Enhanced VXLAN on Cisco Nexus 1000V Switch for VMware vSphere

Deployment Guide

July 2013

Overview

This document provides guidelines for deploying Virtual Extensible LAN (VXLAN) technology on the Cisco Nexus® 1000V Switch for VMware vSphere. For detailed configuration documentation, please refer to the configuration guides at http://www.cisco.com.

Audience

This document is intended for network architects, network engineers, virtualization administrators, and server administrators interested in understanding and deploying Cisco Nexus 1000V VXLAN technology on the Cisco Nexus 1000V Switch for VMware vSphere.

Introduction

Many customers today are implementing private or public clouds. A cloud-based deployment model helps customers share computing, network, and storage infrastructure among multiple tenants. With a cloud-based deployment model, customers can increase ROI and utilization of physical resources and quickly provision applications. Customers can benefit immensely from this new deployment model, but they require complete isolation between tenants that share the physical infrastructure.

Traditionally, the physical network has been shared using IEEE 802.1Q VLANs. However, the IEEE 802.1Q standard allows only 4096 unique segments, making these VLANs unsuitable for a cloud deployment. Cloud provider environments must accommodate a huge number of tenants on the same underlying physical infrastructure, and each tenant may implement multiple Layer 2 networks, creating a need for more Layer 2 networks. Another limitation of VLANs is the limited mobility domain of virtual machines deployed on them. Typically, a VLAN is restricted to a few racks in the data center, with the result that the mobility domain is small.

Cisco and a group of industry vendors are working together to address new requirements for scalable LAN segmentation and to expand the placement and mobility domains of virtual machines. The proposed technology, referred to as Virtual Extensible LAN (or VXLAN), defines a 24-bit LAN segment identifier to provide segmentation at cloud-deployment scale. In addition, the proposed solution extends the LAN segment across IP subnets, expanding the placement and mobility domains of virtual machines. More information can be found in the IETF draft: http://datatracker.ietf.org/doc/draft-mahalingam-dutt-dcops-vxlan/?include_text=1.

Background

Cisco Nexus 1000V Networking

The Cisco Nexus 1000V Switch for VMware vSphere provides Layer 2 switching, advanced networking functions, and a common network management model in a virtualized server environment by replacing the virtual switch within VMware vSphere. As Figure 1 shows, the Cisco Nexus 1000V Switch for VMware vSphere manages a data center as defined in the VMware vCenter Server. Each server in the data center is represented as a line card in the Cisco Nexus 1000V Switch for VMware vSphere and can be managed as if it were a line card in a physical Cisco® switch.

The Cisco Nexus 1000V Switch for VMware vSphere implementation has two main components:

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

Figure 1. Cisco Nexus 1000V Switch for VMware vSphere Switches Managing VMware ESX Servers

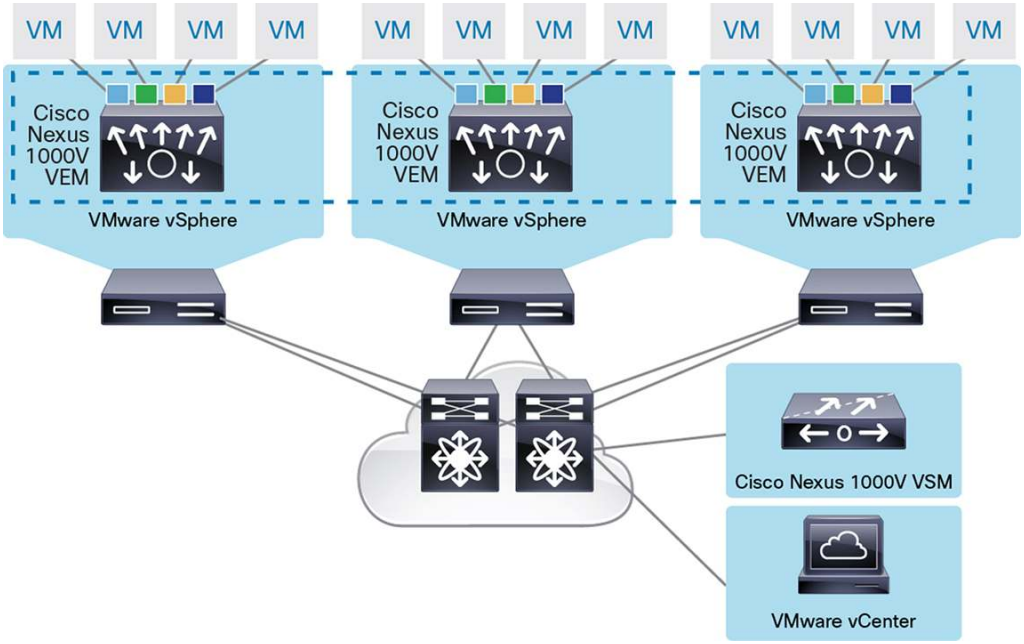

The benefits of deploying the Cisco Nexus 1000V Switch for VMware vSphere include:

- Advanced networking capabilities, such as quality of service (QoS), network statistics collection with Cisco NetFlow Collector, packet mirroring with Cisco ERSPAN, and many others
- Nondisruptive operational model, with the Cisco Nexus 1000V Switch for VMware vSphere fully integrated into VMware vCloud Director and vCenter Server
- Easier regulatory compliance for applications in the cloud because there is complete transparency in both the physical and virtual networks

Overview of Cisco Nexus 1000V VXLAN in Multicast Mode

Cisco VXLAN is a Layer 2 network isolation technology that uses a 24-bit segment identifier to scale beyond the 4096-segment limitation of VLANs. VXLAN technology creates LAN segments by using an overlay approach with MAC-in-IP encapsulation. The VEM encapsulates the original Layer 2 frame from the virtual machine (Figure 2).

Figure 2. VXLAN Encapsulated Frame Format

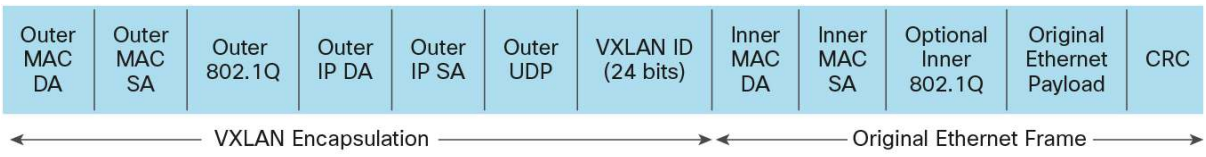

Each VEM is assigned an IP address, which is used as the source IP address when encapsulating MAC frames to be sent on the network. This encapsulation is accomplished by creating virtual machine kernel network adaptors (VMKNICs) on each VEM (Figure 3). You can have multiple VMKNICs per VEM and use them for this encapsulated traffic. The encapsulation carries the VXLAN identifier, which is used to scope the MAC address of the payload frame.

Figure 3. VEM VMKNIC Interface with VXLAN Capability

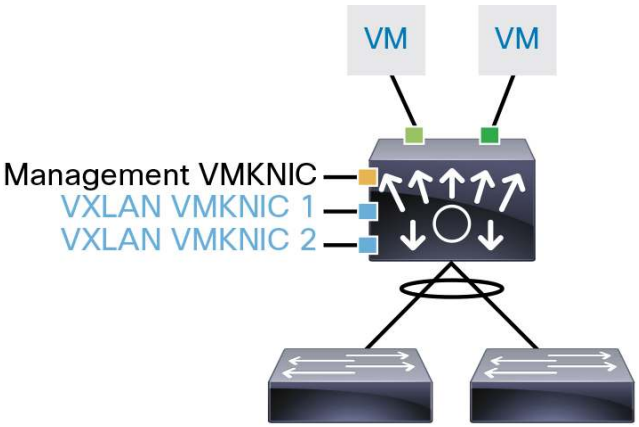

The connected VXLAN is specified within the port profile configuration of the vNIC and is applied when the virtual machine connects. Each VXLAN uses an assigned IP multicast group to carry broadcast traffic within the VXLAN segment.

When a virtual machine attaches to a VEM, if it is the first to join the particular VXLAN segment on the VEM, an Internet Group Management Protocol (IGMP) join is issued for the VXLAN's assigned multicast group. When the virtual machine transmits a packet on the network segment, a lookup occurs in the Layer 2 table using the destination MAC address of the frame and the VXLAN identifier. If the result is a hit, the Layer 2 table entry will contain the remote IP address to use to encapsulate the frame, and the frame will be transmitted within an IP packet destined for the remote IP address. If the result is a miss (as with broadcast, multicast, and unknown unicast traffic, for instance), the frame is encapsulated with the destination IP address set to the VXLAN segment's assigned IP multicast group. To configure a VXLAN in multicast mode, refer to the deployment guide at http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/guide_c07-702975.pdf.

Enhanced VXLAN for Cisco Nexus 1000V Switch for VMware vSphere

The Cisco Nexus 1000V Switch for VMware vSphere provides innovations and improvements for VXLAN. The following enhancements to VXLAN are introduced in the Cisco Nexus 1000V Switch for VMware vSphere:

- VXLAN Unicast-only mode
- MAC Address Distribution
- VXLAN Trunks
- Multi-Mac Capability

VXLAN Unicast-Only Mode

The first implementation of VXLAN in the Cisco Nexus 1000V Switch for VMware vSphere relied on configuration of multicast in the physical network to help ensure that unknown unicast, broadcast, and multicast frames are delivered to all VTEPs that have virtual machines in a VXLAN. However, some customers do not have multicast configured in the network and prefer not to introduce such a configuration. To bring VXLAN to such customers, the Cisco Nexus 1000V Switch for VMware vSphere introduces a unicast-based implementation of VXLAN: unicast-only mode.

When the virtual machine sends a frame in a VXLAN configured in unicast-only mode, a lookup occurs in the Layer 2 MAC address table using the destination MAC address of the frame and the VXLAN identifier. If the result is a hit, the Layer 2 table entry will contain the remote VTEP IP address to use to encapsulate the frame, and the frame will be transmitted within an IP unicast packet destined for the remote VTEP. If the result is a miss (as with broadcast, multicast, and unknown unicast traffic, for instance), the frame is replicated for each VTEP that has active virtual machines in the same VXLAN. The replicated packet is encapsulated and is sent as an IP unicast packet to the destination VTEPs.

The Cisco Nexus 1000V VEM achieves this intelligent replication by maintaining a VTEP table. When a virtual machine is deployed on a VXLAN, the VEM on the host adds an **(VXLAN, VTEP IP)** entry in this table. The Cisco Nexus 1000V VSM continuously aggregates and distributes all the VTEPs for a given VXLAN to all the VEMs. As a result of this behavior, the Cisco Nexus 1000V VEM always has the VTEP IP addresses of all the VEMs interested in traffic for a given VXLAN. Each VEM maintains a VXLAN VTEP table. In the example in Figure 4, VXLAN 5000 has virtual machines in VTEPs 10.10.10.10, 20.20.20.20, and 30.30.30.30. VXLAN 6000 has virtual machines only in VTEPs 10.10.10.10 and 20.20.20.20.

Figure 4. VEM VTEP Table

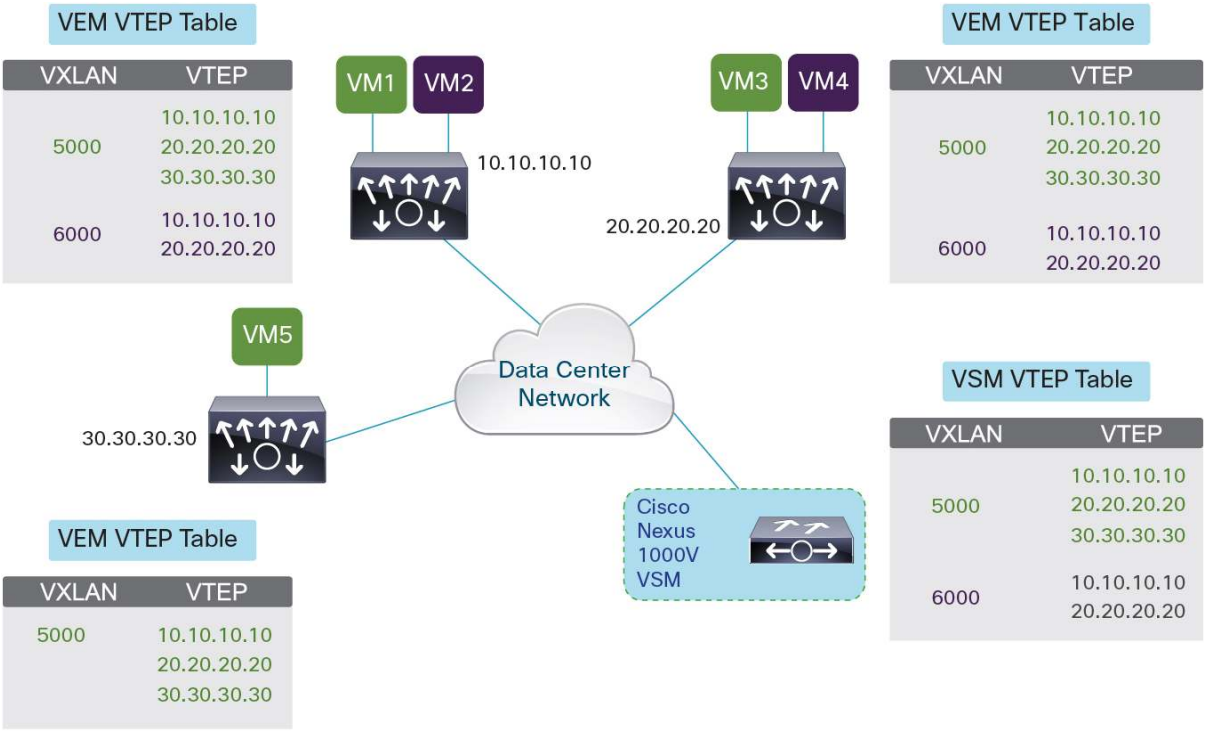

When an unknown unicast or broadcast frame is received on a VXLAN segment, the VEM performs ingress replication and encapsulates the frame, as shown in Figure 5. The replicated frame is then encapsulated with a unicast VTEP IP address and sent directly to the VEMs that have virtual machines in the ingress VXLAN.

Figure 5. Ingress Replication of Unknown Unicast, Broadcast, and Multicast Traffic with Enhanced VXLAN

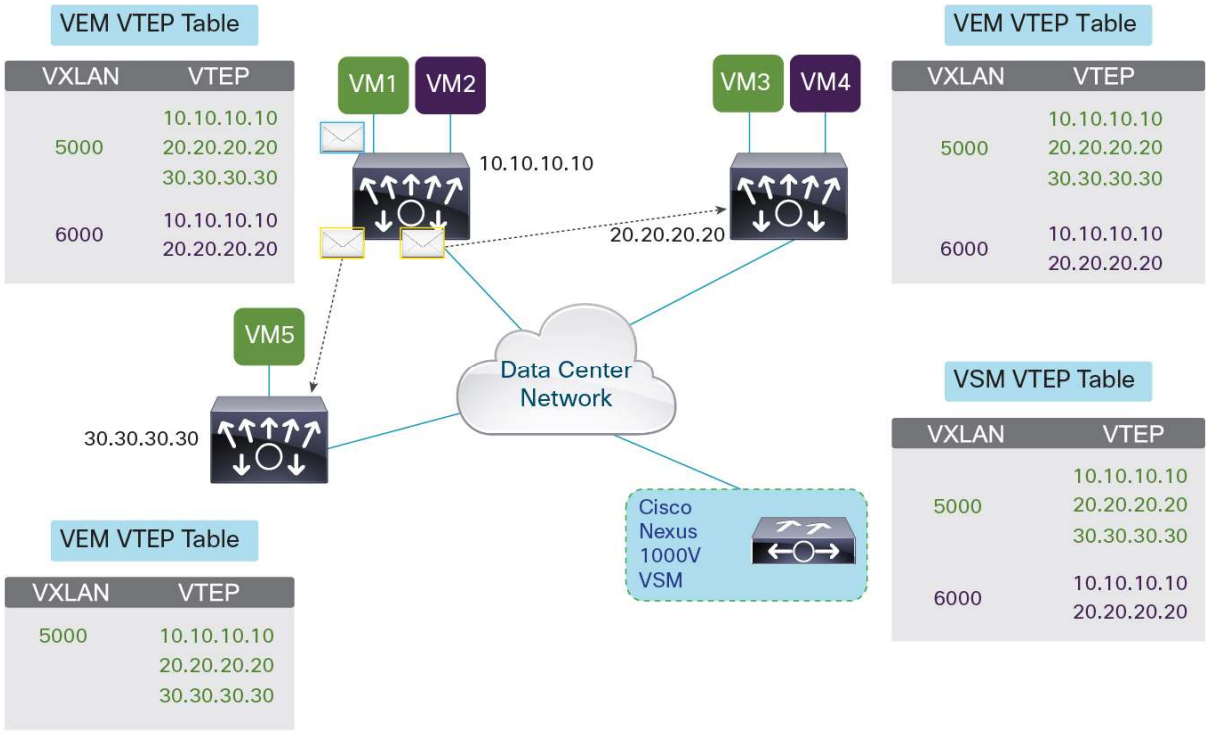

In this implementation, a VXLAN cannot be extended across VEMs managed by multiple Cisco Nexus 1000V VSMs. In other words, the scope of a VXLAN in unicast-only mode is all VTEPs managed by a single Cisco Nexus 1000V VSM. To extend a VXLAN across multiple VSMs, customers can continue to configure VXLAN in multicast mode.

MAC Address Distribution

The Cisco Nexus 1000V introduces MAC address distribution to reduce the amount of flooding in the network. When MAC address distribution is configured on a VXLAN in unicast-only mode, unknown unicast floods are suppressed. On a traditional IEEE 802.1Q bridge, unknown unicast traffic is flooded to all ports in the VLAN. A malicious virtual machine or host can waste replication bandwidth across the entire network by continuously sending unicast frames to a random unicast MAC address. The Cisco Nexus 1000V prevents such an unknown unicast attack with the MAC address distribution enhancement.

With this feature, the VSM continuously learns all MAC addresses in a VXLAN, and the VSM distributes this information to the VEMs. The VEMs use this information to program static MAC address entries in their MAC address tables. At any given time, all VEM MAC address tables contain all the known MAC addresses in the VXLAN domain, as shown in Figure 6.

Figure 6. MAC Address Distribution with Enhanced VXLAN

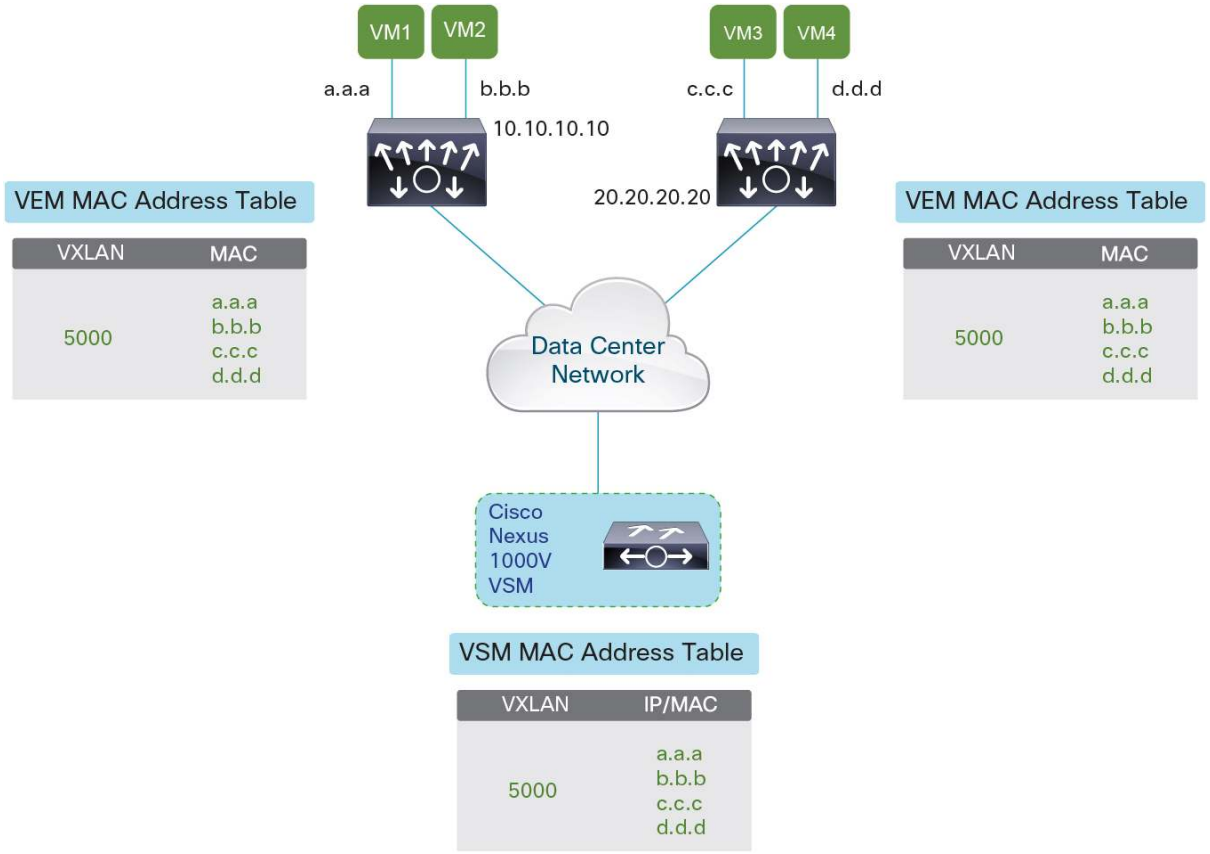

If a malicious virtual machine sends a unicast frame to a random unicast MAC address, the lookup in the MAC address table will fail. Because the VEM knows all the MAC addresses in a VXLAN when the MAC address distribution feature is enabled, the frame can dropped, avoiding the need to flood frames to all VTEPs in the VXLAN. In the example in Figure 7. a malicious virtual machine, VM4, sends a unicast frame to MAC address a.b.c. The lookup on the MAC address table in VTEP 20.20.20.20 for MAC address a.b.c in VXLAN 500 fails. Because MAC address distribution is configured on VXLAN, the frame can be dropped.

Figure 7. Unknown Unicast Frame Dropped When MAC Address Distribution Is Enabled

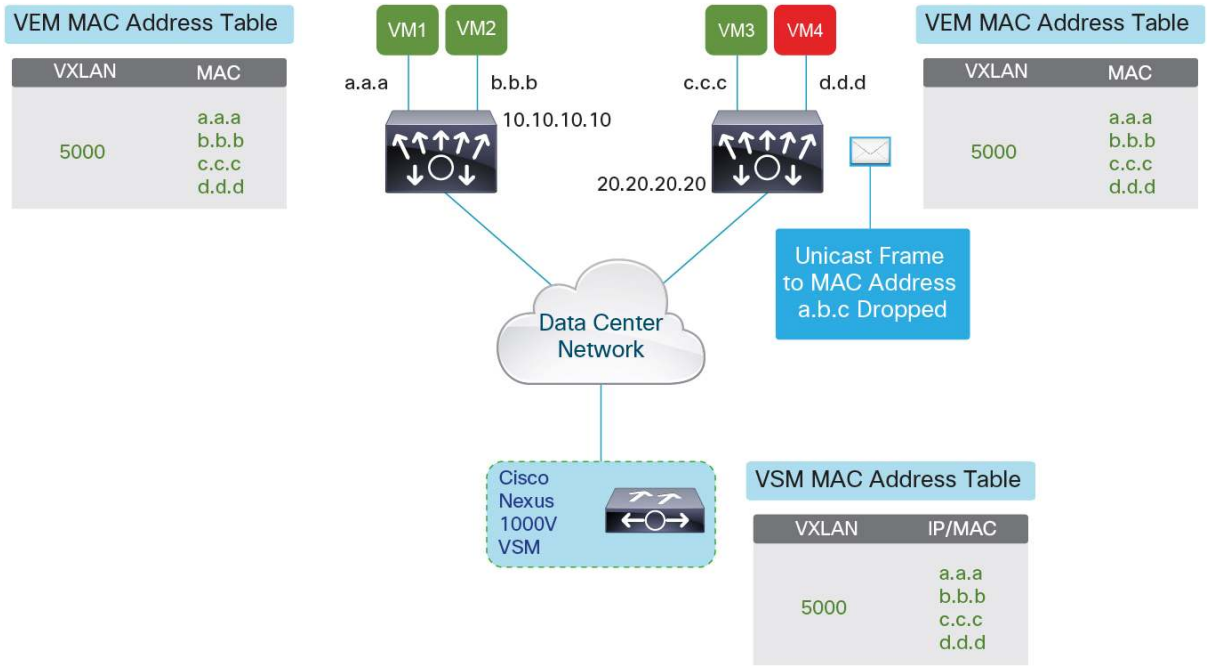

Note: MAC address distribution can be configured only on VXLANs in unicast-only mode. Because VXLAN unicast-only mode is a single VSM solution, MAC address distribution extends only across VEMs managed by a single VSM.

Forward-Publish-Capable and Forward-Publish-Incapable VTEPs

The Cisco Nexus 1000V Switch for VMware vSphere introduces the concepts of forward-publish-capable and forward-publish-incapable VTEPs in a VXLAN overlay. A forward-publish-capable VTEP can always publish MAC addresses of all virtual machines connected to it. The Cisco Nexus VEM is an example of a forward-publish-capable VTEP because the VEM can always identify and publish the MAC addresses of virtual machines deployed on the VLAN or VXLAN.

A forward-publish-incapable VTEP cannot publish the MAC addresses of all the virtual machines and hosts connected it. An example of a forward-publish-incapable VTEP on the VXLAN overlay is the Cisco Nexus 1000V VXLAN Gateway (Figure 8). The Cisco Nexus 1000V VXLAN Gateway is used to connect a VXLAN segment to a VLAN segment at Layer 2. A logical instance of this gateway is a 2-port Layer 2 learning bridge that connects a particular VXLAN segment to an IEEE 802.1Q VLAN. For more information about Cisco Nexus 1000V VXLAN Gateway, refer to the deployment guide at http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/guide_c07-728864.html.

Figure 8. Unknown Unicast Frame Sent to Forward-Publish-Incapable VTEP in VXLAN Overlay

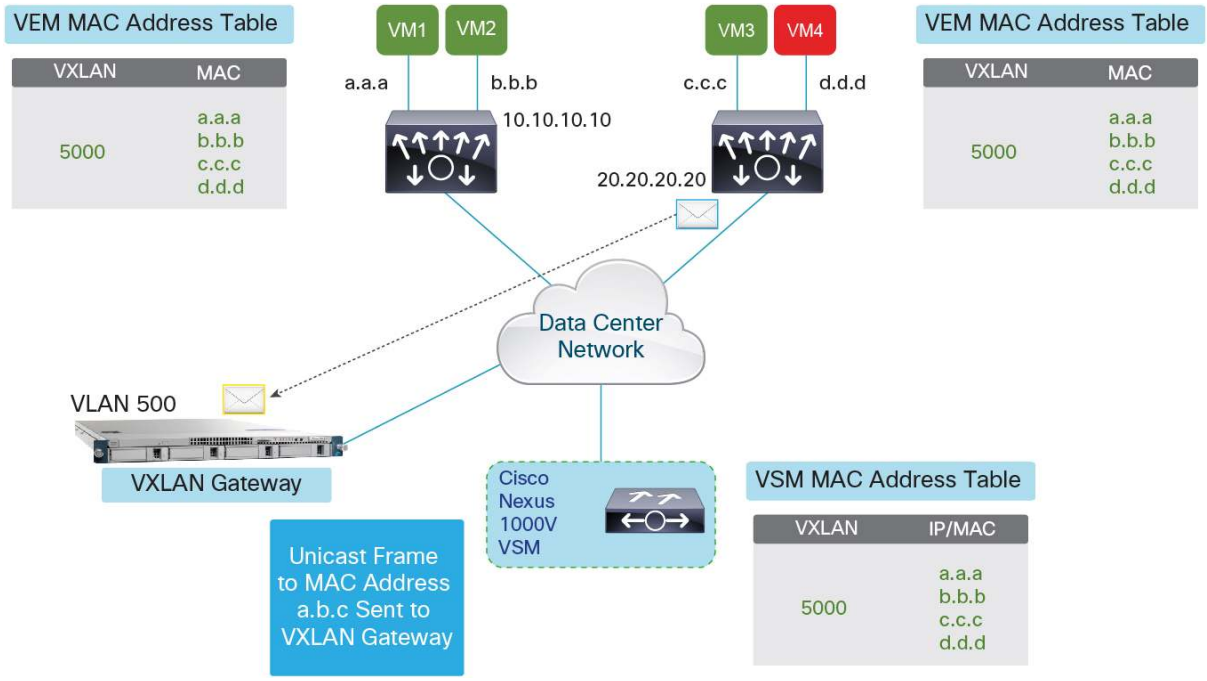

When MAC address distribution is enabled, by default unknown unicast frames are dropped. However, if a forward-publish-incapable VTEP such as the Cisco Nexus 1000V VXLAN Gateway is present on the overlay, when an unknown unicast or broadcast frame is received, the Cisco Nexus 1000V VEM encapsulates the frame with the IP address of the VXLAN gateway and sends it to the gateway device. A VXLAN gateway that implements VXLAN-to-VLAN mapping decapsulates the frames received on the VXLAN and floods them on the VLAN.

Note: Forward-publish-capable and forward-publish-incapable VTEPs are relevant only when a customer implements a VXLAN gateway in the network. By default, all VEMs are forward-publish-capable VTEPs.

VXLAN Trunks

The Cisco Nexus 1000V Switch for VMware vSphere only supports access port mode for VXLAN switch ports. To extend multiple VXLAN segments to a virtual machine, the administrator can create multiple virtual Ethernet interfaces on the virtual machine and connect them to VXLAN access ports on the Cisco Nexus 1000V. However, the drawback of this approach is that the administrator is limited by the number of virtual Ethernet interfaces supported by the hypervisor. On VMware vSphere 5.1, a maximum of 10 virtual Ethernet interfaces can be created on a virtual machine. This restriction severely limits the number of segments extended to a service virtual machine such as a virtual router like the Cisco Cloud Services Router (CSR) 1000V.

The Cisco Nexus 1000V Switch for VMware vShpere introduces VXLAN trunks to scale the number of VXLAN segments exposed to a virtual machine. This feature is implemented by defining an IEEE 802.1Q–to–VXLAN mapping on a port profile. A typical use for VXLAN trunks is to deploy a network service virtual machine such as the Cisco CSR 1000V on a VXLAN trunk interface.

As shown in Figure 9, VXLAN trunks are created by defining an IEEE 802.1Q–to–VXLAN mapping in the port profile service instance configuration. The scope of this mapping extends only to the virtual Ethernet interfaces that inherit this port profile. When a packet is received with an IEEE 802.1Q tag on a port with VXLAN trunk mapping, the packet is mapped to the corresponding VXLAN, where it is then bridged. Currently, up to 4094 mappings can be defined for a port profile service instance. For information about configuring VXLAN trunk mapping, refer the configuration guide.

Figure 9. VXLAN Trunk Interface with IEEE 802.1Q–to–VXLAN Mappings

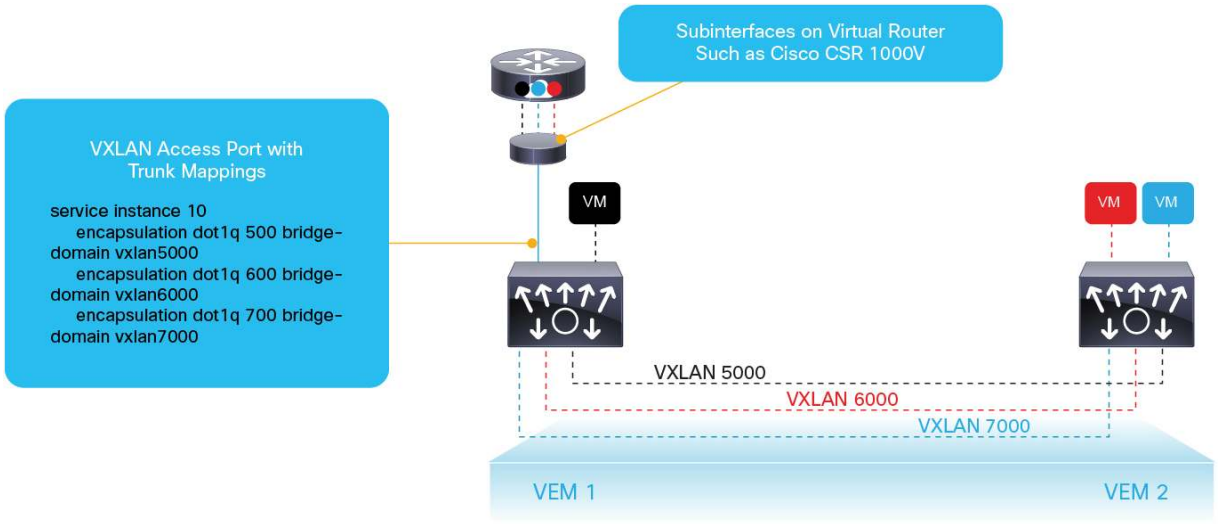

VXLAN Multi-MAC Address Capability

The Cisco Nexus 1000V Switch for VMware vSphere introduces the VXLAN multi-MAC address capability to correctly handle live migration scenarios for service virtual machines such as the Cisco CSR 1000V. When the Cisco CSR 1000V is deployed to the Cisco Nexus 1000V VEM, the VEM learns of remote MAC addresses behind the Cisco CSR as dynamic MAC addresses on the VXLAN. When the Cisco CSR virtual machine is migrated to a new VEM, all the dynamic MAC addresses learned from the Cisco CSR will continue to point to the old VEM. This behavior occurs because the VEM receives MAC address move notifications only for static MAC addresses registered with the virtual Ethernet adapter. This behavior can potentially prevent the dynamic MAC addresses from receiving traffic until the remote MAC addresses are relearned.

To avoid this scenario, the VXLAN multi-MAC address capility was introduced. When a port profile with **capability multi-mac** is applied to a virtual machine, upon live migration of the virtual machine the VEM generates a special flush frame. The flush frame requests all VEMs to flush MAC addresses learned behind the source VEM. When this flush frame with the VXLAN segment ID and the old VEM information is received, all VEMs flush all MAC addresses learned behind the source VEM in the VXLAN segment. Traffic to these remote dynamic MAC addresses is flooded until these MAC addresses are relearned behind the new VEM.

Deployment Considerations

Single VSM Solution

With the VXLAN unicast mode, a VXLAN segment can span any VEM managed by a single VSM; a VXLAN segment defined on one VSM cannot be made available to a VEM managed by another VSM. To create a segment that spans hosts and VEMs managed by multiple VSMs, the administrator needs to configure the VXLAN in multicast mode. This configuration is required because the VSM continuously learns of all the virtual machines in a VXLAN only from VEMs that it manages.

Coexistence of Unicast and Multicast Modes

Starting with Cisco Nexus 1000V Release 4.2(1)SV2(2.1), the Cisco Nexus 1000V supports VXLAN in both unicast and multicast modes. By default, all VXLAN segments are created in unicast mode. MAC address distribution is disabled by default.

The administrator can also set the global default to multicast mode. In addition, the switch supports a segment-level configuration that always overrides the global setting. For example, if the global mode is set to the default unicast mode and the administrator wants to create a VXLAN in multicast mode, the administrator can overide the global setting by setting the mode to multicast in the VXLAN segment configuration. This capability enables the coexistence of unicast and multicast modes (Figure 10).

Figure 10. Setting VXLAN Segments to Multicast Mode

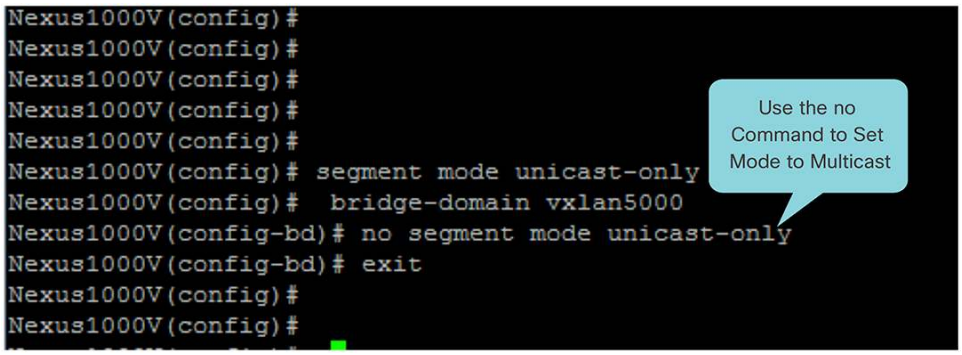

PortChannels

PortChannels use different load-sharing algorithms to divide outgoing traffic among different physical interfaces. When IP encapsulation is used, all outgoing VXLAN traffic carries an outer header that has the source MAC and IP addresses of the VEM's VMKNIC. For optimal load balancing, users must configure a 5-tuple-based hash as the load-sharing algorithm. The use case section of this document discusses how to configure 5-tuple-based hashes.

Maximum Transmission Unit Size

The Cisco Nexus 1000V uplink port profiles and all interconnecting switches and routers between the VMware ESX hosts must have the supported maximum transmission unit (MTU) set to at least 50 bytes larger than the MTU of the virtual machines. For example, the virtual machines default to a 1500-byte MTU (the same as the uplinks and physical devices), so you must set the MTU to at least 1550 bytes. If this configuration is not possible, you should lower all virtual machine virtual network interface card (vNIC) MTU values to 50 bytes less than what the physical network supports, such as 1450 bytes. For more information, see the Cisco Nexus 1000V Port Profile Configuration Guide.

Securing VXLAN in the Physical Network

VXLAN is transported over IP in a physical network, so you should implement some best practices when setting up the transport network for VXLAN.

The preferred option is to have all the VXLAN VMkernel interfaces on the VEM in the same subnet. In this scenario, you can make the interfaces part of the same VLAN and keep that VLAN a strict Layer 2 VLAN. Only the VMKNICs used for VXLAN encapsulation should be attached to this VLAN. This approach will provide protection and limit unwanted exposure to external communications.

In a scenario in which the number of VEMs exceeds the available IP addresses in the subnet, VMKNICs for VXLAN encapsulations may need to be assigned IP addresses in multiple subnets. In this scenario, in which VXLAN VMkernel interfaces belong to two different VLANs, the communications between the multiple subnets has to take place through a Layer 3 switch or router. Both VLANs must have switch virtual interfaces (SVIs). To make sure that VXLAN traffic cannot be attacked or snooped from unauthorized endpoints, use one of following options:

- Use access control lists (ACLs) to prevent unauthorized injection of VXLAN encapsulated traffic into VEM VMKNICs from outside sources.
- Use virtual routing and forwarding (VRF) to segregate the VLANs and SVIs on which VXLAN VMKNICs are assigned IP addresses. For specific configurations of ACLs or VRFs, refer to the configuration guides for your physical Layer 3 switch or router.

The recommended scenarios just described not only reduce external security threats, but also simplify the multicast deployment in the physical network.

Scalability

Today, a single Cisco Nexus 1000V Switch for VMware vSphere supports up to 2048 VLANS and 2048 VXLANs. To deploy more than 2048 VXLANs, additional Cisco Nexus 1000V Switches for VMware vSphere must be deployed.

Routing for VXLANs

To provide external Layer 3 connectivity to a VXLAN, the administrator can use a virtual appliance such as the Cisco ASA 1000V to provide Layer 3 gateway and Network Address Translation (NAT) functions. For more information about how to configure the Cisco ASA 1000V, refer to http://www.cisco.com/en/US/products/ps12233/index.html.

Upgrades

By default, if you upgrade the VSM from an earlier version of the Cisco Nexus 1000V to the current version with the segmentation feature enabled, all the VXLANs continue to operate in multicast mode. If you enable the feature when the VSM is running the current version of the Cisco Nexus 1000V, by default the bridge domains change to unicast-only mode. You must explicitly enable MAC address distribution mode because it is disabled by default.

After completing the software upgrade, you need to explicitly set the segment mode to multicast mode.

Setting Up VXLAN

This section presents the steps for setting up VXLAN.

Step 1. Turn on the network segmentation manager (NSM) and VXLAN feature on the Cisco Nexus 1000V.

```
N1KV-VSM (config)# feature segmentation
```

Verify that the feature is enabled on the Cisco Nexus 1000V

```
N1KV-VSM(config)# show feature
Feature Name           Instance    State
------------           --------    ----------------
dhcp-snooping          1           disabled
http-server            1           enabled
lacp                   1           disabled
netflow                1           disabled
network-segmentation   1           disabled
port-profile-roles     1           disabled
private-vlan           1           disabled
segmentation           1           enabled
sshServer              1           enabled
tacacs                 1           disabled
telnetServer           1           disabled
```

Verify that the default segment mode is unicast and that the MAC address distribution is disabled.

```
Nexus1000V(config)# show running-config bridge-domain
!Command: show running-config bridge-domain
!Time: Thu Feb 14 22:29:28 2013
version 4.2(1)SV2(2.1)
feature segmentation
segment mode unicast-only
no segment distribution mac
```

Step 2. Create a port profile with VXLAN capability.

```
port-profile type vethernet VXLAN-PP
   vmware port-group
  switchport access vlan 300
  capability vxlan
  no shutdown
  state enabled
```

Verify that the VXLAN is enabled on this interface by entering this command:

```
Nexus1000V(config)# show port-profile name VXLAN-PP

port-profile VXLAN-PP
 type: Vethernet
 description:
 status: enabled
 max-ports: 32
 min-ports: 1
```

```
inherit:
config attributes:
 capability vxlan
 switchport access vlan 300
 switchport mode access
 no shutdown
evaluated config attributes:
 capability vxlan
 switchport access vlan 300
 switchport mode access
 no shutdown
assigned interfaces:
 Vethernet4
 Vethernet5
port-group: VXLAN-PP
system vlans: 300
capability l3control: no
capability iscsi-multipath: no
capability vxlan: yes
capability l3-vservice: no
port-profile role: none
port-binding: static
```

Step 3. Create a VMkernel interface on the VMware ESX host.

Attach a VMkernel interface to each VMware ESX host for the cluster in VMware vCenter.

First, choose Home > Inventory > Host and Clusters in VMware vCenter. Next, choose Configuration > Networking > vSphere Distributed Switch and select Manage Virtual Adapters, as shown in Figure 11.

Figure 11. Selecting Manage Virtual Adapters

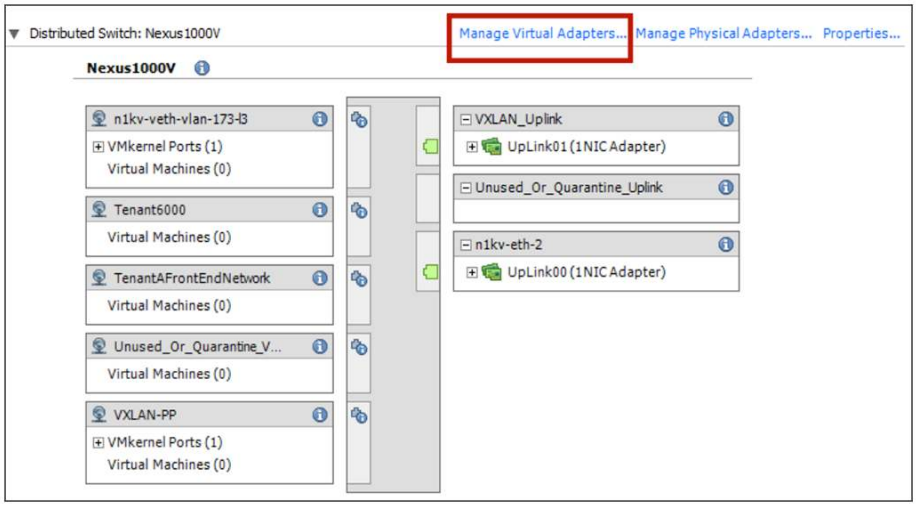

Add a new VMkernel interface, as shown in Figures 12 through 14.

Figure 12. Selecting Add in the Manage Virtual Adapters Window

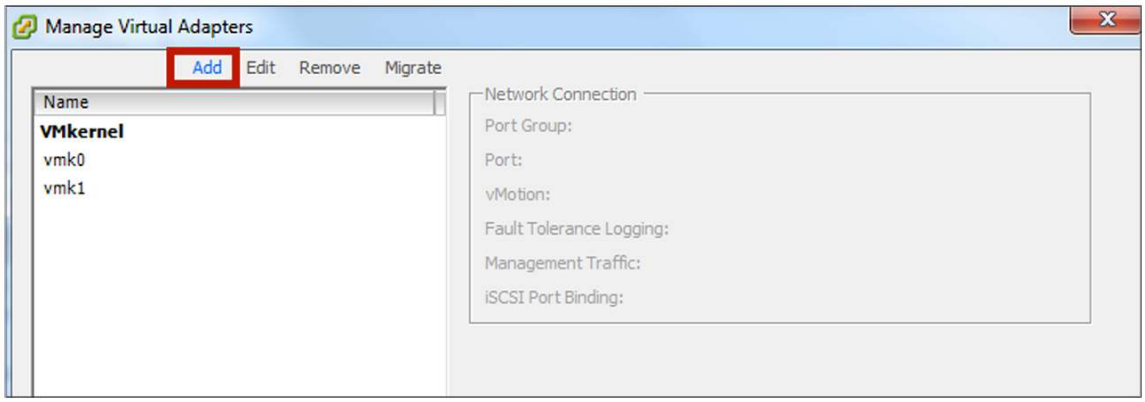

Figure 13. Selecting a New Virtual Adapter

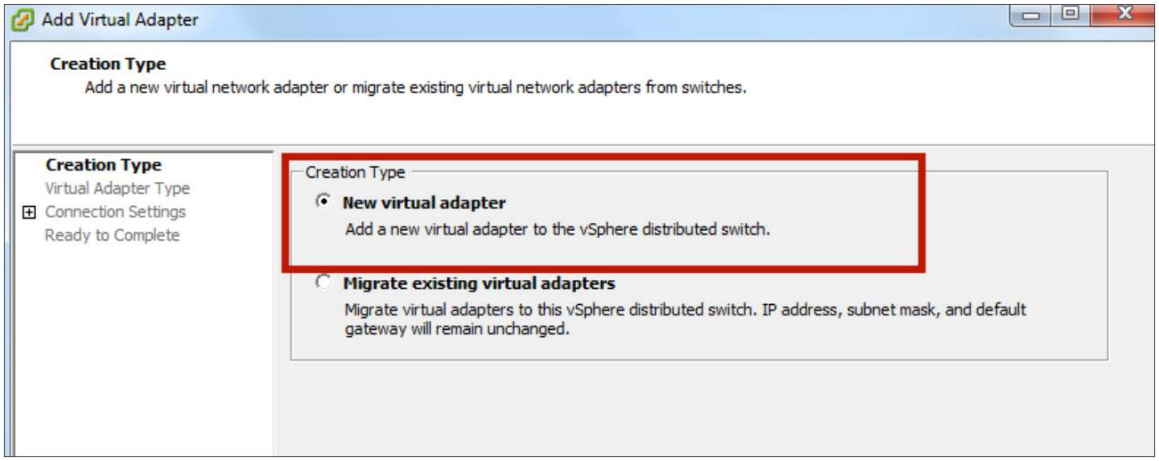

Figure 14. Selecting a New VMkernel Interface

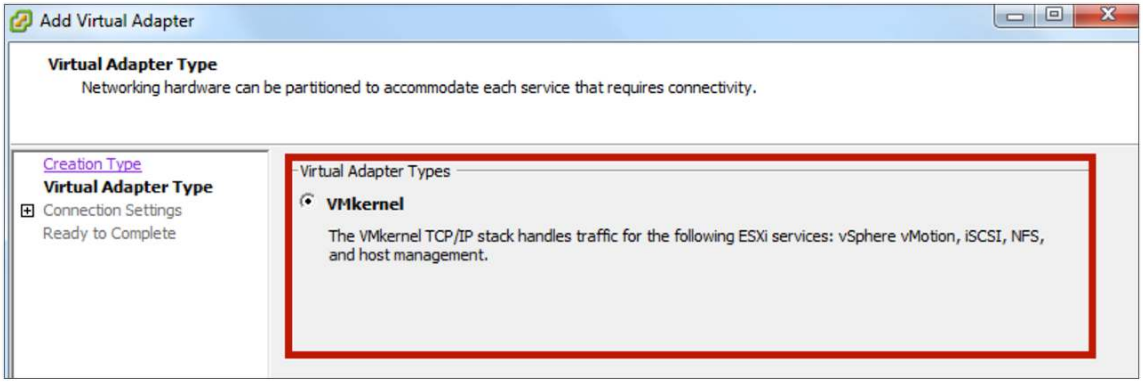

Finally, select the VXLAN-enabled port profile and configure the IP-to-VMkernel interface used to encapsulate the VXLAN, as shown in Figures 15 and 16.

Figure 15. Selecting the VXLAN-Enabled Port Profile

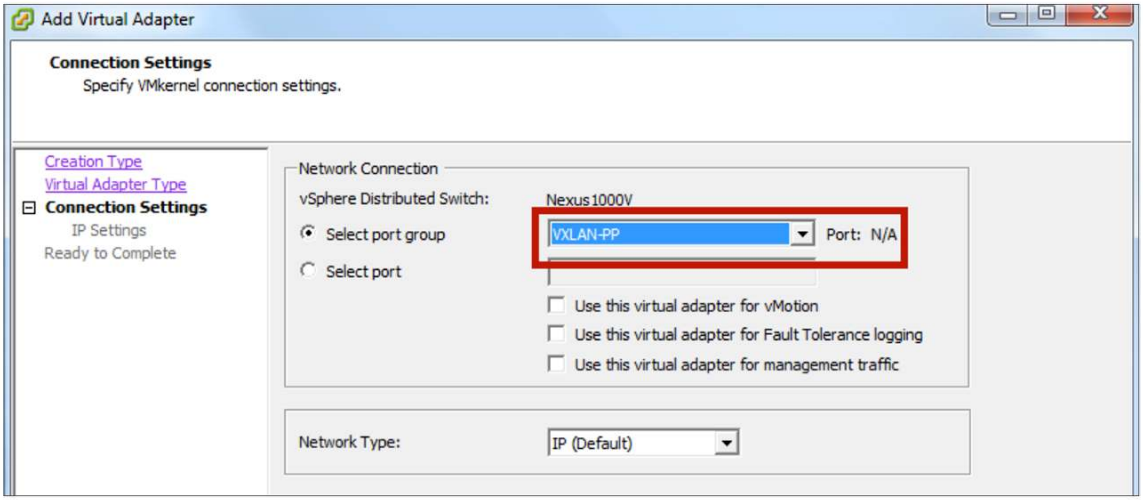

Figure 16. Configuring the IP-to-VMkernel Interface Used to Encapsulate VXLAN

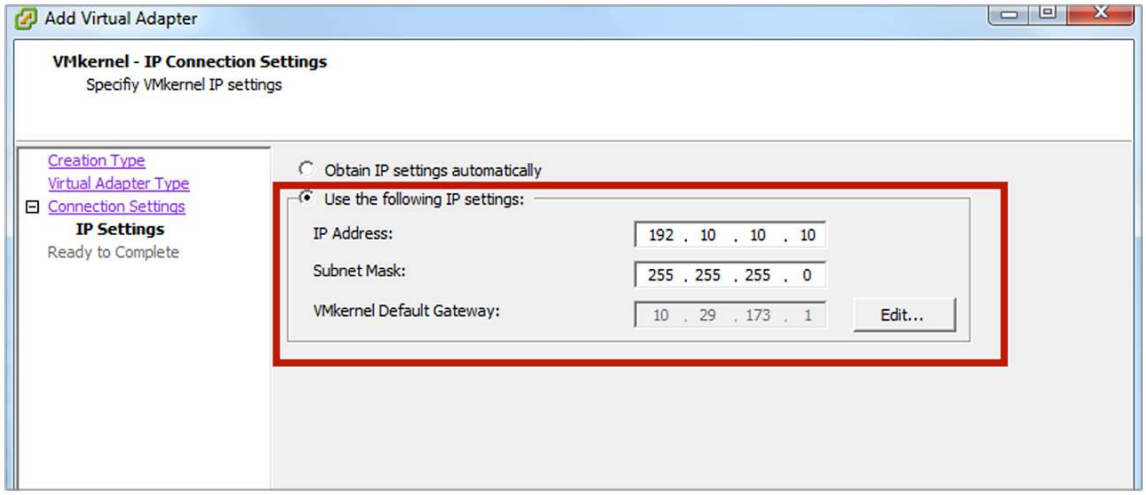

Figure 17 shows a summary of the new VMkernel interface.

Figure 17. Summary of VXLAN VMkernel Interface

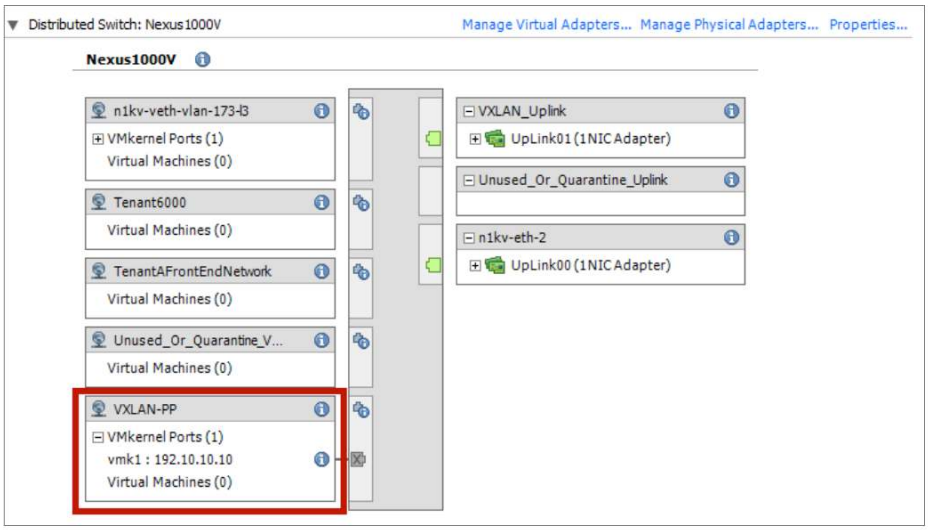

Repeat the preceding steps for other VMware ESX hosts. The only difference is that you need to assign a unique IP address for each interface created on the host.

On the VSM, verify that the interfaces are up on that Layer 3 VMkernel interface by entering the following command:

```
Nexus1000V(config)# show port-profile name VXLAN-PP

port-profile VXLAN-PP
 type: Vethernet
 description:
 status: enabled
 max-ports: 32
 min-ports: 1
 inherit:
 config attributes:
  capability vxlan
  switchport access vlan 300
  switchport mode access
  no shutdown
 evaluated config attributes:
  capability vxlan
  switchport access vlan 300
  switchport mode access
  no shutdown
 assigned interfaces:
  Vethernet4
  Vethernet5
 port-group: VXLAN-PP
```

```
 system vlans: 300
 capability l3control: no
 capability iscsi-multipath: no
 capability vxlan: yes
 capability l3-vservice: no
 port-profile role: none
 port-binding: static

Nexus1000V(config)#
```

The two virtual VMkernel interfaces (vEthernet 4 and vEthernet 5) belong to two different VMware ESX hosts in the example.

Step 4. Change the MTU on the uplink interface.

To avoid fragmentation, you should increase the MTU of the uplink interfaces of the Cisco Nexus 1000V and the physical interfaces of the upstream switches that are connected the Layer 2 domain of the VMware vSphere environment.

Configure the following command on the uplink port profile to increase the MTU:

```
port-profile type ethernet VXLAN_Uplink
  vmware port-group
  switchport trunk allowed vlan 300
  switchport mode trunk
  switchport trunk native vlan 300
  mtu 1550
  no shutdown
  system vlan 300
  state enabled
```

Refer to the system configuration guides for your upstream switches to increase the MTU values of the physical interfaces for all the transit switches and routers.

Step 5. Create the VXLAN (bridge domain) in VSM.

Now you are ready to create VXLAN IDs in the VSM and place the virtual machines on the VXLAN. On the VSM, configure a new bridge domain as follows:

```
config t
bridge-domain vxlan5000
  segment id 5000
```

Step 6. Turn on MAC address distribution on the VXLAN.

This step is required only if MAC address distribution is disabled at the global level.

```
config t
bridge-domain vxlan5000
   segment distribution mac
```

Verify the bridge domain status using the following show command:

```
Nexus1000V(config-bd)# show bridge-domain vxlan5000
```

```
Bridge-domain vxlan5000 (0 ports in all)
Segment ID: 5000 (Manual/Active)
Mode: Unicast-only
MAC Distribution: Enable
Group IP: NULL
State: UP                Mac learning: Enabled
```

Step 7. Create a VXLAN-backed port profile.

```
port-profile type vethernet TenantAFrontEndNetwork
  vmware port-group
  switchport mode access
  switchport access bridge-domain vxlan5000
  no shutdown
  state enabled
```

Step 8. Assign the virtual machine to the newly created port profile.

Choose Home > Inventory > Host and Clusters in VMware vCenter.

Select the virtual machine and right-click to set the properties. Assign the virtual machine to the newly created port profile, as shown in Figure 18.

Figure 18. Assigning the Virtual Machine to the Port Profile

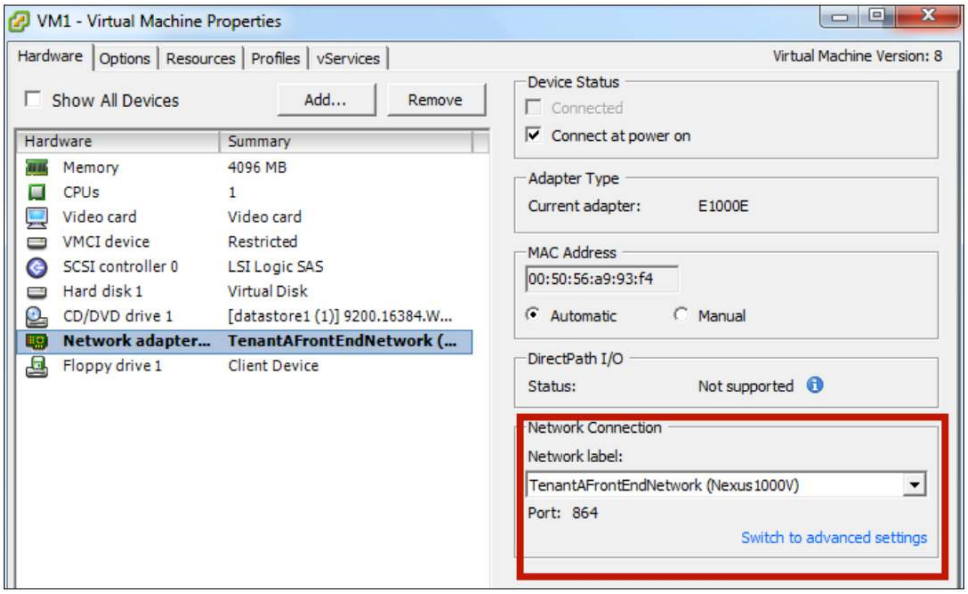

Verifying VXLAN

Enter the commands shown here at the command-line interface (CLI) to display the VXLAN status on the Cisco Nexus 1000V Switch for VMware vSphere.

Command: show bridge-domain

```
vsm-n1kv# show bridge-domain
Global Configuration:
Mode: Multicast/Multicastless
MAC Distribution: Enabled/Disabled
```

```
ARP
Bridge-domain BlueVXLAN (3 ports in all)
Segment ID: 8888 (Manual/Active)
Group IP: 0.0.0.0
Mode: Multicastless
MAC Distribution: Enabled
State: UP               Mac learning: Enabled
Veth7, Veth13, Veth16

Bridge-domain RedVXLAN (1 port in all)
Segment ID: 9999 (Manual/Active)
Group IP: 239.9.9.9
Mode: Multicast
State: UP               Mac learning: Enabled
Veth4
vsm-n1kv#
```

Command: show bridge-domain <<bridge_domain_name>> vteps

```
vsm-n1kv # show bridge-domain vteps
D: Designated VTEP     I:Forwarding Publish Incapable VTEP

Bridge-domain: mytestbd1
Ifindex     Module     VTEPs
---------   ---------- --------
Veth1       3          10.1.1.2(D),
Veth2       3          60.60.60.1 (DI),
Veth3       3          10.1.1.3,
```

Command: show bridge-domain mac

```
vsm-n1kv (config)# sh bridge-domain mac
Bridge-domain: BlueVXLAN
MAC Address      Mod    Port     VTEP-IP          Local-IPs
-------------------+-----+--------+------------+------------------------------+
0002.3d11.4102    3    VEth16   10.3.3.44        67.65.12.3, 1.1.1.1
0002.3d21.4100    3    VEth12   20.2.2.88        2.2.2.2, 3.3.3.3

Bridge-domain: RedVXLAN
MAC Address      Mod    Port     VTEP-IP          Local-IPs
-------------------+-----+--------+------------+------------------------------+
0002.3d11.4102    3    VEth16   10.3.3.44        67.65.12.3, 1.1.1.1
0002.3d21.4100    3    VEth12   20.2.2.88        2.2.2.2, 3.3.3.3
```

Conclusion

The Cisco VXLAN solution enables scalable cloud architecture with replicated server pods in different subnets. Because of the Layer 3 approach of User Datagram Protocol (UDP), virtual machine migration extends even to different subnets. The Cisco Nexus 1000V Switch for VMware vSphere with VXLAN support provides numerous advantages for customers, enabling customers to use LAN segments in a robust and customizable way without disrupting existing modes of operation.

For More Information

For more information about the Cisco Nexus 1000V Switch for VMware vSphere, please refer to the following links:

- Cisco Nexus 1000V Switch for VMware vSphere product information: http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html
- Cisco Nexus 1000V Switch for VMware vSphere technical documentation: http://www.cisco.com/en/US/products/ps9902/products_installation_and_configuration_guides_list.html
- Cisco Nexus 1000V Switch for VMware vSphere community: http://www.cisco.com/go/1000vcommunity

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C07-728863-00 07/13