Realtime publishers

The Essentials Series: Best Practices in Virtual Network Management

Integrating Daily Network Management into Virtualization

sponsored by

by Eric Beehler

Copyright Statement

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.