

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. To find out, 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...

Contents

QUESTION ONE: (tied for 3rd most popular)

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? 4

QUESTION TWO: (10th most popular)

I know 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 which virtual applications/servers can safely share the same physical server? 5

QUESTION THREE: (9th most popular)

How do we task our service desk with supporting mission critical cloud applications that are not on our network and not owned by us? 5

QUESTION FOUR: (tied for 5th most popular)

What is my upgrade path for all the software agents I previously purchased for managing physical servers and physical applications? 6

QUESTION FIVE: (tied for 5th most popular)

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)? 6

QUESTION SIX: (most popular)

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? 7

QUESTION SEVEN: (tied for 3rd most popular)

Doesn't cloud computing increase the problem area considerably for troubleshooting when users experience an application performance problem? 8

QUESTION EIGHT: (tied for 3rd most popular)

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? 8

QUESTION NINE: (2nd most popular)

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? 9

QUESTION TEN: (tied for 5th most popular)

What kind of visibility is my cloud application provider going to give me to help separate out performance issues related to them from ones related to me? 9

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 to-the-second 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 be forced to invest in additional probes to take advantage of it?

(2) Even with NetFlow, a new approach is required to produce real-time intelligence for virtualized applications and their interactions. Other NetFlow based management solutions generate reports to manage network interface utilization. These legacy solutions do not provide a specific, granular and extensible view into applications and all of their interactions. Xangati takes a new approach with NetFlow which:

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

QUESTION TWO: I know 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 which virtual applications/servers can safely share the same physical server?

A: The question is correct in implying old capacity planning models don'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.

One approach is to answer these four critical questions and with the answers 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 faster project win. Best of all, Xangati can help you answer these questions *without* investments in probes, agents or a CMDB.

QUESTION THREE: How do we task our service desk with supporting mission critical cloud applications that are not 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.

Xangati believes that 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 one should be able to DVR record this activity as well for later playback. Then whichever team winds up owning problem resolution has access to the DVR recording and exactly what was happening – drastically reducing mean time to repair (especially if it's your cloud computing provider). This DVR recording of streaming visuals 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 – which is why DVR recordings are standard Xangati functionality.

QUESTION FOUR: What is my upgrade path for all the software agents I 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 vendors—after all, the agent world is heavily tied to discrete physical elements like CPU and memory, which lose relevance 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—which is where Xangati comes in. The way Xangati consumes NetFlow allows an endpoint like a server to be unfettered from a physical definition and allows for near instant visibility after a quick implementation (under one hour). Any and every virtual IP endpoint (server, desktop, storage, etc.) on your infrastructure will be tracked with a fixed identity through identity mapping to one of your directories e.g., DNS or Active Directory. By helping to close your emerging visibility/information gaps, a solution like Xangati serves as a complementary addition to your existing management software.

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 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 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 a Xangati solution that consumes 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: By LinkedIn, this was the most popular question. And 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 capabilities of Xangati comes in handy and 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 at one location, as well as provide the same information after the virtual server has moved elsewhere. This way a visual comparison can be made 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.

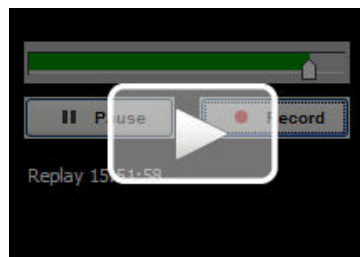


Exhibit 2: The DVR-recording of an Application Management 2.0 solution improves change management through recordings of a before/after analysis.

QUESTION SEVEN: Doesn't cloud computing increase the problem area considerably for troubleshooting when users experience 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 could be with the cloud computing provider and their applications and servers; it could be somewhere in the Internet between your company and the cloud data center; it could be a matter of the application contending for constrained resources across an Internet connection; it could be because of a contention for resources across the WLAN—the list goes on and on.

Recommendations to shrink part of the problem area were provided in the answer to Question #5, but those will only get you so far. What will be absolutely essential 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.

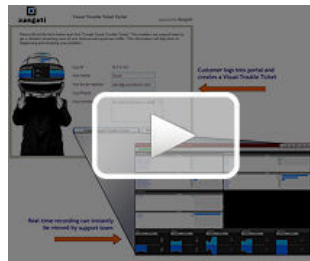


Exhibit 3: The Xangati Visual Trouble Ticket™ captures a full breadth of a given end-user's experience including their to-the-second communication workload to a specific cloud server.

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 want you to 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 you don't find yourself locked into a one vendor solution. At a management construct-level, this 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 – the approach Xangati uses.

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 need better visibility into what is happening to prove the problem isn't the new architecture. Xangati does this by enabling collaboration across multiple departments and by representing and linking together each of their different vantage points. If a collaborative framework is not deployed, you will find yourself in more inter-departmental finger-pointing sessions than you were previously. And you know how that game works.

Therefore, before you complete your implementation, you want to make sure you've championed a management solution that can be leveraged across all key IT disciplines. By having a solution that allows each team to start from their area of interest, work in and tie performance 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 related to them from ones related to me?

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 their specific servers. This is why you will want to have solutions that give you as much visibility as possible into what each of your end-users is doing with the cloud application. Once again, this is where Xangati and its Visual Trouble Ticket portal and DVR recordings of activity come into play. The DVR recordings triggered through the portal will allow IT operations to see everything an end-user is doing with a cloud 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 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 the recording can be played back for review with that provider's support organization – and getting the monkey off your back.

About Xangati

Xangati develops virtual management software for today's virtual world. Xangati provides live, video-like visibility into the performance of each and every application, server, network, and end-user device without agents or probes. This live visibility eliminates blind spots left by existing management tools and accelerates savings from virtualization initiatives. For more information, visit www.xangati.com.