



# **SAP On VMware Best Practices**

Version 1.1  
December 2015

© 2015 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.  
3401 Hillview Ave  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

## Contents

1. Overview .....	5
2. Introduction .....	6
2.1 Support.....	6
2.2 SAP .....	6
2.3 vCloud Suite.....	9
2.4 Deploying SAP with vCloud Suite .....	11
3. Virtualization Overview.....	12
3.1 Memory .....	12
3.2 Virtual CPU, NUMA, and Sizing.....	14
3.3 Storage.....	18
3.4 Network.....	21
4. High Availability.....	21
4.1 vSphere High Availability .....	23
4.2 vSphere Fault Tolerance.....	25
4.3 In-Guest Application-Level Failover Solutions .....	27
4.4 Disaster Recovery with Site Recovery Manager .....	31
5. Best Practices .....	32
5.1 General .....	33
5.2 Guest Operating System.....	35
5.3 Memory .....	35
5.4 CPU.....	37
5.5 Storage.....	38
5.6 Network .....	39
5.7 High Availability and Disaster Recovery .....	40
5.8 Sizing and Architecture .....	45
5.9 BusinessObjects .....	47
5.10 SAP HANA .....	47
5.11 SAP ASE.....	48
5.12 SAP on Oracle .....	48
5.13 SAP on MS SQL Server.....	50
5.14 SAP on IBM DB2.....	52
5.15 Operations Management and Automation .....	52
5.16 Backup .....	53

- 6. References ..... 55
  - 6.1 VMware References .....55
  - 6.2 External References .....55
- Appendix A: High Availability Quantitative Analysis ..... 57
- Acknowledgments..... 61

# 1. Overview

This document provides guidelines and best practices for virtualizing SAP applications on VMware vSphere. The content is organized as follows:

- An overview of the VMware vCloud Suite® and how SAP can interact with this stack
- An overview of the following virtualization components:
  - Memory
  - CPU
  - Network
  - Storage
- High availability scenarios and considerations in the virtual environment. Quantitative analysis of high availability is discussed in Appendix A: High Availability Quantitative Analysis, where an example is provided.
- Best practices and guidelines organized into the following categories:
  - General
  - Guest Operating System
  - Memory
  - CPU
  - Storage
  - Network
  - High Availability and Disaster Recovery
  - Sizing and Architecture
  - Business Objects
  - SAP HANA
  - SAP ASE
  - SAP on Oracle
  - SAP on SQL Server
  - SAP on IBM DB2
  - Operations Management and Automation
  - Backup

Some of the topics, such as SAP HANA, have separate white papers dedicated to them. Consult these documents as necessary because their content is not repeated here.

For more details, refer to the latest versions of any mentioned VMware Knowledge Base articles and SAP Notes.

## 2. Introduction

### 2.1 Support

The current support status of SAP applications on VMware is as follows:

- All SAP solutions and applications, except for HANA databases, are supported on VMware vSphere® 6. Details are available in the SAP Community Network:
  - *SAP on VMware vSphere* (<http://scn.sap.com/docs/DOC-27384>)
  - *SAP HANA virtualized – Overview* (<http://scn.sap.com/docs/DOC-60329>)

**Note** Third-party Web sites are not under the control of VMware, and the content available at those sites might change.

- HANA support status is summarized in SAP Note 1788665 – *SAP HANA Support for VMware vSphere Environment* (<http://service.sap.com/sap/support/notes/1788665>). Details include the following:
  - General support for SAP HANA on vSphere 5.1 in non-production environments.
  - General support for single SAP HANA virtual machines on vSphere 5.5 in production and non-production environments.
  - Controlled availability for multi-virtual machine and Business Warehouse (BW) on HANA scale-out scenarios on vSphere 5.5 in production environments.
  - Non-production support for HANA on vSphere 6.0.

### 2.2 SAP

SAP creates enterprise software to manage business operations and customer relations. The company's best-known software products are its Business Suite of solutions (such as SRM, CRM, SCM, PLM, and ERP), its Enterprise Data Warehouse solution (SAP Business Warehouse - SAP BW), SAP BusinessObjects, and the Sybase and HANA data platform. For further information on SAP solutions, see <http://go.sap.com/solution.html#productcategories>.

SAP NetWeaver is the technical foundation for many SAP applications. It is a solution stack of SAP's technology products. The SAP Web Application Server is the runtime environment for SAP applications on which all of the Business Suite solutions run.

Production support for SAP NetWeaver and the Business Suite of SAP solutions on vSphere has existed since 2008, and this support now includes Business Objects and the Sybase and HANA database platform.

#### 2.2.1 SAP Landscape Virtualization Management

SAP has a solution to manage SAP landscapes called SAP Landscape Virtualization Management (LVM). SAP LVM is a central management point for SAP Basis administrators that allows mass operations, automation of day-to-day administrative and lifecycle management tasks, and automation of the copy, clone, and refresh of SAP systems. VMware has an adapter that integrates with SAP LVM to enable these tasks for a virtualized SAP system. For more information on the VMware Adapter for SAP LVM, see <http://www.vmware.com/products/adapter-sap-lvm.html>.

## 2.2.2 Certified Benchmarks and Sizing

There are certified, virtualized SAP benchmarks that validate the performance of SAP workloads on VMware. These benchmarks are available for the OLTP and the Business Warehouse OLAP workloads. The following table lists some of these benchmarks, and up-to-date benchmarks can be found on the *SAP Standard Application Benchmarks* page (<http://sap.com/benchmark>).

**Table 1. SAP Benchmarks on VMware**

SAP Benchmark	Result
Two-Tier SD Benchmark Cert 2015054 <ul style="list-style-type: none"> <li>vSphere 6.0</li> <li>72 vCPU virtual machine on 36 core, 72 thread server</li> </ul> <a href="http://www.sap.com/solutions/benchmark/sd2tier.epx">http://www.sap.com/solutions/benchmark/sd2tier.epx</a>	82,680 SAPS
Two-Tier SD Benchmark Cert 2014043 <ul style="list-style-type: none"> <li>vSphere 5.5</li> <li>48 vCPU virtual machine on 24 core, 48 thread server</li> </ul>	51,400 SAPS
Three-Tier SD Benchmark Cert 2011044 <ul style="list-style-type: none"> <li>vSphere 5.0</li> <li>DB – 20 vCPU virtual machine. App - 20x 12 vCPU virtual machines</li> <li>11x 12 core, 24 thread servers</li> </ul> <a href="http://global.sap.com/solutions/benchmark/sd3tier.epx">http://global.sap.com/solutions/benchmark/sd3tier.epx</a>	175,320 SAPS
BW Enhanced Mixed Load (BW-EML) Cert 2014021 <ul style="list-style-type: none"> <li>vSphere 5.5</li> <li>Scale-up HANA DB - 64 vCPU virtual machine on 60 core, 120 thread server</li> </ul> <a href="http://global.sap.com/solutions/benchmark/bweml-results.htm">http://global.sap.com/solutions/benchmark/bweml-results.htm</a>	111,850 ad-hoc navigation steps per hour
BW Enhanced Mixed Load (BW-EML) Cert 2015016 <ul style="list-style-type: none"> <li>vSphere 5.5</li> <li>Scale-out HANA DB - 4x 64 vCPU virtual machines on 4x 40 core, 80 thread servers</li> </ul>	124,570 ad-hoc navigation steps per hour

SAP sizing is conducted using the SAPS (SAP Application Performance Standard) metric. SAPS is a hardware-independent unit of measurement that describes the performance of a system configuration in the SAP environment. It is derived from the Sales and Distribution (SD) benchmark. The SAPS metric is used to size in the virtual environment to make x86 server purchasing and design decisions similar to physical.

### 2.2.3 SAP Architecture

This section summarizes SAP architecture concepts and terminology used in this document.

SAP uses the term *system landscape*, which contains all the SAP systems that have been installed. It can consist of several system groups where SAP systems are linked by transport routes. *Transport routes* refer to the path of code migrations between SAP systems, for example from Development (DEV) to Quality Assurance (QAS) to Production (PRD)

([https://help.sap.com/saphelp\\_nw74/helpdata/en/63/a30a4ac00811d2851c0000e8a57770/content.htm](https://help.sap.com/saphelp_nw74/helpdata/en/63/a30a4ac00811d2851c0000e8a57770/content.htm)).

The architecture of a single SAP system is multi-tier and consists of the following components:

- Application Servers (SAP Web Application Servers) – These are ABAP and/or Java (J2EE) based, depending on the specific SAP product or module. Two types exist:
  - Primary Application Server (PAS) – An application server instance that is installed with SAP Central Services in newer NetWeaver releases and is part of the base installation.
  - Additional Application Servers (AAS) – Applications servers installed as required for horizontal scalability.
- SAP Message Service – The SAP Message Service is used to exchange and regulate messages between SAP instances in a SAP system. It manages functions such as determining which instance a user logs onto during client connect, and scheduling batch jobs on instances configured for batch.
- SAP Enqueue Service – The SAP Enqueue Service manages the locking of business objects at the SAP transaction level. Locks are set in a lock table stored in the shared memory of the host on which the SAP Enqueue Service runs.
- Database Server – SAP supports several databases. The most common databases include Sybase, HANA, Microsoft SQL Server, Oracle, and IBM DB2.

The following SAP services are defined based on the Message and Enqueue Services:

- Central Instance (CI) – Comprised of the Message and Enqueue Services and other SAP work processes that allow the execution of online and batch workloads. In newer NetWeaver releases, the CI is replaced with SAP Central Services and the Primary Application Server.
- SAP Central Services – In newer versions of SAP, the Message and Enqueue Services have been grouped into a standalone service. Separate Central Services exist for ABAP- and Java-based application servers. For ABAP variants, it is called ABAP SAP Central Services (ASCS), and for J2EE variants is called SAP Central Services (SCS).
- Replicated Enqueue Server – This component consists of the standalone Enqueue Server and an Enqueue Replication Server. The Replicated Enqueue Server runs on another host and contains a replica of the lock table (replication table). If the standalone Enqueue Server fails, it must be restarted on the host on which the Enqueue Replication Server is running, because this host contains the replication table in a shared memory segment. The restarted Enqueue Server uses this shared memory segment to generate the new lock table, after which the shared memory segment is deleted.

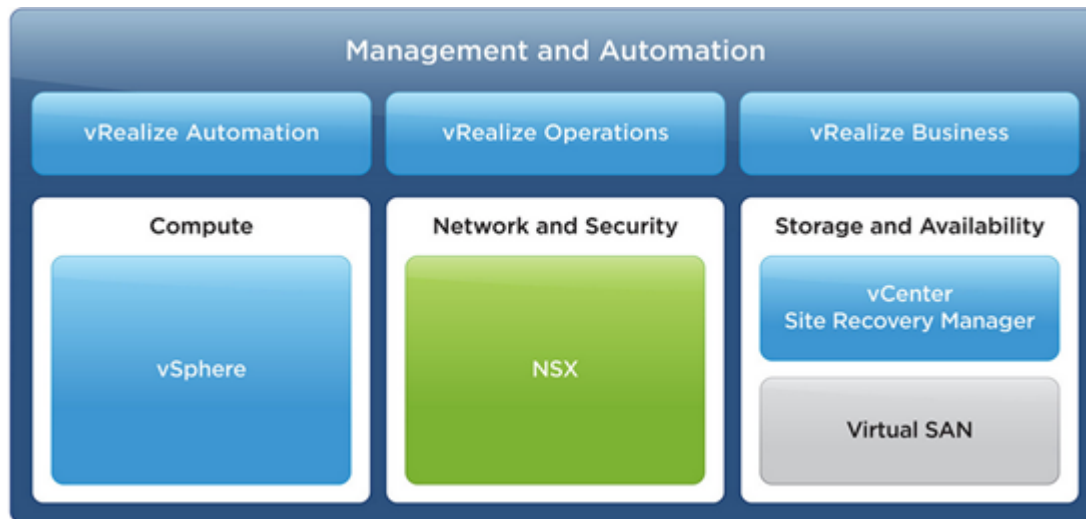
Central Services and the database are both single points of failure, and therefore require considerations for high availability, which are covered in Section 5, Best Practices.



## 2.3 vCloud Suite

vCloud Suite is an integrated set of products that provide infrastructure virtualization, disaster recovery and automation, and cloud management for on-premises vSphere environments. The following figure shows an overview of the vCloud Suite components.

**Figure 1. vCloud Suite Integrated Offering**



vCloud Suite contains the following components:

- Infrastructure Platform:
  - vSphere – Server virtualization platform.
  - VMware Site Recovery Manager™ – Policy-based disaster recovery and testing for all virtualized applications.
- Cloud Management Platform:
  - VMware vRealize® Operations™ – Intelligent performance, capacity, and configuration management for vSphere environments.
  - VMware vRealize Automation™ – Self-service and policy-based infrastructure and application provisioning for vSphere environments.
  - VMware vRealize Business™ – Automated costing, usage metering, and service pricing of virtualized infrastructure for vSphere environments.

Additional offerings available as add-ons to vCloud Suite include the following:

- VMware Virtual SAN™ – Software-defined storage platform that extends vCloud Suite by abstracting and pooling storage to deliver data center virtualization and standardization.
- VMware NSX® – Security and network virtualization that is fully decoupled from hardware. VMware NSX extends vCloud Suite by virtualizing networking to deliver data center virtualization and standardization and security controls native to infrastructure.

The following is a summary of terminology, products, and features of vCloud Suite:

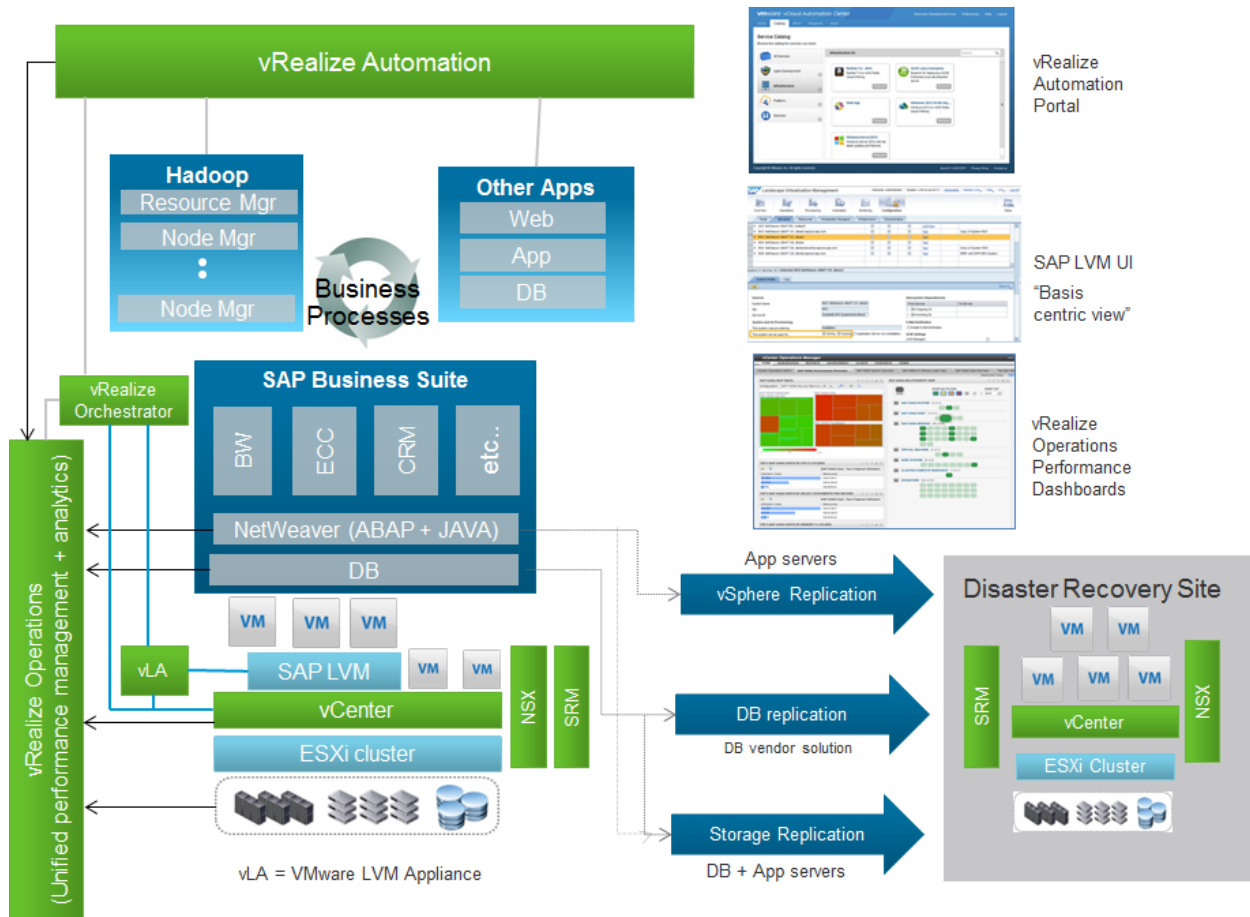
- vSphere – A virtualization platform that allows pooling resources to deliver IT as a service.
- VMware ESXi™ Host – An x86 server running the VMware bare-metal hypervisor ESXi, which allows virtual machines to be run.
- VMkernel – A component of the ESXi hypervisor, which is a POSIX-like operating system developed by VMware. VMkernel is designed specifically to support running multiple virtual machines, and provides core functionality such as resource scheduling, I/O stacks, and device drivers.
- Virtual CPU (vCPU) – VMware uses the terms virtual CPU (vCPU) and physical CPU to distinguish between the processors within the virtual machine and the underlying physical x86-based processors. The number of vCPUs assigned to a virtual machine can be regarded as the number of physical cores of processing power a virtual machine can consume.
- Virtual Machine (VM) – A software implementation of a computer that executes programs like a physical machine. It is configured with virtual CPU, memory, and disk. The disk encapsulates both the operating system and application software. vSphere 6 supports up to 128 vCPUs and 4 TB of RAM per virtual machine.
- Guest Operating System (Guest OS) – The operating system that you install and run in a virtual machine.
- VMware vCenter Server™ – Provides centralized management of virtualized hosts and virtual machines from a single console.
- VMware vSphere Client™ – Provides GUI-based access to vCenter Server to perform management tasks.
- VMware vSphere Web Client – Provides access to a vCenter Server system to manage an ESXi host through a browser.
- VMware vSphere High Availability – Provides easy-to-use, cost-effective high availability solutions for applications running in virtual machines. In the event of server failure, affected virtual machines are automatically restarted on other servers with spare capacity.
- VMware vSphere Fault Tolerance – Provides continuous availability for applications in the event of server failures by creating a live shadow instance of a virtual machine that is always up to date with the primary virtual machine. In the event of a hardware outage, vSphere Fault Tolerance automatically triggers failover, providing zero downtime and preventing data loss. As of vSphere 6, vSphere Fault Tolerance supports up to four virtual CPU virtual machines and is a viable solution for SAP Central Services.
- ESXi Cluster – A group of ESXi hosts defined in vCenter that enables the group to behave as a cluster, so that if one ESXi host fails, all the virtual machines running on the failed ESXi host are restarted on the remaining ESXi hosts in the cluster (using vSphere HA).
- VMware vSphere vMotion® – Enables the live migration of running virtual machines from one physical server to another with zero downtime and continuous service availability.
- VMware vSphere Distributed Resource Scheduler™ (DRS) – Dynamically balances computing capacity across ESXi hosts grouped in an ESXi cluster. This is achieved by live migrating (with vSphere vMotion) virtual machines between ESXi hosts to optimize the utilization of memory and CPU.
- VMware vSphere Replication™ – A hypervisor-based, asynchronous replication solution for vSphere virtual machines. It is fully integrated with vCenter Server and the vSphere Web Client.
- VMware vRealize Orchestrator™ – A development and process automation platform that provides a library of extensible workflows that allow you to create and run automated, configurable processes to manage the vSphere infrastructure and other VMware and third-party technologies. vRealize Orchestrator is installed as a virtual appliance.

- Virtual Disk – Used by a virtual machine to store its OS and application data. A virtual disk is a large physical file, or a set of files, that can be copied, moved, archived, and backed up.
- VMware Tools™ – A suite of utilities installed in the guest operating system that enhances the performance of the virtual machine’s guest operating system and improves management of the virtual machine.

## 2.4 Deploying SAP with vCloud Suite

The following figure depicts an example architecture of the virtualized SAP landscape along with components of vCloud Suite.

Figure 2. SAP and vCloud Suite



The figure shows the following:

- The SAP Business Suite of applications, in addition to other applications, that SAP typically might need to integrate with to deliver complete business processes (for example, integration between HANA and Hadoop to combine structured and unstructured data for data analytics requirements).
- SAP LVM with the VMware Adapter for SAP LVM handles SAP provisioning. This adapter consists of three components:
  - The VMware LVM Appliance (vLA) is a virtual machine that contains all the application code.
  - An adapter that the vLA deploys to SAP LVM. The adapter tells SAP LVM how to communicate with the vLA.

- vRealize Orchestrator workflows that execute the commands SAP LVM sends to the vLA. vRealize Orchestrator executes these instructions on vCenter Server.
- For more information on the VMware Adapter for SAP LVM, see the *User Guide for VMware Adapter for SAP LVM* (<http://www.vmware.com/files/pdf/vmw-adapter-sap-lvm-user-guide.pdf>).
- The automated provisioning of other applications that integrate with SAP is managed by vRealize Automation. For example, with vRealize Automation you can run Hadoop as a service to support any integration requirements with SAP landscapes.
- vRealize Operations management packs are available to extract performance metrics from SAP CCMS, databases, and storage arrays. This enables performance data from multiple layers to be coordinated with vSphere metrics into unified dashboards.
- For database and application virtual machines deployed on the same ESXi cluster, depending on security requirements, VMware NSX micro-segmentation can provide security firewalls between the database and application tier. For more information, see *NSX Micro-segmentation* (<https://www.vmware.com/files/pdf/products/nsx/VMware-Microsegmentation-Solution-Overview.pdf>).
- Site Recovery Manager manages disaster recovery testing and execution. Site Recovery Manager workflows manage the startup order of virtual machines at the disaster recovery site, which helps reduce the recovery time objective (RTO). Three methods are possible to replicate data and virtual machines to the disaster recovery site that address the recovery point objective (RPO):
  - vSphere Replication
  - Database vendor replication solutions
  - Storage array replication

### 3. Virtualization Overview

This section provides background on the following components in the virtual environment:

- Memory
- vCPUs and NUMA
- Storage
- Networking

#### 3.1 Memory

The ESXi hypervisor creates a contiguous addressable memory space for a virtual machine when it runs. This allows the hypervisor to run multiple virtual machines simultaneously while protecting the memory of each virtual machine from being accessed by others. For details on how memory is managed, see *Security of the VMware vSphere Hypervisor* (<http://www.vmware.com/files/pdf/techpaper/vmw-wp-secrty-vsphr-hyprvsr-uslet-101.pdf>).

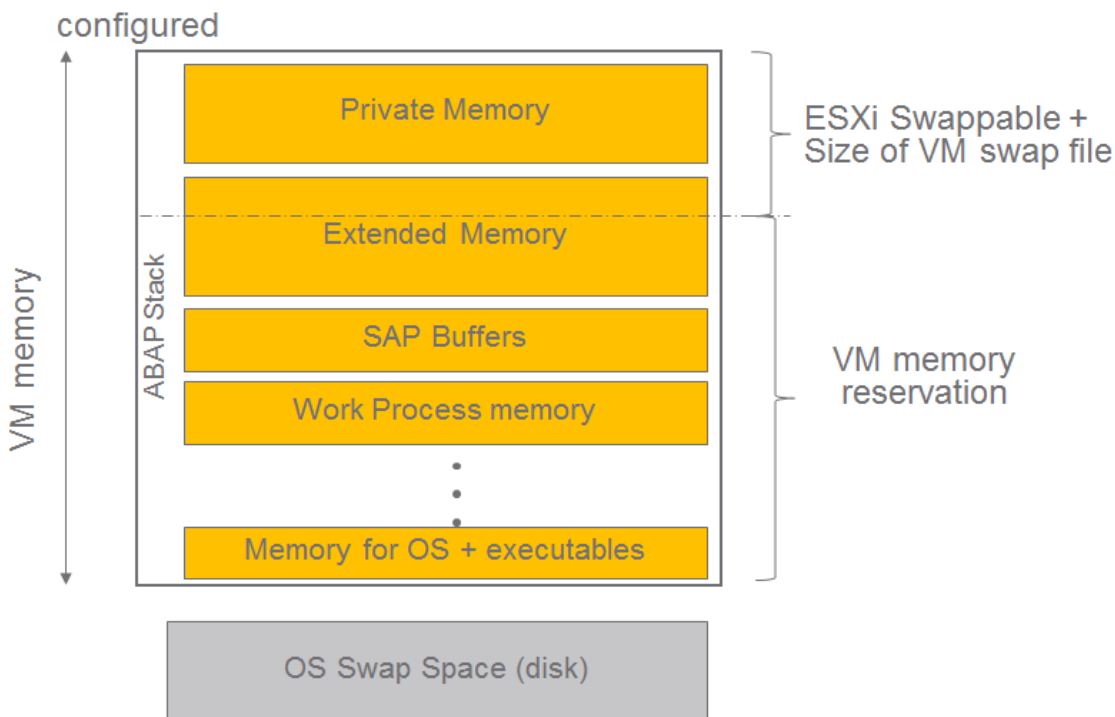
When a virtual machine requires memory, the VMkernel zeros out each memory page before providing them to the virtual machine. Memory isolation is maintained because any attempt by the OS, or any application running inside a virtual machine, to address memory outside of what has been allocated by the hypervisor causes a fault to be delivered to the guest OS. Such a fault typically results in an immediate system crash, panic, or halt in the virtual machine, depending on the OS. Under normal circumstances, the virtual machine has exclusive use of the memory page, and no other virtual machine can touch it or even detect it. The exception is when Transparent Page Sharing is in effect. This is a proprietary ESXi technique to transparently share memory pages between virtual machines, thus eliminating redundant copies of memory pages. This can help with memory overcommit to increase virtual machine consolidation.

Beginning with vSphere 6 (and included in patches or updates to some earlier releases), Transparent Page Sharing is enabled between virtual machines only when those virtual machines have the same salt value. For more information, see VMware KB *Additional Transparent Page Sharing management capabilities and new default settings (2097593)*

([http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2097593](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2097593)). By default, in ESXi 6 salting is enabled (`Mem.ShareForceSalting=2`), and each virtual machine has a different salt value, meaning that Transparent Page Sharing will not happen across the virtual machines by default. This is not an issue for production SAP environments because this provides the highest security between virtual machines. Maximizing virtual machine consolidation through memory overcommit is not a priority, and the best practice for databases is to use large pages that are not shared.

Each virtual machine consumes memory based on its configured size, plus additional overhead memory for virtualization. The configured size is the amount of memory that is presented to the guest OS. The following figure shows the memory settings for a virtual machine running an ABAP-based application server.

**Figure 3. Virtual Machine Memory Settings**



The different memory areas inside the virtual machine correspond to the ABAP stack. Consult SAP documentation for an exact breakdown of the different memory areas of an ABAP-based application server (such as [http://help.sap.com/saphelp\\_nw74/helpdata/en/49/32f2b1e92e3504e1000000a421937/content.htm](http://help.sap.com/saphelp_nw74/helpdata/en/49/32f2b1e92e3504e1000000a421937/content.htm)).

The following settings are shown in the figure:

- Configured memory – Memory size of the virtual machine assigned at creation.
- VM memory reservation – Guaranteed lower bound on the amount of memory that the host reserves for the virtual machine, which cannot be reclaimed by ESXi for other virtual machines.
- ESXi swappable – Virtual machine memory that can be reclaimed by the balloon driver (hypervisor memory reclamation technique) or, in the worst case, by ESXi swapping. This is the automatic size of

the per-virtual-machine swap file that is created on the VMware file system (.vswp file). ESXi swapping can cause severe performance degradation and should be avoided.

- OS swap space – The guest OS requires its own swap file configured according to SAP guidelines, and it resides in the virtual disk. This will function in the same manner as in physical environments.
- For production SAP, it is recommended to set the reservation to the configured size. The size of the ESXi swap file is zero. The virtual machine will only start on or be live migrated to another ESXi host if there is enough free memory equal to the reservation plus its overhead.

You can allocate virtual machines on a single ESXi host based on the following formula:

$$\text{Memory Available for SAP Virtual Machines} = [\text{Total ESXi Server Physical Memory}] - [\text{Memory Required by ESXi}]$$

Memory required by an ESXi host is comprised of memory required by VMkernel plus memory required for each virtual machine (which depends on the size of the virtual machine). The *vSphere Resource Management* guide (<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-resource-management-guide.pdf>) provides more details about this overhead memory.

The memory guidelines are purposely conservative to avoid ESXi kernel swapping. This is important due to the mission-critical nature of SAP business processes.

## 3.2 Virtual CPU, NUMA, and Sizing

For more background on the concepts in this section, see *The CPU Scheduler in VMware vSphere 5.1* (<http://www.vmware.com/files/pdf/techpaper/VMware-vSphere-CPU-Sched-Perf.pdf>).

### 3.2.1 Virtual CPU

VMware uses the terms virtual CPU (vCPU) and physical CPU to distinguish between the processors within virtual machines and the underlying physical x86-based processors. Virtual machines with more than one vCPU are also called SMP (symmetric multiprocessing) virtual machines.

VMware Virtual Symmetric Multi-Processing (Virtual SMP) enhances virtual machine performance by enabling a single virtual machine to use multiple physical processors simultaneously. The biggest advantage of an SMP system is the ability to use multiple processors to execute multiple tasks concurrently, thereby increasing throughput (for example, the number of transactions per second). Only workloads that support parallelization (including multiple processes or multiple threads that can run in parallel) can really benefit from SMP. The SAP architecture (application and database tiers) includes multiple processes that can take advantage of multiple threads, making it a good candidate to take advantage of Virtual SMP. This is demonstrated in the certified SAP benchmarks mentioned in Section 2.2.2, Certified Benchmarks and Sizing.

### 3.2.2 Hyper-Threading

Hyper-threading (also called simultaneous multithreading, or SMT) allows a single physical processor core to behave like two logical processors, essentially allowing two independent threads to run simultaneously. Unlike having twice as many processor cores that can roughly double performance, hyper-threading can provide anywhere from a slight to a significant increase in system performance by keeping the processor pipeline busier.

A completely idle processor, which has both hardware threads idle, provides more CPU resources than only one idle hardware thread with a busy sibling thread. Therefore, the ESXi CPU scheduler prefers to schedule a vCPU of a virtual machine on an idle processor when possible.

The performance impact of hyper-threading for a SAP OLTP workload was derived from SAP OLTP testing, documented in the VMware blog post *SAP Three-Tier Shows Excellent Scaling on vSphere* (<http://blogs.vmware.com/performance/2010/03/sap-threetier-shows-excellent-scaling-on-vsphere.html>). These results show that hyper-threading increased OLTP throughput by approximately 24 percent. Note that a different workload might yield different results.

### 3.2.3 NUMA

In a NUMA (non-uniform memory access) system, multiple NUMA nodes consist of a set of processors and the memory. The access to memory in the same node is local and the access to other nodes is remote. The remote access requires more cycles because it involves a multihop operation. Due to this asymmetric access latency, keeping the memory access local or maximizing the memory locality improves performance. The NUMA load balancer in ESXi assigns a home node to a virtual machine so that all vCPUs of the virtual machine are scheduled within the home node. For the virtual machine, the memory is allocated from the home node. Because the virtual machine rarely migrates away from the home node, the memory access from the virtual machine is mostly local.

This is feasible with SAP application servers where there is flexibility in the virtual machine size. That is, multiple smaller application server virtual machines can fit within a NUMA node because SAP can scale out horizontally in the application tier. However, depending on sizing requirements, the SAP database virtual machine might need to scale vertically beyond the size of a NUMA node.

Virtual machines with more vCPUs than the number of cores in each physical NUMA node are called *wide virtual machines*. These virtual machines are assigned to two (or more) NUMA nodes and are preferentially allocated memory local to those NUMA nodes. Because vCPUs in these wide virtual machines might sometimes need to access memory outside their own NUMA node, they might experience higher average memory access latencies than virtual machines that fit entirely within a NUMA node. This potential increase in average memory access latencies can be mitigated by appropriately configuring virtual NUMA (vNUMA), a feature available since vSphere 5. vNUMA requires virtual hardware version 8, thus allowing the guest OS and database application to take on part of the memory-locality management task. The benefits depend heavily on the level of NUMA optimization in the guest OS and SAP application, which is similar to running SAP instances (application or database) on large NUMA-based physical servers.

Since vSphere 5, when creating a wide virtual machine, you have the option to specify the number of virtual sockets and the number of cores per virtual socket. The virtual machine should be configured so that it aligns with the physical NUMA boundaries. For further information, refer to *Performance Best Practices for VMware vSphere 6.0* (<http://www.vmware.com/files/pdf/techpaper/VMware-PerfBestPractices-vSphere6-0.pdf>).

### 3.2.4 Sizing vCPUs and SAPS

This section provides a sizing and design example based on the vCPU and NUMA concepts previously described.

SAP sizing requires the SAPS rating of the hardware, which for estimation purposes can be obtained from published certified SAP benchmarks. This example uses certification 2014043 (<http://download.sap.com/download.epd?context=40E2D9D5E00EEF7C5D9D82F720182A6DF74F44A0859F1B8BF003A02830D8027B>) and assumes the deployment will use hardware similar to that used in this benchmark. This is a virtual benchmark on vSphere 5.5 with the following result:

- 51,400 SAPS (at ~100 percent CPU) for 48 vCPUs running on a server with two processors.
- 12 cores per processor and 48 logical CPUs because hyper-threading was enabled. Each processor is referred to as a NUMA node.

**Note** Certification 2014043 is an older benchmark. The SAPS values for vSphere on newer servers with faster processors would be different and higher. Work with the server vendors to utilize the most recent and accurate SAPS ratings.

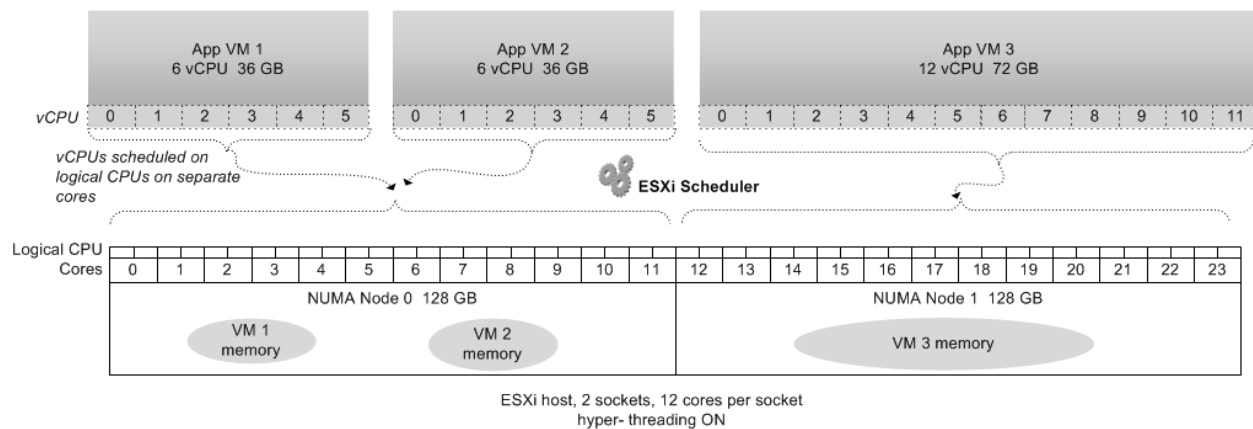


This example is designed for application server virtual machines, which as they scale out horizontally give the flexibility of choosing the number of vCPUs per virtual machine. Typically, one must decide whether the number of vCPUs should equal the number of cores, or if the number of vCPUs should equal the number of logical CPUs. This section provides an example for both. The following will be considered:

- SAP sizing is typically conducted at 60–70 percent CPU, and normal practice is to linearly scale down the benchmark SAPS results. This example does not bother with this, and goes with 51,400 SAPS at 100 percent CPU.
- Size within the NUMA boundaries. Examples for 6- and 12-way virtual machines are provided.
- It is assumed that workloads for all the virtual machines peak at the same time.

The design in the following figure is based on the number of vCPUs equaling the number of cores (no vCPU overcommit).

**Figure 4. Example: Number of vCPUs = Number of Cores, 6-way and 12-way Application Servers**



With all the virtual machines under load simultaneously, the ESXi scheduler by default, with no specific tuning will:

- Allocate a home NUMA node for the memory of each virtual machine.
- Schedule the vCPUs of each virtual machine on its home node, thus maintaining local memory access.
- Schedule each vCPU on a dedicated core to allow exclusive access to core resources.

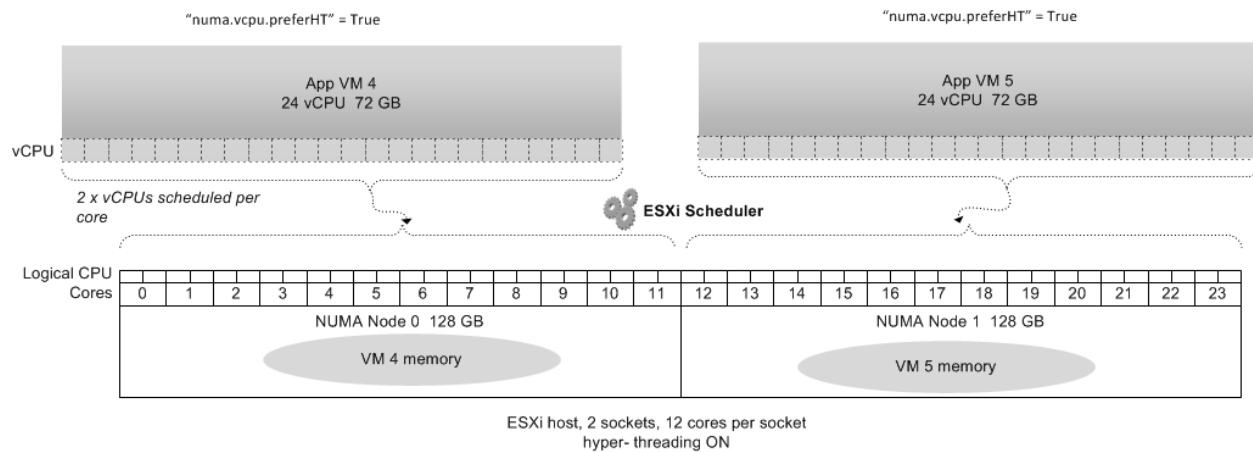
**Note** In physical environments, such NUMA optimizations would require OS commands to localize the processing (for example, the Linux command `numactl`).

However, the above configuration does not provide 51,400 SAPS because not all the logical CPUs are being utilized as in the benchmark. If one assumes the hyper-threading performance boost for a SAP OLTP workload is about 20 percent, so when the number of vCPUs is equal to the number of cores, there should theoretically be approximately  $51,400/1.20 = 42,833$  SAPS. Also, approximately  $42,833/24 = 1,785$  SAPS per vCPU (exclusively scheduled on one core) can be estimated, so the 6-way virtual machine is rated at  $1,785 \times 6 = 10,710$  SAPS, and the 12-way is  $1,785 \times 12 = 21,420$  SAPS (at 100 percent CPU in this example). Are SAPS being wasted by not utilizing all the logical CPUs? Technically yes, but for practical purposes, this is not a major issue because there is some CPU headroom which can be claimed back after go-live, when the virtual machines can be rebalanced based on the actual workload. At this point, vCPU overcommit might be possible because virtual machine workloads might not peak at the same time.



To drive the maximum possible SAPS from this ESXi host, configure the number of vCPUs to equal the number of logical CPUs. The following figure shows an example of this.

**Figure 5. Number of vCPUs = Number of Logical CPUs**



In this design, the virtual machine-level parameter `numa.vcpu.preferHT` must be set to true to override default ESXi scheduling behavior. Default behavior is that ESXi schedules the virtual machine across NUMA nodes when the number of vCPUs for a single virtual machine is greater than the number of cores in the NUMA node. This results in vCPUs of a virtual machine being scheduled on a remote node relative to its memory location. In the example, this is avoided and performance is maximized for the following reasons:

- ESXi schedules all vCPUs of each virtual machine on the same NUMA node that contains the memory of the virtual machine, thus avoiding the penalty of any remote memory access.
- All logical CPUs are being used, thus leveraging the hyper-threading benefit.

**Note** vCPUs are sharing core resources, so the SAPS per vCPU in this case is  $51,400 / 48 = 1071$  SAPS at 100 percent CPU.

This configuration is commonly used in the following situations:

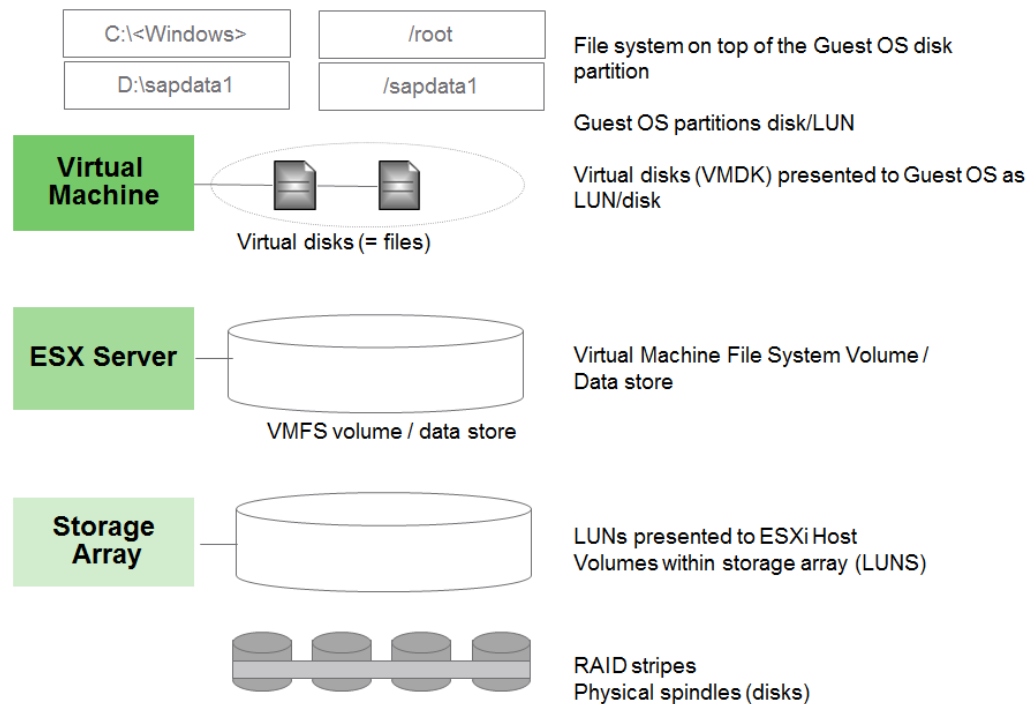
- Running a benchmark to achieve as much performance as possible (as was done for the app server virtual machines in the three-tier vSphere SAP benchmark certification 2011044).
- Conducting physical versus virtual performance comparisons.

For practical purposes, designing so that the number of vCPUs equals the number of logical CPUs might not be so critical. If you were to design for a 24-way app server (as in Figure 5), and actual workloads were less than planned, with lower CPU utilization, there would be plenty of vCPUs without the added gain from hyper-threading.

### 3.3 Storage

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

**Figure 6. Storage Virtualization Layers**



VMware vSphere VMFS (Virtual Machine File System) is a cluster file system that provides storage virtualization optimized for virtual machines. Each virtual machine is encapsulated in a small set of files, and VMFS is the default storage system for these files on physical SCSI disks and partitions. VMware supports Fibre Channel and iSCSI protocols for VMFS.

VMware also supports raw device mapping (RDM). RDM allows a virtual machine to directly access a volume on the physical storage subsystem, and can only be used with Fibre Channel or iSCSI. RDM can be thought of as providing a symbolic link from a VMFS volume to a raw volume. The mapping makes volumes appear as files in a VMFS volume. The mapping file, not the raw volume, is referenced in the virtual machine configuration.

The following table summarizes some of the options and trade-offs when considering between VMFS and RDM.

**Table 2. VMFS Versus RDM**

VMFS	RDM
A volume can host many virtual machines, or can be dedicated to one virtual machine.	Maps a single LUN to one virtual machine so only one virtual machine is possible per LUN.
Increases storage utilization, provides better flexibility, easier administration and management.	More LUNs are required, so it is easier to reach the LUN limit of 256 that can be presented to an ESXi host.
Does not support SQL Server failover cluster for shared disk in a multi-host cluster setup	Might be required to leverage array-level backup and replication tools
	Required for SQL Server failover cluster for shared disk in a multi-host cluster setup

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

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

### 3.3.1 Example Virtual Storage Design

This section applies the virtual storage concepts to a storage design for an example database. Following table shows the virtual SCSI controller assignments.

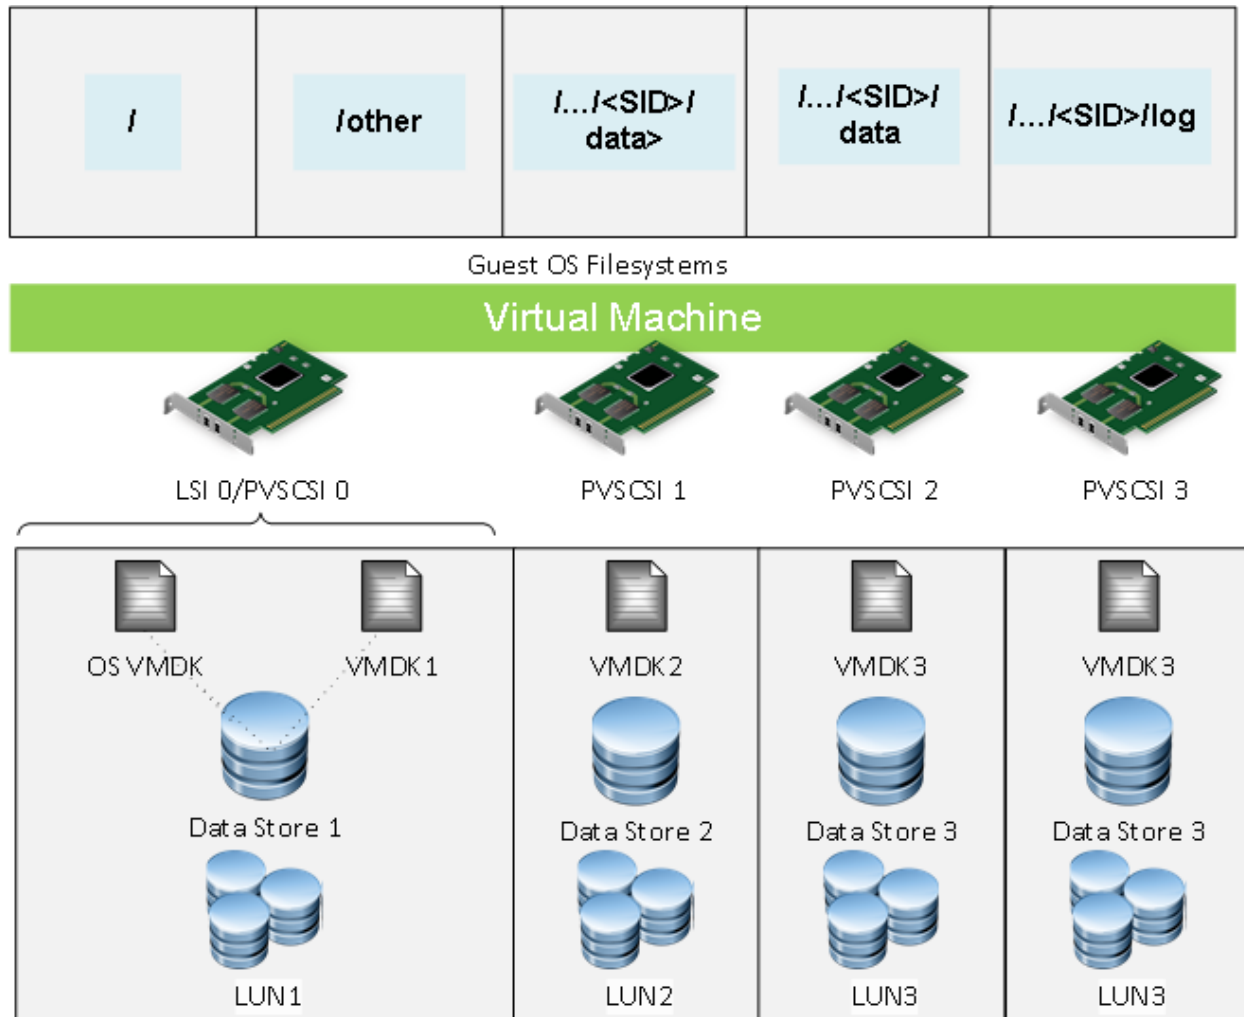
**Table 3. Virtual SCSI Controllers for Example Database**

Virtual SCSI Driver	Virtual Device	File System
LSI Logic/Paravirtual	0:0	/ , other
Paravirtual	1:0	../<SID>/data
Paravirtual	2:0	../<SID>/data
Paravirtual	3:0	../<SID>/log

You can configure the boot drive with the paravirtual driver, but configuration depends on the guest OS version. For details, see *Configuring disks to use VMware Paravirtual SCSI (PVSCSI) adapters (1010298)* (<http://kb.vmware.com/kb/1010398>).

The following figure shows an example storage layout:

**Figure 7. Example Storage Layout for Database**



The storage design implements the following guidelines:

- Spread the virtual disks across all four of the available virtual SCSI controllers to a virtual machine to maximize performance.
- Spread the database files across multiple LUNs and LUN queues to maximize I/O performance.
- Separate database log and data files into different LUNs to minimize I/O contention.

These guidelines are similar to best practices followed in physical environments.

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

- File systems can be created under the data and log directories where all the database data and log files reside, because these can be backed by more virtual disks and data stores.
- Sizing and performance requirements will determine the final design. VMware recommends the sizing be jointly conducted between VMware and database and storage administrators, and incorporate the best practices of the storage vendor. The storage vendors have guidelines for the storage design of databases that should be applied to the virtual environment.

### 3.4 Network

The networking platform links virtual machines to each other within a single host, connects virtual machines to the physical network, joins VMkernel services (such as NFS, iSCSI, and vSphere vMotion) to the physical network, and provides networking for the management interface, which runs management services for vSphere hosts.

The virtual networking stack includes the following:

- Virtual Ethernet Adapters – These are presented to the guest OS by the virtual machine hardware. The guest OS sees virtual adapters as common network interface cards and will use standard drivers. Several special-purpose virtual Ethernet adapters are also available to deliver optimized performance in a virtual machine environment, with corresponding OS drivers provided as a part of VMware Tools.
- Virtual Switches – Allow virtual machines on one vSphere host to communicate with each other using the same protocols as physical switches. The virtual switch emulates a traditional physical Ethernet network switch by forwarding frames at the data-link layer. A vSphere host can have numerous virtual switches, each with more than 1,000 internal virtual ports for virtual machines.
- The VMware vSphere Distributed Switch™ – A single virtual switch that spans multiple associated hosts that simplifies virtual machine networking by letting you set up virtual machine access switching for your entire data center from a centralized interface. This makes it easy to provision, administer, and monitor virtual networking across multiple hosts and clusters.
- Network I/O Control (NIOC) – Manages I/O resource pools to prioritize bandwidth of multiple traffic flows that share the same physical link. Many data centers are shifting to 10 gigabit Ethernet (10 GbE) networking, which provides bandwidth for multiple traffic flows to share the same physical link, and it is necessary to verify that traffic flows can access sufficient bandwidth. NIOC addresses this requirement and is based upon resource pools. You can configure network shares to provide predictable network performance when multiple traffic types contend for the same physical network resources. NIOC is only supported with the vSphere Distributed Switch.

The *VMware vSphere Distributed Switch Best Practices* white paper (<http://www.vmware.com/files/pdf/techpaper/vsphere-distributed-switch-best-practices.pdf>) provides guidelines, considerations, and examples of virtual networking designs. 10 GbE networks work well for mission-critical SAP workloads. This paper provides some example designs based on multiple 10 GbE NICs.

## 4. High Availability

This section describes various high availability scenarios designed to protect the SAP application. The focus is on high availability in the same data center. Disaster recovery is briefly covered at the end of the section. Consideration is given to the SAP single points of failure: Central Services and the database. Maximizing availability involves minimizing the following:

- Unplanned downtime due to unexpected events, such as ESXi host failure. Other outages include failure of network, storage, and power. Designing for these events is not in the scope of this document.
- Planned downtime required to patch the guest OS, SAP, and hypervisor, and update the ESXi host BIOS. This is planned downtime because maintenance windows can be negotiated in advance with the business.

The different solutions to address these factors are summarized in the following figure.

**Figure 8. Overview of High Availability Solutions**

	Reduce Unplanned Downtime	Reduce Planned Downtime
INFRASTRUCTURE	<p><b>vSphere HA</b> – in the event of ESXi host failure VMs are restarted on another host.</p> <p><i>Central Services failover time: VM restart + Guest OS boot + Central Services start</i></p> <p><i>Database failover time: VM restart + Guest OS boot + database crash recovery + database start</i></p>	<p><b>vSphere vMotion</b> – live migrate VMs off ESXi host so hypervisor can be patched and BIOS updated.</p> <p>There is no downtime impact to the application. Need “N+1” type ESXi cluster setups for this to work.</p>
	<p><b>vSphere FT</b> – applicable for Central Services in standalone VM and protects against ESXi host failure.</p> <p>vSphere 6 FT is limited to 4 vCPUs .Majority of virtual SAP databases are sized above 4 vCPU.</p> <p><i>Central Services failover time: ~ zero, no loss of SAP locks.</i></p>	<p><b>vSphere vMotion</b> – live migration of primary and secondary VMs possible off ESXi host so hypervisor can be patched and BIOS updated.</p>
	<p><b>VM Monitoring</b> – protects against Guest OS crash. Will restart VM if VMware Tools heartbeats are not received and the VM isn’t generating any storage or network IO.</p>	N/A
INFRASTRUCTURE + APPLICATION	<p><b>Partner solutions</b> which integrate with VMware HA and provide agents to monitor Central Services and the database (<i>not available out-of-the-box in vSphere</i>)</p> <p>Software failures are detected and automatically restarted. After pre-configured number of failed restarts VMware HA event is invoked restarting the virtual machine.</p>	N/A
APPLICATION	<p><b>Database vendor agent</b> – database vendor specific functionality that monitors health of the database services and in case of failure restarts the database automatically</p>	N/A
	<p><b>3rd party Cluster software</b> - active-passive configuration of two VMs on separate ESXi hosts. VMs connected to shared storage.</p> <p>Cluster resource groups configured for database and Central Services to automate service failover or restart in case of ESXi host or software failure. Standalone + Replicated enqueue possible (<a href="http://help.sap.com/saphelp_nwpi711/helpdata/en/47/e023f3bf423c83e10000000a42189c/frameset.htm">http://help.sap.com/saphelp_nwpi711/helpdata/en/47/e023f3bf423c83e10000000a42189c/frameset.htm</a> )</p> <p><i>Central Services failover time: message server restart + restart standalone enqueue and transfer of SAP locks – no loss of locks</i></p> <p><i>Database failover time: database crash recovery + database start</i></p>	<p>Rolling patch upgrade of the Guest OS is possible; patch the passive VM; failover resources patch the other VM.</p> <p>Depending on the vendor, rolling patch upgrade of the database software is possible. Consult the database vendor documentation for details</p> <p>Rolling kernel upgrade of Central Services is covered in <a href="http://service.sap.com/sap/support/notes/953653">http://service.sap.com/sap/support/notes/953653</a></p>
	<p><b>Database Replication*</b> – database vendors have solutions that enable transactions to be shipped/replayed from a primary to a standby database (in two VMs on separate ESXi hosts). Database storage is duplicated/non-shared.</p> <p>Solution from database vendor or third party required to automate failover. Relatively expensive solution for local HA, typically used in DR scenarios.</p> <p><i>Database failover time: DB IP address failover; make standby DB primary; start DB. (no crash recovery is required as standby database is in consistent state)</i></p>	<p>Depending on the database vendor solution it is possible to perform rolling patch upgrades e.g. upgrade standby database first, perform switchover, patch primary and then switch back. Consult the database vendor documentation for details – e.g. for HANA see <a href="http://service.sap.com/sap/support/notes/1917506">http://service.sap.com/sap/support/notes/1917506</a></p>

**Note** The term *database replication* is used in this document to generally apply to the process of transferring and applying database transaction logs from a primary to a standby database. Database vendors have specific terminology for their respective solutions.

The high availability solutions in the figure are grouped into infrastructure and application. This corresponds to different IT groups who would be responsible for the two areas. VMware infrastructure teams are responsible for the setup of vSphere HA and vSphere FT, and the guest OS, SAP and database administrators would be responsible for the application-level solutions that require cluster or database availability solutions installed inside the virtual machine. Consider the following points:

- Overall availability is a function of unplanned and planned downtime of both layers, infrastructure and application (see Appendix A: High Availability Quantitative Analysis for a quantitative discussion on availability).
- The vSphere HA features in the infrastructure layer are available out-of-the-box with vSphere installations.
- ESXi cluster design involves sizing for spare capacity such that spare resources are available so virtual machines can restart in the event of host failure. This is also required for successful vSphere vMotion operations.
- With cluster software installed in the guest OS, there are two levels of clustering that provide high availability: vSphere HA at the ESXi host level, and cluster resources inside the virtual machine enabling failover of application services, such as database and Central Services.
- Application-level solutions can increase availability:
  - Reduce planned downtime with rolling patch upgrades.
  - Faster failover and reduced RTO.
  - Software monitoring of the database and Central Services.
  - Whether these solutions are needed depends on the business requirements. There is a trade-off because deployment introduces an extra level of configuration and complexity.

The solutions in the figure are covered in more detail in the following sections.

## 4.1 vSphere High Availability

vSphere HA leverages multiple ESXi hosts configured as a cluster and provides protection such as:

- Protecting against a server failure by restarting the virtual machines on other hosts within the cluster.
- Protecting against datastore accessibility failures by restarting affected virtual machines on other hosts that still have access to their datastores.

vCenter Server uses admission control to verify that sufficient resources are available in a cluster to provide failover protection and that virtual machine resource reservations are respected. Admission control imposes constraints on resource usage, and any action that would violate these constraints is not permitted. Examples of actions that can be disallowed include the following:

- Powering on a virtual machine.
- Increasing the CPU or memory reservation of a virtual machine.

By disallowing these actions, admission control sets aside capacity so that when a host failure occurs, vSphere HA uses whatever capacity is available for virtual machine restarts. There are three different policies in which admission control calculates capacity. These are configurable by the VMware administrator and are covered in the *vSphere Availability* guide (<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-availability-guide.pdf>).



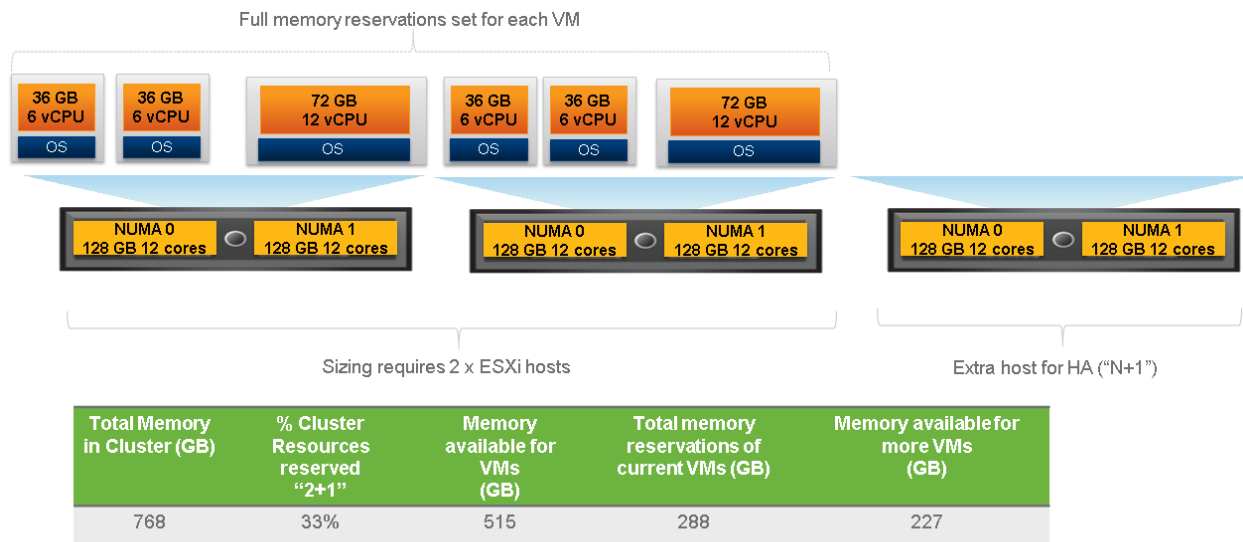
The following table summarizes the policies.

**Table 4. Admission Control Policy Options**

Admission Control Policy	Comments
<p>Percentage of Cluster Resources – Reserve a specific percentage of cluster CPU and memory resources for recovery.</p> <p>Use when reservations vary considerably, and/or there are VMs with large reservations.</p>	<p>Very flexible approach, but manual calculations are required.</p>
<p>Number of Hosts – Calculates slots size based on the largest CPU and memory reservation and determines the number of slots in the cluster.</p> <p>Use when reservations are used sparingly.</p>	<p>Not ideal for SAP VMs that vary in memory size and have full memory reservations set.</p>
<p>Dedicated Failover Hosts – Designate specific hosts as the failover hosts.</p>	<p>Failover hosts are idle and cannot run VMs.</p>

The Percentage of Cluster Resources policy is the most flexible option for SAP virtual machines that vary in size and have large memory reservation settings. The following figure shows an example.

**Figure 9. Percentage of Cluster Resource Admission Control Example**



*Ignoring VMware overhead memory in calculation  
No reservations set for CPU so admission control impacted by memory*

In this example, there is a 2+1 ESXi cluster setup where the initial sizing requires only two ESXi hosts, and a third has been added for vSphere HA failover. The failover host corresponds to one third of the total cluster resources (33%). So, if this is set in vCenter Server under admission control, then 100% - 33% = 67% of the total ESXi cluster memory is available for running virtual machines configured with memory reservations. With six virtual machines running as configured in the figure, ~227 GB is available for more virtual machines (in fact, it will be less than this due to the memory overhead of a virtual machine). The percentage of cluster resources reserved can be increased to provide a more conservative design.



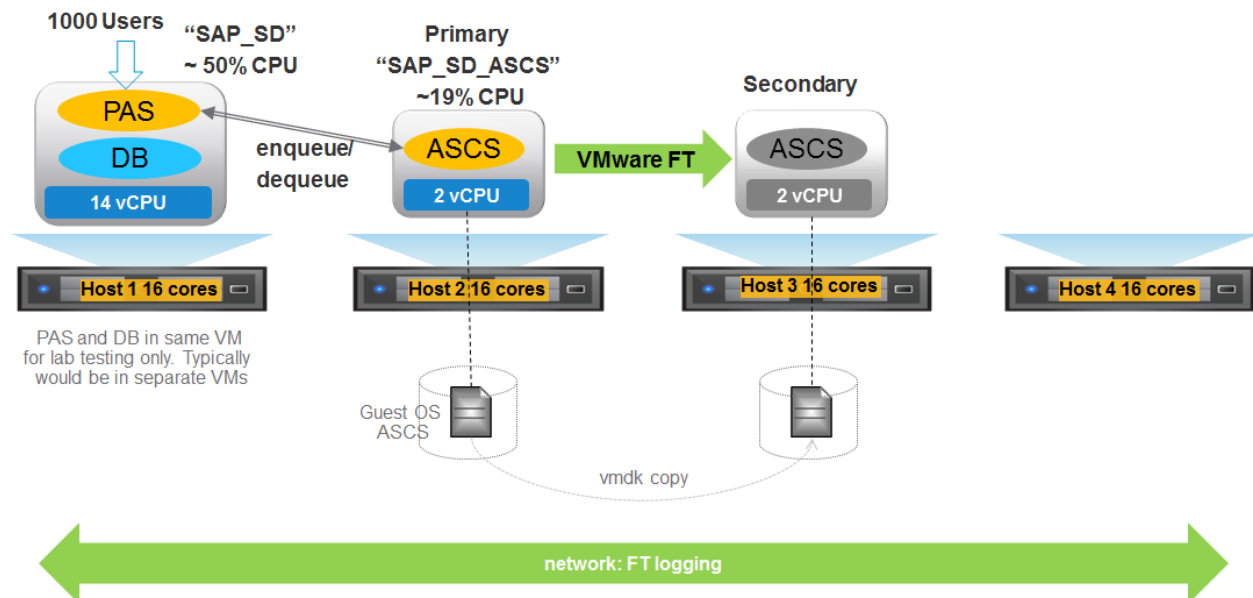
## 4.2 vSphere Fault Tolerance

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

- If the host running the primary VM fails, the secondary VM is immediately activated to replace it with no interruption of service to users. A new secondary VM is started and vSphere FT redundancy is re-established automatically.
- If the host running the secondary VM fails, it is also immediately replaced.
- A fault-tolerant virtual machine and its secondary copy are not allowed to run on the same host.
- The primary VM can be up to 4 vCPUs and 64 GB.
- The secondary VM has its own copy of the primary VM's virtual disks.
- Dedicated 10 GbE NIC is recommended for vSphere FT logging traffic. vSphere FT logging traffic between primary and secondary VMs contains guest network and storage I/O data, as well as the memory contents of the guest operating system.
- For more details on vSphere FT and best practices and limitations see the *vSphere Availability* guide (<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-availability-guide.pdf>) and *Performance Best Practices for VMware vSphere 6.0* (<http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf>).

SAP Central Services is a good candidate for vSphere FT. The following figure describes a lab validation of vSphere FT with Central Services managing an OLTP workload of 1,000 concurrent users executing SAP Sales and Distribution transactions.

**Figure 10. vSphere FT and Central Services Validation**



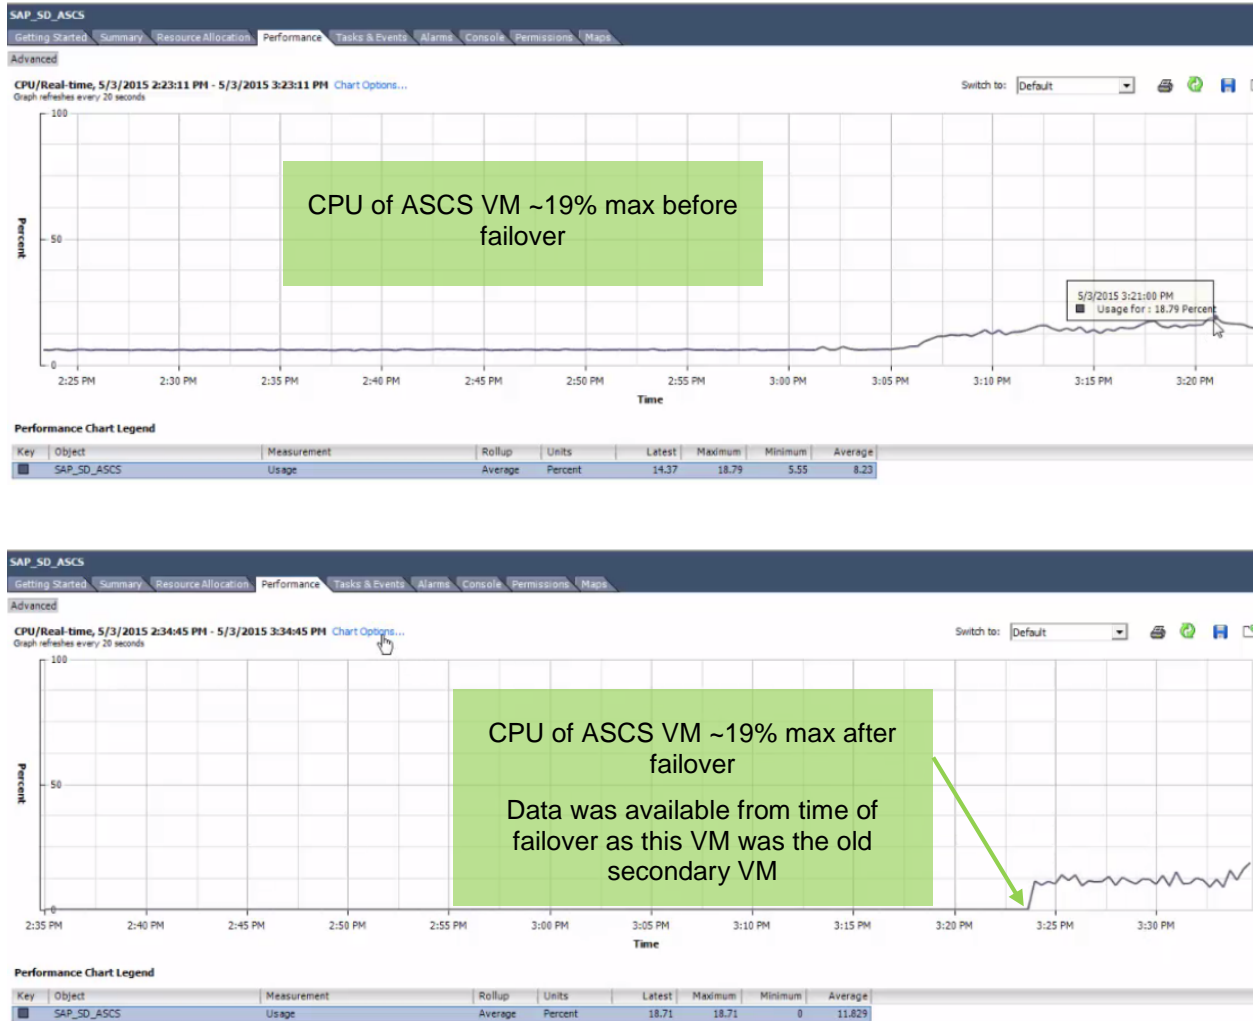
In this validation, the Primary Application Server (PAS) running the ABAP stack and database was installed in the same virtual machine for simplicity, and the ABAP SAP Central Services (ASCS) component was installed in a standalone 2 vCPU, 8 GB virtual machine. No other virtual machines were running, so each virtual machine had full access to host cores based on their vCPU count. SAP lock

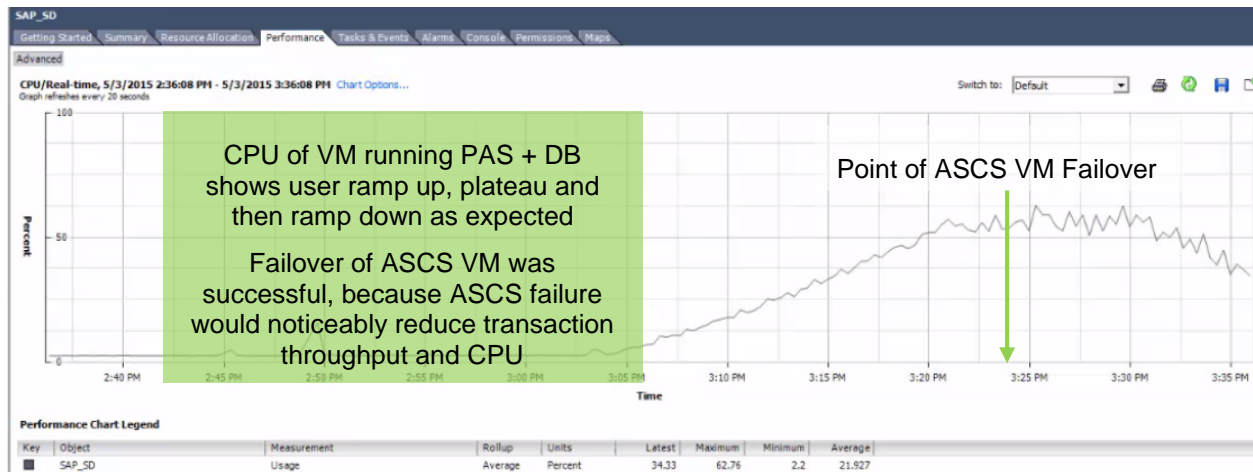
requests and deletes (enqueue and dequeue) generate network traffic between the PAS and ASCS. Network traffic between these virtual machines was measured at ~450 kB per second.

The Test Failover option was used to force a failover (vSphere Client).

The following vCenter Server performance charts show the CPU utilization of the virtual machine running the PAS (on which all the OLTP users are logged on), and the ASCS virtual machine before and after failover.

**Figure 11. VMware FT and Central Services Results**





The CPU utilization of the virtual machine running the PAS confirms that failover of the ASCS virtual machine had no impact. Any loss of Central Services would have caused OLTP transactions to hang and enter a wait status, which would have resulted in a large drop in CPU utilization of the PAS virtual machine. As shown in the chart, this was not the case.

### 4.3 In-Guest Application-Level Failover Solutions

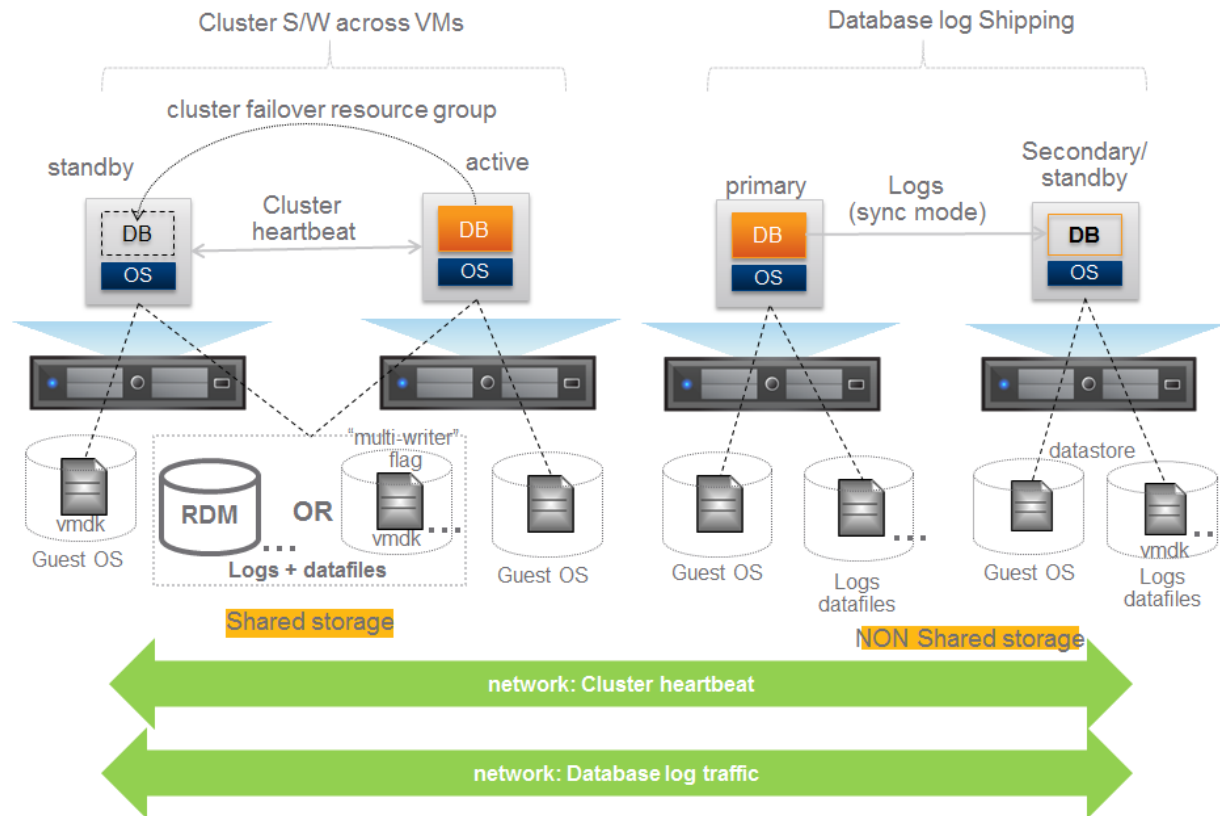
In-guest application-level failover solutions refer to deploying the following in the guest OS:

- Cluster software
- Database replication solutions (common in disaster recovery scenarios, but can also be utilized in a single site)

Third-party clustering software solutions that run on physical hardware are also available to run in the guest OS. This is referred to as in-guest clustering. In the cluster setup, Central Services and the database are installed in cluster resource groups on two virtual machines in active-standby configuration in a similar manner to physical environments. Database replication solutions requires a primary and standby database installed in virtual machines on separate ESXi hosts, and the database vendor technology manages transaction log transfer between the two systems.

The following figure shows the setup of in-guest cluster and database log shipping solutions on an ESXi cluster.

**Figure 12. In-Guest Application Level Failover Solutions**



The in-guest cluster solutions require shared disks attached to the same virtual machine. By default, VMFS disables multiple virtual machines from opening and writing to the same virtual disk (.vmdk file). This prevents more than one virtual machine from inadvertently accessing the same .vmdk file. The multi-writer option allows VMFS backed disks to be shared by multiple virtual machines for cluster-aware applications, where the applications verify that writes originating from two or more different virtual machines do not cause data loss. Details of configuring this option are in the VMware KB <http://kb.vmware.com/kb/1034165>.

Design considerations for the setup in the figure include:

- Separate network for the cluster software interconnect that manages cluster heartbeats.
- Separate network for the database log traffic. In local vSphere HA setups where log transfer is set in synchronized mode, any delay caused by the network can impact performance of the primary VM.
- For active and standby virtual machines in a cluster, and the primary and secondary database virtual machines, ESXi cluster anti-affinity rules are required to keep the virtual machines apart on separate ESXi hosts.
- Database replication solutions are typically used in disaster recovery scenarios where replication is executed across two different sites.

The following table lists some of the vendor failover solutions.

**Table 5. List of Vendor Failover Solutions**

Failover Solution	Windows or Linux	App Monitoring	Enqueue Replication	Comments
VMware HA	Windows Linux	No	n/a	Baseline VMware solution for all VMs
VMware FT	Windows Linux	No	n/a	Up to 4 vCPUs, a good fit for Central Services
<b>In-Guest Cluster Solutions</b>				
↓	Shared storage, active-passive (or active-active) setup. Resource groups failover SAP services in the event of application, OS, and hardware failure			↓
Microsoft Windows Failover Cluster	Windows	Yes	Yes	Requires RDM. vSphere vMotion supported in vSphere 6. <a href="http://kb.vmware.com/kb/1037959">http://kb.vmware.com/kb/1037959</a>
Red Hat Cluster Suite	Linux	Yes	Yes	Option to deploy with VMFS and multi-writer flag. <a href="https://access.redhat.com/articles/29440">https://access.redhat.com/articles/29440</a>
SUSE Linux Enterprise High Availability Extension	Linux	Yes	Yes	Option to deploy with VMFS and multi-writer flag. <a href="https://www.suse.com/partners/alliance-partners/vmware/">https://www.suse.com/partners/alliance-partners/vmware/</a>
Veritas Cluster Services	Windows Linux	Yes	Yes	Option to deploy with VMFS and multi-writer flag. <a href="http://kb.vmware.com/kb/2046035">http://kb.vmware.com/kb/2046035</a>
Oracle Real Application Clusters (RAC)	Windows Linux	Yes	n/a	Active-active setup. Option to deploy with VMFS and multi-writer flag. vSphere vMotion compatible. <a href="http://kb.vmware.com/kb/1034165">http://kb.vmware.com/kb/1034165</a>
NEC ExpressCluster	Windows Linux	Yes	Yes	<a href="http://www.nec.com/en/global/prod/expresscluster/solution/sap/index.html">http://www.nec.com/en/global/prod/expresscluster/solution/sap/index.html</a>
HP Serviceguard	Linux	Yes	Yes	<a href="http://www8.hp.com/h20195/v2/GetPDF.aspx%2F4AA4-2016ENW.pdf">http://www8.hp.com/h20195/v2/GetPDF.aspx%2F4AA4-2016ENW.pdf</a> <a href="http://www8.hp.com/h20195/V2/getpdf.aspx/4AA4-7745ENW.pdf?ver=1.0">http://www8.hp.com/h20195/V2/getpdf.aspx/4AA4-7745ENW.pdf?ver=1.0</a>

Failover Solution	Windows or Linux	App Monitoring	Enqueue Replication	Comments
HANA Host Auto failover	Linux	Yes	n/a	Standby host takes over during failure scenario.
Shared Nothing Architecture				Storage connector handles fencing. Requires validated solution from storage vendor. <a href="http://scn.sap.com/docs/DOC-60470">http://scn.sap.com/docs/DOC-60470</a>
<b>Vendor-Specific Database Replication Solutions</b>				
Transaction log data is transferred from primary to secondary/standby database. Separate storage for secondary database.				
Microsoft AlwaysOn SQL Server	Windows	Yes	n/a	Auto-failover possible in sync mode. <a href="http://blogs.msdn.com/b/saponsqlserver/archive/2012/02/07/sql-server-2012-alwayson-what-is-it.aspx">http://blogs.msdn.com/b/saponsqlserver/archive/2012/02/07/sql-server-2012-alwayson-what-is-it.aspx</a>
Oracle Data Guard	Windows Linux	Yes	n/a	<a href="http://scn.sap.com/docs/DOC-53797">http://scn.sap.com/docs/DOC-53797</a>
VMware Continuent for Replication of Oracle to Oracle	Windows Linux	Yes	n/a	<a href="https://www.vmware.com/files/pdf/products/continuent/VMware-Continuent-Benefits-And-Configurations-Whitepaper.pdf">https://www.vmware.com/files/pdf/products/continuent/VMware-Continuent-Benefits-And-Configurations-Whitepaper.pdf</a>
HANA System Replication	Linux	Yes	n/a	Third-party cluster software required to automate failover or SAP LVM 2.1. <a href="http://scn.sap.com/docs/DOC-47702">http://scn.sap.com/docs/DOC-47702</a>
DB2 High Availability Disaster Recovery (HADR)	Windows Linux	Yes	n/a	<a href="http://scn.sap.com/docs/DOC-60766">http://scn.sap.com/docs/DOC-60766</a>
<b>SAP S/W Monitoring Solutions from VMware Partners</b>				
Agents monitor database and Central Services. Integration with vSphere HA				
Veritas/Symantec ApplicationHA	Windows Linux	Yes	n/a	<a href="https://www.veritas.com/product/business-continuity/applicationha/">https://www.veritas.com/product/business-continuity/applicationha/</a>
SIOS LifeKeeper Single Server Protection	Linux	Yes	n/a	<a href="http://us.sios.com/clustersyourway/products/linux/lifekeeper">http://us.sios.com/clustersyourway/products/linux/lifekeeper</a> <a href="http://docs.us.sios.com/Linux/8.3/LifeKeeperSSPsrc/Output/Finals/Output/ReleaseNotes/index.htm">http://docs.us.sios.com/Linux/8.3/LifeKeeperSSPsrc/Output/Finals/Output/ReleaseNotes/index.htm</a>

## 4.4 Disaster Recovery with Site Recovery Manager

Site Recovery Manager provides disaster recovery protection for virtual environments.

Disaster recovery testing comprises a logistical plan for how an organization will recover and restore partially or completely interrupted critical functions within a predetermined time after a disaster or extended disruption. A disaster recovery plan is only as good as its last successful test. Disaster recovery testing is often difficult because it is usually very disruptive, extremely complex, and expensive in terms of resources. By leveraging virtualization, Site Recovery Manager addresses this problem while making planning and testing simpler to execute.

Two sites are involved when using Site Recovery Manager: a protected site and a recovery site. Site Recovery Manager leverages array-based replication between a protected site and a recovery site to copy virtual machines.

Recovery point objective (RPO) and recovery time objective (RTO) are the two most important performance metrics to keep in mind while designing and executing a disaster recovery plan. RPO is traditionally addressed at the storage layer, where certified storage replication adapters integrate with Site Recovery Manager to enable a fully automated test or real recovery. There are two additional options available for data replication: vSphere Replication (which can also be used to supplement storage replication adapters from the storage vendors) and database vendor-specific replication solutions. These methods are compared in Section 5, Best Practices.

RTO is addressed by the Site Recovery Manager recovery plans that automate the startup sequence of multiple virtual machines that comprise a virtualized SAP landscape and automate network connectivity at the remote site.

## 5. Best Practices

Best practices and guidelines are organized into following sections:

- General
- Guest Operating System
- Memory
- CPU
- Storage
- Network
- High Availability, Disaster Recovery, and Backup
- Sizing and Architecture
- Business Objects
- HANA
- SAP ASE
- SAP on Oracle
- SAP on SQL Server
- SAP on IBM DB2
- Operations Management and Automation
- Backup



## 5.1 General

**Table 6. General Best Practices**

General Best Practices	
1.	<p>Download and follow the guidelines in the following SAP Note:</p> <ul style="list-style-type: none"> <li>2161991 - <i>VMware vSphere configuration guidelines</i> (<a href="http://service.sap.com/sap/support/notes/2161991">http://service.sap.com/sap/support/notes/2161991</a> )</li> </ul> <p>The SAP Note might contain updates, so you should download the latest version.</p>
2.	<p>A list of SAP Notes regarding VMware virtualization is maintained in the SAP Community Network (<a href="http://scn.sap.com/docs/DOC-27321">http://scn.sap.com/docs/DOC-27321</a>).</p>
3.	<p>SAP supports various database vendors. Common vendors on the x86 platform are SQL Server, SAP HANA, Oracle, IBM DB2, and SAP ASE.</p> <p>For database support using SAP applications on VMware, check the support statements in the database-specific notes listed in SAP Note 1492000: <i>General Support Statement for Virtual Environments</i> (<a href="http://service.sap.com/sap/support/notes/1492000">http://service.sap.com/sap/support/notes/1492000</a>).</p>
4.	<p>Together, SAP and VMware provide integrated support. If you have an issue at the VMware level, or if the issue is related to the use of SAP software with VMware features like vSphere vMotion and vSphere Replication, you can log a support call with VMware Support. VMware platform-related support calls should typically be addressed independently of SAP.</p> <p>If your problem is related to the SAP software running in a vSphere environment, you can open a SAP support ticket under the following categories:</p> <ul style="list-style-type: none"> <li>BC-OP-NT-ESX (Windows on VMware ESXi)</li> <li>BC-OP-LNX-ESX (Linux on ESXi)</li> <li>BC-VCM-LVM-VMW (LVM on VMware)</li> </ul> <p>If SAP Support determines that the problem is related to VMware, they will assign the ticket to VMware.</p>
5.	<p>For Windows, apply SAP Note 1409604 – <i>Virtualization on Windows: Enhanced monitoring</i> (<a href="http://service.sap.com/sap/support/notes/1409604">http://service.sap.com/sap/support/notes/1409604</a>). For Linux, apply SAP Note 1606643 – <i>Linux: VMware vSphere host monitoring interface</i> (<a href="http://service.sap.com/sap/support/notes/1606643">http://service.sap.com/sap/support/notes/1606643</a>).</p> <p>Doing so enables some VMware performance counters to be viewed in transaction ST06. The ST06 VMware counters are documented in SAP Note 1482272 - <i>Key Figures of Virtualization on VMware vSphere</i> (<a href="http://service.sap.com/sap/support/notes/1482272">http://service.sap.com/sap/support/notes/1482272</a>).</p> <p>This is required to obtain SAP support, and requires VMware configuration to perform the following operations:</p> <ul style="list-style-type: none"> <li>Activate the host accessory functions on ESXi host (set <code>Misc.GuestLibAllowHostInfo</code>).</li> <li>Activate the accessory functions for the virtual machine (set <code>tools.guestlib.enableHostInfo</code>).</li> </ul>

## General Best Practices

6. Use the latest processor generations from both Intel and AMD, which include hardware features to assist virtualization and improve performance. These features are hardware-assisted CPU virtualization, MMU virtualization, and I/O MMU virtualization.

For more information, see *Performance Best Practices for VMware vSphere 6.0* (<http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf>). In general, review this VMware paper for guidelines to maximize performance on the vSphere platform.

7. BIOS settings for an x86 server can be set to disable unnecessary processes and peripherals to maximize performance. The following table describes some of these settings.

BIOS Feature	Setting	Description
VT-x, AMD-V, EPT, RVI, VT-d , AMD-Vi	Yes	Hardware-based virtualization support.
Node Interleaving	No	Enables NUMA.
Turbo Mode	Yes	Balanced workload over unused cores.
Hyper-Threading	Yes	Hyper-threading is recommended.
Power-Saving	No	Disable if performance is more important than saving power.
Power Management	OS Controlled Mode	Allow ESXi to control CPU power-saving features.
C1E Halt State	No	Disable for performance.
Execute Disable	Yes	Required for vSphere vMotion and vSphere DRS features.

The following settings are not necessary and can be disabled (disabling hardware devices can free interrupt resources): Video BIOS Cacheable, Video BIOS Shadowable, Video RAM Cacheable, On-Board Audio, On-Board Modem, On-Board FireWire, On-Board Serial Ports, On-Board Parallel Ports, and On-Board Game Port.

8. In the vSphere Web Client, select the “High Performance” Power Management policy to maximize performance. (Requires the Power Management BIOS setting to be set to “OS Controlled Mode” as previously indicated).

9. After correctly sizing virtual machines with the memory and vCPUs required for the workload, administration of the SAP application instance within the virtual machine is the same as with physical infrastructure, and the standard SAP administration tasks and procedures apply. Follow parameter settings for SAP application servers as specified in SAP Notes on VMware as per physical deployments.

## 5.2 Guest Operating System

Table 7. Guest Operating System Best Practices

Guest Operating System Best Practices	
1.	Install the latest version of VMware Tools in the guest operating system. VMware Tools is a suite of utilities that enhances the performance of a virtual machine. Although a guest operating system can run without VMware Tools, many VMware features are not available until you install VMware Tools.
2.	To minimize time drift in virtual machines, follow the guidelines in SAP Note 989963 – <i>Linux: VMware timing problem</i> and <i>Timekeeping best practices for Linux guests (1006427)</i> ( <a href="http://kb.vmware.com/kb/1006427">http://kb.vmware.com/kb/1006427</a> ).
3.	From Linux 2.6, set the Linux kernel I/O scheduler to NOOP or Deadline. ESXi uses an asynchronous intelligent I/O scheduler, so virtual guests should experience improved performance by allowing ESX to handle I/O scheduling.  For details, see <i>Linux 2.6 kernel-based virtual machines experience slow disk I/O performance (2011861)</i> ( <a href="http://kb.vmware.com/kb/2011861">http://kb.vmware.com/kb/2011861</a> ).
4.	You can continue to use the Linux OS LVM (Logical Volume Manager) to manage disks for convenience and flexibility.  For example, see <i>Extending a logical volume in a virtual machine running Red Hat or Cent OS (1006371)</i> ( <a href="http://kb.vmware.com/kb/1006371">http://kb.vmware.com/kb/1006371</a> ).

## 5.3 Memory

Table 8. Memory Best Practices

Memory Best Practices	
1.	Set memory reservations equal to the size of the virtual machine for production systems under strict performance SLAs. This avoids memory ballooning, and therefore swapping in the virtual machine, and kernel swapping of the VMware vSphere hypervisor/ESXi host.
2.	Do not disable the balloon driver and memory compression. These are ESXi hypervisor techniques to reclaim virtual machine memory to increase consolidation.
3.	Try to determine the right size for the configured memory of a virtual machine. Use SAP monitoring tools over a reasonable period to determine memory utilization. As memory reservations are used, an over-sized virtual machine wastes memory and reduces server consolidation.

## Memory Best Practices

4. Use large memory pages for databases. Large page support is enabled by default in ESXi versions 3.5 and later and is supported in Linux and Windows. Using large pages can potentially increase TLB access efficiency and improve program performance.

**Note** Large pages can cause memory to be allocated to a virtual machine more quickly, as described in *Use of large pages can cause memory to be fully allocated (1021896)* (<http://kb.vmware.com/kb/1021896>).

5. Follow SAP guidelines for using large pages. Some SAP applications do not support the use of large pages regardless of if those applications are running inside a virtual machine or natively.
  - Not supported for the ABAP stack (see SAP Note 1312995 – *Enabling Large Page Support for the View Memory Model* at <http://service.sap.com/sap/support/notes/1312995>).
  - Supported for NetWeaver Java on Linux (see SAP Note 1681501 – *Configure a SAP JVM to use large pages on Linux* at <http://service.sap.com/sap/support/notes/1681501>).
  - Not supported for Java in Business Objects.

Otherwise, check with SAP Support.

6. In vSphere 5.X, Transparent Page Sharing between virtual machines is generally not enabled by default. Beginning with vSphere 6.0, Transparent Page Sharing is enabled by default within virtual machines (intra-VM sharing), but is only enabled between virtual machines (inter-VM sharing) when those virtual machines have the same salt value, but by default each virtual machine has a different salt value. This provides the highest security between virtual machines, and will only have a small effect in most SAP deployments, because memory overcommit is not a priority. See <http://kb.vmware.com/kb/2097593>.

7. Follow the same SAP Notes as per a physical deployment to configure the size of the OS swap space inside the virtual machine. This is independent of VMware. Any recommendations on OS swap sizing for SAP should be addressed by SAP or SAP Support.

8. If memory reservations are not set, vSphere DRS load-balancing recommendations could be suboptimal for SAP systems with large memory requirements.

Follow the guidelines in *DRS performs unwanted memory load balancing moves or DPM excessively consolidates memory* (<http://kb.vmware.com/kb/2059868>).

## 5.4 CPU

**Table 9. CPU Best Practices**

CPU Best Practices	
1.	<p>Enable hyper-threading if it is available. VMware has seen 20–24 percent improvement in throughput of SAP OLTP workloads with hyper-threading enabled.</p> <p>See <i>SAP Three-Tier Shows Excellent Scaling on VMware</i> (<a href="http://blogs.vmware.com/performance/2010/03/sap-threetier-shows-excellent-scaling-on-vsphere.html">http://blogs.vmware.com/performance/2010/03/sap-threetier-shows-excellent-scaling-on-vsphere.html</a>).</p> <p><b>Note</b> Hyper-threading might provide a lower benefit for CPU-intensive batch jobs compared to OLTP workloads.</p>
2.	<p>Try to size the virtual machine to fit within the NUMA node (non-wide).</p> <p>Wide virtual machine = number of vCPUs of virtual machine &gt; number of cores in a socket/NUMA node</p> <p>The ESXi scheduler has algorithms to keep the virtual machine within a NUMA node to optimize memory access.</p>
3.	<p>Configure vNUMA sockets for wide virtual machines that must cross NUMA nodes, for example, database virtual machines. ESXi exposes NUMA topology to the guest OS, allowing NUMA-aware guest OSes and databases to make the most efficient use of the underlying hardware's NUMA architecture.</p> <p>When creating a virtual machine, you have the option to specify the number of virtual sockets and the number of cores per virtual socket. In general, leave this at the default value of one core per socket with the number of virtual sockets equal to the number of vCPUs. ESXi will automatically configure vNUMA to match the NUMA architecture of the host.</p>
4.	<p>If monitoring the actual workload shows that the SAP application is not benefitting from all the virtual CPUs, the extra vCPUs can impose scheduling constraints, especially under high workload. Therefore, try to reduce the number of vCPUs in the virtual machine.</p>
5.	<p>Generally, setting CPU reservations is not critical. SAP tests on vCPU overcommit show graceful degradation in performance that can be overcome by rebalancing the workloads across an ESXi cluster (using vSphere vMotion). This assumes spare CPU capacity is available in the cluster.</p>
6.	<p>Hot add vCPU disables vNUMA. See <i>NUMA is disabled if VCPU hotplug is enabled</i> (<a href="http://kb.vmware.com/kb/2040375">http://kb.vmware.com/kb/2040375</a>).</p> <p>This is particularly relevant to database virtual machines, which can be NUMA-wide. Generally, databases are sized with vCPUs to handle peak workloads with an additional buffer, so hot add might not be an urgent use case. However, if you need hot add a vCPU for a wide database virtual machine, the loss in vNUMA benefits depends on the workload and NUMA optimization algorithms specific to the database vendor. VMware recommends that you determine the NUMA optimization benefits based on your own workload before setting the hot add vCPU function to determine what the performance trade-off is.</p>

## 5.5 Storage

Table 10. Storage Best Practices

Storage Best Practices	
1.	VMware supports the common storage protocols: iSCSI, NFS, Fibre Channel, and FCoE. Each of these storage protocols can achieve acceptable performance. For a comparison see <i>Comparison of Storage Protocol Performance in VMware vSphere 4</i> ( <a href="http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf">http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf</a> ).
2.	VMware generally recommends VMFS for maximum flexibility. RDMs are typically used when required by backup technologies or third-party clustering software like Microsoft Cluster Server. For background on RDM see <i>Difference between Physical compatibility RDMs and Virtual compatibility RDMs</i> ( <a href="http://kb.vmware.com/kb/2009226">http://kb.vmware.com/kb/2009226</a> ).
3.	VMware has not seen major performance differences for SAP databases on VMFS versus RDM.
4.	Choose storage that supports VMware vSphere vStorage APIs - Array Integration. See <i>Frequently Asked Questions for vStorage APIs for Array Integration</i> ( <a href="http://kb.vmware.com/kb/1021976">http://kb.vmware.com/kb/1021976</a> ) and <i>Abnormal DAVG and KAVG values observed during VAAI operations</i> ( <a href="http://kb.vmware.com/kb/2012288">http://kb.vmware.com/kb/2012288</a> ).
5.	Spread database files across multiple datastores/LUNs. Separate logs from data in separate virtual disks. The creation of dedicated datastores for I/O-intensive databases is analogous to provisioning dedicated LUNs in the physical world. This is a typical design practice for mission-critical enterprise workloads. Validate these recommendations with the database best practices from your storage array vendor.
6.	Use paravirtualized SCSI adapters for database data and log virtual disks. The combination of the paravirtualized SCSI driver ( <code>pvscsi</code> ) and additional ESXi kernel-level storage stack optimizations dramatically improves storage I/O performance. See <i>Configuring disks to use VMware Paravirtual SCSI (PVSCSI) adapters (1010398)</i> ( <a href="http://kb.vmware.com/kb/1010398">http://kb.vmware.com/kb/1010398</a> ).
7.	Spread the database files across all virtual SCSI controllers. This maximizes parallel processing of I/O in the guest OS.
8.	Use eager-zeroed thick disks. These disks have all space allocated and zeroed out at the time of creation. The time to create the disk increases, but it results in the best performance (even on the first write to each block). The use of vSphere vStorage APIs - Array Integration capable SAN storage can speed up eager-zeroed thick disk creation by offloading zeroing operations to the storage array.
9.	Make sure that VMFS is correctly aligned. Like other disk-based file systems, VMFS incurs a performance penalty when the partition is unaligned. Use VMware vCenter™ to create VMFS partitions because it automatically aligns the partitions.

## Storage Best Practices

10. For I/O intensive SAP workloads, you can increase the queue depth to increase performance. See *Large-scale workloads with intensive I/O patterns might require queue depths significantly greater than Paravirtual SCSI default values (2053145)* (<http://kb.vmware.com/kb/2053145>). The queue depth change might not have any impact if the storage array is overloaded or is not correctly configured (the latter has its own queue depths), so work with your storage administrator.

11. When using VMware snapshots, follow guidelines in *Best practices for virtual machine snapshots in the VMware environment (1025279)* (<http://kb.vmware.com/kb/1025279>).

**Note** VMware snapshots are not backups.

12. When deleting a VMware snapshot (for example, during backup operations of a virtual machine running a database in a three-tier setup) there might be a period during which the virtual machine is stunned. See *A snapshot removal can stop a virtual machine for long time (1002836)* (<http://kb.vmware.com/kb/1002836>). The stun can cause application server disconnections, but SAP application servers are configured to auto-reconnect. See SAP Note 98051 – *Database Reconnect: Architecture and function* (<http://service.sap.com/sap/support/notes/98051>).

VMware recommends not performing VMware snapshot delete operations during the period that a batch job is running. Otherwise, the batch job might cancel.

## 5.6 Network

Table 11. Network Best Practices

### Network Best Practices

1. Use the VMXNET family of paravirtualized network adapters. The paravirtualized network adapters in the VMXNET family implement an optimized network interface that passes network traffic between the virtual machine and the physical network interface cards with minimal overhead.

For background, see *Choosing a network adapter for your virtual machine (1001805)* (<http://kb.vmware.com/kb/1001805>).

2. If a physical NIC is shared by multiple consumers (that is, virtual machines and/or the VMkernel), each consumer could adversely impact the performance of others. Thus, for the best network performance, use separate physical NICs for management traffic (vSphere vMotion, vSphere FT logging, and VMkernel) and virtual machine traffic. 10 Gbps NICs have multiqueue support. 10 Gbps NICs can separate these consumers onto different queues. Thus, as long as 10 Gbps NICs have sufficient queues available, this recommendation typically does not apply.

3. Use NIC teaming for availability and load balancing. NIC teams can share the load of traffic among some or all of their members, or provide passive failover in case of a hardware failure or a network outage.

See *NIC teaming in ESXi and ESX* (<http://kb.vmware.com/kb/1004088>).



## Network Best Practices

4. In general, you can follow the networking guidelines in *Performance Best Practices for VMware vSphere 6.0* (<http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf>).

HANA, low latency network, and NIC optimization settings are covered in *Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere* ([http://www.vmware.com/files/pdf/SAP\\_HANA\\_on\\_vmware\\_vSphere\\_best\\_practices\\_guide.pdf](http://www.vmware.com/files/pdf/SAP_HANA_on_vmware_vSphere_best_practices_guide.pdf)).

## 5.7 High Availability and Disaster Recovery

Table 12. High Availability and Disaster Recovery Best Practices

### High Availability and Disaster Recovery Best Practices

1. For vSphere vMotion best practices, see “VMware vMotion Recommendations” section in *Performance Best Practices for VMware vSphere 6.0* (<http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf>).
2. During vSphere vMotion migrations, you might notice a drop in performance of some workloads. This is expected behavior because CPU resources are used to successfully complete the live migration. You can minimize the performance impact by following VMware KB *Virtual machine performance degrades while a vMotion is being performed (2007595)* (<http://kb.vmware.com/kb/2007595> ) Test first in non-production because there is a trade-off. If this KB is applied, certain workloads might take longer to migrate.
3. Central Services can be installed in a standalone virtual machine and protected with vSphere FT which supports up to 4 vCPUs. Follow these guidelines for the virtual machine running Central Services:
  - Use 2 vCPUs for up to 2,000 concurrent users.
  - If more than 2,000 concurrent users, start with 4 vCPUs.

**Note** The above are conservatively extrapolated from a test using a high update workload with a single application server instance. VMware recommends a pre-production performance test to validate the sizing because the utilization of SAP’s Message and Enqueue Servers depends on customer-specific workloads and the number of application server instances.

- Dedicated 10 GbE NIC is recommended for vSphere FT logging traffic.
- For vSphere FT best practices and limitations, see the *vSphere Availability* guide (<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-availability-guide.pdf>) and *Performance Best Practices for VMware vSphere 6.0* (<http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf>).



## High Availability and Disaster Recovery Best Practices

4. When designing for disaster recovery, note that Site Recovery Manager does not currently support the recovery of multi-processor virtual machines that are configured for vSphere FT. Uniprocessor virtual machines are supported with Site Recovery Manager when configured with the legacy method of fault tolerance (vSphere 5.x). For more information, see *Configure vSphere Fault Tolerance on virtual machines that Site Recovery Manager protects* (<http://kb.vmware.com/kb/2109813>).

5. When installing standalone Central Services on Windows in a dedicated virtual machine, follow SAP Note 1678705 – *Installation scenarios for a standalone ASCS instance* (<http://service.sap.com/sap/support/notes/1678705>).

6. vSphere HA and vSphere FT are not application-aware. The following table shows some solutions available to monitor the application health of the SAP single points of failure (SAP Central Services, database) and can be installed in the virtual machine.

App Agent / Monitoring Solution	SAP CS Support	Databases Supported
Veritas/Symantec ApplicationHA	Yes	Various
Database vendor tool	n/a	Check with database vendor
In-guest cluster software – various vendors	Yes	Various

7. In-guest clustering solutions are available on vSphere from third-party vendors. Although these solutions add complexity, they provide additional benefits to complement vSphere HA, such as application-level monitoring and less planned downtime (using rolling patch upgrades of the guest OS and database software).

If you are running third-party clustering solutions on VMware for your SAP deployment and have any installation issues or support questions, follow these guidelines:

- For Windows Failover Cluster and Oracle RAC, contact VMware Global Support Services (GSS).
- For the following, first contact the respective cluster vendor. If the vendor proves that there is a VMware issue, they will contact VMware Support and VMware will work with the vendor and customer to resolve the issue.
  - Red Hat High Availability – *Virtualization Support for RHEL High Availability and Resilient Storage Clusters* (<https://access.redhat.com/articles/29440>)
  - SUSE High Availability – *SUSE Alliance Partners: SUSE + VMware* (<https://www.suse.com/partners/alliance-partners/vmware/>)
  - Veritas Cluster Server – *Symantec Storage Foundation High Availability 6.x in guest storage, high availability, and disaster recovery support (Partner Verified and Supported)* (<http://kb.vmware.com/kb/2046035>)
  - NEC EXPRESSCLUSTER X – *High Availability for SAP by EXPRESSCLUSTER* (<http://www.nec.com/en/global/prod/expresscluster/solution/sap/index.html>)

## High Availability and Disaster Recovery Best Practices

8. The following figure summarizes the high availability options for Central Services and the database.

**Figure 13. High Availability Options for Central Services and the Database**

High Availability Setup		H/W Protect	S/W Protect	Relative Failover Time / RTO	Supports App Planned Downtime	Relative Complexity
NO App Agent in VM	SAP CS VMW HA	Yes	No	Hi	No	+
	SAP CS VMW FT (upto 4 x vCPU)	Yes	No	Lo	No	++
App Agent in VM	SAP CS VMW HA	Yes	Yes	Hi	No	++
	SAP CS VMW FT (upto 4 x vCPU)	Yes	Yes	Lo	No	+++
CS Replicated Enqueue + Cluster S/W		Yes	Yes	Lo	Yes	++++
DB VMW HA - No App Agent		Yes	No	Hi	No	+
DB VMW HA - App Agent		Yes	Yes	Hi	No	++
DB + Cluster S/W		Yes	Yes	Lo-Med	Yes	++++
DB Replication		Yes	Yes	Lo-Med	Yes	++++

App Agent refers to a third-party agent like Veritas/Symantec ApplicationHA or a database vendor-specific monitoring/watchdog tool for the database.

**Note** The column S/W Protect refers to the ability to automatically detect the health of the database or Central Services (for example, check that processes are running) and automatically restart the service in case of a problem. How important this is depends on several factors. For example:

- Some environments running SAP have not experienced any software-related issues with the database or Central Services. That is, failover incidents have been related only to hardware failures.
- Automatic restart is not a requirement because manual intervention is preferred in case of software problems to determine the root cause.

9. Site Recovery Manager can help decrease RTO during disaster recovery scenarios. Site Recovery Manager recovery plans can start up virtual machines of a multitier SAP landscape in the correct sequence as determined by the business and SAP administrator.

10. RPO during disaster recovery site failover is determined by the replication technology. Options include storage array replication, vSphere Replication, and database-specific replication techniques such as log shipping.

---

**High Availability and Disaster Recovery Best Practices**

---

11. Solutions involving storage array or vSphere Replication of databases can result in a crash-consistent copy of the database. This is generally satisfactory for most databases. Work with the database administrator. The process can be supplemented with database-consistent backups.

---

  12. Certified Site Recovery Manager storage array partners are listed in *VMware vCenter Site Recovery Manager Storage Partner Compatibility Matrix SRAs for SRM 5.x SRAs for SRM 4.x* (<https://www.vmware.com/support/srm/srm-storage-partners.html>). For example, instructions to configure Site Recovery Manager with partner storage arrays are in *Setting up VMware vCenter Site Recovery Manager with Partner Storage Arrays* (<http://kb.vmware.com/kb/1014610>).

---

  13. Configure scripts and SAP profiles to auto-start SAP and database instances after an OS restart. This helps with RTO during vSphere HA events and Site Recovery Manager recovery plans in disaster recovery events.

---

  14. As per [kb.vmware.com/kb/2109813](http://kb.vmware.com/kb/2109813), Site Recovery Manager currently supports the recovery of uniprocessor virtual machines that are configured for legacy vSphere FT, not the newer multi-vCPU vSphere FT in vSphere 6.
-

## High Availability and Disaster Recovery Best Practices

15. The following table compares some of the disaster recovery solutions with Site Recovery Manager.

Site Recovery Manager DR Solution	Relative RPO	RTO Considerations	vSphere FT Support
Database replication solution (for example, log shipping) with vSphere Replication for application tier	Low	<ul style="list-style-type: none"> <li>Site Recovery Manager recovery plan for application tier.</li> <li>Database replication is not supported by Site Recovery Manager. Additional scripting is required to handle database-specific recovery steps.</li> <li>Database is application consistent.</li> <li>No consistency between application server VMs and database. This might be required for some SAP modules, such as Portal, which has state on the app server filesystem.</li> </ul>	No
Storage array replication (database and application tier)	Low	<ul style="list-style-type: none"> <li>Site Recovery Manager recovery plans for database and application tier.</li> <li>Database is crash consistent. Check with the storage vendor for a database application-consistent solution.</li> <li>Supports write order fidelity within and across multiple VMs in the same consistency group.</li> <li>Storage vendors have solutions to manage SAP landscape consistency during replication.</li> </ul>	Yes
vSphere Replication (database and application tier)	High Min RPO = 15 mins	<ul style="list-style-type: none"> <li>Site Recovery Manager recovery plans for database and application tier.</li> <li>Supports write order fidelity on the disks/VMDKs that make up a VM. Consistency cannot be guaranteed across multiple VMs.</li> <li>Database is crash consistent, except for SQL Server (Microsoft VSS integration available for application consistency. See VMware KB 2041909 <i>VSS Quiescing Support with vSphere Replication</i>) (<a href="http://kb.vmware.com/kb/2041909">http://kb.vmware.com/kb/2041909</a>)</li> </ul>	No

vSphere Replication does not support vSphere FT or physical RDM-based disks.

See *vSphere Replication FAQ (2005776)* (<http://kb.vmware.com/kb/2005776>).

For a comparison of vSphere Replication and storage array replication, see <https://blogs.vmware.com/vsphere/2015/04/srm-abrvsvr.html>. Storage array replication has fewer constraints.

## High Availability and Disaster Recovery Best Practices

16. If application server VMs have no state that needs to be in sync with the database, the application and database tier can be replicated with separate technologies (for example, vSphere Replication and database replication) or with just vSphere Replication. Some SAP modules based on the Java (J2EE) NetWeaver Application server (for example, Enterprise Portal and Process Integration) persist data to the application server file system that must be time synchronized with corresponding entries in the database. VMware recommends verifying this with the SAP administrator during disaster recovery design. For these situations, storage array replication is recommended because it has consistency groups that can support consistency across virtual machines.
17. When designing stretched cluster solutions that use vSphere vMotion or vSphere HA between two closely located sites, note that the recovery process in case of site failure differs from disaster recovery solutions using Site Recovery Manager. The latter enables the orchestrated restart of different SAP virtual machines in a predefined order. The differences are described in *Stretched Clusters and VMware vCenter Site Recovery Manager* ([http://www.vmware.com/files/pdf/techpaper/Stretched\\_Clusters\\_and\\_VMware\\_vCenter\\_Site\\_Recovery\\_Manage\\_USLTR\\_Regalix.pdf](http://www.vmware.com/files/pdf/techpaper/Stretched_Clusters_and_VMware_vCenter_Site_Recovery_Manage_USLTR_Regalix.pdf)).

## 5.8 Sizing and Architecture

Table 13. Sizing and Architecture Best Practices

### Sizing and Architecture Best Practices

1. As with physical deployments, get the SAPS rating of the desired x86 server from the hardware vendor. The SAPS rating for a virtual machine depends on the core on which the vCPUs are scheduled. The SAPS rating per core, in turn, depends on the server and CPU specifications.  
The hardware vendor has the most current SAPS ratings for their servers. SAP benchmark certifications provide server SAPS capabilities and are available at *SAP Standard Application Benchmarks* (<http://global.sap.com/campaigns/benchmark/index.epx>).
2. Do not reuse any SAPS-per-vCPU values in any VMware sizing documentation or VMworld® presentation (<http://www.vmworld.com>) for your implementation-specific sizing because these values are used only as examples to demonstrate sizing concepts.
3. In general, you can estimate the virtual SAPS rating of a server as the physical SAPS rating minus 10 percent.
4. When sizing, factor memory overhead of the virtual machines and system-wide memory overhead.  
See *vSphere Resource Management* (<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-resource-management-guide.pdf>) for example memory overhead values.
5. For storage sizing based on IOPS requirements, work with the SAP practice of the storage vendor to determine LUN and disk requirements. The latter are dependent on storage array features, such as in-memory cache and vendor-specific tiered storage capabilities.

**Sizing and Architecture Best Practices**

6. Generally, the principles of IOPS sizing for virtual machines are not different from physical deployments. Where multiple virtual machines exist per datastore or LUN, consider the combined IOPS of all the virtual machines.  
  
See *Poor performance and high disk latency with some storage configurations (1031773)* (<http://kb.vmware.com/kb/1031773>).

---

7. Where possible (depending on the SAP product), install the ABAP and Java stacks in separate virtual machines. This is in line with the SAP preferred practice to separate ABAP and Java, as described in *SAPs Dual Stack Strategy* (<http://wiki.sdn.sap.com/wiki/display/SI/SAPs+Dual+Stack+Strategy>).

---

8. Scale-out application servers in virtual machines that fit within a NUMA node (see Section 3.2.3, NUMA).  
  
Database virtual machines can be wide, depending on the sizing requirements.

---

9. In accordance with physical deployments, work with the SAP practice of the x86 vendor to make server purchasing decisions.  
  
Note that some batch jobs are sensitive to single-threaded performance, so in hardware migration or upgrade scenarios, you can compare the SAPS per core between old and new environments. In SAP terminology, this is the Single Computing Unit of performance. See SAP Note 1501701 – *Single Computing Unit Performance and Sizing* (<http://service.sap.com/sap/support/notes/1501701>).

---

10. If the SAP user license covers the database license, you can run application server and database server machines in the same ESXi cluster without any database licensing impact.

---

11. If the database license is obtained separately from the database vendor, depending on the database vendor virtual licensing policies, you can maximize ROI on database licensing costs with a dedicated ESXi cluster for database virtual machines, or you can restrict database virtual machines to dedicated ESXi hosts using DRS anti-affinity policies.

## 5.9 BusinessObjects

Table 14. BusinessObjects Best Practices

BusinessObjects Best Practices	
1.	Follow guidelines in <i>Technical Whitepaper: Best Practices for Virtualizing SAP BusinessObjects BI 4.x on VMware ESXi 5</i> ( <a href="http://scn.sap.com/docs/DOC-42916">http://scn.sap.com/docs/DOC-42916</a> ).
2.	There are no virtualization-specific optimizations of the BI stack beyond what you would already do in a physical environment.
3.	Large pages for Java might not yield any performance benefits. This is based on tests carried out by SAP.  See <i>Evaluating Selected Java Best Practices For SAP BusinessObjects Business Intelligence 4 on vSphere</i> ( <a href="http://scn.sap.com/docs/DOC-29008">http://scn.sap.com/docs/DOC-29008</a> ).
4.	For information on sizing BOBJ see <i>SAP BusinessObjects BI4 Sizing Guide</i> ( <a href="http://scn.sap.com/docs/DOC-33126">http://scn.sap.com/docs/DOC-33126</a> ).

## 5.10 SAP HANA

Table 15. SAP HANA Best Practices

SAP HANA Best Practices	
1.	For SAP HANA support status see SAP Note 1788665 – <i>SAP HANA Support for VMware Virtualized Environments</i> ( <a href="http://service.sap.com/sap/support/notes/1788665">http://service.sap.com/sap/support/notes/1788665</a> ).
2.	SAP supports SAP Business One version for SAP HANA on VMware.  See SAP Note 2020657 – <i>SAP Business One, version for SAP HANA on VMware vSphere</i> ( <a href="http://service.sap.com/sap/support/notes/2020657">http://service.sap.com/sap/support/notes/2020657</a> ).
3.	Virtual HANA guidelines and best practices are covered in the following papers: <ul style="list-style-type: none"> <li>• <i>Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere</i> (<a href="http://www.vmware.com/files/pdf/SAP_HANA_on_vmware_vSphere_best_practices_guide.pdf">http://www.vmware.com/files/pdf/SAP_HANA_on_vmware_vSphere_best_practices_guide.pdf</a>)</li> <li>• <i>Best Practices and Recommendations for Scale-Out Deployments of SAP HANA on VMware vSphere</i> (<a href="http://www.vmware.com/files/pdf/sap-hana-scale-out-deployments-on-vsphere.pdf">http://www.vmware.com/files/pdf/sap-hana-scale-out-deployments-on-vsphere.pdf</a>)</li> </ul>
4.	The vRealize Operations Management Pack for SAP HANA is available from Blue Medora at <a href="http://www.bluemedora.com/products/vrops-management-pack-for-sap-hana/">http://www.bluemedora.com/products/vrops-management-pack-for-sap-hana/</a> .  The management pack includes the adapter and out-of-the-box dashboards to monitor HANA performance metrics.

## 5.11 SAP ASE

Table 16. SAP ASE Best Practices

Sybase Best Practices	
1.	SAP ASE (formerly Sybase) is supported on VMware. See SAP Note 1706801 – SYB: SAP ASE released for virtual systems ( <a href="http://service.sap.com/sap/support/notes/1706801">http://service.sap.com/sap/support/notes/1706801</a> ).
2.	Follow guidelines in SAP Sybase Adaptive Server Enterprise on VMware vSphere Essential Deployment Tips ( <a href="http://www.vmware.com/files/pdf/SAP-Sybase-Adaptive-Server-Enterprise-on-VMware-vSphere.pdf">http://www.vmware.com/files/pdf/SAP-Sybase-Adaptive-Server-Enterprise-on-VMware-vSphere.pdf</a> ).

## 5.12 SAP on Oracle

Table 17. SAP on Oracle Best Practices

SAP on Oracle Best Practices	
1.	Oracle is supported on VMware. See SAP Note 1173954 – Support of Oracle for VMware ( <a href="http://service.sap.com/sap/support/notes/1173954">http://service.sap.com/sap/support/notes/1173954</a> ). For an overview of Oracle support and licensing on VMware, see <i>Understanding Oracle Certification, Support and Licensing for VMware Environments</i> ( <a href="http://www.vmware.com/files/pdf/solutions/oracle/Understanding_Oracle_Certification_Support_Licensing_VMware_environments.pdf">http://www.vmware.com/files/pdf/solutions/oracle/Understanding_Oracle_Certification_Support_Licensing_VMware_environments.pdf</a> ).
2.	If the database license is purchased from Oracle, all cores of ESXi hosts in a cluster running Oracle must be licensed, after which you can run an unlimited number of Oracle virtual machines on those ESXi hosts. Oracle does not support partially licensed hosts. For partially licensed clusters, Oracle has no official stance, so customers should consult with Oracle. A best practice is to have a separate ESXi cluster dedicated to Oracle for licensing and workload segregation.
3.	The VMware <i>Oracle Support Policy</i> is documented at <a href="https://www.vmware.com/support/policies/oracle-support.html">https://www.vmware.com/support/policies/oracle-support.html</a> .  “VMware Support will accept accountability for any Oracle-related issue reported by a customer. By being accountable, VMware Support will drive the issue to resolution regardless of which vendor (VMware, Oracle, or others) is responsible for the resolution.”
4.	See SAP Note 793113 – FAQ: Oracle I/O configuration ( <a href="http://service.sap.com/sap/support/notes/793113">http://service.sap.com/sap/support/notes/793113</a> ) for recommendations concerning maximizing I/O performance, which also apply to virtual environments. Based on these guidelines, create separate VMFS datastores for redo and data files.



**SAP on Oracle Best Practices**

5. For Oracle databases in wide virtual machines, set vNUMA. VMware defers to Oracle NUMA optimizations.  
  
See My Oracle Support Document ID 864633.1 at <https://support.oracle.com> (user credentials required). This support document specifies that NUMA support is disabled by default for Database 11g and later. Oracle recommends testing first before applying to production.

---

6. Oracle Data Guard is supported for SAP. See SAP Note 105047 – *Support for Oracle functions in the SAP environment* (<http://service.sap.com/sap/support/notes/105047>).  
  
A Data Guard standby database can be installed in virtual machines (as per physical) and used with Site Recovery Manager for disaster recovery. You can use vSphere Replication for the application tier and Data Guard for replicating the database (outside of Site Recovery Manager). Custom call-out scripts are required in the Site Recovery Manager recovery plan to activate the standby database on the disaster recovery site.

---

7. To investigate Oracle performance, see SAP Note 618868 – *FAQ: Oracle performance* (<http://service.sap.com/sap/support/notes/618868>). If slow I/O performance is suspected, check the vSphere I/O latency metrics according to *Using esxtop to identify storage performance issues for ESX / ESXi (multiple versions) (1008205)* (<http://kb.vmware.com/kb/1008205>).

---

8. Set huge pages, following SAP Note 1672954 – *Oracle 11g: Usage of hugepages on Linux* (<http://service.sap.com/sap/support/notes/1672954>).

---

9. The Oracle Enterprise Manager vRealize Operations Management Pack is available from VMware partner Blue Medora. A license from Blue Medora is required. The management pack includes the adapter and out-of-the-box dashboards.  
  
Details are available at <http://www.bluededora.com/products/vrops-management-pack-oracle-em/>.

---

10. Oracle RAC is supported by SAP. See SAP Note 527843 – *Oracle RAC support in the SAP environment* (<http://service.sap.com/sap/support/notes/527843>).  
  
RAC configurations are supported on VMware.

---

11. For a consistent online backup of the Oracle database, the following options are available:
  - Use Oracle utility RMAN (this does not require a storage-based snapshot).
  - First, place the Oracle database in backup mode, then take a snapshot (VMware or storage array level).

---

12. For Oracle on Windows, you can use Microsoft VSS services to perform online database backup, which is database consistent.  
  
See [http://docs.oracle.com/cd/B28359\\_01/win.111/b32010/vss.htm](http://docs.oracle.com/cd/B28359_01/win.111/b32010/vss.htm). This process is not supported with Oracle ASM.

### SAP on Oracle Best Practices

13. My Oracle Support Document 604683.1 describes a supported scenario for snapshots when Oracle is not in backup mode, which includes that the Oracle-recommended restore and recovery operations are followed, the database is crash consistent at the point of the snapshot, and write ordering is preserved for each file within a snapshot. Such requirements are typically satisfied by third-party storage array replication solutions, which can be used with Site Recovery Manager for disaster recovery. Check with your replication vendor.

## 5.13 SAP on MS SQL Server

Table 18. SAP on MS SQL Server Best Practices

### SAP on MS SQL Server Best Practices

1. SQL Server is supported on VMware with four options:
  - Microsoft support policy (<http://support.microsoft.com/kb/897615/en-us>)
  - Microsoft Server Virtualization Validation Program (Windows Server 2008 and above) (<http://www.windowsservercatalog.com/svvp.aspx?svvppage=svvpwizard.htm>)
  - Server OEM ([http://www.vmware.com/support/policies/ms\\_support\\_statement.html](http://www.vmware.com/support/policies/ms_support_statement.html))
  - VMware Global Support Services (GSS) and TSANet (<http://www.tsanet.org/>)
2. For SQL Server licensing on VMware, see *Microsoft SQL Server 2012 Virtualization Licensing Guide* ([http://download.microsoft.com/download/C/3/7/C37F243B-0246-493E-ABFC-41A7FFD6DE38/SQL\\_Server\\_2012\\_Virtualization\\_Licensing\\_Guide.pdf](http://download.microsoft.com/download/C/3/7/C37F243B-0246-493E-ABFC-41A7FFD6DE38/SQL_Server_2012_Virtualization_Licensing_Guide.pdf)).
3. For guidelines on Microsoft clustering with vSphere, see *Microsoft Clustering on VMware vSphere: Guidelines for supported configurations (1037959)* (<http://kb.vmware.com/kb/1037959>). Physical RDM and in-guest iSCSI are supported. vSphere vMotion is supported on vSphere 6 (but not prior versions). Follow these guidelines:
  - The vSphere vMotion network must be a 10 GbE link.
  - vSphere vMotion is supported for Windows Server 2008 SP2 and later releases. Windows Server 2003 is not supported.
  - The Microsoft cluster heartbeat timeout must be modified to allow 10 missed heartbeats.
  - The virtual hardware version for the clustered virtual machine must be version 11.
  - For more details, see *Setup for Failover Clustering and Microsoft Cluster Service* (<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-setup-mscs.pdf>).

## SAP on MS SQL Server Best Practices

4. SAP supports databases protected by SQL Server AlwaysOn. See SAP Note 1772688 – *SQL Server AlwaysOn and SAP applications* (<http://service.sap.com/sap/support/notes/1772688>).

AlwaysOn minimizes RTO and RPO in case of failure of the primary SQL Server. For more background, see *Running SAP Applications on the Microsoft Platform* (<http://blogs.msdn.com/b/saponsqlserver/archive/2012/02/07/sql-server-2012-alwayson-what-is-it.aspx>). AlwaysOn configurations can be installed in virtual machines similarly to physical deployments.

5. For wide virtual machines with vNUMA configured, VMware defers to SQL Server NUMA optimizations for extra performance benefits. SQL server process and memory are optimized for NUMA architecture as described in *How SQL Server Supports NUMA* (<http://technet.microsoft.com/en-us/library/ms180954%28v=sql.105%29.aspx>).

6. Place data and log files in separate datastores. This is similar to physical best practices. See *SAP with Microsoft SQL Server 2008 and SQL Server 2005: Best Practices for High Availability, Maximum Performance, and Scalability* (<http://www.sdn.sap.com/irj/scn/go/portal/prtroot/docs/library/uuid/4ab89e84-0d01-0010-cda2-82ddc3548c65?overridelayout=true>).

7. For SQL Server configuration parameters, follow the SAP Notes as per physical:

- 1237682 – *Configuration Parameters for SQL Server 2008*
- 1702408 – *Configuration Parameters for SQL Server 2012*
- 1986775 -- *Configuration Parameters for SQL Server 2014*

These SAP Notes specify the memory requirements for SQL Server, specifically how much extra memory is required beyond SQL Server for the OS. For memory sizing, treat the virtual machine as if it were a physical server and calculate memory allocation within the virtual machine accordingly.

8. To investigate SQL Server I/O performance within the guest OS, see SAP Note 987961 – *FAQ: SQL Server I/O performance* (<http://service.sap.com/sap/support/notes/987961>).

If slow I/O performance is suspected, check the vSphere I/O latency metrics as described in *Using esxtop to identify storage performance issues for ESX / ESXi (multiple versions) (1008205)* (<http://kb.vmware.com/kb/1008205>).

9. SQL Server performance metrics are available in VMware vCenter Operations Management Suite™ using the VMware vRealize Hyperic® (SQL Server Plug-In) and Microsoft SCOM adapters.

10. Backup considerations for SAP SQL Server databases are described in SAP Note 1878886 – *Backup Strategies for SQL Server* (<http://service.sap.com/sap/support/notes/1878886>). Simple recovery mode is not acceptable for SAP databases.

SAN vendor solutions that integrate with Microsoft Volume Shadow Copy Service (for database consistency) are recommended for environments requiring high restore SLAs. Such solutions work on VMware similarly to physical.

## 5.14 SAP on IBM DB2

Table 19. SAP on IBM DB2 Best Practices

SAP on IBM DB2 Best Practices	
1.	<p>IBM supports DB2 on VMware. See SAP Note 1130801 – <i>DB6: Virtualization of IBM DB2 for Linux, UNIX, and Windows</i> (<a href="http://service.sap.com/sap/support/notes/1130801">http://service.sap.com/sap/support/notes/1130801</a>).</p> <p>For an older performance study of SAP on DB2 on VMware, see <i>SAP Performance on vSphere with IBM DB2 and SUSE Linux Enterprise</i> (<a href="http://www.vmware.com/files/pdf/techpaper/vsp_41_perf_SAP_SUSE_DB2.pdf">http://www.vmware.com/files/pdf/techpaper/vsp_41_perf_SAP_SUSE_DB2.pdf</a>).</p>
2.	<p>IBM supports sub-capacity licensing on VMware. For details, see <i>Sub-capacity licensing FAQs</i>, item 13 (<a href="http://www-01.ibm.com/software/passportadvantage/subcapfaqov.html#13">http://www-01.ibm.com/software/passportadvantage/subcapfaqov.html#13</a>).</p> <p>IBM uses the PVU metric for licensing, and for VMware, the IBM policy is “Each vCPU is equal to one processor core for PVU licensing. We license to the lower of the sum of vCPUs or full (physical) capacity of the server”. Given this, here are some licensing examples:</p> <ul style="list-style-type: none"> <li>• 1x 8 vCPU VM running DB2 on 16 core ESXi host, license – 8 cores for PVU licensing.</li> <li>• 3x 8 vCPUs VMs or 1x 32 vCPU VM on 16 core ESXi host (that is, vCPU overcommit), license – 16 cores for PVU licensing.</li> <li>• 1x 16 vCPU VM (running DB2) on an ESXi cluster with 2x 16 cores ESXi hosts (extra host added for vSphere HA failover), license – 16 cores for PVU licensing.</li> </ul>
3.	<p>The vRealize Operations Management Pack for IBM Tivoli Monitoring is available from VMware partner Blue Medora. This enables monitoring of DB2.</p> <p>A license from Blue Medora is required. The management pack includes the adapter and out-of-the-box dashboards.</p> <p>Details are available at <a href="http://www.bluemedora.com/products/vrops-management-pack-for-ibm-tivoli-monitoring/">http://www.bluemedora.com/products/vrops-management-pack-for-ibm-tivoli-monitoring/</a>.</p>

## 5.15 Operations Management and Automation

Table 20. Operations Management and Automation Best Practices

Operations Management and Automation Best Practices	
1.	<p>To investigate performance problems that might require escalation with VMware Support, you can run a VMware tool to track and record virtual performance counters while the workload runs. The performance data can be replayed later for debugging.</p> <p>Follow the procedures in <i>Collecting performance snapshots using vm-support in ESX and ESXi</i> (1967) (<a href="http://kb.vmware.com/kb/1967">http://kb.vmware.com/kb/1967</a>).</p>
2.	<p>To troubleshoot performance in the virtual environment, see <i>Troubleshooting ESXi virtual machine performance issues</i> (2001003) (<a href="http://kb.vmware.com/kb/2001003">http://kb.vmware.com/kb/2001003</a>).</p>

### Operations Management and Automation Best Practices

3. The guest OS monitoring tools can be used to determine memory usage of the SAP instances (database and app instance). If swapping occurs within the guest OS, then as with physical deployments, evaluate increasing the memory of the virtual machine.

vSphere active memory counter provides some indication of VM memory usage, but the in-guest OS memory counters provide a more accurate view of memory consumption by databases and SAP application servers. For more information, see *Understanding vSphere Active Memory* (<https://blogs.vmware.com/vsphere/2013/10/understanding-vsphere-active-memory.html>).

4. The vRealize Operations Management Pack for SAP is available from VMware partner Blue Medora. This enables CCMS data (SAP transaction RZ20) to be available in vRealize Operations. The latest version connects to the SAP Control WebService Interface version 7.10+, through HTTP or HTTPS, and requires a guest OS username and password.

Details are available at <http://www.bluededora.com/products/vrops-management-pack-for-sap/>.

5. To access database-level metrics in VMware vRealize Operations Manager, you can use the database-specific adapters/management packs (see the respective preceding database sections).

6. SAP Landscape Virtualization Management (LVM) is a central management tool to manage and automate operations in a SAP landscape, such as start/stop and copy/clone of SAP systems. Details of LVM integration with vSphere are available at <http://www.vmware.com/products/adapter-sap-lvm.html>.

7. SAP has its own tools for managing and transporting code changes between different SAP systems (SAP Correction and Transport and BusinessObjects Lifecycle Manager). These tools are managed and configured within the SAP software inside the virtual machine and are independent of VMware.

## 5.16 Backup

Table 21. Backup Best Practices

### Backup Best Practices

1. Databases have their own native functionality to backup to the file system that manages database consistency, such as Oracle RMAN and HANA backup. These can be run on VMware in the same manner as physical.

- 
2. VMware vSphere Storage APIs – Data Protection are available for backup in the virtualized environment.

These APIs allow backup vendors to do a centralized, off-host, LAN-free backup of vSphere virtual machines without the need to install an agent or do backup processing within the guest OS. For more information on vSphere Storage APIs – Data Protection, see <http://kb.vmware.com/kb/1021175>.

vSphere Storage APIs – Data Protection backup software leverages VMware snapshots to create a read-only, consistent state of the virtual disk. Therefore, the APIs support virtual disks only or virtual mode RDM.

On Windows, the APIs will leverage VMware Tools, which will request Microsoft VSS services to manage database consistency for databases that have a VSS writer, SQL Server and Oracle).

Otherwise, it is generally the responsibility of the third-party backup software for placing the database in a consistent state before backup.

---

3. You can use storage vendor backup solutions based on storage array snapshots. These are similar to physical environments. Some considerations include:

- Storage vendor solutions are responsible for placing the database in a consistent state before the backup.
  - Dedicated datastores might be required for the database.
  - RDM might be required. Check with storage vendor.
  - Leverage storage-level, split-mirror techniques for faster backup and recovery.
- 

4. VMware vSphere Virtual Volumes™ are a new type of virtual machine object in vSphere 6, which are created and stored natively on the storage array. vSphere Virtual Volumes are stored in storage containers and mapped to virtual machine files and objects, such as VM swaps and VMDKs.

The storage system must be compatible with vSphere Virtual Volumes and able to integrate with vSphere 6 through the VMware APIs for Storage Awareness™ 2.0.

Work with your storage array vendor to determine vSphere Virtual Volumes compatibility.

vSphere Virtual Volumes with certified storage arrays allow for storage array snapshots and backups at the granularity of the virtual machine level (VMware snapshots are not required for virtual machine granularity).

For more information, see *VMware vSphere Virtual Volumes: Getting Started Guide* (<https://www.vmware.com/files/pdf/products/virtualvolumes/vmw-vsphere-virtual-volumes-getting-started-guide.pdf>).

For an example how virtualized SAP databases can leverage vSphere Virtual Volumes, see *Architecting Virtual SAP HANA Using VMware Virtual Volumes And Hitachi Storage* (<https://blogs.vmware.com/vsphere/2015/10/architecting-virtual-sap-hana-using-vmware-virtual-volumes-and-hitachi-storage.html>).

---

## 6. References

### 6.1 VMware References

- *Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere Deployment and Technical Considerations Guide*  
[http://www.vmware.com/files/pdf/SAP\\_HANA\\_on\\_vmware\\_vSphere\\_best\\_practices\\_guide.pdf](http://www.vmware.com/files/pdf/SAP_HANA_on_vmware_vSphere_best_practices_guide.pdf)
- *Monitoring Business Critical Applications with VMware vCenter Operations Manager*  
<http://www.vmware.com/files/pdf/solutions/Monitoring-Business-Critical-Applications-VMware-vCenter-Operations-Manager-white-paper.pdf>
- *Oracle Databases on VMware Best Practices Guide*  
[http://www.vmware.com/files/pdf/partners/oracle/Oracle\\_Databases\\_on\\_VMware\\_-\\_Best\\_Practices\\_Guide.pdf](http://www.vmware.com/files/pdf/partners/oracle/Oracle_Databases_on_VMware_-_Best_Practices_Guide.pdf)
- *Performance Best Practices for VMware vSphere 5.5*  
[http://www.vmware.com/pdf/Perf\\_Best\\_Practices\\_vSphere5.5.pdf](http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.5.pdf)
- *SAP Sybase Adaptive Server Enterprise on VMware vSphere Essential Deployment Tips*  
<http://www.vmware.com/files/pdf/SAP-Sybase-Adaptive-Server-Enterprise-on-VMware-vSphere.pdf>
- *SQL Server on VMware Best Practices Guide*  
[http://www.vmware.com/files/pdf/solutions/SQL\\_Server\\_on\\_VMware-Best\\_Practices\\_Guide.pdf](http://www.vmware.com/files/pdf/solutions/SQL_Server_on_VMware-Best_Practices_Guide.pdf)
- *The CPU Scheduler in VMware vSphere 5.1 Performance Study*  
<http://www.vmware.com/files/pdf/techpaper/VMware-vSphere-CPU-Sched-Perf.pdf>
- *VMware Performance VROOM! Blog – SAP Three-Tier Shows Excellent Scaling on vSphere*  
<http://blogs.vmware.com/performance/2010/03/sap-threetier-shows-excellent-scaling-on-vsphere.html>
- *VMware vSphere Resource Management Guide*  
<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-resource-management-guide.pdf>
- *SAP on VMware Sizing & Design Example Blog*  
<http://blogs.vmware.com/apps/2013/10/sap-on-vmware-sizing-design-example.html>
- *VMware vCloud Suite*  
<https://www.vmware.com/files/pdf/products/vCloud/VMware-vCloud-Suite-Datasheet.pdf>
- *Understanding Memory Resource Management in VMware vSphere 5.0*  
[http://www.vmware.com/files/pdf/mem\\_mgmt\\_perf\\_vsphere5.pdf](http://www.vmware.com/files/pdf/mem_mgmt_perf_vsphere5.pdf)
- *Security of the VMware vSphere Hypervisor*  
<http://www.vmware.com/files/pdf/techpaper/vmw-wp-secrty-vsphr-hyprvsr-uslet-101.pdf>
- *vSphere Availability*  
<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-availability-guide.pdf>

### 6.2 External References

The following third-party Web sites are not under the control of VMware and the content available at these sites might change.

#### 6.2.1 SAP

- *Best Practices for Virtualizing SAP Business Objects BI 4.x on VMware ESXi 5*  
<http://scn.sap.com/docs/DOC-42916>



- *SAP BusinessObjects BI4 Sizing Guide*  
<http://scn.sap.com/docs/DOC-33126>
- *SAP with Microsoft SQL Server 2008 and SQL Server 2005: Best Practices for High Availability, Maximum Performance, and Scalability*  
<http://www.sdn.sap.com/irj/scn/go/portal/prtroot/docs/library/uuid/4ab89e84-0d01-0010-cda2-82ddc3548c65?overridelayout=true>
- SAP Knowledge Base Article 161283 *Hardware Configuration Standards and Guidance*  
<http://service.sap.com/sap/support/notes/161283>
- *Certified SAP HANA Hardware Directory*  
<https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html>

## 6.2.2 EMC

- *SAP Redefined: EMC IT Virtualizes SAP HANA with VMware*  
<http://www.emc.com/collateral/emc-perspective/h112853-sap-redefined-ep.pdf>
- *EMC Cloud Enabled Infrastructure for SAP Business Continuity Series: Data Protection Bundle*  
<https://www.emc.com/collateral/white-papers/h12122-cloud-enabled-infrastructure-sap.pdf>

## 6.2.3 Cisco

- *SAP Applications Built on FlexPod*  
[http://www.cisco.com/en/US/docs/unified\\_computing/ucs/UCS\\_CVDs/ucs\\_flexpod\\_sap.html](http://www.cisco.com/en/US/docs/unified_computing/ucs/UCS_CVDs/ucs_flexpod_sap.html)

## 6.2.4 Hewlett-Packard

- *HP reference architecture for SAP Business Suite 7 Using HP ProLiant DL560, HP 3PAR StoreServ 7400, HP Serviceguard Extension for SAP, VMware and Red Hat Enterprise Linux*  
<http://h20195.www2.hp.com/v2/GetPDF.aspx%2F4AA4-7745ENW.pdf>

## 6.2.5 HDS

- *Deploy SAP ERP 3-tier Using Hitachi Unified Storage VM in a Scalable Environment*  
<http://www.hds.com/assets/pdf/deploy-sap-erp-3-tier-using-hus-vm.pdf>

## 6.2.6 VCE

- *VCE Solutions For SAP*  
<http://www.vce.com/asset/documents/sap-portfolio-solution-overview.pdf>



## Appendix A: High Availability Quantitative Analysis

This appendix introduces some concepts behind the quantitative analysis of high availability of applications running on an ESXi cluster. This enables the quantitative comparison of different high availability options and can assist with risk assessment.

Availability is defined as a number that specifies the degree of uptime expressed as a percentage or fraction. See the following table.

**Table 22. Availability Translation from Percentage to Actual Time**

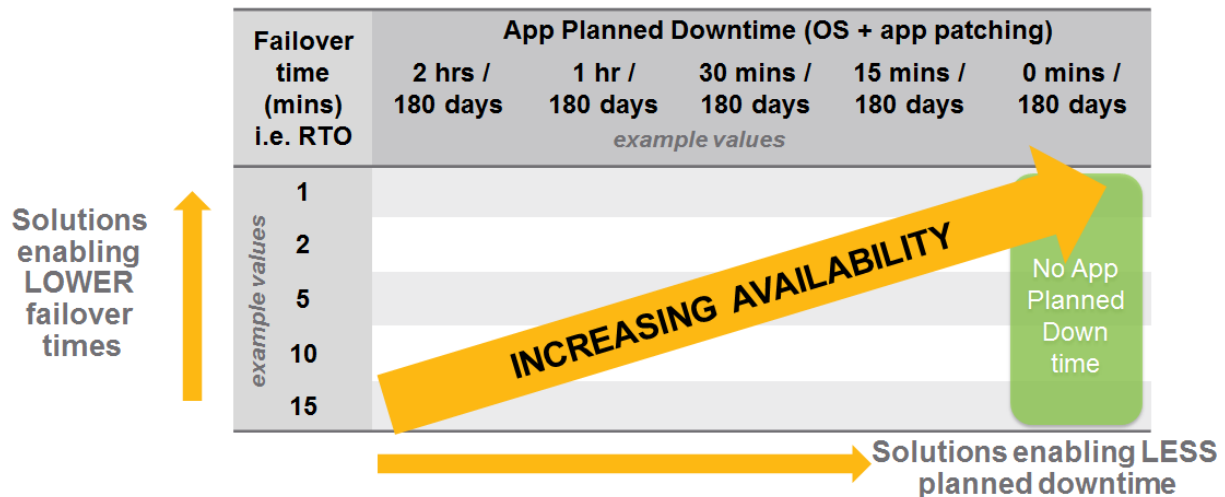
Level of Availability	Percent of Uptime	Downtime per Year
3 Nines	99.9%	8.76 hrs.
4 Nines	99.99%	52.6 min.
5 Nines	99.999%	5.25 min.

Availability is a probability, the likelihood that a system is up and running. Mathematical techniques are available to calculate probability and these are applied to calculate system availability.

The high availability chapter highlights different solutions, categorized into infrastructure and application, that can help protect against unplanned downtime and reduce planned downtime. Overall availability is a function of unplanned and planned downtime experienced in both layers, infrastructure and application.

The high availability options range from vSphere HA and vSphere FT to in-guest clustering and database replication solutions. These solutions affect availability by influencing the following parameters: failover time in case of ESXi host failure, and planned downtime of the guest OS, database, and Central Services software. The following figure shows the relationship.

**Figure 14. Failover Time and Application Planned Downtime**



An availability value (or probability expressed as a percentage) can be calculated based upon the following considerations:

1. Formulae that leverage probability theory as described in the book *Breaking the Availability Barrier: Survivable Systems for Enterprise Computing* by Bill Highleyman. An example can be found in the blog post *Estimating Availability of SAP on ESXi Clusters – Examples* (<https://blogs.vmware.com/apps/2013/07/estimating-availability-of-sap-on-esxi-clusters-examples-2.html>).
2. ESXi host mean time between failures (MTBF) and mean time to repair (MTTR). These values are not publicly documented, but can be estimated based upon empirical results in the data center.
3. N+S cluster considerations (N = number of hosts to satisfy CPU and memory of all the virtual machines, S = number of spare hosts added to the ESXi cluster for high availability). For example, a 4+1 ESXi cluster means that any two simultaneous host failures will result in severe resource limitations for the cluster, which would be classified as downtime. The probability of two simultaneous host failures is small, which increases availability. Consequently, the overall availability becomes more sensitive to failover times and planned downtime.
4. The failover time for the database and Central Services after an ESXi host failure varies depending on the workload and the type of solution (for example, vSphere HA versus an in-guest cluster solution). Values are not publicly documented, but this can be estimated based on proof-of-concepts and actual measurements in the data center.
5. The planned downtime for the application depends on the rolling patch upgrade solutions that are deployed and the downtime windows negotiated between IT and the business.

Taking these into consideration, three availability values can be calculated:

- Infrastructure/VMware level – Incorporates 1–4. (The values correspond to the highlighted area on the right in the figure).
- Application level – Incorporates 5. (Reflected in columns two to five in the above figure).
- Overall availability – A combination of infrastructure and application levels.

Availability calculations provide estimates and not guarantees, because the actual downtime experienced by users is impacted by other events that might not be factored into the calculations, such as application, network, storage, and power outages. The estimates can be used to quantitatively compare different options, such as vSphere HA as opposed to an in-guest cluster solution, and help with risk assessment.

The following figure shows an example availability calculation.

**Figure 15. Example Availability Calculation**

"4+1" ESXi Cluster Availability Calculation ( <i>Example</i> )	
<b>INFRASTRUCTURE</b>	<p><b>Assumption</b> DB failover time = 5 mins, CS failover time = 2 mins</p> <p><b>Assumption:</b> mtbf = 180 days; mtrr = 4 hrs For different failover times, mtbf, mtrr and "N+S" setup -&gt; recalculation required</p> <p><b>Ignoring</b> : ESXi host maintenance; power+network+storage outages (requires further analysis)</p> <p>Use ESXi cluster availability formula and example results table at: <a href="https://blogs.vmware.com/apps/2013/07/estimating-availability-of-sap-on-esxi-clusters-examples-2.html">https://blogs.vmware.com/apps/2013/07/estimating-availability-of-sap-on-esxi-clusters-examples-2.html</a></p> <p>From table in above blog: availability = 99.9964 % = 0.999964</p>
<b>APPLICATION</b>	<p><b>Assumption:</b> 1 hr of planned downtime per 3 months to patch OS, database and SAP</p> <p><b>Ignoring</b> application errors (requires further analysis)</p> <div style="text-align: center;"> <p>Mean Time Between Planned Downtime = 3 months</p> <p>UP</p> <p>DOWN</p> <p>Mean Time To Patch OS, DB, SAP ( 1 hr)</p> </div> <p>3 months ~ 90 days Availability=probability App is up and not being patched = proportion of time App is running and NOT being patched availability = (90 days - 1hr) / 90 days = 0.999537</p>
<b>FINAL AVAILABILITY</b>	<p>Final availability = (availability of infrastructure layer) x (availability of application layer)</p> <p>final availability = 0.999964 x 0.999537 = 0.999501 = 99.9501%</p> <p>i.e. between three to four nines</p> <p><b>CONCLUSION</b></p> <ul style="list-style-type: none"> <li>* <b>Result based on assumptions and input data which varies from customer to customer i.e mileage will vary</b></li> <li>* <b>Result is not a basis for a guarantee as input data are estimates and other factors impact downtime</b></li> <li>* The input data can be adjusted to account for different solutions that yield different failover times and planned downtime. This has to be estimated based on POCs , empirical data and duration of negotiated maintenance windows. This will generate different results that will help to quantitatively compare different HA solutions.</li> </ul>

Note that the calculation includes planned downtime in the final availability estimate. There are no formal rules on how to approach the availability analysis. For example, if the business has agreed to some planned downtime for application patching, then only unplanned events might need to be considered in the availability calculation. In such cases, solutions that help reduce planned downtime (like in-guest cluster software) might not be as critical.

## Acknowledgments

Author: Vas Mitra, SAP Solutions Architect

Thanks to the following for their reviews and inputs

- Arne Arnold - Field Architect/Staff Engineer
- Bob Goldsand – Staff Partner Architect
- Dominic Uliano - SAP Global Technical Alliance Manager
- Erik Rieger – SAP Sr. Systems Engineer / Alliances Manager
- Lars Fiechtner - Senior Technical Account Manager
- Sebastian Lenz - SAP Technical Support Engineer
- Todd Muirhead – Staff Engineer 2, Performance

Thanks to the following for their input on Oracle and Microsoft: Deji Akomolafe; Kannan Mani; Niran Even Chen; Sudhir Balasubramanian; Scott Salyer.