

SAP Solutions on VMware Best Practices Guide

© 2011 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

Contents

1. Introduction	5
2. VMware vSphere	5
3. SAP Platform Overview.....	6
4. Production Support of SAP Solutions on vSphere.....	7
5. Memory and Virtual CPU	8
5.1 Virtual Machine Memory	8
5.2 Virtual CPU	9
6. Storage and Networking.....	11
6.1 Storage.....	11
6.2 Networking	12
7. High Availability.....	13
8. Performance and Sizing.....	16
8.1 Performance.....	16
8.2 Performance Monitoring.....	18
8.3 Sizing	19
9. Timekeeping in Virtual Machines	20
10. Summary of Best Practices.....	21
11. Resources.....	26
11.1 Web Resources.....	26
11.2 SAP Notes.....	27
Appendix A: Virtual Machine Memory Settings	28
Appendix B: OS07N ESX/ESXi Performance Counters	29

1. Introduction

This paper provides best practice guidelines for deploying SAP software solutions on VMware vSphere®. These guidelines only provide general recommendations and do not target any specific size or type of SAP solution implementation.

VMware has created separate best practice documents for the individual areas of storage, networking, and performance. (See Section 11, Resources, for a list of these publications.) SAP also has created a variety of technical notes, published in the SAP Marketplace Web site at <http://service.sap.com/support>, with information about virtualizing SAP solutions on VMware virtual infrastructure. Reference numbers of these notes are identified in this document, and you can refer to these additional technical notes in conjunction with the information provided here.

Section 10, “Summary of Best Practices,” provides an overview of recommendations discussed in this document.

2. VMware vSphere

VMware vSphere is the virtualization platform that creates a layer of abstraction between the resources required by an application and operating system, and the underlying hardware that provides those resources. A summary of the value of this abstraction layer includes the following:

- Consolidation – VMware technology allows multiple application servers to be consolidated onto one physical server, with little or no decrease in overall performance.
- Ease of Provisioning – VMware virtualization encapsulates an application into an image that can be duplicated or moved, greatly reducing the cost of application provisioning and deployment.
- Manageability – Virtual machines can be moved from server to server with no downtime using VMware vSphere® vMotion®, which simplifies common operations like hardware maintenance and reduces planned downtime.
- Availability – VMware High Availability (HA) makes sure that in the case of an unplanned hardware failure, affected virtual machines are restarted on another host in a VMware cluster. With HA you can reduce unplanned downtime and provide higher service levels to an application. VMware vSphere® Fault Tolerance (FT) features zero downtime, zero data loss, and continuous availability in the case of server hardware failures for any application running in a virtual machine.

The latest version of vSphere is vSphere 5, which supports large capacity virtual machines that are especially well-suited to the memory-intensive and CPU-intensive footprint of SAP applications. vSphere 5 host and virtual machine specifications are as follows:

- Each VMware ESXi™ host supports up to 1TB RAM, 160 logical CPUs, and 2048 virtual CPUs.
- Each virtual machine can support up to 32 vCPUs and 1TB of RAM.

3. SAP Platform Overview

SAP ERP (Enterprise Resource Planning) is the SAP flagship product. In addition to ERP software, other key SAP products and solutions include business intelligence, customer relationship management, supply chain management, supplier relationship management, human resource management, product life cycle management, enterprise portal software, and knowledge warehouse software.

Most SAP applications are based on the SAP NetWeaver technology platform. SAP enterprise applications can be deployed in a two or three-tier architecture. The three-tier client/server architecture generally consists of a presentation layer, an application layer, and a database layer. These three layers can run separately on different computers or can all run together on the same computer, depending on the requirements and size of the SAP solution being deployed. In three-tier configurations, the database and application services reside on separate operating system (OS) images, whereas in two-tier configurations, they co-exist on the same OS image. The three-tier architecture scales to support large numbers of users. The two-tier architecture is usually sufficient for many smaller and midsize companies, as well as for sandbox, development, training, and test systems.

4. Production Support of SAP Solutions on vSphere

vSphere (including version 5) is supported by SAP on Linux and Windows guest operating systems as documented in SAP Notes 1122387 and 1409608 respectively. SAP Note 1492000 covers the general SAP support statement for virtual environments. As of Q4 2011, SAP does not support Solaris x64 as a guest OS.

SAP supports running the following 64-bit versions of SAP NetWeaver on VMware virtual infrastructure:

- SAP NetWeaver 2004 (SAP Kernel 6.40) and above (ABAP and/or JAVA stack).
- Older SAP NetWeaver and application versions and 32-bit systems are supportable only during an SAP upgrade.

For Windows, all SAP-certified hardware is supported as long as it is also on the VMware hardware compatibility list (HCL). (See the Resources section for the web location of the HCL.) A list of hardware SAP-certified for Windows is available at <http://www.saponwin.com/>.

Note SAP has removed the requirement to certify servers specifically having VMware virtual infrastructure, and separate server certificates are no longer necessary.

For Linux, hardware vendors have to explicitly support their hardware for Linux running on VMware virtual infrastructure. The Supported Platforms link on the SAP Linux support homepage (<http://www.sap.com/linux>) provides details. In addition, the hardware needs to be on the VMware Compatibility Guide (<http://www.vmware.com/resources/compatibility/search.php>).

SAP acquired Business Objects, and Business Objects applications, are also supported on VMware virtual infrastructure (see SAP Note 1206126).

VMware has worked with SAP to include VMware performance counters in the SAP OS collector program `saposcol`. These VMware ESX[®]/ESXi counters can be viewed in SAP transaction OS07N, and require application of SAP Note 1409604, *Virtualization on Windows: Enhanced monitoring*. (Information in this note also applies to Linux.) This note also must be applied to obtain SAP support.

5. Memory and Virtual CPU

5.1 Virtual Machine Memory

This section provides guidelines for determining the number of virtual machines on a single ESX/ESXi host system based on memory requirements. See Appendix A: Virtual Machine Memory Settings for a description of virtual machine memory settings discussed in this section. For further background on VMware memory management concepts, see the *VMware vSphere Resource Management Guide* at <http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-resource-management-guide.pdf>.

Because SAP applications are generally memory-intensive, and to account for situations where performance is a key factor (for example, in mission critical production environments), VMware recommends the following:

- Do not over-commit memory on ESX/ESXi host servers. For production systems, it is possible to enforce this policy by setting the memory reservation to the configured size of the virtual machine. Also note that:
 - Setting reservations may limit vMotion migration capabilities. A virtual machine can only be migrated if the target ESX/ESXi host has free physical memory equal to or greater than the size of the reservation.
 - Setting the memory reservation to the configured size of the virtual machine results in a per-virtual machine VMkernel swap file of zero bytes (which consumes less storage).
- It is important to *right-size* the configured memory of a virtual machine. Memory is wasted if the SAP applications are not utilizing the configured memory. ESX/ESXi performance counters can be used to determine actual memory usage (see Appendix B: OS07N ESX/ESXi Performance Counters”).
- The guest operating system within the virtual machine still needs its own separate swap/page file, per standard SAP recommendations.
- Do not disable the balloon driver.
- Allocate virtual machines on a single ESX/ESXi host based on the following formula:

$$\text{Memory available for SAP virtual machines} = [\text{total ESX server physical memory}] - [\text{memory required by ESX/ESXi}] - [\text{user-defined "memory buffer"}]$$

Memory required by an ESX/ESXi host comprises memory required by the Console Operating System (COS), plus memory required by VMkernel, plus memory required for each virtual machine (which depends on the size of the virtual machine). The *VMware vSphere Resource Management Guide* provides more detail about memory requirements.

ESX/ESXi transparent page sharing makes more physical memory available, but this additional memory is not counted here to provide a more conservative estimate.

The *memory buffer* is not a VMware parameter, but is a user-defined value designed to provide headroom and flexibility to manage more virtual machines than initial estimates call for (for example, for virtual machines migrated, using vMotion, from another ESX/ESXi host). Actual memory buffer sizes will depend on specific customer design requirements.

These guidelines are purposely conservative to avoid kernel swapping between ESX/ESXi and the guest OS, important due to the mission-critical nature of SAP business processes, which must meet stringent SLAs and the memory intensive requirements of the ABAP and JAVA stack. This best practice can also apply to non-production systems with high performance SLAs for developers and testers who support production environments.

Note that vSphere 5 introduces two new memory reclamation techniques to allow users to overcommit host memory (see *Understanding Memory Management in VMware vSphere 5* at <http://www.vmware.com/resources/techresources/10206>). Memory compression reduces the number of host-swapped pages by storing the compressed format of the pages in a per-VM memory compression cache. Swap-to-SSD leverages SSD's low read latency to alleviate the host swapping penalty. While these two new techniques can help to reduce the performance impact in memory overcommit scenarios, the above guidelines not to overcommit memory for SAP production virtual machines still stands.

To minimize guest operating system swapping, the configured memory size of the virtual machine should be greater than the average memory usage of the SAP application running in the guest. If the SAP application in the virtual machine needs more memory than it has been allocated, the guest operating system paging/swapping mechanisms are invoked.

Memory and swap/page file configuration of the SAP application in the virtual machine follow the same guidelines as for native environments, and generally, you should set them to minimize guest operating system swapping. Follow existing SAP documentation and recommendations as provided in these SAP notes:

- 88416 – *Zero Administration Memory Management as of 4.0A/Windows.*
- 1009493 – *abap/heap_area* parameter Defaults Changed (64-Bit Windows).*
- 723909 – *Java virtual machine settings for J2EE 6.40/7.0.*
- 941735 – *SAP memory management for 64-bit Linux systems (or: STD memory model).*
- 386605 -- *SAP memory management for 32-bit Linux systems (or: MAP memory model).*

5.2 Virtual CPU

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. Virtual machines with more than one virtual CPU are also called SMP (symmetric multiprocessing) virtual machines.

VMware Virtual Symmetric MultiProcessing (Virtual SMP) enhances virtual machine performance by enabling a single virtual machine to use multiple physical processors simultaneously. vSphere supports use of up to 32 virtual CPUs per virtual machine. 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 is multithreaded (NetWeaver JAVA stack) and includes multiple processes (NetWeaver ABAP stack comprises multiple “disp+work” C processes) which makes it a good candidate to take advantage of Virtual SMP.

In the latest versions of ESX/ESXi, the CPU scheduler has undergone several improvements to provide better performance and scalability; for details, see the paper *VMware vSphere: The CPU Scheduler in VMware ESX 4.1*. For example, the relaxed co-scheduling algorithm was refined so that scheduling constraints due to co-scheduling requirements are further reduced. These improvements have resulted in better scalability and performance of SAP workloads, as described in Section 8, “Performance and Sizing.” Consequently, in vSphere, the larger 4-way and 8-way virtual machines exhibit great scalability, so that running multiple smaller 2-way virtual machines for better performance is not required as recommended with ESX/ESXi 3 versions.

While larger virtual machines are possible in vSphere, VMware recommends reducing the number of virtual CPUs if monitoring of the actual workload shows that the SAP application is not benefitting from the increased virtual CPUs. For more background, please see the “ESXi CPU Considerations” section in the whitepaper *Performance Best Practices for VMware vSphere 5*.

Note that SAP has conducted tests on virtual CPU overcommitment (documented in Note 1122388), and shows the performance degradation inside the virtual machines is linearly reciprocal to the overcommitment. As the performance degradation is “graceful,” any virtual CPU overcommitments can be effectively managed by using vMotion to migrate virtual machines to other ESX/ESXi hosts to obtain more processing power.

Hyperthreading technology (recent versions of which are called symmetric 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—which 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. For example, an ESX/ESXi host system enabled for SMT on an 8-core server sees 16 threads that appear as 16 logical processors. Recent SAP benchmarks were conducted on SMT-enabled servers—these are covered in Section 8.

6. Storage and Networking

6.1 Storage

It is preferred practice to deploy virtual machines files on shared storage to take advantage of vMotion and VMware HA. This practice aligns well with SAP solution-based deployments, which are typically installed on third-party shared storage management solutions.

The following two methods of storage configuration are covered here:

- VMware Virtual Machine File System (VMFS) is a clustered file system that provides storage virtualization optimized for virtual machines.
- Raw Device Mapping (RDM) provides a mechanism for a virtual machine to have direct access to a volume on a physical storage subsystem. RDM can only be used with Fibre Channel or iSCSI.

VMware generally recommends the use of VMFS. You can use RDM where existing systems already make use of third-party storage management software so that you can use RDM to leverage existing practices and tools, for example:

- Storage-based backups to disk.
- Database-consistent replication in DR scenarios.
- Where RDM is required when using third-party clustering software.

RDM enables quicker migration between physical server and virtual environments. Database files in either physical or virtual environments can be accessed just as they are, without the need for a data conversion to or from VMFS format.

A mixed storage configuration is viable for an SAP virtual machine—the guest operating system is installed with VMFS and the SAP database files with RDM. VMware template cloning can be used for the guest operating system and database files can be managed by third party storage management software.

Generally, for performance-critical production SAP databases, follow these recommendations:

- Spread out database data files over multiple LUNs, similar to those in native setups, following the storage vendor array guidelines for database layout, LUN, and spindle configuration.
- Use multiple virtual SCSI controllers for database virtual machines. The use of multiple virtual SCSI controllers allows the execution of several parallel I/O operations inside the guest OS. It is recommended to use one controller for OS and swap, one controller for redo/log, and one or more controllers for data files.
- Use the PVSCSI adapter for database files – this offers a significant reduction in CPU utilization as well as potentially increased throughput.
- Use eager-zeroed thick disks for database files. An eager-zeroed thick disk has all space allocated and zeroed out at the time of creation. This increases the time it takes to create the disk, but results in the best performance during database operations, such as when the database autogrows its data files.

Follow the guidelines in the “ESXi Storage Considerations” and “Guest Operating Systems” sections of *Performance Best Practices for VMware vSphere 5*.

6.2 Networking

The standard VMware networking best practices apply to running SAP applications on vSphere:

- Allocate separate network adapters/networks for vMotion, VMware FT logging traffic, and ESX/ESXi console access management.
- Allocate at least two network adapters for SAP data traffic to leverage VMware NIC teaming capabilities. Generally, at least four network adapters are recommended for each ESX/ESXi host.
- Use the VMXNET3 network adapter - this is a paravirtualized device that works only if VMware Tools is installed on the guest operating system. The VMXNET3 adapter is optimized for virtual environments and designed to provide high performance.
- To support VLANs in vSphere, the virtual or physical network must tag the Ethernet frames with 802.1Q tags using virtual switch tagging (VST), virtual machine guest tagging (VGT), or external switch tagging (EST). VST mode is the most common configuration.
- Follow the networking design guidelines in VMworld 2011 session VSP3864 Best Practices for Virtual Networking. This includes designs to efficiently manage multiple networks and redundancy of network adapters on ESX/ESXi hosts.
- Follow the guidelines in the “ESXi Networking Considerations” and “Guest Operating Systems” sections of *Performance Best Practices for VMware vSphere 5*.

7. High Availability

The VMware Fault Tolerance (FT) and VMware High Availability (HA) features together can provide high availability options for SAP single points of failure in the virtualized environment.

VMware FT protects a virtual machine by maintaining a second virtual machine that runs in lockstep with the primary virtual machine. If the primary virtual machine goes down, the secondary machine takes over with no downtime. Currently, VMware FT supports only single-CPU virtual machines and is a viable solution for lightweight components of the SAP architecture such as Central Services. When deploying SAP Central Services standalone in a virtual machine, note the following:

- Linux-based guest OS is supported by SAP and there are no caveats.
- For Windows-based guest OS, see SAP note 1609304, *Installing a standalone ASCS instance*. To obtain support on Windows for a standalone Central Services deployment in a virtual machine follow these guidelines:
 - Use a *sapinst* that allows installation of standalone Central Services (available from NetWeaver 7.3, but also possible with some earlier versions).
 - Take care of RFC destinations that point to the virtual hostname of the Central Services by maintaining RFC group destinations or implementing a standalone gateway.
 - In case of an upgrade, choose the correct upgrade tools (if you need advice, open a SAP message under support component BC-UPG).
 - For clarification, open a SAP ticket under support component BC-OP-NT-ESX before proceeding with an installation.

VMware HA continuously monitors all ESX/ESXi hosts in a cluster and, in case of an ESX/ESXi host failure, restarts all affected virtual machines on the remaining hosts.

Though VMware HA and VMware FT can provide ESX/ESXi server hardware protection to SAP single points of failure, it does not monitor the health of the application (that is, SAP database and Central Instance or Central Services). If application level monitoring is also required, then investigate using third-party clustering software in virtual machines (in-guest clustering), or a solution from Symantec, called ApplicationHA, which integrates with VMware HA. Table 1 summarizes the HA solutions.

Table 1. Summary of High Availability Scenarios

HA Scenario	Hardware Protection	Application Aware	Rolling Patch Upgrade Support	Guest-OS boot after H/W failure	Complexity Cost
VMware HA VMware FT (SCS)	Yes	No	No	Yes	Low
Symantec ApplicationHA + VMware HA	Yes	Yes	No	Yes	Medium
In-guest Clustering	Yes	Yes	Yes	No	High

Table 2 identifies the different in-guest clustering solutions available on VMware.

Table 2. Cluster Solutions Supported on VMware by Vendors

Cluster Solution	Vendor Support	RDM	VMFS	Guest OS	Comments	
Microsoft Cluster	Yes	Yes	No	Windows	<ul style="list-style-type: none"> Requires RDM, cannot vMotion migrate clustered virtual machine. See VMware guide <i>Setup for Failover Clustering and Microsoft Cluster Service</i>. 	
Symantec/Veritas Cluster Services	Yes	Yes	No	Windows, SUSE Linux, Red Hat Linux	<ul style="list-style-type: none"> For iSCSI and FC SAN requires RDM. Cannot use vMotion to migrate clustered virtual machines. http://www.symantec.com/connect/articles/clustering-configurations-supported-vcs-vsphere 	
SUSE Availability Extension	High	Yes	Yes	SUSE Linux	<ul style="list-style-type: none"> For VMFS, need to use “multi-writer flag”. See VMware KB article 1034165. Enables vMotion. http://www.cc-dresden.de/en/whitepaper 	
Red Clustering	Hat	Yes	Yes	Yes	Red Hat Linux	<ul style="list-style-type: none"> Supported by Red Hat from 5.7 and later. For VMFS, need to use “multi-writer flag”, see VMware KB article 1034165. Enables vMotion.
Oracle RAC	Yes	Yes	Yes	SUSE Linux, Red Hat Linux	<ul style="list-style-type: none"> Supported by Oracle from 11.2.0.2 and later as per MyOracleSupport, Document ID #249212.1. For VMFS, need to use “multi-writer flag”, see VMware KB article 1034165. Enables vMotion. 	

Given that there are different high availability design choices available for SAP installation on VMware virtual infrastructure, the final approach taken depends on your specific business and sizing requirements and Service Level Agreements (SLAs). The following considerations may influence your choices:

- If only hardware protection is required, VMware HA and VMware FT provide an economical choice, as it is easy to configure VMware “out-of-the-box” functionality without the complexity of installing clustering software.
- The decision not to go with application level monitoring may depend on your previous failover experiences with clustering software. For example how often a failover occurs due to application failure only (for example, OS, database, Central Instance) and hardware was not the source of the problem.
- Many customers who run SAP solutions on VMware virtual infrastructure have fulfilled their high availability SLAs with VMware HA, which has lowered their total cost of ownership (TCO). You can find examples in a detailed study of three customer implementations documented in the whitepaper TCO and ROI Analysis of SAP Landscapes using VMware Technology (see Section 11, Resources, section for the document link).
- If sizing of the SAP system is such that all SAP locking and messaging activities can be satisfied by one core of the latest x86-technology based processor, then zero-downtime protection against hardware failure for Central Services is possible with VMware FT without the complexity of configuring replicated enqueue in a clustered environment. (This scenario assumes Central Services is installed in a single virtual CPU virtual machine.) For larger systems, an in-house performance test may be required to determine suitability.
- If you require application level monitoring for the database, Central Instance or Central Services:
 - Symantec Application HA helps to bridge the gap between VMware HA and in-guest clustering.
 - Clustering software in virtual machines can provide application monitoring and rolling patch upgrade capability however, you will need cluster configuration skills and may have to pay for additional cluster software license costs.
- The final design choice depends on how much downtime a business can realistically tolerate, and the cost they are willing to invest in the extra resources and skills to install and operate software that provides application monitoring. It is a trade-off.

8. Performance and Sizing

8.1 Performance

Benchmark test results for SAP ERP running on vSphere are available at <http://www.sap.com/solutions/benchmark/index.epx>.

The following table summarizes the vSphere certifications as of Q4 2011.

Table 3. SAP on vSphere Benchmarks (+ one Native for comparison) Q4 2011

Cert	2-tier/ 3-tier	vSphere Version	Virtual Machines	Users	SAPS	Server Configuration
2009028	2-tier	vSphere 4	4-way	1144	6250	8 cores 16 threads
2009029	2-tier	vSphere 4	8-way	2056	11230	8 cores 16 threads
2010016	3-tier	vSphere 4	15 x 8-way	16000	87800	8 cores 16 threads
2011044	3-tier	vSphere 5	20 x 12-way + 20-way	32125	175320	12 cores 24 threads
2011027	2-tier	vSphere 5	24-way	4600	25120	12 cores 24 threads
2011028	2-tier	NATIVE	N/A	4875	26630	12 cores 24 threads

- “SAPS” stands for *SAP Application Performance Standard*, a hardware-independent unit that describes the performance of a SAP system configuration (further details are available at <http://www.sap.com/solutions/benchmark/measuring/index.epx>).
- All the certifications above are on servers with hyperthreading enabled.
- Certifications 2009028 and 2009029 – The total number of vCPUs is <= number of physical cores. In these cases each vCPU is scheduled on a thread on a separate core, that is, the vCPU has full access to the core.
- Certification 2209029 – The total number of vCPUS equals the number of cores but is less than the total threads. The SAPS result is not an indicator of the full capacity of the server as there are unused threads.
- Certification 2011027 – The total number of vCPUs is greater than the number of cores and equal to the total number of threads so that all the server threads are busy and pairs of vCPUs are being scheduled on threads on the same core simultaneously, that is, they are sharing core resources. This result is effectively the full virtual SAPS rating of the server. The performance of the 24-way virtual machine would be higher if it was run on a 24-core host where each vCPU would be scheduled on its own core.
- Certifications 2011027 and 2011028 – These two are a direct comparison between virtual and physical, and the virtual result is within 6% of physical/native.

Separate 3-tier SAP OLTP tests on ESX/ESXi servers with hyperthreading enabled on Intel Xeon 5500 or higher has shown up to a 24% gain in performance from hyperthreading (for background, refer to the VMware performance blog listed in the Resources section).

To maximize performance of SAP applications in the virtual environment, VMware recommends the following:

- Use the latest hardware to exploit vSphere support of hardware nested page tables to obtain the best SAP application performance on vSphere.
- 4-way and 8-way application server virtual machines are viable for environments running on vSphere. Note that in ESX/ESXi 3.X versions, 2-way virtual machines are still optimum.
- Where possible size the virtual machine so that it fits within a NUMA node (virtual memory is less than NUMA node memory and vCPU count that is less than or equal to the number of processors per node). This configuration allows vSphere to employ NUMA optimizations for memory and CPU scheduling and ensures that all memory accesses will be satisfied by the memory closest to the processor
- vSphere 5 introduces virtual NUMA whereby ESXi 5.0 exposes NUMA topology to the guest operating system, allowing NUMA-aware guest operating systems and applications to make the most efficient use of the underlying hardware's NUMA architecture. For SAP systems, where sizing dictates, large single instance databases can be installed in a wide virtual machine that crosses NUMA nodes and virtual NUMA sockets can be defined for the virtual machine. In this setup, the database will behave in similar manner to when installed on a large NUMA based physical server and the NUMA optimizations of the database will determine the performance benefits.
 - See the paper *Performance Best Practices for VMware vSphere 5* for further virtual NUMA guidelines and configuration steps.
 - Consult your database vendor documentation to determine NUMA optimization settings for the database.
- Install the latest version of VMware Tools in the guest operating system.
- Download and follow guidelines in the following SAP Notes
 - 1056052 – *Windows: VMware ESX/ESXi 3.x or vSphere configuration guidelines.*
 - 1122388 – *Linux: VMware ESX Server 3 configuration guidelines.*
- For vSphere 4, follow the guidelines in VMware KB article 1020233 *HaltingIdleMsecPenalty Parameter: Guidance for Modifying vSphere's Fairness/Throughput Balance* to maximize benefits of hyperthreading. This KB article is *not* required for vSphere 5.

8.2 Performance Monitoring

For performance monitoring, ESX/ESXi performance counters are available in SAP transaction OS07N (see SAP Note 1409604 for Windows and Note 1102124 for Linux). Appendix B: OS07N ESX/ESXi Performance Counters includes a screen capture and a description of the virtual counters. OS07N provides a starting point from which you can monitor the virtual environment. For performance troubleshooting, VMware recommends following the guidelines in the paper *Performance Troubleshooting for VMware vSphere 4*. (See Section 11, Resources, for Web site location.) This paper provides a guide for checks including:

- ESX/ESXi host CPU saturation.
- ESX/ESXi host server swapping.
- Virtual machine guest CPU saturation.
- Network and storage issues.

Access to VMware vCenter Server™ through the vSphere client is required to view the major ESX/ESXi performance counters necessary for troubleshooting CPU, memory, storage, and network issues.

Table 4 highlights some useful counters (for a full list of counter, see *VMware Communities: Interpreting esxtop Statistics* at <http://communities.vmware.com/docs/DOC-9279>).

Table 4. Some Useful ESX/ESXi Performance Counters

Subsystem	esxtop Counters	vCenter Counter	Description (Threshold)
CPU	%RDY	Ready (ms in a 20s window)	Measures CPU contention between virtual machines (< 5% per vCPU).
	%USED	Usage	CPU utilization of virtual machine.
Memory	%ACTV	Active	Guest memory being actively used.
	SWW/s SWR/s	Swapin Rate Swapout Rate	ESX/ESXi host swap counters: with no memory overcommit these should be zero.
Storage	DAVG/cmd	deviceWriteLatency and deviceReadLatency	I/O latency from ESX/ESXi to storage array (10-20ms).
	KAVG/cmd	kernelWriteLatency and kernelReadLatency	I/O latency in hypervisor (< 1 ms).
Network	%DRPTX	Dropped Packets	Percentage of transmit packets dropped (0).

8.3 Sizing

SAP has established a sizing process with its hardware partners to determine the hardware requirements to implement an SAP system. The sizing process uses the Web-based Quick Sizer tool, which calculates SAP requirements based on throughput numbers, and the number of users working with the different SAP Business Suite components, in a hardware and database independent format. For more information, go to <http://service.sap.com/sizing>. (SAP Marketplace access is required to reach this site.)

No changes to the Quick Sizer process are needed for a virtualized system configuration—follow the same process when sizing for either a virtual or a physical environment. After obtaining the results of a Quick Sizer project, work with the SAP Competency Center of your specific hardware vendor for sizing and architecture services as you would for physical environments. VMware works closely with the same SAP hardware partners so that they have the equivalent SAPS ratings for virtual machines, as these are directly related to the speed of the processor.

The following considerations are applicable for sizing SAP solutions on VMware virtual infrastructure:

- Consult the SAP Competency Center of a specific hardware vendor for a detailed architecture design and official sizing estimate.
- Use the Quick Sizer tool in the normal manner to obtain SAPS requirements for SAP business modules.
- Note that generally, a direct certified benchmark comparison between a VMware virtual platform and the equivalent physical implementation on the same server may not be possible. This is because the virtual result is based on 90 percent or greater CPU utilization within the virtual machine and the underlying physical ESX/ESXi host may not be fully utilized.
- Hardware vendors have an approximate general memory requirement per core for SAP applications. The same requirement is applicable for virtual CPUs, so the per-core memory rating is equivalent to the virtual CPU rating. For example, a 4-8GB per-core requirement translates to the same 4-8GB requirement per virtual CPU.
- To determine the number of virtual machines for an ESX/ESXi host from a memory standpoint, follow the guidelines provided in Section 5, “Memory and Virtual CPU.”

9. Timekeeping in Virtual Machines

Most operating systems track the passage of time by configuring the underlying hardware to provide periodic interrupts. The rate at which those interrupts are configured to arrive varies for different operating systems. High timer-interrupt rates can incur overhead that affects a virtual machine's performance. The amount of overhead increases with the number of vCPUs assigned to a virtual machine.

For many Linux operating systems, the default timer interrupt-rate is high and can lead to time synchronization errors in SAP applications running in virtual machines:

- Error messages in the SAP Syslog: "System time was set externally to a time in the past. Wait 1 second(s)."
- At the operating system level, the clock of the virtual machine might run too quickly or too slowly.
- Time drift between the application and database server can cause ABAP short dumps with the error message "ZDATE_LARGE_TIME_DIFF."
- To address timekeeping issues when running SAP solutions on Linux guest operating systems:
 - Use Novell SLES 9 and later versions, or Red Hat RHEL 5.1 and later, because these operating system versions allow the frequency of timer interrupts to be reduced.
 - Follow the guidelines in SAP Note 989963 – *Linux: VMware timing problem*.

10. Summary of Best Practices

This section summarizes best practices. See Section 11.2, SAP Notes, for the location of SAP notes mentioned in the following recommendations.

Item	Comments
Recommendation	Download and follow guidelines in the following SAP Notes: 1056052 Windows: VMware ESX 3.x or vSphere configuration guidelines; 1122388 - Linux: VMware ESX Server 3 configuration guideline.
Justification	These notes are updated, so download the latest versions. The different SAP memory models are described in these notes.

Item	Comments
Recommendation	Set memory reservations equal to the size to the virtual machine
Justification	To avoid kernel swapping between ESX/ESXi and the guest OS because SAP application and database servers can be memory-intensive.

Item	Comments
Recommendation	Use large memory pages for databases.
Justification	Large page support is enabled by default in ESX/ESXi versions 3.5 and later, and is supported in Linux and Windows. Enable large pages in the guest OS to improve the performance of SAP applications on vSphere. SAP application servers do not support large pages.

Item	Comments
Recommendation	Use as few virtual CPUs (vCPUs) as possible.
Justification	If monitoring of the actual workload shows that the SAP application is not benefitting from the increased virtual CPUs, the excess vCPUs impose scheduling constraints and can degrade overall performance of the virtual machine.

Item	Comments
Recommendation	Spread database files out over multiple datastores/LUNs.
Justification	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.

Item	Comments
Recommendation	Use eager-zeroed thick disk for database and SAP application server files.
Justification	Provides better performance during database operations, for example, when the database autogrows its data files there is no first zero write penalty.

Item	Comments
Recommendation	Make sure VMFS is correctly aligned.
Justification	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.

Item	Comments
Recommendation	Avoid silos when designing the storage architecture.
Justification	When designing an optimized architecture, involve the database, storage, VMware, and SAP Basis administrator.

Item	Comments
Recommendation	Use paravirtualized SCSI adapters for database datafiles
Justification	The combination of the new paravirtualized SCSI driver (pvscsi) and additional ESX/ESXi kernel-level storage stack optimizations dramatically improves storage I/O performance.

Item	Comments
Recommendation	Spread the database files over all the virtual SCSI controllers.
Justification	Maximizes parallel processing of I/O in the guest OS.

Item	Comments
Recommendation	Use the VMXNET family of paravirtualized network adapters.
Justification	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.

Item	Comments
Recommendation	Separate infrastructure traffic from virtual machine traffic for security and isolation.
Justification	Virtual machines should not see infrastructure traffic (security violation) and should not be impacted by infrastructure traffic bursts (for example, by vMotion).
Item	Comments
Recommendation	Use NIC teaming for availability and load balancing.
Justification	NIC teams can share the load of traffic among some or all of its members, or provide passive failover in case of a hardware failure or a network outage.
Item	Comments
Recommendation	When installing standalone Central Services on Windows, follow SAP note 1609304, <i>Installing a standalone ASCS instance</i> . For clarification open an SAP ticket under support component BC-OP-NT-ESX.
Justification	If sizing permits standalone, Central Services in a one-way virtual machine can be protected under VMware Fault Tolerance.
Item	Comments
Recommendation	When deploying clustering software from SUSE, Redhat, or RAC, follow VMware KB article 1034165, <i>Disabling simultaneous write protection provided by VMFS using the multi-writer flag</i> at http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&md=displayKC&externalId=1034165 . This uses the “multi-writer flag” to enable sharing of vmdks by multiple virtual machines.
Justification	In guest clustering software may be required to satisfy some customer SLAs. Clustering with VMFS provides compatibility with vMotion.

Item	Comments
Recommendation	Use the virtual counters in OS07N to monitor the virtual environment as a starting point. For advanced troubleshooting, use VMware vCenter and/or the esxtop/resxtop utility.
Justification	CPU and memory usage reported within the guest OS can be different from what ESX/ESXi reports. No ESX/ESXi disk counters are available in OS07N.

Item	Comments
Recommendation	Configure virtual NUMA sockets for wide virtual machines that need to cross NUMA nodes. Follow guidelines in <i>Performance Best Practices for VMware vSphere 5</i> at http://www.vmware.com/resources/techresources/10199 .
Justification	Depending on sizing requirements, large database virtual machines might need to be larger than a NUMA node. ESXi 5.0 exposes NUMA topology to the guest operating system, allowing NUMA-aware guest operating systems and databases to make the most efficient use of the underlying hardware's NUMA architecture.

Item	Comments
Recommendation	To minimize time drift in virtual machines, follow the guidelines in SAP Note 989963 – <i>Linux: VMware timing problem</i> .
Justification	The impact of high timer-interrupts in some operating systems can lead to time synchronization errors.

Item	Comments
Recommendation	Use vShield App to protect and isolate SAP virtual machines in consolidated ESX/ESXi clusters.
Justification	Enable security compliance by providing network isolation between production and non-production landscapes, different SAP products, database and application tier, and the like. (See paper in Resources section).

Some further general considerations include the following:

- VMware virtual machine cloning from templates that include a previously-installed SAP instance can drastically decrease the time required to provision new SAP systems, databases, and application servers by making reinstallation of the guest OS, database, and SAP software unnecessary. You will have to make further SAP application-specific changes after provisioning a new virtual machine to change the SAP SID and hostname to make the new instance unique.
- Generally, after you have correctly sized virtual machines with the memory and virtual CPUs 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 Basis administration tasks and procedures apply.

An *SAP on VMware* forum is available and hosted by the SAP Developer Network. This forum is moderated by SAP and VMware engineers and is a good place for discussions and technical questions.

Consult papers and Web sources listed in Section 11, “Resources,” to provide more background and details on the recommendations specified in this document.

The guidelines documented here have enabled VMware partners and customers to successfully size, architect, and deploy SAP solutions on VMware virtual infrastructure. Published success stories are available at the following URL:

<http://www.vmware.com/partners/alliances/technology/sap.html>.

11. Resources

This section lists Web sites where you can obtain more information about using VMware and SAP solutions.

11.1 Web Resources

- *VMware and SAP* Web site (success stories, whitepapers, technical case studies):
<http://www.vmware.com/partners/alliances/technology/sap.html>
- *SAP on VMware* forum at the SAP Community Network:
<http://forums.sdn.sap.com/forum.jspa?forumID=471>
- *Performance Best Practices for VMware vSphere 5*:
<http://www.vmware.com/resources/techresources/10199>
- *Understanding Memory Management in VMware vSphere 5*:
<http://www.vmware.com/resources/techresources/10206>
- *VMware Compatibility Guide*:
<http://www.vmware.com/resources/compatibility/search.php>
- *VMware vSphere Resource Management Guide*:
<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-resource-management-guide.pdf>
- *Configuration Maximums VMware vSphere 5.0*:
<http://www.vmware.com/pdf/vsphere5/r50/vsphere-50-configuration-maximums.pdf>
- *VMware vSphere: The CPU Scheduler in VMware ESX 4.1*:
<http://www.vmware.com/resources/techresources/10131>
- *Performance Troubleshooting for VMware vSphere 4.1*:
<http://www.vmware.com/resources/techresources/10179>
- VMworld 2011 session *VSP3864 Best Practices for Virtual Networking*:
<http://www.vmworld.com/community/sessions> (VMworld account required)
- *vSphere Guest Programming Guide, VMware vSphere Guest SDK 4*:
http://www.vmware.com/support/developer/guest-sdk/guest_sdk_40.pdf
- *SAP SD Standard Application Benchmark Results, Two-Tier and Three-Tier Configurations* (includes benchmarks conducted on vSphere):
<http://www.sap.com/solutions/benchmark/index.epx>
- *TCO and ROI Analysis of SAP Landscapes using VMware Technology*:
http://www.vmware.com/files/pdf/partners/sap/SAP_TCOROI_Customers_Final.pdf
- VMware KB article 1020233 *HaltingIdleMsecPenalty Parameter: Guidance for Modifying vSphere's Fairness/Throughput Balance*:
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1020233

- 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>
- *Ready Time* Blog:
<http://communities.vmware.com/docs/DOC-7390>
- VMware KB article 1008205 – *Using esxstop to identify storage performance issues*:
<http://kb.vmware.com/kb/1008205>
- *VMware vShield App Protecting Virtual SAP Deployments*:
<http://www.vmware.com/resources/techresources/10213>

11.2 SAP Notes

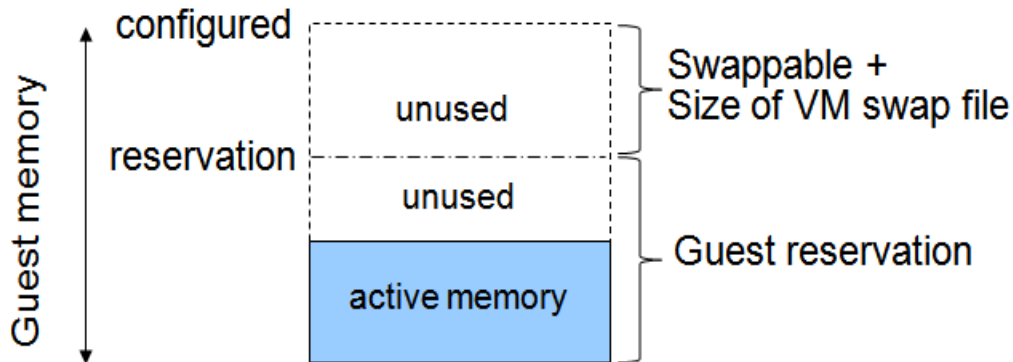
A list of SAP on vSphere notes is available at the following location:
<http://forums.sdn.sap.com/thread.jspa?threadID=1524523>

More notes are available at the SAP Marketplace Web site at <http://service.sap.com/support>. (This site requires login credentials.)

Appendix A: Virtual Machine Memory Settings

Figure 1 shows the memory settings used for a virtual machine.

Figure 1. Virtual Machine Memory Settings



- Configured memory – Memory size of virtual machine assigned at creation.
- Active memory – Memory recently accessed by applications in the virtual machine.
- Reservation – Guaranteed lower bound on the amount of memory that the host reserves for the virtual machine, which cannot be reclaimed by ESX/ESXi for other virtual machines.
- Swappable – Virtual machine memory that can be reclaimed by the balloon driver or, worst case, by ESX/ESXi swapping. This is the automatic size of the per-virtual-machine swap file that is created on the VMFS file system (.vswp file).

For more information about VMware ESX/ESXi memory management concepts and the balloon driver, see the *VMware vSphere Resource Management Guide* at <http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-resource-management-guide.pdf>.

Appendix B: OS07N ESX/ESXi Performance Counters

The vSphere Guest API provides functions that management agents and other software can use to collect data about the state and performance of a VMware ESX/ESXi virtual machine. The API is part of VMware Tools that is installed in the guest operating system of the virtual machine. The SAP operating system collector agent `saposc01` has been updated to call the vSphere Guest APIs to extract virtual information for presentation in transaction OS07N. Figure 2 shows a screen capture of OS07N with the virtual counters.

Figure 2. Screen Capture of OS07N after Application of SAP Note 1409604

The screenshot shows the 'Operating System Monitor' interface. The main window displays a 'Snapshot Overview' for 'WHO tardis_WHO_00 Windows NT 5.2.3790 SP 1 TARDIS AMD64 Level 6 (Mod)'. The table below lists various performance counters categorized by 'Monitoring Category'.

Monitoring Category	Description	Value	Unit
CPU Virtualization Host	Physical CPUs Used for Virtualization	0,09	
	CPU Time Spent on Virtualization	849.137,29	s
CPU Virtualization Virtual System	Minimum CPUs Available	0,00	
	Maximum CPUs Available	Unlimited	
	CPU Time Spent for This Virtual Server	15.758,72	s
	Time Virt. CPU Not Backed by Hypervisor	4.399,74	s
	Physical CPUs Used for Virtualization	0,02	
	Physical CPUs Used for Virtualization	0,02	
Memory	Physical memory	7.929.096	KB
	Free physical memory	3.187.728	KB
	Pages in	0	/s
	Paged in	0	KB/s
	Pages out	88	/s
	Paged out	352	KB/s
	Configured swap size	11.265.732	KB
	Free swap size	1.246.352	KB
Memory Virtualization Host	Phys. Memory Used by Virtual Systems	10.188	MB
	Memory Shared Between Virtual Servers	5.302	MB
	Memory Paged by Virtualization Platform	0	MB
	Memory Available	20.965	MB
Memory Virtualization Virtual System	Minimum Memory Available	8.000	MB
	Physical Mem. Allocated to Virt. System	7.742	MB
	Memory Used by Virtual System	560	MB
	Shared Memory Used by Virtual System	4.349	MB
	Desired Virtual Server Memory Size	4.065	MB

The following table provides an overview and description of the virtual counters that can be seen in transaction OS07N. For more information on counters listed in the table below, see SAP Note 1260719 and the vSphere Guest Programming Guide.

Table 5. Description of VMware Performance Counters

OS07N Counter	VMware Counter Description	Nearest vCenter Counter
Monitoring Category: CPU Virtualization Host		
Physical CPUs used for virtualization.	Number of logical CPUs used by the virtual machines and hypervisor.	Select host > Summary > Resources pane > CPU usage (MHz) .
CPU time spent on virtualization (cumulative time in seconds).	Total time used for executing guest OS and virtualization code for all virtual machines.	Select host > Performance tab > Advanced > Chart Options > CPU > Real-Time > Used (ms) ; measured per time slice.
Monitoring Category: CPU Virtualization Virtual System		
Minimum CPUs available.	CPU reservation setting.	Select virtual machine > Edit Settings > Resources > CPU > Reservation (MHz) . Generally not recommended, value should be zero.
Maximum CPUs available.	Max number of logical processors allowed to be used for virtual machine.	Select virtual machine > Edit Settings > Resources > CPU > Limit (MHz) . By default, this is set to "unlimited" so limit is number of vCPUs assigned to virtual machine.
CPU time spent for this virtual server (cumulative time in seconds since virtual machine start).	Time used by the guest OS and virtualization code for this virtual machine.	Select virtual machine > Performance tab > Advanced > Chart Options > CPU > Real-Time > Used (ms) ; measured per time slice.
Time virtual CPU not backed by hypervisor (cumulative time in seconds since virtual machine start).	Time virtual machine is scheduled, but no CPU was available to let it run.	Select virtual machine -> Performance tab > Advanced > Chart Options > CPU > Real-Time > Ready (ms) ; measured per time slice. Use the vCenter metrics for Ready and CPU Usage to troubleshoot ESX/ESXi host CPU saturation as per VMware performance troubleshooting whitepaper.
Physical CPUs used for virtualization.	Number of logical CPUs used by the virtual machine.	Select virtual machine > Summary tab > Resources pane, Consumed Host CPU (MHz) .

OS07N Counter	VMware Counter Description	Nearest vCenter Counter
Monitoring Category: Memory Virtualization Host		
Physical memory used by virtual systems (MB).	Total physical memory allocated by all virtual machines.	Select host > Summary > Resources pane > Memory usage (MB) .
Memory shared between virtual servers (MB).	Memory savings obtained through transparent page sharing on ESX/ESXi host.	Select host > Performance tab > Advanced > Chart Options > Memory > Real-Time > Shared (KB) .
Memory paged by virtualization platform (MB).	Total amount of memory swapped out for all virtual machines on the ESX/ESXi host.	Select host > Performance tab > Advanced > Memory > Chart Options > Memory > Real-Time > "Swap out."
Memory available (MB).	Total amount of free physical memory that is available on ESX/ESXi host for virtual machines.	Select host > Configuration tab > Memory > Virtual Machines (on right)
Monitoring Category: Memory Virtualization Virtual System		
Minimum memory available (MB).	Memory reservation for the virtual machine.	Select virtual machine > Edit Settings > Resources > Memory > Reservation (MB) . To enforce no memory over-commitment, this can be set to the configured memory size of the virtual machine.
Physical memory allocated to virtual system (MB).	User configured memory size of virtual machine.	Select virtual machine > Summary > General pane > Memory (MB) .
Memory used by virtual system (MB).	Amount of memory recently accessed by virtual machine.	Select virtual machine > Summary > Resources pane > Active Guest Memory (MB) . If sum of active memory of all virtual machines on ESX/ESXi host is greater than physical memory available for virtual machines, then high likelihood of performance degradation.
Shared memory used by virtual system (MB).	Amount of physical memory associated with this virtual machine that is shared via Transparent page Sharing.	Select virtual machine > Resource Allocation > Memory pane > Guest Memory > Shared (GB) .

OS07N Counter	VMware Counter Description	Nearest vCenter Counter
Desired virtual server memory size (MB).	Size of the target memory allocation for this virtual machine.	Target value set by VMkernel for virtual machine's memory balloon size. Used by VMkernel to inflate and deflate the balloon driver for a virtual machine.
Rate of pages paged In (/s).	Amount of memory reclaimed from virtual machine by swapping guest memory to virtual machine swap file.	Select virtual machine > Resource Allocation > Memory pane > Guest Memory > Swapped (MB) . If there is no memory over-commitment, this will be zero.
Other Useful Counters:		
Not available.	System-wide memory used by hypervisor kernel (System)	Select host > Configuration > Hardware pane > Memory > System (on right)
Not available.	Memory used by each virtual machine (depends on size of virtual machine)	Select virtual machine > Summary > General pane > Memory Overhead (MB) .
Not available.	ESX/ESXi Network + Disk statistics.	See performance troubleshooting white paper: http://communities.vmware.com/docs/DOC-10352