

Better Network Security under VMware

By: Paul Imhof

A few months ago I had a friend who is an IT Security Officer at a large institution ask, “Paul, should I allow VMware servers in my DMZ (demilitarized zone)?” After obtaining details on their security implications, I found a solution that would address his concerns and allow his organization to realize server consolidation in their public facing portal.

Risks with virtualizing servers

Server virtualization is not new to organizations. The benefits of cost savings, while enabling disaster recovery, are well understood in most data centers. With a new breed of performance in Intel’s Nehalem chip set, organizations are discovering even higher server consolidation is possible. Obviously this leads to even more cost savings, but with this virtualization proliferation comes certain network vulnerabilities.

By deploying a virtual environment, you are essentially removing the link between hardware and software and in doing so, you blur the lines when it comes to securing your infrastructure. You also make it easier for security hacks to gain undetected network access (especially if you put your virtual machines in your DMZ). Virtualization streamlines the process of provisioning and patching servers, but it also adds complications that IT professionals may not be thinking about.

The explosive growth of virtual servers, set on by reducing IT budgets, has given way to adhering to IT and corporate policies. This is evident in the accepted practice of having Server Administrators (Server Admins) manage virtual switches in VMware. Convenience has trumped security! Would you allow your Server Admin access to your edge switches? This is effectively what you are doing when your Server Admin is managing your VMware virtual switches. I’m sure this is something your IT Security officer would not approve.

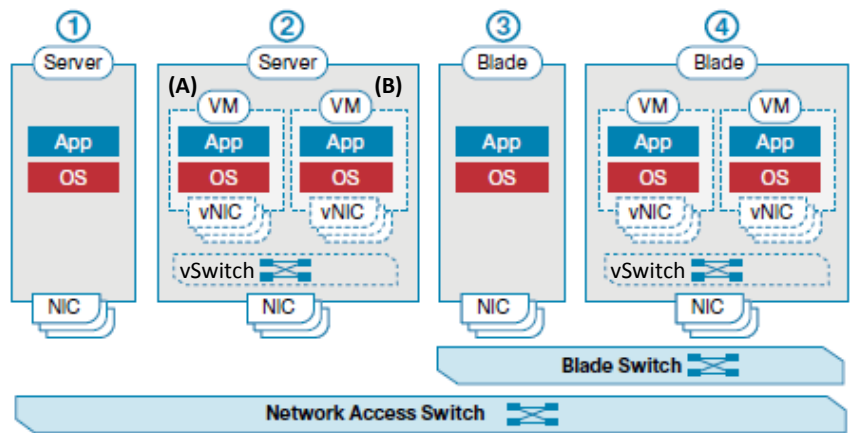
Another area of concern is what happens when a virtual machine (VM) is moved from one physical server to another. Figure 1 shows a comparison of the access layer connections. VM mobility (vMotion) can break when moving between server (2) and Blade chassis (4). When a VM move involves moving from one physical server to another, all the policies in the network for this machine (ACLs, VLANs, etc) should be moved but this can be a problem depending on the vSwitch (virtual software switch) and network access switch configurations. On Server (2), communication between VM (A) and VM (B) on the hypervisor vSwitch is invisible to the Network Administrator (Network Admin). That is a good thing for performance but a bad thing for security.



About the author:

Paul Imhof is a Data Center Program Manager at Evolve Technology Group in Rocklin, CA. He has 20+ years combined experience working for Evolve Technology, Computer Associates and Dell, specializing in systems management, server, storage and virtualization practices. He is a Cisco Certified Design Associate (CCDA) and has an ITIL Foundation v3 certification.

Figure 1: Comparison of Access Layer Connectivity Options in (1) Nonvirtualized Rack-Optimized Server, (2) Virtualized Rack-Optimized Server, (3) Nonvirtualized Blade Server, and (4) Virtualized Blade Server



There are basically three network security issues related to VMs:

Issue1: The network policy must follow a vMotion move

It would be ideal to have a security policy that is attached to the VM as it moves. Unfortunately, today's tools only allow for network policy to be attached to the physical server. In fact, VMware has a tool called DRS, or Dynamic Resource Scheduler, that automatically migrates the VM depending on CPU, memory loads and/or other factors. This does little for the network policies. What Network Admins really need is mobile security policy attached to the VMs.

Issue2: Impossible to view or apply network policy to locally switched traffic

The second issue with server virtualization concerns the vSwitch, located inside the hypervisor, that switches packets blindly (without the ability to view packets) between VMs. It is fairly difficult to see which VM is actually talking to other VMs inside the server. Troubleshooting and debugging capabilities inside the ESX¹ server are needed especially when you want to view VM to VM network conversations.

Issue3: Collaboration needed between Network and Server Admin

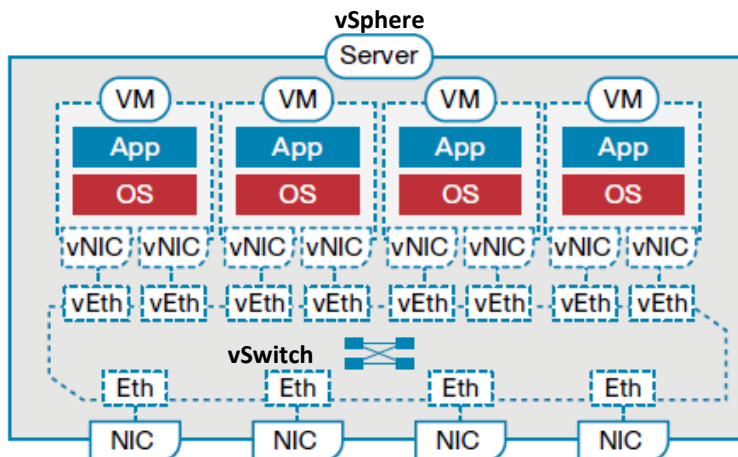
Who exactly owns the management of the vSwitch? Typically Server Admins manage the virtual switch, and they need constant communication with their Network Admin to configure the vSwitch. While Server Admins want their network team to configure the virtual network, Network Admins are demanding network tools to configure the vSwitch. These Network Admins also need network control and visibility down to the VM. From a corporate (policy and security) standpoint, these vSwitches should be owned and managed by the network group.

¹ ESX stands for **Elastic Sky X** - <http://vmfaq.com/entry/32/> and <http://communities.vmware.com/thread/20538>

Virtual Switch Concepts & Definitions

There are several concepts that most Network Admins know, but in the realm of server virtualization, some need a little explaining and further defining before I lay out all the security benefits.

Figure 2: Relationship between Virtual and Physical network constructs with a Nexus 1000v distributed virtual (software) switch



Virtual Switch (vSwitch)

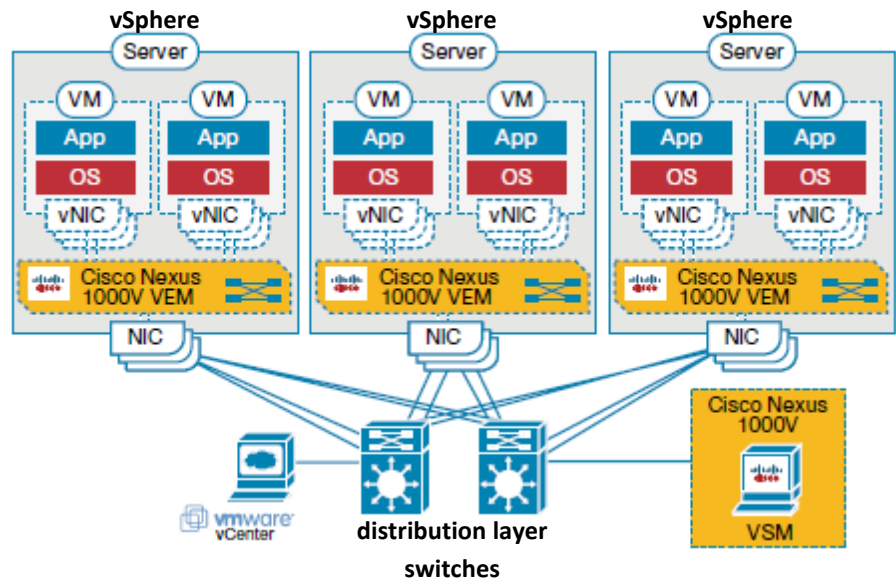
Each vSwitch represents an independent point of configuration and the Server Admin now must maintain and secure a portion of the network without the use of the best practices, diagnostic tools, and management and monitoring available throughout the rest of the infrastructure. All VMware servers (ESX v3 and v4) have this embedded in their kernels.

Distributed Virtual Switch

New to vSphere 4 (ESX v4) is the distributed vSwitch – called vNetwork Distributed Switch (vDS). VMware and Cisco jointly developed the vDS in vSphere 4. The vDS decouples the control and data planes of the embedded switch and allows multiple, independent vSwitches (data planes) to be managed by a centralized management system (control plane). Think of the vDS like you would think of a director class chassis (much like Cisco's 6509 chassis) where this chassis extends to all of your ESX servers in your enterprise (to a point) and each line card is basically the vSwitch contained in each of the ESX Servers. The vDS provides for global - network provisioning, monitoring and management by way of a single distributed switch, represented in multiple VMware vSphere servers.

Cisco's Nexus 1000v is also a software distributed vSwitch for vSphere 4 but it actually extends the capabilities of the vDS further. In fact, all vSphere servers have Cisco's Nexus 1000v code embedded in the kernel – just not turned on. It's enabled with a software key. The differences between the VMware and Cisco vSwitch implementations are available at <https://www.myciscocommunity.com/docs/DOC-8666>.

Figure 3: Relationship between Virtual and Physical network environments with a Nexus 1000v distributed virtual (software) switch as it relates to multiple ESX (vSphere) servers.



VnLink

Cisco uses the vDS framework to offer a feature set and operational model that is consistent with other Cisco networking products. These features are collectively referred to as Cisco Virtual Network Link (VN-Link). Essentially, this allows the network monitoring and control to extend from the core right down to the vSwitch.

vEth (Virtual Ethernet Interfaces)

Virtual Ethernet Interfaces (vEth), like physical NICs, are dynamically provisioned based on network policies stored in the switch as a result of VM provisioning operations at the hypervisor management layer.

VEM (Virtual Ethernet module) & VSM (Virtual supervisor module)

Figure 3 shows the Cisco Nexus 1000v **VEM** in each vSphere server and the Cisco Nexus 1000v **VSM** as a separate server connected to one of the distribution layer switches. The VEM replaces the vSwitch in each of the vSphere servers. The VSM configures and manages the VEMs and has a tight integration with VMware vCenter.

NOTE: The VSM can be a VM or (as in Figure 3) a physical appliance.

Port Profiles

Port profiles are a collection of interface configuration commands for physical OR virtual interface cards. A port profile can have attributes like VLAN ID, PVLAN, ACL and port security.

Bringing it all together

As you probably have discovered, the Cisco Nexus 1000v is at the core of the security and network management paradigm. Installing it on your vSphere 4 servers allows you to gain better control, improve system management and safely secure your network infrastructure from your firewall and core switch right down to your virtual servers.

So, how is this done? Let's start with the nuts and bolts of the Nexus 1000v. As I've mentioned previously the Cisco Nexus 1000V consists of two components – the virtual supervisor module (VSM) and the virtual Ethernet module (VEM). The networking and policy configurations are performed on the VSM and applied to the ports on each VEM. The VEM provides the ports for VM connectivity while the VSM is the clearing house for the overall multi-server distributed switch configuration. Fortunately the VSM doesn't have to be configured in a redundant fashion because once the port profiles and policies are set, they are delivered to each of the vSphere v4 servers running the Nexus 1000v software. If the VSM is shut down (or goes off line), all existing configurations and policies are enforced at each vSphere v4 server. The VSM is needed for network updates and when other vSphere servers are added.

Requirement for the Nexus 1000v: It only works with vSphere 4 servers. Older ESX v3 servers must be upgraded to vSphere 4 servers for the distributed vSwitch features.

Security features available with Cisco Nexus 1000v

Port profiles are used to map features and policies to specific virtual ports. Port profiles move seamlessly as virtual servers are moved from one physical Vsphere server to another. So how does this technology work and provide a higher level of network security in VMware? Basically you are getting all the features of a Cisco core switch right in your virtual infrastructure. This allows your virtual servers to have all the security features of a physically connected server. Conversations between virtual servers that couldn't be monitored inside of ESX, can now be monitored with the Nexus 1000v. QoS can now be applied to an individual VM and all these settings can migrate as VMs move from server to server.

Here's a small list of what the Cisco Nexus 1000v provides:

- VLAN and Private VLANs
- Port mirroring (Switched Port Analyzer [SPAN] and Encapsulated Remote SPAN [ERSPAN])
- Access control lists (ACLs)
- Anti-spoofing
- Quality of Service (QoS)
- NetFlow Version 9
- Port management, Port-channel, Port Security

An exhaustive comparison of ESX v3, vSphere 4 distributed switch and vSphere 4 with Cisco Nexus 1000v can be found at <https://www.myciscocommunity.com/docs/DOC-8666>

Virtual Infrastructure – suggested duties

Have your Server Admin continue to do what they do best but offload their vSwitch configuration duties to the Network Admins. VM workflow doesn't change – Server Admin (in charge of VMs) continues to leverage vCenter for VM creation, maintenance and monitoring. The change should involve the Network manager assuming the responsibility for vSwitch configuration and management (Port Profiles). Offloading this to the Network Admins ensures consistency with the physical network infrastructure. The Server Admin only subscribes VMs to available Port Groups.

By giving the Network Admins more control of the virtual networking you'll:

- Unify network management and operations
- Improve operational security
- Enhance VM network features
- Ensure policy persistence
- Enable VM-level visibility

What to do Next:

- Migrate your ESX v3 servers to vSphere 4 (if you haven't already done so)
- Install Cisco's Nexus 1000v on all of your vSphere 4 servers
- Have your vSwitches managed and monitored by your security and network group
 - Offloading this task from your server and ESX administrators.
- Audit your network often (use an outside specialist).
- Review your network and security policies – changes might be needed with virtual servers

Bottom line – Tell your IT Security Officer, like I've done with my friend, that you can put your VMware servers securely and confidently in your DMZ as long as you deploy Cisco's Nexus 1000v.

If you have any questions or would like further information on this and other data center related topics please feel free to contact me directly – pimhof@go-evolve.com

More Info

- Virtualizing DMZ with Nexus 1000V
<https://www.myciscocommunity.com/docs/DOC-8665>
- Feature Comparison: vSwitch vs. Nexus 1000V
<https://www.myciscocommunity.com/docs/DOC-8666>
- Cisco VN-Link Virtual Machine – Aware Networking (At-a-Glance)
http://www.cisco.com/en/US/products/ps9902/product_at_a_glance_list.html
- Cisco Nexus 1000v (Cisco site)
<http://www.cisco.com/go/nexus1000v>
- Whiteboard Tutorial
www.thepartnergrid.com/wb/nam
- I Love Nexus 1000V Because...
www.VMwaregrid.com/iheartnexus/xyz
- Introductory Presentation of the Nexus 1000V
<https://www.myciscocommunity.com/docs/DOC-8454>
- ROI Presentation for Nexus 1000V
<https://www.myciscocommunity.com/docs/DOC-11124>
- 60-Day Free Evaluation and Download
www.cisco.com/go/1000veval
- Nexus 1000V Test Plan
<https://www.myciscocommunity.com/docs/DOC-12166>