

# Virtual Data Center Management Challenges

## What You Need to Know

**Behavior learning or self-learning technologies**



## What You Will Learn

If you or your team is tasked with building an enterprise-scale virtual data center, you may feel like you are entering uncharted territory. You expect obstacles along the road, but you may be unaware of where they are, or how they can be avoided.

This paper offers tips and the benefit of hindsight from companies that have already done it. Corporate virtualization teams that have achieved tremendous IT cost take-out without sacrificing performance. This is the story of challenges they faced, and what they did to succeed.

## Challenges of Performance Management as Virtualization Scales

When embarking on a mission to virtualize your corporate data centers, you and your team will face a number of new challenges for the first time – both technical and organizational.

What is the new process for provisioning? Who will control, manage and allocate virtual computing power? How can we migrate from physical to virtual and back if there is a problem? Unfortunately, one challenge is often overlooked until late in the game – when it might be too late.

- *How will we monitor performance in the virtual data center?*

Why is it that people put this issue on the backburner? Often it is because early virtualization efforts only target “low-hanging fruit”, such as file servers and print servers. These servers don’t run mission-critical applications, and they don’t stress the virtual infrastructure. It is only when you begin to address the virtualization of “Tier 1” applications that you will face your biggest challenges.

The following sections detail three performance-related issues you need to understand and address to succeed in rolling out and scaling your virtual data center.

### Lack of visibility across virtual and physical infrastructure

VMware’s vCenter is a good tool that collects performance metrics and provides monitoring for VMware infrastructure, including VMs, hosts, clusters, resource pools, and data stores. If you want additional reporting or prettier consoles, any number of simple “tools” that integrate to vCenter may serve your purposes (assuming they scale). But what about other key physical components of the virtual data center? What about network and storage?

Free donuts and coffee may bring the enterprise storage team to your meetings, but it won’t change the fact that they may be unprepared for the change you are about to bring to their world. Instead of standard 60 gig LUNs (Logical Unit Number) for offline server storage, you could be asking for 300 gig LUNs for your VMware hosts. When acting primarily as file servers, enterprise storage is boringly reliable – with maybe two incidents in three years. So your storage team may not realize they are about to be hammered by the I/O load from the new virtual infrastructure. Simply put, they probably won’t be prepared for true performance management.

What about the network monitoring team? Everyone knows systems management evolved from network management, right? You will probably find the network team much more confident in their ability to build out and monitor performance of the network infrastructure needed to support your virtual data center. One of the challenges of course, is that this team probably has a handful of their own specialized monitoring tools – which you are likely to get a peek at if you promise them your first-born.



Just when you think three monitoring silos is enough, you learn about database monitoring tools, application server monitoring tools, web server monitoring tools. So even if everyone were to give you access to their monitoring consoles, do you really want to navigate that maze?

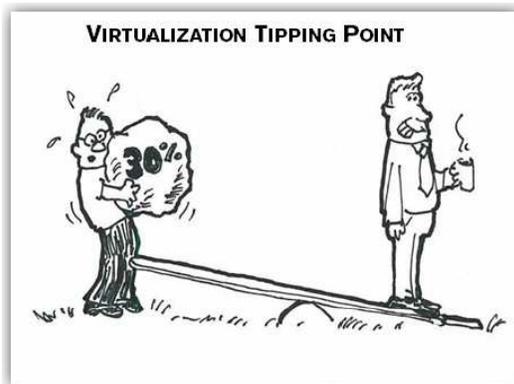
You will come to the realization that you need a solution that gives you a multi-layer view of your virtual data center – across physical and virtual infrastructure – to help you answer the simple question “how is my virtual data center doing *right now*?”

### **The “Tipping Point”: when infrastructure complexity and scale can overwhelm you**

But is better visibility enough to monitor your virtual data center?

When running smaller, lab-scale deployments of VMware, like most people you probably trust the “Mark I” eyeball as your best tool for monitoring the infrastructure. vCenter and integrated 3rd party tools may provide you with impressive dashboards and serviceable reports. But companies that have successfully deployed large scale data centers warn of a “tipping point” at which your performance monitoring capabilities may become overwhelmed.

This occurs at about the 20 to 30% virtualization mark for server or desktop virtualization.

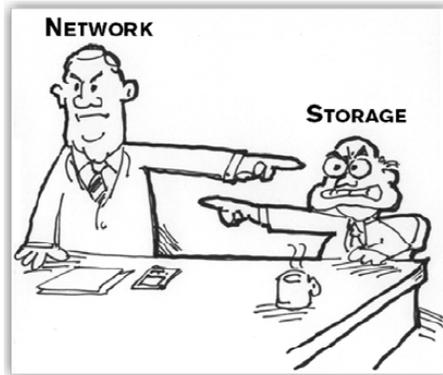


At this milestone, as you expand your virtualization efforts, you *will* get the “evil eye” from “Tier 1” application owners concerned about performance. And their concerns are not entirely unwarranted: some applications will generate I/O loads that will hammer your virtual data center architecture – probably more so than you had planned.

So in this complex environment, you have applications running on VMs, VMs moving from host to host, and hosts living in clusters with

dynamically allocated resource pools. When performance issues hit, can you *imagine* how hard problem diagnostics will be with conventional monitoring?

Is the VM the problem or the host? Or is there a problem in storage? Or in the network? Is the problem on one VM related to another? Can a series of minor recurring problems mask a looming major one?



Invariably, what happens is a phenomenon called “it’s not my fault”, where each IT silo proclaims that they do not see a problem in their monitoring tool. This phenomenon is a time honored tradition in physical data centers as well – with the difference being that your team is now invited to join the endless firefighting calls.

The reason some incidents are so difficult to resolve is that they don’t always have obvious root causes, and big problems can be caused by a number of small things that occur across different silos in the infrastructure.

The dynamic nature of resource allocation in virtual data centers further complicates the forensic process.

At this point, many virtualization teams come to the conclusion that they need a better monitoring solution. Given the scale and complexity of the problem, many go looking for technology that looks across multiple IT silos to provide the performance analytics and correlation necessary to more rapidly isolate problems.

Early investigations invariably lead virtualization teams to explore systems management and event correlation technologies from their current vendors such as IBM, BMC, HP, CA and Microsoft. But can these proverbial industry heavyweights solve the problem?

### **Failure of traditional performance monitoring tools and methods**

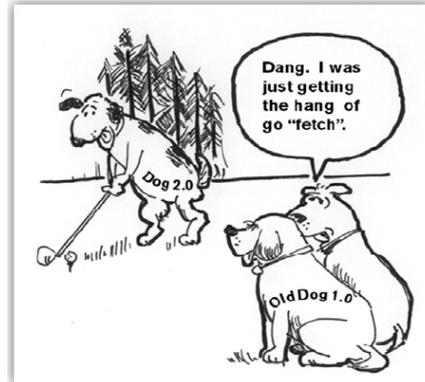
Each year, for three years running, IT operations managers attending the data center conference hosted by a major analyst firm were asked to grade their level of satisfaction with performance monitoring tools from the major vendors. And for three years running, 70% or more stated that they were “less than satisfied” with their solution.

While they excel at availability monitoring, traditional systems management tools fall short when it comes to performance monitoring. That is because they are dependent on the time consuming and error-prone process of setting manual thresholds – on metrics such as memory, storage and CPU – to alert operations staff to possible performance issues.

When multiple alerts happen, event correlation technologies try to isolate root cause by going through a set of correlation rules to arrive at a possible conclusion. This approach is flawed, however. The threshold-based alerts being correlated may *all* be false alarms, and the correlation “rules” can be too inflexible for dynamic environments. The resulting solution doesn’t work. The Operations staff gets frustrated with too many false alerts or missed alerts, and incident management remains reactive – initiated only after users call the help desk.

The problem is exacerbated by the dynamic nature of virtual data centers. To optimize resource utilization, VMware allows for dynamic resource allocation (with DRS), and the shifting of VM loads from host to host (with VMotion). It's easy to conclude that, for a large scale virtual data center, the volume and variability of performance metrics is humanly impossible to analyze using manual rules-based tools.

According to Gartner, what is needed to solve the problem is a new technology gaining broad acceptance which Gartner refers to as "behavior learning" or self-learning technologies<sup>1</sup>. These technologies use advanced mathematics and analytics to automate performance management in complex, dynamic environments, often also adding forecasting capabilities.



### **Behavior learning or self-learning technologies**

A large telecommunications company achieves significant, measurable cost savings and increased IT operations management efficiencies through the adoption of a new, emerging availability and performance management (APM) solution using behavior learning technology. IT operations stakeholders should use these tools to optimize value, for faster returns on investment (ROI) and to reduce tool maintenance costs.

#### **Key Findings**

- Behavior learning technologies have matured, allowing IT organizations to move from a reactive state focused on mean-time to repair to a proactive state where outage avoidance becomes a more realistic objective.
- Behavior learning tools augment and enhance existing APM tools.
- Behavior learning tools enable organizations to effectively transition to service-based monitoring to improve IT business alignment, versus a traditional (and less-relevant) IT element uptime focus.

#### **Recommendations**

- Consider using behavior learning products if your IT organization is looking to move from a reactive, event-driven IT operations state to a proactive one where events are detected before they impact IT services.
- Focus on a specific set of applications when implementing behavior learning software, proving its value; then use the value to justify broadening its use. Test behavior learning tools in a lab environment before purchase, using real production data to prove their abilities.
- Don't assume that IT management behavior learning tools will allow you to eliminate or replace existing APM tools; APM products are typically needed to provide the data from which learning behavior can be ascertained.

Be aware that even if you're not ready to manage end-to-end IT services, implementing behavior learning tools can help you optimize static element performance monitoring investments.

<sup>1</sup> Williams, David; Gartner Research, 15 April 2009; "Behavior Learning Software Enables Proactive Management at One of the World's Largest Telecom Companies"; <http://mediaproducts.gartner.com/reprints/netuite/167307.html>