best practice to create a primary adapter for use with a disk that hosts the system software (OS) and SAP HANA binaries, and multiple, separate SCSI adapters for the SAP HANA data, log, and optional backup devices. The primary purpose for using multiple virtual SCSI controllers is to parallelize the units of work in a database transaction or query. In this case, carefully consider the implications when using multiple SCSI controllers to parallelize a single unit of work within a transaction. For instance, creating several SCSI controllers for data files increases throughput, but it may also increase latency.

File System Considerations and Alignment

As in the physical world, file system misalignment can severely impact performance. File system misalignment not only manifests itself in databases, but with any high I/O workload. VMware makes these recommendations for VMware VMFS partitions:

Similar to other disk-based file systems, VMFS suffers a penalty when the partition is not aligned. Use VMware vCenter to create VMFS partitions because it automatically aligns the partitions along the 64 KB boundary. To manually align your VMware VMFS partitions, check your storage vendor's recommendations for the partition starting block (for example, EMC VNX uses 128K offsets).

When an existing SAP system is migrated to SAP HANA, use an SAP HANA sizing report to determine the needed storage capacity for the data and log. For new systems, use the SAP HANA Quick Sizer.

If this report is not available, then you can use the values shown in Table 7. Table 7 shows the suggested virtualized storage configuration based on SAP's HANA storage sizing guidelines and summarizes the sizing information on how big the volumes/ virtual disks should be configured, and if the usage of NFS instead of VMDK is possible. For a precise definition and calculation of the net data size on disk and disk space required for mergers, refer the SAP HANA storage document.

Table 8 shows the suggested SCSI controller configuration and a VMDK naming convention to easily track VMDKs by use case, disk number, and HANA SID; e.g., vmdk01-SAP-SIDx, where SAP stands for the usr/sap volume and SIDx for the SID number of the system.

GROUP	VOLUME	VMDK	NFS	SCSI CONTROLLER	VMDK NAME	SCSI ID	SIZING AS OF SAP DOCUMENT
1	/(root)	✓	✗	PVSCSI Contr. 1*	vmdk0-OS-SIDx	SCSI 0:0	10 GiB for OS
1	usr/sap	✓	✓	PVSCSI Contr. 1*	vmdk01-SAP-SIDx	SCSI 0:1	50 GiB for SAP binaries
2	data/	✓	✓	PVSCSI Contr. 2	vmdk1-DAT1-SIDx vmdk1-DAT2-SIDx vmdk1-DAT3-SIDx	SCSI 1:0 SCSI 1:1	1.2 x net disk space** **NOTE**: If you use multiple VMDKs, use Linux LVM to build one large data disk, for example.
2	log/	✓	✓	PVSCSI Contr. 3	vmdk2-LOG1-SIDx vmdk1-LOG2-SIDx	SCSI 2:0 SCSI 2:1	[systems <= 512GiB] log volume (min) = 0.5 x RAM [systems >= 512GiB] log volume (min) = 512GiB
3	shared/ (Scale-Up)	✓	✓	PVSCSI Contr. 4	vmdk3-SHA-SIDx	SCSI 3:0	[systems <= 1 TiB] shared volume = 1 x RAM [systems >= 1 TiB] shared volume = 1 TiB

GROUP	VOLUME	VMDK	NFS	SCSI CONTROLLER	VMDK NAME	SCSI ID	SIZING AS OF SAP DOCUMENT
	shared/ (Scale-Out)	✗	✓	-	-	-	Scale-Out 1 x RAM per 4 active worker nodes (e.g. 4+1 1 TiB nodes = 1 TiB, 5+1, 1 TiB nodes = 2 TiB, 9+2, 1 TiB nodes = 3 TiB)
4***	Backup (Scale-Up)	✓	✓	PVSCSI Contr. 4*	vmdk4-BAK-SIDx	SCSI 0:2	Size backups >= Size data+Size redo log Default path for backup is (/hana/shared); this needs to get changed when a dedicated backup volume is used.
4***	Backup (Scale-Out)	✗	✓	-	-	-	Size backups >= Size of data+ size of redo log Default path for backup is (/hana/shared); this needs to be changed when a dedicated backup volume is used. When Scale-Out, then it must be installed on a shareable volume.

* The volumes root, sap, shared, and backup can be operated by one SCSI controller because I/O load is lower on log or data volumes.

** Net disk space = Net disk space for data is the amount of data stored on disk. Rule of thumb: Data = 1.1 x RAM for SoH or 0.84 RAM for BWoH (for simplification) because space usually is not the issue for an SAP HANA TDI storage configuration. If the workload is not defined yet, use 1 x RAM for DATA.

***A dedicated backup volume is recommended to mitigate risk of running out of space on this volume. If it is a Scale-Out configuration, then the backup volume must be accessible from all nodes.

Table 8: Virtualized SAP HANA Storage Configuration Based on SAP HANA Storage Requirements Document

Note: Raw device mappings can be used and are supported, but need to be verified and planned by the selected storage hardware vendor. From a performance point of view, there is very little benefit in using raw device mappings; therefore, VMware recommends using VMDK based disks instead of raw device mappings.

Figure 29 shows the virtualized SAP HANA disk layout of a Scale-Up system. A Scale–out system must use shared volumes for /hana/shared and the optional backup volume. Leveraging VMDK disks for /hana/shared is only possible for Scale-Up systems because only one system is accessing this volume.

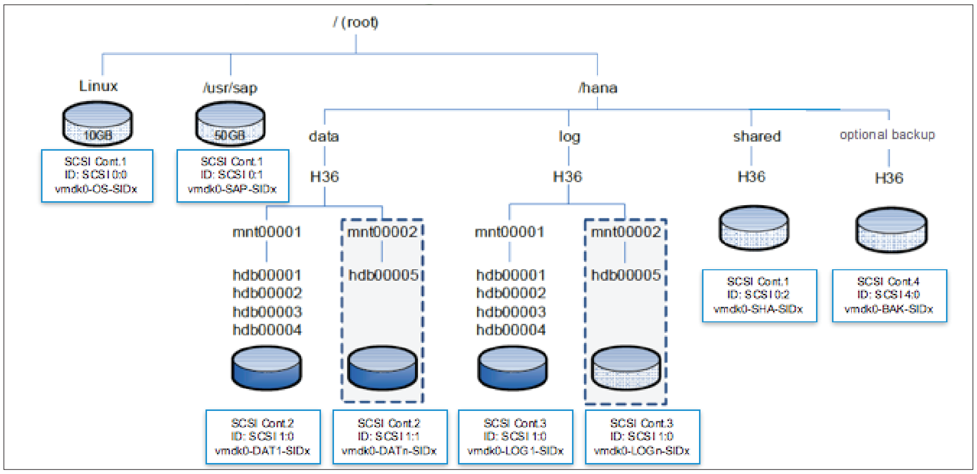

Figure 29: Virtualized SAP HANA Disk Layout of a Scale-Up System

When calculating the capacity needs for Scale-Up or a Scale-Out SAP HANA system, follow the already referenced SAP published storage guidelines. The SAP HANA storage requirements change from time to time; therefore, we recommend downloading the latest SAP HANA storage requirements white paper and adapting the configuration to the latest published SAP guidelines. We also recommend involving the storage vendor offering and SAP HANA TDI storage solution for SAP HANA configurations to ensure they have followed the correct spacing and IOPS sizing guidelines[51].

The storage volume size figures in Table 7 summarize the current guidelines and should be used only as an initial sizing guideline, which will have to be verified by the specific SAP BW or Suite on HANA sizing reports, Quick Sizer, or storage vendor calculation. Again, this should be seen as an initial sizing; changing the volume sizes may be required once the system is in production for some time.

As a default, the backup volume will be stored under the /hana/shared volume and is calculated with the site of the data volume + redo logs. As shown in Figure 27 and in Table 7, we recommend creating a dedicated VMDK volume or NFS mount point to store backup data; otherwise, the SAP HANA system may stop working because a backup may exceed the volume capacity.

When selecting an SAP HANA TDI storage configuration, review the "SAP HANA Server Installation and Update Guide,"[52] which describes the supported storage options for SAP HANA and the Certified Enterprise Storage list[53] on the Certified SAP HANA Hardware Directory.

[51] SAP HANA Storage Requirements, http://go.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html

[52] http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf

[53] https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/enterprise-storage.html

Note: When selecting a storage configuration, keep in mind that the number of supported SAP HANA VMs, which will run in parallel on a single vSphere host, is limited by the I/O capacity of the used storage. As with RAM and CPU resources, overcommitting of I/O is not supported. SAP has defined KPIs for Data Throughput and Latency for Production SAP HANA systems, which are the minimum values required for successful certification of production SAP HANA systems. For development and test, SAP HANA databases there are relaxed storage KPIs.

For an example of how a physical SAP HANA TDI storage can be configured to support a specific SAP HANA VM size, refer to the example below, which represents the Hitachi Data Systems (HDS) architecture and storage layout of an FC–based storage system used to certify SAP HANA on vSphere 6.0 at SAP's facility in Walldorf, Germany. Although HDS equipment was used for certification purposes, SAP HANA deployments are not limited to HDS. All configurations contained in the Certified SAP HANA Hardware Directory[54] and VMware Hardware Compatibility Guide[55] are fully supported.

Example: To meet the SAP HANA KPIs, Hitachi's general approach is to deploy Raid6 (6D+2P) raid groups using 600 GiB 10k SAS drives. The number of raid groups required to support an SAP HANA VM depends on the size of the virtual machines. Table 9 shows the approximate design.

HANA VM SIZE	TOTAL # RAID GROUPS	OS & SHARED	LOG	DATA
VM ≤ 512 GiB	3	1	1	1
512GiB<VM≤3 TiB	5	1	1	3
3 TiB < VM	7	2	1	4

Table 9: Example TDI Configuration, Number of RAID Groups–All of Type RAID6(6D+2P), Based on a HDS Storage System

To better understand the infrastructure needs and setup, Figure 30 shows the hardware architecture of the HDS systems used in Walldorf during the SAP HANA on vSphere 6.0 testing.

The HDS systems stand here as an example and all other VMware and SAP HANA certified and supported TDI storage solutions can be used. Other vendor configuration may differ from this configuration, so treat it as an example as it is intended for.

[54] http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/appliances.html

[55] http://www.vmware.com/resources/compaTiBility/search.php

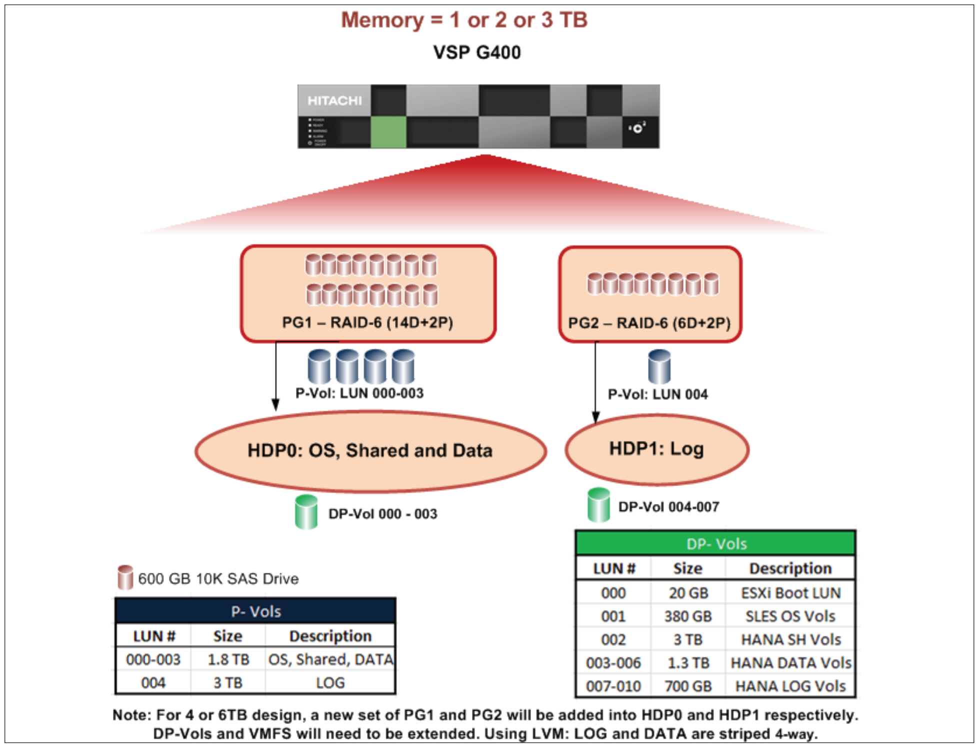

Figure 30: SAP HANA on vSphere 6.0 Certification Architecture

To summarize, the major difference between a virtualized and nonvirtualized SAP HANA environment–especially from the application view–is that inside a VM, the OS and SAP HANA application "sees" only "standard" hardware and SCSI devices, regardless of which storage technology is on the host. The performance and capacity requirements are the same, and when measured and verified inside the SAP HANA VM, the KPIs will be met. When NFS storage is directly accessed inside a VM instead of the host, there is no difference between a virtualized or nonvirtualized system.

Note: When configuring storage of a virtualized SAP HANA storage, follow the sizing and configuration guidelines of the selected storage vendor. The information provided in this guide is only a starting point for your planning. If no vendor-specific virtualized SAP HANA storage guidelines are available, then use the physical SAP HANA configuration guidelines and work with the vendor how to adapt these guidelines to a virtualized vSphere environment.

Server and Network Layer

The next layer to discuss is the server and network layer. The reason both layers are discussed here is how a virtualized network gets built and configured and that a virtual network could for instance gets limited to allow network traffic only inside a single host. Only when systems that run not on the same host need to be accessed, it is required for network packages to pass external network devices. If two VMs that run on the same host communicate, no external network communication is needed.

SAP HANA on vSphere Server Considerations

Let's discuss compute first because it is the technical foundation of a virtualized infrastructure and provides all necessary local compute resources such as CPU and RAM, and also network resources to operate an SAP HANA system securely, reliably, and with cost-efficient performance.

Selecting the right server platform for a virtualized SAP HANA system is critical because only VMware vSphere and SAP HANA certified server systems are supported. This is the same as other SAP applications, where a server needs to be SAP on Microsoft® Windows® or SAP on Linux supported.

As of December 2016, VMware vSphere is the only on-premises hypervisor that is supported by SAP to run SAP applications virtualized with Microsoft Windows and Linux OS. Other hypervisor solutions are supported only to run Linux or Windows based SAP applications. For details, see SAP Note 1122387–Linux: SAP Support in virtualized environments, and Note 1409608–Virtualization on Windows.

Beside the fact that the selected server and storage components must be SAP HANA certified or supported, for the VMware vSphere release it is important to get full SAP production support. For details, check out the tables on which vSphere version is supported with which deployment options/scenarios in the first section of this document, "Production Support."

Note: Only SAP HANA certified or supported server or storage systems, which also have passed the VMware vSphere hardware certification, are supported for SAP HANA virtualization.

The next section provides information on virtualization specifics and explains some of the concepts, such as vCPU or NUMA node locality.

Virtual CPUs (vCPUs) and Hyperthreading

Virtual CPUs (vCPU) in VMware virtual machines appear to the operating system as single core CPUs. It maps onto a logical thread on the physical CPU. When hyperthreading is enabled, each physical core has two logical threads. In order to assign all of the CPU resources to virtual machines, the number of virtual CPUs assigned needs to equal the number of logical threads on the server (for example, an Intel Xeon processor E7 series based server that has 18 cores per socket). In a 4-socket version, it has a total of 72 physical cores and 144 logical threads. The vSphere scheduler prefers to assign a vCPU to an idle core. Therefore, hyperthreads (thread 1) are used only when all thread 0 of all cores are already utilized.

For memory latency sensitive workloads with low processor utilization, such as SAP HANA, or high interthread communication, we recommended using hyperthreads with fewer NUMA nodes instead of full physical cores spread over multiple NUMA nodes. To ensure the usage of hyperthreads and to enforce NUMA node locality, set the **Numa.PreferHT=1** parameter as documented in KB 2003582.

Hyperthreading for processors allows for multiple instruction threads to execute on a single physical core. While many of the core's resources are actually shared between cores, the additional logical thread allows for an increase in performance, usually in the range of 10 percent to 20 percent. We recommend enabling hyperthreading on any system running vSphere and SAP HANA. For HANA sizing purposes, we calculate with a hyperthreading gain of 15%.

When configuring SAP HANA virtual machines for production environments, ensure that the total vCPU resources for the virtual machines running on the system do not exceed the CPU capacity of the host. Do not overcommit CPU resources on the host. If the host CPU capacity is overloaded, the performance of the virtual database might degrade.

The VMware "Guest and HA Application Monitoring Developer's Guide"[56] provides some information on how to determine if a VM was sized correctly by leveraging the vm.cpu.contention.cpu metric, which is described in this guide. As stated, a contention value of < 5 percent is normal "undercommit" operating behavior, representing minor hypervisor overheads. A contention value > 50 percent is "severe overcommit" and indicates CPU resource starvation: the workload would benefit from either adding CPUs or migrating VMs to different hosts. A contention value between 5 percent and 50 percent is "normal overcommit."

The VMware "Guest and HA Application Monitoring Developer's Guide" further describes that when actual usage is below compute capacity, the hypervisor is "undercommitted." In this case, the hypervisor is scaling linearly with the load applied, and there is wasted capacity. As actual usage exceeds available compute capacity, the hypervisor begins utilizing hyperthreads for running virtual machines to ensure performance degradation is not disruptive; maximum aggregate utilization occurs during this "normal overcommit" (between 5 percent and 50 percent contention), where each virtual machine sees somewhat degraded performance, but overall system throughput is still increasing. In this "normal overcommit" region, adding load still improves overall efficiency, though at a declining rate. Eventually, all hyperthreads are fully utilized. Efficiency peaks and begins to degrade; this "severe overcommit" (> 50 percent contention) indicates the workload would be more efficient if spread across more hosts for better throughput.

Often in with SAP HANA workloads, when a virtual machine is configured to match the number of host cores (not including hyperthreads), it will peak at the capacity of those cores (with < 5 percent contention), but at a performance between 10 percent and 20 percent lower than an equivalent physical machine utilizing all cores and hyperthreads at all times. A virtual machine configured to match the number of host threads (2 x host cores) will peak at a performance level more analogous to a physical machine. A recent SAP HANA BW-EML benchmark[57] has shown that the difference can be as little as under 4 percent, but will show ~40 percent contention (the upper end of "normal overcommit") from running half the cores on hyperthreads. This contention metric indicates the load would run better when run on a larger host with additional cores, so it is technically "overcommitted," even though performance is better than a hypervisor running at full commit.

Typical SAP HANA VMs should be configured to be within the "normal overcommit" range, as shown in Figure 31.

[56] http://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vs600_guest_HAappmon_sdk.pdf, page 25++

[57] SAP BW-EML VCE vSphere SAP HANA Scale-out Benchmark, certification number 2015016

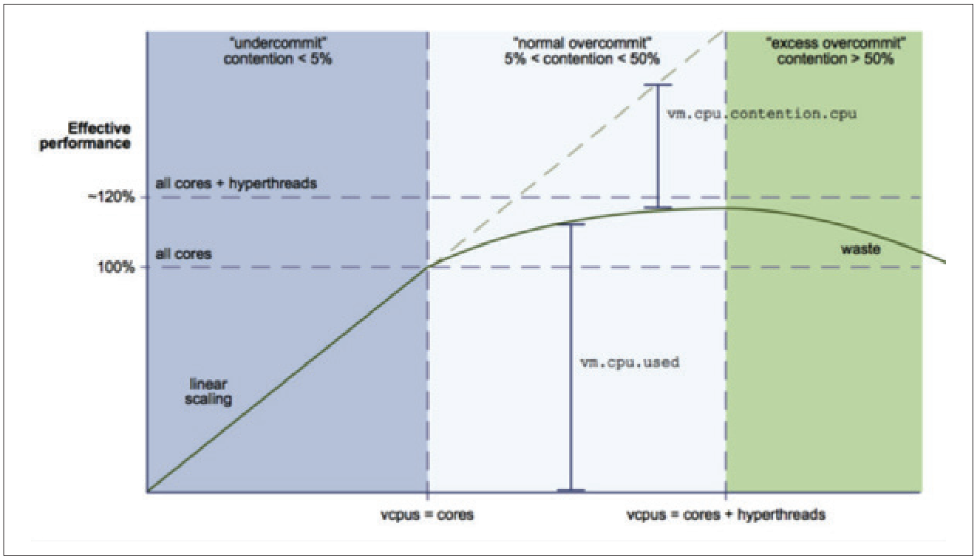

Figure 31: CPU Utilization Across All Virtual Machines

Configuring virtual SAP HANA with excess vCPUs can impose a small resource requirement on vSphere because unused vCPUs continue to consume timer interrupts. vSphere co-schedules virtual machine vCPUs and attempts to run the vCPUs in parallel to the best extent possible. Unused vCPUs impose scheduling constraints on the vCPU being used and can degrade its overall performance.

Note: Right sizing of SAP HANA VMs is important; configuring virtual SAP HANA with overcommitted or an excess of vCPUs may degrade a VMs overall performance.

Memory Considerations

Regarding memory, the best practice is to configure memory reservations equal to the size of the SAP HANA configured RAM. Do not overcommit memory. When consolidating multiple non-production SAP HANA instances on the same host, vSphere can share memory across all the virtual machines that are running the same operating system. In this case, vSphere uses a proprietary, transparent page-sharing technique to reclaim memory. This allows databases to run with less memory than when physical RAM is used.

In this case, leave the memory for overhead of the vSphere ESXi kernel and for the virtual machines. As a conservative estimate, use up to 3 percent of system memory for this overhead. Do not assign memory overhead to virtual machines. A more optimistic overhead figure is 0.5 percent, which we have also seen during the SAP HANA CA phases at customer sites. The actual memory overhead depends on the actual VM configuration, such as the number of SCSI or network adapters.

Note: The real memory overhead figure can be defined only when verifying the actual vSphere host and all of its running VMs. It may be 3 percent or much lower. For sizing calculations, we use an average figure of 1.5%.

For more information about memory overheads, see the "vSphere Resource Management Guide 6.0"[58] in the "Administering Memory Resources" section.

Non-Uniform Memory Access

Beside the VM memory limitations, consider the actual server architecture in terms of CPU and NUMA architecture. This is important for optimizing virtualized SAP HANA performance and ensuring the best response time for users.

Note: Unlike with physically deployed SAP HANA systems where all available compute resources such as CPU (NUMA nodes) and RAM resources will be allocated and shared by the installed HANA instance and its tenants or possibly (A)SCS instances (SoH deployments) in a virtualized environment, you can optimize the resource commitments and NUMA node placement and usage for every HANA instance.

The non-uniform memory access architecture is common on servers with multiple processor sockets. It is important that you understand this memory access architecture when sizing memory-bound and latency-sensitive systems.

A NUMA node is equivalent to one CPU socket. For a 2-socket server there are two NUMA nodes, for a 4-socket server there are four NUMA nodes, and so on. Therefore, the available number of physical CPU cores and RAM can be divided equally among NUMA nodes. This is critical when sizing virtual machines for optimal NUMA node placement and utilization.

For example, a 4-socket, 72-core (18 cores per socket) Intel EX E7-88x0v3 based server system with 2 TiB RAM has four NUMA nodes, each with 18 CPU cores, 36 CPU threads, and 512 GiB RAM (4096 GiB divided by four) per NUMA node/CPU socket.

When you size virtual machines, carefully consider NUMA node boundaries. Exceeding the CPU and RAM size boundaries of a NUMA node causes the virtual machine to fetch memory from a remote location, and this can diminish performance. Figure 32 shows a small, single NUMA node VM and a VM that spans more NUMA nodes, which is also called a wide VM. An SAP HANA VM can be as large as 8 NUMA nodes on Intel Xeon E7-v3 (Haswell) server systems by maintaining the SAP defined RAM to NUMA node ratio.

[58] http://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-resource-management-guide.pdf

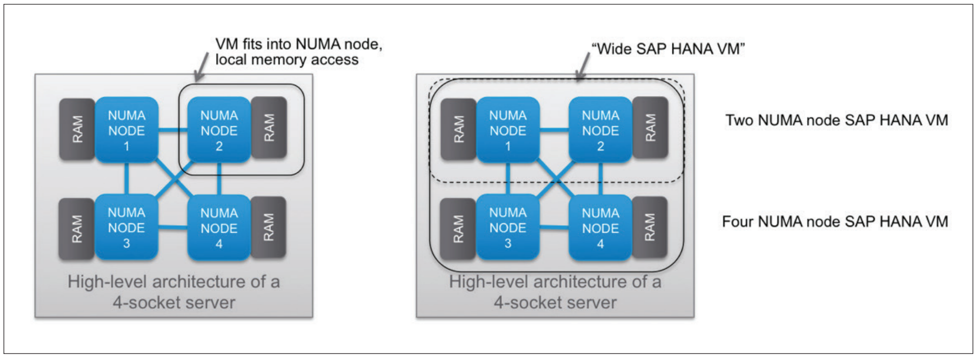

Figure 32: Node Boundaries and VM Placement

Figure 33 shows a graphical view of the node boundaries and highlights the local and remote memory access. As with any other application, SAP HANA VMs should be sized to stay within a NUMA node for as long as possible to optimize performance.

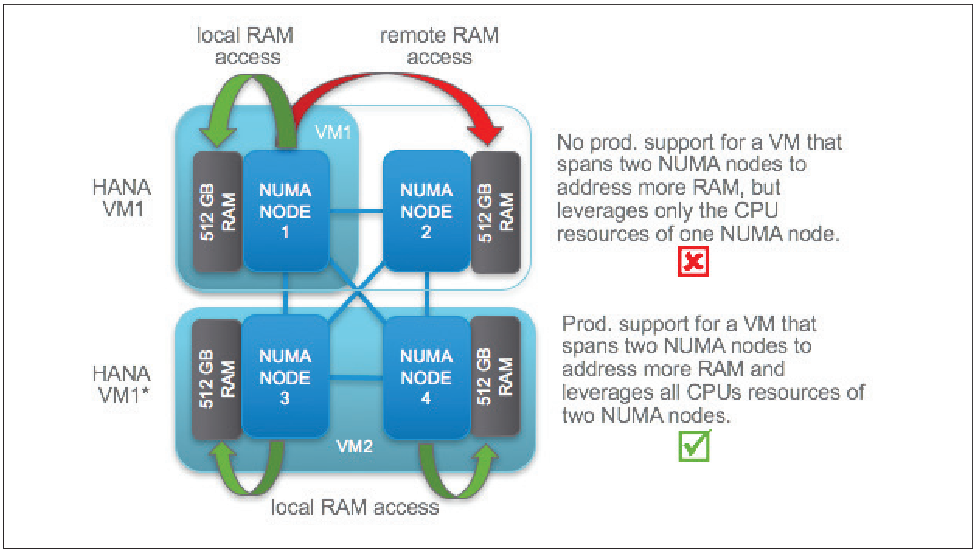

Figure 33: SAP HANA on vSphere – NUMA Best Practices

Note: As with any other application, whenever possible, size an SAP HANA VM with the least number of NUMA nodes.

When an SAP HANA VM needs to be larger than a single NUMA node, allocate all resource of this NUMA node; and to maintain a high CPU cache hit ratio, never share a NUMA node for production use cases with other workloads. Also, avoid allocation of more memory than a single NUMA node has connected to it because this would force it to access memory that is not local to the SAP HANA processes scheduled on the NUMA node. Figure 34 provides an overview of NUMA node sharing and support status.

Note: For production workloads, NUMA node sharing is yet not supported because it may impact memory performance. Future joint testing with SAP will investigate the real impact of NUMA node sharing on SAP HANA, and will provide additional guidelines on how to take advantage of virtualization and its resource isolation capabilities by utilizing CPU and memory resources aligned to the sized SAP HANA VM configuration.

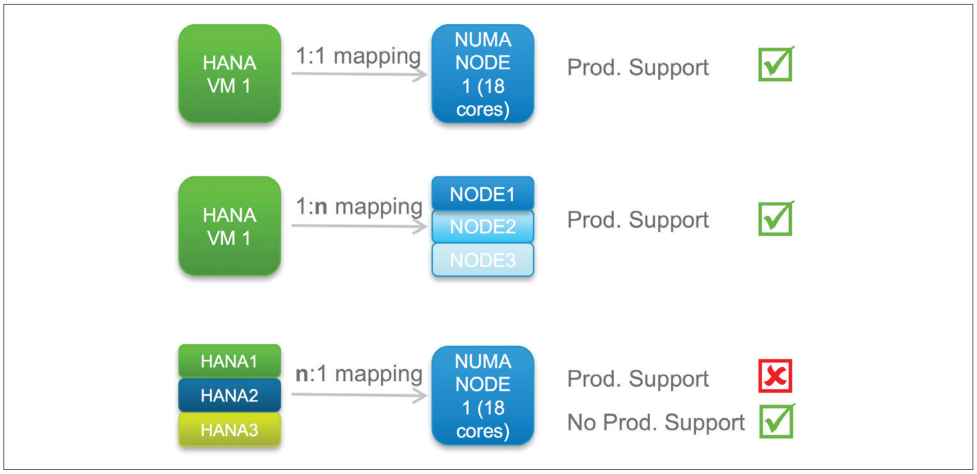

Figure 34: NUMA Node Sharing Example with Haswell 18-Core System

Figure 35 shows a small 18-CPU core SAP HANA VM that fits into a NUMA node of an Intel Xeon E7-v3 (Haswell) based server system, and shows the correlation between CPU cores and threads. To benefit from hyperthreading, the hyperthreads have to be used as well and 36 vCPUs have to be configured for this VM. The VM has less than the available 512 GiB configured, ensuring that the hypervisor is always able to start this VM on a single NUMA node.

We recommend configuring SAP HANA VMs with at least 10 vCPUs. Using a 18-core, 36-thread CPU, this would allow for up to three 12-vCPU SAP HANA non-production VMs on one CPU socket.

Note: Use at least 10 vCPUs (5 cores and 5 hyperthreads) for an SAP HANA VM. With fewer than 10 vCPUs, the SAP HANA system may become slow and unresponsive because it does not have enough threads to support all needed OS and SAP HANA processes.

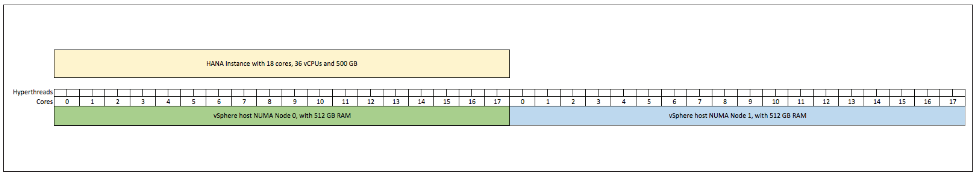

Figure 35: SAP HANA VM with 36 vCPUs and 500 GiB RAM, Sized to Fit into a Single NUMA Node (CPU Socket) of an 18-Core CPU with 512 GiB RAM Attached

To ensure the 18 additional threads are scheduled on the local hyperthreads (NUMA node), add the parameter **Numa.PreferHT=1** to the VM or, alternatively, to the host configuration to ensure only the threads of the local NUMA nodes are accessed, and not the threads of another NUMA node. This will break the memory locality. A risk is that when all 36 vCPUs are heavily utilized, the CPU contention value may increase (vm.cpu.contention.cpu). View details about CPU contention in the VMware vSphere SAP HANA sizing guidelines.

Large SAP HANA VMs that consume more than two NUMA nodes (CPU sockets) may not be able to use all available CPU threads (64 or 128 vCPU limit). In this case, where more CPU threads are available as vSphere can address per VM, the ESXi kernel will always prefer the physical CPU core over the hyperthread.

Note: Use parameter Numa.PreferHT=1 per VM to ensure that hyperthreads are used instead of potentially idle CPU cores of other NUMA nodes. This ensures local thread local memory access.

In the sizing section of this document, you will find configuration examples on how to configure SAP HANA VMs on 4- and 8-socket server systems.

SAP HANA on vSphere Network Considerations

For SAP HANA, the networking configuration includes virtual switch and virtual network card optimization. The virtual switch configuration is done on the hypervisor and the virtual network card configuration is done inside the VM. Each is described in the sections below.

vSphere offers “standard and distributed” switch configurations. Both switches can be used when configuring an SAP HANA in a vSphere environment. A vSphere standard switch is very similar to a physical Ethernet switch. It can bridge traffic internally between virtual machines in the same virtual LAN (VLAN), and link to external networks by connecting the physical network interface cards (NICs) of the hosts to uplink ports on the standard switch.

Virtual machines have network adapters you connect to port groups on the standard switch. Every port group can use one or more physical NICs to handle their network traffic. If a port group does not have a physical NIC connected to it, virtual machines on the same port group can communicate only with each other, but not with the external network. Detailed information can be found in the vSphere networking guide[59]. Figure 36 shows the vSphere standard switch architecture and how virtual machines are connected to a virtual switch.

[59] https://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-networking-guide.pdf or https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-networking-guide.pdf

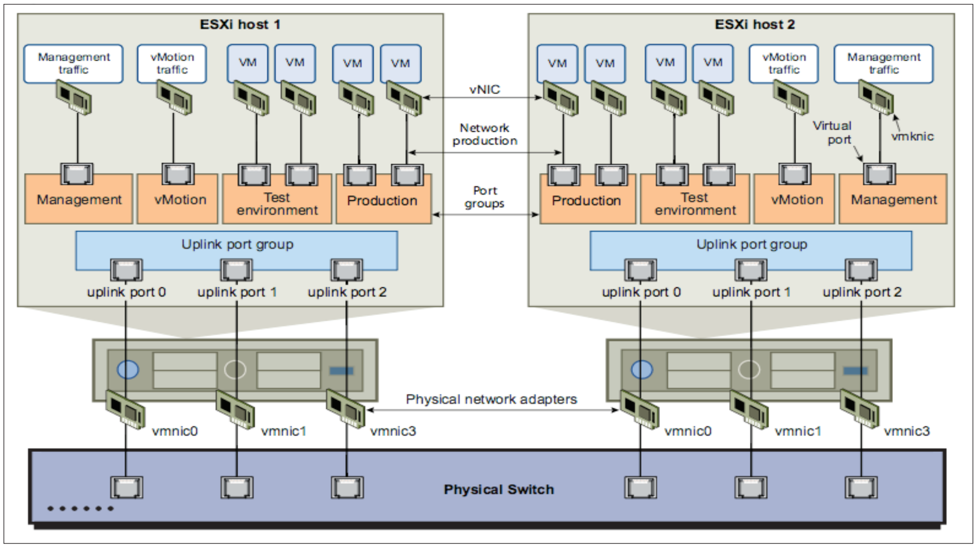

Figure 36: vSphere Standard Switch Architecture

We recommend using a vSphere distributed switch in SAP HANA environments because it provides centralized management and monitoring of the networking configuration of all SAP HANA hosts that are associated with the switch. You can set up a distributed switch on a vCenter Server system and its settings are propagated to all hosts associated with the switch. This ensures network consistency and an error-free configuration. It also provides enhanced security and monitoring for virtual machines migrated via VMware vMotion through maintenance and migration of port run-time state and enhanced provisioning and traffic management capabilities through private VLAN support and bidirectional virtual machine rate-limiting.

The type of virtual switch an IT organization decides to use for SAP HANA depends on the VMware vSphere license the organization has purchased. The minimum requirements for setting up SAP HANA on vSphere are fulfilled with a standard switch configuration.

Note: For easier management and host consistent configuration, we recommend the use of virtual distributed switches; Scale-Out configurations particularly benefit from virtual distributed switches.

The actual network configuration inside the VM and for the SAP HANA instance is 100%—the same as it would be for a physically deployed system. The only difference is that by leveraging the possibility of a virtual switch and virtual NIC, fewer physical network cards can be used.

Figure 37 shows the SAP HANA logical network connections as described in the "SAP Network Requirement"[60] document.

[60] http://scn.sap.com/docs/DOC-63221

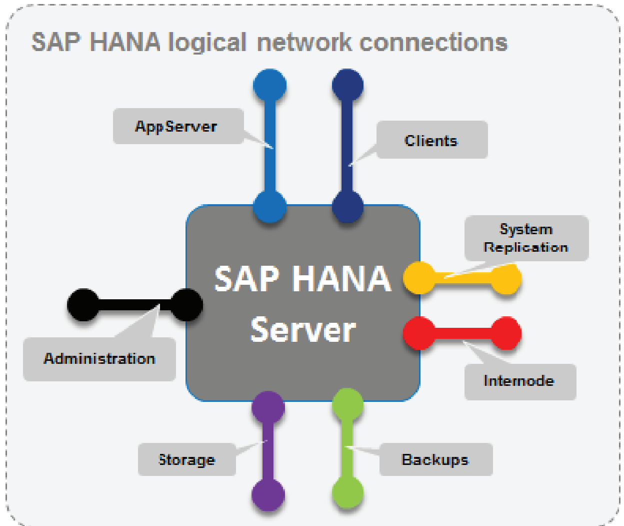

Figure 37: Logical Network Connections per SAP HANA Server

SAP divides these networks into three groups:

- Client Zone – External (e.g., application server and client-facing network)
- Internal Zone – Internode Network (e.g., Scale-Out network and replication network)
- Storage Zone – All storage-related networks (e.g., NFS shared SAP HANA or backup volume network and NFS Data and Log when NFS is used)

Based on this network role grouping, and the SAP network bandwidth and isolation recommendations in the "SAP HANA Server Installation and Update Guide"[61] and the "SAP Network Requirement" document, Table 10 shows the recommended network configuration for virtualized SAP HANA systems.

The application server virtual network shares the same physical NIC as the admin, VMware HA, and backup network. If required, use dedicated physical NICs for these networks per Table 10. The backup and replication networks are optional. In the table, the backup network uses the shared pNIC 0.

[61] http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf, page 15

		ADMIN AND VMWARE HA NETWORK	SAP HANA SHARED (NFS)	APPLICATION SERVER NETWORK	BACKUP NETWORK*	vMOTION NETWORK	SAP HANA INTERNODE NETWORK	SAP HANA NFS STORAGE-N**	SYSTEM REPLICATION NETWORK
Host Network Configuration	Network Label	Admin	HANA-Shared	App-Server	Backup	vMotion	HANA-Intra	HANA-NFS-n	HANA-Repl
	Bandwidth	1GbE	1GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE
	MTU Size***	Default	Default	Default	9,000	9,000	9,000	9,000	9,000
	vmNIC (pNIC)	0			1		2	3**	Optional
	VLAN	0	1	2	3	4	5	****	Optional
VM Configuration	vNIC	0	1	2	3	-	4	****	Optional
Network required for deployment option is marked with an X.									
Deployment Option	Scale-Up	X	Optional	X	Optional*	X	Not required	****	Optional
	Scale-Out NFS	X	X	X	Optional*	X	X	X	Optional
	Scale-Out FC	X	X	X	Optional*	X	X	Not required	Optional

*Backup network is optional when other backup methods, such as storage-based backups, are used.

**One dedicated 10GbE NIC per NFS SAP HANA VM system that runs on a vSphere host (data+log combined); for example: One HANA VM per host -> 1 x 10GbE; two HANA VMs per host.

-> 2 x 10GbE NICs for NFS storage.

*** Default MTU size as defined by IT provider.

**** Needed when NFS storage is selected. FC- or NFS-based NAS can be selected for Scale-Up.

Table 10: SAP HANA vSphere Minimum Recommended Host and VM Network Configuration

Table 11 shows the network configuration when dedicated network cards are used for backup and replication, which may be required when SAP HANA system replication is used. If storage replication is used, this network is not required.

		ADMIN AND VMWARE HA NETWORK	SAP HANA SHARED (NFS)	APPLICATION SERVER NETWORK	BACKUP NETWORK*	vMOTION NETWORK	SAP HANA INTERNODE NETWORK	SAP HANA NFS STORAGE-N**	SYSTEM REPLICATION NETWORK
Host Network Configuration	Network Label	Admin	HANA-Shared	App-Server	Backup	vMotion	HANA-Intra	HANA-NFS-n	HANA-Repl
	Bandwidth	1GbE	1GbE	10GbE	10GbE	10GbE	10GbE	10GbE	10GbE
	MTU Size***	Default	Default	Default	9,000	9,000	9,000	9,000	9,000
	vmNIC (pNIC)	0			1	2	3	4**	Optional
	VLAN	0	1	2	3	4	5	****	Optional
VM Configuration	vNIC	0	1	2	3	-	4	****	Optional
Network required for deployment option is marked with an X.									
Deployment Option	Scale-Up	X	Optional	X	Optional*	X	Not required	****	Optional
	Scale-Out NFS	X	X	X	Optional*	X	X	X	Optional
	Scale-Out FC	X	X	X	Optional*	X	X	Not required	Optional

*Backup network is optional when other backup methods, such as storage-based backups, are used.

**One dedicated 10GbE NIC per NFS SAP HANA VM system that runs on a vSphere host (data+log combined); for example: One HANA VM per host -> 1 x 10GbE; two HANA VMs per host.

-> 2 x 10GbE NICs for NFS storage.

*** Default MTU size as defined by IT provider.

**** Needed when NFS storage is selected. FC- or NFS-based NAS can be selected for Scale-Up.

Table 11: SAP HANA vSphere Bandwidth Optimized Host and VM Network Configuration

Figure 38 shows the network diagram for an SAP HANA Scale-Out system running on vSphere by leveraging a virtual distributed switch. The network diagram of a Scale-Up system that runs on a shared vSphere resource cluster is similar because all the vSphere hosts need access to the same networks. Only the Scale-Out internode communication can be removed. In a virtualized network, the VM network configuration complexity is hidden from the outside world because only the physical network links of the server NICs and the VLAN configuration are exposed.

The SAP HANA-NFS-Storage-x network is optional, and is needed only when NAS-NFS storage is used for SAP HANA data and log files. If FC-SAN storage is used, this network is not required.

Note: We recommend that each vSphere host has redundant network connections. We also recommend jumbo frames (with an MTU size of 9,000) for all networks but the client network. The client network MTU size is determined by the IT standards of your organization.

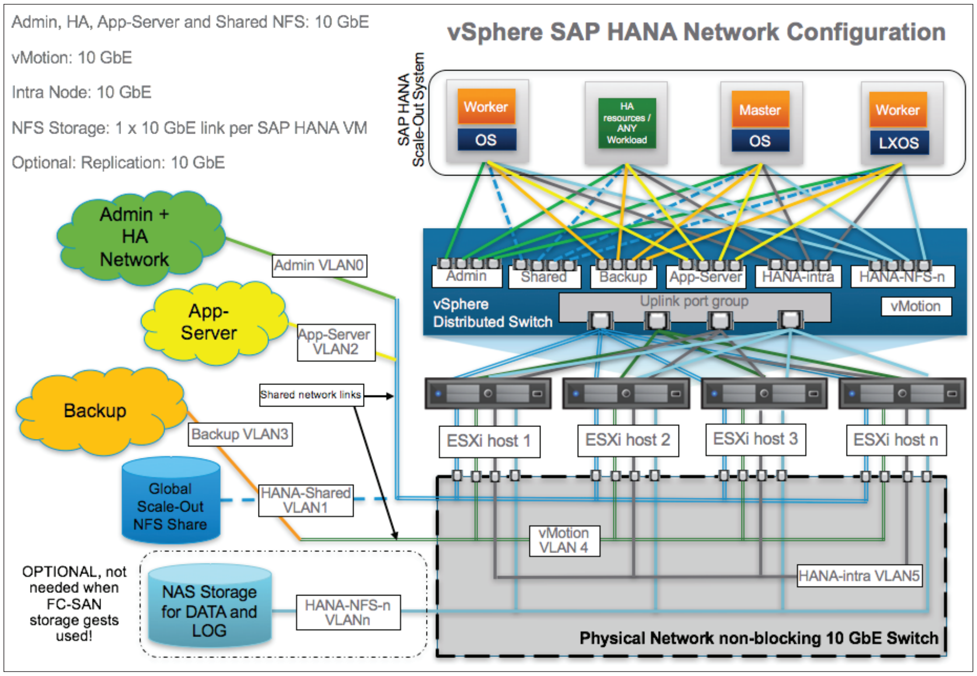

Figure 38: vSphere SAP HANA Scale-Out Network Configuration of a Virtualized 3+1 SAP HANA Scale-Out System

Note: The Scale-Out system in Figure 35 is very similar to an SAP HANA Scale-Up system because a vSphere resource cluster is always a group of servers that share the same network and need therefore access to these networks.

Depending on the used storage, additional 10GbE NICs or FC host bus adapters and switches may be required to satisfy the storage network bandwidth that additional SAP instances require. In Table 10 and Table 11, the vmNIC (pNIC) row shows the physical network cards, which must be adapted to the unique customer configuration and, especially when NFS is used, must follow the NFS storage vendors network sizing recommendations.

Note: When adding new SAP HANA Scale-Out nodes or additional SAP HANA VMs on an existing node, close network monitoring and planning is required before adding these new nodes because the network card may become a performance or latency bottleneck.

For more information on the network requirements of NAS-NFS storage based SAP HANA systems, review VMware the documents "Best Practices for Running VMware vSphere on Network Attached Storage"[62] or "SAP HANA on NetApp FAS Systems with NFS Configuration Guide."[63]

[62] http://www.vmware.com/files/pdf/techpaper/VMware-NFS-BestPractices-WP-EN.pdf

[63] http://www.netapp.com/us/media/tr-4338.pdf

Virtualization Layer

VMware vSphere, the technical foundation of a virtualized infrastructure, has been described in detail already. In addition, relevant information on how to size and configure SAP HANA ready VMs has been discussed. This section provides additional information on the different vSphere versions and provides guidelines on which version you should use.

Customers can choose from three editions of vSphere and vSphere with Operations Management, which are supported for SAP virtualization. As of June 30, 2016, VMware vSphere® Enterprise Edition™ and VMware vSphere® with Operations Management™ Standard/Enterprise are no longer available[64]. All other editions, such as VMware vSphere® Essentials Kit, are not supported by VMware for use with SAP applications.

VMware vSphere® Standard Edition™ provides an entry-level solution for basic server consolidation to slash hardware costs while accelerating application deployment.

VMware vSphere® Enterprise Plus Edition™ offers the full range of vSphere features for transforming data centers into dramatically simplified cloud infrastructures for running today's applications with the next generation of flexible, reliable IT services.

VMware vSphere® with Operations Management™ Enterprise Plus Edition offers the full range of vSphere features for transforming data centers into dramatically simplified virtualized infrastructures for running today's applications with the next generation of flexible, reliable IT services. Best-in-class management capabilities allow IT to gain visibility into virtual environments, proactively identify and remediate emerging performance issues, and optimize resource utilization through a unified console.

For more information on the different VMware vSphere editions, check out: https://www.vmware.com/products/vsphere/compare. For local currency prices for vSphere and vSphere with Operations Management Editions, visit: http://www.vmware.com/products/datacenter-virtualization/vsphere/pricing.html.

SAP HANA relevant features such as DRS, I/O control, SR-IOV, host profiles, or the distributed switch feature are available only in Enterprise Plus or later.

Note: Recommended and supported vSphere versions for SAP HANA are vSphere Enterprise or later. As of June 30, only vSphere Enterprise Plus with and without Operations Management are available.

Figure 39 shows the supported versions and provides an overview of the features supported with each of the editions. VMware vSphere gets licensed per CPU, with a vCenter Server sold separately.

[64] https://www.vmware.com/files/pdf/products/vsphere/VMware-vSphere-vSOM-Pricing-FAQs.pdf

	vSphere Standard	vSphere Enterprise Plus	vSphere with Operations Management Enterprise Plus
Overview	Server consolidation and business continuity	Resource management, enhanced application availability and performance	Intelligent operations, consistent management and automation with predictive analytics
License entitlement	Per 1 CPU	Per 1 CPU	Per 1 CPU
vCenter Server (sold separately)	vCenter Server Standard	vCenter Server Standard	vCenter Server Standard
VMware Integrated OpenStack		Support for VMware Integrated OpenStack is sold separately	Support for VMware Integrated OpenStack is sold separately

Product Features	Standard	vSphere Enterprise Plus	vSphere with Operations Management Enterprise Plus
Business Continuity and Security			
vMotion	• (+ Cross vSwitch)	• (+ Cross vSwitch / Cross vCenter / Long Distance)	• (+ Cross vSwitch / Cross vCenter / Long Distance)
Storage vMotion	•	•	•
High Availability	•	•	•
Data Protection	•	•	•
Fault Tolerance	2-vCPU	4-vCPU	4-vCPU
vShield Endpoint	•	•	•
vSphere Replication	•	•	•
Hot Add	•	•	•
Resource Prioritization and Enhanced Application Performance			
Virtual Volumes	•	•	•
Storage Policy-Based Management	•	•	•
Reliable Memory		•	•
Big Data Extensions		•	•
Virtual Serial Port Concentrator		•	•
Distributed Resources Scheduler (DRS), Distributed Power Management (DPM)		•	•
Storage DRS		•	•
Storage I/O Control		•	•
Network I/O Control		•	•
Single Root I/O Virtualization (SR-IOV) Support		•	•
NVIDIA GRID vGPU		•	•
Automated Administration and Provisioning			
Content Library	•	•	•
Storage APIs for Array Integration, Multipathing	•	•	•
Distributed Switch		•	•
Host Profiles and Auto Deploy		•	•
Operations Management			
Consistent Management			•
Intelligent Operations			•
Operations Automation			•
Workload Balancing			•

Figure 39: vSphere Versions–Compared[65]

Besides selecting the right vSphere license, it is important to select the **right service contract**. For business-critical SAP environments, we recommend selecting at least a business-critical (BCS) or mission-critical (MCS) support contract. For details about our support contracts, see the webpage: https://www.vmware.com/support/services.html.

[65] http://www.vmware.com/products/vsphere/compare.html

Application and OS Layer

SAP also requires certified operating systems (OS) similar to the hardware certification SAP performs rigid tests and demands from their OS partner's joint testing's and SAP HANA certification.

As of September 2016, only SUSE and Red Hat Linux operating systems are supported for use with SAP HANA. Released Linux Distributions for HANA can be found in SAP Note 2235581 - SAP HANA: Supported Operating Systems. Please check linked SPS release notes. Note that not all Linux versions and patch versions are supported with all available SAP HANA certified systems. Hardware vendor-specific support status and limitations are listed in the SAP HANA Hardware Directory. Use the OS filter settings to verify if the OS version and support pack is supported by the selected SAP HANA appliance as shown in Figure 40.

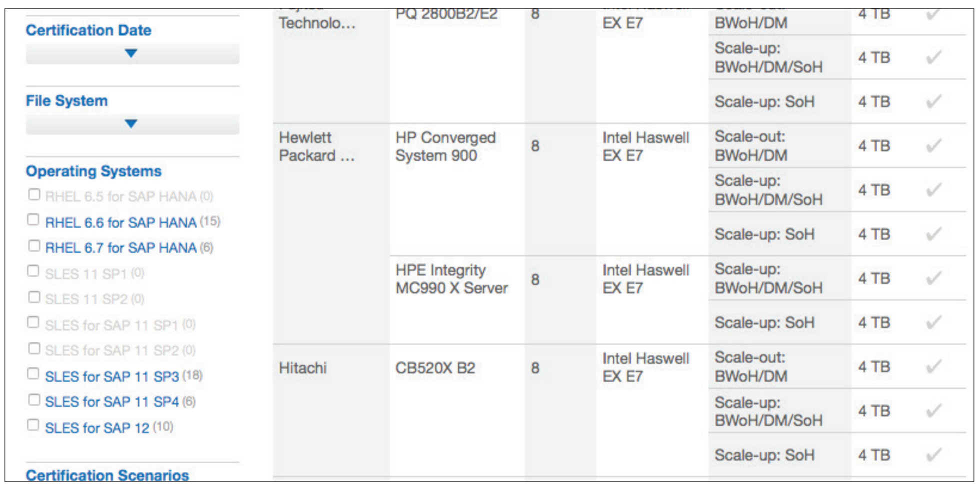

Figure 40: SAP HANA Hardware Directory–OS Filter Criteria

Note: The selected OS release must be supported by all partners—released for SAP, and supported by the hardware partner and by VMware to ensure full virtualized SAP HANA support.

Also ensure that the SAP HANA version and service pack level you want to use are supported with the entire stack (HW, vSphere, and OS).

Management Layer

As shown in Figure 24. VMware vCenter Server, the main component of the management layer of a VMware virtualized infrastructure is the vCenter Server. VMware vCenter Server provides a centralized and extensible platform for managing virtual infrastructure. vCenter Server manages VMware vSphere environments, giving IT administrators simple and automated control over the virtual environment to deliver infrastructure with confidence.

More information on the vCenter features: https://www.vmware.com/files/pdf/products/vCenter/VMware-vCenter-Server-Datasheet.pdf

Availability Considerations for vCenter Server

vCenter Server can be a single point of failure in the environment. Many VMware solutions are layered on top of vCenter Server. Its availability impacts the usability of these components. Loss of vCenter Server curtails the ability of these solutions to perform many critical functions, such as making changes and creating new virtual machines (VMs).

There are multiple high-availability options for VMware vCenter Server 6.0. VMware vSphere High Availability, VMware vSphere Fault Tolerance, and watchdog processes can all be leveraged to protect vCenter Server services. Based on customer requirements, multiple deployment modes can be leveraged for availability in local and multisite configurations.

For more information on availability solutions for vCenter, read the guidelines published in the "VMware vCenter Server™ 6.0 Availability Guide."[66]

SAP HANA on vSphere High Availability (HA)

SAP HANA offers several methods for high availability (HA) and disaster recovery. There are auto failover, service restart options, backups, system replication, and standby host systems. In VMware virtualized environments, additional VMware HA can be invoked to minimize unplanned downtime due to faults.

In the introduction section of VMware HA it got already discussed how vSphere is able to provide fault tolerance by supporting redundant components such as dual network and storage pathing or the support of hardware solutions such as UPS, or the support of CPU built-in features that allow to tolerate failures in memory models or that ensure CPU transaction consistency. But what happens when a fault or a disaster happens? This section provides the answer to these questions and provides an overview on HA solutions available for VMware virtualized SAP HANA systems.

High availability support can be separated in two different areas: fault recovery and in disaster recovery. As an overview, following summary:

High availability, by providing "fault recovery":

- SAP HANA Service Auto-Restart
- Host Auto-Failover (standby host)
- VMware HA
- SAP HANA system replication

High availability, by providing "disaster recovery":

- Backup and Restore
- Storage Replication
- VMware Replication
- SAP HANA system replication

[66] https://www.vmware.com/files/pdf/techpaper/VMware-vCenter-Server-6-0-Availability-Guide.pdf

In the above list, we see SAP HANA system replication mentioned in both recovery scenarios. Depending on the customer requirements for very low RTO (recovery time objective–in minutes), SAP HANA system replication may be the only possible solution to recover quickly from disasters.

Different recovery point objectives (RPOs) and recovery time objectives (RTOs) can be assigned to different faults disaster recovery solutions. SAP describes[67] the phases of high availability in their document and Figure 41 shows a graphical view of these phases.

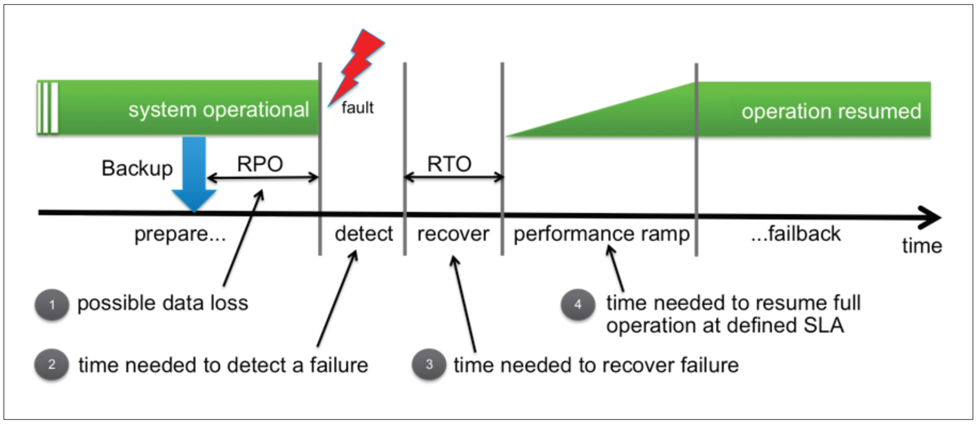

Figure 41: SAP HANA System High Availability Phases

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

Customers have different requirements on allowed downtime (RTO) and amount of possible data loss (RPO). Minimizing RTO and RPO with the available IT budget and resources should be the goal, and is the responsibility of the IT team operating SAP HANA.

Note: Selecting the right high availability solution is a typical conflict between cost, speed, and quality; you need to find a compromise that meets the customer requirements.

---

[67] http://scn.sap.com/docs/DOC-60334

SAP HANA High Availability Solutions Overview

Table 12 provides an overview of the different available SAP HANA failure and disaster recovery solutions when running on VMware vSphere, and allows you to select the right solution for the unique customer requirements.

The complexity and cost increase with certain solutions and with high RPO and RTO requirements. The VMware HA solution is a good choice for most fault recovery scenarios because it is included in vSphere licensing, and is very easy to implement and operate. From the feature side, it is comparable to the SAP HANA Host Auto-Failover, which is one option when running SAP HANA in a non-virtualized environment.

When implementing an HA strategy, the execution of a backup and restore solution is the first step. It is the basis for any disaster recovery, such as recovering lost data due to logical or physical hardware faults, which no other HA solution can provide.

	Failure Recovery			**Disaster Recovery**				
HIGH AVAILABILITY SOLUTION	VMWARE HA + SAP HANA AUTO-Service RESTART FEATURE	VMWARE HA + EXTERNAL MONITORING SOLUTION	SAP HANA HOST AUTO-FAILOVER (STANDBY VM)	SAP HANA BACKUP	SAP HANA SYSTEM REPLICATION	STORAGE SYSTEM REPLICATION	VMWARE VSPHERE REPLICATION	STORAGE OR VMWARE SYSTEM REPLICATION + VMWARE SRM
Scenario description	Standard VMware HA combined with SAP HANA Service Auto-Restart watchdog running inside a VM to monitor SAP HANA application status and triggers an SAP HANA process restart. OS and HW failures will get handled by VMware HA.	VMware HA combined with an external SAP HANA application monitoring solution that detects SAP HANA failures not covered by the SAP HANA watchdog. In the failure situation the application gets re-started or he whole VM gets re-booted.	SAP HANA Standby VM to take automatically over the role of another SAP HANA VM, in case of a detected failure. Only possible when a storage vendor provides the required storage connector API and STONIT script.	Allows Disaster Recovery, Lowest cost, simplest Supports point-in-time recovery Can also be used to "clone" or copy systems, is the only method to recover from logical failures like unintentional data deletion.	Data replication between primary and secondary sites by leveraging SAP HANA System Replication functionality; there is no automated failover process. Third-party cluster solutions can be used to automate site failover.	Data replication between primary and secondary site by leveraging storage system replication functionality. No automated failover process. Third-party cluster solution can get used to automate site failover.	VMware Replication to replicate VMs to a DT site. VMware Replication is not recommended for the usage of SAP HANA or other applications with high RPO requirements. System with low change volume can get replicated with VMware replication.	VMware SRM to automate the VMware or storage replicated SAP HANA system VMs and all related servers to another site. No SAP HANA System Replication integration yet, VMware replication only for non-prod systems!
Operating system failures	Yes	Yes	Yes	-	Yes	Yes	Yes	Yes
Hardware failures	Yes	Yes	Yes	-	Yes	Yes	Yes	Yes
Application failures	Yes, via HANA watchdog	Yes, via external solution	Yes, via HANA watchdog	-	Yes	Yes	Yes	Yes
IP redirect/DNS update	Not necessary	Not necessary	Not necessary	Not necessary	Yes, manual or third-party software, such as cluster manager solution, is needed	Not necessary as complete VMs with all data disks get replicated	Not necessary as complete VMs with all data disks get replicated	Not necessary as complete VMs with all data disks get replicated
Logical Failures	No	No	No	Yes	No	No	No	No
Split Brain Situation	Not possible	Not possible	Possible, STONIT script required	Not possible	Possible, initial system must get turned off	Possible, initial system must get turned off	Possible, initial system must get turned off	SRM manages the failover and ensures power off of old systems!
SAP HANA Storage Connector API	Not necessary	Not necessary	Required	Not necessary	Not necessary	Not necessary	Not necessary	Not necessary
RTO	Medium (crash recovery of DB)	Medium (crash recovery of DB)	Short to Medium (crash recovery of DB, short if in-memory data loading)	Long and manual restart (restore of DB)	Shortest to medium, depending on IP redirect solution (crash recovery of DB or in-memory data loading)	Medium to long, depending on the storage failover process (crash recovery of DB)	Medium to long, depending on the storage failover process (crash recovery of DB)	Fast, automated system restart via SRM (crash recovery of DB)
RPO	0	0	0	Depends on the backup interval	0 (synchrony) >0 (asynchrony)	0 (synchrony) >0 (asynchrony)	> 15 minutes (asynchrony)	0 (synchrony) >0 (asynchrony
Performance ramp	Minutes to hours, depending on the database size (bootstrap + loading)	Minutes to hours, depending on the database size (bootstrap + loading)	Minutes to hours, depending on the database size (bootstrap + loading, short if in-memory data loading)	Hours to days	Seconds to minutes if synchronous replication into memory is selected, depending on the selected modus and IP redirect solution.	Hours (bootstrap + loading)	Hours (bootstrap + loading)	Hours (bootstrap + loading)
Complexity	LOW	MEDIUM	HIGH	LOW	HIGH	HIGH	MEDIUM	MEDIUM
Cost	INCLUDED	MEDIUM	INCLUDED	INCLUDED	HIGH	HIGH	HIGH	HIGH

Table 12: SAP HANA on vSphere High Availability Solutions

Fault Recovery

SAP HANA Service Auto-Restart

"In the event of an SAP HANA software failure, that disables one of the configured SAP HANA services (Index Server, Name Server, etc.), the service will be restarted by the SAP HANA **Service Auto-Restart** *watchdog* function, which automatically detects the failure and restarts the stopped service process. Upon restart, the service loads data into memory and resumes its function. While all data remains safe (RPO=0), the service recovery takes some time." (Source: "SAP HANA–High Availability" document[68]).

Host Auto-Failover (standby host)[69]

Host Auto-Failover is a local "N+m" fault recovery solution, similar to a traditional cluster solution. Typically, one (or more) standby SAP HANA host (or when running on vSphere, one VM) is configured to work in standby mode and is idle until the failover event. These VMs do not contain any data and do not accept requests or queries, but consume compute resources and must be maintained to ensure that, for example, the OS has the same level as the working VMs.

The standby host solution with virtualized SAP HANA systems works exactly the same way as when it is deployed in nonvirtualized environments and requires the implementation of the storage connector API and a STONIT script. Because of this the implementation of the standby host, the solution is limited to the storage vendors providing a vSphere version of their standby host implementation.

From the RTO and RPO perspectives, this solution is comparable to VMware HA, with the issues that above scripts and storage API needs to get implemented. With VMware HA, the affected VM will restart and is not required to implement, for example, a STONIT script. The difference between restarting a VM and failing over to a standby node may be the reboot time of a virtual server, which takes on average 40–60 seconds.

Note: Details on how to implement a Host Auto-Failover solution with vSphere are available from storage vendors supporting this feature with VMware vSphere, such as EMC.

VMware vSphere HA for SAP HANA

VMware vSphere High Availability delivers the availability required by most applications running in virtual machines–such as SAP HANA–independent of the operating system or applications running on them. VMware HA provides uniform, cost-effective failover **protection against hardware** and **operating system outages** within your virtualized IT environment. VMware HA can:

- Monitor VMware vSphere hosts and virtual machines to detect hardware and guest operating system failures, and restart the VM on the same or another host in the vSphere cluster.
- Restart virtual machines without depending on the applications running on the VM or on other vSphere hosts in the cluster without manual intervention when a server outage is detected.

[68] http://scn.sap.com/docs/DOC-60334

[69] http://scn.sap.com/docs/DOC-60334

- Reduce SAP HANA downtime by automatically restarting virtual machines on detection of an operating system failure.
- Provide 99.9 percent[70] service availability by protecting against hardware and OS failures.

VMware HA protects SAP HANA without any dependence on external components, such as DNS servers, or solutions, such as the SAP HANA Storage Connector API. In the event of a failover, the entire SAP HANA VM is restarted either on the same or another host in the vSphere cluster when the hardware of the initial host is defective. Since all virtualized disks, such as OS, data and log VMDKs, and in-guest mounted NFS volumes are failed over to a new ESXi host as well, no specific storage tasks, STONIT scripts, or extra cluster solutions are required.

Note: Protecting SAP HANA with VMware HA does not require a specific storage API or STONIT script.

By using VM restart polices and affinity and anti-affinity rules, it is possible to decide on which server the failed VM will be restarted. This allows a mix of SAP HANA and non-SAP HANA VM on a vSphere resource cluster and selective HA protection for independent SAP and non-SAP HANA VMs. The only requirement is that the failover host has enough resources left to start the failed VM, and that SAP HANA services are configured to start automatically after an OS reboot.

By leveraging the **SAP HANA Service Auto-Restart function**, HANA itself provides some protection against application failures that VMware HA normally would not be able to detect. SAP HANA service failures, for instance the Index or Name Server, will be restarted by the SAP HANA Service Auto-Restart watchdog function, which automatically detects the failure and restarts the stopped service process. Upon restart, the service loads data into memory and resumes its function. Figure 42 shows how the VMware vSphere HA and SAP HANA Service Auto-Restart feature work together and what components get protected.

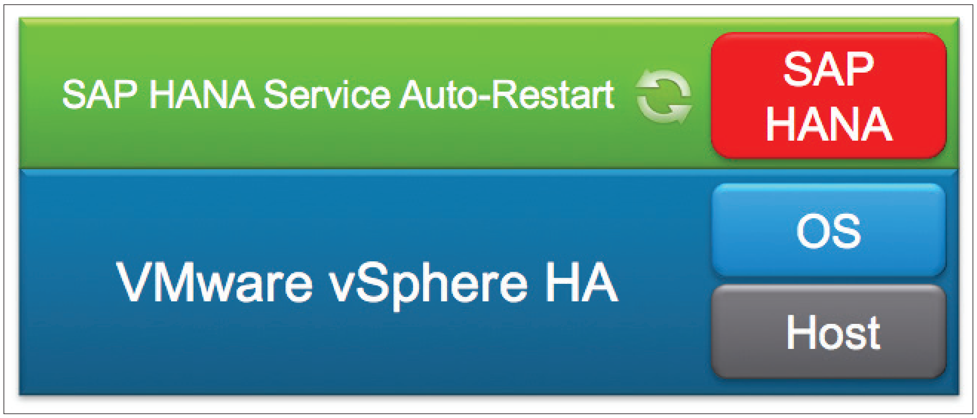

Figure 42: VMware vSphere HA Combined with SAP HANA Service Auto-Restart

[70] EMC IT, 02/14 EMC Perspective, H12853

In an SAP HANA production environment, resource commitments are required and must be enforced to ensure optimal performance.

In the case of a hardware failure, the SAP HANA VM will be restarted on a host with enough resources. In a virtualized environment, it is possible to run more independent systems on a vSphere host; in this case, it must enough system resources must be available to support the failover and restart process of the failed SAP HANA VM. If the host does not provide enough resources and if another VM is already running on the failover/standby ESXi host, then besides failing over the VM, the other VMs that run on the standby ESXi host must be shut down or moved to another host first to free up the needed resources. Another method is to configure the noncritical VM without resource commitments to purposely allow overcommitting of resources for this non-critical VM. In the event of a failover, the SAP HANA VM will request that all needed resources from the ESXi host and the non-critical VM will get only unclaimed resources remaining on the ESXi host.

Note: VMware HA, combined with the SAP HANA auto-restart feature, provides a solid HA solution to minimize failures due to power, HW, OS, or SAP HANA application faults. Combining these two solutions provides a high–availability solution for SAP HANA with over 99.9 percent[71] service availability.

Crash recovery steps are performed automatically when SAP HANA starts after an unexpected crash or failure. To minimize downtime, it is required to enable the automated SAP HANA Autostart feature. If Autostart is enabled, a VM–which got protected with VMware HA–will automatically start not only the OS, but also the SAP HANA database. Use the SAP HANA instance profile parameter “Autostart = 1” to enable Autostart.

Note: To minimize downtime, enable SAP HANA Autostart by setting the profile parameter “Autostart = 1”.

Figure 43 shows a vSphere resource cluster that has HA enabled for the SAP HANA VMs. The example configuration shows two SAP HANA Scale-Out systems running on an n+1 cluster, where the 1 represents the resources needed to protect against a complete physical host failure. In this 3+1 example, 25% of the compute resources are used for non-critical workload or are not utilized to speed up the failover and restart time needed in the case of a failover.

[71] EMC IT, 02/14 EMC Perspective, H12853

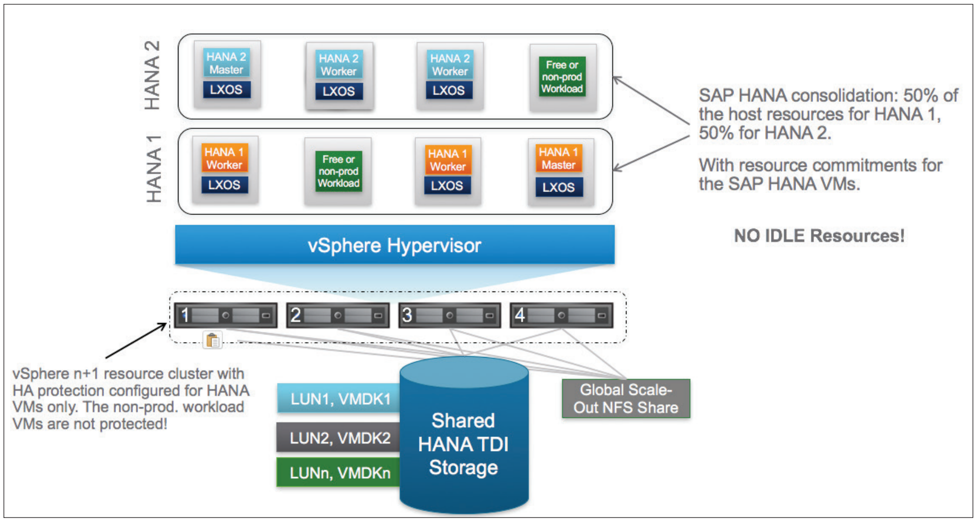

Figure 43: SAP HANA Protected Systems via vSphere HA

Figure 41 through Figure 43 show three different HA events.

Figure 44 shows failure 1, where the OS of HANA Master 1 is failing. Once the failure gets detected by the VMware HA, the VM will be restarted on the same host. On average, rebooting a virtual server takes between 40–60 seconds, which is very fast compared to the start time of a physical server, which could be from 5–30 minutes due to the preboot tests. After the reboot, SAP HANA needs to get loaded. The time needed for this depends on the SAP HANA database size.

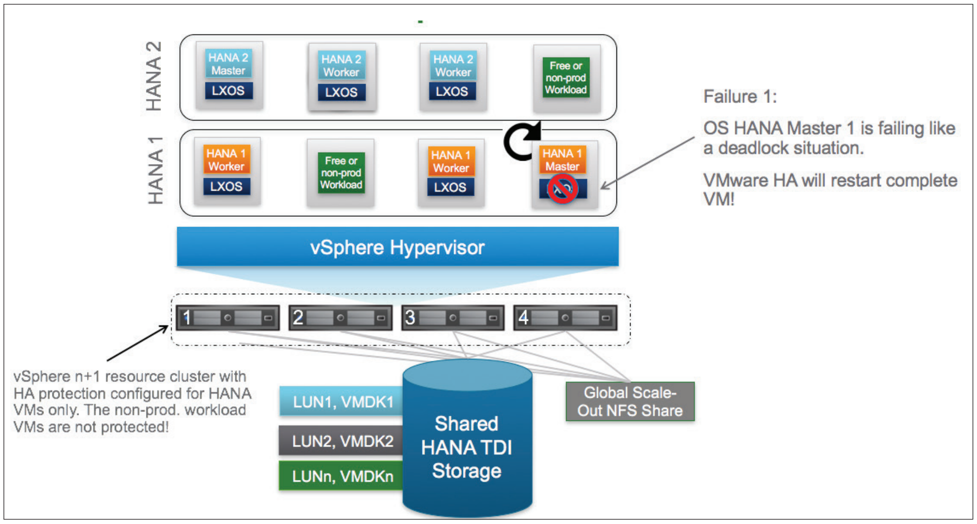

Figure 44: VMware HA–VM Restart After OS Failure

Figure 45 shows failure 2, where host number 4 is failing. Once the failure gets detected by the VMware HA, the VM will be restarted on a host with enough free resources or will be restarted on the preconfigured failover target, in this case host number 2.

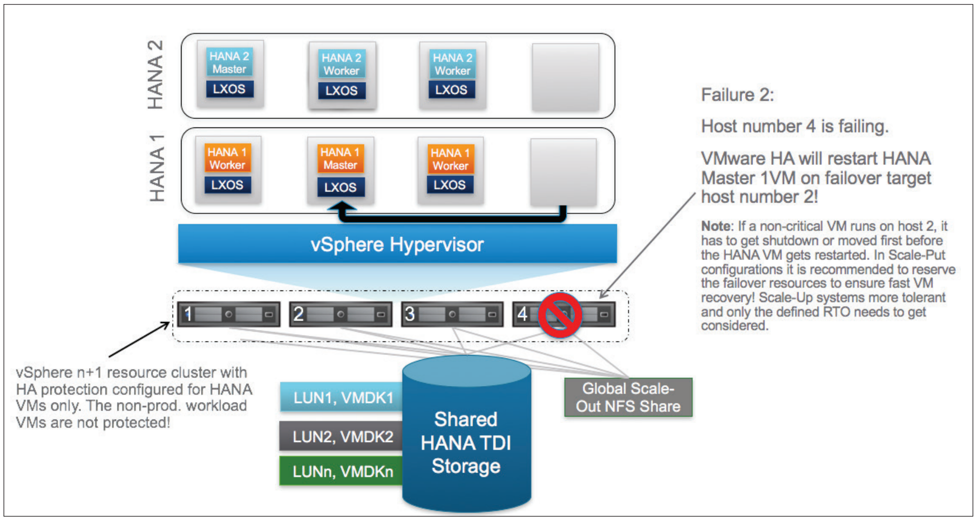

Figure 45: VMware HA - Protection Against Host Failures

Note: In Scale-Out configurations, we strongly recommend reserving the failover resource capacity in the cluster for a complete physical host failure. Using the failover capacity for other workloads would require shutting down these resources first to allow the restart of the failed VM. Because in a Scale-Out cluster all SAP HANA nodes work together, the delay in freeing up the failover resources may take too long, and the Scale-Out system may be affected negatively. Standalone Scale-Up systems are less critical because they are not dependent on other worker nodes.

Figure 46 shows failure 3, where VMware HA is not able to react, but the SAP HANA Service Auto-Restart feature helps to maintain system availability.

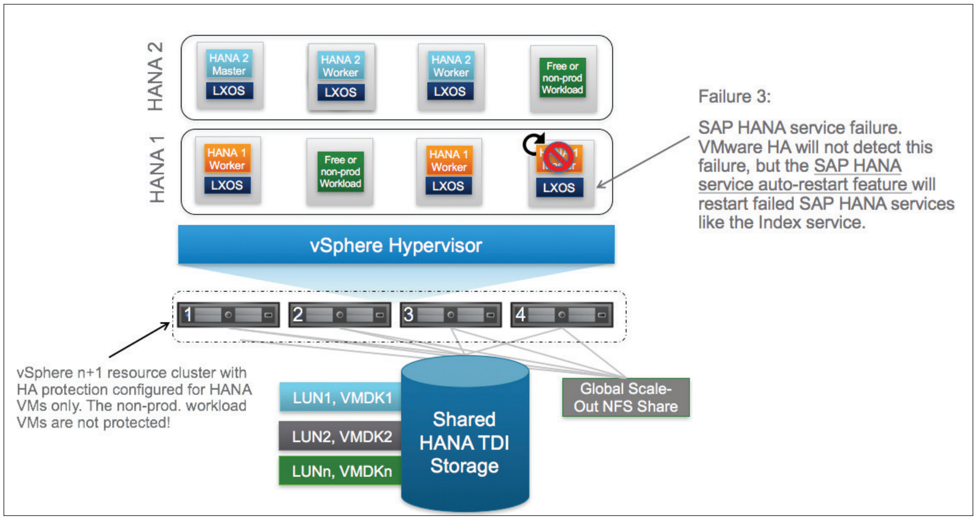

Figure 46: SAP HANA Service Auto-Restart Feature Combined with VMware HA

The SAP HANA Service Auto-Restart/watchdog function will automatically detect failures, whether they are software failures or intentionally stopped SAP HANA processes by an administrator, and restart these stopped HANA processes. All SAP defined HANA services (e.g., Index Server, Name Server, etc.) are protected by this SAP HANA built-in functionality.

The in-memory loaded SAP HANA data not affected by a process failure will not get reloaded. After the processes has restarted, only the affected SAP HANA service process data gets reloaded into memory, and SAP HANA resumes its function.

Note: VMware HA combined with the SAP HANA Auto-Restart feature provides a high-availability solution over 99.9 percent[72] and is not required to implement the SAP standby host HA solution.

Follow the configuration steps to determine how to configure the SAP HANA Auto-Restart feature described in the “Best Practices” section at the end of this document.

VMware HA can be the standard, unified HA solution for all business-critical applications in a datacenter, whether SAP or non-SAP applications need to be protected. Because of this, all HA IT processes and solutions can be standardized and optimized.

VMware vSphere HA with External Cluster or Monitoring Solution for SAP HANA

Additionally, SAP HANA application failures—such as process crashes or OS failures—may also get protected via third-party solutions, such as those offered by Veritas

[72] EMC IT, 02/14 EMC Perspective, H12853

Application HA for SAP HANA on vSphere[73] or the SUSE Linux Enterprise Server HA extensions[74] for SAP HANA.

Veritas ApplicationHA provides visibility to applications running inside VMs to ensure high availability of SAP HANA. ApplicationHA minimizes risks associated with application downtime by monitoring not only the VMs, but also the applications running inside them. When appropriate, it coordinates with the VM to restart through its tight integration with VMware HA.

A benefit of such a third-party solution is that the monitored application is separate from the solution that is monitoring it and that in addition to application failures, OS failures also can be detected by a single HA solution. ApplicationHA also leverages the vSphere HA API to initiate a VM restart or failover, whereas the SAP HANA watchdog is only monitoring and restarting the SAP HANA processes.

Applications can be monitored directly from vCenter Server or Operations Manager. ApplicationHA can monitor SAP HANA as well as hundreds of other applications and their services.

Please check out the third-party vendor product pages[75] and the previously referenced white paper for product availability and more information on this solution.

SAP HANA system replication can be used to recover from hardware or software faults, and may be used in addition to VMware HA to protect the following: for instance, the SAP HANA production VM of a system landscape when RTO is very low and aggressive. Remember, when HANA system replication is used, then an online Linux VM target is required and the actual switch to this HANA replication target needs to be done manually or by leveraging an SAP HANA ready cluster solution.

The details of system replication is discussed in the next section where the available disaster recovery methods get explained. Figure 47 shows how to combine VMware with SAP HANA system replication.

[73] https://www.veritas.com/content/dam/Veritas/docs/data-sheets/ds-virtualizing-sap-hana-confidently-with-veritas-applicationha-en.pdf

[74] https://www.veritas.com/product/business-continuity/applicationha.html

[75] https://www.veritas.com/product/business-continuity/applicationha.html

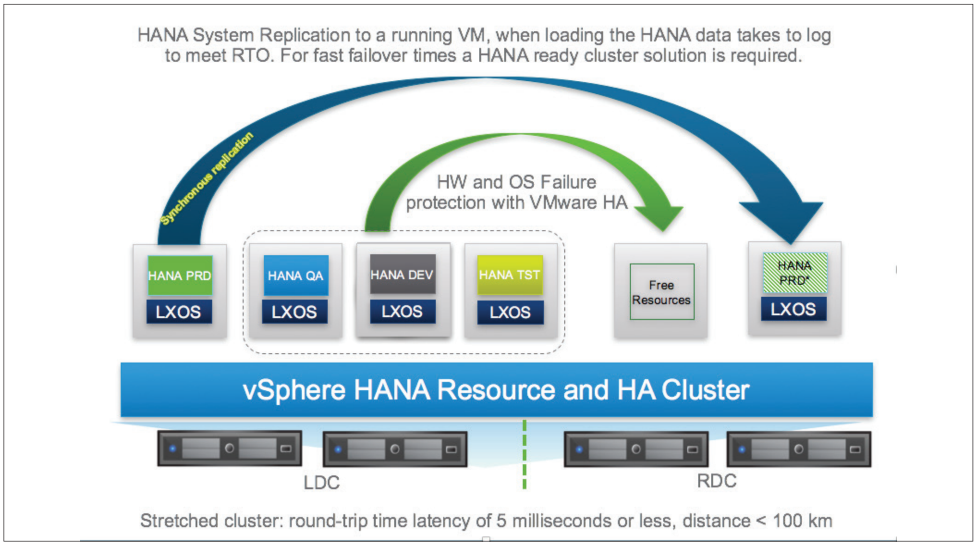

Figure 47: SAP HANA System Replication as Fast RTO Cluster Solution

Disaster Recovery (DR)

As noted earlier, SAP HANA Standby or VMware HA solutions are able to restart the SAP HANA services or the complete SAP HANA VM, but these solutions are not able to protect against file corruption, user errors, logical or site failures, or failures affecting the SAP HANA storage subsystem. To protect an SAP HANA system against such disastrous failures and situations, backup and system replication solutions should be implemented to protect the data.

Four disaster recovery following solutions are available to protect a VMware virtualized SAP HANA environment:

- Backup and Restore
- Storage Replication
- VMware Replication
- SAP HANA system replication

SAP HANA Data Persistence

Before we focus on Backup and Restore solutions, it is important to understand how SAP HANA maintains data persistence.

HANA is an in-memory database, storing all its data in volatile memory, which maintains data only for as long as the server is powered. To secure the memory content and all possible changed memory blocks, SAP HANA tracks all changes in the transaction log. When a transaction is committed, or when the transaction log becomes full, the log is written—using sequential writes in 1 MB blocks—to persistent storage on a storage device. Point (2) in Figure 48 shows log writes. During normal operation, SAP HANA also automatically saves in full 5-minute increments (default value) the complete memory image representing the date-to-disk by creating a savepoint (see point [1]). In the event of a failure, such as a power outage or hardware failure (point [3] in the figure), SAP HANA restarts and returns to its last consistent state by replaying the log since the last savepoint (crash recovery).

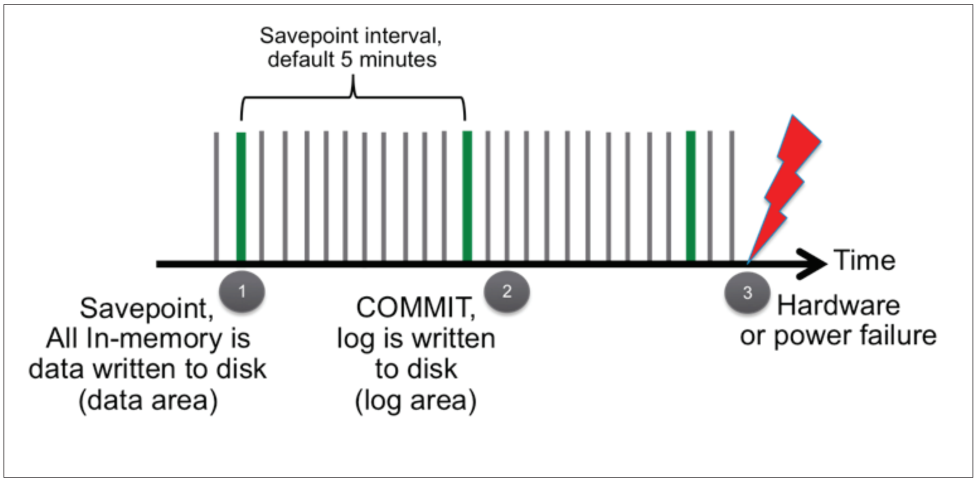

Figure 48: SAP HANA Data Persistence

SAP HANA on vSphere Backup and Restore

After refreshing 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. Which files need to get backed up to ensure that a restore process can finish successfully?

An SAP HANA database backup consists of data backup (snapshots) and transaction log backups. The data backup can be started scheduled or 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. What's the content and role of these files?

- Transaction redo **logs** are used to record any changes done 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 and represent 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. 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 will ensure 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 physical 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. Another option is to use storage built-in snapshot functionality to create backups. This method is the fastest way to create a backup. Some vendors work on VMware vSphere snapshot compatible backups. Storage systems that support today already the new VVOL standard allow already snapshotting VMs with the full awareness of the virtual disks belonging to a VM. Figure 49 provides an overview of the SAP HANA backup and recovery methods.

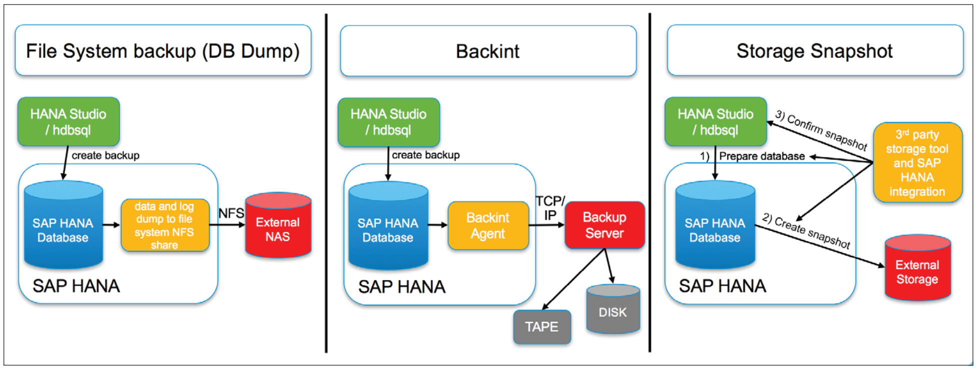

Figure 49: SAP HANA Backup and Recovery Methods

Recovery works the same way and is differentiated only by the way the backup files need to get handled. Recovery via backing, for instance, can be done in an automated way by selecting the right recovery set inside SAP HANA Studio. If a storage vendor has implemented its backup snapshot based backup into SAP HANA, then the recovery would be similarly easy.

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

SAP HANA on vSphere Storage Replication

Beside taking backups of SAP HANA, data replication provides a suitable method to protect against disasters that negatively impact local data, storage subsystems, or complete sites. With data replication, the SAP HANA primary data on a logical or on the storage level is replicated to a secondary site using either software or storage system based replication technologies.

The following methods for data replication are available to customers:

- Storage based replication
- VMware based replication
- SAP HANA system replication

The solutions differentiate primary in how data get replicated and the RTO and RPO times that are possible with these solutions. For instance, VMware replication would not be the first choice when RPO should be 0 because VMware replication allows only asynchronous replication.

Note: Just as with physically deployed SAP HANA systems, enabling synchronous replication that ensures an RPO of 0 will negatively impact the responsiveness of the SAP HANA database on user transactions that involve writing data to disk.

Storage-based Replication

The first replication solution we're introducing is storage system-based replication solutions, sometimes called array-based storage replication, which replicates the data of primary storage to a secondary storage system, typically located in another datacenter. All major storage system vendors support their replication solutions in combination with VMware vSphere. Most of the vendors also support VMware Site Recovery Manager™ (SRM). SRM will be covered later.

Depending on the RTO and RPO requirements and the distance between the primary and secondary sites, implementation of synchronous storage replication (RPO=0) or asynchronous storage replication (RPO≥0) is possible. If a disaster occurs, the RTO is typically the time that is required to start up the SAP HANA database at the secondary site.

Unlike with SAP HANA system replication (SAP HANA software-based replication), storage-based replication can be used as a unified data replication solution, which can include OS boot volumes, the SAP HANA shared file system, and data from non SAP HANA applications, in addition to the SAP HANA database persistence. This ensures that after a disaster the all data of all replicated systems are available at a remote site, allowing a fast recovery from primary site disasters at the secondary site.

Note: Storage-based replication enables a unified and federated DR strategy for all business-critical applications.

Vendor-specific documentation, such as from EMC, NetApp, HDS, or HPE are available to plan, implement, and configure a storage block level-based replication method that is unified between all business-critical applications running in VMs on VMware vSphere.

- EMC: https://www.emc.com/collateral/hardware/white-papers/h7118-using-vmware-virtualization-platforms-vplex.pdf
- HDS: https://www.hds.com/en-us/pdf/best-practices/deploy-site-recovery-manager-with-vmware-vsphere-on-hitachi-vsp-family-systems.pdf
- HPE: http://www.vmware.com/files/pdf/partners/hp/vmware-vsphere-hp-3par-storeserv-so.pdf
- IBM: http://www.redbooks.ibm.com/redbooks/pdfs/sg248158.pdf
- NetApp: http://www.netapp.com/us/media/tr-3606.pdf

VMware vSphere Replication[76]

VMware vSphere Replication is a hypervisor-based, asynchronous replication solution for vSphere virtual machines. It is fully integrated with VMware vCenter Server and the vSphere® Web Client, and is fully transparent for the applications running inside VMs. vSphere Replication operates at the individual virtual machine disk (VMDK) level, allowing replication of individual virtual machines between heterogeneous storage types supported by vSphere. Because vSphere Replication is independent of the underlying storage, it works with a variety of storage types including VMware Virtual SAN[77]; vSphere Virtual Volumes; traditional SAN and NAS; and direct-attached storage (DAS). This allows you to:

- Use differing storage technologies at opposite sites (e.g., SAN to Virtual SAN, FC to iSCSI).
- Repurpose older storage at the recovery site to reduce cost.
- Use secondary storage only for protected virtual machines, not the entire environment.
- Recovery point objective (RPO) of minimal 15 minutes.

vSphere Replication copies only changed data to the recovery site to lower bandwidth utilization, improve network efficiency, and enable more aggressive recovery point objectives (RPOs) than a manual, full-system copy. With vSphere Replication, you can:

- Utilize a "seed copy" of virtual machine data for the initial synchronization to reduce the amount of time required to create the initial copy.
- Ensure efficient network utilization by tracking changed disk areas and replicating only those changes.
- Optionally enable data compression to further reduce network bandwidth consumption.

The virtual machine replication process is nonintrusive and takes place irrespective of the type of operating system or applications in the virtual machine. It is transparent to protected virtual machines and requires no changes to their configuration or ongoing management.

Once SRM got installed in a vSphere environment, it is easy to configure replications on per VM basis. Figure 50 shows how to configure a replication per VM.

[76] Source: https://www.vmware.com/products/vsphere/features/replication

[77] As of writing of this guide, VMware virtual SAN is not supported by SAP as SAP HANA TDI storage solution

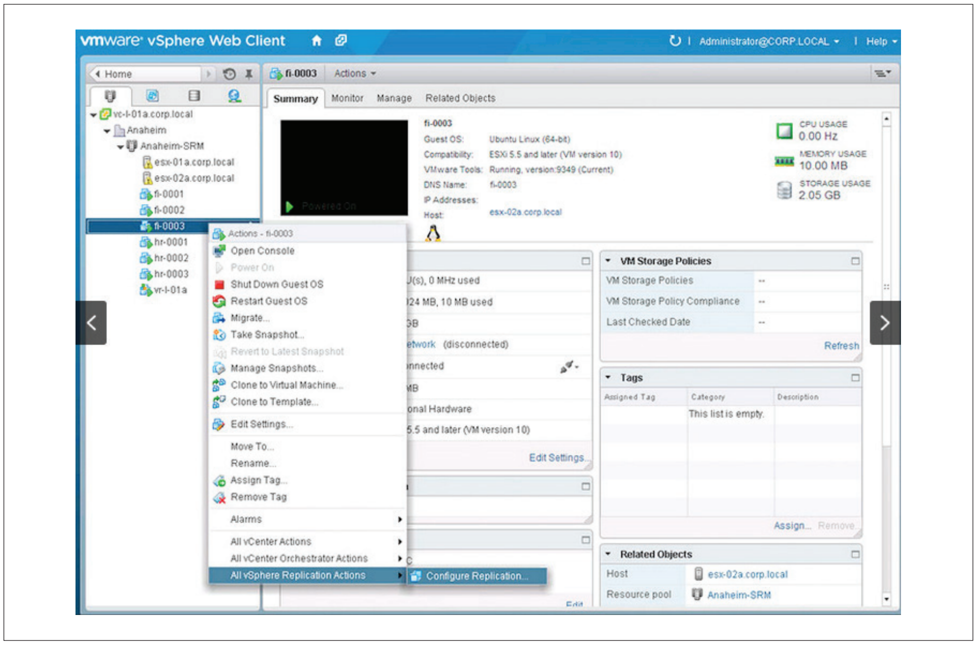

Figure 50: SRM–Configure Replication per VM

Figure 51 shows the recovery settings screen and how it is possible to configure the recovery settings on a per-VM level.

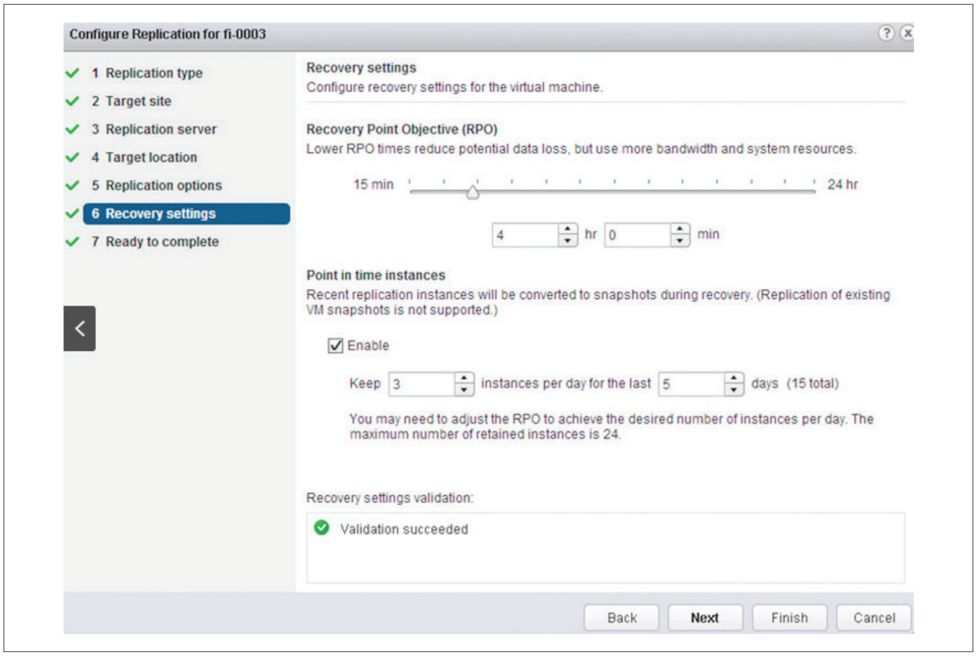

Figure 51: Configure RPO per VM

Important: VMware Replication provides a Recovery Point Objective (RPO) with a maximum 15 minutes. Using VMware Replication with production-level SAP HANA systems, which require less than RPO time, are not supported. For RPO times less than 15 minutes, use storage or SAP HANA system-based replication.

---

vSphere Replication in the SAP HANA context should been seen as an extension of available data replication such as SAP HANA system replication because it can be used to replicate the SAP HANA OS and shared disk/volumes and all SAP systems around an SAP HANA system such as SAP application servers, solution manager, SAP gateway, also. This allows, in combination with RPO 0 replication solutions, the setup of disaster-tolerant SAP environments, which include all systems and not only the SAP HANA data.

Note: vSphere Replication delivers flexible, reliable, and cost-efficient replication to enable data protection and disaster recovery for RPO≥15 virtual machines in your environment. Business-critical RPO=0 SAP HANA systems must be replicated using SAP HANA system or storage-based system replication solutions.

SAP HANA system replication

Unlike other solutions, SAP HANA system replication is part of the SAP HANA solution. It provides a robust solution to replicate the SAP HANA database content to a secondary disaster site, and it is even possible to use it to pro-load data on a second running VM that receives the data of the primary system. This allows fast recovery times and is, when only SAP HANA needs to get protected a RTM of minutes instead of sometimes hours compared to other solutions.

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 different modes for system replication:

- Synchronous
- Synchronous in-memory
- Asynchronous

Depending on requirements, the disaster recovery VMs can consume greater or fewer resources on the disaster recovery VMware vSphere cluster. For example, selecting the synchronous in-memory mode will consume the same amount of RAM as 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[78] that by only replicating the data, around 10 percent of the system resources are required, allowing up to 90 percent of the resources to continue being used by other systems, such as test or QA systems.

Figure 52 shows the scenario of a co-deployed configuration where all SAP HANA Scale-Out VMs or Scale-Up VMs (not shown) are replicated to a disaster recovery site. If SAP HANA system replication is used, the same number of VMs have to be installed

[78] http://scn.sap.com/docs/DOC-47702

and online (as shown in the figure) to receive the replicated data. If storage based replication is used, then only the storage is replicated and no online VMs that receive the data are needed. This frees up server resources that could then be used for QA or other systems.

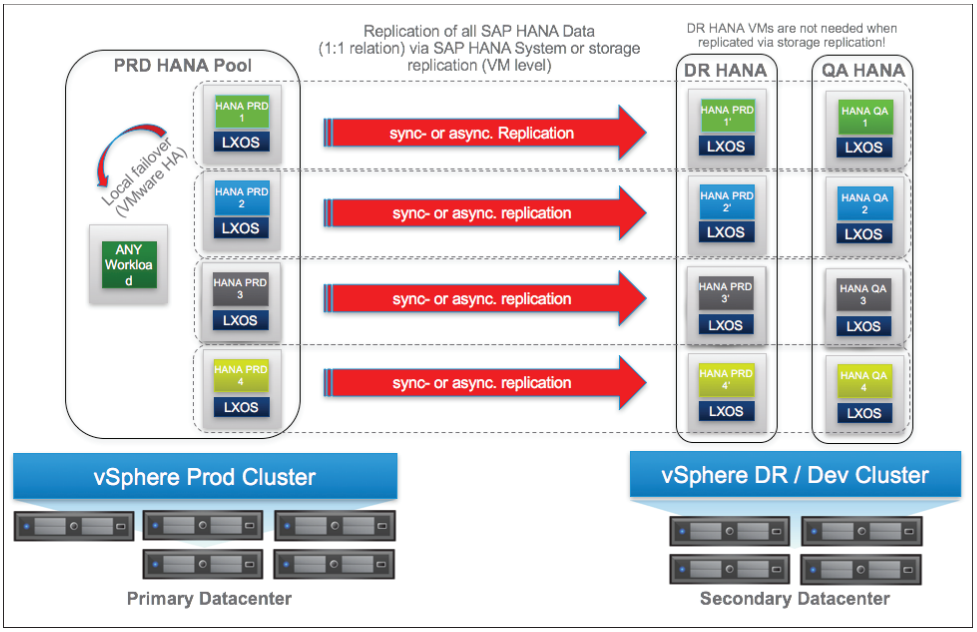

Figure 52: SAP HANA on vSphere Replication

In the scenario shown in the figure, resource overcommitments are used to allow the co-deployment of such an environment. By using resource pools and resource shares, it is possible to provide the needed resources to the disaster recovery SAP HANA VMs. The co-deployed system, with fewer resource shares, will experience performance degradation after the disaster recovery systems are used after 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 rather than running both systems in parallel—with resource limitations—on the same platform.

With the option Synchronous in-memory, which is the default setting of SAP HANA system replication, an HA solution also can be created that delivers very fast RTO figures because after the switch to the new systems, the data doesn't need to get loaded into RAM; they are already pre-loaded. Therefore, system ramp-up times can be minimized and may be the only solution available when RTO times below 15-30 minutes are required (depending on the size of the HANA in-memory database).

The issue with system replication is that an external tool or an administrator has to decide when to switch the system. If the decision was made, then additional steps during and after a site failover has taken place for instance switch of network identity (IP redirect) of the replicated systems from the disaster recovery configuration to the production configuration and it has also to be ensured that old primary system is deactivated and got shutdown. Otherwise "split brain" situations may occur.

This can be done manually or via automated tools such as SUSE HA or other SAP HANA supported cluster managers. The configuration of such a solution in a virtualized environment is similar to the configuration of natively running systems.

When replicating the SAP HANA as well as all other systems to the secondary site with storage or, when appropriate, VMware replication, then you can use VMware Site Recovery Manager™ (SRM) to perform the necessary failover steps automatically. All that is required is the initial switchover command.

VMware vSphere Site Recovery Manager (SRM)

VMware site recovery manager can help to reduce the complexity of an SAP HANA disaster recovery (DR) solution by automating the complex disaster recovery steps on any level.

Site Recovery Manager is designed for disaster recovery of a complete site or data center failure. Site Recovery Manager supports both uni- and bi-directional failover. In addition, it also supports "shared recovery site," allowing organizations to failover multiple protected sites into a single, shared recovery site.

The key elements that make up a Site Recovery Manager deployment for SAP are:

- 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 Servers are independent, with each managing its own site. Site Recovery Manager makes them aware of the virtual machines they will need to recover if a disaster occurs.
- The Site Recovery Manager service takes care of managing, updating, and executing disaster recovery plans. Site Recovery Manager is managed via a vCenter Server plug-in.
- Site Recovery Manager relies on vSphere replication, storage vendors' array-based replication Fibre Channel, or NFS storage, which supports replication at the block level to replicate SAP HANA data and log files to the DR site. Site Recovery Manager communicates with the replication via storage replication adapters that the storage vendor offers and that have been certified for Site Recovery Manager.
- vSphere Replication has no restrictions on the use of storage types or adapters and can be used for non-performance-critical or static VMs, such as infrastructure services or SAP application servers with RPO time objectives of 15 minutes or longer.
- Site Recovery Manager cannot be used to switch an SAP system replication replicated system because SRM manages the failover and switch process on a VM and not on an application level. This must get performed manually or via a cluster solution such as SUSE Linux Enterprise High Availability Extension for SAP HANA.

Figure 53 shows an example of an SRM- and replication-protected SAP landscape. The VMs running on the primary site contain all the needed infrastructure and SAP components such as LDAP, SAP HANA database, and SAP application servers. The VMs can be replicated, depending on the RPO needs via vSphere or storage-based replication. Again, vSphere replication can be used with VMs that tolerate an RPO time of 15 minutes or longer.

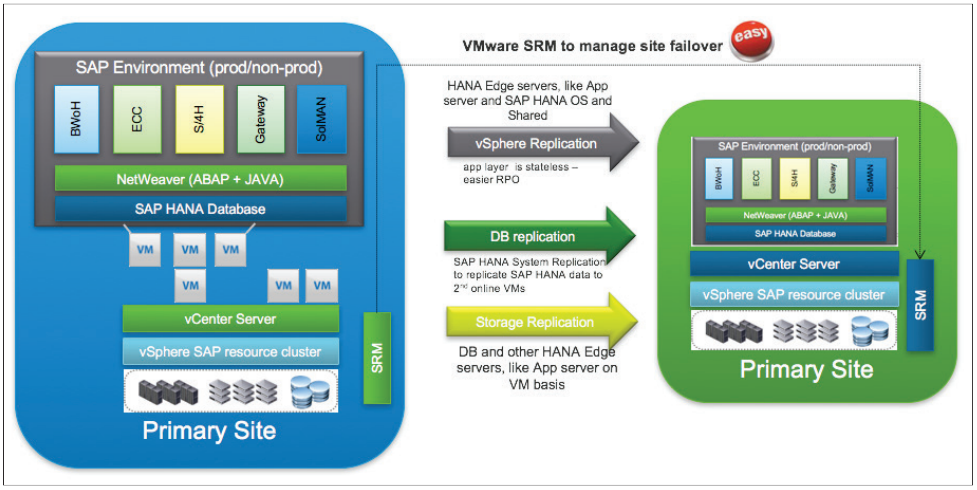

Figure 53: SAP HANA Site Failover Managed by VMware SRM

With VMware vCenter Site Recovery Manager combined with SAP HANA system or storage replication and an SAP HANA ready cluster manager (when HANA system replication is used), it is possible to build a very solid, automated, disaster-tolerant SAP landscape.

Note: Currently, the automated switch of an SAP HANA system replicated system is not possible because SRM works on a VM level, not an application level. An additional SAP HANA ready cluster solution such as SUSE Linux Enterprise High Availability Extension for SAP HANA has to be used instead.

Virtual SAP HANA Sizing

Virtual SAP HANA Sizing Considerations

Before starting SAP HANA sizing, you must understand the VMware vSphere VM resource limitations and how to translate the sized resource needs of an SAP HANA database system and possible required SAP Business Suite application servers that may get installed beside the HANA database inside the same VM.

Table 13 provides the average virtualized SAPS capacity of the named CPU and the CPU socket-to-RAM ratios important for Business Warehouse powered by SAP HANA configurations because no SAPS sizing exists today. To translate the CPU needs of a BWoH workload, which currently is defined by the RAM requirement, only the CPU socket-to-RAM ratios can be used to determine how many CPU resources are required as defined by SAP to run a specific memory configuration.

The SAPS figures in the table are based on published 2-tier SD benchmarks of SAP enhancement package 5 for SAP ERP 6.0 on Sybase DB and Linux OS configurations. The socket-to-RAM ratios are based on the SAP HANA certified appliance configurations.

CPU	AVERAGE SAPS PER CPU/ SOCKET	AVERAGE SAPS PER CPU CORE	HANA WORKLOAD	HANA SPS VERSION	CPU SOCKET-TO-RAM RATIO
Intel Xeon EX E7-v4	56575	2357	SoH	SPS <= 11 SPS >= 12	<=768 GB/ socket <=1024 GB/ socket
			BWoH	SPS >= 11	<=512 GB/ socket
Intel Xeon EX E7-v3	41312	2295	SoH	SPS >= 11	<=768 GB/ socket
			BWoH	SPS <= 10	<=384 GB/ socket
				SPS >= 11	<=512 GB/ socket
Intel Xeon EX E7-v2	33894	2260	SoH	SPS >= 10	<=768 GB/ socket
			BWoH	SPS >= 10	<=256 GB/ socket
Intel Xeon EP E5 26xx v4 (maximum 2- socket supported)	57025	2592	SoH	SPS <= 11 SPS >= 12	<=768 GB/ socket <=1024 GB/ socket
			BWoH	SPS >= 10	<=384 GB/ socket

Table 13: SAP HANA Platform Sizing Figures for Selected CPU Types

Because in a physical appliance server all installed CPU resources (cores and hyperthreads) will be available for SAP HANA and the installed RAM, it is important to understand that a virtual machine may not be able to utilize all resources due to the 64 vCPU limitation of vSphere 5.5 and the 128 vCPU limitation of vSphere 6.0.

By default, vSphere will always prefer a CPU core before it will use a CPU hypterthread and, if the number of CPU cores exceeds the number of supported vCPUs (64 or 128), then hypterthreads cannot be used. As described in a previous section, leveraging hyperthreads will provide some additional CPU calculation capacity. For calculation reasons, we defined a hypterthread gain of 15 percent. All of these additional 15 percent CPU cycles will not be available when a single 128 VM runs, for example, on a 4-socket Broadwell server with a total of 96 CPU cores and 192 CPU threads. In this case, 64 CPU threads cannot be used. This effect is not relevant if in this configuration two SAP HANA 96 vCPU VMs run on the server. In this case, 100 percent of the available CPU can be used.

Note: In typical SAP HANA VM consolidation (multi-VM) configurations, this effect is less critical because the sum of all available vCPUs may be able to consume all available CPU cores and threads, and is only critical when one large VM runs on a server.

The formulas below use the certified appliance system CPU socket-to-RAM ratios to calculate how many CPU resources are needed to support different RAM configurations. Again, this is required especially when a BWoH SAP HANA database server needs to be configured as a VM because the currently available SAP sizing tools do not provide information on how many CPU resources are needed, and provide information only on the memory configuration. To configure a virtualized SAP HANA system–in addition to the vRAM size–the number of vCPUs are configured. The provided formulas help to determine how many CPU resources are required for a specific server platform and HANA SPS by using the SAP defined socket-to-RAM ratios.

For the **SAP Business Suite application**, the SAPS figure of an existing database that will get migrated to SAP HANA can be **multiplied by 3** to estimate the required CPU resources instead of using the formulas. If the SAPS performance is known, such as 50,000 SAPS for an old Oracle SAP ERP database, then for SAP HANA it is possible to calculate with 150,000 SAPS. See **SAP Note** 1793345[79] –**Sizing for SAP Suite on HANA** for details.

The formulas can also be used to determine the maximum supported RAM per single wide VM and will ensure that there is no mismatch in the SAP defined CPU resources-to-RAM ratios.

Virtualized SAP HANA BWoH Sizing Based on vCPUs and vRAM

Maximum VM RAM formula for BWoH:

$$\text{maximum VM RAM} = \frac{max.\,Appliance\ RAM\ x\ vRAM\ overhead}{max.\,Appliance\ CPU\ cores\ x\ 1.15}\, x\ vCPUs\ (\text{max. 128 vCPUs})$$

vCPU formula for BWoH:

$$\text{needed vCPUs} > 10 = \frac{HANA\ VM\ RAM}{max.\,Appliance\ RAM\ x\ vRAM\ overhead}\, x\ \boldsymbol{max.\,usable\ CPU\ cores\ system}\ x\ 1.15$$

Legend:

- Maximum VM RAM = RAM maximum supported by a certain amount of vCPU
- Maximum appliance RAM = the RAM supported by either BWoH or SoH and Scale-Up or Scale-Out appliances
- vRAM overhead = RAM cost of virtualization, this depends on the actual VM configuration and is between 0.5% and 3%. On average, we have seen an overhead

[79] https://launchpad.support.sap.com/#/notes/1793345

below 3%; therefore, a factor between 0.995 (-0.5%) and 0.97 (-3%) can be used, but must get verified on the actual system.

- Maximum appliance CPU cores = amount of total available CPU cores of an appliance system
- Needed vCPUs = amount of needed vCPUs (cores + hyperthreads) to support the needed HANA VM RAM
- Maximum usable CPU cores = amount of available CPU cores of the NUMA nodes to be used (e.g., three NUMA nodes for the VM of a 4-socket server)
- HANA VM RAM = RAM size of the SAP HANA overall system (includes all needed RAM, not only HANA DB tables)

SAP HANA is a multi-threaded application that requires a multi-CPU VM configuration. Testing has shown that the smallest non-production HANA instance should be configured with at least 10 vCPUs. When a Haswell 18-core, 36-thread CPU is used, it could support up to three 12 vCPU VMs **non-production** VMs on one CPU socket, with a 24-core Broadwell CPU up to 4 non-production VMs on one CPU socket would be possible.

Note: Use at least 10 vCPUs for an SAP HANA VM. Below 10 vCPUs, the SAP HANA system may become slow and unresponsive.

BWoH NUMA Socket Sizing:

As an alternative to the formulas and an easy rule of thumb, it is possible to estimate the required CPU resources by dividing the sized RAM by the socket-to-RAM ratios. This will indicate how many CPU cores should be assigned to an SAP HANA VM without leveraging the available hyperthreads.

Example:

Sized BW RAM is 2,5 TiB, the server is a Broadwell server with a socket-to-RAM ratio of 512 GiB per CPU socket, the needed CPU socket for the RAM and CPU is 5 -> 2,5 TiB/512 GiB ~5. Five (5) CPUs have to be used for this VM and, because a CPU socket has 24 CPU cores, the SAP HANA BWoH VM will have 5 x 24 = 120 vCPUs configured. To take advantage of some hyperthreads, we recommend configuring the VM with the maximum number of vCPUs available. In this case, it would be 128 vCPUs.

Note: With the specified formulas and the NUMA socket sizing guideline, it is now possible to calculate the CPU needs of an SAP HANA instance when no SAPS figures for the database part are available. The actual CPU resource needs of the SAP HANA database may be lower as calculated. The formulas are used only to ensure that the SAP defined CPU socket-to-RAM ratios are maintained and that the first VM configuration is not too small. Once SAP publishes CPU sizing rules for BWoH systems, these ratios and formulas are not required any longer.

The RAM calculated with the provided formula is a starting point for the VM configuration and should be aligned to the real resource utilization, which should be measured once the system is online and used for a while.

Table 14 summarizes the currently maximum utilizable memory sizes and VM configurations based on the published SAP appliance system configurations. These maximums can be lifted if a server has very low CPU utilization documented in several SAP early watch service reports.

CPU	HANA WORKLOAD	HANA SPS VERSION	VSPHERE VERSION	MAXIMUM "THEORETICAL SIZING" RAM SIZE*	THEORETICAL MAXIMUM VM SIZE AND CONFIGURATION**
Intel Xeon E7-v4	SoH	SPS >= 12	6.0	4748.99 GB	128 vCPUs, 4080 GB vRAM
	BWoH			3165,99 GB	128 vCPUs, 3166 GB vRAM
Intel Xeon E7-v3	SoH	SPS <= 10	5.5	2374.49 GB	64 vCPUs, 1024 GB vRAM
		SPS >= 11	6.0	4748.99 GB	128 vCPUs, 4080 GB vRAM
	BWoH	SPS <= 10	5.5	1582.99 GB	64 vCPUs, 1024 GB vRAM
		SPS >= 11	6.0	3165.99 GB	128 vCPUs, 3166 GB vRAM
Intel Xeon E7-v2	SoH	SPS <= 10	5.5	2374.49 GB	64 vCPUs, 1024 GB vRAM
		SPS >= 11	6.0	4748.99 GB	120 vCPUs, 4080 GB vRAM
	BWoH	SPS <= 10	5.5	1582.99 GB	64 vCPUs, 1024 GB vRAM
		SPS >= 11	6.0	3165.99 GB	120 vCPUs, 2048 GB vRAM

*Memory size that a maximum vCPU sized VM would be able to utilize when applying the SAP HANA core- to-RAM ratio. The vRAM configuration may be changed after initial deployment and real CPU utilization based on SAP early watch reports. If very low CPU utilization is documented, larger RAM configurations may be possible (within the vSphere limitations).

**The maximum VM CPU and RAM sizes depend beside the vSphere limits also and the current certification status on the underlying HW, Ivy Bridge systems support less hardware resources as Haswell or Broadwell systems!

Table 14: Maximum SAP HANA VM Configuration

Note: Sizing experts expect that the maximum RAM configured for an SAP HANA VM can be increased to the maximum RAM attached to all utilized NUMA nodes minus the virtualization overhead between 0.5%–3%. Regular SAP HANA early watch reports will provide the needed information on how intensive the CPUs of the system are utilized. If there is a low CPU utilization, more RAM can be supported.

SAP HANA Sizing Procedure

Sizing the SAP HANA database starts with the determination of the database size. This can be done by either using the SAP HANA Quick Sizer[80] application (new system) or by running a specific SAP HANA sizing report (existing systems), as

80 https://websmp109.sap-ag.de/Quick Sizer

documented in SAP Note 1872170[81] (S4/HANA) and SAP Note 2296290[82] (BW). Besides determining the database size of HANA that will be operated in RAM, it is also necessary to size the needed resources for the application server stack. This can be done by using the SAP HANA Quick Sizer because this tool will provide the necessary system configuration required for the application part of an SAP HANA based business application.

Figure 54 shows the general sizing procedure as described in the SAP HANA sizing guidelines, and should be followed when sizing virtualized or non-virtualized SAP HANA systems.

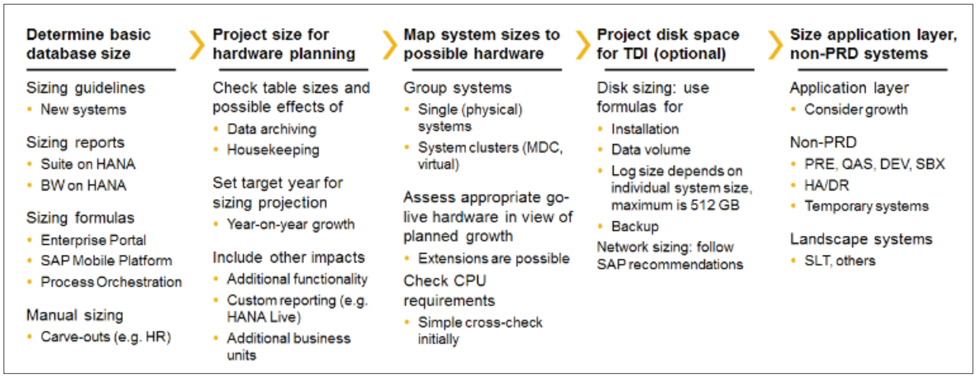

Figure 54: General SAP HANA Sizing Procedure with Building Blocks, Sizing Guide

SAP Note 1793345–Sizing for SAP Suite on HANA, provides important information on how to size a Suite on SAP HANA system and provides some tips on first rough estimations on needed memory, disk space, and CPU resource needs. For the final sizing, use the Quick Sizer or the ABAP reports. The following guidelines are estimations, but they provide a good first understanding on how an SAP HANA VM configuration will look such as and if the estimated sizes are too large for the current VMware limitations, source SAP Note 1793345:

Memory estimation

- For a first estimation for your required Suite on HANA main memory, check the size of your current well-maintained and uncompressed database. The recommendation from SAP is that you take half the size of your disk-based database, include a safety buffer of 20%, and add 50 GB fixed size for code, stack, and other services. This means that if the database is approximately 2000 GB in size (tables plus indexes), the maximum memory consumption will be 1250 GB (2000 GB/2 x 1.2 + 50 GB).
- Don’t forget to consider future data growth; add this to the estimated memory. SAP doesn't expect changes in memory requirements for applications using liveCache. If you are using SAP HANA Enterprise Search, you should add 20% memory resources on top.

[81] https://launchpad.support.sap.com/#/notes/0001872170

[82] https://launchpad.support.sap.com/#/notes/0002296290

CPU estimation

- To estimate the Suite on HANA CPU requirements, **check** the current **CPU consumption of the existing database.** To fully support the parallel processing capabilities of HANA for optimal response times for analytical applications, SAP prefers a factor of **3 more CPU power for HANA than for disk-based databases** without the parallelization of single statements. This also considers the load for running OLTP and Reporting simultaneously, and includes a moderate use of SAP HANA Enterprise Search. Extensive usage of SAP HANA Enterprise Search capabilities might require additional CPU resources.

Storage space estimation

- To estimate the Suite on HANA disk space (or total net disk space), you have to take into consideration the net data size on disk (anyDB) and the disk space required for delta merges. During a delta merge, the affected tables are temporarily duplicated on disk for a short period of time. The disk space for merges is calculated by taking the sum of the two biggest tables and including this value into the calculation. On top, you should add 25 GB for the space required by the statistics server, HANA metadata, etc. The final calculation looks as follows:
  - HANA disk space (Total Net Disk Space) = (Net Data Size anyDB + Disk Space for Merges)/4 (compression) x 1.2 (20% safety buffer) + 25GB
  - If you are using SAP HANA Enterprise Search, you should add 20% disk resources on top
  - Please be aware that space for backup and

Unfortunately, there is no similar information available for BW on HANA to perform some quick estimates. For BW systems on HANA, use the Quick Sizer tool or the ABAP sizing reports for BW on HANA.

The most important sizing factor for SAP HANA is the needed RAM. A typical SAP HANA system needs approximately 2 times the RAM of the sized database. The Quick Sizer tool and the ABAP sizing reports show the total RAM needed for an SAP HANA system, and you can use this size for the server selection process.

SAP has introduced with its hardware partner's appliance configurations that help to select a server configuration that allows you to deploy these systems easy and quickly. Defining a virtualized SAP HANA VM size works the same way as selecting a physical server system.

As stated previously, when you are sizing a virtualized SAP HANA instance, we recommend following the same sizing guidelines that were published for physical HANA systems. The only difference is the maximum VM sizes of VMware vSphere 5.5 (1 TiB RAM) and vSphere 6.0 (4 TiB RAM). For completeness, memory virtualization costs of about 0.5%–3% need to be considered. When using the Quick Sizer tool to determine the application server stack and the required SAPS performance, between 5%–0% of CPU resources should be added consider virtualization costs, when defining the SAP HANA VM configuration.

Note: Generally, you would size new systems using the SAP HANA Quick Sizer and not by using the CPU core/socket-to-RAM ratio that is used with the SAP HANA appliances.

Using SAP Quick Sizer HANA version (new systems)

Using the HANA version of the SAP Quick Sizer allows you to determine the memory, CPU, and storage requirements of an SAP HANA system or application that is deployed, and should be the leading tool for all SAP HANA native applications or for greenfield sizing.

When you perform a Quick Sizer sizing, then you can select between two different sizing methods: a user-based sizing and a throughput-based sizing. The throughput-based sizing is the preferred method because it will provide a more precise result.

Please follow the Quick Sizer documentation for details how to use this tool and how to achieve usable results. A good start is the SAP HANA Academy - Sizing SAP HANA: Quick Sizer overview video.

As an example, the SAP published Quick Sizer project '**FI-CA HANA**', customer **188213** was updated to version 244 and modified to require less RAM as the initial example. The modification was done in Table 3, throughput disk sizing BIT Upload, and the length per object got reduced from 1024 to 512.

Scenario description source SAP Quick Sizer project, SAP Simple Finance on HANA application:

"This example sizing is for a cable company that provides landline phone and cable TV services to their customers. They sell however packages including mobile phones which are operated by a partner.

From their TV-decoder, customers may subscribe to TV series or events for which an additional fee is required.

Shows to be broadcasted in the future will lead to a deferred revenue because the money can be posted only once the show has started. A small online shop permits customers to buy gadgets. Landline phone calls come in as BITs and CITs, TV subscriptions only as CITs because out of 10 subscriptions the cheapest one will be given free of charge, so charging is not possible at the time of purchase (as each item may be the cheapest one in a row of 10). The mobile operators send billing documents per business partner to the company and from the online shop the system receives invoices.

For the sizing we assume: 1,000,000 business partners doing in average 20 calls per day and purchasing 2 TV shows per day of which one is broadcasted immediately and one the upcoming weekend or later which leads to deferred revenues. This gives an average of 682 BITs per month per business partner (31*22). Phone calls are already rated when they arrive in the system but TV is not as the company offers a discount on all TV shows if you buy more than 10 per month. While BITs and CITs come in 365 days a year, billing is done only at 20 days per months. This gives 50,000 invoices to be produced daily. Billing documents from the mobile operators are loaded into the system and taken into account on the invoice.

85% of the customers opt for direct debit, 15% pay once the bill arrives.

0.1% of the bills are wrong and need to be redone and

10% of the customers require a reminder before they meet their bill.

Sales from the online shop must be paid in advance by credit card; there are 1000 sales per day. They are loaded into the system for revenue recognition purposes. All documents are to be kept 24 months except BITs and CITs which are kept for 6 months only due to size and legal requirements.

In this example, frequently documents from different sources are sized separately. Sometimes this is for educational reasons only, so feel free to add up values before you enter them in the Quick Sizer."

Figure 55 shows the sizing result of the SAP Simple Finance on HANA application example as described previously. The SAP HANA system RAM size is classified as “L” with a total RAM size of 3,531,775 MB of RAM, which is approximately 3.4 TiB of RAM. The disk space for the database files alone were sized with 4,020 GB, which is around 4 TiB. Here the space for OS, log, shared, and backup needs to be added, and the CPU requirements were calculated as category “XS” with 6,000 SAPS for the application and 14,000 SAPS for the HANA database part. Total SAPS in a 2-tier configuration were calculated with 16,000 SAPS.

This sizing shows that the RAM is the leading sizing factor of the system and that the CPU requirements are relatively low for this SAP Simple Finance example. Other SAP HANA applications, especially analytic applications, will require higher CPU bandwidth and typical SoH and S/4HANA applications.

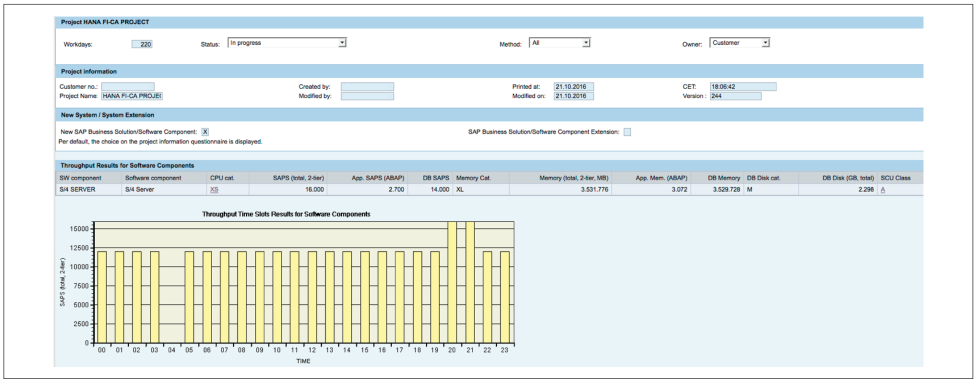

Figure 55: Sizing Result of an SAP Simple Finance on HANA Application, Quick Sizer Version 244

Translating the Quick Sizer result into a VM configuration is relatively easy. It is just a matter of adding the resources needed by the hypervisor; in this case, 1% for RAM and 5% for compute because overhead is used:

SoH Server:

Total RAM needed on host: 3,531,775 MB + 1%	=	3,567,093 MB
Total RAM VM:	=	3,531,775 MB
Total SAPS VM:	=	16,000 SAPS
Total CPU sockets:	=	4 (3484 GiB/1024 GiB)
Total CPU cores:	=	7 (16,000 SAPS/2357 SAPS per core)
Total vCPU:	=	14 (including hyperthreads)

Since we have to allocate close to 3.5 TiB of vRAM, we have to allocate all CPU sockets/NUMA nodes for this VM when running in production. The unused CPU and RAM resources cannot be used for other workloads because NUMA node sharing is yet not certified for production-level SAP HANA workloads. With SoH, it is allowed to install the ABAP SAP Central Services (ASCS) instance on the same OS as the SAP HANA database. This allows a 2-tier configuration.

For non-production environments, you can use the free CPU resources (172 threads and 450 GiB vRAM) for other VMs such as the primary application servers (PAS) and additional application servers (AAS). A 3-tier configuration, where additional application server instances are needed to scale the application tier, can be installed on the same server. See Figure 56 (right side) for details.

Note: NUMA node sharing, which is standard for non-HANA workloads, is not yet supported with HANA. Therefore, for production-level SAP HANA VMs, all CPU sockets should be assigned to this VM; running other workloads on these NUMA nodes is not supported to mitigate possible performance problems of the production SAP HANA instance.

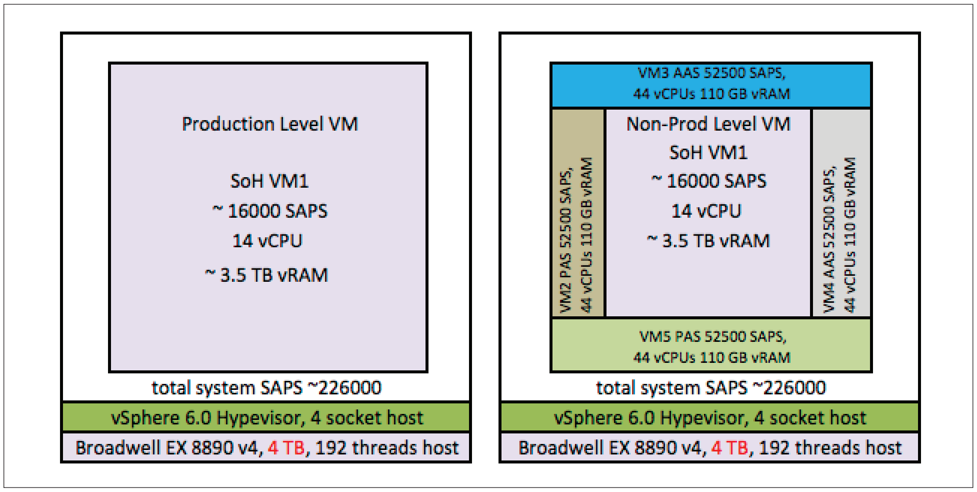

Figure 56: SoH Sizing Example Based on Quick Sizer Result

Figure 57 shows a Quick Sizer sizing example of an SAP Business Warehouse powered by an SAP HANA application.

The sized BW application is based on a customer example for approximately 240 business planning and simulation users (H-PLANN-1, 150, -2, 50, and -3 40) and around 900 reporting and analysis users (e.g., H-BW-INFO, or H-BW-BUSI).

For the sizing example, a BW on HANA system with several billion data records, 11 InfoCubes, and 16 DSO objects were used for the calculation. For details, review the example online by entering the customer number **188213** and project **TBTEST_ LARGE**. Click **Display Project**.

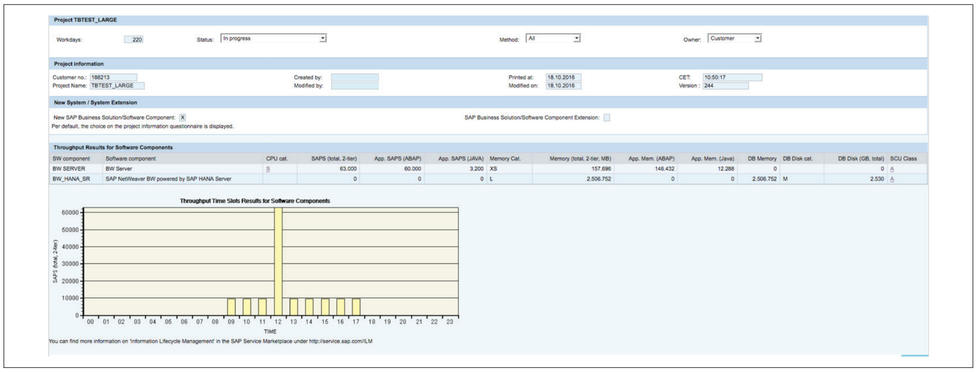

Figure 57: Sizing Result of an SAP Business Warehouse Powered by SAP HANA Application, Quick Sizer Version 244

The Quick Sizer result now shows two different lines, one for the BW application part and one for the SAP HANA database that stores the BW data used by the BW server.

Translating this result into a VM configuration will display the following figures:

BW Server VM

Total RAM need on host: 157696 MB + 1%	=	159273 MB
Total RAM VM:	=	157696 MB
Total SAPS VM:	=	63000 SAPS
Total CPU cores:	=	28 (≈63000/2295)
Total vCPU:	=	56

HANA Server VM

Total RAM need on host: 2506752 MB + 1%	=	2531793 MB
Total RAM VM:	=	2506752 MB
Total CPU Sockets:	=	5 (≈2448 GiB/512 GiB) (CPU/RAM ratio)
Total CPU cores:	=	120 (5 x 24)
Total vCPUs	=	128 (vSphere 6.0 maximum)

Figure 58 shows how a production-level environment would look, based on an 8-socket Intel Xeon E7-v3 server system supported for use with HANA.

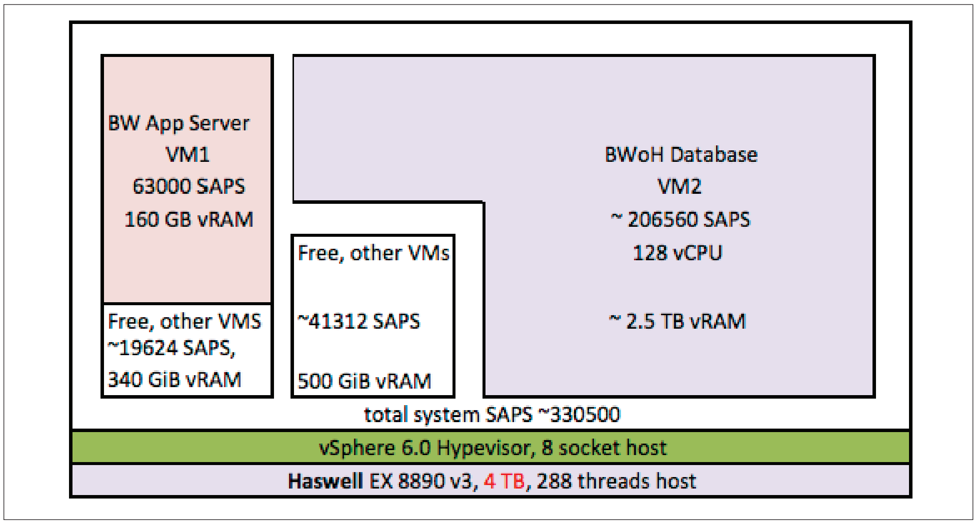

Figure 58: BWoH Sizing Example Based on Quick Sizer Result

Determining a virtualized SAP HANA server and VM configuration based on the Quick Sizer results

The Quick Sizer results for the SAP Simple Finance on HANA application example are:

Total RAM needed on host: 3,531,775 MB + 1%	=	3,567,093 MB
Total RAM VM:	=	3,531,775 MB
Total SAPS VM:	=	16000 SAPS

Using the SAPS figures in Table 13 allow us a much better sizing as the CPU Socket-to-RAM ratio and the formula of pervious section.

Table 15 is an extract of Table 13 showing the E7–v3 and E7–v4 CPU figures.

CPU	AVERAGE SAPS PER CPU/ SOCKET	AVERAGE SAPS PER CPU CORE	HANA WORKLOAD	HANA SPS VERSION	CPU SOCKET-TO-RAM RATIO
Intel Xeon EX E7-v4	56575	2357	SoH	SPS <= 11 SPS >= 12	<=768 GB/socket <=1024 GB/socket
			BWoH	SPS >= 11	<=512 GB/socket
<=1024 GB/ socket	41312	2295	SoH	SPS >= 11	<=768 GB/socket
			BWoH	SPS <= 10	<=384 GB/socket

Table 15: SAP HANA Platform Sizing Figures for E7-v3 and E7-v4 CPUs

The Quick Sizer result of the total SAPS requirements for the SAP Simple Finance on HANA application was 63,000 SAPS. Dividing the required SAPS with the SAPS capacity of a CPU core provides us with the amount of CPU cores required for this workload.

Virtualized SoH or S4/H SAP HANA SAPS formula

Note that this formula does not consider hyperthreading. If you want to take advantage of hyperthreading, then the VM configuration parameter numa.vcpu.preferHT=TRUE has to be set and the number of the calculated vCPUs needs to be doubled.

$$\text{needed vCPUs} < 10 = 10, > 128 = 128$$

$$\text{needed vCPUs} > 10 = \frac{total\ SAP\ HANA\ Quick\ Sizer\ SAPS}{SAPS\ per\ CPU\ core\ x\ CPU\ overhead}$$

SAP Simple Finance on HANA application VM on Broadwell configuration

Total RAM VM:	=	3.37 TiB
Total SAPS VM:	=	16000 SAPS
Total vCPUs VM	=	7 vCPUs (rounded up)

$$\text{needed vCPUs} > 10 = \frac{16{,}000}{2357\ x\ 0.9}$$

By doubling the vCPU count and by using the numa.vcpu.preferHT=TRUE VM configuration parameter, we will also leverage the available hyperthreads and gain around 15% more CPU performance for this VM. Use the following formula when hyperthreading must be considered:

$$\text{needed vCPUs} > 10 = \frac{total\ SAP\ HANA\ Quick\ Sizer\ SAPS}{SAPS\ per\ CPU\ core\ x\ CPU\ overhead\ x\ 1{,}15\ HT\ gain}\ x\ HT$$

$$\text{needed vCPUs} > 10 = \frac{16{,}000}{2357\ x\ 0.9\ x\ 1.15}\ x\ 2$$

SAP Simple Finance on HANA application VM configuration with numa.vcpu.preferHT=TRUE:

Total RAM VM:	=	3,37 TiB
Total SAPS VM:	=	16000 SAPS
Total vCPUs VM	=	14 vCPUs (with HT)

Using the Quick Sizer example result for the **SAP Business Warehouse powered by SAP HANA** application will provide following SAP HANA and SAP BW Server VM on Haswell configuration:

BW Server VM:		
Total RAM VM:	=	154 GiB
Total SAPS VM:	=	63,000 SAPS
Total vCPUs VM:	=	28 vCPUs (rounded up)
Alternative with HT vCPUs VM:	=	56 vCPUs (rounded up)

To calculate the needed vCPUs to support the HANA BW database VM, we have to use the CPU socket-to-RAM ratio or, when it is a non-production HANA VM, use the formulas introduced on page 73.

Production HANA database VM calculation:

For E7-v3 and E7-v4 based servers, the CPU socket-to-RAM ratio is 512 GB per CPU socket. E7-v3 server systems have 18-core CPUs and E7-v4 CPUs have 24 cores. When calculating the needed vCPU, base it on the NUMA node allocation and then all available CPU resources should be allocated.

HANA Server VM:		
Total needed RAM VM:	=	2448 GiB
Total needed NUMA nodes VM	=	2448 GiB/512 GiB = 4.7 = 5
Total to configure vCPUs VM	=	between 5 x 18 and 5 x 24 = 100-120 CPU cores or 128 vCPUs

Non-production HANA VM calculation:

If the vCPU-to-RAM formula is used, then it depends on which server is used. Because more CPU cores are available on Broadwell, but the RAM ratio of 512 GB is still the same as on Haswell, an average can be used for the calculation. Otherwise, there should be a very low CPU resource utilized system:

$$\text{needed vCPUs} > 10 = \frac{HANA\,VM\,RAM}{max.Appliance\ RAM\ x\ vRAM\ overhead} x\ max.usable\ CPU\ cores\ x\ 1.15$$

$$\text{needed vCPUs} > 10 = \frac{2448\ GiB}{4096\ GiB\ x\ 0.99} x \frac{(18\ (Haswell) + 22(Broadwell) x\ 8}{2}\ x\ 1.15$$

Total to configure vCPUs VM = 112 (rounded up)

* The max. usable CPU cores of this example calculation is based on an 8-socket system, and because both the Broadwell (E7-v4) and Haswell (E7-v3) platform would be OK from a sizing perspective, we use a mix of them to calculate the needed vCPUs of the example VM.

Note: As of January 2017, 8-socket Intel E7-v4 systems are currently not supported, but are in certification.

Note: When planning a complete landscape, then the development, test, and QA systems, as well as the HA systems, have to be planned and configured. All of these systems can be consolidated on a vSphere SAP HANA resource cluster and the production-level VMs can be protected with VMware HA. This would allow a massive reduction on asset, maintenance, and operational costs.

Final comment on sizing: The sizing of virtualized SAP HANA works just the same as sizing a natively installed SAP HANA system within the VM resource limitations. The benefit is that an SAP HANA VM can be configured as sized and it is not required to use all the resources of a host. Below is a general statement on SAP HANA sizing according to the SAP document, "Sizing Approaches for SAP HANA"[83]:

> "Both the advantage and the disadvantage of the sizing reports is that they provide a number, that is, the projected RAM requirements for SAP HANA. Some people take this number, ask one or several hardware vendors for a configuration and then they are done. By doing this they are not always doing themselves a favor, as both reports offer sufficient food for thought that may alter an initially conceived target configuration. Not always towards the higher end, sometimes also to a smaller configuration."

[83] SAP HANA Sizing Approaches

As of October 2010, with the current vSphere version 6.0, an SAP HANA single, Scale-Up VM cannot be virtualized once it is larger than 4 TiB of RAM. If the application type is BWoH, then a Scale-Out configuration can be used. SAP recommends using Scale-Out configurations for physical systems that are larger than 6 TiB.

Use dynamic data tiering or constant data archiving to lower the memory requirements of an SAP HANA instance.

Note: SAP HANA VMs that are larger than 4 TiB cannot be virtualized with the currently available vSphere and supported versions.

Using SAP HANA sizing reports to determine the memory and disk requirements of SAP HANA (existing systems)

The SAP HANA sizing of existing systems can be done with the already mentioned by running specific SAP HANA sizing reports (BW: /SDF/HANA_BW_SIZING and S/4: /SDF/HDB_SIZING). Review SAP Note 1872170 (S4/HANA) and SAP Note 2296290 (BW) for downloading and using these reports.

For example, Figure 59 shows an excerpt of a BW sizing report and compares the target database size of 398 GiB with available 512 GiB and 1 TiB RAM HANA appliance systems. When data growth is considered, then the 512 GiB system won't be able to support data growth for a long time; therefore, the report automatically recommends the next larger system (in this case, a 1 TiB RAM appliance system). Besides the needed memory, the storage requirements for data and log is calculated, also. Please be aware that the space for log is shown only for BW systems. The ABAP report for Suite systems calculates only the needed space for data. Here it is required to use the introduced storage sizing guidelines that SAP has published for SAP HANA.

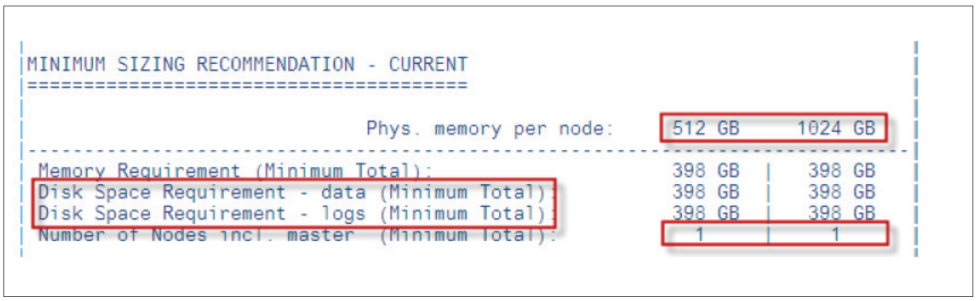

```
MINIMUM SIZING RECOMMENDATION - CURRENT
=======================================

                         Phys. memory per node:     512 GB       1024 GB
--------------------------------------------------------------------------
 Memory Requirement (Minimum Total):                398 GB   |   398 GB
 Disk Space Requirement - data (Minimum Total):     398 GB   |   398 GB
 Disk Space Requirement - logs (Minimum Total):     398 GB   |   398 GB
 Number of Nodes incl. master  (Minimum Total):       1      |     1
```

Figure 59: Excerpt from a BW Sizing Report[84]

Before running the SAP HANA sizing report, it is a good practice to think about which data needs to get migrated and which data can be archived or removed from the database. Archiving old data will help to reduce the memory requirements and the resource needs of the new SAP HANA server system or VM.

[84] SAP HANA Sizing Approaches, page 12

The ABAP sizing report provides information on the SAP HANA system memory requirements only, so the SAP application server stack (Suite or BW) needs to be added to the overall landscape configuration. Use the BWoH sizing formulas or the NUMA socket sizing guideline to determine the needed CPU resources for this server as explained earlier.

If, for instance, the old environment had five application servers with a total of 50.000 SAPS and 8 GB RAM each, then it is a good practice to either consolidate the application server to fewer systems (less admin and operation effort) or sing the same amount of applications server as VMs. In this case, five application server VMs with 10.000 SAPS and 8 GB RAM.

Future SAP HANA native applications may reduce the need for additional application servers, but when migrating existing systems, follow this approach:

1. Determine the SAP HANA database server configuration.
2. Migrate existing application server systems to VMs, and align the new configuration to the new server capabilities because newer systems may provide higher SAPS capacity per CPU core.

HANA database server selection process based on memory requirements

Once the SAP HANA database server and application server resource requirements are defined and calculated, it is possible to select the right server systems for the new SAP HANA based system landscape.

Remember, for SAP HANA only certified or supported Intel Xeon E5 and E7 based server systems can be used.

Here we will present two examples. The first example shows you how to select the right server and VM configuration based on the ABAP sizing reports. The second example uses the Quick Sizer example discussed previously.

Another example of determining a virtualized SAP HANA server and VM configuration based on ABAP sizing reports

Assumption: A customer wants to migrate three existing SAP database systems to SAP HANA. The customer has executed the SAP HANA ABAP sizing reports for two existing SAP ERP systems and one existing BW system with following results and recommendations in terms of needed SAP HANA appliance RAM sizes (total system RAM). SAP Quick Sizer is not used because the CPU resources are not the limiting factor for this configuration, and it is expected that the CPU utilization level is very low for the total server system.

The following three SAP HANA systems will be migrated to virtualized SAP HANA:

1. ERP SID1 – 300 GiB system RAM, data growth rate per month 4 GiB (per year 48 GiB)
2. ERP SID2 – 140 GiB system RAM, data growth rate per month 0.7 GiB (per year 8.4 GiB)
3. BI – SID3 – 650 GiB system RAM, data growth rate per month 9.97 GiB (per year 119.67 GiB)

If the system configuration has to support the environment for three years without data archiving, the new sizes would be:

1. ERP SID1 – 444 GiB
2. ERP SID2 – 165.2 GiB
3. BI – SID3 – 1009 GiB

The total memory needed over three years would be 1618.2 GiB + 1 percent RAM overhead for vSphere (~ total 1635 GiB).

Because Business Suite and Business Warehouse database systems will get consolidated on the same HANA server, the selected x86 server configuration has to be certified to run BW powered by HANA (BWoH). The memory configuration and CPU socket-to-RAM ratio for BW (512 GiB/socket) will be used for this example and will define the whole configuration because it is not allowed to exceed this RAM-to-socket ratio.

With this information, you can use the Certified SAP HANA Hardware Directory to map the sum of all memory-sized SAP HANA systems to a listed server configuration.

When selecting a server from the catalog, ensure that you click on the XFS file system because GPF is not supported in VMware virtualized environments. Figure 60 shows only some of the certified SAP HANA certified hardware vendors. We also pre-selected the BWoH application type because BWoH will get deployed and 2 TiB RAM, because the smaller appliance configurations would not be able to support the required total memory need of ~1.6 TiB.

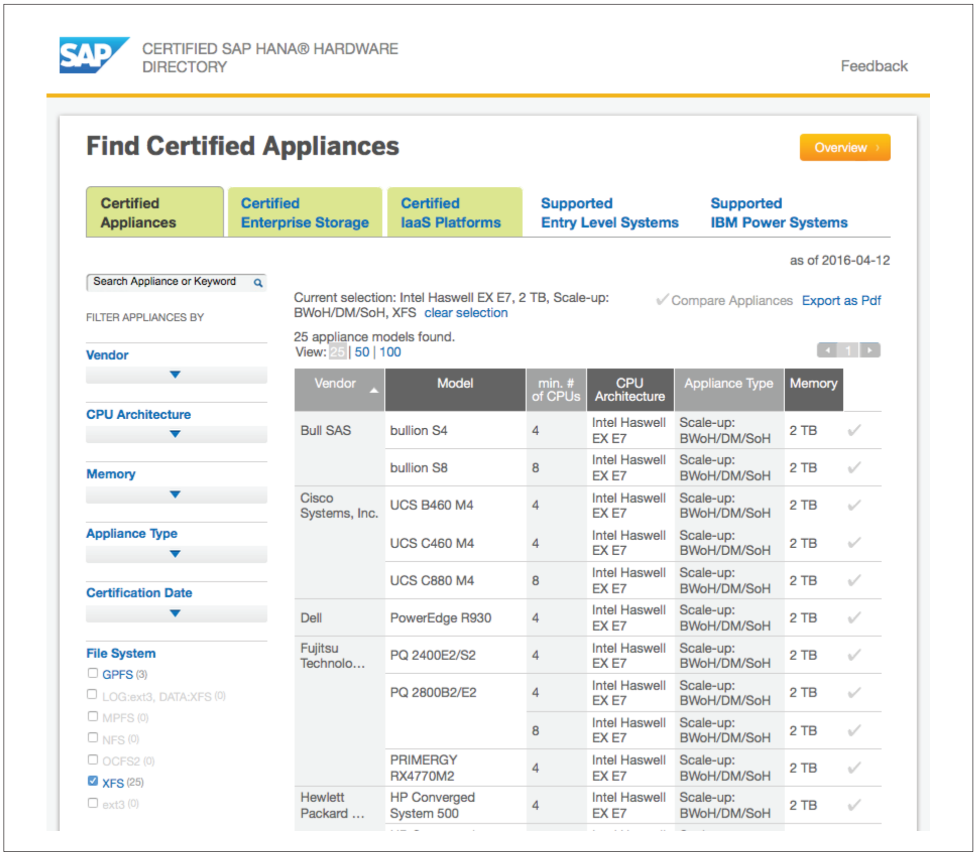

Figure 60: Certified SAP HANA Hardware Directory

The 2 TiB and larger server systems support the example configuration. We selected a 2 TiB BWoH SAP HANA appliance server configuration from the list. Now you must verify that the selected server is VMware vSphere certified for the desired vSphere version. Otherwise, the configuration is not supportable.

Access the VMware vSphere HCL[85] and use following filter settings to view a list of servers that support the selected RAM configuration. See Figure 61 for details.

The suggested filter settings are highlighted in Figure 59:

• vSphere release, ESXi 6.0 U2 and ESXi 5.5 U3
• CPU Series, Intel Xeon E7-8800/4800-v3

Note: Do not enter any data in the **Min Certified Memory** field.

[85] https://www.vmware.com/resources/compaTiBility/search.php

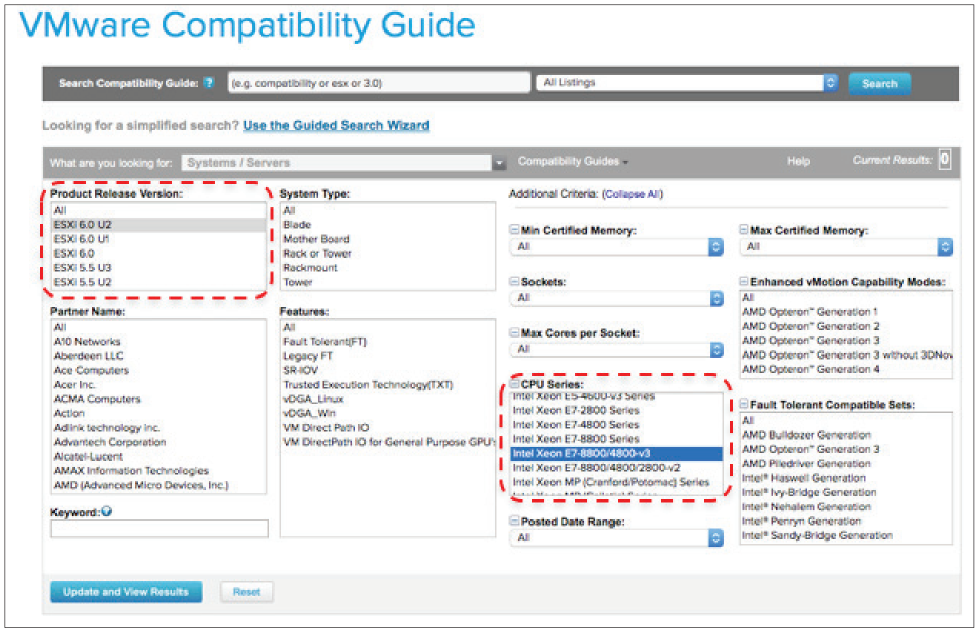

Figure 61: VMware Compatibility Guide

After the filters are configured, click **Update and View Results**.

Figure 62 lists the results of this query in alphabetical order, and allows verification if the selected SAP HANA supported server is also VMware vSphere supported. In this example, 29 systems comply with the search criteria and can be used with SAP HANA.

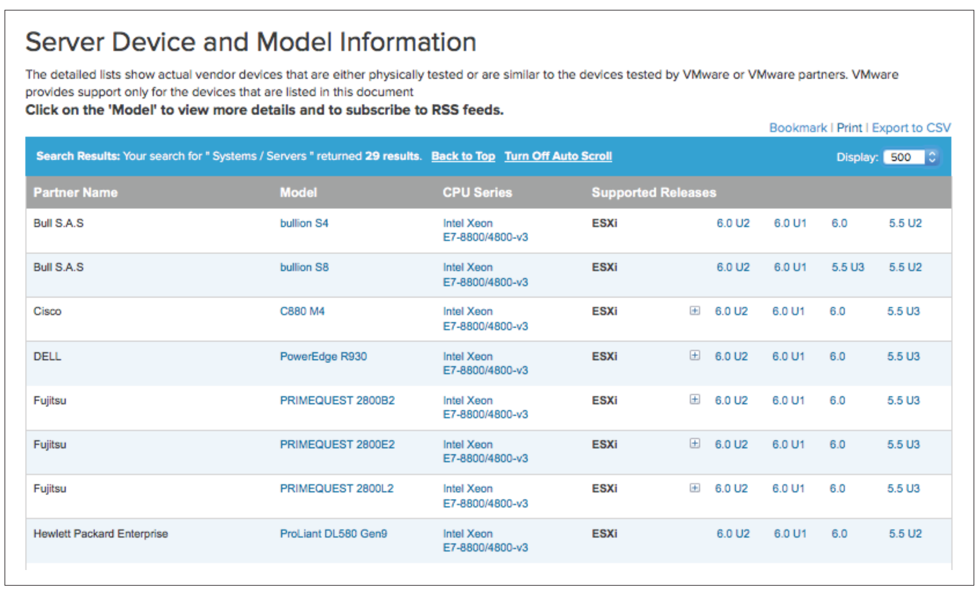

Partner Name	Model	CPU Series	Supported Releases				
Bull S.A.S	bullion S4	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	6.0	5.5 U2
Bull S.A.S	bullion S8	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	5.5 U3	5.5 U2
Cisco	C880 M4	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	6.0	5.5 U3
DELL	PowerEdge R930	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	6.0	5.5 U3
Fujitsu	PRIMEQUEST 2800B2	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	6.0	5.5 U3
Fujitsu	PRIMEQUEST 2800E2	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	6.0	5.5 U3
Fujitsu	PRIMEQUEST 2800L2	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	6.0	5.5 U3
Hewlett Packard Enterprise	ProLiant DL580 Gen9	Intel Xeon E7-8800/4800-v3	ESXi	6.0 U2	6.0 U1	6.0	5.5 U2

Figure 62: VMware HCL Results View

Let's have a closer look at one of the servers, such as the Dell PowerEdge R930, which is also SAP HANA supported and certified. The detail view in Figure 63 shows which vSphere features were validated and the maximum amount of RAM with which the server was certified. In this case, all SAP HANA relevant vSphere features can be used, and up to 8 TiB of RAM can be installed in the server.

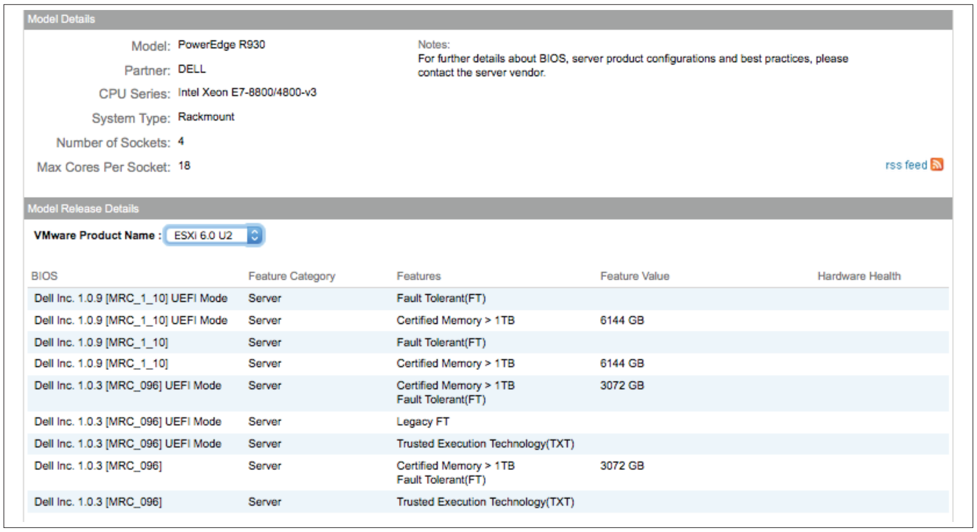
Model Details

Model: PowerEdge R930
Partner: DELL
CPU Series: Intel Xeon E7-8800/4800-v3
System Type: Rackmount
Number of Sockets: 4
Max Cores Per Socket: 18

Notes:
For further details about BIOS, server product configurations and best practices, please contact the server vendor.

rss feed

Model Release Details

VMware Product Name : ESXi 6.0 U2

BIOS	Feature Category	Features	Feature Value	Hardware Health
Dell Inc. 1.0.9 [MRC_1_10] UEFI Mode	Server	Fault Tolerant(FT)		
Dell Inc. 1.0.9 [MRC_1_10] UEFI Mode	Server	Certified Memory > 1TB	6144 GB	
Dell Inc. 1.0.9 [MRC_1_10]	Server	Fault Tolerant(FT)		
Dell Inc. 1.0.9 [MRC_1_10]	Server	Certified Memory > 1TB	6144 GB	
Dell Inc. 1.0.3 [MRC_096] UEFI Mode	Server	Certified Memory > 1TB Fault Tolerant(FT)	3072 GB	
Dell Inc. 1.0.3 [MRC_096] UEFI Mode	Server	Legacy FT		
Dell Inc. 1.0.3 [MRC_096] UEFI Mode	Server	Trusted Execution Technology(TXT)		
Dell Inc. 1.0.3 [MRC_096]	Server	Certified Memory > 1TB Fault Tolerant(FT)	3072 GB	
Dell Inc. 1.0.3 [MRC_096]	Server	Trusted Execution Technology(TXT)		

Figure 63: Detailed View of a Server in the VMware HCL

Verifying the server configuration with the HANA certified directory provides the following information as shown in Figure 64. The detailed view provides information on CPU type and supported memory, and includes the local appliance storage configuration, which is not required when the server is used in a TDI configuration.

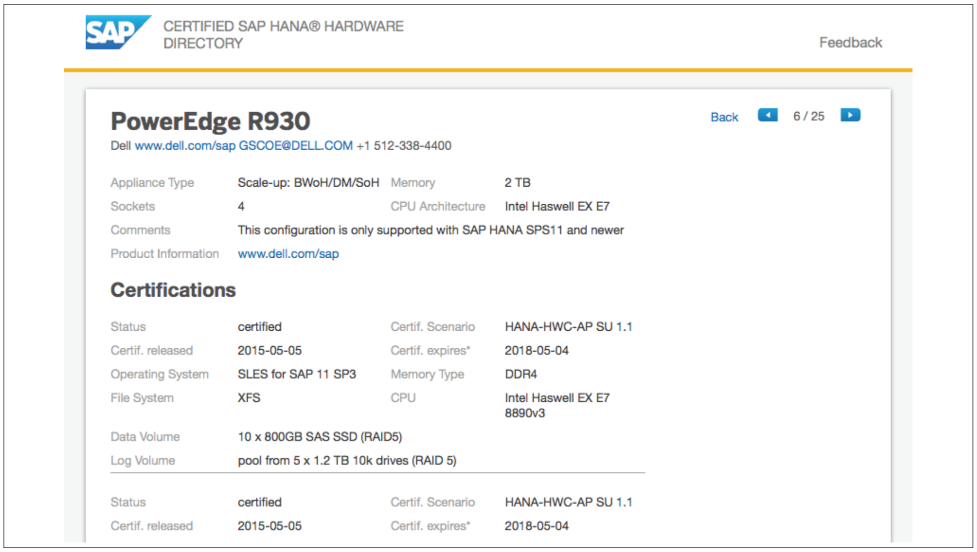

Figure 64: HANA Certified Directory Server Detailed View

Matching the SAP HANA and VMware HCL results provides the necessary information to ensure that a system is supported by both companies.

Selecting this server now allows us to fine-tune the SAP HANA VMs because we want to ensure optimal performance by maintaining the NUMA node locality.

Back to the example configuration, these are the initial VM sizes to support data growth over the next three years:

1. ERP SID1 – 444 GiB
2. ERP SID2 – 165.2 GiB
3. BI SID3 – 1009 GiB

The selected server has four 18-core Xeon E7 8890-v3 CPUs (72 cores and 144 threads total) and 2 TiB RAM installed, which we now have to assign to the planned VMs.

In production environments, SAP requests that the NUMA node not be shared with another VM; therefore, the first two instances can be all available compute resources; that is, 18 CPU cores, 36 threads, and 512 GiB of RAM. The third VM, the BI SID3 instance, needs to be configured with two NUMA nodes and all available resources because the RAM requirements are bigger as on one CPU socket is connected.

If no SAPS figures exist on the old systems and if no Quick Sizing was done, then the easiest process to define production-level SAP HANA VM configurations is to map the VMs to the SAP HANA CPU socket-to-RAM ratio by using the **NUMA node sizing guidelines** (see "Table 13. SAP HANA Platform Sizing Figures for Selected CPU Types"). This ratio depends on the used CPU version, the HANA SPS level, and the SAP HANA workload type that will be deployed. Calculating the CPU socket-to-RAM ratio can be done just by dividing the maximum RAM supported in an SAP HANA certified server for BWoH or SoH and a specific SAP HANA SPS level.

Using the NUMA node sizing guidelines and the Intel Xeon E7-v3 system values listed in "Table 13. SAP HANA Platform Sizing Figures for Selected CPU Types" for SoH = 768 GB/socket and BWoH = 512 GB/socket allows easy mapping of the sized systems to a virtualized environment. In this case, we have to use the BWoH 512 GB/socket configuration because the BWoH system is the leading system in this configuration and the systems usually have a homogeneous memory configuration, where all CPU sockets have the same memory configuration.

With these figures, defining the size of the VM (CPU and RAM resources) is actually just a matter of dividing the sized HANA memory size with the RAM connected on a NUMA node (as outlined in Table 13), and then by rounding up the result:

1. ERP SID1 – 444 GiB/512 GiB RAM = 0.87 x RAM of CPU socket
2. ERP SID2 – 165.2 GiB/512 GiB RAM = 0.32 x RAM of CPU socket
3. BI – SID3 – 1009 GiB/512 GiB RAM = 1.97 x RAM of CPU socket

Note: The CPU socket-to-RAM ratios depend on the used CPU and represent the currently supported RAM maximums per socket as defined by SAP. Actual server systems may have less RAM connected on a CPU socket /NUMA node. If the maximums don't get used, then the actual RAM attached to a socket needs to get used.

The example configuration would be able to run on vSphere versions 5.5 and 6.0 because the largest VM is below 1 TiB.

The final VM configurations would look like this, where all resources of the selected Haswell server touched NUMA nodes are used for the VM. In this configuration, we use all available CPU resources (36 cores per socket for Haswell) as for production-level SAP HANA VMs it is yet not supported to share a NUMA node with other instances. Not leveraging the CPU resources available would waste the available and not usable CPU resources:

1. ERP SID1 = 1 NUMA node with 36 vCPUs (18 cores plus 18 threads) and ~ 512 GiB vRAM
2. ERP SID2 = 1 NUMA node with 36 vCPUs and ~ 512 GiB vRAM
3. BI SID3 = 2 NUMA nodes with 72 vCPUs and ~ 1024 GiB vRAM

Figure 65 shows a graphical view of this configuration. The CPU resources of the large BW VM should get pinned. For the single NUMA node VMs, this is not necessary because these VMs do not span a NUMA node.

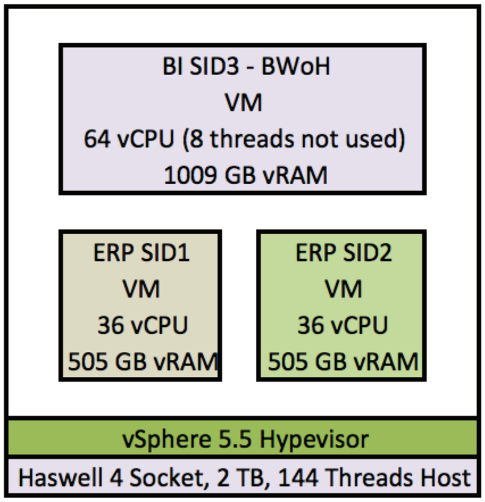

Figure 65: Server Configuration Example with Three Independent SAP HANA Instances and 1.5% RAM Overhead

Note: This calculation provides the amount of needed NUMA nodes based on RAM per CPU socket, but does not consider the vCPU limits of the current vSphere hypervisor; therefore, we recommend that you verify this calculation with the following formulas, as it may happen, that with BW workloads not all assigned RAM of a NUMA node can be used due to higher CPU demands of an OLAP such as workload. See Table 14 for the maximum usable RAM per VM and HANA workloads.

The example shows how an SAP HANA VM can be sized using the RAM sizes provided by the HANA ABAP report and the SAP HANA CPU socket-to-RAM ratios.

Using fixed CPU-to-RAM ratios specified by SAP does not consider how an application utilizes CPU resources during its lifetime. Selected SAP HANA appliance server systems and systems sized with this approach may result in systems with very low average CPU utilization.

The next example shows how the previous example could be configured to run on as few resources as possible to achieve better server utilization ratios.

Note: For non-production SAP HANA systems, strict resource commitments are required. Overcommitting system resources for SAP HANA VMs is not supported.

Sizing the development and test systems of the example landscape

Redefine the SoH and BWoH VM configuration based on the RAM needed, the defined BWoH formulas, and NUMA node sharing to better leverage the system resources. In this example, we use 1.5% as memory overhead (0.985). Again, memory overhead could be as little as 0.5% and as much as 3%. This depends on the configured virtual hardware configuration.

Again, you will receive the best results for CPU requirements once a Quick Sizer calculation is done or the SAPS figure of an existing SAP environment is known. Until then, the calculation example below helps to determine VM configuration by calculating the CPU needed based on the SAP defined socket-to-RAM ratios, which result for SoH in a CPU over sizing. For BWoH, it is currently the only method to calculate the CPU requirements.

1. ERP SID1 VM = 444 GiB vRAM and 18 vCPUs

$$\text{needed vCPUs} = \frac{444\text{ GiB}}{2048\text{ GiB x }0.985}\text{x }72\ x\ 1.15$$

$$\text{needed vCPUs} = 18.22$$

2. ERP SID2 VM = 165.2 GiB vRAM and 10 vCPUs

$$\text{needed vCPUs} = \frac{165.2\text{ GiB}}{2048\text{ GiB x }0.985}\text{x }72\ x\ 1.15$$

$$\text{needed vCPUs} = 6.78$$

3. BI SID3 VM = 1009 GiB vRAM and 42 vCPUs

$$\text{needed vCPUs} = \frac{1009\text{ GiB}}{2048\text{ GiB x }0.985}\text{x }72\ x\ 1.15$$

$$\text{needed vCPUs} = 41.41$$

Free resources: 74 CPU threads and 430 GB RAM for other VMs

Sizing the non-production VMs with the calculated figures allows you to use the non-utilized 74 CPU threads and 430 GB RAM (-1.5% = 420 GB) for other workloads, such as for other SAP HANA VMs or to reserve these free resources to provide VMware HA resources for potentially failing VMs of other host systems that now may get restarted on this host.

The left side of Figure 66 shows the newly configured VMs with the number of vCPUs required to support the needed RAM sizes. The right side of the figure shows how this space can be used by other VMs, such as the application servers of an additional SAP HANA QA instance. Strict resource commitments will ensure a stable and supportable SAP HANA non-production environment.

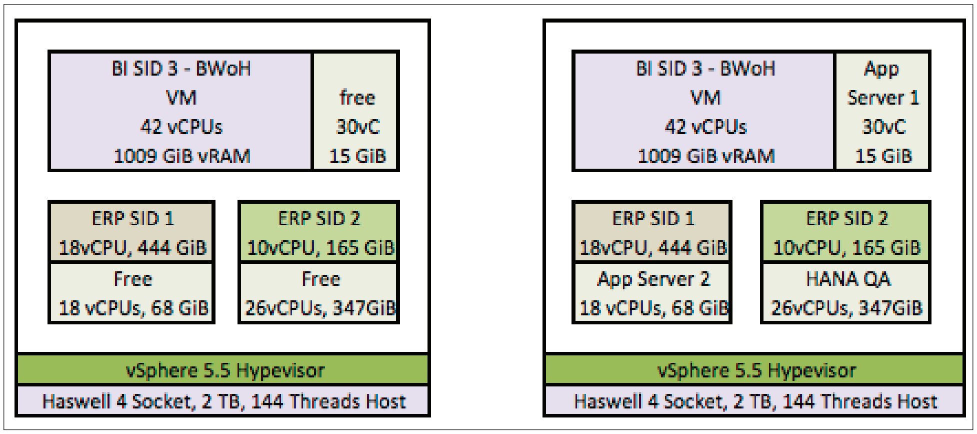

Figure 66: Example Server Non-production Configuration with Five Independent SAP HANA Instances

As stated above, with this configuration it is important that system usage and utilization is monitored and, if required, the different VM configurations have to be modified to ensure best performance. While operating in this environment, it may be required to increase the CPU resources of BI SID3. A CPU resource reduction of the App Server 1 VM would be required. If this is not possible, then the App Server 1 VM has to be moved to another server to free up resources for the BI SID3 VM.

Virtual HANA Configuration Examples

The following section provides some example configurations that follow the SAP HANA CPU socket/core-to-RAM ratios and use the all available vCPUs or the maximum addressable CPU cores available (128 vCPU limit). Please be aware that the SoH examples will have a particularly low CPU utilization because memory is the leading sizing factor for an in-memory database, and memory needs to be assigned strictly to an SAP HANA VM.

The examples in this section are based on the latest Intel Xeon EX E7-v4 systems, and Figure 67 and Figure 68.

CPU	AVERAGE SAPS PER CPU/SOCKET	AVERAGE SAPS PER CPU CORE	HANA WORKLOAD	HANA SPS VERSION	CPU SOCKET-TO-RAM RATIO
Intel Xeon EX E7-v4	56575	2357	SoH	SPS <= 11 SPS >= 12	<=768 GB/socket <=1024 GB/socket
			BWoH	SPS >= 11	<=512 GB/socket

Table 16: Virtual HANA Configuration Examples

Note: The following examples show virtualized SAPS (physical SAPS -10% for VM sizing). From the RAM shown, between 0.5% and 3% on RAM must be subtracted from the total available physical RAM for the virtualization memory costs. The actual RAM overhead can be defined only once the VMs have been configured.

Intel Xeon E7-v4 CPU 4-Socket System T-shirt HANA Examples (Figure 67 and Figure 68.)

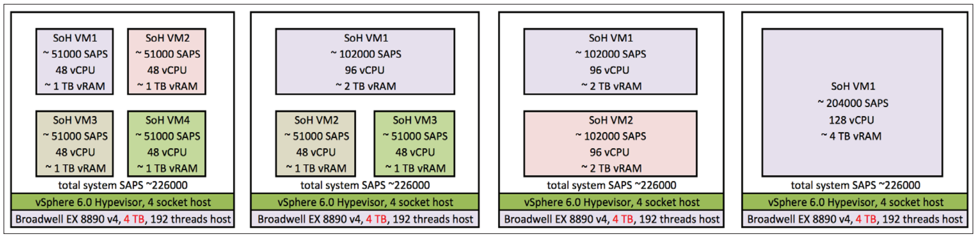

Figure 67: Example Server Configuration for SoH Instances on 4-Socket 4 TB E7-v4 Systems

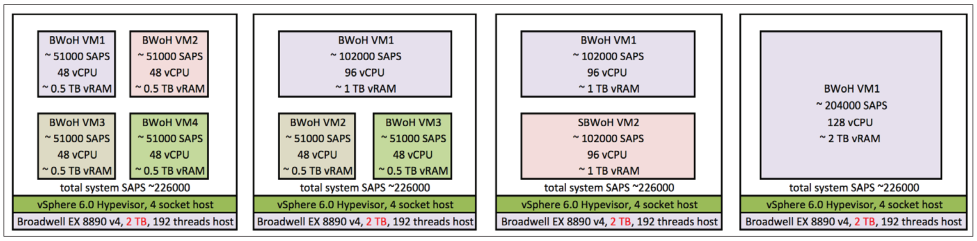

Figure 68: Example Server Configuration for BWoH Instances on 4-Socket 2 TB E7-v4 Systems

Best Practices and Configuration Settings

This section provides information on best practices and suggested settings for SAP HANA on VMware vSphere. The provided guidelines were tested and validated during the joint SAP HANA on vSphere validation and performance tests done with SAP. Some of the settings are recommended only when the highest achievable performance or the lowest latency is required.

Reference documents and generic guidelines

- "VMW-Tuning-Latency-Sensitive-Workloads"
- "Timekeeping in VMware Virtual Machines"
- SAP Note 2161991–VMware vSphere configuration guidelines

Deployment Options and Support Status

When deploying SAP HANA on vSphere, the main consideration is if the SAP HANA virtual machine will be used for production or non-production workloads. Table 16 summarizes the capabilities, supported deployment options, and best practices, such as minimal vCPU count or maximum vRAM sizes for SAP HANA VMs on VMware vSphere. This table is based on the support status as of May 2016 and may be outdated. Check out the SAP HANA on VMware vSphere SCN page for updated support status and deployment options.

CAPABILITY/OPTION	VSPHERE 5.5		VSPHERE 6.0/6.5[1]	
	SUPPORTED IN PRODUCTION	NO SAP SUPPORT, VMWARE BEST EFFORT SUPPORTED	SUPPORTED IN PRODUCTION	NO SAP SUPPORT, VMWARE BEST EFFORT SUPPORTED
SAP HANA Scale-Up VM ≤ 1024 GiB	Yes			
SAP HANA Scale-Up VM ≤ 4080 GB	No		Yes	
SAP HANA Scale-Out VM ≤ 1024 GB[2]	Yes		No	Yes
SAP HANA Scale-Out VM ≤ 4080 GB[2]	No		No	Yes
SAP HANA Multi-VM[3]	Yes		Limited[4]	Yes
SAP HANA Version	SPS07 and above		SPS11 and above	SPS09 and above
VMware vSphere and SAP HANA HA and Operation Features[5]				
VMware HA	Yes			
SAP HANA Host Auto-Failover[6]	Yes			
SAP HANA System Replication	Yes			
VMware FT for HANA	No			
VMware SRM	Yes			
vSphere vMotion	Yes			
VMware DRS[7]	Yes			
Supported Hardware Configurations[8]				
Supported SAP HANA Systems for VMware Virtualization	Only SAP HANA and VMware certified 2-, 4- and 8-socket Intel E7 v2 and later Intel processor-based server systems and Intel Xeon E5 v3 and v4 based 2-socket single node SAP HANA entry level systems, with a minimum of 8 cores per CPU are supported.			
Intel Xeon E7 CPU Support	v2, v3		v2, v3, v4	
Maximum E7 NUMA Nodes	4	up to 8 (4 with E7 v4)		
Intel Xeon E5 CPU Support	v2, v3	v2, v3, v4	v2, v3, v4	

CAPABILITY/OPTION	VSPHERE 5.5		VSPHERE 6.0/6.5[1]	
	SUPPORTED IN PRODUCTION	NO SAP SUPPORT, VMWARE BEST EFFORT SUPPORTED	SUPPORTED IN PRODUCTION	NO SAP SUPPORT, VMWARE BEST EFFORT SUPPORTED
Maximum E5 NUMA Nodes	2			
NUMA Node Sharing[9]	No	Yes	No	Yes
Enable Hyperthreading	Yes			
Maximum RAM Installed in Server[10]	6 TiB		12 TiB	
Supported Storage Configuration				
Supported SAP HANA Storage Systems for Virtualization[11]	All SAP HANA TDI and VMware certified/supported storage solutions can be used.			
TDI Storage KPI Requested	Yes, per HANA VM	No	Yes, per HANA VM	No
SAP HANA Virtual Machine Configuration[12]				
Max. VM Size #vCPU	64		128	
Max. VM Size #vRAM	1024 GB		4080 GB	
Minimal VM Size #vCPU	All threads of a single CPU socket	10	All threads of a single CPU socket	10
Minimal VM Size #vRAM	RAM of CPU socket	As sized	RAM of CPU socket	As sized
Minimal VM Virtual Hardware Version[13]	10		11	
Supported OS for Virtualized SAP HANA[14]				
SLES 11 and 12	Yes			
RHEL 6 and 7	Yes			

1 vSphere 6.5 only SAP HANA certified/supported 2 and 4 socket Intel E7 Broadwell EX or 2 socket Intel E5 Broadwell EP processor-based configurations in single-node, scale up configurations only, as of SAP HANA SPS12 or later.

2 SAP support for SAP HANA scale-out deployments of SAP Business Warehouse (BW), powered by SAP HANA, on VMware vSphere as outlined in SAP Note 215787.

3 The defined KPIs for data throughput and latency for production SAP HANA systems have to be fulfilled for each VM; see SAP Note 1943937.

4 Just as with the vSphere 5.5 SAP HANA support release in 2014, vSphere 6.0 currently supports only one production-level VM that may be co-deployed with nonproduction-level SAP HANA VMs. No resource sharing with other VMs is supported for production-level SAP HANA VMs.

5 Generally, all VMware vSphere features, such as distributed switches, are supported; no support for VMware FT for SAP HANA, FT may be used with an (A)SCS instance. vMotion and DRS are listed specifically because running SAP HANA VMs can be moved from one host to another host while the application is online and users are connected, and any negative impact needs to get excluded.

6 Only available with specific SAP HANA TDI Storage vendors; see SCN page DOC-60470.

7 VMware DRS should be configured in "manual mode."

8 Both the SAP HANA appliance and SAP HANA Tailored Datacenter Integration (TDI) delivery methods are supported for SAP HANA on VMware vSphere.

9 As of writing this guide NUMA node sharing (more HANA VMs run on a single CPU socket) is not supported for production level workload. Future testing's will investigate supportability.

10 Maximum RAM installed in physical server as specified by SAP, in the Certified SAP HANA Hardware Directory. The listed memory figures are the maximum memory configurations supported by VMware vSphere. Up to 6 TB is supported for ESXi 5.5 Update 2 and later, and up to 12 TB is supported with vSphere 6.0 on specific OEM certified platforms.

[11] For details, refer to the Certified SAP HANA Hardware Directory and go to "Certified Enterprise Storage" configurations.

[12] The maximum size of a virtual SAP HANA instance is limited by the maximum size a virtual machine can be with the different VMware vSphere releases. Each SAP HANA instance/virtual machine is sized according to the existing SAP HANA sizing guidelines and VMware recommendations. CPU and memory overcommitting must not be used. The maximum size of a virtual SAP HANA instance is limited by the maximum size of a virtual machine in the VMware vSphere release. Each SAP HANA instance/virtual machine is sized according to the existing SAP HANA sizing guidelines and VMware recommendations. CPU and memory overcommitting must not be used. The maximum usable memory for virtualized BWoH systems may be lower as the vSphere memory maximums because the vSphere 6.0 128 vCPU support may limit the usable RAM size of a BW instance. For details, refer to the SAP HANA on VMware architecture guide.

[13] For best performance and full CPU feature support, the virtual VM hardware version must be aligned to the vSphere version. For example, the vSphere 5.5 VM virtual hardware version should be upgraded to version 10, a vSphere 6.0 VM should be at HW level 11.

[14] Both SUSE Linux Enterprise Server (SLES) 11 and 12 and Red Hat Enterprise Linux (RHEL) 6 are supported operating systems when virtualized. Check for an updated support status via the corresponding SAP support notes.

Table 17: SAP HANA on vSphere Supported Capabilities and Options

Best Practices–Configuration

Installation Steps for an SAP HANA Scale-Out Installation Without a Standby Host

Unlike in physical SAP HANA environments, it is possible to operate SAP HANA as highly available without a standby host. VMware HA can substitute for host/VM failover functionality.

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

1. Install SAP HANA as a single host installation on the first VM. After a successful installation, go to „/usr/ sap/<SID>/SYS/global/hdb/custom/config" and add the following entries to the „global.ini":

    [communication]
    listeninterface = .global
    [persistence]
    basepath_shared = no

    ----

2. After editing the global.ini file, restart this SAP HANA instance.
3. For the next step, you need the ID number of the <SID>adm (uid) and the sapsys group (gid) of the first SAP HANA node. To find these IDs, execute "id <SID>adm" on the first VM. For example:

    vm1 # id vmwadm

    uid=1001(vmwadm), gid=79(sapsys), groups=16(dialout), 33(video), 1001(vmwshm), 79(sapsys)

4. On the next VM, you want to add a worker node to the SAP HANA installation by creating "/hana/log/<SID>" and "/hana/data/<SID>". After creating the two directories, change the owner of those two directories to the previously retrieved <SID>adm (uid) and the sapsys group (gid) by executing:

    "chown <UID>:<GID><directory>",

    Example: vm2 # chown 1001:79 /hana/data/VMW

5. Now go to "/hana/shared/<SID>/global/hdb/install/bin" on the VM you want to add as a worker and execute "./hdbaddhost --storage_partition=<X> -H <Hostname>", where <X> is a number greater than 1 and <Hostname> refers to the internal hostname (the hostname used for the SAP HANA interconnect link) of the host you want to add. For example:

   vm2 # ./hdbaddhost --storage_partition=2 -H hana2int

6. Repeat these steps for every worker you want to add to your SAP HANA installation.
7. Check the status of added "Stand by" nodes in the "Landscape tab" of SAP HANA Studio.

Deactivate Scale–out Master Polling

Because we recommend using VMware HA instead of the standby node SAP HANA HA solution when SAP HANA was virtualized, we also recommend turning off the SAP HANA Scale-Out master node polling. This is not required because VMware HA will restart the existing master node fast enough so that promoting a new node to the master role is not necessary.

Enable "Automatic Start" of the SAP HANA Database

By default, the SAP HANA database is configured so it does not start automatically when the SAP VM is started or rebooted.

If VMware HA is selected as the HA solution for SAP HANA, we recommend enabling the automatic start feature of SAP HANA.

1. To do this, modify the SAP HANA database profile with the following procedure:
2. Log on to the SAP HANA host as a user with root authorization.
3. Change to the system profile directory /usr/sap/<SID>/SYS/profile.
4. Edit the profile <SID>_HDB<instance number>_<host name>.
5. Change the parameter setting Autostart = 0 to Autostart = 1.
6. Save the profile <SID>_HDB<instance number>_<host name>.

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

vMotion Best Practices

vMotion between different hardware generation of a CPU type is possible, but in the context of a performance critical application such as SAP HANA it is important to follow these two best practices when using this feature:

- Customer shall run SAP HANA VMs within vSphere cluster on identical HW (with same CPU clock speed & synchronized TSC) only.
- Customer may use vMotion during HW refresh (non-identical source and destination host/clock speed), but should plan for restart of the VM afterwards.

Best Practices–Optimizing SAP HANA on vSphere

VMware vSphere can run multiple virtual machines in a heterogeneous clustered environment. To create SAP landscapes, application owners and virtual infrastructure administrators need to weigh the advantages of absolute performance versus operation and heterogeneity.

This section describes how to optimally configure SAP HANA on vSphere.

The **listed settings** show the parameter settings recommended for the physical server, the ESXi host, the VM, and the Linux OS to achieve optimal operational readiness and stable performance for SAP HANA and on vSphere, and are **the minimal settings that should be configured for SAP HANA**.

Table 17 through Table 20 and the next sections show additional configuration settings for performance or optimal resource utilization.

HOST BIOS SETTINGS	DESCRIPTION
Enable Intel VT technology	Enable all BIOS virtualization technology settings
Enable CPU Intel Turbo Boost	Allow Intel automatic CPU core overclocking technology (P-States)
Disable QPI Power Management	Static high power for QPI links
Enable hyperthreading	Double the logical CPU cores
Enable execute disable feature	Enable the Data Execution Prevention bit (NX-bit), required for vMotion
Disable node interleaving	
Disable C1E Halt state	Disable Enhanced C-States in Bios
Set Power Management to High Performance	No power-saving modus (C-States)
Disable all unused BIOS features, such as:	Video BIOS Shadowable, Video RAM Cacheable, on-board audio, on-board modem, on-board serial ports, on-board parallel ports, on-board game port, floppy drive, CD-ROM, USB

Table 18: vSphere Physical Host Server Settings

VSPHERE ESXI HOST	DESCRIPTION
Physical NIC Settings	
Networking	Use: Virtual Distributed Switches (vDS) to connect all hosts that work together Define port groups that are dedicated to SAP HANA, management and vMotion traffic
Storage	
Log	Use dedicated HBAs/SCSI controller for DATA, LOG and ESXi Swap/Scratch partition, separate log from data! See storage section of this guide for more details and align the physical storage configuration to the logical VM storage configuration.
Data	
Swap	
Memory	
Configure RAM hemisphere mode	Distribute DIMM modules in a way to achieve best performance (hemisphere mode), use fastest DIMM modules available for the selected RAM size
CPU	
Populate all available CPU sockets, use glueless NUMA architecture	To avoid Timer Synchronization Issues, use a multi-socket server that ensures NUMA node timer synchronization. NUMA systems that do not run synchronized will need to 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. Details: http://service.sap.com/sap/support/notes/1409604
Virtual Machine Monitor	Allow vSphere to choose the best virtual machine monitor based on the CPU and guest operating system combination.
CPU/MMU Virtualization Option = Automatic	Hardware Assisted Memory Virtualization

Table 19: vSphere ESXi Server

SAP HANA VIRTUAL MACHINE	DESCRIPTION
UEFI BIOS	Use UEFI BIOS as the standard BIOS version for vSphere hosts
SAP monitoring	Enable SAP monitoring on the guest VM tools.guestlib.enableHostInfo = true Details: http://service.sap.com/sap/support/notes/1409604
vCPU hotplug	Ensure that vCPU hotplug is deactivated (KB 2040375), otherwise vNUMA gets disabled.
Tips how to edit the .vmx file	See Tips for editing a .vmx file (1714).
VMXNET3	Use paravirtual VMXNET 3 virtual NICs for SAP HANA virtual machines
Set"ethernetX.coalescingScheme" to disable	Disable virtual interrupt coalescing for VMXNET 3 virtual NICs that communicate with the App servers or front end
Set halt_in_monitor = "TRUE"	Edit VMX and add: monitor_control.halt_in_monitor = "TRUE" In default configuration of ESX6.0, idle state of guest HLT instruction will be emulated w/o leaving VMM if a vCPU has exclusive affinity. If affinity is non-exclusive, guest HLT will be emulated in vmkernel, which may result in having vCPU descheduled from the physical CPU and can lead to longer latencies, therefore it is recommend to set this parameter to "TRUE", to ensure that the HLT instruction gets emulated inside the VMM and not in the vmkernel.
Set idleLoopSpinBeforeHalt = "TRUE"	The guest operating system issues a Halt instruction which stops (or de-schedules) the vCPU on the ESX host. Keeping the virtual machine spinning longer before Halt negates the number of inter-processor Wake up requests. Setting this parameter to true help to improve performance of a VM. Edit VMX and add: monitor.idleLoopSpinBeforeHalt = "TRUE". idleLoopSpinBeforeHalt option tells VMM to spin before letting vCPU be descheduled.
Set Lat. Sensitivity = normal	Use the vSphere client and set in the VM configuration Latency Sensitivity Settings to "normal. Alternatively edit the VMX file and add: sched.cpu.latencySensitivity = "normal".
Set numa.vcpu.preferHT=TRUE	Important for Multi-VM configurations (not needed or recommended for 1:1 configurations) to ensure that an SAP HANA VM leverages the hyperthreads of the local NUMA node and does not try to schedule vCPUs on another NUMA node.

SAP HANA VIRTUAL MACHINE	DESCRIPTION
Dedicated CPU and memory resources for HANA instances	Do not overcommit CPU or memory resources and use as few NUMA nodes as possible to optimize memory access.
Align virtual CPU VM configuration to actual server HW	Example: A VM running on a 4-socket server with 10-core CPUs, should get configured with 10-cores per socket, such as 2-CPU socket and 20 vCPUs.
Paravirtualized SCSI driver for I/O devices	Use dedicated SCSI controllers for OS, log, and data, to separate disk I/O streams. For details see SAP HANA disk layout section
Use Virtual Machine File System	Use VMDK disks whenever possible to allow optimal operation via the vSphere stack.
Create dedicated and isolated datastores for SAP HANA data and log files	Ensure the storage configuration passes the SAP HANA hardware check tool storage and file system requirements
Use Eagerzeroed thick virtual disks for data and log disk	This avoids lazy zeroing
Remove unused devices	Such as floppy disks or CD-ROM to free up resources and to mitigate possible errors.

Table 20: SAP HANA VM Settings

LINUX OS	DESCRIPTION
SAP HANA DB: Recommended OS settings for SLES	SLES 11 for SAP - SP 2 http://service.sap.com/sap/support/notes/1824819 SLES 11 for SAP - SP 3 http://service.sap.com/sap/support/notes/1954788 SLES 11 for SAP - SP 4 http://service.sap.com/sap/support/notes/2240716 SLES 12 for SAP - SP 0/1 http://service.sap.com/sap/support/notes/2205917
SAP HANA DB: Recommended OS settings for RHEL	RHEL 6.5 http://service.sap.com/sap/support/notes/2013638 RHEL 6.6 http://service.sap.com/sap/support/notes/2136965 RHEL 6.7 http://service.sap.com/sap/support/notes/2247020 RHEL 7.2 http://service.sap.com/sap/support/notes/2292690

LINUX OS	DESCRIPTION
Disable transparent HugePages	THP is not supported for the use with SAP HANA DB, as it may lead to hanging situations and performance degradations. To check the current configuration run the following command: # cat /sys/kernel/mm/transparent_hugepage/enabled Its output should read: always madvise [never] If this is not the case, you can disable the THP usage at runtime by issuing the following: # echo never > /sys/kernel/mm/transparent_hugepage/enabled For details read above referenced support notes.
Configure C-States and BIOS EPB 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 Set Energy Performance in BIOS to "maximum performance", optional: cpupower set -b 0
GlibC update to avoid potential problem in the index server:	Update your GlibC to version glibc-2.11.3-17.56.2 or later.
Optional: Disable Large Receive Offload (LRO) in the Linux guest operating system to lower latency for client/application server-facing NIC adapter	This helps to lower network latency of client/application server facing NIC adapters run: “ethtool -K ethY lro off”. Do not disable LRO for throughput NIC adapters such as for Backup, replication or SAP HANA internode communication networks! Works only with Linux kernel 2.6.24 and later and uses a VMXNET3. Additional details: http://kb.vmware.com/kb/2055140
Linux kernel	VMware strongly recommends using the latest SAP HANA supported kernel version
Minimal supported SUSE Linux kernel 3.0.101-0.35.1se	Otherwise, customers may experience unplanned crashes/downtimes when many CPUs and much memory is used. For details: http://service.sap.com/sap/support/notes/1557506
SLES Linux page cache limits	Please follow SAP Note: http://service.sap.com/sap/support/notes/1557506 and VMware vSphere configuration guidelines: http://service.sap.com/sap/support/notes/2161991 to use the optimal configuration.
Install the latest version of VMware Tools	VMware Tools is a suite of utilities that enhance the performance of the virtual machine’s guest operating system and improves management of the virtual machine. For details: http://kb.vmware.com/kb/1014294

LINUX OS	DESCRIPTION
Disable AutoNUMA	Newer Linux kernel (RHEL 7 and SLES 12) supporting auto-migration according to NUMA statistics. For SLES: # yast bootloader, choose "Kernel Parameters" tab (ALT-k) and edit the "Optional Commandline Parameters" section by appending numa_balancing=disabled For RHEL add "kernel.numa_balancing = 0" to /etc/sysctl.d/sap_hana.conf and reconfigure the kernel by running: # sysctl -p /etc/sysctl.d/sap_hana.conf
Disable I/O scheduling	Set kernel parameter "elevator=noop" for Virtual Machines. This disables I/O scheduling. The VMware hypervisor has its own I/O scheduling mechanisms, therefore scheduling I/O inside the Guest Operating System can cause unnecessary overhead. For details, see Red Hat KB http://kbase.redhat.com/faq/docs/DOC-5428 and SUSE KB http://www.novell.com/support/kb/doc.php?id=7009616 (applies also to other SLES versions than the mentioned).
Disable all unnecessary Linux services, such as:	anacron, apmd, atd, autofs, cups, cupsconfig, gpm, isdn, iptables, kudzu, netfs, and portmap, disable SUSE scheduler
Install the following packages:	If you want to runthe SAP HANA Studio on the SAP HANA system, then following packages have to get installed: gtk2 (use version as provided with operating system distribution) java-1_6_0-ibm or later 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, tcsh, libssh2-1, expect, autoyast2-installation. yast2-ncurses (use version as provided with operating system distribution) To run SAP HANA Database (revision 110 or higher) on RHEL 6 or SLES 11, the required compiler runtime libraries (GCC x.y version libraries) need to be installed.
Turn off the Linux kernel dump function	Turn off the Linux kernel dump function (kdump) if not needed for specific reasons.
To optimize large-scale workloads with intensive I/O patterns, change the queue depths of the SCSI default values (KB2053145)	To optimize throughput, set through yast in the boot loader section, or just copy and append these to kernel boot arguments (for example, on Red Hat Enterprise Linux edit /etc/grub.conf): elevator=noop vmw_pvscsi.cmd_per_lun=254 vmw_pvscsi.ring_pages=32 Details: https://doc.opensuse.org/documentation/html/openSUSE_121/opensuse-tuning/cha.tuning.io.html and https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2053145

LINUX OS	DESCRIPTION
Change the following parameters in /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
Check the following shared memory settings if they are set correctly for SAP HANA:	Small: Shmmni value = 4096 if RAM < 64 GB Medium: Shmmni value = 65536 if RAM >= 64 GB Large: Shmmni value = 524288 if RAM >= 256 GB
Install XFS patches	If you encounter performance problems when using XFS as filesystem, consider using the following kernel 3.0.101-0.15.1 from SUSE: http://download.novell.com/Download?buildid=5Vh7boLCyOQ~ Use minimum SLES SP3 to avoid the XFS locking problem.
RHEL 7.2, disable SELinux	RHEL 7.2 is delivered with SELinux enabled by default. Since there is no SELinux policy available for SAP HANA, leaving SELinux enabled can lead to problems when running SAP HANA on RHEL. To disable SELinux set SELINUX=disabled in /etc/sysconfig/selinux and reboot the system.
Configure NTP time server	Linux VMware timing problem. For details: http://service.sap.com/sap/support/notes/989963

Table 21: Virtual Machine Guest Operating System

Monitor and verify an SAP HANA installation

- SAP Note 1698281 provides information on how you can monitor the data growth and utilization of actual memory. With this, it is also possible to detect and diagnose possible memory leaks during operation.
- SAP Note 1969700 covers all the major HANA configuration checks and presents a tabular output with configurations that were changed. The collection of SQL statements are very helpful in checking and identifying parameters that are configured and conflict with the SAP recommended configuration parameters.

Performance Optimization for Production-Level SAP HANA VMs

Optimizing SAP HANA for performance is required when SAP HANA should perform as close bare-metal SAP HANA performance and with the shortest latency in terms of database access times.

In the section, Memory and NUMA Node Considerations, we discussed the impact of local and remote memory, and NUMA node placements. When optimizing SAP HANA for best performance, we recommend sizing an SAP HANA VM with the least number of NUMA nodes. When an SAP HANA VM needs more CPU or RAM than a single NUMA node provides, configure an additional NUMA node and its resources so that SAP HANA has what it needs.

To achieve optimal performance for an SAP HANA virtual machine, use the settings as described in the following pages in addition to the previously described settings. In terms of CPU scheduling and priority, these settings improve performance by reducing the amount of vCPU and vNUMA migration, while increasing the priority of the SAP HANA production virtual machine.

vCPU Affinities _ Pinning Virtual NUMA to Physical NUMA

Setting vCPU affinities is required for production-level VM to ensure the lowest possible CPU latency and optimal use of CPU caches when computing SAP HANA processes, and is important for wide SAP HANA VMs that span more than one NUMA node to avoid possible virtual NUMA node migrations to different physical NUMA nodes. Setting affinities is not necessary for single NUMA-sized SAP HANA VMs becausae the VMware kernel will keep the SAP HANA VM on the NUMA node on which it was started initially.

The VM configuration parameters halt_in_monitor = "TRUE" and idleLoopSpinBeforeHalt = "TRUE" ensure that the CPU latency is as little as possible. Setting affinities is for this is not required.

If vCPU affinities is used, set affinities for each vCPU to a physical NUMA node using the following configuration settings in the virtual machine's .vmx file. Add these settings directly to the .vmx file via an editor and not via the Web Client because this would result in a .

Below is an example of a 8-socket E7-v3 server system, optimized for use of 128 vCPUs. When more systems are consolidated on this server, then all CPU threads should be assigned. This example does not allocate all CPU threads due to the 128 vCPU limit, therefore only the 1st 16 cores of the in total 18 cores got used in this performance optimizing example.

sched.vcpu0.affinity = "0-15"

sched.vcpu1.affinity = "0-15"

sched.vcpu2.affinity = "0-15"

…

sched.vcpu6.affinity = "0-15"

sched.vcpu7.affinity = "0-15"

sched.vcpu8.affinity = "0-15"

sched.vcpu9.affinity = "0-15"

…

sched.vcpu14.affinity = "0-15"

sched.vcpu15.affinity = "0-15"

```
sched.vcpu16.affinity = "16-31"
sched.vcpu17.affinity = "16-31"
…
sched.vcpu23.affinity = "16-31"
sched.vcpu24.affinity = "16-31"
sched.vcpu25.affinity = "16-31"
…
sched.vcpu29.affinity = "16-31"
sched.vcpu30.affinity = "16-31"
sched.vcpu31.affinity = "16-31"

sched.vcpu32.affinity = "32-47"
sched.vcpu33.affinity = "32-47"
sched.vcpu34.affinity = "32-47"
sched.vcpu35.affinity = "32-47"
…
sched.vcpu39.affinity = "32-47"
…
sched.vcpu40.affinity = "32-47"
sched.vcpu41.affinity = "32-47"
sched.vcpu42.affinity = "32-47"
…
sched.vcpu47.affinity = "32-47"

sched.vcpu48.affinity = "48-63"
sched.vcpu49.affinity = "48-63"
sched.vcpu50.affinity = "48-63"
…
sched.vcpu55.affinity = "48-63"
sched.vcpu56.affinity = "48-63"
sched.vcpu57.affinity = "48-63"
```

```
…
sched.vcpu62.affinity = "48-63"
sched.vcpu63.affinity = "48-63"

sched.vcpu64.affinity = "64-79"
sched.vcpu65.affinity = "64-79"
…
sched.vcpu70.affinity = "64-79"
sched.vcpu71.affinity = "64-79"
…
sched.vcpu77.affinity = "64-79"
sched.vcpu78.affinity = "64-79"
sched.vcpu79.affinity = "64-79"

sched.vcpu80.affinity = "80-95"
sched.vcpu81.affinity = "80-95"
sched.vcpu82.affinity = "80-95"
…
sched.vcpu93.affinity = "80-95"
sched.vcpu94.affinity = "80-95"
sched.vcpu95.affinity = "80-95"

sched.vcpu96.affinity = "96-111"
sched.vcpu97.affinity = "96-111"
sched.vcpu98.affinity = "96-111"
…
sched.vcpu109.affinity = "96-111"
sched.vcpu110.affinity = "96-111"
sched.vcpu111.affinity = "96-111"

sched.vcpu112.affinity = "112-127"
sched.vcpu113.affinity = "112-127"
```

sched.vcpu114.affinity = "112-127"

…

sched.vcpu125.affinity = "112-127"

sched.vcpu126.affinity = "112-127"

sched.vcpu127.affinity = "112-127"

Note: Setting vCPU affinities to physical CPU cores may negatively impact performance if you move a VM with vMotion to another ESXi host that has a workload running on that CPU already. The ESXi kernel will migrate potential other workloads, which is not configured with vCPU affinities to another NUMA node after a while. Until then, the two VMs will share the same CPU resources. To avoid this all VMs of such an environment need to get configured with CPU affinities that are aligned and to not overlap.

Low Latency Setting (vSphere 5.5 Only)

To achieve the lowest possible latency for SAP HANA workloads on vSphere 5.5, you can set the latency sensitivity setting to High as follows:

- Go to the Virtual Machine Settings in vCenter Server for the SAP HANA virtual machine.
- Click the Virtual Machine Options tab.
- Select VM Options.
  - Select High from the Latency Sensitivity menu as shown in Figure 44.

Caution: The low latency setting will lead to longer VM boot times due to memory preallocation and longer vMotion VM switchover times. Longer switchover times may lead to a communication interruption between the application server and the database. Also, you may not be able to use all available CPU resources for usage in VM because CPU resources need to be reserved for the hypervisor. Use the low latency setting only if there are performance issues. We recommend that you avoid this setting because of possible negative impacts on application server disconnects during vMotion actions.

This setting will tune a VM for latency-critical workloads automatically. The Latency Sensitivity feature does the following:

- Gives exclusive access to physical resources
- Bypasses virtualization layers
- Tunes virtualization layers
- Preallocates memory

For details, review the VMware document, "Deploying Extremely Latency-Sensitive Applications in VMware vSphere 5.5."[86]

[86] http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/latency-sensitive-perf-vsphere55-white-paper.pdf

NIC Optimization

Client-facing network cards benefit from disabling interrupt coalescing by reducing database-to-client response times. Unlike back-end networks—such as NFS storage or SAP HANA internode communication networks where network bandwidth is the critical factor—optimizing network latency is important for client-facing network interfaces.

Most NICs provide a mechanism for disabling interrupt coalescing, which is usually available via the ethtool command and/or module parameters.

To determine if disabling the physical NIC interrupt moderation on the VMware ESXi host is needed, enter the following command:

 # esxcli system module parameters set -m ixgbe -p "InterruptThrottleRate=0"

This example applies to the Intel 10GbE driver, called ixgbe. To find the appropriate module parameter for the NIC, first find the driver using this ESXi command:

 # esxcli network nic list

Then find the list of module parameters for the driver being used:

 # esxcli system module parameters list -m <driver>

Note: Disable interrupt coalescing of only the client/application server-facing network cards in which you need lower network latency.

Large Receive Offload (LRO) is another feature of VMXNET 3 that helps deliver high throughput with lower CPU utilization. LRO aggregates multiple received TCP segments into a larger TCP segment before delivering it to the guest TCP stack.

However, for latency-sensitive applications that rely on TCP, the time spent aggregating smaller TCP segments into a larger one adds to the latency. It can affect TCP algorithms such as TCP delayed acknowledgment (ACK), which causes the TCP stack to delay an ACK until the two larger TCP segments are received. This also adds to the end-to-end latency of the application.

Review the VMware KB "Understanding TCP Segmentation Offload (TSO) and Large Receive Offload (LRO) in a VMware environment (KB 2055140)" and "Poor TCP performance might occur in Linux virtual machines with LRO enabled (KB 1027511)" to decide if deactivating LRO would help, and how to do it. LRO is enabled by default.

Note: Deactivating LRO to lower network latencies will increase CPU load.

Virtual SCSI Adapter Optimization

To increase the queue depth of the pvSCSI driver inside the Linux-based guest on which SAP HANA runs, add the following Linux kernel boot option to optimize large-scale workloads with intensive I/O patterns the queue depths change the SCSI default values (for example, on Red Hat Enterprise Linux, edit /etc/grub.conf):

 # vmw_pvscsi.cmd_per_lun=254

 # vmw_pvscsi.ring_pages=32

Note: Please review VMware KB article 2088157 to ensure that the minimum VMware patch level is used to avoid possible virtual machine freezes under heavy I/O load, and KB 2053145 for details how to add these values.

Support Process for SAP HANA on VMware vSphere

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

Open an SAP Support Request Ticket

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

Before opening a VMware support ticket, we recommend opening a support request within the SAP support system when the SAP HANA system runs virtualized with VMware. Please follow support note 1158363[87] before you open a support ticket at SAP. This ensures that SAP HANA VMware specialists will work on the case and, if needed, escalate the issue to VMware product support (when it is a VMware product issue) or to SAP support (when it is an SAP HANA issue).

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

- BC-OP-NT-ESX (Windows on VMware ESX)
- BC-OP-LNX-ESX (Linux on VMware ESX)

SAP HANA VMware vSphere-related issues should be escalated directly via the SAP Solution manager to **BC-OP-LNX-ESX**. Figure 69 shows the support process workflow for VMware related SAP HANA issues.

[87] http://service.sap.com/sap/support/notes/1158363

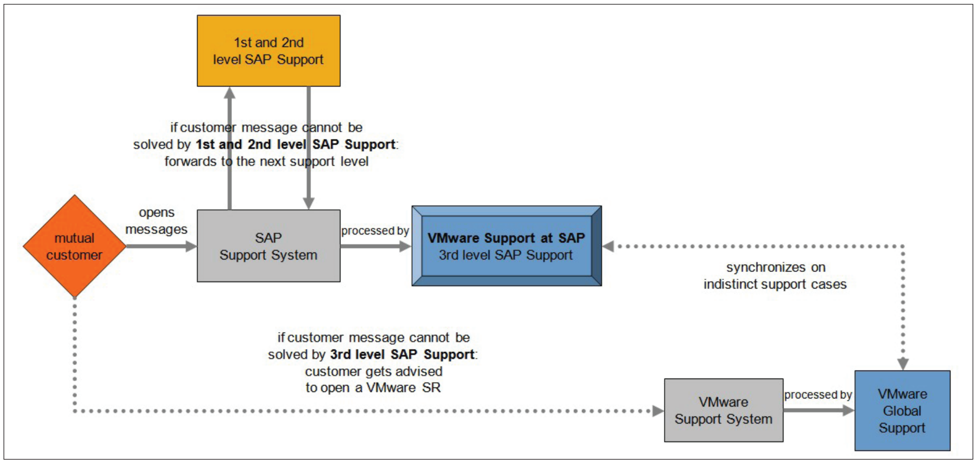

Figure 69: SAP Support Workflow for VMware Related Escalations

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

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

Open a VMware Support Request Ticket

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

In addition:

- Follow the troubleshooting steps outlined in "Troubleshooting ESX/ESXi virtual machine performance issues" (2001003) at http://kb.vmware.com/kb/2001003.
- Run the vm-support utility, and then execute the following command at the service console: vm support-s

This command collects the necessary information that VMware uses to help diagnose issues. It is best to run this command when symptoms first occur.

Summary

A virtualized SAP HANA system can be sized and operated as a natively installed SAP HANA system. Virtualization provides some benefits in terms of operation, such as SAP HANA database live migration with vMotion or strict resource isolation on a virtual server level, and helps to increase the security and standardization of IT processes and operation. Better resource utilization, an easy HA solution via VMware HA, and custom HANA system sizes allow you to operate SAP HANA like the rest of a customer's business critical systems in a fully virtualized environment.

This document concludes with the following quote:

> "I think anything 'software-defined' means it's digital. It means we can automate it, and we can control it, and we can move it much faster."
> —Andrew Henderson, CTO, ING Bank

VMware Acknowledgments

Author:

Erik Rieger, Global Solutions Architect, Global SI/SO & Strategic ISV Technical Alliances, VMware

The following individuals contributed content and/or their knowledge to this guide:

Arne Arnold, Field Architect (SAP)/Staff Engineer, Business Critical Applications, VMware

Bob Goldsand, Staff Partner Architect Global Alliances, VMware

Kevin Christopher, Senior Staff and Tech Lead for Platform Infrastructure Group, VMware

Louis Barton, Staff Engineer Performance Team, SAP Alliances, VMware

Santosh Choudhary Senior Systems Engineer, Business Critical Applications, VMware

Sebastian Lenz, Senior SAP Support Engineer, VMware

Todd Muirhead, Staff Engineer Performance Team, SAP Alliances, VMware

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