

Demystifying Five Myths of Virtualization Management



Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations, Inc.

eG Innovations, Inc., makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Copyright

© Copyright eG Innovations, Inc. All rights reserved. All trademarks, marked and not marked, are the property of their respective owners. Specifications subject to change without notice.

Introduction

While statistics vary about just how pervasive virtualization is today, it seems clear that virtualization is becoming ubiquitous. According to Gartner Inc., International Data Corp. (IDC), and other industry watchers, it might mean that as little as half -- or closer to 60 percent -- of compute, networking, and storage resources will be virtualized.

A 2010 survey conducted by SearchDataCenter.com found that the drivers for increasing virtualization budgets are primarily based on a quick ROI: saving on hardware costs, saving on power and cooling costs and reducing data center floor space. IT shops surveyed listed the following top drivers for new server purchases:

- Increasing capacity (46%)
- Supporting new applications (38%)
- Enhancing virtualization capabilities (42%)
- Improving server efficiency (25%)
- Replacing servers hitting end-of-life (39%)
- Reducing floor space and server consolidation (24%)

While these budget-friendly attributes of server virtualization are widely accepted as fact, there is considerable confusion when it comes to the best way to approach managing virtualized IT resources for maximum availability and performance. After all, what good is saving money on hardware if you're losing control over critical business services and applications through virtualization.

Here's the core of the problem. IT professionals feel challenged to manage their **physical** IT infrastructures as those environments increased in scale and complexity. Adding virtualization to the mix just makes a difficult task even more formidable.

The good news is that IT professionals now realize that managing virtualized resources is far more daunting than managing their physical environment. But as they assess their management challenges, some misconceptions, call them myths, have emerged about the task at hand.

This white paper addresses five myths of Virtualization Management:

1. Virtualization Makes Monitoring Easier
2. Resource reservation can be used to ensure that one VM never interferes with the performance of another.
3. Virtualization is just another infrastructure tier. It can be monitored and managed independent of other software and hardware tiers supporting business services.
4. Virtualization platforms include all the tools and metrics needed to ensure that the platform is well tuned and operating as expected.
5. Tools for monitoring virtual servers are sufficient for monitoring virtual desktops because, after all, Virtual Desktops are just VMs.

We will discuss each of the five myths and then replace them with five truths of virtualization management.

Myth #1: Virtualization Makes Monitoring Easier

Virtualization reduces the number of physical servers in the infrastructure (see Figure 1). So many believe that virtualization makes the infrastructure easier to monitor and manage.

While Virtualization does reduce the number of physical servers, it does not necessarily simplify monitoring. This because of two main reasons:

- Even though virtualization reduces the number of physical servers, the number of VMs running on the physical servers still need to be monitored. The VMs support guest operating systems on which applications are running. The number of guest operating systems and applications remains the same before and after virtualization. So the **total** number of components that need to be monitored is **not** reduced as a result of virtualization.
- In fact, the number of components to be monitored in a virtual infrastructure are more than those in a physical infrastructure. This is because virtualization adds a whole new set of components and functions that need to be managed. Firstly, there is the hypervisor which is the heart of the virtualization solution. Then there are datastores which provide storage for data used by the VMs. Virtual switches offer connectivity




Figure 1 : Architecture of a Virtualization platform. Several VMs can be hosted on one physical machine.

between the VMs. VMs can be grouped into resource pools so resources can be allocated differently for different pools. To support high availability, physical machines are set up as clusters and VMs can migrate depending on the load on the physical machines from one machine to another. A monitoring solution for virtual infrastructures must monitor the hypervisors, the VMs, the datastores, the virtual switches, the resource pools and clusters.

Truth: Virtualization makes monitoring more complex!

Myth #2: Resource reservation can be used to ensure that one VM never interferes with the performance of another.

While virtualization allows the resources of the physical machine to be shared across multiple virtual machines, it also introduces the complication that if one virtual machine has a runaway job, it could impact the performance seen by other virtual machines on the same physical server. To avoid this problem, many early implementations of virtualization used resource reservation – so the minimum amount of resources needed for a virtual machine are reserved in advance.

Most real-world implementations of virtualization do not use resource reservation widely. This is while CPU and memory reservation is supported by most virtualization platforms, reservations for disk I/O and network bandwidth are not available in most virtualization platforms (VMware vSphere supports disk I/O and network bandwidth reservation from v 4.1 onwards). Secondly, statically reserving resources for a VM prevents all other VMs from using the reserved resources even if the VM that has made the reservations is idle. This significantly reduces the resource sharing benefits that virtualization offers. Consequently, in real-world implementations, static resource reservation is not common.

When multiple VMs share the resources of the same physical server, it is possible that one of the VMs may experience a surge in requests that could impact the performance of other VMs hosted on the same physical server. In the example shown in Figure 2, a media server and an Oracle database are hosted on the same physical machine. When a number of users access videos from the media server, the high load could choke the physical server, thereby slowing down all accesses to the Oracle database server as well. Thus, we can see that unlike in a physical infrastructure, virtualization introduces new forms of inter-dependencies that need to be taken into account when managing such infrastructures.

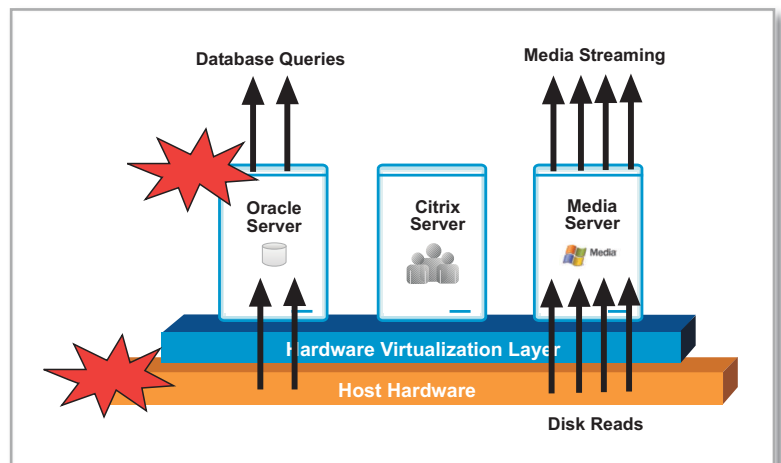


Figure 2: Unexpected load for a media server is checking the physical machine and in turn, is affecting the performance of the Oracle database server.

Truth: Resource reservation in a virtual infrastructure is not always possible nor is it efficient.

Myth #3: Virtualization is just another infrastructure tier. It can be monitored and managed independent of other software and hardware tiers supporting business services.

Enterprises have a tendency to view virtualization as yet another infrastructure silo. The virtualization team manages the physical servers and their resources and is responsible for provisioning the VMs on these servers. When a business team requires a server, the virtualization team provisions a VM and provides access to the business team. The business team is responsible for all the applications that are deployed within the VM. From the virtualization team's perspective, they only care about the resources provisioned for a VM but are not really concerned about what is happening within the VM.

Such a silo-based approach to monitoring leads to significant operational inefficiencies. Because VMs running on the same physical machine may be sharing the physical resources of the machine, a malfunctioning VM assigned to one business unit could impact the business service of another unit. To understand this better, consider the example below. Figure 3 shows the topology of a business service. In this example, the SQL database server is having a problem and as a result is impacting the web front-end.

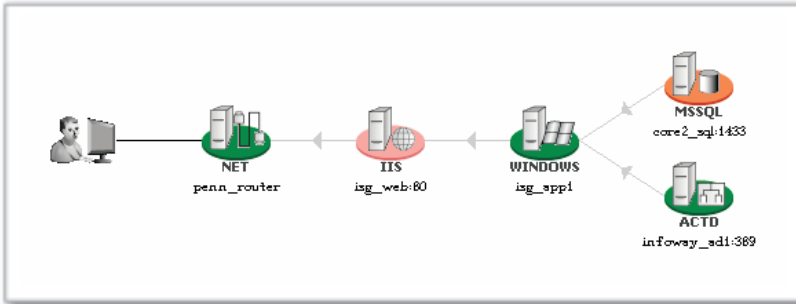


Figure 3: Topology of a business service showing the applications and devices involved in service delivering and the interdependencies between them.

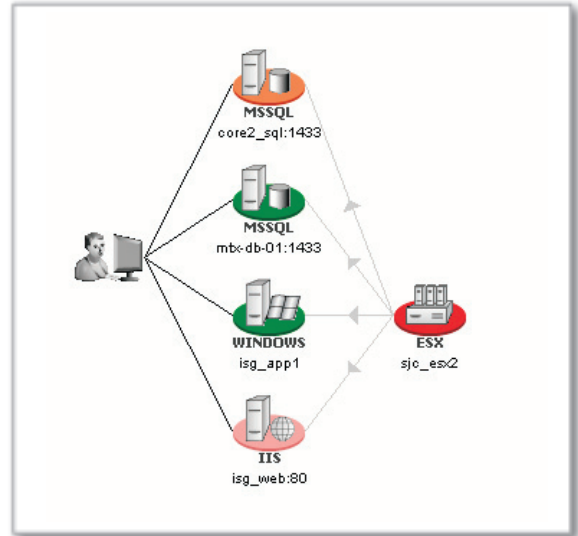


Figure 4: The virtual topology view showing the applications that are running on VMs hosted on the same physical machine.

Figure 4 shows that the SQL database is actually running in a VM that is hosted on a physical machine. This virtual topology highlights that the performance of the VMs is possibly being impacted by a critical issue with the physical machine.

Figures 5 and 6 provide additional details, revealing that a CPU bottleneck on the physical machine caused by a backup job running on the service console VM is causing the SQL database to be slow.

This example highlights the importance of managing virtualization in the context of business services it supports. With only silo views, the business team would be focusing on how to resolve the SQL server performance problem, without being aware that the issue being caused by the backup job running in another VM. On the other hand, in isolation, a backup job running on the service console is not an alarming occurrence. The fact that the backup job is running during the day and is impacting the performance seen by all VMs on the physical machine is what makes the problem significant. If the VM administrator had information about the effect the virtual infrastructure had on the business service's performance, he/she would have resolved the problem earlier to avoid any impact on service performance.

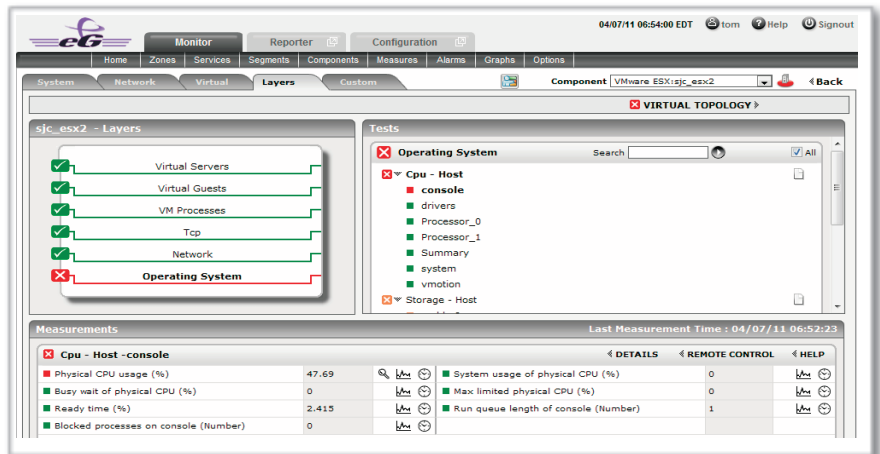


Figure 5: Details of the physical machine's performance. This figure shows that the service console is taking an unusually high 47% of the physical CPU.

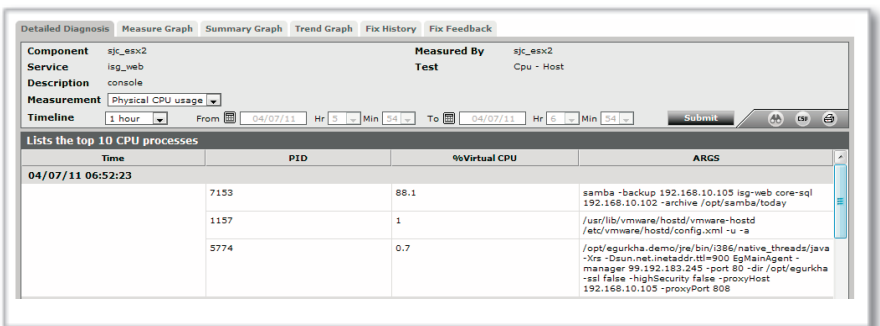


Figure 6: Detailed diagnosis revealing that a samba backup process is responsible for the CPU spike on the console OS.

Truth: Virtualization monitoring must be integrated into business service monitoring to be effective.

Myth #4: Virtualization platforms include all the tools and metrics needed to ensure that the platform is well tuned and operating as expected.

Virtualization platforms have only an “outside” view, which reveals the portion of physical resources that each VM is consuming. What the outside view does not reveal is what is happening within a VM. For example, the outside view does not indicate whether a resource spike of a VM is because of a runaway process or because there are several users accessing the VM. To be able to diagnose performance problems in a virtual infrastructure, it is essential to have an “inside” view of a VM which reveals the portion of a VM’s virtual resources that each application is taking up.

To understand the importance of the “inside view” of a VM, consider Figures 7 and 8. Figure 7 shows that a specific user’s VM is taking up a lot of CPU. This is the outside view.

Figure 8 shows that the CPU has spiked on a virtual desktop because the user was running Windows Media Player. Without an inside view into each VM, will not be able to explain why CPU was high on one desktop but not another.

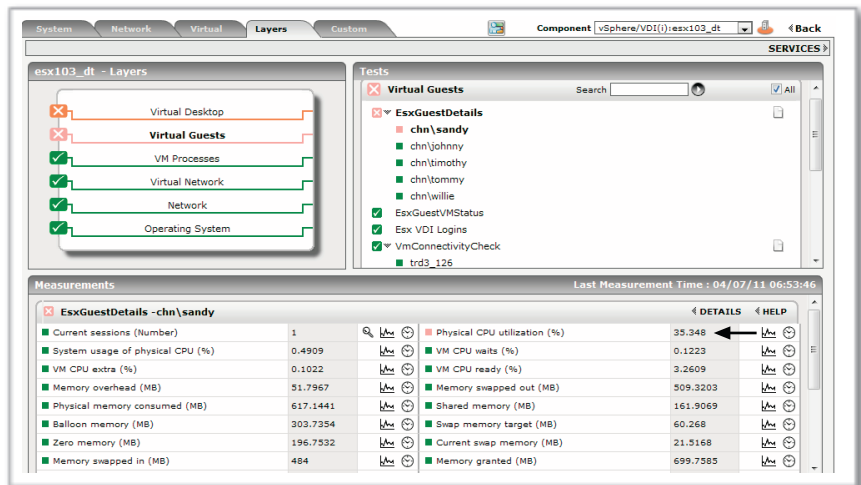


Figure 7: A specific user’s VM (sandy’s) is taking 35% of the physical CPU of a server.

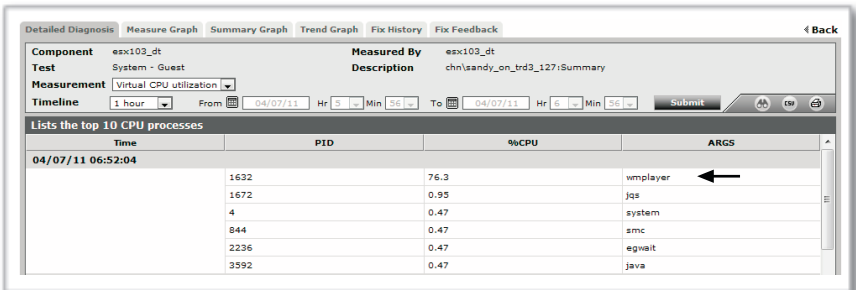


Figure 8: The Windows Media Player application is responsible for the high CPU usage on sandy’s VM.

Truth: Virtualization platform metrics show only the “outside” view of each VM. For effective diagnosis, the outside view has to be complemented with an “inside” view of each VM.

Myth #5: Tools for monitoring virtual servers are sufficient for monitoring virtual desktops because, after all, Virtual Desktops are just VMs.

First, the workload of a virtual desktop is dynamic and depends on the user who is logged in to that desktop. Therefore, VDI monitoring needs to be based on user activity - not VM activity. Second, the virtualization platform is only one of the tiers in a VDI. End-to-end VDI monitoring requires monitoring all of the components supporting the VDI, including connection brokers, terminal servers, profile servers, license servers, and so on.

And finally, while agent-based monitoring is reasonable for virtualized servers, an “agentless” method is more practical for seeing inside thousands of dynamically provisioned desktops. If you are deploying VDI, don’t try to use an existing virtual server monitoring solution as is. Rather look for a monitoring solution that is specialized to handle the unique characteristics of VDI.

Truth: A virtual desktop infrastructure is very different from a virtual server infrastructure. Therefore, to monitor a VDI, you need monitoring tools that understand the complexities of such infrastructures. Virtual server monitoring tools cannot just be reapplied for monitoring VDI.

In summary:

The Five Truths of Virtualization Management

1. Virtualization makes monitoring more challenging.
2. VMs can and will interfere with each other if the infrastructure is not carefully planned and monitored.
3. Virtualization must be monitored in the context of the business services it supports in order to be effective.
4. Management tools included with virtualization platforms must be complemented with tools that look inside the VMs to understand why resources are being used.
5. VDI monitoring is more complex than virtual server monitoring and has different requirements.

Now that you know the truth about virtualization management, you need a solution that matches reality.

The eG VM Monitor: The Complete Solution for Virtualization Monitoring

The **eG VM Monitor™**, part of the eG Enterprise Suite™, is a comprehensive solution for monitoring and managing all aspects of virtual hosts and guests, whether the infrastructure is used to support server or desktop applications. Coupled with the ability of the eG Enterprise Suite to monitor over 125 applications, 10 operating systems and seven hypervisors, the eG VM Monitor – with its patent-pending In-N-Out Monitoring™ technology -- provides a comprehensive end-to-end solution for monitoring and managing the performance of virtual IT infrastructures. Hypervisors supported out of the box include VMware, Citrix, Microsoft Hyper-V, IBM LPARs and Sun LDOMs and Solaris Containers.

Administrators can use the eG VM Monitor to monitor the performance of their physical and virtual infrastructures, troubleshoot problems to determine where the root-cause lies, assess where capacity bottlenecks are, and plan the usage of their servers and applications to optimize the utilization of the physical and virtual resources. The key customer benefits of this solution include higher uptime, better end-to-end performance, and operational cost savings through more effective utilization of key IT staff.

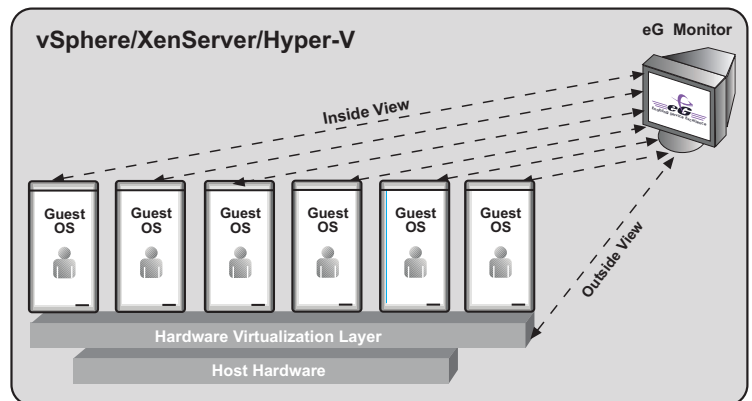
eG Enterprise is the only monitoring solution certified by both VMware and Citrix.



In-N-Out Monitoring of VMs

As mentioned with Myth #4, an “outside” view of the virtual server does not reveal why a VM is taking excessive resources. An “inside” view of each VM is needed to troubleshoot VM performance problems. Using a patent-pending In-N-Out Monitoring™ approach, the eG VM Monitor provides a comprehensive view of a virtualized servers, including the performance of the VM kernel, the service console and all of its virtual machines.

Agent-based and agentless monitoring can be used in user-defined combinations for monitoring virtualized servers. When agent-based monitoring is used, eG agents only have to be installed on the server -- not on individual guests. Using native server APIs, the agents provide an “outside view” of a guest VM’s performance. The relative



Monitoring VM guests: eG agents track the performance of each guest VM relative to shared infrastructure resources (outside view) as well as the workload and application mix of the individual guest VMs (inside view).

resource usage levels of the guest VMs show where the performance hogs exist. To complement the outside view, the eG agent obtains an “inside view” that details the user activity, resource allocation and the application mix running inside the VM guest operating system.

All the capabilities of agent-based monitoring are also available with the agentless monitoring option for VMware vSphere/ESX servers. The eG VM Monitor automatically baselines all the metrics it collects, so that IT administrators can be informed proactively of any deviations from the norm. No other virtualization monitoring solution offers this combination of features.

Monitoring Virtual Desktops

This topic is the focus of Myth #5, that monitoring virtual desktops is essentially the same as monitoring virtual servers. For reasons discussed under the “truth” behind Myth #5, this is simply not the case.

If your IT organization is looking to use virtualization technologies to centralize deployment and management of desktops, you must consider the monitoring and management challenges that you will face as you roll out your VDI. Your customers will expect their virtual desktops to be as reliable and as fast as their physical desktops. While centralizing desktops on large server farms gives you significant benefits, it also opens up a lot of challenges. A single server failure can bring down tens to hundreds of virtual desktops. Furthermore, since desktops share the same physical server resources, a single resource intensive desktop could reduce the resources available for other desktops.

To monitor your virtual desktop infrastructure effectively, you must:

- **Monitor user activity – Not just VMs:** In a virtual desktop infrastructure, the workload of a VM is dependent on which user is logged into the VM and what applications he/she is using. So monitoring virtual machines is not enough, it is important to monitor users and user activity as well.
- **Monitor activity inside the desktop:** From the virtualization platform, administrators can determine what portion of the physical resources a VM is taking up. While this “outside” view of a VM provides an indicator of the level of activity of a VM, it does not provide additional details regarding why a specific VM is taking up excessive resources. In-depth monitoring to determine which process could be responsible for the increased workload of a VM is critical for effective problem diagnosis and for ensuring peak performance.
- **Monitor VDI end-to-end:** While the virtualization platform and the virtual desktops are key components of the VDI infrastructure, the performance of the desktop service is also dependent on all the other infrastructure applications and network devices that function together to enable the service. To successfully monitor a VDI implementation, at a minimum you need to monitor all of the components that are involved in the delivery of a desktop with applications to the end user:

- | | |
|---|--|
| >> Virtual Desktops | >> Terminal servers |
| >> Hypervisors (vSphere, XenServer, Hyper-V) | >> Network routers |
| >> Connection brokers (Citrix XenDesktop, Leostream, VMware View) | >> Firewalls |
| >> Profile servers | >> Web front end servers |
| >> Licensing servers | >> Active directory |
| | >> Enterprise applications and databases |

The eG VM Monitor monitors all of these essential components to alert you to impending problems before users notice and complain.

About eG Innovations

eG Innovations, Inc. (<http://www.eginnovations.com>) is a global provider of performance monitoring and triage solutions for virtual, physical and cloud-based IT infrastructures. The company's patented technologies provide proactive monitoring of every layer of every tier in the infrastructure, thereby enabling rapid diagnosis and recovery in enterprise and service provider networks. By ensuring high availability and optimum performance of mission-critical business services, eG Innovations' solutions help enhance customers' competitive positioning, lower operational costs and optimize the performance of their infrastructures.

USA

eG Innovations, Inc.

33 Wood Ave. South, Suite 600
Iselin, NJ 08830
Ph: (866) 526 6700

SINGAPORE

eG Innovations Pte Ltd

33A Tanjong Pagar Road
Singapore 088456
Ph : (65) 6423 0928
Fax : (65) 6423 1744

UK

eG Innovations UK Ltd.

3 Grange Road, Camberley
Surrey, GU15 2DH
Ph: +44 (0) 1276 501590

Rest of Europe

eG Innovations, Europe

Montval
29 rue Leonard de Vinci
59700 MARCQ-EN-BAROEUL
France
Ph: +33 (0)3 66 64 06 16

INDIA

eG Innovations Pvt Ltd

2, Murali Street, Mahalingapuram
Chennai 600 034
India
Ph : (91) 44 2817 2801
Fax : (91) 44 2817 9041

Email : sales@eginnovations.com

Web : www.eginnovations.com

