

Realtime
publishers

Resolving VMware vSphere's Six
Biggest Performance Issues
The Essentials Series

What Six vSphere Issues Most Impact VM Performance?

sponsored by



Greg Shields

Introduction to Realtime Publishers

by **Don Jones, Series Editor**

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers..... i

What Six vSphere Issues Most Impact VM Performance?..... 1

 Monitoring Behaviors to Find Performance Issues 1

 vSphere’s Six Biggest Performance Issues..... 3

 Performance Issue #1: CPU Utilization..... 3

 Performance Issue #2: Memory Utilization 3

 Performance Issue #3: Storage Utilization and Disk I/O 4

 Performance Issue #4: Application Issues 4

 Performance Issue #5: Hypervisor Problems..... 4

 Performance Issue #6: Overhead Utilization and Scalability 5

Actionable Intelligence Is More than Monitoring..... 5

Copyright Statement

© 2011 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

What Six vSphere Issues Most Impact VM Performance?

VMware vSphere is a complicated beast. vSphere is full of moving parts, deep integrations, and elements that sometimes do and sometimes don't impact each other, which makes keeping it running with best performance a never-ending exercise in capacity management. Assisting you with the task is a long list of counters. These counters, measured by vCenter Server and exposed in the vCenter Client, help you quantify the behaviors experienced by your virtual machines.

The hard part, however, is in recognizing how those counters translate into actual performance issues. With hundreds at your disposal, which counters present actually-useful information? Which ones unnecessarily muddy the waters by delivering too much? And which help you truly understand the capacity and performance issues that are truly impactful. This Essentials Series exposes ten counters at the center of those activities.

But counters themselves aren't the only factor in successful capacity management. They are, in fact, only one piece of a much greater puzzle. Also critical to maintaining your environment are effective processes; daily, monthly, and yearly activities that preserve vCenter's overarching health. This guide delivers a set of important processes you'll want to implement.

Most important, an understanding of VMware's biggest performance issues can only be achieved when you can translate the raw data it supplies into actionable intelligence. What does it really mean when *mem.active.average* reads 3822 today? Should you do something? If so, what? This guide concludes with a look at exactly that actionable intelligence every administrator really wants, showing you why you feel overloaded with data and how to glean real resolutions from raw data. Here's a hint: *The answer isn't completely in the numbers.* It's also in the human processes you must implement to control and stabilize your vCenter environment.

Monitoring Behaviors to Find Performance Issues

Before we can delve into the technical information, it is necessary to recognize the biggest issues vCenter environments face. When you step into your office on a Monday morning to find a dozen work orders and voicemails, it's your job to figure out why "the server is slow today."

That troubleshooting process has for too long been a subjective activity. Part of the reason for our gut-feeling approach to performance and capacity management has centered on the servers' lack of instrumentation. In the physical world, instrumenting a server required extra effort, additional and sometimes expensive software, and an advanced degree in statistics and data analysis. Today, even as virtualization complicates these activities through its collocation of virtual machines, it also eases performance and capacity management by automatically instrumenting virtual machine activities with a range of behavioral monitors.

Behavioral monitors for CPU utilization, memory, disk, and storage, at the host level as well as inside virtual machines create an endless supply of beautifully-complex graphs that highlight vSphere's raw experiences. Yet although vSphere's graphs—Figure 1 provides an example—are academically interesting, can you as a human divine actionable information from its jagged and overlapping lines?




Figure 1: Raw data from vSphere's performance tab.

There's an argument that you can't. With even just a few data points being collected per server, monitoring 100 virtual machines requires analyzing over a million of them per day. As an environment scales up, no human unaided can use those graphs to answer the question: *What should I do?*

vSphere's Six Biggest Performance Issues

That's why this look at vSphere's biggest performance issues must start with that reality check. Notwithstanding its size or complexity, almost every vSphere environment suffers from a similar series of issues. They exist not because there's a lack of visibility into key capacity and performance metrics. Far from it. Rather, these issues cause pain because, unaided, people can't turn that data into *actionable intelligence*.

Keep that term fresh in your mind as you read on. It's the intelligence you can perform action upon that supplies you with the answer to the question: *What should I do?*

So what are the six biggest performance issues you and your virtual administrator peers are experiencing? Take a look through the following issues to see if they relate to the behaviors you suspect are sapping performance out of your vSphere environment.

Performance Issue #1: CPU Utilization

The first of these issues is likely the one you're most familiar with. vSphere virtual machines don't work very well when they run low on CPU capacity. The linkage between CPU supply and virtual machine performance is well known and seemingly easy to track down. For most of us, seeing a host whose CPUs are constantly pegged or where CPU Ready time is high immediately points us to an overuse condition.

Yet what many don't realize is that CPU oversubscription is as much a capacity issue as one of performance. When all eight of an eight-way host's CPUs run consistently at 90%, that host *does not have the capacity to support its workloads*. Needed is additional hardware to offload virtual machines and rebalance the load.

Most environments lean on vSphere's built-in Dynamic Resource Scheduler (DRS) to automate the rebalancing on their behalf. Yet doing so without appropriate monitoring will quickly create a distributed capacity shortfall as the environment grows. In essence, you'll load balance yourself to a cluster-wide capacity shortfall as you keep adding virtual machines. VMware vSphere's number one performance issue happens when highly-automated environments can't plan for that situation before it happens.

Performance Issue #2: Memory Utilization

Issue number two is slightly more difficult to spot. The reason for this difficulty lies in VMware's much-touted memory over-commitment capabilities. With them, virtual machines on a host can be assigned more memory than is physically available on that host.

Although great for consolidation and an absolute boon to DRS' rebalancing activities, over-commitment is in reality a situation you should avoid whenever possible. Avoiding over-commitment means not forcing vSphere to engage in its "extra" memory management activities that facilitate the sharing of memory. Those activities consume unnecessary resources that will impact performance.

A much better approach is to right-size assigned memory to actual virtual machine requirements, particularly taking into account the memory that's needed at peak usage times. This is obviously a difficult task without effective monitoring that provides actionable intelligence to tell you how to adjust your configuration.

Performance Issue #3: Storage Utilization and Disk I/O

A growing source of concern in the virtual world is the impact of storage on overall virtual machine performance. IOPS and total storage throughput are measurements you've probably been hearing about very recently; yet, its bearing on performance is only now being recognized as extremely important.

Today's vSphere fails in this regard. Its counters do not deliver information about storage performance in an easy-to-understand format. Unless you're skilled in reading vSphere counters and relating them to your environment composition, you're not likely to glean information from their data that can tell you what to do.

Storage utilization is a better-understood topic, although not one that is necessarily well-alerted inside vSphere's interface. Complicating the situation is the vast array of storage options available inside the typical data center. vSphere alone does not do a good job of helping you understand which of your array of storage options makes best sense for virtual machine location. Factoring in cost, capacity, and even IOPS into this calculation requires extra effort or outside support.

Performance Issue #4: Application Issues

CPU, memory, and storage are often treated as aggregate counters once virtual machines are virtualized. Knowing that you have some number of megahertz of processing capacity is useful for planning what future date more supply must be purchased. But, as is often said, *virtualization rewards smart administration*, meaning that a smarter workload configuration results in needing to buy less hardware.

vCenter's instrumentation into a virtual machine's behaviors is limited to the aggregate behaviors exposed at the host. However, sometimes a drain on capacity has little to do with the virtual machine itself and more to do with the workload on top.

Untamed applications, particularly when combined with over-allocated resources, can have a deleterious effect on total capacity. Consider the poorly-tuned database or middleware application that consumes every resource available. Leveraging smart monitoring tools that inform when these situations occur can help the virtual administrator finely tune the application rather than resorting to other, more brute-force approaches.

Performance Issue #5: Hypervisor Problems

Although today's hypervisors are mostly bomb-proof, they aren't completely devoid of issues. Many of those issues are in fact created by well-meaning administrators. A hypervisor that has been asked to do too much will suffer an excessive loss of performance. One whose virtual machine communication channels (*vis-à-vis* the VMware Tools among others) are severed, disabled, or non-present causes extra unnecessary work.

The kinds of actionable intelligence an administrator desires helps identify when hypervisors aren't configured appropriately. That same information alerts them when too many CPUs, memory, or other resources are assigned to virtual machines. It also sends up red flags when communication paths are not available or optimized.

Performance Issue #6: Overhead Utilization and Scalability

Lastly and most importantly is getting one's arms around the activities of the entire data center. Today's "sweet spot" for vSphere clusters is said to range between 16 and 24 hosts. Many data centers require far more hardware than that. Optimizing performance across multiple hosts and entire clusters is a task not well visualized inside vSphere's Performance tabs alone.

When working in such a distributed environment, you will need to have visibility outside the boundary of the individual cluster to best optimize your resources. Seeing performance and capacity information that spans their boundaries helps you determine when virtual machines are best rebalanced across clusters (or even data centers).

Scalability isn't only a cluster-specific calculation. There is a certain quantity of "extra" resources that are required to manage the assigned resources of each virtual machine. These extra resources represent a drain on those that can be assigned elsewhere. As a result, oversized virtual machines tend to consume a greater level of overhead than those that are properly configured. In short, oversizing virtual machines pays a kind of double tax on available resources.

Actionable Intelligence Is More than Monitoring

Solving these six big issues requires a superior analysis of the data VMware vCenter exposes. It requires the assistive support of external services that watch the data for you, crunch its numbers on your behalf, and deliver to you actionable intelligence instead of just raw data.

Truly appreciating this statement, however, requires first a look at the counters themselves. Only by seeing the intrinsic complexity within just ten of VMware's most important counters can you truly recognize that you'll need help to answer the question *What should I do?*