



THE  
POWER  
TO KNOW.

---

*Technical Paper*

**Moving SAS Applications from a Physical to a  
Virtual VMware Environment**

*Best Practices and Performance Expectations  
March 2011*

---



---

## Table of Contents

---

<b>Introduction .....</b>	<b>1</b>
VMware Virtual Infrastructure Overview .....	1
Virtual Performance Expectations .....	1
SAS Performance Caveats.....	2
<b>CPUs and vCPUs.....</b>	<b>2</b>
VMware Virtual Machine Limits and Restrictions .....	3
<b>Storage (and I/O Throughput).....</b>	<b>3</b>
VMware Volume/LUN and File Size Limits.....	4
<b>RAM .....</b>	<b>4</b>
VMware RAM Configuration Limits.....	5
<b>SAS Tier Performance on VMware .....</b>	<b>5</b>
<b>VMWare vMotion .....</b>	<b>6</b>
<b>Diagnosing Performance.....</b>	<b>6</b>
<b>Best Practices Summary .....</b>	<b>7</b>
General Best Practices .....	7
Hardware Level and Operating System Best Practices .....	7
<b>References and Additional Information .....</b>	<b>9</b>

---

## Introduction

---

A physical SAS environment can be migrated, with a properly planned architecture, to use virtual infrastructure provided by VMware vSphere 4. When planning a move to a VMware virtual infrastructure, there are critical system details to consider. Thus, the following best practices in implementing a virtual solution and setting clear expectations for its operation and performance, are essential for successfully deploying SAS®9 solution in a VMware xSphere environment.

---

## VMware Virtual Infrastructure Overview

---

VMware vSphere 4 provides ESX and ESXi “bare-metal” hypervisors that install directly on top of a physical server and partition it into multiple virtual machines that can run simultaneously, sharing the physical resources of the underlying server. Each virtual machine represents a complete system, with processors, memory, networking, storage and BIOS, and can run an unmodified operating system and applications.

**NOTE:** The functionality and performance of VMware ESXi and ESX hypervisors are the same; the difference between the two hypervisors resides in their architecture and operational management. VMware ESXi is the latest hypervisor architecture from VMware and as of the vSphere 4.1 release, VMware’s recommended best practice when deploying VMware vSphere. It has an ultra thin footprint with no reliance on a general-purpose OS, setting a new bar for security and reliability. The small footprint and hardware-like reliability of VMware ESXi hypervisor architecture enable it to also be available preinstalled on industry standard x86 servers.

Virtual machines that reside on the same physical server share underlying hardware resources, but are completely isolated from each other as if they were physically separated. This means that if one virtual machine experiences problems, it will not affect applications running on the other virtual machines on the same server. In the event that the underlying hardware itself experiences performance or other problems running an application, VMware provides the ability to migrate entire virtual machines automatically from one physical server to another with little or no interruption in service. This capability holds enormous advantages over physical infrastructures when it comes to disaster recovery, but VMware also leverages virtual machine migration to enable exclusive capabilities such as dynamic policy-based allocation of hardware resources.

---

## Virtual Performance Expectations

---

There are additional resource requirements when moving from a pure physical to a virtualized environment. These can impact overall performance if not compensated for in provisioning virtual infrastructure to handle specific application workloads. For example, Foundation SAS 9.2 performance testing has shown that there may be as much as a 3 to 7% overall performance reduction when moving to a virtual environment with the same hardware and storage configuration. When multiple virtual machines are operating on the same host, performance reductions may also occur, for example, when the applications in all the virtual machines are active at the same time and the demand for the resources exceeds the available resources. Configuration tuning, following best practices, and using the latest versions of VMware software (version 4.1 as of this paper’s publication) can eliminate or help minimize any reductions in performance that might otherwise occur for whatever reason.

**Performance Implications for CPUs** – VMware solutions have the ability to balance processor loads in a highly efficient manner, and VMware virtual machines can also fully leverage multi-cores and multi-processor configurations, usually making it possible to run processor-intensive workloads on virtual machines without adversely affecting application performance. Also, using the latest server processors with hardware virtualization assist technologies also decreases any

impact of CPU virtualization on performance.

**Performance Implications for Memory (RAM)** – Virtualization requires only small amounts of additional memory, so it does not significantly increase the amount of RAM required to run an application, its host operating system, or any other supporting software running in a virtual machine. However, memory is often a deciding factor in determining the total number of virtual machines that you can consolidate on to a single physical server, since memory is allocated separately to each virtual machine. VMware technology offers advanced memory management mechanisms such as transparent page sharing, ballooning, memory compression and swapping that automatically expand or contract the amount of memory allocated to each virtual machine as application loads increase and decrease. This capability lets you achieve a higher level of server consolidation than is possible with traditional static virtual memory.

**Performance Implications for Storage** – When multiple virtual machines are consolidated on a single physical server, they can impact I/O performance and throughput of individual virtual machines with their simultaneous need for rapid access to stored data. VMware vStorage provides near native I/O performance by providing virtual machines with simultaneous access to shared data stores. Centralized storage helps reduce latency and increase throughput, and provides the foundation for unique capabilities such as virtual machine migration and consolidated backup.

**Performance Implications for Networking** – Networking sizing and performance considerations in a virtual infrastructure are very similar to networking considerations in physical IT environments. In most cases, network throughput of virtualized workloads is comparable to the network throughput of the physical workloads. VMware offers an ideal platform for secure, high-speed networking between virtual machines on a single physical server, supporting network topologies that normally depend on the use of additional hardware to provide security and isolation. You can also network virtual machines across physical servers with transparency and high throughput, as each virtual machine gets its own IP address and can utilize up to four virtual network interface cards (NICs).

## SAS Performance Caveats

---

SAS applications are supported in VSphere-based virtual environments and vSphere can support the performance requirements of most SAS applications. Refer to [SAS® Product Support for Virtualization Environments](http://support.sas.com/techsup/pcn/virtualization.html) (<http://support.sas.com/techsup/pcn/virtualization.html>) for additional information. Any sizing done by SAS staff is based on physical systems **only** and does not take into account any additional resources that can be required to support the virtualized environment. So, to achieve optimal SAS software performance, it is critical that you follow the best practices in this document as well as those published in other VMware publications that provide performance best practices and tuning guidelines.

The following sections describe specific best practice guidelines for the main computer resources used by SAS – CPUs, Storage and RAM.

---

## CPUs and vCPUs

---

In a VMware virtual environment, there are virtual CPUs (vCPUs) that serve as the equivalent to physical CPUs(not threads)provided on the physical hardware. When trying to determine the number of vCPUs needed for deployment of a SAS application in a virtual environment, you need primarily to take into account the number of virtual machines that reside on the same physical server. Table 1 provides an estimate of the physical to virtual CPU core requirements for SAS applications. (Actual requirements are dependent on specific deployment scenarios and workloads.) Increasing the numbers of active virtual machines on the same host might actually increase the number of physical cores needed to

support a particular workload.

vCPUs	Physical Core Equivalent
4 vCPUs	~3.7
8 vCPUs	~7.5

Table 1: Converting from Physical CPU Cores to Virtual CPU Cores

These numbers are provided as an estimate for resource sizing, not a prediction of actual performance. CPU and I/O performance can be affected by activity of other virtual machines on the same physical server. So, for example, when the physical environment for a SAS application requires four physical cores, you may want to increase the vCPU count for VMware virtual machines to at least five. When an application needs more than 7.5 physical cores, you will not be able to place your application in a single virtual machine as the current maximum vCPUs allowed per virtual machine is eight.

## VMware Virtual Machine Limits and Restrictions

---

**Maximum number of virtual CPU in a virtual machine (guest operating system)** is eight for ESX 4.x (also called Virtual SMP).

**Maximum number of virtual CPU per physical piece** of hardware is 512 per host.

**Maximum number of virtual machines per host** is 320.

**Virtual CPU per physical core** is 20.

For more details about the above, please refer to the VMware “Configuration Maximum” paper at: [http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_config\\_max.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_config_max.pdf)

---

## Storage (and I/O Throughput)

---

The most significant reduction in overall system performance is the I/O throughput capabilities of the underlying storage subsystems and this is especially true in a virtualized environment. To design a storage architecture that provides the necessary throughput needed for optimal SAS solution performance, you have to understand how to properly size and configure it in a virtualized environment. Moving to a virtual environment does not reduce the I/O throughput requirement for SAS software.

As with all SAS implementations, not overloading storage subsystem in a virtual environment is critical. For example, Foundation SAS® commonly requires 50-100 MBps of sustained I/O (read and write activity to storage). SAS and VMware staff tested various VMware storage devices and different layouts. Other than following standard VMware, SAS, and operating system best practices, minimal or no extra tuning was required to meet requirements, as long as the storage was not overlooked.

NOTE: Refer to the “[Best Practices Summary](#)” and the “[References and Additional Information](#)” sections for additional information about storage best practices, performance tuning, and troubleshooting.

When alternating between RDM and VMFS storage access, there was little to no difference observed in performance. Use the recommended storage type that best fits specific deployment scenarios and requirements, for example; maximum volume and file size.

## VMware Volume/LUN and File Size Limits

---

**Maximum file size VMFS-2:** 64TB (with default 256MB block size)

**Maximum file size VMFS-3:** 2TB (with default 8MB block size)

**Raw device mapping (RDM) size:** 2GB (minus 512B)

**Volume size:** 64TB (minus 16K)

For more details about the above, please refer to the VMware “Configuration Maximum” paper at: [http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_config\\_max.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_config_max.pdf)

---

## RAM

---

Each virtual machine running in a virtual environment on a physical server consumes memory based on its configured size, plus a very small amount of additional virtualization overhead memory.

To reduce the amount of physical memory used on the underlying server host VMware ESX/ESXi provide support for various memory optimizations to dynamically reduce the amount of machine physical memory required for each virtual machine. However, with each of these techniques, there are potential impacts to overall, guest operating system and application performance.

**Page sharing** – transparently and securely share memory pages between virtual machines, thus eliminating redundant copies of memory pages.

**Ballooning** – reduces a virtual machine’s memory usage (as it approaches memory limits) by relinquishing memory pages considered least valuable. Ballooning provides performance closely matching that of a native system under similar memory constraints. However, a virtual machine’s guest operating system must be configured with sufficient swap space.

**Memory compression** – reclaims memory by compressing the pages that need to be swapped out. If the swapped out pages can be compressed and stored in a compression cache located in the main memory, the next access to the page only requires a page decompression, which can be an order of magnitude faster than disk access.

**Memory swapping** – forcibly reclaims memory from virtual machines, writing memory pages to allocated disk swap space. Note that because this may swap out active pages used by the virtual machine, it can cause virtual machine performance to degrade significantly.

When configuring virtual machines, carefully select the amount of memory that you allocate to your virtual machines. As with other computer resource best practices mentioned (CPU and I/O), over-committing physical RAM can cause significant performance degradation in all environments. This happens when all the virtual machines ask for their allotted memory maximums and this exceeds the amount of physical RAM in the virtual host.

While ESX/ESXi allows significant memory over-commitment, usually with little or no impact on performance, you should avoid over-committing memory to the point that it results in heavy memory reclamation. If the memory savings provided by transparent page sharing reach their limit and the total memory demands exceed the amount of machine memory in the system, ESX/ESXi will need to use additional memory reclamation techniques. As described above, ESX/ESXi next uses ballooning to make memory available to virtual machines that need it, typically with little or no performance impact. If still more reclamation is needed, ESX/ESXi uses memory compression and then host-level swapping, which can result in more noticeable performance degradation.

Below are the limits on virtual machine and host machine RAM configurations. For more information about ESX/ESXi memory considerations and best practices refer to the *Performance Best Practices for VMware vSphere® 4.0 Guide* available at: [http://www.vmware.com/pdf/Perf\\_Best\\_Practices\\_vSphere4.0.pdf](http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0.pdf)

## VMware RAM Configuration Limits

---

**Maximum RAM in a virtual machine:** 256GB.

**Maximum host memory:** 1TB.

For more details about the above, please refer to the VMware “Configuration Maximum” paper at: [http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_config\\_max.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_config_max.pdf)

---

## SAS Tier Performance on VMware

---

Virtualized SAS application performance is highly dependent on hardware resource availability from the host server(s) on which virtual machines are running. As such, over-committing resources can affect SAS tiers in different ways. Table 2 highlights individual resources and the impact that over-commitment of those resources will have on Foundation SAS applications and other tiers and products.

SAS Product Tier	Over-committed resource			Notes
	CPU	RAM	I/O	
Foundation SAS <sup>®</sup> (Base SAS <sup>®</sup> , SAS/GRAPH <sup>®</sup> )	Some reduction	Reduced performance	Reduced performance	Performance decreases quickly as more I/O throughput and available physical RAM is required and resources to match are not available. Entire file cache for the operating system must be able to be completely memory resident.
SAS <sup>®</sup> Advanced Analytics (SAS <sup>®</sup> Enterprise Miner <sup>™</sup> , SAS/STAT <sup>®</sup> )	Medium to heavy reduction			
SAS Middleware (includes application server requirements)	Some reduction	Some affect	Little to no affect	Ensure that it is possible to hold the entire Java Heap size and operating system in available physical RAM
SAS <sup>®</sup> Metadata Server	Some reduction	Reduced performance	Little to no affect	Make sure you can hold entire SAS metadata server repository and operating system in available physical RAM

Table 2: Impact of Over-committed Resources on Various SAS Software and Tiers

---

## VMWare vMotion

---

Overall, Foundation SAS servers require good I/O performance. When planning to move Foundation SAS application virtual machines and between different vSphere hosts using vMotion feature, ensure that the SAN or NAS system provides enough storage bandwidth for any hosts involved without over-committing hardware resources. This is particularly important when using vMotion to move virtual machines around on a vSphere cluster of physical hosts for ESX/ESXi). Using VMware monitoring and other third-party tools, pay close attention to the resources required by all the active virtual machines. These tools can help track utilization across all physical hosts and also utilization of resources on individual host servers and virtual machines.

---

## Diagnosing Performance

---

This can be difficult at times since most internal operating system tools report on the total virtual host, and not the individual virtual machines. When working with a virtual environment, you need to first look at overall performance of the virtual machine and then look at the overall performance of the virtual host. vSphere has a tool (esxtop) that can be used to monitor the performance metrics on each virtual machine. Information about how to use this tool can be found in Appendix A (page 89) of the “vSphere Resource Management Guide” ([http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_resource\\_mgmt.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_resource_mgmt.pdf)). This tool provides a window into overall VMware performance that was not required in a stand-alone, non-virtual environment.

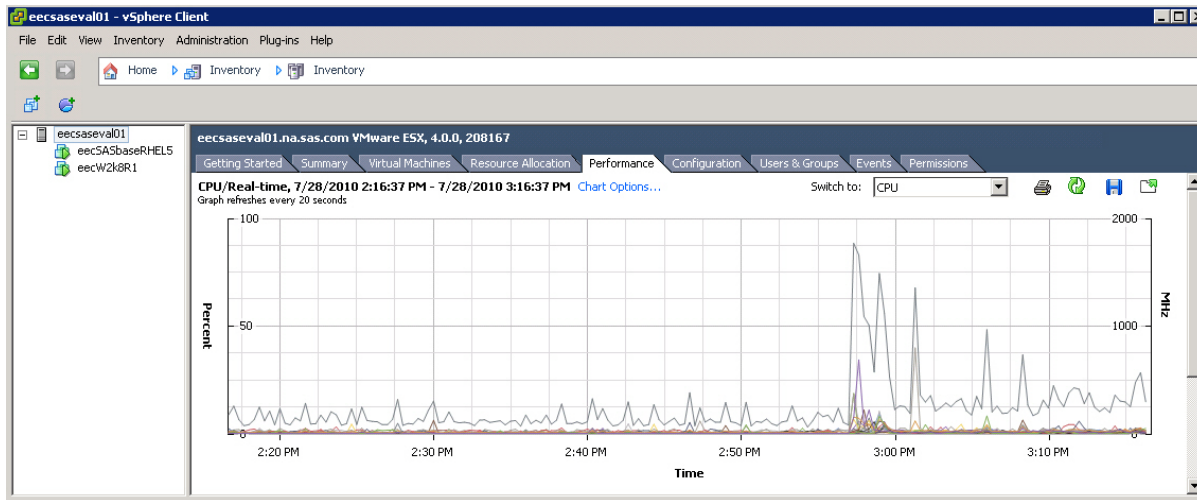


Figure 1: Performance chart on VMware Vsphere Client is a sample of the tool used to monitor VMware cloud system.

Please refer to the VMware Performance Community (<http://communities.vmware.com/docs/DOC-10352>.) for papers such as "Performance Troubleshooting for VMware xSphere 4".

---

## Best Practices Summary

---

### General Best Practices

---

Here is a summary of high level best practices for SAS applications running on VMware virtual infrastructure using VMware vSphere virtual machines.

- Do not over-commit physical RAM such that the total demand for virtual machine memory is greater than total physical RAM available on the server.
- Properly size storage environments to accommodate worst-case scenarios; like maximum required MB per second I/O throughput for a virtual host.
- Over-committing CPU resources is acceptable but closely monitor utilization and job response time.
- Design and provision physical servers and other infrastructure hardware to provide extra resources to account for virtualization workloads.
- Use VMware monitoring tools to help diagnose performance issues at the VMware ESX/ESXi host level.

### Hardware Level and Operating System Best Practices

---

Here are some additional hardware and OS-level best practices for running SAS applications on VMware virtual infrastructure.

- **Use newer generation processors (CPUs) that provide support for hardware virtualization** such as the

newer Intel processors (for example, Nehalem 2) or AMD SVM-V processors. These CPUs provide hardware-assist virtualization features that reduce virtualization overhead.

- **Turn on hardware virtualization in the BIOS** (using the BIOS setup program available on individual server computers; exact configurations varies by system and vendor). CPU hardware virtualization features, by default, might not be turned on.
- **Set CPU power management for maximum performance.** Typically, operating system parameters are adjustable with exact configuration settings varying by operating system. Newer processor can auto-adjust their clock speed to save power, but this can also impact performance for both virtual and physical systems.
- **In larger multi-CPU socket systems, VMware recommends disabling NUMA interleave.** Typically, this option is set in the BIOS. VMware ESX/ESXi also manages the NUMA non-uniform (memory access).
- **Use eager zero on the VMDKs.** This is a format for a virtual disk.(The eager zero option is selected via the command line, during virtual disk creation.) Formatting force a zero setting on all the blocks in the file system, instead of quickly allocating the disk space. Pre-allocating the blocks to zero helps performance downstream.
- **Follow all SAS best practices for the guest operating system.** (Windows 2008 R2, Linux Red Hat 5.5.) and File System (NTFS, ext3, XFS). For current best practices for various operating systems visit: <http://support.sas.com/>

---

## References and Additional Information

---

SAS Institute Inc. 2010. *Windows 2008 Configuration and Tuning Guide*. Cary, NC: SAS Institute Inc. Available at:  
<http://support.sas.com/resources/papers/WindowsServer2008ConfigurationandTuning.pdf>

This paper is a high-level guide for Windows 2008 performance tuning. When using Windows 2008 in SAS deployments in physical and virtual environments, performance tuning is especially important.

-----

Brown, Tony, Crevar, Margaret, and Ihnen, Leigh; 2007. *Best Practices for Configuring your IO Subsystem for SAS 9 Applications* Proceedings of the SAS Global Forum 2007 Conference. Cary, NC: SAS Institute Inc. Available at:  
<http://support.sas.com/rnd/papers/sgf07/sgf2007-iosubsystem.pdf>

This paper provides general guidelines for configuring I/O and storage for SAS 9 applications.

-----

Brown, Tony, and Crevar, Margaret. 2010. *Frequently Asked Questions Regarding Storage Configurations*. Proceedings of the SAS Global Forum 2010 Conference. Cary, NC: SAS Institute Inc. Available at:  
<http://support.sas.com/resources/papers/proceedings10/FAQforStorageConfiguration.pdf>

This paper addresses questions from the Best Practices paper listed above.