Realtime publishers

Resolving VMware vSphere's Six Biggest Performance Issues

The Essentials Series

sponsored by

VKERNEL

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

Copyright Statement

Article 1: 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?

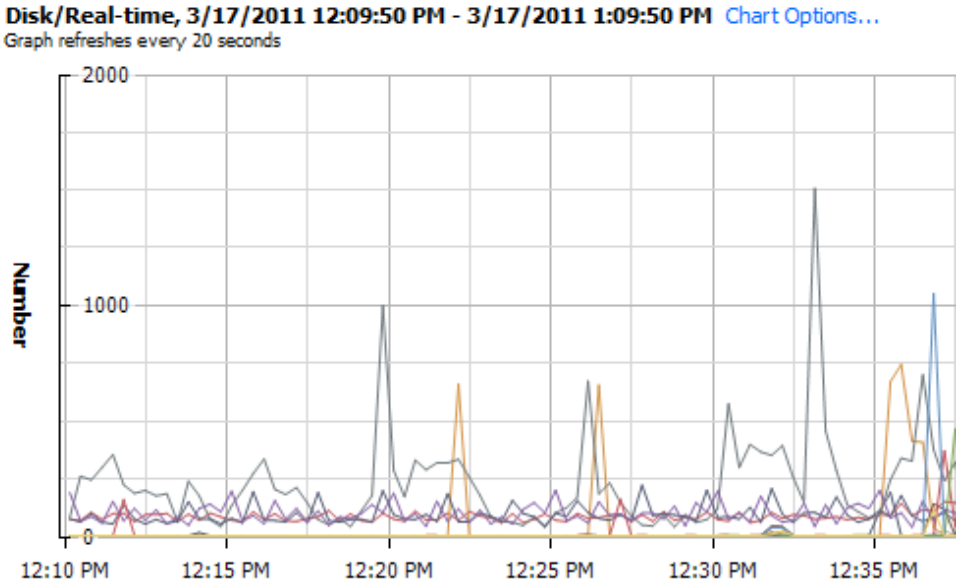

Figure 1: Raw data from vSphere's performance tab.

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?*

Article 2: What Ten Counters Quantify those Behaviors?

A virtual environment is by nature an invisible environment. You simply can't crack the case on a vSphere host and expect to "see" the behaviors going on inside. That's why its counters are so important. They represent your only way to understand the behaviors and quantify potential resolutions.

But counters by themselves are very scary things. A counter is by definition just a number. Put together enough of those numbers, and you'll create a graph not unlike Figure 1 in the previous article. Divining meaning from the points on that graph, however, is another thing entirely. Studying charts and graphs is an activity that can consume every part of your workday. With those graphs constantly evolving with a virtual environment's behaviors, just keeping up is a challenge all its own.

Yet monitoring virtual machine performance is a virtual environment's most important activity. That's why this series' previous article suggested that an unaided person can never effectively convert raw data into actionable intelligence. Oh, yes, in a tiny environment with just a few interdependencies, you probably might. But most of our VMware vSphere data centers are large and distributed. Finding the source of a performance issue isn't easy when you're starting at its unending integers.

Relating those numbers to the actions you should take is what you really want. Figure 1 shows that line of thinking in relation to the first article's notion of actionable intelligence. In it, you can see how the raw monitoring data from a vSphere system can flow through some kind of built-by-somebody-else mathematical model that converts raw data into suggested actions.

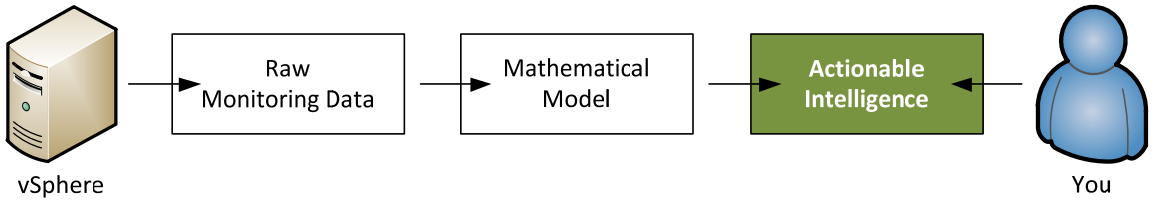

Figure 1: How actionable intelligence is generated.

What does that raw data look like? What top-ten counters might you plug-in to such a model to represent virtual machine and virtual host behavior? These ten are the topic of this second article. In the next section, you'll gain an appreciation for the amount of effort it takes to convert just ten integers into answers.

Ten Behavior-Quantifying Counters

What follows are ten counters considered to be most important in determining the rough behaviors within a vSphere environment. Think of them as an equation with ten variables. With the right equation, plugging in these values will net you an approximation of your vSphere environment's behaviors.

Counter #1: CPU.ready.summation

CPU ready time relates to the percentage of time a virtual machine was ready to use a physical CPU but could not get scheduled to run on it. A high CPU.ready.summation time means that virtual machines are waiting for physical CPU resources that aren't available. A high count here tends to indicate that physical CPU is a bottleneck to performance.

Counter #2: CPU.usagemhz.average

The CPU.usagemhz.average counter measures average CPU usage in megahertz during a configured interval. Measured over all physical processors, this counter is a primary indicator of the amount of CPU load being placed on the host.

Counter #3: mem.active.average

mem.active.average represents an estimation of how much memory is actively being used by virtual machines. This estimation is made by the VMkernel and is based on recently-touched memory pages. This memory counter references the quantity of guest memory the guest is actually using to accomplish its workload requirements.

Counter #4: mem.consumed.average

Slightly different than the mem.active.average counter, mem.consumed.average measures the amount of guest physical memory consumed by a virtual machine. When measured for a virtual machine, this level of memory includes shared memory and memory that is reserved but not used. This counter can also be measured for hosts and clusters. When measured for a host, the counter measures the amount of machine memory used on the host. For a cluster, it measures the amount of memory used by all powered on virtual machine in the cluster.

Counter #5: mem.swapped.average

Host swapping is a last-ditch approach used by the VMkernel during periods of contention to ensure virtual machines never run out of available memory. Host swapping requires transferring memory from RAM to disk, which significantly reduces its performance. The mem.swapped.average counter references the current amount of guest physical memory that has been swapped out to its swap file. A value here greater than zero can indicate that memory is a bottleneck to performance.

Counter #6: mem.vmmemctl.average

Over-commitment in vSphere environments is handled through a process called ballooning. The process reclaims unused memory from running virtual machines to make it available for others. An excessive value for mem.vmmemctl.average means that virtual machines have been assigned too much memory that they are not using but is needed by other virtual machines. A large amount of ballooning can have an impact on overall host performance.

Counter #7: disk.busResets.summation

SCSI bus resets occur when a read or write command cannot be completed within an acceptable amount of time. These resets often indicate an underlying performance issue within storage hardware and are measured using the disk.busResets.summation counter. A non-zero value here can indicate that storage is a bottleneck.

Counter #8: disk.totalLatency.average

Total latency refers to the absolute quantity of time elapsed between submitting a command to storage, processing that command, and receiving the anticipated response. Measured by the disk.totalLatency.average counter, its information helps to identify the total amount of virtual machine processing delay caused by storage hardware, and is another indication that storage is a bottleneck.

Counter #9: disk.usage.average

Related to the disk.totalLatency.average counter, disk.usage.average measures the total disk I/O rate. This information useful for identifying when storage—either within the storage itself or the connection to that storage—is a bottleneck to virtual machine performance.

Counter #10: net.usage.average

This final counter measures the combined send and receive rates for network traffic during a configured interval. The net.usage.average counter can be measured against a single virtual machine or the entire host, and is used to identify how much data traffic is passing in and out of the measured host or virtual machine.

Counters Aren't Everything

These counters might create that ten-variable equation, but they by no means fully approximate every behavior your vSphere environment experiences. The actual set of counters is far higher, making the model far more complex. That said, knowing these critical ten gets you started down the road of quantifying vSphere's behaviors. It also gives you a much better appreciation for how necessary assistive support is in resolving VMware vSphere's biggest performance issues.

That said, counters aren't everything—nor are mathematical models. Both counters and the models they feed only work when the environment remains predictably stable. Although virtualization is by nature a dynamic architecture, you can gain the right kinds of stability by observing a set of good practices. The third and final article in this series explains twelve practices that will ensure your actionable intelligence is in fact intelligent.

Article 3: What Twelve Practices Convert Raw Data into Resolutions?

Keeping eyes on ten counters for a single virtual machine isn't easy. Doing the same for dozens or hundreds of virtual machines is functionally impossible for any human being. That's why assistive tools are necessary to convert those counters' raw data into actionable intelligence. Answering that all-important question of *What should I do?* requires aligning what's going with the range of possible resolutions.

This last article was written specifically to highlight how difficult that process is with counters alone. If net.usage.average is high today but so is disk.busResets.summation, what should you do? Is the bottleneck related to network oversubscription or to a situation in your storage layer? Even worse, are both subsystems experiencing a problem, or is one problem causing the other?

Even more insidious is the situation where the issue isn't a problem at all. Instead of sourcing from some hardware shortcoming, perhaps the problem relates to another administrator's storage or networking activities. Maybe they've just begun a large and unthrottled migration of data over the network. *Numbers lie*. They do so particularly when no governance exists over the activities those numbers are measuring.

Thus, this final article intends to bring stability to your vSphere environment. Indeed that environment is highly dynamic. That's the nature of virtualization and its technologies that aggregate IT workloads. But taming its dynamics requires a set of stabilizing practices that ensure counter data retains meaning (see Figure 1). Governing your vSphere environment's activities goes a long way in ensuring its behaviors can be predictably categorized.

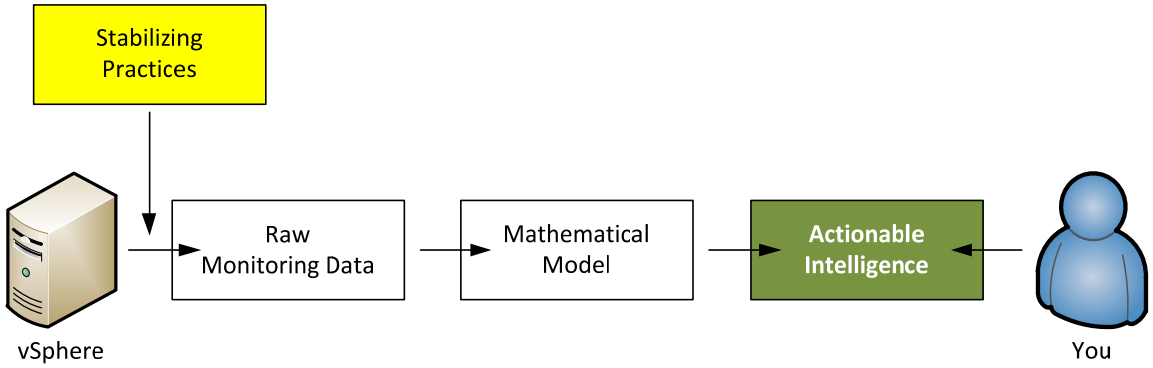

Figure 1: Stabilizing practices are the final piece to resolving vSphere's biggest performance issues.

The eleven processes described in this article assist with this task. Each deals more with the "people" side of virtualization's technology, but all are necessary to instill that predictable stability the mathematical models require.

Along with the first eleven is a twelfth and final process that involves the model itself. You could absolutely create your own model, one that takes into account each behavior's range of possibilities and its impacts on counters—but why do that? Leveraging one built by experts means immediately incorporating their experience into your vSphere environment. *In a way, it's a lot like having the world's greatest performance and capacity management experts right at your fingertips.* Consider these twelve processes as your final piece in resolving VMware vSphere's biggest performance issues.

The Daily Practices

VMware's activities represent an always-on service. Unlike a file server that can go down once in a while and not harm the bottom line too much, VMware's services are the foundation upon which all other data center activities reside. When VMware goes down, everything goes down.

That's why vSphere's daily practices deal most with monitoring. Discovering inappropriate behaviors early and before they impact users is of greatest importance in these daily practices. Doing so via a dashboard that incorporates hardware and software behaviors beneath a single pane of glass should be a desired goal. Get there by incorporating the first three practices:

Practice #1: SNMP Monitoring

VMware by default doesn't do a terrifically good job with its SNMP exposure. Enabling and tuning such monitoring requires extra steps that aren't immediately obvious within its interface. But SNMP monitoring is critically important when that single-pane-of-glass management is your desired end state. For the first practice, brush up on your SNMP technologies, or find a solution that'll automate their implementation (and, more importantly, their tuning once those technologies are in place!).

Practice #2: Resource Utilization Monitoring

Resources are constantly in flux inside vSphere. Virtual machines use more CPU for a while, then use less. Their use depends on the needs of processes and users. Monitoring that resource utilization across virtual machines, hosts, and clusters, is fundamentally important—even in the fully HA/DRS-automated environment—to get a handle on capacity issues before they impact users.

Practice #3: Alert Monitoring

vSphere by default will alert you when preconfigured conditions occur. But when was the last time you looked through its alerts? Do you know which are enabled and which are properly tuned? Have you also integrated them into your greater alert management system? If not, you'll be missing them when they're announced in the vCenter Client but nowhere else, or you'll never get them at all because they were never enabled. The practice of checking vCenter alerts on a daily basis is your first line of defense against a vSphere environment that isn't meeting the needs of its virtual machines.

The Monthly Practices

In addition to the daily practices are a set that require less-frequent attention. This reduction in frequency does not insinuate that these tasks are less meaningful, only that they require less frequent attention. Arguably, the monthly tasks comprise the more important group because they are more likely to be forgotten over time. Set up a scheduled activity on your calendar, or incorporate a solution to help you remember these five in-depth practices.

Practice #4: Disk Space Utilization

Another area where VMware vSphere has never done a terrifically good job is alerting when available disk space is low. Yet at the same time, VMware warns that a datastore that fills completely is one of the worst situations any environment can experience. You never want the situation where disk space runs out, particularly in thin provisioned environments where virtual machines think they have more disk space than they really do. That's why Practice #4 reminds you to verify your disk space on at least a monthly basis—if not every day during times when space is low.

Practice #5: Application Restarts

The virtual administrator spends so much time worrying about resources and hardware that they sometimes forget vSphere is really about the applications. Those applications sometimes experience bad behaviors like restarts during inappropriate times. Others need restarting from time to time to return them to health. Taking a monthly look at application histories and behaviors helps keep your applications in-line.

Practice #6: Server Reboots

Virtual servers too can have odd reboot requirements and behaviors. Some reboot spontaneously, giving little warning that they're about to incur a service outage. Others need regular reboots to clear memory and collected processes. Maintaining a server reboot log on a monthly basis and monitoring for resource oddities keeps their health at maximum.

Practice #7: Offline Maintenance

One class of reboots no one likes but everyone does are those surrounding updates. The monthly update cycle has become de rigueur in most data centers, with patches themselves often released on a monthly basis. Use that downtime as your opportunity to right-size assigned resource levels to the values you've determined over the past month. Also use that time to handle any special backups, snapshots, and other maintenance activities that work best while the virtual machine is offline.

Practice #8: Overall Health

Last, an overall health check is in order on a monthly basis. This health check isn't necessarily just to the virtual machines running atop vSphere but also to the vSphere environment itself. Review logs, validate load balancing effectiveness (including moving virtual machines between clusters), verify data center health, and perform all the care-and-feeding activities on a monthly basis you keep promising to do. Leaning on assistive technology here that reminds you of such health checks can be useful, particularly when that technology gives you specific advice for what steps to take.

The Yearly Practices

Your VMware vSphere yearly practices don't come around that often, but they're no less important. They center on evaluating future initiatives, incorporating feedback from monitoring solutions and users alike, and reflecting on and optimizing the processes you've laid into place. Your yearly checks are important opportunities for you to improve upon the governance activities you incorporated in the past year.

Practice #9: Budgeting for Replacement or Augmentation

The first task of capacity management is resource assurance, ensuring that virtual machines have the physical resources they need to do their jobs. The second task is planning, analyzing resource utilization over the long term to look for trends. A well-managed environment should be able to draw a straight line that begins with their historical usage and ends at the date they'll need more resources. With the right daily and monthly practices in place, Practice #9 becomes easy when your annual budget numbers are due.

Practice #10: User Feedback

That planning activity also requires an understanding of the users' experience. Your numbers may, for example, show that servers and hosted desktops are performing to expectations; however, your expectations and those of your users might be mismatched. Interviewing them and incorporating their feedback ensures that the services you're delivering are meeting their needs.

Practice #11: Large-Scale Change Initiatives

The eleventh task goes one step further. Just like your virtual environment, your business is a highly-dynamic entity. Planning for the future in a void of business initiatives means you won't have the capacity in place when the business requires it. Planning for and implementing that capacity beforehand means avoiding performance conflicts—ones that will have you updating your resume as you search for new employment—down the road.

Practice #12: The Ongoing Practice

Last is the entire reason for this Essentials Series. You as a human can't monitor all those counters alone. You need the assistive support of software solutions that translate raw data into actionable intelligence. The final practice suggests that you look to solutions that accomplish that task. Only with a little help can you tame VMware's complicated beast, full of moving parts, deep integrations, and unexpected impacts.