Integrating Networks into Virtual Environments with Virtual Server Networking

September 2011

Contents

Figures

Automating the Network

As organizations move toward the virtualized datacenter, many are contemplating the changes that must take place at the network level. Server and storage virtualization technology is well in place, but those technologies, and the ways that they are used are driving a profound need for network automation. In fact, network automation is thought by many to be essential to realizing the full benefit of virtualization technology, and vital for producing a tangible return on investment from that technology.

At the heart of this transition is a fundamental change in the ways that computational power is obtained, deployed, and configured. Physical systems are relatively static, being selected, purchased, configured, and deployed at a pace that allows some degree of manual intervention. Likewise, physical systems are connected to physical networks in a fairly static and manual fashion, with network and system administration remaining essentially separate disciplines. With virtualization, those lines are starting to blur.

Virtualized environments are built to be completely dynamic, demanding robust network automation – elements of which are only now starting to emerge. While very small virtualized environments may lend themselves to manual configuration, they are now the exception. To connect virtual machines (VMs), virtual software switches now run on the systems that host virtual machines – augmenting physical Layer-2 switching. Virtual switches are under control of hypervisor management tools, and network administrators typically have low visibility into their configuration. Those same management tools now allow administrators to create, move, destroy or redeploy hundreds or thousands of VMs on demand. For those actions to be meaningful to virtualized applications, however, corresponding actions must be taken on networking switching infrastructure.

To realize the potential of virtualization, proper coordination must take place between virtual network components and connected physical network devices. Specifically, connected physical devices need to be able to anticipate the dynamics of virtual server deployment in a timely fashion. Physical network infrastructure must become aware of pending actions on the part of hypervisors or management tools, so that the appropriate conditions can be set in advance of changes to virtualized environments. This coordination is essential in order for virtualized applications to continue to operate in an uninterrupted fashion – even if application components are moved from one system to another. For example, virtual local area networks (VLANs) must be configured appropriately on the switch ports that connect systems where VMs are being moved or established – maintaining appropriate logical grouping, security, and access for dependent applications.

Dell Force10 Network's Virtual Server Networking technology – a part of Dell Force10's Open Network Automation Framework – is designed to provide this essential coordination between the hypervisor and the physical network fabric. Virtual Server Networking accomplishes this functionality without requiring a host-based agent on the servers where hypervisor software operates. As of this writing, Virtual Server Networking supports:

- VMware 4.0 and 4.1
- Citrix XenServer v5.6

Dell Force10's Open Automation Framework

Dell Force10's Open Automation Framework is made possible through the ubiquitous Dell Force10 switch operating system – FTOS – that runs across all Dell Force10 switches and routers. By delivering the same operating system across its entire switch and router line, Dell Force10 ensures that organizations benefit from stable code, a consistent feature set, and simpler software management. FTOS also gives Dell Force10's switches an extensible and autonomous operational model that is essentially more like a server than a traditional switch. FTOS adds server-style intelligence and general programmability to Dell Force10's switches and routers, greatly extending their capabilities. The NetBSD kernel at the heart of FTOS provides a stable operating system, handling memory allocation and process scheduling, while all other applications run as independent and modular processes in their own protected memory space.

Running on FTOS, Dell Force10's Open Automation Framework is designed to transform the data center network into a policy-driven cloud computing fabric, and to provide data center network managers with greater visibility into how the network is performing.

Open Automation incorporates the following elements:

- Bare Metal Provisioning – reduces installation time, eliminates configuration errors and enforces standard configurations by automatically configuring network switches.
- Smart Scripting – improves network monitoring and management with a robust, Perl/Python scripting environment.
- Virtual Server Networking – increases network flexibility by automatically provisioning VLANs when VMs are migrated
- Programmatic Management – simplifies network management by integrating with multiple third party system management tools.
- Web GUI Interface – extends web connectivity to Dell Force10 platforms via an extensible web GUI enabling custom GUI solutions while simplifying management.

Virtual Server Networking Architecture

As a part of Dell Force10's Open Automation framework, Virtual Server Networking provides real-time communication between the Dell Force10 network fabric and virtualized server technologies to automate network management and configuration tasks throughout the data center infrastructure. Through a switch-resident agent, Virtual Server Networking provides a closed-loop provisioning system to enable, for example, the automatic reprovisioning of VLANs across interconnected Dell Force10 switches simultaneously, thereby increasing employee productivity, and minimizing human error. Because Open Automation supports hypervisors from multiple vendors, IT managers can use a single mechanism to simultaneously support multiple hypervisors and their management tools.

Figure 1 provides an overview of an architecture where an FTOS based Dell Force10 switch is connected to multiple servers, each hosting either VMware ESX or Citrix Xen hypervisors. Depending on implementation, these hyper-visors may or may not have a centralized hypervisor management component. For instance, when the number of servers with the same type of hypervisor is small – less than five – a hypervisor management system may not be present. When there is a large number of servers with the same type of hypervisor, there is likely a coordinating entity managing the virtual machines and virtual switches hosted on the servers. In the illustration, VMware vCenter is shown using a management network to control VMware ESX nodes. Virtual server networking can work directly with hypervisors running on individual systems, or through an aggregation point such as a hypervisor management system (VMware) or master hypervisor (Xen).

Figure 1. **Virtual Server Networking can register for events with multiple individual hypervisors, or a master hypervisor, or a hypervisor manager as appropriate to the hypervisor platform.**

Hypervisor Management APIs

For applications to operate seamlessly across a migration event, network configurations must become a part of the process. Associated VLANs must be made available on the physical switches where VMs will be located in advance of instantiating VMs, or moving VMs from one system to another.

Most hypervisor frameworks present a hierarchical object model that allows a third-party application such as Virtual Server Networking to express interest in one of more objects, such as the VM itself, the CPU, storage access, the virtual network, and object state. In the case of VMware and Citrix Xen, APIs are provided to allow registration for events related to pending actions on specific objects.

Since FTOS represents a generally-programmable environment, Virtual Server Networking implements this registration as a program running directly on the Dell Force10 switch.

After registering interest in a set of objects, Virtual Server Networking is notified of events related to those objects via the hypervisor API. Virtual Server Networking can communicate directly with hypervisors hosted on servers throughout the network or it can communicate with centralized hypervisor management systems (VMware) or master hypervisors (Xenserver) as required by the particular management software. Upon notification of a pending action – such as movement of a VM from one server to another – Virtual Server Networking can perform the necessary changes on the physical network, enabling ports and setting VLANs as appropriate.

Virtual Server Networking Support for VMware

For communication with the VMware ESX hypervisor, Virtual Server Networking utilizes the vSphere API. Unlike competing solutions that require the most expensive VMware Enterprise Plus Edition, Virtual Server Networking works with the more economical Standard, Advanced, or Enterprise VMware editions. Virtual Server Networking connects to the hypervisor and queries its attributes, and can also communicate directly with VMware vCenter to register for events related to groups of VMs.

Also, VMware vCenter is a centralized server management system that can manage more than 1,000 hosts and up to 10,000 VMs. VMware vCenter is accessed by Virtual Server Networking using the same vSphere API that is used to access a single VMware server. Only the vCenter access attributes need to be configured, and only the Virtual Server Networking connection with vCenter needs to be established in order to access the configurations of all of the VMware ESX servers managed by vCenter.

To connect with vCenter or an individual ESX hypervisor, a secure connection is established via HTTPS and user authentication is performed. These parameters are specified from the FTOS command line when the Virtual Server Networking command is issued to establish the connection. If vCenter is the target, the user name and password for vCenter are specified. If monitoring multiple hypervisors is desired, a separate Virtual Server Networking command is issued for each of them from the FTOS command line.

Figure 2 provides an architectural overview of Virtual Server Networking in the case of VMware. The right side of the figure depicts an example host system containing a virtual network that includes three virtual switches (vSW1, vSW2, and vSW3). In this configuration, each vSwitch connects to the physical network via a physical NIC. The left side of the diagram depicts a Dell Force10 switch that communicates with vSphere, using a secure vSphere session to pull configuration information via the Virtual Server Networking agent.

Figure 2. For VMware, Virtual Server Networking pulls configuration information via a secure vSphere connection.

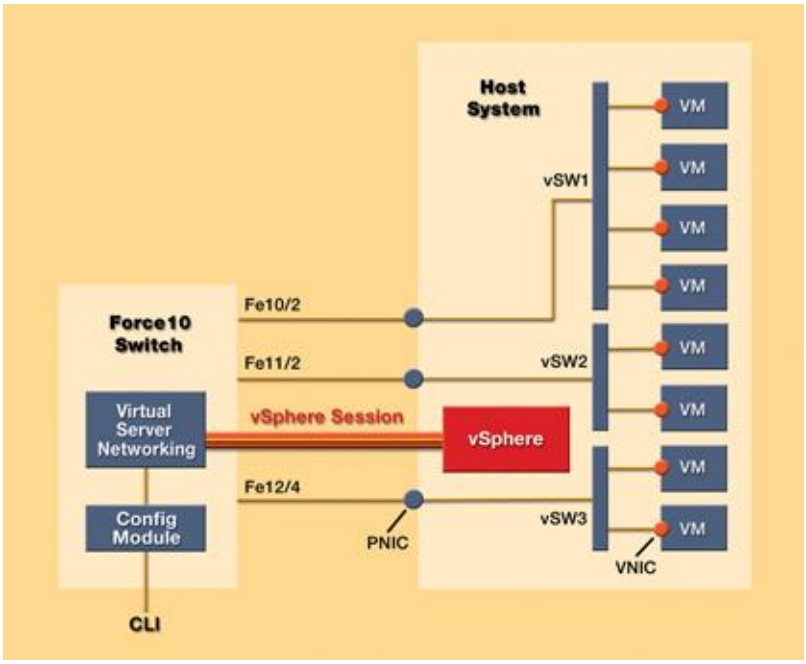

Since a Dell Force10 switch can be connected to multiple host systems, at least one Virtual Server Networking session is established with each of the host systems to be monitored. A Virtual Server Networking agent can also be configured to communicate directly with VMware vCenter to pull network configurations of all of the VMware hosts that are managed by that vCenter instance.

Once the connection is established, the Virtual Server Networking agent starts to pull all of the virtual network configurations from the host (or vCenter), sending configuration commands to the Config

Module based on the information obtained. For example, if there are VLANs configured on various vSwitches, those VLANs are then configured by the Config Module on the physical switch port.

vSphere APIs provide a notification mechanism whereby Virtual Server Networking can be notified whenever a monitored object is changed on a VMware host. In this way a VM or a VLAN being added or removed would generate a notification, and the FTOS Config Module could take the appropriate action. This model scales well with a large number of switches or a large number of ports.

Virtual Server Networking support for Citrix Xen

Figure 3 illustrates an architectural overview for Virtual Server Networking in the case of Citrix Xen. As before, this illustration shows several VMs connecting through three virtual switches, each associated with a physical NIC that connects to a port on the Dell Force10 switch. Similar to vSphere, the XenServer API also provides a notification mechanism by which Virtual Server Networking can be notified whenever a registered object is changed on a host. Rather than a centralized management application, Citrix Xen provides the ability for XenServer to be grouped into a loosely coupled cluster. The cluster of VMs consists of a master server and multiple slave servers.VLANs. The result is truly open network automation that allows organizations to choose the technology that best fits their needs while enjoying the benefits of automated virtualization technology.

Figure 3. To communicate with Citrix Xen, the Virtual Server Networking agent establishes a secure XenServer session.

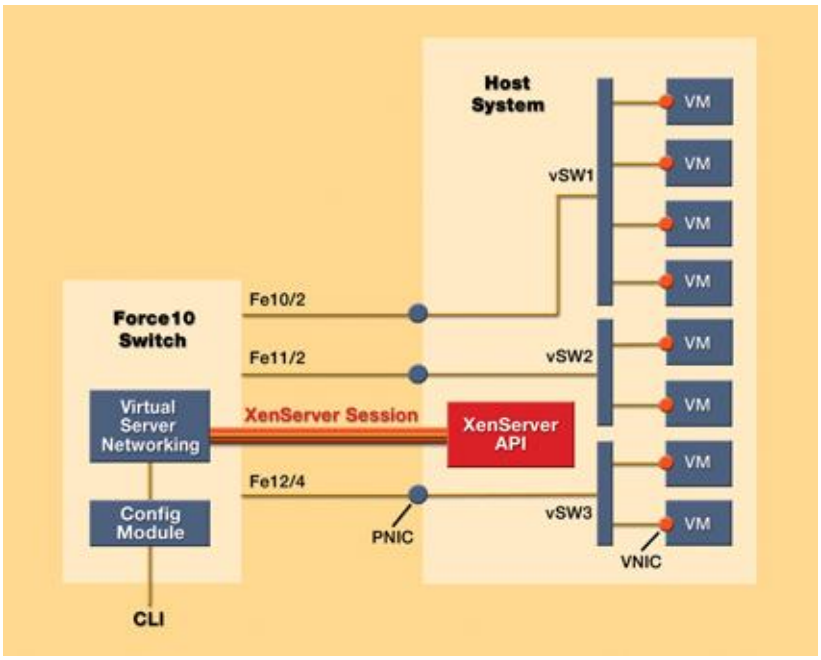

Upon initialization, the Virtual Server Networking agent initially connects to one of the XenServers in a group. The initially connected node will inform Virtual Server Networking who the master node is, and Virtual Server Networking will reconnect to the IP address of the master server. From the master server, the Virtual Server Networking agent can retrieve virtual network configurations for all of the servers in the group. Since the master server in a group may change, the Virtual Server Networking agent continues to contact the current master server, and only connects to a new master server if the current master server becomes a slave server.

Conclusion

As interest in virtualization technology continues to grow, network automation only becomes more important. Virtualized environments such as cloud computing depend on the ability to move VMs at will between systems, with networking components responding as necessary with appropriate configuration actions. While some vendors have implemented simplified and proprietary approaches to automation, Dell Force10 believes that interoperating with a wide variety of hypervisors requires an open approach, and Virtual Server Networking is a component of that strategy.

Through its switch-resident agent, Virtual Server Networking allows Dell Force10 switches to register interest in system and network objects hosted within virtualized systems from multiple vendors. This scaleable approach gives Dell Force10 switches key visibility into hosted networking components such as virtual switches – without requiring a host-resident agent – and allows physical network elements to be configured in anticipation of pending dynamic actions on virtual machines or VLANs. The result is truly open network automation that allows organizations to choose the technology that best fits their needs while enjoying the benefits of automated virtualization technology.