Realtime publishers

The Essentials Series

Best Practices in Virtual Network Management

sponsored by

by Eric Beehler

Copyright Statement

Article 1: Correctly Managing Virtual Switches

Virtualization is changing the data center. The traditional bounds of network management and server management are becoming blurred as virtual machines are deployed across enterprises to save money and increase efficiencies. The simplicity presented by virtualization—that a virtual machine is no different than a physical machine—only gets you so far. This simplistic view may apply to the operating system (OS) and applications running in a virtual machine, but the mechanics of the system require a deeper look.

Virtual servers are introducing new devices into the typical inventory of managed components. Thus, although virtualization reduces the complexity from one vantage point, it adds new layers for others. What this means for network managers is the need for an understanding of the familiar technology of switches from a new perspective. The switch is the backbone of our networks; virtualization adds a new wrinkle by putting a virtual switch inside the virtual machine. Now the server is not just an endpoint but a complex network device that must be integrated into the entire network solution. In addition, it's not exactly the same as other physical network devices on your network. Rather than ignore this addition to the network environment, a much better option is to manage the virtual switch as a part of your infrastructure.

The Virtual Switch

The key to understanding the VMware ESX virtual switch, also known as a vSwitch, is to understand that the switch is pure software. To communicate with the outside world, the virtual switch uses physical ports on a server, but even that is not necessary. The switch could just allow virtual machines to communicate with each other within the virtual host. The vSwitch itself is necessary because all networking is virtualized to the server. Thus, when looking at a virtual machine, the server's network interface card (NIC) is actually a virtualized software network card whose traffic is either sent through the virtual switch to interface with the internal network of the virtual machine host or with the outside IP infrastructure.

These virtual network components are all part of the kernel of VMware's hypervisor. In fact, any modern hypervisor from other vendors such as Sun, Microsoft, or Citrix have similar setups. So, the normal parameters of the physical switch are now variables in the virtual environment. For example, when managing a 48-port switch, you know there will only ever be 48 ports. To make it any different would require a hardware change of some sort, if it's even possible with that particular switch. In a virtual host, the number of virtual ports on a vSwitch can change. A vSwitch can have a minimum of 8 ports and a maximum of 1016 ports, with a default port count of 52. In fact, the number of vSwitches can change dynamically as well.

However, a virtual switch is a switch, so there are some consistent comparisons to a physical switch. A vSwitch maintains a MAC address table, checks each frame's MAC address for destination, and forwards those frames to avoid unnecessary traffic. There are still ports that network cards connect to; they are just virtual ports and virtual network cards.

VMware Port Groups

The virtual switch handles all traffic inside the virtual host. There are several port groups set up by default in different port groups on the vSwitch. The administration network port group is designed to connect admin-only functions such as the virtual console, remote access (typically known as iLO or DRAC cards on a physical server), iSCSI storage, and the virtual server host itself. Some sensitive data on this network must travel unencrypted, and this network provides the segmentation to accomplish that goal.

There is also the VMotion network port group. VMotion allows machines to move between different virtual hosts. This process moves the storage and memory contents of a virtual machine across the network to another computer, running hot and moving without an interruption of service. The need for security in this instance is obvious. Allowing traffic to traverse a normal network path not only is a possible drain on bandwidth but also introduces the very real potential for a security leak. Best practice is to keep this switch on a private physical network or VLAN.

The latest versions of VMware ESX include the VMKernal network, which carries NFS storage traffic. Finally, the VM network carries all the traffic for the virtual machines to the outside world by connecting to physical ports.

The ports of a virtual host server will not physically correspond to the number of virtual ports on the vSwitch. Configuring a separate physical port for each port group is an option; however, the consideration of having enough NIC ports comes into play. As you've seen, there are at least four distinct and separate networks to support. Some servers may have sufficient capacity to support dual NIC ports per port group, but many will be limited to two or four ports, especially smaller 1U or blade servers. The ports of a NIC do not need IP addresses. Instead, the NICs are placed into bridge mode and the IP addresses are assigned to the service console and the virtual machines attached to each vSwitch. You do not assign IP addresses to vSwitches.

VLANs and Trunks from the Virtual Host

The most common method to address limited port capacity is to configure the ports to support VLANs and apply the best practice of segmentation of the port groups on the vSwitch. VLANs are nothing new to switches. In fact it's a way to provide a form of virtualization inside the switch; but VLAN trunks on servers are still not a common occurrence. Conceptually, you are connecting to a VLAN trunk port on a switch, not just another NIC on a server.

Of course, the trunking protocol used should be selected with intent. Leaving the default in place could bring incompatibilities between the network and server or may not otherwise serve the needs of the network or the virtual host.

The vSwitch supports VLAN tagging and has the option of using the 802.1q standards of external switch tagging (EST), virtual guest tagging (VGT), or virtual switch tagging (VST). External switch tagging sends each network in the trunk over a pair of network cards. This limits the number of networks the trunking can support. VGT requires that the guest OS on the virtual machines handle the trunking. As this requires special drivers and configuration and puts trunking into the guest OS, it's usually avoided.

The most common and recommended method used is VST. This acts as a typical trunk, tagging all the packets and running all port groups through a single set of network ports. This allows you to configure two ports and give the host 2GB of bandwidth on a typical gigabit switch, and simplifies the configuration on the physical network switch. The configuration complexity now falls to the virtual server administrator, requiring them to configure the vSwitch and port groups properly.

Redundancy and Load Balancing to the Physical Network

As with a regular server, the virtual host should be set up with redundant connections to the physical network. Two NICs can be set up in a team in order to provide for the following options:

- Redundant NICs—The virtual host has the option to be set to use redundant NICs, providing failover support.
- Load-balanced NICs—The virtual host can be set to be load balanced inside the NIC teaming settings, spreading the traffic between the NICs and providing failover.

The vSwitch gives the ability to fail over a network connection if more than one physical NIC is connected to it. As with a physical switch, best practice is to connect two virtual NICs to a vSwitch to provide for failover and load balancing. The physical connection is set up in the same way a connection on any production server should be set up—with a separate connection to two physical switches.

Best practice dictates that 2GB ports from two physical NICs be assigned to the VM network for full performance. If there is an issue with one of the switches, the other switch will take over. In fact, if additional pairs of network ports are available, setting up a separate set of physical NIC ports for the VMotion network or IP storage network such as iSCSI will help with bandwidth and security. This also translates into your physical switch, where you can have one physical trunk connected to one switch and the other to a separate physical switch. It's also common to have a separate NIC for the service console network.

The VMotion port group should be separated from other networks to avoid a virtual machine's memory contents. In fact, the VMotion network should be separate from the production network and put onto separate physical switches with its own physical NICs as a whole to avoid bandwidth and security concerns if possible.

Load balancing is often used so that you don't waste a port on just failover but instead provide additional bandwidth. There are several types of load balancing available to the virtual host. The option to route based on the originating virtual port ID is the default and uses the rarely changing port ID on the virtual switch. Routing based on the source MAC hash (out-mac) is another option that uses the MAC address of the VM but is limited to the bandwidth of the virtual NIC and cannot take advantage of physical NIC teaming bandwidth. Routing based on IP hash (out-IP) uses source and destination IP utilizing Etherchannel or 802.3AD bonding, making it able to take advantage of additional bandwidth but also making it more CPU intensive. The default of Port ID is preferred or out-mac in previous versions of ESX. These are preferred because the other option of out-ip is a sequential method based on a fixed volume of packets. If IP load balancing will be used, Etherchannel or 802.3AD must be configured on the physical switch.

vSwitch Limitations Compared with Physical Switches

Virtual switches have several limitations over a standard physical switch. For example, they are not able to operate above layer 2. This serves to isolate networks but also hinders the ability of vSwitches on the same virtual host to talk to one another. Virtual switches cannot be connected to one another. You might just drop an Ethernet cable from one port to another on a physical switch; this is not possible in the virtual switches. This isolates any possible loops that could have occurred, but this limitation brings the issue of routing and connection of different vSwitches to the outside network devices.

Another point to note is the vSwitch does not understand or negotiate the Spanning Tree Protocol. Thus, it is best practice to disable Spanning Tree Protocol on any port connecting virtualized switches on a virtual host and just use PortFast to reduce the downtime associated with failing over to another port if failover options are set up. The forwarding table of a virtual switch is unique to that switch.

Security Features

There are policies that can be applied to virtual switches that allow for some optimization. To monitor the traffic of a physical switch, you can put a port into promiscuous mode and attach a monitor. This can be accomplished on a vSwitch by setting promiscuous mode and passing the traffic to a virtual machine, where a network monitor can be installed. MAC addresses are unique on each virtual server, but like many server network drivers, the MAC can be changed. You can set the security policy of MAC address changes to deny forcing a match of the MAC address received to the MAC address specified in the virtual machine configuration file. Forged transmits is another security option that, when set to deny, will deny outgoing IP traffic if the MAC address doesn't match the configuration file.

Traffic Shaping

The vSwitch does have the ability to limit traffic but in a very limited capacity. These are basic settings that can be applied to a vSwitch on a port group or on a per virtual machine basis in ESX versions before 3.5. Average bandwidth, peak bandwidth, and burst size can all be specified. Average bandwidth can be set between 0Kbps and 102,400Kpbs, which is a 1Gb network connection. This setting allows the burst above the set speed, but overall if you want to limit a port group to half a gigabit, you can. Peak bandwidth will set the upper limit and will never be allowed to rise above that limit. Finally, burst size will limit the size of a burst of data, including not just the data but also the packet size.

So why would someone want to traffic share? Well, the answer depends on the situation, but several examples exist where traffic sharing is used to simulate the effect of WAN links on a server environment. This is one example of possibly many reasons to utilize the feature. For most applications though, this will likely remain disabled in production scenarios.

Detecting Failures

Integrated into the vSwitch are basic features to ensure internal monitoring and management of failure. The first method relied upon is link state, which basically tells the vSwitch if the associated link is up or down. This will initiate any failover that needs to happen based on the configurations of load-balancing or failover settings.

Another option that can help determine a failure under the NIC teaming settings is called *beacon probing*. A beacon is sent down all network paths and if the beacon does not return down one of those paths, the path is then considered dead. Be careful with this setting. If the beacon is trapped by the switch and not returned, the path will be considered dead. Testing with existing network gear should take place before deploying this setting. In fact, there is little reason to use beacon monitoring if network monitoring information is available from other sources. If beacon monitoring will be used, don't connect multiple connections to the same physical switch, which would could cause the same data to be sent to the same device. Consider using it with multiple physical switches that exist in the same broadcast network.

The vSwitch does have the ability to talk the language of the Cisco Discovery Protocol. This protocol is recognized by enabled network gear to identify directly-connected CDP devices. This information may be basic, but it is one of the keys to gaining information on these virtual network devices remotely and automatically. CDP might not be turned on by default, especially if the hosts have been upgraded to ESX 3.5 from a previous version. Once enabled, it will be able to either just listen or listen and advertise. This can greatly reduce the need to translate information from the ESX administrator to the network administrator. Details such as ports, IP addresses, and VLANs will all be visible to network management and monitoring. Previous to 3.5, CDP was not available, but it can now be used to help address the network devices as a whole.

If additional management needs are required, VMware makes an additional product called the vNetwork Distributed switch. This is the framework to manage virtual network components centrally. This requires the data center-centric vSphere and allows the addition of virtual appliances on the virtual host such as the Cisco Nexus 100V virtual switch, which allows for more seamless management between physical and virtual network devices. This is not an integrated feature, so it is a significant cost to an environment to have this integration. This product allows management across VMware servers centrally.

Conclusion

Although a virtual server is really just multiple computers running on top of the same hardware, no virtual server will run without the virtual network components. The vSwitch is a necessary component to configure before any virtual machines can get on the network. It adds a layer of management for the network administrator as well as the virtual server administrator. Knowing what configuration and management options are available is the first step to ensuring network administrators understand and provide proper information to properly configure and manage the network connected to the VM. The network will appear the same when looking at the cables hooked up to any server, but the required configurations between the physical and virtual network are certainly different.

Article 2: 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.

Article 3: Integrating Daily Network Management into Virtualization

Virtualization is adding a layer of management for network administrators. The effort to keep the network running means exposure to these virtual hosts and their virtual networks. Getting to those virtual machine endpoints means going through the virtual switches that connect every virtual machine to the traditional network. It's not just an abstract method of connectivity; these switches are fully customizable, much as their physical cousins are. When systems administrators set up these virtual machines, they may not have a familiarity with the importance of certain network settings. Features such as VLANs, trunking, naming standards, and redundancy are areas where the network administrators need to interject their expertise. The systems administrators' priorities often lie with the details of the virtual servers and their important features and settings. Recognizing that vSwitches should be managed properly and made part of daily management will maintain SLAs and standards to keep the systems running smoothly.

The Importance of Port Groups

When discussing vSwitches, the VM administrator is likely to talk to you about their vSwitch's port group. Port groups allow a VM admin to take a group of virtual machines that need the same network configuration and group them together. It doesn't imply a specific number or type of ports on the switch. It is essentially another way to term a vSwitch template for certain definable ports. This can help the administrator to define certain kinds of machines and apply vSwitch network settings across those machines. From a network perspective, port groups are a useful method to group machines for a specific network segment. Defining port groups based upon the physical network segments is smart because this is where VLANs, security settings, NIC teaming, and traffic shaping are defined. Port group settings will apply over and above what has been defined as policy for a vSwitch, so this is the perfect place to put like virtual servers that exist in the same segment or need the same settings.

It would seem simple to use port groups correctly, but the concept is not considered very critical to some people and even some tool sets. This will be a problem if you plan to use features such as VLAN tagging to connect the physical and virtual networks. For example, Microsoft's System Center Virtual Machine Manager supports managing both Hyper-V and VMware virtual machines and is designed to provision new virtual hosts and virtual machines quickly with automation.

The problem is that the software only has support for virtual switches not port groups. This could lead to a design where each segment or even every machine gets its own switch. This design is not hard to imagine, as those vSwitches can seem to have a small impact on the virtual machines, but sending unnecessary traffic across switches can cause headaches for the network. Thus, this design is certainly not the most flexible when considering that vSwitches only operate at layer 2 and cannot communicate. Any traffic that needs to traverse the internal virtual network needs to be handled outside the vSwitch in the physical network because there is no vSwitch to vSwitch connectivity available within the virtual host.

There is also the possibility that port groups are over-engineered. As most network administrators know, over-engineering a network adds overhead and makes troubleshooting difficult to manage. Take, for example, a set of application servers that exists on the same subnet on the physical network and need to communicate with the infrastructure servers on that same segment. If someone decided to put each type of application on its own port group, this setup would require separate configurations for each group when the only difference between those servers is the application running on them—not any network- or feature-setting difference. This kind of over-engineering is easy when someone isn't thinking from a network perspective when engineering the virtual switch. Standards here are helpful and avoid the over-engineering problem. Port groups should include virtual machines that require like network settings and VLAN settings.

Dealing with VMotion

The VMotion service is a key feature of VMware and one that helps avoid downtime by allowing machines to move to another physical host without being offline. This kind of technology is truly impressive but can be network intensive for the requirements of bandwidth and the security settings. As the full contents of memory for a machine need to be quickly copied from one host to another over the network, an appropriate network design is needed. It's recommended that this network connection have at least a 1Gb speed connection and exist on a separate physical switch than the rest of the virtual machine network. The reason is security and performance. Many modern attacks go straight for memory registers. In fact, a recent demonstration that unlocked the keys to disk-based encryption relied on access to the memory of a machine because the keys to the encryption, once loaded into memory, are available to unlock the data afterwards if someone can get access to the content of RAM. VMotion sends those in-memory bits of a server across the wire to another computer, which is likely to include sensitive data. It is this kind of exposure that makes the possibility of sniffing that data on the network so dangerous.

The VMotion network not only needs to connect the virtual host to a private network but also any machine that is part of the VMotion group. A single, separate NIC should be used for this purpose alone, and if high availability is required, the recommended configuration is to set up failover for two NICs on the VMotion network. The VMotion network also requires its own IP subnet that is not routed to the rest of the network. Using separate physical switches is preferred, but VLAN tagging is the must-have option if physical switches or sufficient network adapter connections aren't available. Normally, the configuration of the VM host will include a separate vSwitch for the connection to the VMotion network.

Extend Naming Standards to the Virtual Switch

Many of the standards set by network administrators are transparent to other IT groups in an organization. Switches and routers are usually named in a logical fashion, giving their names purpose. This also applied to the VLAN numbering scheme. When dealing with virtual switches, even though they are only software, they will respond as any other network device when called by the Cisco Discovery Protocol (CDP). Giving the name purpose can help identify a directly-connected vSwitch. This method of discovery is only useful if the group that provisions the vSwitches follows the proper protocol for naming.

When discussing VLANs, naming isn't just an option, it's a requirement. Now that vSwitches participate actively in trunking networks to the physical LAN, a proper set of VLANs should be reserved for the virtual switches. This requirement doesn't apply only to local switches because the VMotion feature will have machines moving from host to host, bringing their VLANs with them to completely different hosts.

With the most common method of VLAN tagging, virtual switch tagging (VST), the configuration of VLAN ID to the proper port group is required. Of course, the port on the physical switch will have to be configured for VLAN trunking as well. Many server administrators will consider a port to be a port, so definition of this requirement should be up front during provisioning and the proper information on VLANs and proper port connection made clear. As in the physical world, layer 2 devices such as a vSwitch require a router in order to allow the different networks to communicate. If network routing is required, you will have to provide that functionality. It will not happen within the vSwitch. It is also not possible to trunk between vSwitches. Consider the vSwitch very rudimentary in this regard. In order to get the vSwitch to do the tricks of modern switches, a connection to the physical network is often required.

From the administrator setting up the virtual machine to the network administrator managing the ports to Tier-I support personnel required to respond to daily issues, everyone is required to understand what pieces are where in your network. This means everyone needs to come together to understand the naming standards. When a systems administrator and a network administrator talk, they need to be able to discuss the same thing in the same way. The days of getting a systems administrator to just provide an IP address are done. You'll need much more from an administrator of a virtual host.

The default naming pattern of a vSwitch is vSwitch 1, vSwitch2, and so on. The names of these switches may not be important in the short term; future management and monitoring may make it very important for them to have unique names. Take the current naming standard of your physical switches and extend that into virtual switches. You might want to take the name of the virtual host into consideration to easily ID the location of a vSwitch.

The port groups come named for their functions; for example, the default port group for virtual machines is, in fact, Virtual Machines, and the name of the default for the service console is Service Console. This is certainly functional for the virtual systems administrator but may make little sense when discussing specific networks or VLANs. One possible direction is to name the port groups according to their specific function and include key networking information. For example, name the port group that includes the Microsoft Exchange messaging servers using a VLAN of 2050 on the port group as Exchange_Prodnet_VLAN2050. This definition makes it very obvious to a support person what function it provides as well as key VLAN information. Thus, the default Virtual Machines port group would actually be broken out to separate port groups based on this scheme. You would also rename any other port group, such as the service console or VMotion port groups, to include a VLAN ID. In the process, avoid spaces to prevent possible scripting issues down the road for tools that may do automatic provisioning and maintenance.

Addressing Faults, Performance Issues, and Troubleshooting

Virtualization complicates not just network provisioning but also incident management. First, for most tools, the virtual switch is going to be a phantom; unseen by most monitoring tools that will only look to the endpoint as the NIC on the server as the final destination of the network. Thus, troubleshooting a virtual host issue will require considerations not normally associated with testing any other endpoint. Check the vSwitch for the correct VLAN IDs assigned. Check for trunks enabled on the right switch port as well as having trunking enabled on the virtual host. You can also ensure Etherchannel and teaming settings are properly configured where applicable. Remember that duplicate MAC addresses can exist because they can be changed and duplicated when copying a virtual machine, so check layer 2 for MAC inconsistencies. Also, you can use the CDP to find configuration information on the port (this feature is available on ESX 3.5 but not on earlier versions or on ESXi).

The vSwitch has the ability to detect some of its own faults and failover if it is configured with multiple paths. In a way, this is much like failover NICs, but instead of the failover integrated into a driver and application on the server, it is now integrated into the vSwitch. The virtual machine only needs one connection because the vSwitch will handle any failover scenario when configured to do so.

Beacon monitoring is a method that sends broadcasts down every VLAN to check for the state of connectivity down all paths beyond the immediate port connected. This functionality can be problematic if multiple ports from the virtual host are connected to the same switch—beacon monitoring can send the same MAC address, causing classic duplication issues on the switch.

The link state monitoring is also beyond the standard up/down monitoring often done by teamed NICs. If the physical switches support Link State Tracking, the vSwitch will be able to look upstream for path failures and switch over to the other path. You can also set the Notify Switches option to Yes in order to send out notification to connected switches that they need to update lookup tables whenever a failover event occurs or a machine has been VMotioned to a new host.

Beyond a virtual switch's ability to deal with failures on its own, you need to be able to respond properly to a virtual switch incident. Gaining access to the service console gives you the command line access into the virtual machine. Gaining access to the service console can be accomplished by using SSH, remote console, or the physical console. There are a slew of useful commands you can issue. For example, to make sure the virtual host is broadcasting and listening to CDP, type the command

```
esxcfg-vswitch –B both vSwitch1
```

Much of this discussion is based on the fact that vSwitches are quite a bit different than the standard Cisco switch, but the VMware hypervisor allows for virtual appliances. This ability to reach into VMware with an API allowed Cisco to release a virtual switch. This Cisco software switch allows for the traditional Cisco management methods to be applied to the virtual switch. Possibly the most exciting thing is you will now get a traditional Cisco IOS command-line interface to the switch, making network management quite a bit easier. The policy integration into vCenter allows you to define policy such as VLANs across many virtual hosts and port groups. When a VMotion occurs, this policy moves with the virtual machine, including the same virtual port on the virtual switch of the new virtual host.

So how do you handle the full life cycle of a network event in your organization? That is largely going to depend on how well the virtual systems administrators and network administrators work together and share information. The ability to access the virtual switch is also a key factor, but a good working relationship is crucial, especially if you need to rely on the systems administrator to feed information back during an incident.

The virtual machine concept bases many of its strengths in the ability to automate. Virtual NICs, vSwitches, and all the details surrounding them can be scripted for full automation. Some may choose to use a central automation tool, but simple scripting can work pretty well for many provisioning tasks that are not large scale. A good set of network-creation scripts may be just what administrators need when provisioning, giving them a simple set of commands that have been vetted and tested by the network team to configure all the right settings. For those that need centralized configuration and monitoring capabilities, look to your network change and configuration management systems. These can help align your configurations and processes between groups and give you proper control of standards in one place. They can also help you understand performance, risk, and security compliance of the virtual network.

The challenge that faces network administrators is better automation across platforms. Each vendor has its own solution, including the virtualization vendors, but consider the need to manage network devices across brands, types, and logical segments of the data center; we are now integrating devices that fall into the grey area dividing servers and networks. The tools promising integration now have to step up and fulfill that promise further. It is important that your network management system has the ability to understand how your physical switches and virtual switches integrate. That single management point will be crucial for any sizable deployment of virtual network infrastructure. Without it, troubleshooting and configuration management become a real chore that is added to the need to deal with unnecessary problems that arise without the process and management that these tools provide.

Best Practices for Virtual Server vSwitches and vNetworks

The following list provides a summary of the best practices for managing virtual server switches:

- Ensure redundant physical switches for each critical network on a vSwitch. Doing so provides for the necessary bandwidth using Gigabit ports and allows for redundancy in failover.
- Use at least four physical NICs on a virtual host. One for the service console, one for VMotion, and another two used for the virtual machines. Six physical NICs allow for redundancy for all these networks.
- For additional vSwitches in use, provide dual NICs and separate physical switch port connections.
- The service console requires an IP address, network mask, gateway, DNS servers, and a redundancy method. Having these pieces at the ready for a new host will reduce procurement time.
- Trunking standards need to be set for the vSwitches' connections to the physical ports. Provide the proper 802.1q trunking settings for the virtual switches.
- Assign and document IP addresses for all virtual machines. This information should include proper network mask, gateway, and DNS settings.

- Set naming standards for the virtual switches and assign names to those switches if appropriate to your procedures.
- Note whether PortFast is enabled and spanning tree is disabled on the physical switch. These settings are best practice, but status of those port settings also aid in troubleshooting.
- Connect the VMotion network NIC(s) to a separate physical switch or separate network that is not routable to the production network for security and maximum performance.
- Integrate network toolsets such as network change and configuration management systems when possible to enforce standards and monitor virtual network devices.

Conclusion

The best approach to managing the virtual networks inside virtual hosts will depend on each organization. They key is to understand that ignoring vSwitches will only lead to long nights of troubleshooting without any documentation or standard to fall back on. Your application owners will not be forgiving when they find out a VMotioned server went offline due to missing VLAN configurations. Now is the time to capture the configurations, set standards, and make systems administrators aware of the configurations you need and the harmony between the physical networks and virtual switches to make everything operate smoothly. Virtualization doesn't work without the network and network operations that address issues with best practices. With the active participation of the network administrators, best practices will come to the vSwitch.