

TECHNICAL PAPER

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

Last update: April 2015



Contents

- Introduction.....3**
 - Important Definitions 3
 - Performance Expectations..... 3
 - SAS Performance Disclaimer..... 4
- CPUs4**
 - Estimating the Number of VMWare vCPUs Needed 4
 - Important Values and Notes on VMware vCPUs..... 5
- I/O5**
 - Important VMware Values that are Related to I/O Throughput..... 6
 - Important VMware Values that are Related to Storage Configuration..... 6
- RAM6**
 - Important VMware Values that are Related to RAM 6
- SAS Performance on VMware7**
- VMware vMotion and Virtual Cloud Performance8**
- Diagnosing Performance.....8**
- Advanced Tuning Techniques.....8**
- Best Practices Summary.....9**
 - High-Level Best Practices 9
- Resources10**

Introduction

SAS® software can be deployed in a virtual environment based on VMware® vSphere technology. During the planning phase, several architecture considerations and details are critical factors for establishing a smooth and successful deployment. This paper summarizes best practices, and outlines key considerations based on testing SAS® 9 with VMware 5.

Important Definitions

Virtual Host -- The physical hardware on which VMware ESXi software is installed. This is the machine that ESXi uses to share resources in the form of compute (RAM and CPU), network and storage with virtual machines.

Virtual Machine – This is the instance of the operating system on the Virtual Host that shares resources provided by ESXi provide other virtual machines. It might be referred to as a guest operating system instance. The physical hardware is shared with other virtual machines. This is not visible to the end user.

Performance Expectations

It has been our experience that your SAS environment may experience a slight increase in performance overhead when moving from a pure physical environment to a virtualized environment with the same hardware and storage configuration. General tuning of SAS that follows best practices for using VMware can help reduce that overhead. For example, performance testing of Foundation SAS, conducted with VMware engineers has shown that there is a 3 to 7% overall reduction in performance when moving to a virtual environment. This is due to the ESXi hypervisor that runs on the bare metal physical machine and is used to pool resources for sharing amongst all virtual machines running on that instance of the ESXi hypervisor. The physical hardware on which VMware ESXi software is installed. This is the machine that ESXi uses to share resources in the form of compute (RAM and CPU), network and storage with virtual machines.

Note that adding virtual machines to the same physical virtual host environment may trigger performance competition for resources that could result in performance degradation. This is due to phenomenon often called 'noisy neighbor' due the ESXi capability to 'overcommit' these resources since many applications require access to these resources in a random fashion. If all the applications on the virtual machines become active at the same time and the demand for resources exceeds the available resources in the virtual host, you might see symptoms of this competition for resources such as an increase in response time for an average job or during multiuser test scenarios when additional virtual machines become active. This performance reduction occurs when the hardware does not have the compute resources to handle all of the virtual machines at the same time.

VMware vSphere provides mechanisms to avoid this physical resource contention including [Resource Allocation](#), physical host and CPU affinity or simply committing only those resources that would be used by all virtual machines in a fully loaded situation. Migrating workloads (vMotion) manually to evacuate potential competition is another potentially useful technique where it can be supported recognizing that fully automated migration may cause [issues](#) in the case of CPU affinity. Newer versions of VMware are designed to help mitigate this overhead. It is best to use the latest version of VMware (at a minimum version 5.5) to enhance performance of your virtual environment.

SAS Performance Disclaimer

SAS applications execute and are supported within a virtualized environment. Due to the nature of virtualized environments, users should not expect performance equivalents in these environments. Any sizing done by SAS staff is based on physical systems only and does not take into account any overhead that can be required to support the virtualized environment. The physical hardware on which VMware ESXi software is installed. This is the machine that ESXi uses to share resources in the form of compute (RAM and CPU), network and storage with virtual machines. Refer to [SAS® Product Support for Virtualization Environments](#) for additional information. Due to the nature of virtualized environments, users should not expect performance equivalence in these environments. It is critical to follow the best practices in this document to mitigate conflicts for system resources and to achieve optimal SAS software performance. Now that we have set the basis for this paper, we will discuss some best practices for the main computer resources used by SAS – CPUs, I/O and RAM.

Virtual Machine – This is the instance Virtual machines running in VMware can support SAS applications with high-performance demands. Understanding how the ESXi hypervisor presents physical resources to the virtual machines guest operating systems is paramount in designing the vSphere infrastructure that will run SAS., As with physical devices, the stronger the technical assessment done before laying out the hardware, the better the final implementation.

CPUs

In a VMware environment, there are virtual CPUs (vCPUs) that are equivalent to physical CPU cores (not threads) on the host. Testing has shown that Intel HyperThreading, which allows for a single physical core to represent 2 threads to the operating system, provided no noticeable gains in performance and therefore should remain disabled. When trying to determine the number of vCPUs needed for a virtual environment, take into account the following:

- overhead of VMware
- management requirements
- multiple virtual machine instances

Estimating the Number of VMWare vCPUs Needed

CPU and IO performance can decline when other virtual machines in the same physical environment begin increasing activity. Over time CPU manufacturers such as Intel release processors with increasing numbers of cores as described in this intel blog posting: [How Many Virtual Processors \(vCPUs\) are Within an Intexl® Xeon® Processor?](#)

Table 1 below provides estimates of how to convert from physical cores to vCPU cores.

Table 1. Converting from Physical to vCPU Cores

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

These numbers are an estimate for budgetary sizing, not an exact prediction of true performance. These estimated values depend on the specifics of your environment. Increasing numbers of active virtual machines on the same virtual host might increase the number of physical cores needed to support a particular workload. For example, when the environment requires four physical cores, increase the virtual machines vCPU count to at least five. As more cores per CPU are introduced some applications may not appreciate linear scalability within these cores as some CPU functions are shared across all cores.

Important Values and Notes on VMware vCPUs

We have highlighted several values and notes based on VMware software that can impact your planning. It is important to keep in mind how many vCPU the virtual machine guest operating system can handle as well as how many vCPU should be provided for optimal performance.

For more details about the above, please refer to the [VMware Configuration Maximums](#) resource site.

- Maximum number of virtual CPU in a virtual machine (guest operating system) is 128 for ESX 6.x (also called Virtual SMP).
- Maximum number of virtual CPU per physical piece of hardware is 4096 per host.
- Maximum number of virtual machines per host is 1024.
- Maximum number of virtual CPUs per physical core is 32.

I/O

I/O throughput and storage configuration are crucial components for good performance with SAS. Moving to a virtual environment does not reduce the I/O throughput requirement for SAS software. Running in a virtualized environment slightly increases the I/O throughput requirement to overcome virtualization overhead.

As with all SAS implementations, avoid overloading the throughput of your I/O subsystem. Understanding application I/O throughput requirements before deploying in a virtual environment is critical. For example, Foundation SAS commonly requires 75-100MB per second of sustained I/O (read and write activity to storage). SAS and VMware staff tested various VMware storage devices and different layouts. Other than standard VMware, SAS software, and operating system best practices, minimal or no extra tuning was required. Refer to the [Best Practices Summary](#) for additional support materials. When using RDM and VMFS, there was little to no difference in

performance. Use the recommended storage type that best fits specific needs (e.g., maximum volume size, maximum file size). Understanding the application's I/O throughput requirements before deploying in a virtual environment is critical.

Important VMware Values that are Related to I/O Throughput

- Maximum file size VMFS-3: 2TB minus 512 bytes (with default 8MB block size)
- Maximum file size VMFS-5: 62TB (with default 1MB block size)

For more details about the above, please refer to the [VMware Configuration Maximums](#) resource site.

Important VMware Values that are Related to Storage Configuration

- Maximum volume size / LUN size / raw device size: 64TB (2TB minus 512 bytes Raw Device on VMFS-3)
- Maximum file system / volume size: 64TB minus 16K (VMFS-3 volumes with 1 MB block size maximum is 50 TB)

RAM

In a virtualized environment, memory is logically allocated to virtual machines as it is needed. When all the virtual machines ask for their allotted memory maximums, the demand can exceed the amount of physical RAM in the virtual host. As with CPU and IO resources, over-committing physical RAM can cause significant performance degradation in all environments. It is important to understand how the architecture of the computer used for the host allocates RAM in DIMM sockets local to the CPUs on the motherboard. This is because most modern chipsets include a provision to allow access of non-local memory, that is, memory from another CPU socket. Unless specifically provided for by the server manufacturer these non-local accesses will incur more latency and lead to performance overhead due to Non-Uniform Memory Access (NUMA). The following best practices will allow users to get the best performance from their virtual memory configuration:

- Disable node interleaving from the BIOS level of your vSphere host, which allows local memory access to have the minimum latency.
- Where possible keep the amount of memory allocated to the VM small enough that it fits within the physical memory of the sockets being used by the virtual machine. This allocation will typically happen by default, but could be more difficult if a vSphere host is configured with less memory than optimal.
- Finally, to ensure you are getting memory accesses with the minimum latency, use `esxtop` to monitor the “%local” counter on the memory screen. % local should always be 100%, or very nearly 100%. Setting memory affinity can be used in some cases to eliminate non-local NUMA access that is observed in `esxtop`, but should be used only when truly needed.

Important VMware Values that are Related to RAM

- Maximum RAM in a virtual machine: 4TB
- Maximum host memory: 6 TB (12 TB is supported on specific OEM certified platforms)

For more details about the above, please refer to the [VMware Configuration Maximums](#) resource site.

SAS Performance on VMware

SAS software performance is highly dependent on hardware resource availability on the virtual host. As described in the CPU, I/O, and RAM sections, only CPU can be over-committed effectively for SAS software. Because the SAS tiers have varied dependencies on resources, over-committing a resource can affect SAS architectural tiers in different ways. Understanding how to implement virtual machines that support SAS based on the underlying physical capabilities of the system is paramount due to the fact that any of the resources that are starved will impact the entire system. Scaling up, scaling out, scaling out 'in a box', etc. can all be viable options as there are gains to be made by creating virtual paths, such as networks, that never have to serialize data to the physical network. All of these things must be taken into consideration as the target architecture, both physical and virtual, to run SAS is planned.

Table 2 highlights the resources and the impact to SAS and other tiers.

Table 2. Performance expectations of various SAS products and tiers when certain resources are over-committed

SAS Product Tier	Over-Committed Resource			Notes
	CPU	RAM	I/O	
Foundation SAS® (Base SAS®, SAS/GRAPH®)	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® Analytics (SAS® Enterprise Miner™, SAS/STAT®)	Medium to heavy reduction			
SAS middle-tier (includes application server requirements)	Some reduction	Some affect	Little to no effect	Ensure that it is possible to hold the entire Java Heap size and operating system in available physical RAM
SAS® Metadata Server	Some reduction	Reduced Performance	Little to no effect	Make sure you can hold entire SAS metadata server repository and operating system in available physical RAM

VMware vMotion and Virtual Cloud Performance

Foundation SAS servers require I/O performance. When planning to move to virtual machines and between VMware hosts, ensure that the SAN or NAS system provides identical storage bandwidth for any hosts. Do not over-commit hardware resources.

This precaution is important when moving virtual machines around on a virtual cloud (group of virtual machines). When moving virtual machines around the cloud, pay close attention to the required resources for all the active virtual machines. It is difficult, by nature, to continue running total resources in a fluid virtual environment. Monitor the overall VMware architecture with VMware monitoring and other third-party tools. These tools help track utilization across the cloud for individual machines and the hardware resources as a whole.

Diagnosing Performance

Because most internal operating system tools report on the total virtual host and not the individual virtual machines, assessing performance and diagnosing problems can be difficult at times. When working with a virtual environment, you need to look first at overall performance of the virtual machine and then look at the overall performance of the virtual host. VMware 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 the “[vSphere Resource Management Guide](#)” (zip file). These tools will provide a window into overall VMware performance that was not required in a stand-alone non-virtual environment. Foundation SAS servers require I/O performance. When planning to move to virtual machines and between VMware hosts, ensure

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 for papers such as “[How to Troubleshoot vSphere 5.x Performance Issues Using vCenter Operations](#)” and “[Resource Management in vSphere 5](#)”.

Advanced Tuning Techniques

In very large and performance critical deployments, additional tuning may be warranted to reduce virtualization

overhead to the absolute minimum.

- In such cases, it is possible to allocate additional resources persistently to virtual machines running SAS. This step can reduce the benefits of virtualization but potentially improve overall system throughput. Persistent allocation of resources is typically done by first fully reserving system CPU and memory for a given VM (i.e. reserve the entire CPU and RAM amount allocated to the VM). The hypervisor will no longer share those resources with other VMs on the system, ensuring other VMs running will have other resources to exploit. It is also recommended that thick-provisioning be used for storage allocation, so storage is fully allocated at provisioning time rather than use time, hence reducing slightly the overhead associated with IO activities.
- In most cases, running with the VMXNET3 and PVSCSI drivers will give high network and storage throughput respectively.
- It is also possible to fully allocate a hardware NIC to a given VM via NIC pass-through (also called [VMDirectPath I/O](#)). Clearly this will reduce the number of physical NICs available to the system, so should only be done in systems with a sufficient number of NICs.
- Using [SR-IOV \(Single Root IO Virtualization\)](#) on approved NICs (including Infiniband) is another way to allocate a channel of a NIC to a VM, but still allow other VMs to access parts of that NIC, though not all NICs currently support SR-IOV. Note that NIC pass-through will disable vMotion capabilities, so may make maintenance and load balancing on the system more difficult.
- Finally, for workloads that require immediate access to CPU resources, the “latency sensitive” flag of systems on vSphere 5 and above can be set too high. This does not benefit all workloads, but for those that require intermittent but immediate access to small amounts of CPU it may reduce the latency of CPU and potentially network access.

Best Practices are detailed here [Performance Best Practices for VMware vSphere® 5.5](#) and [Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs](#).

Best Practices Summary

High-Level Best Practices

- Do not over-commit physical RAM (total virtual machine RAM should never be greater than total physical RAM).
- 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. Closely monitor utilization and job response time.
- Design the system to incorporate extra resources to overcome virtualization overhead.
- Use VMware monitoring tools to help diagnose performance issues at the virtual host level.
- Use CPUs with support for hardware virtualization, common in most modern server processors. These CPU features help mitigate virtualization overhead.
- Turn on hardware virtualization in the BIOS (use BIOS set up - exact directions varies by system). The CPU feature, by default, might not be turned on.
- **Set CPU power management for maximum performance.** Typically, operating system parameters adjust. Exact setting directions vary by operating system. Newer chips can auto adjust their clock speed to save power. This can hurt performance for virtual and physical systems

Resources

For SAS best practices for the guest operating systems (Windows, Linux Red Hat Enterprise Linux), file systems (NTFS, EXT4, XFS) and storage arrays, please visit [SAS Usage Note 53873: Operating System Tuning Papers](#).

For VMware best practices please visit [Performance Best Practices for VMware vSphere® 5.5](#).

For all other hardware infrastructure and IO subsystem tuning guidelines and best practices for SAS, please visit [SAS Usage Note 42197: A list of papers useful for troubleshooting system performance problems](#).

Release Information

Content Version: 1.0 April 2015

Trademarks and Patents

SAS Institute Inc. SAS Campus Drive, Cary, North Carolina 27513

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. R indicates USA registration. Other brand and product names are registered trademarks or trademarks of their respective companies.

To contact your local SAS office, please visit: sas.com/offices

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries.
® indicates USA registration. Other brand and product names are trademarks of their respective companies. Copyright © SAS Institute Inc. All rights reserved.

