



The Challenges of Managing Virtualized Server Environments

Jim Metzler

Introduction

Virtualization is currently one of the hottest topics in the information technology (IT) field and it is likely to remain that way for the foreseeable future. While virtualization is important to virtually all IT organizations, it certainly isn't new. IT organizations have been implementing virtualized technologies such as Virtual Local Area Networks (VLANs) and virtual private networks (VPNs) for at least twenty years. In recent years, many network hardware vendors have virtualized technologies that had previously been deployed as dedicated hardware devices and placed the virtualized instances of these technologies into routers and switches.

In the context of virtualization, almost every component of IT systems can be virtualized. This includes:

- Servers
- Desktops
- Applications
- Wide Area Networks
- Local Area Networks
- Storage
- Appliances such as WAN optimization controllers, application delivery controllers and firewalls

Virtualization enables IT organizations to get better control over IT resources, reduce the cost of network equipment and also reduce the organization's power and space requirements. In addition to the many benefits provided by virtualization itself, virtualization is often associated with cloud computing. There is no universal definition that identifies the characteristics of a cloud computing solution. Virtualization, however, is closely linked with cloud computing because most cloud computing solutions involve one or more forms of virtualization, with server virtualization being the most common form.

As will be discussed in this paper, in a virtual data center environment there is still a physical IT infrastructure that is comprised of devices such as servers, LAN switches and firewalls. As such, IT organizations still have the same traditional management challenges with these devices that they have always had; e.g., application discovery, baselining, application profiling and response time analysis, troubleshooting, etc. IT organizations also have the same traditional management challenges inside of a virtualized server, as they now have the requirement to manage virtual machines (VMs), virtual switches and virtual appliances such as firewalls. In addition, IT organizations have some new management challenges such as those that are created by the ability to dynamically move a virtual machine from one physical server to another.

This paper will explore the application and service delivery management challenges created by server virtualization. It will also describe how the IT organization must deploy management functionality inside of a virtualized server in order to regain the visibility that was lost due to virtualization.

The Definition of Server Virtualization

One of the most common ways to characterize virtualization is to think of it as involving a logical abstraction of physical systems that allows one of the following:

- Multiple physical systems to appear as a single logical system; e.g., a compute cluster with a single system image. When a host is added to a compute cluster, the host's resources become part of the cluster's resources and the cluster manages the resources of all hosts within it.
- A single physical system that is partitioned to appear as multiple independent logical systems; e.g., multiple VLANs defined on a single physical LAN or multiple VPNs on a single WAN link.

Most of the current interest in virtualization revolves around the virtualization of servers. Referring back to the preceding characterization of virtualization, server virtualization is an example of partitioning a single physical system to appear as multiple independent logical systems. In particular, as a result of server virtualization, a single physical server appears as multiple VMs.

In a recent survey conducted by Ashton, Metzler and Associates, 339 IT professionals were asked to indicate the percentage of their company's data center servers that have either already been virtualized or that they expected would be virtualized within the next year. Their responses are shown in Table 1.

Table 1: Deployment of Virtualized Servers¹

| | None | 1% to 25% | 26% to 50% | 51% to 75% | 76% to 100% |
|--|-------|-----------|------------|------------|-------------|
| Have already been virtualized | 21.6% | 33.0% | 18.9% | 15.1% | 11.3% |
| Expect to be virtualized within a year | 12.4% | 25.6% | 21.9% | 21.9% | 18.2% |

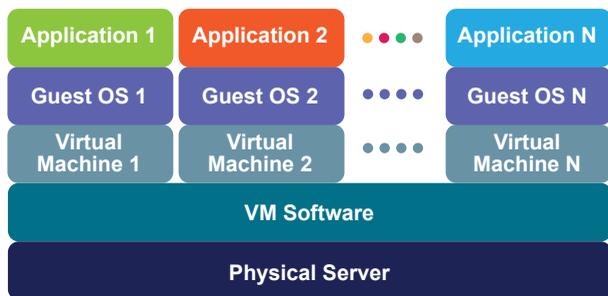
The data in Table 1 shows the ongoing interest that IT organizations have relative to deploying virtualized servers. Two observations that can be drawn from Table 1 are that within the next year:

- The number of IT organizations that have virtualized the majority of their servers will grow by sixty percent.
- The number of IT organizations that have not implemented server virtualization will be cut almost in half.

Server virtualization is not a fundamentally new concept. IBM was shipping virtualized mainframes roughly thirty-five years ago. In the current environment, VM software that is alternatively referred to as a hypervisor or as a virtual machine monitor (VMM) runs on a server. As shown in Figure 1 (page 3), with VM software, a single physical machine can support a number of guest operating systems (OSs), each of which runs on its own complete virtual instance of the underlying physical machine. The guest OSs can be instances of a single version of one OS such as Linux, different releases of the same OS, or completely different OSs. For example, an IT organization could choose to run Linux, Windows, Mac OS-X, and Solaris on separate VMs within a single physical virtual server platform.

¹<http://webtorials.com/abstracts/understanding-cloud-computing.htm>, to be published by November 2009

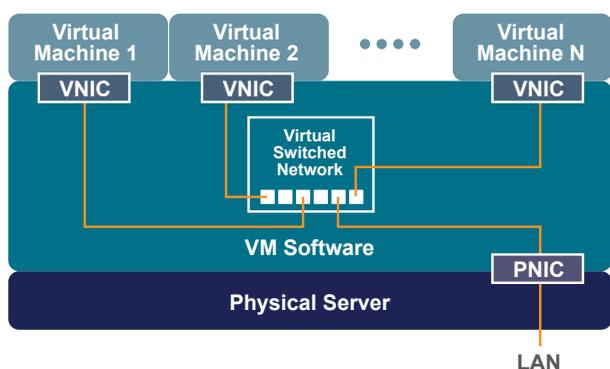
Figure 1: Overview of Virtual Machine Technology



The VM software creates and controls the virtual machine's virtual subsystems. It also takes complete control of the physical machine, and provides resource guarantees for CPU, memory, storage space, and I/O bandwidth for each guest OS. The VM software can provide a management interface that allows server resources to be dynamically allocated to match temporal changes in user demand for different applications. The majority of IT organizations that have currently implemented server virtualization have based that implementation on software from VMware®. A number of other vendors, however, have recently entered the market, including Citrix, IBM, Microsoft, and Sun.

Figure 2 shows how Ethernet Networking Input/Output (I/O) is typically virtualized by VM software. The VMs within a virtualized server typically share a conventional physical Ethernet Network Interface Card (NIC) to connect to the data center LAN. The VM software provides each VM with a virtual NIC (vNIC) instance complete with MAC and IP addresses and creates a virtual switched network to provide the connectivity between the vNICs and the NIC. The virtual switch (vSwitch) inside of the server performs the I/O transfers using shared memory and asynchronous buffer descriptors in similar fashion to a shared memory Ethernet switch. With this software-based I/O virtualization, the VM software provides the data path for inter-VM traffic as well as for traffic destined for the external LAN.

Figure 2: Software-Based Virtual Ethernet I/O



In the data center, an entire n-tier application architecture can be implemented on a single high performance, multi-core virtualized server. One or more VMs can be allocated to each server function and additional VMs can be dedicated to virtual appliances such as a firewall. In instances like this, the vast majority of the inter-VM traffic stays within the server and hence does not traverse

the physical data center network. This is in contrast to the traffic flow of a typical multi-tier application in which each tier runs on a dedicated physical server. As will be discussed below, there are some significant management challenges associated with managing inter-VM traffic flows.

The Benefits of Server Virtualization

There are numerous benefits that can be derived from server virtualization, including:

- **Server Consolidation**—A single physical virtual server can support 10 or more VMs, allowing numerous applications that normally require dedicated servers to share a single physical server. This facilitates reducing the number of servers in the data center while simultaneously increasing average server utilization from as low as 5-10% up to 60-70%.
- **Flexible Server and Application Provisioning and Movement**—Whereas it can take weeks or months to provision a new physical server, a VM can be established in seconds. In addition, a production VM can be transferred to a different physical server, either to a server within the same data center or to a server in a different data center, without service interruption. This capability enables workload management and optimization across an IT organization's virtualized data center(s). This capability also helps to:
 - Streamline the provisioning of new applications
 - Improve backup and restoration operations
 - Enable zero-downtime maintenance
- **Lower TCO**—Server virtualization allows significant savings in the both CAPEX (i.e., costs of server hardware, SAN Host bus adapters, and Ethernet NICs) and OPEX; i.e., server management labor expense, plus facility costs such as power, cooling, and floor space.
- **Green IT**—In some instances, part of the motivation to deploy virtualized servers is to reduce the number of servers, not just to reduce the cost of those servers, but to also reduce the associated carbon footprint as part of a Green IT initiative.
- **Enable Other Initiatives**—In order to more effectively manage and control user desktops, many IT organizations are implementing virtual desktops. Virtualized servers are a prerequisite for virtualized desktops because a virtual desktop implementation requires that a VM in a data center server hosts a complete user desktop including all its applications, configurations, and privileges.

The Challenges of Server Virtualization

The Fractal Data Center

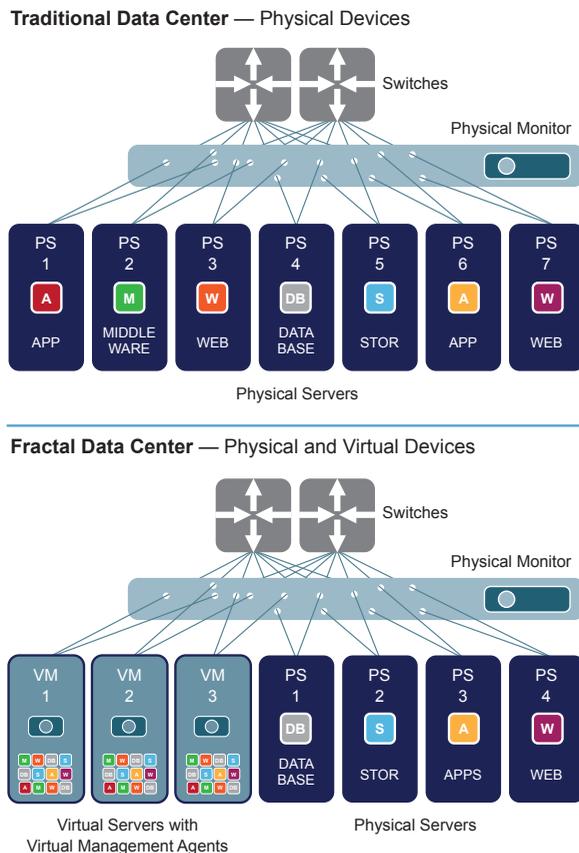
One way to think about the current generation of data centers, and the related management challenges, draws on the concept of a fractal². A fractal is a geometric object that is similar to itself on all scales. If you zoom in on a fractal object it will look similar or exactly like the original shape. This property is often referred to as self-similarity.

The relevance of fractals is that the traditional data center is comprised of myriad physical devices including servers, LAN

²<http://www.pha.jhu.edu/~ldb/seminar/fractals.html>

switches and firewalls. The virtualized data centers that most IT organizations are in the process of implementing are still comprised of physical servers, LAN switches and firewalls. In addition, these data centers house virtualized servers, which themselves are comprised of myriad virtualized devices including virtual machines, a virtual LAN switch and in many cases virtual firewalls. Hence, if you take a broad overview of the data center you see certain key pieces of functionality. As shown in Figure 3, if you were to then zoom inside of a virtualized data center server you would see most of that same functionality.

Figure 3: The Fractal Data Center



Traditional Management Challenges

In order to successfully deploy a fractal data center, IT organizations need to respond to a set of traditional management challenges. These traditional management challenges apply at two levels. Those levels are:

- The physical data center level**
 IT organizations will have the same management challenges with the physical components of their data center as they have always had.
- The virtual data center level**
 IT organizations will have the same management challenges with the virtual devices inside of a virtualized server as they have typically had with the physical components of their data center.

An example of a traditional management challenge is the growing requirement to discover and monitor all types of applications and services. From an application-focused perspective, this includes complex applications such as SAP as well as custom-built applications. It also includes broadcast, multicast and unicast applications as well as peer-to-peer applications. From a service-focused perspective, this includes services such as voice and email.

Another example of a traditional management challenge is the requirement to baseline performance. Baselineing is important because it provides a reference from which service quality and application delivery effectiveness can be measured. It does so by quantifying the key characteristics (e.g., response time, utilization and delay) of applications and services as well as various IT resources including servers, WAN links and routers.

Application profiling and response time analysis provides the capability for IT organizations to understand the behavior of critical applications and services. That capability helps an IT organization detect and resolve performance issues before they impact end users, which is important in both a traditional as well as a virtualized environment. This capability also helps an IT organization make the transition to a virtualized environment by identifying which applications make sense to run in a virtualized environment. In particular, an IT organization will not be successful with server virtualization if it loads onto a physical server a large number of applications, all of which are either CPU-intensive or I/O-intensive. Application profiling provides the insight IT organizations need to avoid this situation.

Activities such as baselining and application profiling can help an IT organization identify and remedy problems before they impact end users. However, there will always be instances in which IT organizations need to respond to a problem reactively. In order to troubleshoot effectively, IT organizations need to monitor and alert on key performance indicators such as responsiveness, packet loss and errors. All of these activities must be performed whether or not the server has been virtualized. If the server has been virtualized, IT organizations also need to be able to troubleshoot on a per-VM basis.

New Management Challenges

In addition to traditional management challenges, as IT organizations implement fractal data centers they also encounter a new set of management challenges. For example, as was previously noted, one of the advantages of a virtualized server is that a production VM can be transferred to a different physical server, either to a server within the same data center or to a server in a different data center, without service interruption. While this capability provides significant value, it can be challenging for IT organizations to ensure that the migrated VM retains the same security, storage access, and QoS configurations and policies. Transferring production VMs can make it difficult for IT organizations to keep track of where each VM is located and can also cause significant shifts in network traffic. All of these factors can significantly complicate the troubleshooting process.

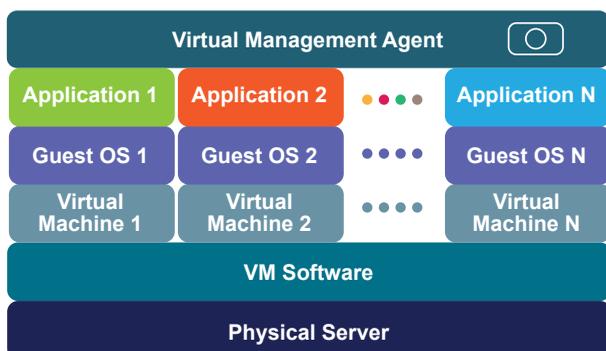
Another emerging management challenge stems from the fact that in the traditional data center environment, IT organizations typically leverage the management functionality in their physical LAN switches to manage the traffic flows between servers. Unfortunately, in most instances the vSwitch that resides in a virtualized server provides little if any management insight. As such, IT organizations lose visibility into the traffic flows between and among the VMs on the same physical server.

For example, while most embedded virtualization management tools can identify the total volume of traffic within the entire virtual environment, they cannot provide information on individual network services such as HTTP or FTP. This lack of management insight can dramatically limit the ability of the IT organization to be able to effectively monitor performance or resolve performance problems. In addition, the current generation of vSwitches provides little if any inherent security visibility. As a result, it is typically not possible for an IT organization to know if unsanctioned network services have been enabled, or if worms or other forms of malware are propagating within the virtual environment.

Managing Application Performance on Virtual Machines

IT organizations need the same visibility into the performance of an application running on a VM as they have into the performance of an application running on a physical server. A key requirement for regaining visibility into the performance of virtualized applications is to be able to see applications within their virtual environment in a manner that is consistent with how these applications were viewed in a physical server environment. As shown in Figure 4, this is best achieved by instrumenting the virtualized server with a data source that can both discover and see application traffic. Ideally this instrumentation would look at actual packets to provide the same perspective as a physical network-based packet-flow data source would. This requires the virtualization of physical probe technology into a virtual appliance that can be deployed within a virtual server. An effective virtual probe should behave exactly like its physical counterpart, except that it resides within the virtual server and looks at the interaction of all applications over the server's virtual network.

Figure 4: Embedding Management Data Sources into Virtual Machines



One of the characteristics of a virtualized server is that each virtual machine only gets a fraction of the virtual server's compute resources (i.e., CPU, memory, storage). As a result, in order for a virtual probe to be effective it should not consume significant resources. In addition, the virtual probe should reside in its own VM and be protected from impacting, and being impacted by other applications.

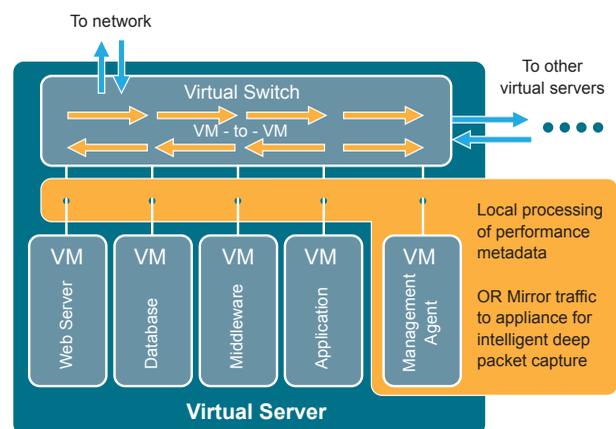
The way that a virtual probe works is identical to how many IT organizations monitor a physical switch. In particular, the vSwitch has one of its ports provisioned to be in promiscuous mode and hence forwards all inter-VM traffic to the virtual probe. As a result, the use of a virtual probe enables the IT organization to reestablish the necessary visibility into the inter-VM traffic to effectively and proactively manage application and service traffic within a virtual server.

Not all virtual probes or virtual monitoring technology is the same. They can differ in a variety of ways. For example, they can provide:

- Near real-time monitoring of application layer details within the virtual machine using flow-based data as the source
- Real-time monitoring of all transaction data within the virtual machine using packet-based traffic as the source
- Streaming of the packet-based traffic to an external hardware collection appliance for post-event forensics analysis
- A combination of both real-time monitoring and streaming packet flow analysis for complete coverage

Having the choice to combine both real-time monitoring with post-event forensics against the same stream of packet flow traffic offers the most flexible solution to performance issues in both a traditional as well as a virtualized environment.

Figure 5: Virtual Monitoring Technology



As noted, what is true in the traditional data center environment is also true in a virtual environment –packet level data provides the richest insight into the dynamics of service delivery that encompasses the applications, the network and the users. For example, packet-level data can provide insight into the performance of a wide range of applications, including well-known applications such as Lotus Notes, custom-developed applications, peer-to-peer applications, and complex applications such as SAP,

Exchange and Citrix. Monitoring packet-based traffic also provides a rich set of data that can be used to:

- Generate targeted, customizable, scheduled and ad hoc reports
- Support critical performance management processes, including application monitoring, network monitoring, capacity planning, troubleshooting, fault prevention, service-level management, modeling, and billing.

This data can also be retained in an external deep packet capture platform and used to perform forensic analysis of all of the traffic between the virtual applications that transits the vSwitch.

While a virtual probe that monitors packet-level data on a per-virtual machine basis is helpful, it needs to be part of a broader solution. That solution should also include the ability to store significant amounts of data and utilize sophisticated management analysis of the data. This enables IT organizations to troubleshoot an intermittent problems based on real historical data.

As noted, the deployment of a virtual probe that provides packet-based data allows an IT organization to have the same management capabilities in a fractal data center that it has in a traditional data center. There is, however, one characteristic of the traditional IT management environment that should not be extended into the fractal data center. That characteristic is that the traditional management environment is characterized by multiple discrete management tools with little or no interaction. That siloed approach to management increases cost and elongates the time it takes to repair complex problems. As IT organizations develop a strategy for managing their fractal data centers and virtualized applications, they should look to acquire the tools they need from their existing vendors to gain a more unified perspective of their application delivery environment and avoid the temptation to add yet one more management stovepipe.

Summary

The virtual data centers that most IT organizations are implementing can be thought of as being fractal. By that is meant that if you take a high level view of a virtual data center environment you see a physical IT infrastructure that is comprised of devices such as servers, LAN switches and firewalls. It also means that if you drill down inside of a virtualized server you will see most of that same functionality - only the functionality is virtualized; e.g., VMs, virtual switches and virtual firewalls.

IT organizations have the same management challenges with the traditional physical devices in their data centers that they have always had; e.g., application discovery, baselining, application profiling and response time analysis, troubleshooting, etc. IT organizations also have the same traditional management challenges inside of a virtualized server, as they now have the requirement to manage the virtual devices that reside inside of the virtualized servers. In addition, IT organizations have some new management challenges, such as those that are created by the loss of management insight into the inter-VM traffic that typically occurs when servers are virtualized.

In order to overcome the challenges associated with server virtualization, IT organizations need to deploy monitoring instrumentation within their virtualized servers. The required instrumentation is a virtual probe that helps to provide the IT organization with the visibility that was lost when the servers were virtualized. As a result, the virtual probe enables consistent manageability of application performance within a virtual environment. In order to provide the greatest value this virtual probe needs to provide packet-level data on a per-VM, per-application basis. This level of detail is necessary in order to provide the IT organization with a unified view into the performance of the application as seen by the end user, as well as into the performance of the services and the network—both physical and virtual.

Data from a virtual probe must feed into a broader, end-to-end solution that looks at performance from an application, network and user perspective. Since most IT organizations will likely deploy virtual probes in most, if not every virtualized server, the overall solution needs to have the scale and capacity to track, store and analyze huge amounts of data from a multitude of data sources. A solution of this type enables IT organizations to perform forensics in order to troubleshoot a problem based on real historical data and it also enables the IT organization to know when a VM has been moved from one physical server to another. As IT organizations deploy this type of solution they should look to acquire the tools they need from their existing vendors and avoid the temptation to add yet one more management stovepipe.



Ashton, Metzler & Associates
P.O. Box 1640
Sanibel, FL 33957