Realtime publishers

The Essentials Series: Best Practices in Virtual Network Management

Addressing the Complexities of Virtual Network Devices

sponsored by

by Eric Beehler

Copyright Statement

Addressing the Complexities of Virtual Network Devices

Typically, the lines between IT support groups are fairly clear. The server administrator takes care of the servers, network operations monitors the network, the data center personnel cable everything, and when a trouble call comes in, the triage sends it to the right group. One group handles the issue until another, with a different set of skills and expertise, needs to get involved. It's a hand-off, one-group-responsible-at-a-time type of workflow. As with most new technologies, virtualization has been a disruptor to this system within IT centers where virtualization has been implemented. The technology that is so beneficial also shakes up the normal mode of operations.

When server virtualization was first introduced, many companies deployed it in lab environments, allowing specific groups to test and understand how the technology could be implemented in a larger production environment. Initially used in low-impact ways such as for development servers and test deployments, the technology was proven but not necessarily the processes surrounding it. VMware ESX and other vendors' virtualization products are really an operating system (OS) that facilitates running other OSs, systems administrators—sometimes specialized in the virtual machine technology—would take the lead in virtualization projects. This is no problem normally, except in this case, the virtual host contains virtual networking devices not just virtual server machines. For network administrators, the line of responsibility usually stops at the port and cable connecting the server to the network. Those network devices that are normally managed by network operations exist inside those very servers. The lines of responsibility have been blurred. Virtualization is a useful but potentially disruptive technology that requires management to rethink how the current structure of the technology teams affects the IT organization's efficiencies.

Addressing the Network Beyond the Network

When there are unmanaged network devices connected to the network, the result is usually lack of visibility and time wasted troubleshooting. One example is unauthorized wireless access points. They crept up in many companies over recent years and had to be removed because they were a security threat and made troubleshooting network problems nearly impossible without removing these unauthorized devices. The same goes for unauthorized switches under desks to expand ports. They are a convenient way for a user to expand ports without having to go through a provisioning process, but as soon as that switch malfunctions, it will take all those connected computers down, maybe even a whole segment, and tracking that problem can take time because it is an unknown variable.

The same could be said for virtual switches that are a part of every virtual host. If network operations needs to troubleshoot a problem on a virtual server, the ability stops at the port of the virtual host server, but there is still another network device in the path before you can reach the virtual server endpoint of that network.

Network administrators need to be at least aware that these vSwitches exist and what technology they bring to the table. The network administrators may already be aware that these virtual servers require different types of connections than a normal switch needs. In fact, the virtual server administrator has likely asked for 802.1q trunking ports or possibly a new load-balancing configuration for the server. The server administrators may have a requirement to route traffic for different network segments that exist within the virtual server. Network administrators needs to offer more than just the right ports for the job; they need to be able to ensure that the traffic will make it to and from these virtual servers. In addition to the fact that multiple networks and possibly several virtual switches can exist inside the virtual host, there is the issue of VMotion, which allows virtual machines to jump from one host to another. Thus, a virtual server is not limited to a single host and all of the connectivity to all the virtual hosts needs to be configured to take the traffic of any of those servers.

To have a proper picture of connectivity, monitoring tools that just look at the status of the switch port connection that connects to a virtual server is inadequate. The virtual switch is not just connection aggregation for virtual machines. It also has its own settings that can affect connectivity. For example, an incorrect setting for port groups can cause a server to be in the incorrect network segment. Standard monitoring methods of physical switches will not allow network operations to detect and help fix this problem. Since the virtual server itself isn't aware of the virtual network and the physical network isn't aware of the virtual network, what tools are available to bridge the gap and provide visibility into that virtual network? A review of the existing monitoring solution should be undertaken to see what it takes to get that visibility into the virtual system. VMware has added support such as the Cisco Discovery Protocol, which helps when troubleshooting, but a review of the current tools and processes is necessary to make sure you can continue to meet SLAs without a lot of escalations and manual effort. Several vendors including VMware sell additional products designed to manage and monitor the virtual environment. If the virtual administrator is not thinking about this kind of monitoring yet, sit down and come up with a plan to expand network monitoring into the virtual server and the vSwitches.

Automation tools are critical, especially in larger environments. Without the proper toolset, there is little the network team can do to properly manage those virtual network devices. Many of the virtualization-specific vendor tools are centered on a virtual system administrator; this can leave the network administrators without a method to access those devices. Network change and configuration management (NCCM) systems have been in a state of growth and change the past few years. You should be reviewing these kinds of solutions for completeness in the area of virtual networks and understand how you will work with virtual systems administrators to take control of those configurations without conflicting with the tools used by virtual administrators. In addition, consider any issue of a mixed environment. Even though VMware is the dominant player in standard server virtualization, products from Microsoft and Citrix, among others, may add complexity without a toolset that can address configurations across all those platforms.

When considering network management automation applications such as NCCM, the issues of virtual networks need to play with the bigger questions around service level delivery, performance management, configuration tracking, deployment planning, and measurement of services provided. These topics are not new but now extend across device types. The simplistic answers of more servers or faster switches will not be as easy to give when the servers are intertwined with the network. These tools might have seemed like something you could do without before, but you must now consider the complexities of getting successful root cause troubleshooting with just up/down monitoring on your switches and ping tests to servers. There are multiple network segments existing in virtual servers, which multiplies out when virtual machines move between hosts, making for a confusing and hard-to-troubleshoot scenario using traditional, simple monitoring tactics. Now bridge that concept into performance, security, or configuration management and the scenario of managing the network becomes even more troublesome without the right centralized solution. A package that can address the entire network, both physical and virtual, managing configurations as well as incidents becomes more of a necessity than a nice thing to have.

Determining Ownership

The old saying in every network team is “Everyone blames the network,” and a similar saying is known amongst the server administrators that “They always blame the servers.” There is some truth here, and as these two groups don’t tend to tread into each other’s waters often, the interaction between the groups often only happens when necessary. Sometimes these two groups are only interacting when something goes wrong, maintaining separate operations for the most part during other times. A symptom of this split is the finger pointing that can occur when these chasms run in an IT organization.

This kind of existence may be adequate when dealing with a network where there are physical hosts at the end of each physical switch port, but virtualization is complicating the norm. In order to provision, monitor, and track the network properly, network operations needs to understand what they are connecting to in a virtual host. Connecting to this network requires access to the virtualization service console and the associated tools. This is usually the domain of the systems administrator of that server. However, this connectivity requires technologies that are normally the domain of network administrators to be configured properly.

As there is a blur in the lines of ownership, the real question is, "Which group owns what piece of the network and at what point?" The answer could go a couple different ways. The first possible road is to continue to let the network group control whatever is network related. They can certainly understand the technologies within the virtual switch. Port groups, Etherchannel, and trunking, among other network technologies, are all familiar to any network administrator who manages physical switches. In fact, this approach could reduce the need for cross training on network technologies and thorough documentation and configuration management.

The problem with this approach is that the networking group will have to be involved heavily in provisioning new virtual servers and administrating the virtual systems. This kind of division results in loss of control for the virtual server administrator by segmenting the technologies on those servers. Essentially, this is like drawing a line across the room of two roommates, giving one side to the network administrators and the other side to the server administrator. The problem is obvious; there will still be some need to cross the line and access the other's area, so as much as the two groups may try to separate the responsibilities by technology, they still need to work together.

Some organizations have broken out the virtual administrator, or ESX admin, as a separate skill set. This is often an appropriate step. A group that addresses only the virtual server may be appropriate, especially if IT relies heavily on those virtual servers throughout the environment and have deployed the more complex toolsets related to the virtual platform. Oftentimes the responsibility will rest with the systems administrators that have training or experience with the virtual platform.

The virtual servers are often moving from the lab to production, and the deployment of production servers is often considered an internal issue to the systems administrator team because these engineers have been working on the technique of deployment and support in the lab. In this case, the consideration of other groups' lead times and all the specifics required to bring up a virtual host quickly and efficiently isn't considered fully. In the lab, all the decisions are made by the individual, but that doesn't integrate into the overall process. Leaving virtual host management to systems administrators may result in a configuration that hasn't been fully thought out. Missing VLANs, inadequate support for VMotion, and incorrect trunk settings are examples of configuration issues that could arise when the physical switch and the virtual switch are not configured to be on the same page.

A better option is to have proper configuration management, process, and integration between the groups. This is an effort on the part of both groups, and since a true, integrated configuration management database of all IT infrastructures is still more of a goal than a reality for many organizations, communication needs to come into play. The server administrators need to realize that the network settings are thought out and best practices are set by the network administrators. The network team must realize the requirements for a successful network for a virtual machine by understanding the technology available, how it integrates with the existing network, and best practices that apply to the virtual networks inside the virtual hosts. Cross-training in some areas of virtualization and networking will be necessary. In addition, the toolset may be able to be shared. When using a virtualization management tool, permissions and roles can be set so that the network administrators can view and edit only the network settings of a virtual host. Network management tools should be able to integrate virtual devices into the existing network management process and avoid issues around sharing virtual server toolsets.

Provisioning Harmony

One of the big benefits to virtualization is the ease in provisioning a new server. What used to be a labor-intensive task of purchasing, racking, cabling, staging, and configuring a new server is now just a few mouse clicks. Systems administrators don't even need to load the OS anymore; they can take a server image and change the name and IP address for a brand new server.

When looking at a VMware ESX virtual host, there is an internal network that must be built in order to connect any virtual machines, so the initial build of a virtual host is critical to the network administrator. Those devices are usually built by server administrators, and if those administrators are not familiar with the standard and methods of the network team, it can cause several issues. Connecting the virtual server to the physical network can cause immediate issues due to the fact that several segments (also known as port groups) exist inside the virtual switch, and several protocol configurations will come into play that don't normally exist with a regular host. There can also be issues down the road if the entire virtual environment is not thought out correctly—considerations such as security and performance will be obvious if, for example, the VMotion network traffic is allowed to traverse the same wire as the production network.

The requirements needed to provision virtual servers from both the network and server camps necessitate a tighter integration during provisioning. In order to avoid unnecessary delays as well as misconfiguration issues that can affect production applications and SLAs, the back-and-forth workflow and communication needs to be addressed. Normal server provisioning usually requires two things from the networking group: a cable connection to the switch and a static IP address. When spinning up a new virtual host, not only will there need to be switch port connectivity but also likely several connections needed to different switches or VLANs and several distinct networks for virtual machines. The administrator network, the VMotion network, and others such as IP-based storage (such as iSCSI) are the basic pieces that need connectivity. There can be more if additional vSwitches are created. In addition, each virtual machine needs at least one IP address and must be put in the proper network segment. Some of the ports will be trunk ports, running multiple VLANs across the connection instead of regular IP traffic. Then the correct options for items such as Etherchannel or VLAN tagging must be selected on both the physical switch and the virtual switch. Thus, a virtual server administrator cannot just pick a VLAN name out of thin air and a network administrator cannot just plug the NICs of the server into any port and light them up.

A proper process is essential to reduce the time needed to bring up a new virtual host as well as a new server. If the organization currently has a lead time of 3 days, for example, to get an IP address assigned to a new server, that kind of response will reduce the efficiency of the virtualization technology because the bottleneck is now in the process, not the ability to activate a new server quickly. Instead, work with the server team to find out all the necessary information they will need for a server and turn it into a single request that can be responded to quickly.

With a virtual host, the easy thing to do is to hook up the virtual host to the switch and run all traffic across a single NIC, or possibly set up fault tolerance between two NIC as you would a traditional server. This kind of configuration will work, but consider first the performance implications. With a Gigabit connection, the VMotion network will transmit the entire machine including contents of memory over the network at once. Then any other traffic such as IP-based storage will be traversing the same network connection. Considering that a typical VMware host may have 8 to 16 virtual machines on a single host, you can see that the network can become a bottleneck. If there are additional network interface cards (NICs) available, the IP storage, VMotion, management, and virtual machine networks should all be separated as much as possible. At the least, put the VMotion network and administration network on completely separate networks on separate physical switches.

There is another reason to do so: security. The fact that the VMotion transmits the memory contents of a server over the network means there is potential for that data to touch where it shouldn't, which is a very real concern. The administration network includes the virtual host, the remote access cards, and the ESX service console, and only administrators should have access to this network. Overall, separating networks and setting up load balancing across multiple NICs is best practice, so consider moving this direction instead of using a low common denominator for configuration of the virtual hosts.

Virtual systems administrators will surely have methods for provisioning new servers, so network managers should consider what toolsets they will use. Some smaller shops tend to use basic scripting automation, but those that need to track configurations will utilize their integrated network provisioning tools. If you are running or are considering an NCCM system, this is one area that those systems can shine. As a team, decide the standards and best practices to follow for configurations, and integrate provisioning of vSwitches and virtual ports for new servers into the overall process. Also consider the integration with your overall data center management strategy. Ask if this enters your workflow at the correct point to avoid confusion and reduce churn between groups.

The New Troubleshooting Process

When a trouble call comes in, what is the process for determining how it's managed? In many organizations, the call comes in at a first-level call center, and when the call cannot be resolved, it is escalated to the proper group. A key question during a server outage is "Does this affect more than one computer?" If the answer is yes, it moves down the path of a network issue. If not, it moves down the path of a server issue. With virtual machines, where does the ticket go? A multiple-server outage could be related to a virtualization issue or could be a network problem. When trying to meet SLAs and bring the applications and services back as quickly as possible, the correct process is crucial.

Instead of splitting the issue into server and network, the Help desk should have a configuration database available that will tell them if the application or server is virtual. Too often this information is kept completely hidden from anyone but the administrators, but there is good reason to give this information to the first-level Help desk. Consider complaints that can come in from a slowdown of an application to a full-blown outage: the Help desk can use information about virtualization to help possibly narrow the scope of an issue, especially when there are multiple issues related to a common problem. Putting these together increases the efficiency of the Help desk and routes those problems more quickly. Consider how the service desk configuration can integrate into a network management platform to create a workflow that will be much easier to use. Some products offer integration of the configuration and provide help and possibly root cause analysis when presenting an incident, greatly reducing the need for extensive discovery before escalation.

When the issue makes it to the second and third levels of escalation, proper configuration documentation is crucial. Knowing what standards are in place for all the vSwitch settings is critical, and if there are variances in those standards, the network operation group should know those as well. Being able to reference what the configuration should be is always a good way to figure out if something changed to cause a problem. If the systems and network administrators need to troubleshoot a network issue, it should be approached as a combined effort. The typical isolation of an escalated issue should be bypassed as soon as the issue can be tracked to a vSwitch to physical switch issue. Throwing the issue across the wall to the other guy will only delay resolution. In fact, allowing the network administrators access to the settings of the vSwitches can enable them to quickly double-check that settings between the vSwitch and physical switch are correct.

Of course, automation can play a big part in avoiding issues by keeping configurations consistent and following best practices. For example, automation can police the configuration of vSwitches to ensure that certain settings are kept unchanged and set to a standard. This will go far in preventing some of the most common break/fix issues you are likely to see on virtual networks.

Conclusion

Virtualization reduces the time it takes to provision new servers and reduces the resources necessary to run multiple servers and applications, but the procedures to manage virtualization can stand in the way of those efficiencies. Organizations need to determine how the network will be addressed inside the virtual machine and allow the network and server teams to gain exposure to network technology that converges in the virtual host with the vSwitch. Don't allow the lowest common denominator to rule the network configuration. Instead, plan out the configuration of the network properly to avoid performance and security issues. When troubleshooting, the model has to change from the "single group or person" troubleshooting to the network and server teams working together, relying on proper documentation of the current configuration, to quickly solve issues. Take a much closer look at network management platforms that bring virtual network devices under the umbrella of day-to-day operations and provisioning. With the theme of virtualization being simplification and consolidation, make sure competing methodologies between groups and toolsets are not standing in the way of realizing the goal. With these changes, IT can then reap the benefits of virtualization.