APM Experts

White Paper:

Infrastructure Performance Management for Virtualized Systems

Bernd Harzog

CEO - APM Experts

March 2011

Abstract

Most enterprises that have deployed it have realized substantial hard dollar savings from virtualization, driven primarily by the benefits derived from server consolidation. These enterprises have also discovered that management activities like provisioning, recovering from server failure, or providing for disaster recovery can be accomplished in a more agile and consistent manner across a variety of systems using virtualization as the underlying technology. However, most large enterprises are only 30% virtualized, with virtualization occurring only on the application systems that are under the direct control of IT Operations.

To extend virtualization to the business-critical and performance-critical applications that are under the ownership of dedicated application support teams, the IT Operations group who owns the virtual environment must provide accurate and credible performance assurances for the virtual infrastructure that will be supporting these applications. Without these assurances, the application support teams and their business constituents have the political power to prevent these systems from becoming virtualized, and will exercise that power.

The IT Operations group needs to put tools in place that can measure the performance of the virtual environment and provide verifiable service-level assurance data to the owners of these key applications. The management tools selected to support virtualization are also essential to the ability of IT to grow the virtual environment without proportionately increasing the staff required to manage all of the new physical host servers and their guest VMs.

Table of Contents

I. Introduction: Virtualization Today and Tomorrow

Most of the enterprises that have adopted it have realized substantial hard dollar savings from virtualization, driven primarily by the benefits derived from server consolidation. These enterprises have also discovered that many management activities like provisioning, recovering from server failure, providing for disaster recovery, backup, and security can be accomplished in a more agile and consistent manner across a variety of systems using virtualization as the underlying technology. This has allowed IT agility to become an enabling driver of business agility, turning virtualization from an IT cost saving exercise, into a potential source of strategic advantage for the business.

However, many barriers remain that prevent large enterprises from virtualizing key systems and applications, including the political skirmishes associated with separate applications and Infrastructure groups within the IT organization and the lack of reliable management data. As a result, most large enterprises are only 30% virtualized, with virtualization occurring only on the application systems that are under the direct control of IT Operations.

Breaking through the Barriers

In order for enterprises to go from 30% virtualized to 80% virtualized significant changes will have to be made to how virtualized systems and virtualized applications are managed. It is the case with the vast majority of business and performance critical applications that run on dedicated physical infrastructure, that no mechanism is in place to measure and assure the responsiveness of these systems to their constituents. Rather, the risk of performance issues is mitigated through over-provisioning the capacity of the underlying hardware, and by using tools that attempt to infer the performance of the systems by monitoring resource utilization metrics throughout the environment.

Given the current state of infrastructure and applications management in the physical world, if these same approaches are used in the virtualized world, virtualization provides no benefits to the applications' owners, and only increases the risk of performance issues – since virtualization interjects a new layer of software into the stack, and interjects sharing of previously dedicated hardware resources between applications.

Therefore, enterprises are not going to be able to get to 80% virtualized, unless the virtualized environment can promise and deliver a higher level of true performance assurance than what is in place in the physical environment. This will only be possible through the use of new tools and new approaches that allow both infrastructure owners, and applications owners to understand the true performance (response time) of their respective layers of the stack in a much more real time, comprehensive, deterministic, and relevant manner than what is in place in the physical world today.

IT as a Service initiatives will compound the problems that stand in the way of 80% virtualization. By introducing fully automated provisioning of customer driven workloads into the environment, ITaaS initiatives cause there to be a constantly changing and growing mix of workloads. This rate of change, combined with the rapid changes caused by dynamic operations in the environment, will again create the need for tools that can keep up with the environment and the applications in a comprehensive and real time manner.

II. Performance Management Challenges Created by Virtualization

Multiple products and solutions are available to provide basic monitoring of resource utilization in a virtual environment. While it is important to know how key resources like CPU load, memory consumption, network load, SAN load, and the load on the storage array are utilized relative to their capacity, this information does not provide an accurate picture of virtual infrastructure performance. And although it is possible to approximate the performance of a *physical* infrastructure based on how its resources are being used, this is not possible in a virtual environment. A much more comprehensive approach is needed.

Collapsed and Centralized Application Infrastructure

Prior to virtualization, most business-critical applications ran on over-provisioned, dedicated servers, with massively over-provisioned LANs handling the traffic between these servers. When these servers are virtualized, this dedicated infrastructure is collapsed into a shared pool of server and network resources. IT is responsible for the technology that enables this sharing, the virtualization platform. When physical servers are consolidated to guests on a shared host, the IT Operations group thus becomes responsible for any application performance problems that arise, issues that are perceived, rightly or wrongly, to have been caused by this higher degree of sharing and/or the new layer of software (the hypervisor) in the software stack. In the absence of accurate knowledge of the root cause of performance issues, IT Operations staff are considered guilty until proven innocent.

Time-Based Metric Measurement

The service-level agreements that bind applications teams to the larger enterprise of end-users are based on deviations from a set of “normal” performance metrics. To define and codify normal operating patterns, many management solutions rely on either a guest virtual machine (VM) or an agent built into the operating system, such as WMI, to collect resource utilization statistics. However, the accuracy of all time-based metrics collected inside of guests, including CPU utilization, network or disk I/O rates, page fault rates, and context switch rates, is susceptible to a time-keeping issue between the guests and their host. As a result of this timing issue, the data that was used to infer infrastructure performance in a physical environment can no longer be used to report on infrastructure performance in a virtual environment.

In physical environments, many performance management solutions calculate resource utilization baselines for time-of-day or workload. Deviations from these baselines were indicative of a performance problem. However, when a server is virtualized, resource baselines are not a reliable predictor because the resources allocated to a guest at any given time are variable and dynamic in nature. Therefore, baseline deviations are no longer reliable indicators of infrastructure or application performance.

Density-Based Interactions

The hard dollar ROI from consolidation comes from achieving a higher utilization rate for server and network resources. As a result, virtualization raises the prospect that isolated workload peaks can now cause resource conflicts—which can, in turn, create performance issues. The monitoring approach that generates alarms based on deviations from baseline utilization values is too limited to tackle the complex factors that contribute to workload peaks in a virtualized environment.

Dynamic Operations

The VMware vSphere platform contains several features, such as VMotion™, HA, and DRS, that enable the movement of workloads among physical hosts. In order for IT and the business to be able to benefit from the flexibility provided by these features, IT Operations must be able to demonstrate that dynamic operations are not interfering with the responsiveness of the infrastructure to applications running on the infrastructure.

IT as a Service

IT as a Service (also known as Private Cloud), adds configuration management, multi-tenancy, provisioning, orchestration, automation and a service catalog from which users and business constituents can order new compute environments and applications. This places two new requirements upon performance management tools. The first is that they be able to keep up with the rate of change in the environment caused by IT as a Service. The second is that are able to automatically instantiate monitoring of new applications as they arrive without requiring manual configuration or intervention into the monitoring system.

Business Demand for Service Level Management

Applications owners and their business constituents will not embrace virtualization of business critical and performance critical applications unless they see benefits to them and their applications that come from the virtualization process. One of those benefits needs to be that the performance of virtualized infrastructure and virtualized applications will be at least as good as, and possibly better than what was the case for their physical predecessors. In order for the performance of these environments and applications to be better, it will have to be managed by different tools, and via different approaches than the incumbent methods and tools in the physical world.

Applications owners and their business constituents are therefore going to demand visibility into the actual performance of both the infrastructure and the applications (actual performance being defined as response time not resource utilization) in order for virtualization projects to proceed. Therefore, a new approach to infrastructure performance management is needed that is optimized for virtualized infrastructures that addresses the above challenges and allows IT Operations, applications owners, and business constituents to know with certainty how their respective layers of the stack are actually performing.

III. The Current Approach to Virtualization Performance Management

The issue with how to manage performance for virtualized infrastructure and applications is not just an issue of legacy physical tools, vs. new virtualization-aware tools. It is really whether the same approach (looking at resource utilization metrics) that is used in the physical world is appropriate for virtualized environments. There are many monitoring tools available on the market today that can be used to monitor, in particular the VMware vSphere environment.

Most of the vendors of these tools leverage the vCenter Server API, which provides a robust set of resource utilization statistics about the virtual environment. Resource utilization statistics provide insight into the availability and capacity of the virtualized system, but do not lead to a granular or comprehensive ability to monitor the actual performance of the virtual infrastructure.

As shown below, many vendors provide resource and availability management for the VMware vSphere platform. Some of these vendors, like VKernel, Veeam and Quest (vFoglight) only monitor the VMware environment, while others, like eG Innovations, Systar and Zenoss, have been monitoring a broad array of physical server and network resources and have added the collection of VMware vCenter API data to their product suites.

VMware has recently announced vCenter Operations, which combines resource, capacity, and configuration data into an analytical model to automatically calculate Health, Capacity and Workload scores. But the source of the underlying data is just the vCenter API data, which focuses upon resource utilization and not performance.

Therefore it is fair to say that most vendors (including VMware) who offer performance management solutions in the virtualized environment take the same approach of focusing upon resource utilization data that has been used in the physical environment to try to infer the performance of the environment from this data. But as we demonstrate in the next section, this approach is fundamentally flawed.

Virtualization Resource and Capacity Management Solutions and their Use Cases

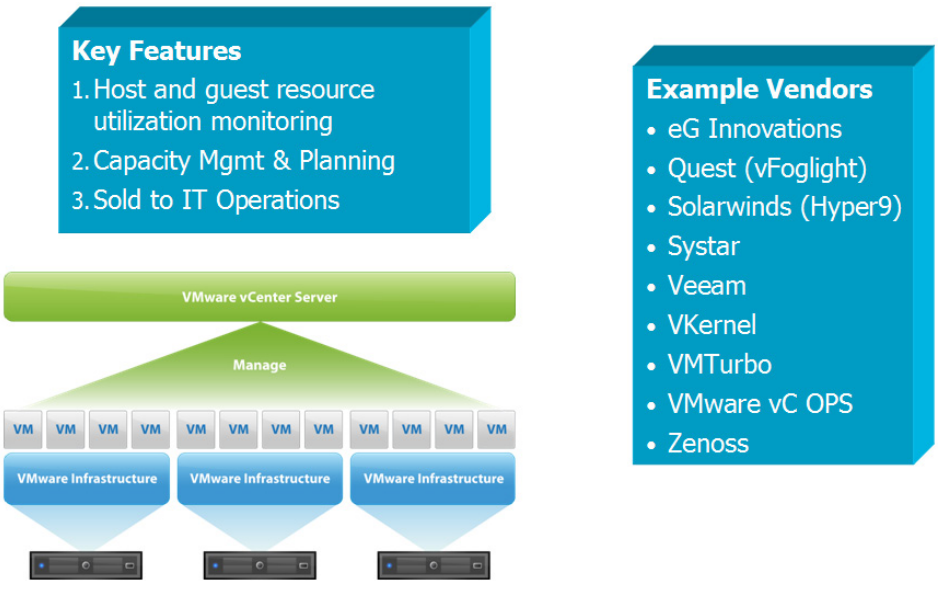

IV. Infrastructure Performance Management – The New Way

Infrastructure Performance Management (IPM) is a new approach to managing the performance of the physical and virtual hardware and software resources that comprise virtualized and cloud-based computing environments. IPM is a superset of the Resource and Availability Management category. When optimized for a virtualized environment, solutions in this category collect vCenter data, but build on this data by collecting unique data of their own, which allows them to provide a response time or latency based perspective on infrastructure performance. This new approach is based on a new performance metric for the infrastructure: Infrastructure Response Time.

Infrastructure Response Time Defined

Infrastructure Response Time (IRT) is defined as the time it takes for any workload (application) to place a request for work on the virtual environment and for the virtual environment to complete the request. Another way to think of IRT is that it is an end-to-end measurement of infrastructure latency. The request could be as simple as a bi-directional exchange of data between two guest VMs on one host over the vSwitch. Or the request could comprise multiple hops among various VMs on multiple hosts and then include a database transaction, which ultimately requires a write to a storage array and a confirmation back to the original requesting component of the application. IRT data is not available via any of the common methods of collecting performance data like the vCenter API's, SNMP, WMI, or SMIS. Therefore IPM solutions that feature IRT, have to collect this data themselves, and not rely upon simply polling for commodity monitoring data.

Infrastructure Performance Management exists as a separate category, not subject to the limitations of the resource utilization approach. The key requirements of an Infrastructure Performance Management solution are discussed below.

Real Time, Comprehensive and Deterministic

The core problem with all approaches that attempt to manage both physical and virtual infrastructures is that they rely upon infrequent (every 5 minutes, or even every 15 minute) polling of commodity data that is worthless when it comes to assessing the actual performance of a physical or virtual infrastructure in support of its applications and workloads.

At the infrastructure performance management layer, the key is to measure the latency of the response of the infrastructure to requests on the part of the workloads, and to do so comprehensively (across every single request), in real time (as the transaction occurs, not some time thereafter), and deterministically (based upon the actual latency information, not an average, nor a synthetically created approximation of the transaction).

We need to move from measurements that attempt to approximate infrastructure and applications performance, to measurements that actually measure how the infrastructure is performing on behalf of the applications, and how applications are actually performing on behalf of business constituents and end users.

In summary we need to move to a system where the actual latency for every request is measured for its responsiveness and when there is a deviation, the affected applications, and the at-fault infrastructure components are indentified at the same time that a latent infrastructure transaction is noticed.

Applications Agnostic

Infrastructure Response Time is relevant for the entire virtual environment to the team supporting that environment. To provide assurance to the application teams that the virtual infrastructure is performing well, IRT must be calculated and reported for every application running on that infrastructure. This means that it needs to work for every application in the environment, not just ones written to a specific applications run time (like Java or .NET), that use a specific applications layer protocol (like HTTP).

Infrastructure Topology Discovery

Continually discover the topology of the infrastructure supporting each application. Once you understand how applications are communicating with each other, it becomes necessary to dynamically identify the chain of virtual and physical resources that are supporting an application at a given moment in time, based on continual discovery.

Take an End-to-End Approach

Calculate IRT from the guests to the spindle and back. Infrastructure Response Time must be calculated across the breadth and depth of the virtual environment. For example, many scaled-out applications have multiple tiers that run within many different guests across a virtual infrastructure, and a significant portion of those guests may make database or other I/O calls that are serviced by a physical disk array. When calculating IRT metrics, the full scope of the virtualized environment must be considered.

Address Performance and Capacity

Provide both performance- and capacity-based analytics around Infrastructure Response Time. IRT is more than just the standard by which virtual infrastructure performance should be understood and the metric on which performance troubleshooting should be based. IRT is also the basis for a new understanding of capacity planning. When an increase in workload is contemplated, the first question that needs to be answered is whether or not this increase in workload will increase IRT beyond acceptable levels. If the metrics indicate that it will reach unacceptable levels, it must be possible to determine whether the increase signals a capacity issue, or whether it is caused by some other factor.

Provide out-of-the-box value

Virtualized environments change too rapidly for an approach that requires extensive manual configuration to provide value. New applications, and new versions of existing applications, are constantly being added to the environment. These must be automatically discovered, and the management product should start providing IRT information on these new applications immediately after discovery. Similarly the addition of new servers, new vSwitches, new VLANs, or a new storage array should not require manual reconfiguration of the IPM solution; the management system should simply discover the additions and changes and adapt accordingly.

Work across multiple Virtualized and Physical Infrastructures

Even once you get to 80% virtualization, some of what will be left will be applications systems that are 80% virtualized, but where perhaps the back end Oracle database is still running on physical hardware. And effective IPM solution must have visibility into the IRT for not just virtualized infrastructures, but also infrastructures that are partially virtualized. It is also critical that the IPM solution be able to understand resource conflicts upon resources that are shared by physical and virtual systems. For example, if a storage array is serving both physical and virtual infrastructures, and activity from the physical system is creating

contention in the array that is impacting the virtualized system, an effective IPM solution must be able to see the contention and understand both its source and its impacts.

While VMware vSphere is the clear market leader in large enterprises today, especially for business-critical workloads, other virtualization platforms are gaining sufficient technical maturity and market traction so that support for these platforms should be forthcoming in the near future within market-leading IPM solutions. It is therefore important to choose an IPM vendor whose approach is either agnostic to which hypervisor is being used, or choose one whose product roadmap clearly includes whichever hypervisors are going to gain enterprise traction in the near future.

The following diagram shows the flow of an IRT transaction from a guest to the spindle and back, through the layers of the virtual environment, and lists four vendors of IPM solutions who are profiled in the next section.

Infrastructure Performance Management Solutions and their Use Cases

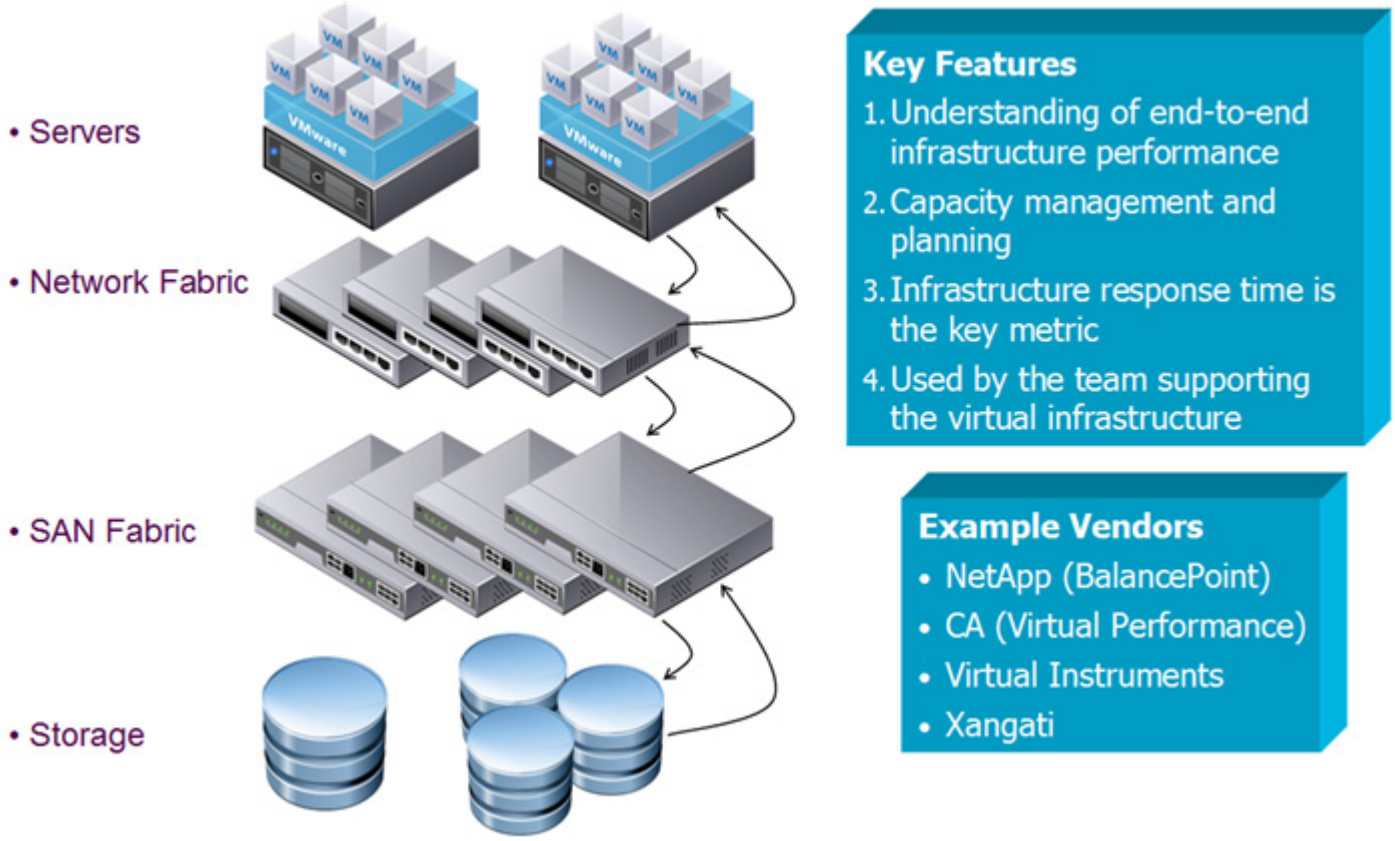

V. Approaches to Measuring Infrastructure Response Time

If IRT is the single most important metric when it comes to assessing the performance of a virtualized infrastructure in support of its hosted workloads, then what are the best ways to collect and measure IRT?

The answer, as has been written about earlier in this paper is not to try to infer IRT by looking at resource utilization statistics like CPU usage, CPU Ready, memory consumption, network I/O rates, disk I/O rates, context switches and page faults. Using these types of metrics to infer system performance worked in the physical world, but it does not work for virtualized infrastructures. In general there are three approaches to getting IRT data.

Use a Queuing Model

It is possible to monitor the CPU, network, SAN, and disk array queues in a virtualized infrastructure and to then use a queuing model to calculate IRT. This is the approach used by NetApp BalancePoint. The benefit of this approach is that it does provide a map from the guests to the storage arrays, and shows the IRT along each path in the map. The downside of this approach is that it is computationally intense, and that therefore cannot occur on a real time and continuous basis – resulting in updated data only every 15 minutes or so. Furthermore, any tool that queries the array for the information that it tracks will provide array specific information about the self-monitored response times. While these times can be useful to find configuration issues within the device such as inappropriate LUN layout, they are vendor specific and can vary in how the measurement is made from one rev of firmware to another, making it difficult to compare different levels of firmware and impossible to compare devices from different manufactures. Because the devices themselves lack key metrics such as queued or pending exchanges and command to first data times, the information lacks detail to truly understand the cause for the latency in many cases. As a result, this approach may not be suitable for larger-scale enterprises or in mission-critical environments where a real-time view of transaction data is required.

TAP the IP Network

Both physical and virtual switches can be configured to have mirror ports (also known as spanned ports). All of the data that flows through a physical and virtual switch is routed on a read only basis to the mirror port. A physical or virtual appliance can then capture this data, and by doing deep analysis of the packets, and their sources and destinations derive IRT for the IP network. This approach is used by CA Virtual Performance. The benefit of this approach is that it provides a real time view of IRT for the entire IP network. The downside is that unless the storage array is IP attached with the ISCSI storage protocol, this approach does not provide visibility into the layer of the infrastructure (the storage layer) that is the most frequent cause of performance problems for virtualized infrastructures. The other problem with mirror ports on switches is that, when the switch gets congested under heavy loads, traffic is prioritized away from the mirror port and IRT data is no longer available.

TAP the Storage Area Network

Like TCP/IP switches, some SAN switches also come with mirror ports. But the SPAN ports on SAN switches are particularly subject to losing valuable performance data at the exact time that it is needed the most (which is when the switch is busy according to this Cisco White Paper). But, it is possible (through Virtual Instruments) to add a TAP to a SAN switch. This then allows all of the Fibre Channel data within that switch to be seen and analyzed for performance. This is extremely useful not only for understanding the performance of the SAN

itself (which ports are congested, which ports are so underutilized that they should be consolidated), but it also provides the only possible real time and continuous view of IRT for the storage arrays attached to the SAN. TAPs have no performance impact on the SAN traffic and have the advantage to scaling to 100,000 ports or more.

Having an independent auditor of the environment that does its own measurement of performance of the environment also enables an apples-to-apples comparison of devices, regardless of manufacturer, firmware or configuration. It enables real-time analysis and monitoring of the environment that isn't possible by querying the devices themselves, as with the queuing model.

This real time and continuous view of storage array performance from the perspective of the SAN is a crucial piece of information for the teams in charge of the performance of the virtualized infrastructure. For the same reasons that VMware teams do not want to give vCenter access to just anyone, storage administrators need to tightly control access to the SRM tools that they use to manage the storage arrays. Therefore this “outside-in” approach to understanding storage performance is the only way for the virtualization team to get this crucial real time and continuous view of storage performance.

VI. Comparison of IPM Solutions

The table below compares the major Infrastructure Performance Management solutions available on the market today.

Feature/Capability	NetApp BalancePoint	CA Technologies Virtual Performance	Virtual Instruments VirtualWisdom	Xangati™
Data Collection Methods	vCenter APIs, direct instrumentation of SAN and storage arrays via the storage console CLI	vCenter APIs, NetFlow data from physical switches, application performance data from virtual and physical switches via mirror ports	vCenter APIs, proprietary taps into SAN fabric, SNMP to fibre channel switches/directors	vCenter APIs, Netflow data from physical and virtual switches
Breadth and Depth of Infrastructure Response Time Data Collected	Infrastructure Response Time is collected end-to-end (from guest to spindle on storage array)	Infrastructure Response Time is collected for each application identified via port and protocol from the guest through the entire IP network (LAN, WAN, and IP storage)	Measures the response time of individual Fibre Channel frames and maps them from VMs to LUNs with unique ability to identify where latency issues exist.	Infrastructure Response Time is collected for each application identified via port and protocol for the IP network (LAN, WAN, and IP storage)
Storage Performance Visibility	Has specific instrumentation for storage arrays. Captures IOPS and storage latency to physical spindles. Maps guests and workloads to spindles	Only for IP-attached storage devices using ISCSI	Taps the SAN data directly for latency and load information for all Fibre Channel traffic	Only for IP-attached storage devices using ISCSI
LAN and WAN Performance Visibility	No visibility to the LAN and the WAN	Deep visibility into all IP traffic (LAN and WAN)	No visibility into the LAN or WAN	Deep visibility into all IP traffic (LAN and WAN)
Server Performance Visibility	Direct calculation of IRT impacts on a per-guest and per-host basis	Sees server impacts from the perspective of end-to-end application response time	Relies on vCenter data to infer server-level performance issues from resource utilization data	See server impacts from the perspective of the network
Visibility into Performance Issues between Guests on one Host	No	Virtual appliance on the mirror port of the vSwitch sees interactions between guests on one host	No	Derived from Netflow data collected from the vSwitch in the host
Level of Application Identification	Pulls process list from guests via WMI. Able to identify certain key applications and workloads	Identifies applications based upon ports and protocols	No ability to tie applications to infrastructure slowdowns	Identifies applications via Cisco IP SLA protocols
Data Collection Interval	Polls the entire virtual infrastructure every 15 minutes	Real-time	Real-time	Real-time
Built-in Analytics	Automatically calculates a performance index which compares IRT against capacity utilization	Automatic baselines and thresholds, Top-N reporting. Optional investigations and notifications when performance degrades.	Basic analytics supplied with reporting function (Views).	Automatic baselining of all traffic volume and response time metrics
Deployment Model	Deployed as one subnet-attached virtual appliance in the VMware resource pool	Deployed as a virtual appliance on the vSwitch in each VMware host, physical appliances on the mirror ports on the LAN switches, and on one management appliance	Deployed as a physical tap on the Fibre Channel SAN	Deployed as a virtual appliance in each VMware host, and as a separate virtual appliance for management

VII. Virtual Instruments VirtualWisdom

Virtual Instruments VirtualWisdom is the only IPM solution on the market today that is able to provide real time, comprehensive, and deterministic Infrastructure Response Time information for Fibre Channel attached storage arrays to the team supporting a virtual (or for that matter a physical) infrastructure.

By tapping into the SAN switch, and calculating the Exchange Completion Time for each storage transaction running over the SAN, VirtualWisdom provides a real time, comprehensive and array independent view of storage performance. This information is not available to the team supporting a virtual infrastructure via any other mechanism or solution.

The table to the right below demonstrates one small example of the benefits of using real time deterministic data to assess the performance of the storage arrays. In the left column the 5 minute average latency for servers running VMware that is collected by vCenter is shown. In the right column the real time and completely accurate data collected by VirtualWisdom is shown. Notice that the averaging process is obscuring the peaks in latency (vCenter is reporting a peak latency of 22ms, and VirtualWisdom is showing the true value of 42ms), and notice that the vCenter data has the server with the worst latency (6020) ranked #3 in its list with a status of green.

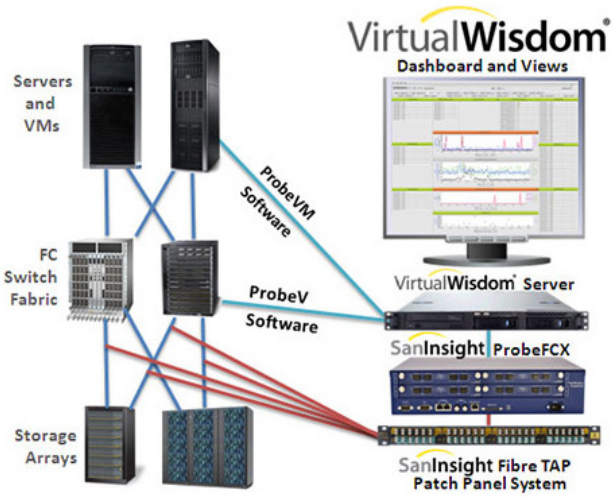

VMware vCenter 5 Minute Average Data		Virtual Instruments VirtualWisdom Real Time Data	
Highest Write Latency	+ !	Top ITL Write ECT	+ !!!
Server_6017	22	Server_6020:Array1_Port1:16	41.483
Server_6021	21	Server_6020:Array1_Port1:12	37.369
Server_6020	19	Server_6024:Array1_Port1:8	33.707
Server_6019	18	Server_6017:Array1_Port1:21	32.335
Server_6022	17	Server_6022:Array1_Port1:20	30.895
Server_6023	13	Server_6024:Array1_Port1:14	29.733
Server_6024	13	Server_6018:Array1_Port1:1	29.666
Server_6018	10	Server_6024:Array1_Port1:6	28.897
Server_6006	4	Server_6017:Array1_Port1:3	28.243
Server_6010	2	Server_6024:Array1_Port1:12	28.221
Server_6011	2	Server_6023:Array1_Port1:12	27.963
Server_7928	2	Server_6024:Array1_Port1:17	27.846
Server_6013	2	Server_6021:Array1_Port1:18	27.445
Server_6014	2	Server_6021:Array1_Port1:7	27.241
Server_6015	2	Server_6017:Array1_Port1:1	26.968
Server_6007	2	Server_6019:Array1_Port1:15	26.874
Server_6038	2		

VIII. VirtualWisdom Use Case

Problem:

The large financial services company is a major international provider of financial services and investment resources that help individuals and institutions meet their financial objectives.

The company had built out a large set of servers, Fibre Channel SANs, and Tier 1 storage in order to set up a flexible environment for implementing virtualized physical servers. By using VMware ESX, time to market for expanding business applications would be faster and lower-cost than the traditional model of one application per physical server.

As the number of virtualized servers increased, the SAN infrastructure and connected Tier 1 storage exhibited serious performance and availability problems to the point of full production outages requiring reboots of storage and/or servers. So much so that non-production test and development servers had to be shut down to prevent congestion and retries every Sunday evening or periodically when problems appeared. Eventually, all production applications were moved to other environments to mitigate the business impact of the SAN problems. For a period of eight months, the application, VMware server and storage teams, along with personnel from their storage and server vendors worked through a litany of problems all impacting the SAN.

Approach:

Virtual Instruments' professional services were engaged to solve these challenges. The team installed the TAPs, the Portable Assessment Kit, and monitored the environment for a week. It then took 3 days to analyze the findings and present the recommendations to the IT team.

Virtual Instruments showed that part of the latency issue was caused by incorrect Queue Depth settings. During configuration, the queue depth settings were set too high on the storage array Fibre Channel ports and this resulted in increasing latencies for various applications. Virtual Instruments demonstrated that the server demand was not evenly distributed across the storage array controller ports. On an array with 32 ports there were two ports that were at or near the upper limit of utilization. A handful of ports had moderate traffic while the remainder had little or no traffic load to speak of (less than 3% on average). With proper layout it was assessed that the same performance could easily be achieved by having only half as many storage ports.

One of the working theories by one of the storage vendors was that incompatibility between the switches was the causing traffic to be dropped and was the source of the poor performance. Virtual Instruments was able to prove conclusively that no frames were being dropped. The physical layer was healthy and error free and all communication was completing successfully even though it was with significant performance issues. This enabled the customer to focus on the real root cause of the performance issues and not rely on the theories put forth by competing storage vendors. Virtual Instruments demonstrated that during higher load times such as virus scans and backups, VMotion and DRS can cause SCSI reservation storms that lead to unacceptable levels of latency and VMotion failures.

As a result of the recommendation, the customer restored use of expensive storage and server components to production use, using best practices to optimize VMware performance and avoid performance issues and outages. They gained a full understanding of the I/O performance impact of VMware VMotion, VMware Distributed Resource Scheduler (DRS) and VMware High Availability (HA), and developed best practices for

configuration and deployment of VMware ESX virtual machines and servers with their SAN. Finally they dramatically reduced the admin resources and effort needed to monitor and troubleshoot the environment from ESX server to storage array.

Moving forward the company's application, storage, and VMware groups now have the confidence that they can meet new business requirements in a timely and cost-effective fashion.

IX. Importance of Virtualization Infrastructure Performance Management

Deploying a Virtualization Performance Management solution built around a broad and deep understanding of individual application Infrastructure Response Time represents the only credible approach to virtualizing the production environment. This approach allows the teams that own and support the virtual environment to virtualize business-critical and performance-critical applications. Without adequate performance monitoring, application owners will effectively resist virtualization because it increases their perceived risk of performance issues and reduces their ability to over-provision the resources that are assigned to these applications. The conflict between the IT group, who wants to deliver more cost savings and technical or business agility to the enterprise, and the application teams, who want to fulfill service level agreements, is natural from the point of view of each constituent.

To resolve this conflict, IT must be able to prove that the portion of the infrastructure that is serving each application at each moment in time is actually performing at a level that does not impede application performance. Offering resource utilization metrics as proof will not be accepted by the application owners, who cannot be expected to mentally translate a variety of component infrastructure metrics into a number that is meaningful for their application. Only a response time-based metric like Infrastructure Response Time can serve as the key point of agreement between these two constituencies. Once an agreement is established, the benefits of virtualization can be extended to additional applications in the organization.

In summary, only through the use of an IPM solution that provides an accurate, real time, comprehensive and deterministic view of Infrastructure Response Time will IT succeed in virtualizing business critical and performance critical applications, and going from 30% virtualized to 80% virtualized.

X. About APM Experts

Bernd Harzog is the CEO of APM Experts™, a consulting and analysis firm focusing on the virtualization infrastructure and application performance management markets, vendor strategies in these markets, and customer use cases in these markets. Bernd was formerly a Gartner Group® Research Director focusing on the Windows Server® operating system, CEO of RTO Software, and VP of Products at Netuitive®, and has been involved in vendor and IT strategy since 1980.