

WHITE PAPER

INDUSTRY: ALL

SAP Solutions on VMware vSphere™ 4

Best Practice Guidelines

December 2009



Table of Contents

Introduction	1
VMware vSphere 4.....	1
SAP Platform Overview.....	2
Production Support for SAP Solutions on vSphere.....	2
Memory and Virtual CPU.....	3
Virtual Machine Memory.....	3
Virtual CPU	5
Storage and Networking.....	7
Storage.....	7
Networking	8
High Availability	8
Performance and Sizing	11
Performance	11
Performance Monitoring	12
Sizing.....	12
Timekeeping in Virtual Machines.....	14
Summary.....	15
Resources	17
Web Resources.....	17
SAP Notes.....	18
About The Author.....	19
Acknowledgments.....	19
Appendix A. Virtual Machine Memory Settings.....	20
Appendix B. OS07N ESX Performance Counters.....	21

Introduction

This paper provides best practice guidelines for deploying SAP® software solutions on VMware vSphere™ 4. 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 the Resources section for a list of these publications.) SAP also has created a variety of technical notes, published in the SAP Marketplace, for more 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.

VMware vSphere 4

VMware virtualization solutions provide numerous benefits to IT administrators and users of SAP solutions-based systems. VMware virtualization 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 may be moved from server to server with no downtime using VMware® vMotion™, which simplifies common operations like hardware maintenance and reduces planned downtime.
- **Availability** - VMware® High Availability (HA) ensures 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® Fault Tolerance (FT) features zero downtime, zero data loss continuous availability in the face of server hardware failures for any application running in a virtual machine.

VMware vSphere supports large capacity virtual machines that are especially well-suited to the memory- and CPU-intensive footprint of SAP applications. vSphere host and virtual machine specifications are as follows:

- Each VMware® ESX™ host supports up to 1 TB RAM, 64 logical CPUs and 512 virtual CPUs.
- Each virtual machine can support up to eight virtual CPUs and 255 GB of RAM.

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.

Production Support for SAP Solutions on vSphere

vSphere is certified by SAP on Linux and Windows guest operating systems as documented in SAP Notes 1122387 and 1409608 respectively. As of Q3 2009, 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 with 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 HCL.

SAP does not support or recommend running a small number of applications on vSphere in production environments:

- BPC – Business Planning and Consolidation, originally Outlooksoft, which was acquired by SAP. (See SAP Note 1098847.)
- Master Data Management. (See SAP Note 1070760.)
- SAPConsole (used to translate SAPGUI Screens to Character Based Screens compatible with handheld Radio Frequency devices).
- CRM Mobile Laptop. (See SAP Note 1336014.)

SAP acquired Business Objects, and Business Objects applications are also supported on VMware virtual infrastructure - see SAP Note 1206126. Supply Chain Management Optimizer and TREX 7.1 can be run in production on VMware virtual infrastructure, but make sure you follow the recommendations in SAP Notes 1223407 and 1303814

VMware has worked with SAP to include VMware performance counters in the SAP OS collector program “saposcol.” These ESX counters can be viewed in SAP transaction OS07N and requires 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.

Memory and Virtual CPU

Virtual Machine Memory

This section provides guidelines for determining the number of virtual machines on a single ESX host system based on memory requirements. See Appendix A for a description of virtual machine memory settings discussed in this section. For further background on VMware memory management concepts, refer to the *VMware vSphere Resource Management Guide*.

Since 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 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. A virtual machine can only be migrated if the target ESX 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 will consume less storage).
- It is important to “right-size” the configured memory of a virtual machine. Memory will be wasted if the SAP applications are not utilizing the

configured memory. ESX performance counters can be used to determine actual memory usage. (See Appendix B.)

- 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 host based on the following formula:

Memory available for SAP virtual machines = [total ESX server physical memory] - [memory required by ESX] - [user-defined “memory buffer”]

Memory required by an ESX 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 transparent page sharing makes more physical memory available, but this additional memory is not counted here in order 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 host machine). Actual memory buffer sizes will depend on specific customer design requirements.

The guidelines described above are purposely conservative to avoid kernel swapping between ESX 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. However, it is feasible that once the SAP workload is known and predictable, if VMware vCenter™ reports that steady state active memory usage is below the amount of memory on the server, then the reservation settings may be relaxed to the steady state active memory value. This scenario is discussed in the VMworld® 2009 presentation, TA2627 - *Understanding “Host” and “Guest” Memory Usage and Related Memory Management Concepts*.

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 will be 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)

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 multi-processing) 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. vSphere supports use of up to eight 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 multi-threaded (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 ESX 4, the CPU scheduler has undergone several improvements to provide better performance and scalability; for details, see the paper *VMware vSphere 4: The CPU Scheduler in VMware ESX 4*. For example, in ESX 4, the relaxed co-scheduling algorithm has been 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 the “Performance and Sizing” section of this document. 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 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 “ESX CPU Considerations” section in the whitepaper *Performance Best Practices for VMware vSphere 4*.

Setting a CPU Reservation sets a guaranteed CPU allocation for the virtual machine. This practice is generally not recommended, since the reserved resources are not available to other virtual machines and flexibility is often required to manage changing workloads. Note that SAP has conducted tests on virtual CPU over-commitment, which is documented in Note 1122388, and shows the performance degradation inside the virtual machines is linearly reciprocal to the over-commitment. As the performance degradation is “graceful,” any virtual CPU over-commitments can be effectively managed by using vMotion to migrate virtual machines to other ESX hosts to obtain more processing power.

Hyper-threading 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 host system enabled for SMT on an 8-core server will see 16 threads that appear as 16 logical processors. Recent SAP benchmarks have been conducted on SMT enabled servers and are covered in the Performance section below.

Storage and Networking

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.

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 in the following scenarios:

- Where existing systems already make use of third party storage management software, 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, you should follow these recommendations:

- Database data files should be spread out over multiple LUNs, similar to those in native setups, following the storage vendor array guidelines for database layout, LUN and spindle configuration.
- Maintain a 1:1 mapping between the number of virtual machines and LUNs to avoid any disk I/O contention.
- A minimum of two HBA adaptors should be configured per ESX host server.

Follow the guidelines in the “Hardware Storage Considerations” and “Guest Operating Systems” sections of *Performance Best Practices for VMware vSphere 4*.

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 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 per ESX 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 2009 session TA2105 - *Virtual Networking Concepts and Best Practices* - this includes designs to efficiently manage multiple networks and redundancy of network adapters on ESX hosts.
- Follow the guidelines in the “Hardware Networking Considerations” and “Guest Operating Systems” sections of *Performance Best Practices for VMware vSphere 4*.

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. The paper *SAP Solutions on VMware vSphere: High Availability* provides a description of VMware availability technologies and scenarios where they are useful. (See the Resources section for the document link.)

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.

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

While VMware HA and VMware FT can provide ESX 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/Central Services). If application level monitoring and automatic failover is also required, then you will want to investigate using third-party clustering software. Given that there are different high availability design choices available for SAP installation on VMware virtual infrastructure, the final approach taken will depend 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. Furthermore:
 - 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 has occurred 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 the Resource 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 and Central Instance or Central Services, then clustering software can address this requirement; however, note the following:
 - Currently MSCS software is the only cluster software officially supported by VMware and MSCS clustered virtual machines cannot be migrated via vMotion or be part of a DRS cluster.
 - Symantec VERITAS Cluster Server (VCS) is supported by Symantec in VMware virtual infrastructure environments. VMware also provides best

effort support for VCS under its Third Party Hardware and Software Support Policy. (See the Resources section for support policy details.)

- You will need personnel with cluster configuration skills and may have to pay for additional cluster software license costs.

Performance and Sizing

Performance

For background on how SAP performs on vSphere, please see the paper *Virtualized SAP Performance with VMware vSphere 4*. (The web location of this document is provided in the Resources section.) This paper describes performance tests including an SAP OLTP user workload run against SAP ECC 6.0 on Windows Server 2008 in a single virtual machine. Results of these tests show:

- Almost linear scalability from one to four virtual CPUs with slightly less scalability to eight virtual CPUs.
- Virtual environments support 5 to 15 percent fewer users than physical environments, depending on the size of the virtual machine.
- An improvement of 15 to 20 percent in the virtual environment, depending on the memory model, in supported SAP users when using a server with hardware nested page tables (NPT). Hardware NPT takes care of the translation between the guest address of a virtual machine and the physical address to increase virtual performance. Implementation of this feature on AMD chips is called Rapid Virtualization Indexing (RVI) and on Intel chips is called Extended Page Tables (EPT).

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 in order to obtain the best SAP application performance on vSphere.
- If using a processor with hardware nested page tables (RVI or EPT) and Linux, choose the “STD” memory model. (See SAP Note 941735 for details of Linux memory models.)
- If using a processor with hardware nested page tables (RVI or EPT) and Windows 2008, the choice of memory model has only a minor effect on performance. In that case, follow the memory model guidelines detailed in SAP Note 1002587.
- SAP OLTP workloads scale well from one to eight virtual CPUs in a single virtual machine. Hence, 2-way, 4-way and 8-way virtual machines are recommended for environments running on vSphere. Note that in ESX 3.X versions, 2-way virtual machines are still optimum.
- Install the latest version of VMware tools in the guest operating system.
- Download and check for additional guidelines in the following SAP Notes: 1056052 - *Windows: VMware ESX 3.x or vSphere configuration guidelines*; 1122388 - *Linux: VMware ESX Server 3 configuration guideline*.

Benchmark test results for SAP ERP running on vSphere are available here:

<http://www.sap.com/solutions/benchmark/sd2tier.epx>

At that location, search for “ESX.”

As of October 2009, the following certifications exist:

- Certification 2009028: 6250 SAPS for a 4-way virtual machine
- Certification 2009029 : 11230 SAPS for an 8-way virtual machine

Note: SAPS stands for SAP Application Performance Standard, a hardware-independent unit that describes the performance of a system configuration in the SAP application environment.

Performance Monitoring

For performance monitoring, ESX performance counters are available in SAP transaction OS07N (after application of SAP Note 1104578). Appendix B 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 the Resources section for this document’s web site location.) This paper provides a guide for checks including:

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

Access to VMware vCenter Server via the VMware® vSphere client is required to view the major ESX performance counters necessary for troubleshooting of CPU, memory, storage, and network issues.

Sizing

SAP has established a sizing process with its hardware partners to determine the hardware requirements necessary to implement an SAP system. The sizing process uses the web-based Quick Sizer tool, which calculates SAPS 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. You can find further information about the SAP sizing process at the following location:

<http://service.sap.com/sizing>

Note that 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 they will 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.
- You can use public-certified 4-way or 8-way vSphere benchmark results for an approximate sizing estimate, but note the following:
 - Generally, a direct benchmark comparison between a VMware virtual platform and the equivalent physical implementation on the same server may not be possible, as the virtual result is based on 90 percent or greater CPU utilization within the virtual machine and the underlying physical ESX host may not be fully utilized.
 - A certified 8-way virtual machine benchmark on an SMT enabled 8-core/16 thread server uses 90 percent or more of the CPU cycles within the virtual machine, but the physical ESX host CPU capacity may be less than 90 percent utilized due to the increased number of threads available with SMT. In such a scenario, the ESX host system can support more virtual machines.
- 2, 4, and 8-way virtual machines are common with vSphere.
- For a non-SMT enabled server, the virtual CPUs are sized on the ESX host so that:

[total number of virtual CPUs for initial sizing estimate] <= [the total number of physical cores on non-SMT enabled server]

The rule above applies to the initial sizing estimate only. After monitoring the actual workload, virtual CPU over-commitment is possible if CPU utilizations are low, that is, the total number of virtual CPUs can exceed the total number of physical cores on the server.

- For an SMT-enabled server (for example, an 8-core server with 16 threads/logical cores), performance improvement due to SMT allows for more virtual CPUs to be used than the number of physical cores that exist. In this case:

[total number of virtual CPU for initial sizing estimate] <= [effective number of logical cores on SMT-enabled server]

The effective number of logical cores is not automatically twice the number of physical cores, but will depend on the actual SAPS performance improvement due to SMT. This guideline applies to the initial sizing

estimate only. After monitoring the actual workload, further virtual CPU over-commitment is possible.

- 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 4-8 GB per core requirement translates to the same 4-8 GB requirement per virtual CPU.
- To determine the number of virtual machines for an ESX server host, from a memory standpoint, follow the guidelines provided in the “Memory and Virtual CPU” section.

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 may run either 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, since these operating system versions allow the frequency of timer interrupts to be reduced.
- Follow the guidelines in SAP Note 989963 - *Linux: VMware timing problem*.

Summary

This document provides best practices for running SAP solutions on VMware virtual infrastructure, organized by category: memory, virtual CPU, storage, network, performance and sizing, and timekeeping, as summarized here:

- **Memory** – the memory-intensive nature of SAP applications on the ABAP and JAVA stack warrants a conservative approach to virtual memory sizing for mission-critical installations. Memory over-commitment is not recommended; this policy can be enforced by setting memory reservations for virtual machines.
- **Virtual CPU** – enhanced scheduling functionality in vSphere has produced almost linear scalability of SAP workloads in a virtual machine using from one to eight virtual CPUs. CPU reservations are not required and CPU over-commitment is possible. SAP's own tests (as per SAP Note 1122388) describe graceful and predictable behavior when over-committing virtual CPUs.
- **Storage** – It is possible to mix RDM and VMFS disks in SAP environments. Mission-critical production SAP virtual machines should follow a 1:1 LUN mapping to avoid disk I/O contention and LUN and spindle design should follow the same guidelines as in physical environments.
- **Network** – follow standard VMware best practice guidelines documented in VMware whitepapers and VMworld presentations. Use the recommended minimum of four network adaptors per ESX host and the VMware VMXNET adaptor for best performance.
- **Availability and Fault Tolerance** – VMware HA and VMware FT can provide economical options to protect against ESX host server hardware failure as they are easy to configure without the complexity of installing clustering software. The lightweight SAP Central Services component is a good candidate to be protected with zero-downtime by using VMware FT in a 1-way virtual machine.
- **Performance and Sizing** – 2, 4, and 8-way virtual machines can be used in vSphere with good performance scaling. Use the SAP Quick Sizer tool as you would with physical infrastructure to generate the business requirements in SAPS. Work with the SAP Competency Center of your hardware vendor for a detailed architecture design. Use only the number of virtual CPUs needed in a virtual machine – the number of virtual CPUs used is initially based on sizing calculations, but can be adjusted after monitoring the actual workload.
- **Virtual Counters** - Transaction OS07N enables monitoring of virtual counters and provides an initial overview of the virtual environment; for more thorough performance monitoring of ESX host servers, access to counters in vCenter is required.

- **Timekeeping** - Time drift can occur in Linux based guest operating systems. To avoid this issue, use Novell SLES 9 or later, or Red Hat RHEL 5.1 or later versions of Linux; follow the procedure in SAP Note 989963.

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, once 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 standard SAP Basis administration tasks and procedures apply.
- The following SAP Notes provide an overview of technical best practices: 1056052 - *Windows: VMware ESX 3.x or vSphere configuration guidelines*; 1122388 - *Linux: VMware ESX Server 3 configuration guideline*.
- 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 an ideal place for discussions and technical questions.

Additional papers and web sources listed in the Resources section should be consulted to provide more background and details to 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 here:

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

Resources

You can find more information about using VMware and SAP solutions via the links listed below.

Web Resources

- VMware and SAP Web site (success stories, whitepapers, technical case studies):
<http://www.vmware.com/partners/alliances/technology/sap.html>
- *SAP Solutions on VMware vSphere: High Availability*.
http://www.vmware.com/files/pdf/SAP_vsphere_high_availability.pdf
- *Virtualized SAP Performance with VMware vSphere 4*:
http://www.vmware.com/files/pdf/perf_vsphere_sap.pdf
- “SAP on VMware” forum at the SAP Developer Network:
<http://forums.sdn.sap.com/forum.jspa?forumID=471>
- *Performance Best Practices for VMware vSphere 4*:
http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0.pdf
- *VMware Systems Compatibility Guide*:
<http://www.vmware.com/resources/compatibility/search.php>
- *VMware Resource Management Guide*:
www.vmware.com/pdf/vsphere4/r40_u1/vsp_40_u1_resource_mgmt.pdf
- *Configuration Maximums VMware® vSphere 4.0 and vSphere 4.0 Update 1*
http://www.vmware.com/pdf/vsphere4/r40/vsp_40_config_max.pdf
- *VMware vSphere 4: The CPU Scheduler in VMware ESX 4*:
http://www.vmware.com/files/pdf/perf-vsphere-cpu_scheduler.pdf
- *Performance Troubleshooting for VMware vSphere 4*:
<http://communities.vmware.com/docs/DOC-10352>
- VMworld 2009 session TA2627 – *Understanding “Host” and “Guest” Memory Usage and Related Memory Management Concepts*:
<http://www.vmworld2009.com/docs/DOC-3817> (VMworld account required)
- VMworld 2009 session TA2105 *Virtual Networking Concepts and Best Practices with VMware*:
<http://www.vmworld2009.com/docs/DOC-3800> (VMworld account required)

- *vSphere Guest Programming Guide*, VMware vSphere Guest SDK 4:
http://www.vmware.com/support/developer/guest-sdk/guest_sdk_40.pdf
- *Symantec VMware Partner Product Catalog - VERITAS Cluster Server*
<http://wwwa.vmware.com/partnercatalog/catalog/ViewProduct.php?pubproductid=a0450000008ZdyrAAC>
- VMware Third Party Hardware and Software Support Policy:
<http://www.vmware.com/support/policies/ThirdParty.html>
- *News on the SAP Adaptive Computing Controller: Integration with VMware Infrastructure:*
<http://communities.vmware.com/blogs/SAPsolutions/2008/10/07/news-on-the-sap-adaptive-computing-controller-integration-with-vmware-infrastructure>
- *SAP SD Standard Application Benchmark Results, Two-Tier Configuration* (includes benchmarks conducted on vSphere):
<http://www.sap.com/solutions/benchmark/sd2tier.epx>
- *TCO and ROI Analysis of SAP Landscapes using VMware Technology*
http://www.vmware.com/files/pdf/partners/sap/SAP_TCOROI_Customers_Final.pdf

SAP Notes

SAP on vSphere Notes (available at SAP marketplace - <http://service.sap.com/support>):

- 1409608 - *Virtualization on Windows*
- 1056052 - *Windows: VMware ESX 3.x or vSphere configuration guidelines*
- 1409604 - *Virtualization on Windows: Enhanced monitoring*
- 1122387 - *Linux: SAP Support in virtualized environments*
- 1122388 - *Linux: VMware ESX Server 3 configuration guideline*
- 989963 - *Linux: VMware timing problem*
- 1260719 - *SAPOSCOL: Detailed virtualization data*
- 1158363 - *"vm-support" - Exporting Diagnostic Data from VMware*
- 1303814 - *TREX 7.1: Usage of TREX on Virtual Machines*
- 1223407 - *Using the SCM Optimizer in virtual environments*
- 1070760 - *Running a Virtual Machine and MDM*
- 1173954 - *Support of Oracle for VMware and XEN*
- 1142243 - *MaxDB release for virtual systems*
- 1130801 - *DB2 LUW, DB2 z/OS release for VMware, XEN, and Hyper-V*

- **1098847** - *Virtual Machine Support (BPC)*
- **1336014** - *CRM Mobile Laptop: MDW and MRS usage on Virtual Machines*
- **1206126** - *Business Objects Support for VMware Virtualized Environments Support Policy*
- **1374671** - *High Availability in Virtual Environment on Windows*
- **1002587** - *Flat Memory Model on Windows*
- **941735** - *SAP memory management for 64-bit Linux systems*

About The Author

Vas Mitra is a Senior Solutions Engineer in the SAP Alliances organization at VMware. Vas has worked on SAP projects since 1993 as an ABAP programmer, Basis Administrator and technical SAP Solutions Architect. These roles have been with a large system integrator, end-user companies in the chemical/pharmaceutical industries and a large OEM vendor.

Acknowledgments

The author would like to thank the following for their invaluable technical contributions:

- Ken Barr - Staff Engineer, Performance
- Michael Hesse - Staff Systems Engineer, SAP Alliances
- Todd Muirhead - Staff Engineer, Performance
- Matthias Schlarb - Technical Alliance Engineer, SAP Alliances

Appendix A. Virtual Machine Memory Settings

The following figure illustrates the memory settings used for a virtual machine.

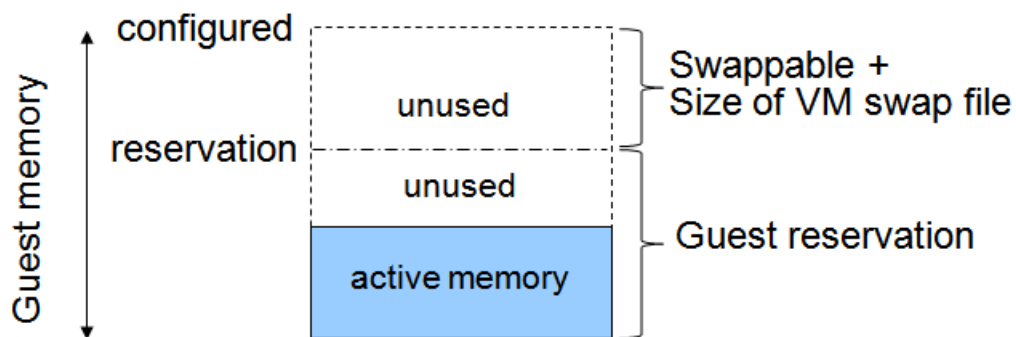


Figure 1. Virtual Machine Memory Settings

Definition of the terms used in Figure 1:

- **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 VMware ESX for other virtual machines.
- **Swappable** - virtual machine memory that can be reclaimed by the balloon driver or, worst case, by ESX 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 memory management concepts and the balloon driver, please consult the VMware *Resource Management Guide*.

Appendix B. OS07N ESX 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 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 “saposcol” has been updated to call the vSphere Guest APIs to extract virtual information for presentation in transaction OS07N. Figure 2 below shows a screen capture of OS07N with the virtual counters.

Operating System Monitor Edit Goto Environment System Help

WHO tardis_WHO_00 Windows NT 5.2.3790 SP 1 TARDIS AMD64 Level 6 (Mod)

Snapshot Overview Tue Sep 29 15:29:37 2009 Interval 10 sec.

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
Memory	Physical CPUs Used for Virtualization	0,02		
	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	
	Maximum swap size	4.190.208	KB	
	Actual swap size	3.771.256	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	

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

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, please consult SAP Note 1260719 and the vSphere Guest Programming Guide.

Table 1. 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 > CPU > Chart Options> 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 > CPU > Chart Options > 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 > CPU > Chart Options > Ready (ms); measured per time slice. Use the vCenter metrics for Ready and CPU Usage to troubleshoot ESX 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).
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).

OS07N Counter	VMware Counter Description	Nearest vCenter Counter
Memory shared between virtual servers (MB).	Memory savings obtained via transparent page sharing on ESX host.	Select host > Performance tab > Advanced > Memory > Chart Options > Shared (KB).
Memory paged by virtualization platform (MB).	Total amount of memory swapped out for all virtual machines on the ESX host.	Select host > Performance tab > Advanced > Memory > Chart Options > "Swap out."
Memory available (MB).	Total amount of free physical memory that is available on ESX host for virtual machines.	Select host > Summary > Resources pane > Capacity (MB) & Memory usage (MB).
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 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).
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.
Some Other Useful Counters:		
Not available.	System-wide memory used by hypervisor kernel (System) and Service Console.	Select host > Configuration > Hardware pane > Memory > Memory pane.

OS07N Counter	VMware Counter Description	Nearest vCenter Counter
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.	Network + Disk statistics.	See performance troubleshooting whitepaper: http://communities.vmware.com/docs/DOC-10352

