Cisco VN-Link Virtual Machine–Aware Networking

Effects of Virtualization

Server virtualization invalidates a common assumption of server access network design: that each network access port corresponds to a single physical server running a single image. Server virtualization also invalidates a second assumption: the static nature of the relationship between an image and the network. Server virtualization effectively enables OS images to become mobile. The consequences of this new level of mobility on the network are not trivial, and their effects may go beyond just the access layer, as, for example, some of the services deployed in the aggregation layer may need to be modified to support virtual machine mobility. Figure 1 provides a visual comparison of the access layer connectivity options.

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

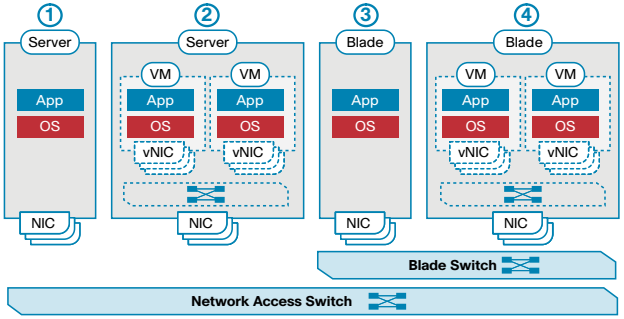

Virtual machine mobility also breaks several other features that have been implemented in the network under the assumption that computing is relatively static, and moving a physical server in the data center is not practical to do very often. Further, as virtual machines move from one physical server to another, all the network policies defined in the network for the virtual machine (for example, access control lists [ACLs]) should be consistently applied, no matter what the location of the virtual machine in the network.

Hypervisor-Embedded Virtual Switch

The easiest way to network virtual machines is to implement a software switch as part of the hypervisor. This is what VMware did with the virtual switch (vSwitch). Each virtual network interface card (vNIC) logically connects a virtual machine to the vSwitch and allows the virtual machine to send and receive traffic through that interface. If two vNICs attached to the same vSwitch need to communicate with each other, the vSwitch performs the Layer 2 switching function directly, without any need to send traffic to the physical network.

The primary benefit of the embedded vSwitch approach is its simplicity: each hypervisor includes one or more independent instances of the vSwitch. Unfortunately, this strength becomes a weakness when several VMware ESX servers are deployed in the data center. Each embedded vSwitch represents an independent point of configuration and the server administrator 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.

Furthermore, the administrator must manually make sure that the migration of the virtual machine can take place without breaking network policies or basic connectivity.

To overcome the limitations of the embedded vSwitch, VMware and Cisco jointly developed the concept of a distributed virtual switch (DVS), which 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.) VMware has branded its own implementation of DVS as the vNetwork Distributed Switch (vDS), and the control plane component is implemented within VMware vCenter. This approach effectively allows virtual machine administrators to move away from host-level network configuration and manage network connectivity at the VMware ESX cluster level.

Cisco VN-Link

Cisco is using the DVS framework to deliver a portfolio of networking solutions that can operate directly within the distributed hypervisor layer and offer a feature set and operational model that are familiar and are consistent with other Cisco® networking products. These features are collectively referred to as Cisco Virtual Network Link (VN-Link).

Figure 2: Relationship Between Virtual and Physical Network Constructs in a VN-Link Enabled Switch (Cisco Nexus™ 1000V Series Switches)

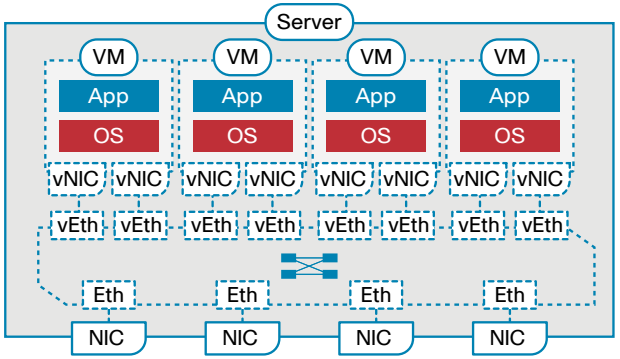

Virtual Ethernet Interfaces

Virtual Ethernet interfaces are the virtual equivalent of physical network access ports. These virtual interfaces are dynamically provisioned based on network policies stored in the switch as the result of virtual machine provisioning operations at the hypervisor management layer.

Port Profiles

Port profiles are a collection of interface configuration commands that can be dynamically applied at either

C45-532467-00 04/09

physical or virtual interfaces. A port profile can define a sophisticated collection of attributes such as VLAN ID, private VLAN (PVLAN), access control list (ACL), and port security. Port profiles are tightly integrated with the management layer for the virtual machines and allow virtual machine administrators to simply choose among a menu of profiles as they create virtual machines. When a virtual machine is powered on or off, its corresponding profiles are used to dynamically configure the vEth in the VN-Link switch.

Deploying VN-Link in Existing Networks with Cisco Nexus 1000V Series Switches

The Cisco Nexus™ 1000V Series distributed virtual switch not only maintains the virtualization administrator's regular workflow; it also offloads the vSwitch and port group configuration to the network administrator. These features reduce the virtualization administrator's workload while also reducing network configuration mistakes.

The Cisco Nexus 1000V Series consists of two components that can virtually emulate a modular Ethernet switch with redundant supervisor functions:

- **Virtual Ethernet module (VEM) (data plane):** This software replaces the vSwitch in each hypervisor. It performs switching between directly attached virtual machines, and provides uplink capabilities to the rest of the network.
- **Virtual supervisor module (VSM) (control plane):** This standalone physical or virtual appliance is responsible for the configuration and management, of the overall Cisco Nexus 1000V Series as well as the integration with VMware vCenter.

Figure 3: Cisco Nexus 1000V Series Distributed Switching Architecture

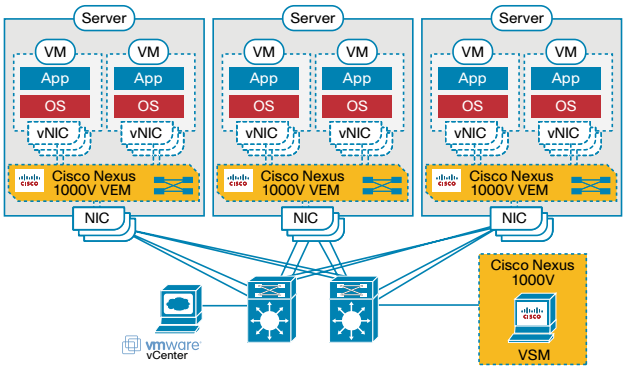

Deploying VN-Link with Network Interface Virtualization

In addition to the software-based distributed virtual switch model, Cisco has developed a hardware approach based on the concept of network interface virtualization (NIV). NIV completely removes any switching function from the hypervisor and locates it in a hardware network switch. NIV still requires a component on the host, called the interface virtualizer. The purpose of the interface virtualizer is twofold:

- For traffic going from the server to the network, the interface virtualizer explicitly tags each of the packets with a unique tag, known as a virtual network tag (VNTag).
- For traffic received from the network, the interface virtualizer removes the VNTag and directs the packet to the specified vNIC.

Switching is always performed by the network switch to which the interface virtualizer connects, which in this case is called the virtual interface switch (VIS) to indicate its capability not only to switch between physical ports, but also between virtual interfaces (VIFs) corresponding to vNICs that are remote from the switch. Said in a different way, each vNIC in a virtual machine corresponds to a VIF in the VIS, and any switching or policy enforcement function is performed within the VIS and not in the hypervisor. (Figure 4).

Figure 4: Architectural Elements of the NIV Model

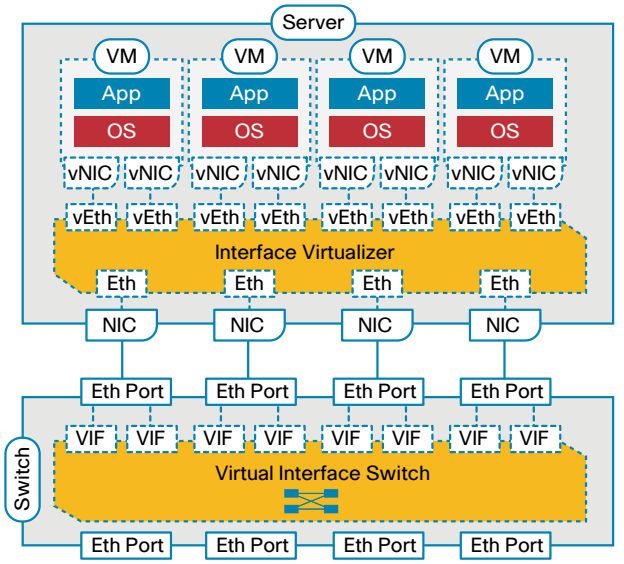

Conclusion

Cisco VN-Link provides an immediate solution to virtual machine networking requirements, while laying the foundation for future enhanced and simplified connectivity options in virtualized data centers.

For More Information

For more information, visit http://www.cisco.com/go/vn-link and http://www.cisco.com/go/nexus1000v.

C45-532467-00 04/09