3 Steps to Server Virtualization Visibility

Each enterprise has its own reasons for moving to virtual infrastructure, but the common benefit is better and more efficient server utilization. Ensure comprehensive visibility with three practical steps.

Virtualization is a money-saving technology that lets the enterprise stretch its IT budget much further by better using server assets. Consider the immediate reduction in data center footprint, maintenance, and capital and operating expense overhead. Then add the promise to dynamically adjust server workloads and service delivery to achieve optimal user experience—the vision of true orchestration. It's easy to see why server virtualization is key to many organizations' operational strategy.

However, what if something goes wrong? With network infrastructure, you can usually north/south track the root cause back to one location via careful instrumentation of the resources. Troubleshooting is then facilitated with any number of free- and commercial-ware monitoring tools. How can you get the same visibility you need to validate service health within the virtual server hypervisor and vSwitch east/west traffic?

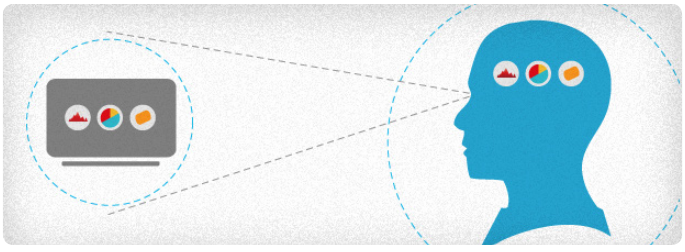

Virtual Eyes

Network teams are the de facto first responders when application performance degrades. For this reason, it's critical to maintain visibility into and around all virtual constructs for effective troubleshooting and optimal service delivery. Otherwise, much of the value of server virtualization and consolidation efforts may be offset by sub-par application performance.

Fundamentally, achieving comprehensive visibility into a virtualized server environment requires an understanding of the health of the underlying resources, including host, hypervisor, and virtual switch (vSwitch) along with perimeter client, and application traffic.

In addition, communication technologies like VXLAN and Cisco® FabricPath must also be supported for full visibility into the traffic in these environments. Without this support, network analyzers cannot get comprehensive views of virtual data center (VDC) traffic.

An effective monitoring strategy for the VDC can be summarized in three steps:

- **Step One:** Get status of host and virtualization components
- **Step Two:** Monitor vSwitch traffic
- **Step Three:** Inspect perimeter and client conversations

Step One

Get Status of Host and Virtualization Components

The host, hypervisor, and vSwitch are the foundation of the entire virtualization effort so their health is crucial. Polling technologies such as SNMP, WSD, and WMI can provide performance insight by interrogating the host and various virtualized elements. A fully-integrated performance management platform can not only provide these views, but also display relevant operating metrics in a single, user-friendly dashboard.

Metrics like CPU utilization, memory usage, and virtualized variables like individual VM instance status are examples of accessible data. Often, these parameters can point to the root cause of service issues that may otherwise manifest themselves indirectly.

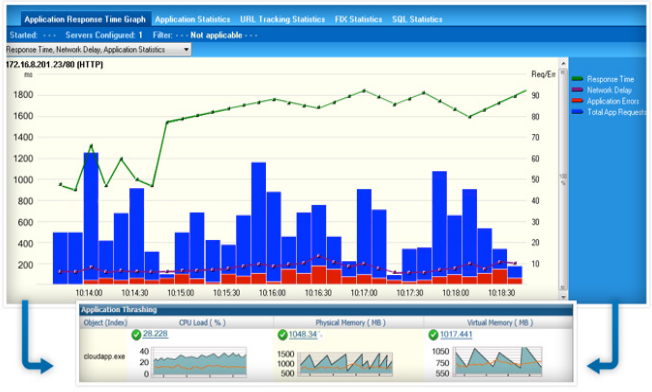

For example, poor response time of an application hosted on a virtualized server may have nothing to do with the service or the network, but may instead be tied to excessively high CPU utilization. Without this monitoring perspective, troubleshooting will be more difficult and time consuming.

Step Two

Monitor VSwitch Traffic

Monitoring how east/west traffic flows within the vSwitch is critical to understanding the overall virtualized server health. This is particularly true when multiple tiers of the application are deployed within the same host.

The implications can be significant. Without comprehensive visibility, network engineers can only guess what happens to the traffic once it passes into the virtual server. Things happen to the user request inside (good or bad), and then, assuming the situation is salvageable from a functional standpoint, out comes a service response.

Without vSwitch or hypervisor east/west visibility, the engineer on the outside can only view server application access layer conversations at the edge from this vantage point. To address this deficiency, there are two options. Depending upon the monitoring team's preferences, unique resource requirements, and virtual server solution deployed, an inside or outside model may be chosen.

In either case, to minimize the impact to user traffic, either a port SPAN session (which typically requires a dedicated physical NIC) or ERSPAN should be employed for bringing traffic out of the host that would otherwise remain inside the virtualized server. If this is not an option, be alert to the potential impact of production traffic by monitoring bandwidth consumption exiting the server. Alternatively, some network monitoring switch vendors offer hypervisor-based visibility into east/west which allows copies of traffic to be transmitted out-of-band outside the server or within the existing vSwitch. Finally, be aware that some virtual server vendors now support NetFlow, another rich source of vSwitch activity and health which can be transmitted to and analyzed by many monitoring tools.

The Inside Virtual Monitoring Model

The process begins by creating a dedicated VM "monitoring" instance and installing a monitoring tool on it. Next, you must deliver the packets to this VM. There are two methods to achieve this; the first is to exploit port-spanning capabilities available from most server virtualization vendors within their vSwitch. The second, as mentioned above, is to install hypervisor-based visibility from third-party vendors within the hypervisor, capturing and transmitting relevant traffic to the monitoring VM. Summary or select packet data can then be transmitted to an external central analysis solution for rollup with other network infrastructure data.

Since the monitoring processing is done within the server, there can be concerns as to the amount of compute resource overhead—this is something that will need to be assessed for each unique environment. Regardless, there are several distinct advantages to this method. First, it is simpler conceptually to keep the monitoring of a given virtualized server self-contained. Second, depending on the monitoring strategy, it also has the potential to reduce the amount of outbound monitoring traffic exiting the server (since processing is performed inside and hence only summary data is often transmitted). The following diagram illustrates the inside monitoring process:

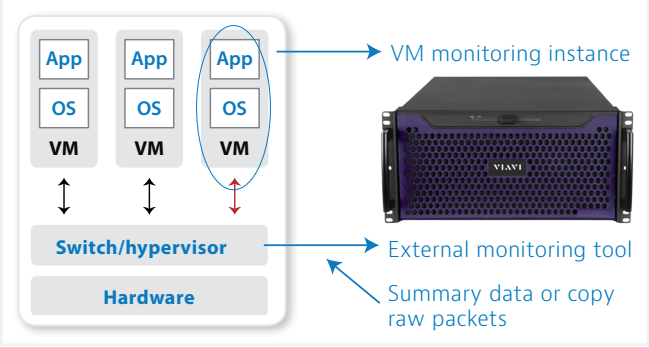

Figure 1. The inside virtual monitoring model: virtualized server monitoring using a VM monitor instance

A significant disadvantage is the potential for considerable resource overhead (such as processing and storage) by the monitoring instance on its VM and the entire virtual server, and thus reducing the amount available for the actual services running on the host. Depending on load-levels and assuming VMware or Hyper-V is deployed, a more compelling instrumentation method may be the outside virtual monitoring model described below.

Did You Know?

Virtual servers are often highly provisioned and operating at elevated utilization levels. Assessing their underlying health and adding additional resources, when necessary, is essential for peak performance.

The Outside Virtual Monitoring Model

In this method, performance monitoring is achieved by pushing a copy of "raw" (unprocessed) vSwitch east/west traffic out of the virtualized server via the built-in port mirror functionality or as described below by using hypervisor-based visibility available from some network monitoring vendors.

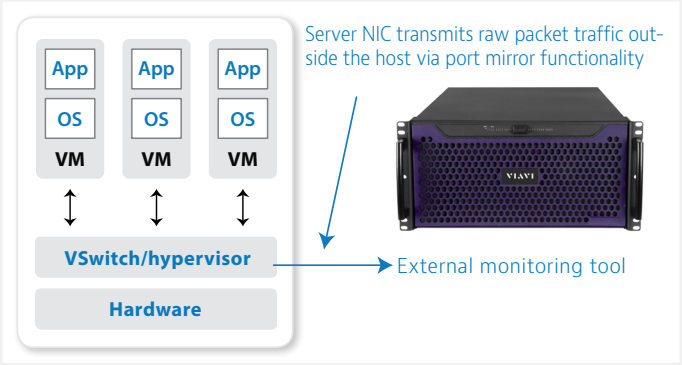

Figure 2. The outside virtual monitoring model: all analysis is performed externally

The method of vSwitch port mirroring offers an advantage in that it requires less processing, or resource overhead, since data analysis is not done within the server. This method also offers some flexibility as to the placement of analysis tools.

To transmit data outside the virtual server to an external device, there are two options. The first is to send copies of the raw packets to an external long-term capture device via a dedicated network card. Given the limited availability of dedicated cards, another option is to utilize ERSPAN. ERSPAN takes the traffic between the two VMs and wraps it in a tunnel which is sent to an IP address that can be either virtual or physical. The tunnel is then pushed onto the physical network to the external monitoring device.

Hybrid Strategies

To scrutinize east/west traffic, some organizations employ a hybrid model using monitoring strategies from both of these methods based on their specific needs. For example, it may be advantageous in certain circumstances where there are I/O limitations into and out of a specific virtualized server to perform some of the preprocessing within a VM in order to reduce outbound monitoring traffic from the server. You would still need to offload the remaining analysis externally to keep host overhead at acceptable levels.

A Final Consideration

One way to simplify the monitoring process and effectively eliminate the need for monitoring the vSwitch is to restrict processing of multi-tier applications to an individual tier per host. Doing so will ensure visibility for all points in an application as traffic between two particular tiers will be accessible at the perimeter. Note however, that this will demand planning and agreement with at least the server and applications teams to guarantee operational compliance. The downside is a loss of some hosting flexibility and a potential for possible reduced application performance as inter-tier communications are more distributed.

Multiple Hosts and Encapsulated Traffic

To effectively manage compute resources, companies often spread the virtual machines which house application components across multiple physical hosts. This also improves the reliability and fault tolerance of multi-tiered applications. As traffic travels from a VM on one host to a VM on another, it will often hit a physical link where traffic can be passively monitored.

This traffic is frequently sent using communication technologies like VXLAN and Cisco FabricPath. For full visibility, ensure that the monitoring solution you deploy is capable of providing the support and intelligence to decapsulate and analyze this traffic.

Step Three:

Inspect Perimeter and Client Conversations

Frequently referred to as the application access layer, these north/south links are often heavily utilized, running at least gigabit, frequently 10 G, and in some of the highest performance environments, even 40 G network speeds. Here, there can be considerable congestion resulting in service bottlenecks if the network or applications are not properly architected for optimal inter-tier communications or the resources are oversubscribed. This is a particularly good vantage point for understanding the behavior of client traffic as it enters the virtual server environment, as well as a source of insight into end-user experience.

This north/south traffic visibility provides a crucial view into a potential service pinch point. Therefore, beyond the core and distribution layers, it is another area of the network infrastructure that should be instrumented with back-in-time analysis tools with in-depth knowledge of the application layer.

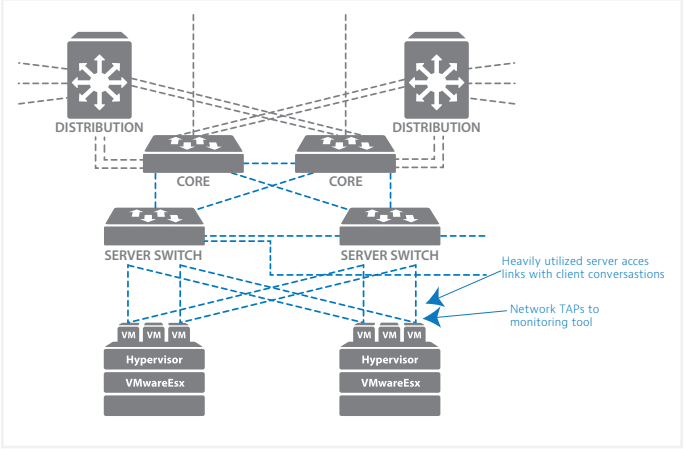

Figure 3. Instrumenting the application server access layer to avoid traffic congestion

Three Steps to Server Virtualization Visibility: A Cheat Sheet

Step One Get status of host and virtualization components	Use polling technologies such as SNMP, WSD, and WMI to provide performance metrics like CPU utilization, memory usage, and virtualized variables like individual VM instance status to find the real cause of service issues. Poor application response time and other service issues can be tied to unexpected sources.
Step Two Monitor vSwitch east/west traffic	To the network engineer, everything disappears once it hits the virtual server. To combat this "black box effect," there are two methods to maintain visibility, keeping in mind that any monitoring solution should also provide support for communication technologies like VXLAN and Cisco FabricPath: · Inside Virtual Monitoring Model 1. Create a dedicated VM monitoring instance. 2. Transmit relevant data to this instance for analysis. 3. Analyze traffic locally with a monitoring solution. 4. Transmit summary or packet data to an external central analysis solution. · Outside Virtual Monitoring Model Push copies of raw, unprocessed vSwitch east/west traffic out of the virtualized server.
Step Three Inspect perimeter and client north/south conversations	Instrument highly-saturated application access layer links with a packet capture device like Observer GigaStor™ to record conversations and rewind for back-in-time analysis.

Tech Tip

Be sure the packet-capture appliance you select is up to the task of keeping up with traffic that can regularly operate at loads up to 70–80% utilization.

Conclusion

Virtualization and consolidation offers significant upside for today's dynamic data center model and in achieving optimal IT business service delivery. However, monitoring visibility must be maintained so potential application degradation issues can be detected and resolved before impacting the end user. To do so, care must be given to properly instrument virtualized server deployments and the supporting network infrastructure.

Each data center environment is unique, with distinct challenges and demands. These are the variables that will ultimately drive the final monitoring architecture so the concepts provided in this white paper may need modification for each unique real-world deployment.

For best results, collaborate with network, server, application, operations, and potentially even storage teams to maximize the probability that the agile, service orchestration vision, of which server virtualization and consolidation plays an integral part, will translate into true business value and outstanding application performance.

For a fully-integrated performance management platform that can provide these views and display relevant operating metrics in a single, user-friendly dashboard, check out www.viavisolutions.com.

Use Three Steps to Server Virtualization Visibility: A Cheat Sheet to easily share key concepts with your team.

VIAVI

Contact Us **+1 844 GO VIAVI**
(+1 844 468 4284)

To reach the Viavi office nearest you,
visit viavisolutions.com/contacts.

threestepsservervirtualizationvisibility-wp-ec-ae
30176206 902 0915