

Bridging the Virtual Divide



With HQ, you finally have the visibility you need to achieve virtualization optimization

Host OS

Virtual Machine (Guest OS)

Server

SERVICE

SERVICE

SERVICE

Virtual Machine (Guest OS)

Server

SERVICE

SERVICE

SERVICE

Hypervisor



Understanding the Virtual Divide

Virtualization is one of the most disruptive trends in IT today. It's a safe bet that every datacenter has either evaluated or seriously considered the benefits of consolidating infrastructure around fewer physical servers. Power, cooling, floor space, and utilization are some of the key drivers for companies considering whether or not to virtualize infrastructure.

Ironically, the operational efficiency drivers that motivate virtualization create new challenges for managing infrastructure. Virtualization software (or hypervisor) carves out a logical, or virtual, machine from a physical one. This creates an artificially divided relationship between the applications inside the virtual machine and the physical resources they consume. This lack of visibility across the divide increases the propensity for configurations to change as administrators are forced to move virtual machines to improve performance rather than manage the performance of the VMs directly in context with the applications running inside them.

As a result, an IT administrator can wake up to an infrastructure that has shrunk from 100 physical boxes to 25, but grown to 250 new virtual machines (VMs)—with a greater combination of operating systems, applications and changing, complex infrastructure that needs to be managed.

Not only does this system administrator have more to manage, he must also change his management approach. He must manage the physical performance of the host, typically with the virtualization software vendor's tools, and then separately, without the ability to automatically relate it to the virtualization software performance, or manage the operating systems and applications inside the hosts.

Management of IT resources across this divide requires a new type of systems management software -- software that bridges this new virtual divide.



The Quest for Virtualization Optimization


The partitioning of a single physical server through virtualization is not a new concept.

Virtualization, and its accompanying terminology, hails from the early mainframe time-sharing days. While things have gotten quite a bit more complex since then, the core benefit remains the same: squeeze as much performance and utilization out of physical servers while allowing virtual applications to be safely isolated from each other.

The isolation brings benefits, such as the ability to relocate complete VMs seamlessly within a virtualized environment (a key element in virtualization lifecycle management). However, in order to judge a virtualization implementation a success, there must be a

measurable improvement in the utilization of physical resources and, at a minimum, no impact in the overall availability or performance of the entire environment. This is how the same aspects that make virtualization such an attractive proposition to businesses also make measuring its success very difficult.

Other Systems Management Software



Currently, the task of monitoring virtualization performance and efficiency starts and stops with the physical/hypervisor layer itself. Operators typically judge success by measuring what the virtualization software says about how much of its physical resources are being consumed and distributed by the VM.

There are two key pieces of information missing from this analysis. First, the hypervisor-level metrics only show the impact of the VM on the physical hardware, not its overall performance. Second, the overall performance and stability of the virtualized guest and applications are ignored. Consolidating 4 servers into 1 physical server with 4 VMs doesn't mean the applications are higher performing or more stable.

It's imperative to give equal (if not more) importance to the impact virtualization has on application performance as the impact it has on server utilization.

If virtualization has been around for a long time, how has it become so pervasive without an adequate management strategy? Until recently, virtualization was used primarily in QA and development departments because it offered a fluid, highly utilized environment to optimize testing permutations. Since testing environments typically do not have the same management strategies as production software, management has been limited.

Virtualization works. It increases the utilization and fluidity of applications used in production that are CPU bound, are distributed, and are created as reusable services. Popular, proven examples of this include core services such as file and print, DNS, LDAP and other services that aren't powered by complex software stacks, yet provide critical services which need to be easily replicated.

Like the adoption of other disruptive technologies such as Linux, this set of virtualized infrastructure applications is the tip of the iceberg. The next big challenge virtualization is conquering is the software stack, which powers more complex, business applications. This process has already begun with custom applications built on web technologies. The rise of SOA and composite applications has created a new distributed architecture where services are discrete, in many cases stateless, and a perfect fit for a virtualized environment.

Virtualization software has proven to be incredibly effective at marshalling CPU resources across virtual machines. This effectiveness, however, can result in an application that is starved for critical resources that affect its underlying performance.

Virtualization is a victim of its own success. The fact that these applications are so easy to propagate in an environment means this problem could repeat itself anywhere from a dozen to a few hundred times in a given environment.



Hypothetical Case Study: Virtualized Webapp

Imagine building a new CRM application using an Apache web server on the front end of a Tomcat server all running on a Linux platform with a MySQL database collecting all the data. This application is currently running on an off-the-shelf, dual CPU machine that occupies one slot on a rack deep in some sprawling data center.

The scenario is a familiar one in today's fast-moving environment. It ended up there because a project that started off as an experiment or a rapid prototype became a customer facing, business critical application.

While this application is heavily used, it doesn't consume the full set of resources the physical host provides it. This excess of resources goes unnoticed for a year is not noticed. The next year, an assessment of the amount of space and physical resources in the data

center reveals that this same thing has happened 20, maybe 50 times.

Applications get developed, deployed, and racked by the various teams that support them creating a maze of 1U boxes running singleton applications. Time for a new plan: contain the sprawl through virtualization.

The team notices that the CRM application fits the criteria for virtualization, and therefore it is slated for the new virtualized application stack. The team leading this effort creates a nice little VM for the application very quickly and efficiently and it passes basic QA testing. The VM then gets deployed on a machine running virtualization software such as VMware ESX.

In typical virtualization fashion, it gets deployed alongside 4 other VMs each running different applications with different technologies, including a J2EE app, a .NET app and mail server—each of which has its own appetite for resources.

The VMware's internal monitoring tools give the operators the impression that there's adequate capacity for all the VMs that are deployed on it. Until one day, a usage change occurs. Perhaps its quarter end and there's a huge sale, or maybe a new product release has a serious bug and the call center is suddenly flooded.

This little CRM application happens to be a critical function of the company's ERP system. End users are busy creating opportunities, managing accounts, and doing their usual activities, but at a much more rigorous pace. The heat is on and suddenly the application is struggling to keep up. Inevitably, the question becomes: "Is virtualization the cause?"

Unfortunately, this question isn't easy to answer. On one end, the virtualization software is not giving enough information about how it's allocating resources to this now very busy VM. The original assessment of whether the CRM app fit within this particular physical host was based on a completely different load profile.

However, the team lacks visibility into the performance of the different pieces of the supporting stack to determine if the application itself is the problem, or to have anticipated the performance issues before the users started complaining, regardless of whether or not the app was running in a VM or on a dedicated host. The only solution to this problem is to dynamically re-provision the VM to another, less taxed virtualized host machine.

While this alleviates the problem for the moment, the scenario can turn into the IT equivalent of a game of "whack-a-mole"—where lack of end-to-end visibility, the dynamic nature of web application load, and the sheer number of moving parts in an environment force applications to be moved from one side of virtualized infrastructure to another without solving the underlying problems.

This scenario undermines the purpose of virtualization in the data center.



Introducing Hyperic HQ to Virtualization

How are the dual challenges of manageability across virtualized infrastructure and adapting to the rapid sprawl of other virtual machines overcome?

The answer is Hyperic HQ, a new, open source approach to manageability which combines rich application management features like asset discovery, monitoring, and control with built-in support for virtualization technologies in use today, and those which will emerge tomorrow.

The Hyperic software is purpose-built for the web infrastructure powering the adoption of production virtualized environments. Hyperic starts with an inventory model that represents a server, all its applications and components within an intuitive set of relationships that apply to any combination of physical or virtual hardware, software, and services—regardless of the technologies involved.

Hyperic HQ then provides a set of management features essential to the operation of modern infrastructure:

- **Inventory Auto-Discovery** – automatically populate every aspect of the physical and virtual machines, applications, and services running in a given environment as well as the relationships between virtualized components and the physical resources they depend on.
- **Real-time and Historical Monitoring** – collect and store any availability, performance, and utilization metric for any resource in the inventory repository.
- **Event Management** - track logs and configuration changes from any resource in inventory.
- **Control** – take corrective or pre-emptive action on any application, service, or script in inventory.
- **Analyze** – combine as much information from logs, performance metrics, and infrastructure events into a set of correlated views that help you diagnose problems and determine capacity.

These core management functions are essential to the proper operation of any production IT environment, virtualized or not. The challenges of virtualization put added strain on legacy solutions designed for similar purposes.

The CRM application example above illustrates how virtualization allows for rapid provisioning, deployment, and relocation of complete application environments. Hyperic HQ's inventory auto discovery capability can detect every layer of a virtualized

environment and react accordingly when any of those resources are relocated.

In that scenario, Hyperic HQ would have automatically discovered the Linux guest; the Apache server and its virtual hosts; and the MySQL installation and its corresponding tables and indexes. These resources would have been mapped with “parent-child relationships” automatically to their counterparts on the physical hypervisor. This mapping would provide operations with a complete view of the CRM application, as well as how it relates to the physical host on which it lives. In the event that the VM for our CRM application is relocated, Hyperic HQ would be immediately aware of the change and simply update its inventory to reflect the new parent of the VM, leaving its internal components like the apache server intact.

Once the virtualized environment is properly mapped, the next challenge is to provide monitoring visibility into every layer. Hyperic HQ enables this by providing out-of-the-box support for collecting real-time and historical values of all performance and health statistics at every layer of the infrastructure. In the CRM example, Hyperic HQ would provide monitoring for the virtualization software itself.

Table - a sampling the monitoring capabilities HQ provides at each layer in our CRM application example

Layer	Supported Types	Metrics
Physical Host	- x86 - x86_64 - Sparc	- Load average - CPU - Physical Memory - Disk I/O - Disk Utilization/Capacity - Network I/O
Virtualization Software	- VMware ESX Server - Solaris 10	- Memory available for VMs - VM CPU utilization - Virtual Disk I/O - Virtual Network I/O - VM Memory Sharing
Guest Operating Systems	- Linux - WindowsXP/2003 - Server/Vista - Free/Open BSD - Sun Solaris	- Load Average - CPU - Network I/O - Disk I/O - Memory activity - Disk Capacity
Guest Applications	- Web Servers - MySQL Database - Mail Servers - J2EE Application Servers .NET - Messaging Servers - Tomcat	- Availability - Request activity - Response time - Memory consumption - CPU Consumption - Throughput - eCapacity


In order to properly assess the performance of our CRM application it's important to see all of this data in concert with each other. Before, it was known that the CPU utilization had maxed out, and it was moved to another location with more CPU without any information or analysis on what caused the change. It can only be assumed that the increased transaction load increased the appetite for CPU resources. The only fix suppressed the problem, but it didn't uncover or address the root cause, and certainly didn't optimize the virtualized environment.

By using Hyperic HQ, an administrator could see the trend of the CPU utilization on the CRM VM moved towards the maximum threshold. It's clear that the Tomcat server consumed all the free memory on the OS and stopped responding. In this case, the administrator could correlate the memory usage from the OS with the memory usage from the Tomcat server and the Tomcat webapp services. It's evident that there is one particular webapp service that corresponds to the number of requests served by the webapp. This classic example of a memory leak was always there, and was exacerbated by the increased usage.

Armed with this information, action can be taken. First, the webapp developers can be alerted to a probable memory leak in their application, and are directed to a specific service to fix. Next, instead of moving the VM to a larger location, the Tomcat server can be restarted to clear out the memory.

This is a faster course of action, and will satisfy the users while the development team estimates the fix. Additionally, alerts can be set to beep the administrator's pager in case the Tomcat's memory usage gets too high again. This solution is a good one because the long term and unexpected outages can be limited and the VM can likely continue in its current location once the fix has been applied.

Other Systems Management Software





Virtualization Optimization Success

The team has a complete picture of the entire virtualized stack, and can rationalize performance and availability of applications with the overall goal of optimizing the physical infrastructure. The team can stay ahead of the chaos by implementing a management system that reacts seamlessly to the dynamic provisioning activities common to virtualized environments.

Hyperic HQ provides a quicker path to virtualization optimization, and a better chance at closing that physical to virtual divide in IT management once and for all.