WHITE PAPER

Answering Ten Tough Questions

On Managing Virtualization

by David Messina, VP Product Management, Xangati

www.xangati.com

Answering Ten Tough Questions on Managing Virtualization

In June 2009, Xangati hosted a webinar with guest speaker, Zeus Kerravala, SVP of The Yankee Group, titled "Managing Virtualization in 2009—A Double Edge Sword." When doing initial research for the webinar, Zeus quickly observed the topic of virtualization was entirely rosy—there were no discussions of roadblocks, pitfalls, miscues—and vendors were not answering the tough questions. Not only did Zeus expose and answer those tough questions about virtualization during the webinar, but additional tough questions were asked during the webinar which led to the writing of this paper, "Answering Ten Tough Questions on Managing Virtualization."

Before getting to the Q&A portion, let's look at the popularity of the individual questions to see what your peers found to be the most interesting or "toughest" question. In one promotional tactic, Xangati posted the individual questions on LinkedIn to see which were the most popular among your virtualization peers.
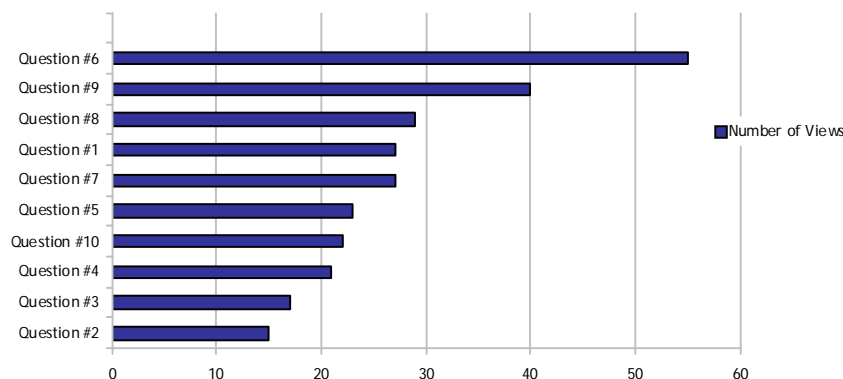
Exhibit 1: Number of page views for each tough question on LinkedIn.

The three LinkedIn groups where this Q&A series was posted were VMWare Virtualization and Cloud Computing Group (13,779 members), Virtualized Solutions (2,361 members) and Network and Application Performance Management (510 members). A quick glance shows that Question #6 was the most viewed Q&A from these groups.

Now, onto the tough questions and their answers…

QUESTION ONE: In a virtualized environment, applications can move anywhere at any time. How do we continue to keep track of that application's activity even with these sudden shifts? How do we know who or what that app's instance is dealing with at any moment?

A: Not only must we track applications and relationships in real-time, we must do it in an extensible, cost effective fashion. Existing solutions rely on software agents or hardware probes to create static, non-real-time, interdependency maps and too-little-too-late reports. Up to the minute tracking of applications and their inter-relationships requires both (1) a new data object, and (2) a new application management component within your existing IT management framework:

(1) Cisco NetFlow is *the best and only solution* for up to the minute, extensible application tracking in a virtualized environment because it is:

- Real time with <u>to-the-second</u> granularity
- Interaction-aware
- Supports all virtualized applications
- A push solution instead of a polling model
- Zero-footprint (No software agents, No hardware probes)

And just to underscore the flexibility of where NetFlow can be used I suggest you take a look at the Cisco Nexus 1000V which is a software virtual switch for a VMware world. Would you really want to have to invest in probes to tap into each one of these?

(2) An Application Management 2.0 solution is required to align with NetFlow and produce real-time intelligence for virtualized applications and their interactions. Other NetFlow based management solutions generate reports to manage network interface utilization. These 1.0 solutions do not provide a specific, granular and extensible view into applications and all of their interactions. The Application Management 2.0 solution must:

- Process and analyze NetFlow records in real-time instead of just collecting them and reporting on things later.
- Provide a streaming (to the second) view for each of your applications and all of their users and interactions.
- Provide an identity for every user, every server (virtual or not) and every application (including cloud applications).

QUESTION TWO: I know that the whole idea of virtualization is taking many and consolidating them into one, but what is a methodology that will help us programmatically plan out what virtual applications/servers can safely share the same physical server?

A: The question is correct in implying old capacity planning models won't work well in a virtual world. This is a fact substantiated by J.P. Garbani, VP at Forrester Research, in a research report entitled Capacity Management in a Virtual World. In this report, Garbani points out that "many of the proposed [management] dashboards are using traditional parameters such as CPU load and memory use, which are relatively meaningless in a virtual environment." Instead, he notes that the application workload is of paramount importance in understanding how to make the proper assessments for capacity planning. And, having visibility into the application workload from the network perspective down to every physical server will help provide you with the foundation for the methodology you seek.

An Application Management 2.0 solution answers these four critical questions and with these answers you can create a prioritized worksheet of your targets for physical to virtual migration.

1. What is the complete active inventory of my servers?
2. What is the active inventory of applications on those servers?
3. What is the complete application workload down to each one of these servers?
4. What is the user community tied to these servers and when are they active?

Knowing the above will give you your first targets, which are servers with minimal application workload and a small user community. Anchoring on these "low-hanging fruit" will lead to a fast and pitfall-free phase of implementation, which will allow you to show a quick project win. Thereafter, the same information will help you work through the rest of the list.

Reminder: all this can be done *without* an investment in probes, agents or a CMDB.

QUESTION THREE: How do we task our service desk with supporting mission critical cloud applications that are not actually on our network and not owned by us?

A: Adding cloud applications to the long list of what your service desk already has to support is compounding an existing problem—the front-line currently "flies blind" when an end-user has a networked application performance problem. Delivering a

solution means providing the service desk with real-time visibility into the end-user's application experience along with utilization information on the network transport between your enterprise and the cloud application provider.

Given the dynamic fluctuation in performance levels, the service desk needs a streaming visual of both the user's application activity on the network and the utilization of the transport connection to the application cloud. And, they should also be able to see what server(s) the end-user is communicating with (or at least trying to.) Not only should this information be available live, but the service desk should be able to DVR record this activity as well. If the communication is successfully made, but the process is still slow, that would likely give your service desk (or an escalation team) the necessary information to either route the ticket to the internal team that owns the transport or communicate directly with your cloud computing provider.

Now, instead of just randomly communicating a performance issue, it can be presented as a DVR recording forwarded onto the support group of the cloud application provider along with supporting details that the transport link has already been ruled out as problem. This streaming visual and DVR recording information will be an invaluable addition to the communication flow between your IT team and your cloud computing vendor—reducing the finger pointing and speeding the time to resolution for impacted users.

QUESTION FOUR: What is my upgrade path for all the software agents that I have previously purchased for managing physical servers and physical applications?

A: If there is an upgrade path, it is likely to be a long, winding and expensive one— since neither our customers nor key analysts can relay clear-cut answers from the vendor community on this subject. It is not outside the realm of possibility that a "forklift" upgrade will be necessary for some of the Application Management 1.0 vendors—after all, the agent world is heavily tied to discrete physical elements like CPU and memory, which begin to be obfuscated when those elements migrate to shared virtual resources.

Given the uncertainty noted above, it is important to find a way to get quick and broad coverage of the virtual elements that are on your infrastructure—that is where an Application Management 2.0 solution fits in. Xangati's consumption of NetFlow allows an endpoint like a server to be unfettered from a physical definition and allows for near instant visibility post a quick one hour install. Any

virtual IP endpoint (server, desktop, storage, etc.) on your infrastructure will be tracked with a fixed identity—identity mapping comes through a linkage to one of your directories e.g., DNS or Active Directory. By helping to close your emerging visibility/information gaps, an Application Management 2.0 solution like Xangati serves as a complementary addition to your management portfolio.

QUESTION FIVE: Although cloud computing is theoretically a centralization effort, doesn't it have the effect of decentralizing management—when all my end-users can access a cloud application from their local Internet connection (either in their regional office or home office)?

A: In short, yes cloud computing decentralizes manageability if you let it. That is why before you begin the broad adoption of cloud services, it is imperative you make sure both your networking team and your IT operations team are involved in the discussion. If they are not, you are going to find yourself blind-sided with a tremendous visibility gap. In your discussion with these other teams, the first area of focus should be to make sure that you have a way to fully account for your mobile/remote access employees. This means a potential shift in your VPN model if your VPN clients are configured to support a "split tunnel." A split tunnel allows Internet traffic to be routed out locally, while corporate traffic is shuttled back through the corporate network and out its big Internet pipe. Given cloud applications are largely Internet-based, removing split tunneling will go a long way in helping eliminate support guess work.

The other visibility gap to address is somewhat ironic given enterprises just went through major data center consolidation efforts. The migration to cloud computing means that key applications will no longer be accessed centrally across the WAN, but rather directly through the Internet connection of each remote office. This communication shift renders your existing management instrumentation insufficient to give you the full picture and creates a need for a cost effective way to get visibility out to those remote offices. This is where an Application Management 2.0 solution consuming NetFlow can be of great value. NetFlow can be pushed from each one of the remote office routers to a central management solution back at headquarters, which allows you to fill in any visibility gaps. In the end, manageability can be maintained even for cloud computing if the right steps are taken.

QUESTION SIX: If configuration issues are generally the biggest cause of complex problems, isn't configuration management going to be a bigger/broader problem in an environment where applications and servers can be moved dynamically?

A: Yes, it is certainly the case that configuration management has the potential to become an even thornier challenge. There appears to be a lot of promise in emerging technologies like Cisco's Unified Computing System (UCS), which is supposed to tackle this challenge by integrating everything in a single solution. However it is too early in the lifecycle of this technology to assume that all is well in terms of configuration management. In the meantime, the most effective thing that can be done in the transition to virtualization is to test and then retest. As Zeus Kerravala pointed out in the webinar, "for every dollar spent on testing now, it will save five dollars in support later."

This is where the DVR capability of an Application Management 2.0 system will come in handy and can serve as a compliment to existing load-testing solutions. The DVR recording will allow you to have a 360° vantage point of a virtual server's standard interactions when it is one location, as well as provide the same information after the virtual server has moved elsewhere. This way there can be a visual comparison to see if anything has changed within the communication ecosystem of that server: Are the same users able to access it? At the same speed? With all the same applications? And, if things are not the same, you can visually see where the differences are.

Exhibit 2: The DVR-recording of an Application Management 2.0 solution improves change management through recordings of a before/after analysis.

QUESTION SEVEN: **In terms of cloud computing, haven't we increased the problem area considerably for troubleshooting when my users are experiencing an application performance problem?**

A: Yes, the problem area has increased considerably much to the dismay of even your most seasoned troubleshooters. There are now many more touch points involved that could adversely affect application performance—a problem can be with the cloud computing provider and their applications and servers; it can be somewhere in the Internet between your company and the cloud data center; it can be a matter of the application contending for constrained resources across an Internet connection; it can be because of a contention for resources across the WLAN—the list goes on.

Part of shrinking the problem area is through recommendations in the answer to Question #5, but that will only get you so far. What will be absolutely essential for helping to automate your cloud computing support process is the direct collaborative involvement of your end-users. This can be achieved by enabling them to initiate DVR recordings of their experience as they are trying to communicate with their cloud application. The Xangati Visual Trouble Ticket™ portal captures the full breadth of a given end-user's experience including their to-the-second communication workload to the specific cloud server, as well as the response time to this device and where the delay in communication is occurring. This kind of visual information will be invaluable in helping to focus and shrink the troubleshooting efforts.

Exhibit 3: The Xangati Visual Trouble Ticket™ captures a full breadth of a given end-user's experience including their to-the-second communication workload to a specific cloud server.

QUESTION EIGHT: If I don't want to go down the path of working with just one virtualization vendor, what kind of solutions am I going to need for heterogeneous environments?

A: It is very important to keep in mind that your virtualization vendor is going to be highly interested in having you select their management solution, or one from a trusted partner. While that solution ultimately may be necessary and right for your organization, it is important that you don't find yourself locked into a one vendor solution. At a management construct-level, that means at the very least ensuring that you have a solution that treats VMware, Citrix and Microsoft virtual servers in the same fashion. Moreover, virtual servers and traditional servers should be able to be managed seamlessly under the same framework.

QUESTION NINE: We have just made a major migration to virtualization and there is a big performance hit to a key application, how do I keep the blame from sitting with the new architecture?

A: You better make sure that you have an Application Management 2.0 framework in place that is focused on enabling collaboration between multiple departments and can represent and link together each of their different vantage points. If a collaborative framework is not deployed, you are going to find yourself in more inter-departmental finger-pointing sessions than you were previously. And you know how the game works, if virtualization is the new addition, then it is surely to be the first blamed.

Therefore, before you complete your implementation, you want to make sure that you've championed a management solution that will be leveraged across all key IT disciplines. By having a solution that allows each team to start from their area of interest in (and can also tie things back to the end-user experience), there will be a much more fluid communication process. The system itself becomes the fabric for collaboration and things like the DVR recordings are the visual notes that can be shared across teams.

QUESTION TEN: What kind of visibility is my cloud application provider going to give me to help separate out performance issues that are related to them from ones related to my network?

A: The reality is the vast majority of cloud application providers are not going to provide you with any visibility into what's going on within their cloud and with their specific servers. This is why you will want to have solutions that will give you as much visibility possible into what each of your end-users is doing with the cloud

computing application. Once again, this is where Xangati's concept of the Visual Trouble Ticket portal comes into play. The DVR recording triggered through the portal will allow IT operations to see everything that an end-user is doing with a cloud computing application at the time of a performance problem. To provide clarity on the end-user's issue, most of our customers instruct that user (through instructions on the portal) to re-try their transaction two or three times to ensure that there is plenty of detail in the recording. In addition, there is an embedded applet within the portal that will help track response times and where the likely delays are located during the 15-minute recording cycle. If the information points the issue in the direction of the cloud provider, then that recording can be played back for review with that provider's support organization.

About Xangati

Xangati is the provider of the industry's first Application Management 2.0 solution for both service providers and IT organizations that want a highly-developed solution for managing the dynamic and unpredictable nature of applications. Xangati has integrated prominent Web 2.0 concepts like streaming, collaboration and user-generated content into a comprehensive system for managing networked application environments.

APPLICATION MANAGEMENT		
2.0 Solution	vs	**1.0 Solution**
Streaming Visuals—To-the-second views of any app and any end user on the network with the ability to "DVR record" their network activity. This results in dramatically faster problem identification and improved change management.	vs	**Report-Centric—**Which means too little information, too late and is why operations departments often find themselves chasing complex performance problems for days, weeks or months.
Proactive Alerting for "Shades of Gray"—an Application Management 2.0 solution profiles every application and understands its workload and user community for every hour of the day. The operations team will be alerted when significant changes in activity (above and below normal) occur not just complete outages.	vs	**Red Lights/Green Lights—**Most products only show when an application or node is completely broken, but that doesn't stop the phone calls coming into the front-line support desk.
Interactions 360°—Real-time view of an app or end user and a 360° pan and zoom of the elements it is interacting with. This capability provides the only way to have complete situational awareness of an issue when the state of change is perpetual.	vs	**Static Maps—**Topology and dependency maps were adequate for traditional enterprise and service provider network architectures where not much can and did change.
Zero Footprint—One central appliance gets "flow data" pushed to it from routers/switches across a network. The result is low TCO, rapid time to value and extensibility across the network.	vs	**Agent/Probe Heavy—**a 1.0 solution is best for a high degree of instrumentation for a few centralized apps— beyond that the investment is too high. Moreover, how do you instrument for apps that are in the cloud and can be accessed from any Internet link?
End-User In Approach—an Application Management 2.0 solution tracks the application experience of every end user on the network. The result is an ability to greatly improve end-user satisfaction with their app performance.	vs	**Bottom Up Approach**—This approach is focused on the health of an element or a node and that information is passed from the bottom up to a manager of managers. However, what is not part of this model is a tracking of the interactivity of the elements/nodes nor the end-user experience.
Collaborative and Multi-Role—Application Management 2.0 provides a framework for various departments to see issues from their own vantage point and work together to analyze and solve problems more quickly.	vs	**Role Specific**—1.0 management is about departmental specific tools that can only be used in the hands of a few. This often leads to problem resolution ping pong and inter-departmental finger pointing.
Installs in Less Than an Hour—The cost of the system is the cost including integration.	vs	**Hidden Systems Integration Costs**—To get a system deployed and integrated can take months and integration costs can be more than the actual software purchased.

For more information on Xangati and its Application Management 2.0 solution, visit the Xangati [website](#).