

## **IS VIRTUALIZED IT STALLED ON DELIVERING PERFORMANCE?**

**Mike Matchett, Sr. Analyst**

**FEBRUARY 2013**

One of the hardest challenges for an IT provider today is to guarantee a specific level of “response-time” performance to applications. Performance is absolutely mission-critical for many business applications, which has often led to expensively over-provisioned and dedicated infrastructures. Unfortunately broad technology evolutions like virtualization and cloud architectures that are driving great efficiency and agility gains across wide swaths of the data center can actually make it harder to deliver consistent performance to critical applications.

For example, solutions like VMware vSphere and Microsoft Hyper-V have been a godsend to overflowing data centers full of under-utilized servers by enabling such high levels of consolidation that it has saved companies, empowered new paradigms (i.e. cloud), and positively impacted our environment. Yet virtualization projects tend to stall when it comes time to host performance-sensitive applications. Currently, the dynamic infrastructures of these x86 server virtualization technologies don’t provide a simple way for applications to allocate a “real performance level” in the same manner as they easily allocate a given capacity of virtual resource (e.g. CPU, disk space). In addition, virtualization solutions can introduce extra challenges by hiding resource contention, sharing resources dynamically, and optimizing for greatest utilization.

The good news is that best practice performance management can help IT virtualize applications that require guaranteed performance, identify when performance gets out of whack, and rapidly uncover where contention and bottlenecks might be hiding. We have identified a few solutions that help address these performance challenges head-on. But first, in order to really understand why these solutions are necessary let’s explore why performance in today’s IT environment is such a challenge.

### ***Performance Is a Result, Not a Resource***

Virtualization is primarily designed to efficiently share expensive IT resources. It strives to extract maximum utilization out of the physical infrastructure. And since virtualized resources can be managed and optimized independently of other resources and IT domains, virtualization represents a dynamically flexible way to allocate and balance resources. Meeting these kinds of efficiency goals at first seems to provide the biggest ROI in terms of the financial investment made in capital assets, but virtualization can introduce significant performance-related challenges including:

- An inability to simply dial-in performance in the same manner as one might allocate a quantity of resource. Performance isn’t a simple resource to allocate, rather it is the result of a complex function of non-linear interactions between contentious clients – and virtualization adds complications with virtual configurations, allocation policies, host scheduling and dynamic load balancing.
- Not knowing what performance actually is. Performance can’t be fully measured or managed through simple utilization metrics. Unfortunately, what many IT performance management solutions analyze are natively measured resource utilization “percentages”.
- Degraded performance from an intentionally busy system. A system’s performance naturally degrades as internal utilization increases. Significantly, performance deteriorates faster if uti-

lization is pushed past an inherent, but hard to discern, inflection point. Simply aiming to maximize utilization can destroy performance.

- Cross-domain blindness in not knowing who is really sharing or contending for which actual resources. Ultimately, ensuring good client performance requires visibility that spans the whole IT infrastructure from end to end, apps to servers to storage. Virtualization can make it hard to discern the causes and impacts of contention both within a shared pool of infrastructure and hidden beneath and between virtualized domains.

IT architectures that blindly optimize for maximum utilization, and subsequently make it hard to figure out who is really doing what, are practically unsuitable for hosting performance-sensitive applications – which is to say, most mission-critical applications. But the situation is not hopeless. There are solutions in the marketplace today that aim to help IT with performance-focused management in dynamic IT infrastructures.

### ***Managing For Desired Performance***

From the point of view of an application, the performance it gets from the infrastructure is the sum of performance from all of the components– CPU, memory, network, and storage. IT Performance management is an inherently cross-domain problem. However, many management solutions approach performance from a ‘silo’ or element perspective. While some performance-impacting issues do derive from hardware faults or operational errors, performance degradation can also stem from non-failure causes including hidden utilization bottlenecks, unintentional sharing and contention, and even dynamic “thrashing”. Many of these performance issues are only identifiable when looking across multiple domains holistically.

Cross-domain visibility and analysis solutions include a variety of APM tools including VMware’s vFabric APM and vCenter Infrastructure Navigator that map applications, transactions or even code back to virtualized resources, virtualization storage performance solutions like NetApp’s Balance that map from within guest vm’s across hosts and SANs into array architectures, deep infrastructure performance solutions like Virtual Instruments that can map vm’s down into SAN protocol level communications, and service optimization solutions like TeamQuest Surveyor that map application processes across vms and hosts to array disks. While none of these solutions serves all IT needs, they each provide a valuable perspective. It’s likely that one or two will align more or less with particular IT staff responsibilities, but any of these solutions enhance cross-domain communication and problem-solving.

Many IT organizations are stuck in a reactive break/fix mindset when it comes to performance management. A good approach to maturing performance management practices is to calculate key performance indicators (KPI) to be used in managing the environment proactively with an eye towards optimization. For example, VMware’s vCenter Operations Manager produces KPI’s for health, risk, and efficiency. By deliberating setting goals to reduce or optimize such KPIs, IT organizations can even track their own relative performance.

Virtualized organizations only really become performance-savvy when the relationship between infrastructure utilizations and delivered performance is well understood. One way to deeply understand performance is with sophisticated analytical modeling that mathematically relates resource utilizations to desired performance. The better the “model”, the better IT can actively and accurately make configuration, deployment, and operational decisions to deliver targeted performance-based service levels. Key to assuring performance is modeling that goes beyond simply determining if an application’s capacity requirements can “fit” to predicting that the resulting “response time” performance will be acceptable.

At one time BMC had the most well-known, if not difficult to master, performance-based server modeling. In the last few years, key elements have been folded into the broader capacity-focused Proactive management suite. With twenty-two years of experience, TeamQuest supports expert predictive performance modeling with their Predictor solution, while the younger NetApp Balance internally leverages cross-domain modeling to automatically analyze performance, but doesn't provide interactive "what-if?" capabilities.

Both NetApp and TeamQuest leverage their analytical modeling engines to produce uniquely valuable infrastructure performance KPIs. These KPIs are useful to IT in proving both ongoing service quality and improvements gained with infrastructure upgrades to business-side clients.

Certainly performance can be improved with a system that automatically balances its resources based on current application needs. However, just because a system does a good job of balancing the capacity it has does not mean it has enough resources or know where to assign them to meet performance-based goals. And perhaps the biggest key to guaranteeing performance is the ability to plan infrastructure in advance of application growth or change. Solutions supporting "what-if?" modeling and scenario planning for virtualized environments help ensure that the right resources are deployed at the right time and place going forward.

### ***Delivering Promised Performance***

For the many reasons described, it's hard to host performance-sensitive applications in capacity-optimizing virtualized environments. And I've heard rumblings about excessive costs from over-provisioning in the supposedly cheap and elastic cloud – perhaps reprising the poor way IT often guaranteed performance in dedicated infrastructures.

In both cases I think performance-based *capacity planning* for performance has been prematurely dropped from many IT organizations in favor of simpler solutions that focus on *capacity management* for efficiency. Not only does this hamper IT's ability to deliver on specific performance goals, but the difference is also one of attitude - managing for internal benefit versus managing to deliver high quality services for clients.

The good news is that additional performance management can help IT revitalize virtualization initiatives and host critical applications that require guaranteed performance. Good solutions pierce through virtualization layers when necessary to identify causes of contention and other performance impacting issues. They help measure and report on delivered performance across IT domains including servers and storage. And key to helping IT deliver on specific performance goals and virtualize important applications is that they can predict the future.

*(A version of this article appeared in Virtualization Review on Feb 20, 2013)*

---

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.