JUNE 2011 / VOL. 2 / NO. 3

NETWORK evolution

BUILDING THE INFRASTRUCTURE TO ENABLE THE CHANGING FACE OF IT

TURF WAR

SYSTEMS VS. NETWORKING: THE STRUGGLE FOR THE VIRTUAL EDGE

PLUS:

IDEA LAB

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

TechTarget

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

idealab

Where evolving network concepts come together

Virtualization Demands a Full IT Infrastructure Audit

TODAY'S NETWORK ENGINEERS have a tough challenge before them: They've got to provide high-availability network infrastructure that can handle traditional applications, as well as virtualized multihost systems and access by smart devices.

Both virtualization and smart device management demand an IT organization that is not broken into separate camps for applications, systems, storage and networks. What most people don't realize is that this type of unified organization starts with an IT infrastructure audit and network documentation strategy that reaches across all of these camps.

Infrastructure inventory

The first step to creating a unified IT department that can build application-aware infrastructure is creating an inventory system that reflects all of these resources that were once held in silos. Your unified IT organization should be able to answer the following questions:

- Does your documentation describe the server interdependencies that exist to support all of the corporation's applications?
- How many different applications utilize the network's bandwidth, and how much capacity do they typically use?
- Have network performance and transactional baselines been performed on core network-dependent applications?
- Does an inventory exist of each server's network service and utilization profile?
- Does an inventory exist of the different operating systems, applications and network service daemons?
- What percentage of applications utilizes data encryption on the network?

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

- If an application or system were compromised on the network, what kind of network layer events would indicate this condition?

If all seven questions can be answered, look at step 4 below. If five or more of the questions above could be answered, look at step 3. If only four or less could be answered, look at step 1. If none of the questions could be answered ... you've got a very large problem.

The following steps help work toward a unified documentation and inventory:

STEP 1: Conduct a documentation review. When implementing network layer services it is critical to know what applications, protocols, data exchange and service dependencies exist between the servers that support applications.

STEP 2: Conduct an audit. Teams need different kinds of data to perform their functions. A collective effort is required here to design a bottom-to-top audit.

STEP 3: Update documentation. Once the audit is complete, that data should be used to update documentation. This sounds obvious, but collected data often collects dust. Furthermore, an accommodation to refresh the data in a timely manner needs to be made. Otherwise, you will quickly find yourself back at step 2.

STEP 4: Use this information to pilot a new technology. With all of this information in place, you can asses a new technology, such as virtualization. Your audit data should serve as your baseline and your success criteria should in part be based on how those baseline elements will be affected by implementation. ■

FACTOID

Good, thorough documentation is important to any IT project or implementation; it's particularly important to server virtualization projects and the impact of those projects on the data center network. For a virtual implementation, planning is the most important step. Plan and then plan some more.

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

vCloud and VEPA Both Fail in Virtualization Networking

THE PROBLEM began with VMware, which had to implement switching functionality in its ESX hypervisors to enable connectivity between virtual machines running within the same physical server. The company could have implemented Layer 3 switching, but it decided instead to take the easy route and build a substandard switch, which after several years still lacks feature parity with SMB products you can buy at Office Depot.

But embedding third-party Layer 2 switches in physical servers hasn't worked as a virtualization networking solution either. Since servers could no longer be treated as end-hosts, physical switches had to peer with them on trunked links. We all know large-scale bridging environments tend to be fragile—and coordinating changes between

ACCOMMODATING SERVER VIRTUALIZATION

What changes is your company making in its network to accommodate server virtualization?

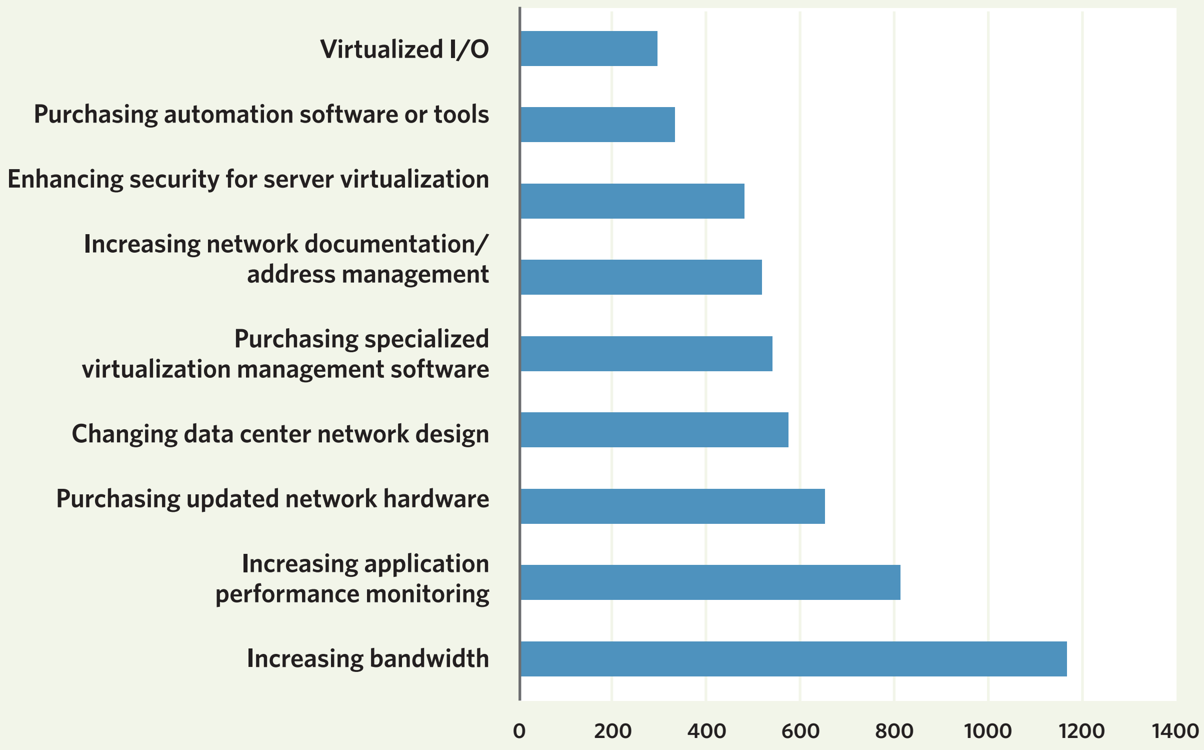

SOURCE: TECHTARGET NETWORKING PRIORITIES SUVREY, NOV. 2010

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

two teams (when the server team handles whole VMware configurations) made deployment and troubleshooting even more fun.

Virtualization networking solutions emerge—but still don't work

VMware has tried to address these virtualization networking problems with the vCloud Director, which implements switching, routing, firewalling, NAT and even DHCP servers within the VMware framework. vCloud could have been a great product, but instead VMware decided to use proprietary protocols and implement most of the networking functions in virtualized appliances. For example, its MAC-in-MAC encapsulation doesn't follow any standard known to the industry, though the 802.1ah standard (specifying MAC-in-MAC encapsulation) was published two years ago.

VMware's GUI definitely satisfies the eye candy requirements (and must look fantastic in demonstrations), but vCloud fails to solve any of the significant problems: VM mobility still requires bridging; bridging is no more scalable than it was before vCloud, and the traffic trombones are still winding their way around your data center.

Meanwhile, Cisco took a stab at networking for virtualization and launched the Nexus 1000V (which became part of the VN-Link/VN-Tag strategy) with a subliminal "take the control back" message to the networking team. And other networking vendors are hastily inventing Virtual Ethernet Port Adapters (VEPA)—a technology that would pull all the traffic out of the hypervisor and let the first-hop switch handle all the necessary aspects of bridging, VLANs, routing, QoS and security. Cisco's VN-Tag story is being considered by the IEEE, along with the VEPA standard (802.1Qbg). The VEPA vendors are quick to tell you how fantastic your life will be after VEPA gets implemented, but they keep quiet about an important detail: VEPA will need hypervisor support. While VMware has supported the Nexus 1000V, I doubt the company will jump at the opportunity to implement VEPA in its vSwitch now that it has a competing product.

> VEPA vendors are quick to tell you how fantastic your life will be after VEPA gets implemented, but they keep quiet about an important detail: VEPA will need hypervisor support.

What's more, VEPA does not solve any of the aforementioned problems; it just gives you a different GUI/CLI—

meaning you're configuring virtual NICs on the physical switches, not in vCenter/vCloud.

With all the confusion and half-baked solutions that vendors are generating, one has to wonder about the real reason for their frantic activities. The answer is simple: All of them are after your limited budget.

If VMware persuades you to go with vCloud, you will need dumb, low-cost commodity switches that are able to perform rudimentary bridging between ESX servers. If the networking vendors sell you on VEPA, you won't buy vCloud and vShield licenses, but you will have to invest in smarter switches.

But in the heat of the marketing arguments, everyone forgot what the customers really want: a scalable, manageable solution supported by all parties. Easy end-to-end orchestration would be almost too much to hope for, but let's add it to the wish list. Would it be too much to ask you, dear vendors, to focus on that? ■

VIRTUALIZATION DRIVES INVESTMENTS

What data center networking technologies do you expect to invest in during the next 12 months?

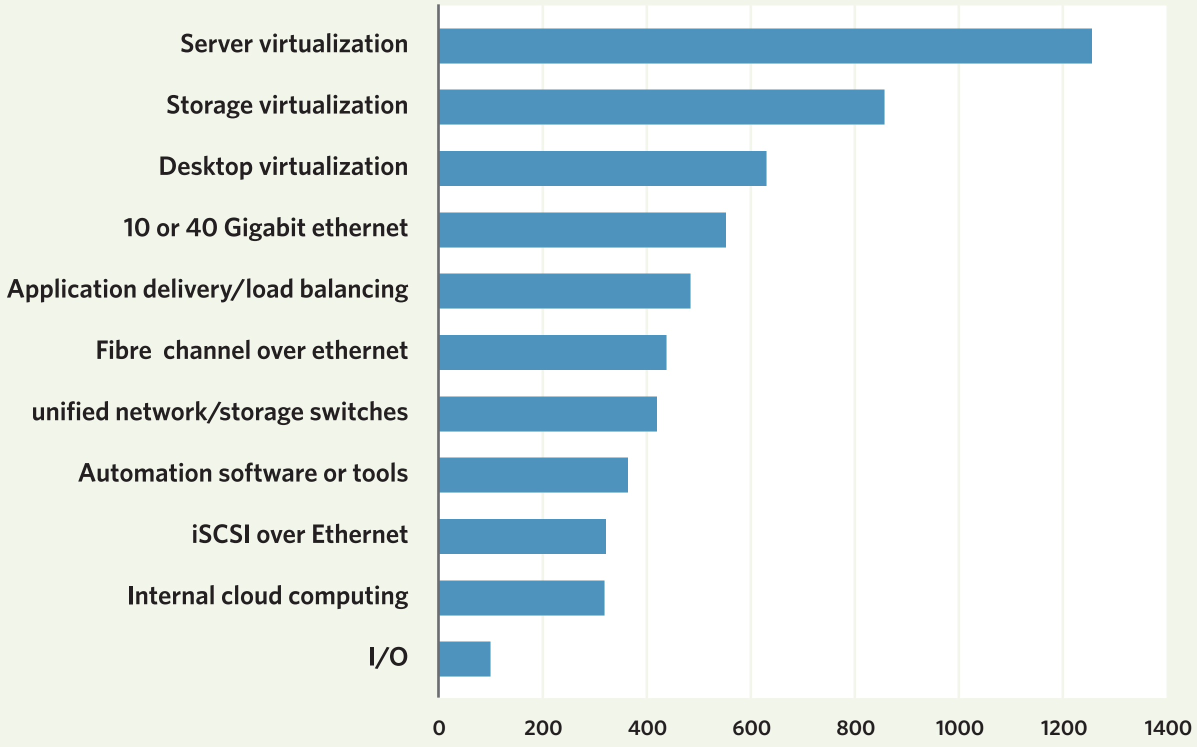

SOURCE: TECHTARGET NETWORKING PRIORITIES SUVREY, NOV. 2010

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

Understanding Desktop Virtualization Clients

NETWORK ENGINEERS must firm up strategies to prioritize traffic, manage bandwidth and ensure application performance to virtual desktops. And the needs can change depending on the type of client used and whether virtual desktop infrastructure (VDI) is server- or desktop-centric.

Desktop virtualization can involve a host of different clients—all of which have different bandwidth and support needs:

- **Thick client:** Thick clients are standard desktops or laptop PCs that use an operating system in addition to running a virtual image or terminal services. In a typical scenario, the users are part of the help desk function for a period of the day, during which time they run a separate virtual desktop image or just terminal services. The downside of this option is that IT still has to service full desktops that require regular updates, patches and support. Many companies use thick clients to pilot desktop virtualization.

- **Thin client:** Thin clients use a stripped-down OS (Windows CE and Linux) and support user interface functionality (video, audio, USB, printer, mouse and keyboard.).

- **Net client:** Net clients are in between thick and thin clients. They run stripped down OSes or can be mobile devices like the iPhone running Citrix receiver for iPhone or Wyse's Apple iPhone App.

- **Zero client:** With a zero client, all of the OS or applications run on the server with nothing at the desktop. The zero client is essentially an extension cord that extends the keyboard, mouse, screen, audio, printer and USB ports to the desktop. This solution works best in hazardous environments and environments where maintenance is a challenge. Zero-client desktops are excellent for factory floors, chemical plants and hospitals. Zero clients also draw much less power than a conventional desktop or even a thin client. ■

FACTOID

The benefits of desktop virtualization are enormous but they can seem clouded by challenges resulting from optimizing application performance using existing bandwidth.

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

How 802.1Q VLAN Tagging Works

THERE ARE SEVERAL methods for tagging vSphere VLANs, but they are differentiated by where the tags are applied. Virtual Machine Guest Tagging (VGT) mode does this at the guest operating system layer, External Switch Tagging (EST) mode does it on the external physical switch, and Virtual Switch Tagging (VST) mode does it inside the VMkernel. The differences among the VLAN tagging modes are outlined below:

1 **Virtual Machine Guest Tagging (VGT mode):** With this mode, the 802.1Q VLAN trunking driver is installed inside the virtual machine. Tags are preserved between the virtual machine networking stack and external switch when frames are passed to and from virtual switches.

2 **External Switch Tagging (EST mode):** With this mode, you use external switches for VLAN tagging. This is similar to a physical network, and VLAN configuration is normally transparent to each individual physical server. The tag is appended when a packet arrives at a switch port and stripped away when a packet leaves a switch port toward the server.

3 **Virtual Switch Tagging (VST mode):** With this mode, you configure port groups on a virtual switch for each VLAN and connect the vNIC of the virtual machine to the appropriate port group. The virtual switch port group tags all outbound frames and removes tags for all inbound frames. It also ensures that frames on one VLAN do not leak into a different VLAN. The VST mode is the one that is most commonly used with VLANs in vSphere because it's easier to configure and manage. It also eliminates the need to install a specific VLAN driver inside a virtual machine, and there is almost no performance impact from doing the tagging inside the virtual switches. ■

FACTOID

VST mode is most commonly used with VLANs in vSphere because it's easier to configure and manage. It also eliminates the need to install a specific VLAN driver inside a virtual machine.

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

TURF WAR:
SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

UNTIL RECENTLY, Matthew Norwood, a network architect with a Tennessee-based healthcare enterprise, simply couldn't ensure application performance on the network once it hit the virtual environment. That's because he had no control or visibility of the virtual traffic between ESX servers and was unable to extend his network application performance monitoring tools into the virtual environment.

Norwood's scenario is not uncommon.

Server virtualization has most network engineers in a tizzy. On one hand, they're forced to build large Layer 2 domains for virtual machine (VM) mobility, as well as high-bandwidth networks with hundreds of 10 gigabit server-facing ports for ballooning VM-to-server ratios. On the other hand, they've lost complete administrative control over the virtual edge to systems administrators and their hypervisor-based virtual switching.

Without that administrative control of the virtual edge, network engineers say they can't effectively enable large-scale virtualization and elastic cloud computing environments.

> Network engineers say they can't effectively enable virtualization and cloud computing environments without control of the virtual edge.

Regaining the edge is crucial for automation, security and basic accountability. What's more, the scenario has networking and systems teams in a mess of finger pointing when things go wrong in

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

the virtual environment.

"If an application group comes to the network people and says, 'This app is slow or this environment is slow,' right now we reach a certain point where we have to throw up our hands and say, 'All I can tell you is that it's fine going from here into this particular device,'" said Norwood. "But I can't see anything between two virtual machine guests that are on the same VLAN. If it never leaves a physical server, I can't see it and I can't do anything about that."

So even though Norwood's company had invested heavily in Application Performance Monitoring (APM) software, his team couldn't instrument that software at the virtualized network edge.

Norwood has struggled with mirroring and spanning out VM traffic into his APM platform. Any VM-to-VM traffic that stays within a VMware ESX host server or hops between ESX hosts is mostly invisible to his APM software. If the hypervisor's virtual switch is the culprit behind application performance degradation, neither the networking team nor the server team troubleshoot it effectively.

"If you don't have a good network monitoring package in place that can actually see some of these things that are going on in virtual environments, there could be a huge amount of finger pointing," Norwood said.

METHODS FOR RECLAIMING THE VIRTUALIZED NETWORK EDGE

Cisco Systems was one of the first vendors to offer a solution to the visibility and control problem created by server virtualization. The Nexus 1000v distributed virtual switch is a software construct that can replace VMware's embedded virtual switch. Network admins can manage the Nexus 1000v just like a physical switch, logging in to it via SSH.

> Regaining control of the network edge is crucial for automation, security and basic accountability.

No other major networking vendor has produced a similar virtual switch, although an open source alternative—Open vSwitch—has emerged, with contributions from XenServer hypervisor maker Citrix and Nicira Networks, a network virtualization vendor that recently emerged from stealth mode.

Multiple vendors are working on methods for extending visibility and control over the virtualized network edge via the nearest hardware switch. Many of these methods depend on emerging IEEE standards,

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

Cisco-supported VN-Tagging (802.1Qbg) and HP-supported Virtual Ethernet Port Aggregation or VEPA (802.1Qbh). VN-Tagging technology adds a new header to an Ethernet frame, which extends visibility and control over virtual ports to the nearest VN-Tag capable switch. VEPA, on the other hand, modifies a hypervisor's virtual switch, forcing it to send all traffic upstream to the nearest switch regardless of whether that traffic is meant to move east-west among VMs on the hypervisor host.

Extreme Networks has embraced VEPA, also known as Edge Virtual Bridging (EVB), as part of its Direct Attach solution. Force10 Networks has also announced support of VEPA, as has Juniper Networks, which is developing Virtual Control, a Web-based app that runs on the company's Junos Space network application platform and integrates with VMware vCenter to give network managers administrative control over virtual networks for VMs.

RETOOLING APPLICATION PERFORMANCE MANAGEMENT FOR VIRTUALIZED ENVIRONMENTS

Norwood plans to put at least two instances of the Nexus 1000v into production to manage networking for more than 100 ESX hosts at his healthcare enterprise in order to extend his APM software's visibility into virtual network traffic on VMware ESX hosts.

"With me being able to manage the 1000v, I can span out traffic [from ESX hosts] and get NetFlow dumps," he said. "Now I can give you an educated response. Here's where we think the problem is."

> Many vendors are working on methods for extending visibility and control over the virtualized network edge.

Norwood has the Nexus 1000v in his development environment and plans to roll it into production in about a month. He also plans to use the Nexus 1010 hardware appliance, which Cisco designed to host the management and control plane of the Nexus 1000v, the Virtual Supervisor Module (VSM). The data plan of the Nexus 1000v, the Virtual Ethernet Module (VEM), is distributed as different instances on each ESX host managed by an individual VSM. Many enterprises run the VSM as a virtual machine on an ESX host, but VSMs demand a large amount of processing and memory, stretching the limits of a VM on an ESX host. With the Nexus 1010, enterprises

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

can run four VSMs on a single box without limitations on computing power.

GAINING BACK CONTROL TO SECURE THE VIRTUALIZED ENVIRONMENT WITH VLANS

Gaining control of the network in the virtual environment is also important when it comes to extending VLANs for security purposes. Generally VLANs are used on physical switches to segregate traffic and apply security policy into the virtual environment. But while virtual switches embedded in hypervisors have basic VLAN functionality, hypervisor-based VLANs aren't secure or scalable enough.

For Bart Falzarano, chief information security officer for Walz Group, a compliance and document management software and services company in Temecula, Calif., using traditional VLANs was not solving the security problem. "If a compromise goes down to the hypervisor level, you have a problem. Even with a separate VLAN, all your virtual machines are there [on the host server]. There's no further segregation," he said.

Many enterprises still use physical VLANs to apply security and policy to virtual infrastructure, restricting virtual workloads to servers within a specific VLAN. However, these physical VLANs are inflexible. If computing demand outstrips the capacity of the server pool in a specific VLAN, applications crash. If a system administrator needs to move VMs to another server rack or data center for maintenance or business continuity reasons, the network engineer needs to be ready to set up a new VLAN to facilitate that migration. This approach is too manual and inflexible, especially for organizations that are building private clouds.

> "It's important that your staff is cross-trained."
> —BART FALZARANO

The Walz Group recently deployed a private cloud with Cisco's Unified Computing System Blade Server Chassis, VMware software and NetApp storage. Knowing that he needed robust and automated VLAN capabilities with visibility and control over a virtualized network edge, Falzarano adopted Cisco's Nexus 1000v.

"If this were a bare metal, non-virtual environment, you'd still have to include segregation between systems. That follows the standard application service provider model," Falzarano said. "The front-end systems are in tier one in the DMZ.

Then you have your tier-two applications servers in the middle tier. Then you have your most inside trusted tier, your database, storage and data repositories. You'd segregate those with a security component. You'd set up VLANs on your physical switches. To take that same type of design and apply it in the virtual space, you need a component like the 1000v."

RELATIONS BETWEEN NETWORKING AND SYSTEMS TEAMS WILL MEAN CROSS-TRAINING

While IT shops scramble to figure out management of virtualization networking, the lack of visibility into the virtualized network edge has degraded relations between network engineers and systems engineers. But good leadership can counter that trend.

"A lot comes down to policies and procedures and change management," Norwood said. "The Nexus 1000v will facilitate that as well. I'm intimately more aware of what's going on in the virtual environment, so when there's a problem it's not just the VMware guys trying to figure it out. The network team can get in there and say, 'Here's what I'm seeing.' So we're able to collaborate a lot more and resolve problems faster than we would have before."

In fact, slowly but surely, the Nexus 1000v and cloud deployment in general have blurred the lines between systems and network engineering at Walz Group, Falzarano said.

"[In a private cloud] there's been a confluence of different technology expertise," Falzarano said. "It's important that your staff is cross-trained. We view the environment as being very flexible. You're not only dealing with VMware virtual machines, but also service profile templates, which means you're not managing individual bare metal anymore. You're managing templates. We don't silo it."

That cross-training is essential when deploying a high-availability pair of Nexus 1000v VSMs on an ESX host within his server blade enclosure, Falzarano said. Network engineers will use traditional networking tools to manage their components. But when those network components exist as VMs on an ESX host, a network engineer must work with systems engineering tools, too.

"He can still manage [the Nexus 1000v] through SSH, but in understanding where the VSMs are—if they've moved—that's where he would use tools like vCenter," Falza-rano said. That won't work without cross-training and cooperation. ■

SHAMUS MCGILLICUDDY is the News Director for TechTarget Networking Media.

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

TROUBLESHOOTING SLOW network performance is difficult enough in a physical environment, but when virtualization comes into the mix, things get even more complicated.

After all, in a bare metal data center, network administrators have lots of tools at their disposal to follow packets from end to end, but in a virtualized data center, east-west network traffic between virtual machines on the same physical server is invisible to physical network infrastructure, blocking a complete view of the network.

Fortunately, there are ways to let the network team peer into the black box of the virtualized environment using a patchwork of tools—and cooperation between the network and virtualization teams.

GETTING INSIDE THE VIRTUAL ENVIRONMENT WITH VIRTUAL SWITCHES

Hypervisors on virtualized servers use built-in virtual switches (vSwitches) to manage traffic among virtual machines. These software vSwitches are designed to replicate the functions of a physical network switch, including support for port monitoring. However, they typically lack the deep functionality of a traditional hardware-based network switch.

Cisco Systems' virtual Nexus 1000v switch and the open source Open vSwitch offer more features and functionality than hypervisor-embedded virtual switches, and they more closely resemble their physical counterparts. As a result, they offer network administrators centralized management and visibility of both physical and virtual network infrastructure.

But virtual switches also have some downsides. Every function they offer adds additional workloads and saps CPU cycles from the virtual server environment. Network administrators need to balance the need for port monitoring and network troubleshooting with the overhead that such monitoring will pose to compute resources.

INTEGRATING NETWORK MANAGEMENT PLATFORMS WITH HYPERVISORS

Some vendors offer data center network management platforms that can simply gather information about virtual machine activity from the hypervisor. VMware offers an application programming interface (API) on its management platform that allows third-party management tools to track virtual machines. Network administrators generally depend on their management tool vendors to do this integration, however. And this approach does not necessarily provide the depth of information that pure packet captures provide. Integrating with the virtualized environment, however, does contextualize the impact that virtual machines are having in the network and helps define the path that an application is taking through the environment.

DEPLOYING NETWORK PROBES INSIDE THE VIRTUAL ENVIRONMENT

Network diagnostic and forensic vendors are developing virtual probe software that can provide instrumentation within a virtualized environment. WildPackets' OmniPeek analysis solution, for instance, now features OmniVirtual, a virtual network probe that installs on each virtual machine as it communicates with both virtual and physical elements of the network. OmniVirtual transmits data to a centralized network analysis appliance.

> The network and server teams have to ensure that virtualized network probes are deployed on every virtual machine in the environment.

Combining virtual and physical network probes delivers complete network visibility. These virtual services, however, are specific to each network analysis product and will only provide visibility to the vendor's network management products.

The network team will also have to work closely with the server team to ensure that virtual network probes are deployed on every virtual machine in the environment. Any workloads running without a probe will remain invisible to the network management tool.

FORCING VIRTUAL MACHINE TRAFFIC BACK INTO NETWORK HARDWARE

A new set of standards pending with IEEE offers administrators virtual network visibility without the CPU overhead associated with virtual switches or virtual probes. Edge Virtual Bridging (EVB), also known as 802.1Qbh, brings together hardware, software and protocol standards to simplify and automate the links between physical and virtual Layer 2 networks in the data center. The standard will allow physical and virtual switches to talk to one another and share configuration information.

The EVB standard will also include a technology called Virtual Ethernet Port Aggregation (VEPA), which instructs virtual switches to send all traffic from virtual machines upstream to the nearest physical network switch. This exposure of the virtual machines to the physical network allows network administrators to apply traditional network analysis and management.

VEPA also includes provisions for communications between virtual machines on the same physical server and network hardware, known as a reflective relay or a hairpin turn. In addition, VEPA allows a physical switch to send data back across the same network port it came from. Administrators can deploy VEPA either through upgraded virtual switch software or within hardware on supported network interface cards. Depending on the workloads and utilization on a given server, most enterprises will deploy VEPA as a mix of software and hardware.

Edge Virtual Bridging can also help automate network configuration and policy management through the Virtual Station Interface Discovery Protocol (VDP). VDP allows the network to know of the movement of a virtual machine in advance of the move and automate network configuration for the destination hypervisor host.

PULLING IT ALL TOGETHER

Each of these solutions gives the networking and server teams the ability to focus on their own domains, but they all require collaboration between the two teams to implement. Enterprise infrastructure can no longer be siloed. It must function as an ecosystem. Efforts to automate the virtual environment will have to be matched with automation on the data center and edge networks. Likewise, silos of IT management will have to be broken down and collaboration between network, server and storage teams will have to be just as agile as the infrastructure they hope to build. ■

MICHAEL BRANDENBURG is the technical editor for TechTarget Networking Media.

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

WHY *ARE* NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

NETWORK ENGINEERS are tired of being viewed as plumbers—especially when it comes to virtualization. After all, the job of supporting virtualized traffic goes so much deeper than providing an always-available pipe. Systems teams understand the complexity of a virtualized environment, but don't always see the network admin's role in managing that process. The split results in ineffective troubleshooting strategies and network architectures that don't always better a virtualized environment.

Virtualization architect **Bob Plankers** recognized that problem amongst his own ranks at a large Midwestern university and set out to change things by opening up conversation—and management tools—between the two teams. The result? A new network architecture and an effective approach to virtualization management.

Is there really a disconnect between networking and systems folks when it comes to managing virtualization?

Absolutely. Virtualization or systems people don't include the network guys in what's going on. In traditional data center models, when workloads stayed in one place and things were static, [the networking team was aware], but to have the ability to do vMotion, moving VMs around in a data center without any notice, is kind of distributing to them. I don't want to liken network guys to plumbers, but if they're maintaining pipes, and all of a sudden you've moved a bunch of water flow from one place to another, they're wondering what's going on. Some network guys don't understand what is possible with virtualization or what their systems guys are doing.

But systems guys don't understand why network guys care. They just look at the network as a pipe. They say, 'Well, there's one across my data center so I'll put my ESX hosts there, and I'll put [one] of my hosts here,' and they have no concept of the infrastructure that's required to connect the switches and how much bandwidth there is. They just see it as this always-on service, which I guess is a credit to network guys in general, but at the same time the two need to talk about what's going on.

There has to be concern among systems guys about whether there's enough capacity, though, right?

Yes, there should be. And there are two types of capacity with virtualization. There's outward facing, from a VM that generates traffic as a server on the network, and then there's the vMotion and inter-cluster communications within a VMware cluster. The vMotion is really taxing on a network. You take a physical host with 256 gigs of RAM and you want to copy that 256 gigs of RAM somewhere else to another host as quickly as possible—that drives quite a lot of traffic. Also, VMware has pretty specific limits on how much latency there can be between ESX hosts when you're vMotioning things. You can't have a router in between.

Virtualization architect Bob Plankers

The problem [becomes] whether the virtualization guys talked to the network guys when they were designing their stuff, or did they just plop it on the network? In a lot of cases [the environment] grew organically, so you had one or two virtualization hosts and you thought, 'Hey, this is pretty cool. I just saved a bunch of money.' So you added a third and then a fourth, but then you didn't have any room, so they're scattered all over the data center.

What does your own environment look like?

We have all Dell servers—just rack mounts, not blades. And we're also a Cisco shop for the network. I've got two VMware vSphere clusters. One has 10 machines in it and the other has eight machines in it, serving as the physical host for about 500 virtual machines.

That's a big environment. Do you have a communications problem with the networking team?

Yes, but there was a Networking Tech Field Day (a conference of networking bloggers) last August where I was the only systems guy in the room with 11 other network guys. One of the Force10 guys was going on about how systems guys get up in the morning and do all this vMotion crap and he's like, 'I don't know why they do it.' So I raised my hand and said, 'Would you like to know?' It became very clear in that moment that network guys have no idea why systems guys do things, and they're a little bitter about not being included sometimes. They're bitter about being seen as plumbers.

I realized that I needed to start talking to my network guys. As a result, we're starting a project right now where we're [changing] 1 gig connections to virtualization hosts. When you're trying to vMotion off VMs that occupy a host of 256 gigs of RAM, or 512 gigs of RAM, hosts get bigger. One of the things about virtualization is that it pays to have fewer larger hosts than more smaller hosts. But as the hosts get bigger, the vMotion process gets longer. If you're clearing it off because that host is having hardware failure, you need to go faster. So we've decided it would be better if we could put all of our gear in one rack column, in one spot in the data center. We will put in a top-of-rack switch that's 10 gig, and all the inter-cluster stuff can be limited to that switch, so you're not taxing other parts of the network. We're making changes that make networking happy and make me happy because I'm getting 10 gig connectivity. That's a testament to what happens when you work with people.

Will that mean the network team has any more control of traffic management inside the vSphere environment?

Not really. They don't manage any of the distributed switches or anything like that, but they will have access to [see] them. One of the other things that came out of [cross-team conversations] is that I've given [networking] access to see where the VM is and on which host. It turned out a few months ago we were having a problem where it became really clear that if they had access to that data they could help diagnose as opposed to watching us diagnose. They've got their own set of tools for monitoring and management and then I've got mine. They are still separate, but now I can see their router logs so it's a much more unified effort.

You gave them access to your VMware tools?
They've got access to the vCenter client, and they can look at the logs. I also showed them how to see the network configurations. They don't have permission to change them because I would like them to talk to me about it—just like I don't have permission to change things on their switches and routers.

Is there a possibility of moving to a joint, third-party management tool that shows what's available in the physical and virtual environments?
Absolutely. Xangati has some [cross-platform] tools that are network oriented, and they are able to pick up data from a variety of sources including physical switch gear, so you can see your VM end-to-end. We've been looking at it, but for us it's a budgetary thing.

Xangati is good, but in a lot of other cases, there are tool vendors saying they can monitor virtualization, but it's a limited add-on compared to what you get natively from VMware products. Then you have to ask, 'Is it better to have one tool that's OK at everything or two tools that are really good at what they do?'

What about the Nexus 1000v, which gives network engineers more control of the virtual environment?
For us it's an added cost; we don't need the functionality of it, so we haven't implemented it. In some places that was the way to appease the network guys, by basically giving them control of the virtual switches, but I guess each organization has its own style and way of dealing with this situation. For some who have tried to implement it, they might have tried talking to each other first.

Application performance is something that network guys are often responsible for. How do they address this if they can't get their hands on the virtual network?
You can't. How can you manage something you can't see? If they are in charge of managing performance, they need the tools to see what's going on or they're not in charge of managing the performance.

Who is in charge of application performance management in your situation?
With us it's a tiered thing. We've got network guys, storage guys, server or virtualization guys —me—sitting in the middle of all this stuff. Then

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

we've got sys admins who are bridging the gap between me and the app people. And we've got app people as well. If an app is having a performance problem, there are a lot of people who need to be involved.

> Any performance tool that I implement needs to be shared with everyone, so the app admin, the storage admin and the network guy all see the data.

In our particular environment, that can get kind of tricky because the app guys point the finger at virtualization when a VM is slow, and I turn around and say the VM is slow because storage is slow, and maybe storage is slow because the network is slow.

For us, any performance tool that I implement needs to be shared with everyone, so the app admin, the storage admin and the network guy all need to see the data.

Traditionally, network engineers use VLANs to segment and secure traffic. In a virtualized environment that's different. How do you address traffic segmenting and security in this environment?

We use the VLAN capabilities in the virtual switches all the time. It's either that or we have to put a ton of network interfaces in our hosts. For us, if the VLAN is enough segmentation to appease security people and enough for network guys on their uplinks and on their back trunks, it's good enough for me as well. Then I just configure my virtual switches to use the VLAN capabilities.

Networking folks don't love automation, especially without granular management. How do you address that?

For systems guys, automation is ridiculous, and for network guys, their attitude seems to be that crap rolls downhill. If a systems guy is having a problem they will blame the network, and automation makes that worse.

Automatic provisioning of VMs can be kind of scary, but certain levels of automation can go a long way toward helping us and saving us time. There needs to be oversight so it's not scary. If firewall rules are automatically changed, somebody like a security guy needs to go back and make sure it's right. Automation doesn't replace audit processes. In fact, it drives the need for more audits.

HOME

IDEA LAB

TURF WAR: SYSTEMS VS. NETWORKING AT THE VIRTUAL EDGE

NETWORK DIAGNOSTICS THAT SEE THROUGH VIRTUALIZATION

WHY ARE NETWORK ENGINEERS SO BITTER WHEN IT COMES TO VIRTUALIZATION?

Do you use the firewalls that are built into VMware or will you look at third-party security?
I am letting my network guys do the firewalling. They've got a really mature solution for firewalling any device on the network [using Cisco ASA firewalls]. I am not going to reinvent anything. As there are replacement cycles, now that we are talking with one another, we can actually have conversations about things like that going forward though. We might go toward some of the virtual firewall vShield stuff. Altor Networks makes a decent firewall. Some of that stuff is interesting because it can do firewalling at the VM level. It can say that VM 'X' absolutely can't talk to VM 'Y' even if they're sitting right next to each other on the same network segment, on the same VLAN. That's cool for shared hosting, multi-tenant environments.■

> Automatic provisioning of VMs can be kind of scary, but it can go a long way toward helping us.

RIVKA GEWIRTZ LITTLE is the senior site editor for TechTarget Networking Media.

TechTarget®
Where Serious Technology Buyers Decide

Network Evolution Ezine is produced by TechTarget Networking Media.

Rivka Gewirtz Little
Senior Site Editor
rlittle@techtarget.com

Shamus McGillicuddy
Director of News and Features
smcgillicuddy@techtarget.com

Michael Brandenburg
Technical Editor
mbrandenburg@techtarget.com

Kara Gattine
Senior Managing Editor
kgattine@techtarget.com

Linda Koury
Director of Online Design
lkoury@techtarget.com

Susan Fogarty
Editorial Director
sfogarty@techtarget.com

FOR SALES INQUIRIES, PLEASE CONTACT:

Tom Click
Senior Director of Sales
tclick@techtarget.com
617-431-9491

Cover photograph byXXX