Deploying Cisco Nexus 1000V Series Switches with VMware vCloud Director and VXLAN 1.0

Deployment Guide

February 2013

For further information, questions and comments please contact ccbu-pricing@cisco.com

Contents

Overview

Today's virtualized data center demands that multivendor solutions integrate and work together. VMware vCloud Director facilitates easier deployment of virtual machines to meet the scaling needs of a cloud-enabled data center. One of the main functions of VMware vCloud Director is to provide networking as a managed, allocated resource. VMware vCloud Director uses the advanced features of the Cisco Nexus® 1000V Series Switches to provide a scalable, highly secure, and agile cloud solution for private enterprises and service providers.

This document provides guidelines for deploying VMware vCloud Director with Cisco Nexus 1000V Series Switches using VLAN-backed, port-group-backed, and Virtual Extensible LAN (VXLAN)–backed network pools.

For detailed configuration documentation, please refer to the respective Cisco® product configuration guides found on http://www.cisco.com. You will find links to the product configuration guides and other related deployment guides in the "For More Information" section of this document.

Audience

This document is intended for network architects, network engineers, virtualization administrators, and server administrators interested in understanding and deploying Cisco Nexus 1000V Series with VMware vCloud Director.

Background

It is essential to understand some key elements of vCloud Director, including vCloud networking, before getting into the details of deploying vCloud Director with the Cisco Nexus 1000V Series.

Cisco Nexus 1000V Series Networking

The Cisco Nexus 1000V Series 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 Series Switches manage a data center as defined in VMware vCenter Server. Each server in the data center is represented as a line card in the Cisco Nexus 1000V Series Switch and can be managed as if it were a line card in a physical Cisco switch.

The Cisco Nexus 1000V Series implementation has two main components:

- Virtual Supervisor Module (VSM)
- Virtual Ethernet module (VEM)

Figure 1. Cisco Nexus 1000V Series Switches Managing VMware ESX Servers

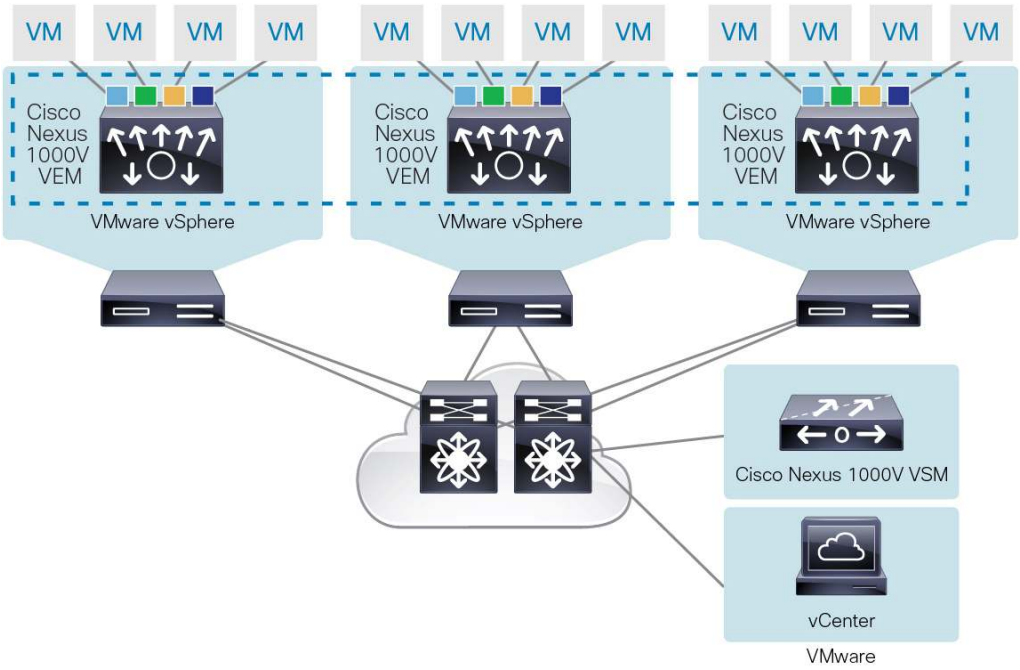

Benefits

The benefits of deploying Cisco Nexus 1000V working with VMware vCloud Director are:

- Advanced networking capabilities, such as quality of service, network statistics gathering with Cisco NetFlow Collector, packet mirroring with Cisco ERSPAN, and many others
- Nondisruptive operational model with Cisco Nexus 1000V fully integrated into vCloud Director and VMware vCenter Server
- Ease of troubleshooting due to the one-to-one physical network mapping of vCloud Director and the organization's application network
- Easier regulatory compliance of applications in the cloud since there is complete transparency in both the physical and virtual networks

VMware VMware vCloud Director

VMware vCloud Director provides the ability to build multitenant clouds by pooling virtual infrastructure resources into virtual data centers and exposing them to users through web based portals (Figure 2).

Figure 2. VMware vCloud Director

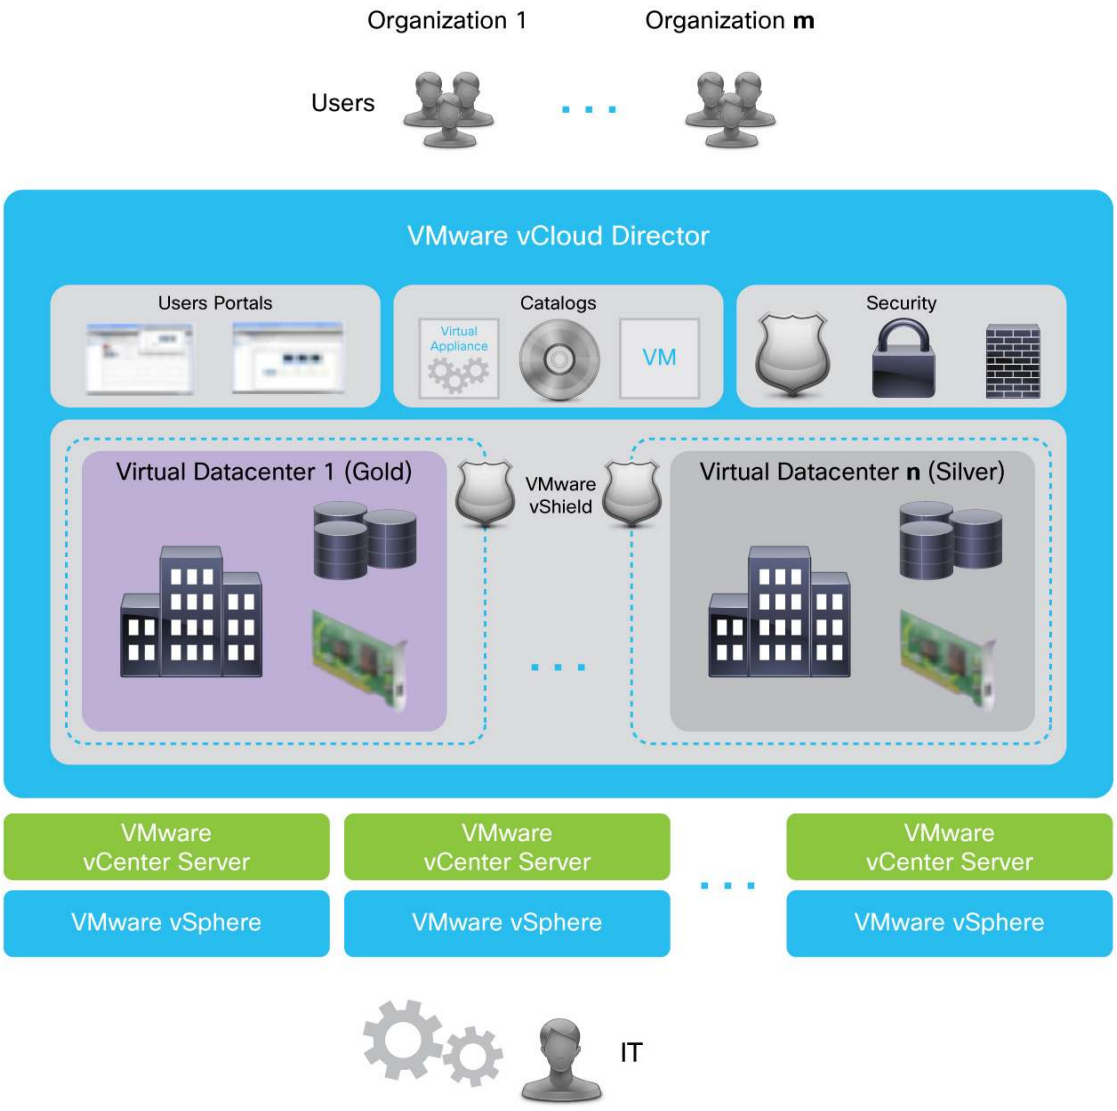

Organizations

vCloud Director supports multitenancy through the use of organizations. An organization is a unit of administration for a collection of users, groups, and computing resources. Users authenticate at the organization level, supplying credentials established by an organization administrator when the user was created or imported. vCloud Director administrators create and provision organizations, while organization administrators manage organization users, groups, and catalogs.

Provider Virtual Data Center (vDC)

A provider vDC combines the compute and memory resources of a single vCenter server resource pool with the storage resources of one or more data stores available to that resource pool.

You can create multiple provider vDCs for users in different geographic locations or business units, or for users with different performance requirements.

Organization Virtual Data Centers (OvDC)

An organization virtual data center (OvDC) provides resources to an organization and is partitioned from a provider vDC. Organization vDCs provide an environment where virtual systems can be stored, deployed, and operated. A single organization can have multiple organization vDCs.

VMware vCloud Director: Networking

VMware vCloud Director provides three classes of networks. The network class defines the boundaries and respective service level for each function within a given cloud's network architecture (Figure 3).

Figure 3. VMware vCloud Director Networking Options

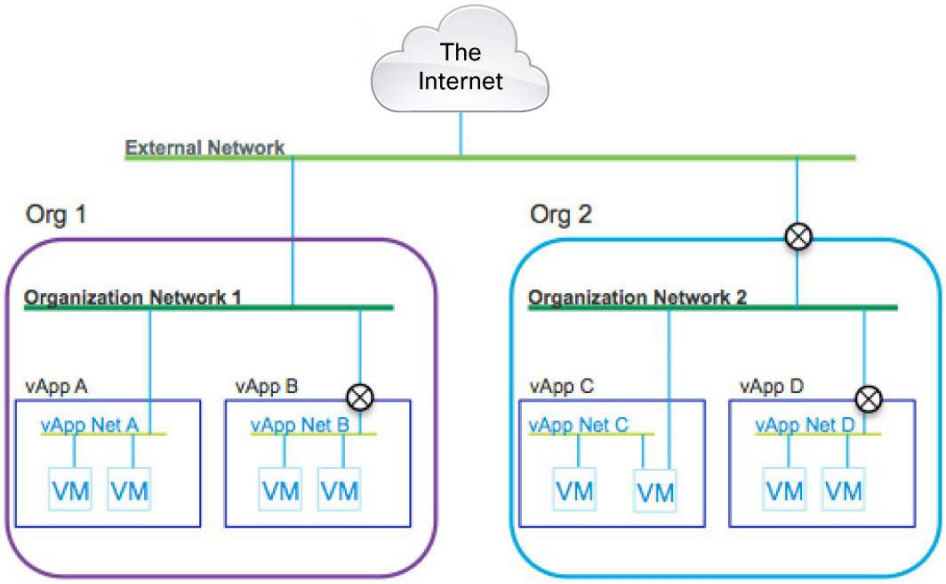

The network classes are external network, organization network, and vApp network.

External Networks

External networks provide transport between organizations or to networks outside a single-tenant network, such as the Internet. External networks are managed by the vCloud Director administrator and are not directly visible to a tenant organization. This network type is also sometimes called a provider or data center network.

Organization Networks

A network allocated to a single organization or tenant and backed by the managed allocation of network resources for that organization. A single organization may have many types of organization networks.

Direct Organization Networks

A direct organization network is “directly connected” to an external network (Figure 4). If a virtual machine is connected to a direct organization network, the VMNIC will be directly attached to the port-group of the external network.

Figure 4. Direct Organization Networks

Routed Organization Network

In routed organization networks, the network is connected to vShield Edge to provide connectivity to the external network. This is used when a vApp needs to be connected to a private network with limited access to the external network (Figure 5).

Figure 5. Routed Organization Networks

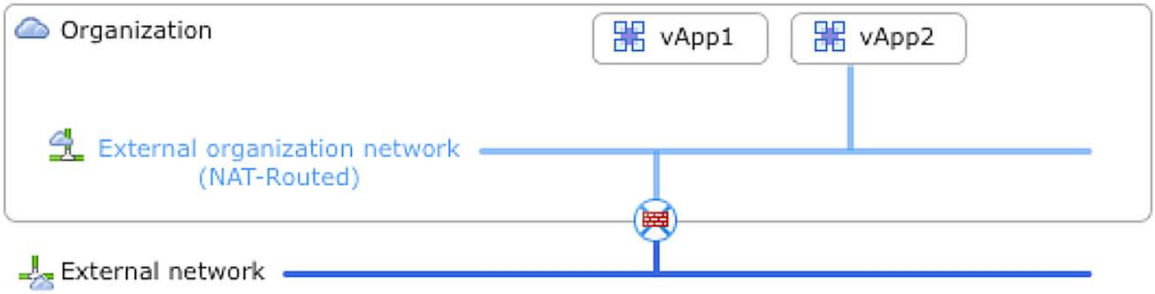

Isolated Organization Networks

In an isolated organization network, the network has no external connectivity. This is typically used for a vApp that does not require external connectivity (Figure 6).

Figure 6. Isolated Organization Networks

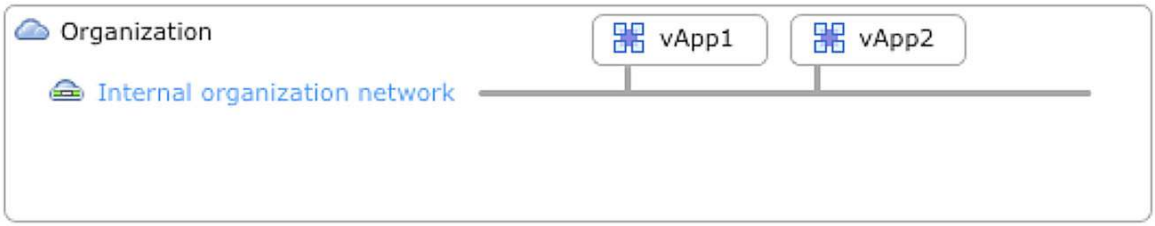

Organization networks provide network segments within a single tenant, and allow connectivity between vApps assigned to the same organization network. vApps that are on different organization networks, even within the same tenant organization, are not in the same broadcast domain.

The resources to create the isolation are managed by the vCloud administrator and are provided to organizations as a managed allocation, to allow the organization administrator to create isolated networks as needed.

vApp Networks

Like an organization network, a vApp network is a segment that is created for the particular application stack within the organization's network to enable multitier applications to communicate with each other, and at the same time to isolate the intra-vApp traffic from other applications within the organization.

All three classes of networks can be backed using the virtual networking features of the Cisco Nexus 1000V Series.

It is important to understand the relationship between the virtual networking constructs, features of the Cisco Nexus 1000V, and the classes of networks defined and implemented in a vCloud Director environment. Most often a network class (organization and vApp, specifically) is described as being backed by an allocation of isolated networks. In other words, in order for an organization administrator to create an isolated vApp network, the administrator must have a free isolation resource to consume and use to provide that isolated network for the vApp.

vCloud Director employs three different networks to create managed pools of isolation that can be allocated between and within tenant organizations. These three network pool types are vSphere port-groups, VLANs, and vCloud network isolation.

Port-Group Backed

A port-group pool type is a network pool created by statically allocating predefined network port-groups on Cisco Nexus 1000V.

VLAN Backed

A VLAN-backed pool type is a network pool created by allocating unused VLAN IDs, which are then dynamically allocated by vCloud Director to back a dynamically created network. When the network is deleted, the VLAN ID is released to the pool for reuse.

Both VMware vSphere port-group-backed network pools and VLAN-backed network pools rely on the VLAN construct to isolate the traffic on the physical segment; the difference is the mechanism by which the port groups are created and associated with a VLAN ID. For port-group-backed network pools, the port groups are created as shown later in this guide, using the Cisco configuration interface (see "Configuring VLAN-Based Isolation on the Cisco Nexus 1000V Series," step 2). The VLAN-backed pool is the mechanism by which both the port groups and the requisite VLANs are created by VMware vCloud Director, by provisioning the same VLANs and port groups on the VMware distributed network switching platform.

vCloud Network Isolation Backed

A vCloud network-isolation-backed network pool provides isolated Layer 2 networks for multiple tenants of a cloud without consuming the VLAN IDs. This isolation-backed network pool does not require pre-existing VLAN IDs in vSphere. It uses port-groups that are dynamically created. A cloud isolated network spans hosts, provides traffic isolation from other networks, and is the best source for vApp networks.

When this option is selected with Nexus 1000V, the isolation technology is VXLAN.

Note: This type of network pool is not supported for VMware vCloud Director 5.1 and later. VMware vCloud Director 5.1 supports VXLAN-backed network pools instead, which are automatically created when provider vDCs are provisioned in VMware vCloud Director. For more information about network-pool compatibility, see Cisco Nexus 1000V and VMware Compatibility Information).

VLAN-Backed Network Pools

VLAN-backed networks in VMware vCloud Director are supported by Cisco Nexus 1000V Series Switches starting with Release 4.2(1)SV1(5a). For more information about network pool compatibility, see Cisco Nexus 1000V and VMware Compatibility Information. This document does not present the details of this type of network deployment. This support is available through the integration of Cisco Nexus Network Segmentation Manager (NSM) with VMware vShield Manager. When selecting a range of VLANs from VMware vCloud Director, make sure that the range is not part of the infrastructure VLANs.

In Figure 7, VLAN range 100 to 200 is used for organization networks. Make sure that these VLANs are not part of the infrastructure VLANs: that is, not part of Cisco Nexus 1000V Series VLANS for control, packet, management, VMware vSphere management console, VMware vMotion, storage, fault tolerance, VXLAN transport, and so on.

Figure 7. VLAN-Backed Network Pool Using Cisco Nexus 1000V Series Switches

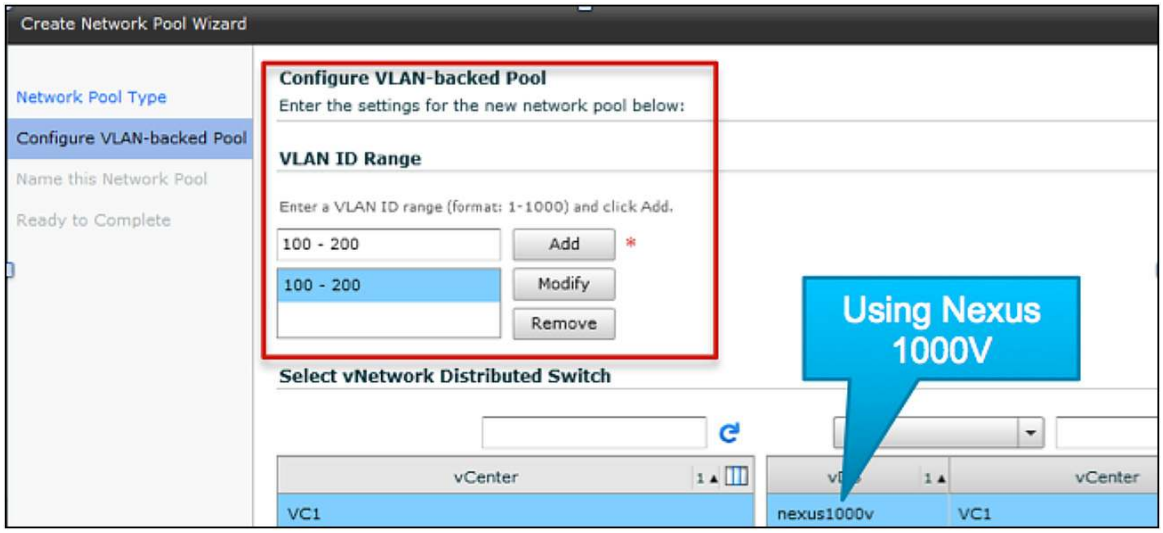

Port-Group-Backed Network Pools

The next example shows how to use port-group-backed network pools with the Cisco Nexus 1000V Series. Each port group will be isolated in its own VLAN ID. The configuration example in Figure 8 shows how to configure and use the three classes of networks for a vApp.

Figure 8. VMware vCloud Director and Cisco Nexus 1000V Series Networking

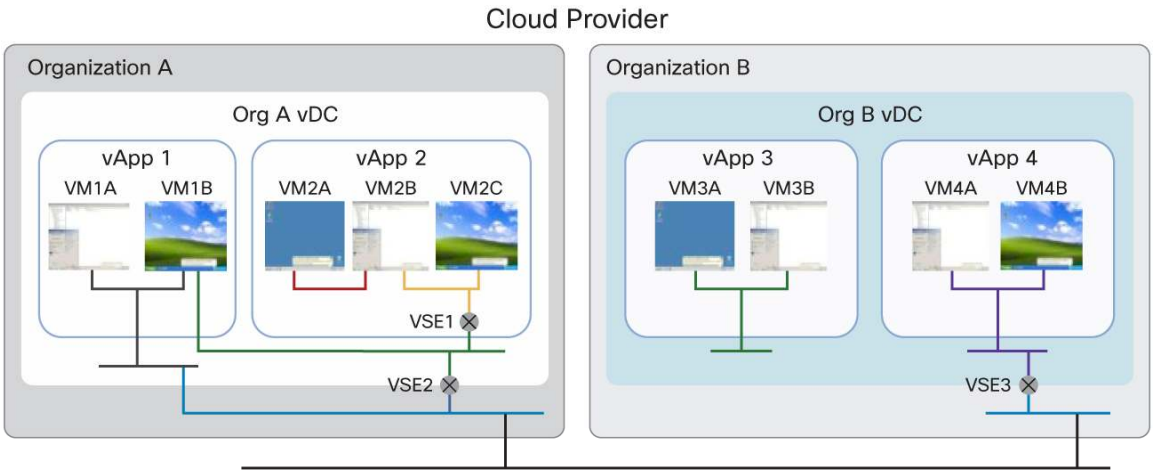

Network Type	Label	Cisco Nexus 1000V Port-Profile
vApp Network		N1KV_vApp_VLAN301 N1KV_vApp_VLAN300
Organization Directly Connected External Network		Connected to N1KV_Provider_Ext
Organization Routed Network		N1KV_Org_VLAN200, N1KV_Org_VLAN201
Provider External Network		N1KV_Provider_VLAN170

The configuration examples in this document show how to create an external network and then an organization vDC and vApp with internal, routed, and directly connected networks. You can use Figure 9 as a reference to see how the Cisco Nexus 1000V Series can be used to create all the available network classes and types in VMware vCloud Director.

Figure 9. Cisco Nexus 1000V Series Port Profiles for Network Pools Example

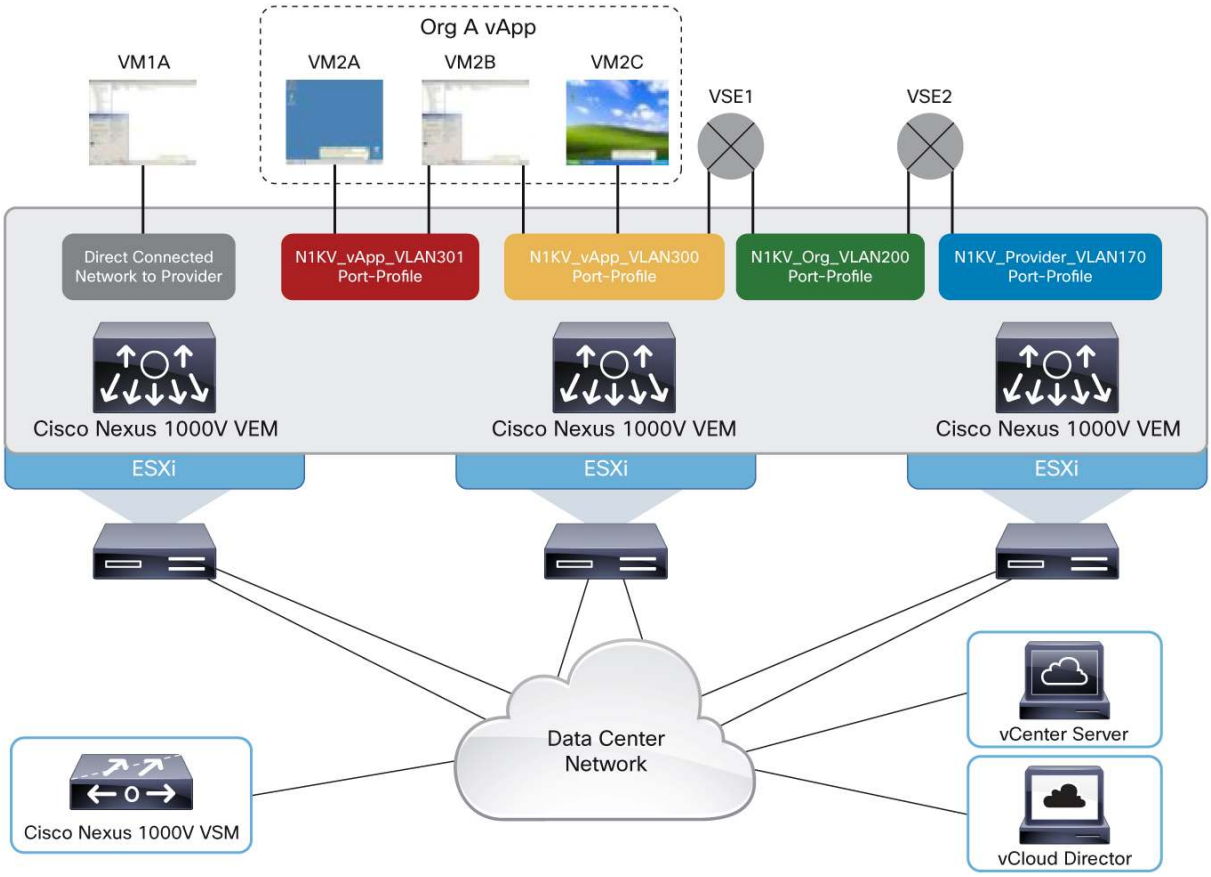

The prerequisites for the examples are as follows:

- All VMware vSphere components have been deployed, including:
  - At least one VMware vCenter Server
  - Two or more hosts running VMware ESX or ESXi 4.0 or later
- The Cisco Nexus 1000V Series VSM is installed and functioning.
- The Cisco Nexus 1000V Series VEM is installed on the VMware ESX and ESXi hosts that are part of VMware vCloud Director.
- The VMware vCloud Director cells and database have been completely installed.
- The VMware vCloud Director provider vDC and organizations have been defined.

Configuring VLAN-Based Isolation on the Cisco Nexus 1000V Series

Because port profiles are represented as port groups in VMware vSphere, they can be used to back VMware vCloud Director network pools. Network pools are used to create all the network types for VMware vCloud Director. Each of the port profiles used should have a unique VLAN assigned to it because each VMware vSphere port group needs to be isolated to Layer 2. This approach helps ensure that each network is isolated based on its VLAN and also provides the capability to offer many of the benefits of the Cisco Nexus 1000V Series (such as security, access control lists [ACLs], encapsulated remote switch port analyzer [ERSPAN], quality of service [QoS], and Internet Group Management Protocol [IGMP] Snooping).

The first step is to provision VLANs on the Cisco Nexus 1000V Series that will be used for the VMware vCloud Director deployment. You should use a meaningful convention to name and describe the VLANs.

For example, the sample configuration here uses the following conventions and ranges:

- VLAN 1–199: External provider networks and infrastructure VLANs
- VLAN 200–299: Organization networks and VLANs
- VLAN 300–399: Internal networks for vApps

Step 1. Define the proper VLANs on the VSM, as in this example:

```
vlan 170
   name Provider_Infra_VLAN170
vlan 200
   name Org_VLAN200
vlan 201
   name Org_VLAN201
vlan 300
   name vApp_VLAN300
vlan 301
   name vApp_VLAN301
vlan 302
   name vApp_VLAN302
```

Step 2. Create port profiles and assign a unique VLAN to each. Other features should also be configured here. It is recommended that you use ephemeral port binding if you are using Cisco Nexus 1000V Series Software Release 4.2(1)SV1(4) or later. Here is an example:

```
port-profile type vethernet N1KV_Provider_VLAN170
   vmware port-group port-binding ephemeral
   switchport mode access
   switchport access vlan 170
   no shutdown
   state enabled
port-profile type vethernet N1KV_Org_VLAN200
   vmware port-group port-binding ephemeral
   switchport mode access
   switchport access vlan 200
   no shutdown
   state enabled
port-profile type vethernet N1KV_Org_VLAN201
   vmware port-group port-binding ephemeral
   switchport mode access
   switchport access vlan 201
   no shutdown
   state enabled
port-profile type vethernet N1KV_vApp_VLAN300
   vmware port-group port-binding ephemeral
   switchport mode access
```

```
    switchport access vlan 300
    no shutdown
    state enabled
port-profile type vethernet N1KV_vApp_VLAN301
    vmware port-group port-binding ephemeral
    switchport mode access
    switchport access vlan 301
    no shutdown
    state enabled
port-profile type vethernet N1KV_vApp_VLAN302
    vmware port-group port-binding ephemeral
    switchport mode access
    switchport access vlan 302
    no shutdown
    state enabled
```

Step 3. Allow the appropriate VLANs on the Ethernet uplinks on the VMware ESX and ESXi hosts that are part of the VMware vCloud Director deployment, as in this example:

```
port-profile type ethernet uplink
    vmware port-group
    switchport mode trunk
    switchport trunk allowed vlan 1-399
    channel-group auto mode on mac-pinning
    no shutdown
    system vlan 170, 254-255
    state enabled
```

Note that the same VLANs need to be defined and trunked on the upstream physical switches for VMware vCloud Director networks to work across multiple physical hosts.

Deployment Example: Using Port-Group-Based Network Pools with the Cisco Nexus 1000V Series

The Cisco Nexus 1000V Series can be used for all network types in VMware vCloud Director. Follow the steps described here to start using the networks and features of the Cisco Nexus 1000V Series.

Note that at least one provider vDC and one organization must have been created before you begin these steps.

Step 1. Identify and verify the VMware vCenter Server (already part of VMware vCloud Director) that will be used.

Step 2. Create an external provider network in the VMware vCloud Director interface.

You can do this by selecting the Manage & Monitor tab, choosing External Networks > Add Network, and selecting the appropriate provider network port group. In this case, it is N1KV_Provider_VLAN170 (Figure 10). By definition, this network should have connectivity to other external networks (routed), such as connectivity to the Internet.

Figure 10. Creating an External Provider Network in VMware vCloud Director

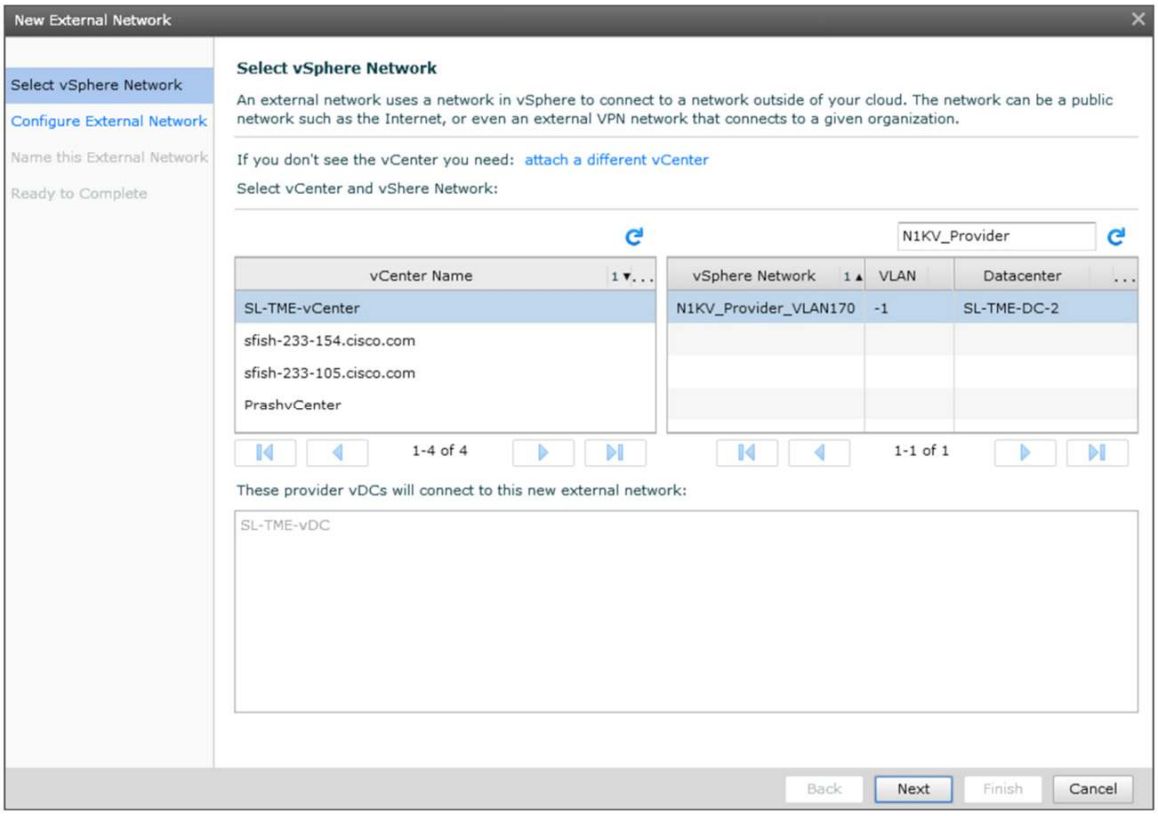

Click Next and complete the New External Network dialog box by specifying the subnet mask, gateway, DNS, and IP pool of addresses on this subnet. These IP addresses will later be dynamically assigned for external communications for directly connected and routed network connectivity for virtual machines and vApps. After successful completion of this process, the newly created external network will appear in the list of external networks (Figure 11).

Figure 11. New External Network

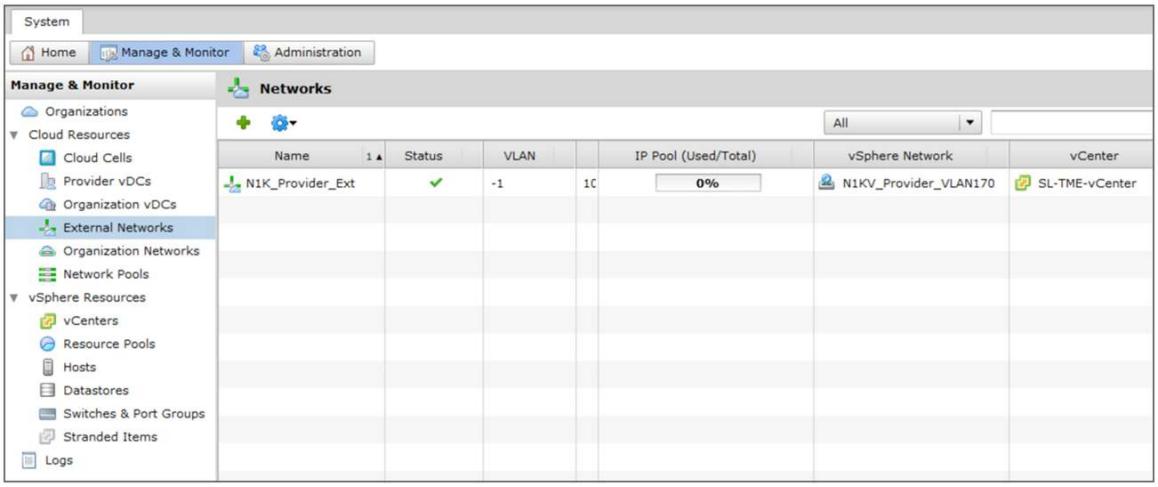

At this point, you have successfully created the first of three VMware vCloud Director network classes, an external network, using a Cisco Nexus 1000V Series virtual switch. In the steps that follow, you will create connections from isolated, organization-specific networks (which will also use Cisco Nexus virtual switching) to this external network.

Step 3a. Create network pools backed by the Cisco Nexus 1000V Series port profiles, which were defined in step 2 of the section “Configuring VLAN-Based Isolation on the Cisco Nexus 1000V Series.”

To create the pools, select the Manage & Monitor tab, choose Network Pools > Add Network Pool, and select the “vSphere Port Group-backed” option (Figure 12).

Figure 12. Creating a Network Pool

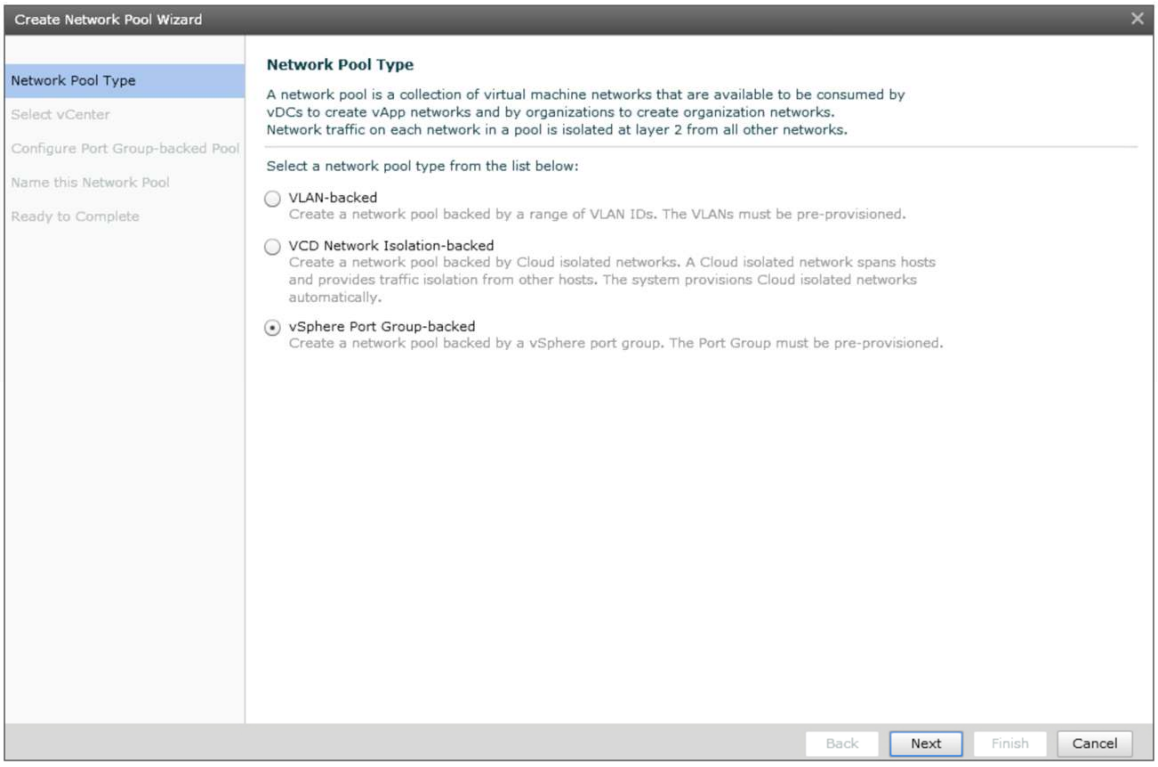

Step 3b. After selecting the appropriate VMware vCenter server instance, click Next and select the appropriate port groups. In this case, the vApp port profiles that were previously configured are used (Figure 13).

Figure 13. Selecting Port Groups

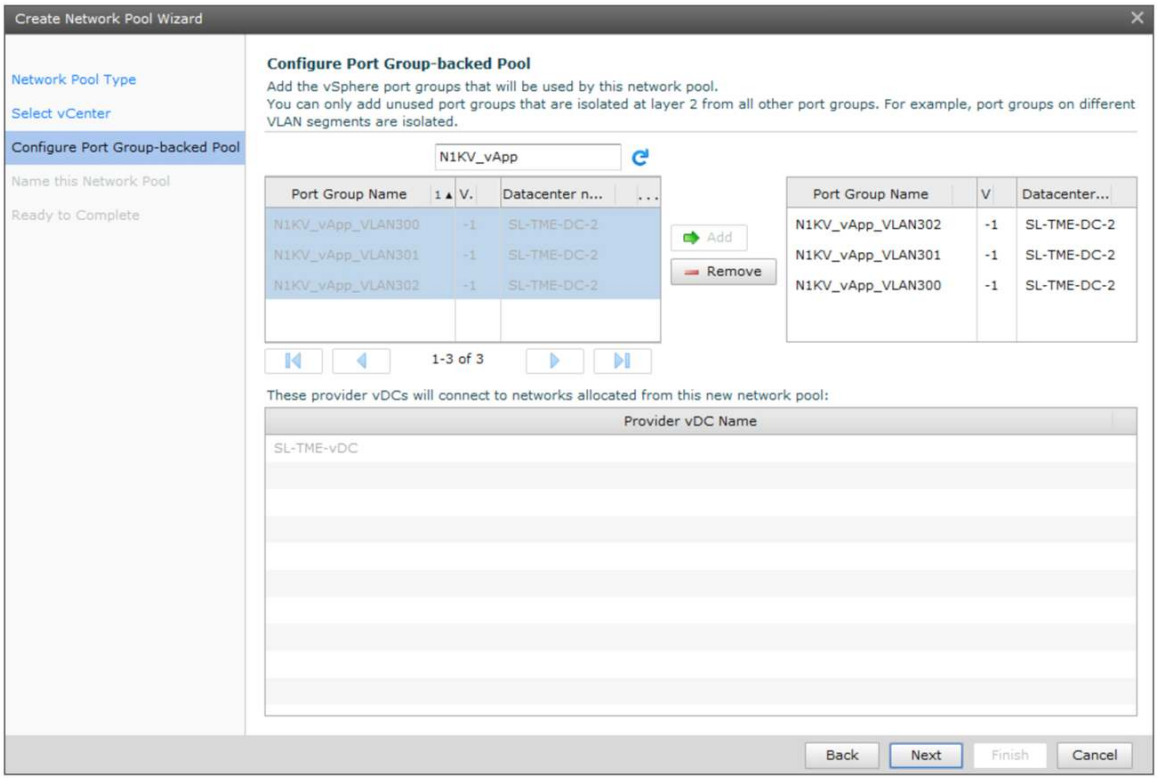

Step 3c. Complete the network pool by naming the pool, reviewing the settings, and clicking Finish (Figure 14).

Figure 14. Completing the Network Pool

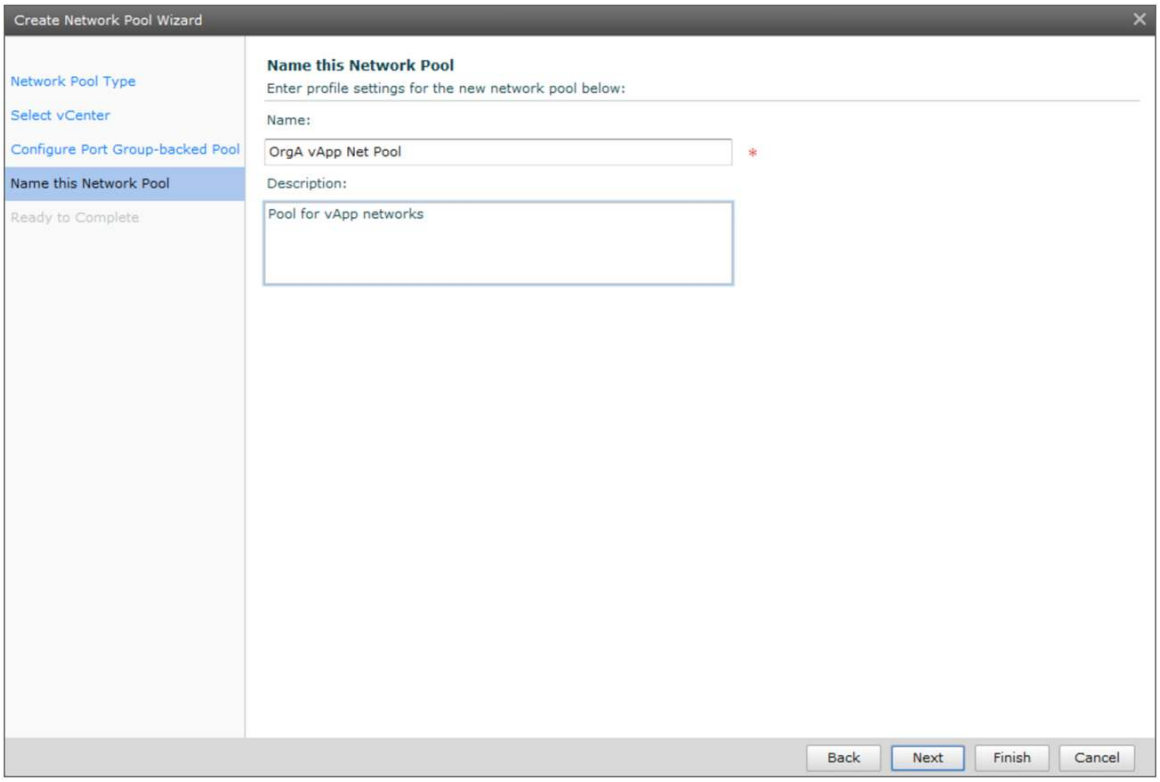

For the examples that follow, step 3 has been repeated to create an additional pool for the organization networks. The organization network pool will be used to create internal, routed, and directly connected networks.

Step 4a. Create an organization vDC and associate it with the external provider network defined in step 2. In this example, a previously configured organization, “Org_A,” is used.

Create the organization vDC by selecting the Manage & Monitor tab and choosing Organization vDCs > Add vDC.

Select the organization (Figure 15).

Figure 15. Selecting the Organization

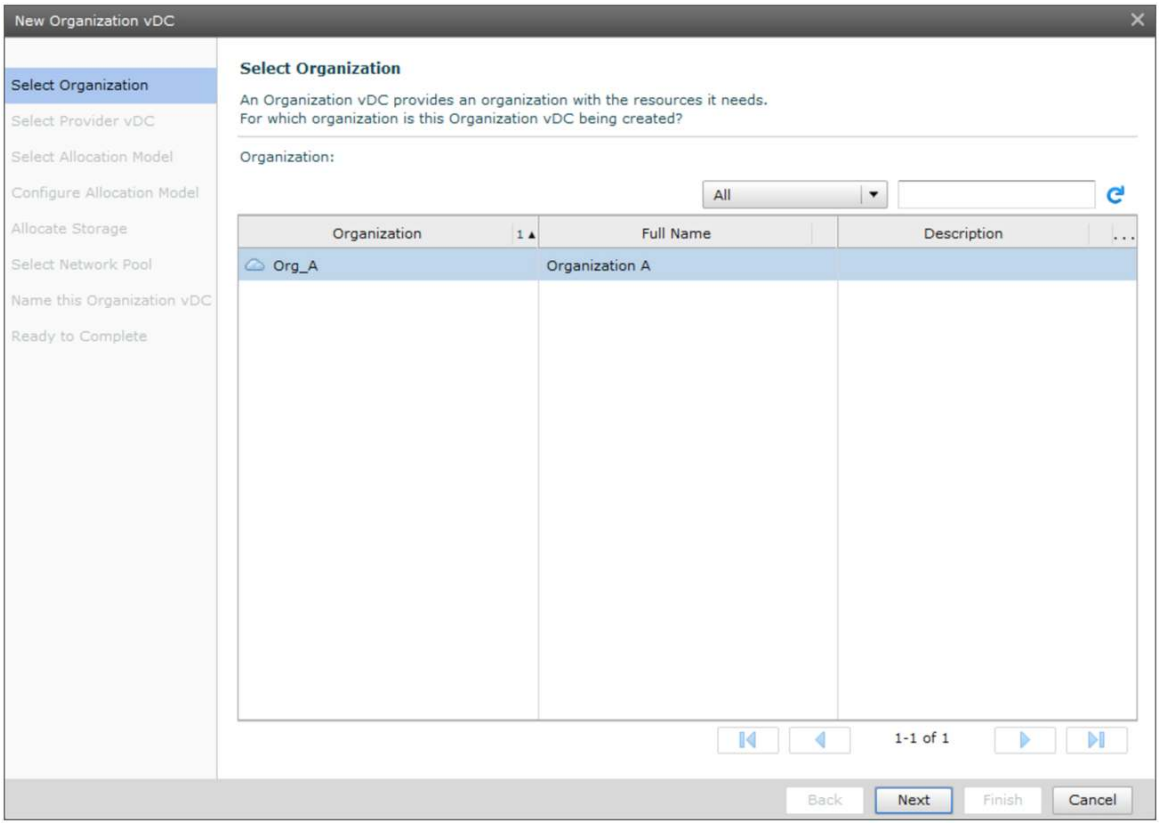

Step 4b. Select the provider vDC and the external network defined in step 2 (Figure 16).

Figure 16. Selecting the Provider vDC

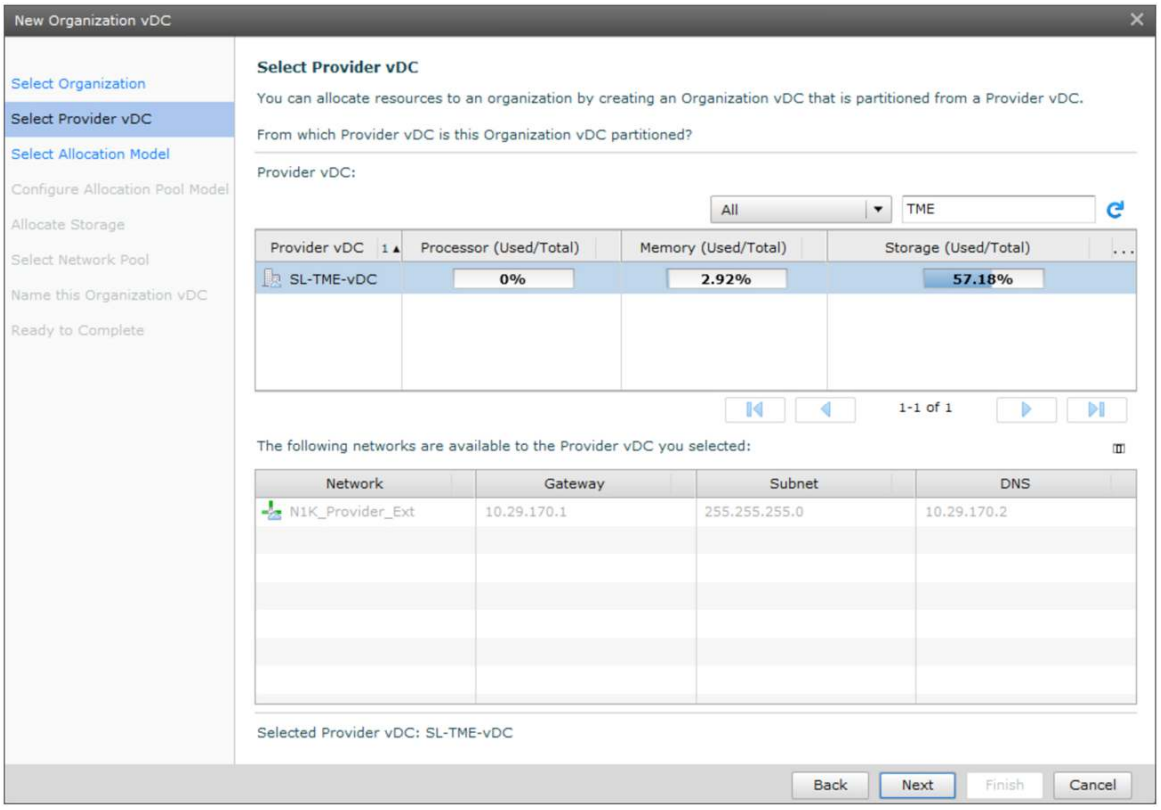

Step 4c. Select the allocation model. Click Next, select Configure Allocation Model, and click Next (screen not shown, for brevity).

Step 4d. Allocate the necessary storage for the organization and click Next (screen not shown, for brevity).

Step 4e. Select the network pool for the organization, which was created during step 3, and click Next (Figure 17).

Figure 17. Selecting the Network Pool for the Organization

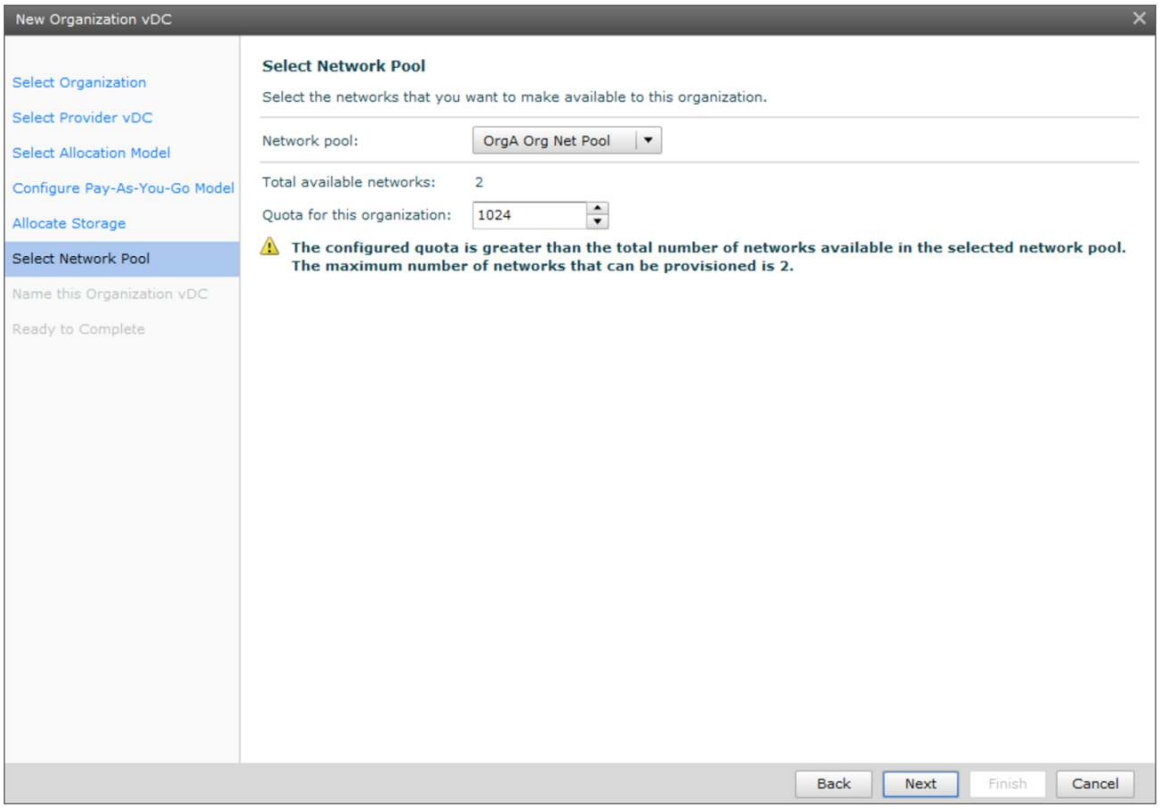

Step 4f. Name the organization vDC, click Next, and complete the new organization vDC (Figure 18).

Figure 18. Completing the Organization vDC

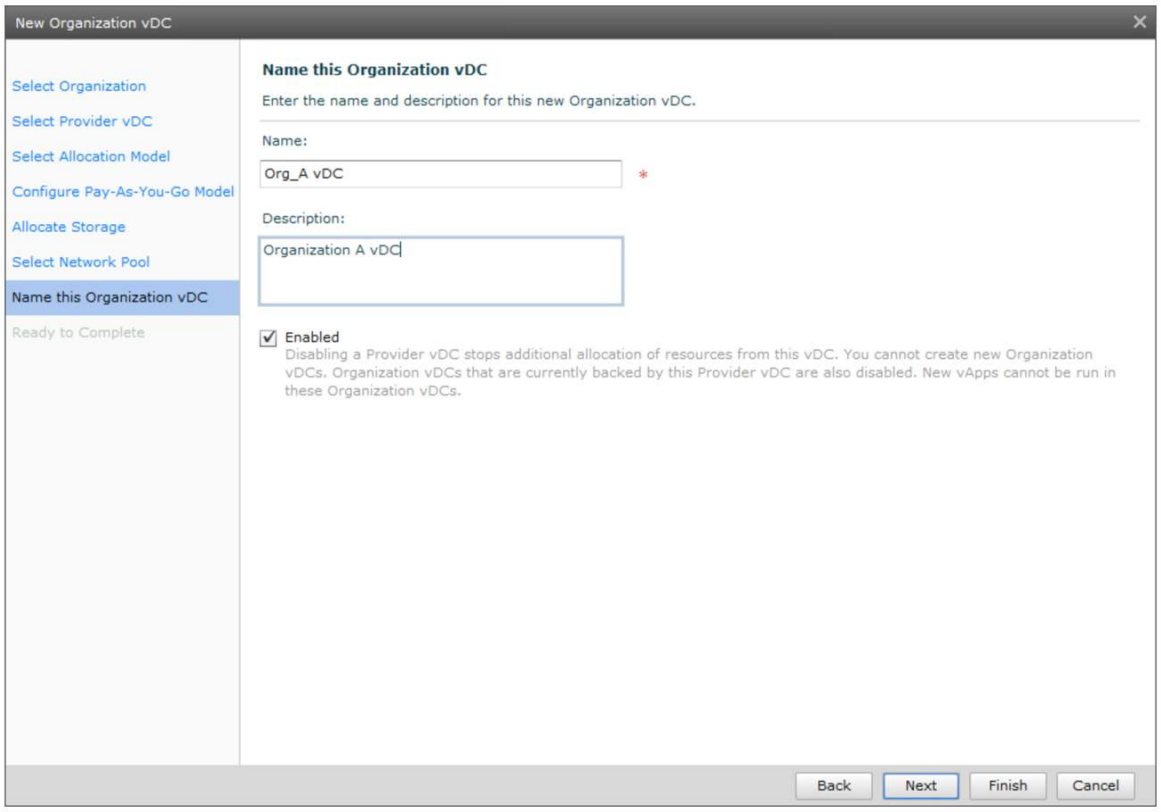

At this point, the organization has been created. It has been given a network pool backed by VMware vSphere port groups, as selected in step 3a and defined in the Cisco VSM in step 2 of the earlier section, “Configuring VLAN-Based Isolation on the Cisco Nexus 1000V Series.”

The steps that follow demonstrate how to create organization networks using this network pool allocation with three different types of connectivity. As you consider these examples, it is critical to understand the different types of connectivity that VMware vCloud Director provides for organization networks. An organization network can be isolated from all other networks: that is, no traffic can leave or enter this broadcast domain, and all connectivity is local to the vApps and vApp networks that are connected to the network. This type of organization network is called an internal network and will be created in step 5b that follows. At the same time that the examples create this internal network, a second organization network will be created: an external organization network. Step 5b also creates an external organization network with a Network Address Translation (NAT) routed profile. This organization network differs from the aforementioned internal network in that traffic can leave and enter this broadcast domain through a NAT routed connection. This routing and NAT service is created and managed by VMware vCloud Director as a VMware vShield Edge appliance and will be discussed later. Finally, after two organization networks are created simultaneously, step 6 repeats the process of creating an organization network but uses a third connectivity profile: direct connection. Whereas the internal network is completely isolated from the external network, and the external organization network with a routed connection is isolated through a NAT gateway, the external organization with a direct connection is not isolated in any way. This organization network and the external network are in the same broadcast domain.

Keep these concepts in mind when considering the examples that follow showing how to create these various organization networks.

Step 5a. Create both an internal organization network and a routed external network for the organization. Both of these networks can be added in one step using the Create Organization Network Wizard.

To begin, select the Manage & Monitor tab and choose Organization Networks > Add Network. Select the previously configured organization and then click Next (Figure 19).

Figure 19. Selecting the Organization

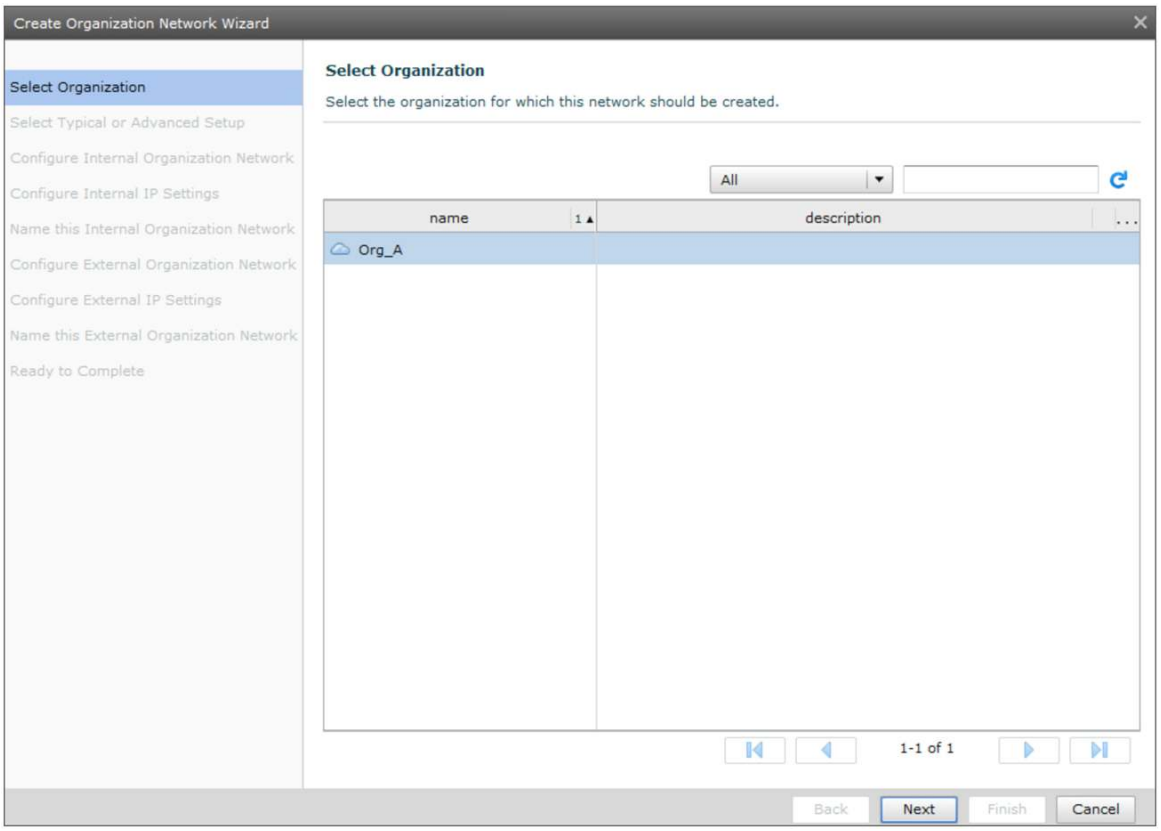

Step 5b. Using the Typical network setup, select both the “Create internal network” and “Create an external network via” options. From the pull-down menu, choose “routed connection” and click Next (Figure 20).

Figure 20. Specifying the Setup Options

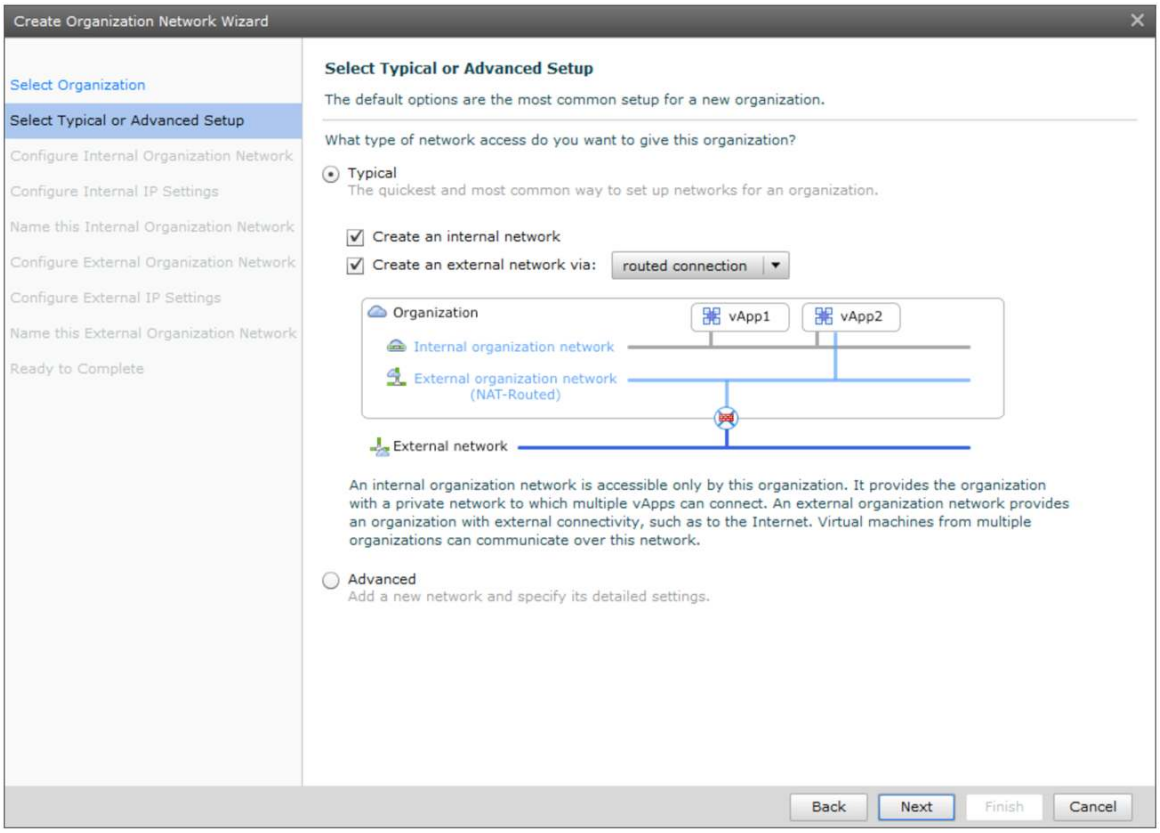

Selecting the “Create an internal network” option will instruct VMware vCloud Director to create a network to be used for vApp and virtual machine communication within the organization. This broadcast domain is not visible to other tenant organizations. Selecting “Create an external network via” with “routed connection” chosen from the pull-down menu will form a network that is routed to an external provider network but is secured from that segment by a NAT service; the external organization network will be backed, or isolated, using a resource from a network pool. An external provider network will need to be associated with this fenced network. The routing is provided by the VMware vShield Edge (vSE) virtual machines (with two virtual interfaces) included with VMware vCloud Director. Therefore, traffic between any vApp or virtual machine connected to the external organization network and the external provider network will flow through the VMware vSE virtual machine on its internal interface and be routed out to the external network through its external-facing interface.

Step 5c. Choose the internal organization network pool that will be used and click Next. This pool can be used by vApps to communicate within the organization, for example. In this case, it is the OrgA vApp Net Pool (Figure 21).

Figure 21. Selecting the Internal Organization Network Pool

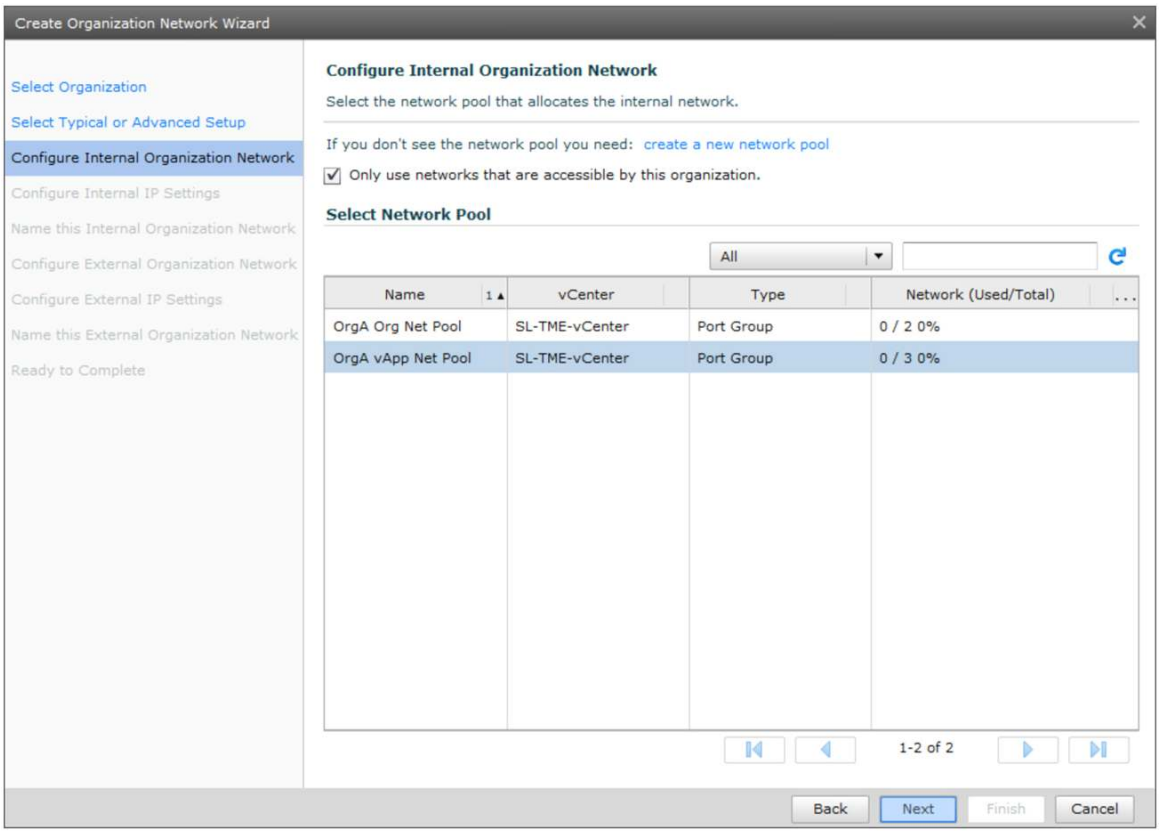

Step 5d. Configure the IP settings, including the range of IP addresses that vApps can use on this internal network, and click Next (Figure 22).

Figure 22. Configuring the IP Settings

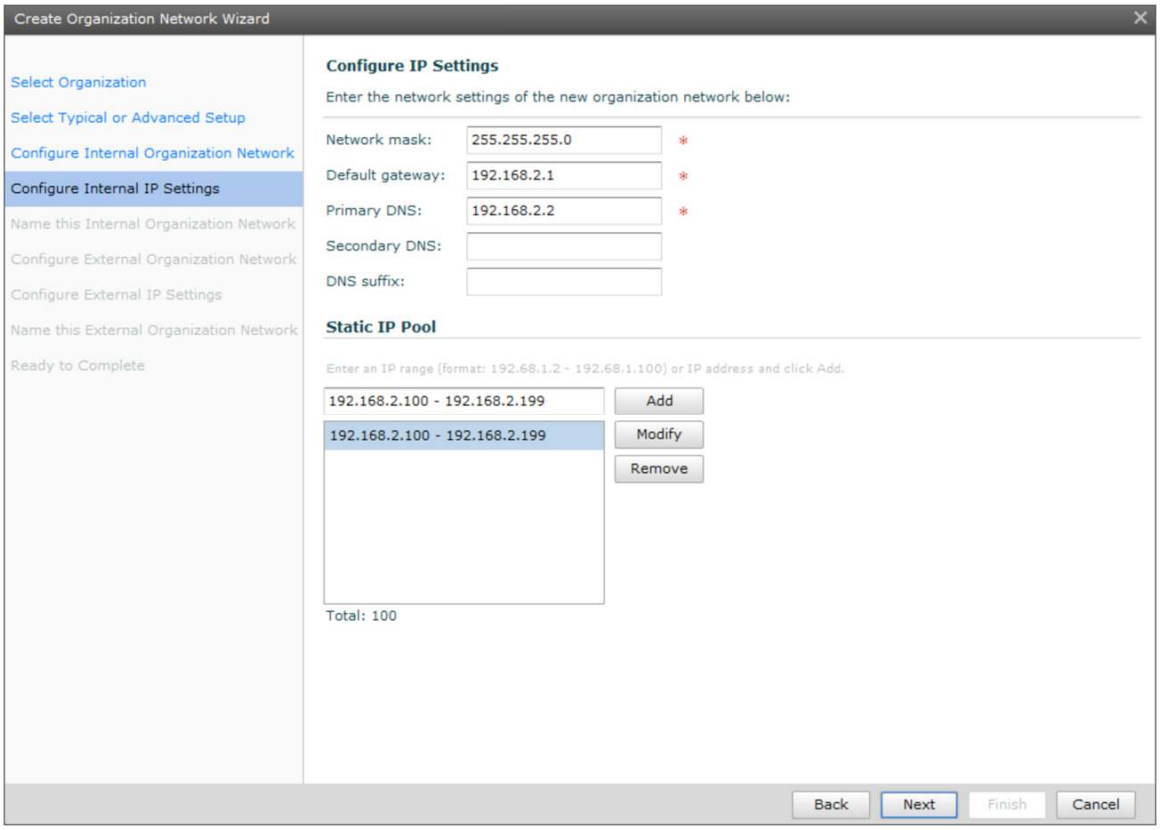

Step 5e. Name this internal network; then click Next (Figure 23).

Figure 23. Naming the Internal Network

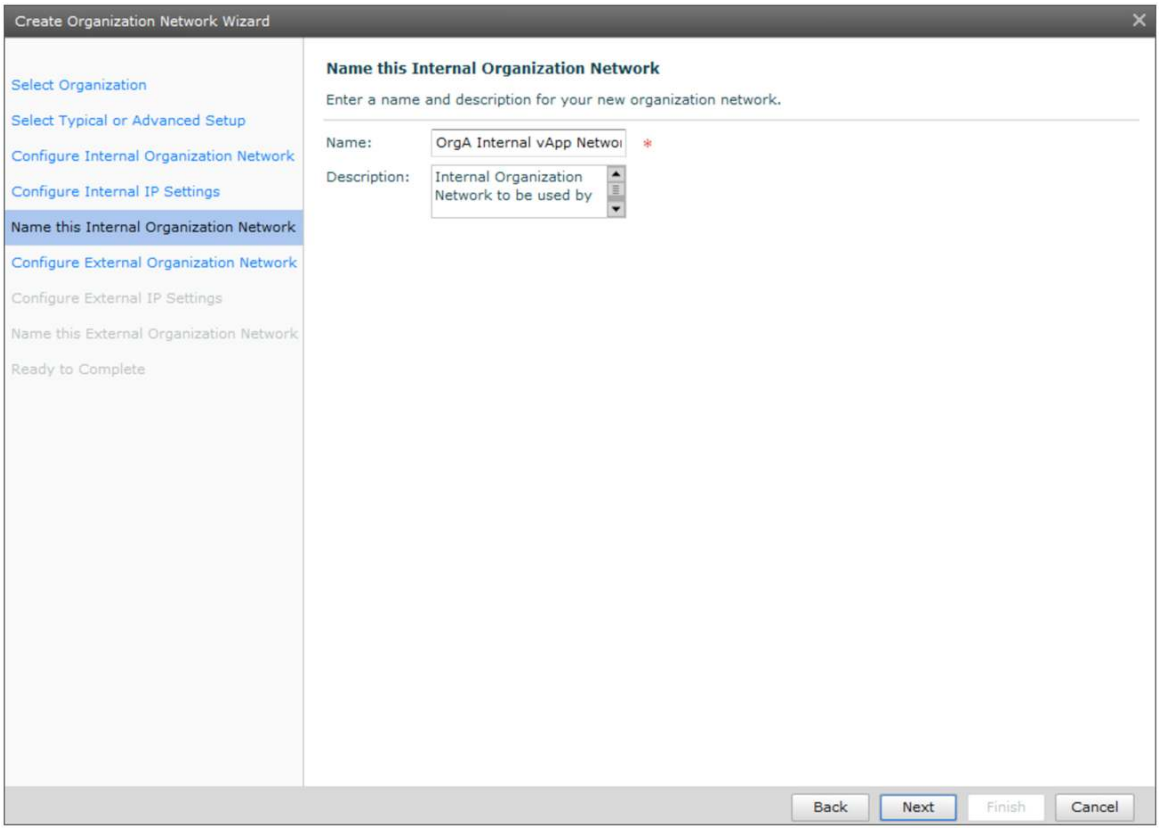

Step 5f. Configure the external organization network. The first step is to specify the external network that outbound traffic will traverse; this provider network will be used for communication outside the organization's network resources, such as to the Internet. The VMware vSE appliance will automatically be deployed with one network interface connected to this network; the VMware vSE appliance virtual machine will route, and apply the NAT service for, all communication between the external organization network being created in this step and the external network specified here. The other interface of the VMware vSE virtual machine will be connected to the external organization network when the operation is complete. Because the traffic on this external organization network is isolated from the other organization networks, select the network pool that will back the isolation (Figure 24).

Figure 24. Configuring the External Organization Network

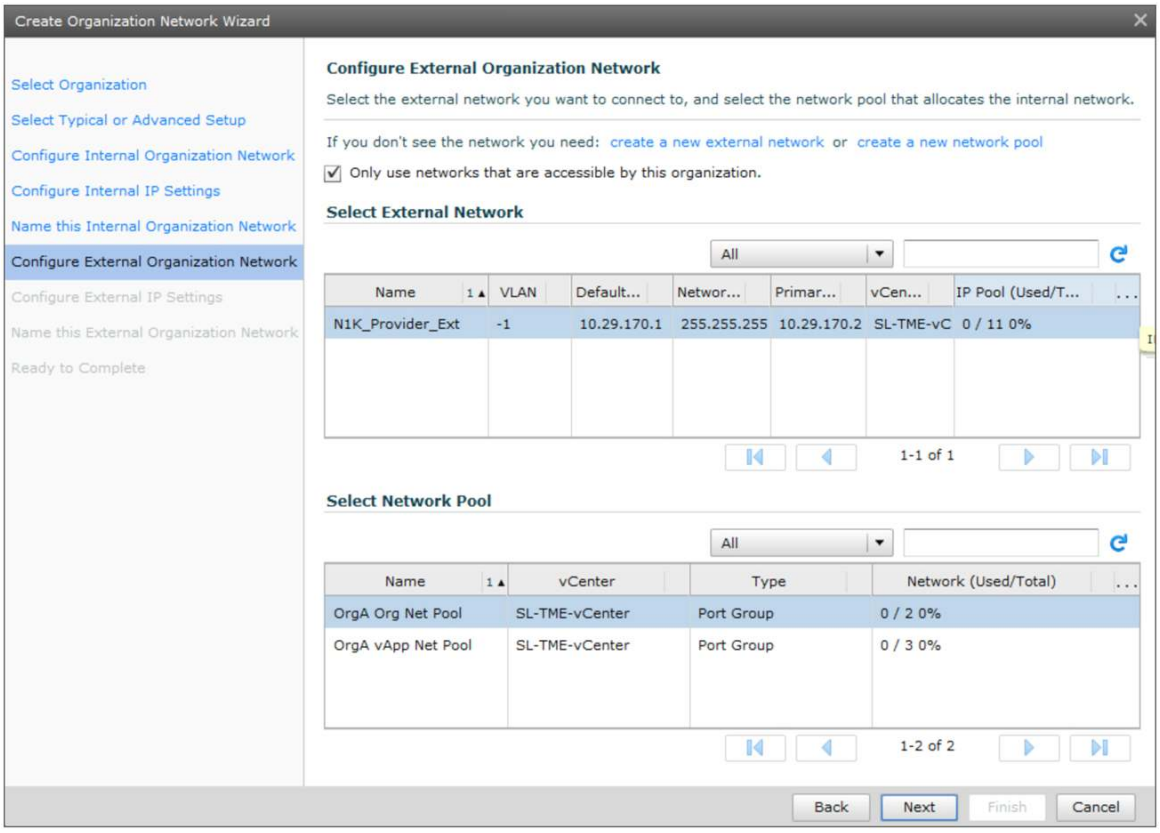

Step 5g. Configure the IP settings for the external organization network; then click Next (Figure 25). These addresses will be dispensed by the VMware vSE appliance virtual machine to clients within the organization connecting to this external organization network. When a vApp is deployed on this segment, VMware guest customization processes can inject these addresses into the virtual machine.

Figure 25. Configuring the IP Settings for the External Organization Network

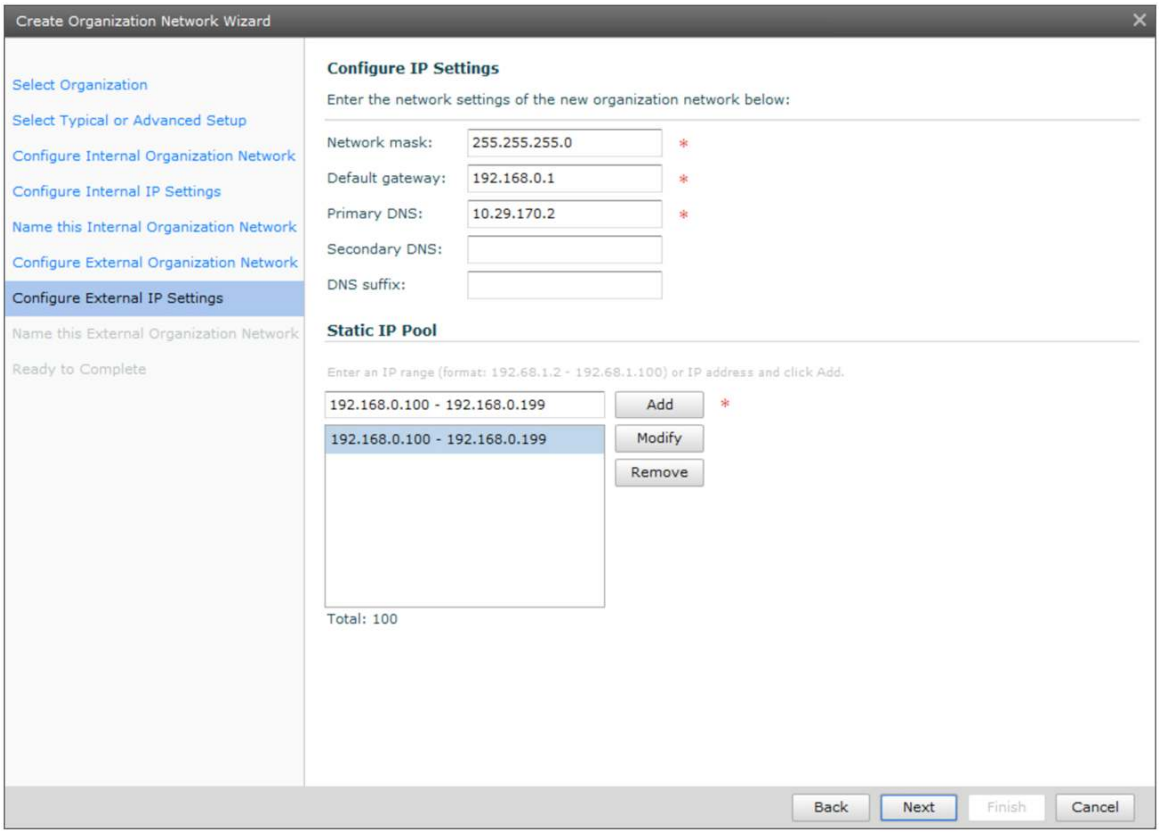

Step 5h. Name this external organization network and click Next to complete the process (Figure 26).

Figure 26. Naming the External Organization Network

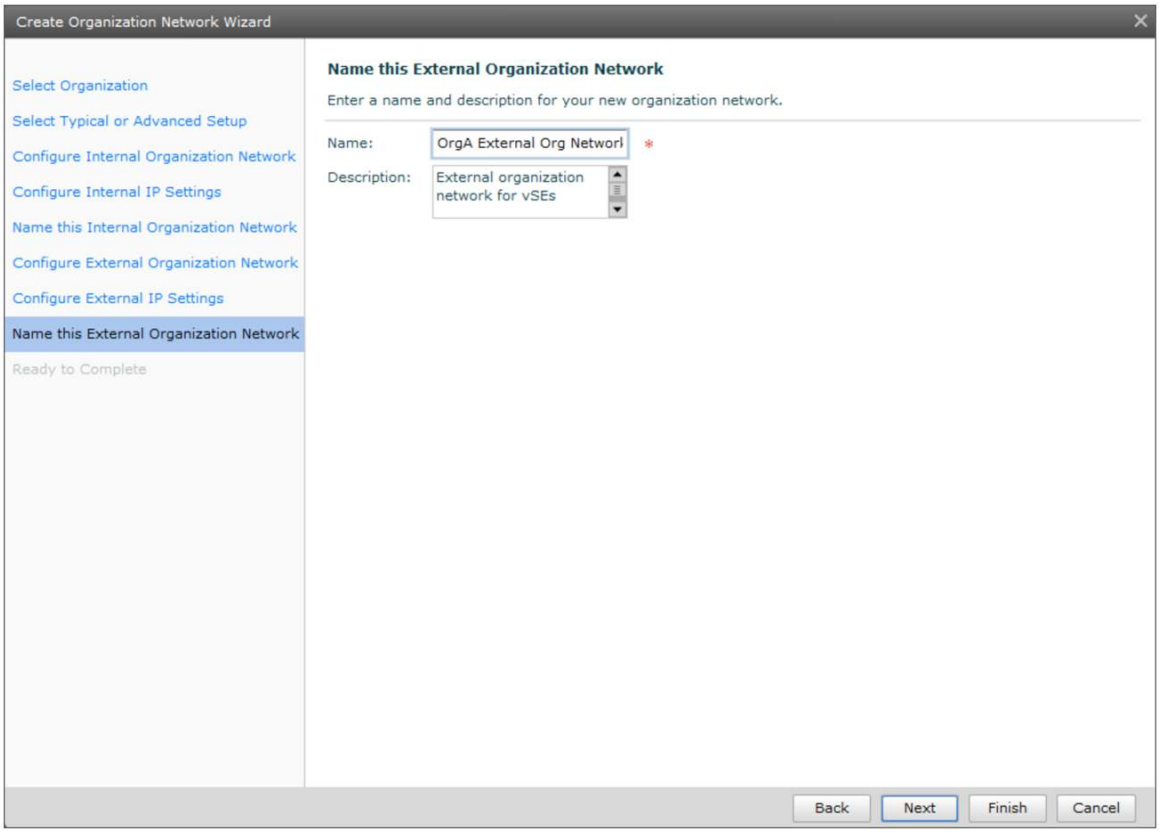

Step 5i. Review the final settings and click Finish (Figure 27).

Figure 27. Completing the Process

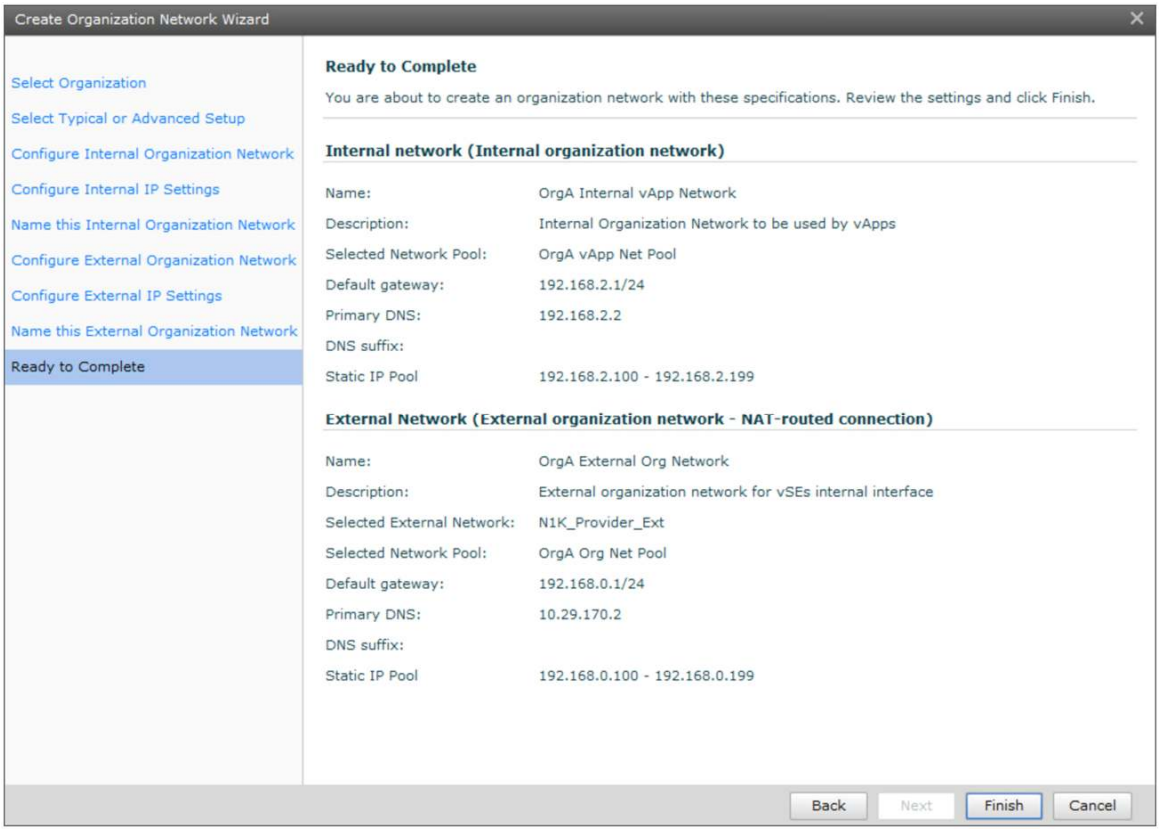

Step 6a. Create an external direct connected network for the organization. To begin, select the Manage & Monitor tab and choose Organization Networks > Add Network.

Select the organization that this network will belong to and then click Next (Figure 28).

Figure 28. Selecting the Organization

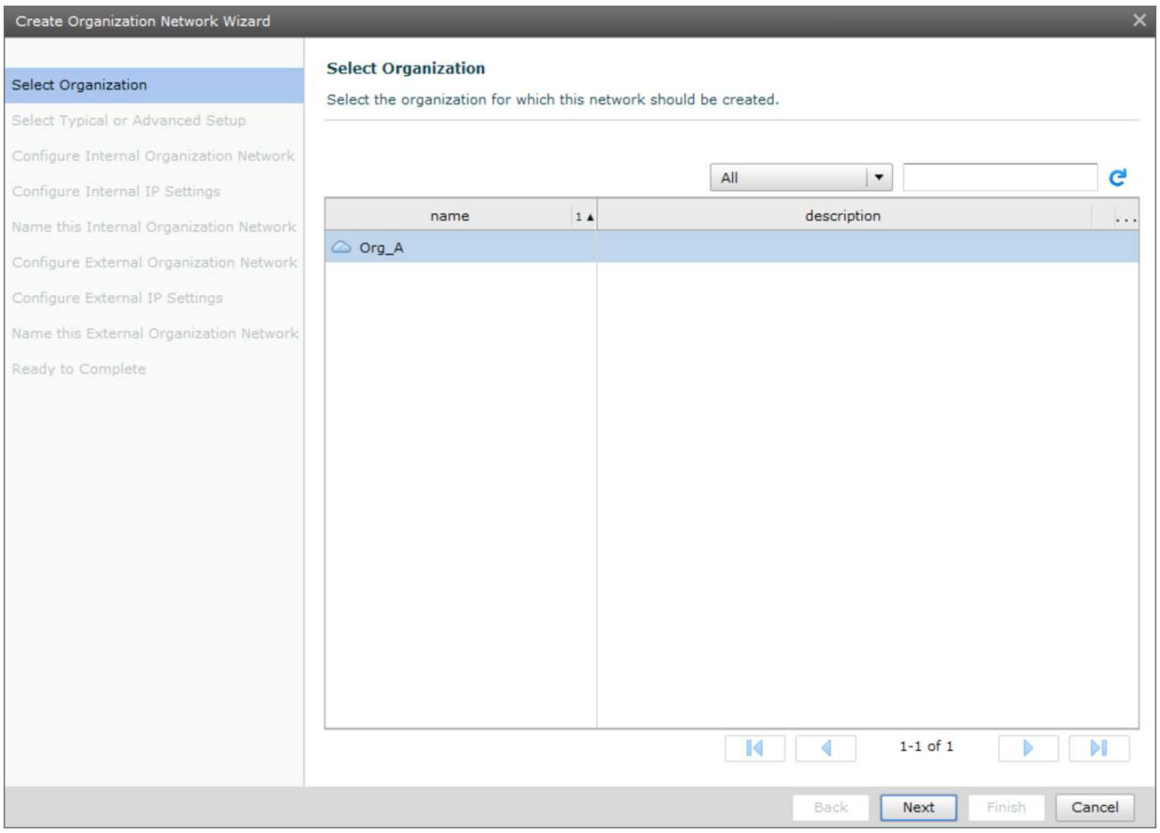

Step 6b. Since you are now creating only a directly connected external network, be sure that the Typical network setup is selected, deselect the “Create an internal network” check box, and select the “Create an external network via” check box. In the drop-down list, choose “direct connection” (Figure 29).

Figure 29. Specifying the Setup Options

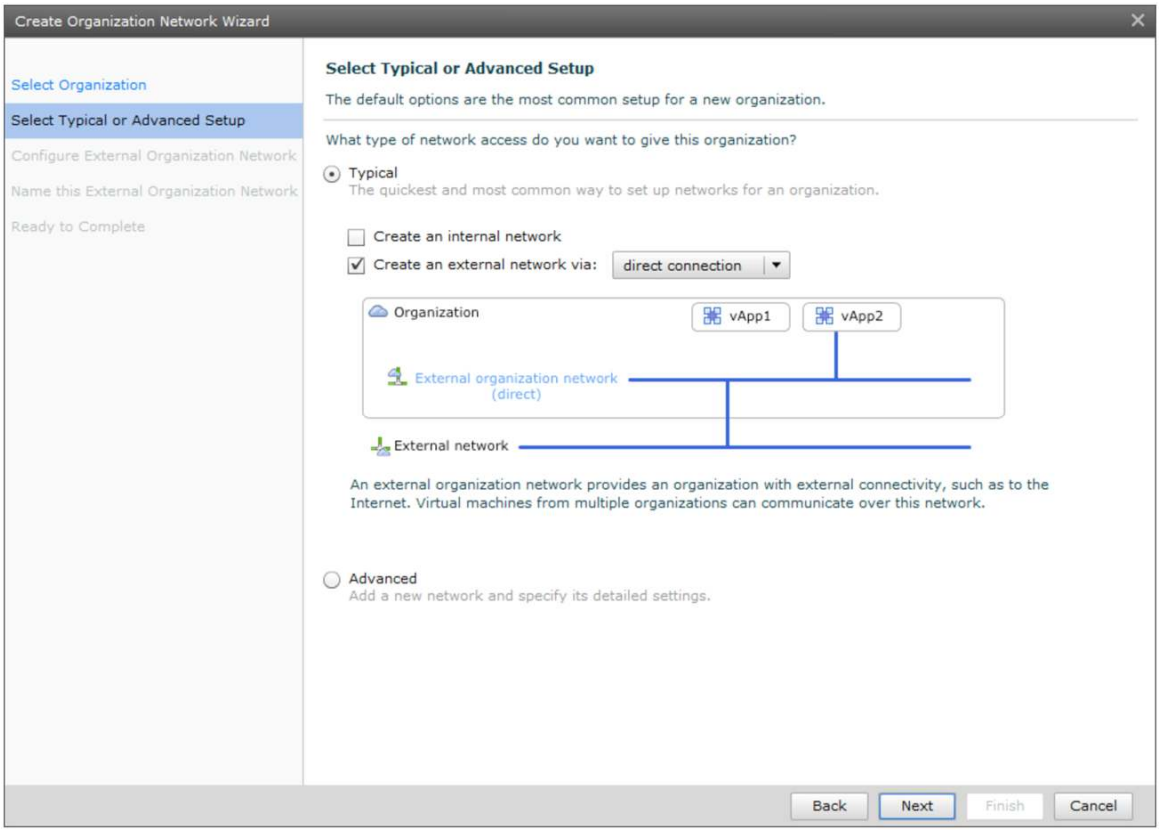

Step 6c. Select the external network to which this organization network will connect; then click Next (Figure 30).

In this case, there is only one choice, which is the external network defined in step 2 in the earlier section "Configuring VLAN-Based Isolation on the Cisco Nexus 1000V Series."

Figure 30. Selecting the External Network to Connect To

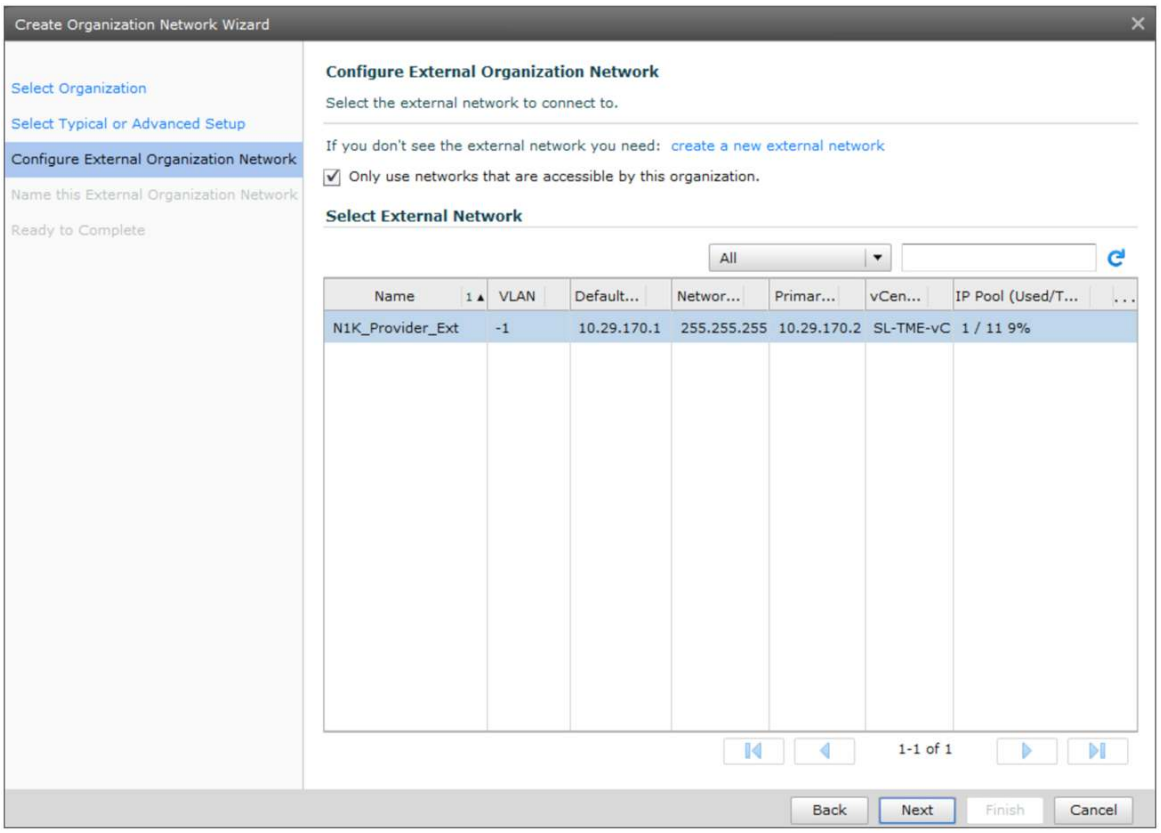

Step 6d. Type a name for this directly connected external organization network and click Next (Figure 31).

Figure 31. Naming the Network

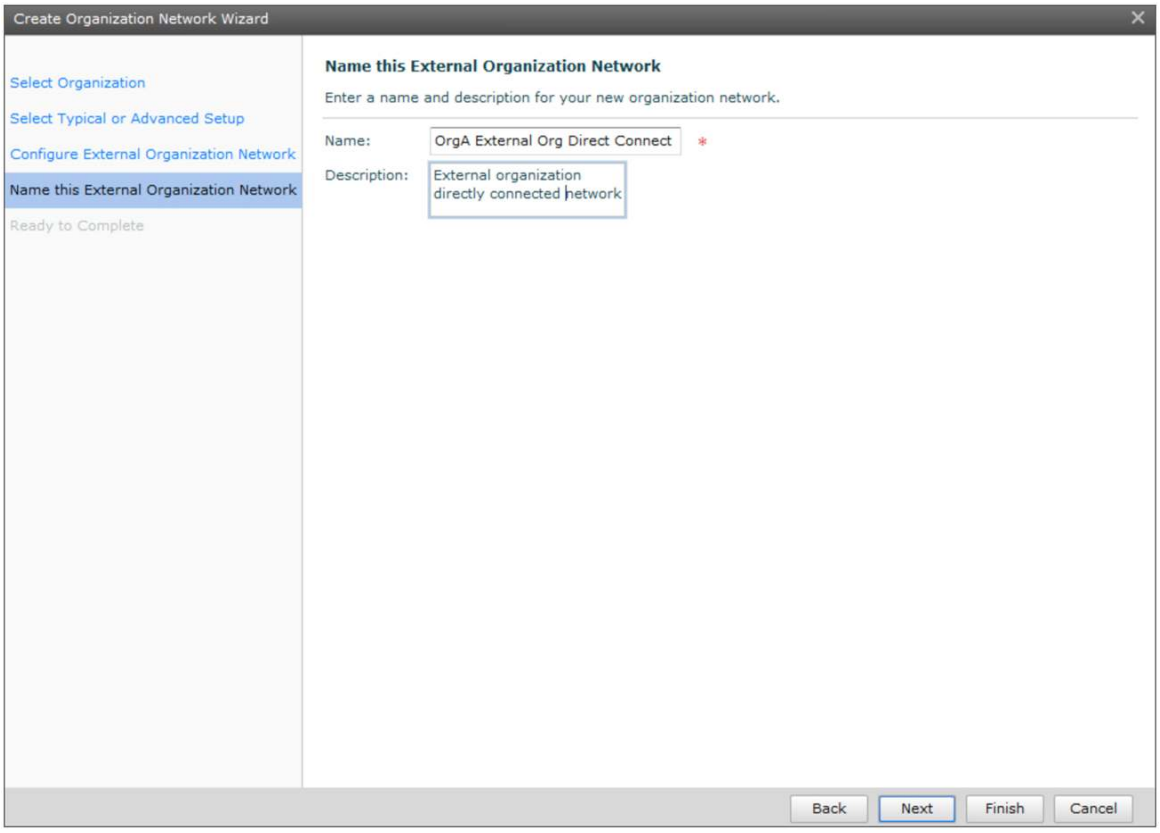

Step 6e. Review the final settings and click Finish (Figure 32).

Figure 32. Completing the Network

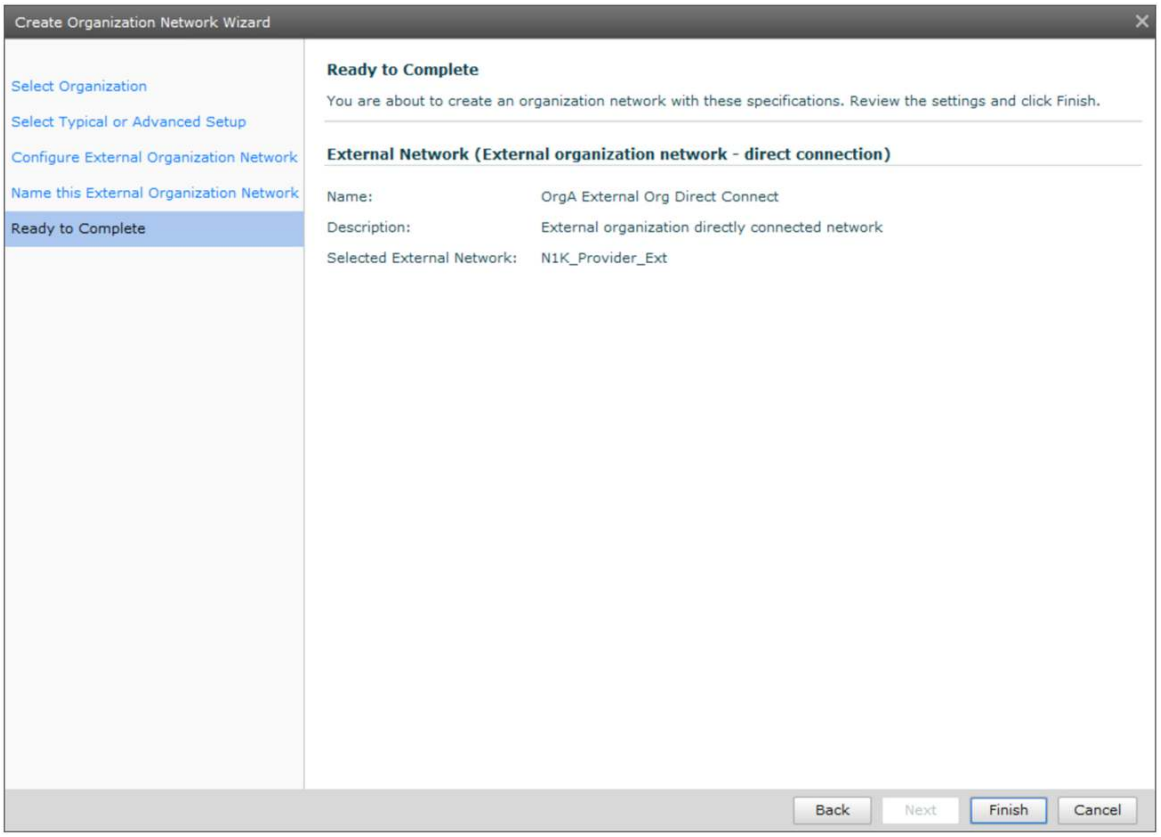

The organization is now prepared to start hosting vApps and virtual machines using all the available types of organization and provider networks. From here, vApp administrators and other administrators with the proper permissions can consume the networks that were created while taking advantage of the network services and features of the Cisco Nexus 1000V Series. For example, a vApp can be deployed, and the vApp network used (specified at the time the vApp is deployed) will use any of the available pools.

VXLAN-Backed Network Pools

Many customers are building private or public clouds. Intrinsic to cloud computing are multiple tenants with numerous applications using the cloud infrastructure. Each of these tenants and applications needs to be logically isolated, even at the networking level. For example, a three-tier application can have multiple virtual machines requiring logically isolated networks between the virtual machines. Traditional network isolation techniques such as IEEE 802.1Q VLAN provide 4096 LAN segments (through a 12-bit VLAN identifier) and may not provide enough segments for large cloud deployments. Cisco and a group of industry vendors are working together to address new requirements of scalable LAN segmentation as well as methods for transporting virtual machines across a broader diameter. The underlying technology, referred to as Virtual Extensible LAN (or VXLAN), defines a 24-bit LAN segment identifier to provide segmentation at cloud scale. More details can be found in the IETF draft: http://www.ietf.org/mail-archive/web/i-d-announce/current/msg39532.html.

In addition, VXLAN provides an architecture for customers to grow their cloud deployments with repeatable pods in different subnets. With Cisco Nexus 1000V Series Switches supporting VXLAN, customers can quickly and confidently deploy their applications to the cloud.

The Cisco Nexus 1000V Series supports VXLAN and provides significant benefits beyond VXLAN's baseline capabilities:

- Fully supports VMware vCloud Director 1.5 and later: See the compatibility matrix at Cisco Nexus 1000V and VMware Compatibility Information.
- Extends existing operational model to the cloud: The Cisco Nexus 1000V Series offers a nondisruptive operational model for network and server administrators. With the Cisco Nexus 1000V Series supporting VXLAN, the same operational model can now be extended to the cloud without disrupting the existing operational model, accelerating cloud deployment.

This guide does not discuss the details or best practices for deploying the Cisco Nexus 1000V Series. For that information, refer to the Cisco Nexus 1000V Series Switches Deployment Guide at http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/guide_c07-556626.html.

Overview of Cisco Nexus 1000V Series VXLAN

VXLAN is a Layer 2 network isolation technology that uses a 24-bit segment identifier to scale beyond the 4000-address limitations of VLANs. VXLAN technology creates LAN segments by using an overlay approach with MAC-in-IP encapsulation. The Cisco Nexus 1000V Series VEM encapsulates the original Layer 2 frame leaving the virtual machine (Figure 33).

Figure 33. VXLAN Encapsulated Frame Format

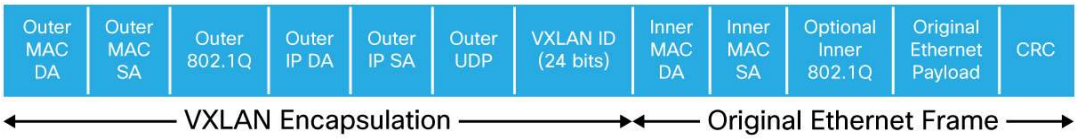

Each VEM is assigned an IP address, which is used as the source IP address when encapsulating MAC address frames to be sent on the network. This assignment [[OK?]] is accomplished by creating VMKNICs on each VEM (Figure 34). You can have multiple VMKNICs per VEM that are used as sources for this encapsulated traffic. The encapsulation carries the VXLAN identifier, which is used to scope the MAC address of the payload frame.

Figure 34. VEM VMKNIC Interface with VXLAN Capability

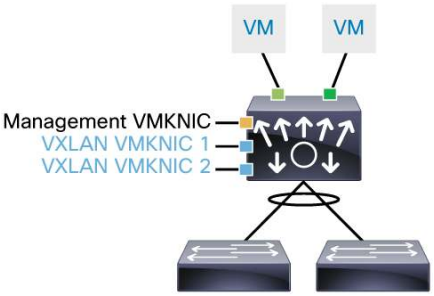

The connected VXLAN is specified within the port-profile configuration of the vNIC and is applied when the virtual machine connects. Each VXLAN uses an assigned IP multicast group to carry broadcast traffic within the VXLAN segment.

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 is performed 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 in an IP packet destined for the remote IP address. If the result is a miss (broadcast, multicast, and unknown unicast traffic falls into this category), the frame is encapsulated with the destination IP address set to the VXLAN segment's assigned IP multicast group.

Solution Architecture

Figure 35 shows the components of the solution. Each component will be discussed in detail with a focus on integration of the Cisco Nexus 1000V Series with VMware vCloud Director to support the VXLAN feature.

Figure 35. Cisco Nexus 1000V Series with VMware vCloud Director Solution Architecture

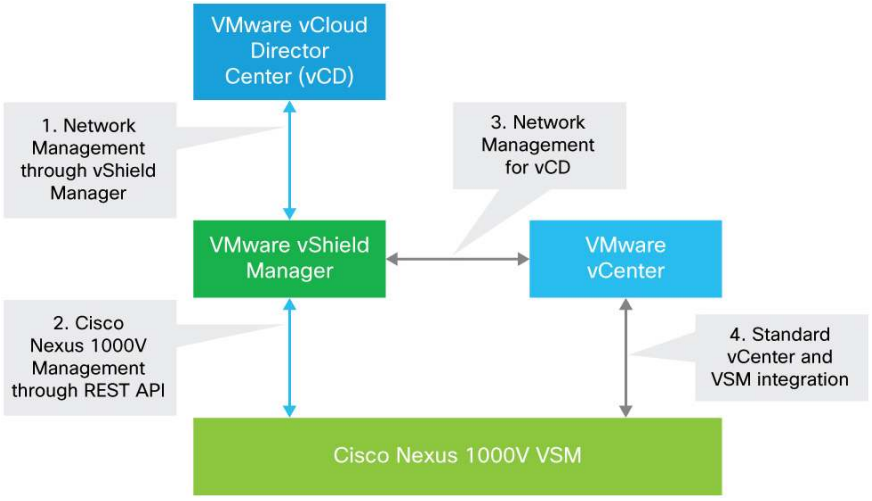

Solution Components

The main components of the solution are:

- VMware vCloud Director and vShield Manager communications
- Cisco Nexus 1000V Series and VMware vShield Manager communications
- VMware vShield Manager and vCenter communications
- VMware vCenter and Cisco Nexus 1000V Series communications

The next sections look at the components from the solution perspective.

VMware vCloud Director and vShield Manager Communications

VMware vCloud Director provides network services to the cloud through VMware vShield Manager. VMware vShield Manager interacts with the Cisco Nexus 1000V Series to make the Cisco Nexus 1000V Series available to VMware vCloud Director to build any type of network when building a tenant cloud. Each VMware vCloud Director cell requires access to a VMware vShield Manager host, which in turn provides network services to the cloud. You must have a unique instance of VMware vShield Manager for each VMware vCenter server you add to VMware vCloud Director vCenter.

Cisco Nexus 1000V Series and VMware vShield Manager Communications

VMware vCloud Director interacts with the Cisco Nexus 1000V Series using VMware vShield Manager. The VSM implements a representational state transfer (REST) API that allows the user to create all types of networks supported by VMware vCloud Director. This feature allows the user to design and implement networks in VMware vCloud Director that then are created on the Cisco Nexus 1000V Series Switch.

This feature is turned off by default in the Cisco Nexus 1000V Series, but it can be enabled by the following command on the Cisco Nexus 1000V Series Switch:

```
N1KV (Config) # feature network-segmentation-manager
```

VMware vShield Manager needs the following information to manage the VSM:

- VSM connectivity details
- Number of multicast addresses available for VMware vCloud Director
- Number of VXLANs that can be consumed by VMware vCloud Director

VMware vShield Manager and vCenter Communications

This communication will occur when an organization routed network is required for an organization. VMware vShield Manager will instantiate a VMware vSE appliance dynamically to provide NAT and IP gateway service for an organization network.

VMware vCenter and Cisco Nexus 1000V Series Communications

VMware vCenter provides centralized control and visibility to VMware vSphere virtual infrastructure. The Cisco Nexus 1000V Series is tightly integrated with VMware vCenter (Figure 36). This integration enables the network administrator and the server administrator to collaborate efficiently. While the networking policies can be enforced in the virtual access layer just as in the physical network, the Cisco Nexus 1000V Series helps maintain the separation of duties for the network and server teams. This communication is part of the initial Cisco Nexus 1000V Series setup, and there is no change in this communication because of VXLAN Implementation.

Figure 36. VSM to VMware vCenter Integration

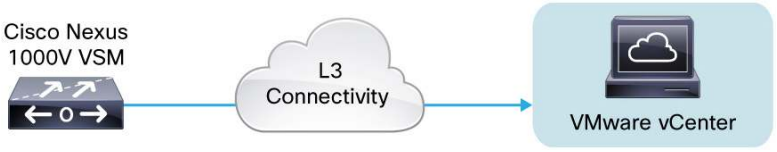

Deployment Steps

Here is a summary of the steps required to deploy VMware vCloud Director with the Cisco Nexus 1000V Series to take advantage of VXLAN:

1. Configure the VMware vSphere environment.
2. Install VMware vCloud Director.
3. Install VMware vShield Manager.
4. Install the Cisco Nexus 1000V Series Switch.
5. Turn on the VXLAN feature (feature segmentation) on the Cisco Nexus 1000V Series Switch.
6. Turn on the Network Segmentation Manager feature on the Cisco Nexus 1000V Series Switch.
7. Create a new port profile with the VXLAN capability.
8. Create a new VMware VMkernel interface to each VMware ESX host and assign the new port profile created in step 7.
9. Turn on multicast on the uplink physical Layer 3 switch or router.
10. Increase the maximum transmission unit (MTU) size on the Cisco Nexus 1000V Series uplink interfaces and the uplink physical interfaces to a minimum of 50 bytes.
11. Add the Cisco Nexus 1000V Series Switch to the list of switches managed by VMware vShield Manager.
12. Create the segment ID and multicast address in VMware vShield Manager.
13. Map clusters to the distributed virtual switch (DVS) in VMware vShield Manager to use VXLAN-based network pools (applicable only to VMware vCloud Director 5.1).

After performing these steps, you are ready to create network pools backed by VXLAN.

Deployment Considerations

Cisco Nexus 1000V Series Deployment

Standard practices of VSM deployment should be followed. VSM can be part of the same cluster in the vCenter where it is providing Layer 2 networking functionality. In vCenter, the additional recommendation is to have a dedicated resource pool for the VSMs that are excluded from the resources available to vCloud Director. The alternate option is to host VSMs on a dedicated Cisco Nexus 1010 Appliance.

Multicast

In a typical Layer 2 network using VLANs, if a frame is received with an unknown destination MAC address, it is flooded out of every interface (except the one it came from). In VXLAN, multicast/broadcast (and unknown unicast) frames will be sent encapsulated with a multicast destination IP address. Ideally, each VXLAN should use a different IP multicast group address to avoid flooding frames to VEMs that do not need them. When VXLAN encapsulation is used, a multicast IP address must be assigned to each VXLAN. It is up to vCloud Director to decide which VXLANs will share the same IP multicast group address.

If the VXLAN VMKNICs on different VEMs belong to the same subnet, you need to enable IGMP snooping only on the VLAN on upstream switching to provide Layer 2 optimization for multicast traffic.

If the VXLAN VMKNICs on different VEMs are in different subnets, Layer 3 multicast routing must be configured on the upstream routers. The recommended protocol to deploy is Bidirectional Protocol Independent Multicast (PIM), since the VEMs act as both multicast speakers and receivers at the same time. For more information on deploying multicast on Cisco switches and routers, please refer to the link below: http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/ps6592/whitepaper_c11-474791.html.

Proxy Address Resolution Protocol (ARP)

The VXLAN encapsulated packet uses the source address of the dedicated VMkernel interface IP address on the VEM. The destination address is initially the Multicast Group address. Once the actual VEM destination is determined, the subsequent packets have the actual unicast address for the destination. In the event that the destination is on a different subnet, the VEM will still ARP for the destination instead of sending it to the default gateway. To address this behavior, you need to enable the Proxy ARP feature on the Layer 3 gateway (upstream Layer 3 switch or router), so it can respond to off-subnet ARPs.

Communications Outside the VXLAN

The VXLAN format is supported only by the Cisco Nexus 1000V. If communications have to be established outside the VXLAN, there are two options available today.

Virtual Machine with One Interface in VXLAN and One in VLAN

Multihomed virtual machine with one interface in VXLAN and one interface in a VLAN. Any communication to be established to VXLAN from outside has to traverse the multihomed virtual machine through the VLAN interface. Now consider an example of a vApp with the following virtual machines:

- Dual-homed client machine with one interface in VXLAN 4400 and one interface in VLAN 44
- Web server with one interface in VXLAN 4400 and one interface in VXLAN 4401
- Database server with one interface in VXLAN 4401

In this scenario, illustrated in Figure 37, you can remote desktop to a client machine on the interface, which is on VLAN 44, and then browse to the web server in VXLAN 4400. Now the web server can communicate to the database server on VXLAN 4401, which is totally isolated - that is, the client machine has no access to the database server directly.

Figure 37. External Connectivity with Dual-Homed Virtual Machine

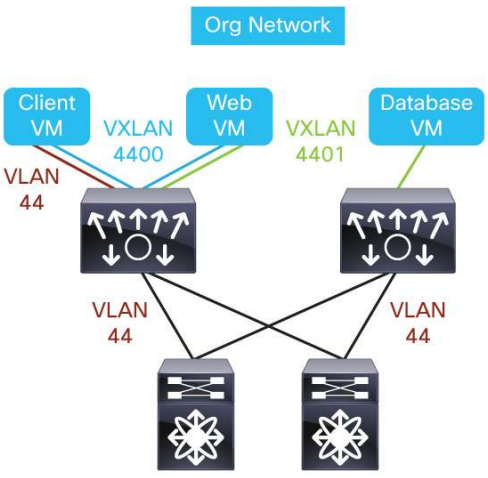

vShield Edge Providing NAT/Gateway Functions

vCloud Director deploy vShield Edge virtual device to provide NAT and IP gateway connectivity between a vApp network and an organization network or between an organization network and an external network. This option is available if you select routed mode for external connectivity when you create an organization network. Figure 38 shows how the configuration above connected to external network using vShield Edge. The external interface will be in VLAN 44 and internal interface on VXLAN 4400.

Figure 38. External Connectivity with VMware vShield Edge

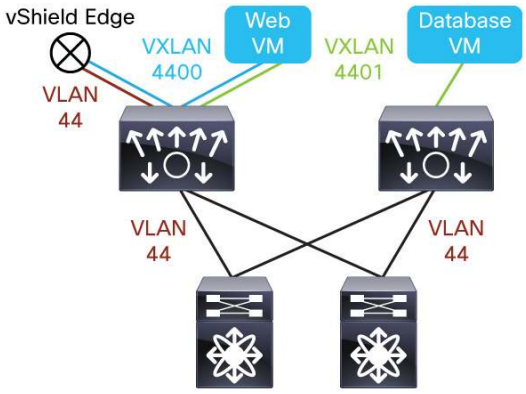

VXLAN Working with OTV/LISP

VXLANs are intended for creating a large number of logical networks in a cloud environment within a data center. Overlay Transport Virtualization (OTV) is a data center interconnect technology extending VLANs to different data centers over Layer 3. Unlike VXLAN, OTV has simpler deployment requirements since it does not mandate a multicast-enabled transport network. Locator ID Separation Protocol (LISP) goes a step further by providing IP address mobility between data centers with dynamic routing updates. While VXLAN, OTV, and LISP all use UDP/IP encapsulation, they serve very different networking purposes and are hence complementary to each other.

Scalability with VXLAN

Today, a single Cisco Nexus 1000V Series Switch supports a total combination of up to 2000 VXLAN and VLAN Layer 2 logical networks. In order to scale beyond 2000 Layer 2 logical networks, you need to deploy additional Cisco Nexus 1000V Series Switches.

In case where a single organization has spanned multiple Cisco Nexus 1000V Series Switches, the organization administrator will be presented with same organization networks with different names.

Securing VXLAN in the Physical Network

Since VXLAN is transported over IP in physical network, some best practice recommendations should be implemented 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 them part of the same VLAN and keep that VLAN a strict Layer 2 VLAN. Only the VMKNICs used for VXLAN encapsulation should attach to this VLAN. This approach provides natural protection and limits unwanted exposure to external communication.

In the scenario where number of VEMs has exceeded the available IP addresses in the subnet, VMKNICs for VXLAN encapsulations may need to be assigned IP addresses in multiple subnets. In this scenario, where VXLAN VMkernel interfaces belong to two different VLANs, the communication between the multiple subnets has to take place through a Layer 3 switch or router. Both VLANs must have SVI interfaces.

To make sure that VXLAN traffic cannot be attacked or snooped from unauthorized endpoints, you can use one of two options:

- Use ACLs to prevent unauthorized injection of VXLAN encapsulated traffic to VEM VMKNICs from outside sources.
- Use a VRF to segregate the VLANs and SVIs on which VXLAN VMKNICs are assigned IP addresses.

For specific configurations of ACLs or VRFs, please refer to the configuration guides of you physical Layer 3 switch or router.

The options just described not only reduce external security threats, but also keep the multicast deployment simpler in the physical network.

Port Channels

Port channels use different load-balancing algorithms for dividing outgoing traffic among different physical interfaces. IP encapsulation results in all outgoing VXLAN traffic carrying an outer header that has the source MAC/IP address of the VEM's VMKNIC. For optimal load balancing, users must configure either a 5-tuple-based hash as the load-sharing algorithm. The following command is recommended for optimizing the VXLAN traffic on multiple uplinks:

Example:

```
n1000v(config)# port-channel load-balance source-dest-ip-port
```

MTU Size

VXLAN traffic is encapsulated in a UDP packet when sent out to the physical network. This encapsulation imposes the following overhead on each packet:

Outer Ethernet Header (14) + UDP header (8) + IP header (20) + VXLAN header (8) = 50 bytes

To avoid fragmentation and possible performance degradation, all the physical network devices transporting the VXLAN traffic need to handle 50 bytes greater than the max MTU size expected for the frame. Therefore, you must adjust the MTU settings for all the devices that will transport the VXLAN traffic. This includes the uplink port-profiles of Cisco Nexus 1000V Series Switches carrying the VXLAN traffic.

Some switches take a global setting for the MTU and others must be set per port. Refer to the system configuration guides of your upstream switches to increase the MTU of the physical interfaces of all the transit switches and routers.

VXLAN Deployment Use Cases

Deploying Two-Tier Web Development vApp

In the first use case, illustrated in Figure 39, we are deploying a web development vApp using VXLAN for an isolated Layer 2 network. The vApp consists of three virtual machines: web, database, and Windows client. The goal is to provide the web developer with an isolated test environment in where he can remote desktop to the Windows client to access the web server that resides in the VXLAN. Only the Microsoft Windows client has northbound connectivity via the external network.

Figure 39. Deploying a Two-Tier Web Development vApp

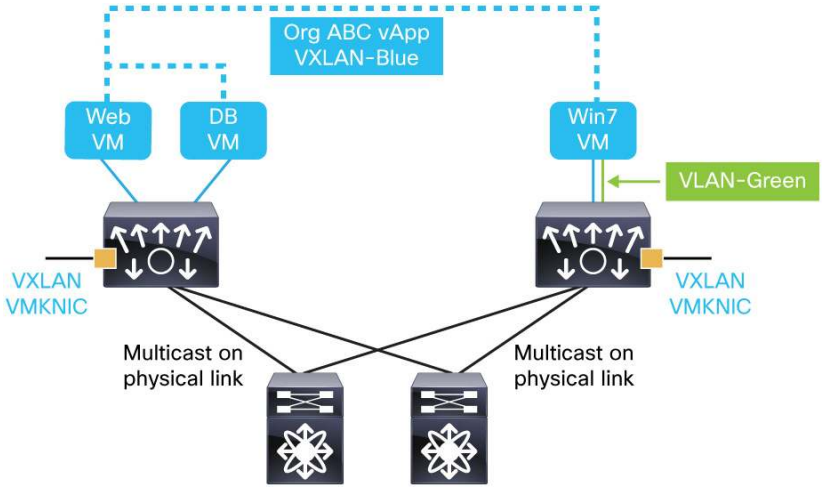

This document does not discuss the installation and setting up VMware vSphere environment with Cisco Nexus 1000V Software Release 1.5 and VMware vCloud Director 1.5.

Setting Up the Cisco Nexus 1000V Series for VXLAN

Step 1. Turn on the NSM and VXLAN feature on Cisco Nexus 1000V Series.

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

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

```
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         enabled
port-profile-roles           1         disabled
private-vlan                 1         disabled
```

```
segmentation                    1         enabled
sshServer                       1         enabled
tacacs                          1         disabled
telnetServer                    1         disabled
```

You can get more details around NSM configuration by issuing the following command:

```
N1KV-VSM(config)# show network-segment manager switch
switch: default_switch
  state: enabled
  dvs-uuid: 3a 7d 2d 50 d7 89 d3 29-3e 62 ad d6 9a af e2 e0
  dvs-name: N1KV-VSM
  mgmt-srv-uuid: 1190CA35-488B-45B1-A2C3-B43815F49534
  reg status: unregistered
  last alert: - seconds ago
```

The registration status will change when we integrate the NSM with vShield Manager.

Step 2. Create a port-profile with capability VXLAN.

```
port-profile type vethernet VMK-FI-A
   vmware port-group
  switchport access vlan 10
  capability vxlan
  no shutdown
  state enabled
```

You can verify that the VXLAN is enabled on this interface by issuing the command:

```
N1KV-VSM(config)# show port-profile name VMK-FI-A

port-profile VMK-FI-A
 type: Vethernet
 description:
 status: enabled
 max-ports: 32
 min-ports: 1
 inherit:
 config attributes:
  switchport access vlan 10
  capability vxlan
  no shutdown
 evaluated config attributes:
  switchport access vlan 10
  capability vxlan
```

```
 no shutdown
assigned interfaces:
port-group: VMK-FI-A
system vlans: none
capability l3control: no
capability iscsi-multipath: no
capability vxlan: yes
port-profile role: none
port-binding: static
```

Step 3. Create a VMkernel interface on each ESX host.

Attach a VMkernel interface to each ESX host of the cluster in vCenter, which will be managed by vCloud Director.

Navigate to Home > Inventory > Host and Clusters in the vCenter. Select the host of the cluster, which will be managed by vCloud Director. Under Configuration > Networking > vSphere Distributed Switch, select Manage Virtual Adapter, as shown in the Figure 40.

Figure 40. Adding VMkernel Interface on VMware ESX Host: Screen 2

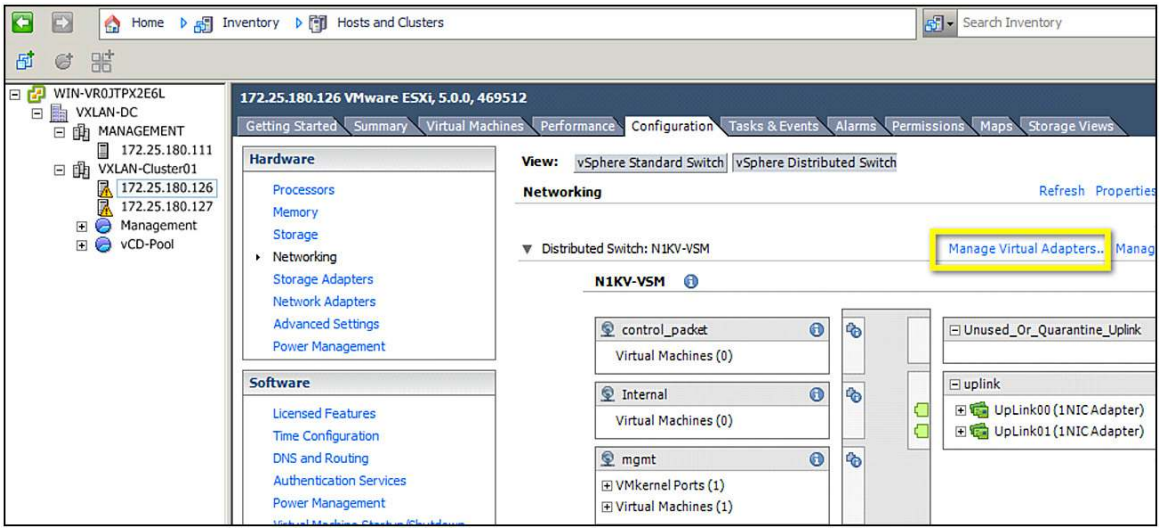

Now add a new VMkernel interface, as shown in Figure 41.

Figure 41. Adding VMkernel Interface on VMware ESX Host: Screen 2

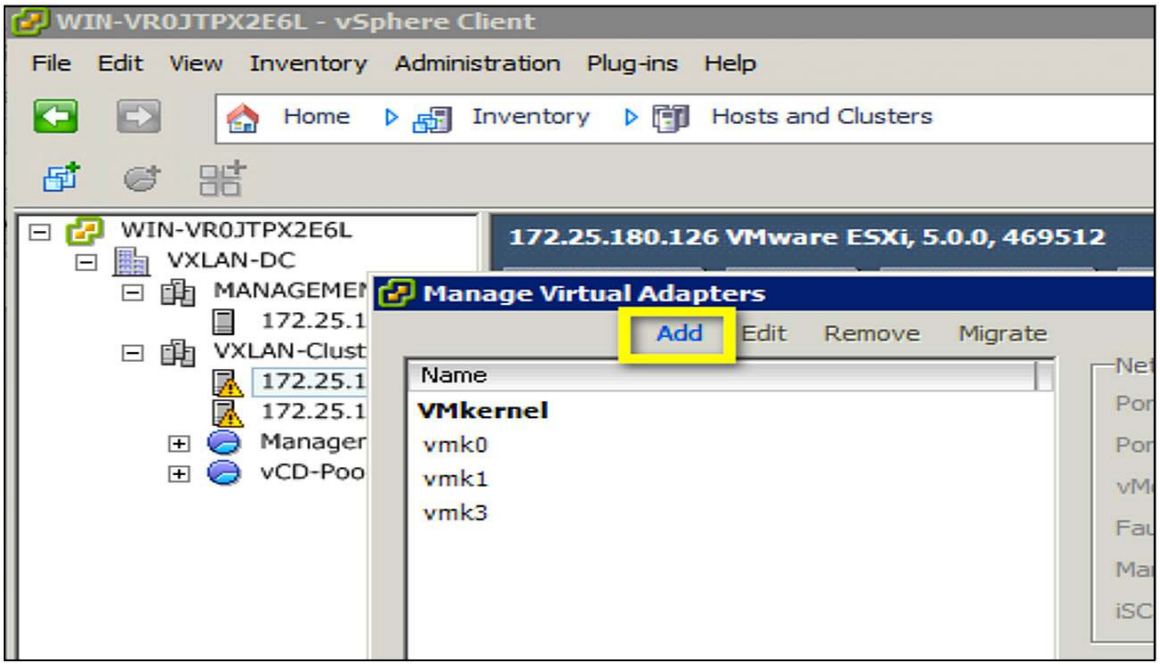

Complete the process by following the steps in Figures 42 through 46.

Figure 42. Selecting a New Adapter

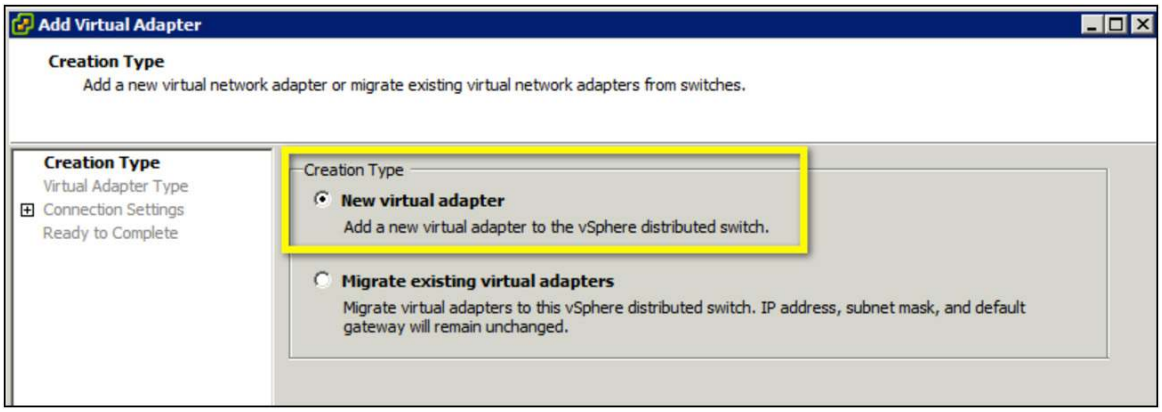

Figure 43. Selecting a New VMkernel Interface

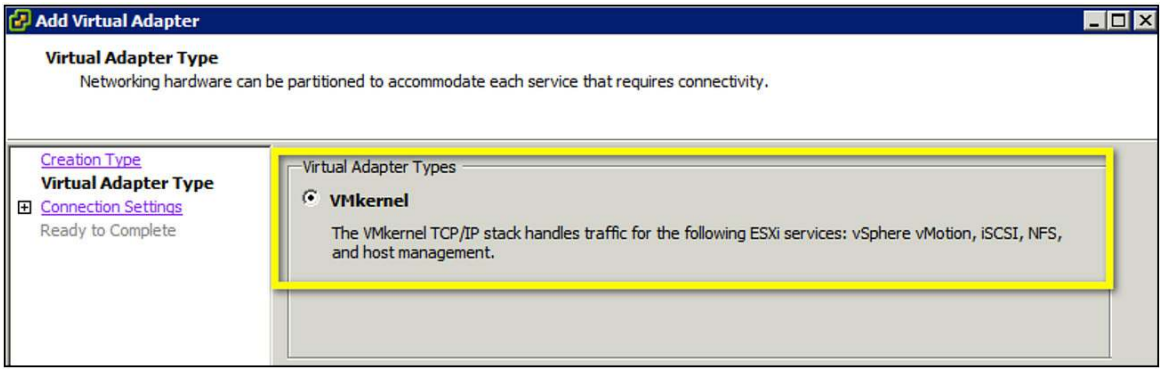

Figure 44. Selecting a VXLAN-Enabled Port Profile

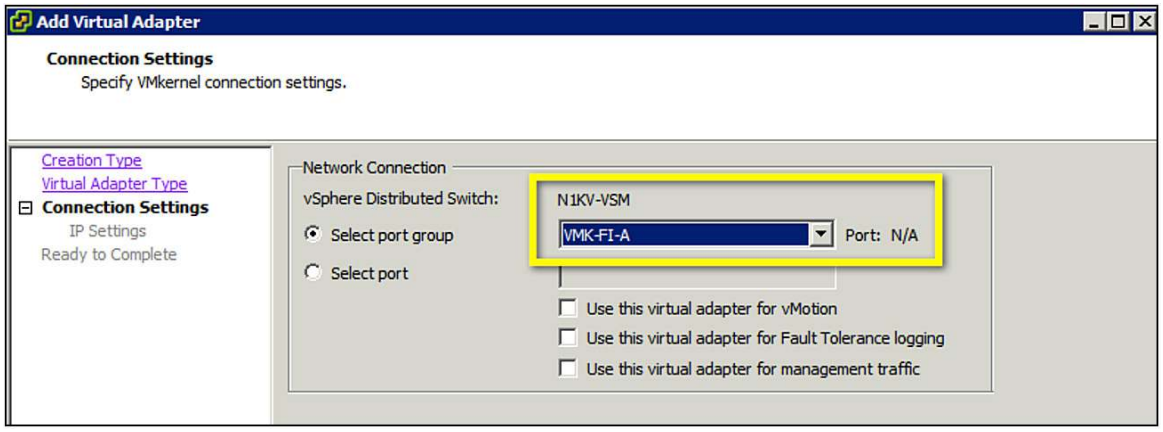

Figure 45. Configuring an IP to VMkernel Interface Used to Encapsulate VXLAN

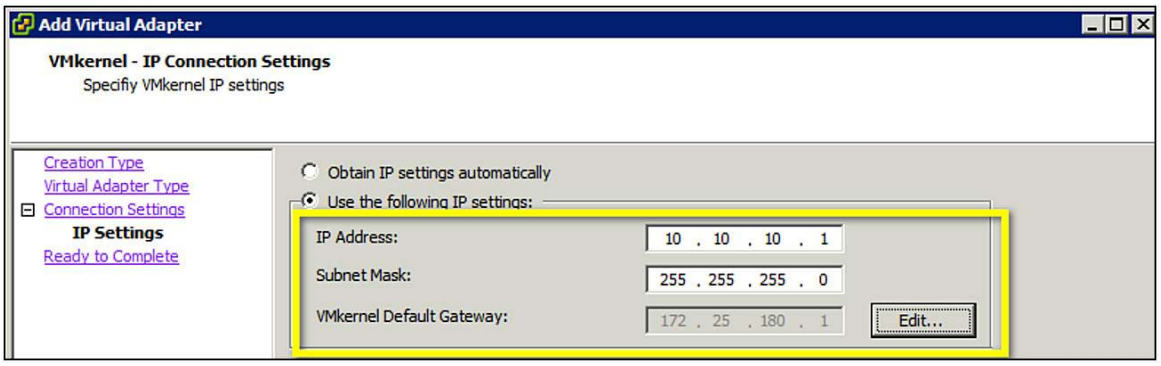

Figure 46. Summary of VXLAN VMkernel Interface

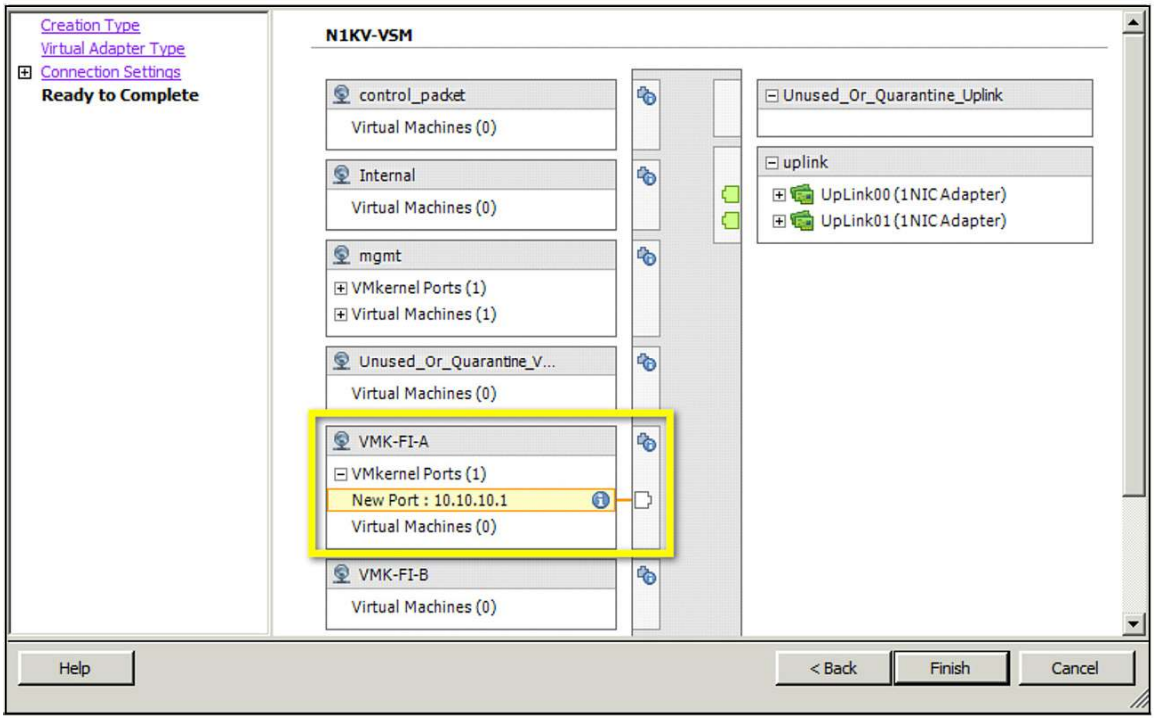

Repeat the steps for other ESX hosts. The only difference is that you need to assign a unique IP address for each VMkernel interface created on the host.

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

```
N1KV-VSM(config)# show  port-profile name VMK-FI-A

port-profile VMK-FI-A
 type: Vethernet
 description:
 status: enabled
 max-ports: 32
 min-ports: 1
 inherit:
 config attributes:
  switchport access vlan 10
  capability vxlan
  no shutdown
 evaluated config attributes:
  switchport access vlan 10
  capability vxlan
  no shutdown
 assigned interfaces:
```

```
  Vethernet4
  Vethernet5
 port-group: VMK-FI-A
 system vlans: none
 capability l3control: no
 capability iscsi-multipath: no
 capability vxlan: yes
 port-profile role: none
 port-binding: static
```

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

Step 4. Change the MTU on uplink interface.

To avoid fragmentation, it is highly recommended to increase the MTU of the uplink interfaces of Cisco Nexus 1000V and the physical interfaces of the upstream switches, which are connected the Layer 2 domain of the vSphere environment.

The following command needs to be configured on the uplink port-profile to increase the MTU:

```
port-profile type ethernet uplink
  vmware port-group
  switchport trunk allowed vlan 10,180
  switchport mode trunk
  switchport trunk native vlan 180
  mtu 1550
  no shutdown
  system vlan 10,180
  state enabled
```

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

Enabling Multicast on the Upstream Physical Switch

In this example, all the VEM VXLANs are in the same VLAN, and we are enabling the IGMP snooping querier on the VLAN:

```
vlan 10
  ip igmp snooping querier 10.45.46.45

5K-B# show ip igmp snooping querier

Vlan  IP Address       Version   Expires     Port
10    10.45.45.45      v3        00:02:45    Ethernet1/8
```

Integrating with VMware vCloud Director 1.5.1 and vShield Manager 5.0.1

Integrating VSM (Cisco Nexus 1000V Series) with VMware vShield Manager

Now we will integrate vShield Manager with Cisco Nexus 1000V Series. The following information is required to add the Cisco Nexus 1000V Series as a managed switch in VMware vShield Manager.

- VSM connectivity details
- Multicast addresses
- Number of VXLANs

As shown in Figure 47, start by logging into the vShield Manager web interface: https://vShield-Manager-IP

Figure 47. VMware vShield Manager Login GUI

Select Settings & Reports > Configurations > Networking, as shown in Figure 48.

Figure 48. Adding Cisco Nexus 1000V Series Switch in VMware vShield Manager

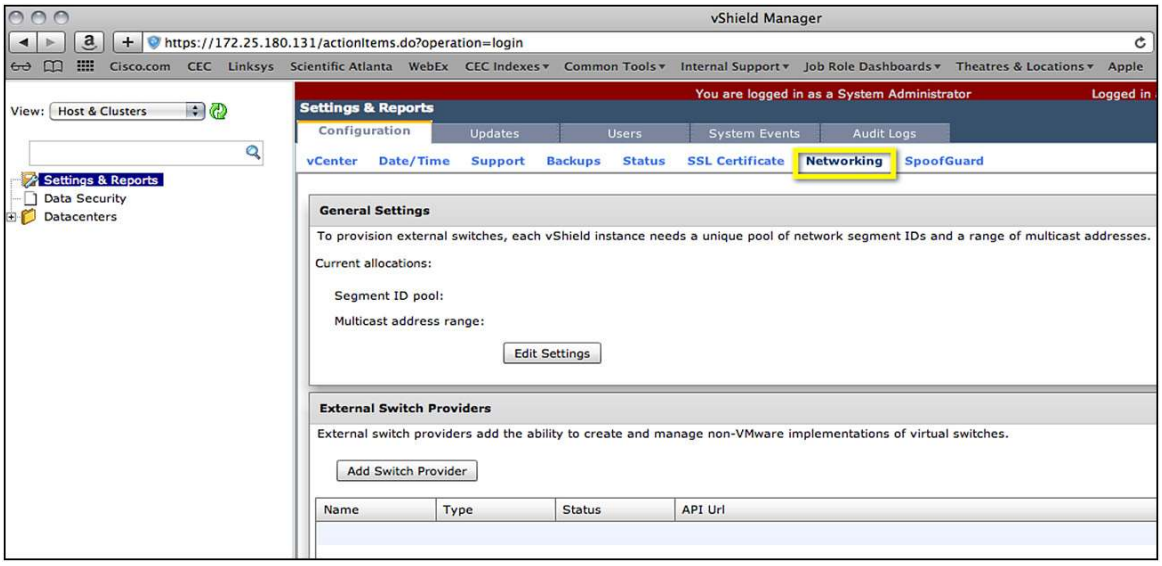

Next, provide the VXLAN ID range and multicast address range (Figure 49). Try to put as many multicast groups as available that will not exceed the state in the transit switches and routers. In case you have multiple vShield Managers, the VXLAN IDs and groups must not overlap between them.

Figure 49. Providing VXLAN ID Ranges and Multicast Address Ranges

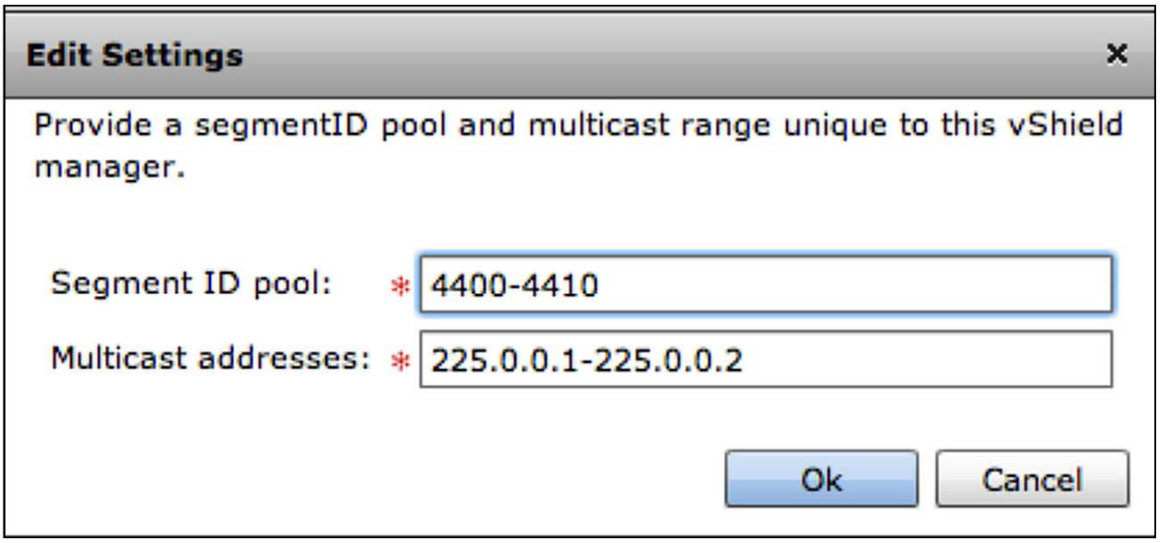

Now add the Cisco Nexus 1000V as a switch in vShield Manager. Figure 50 shows the specific settings.

Figure 50. Cisco Nexus 1000V Series Specific Settings

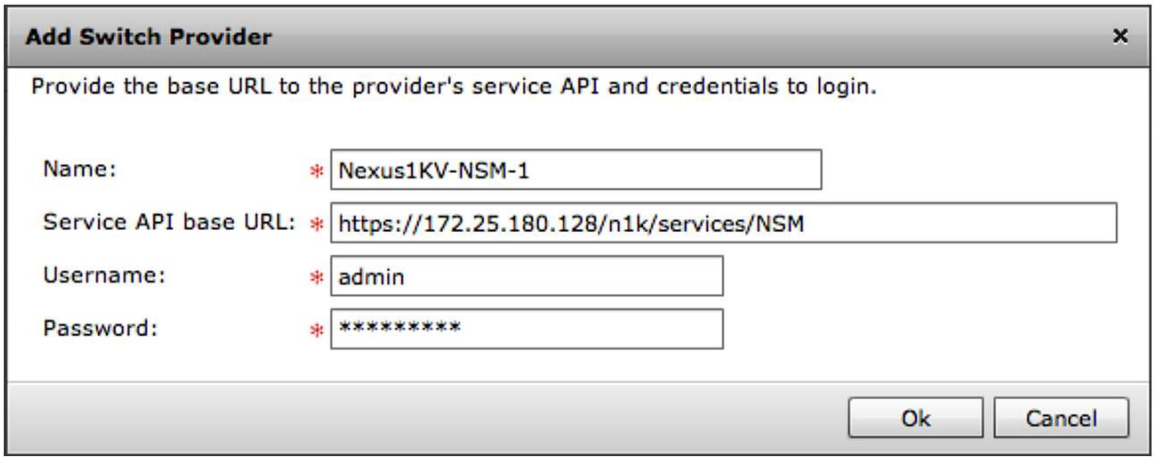

A green checkmark should appear under Status, as shown in Figure 51.

Figure 51. Cisco Nexus 1000V Switch Added to vShield Manager List of Managed Switches.

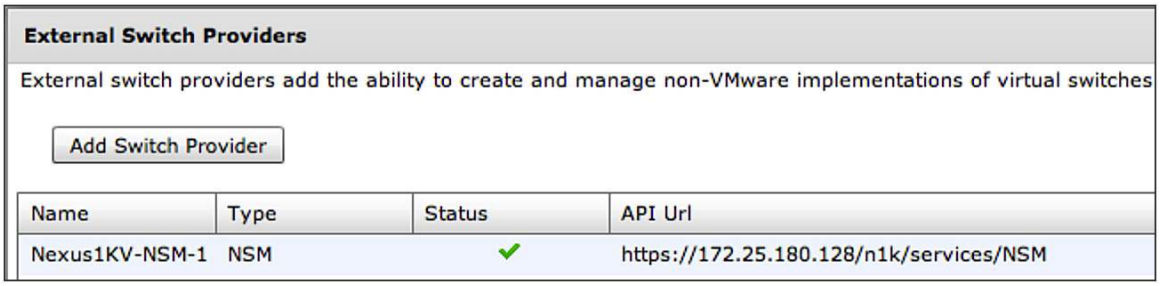

VMware vCloud Director Settings

The following section will focus on the vCloud Director network configuration in our use case. Please refer to the VMware documentation for non-network settings, including creating virtual data centers, organizations, organization virtual device contexts, and so on.

Building an External Network for Provider vDC

The external network provides northbound connectivity to an organization within vCloud Director. This will always be a port-group backed network. Figures 52 through 55 show the steps involved in associating the existing port-group with the newly created external network.

Navigate to System > Home > Guided Tasks. Click Create External Network. Select the port-group which will provide the external connectivity.

Figure 52. Selecting the Preconfigured Port Group (Port Profile)

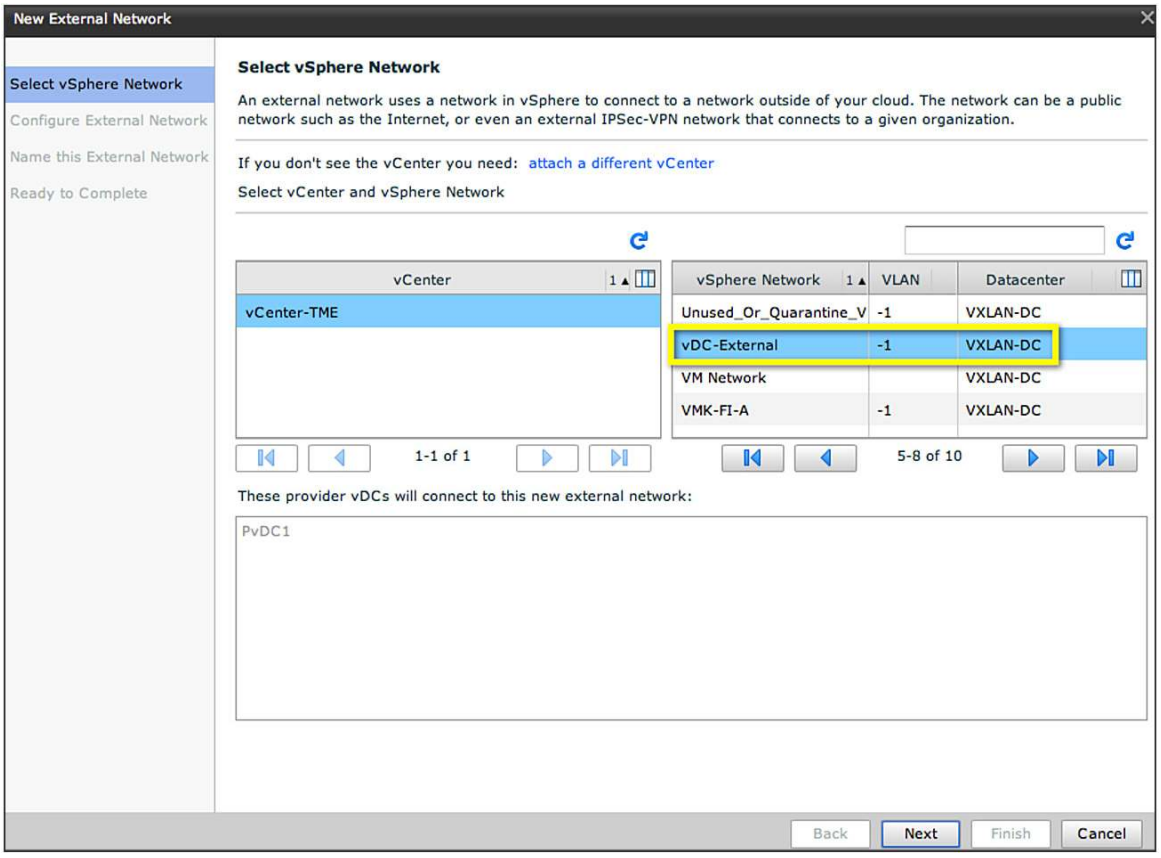

Next, specify a pool of IP addresses which can be consumed by the virtual machines requiring external connectivity

Figure 53. Network Settings

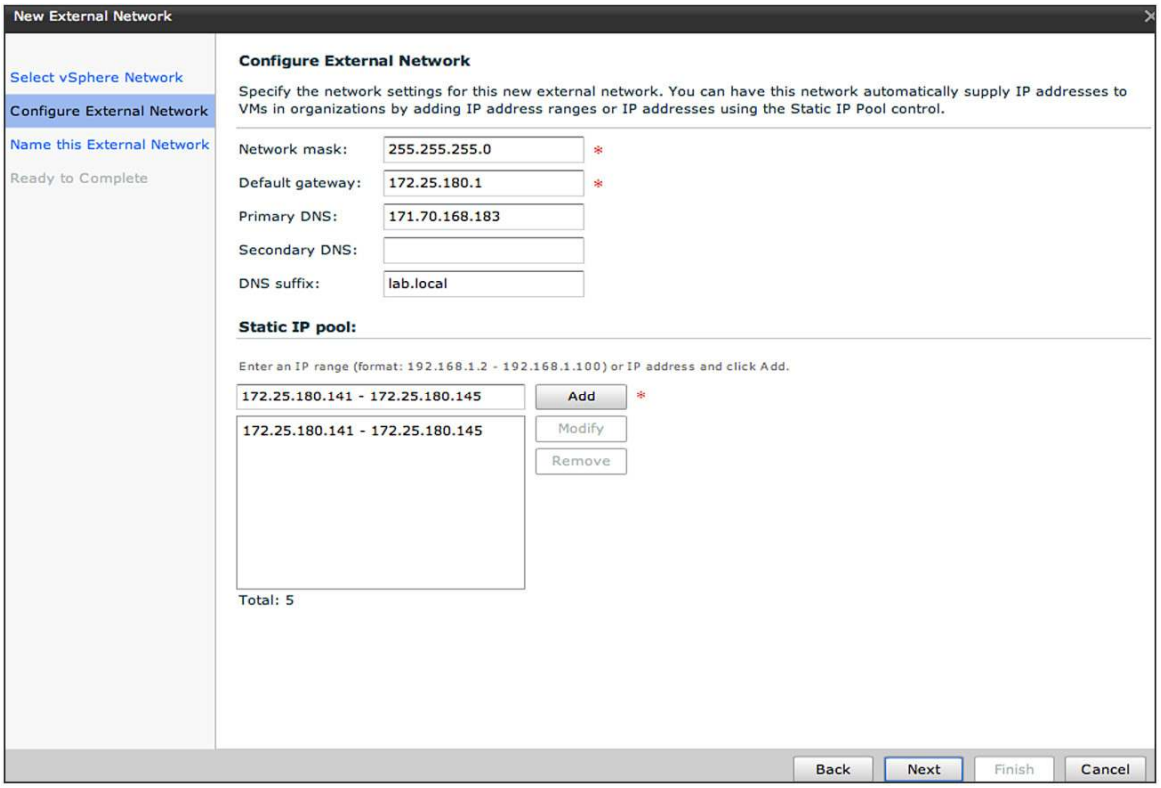

Provide the name for this external network.

Figure 54. Naming the External Network

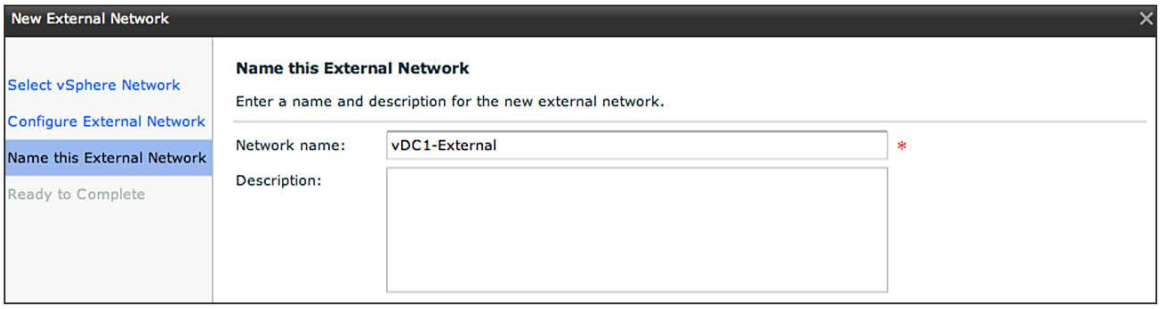

Figure 55. Summary of the External Network

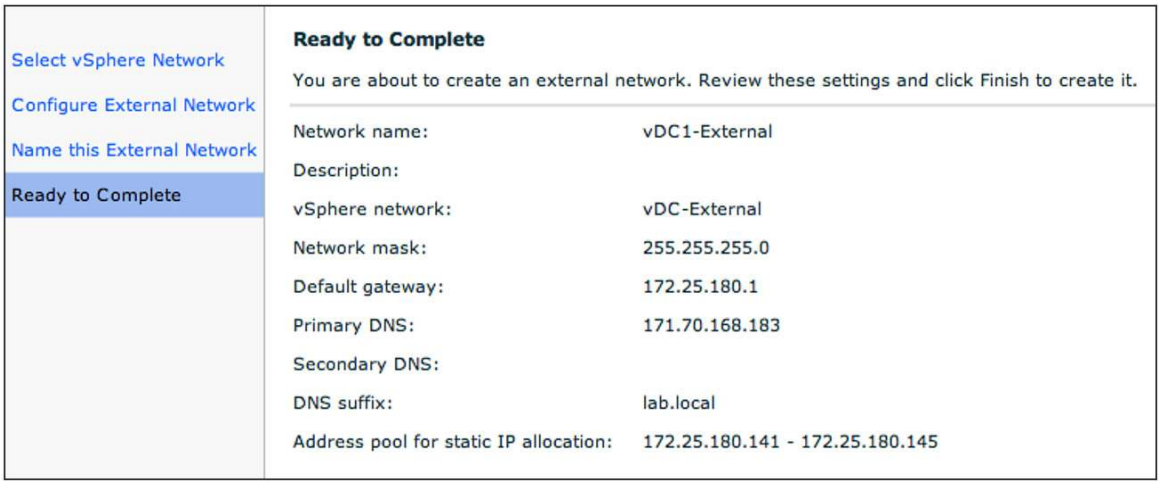

Creating the VXLAN Network Pool

This section shows the steps for creating network pool, which provides VXLAN segments.

Navigate to System > Guided Tasks > Create Network Pool. Select Create Network Pool. The wizard will appear which will guide you through the steps required (Figures 56 through 59).

Figure 56. Select “Network Isolation-backed”

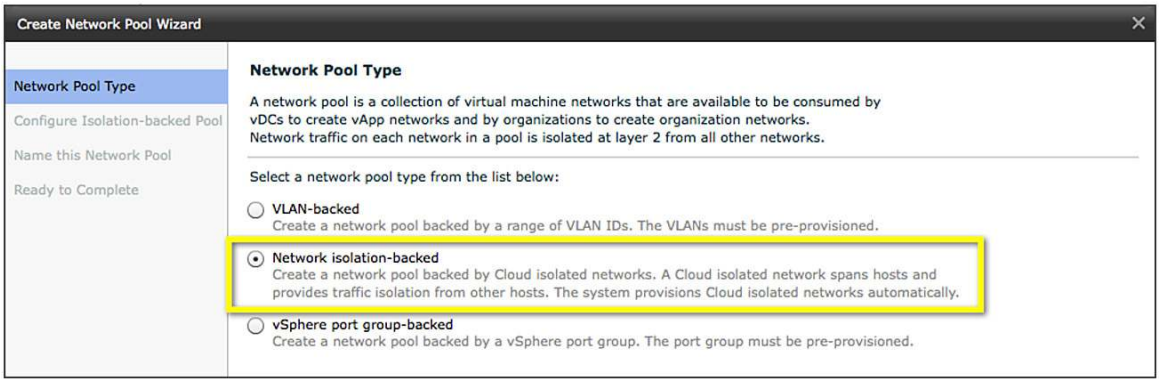

Figure 57. Enter the Number of VXLAN Segments Available in the Pool

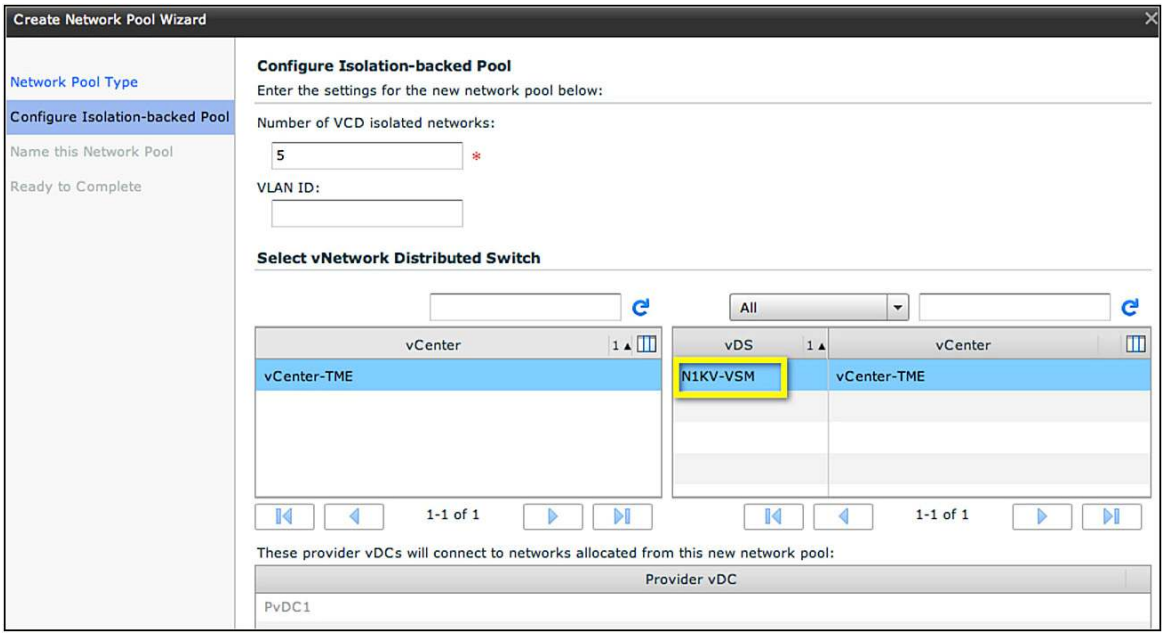

Figure 58. Enter the Name of the Network Pool

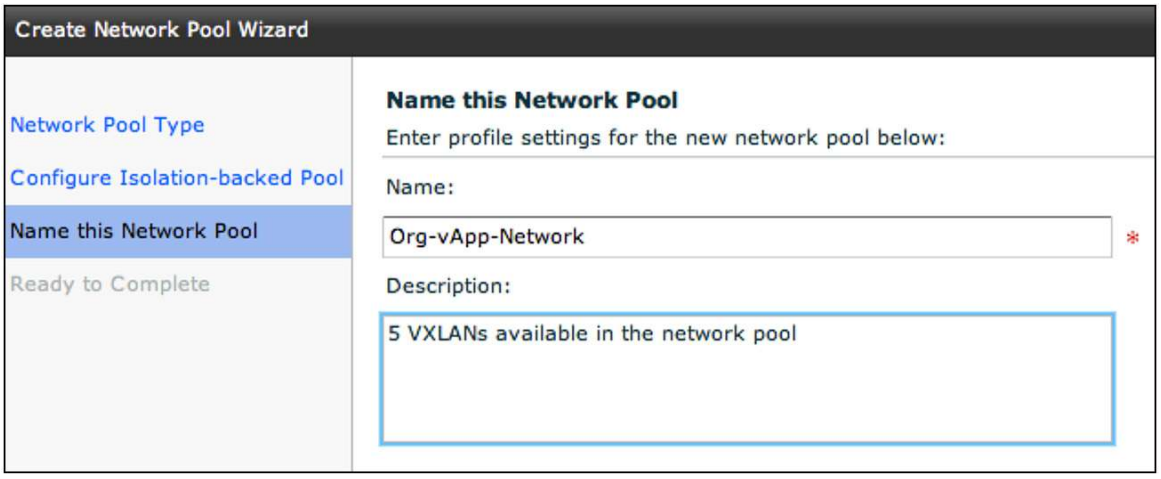

Figure 59. Wizard Summary Screen

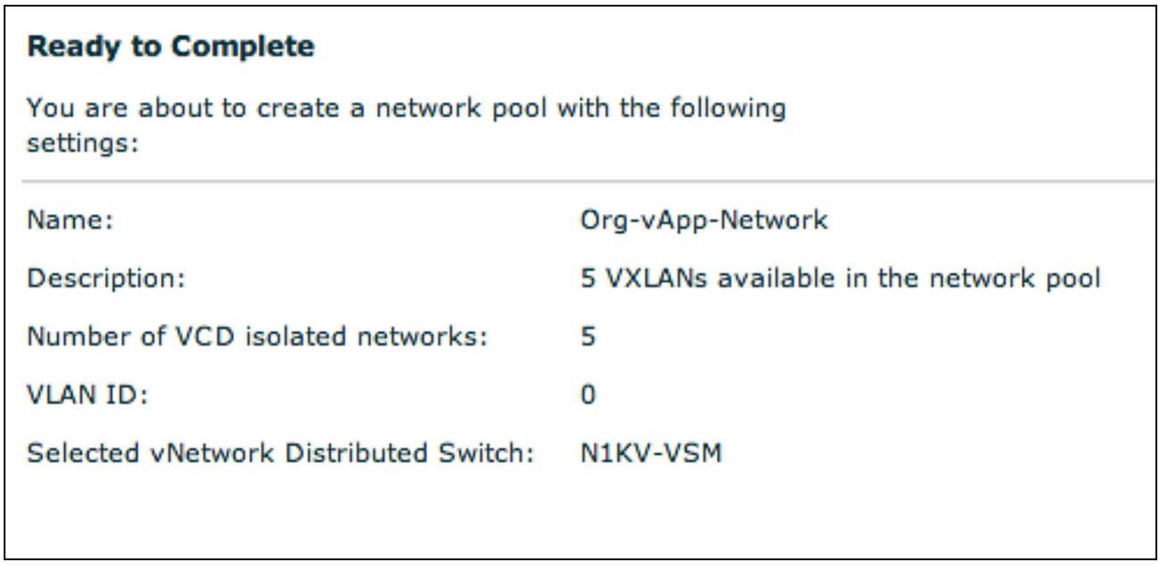

This completes the creating of network pool which is VXLAN backed. Now, we can consume the VXLAN backed Layer 2 segment for Organizations.

Assigning Network Resources to an Organization

When you are creating an organization network in vCloud Director, you need to select a network from the Network pool. In this case, you will select a network from the network pool, which was defined in the previous section (Figure 60).

Figure 60. Selecting the VXLAN-Backed Network Pool for the Organization

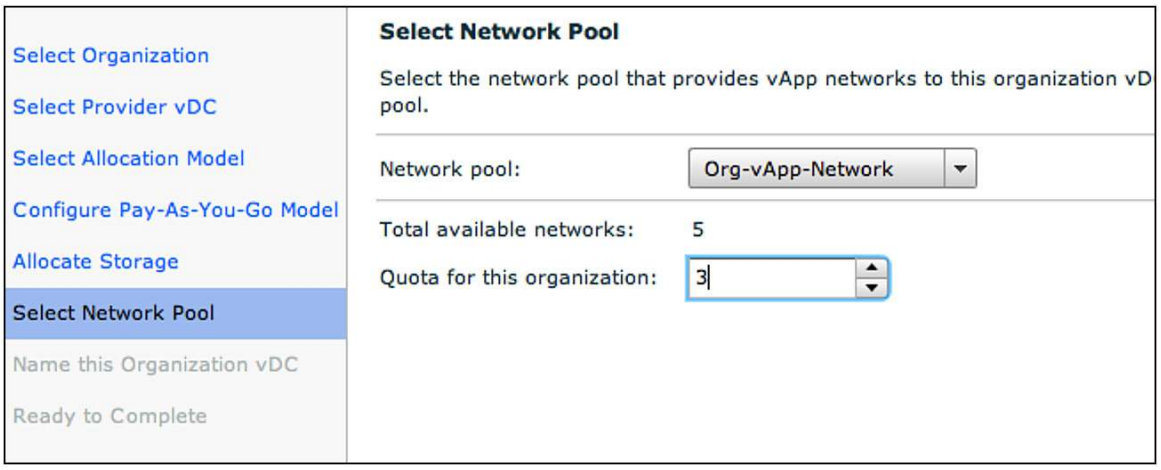

Creating Organization Networks

In this section, you will create organization networks, which will consume Layer 2 network segments from the network pool. Navigate to System > Home > Add network to an organization (Figure 61).

Figure 61. Select Organization

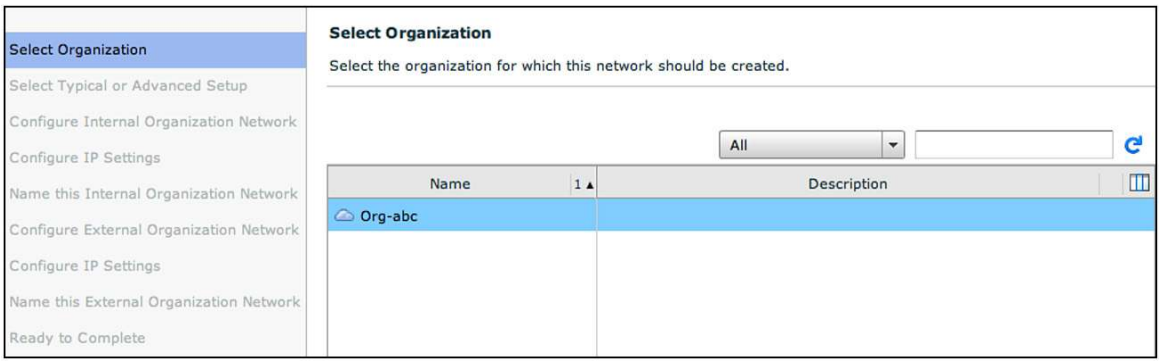

In this example, you will create both internal and external network for the organization (Figures 62 through 66).

Figure 62. Create Both Internal and External Networks

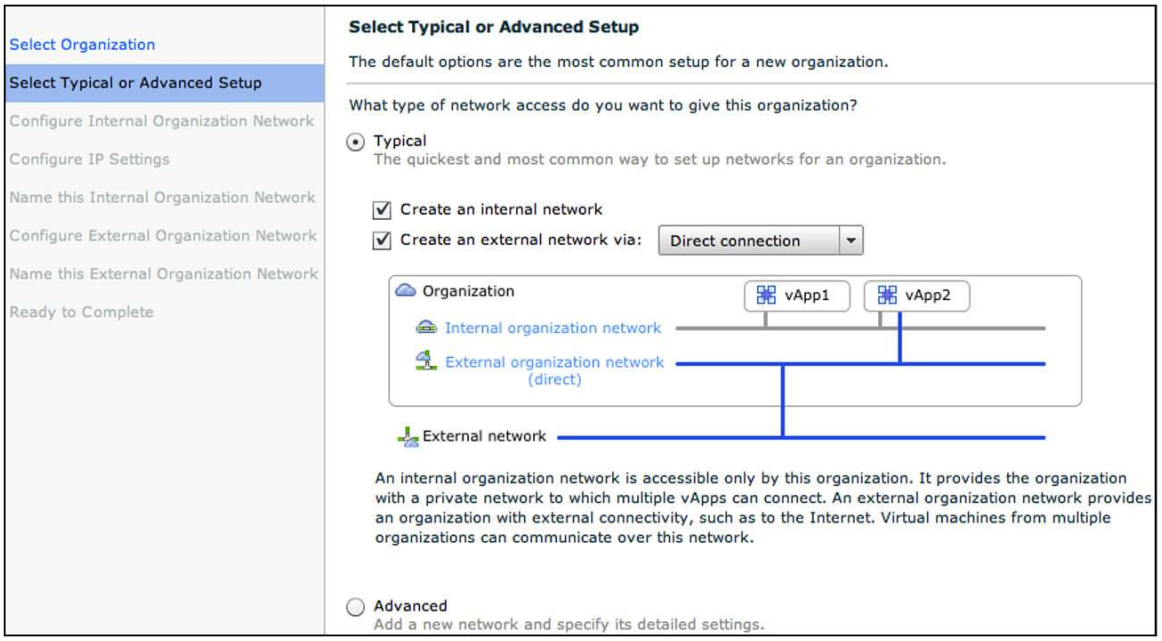

Figure 63. Configure IP Settings

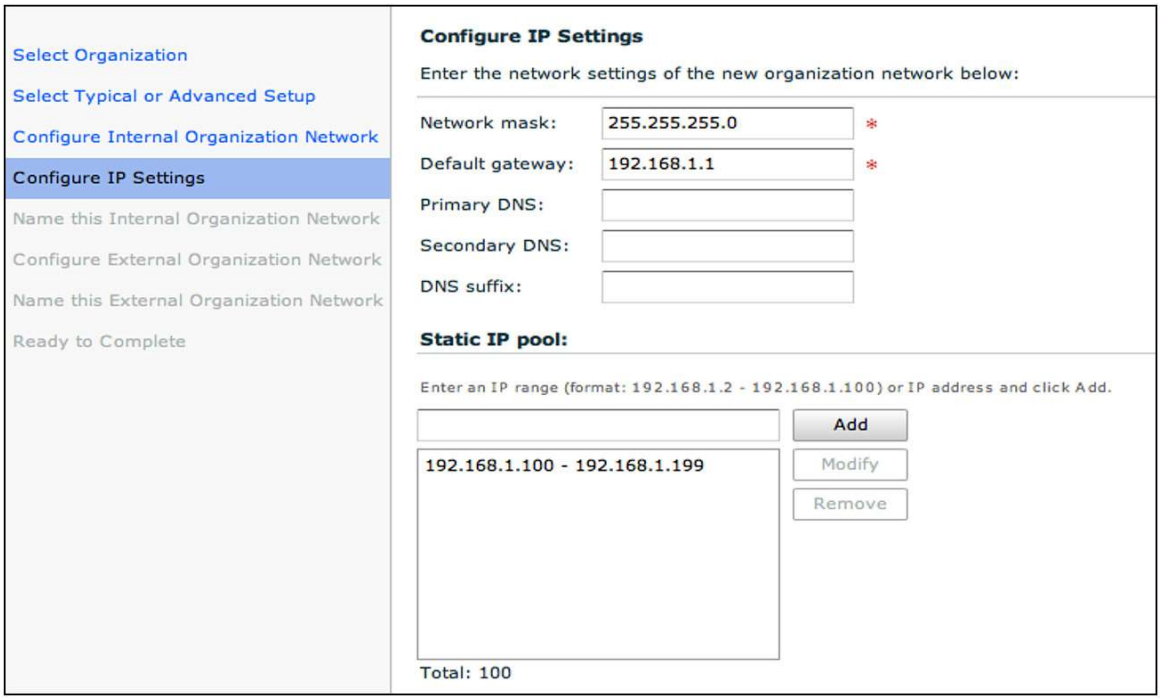

Figure 64. Name the Internal Organization

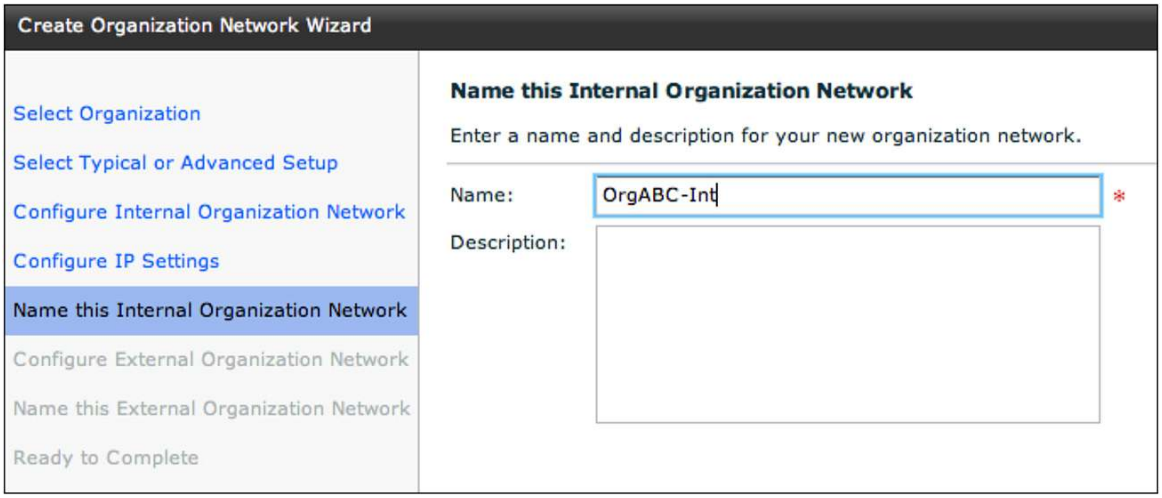

Figure 65. Select the Network for Organization External Connectivity

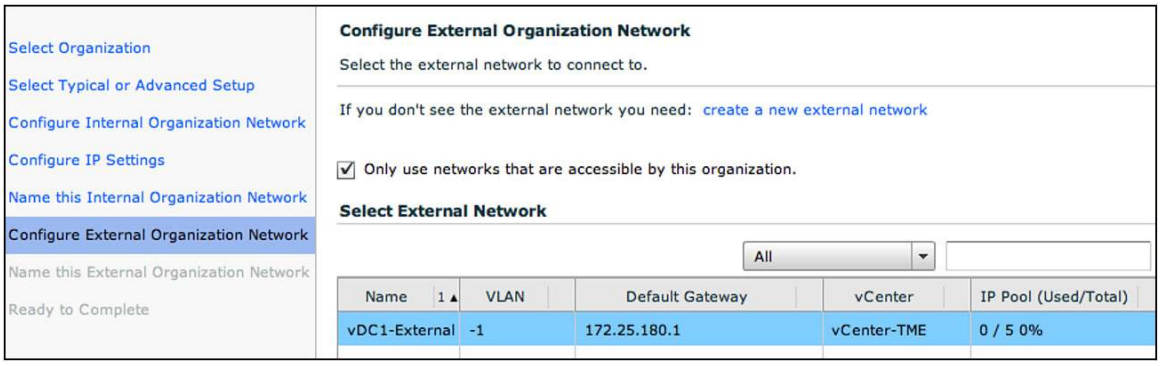

Figure 66. Organization Network Summary

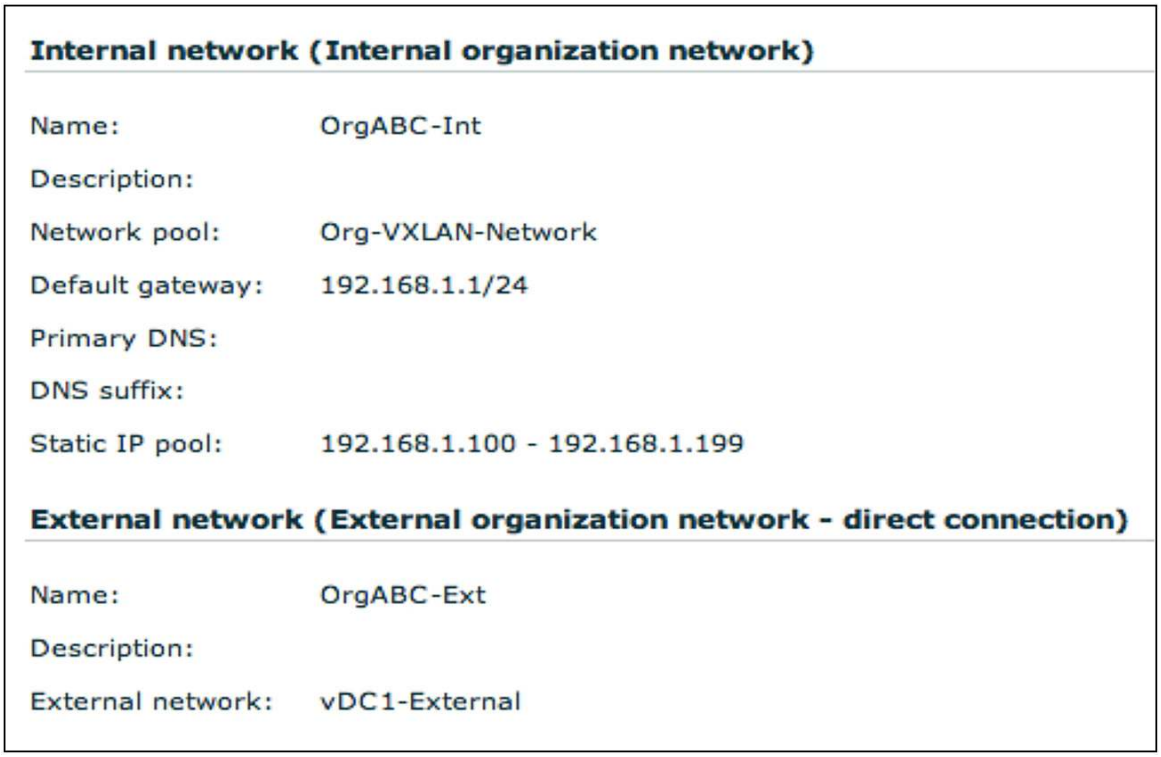

After creating the organization networks, we are ready to deploy a vApp that uses both internal and external organization networks. Figure 67 shows the vApp diagram which contains three virtual machines deployed. One of the virtual machine (client virtual machine) has network interfaces in both Internal and external network.

Figure 67. vApp Diagram Showing Two Organization Networks

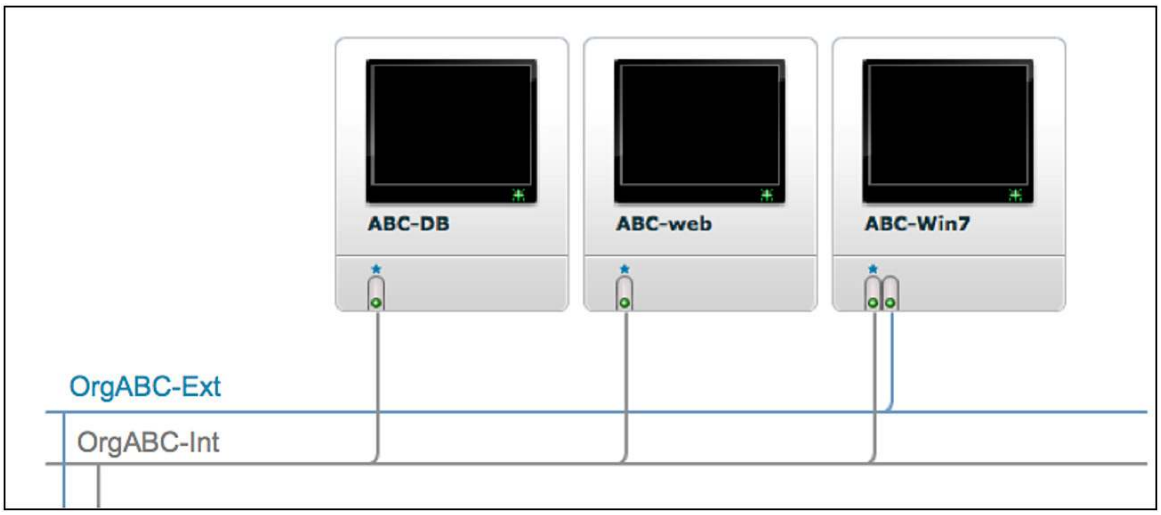

Integrating with VMware vCloud Director 5.1 and vShield Manager 5.1

Integrating VSM (Cisco Nexus 1000V Series) with VMware vShield Manager

This section discusses how to integrate VMware vShield Manager with the Cisco Nexus 1000V Series VSM. The following information is required to add the Cisco Nexus 1000V Series Switch as a managed switch in VMware vShield Manager:

- VSM connectivity details
- Multicast addresses
- Number of VXLANs

As shown in Figure 68, start by logging into the VMware vShield Manager web interface: https://vShield-Manager-IP.

Figure 68. VMware vShield Manager Login GUI

Choose Settings & Reports > Configurations > Networking, as shown in Figure 69. Select Add Switch Provider to configure the Cisco Nexus 1000V Series VSM settings.

Figure 69. Add Switch Provider in VMware vShield Manager

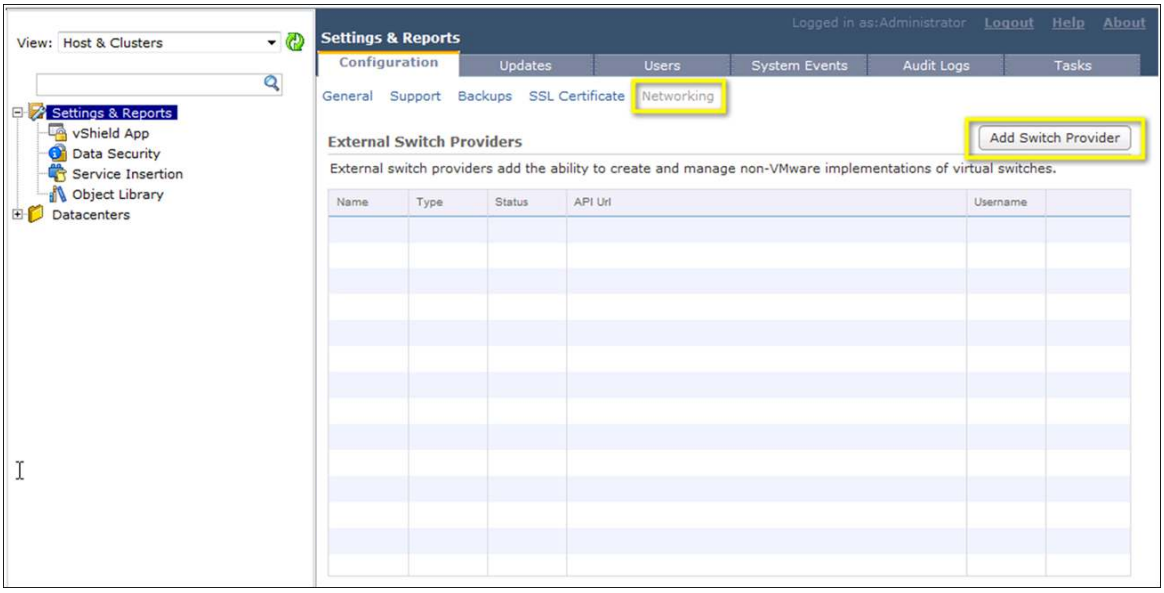

Enter the details for the Cisco Nexus 1000V Series NSM to register as a switch provider (Figure 70).

Figure 70. Cisco Nexus 1000V Series Switch Configuration

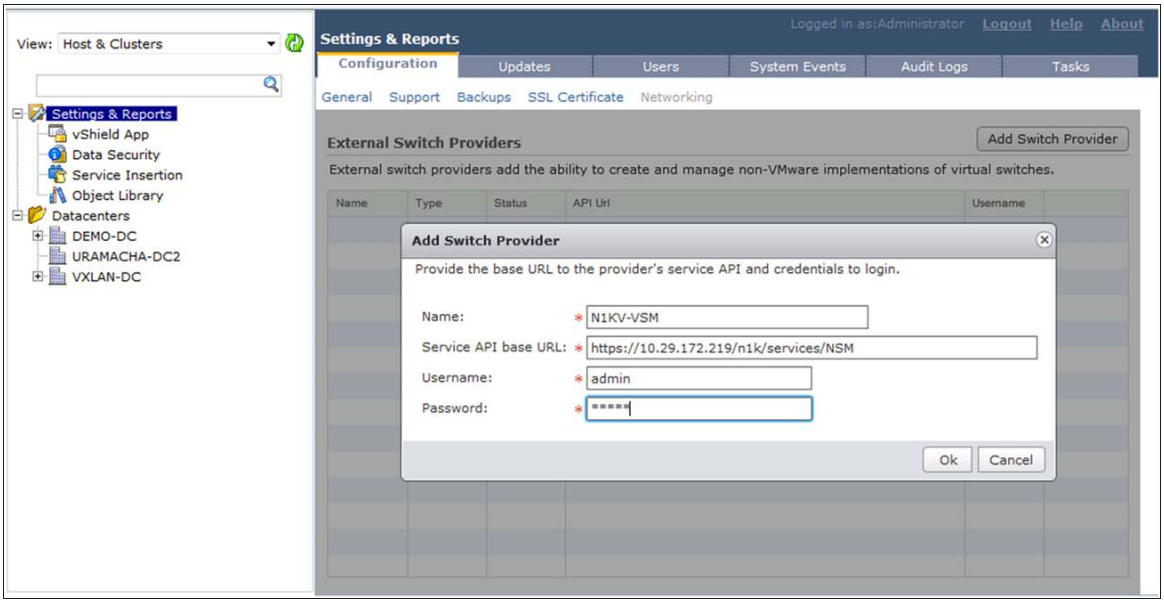

After the Cisco Nexus 1000V Series VSM is registered successfully, it will be added to the list of switch providers, and a green check mark will appear in the Status column (Figure 71).

Figure 71. Cisco Nexus 1000V Series Switch Added to VMware vShield Manager List of External Switch Providers

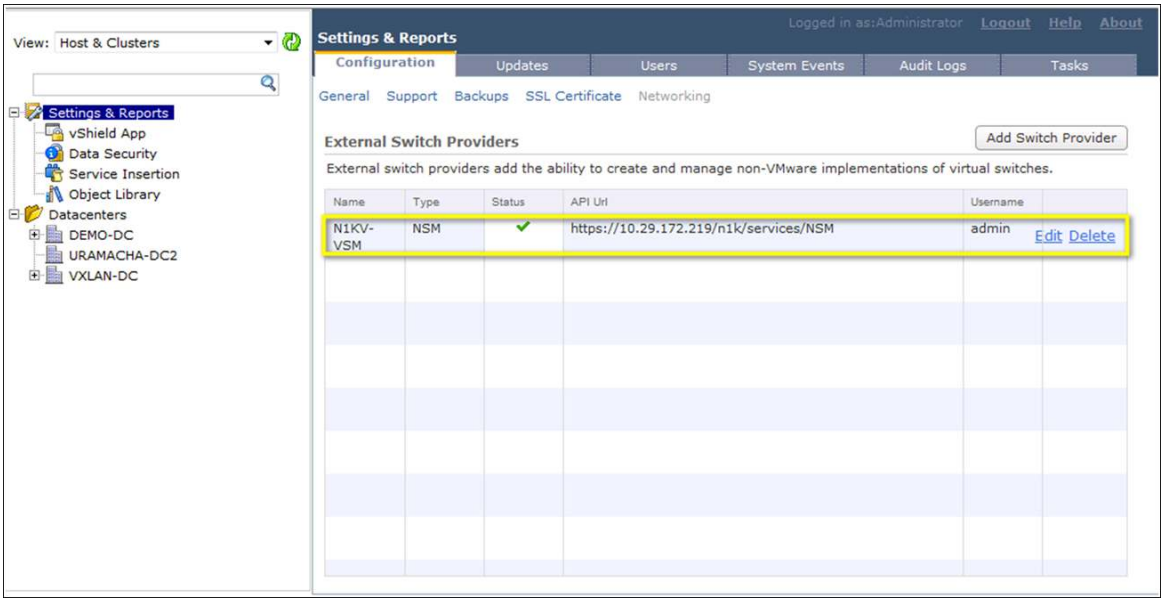

With VMware vShield Manager 5.1, the hosts and clusters need to be prepared for VXLAN by assigning a distributed virtual switch to a cluster. This distributed virtual switch will be the provider for the VXLAN network pool.

Select the data center from the list of data centers associated with the VMware vCenter server registered with VMware vShield Manager. Choose Network Virtualization > Preparation > Connectivity as shown in Figure 72.

Figure 72. Prepare Clusters for VXLAN Using Cisco Nexus 1000V Series

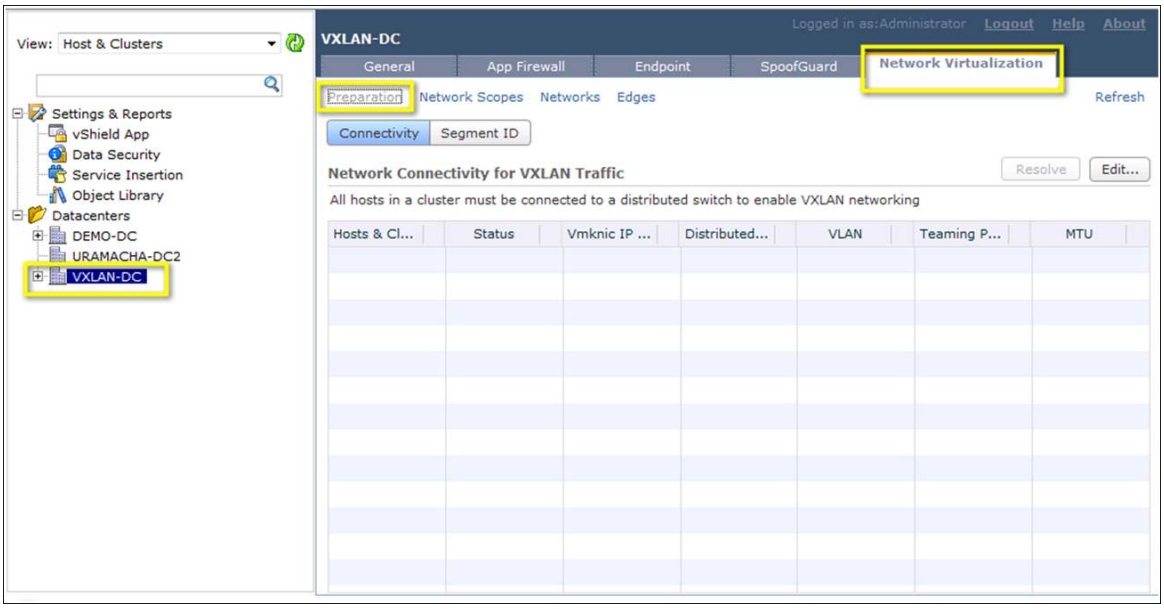

Click Edit to configure the clusters for VXLAN networking. In this example, one cluster is defined for the data center, and one is associated with the Cisco Nexus 1000V Series distributed virtual switch called VSM-CX. Select the cluster and the distributed switch from the drop-down menu and click Next (Figure 73).

Figure 73. Associate Cisco Nexus 1000V Series Switch with a Cluster

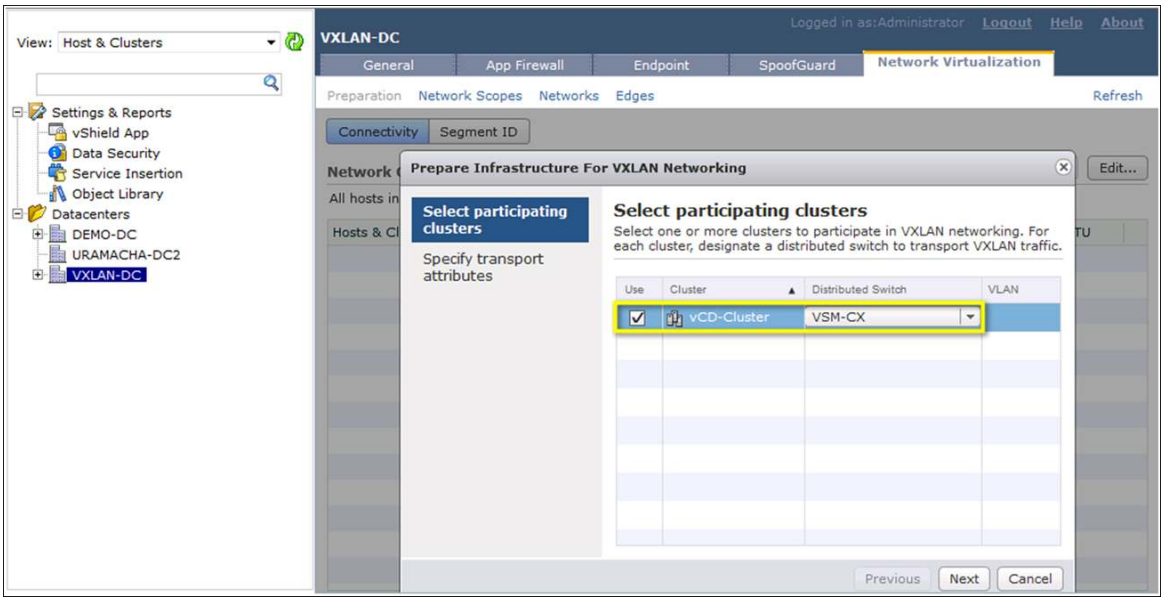

The transport attributes are not configurable when the Cisco Nexus 1000V Series Switch is selected as the distributed switch and will always be listed as Static EtherChannel. Click the Finish button to complete the configuration (Figure 74).

Figure 74. Configure Transport Attributes for VXLAN Traffic

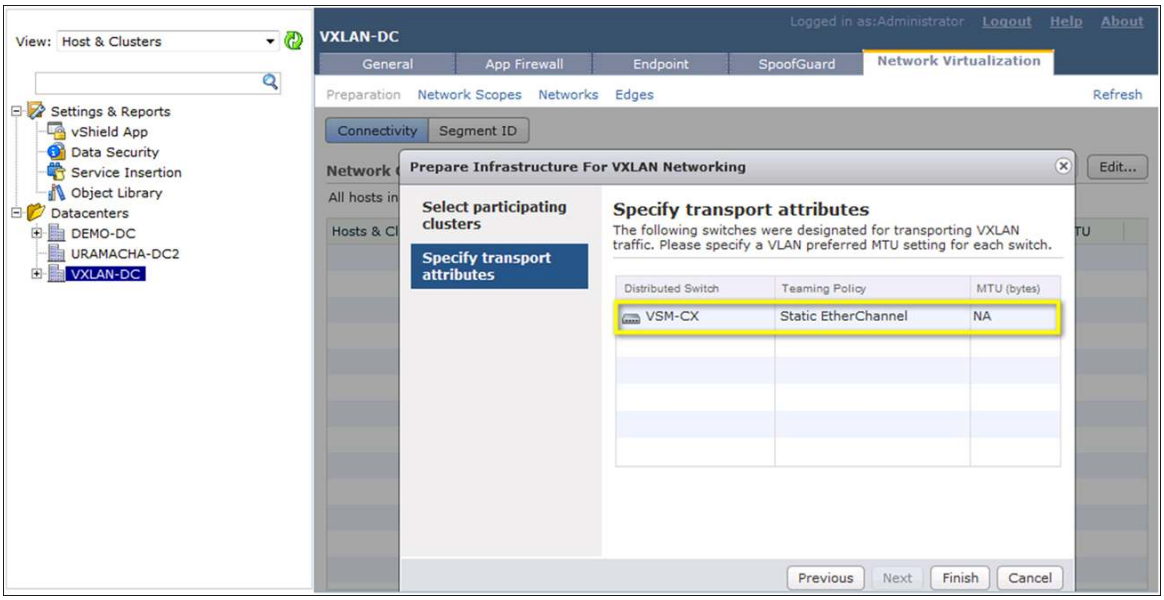

After all the hosts in the cluster are prepared for network connectivity through the Cisco Nexus 1000V Series distributed switch, the cluster will be marked Ready (Figure 75).

Figure 75. Network Connectivity for VXLAN Traffic Using Cisco Nexus 1000V Series Switch

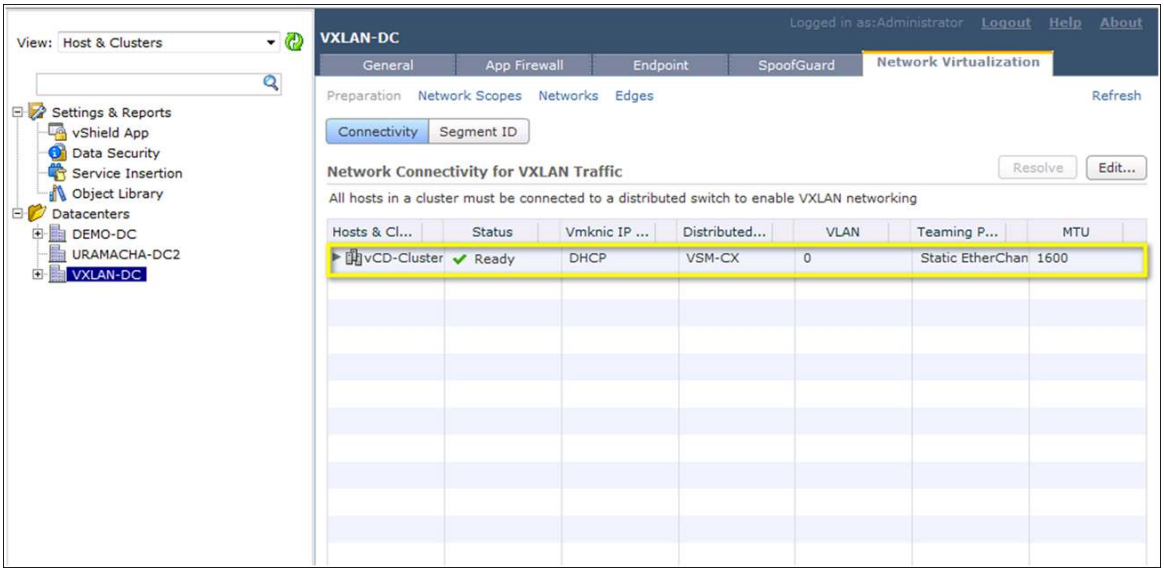

Next, provide the VXLAN ID range and multicast address range (Figure 76) by choosing Segment ID > Edit. Try to put as many multicast groups as available that will not exceed the state in the transit switches and routers. If you have multiple VMware vShield Managers, the VXLAN IDs and groups must not be duplicated among them.

Figure 76. Providing VXLAN ID Ranges and Multicast Address Ranges

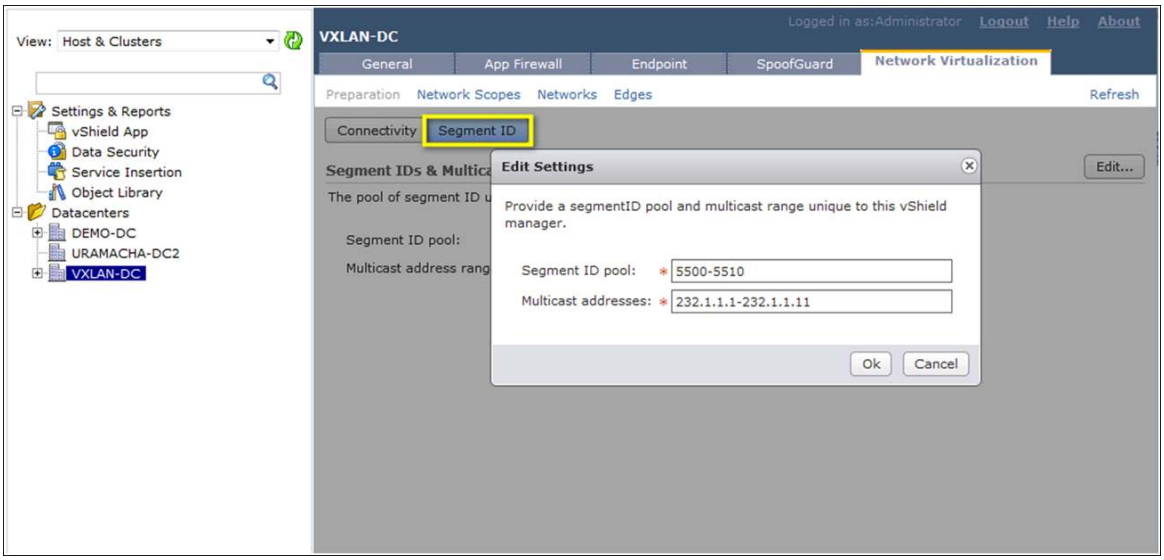

VMware vCloud Director Settings

This section focuses on the VMware vCloud Director network configuration in our use case. Please refer to the VMware documentation for settings other than network settings, including settings to create virtual data centers, organizations, organization virtual device contexts, and so on.

Creating a Provider vDC

Starting with VMware vCloud Director 5.1, the VXLAN network pool is created automatically when the provider vDC is configured. Configure the provider vDC as usual based on the VMware documentation.

The cluster that is chosen for the resource pool will determine which distributed virtual switch will be used to create the VXLAN network pool. Here the vDC cluster resource pool is selected; this pool was previously associated with VSM-CX, the distributed virtual switch, in VMware vShield Manager (Figure 77).

Figure 77. Add a Provider vDC

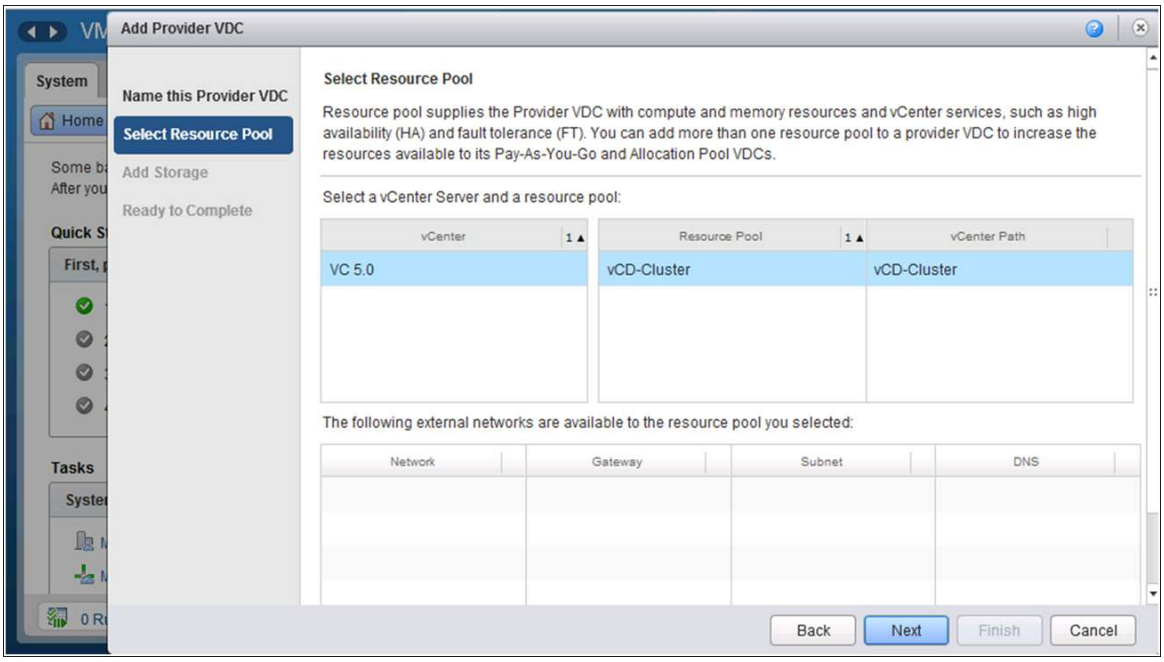

After the provider vDC is configured, you can verify the network pool by choosing Manage & Monitor > Network Pools. In Figure 78, you can see that the VXLAN pool is created and associated with the VSM-CX vDS.

Figure 78. VXLAN Pool Created and Associated with VSM-CX vDS

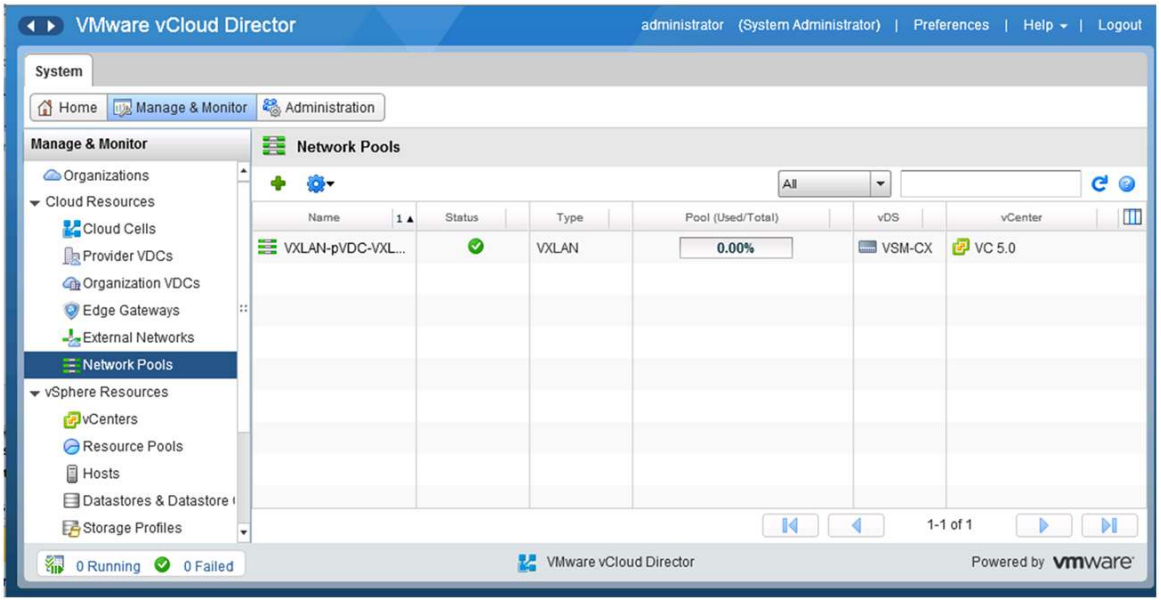

Building an External Network for Provider vDC

The external network provides northbound connectivity to an organization in VMware vCloud Director. This network will always be a port-group-backed network. There are no changes to the configuration steps for external networks for VMware vCloud Director 5.1.

Choose System > Home > Quick Start. Click External Network. Select the port group that will provide the external connectivity (Figure 79).

Figure 79. Select the Preconfigured Port Group (Port Profile)

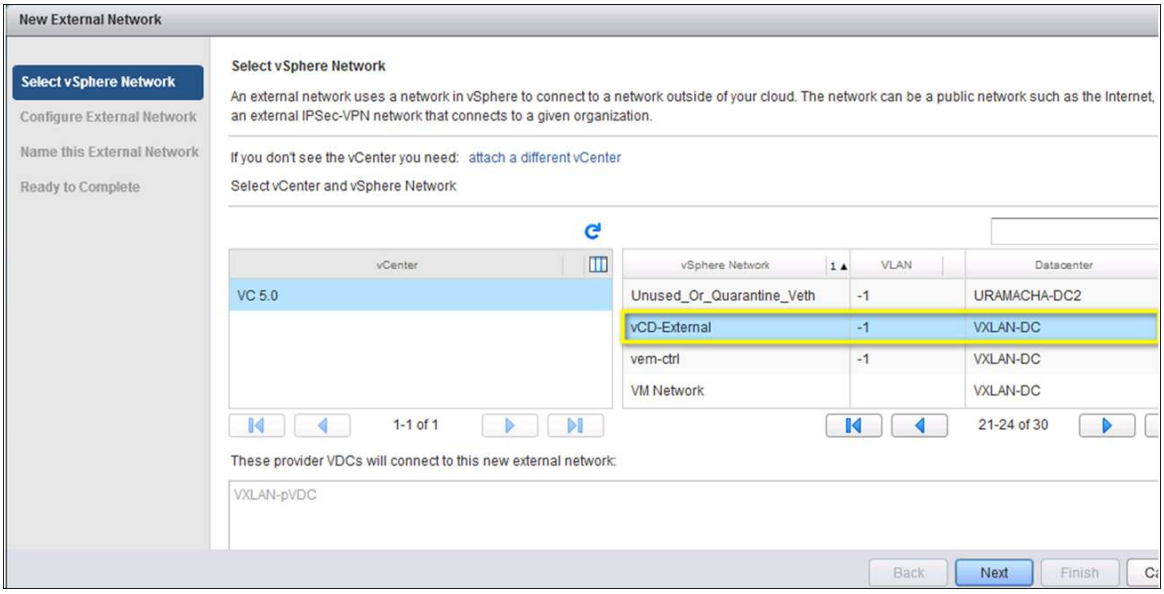

Next, specify a pool of IP addresses that can be consumed by the virtual machines requiring external connectivity. Click Add to add a new address range (Figure 80).

Figure 80. Network Settings

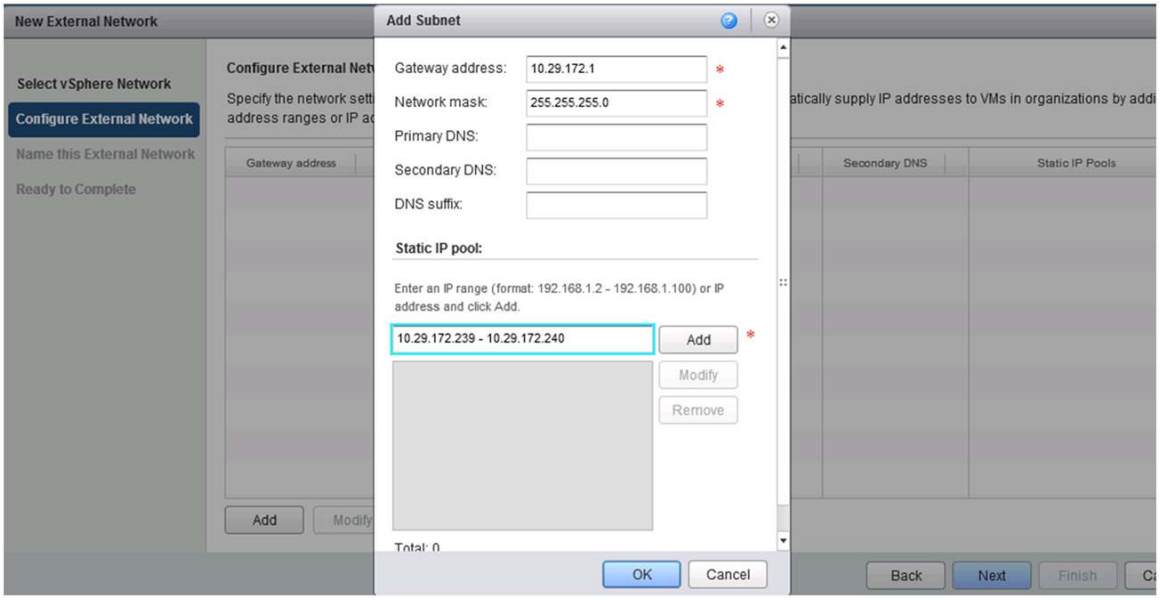

Provide the name for this external network (Figure 81).

Figure 81. Name the External Network

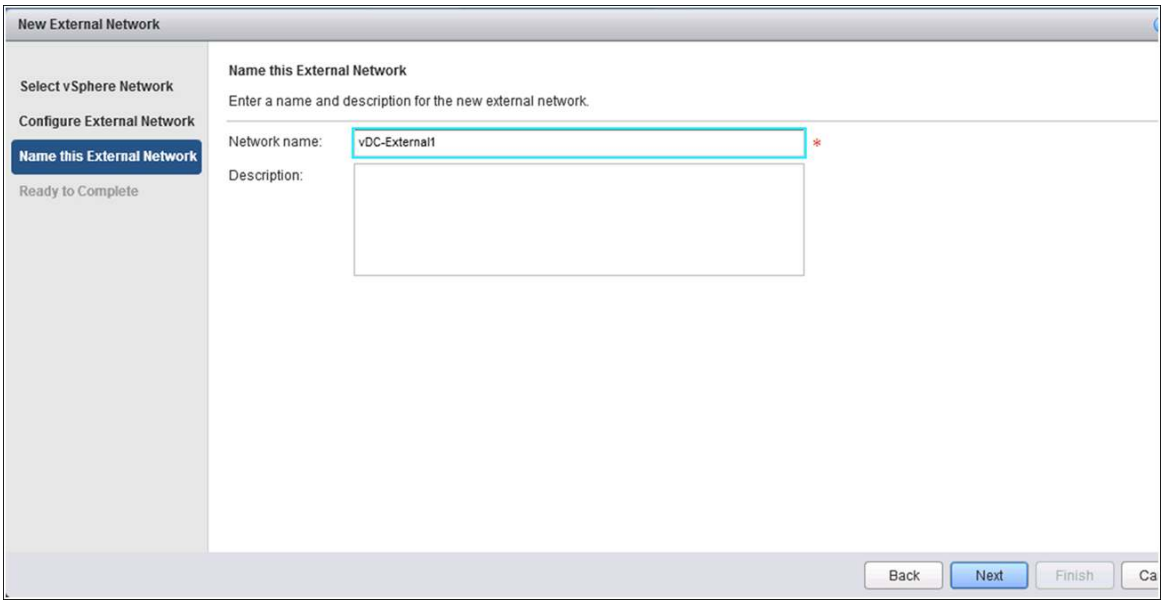

Verify the network settings and click Finish to complete the configuration (Figure 82).

Figure 82. Summary of the External Network

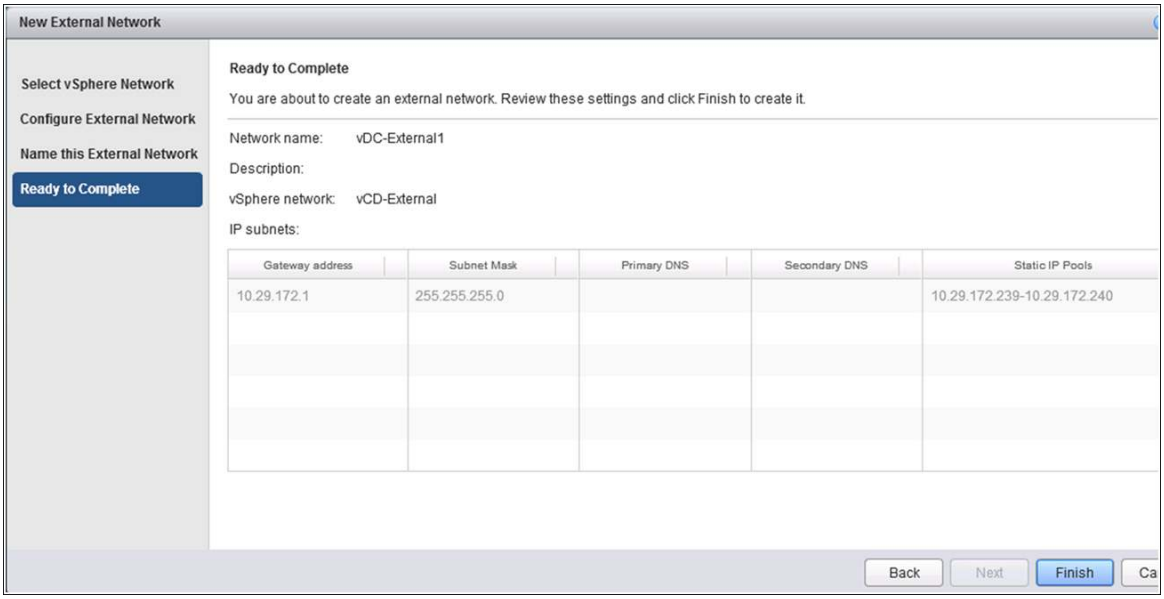

Assigning Network Resources to an Organization

After an organization is created, you allocate resources to it by choosing System > Home > Quick Start > Allocate resources to an organization. On the Select Network Pool & Services screen, the network pool that was automatically created when the provider vDC was created will appear in the drop-down menu. This is the VXLAN-backed network pool that is being supported by the Cisco Nexus 1000V Series DVS. Select this pool and assign a quota for the organization (Figure 83).

Figure 83. Select the VXLAN-Backed Network Pool for the Organization

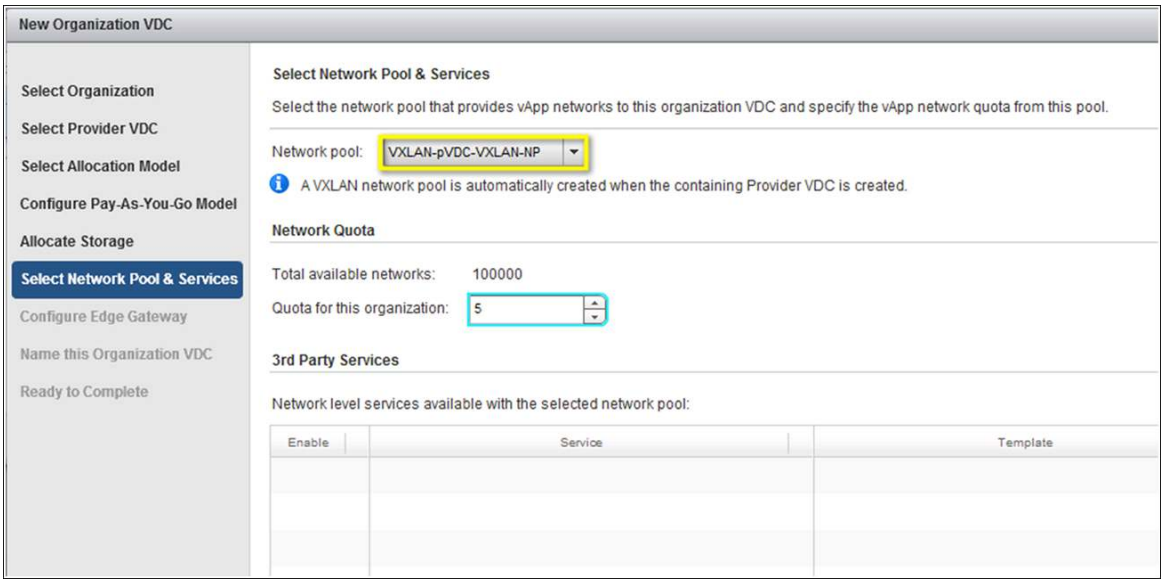

Configuring the Organization vDC

With VMware vCloud Director 5.1, an organization vDC (Org vDC) network model has replaced the organization network model. Org vDC networks tie the network resources to an organization. In this section, we will create an Org vDC and configure an isolated Org vDC network, which will consume Layer 2 network segments from the network pool. Choose Manage & Monitor > Organization vDCs and select the vDC that was created for the organization (Figure 84).

Figure 84. Select Organization vDC

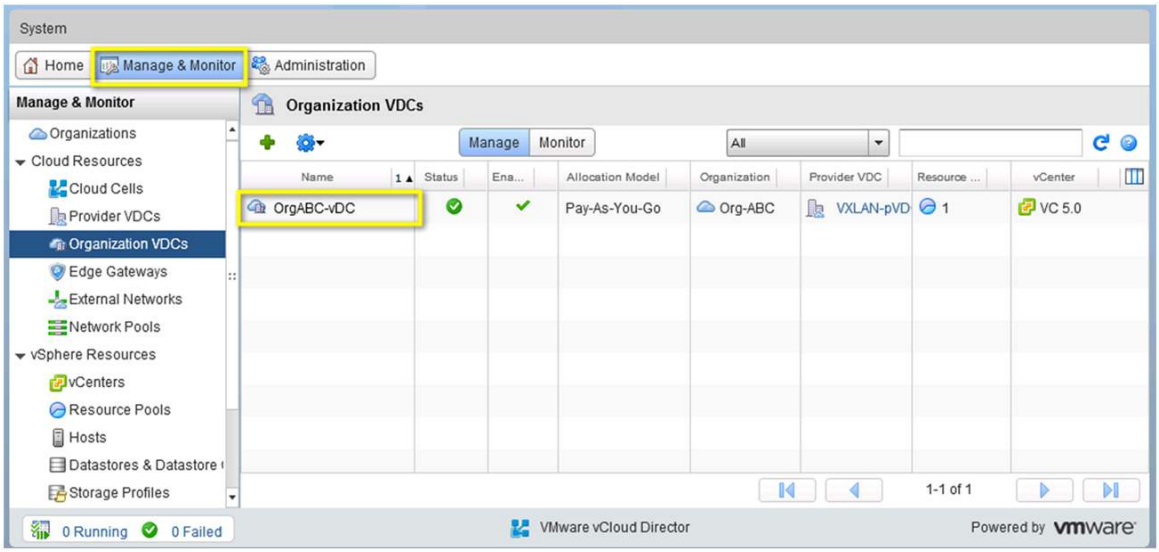

After the organization vDC is opened, navigate to the Org vDC Networks tab and click the “+” button to add a new network. In this example, we will create an isolated network for the internal communication among the web, client, and database virtual machines, and we will create an external network for external communication to the client virtual machine (Figures 85 through 88).

Figure 85. Create a New Isolated Org vDC Network

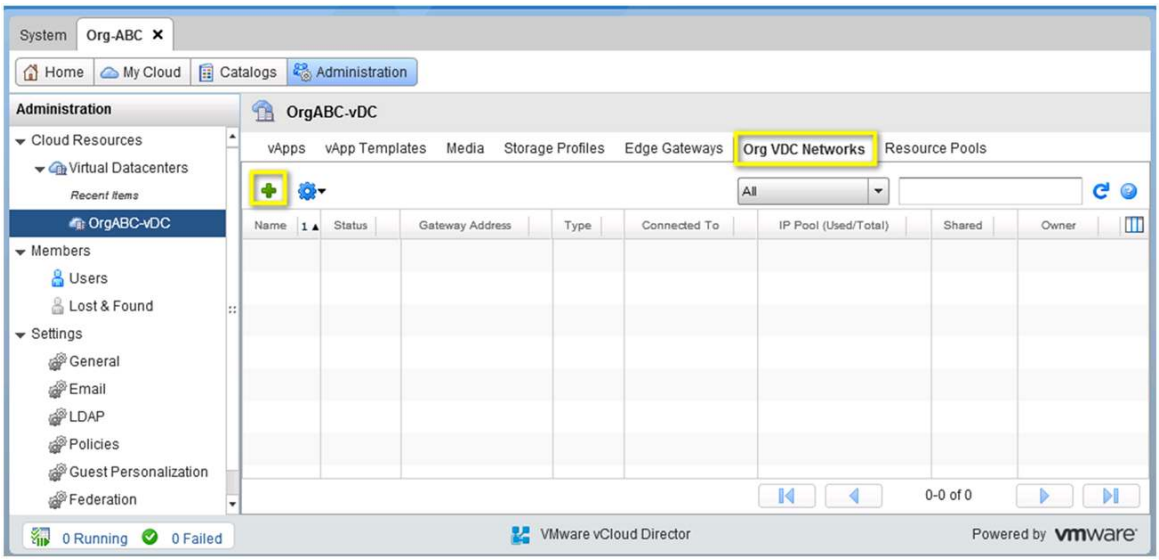

Figure 86. Select the Network Type

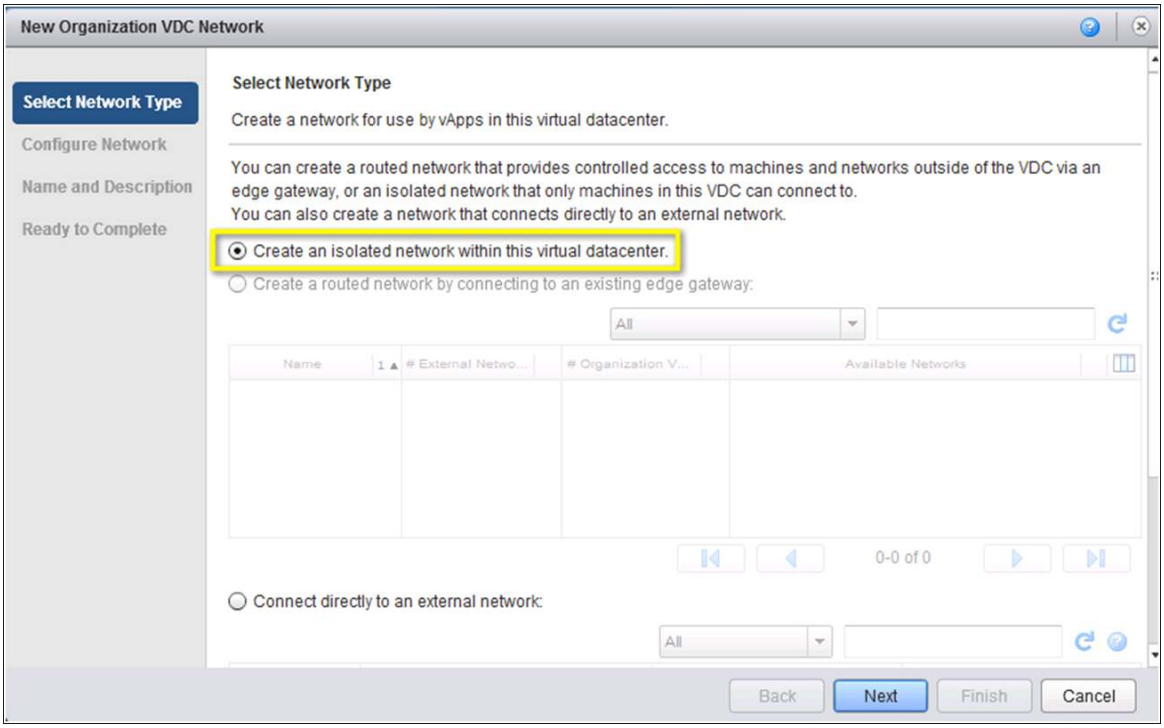

Figure 87. Configure the Network for the Isolated Org vDC Network

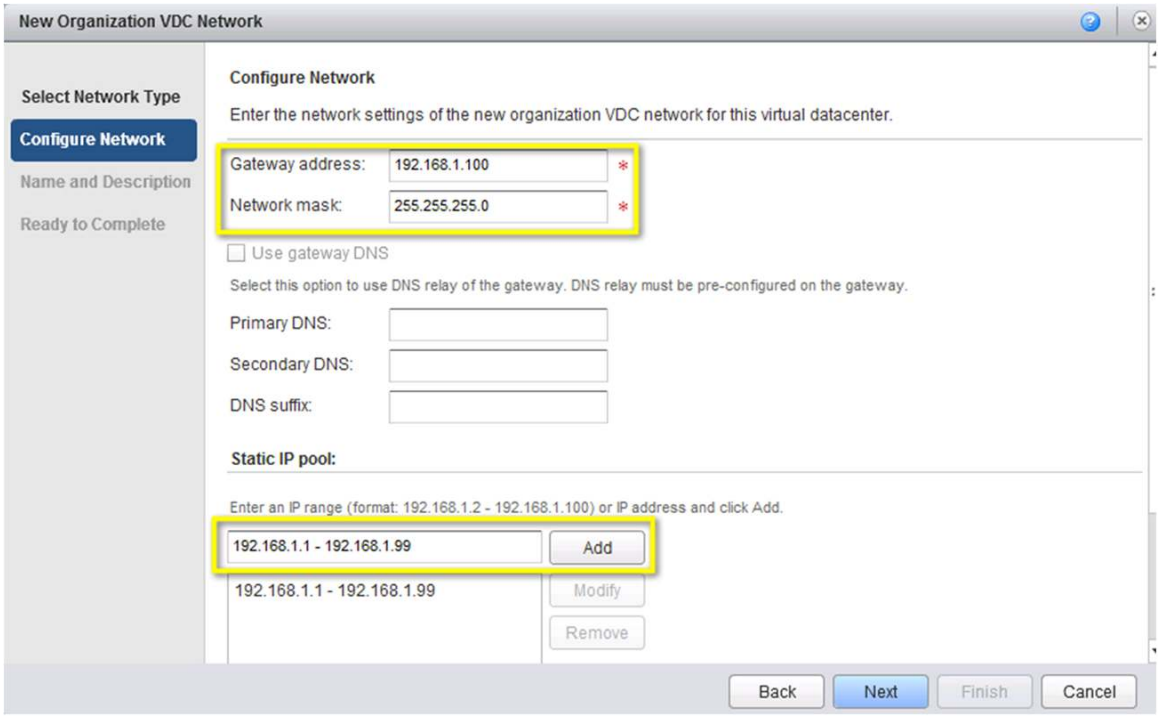

Figure 88. Name This Org vDC Network

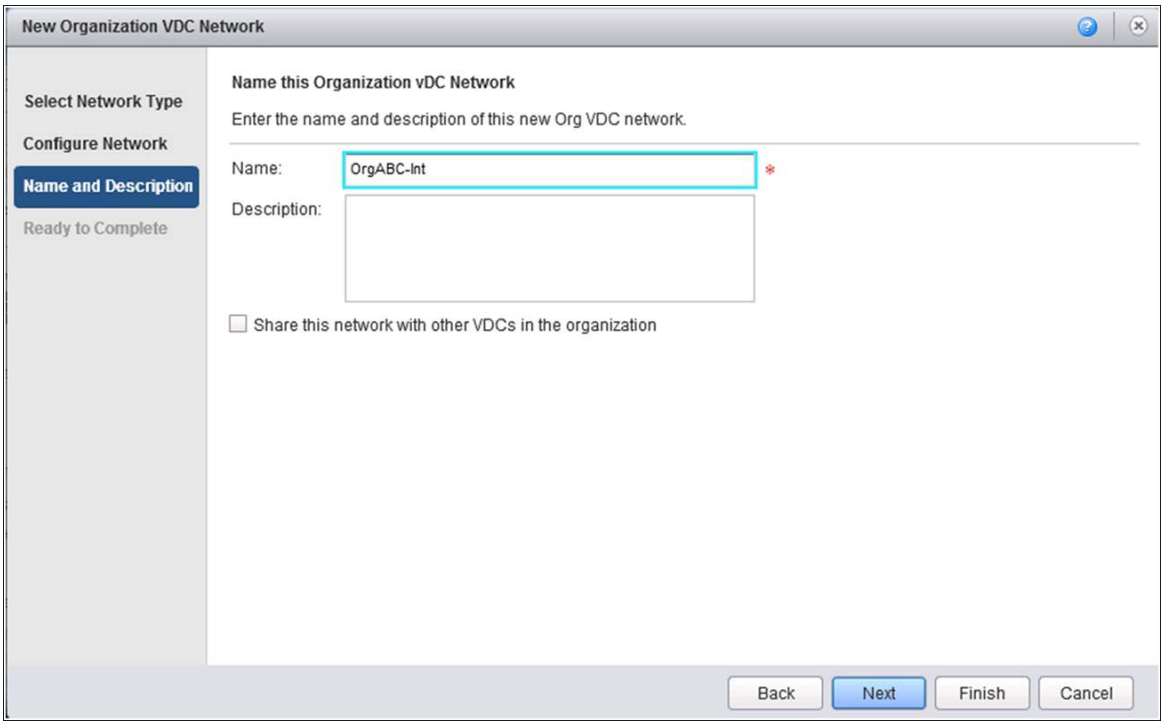

Now that isolated network has been created successfully, the next step is to create an external network that will use the vCD-External port profile configured on the Cisco Nexus 1000V Series Switch. Click on the “+” button to add another Org vDC network (Figure 89).

Figure 89. Create a New Isolated Org vDC Network

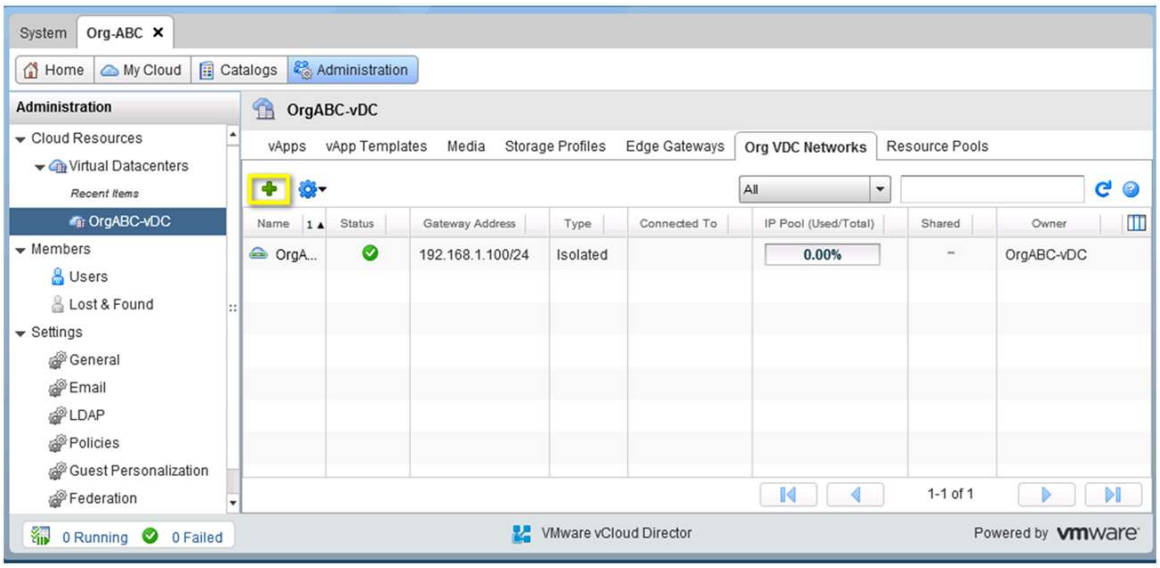

Select the option to connect directly to an external network. The available external networks in the provider vDC will be displayed. In this example, we have the vDC-External pool that is being supported by the vCD-External port profile on the Cisco Nexus 1000V Series Switch. The IP addresses for the external network will be assigned from the range that was provided in the configuration of the vDC-External network (Figures 90 and 91).

Figure 90. Select the Network Type

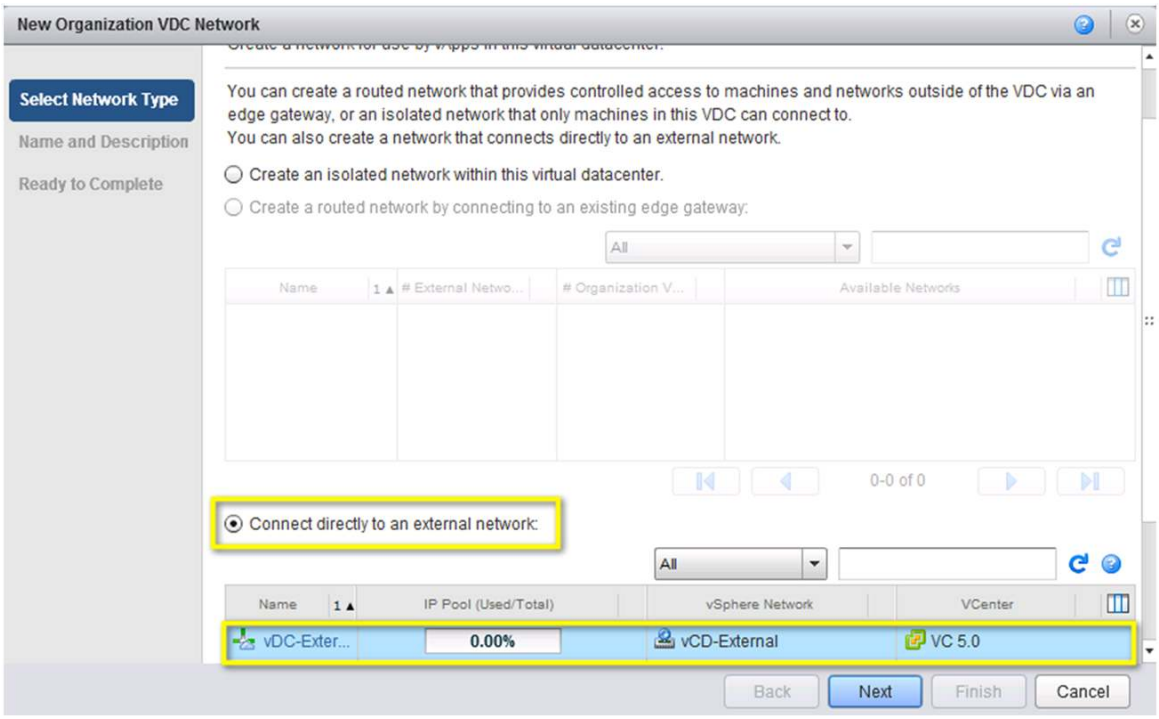

Figure 91. Name This Organization Network

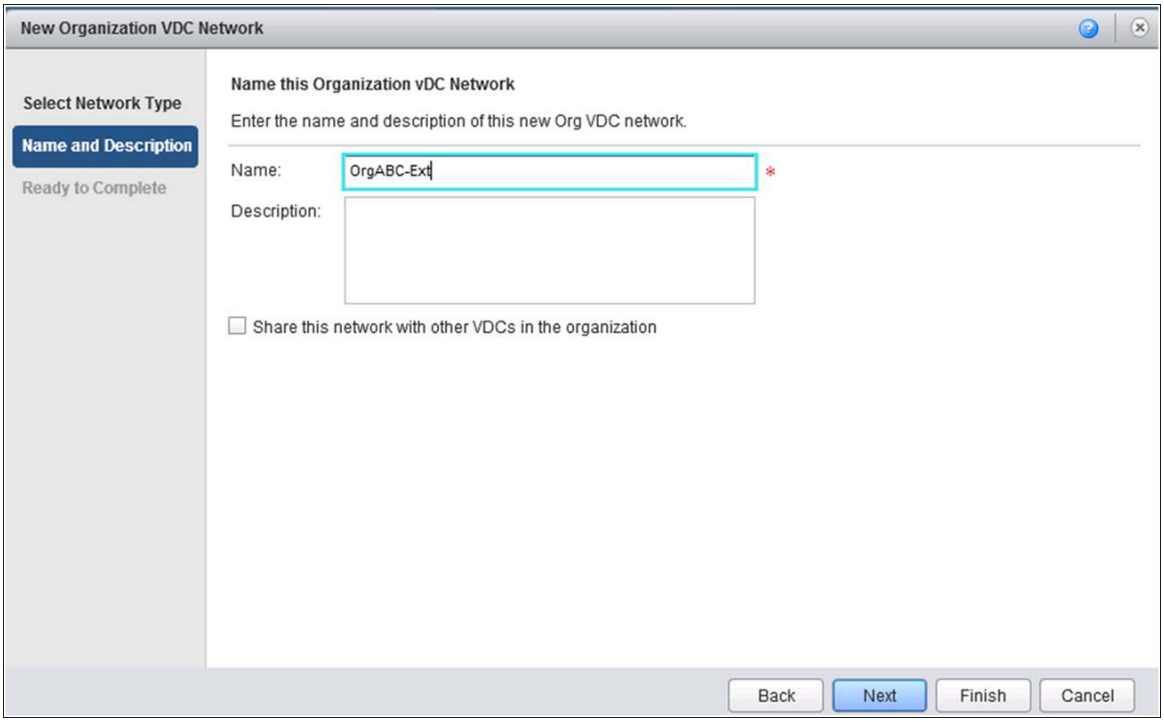

After creating the Org vDC networks, we are ready to deploy a vApp that uses both internal and external organization networks. Figure 92 shows a vApp that contains three deployed virtual machines. One of the virtual machines (the client virtual machine) has network interfaces in both the Internal and external networks.

Figure 92. vApp with Two Organization Networks

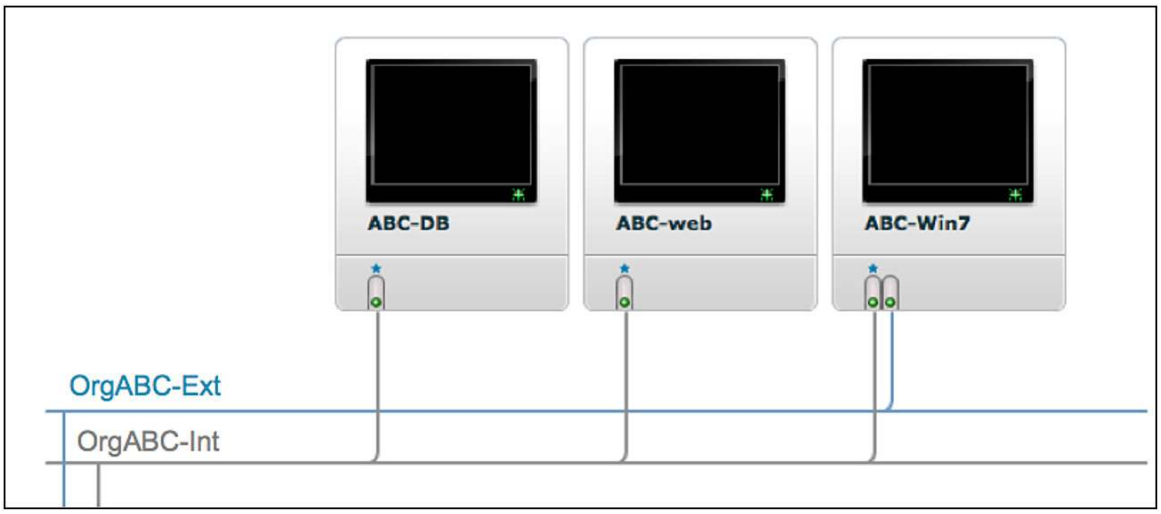

Applying Cisco Virtual Security Gateway Service with VXLAN and VMware vCloud Director

The Cisco Nexus 1000V Series with Cisco Virtual Services Data Path (vPath) makes it possible to configure network services for an organization network that is backed by a VXLAN pool in VMware vCloud Director. Cisco Virtual Security Gateway (VSG) is a virtual firewall for Cisco Nexus 1000V Series Switches that delivers security and compliance for virtual computing environments. In a VMware vCloud Director environment, Cisco VSG can be inserted to provide tenant-level security when the organization network is backed by a VXLAN pool provided by a Cisco Nexus 1000V Series Switch. The white paper Enable Cisco Virtual Security Gateway Service on a Virtual Extensible LAN Network in VMware vCloud Director describes how to deploy Cisco VSG in a VXLAN and VMware vCloud Director environment.

Conclusion

This guide demonstrated how to integrate the capabilities and features provided by the Cisco Nexus 1000V Series into a VMware vCloud Director environment. The examples showed the creation of external and organization networks. Both types of network used the Cisco Nexus 1000V Series port profiles and associated port groups to provide isolation and connectivity for internal networks and external organization networks with both the routed and direct connection profiles. The same concepts and capabilities translate directly to the third type of VMware vCloud Director network - the vApp network. The vApp network type is, in terms of connectivity profiles, functionally equivalent to the organization network and can be created as an internal vApp network or external vApp network, with the latter connecting to an organization network using either the routed or direct connectivity profile. All network types and connectivity profiles consume and manage the isolation of the Cisco Nexus 1000V Series port profiles and VLAN and VXLAN isolation.

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

Glossary

VMware vCenter

VMware vCenter provides centralized control and visibility to VMware vSphere virtual infrastructure. The Cisco Nexus 1000V Series is tightly integrated with VMware vCenter. This integration enables the network administrator and the server administrator to collaborate efficiently without each having to learn a different management tool. The network administrator uses the Cisco NX-OS CLI on the VSM, and the server administrator continues to use VMware vCenter.

VMware vCloud Director

VMware vCloud Director is a cloud computing management platform. It abstracts the virtualized resources to enable users to gain self-service access to them through a services catalogue.

VMware vShield Manager

VMware vShield Manager provides a central point of control for managing vShield products. For the purposes of this document, vShield Manager is acting as an integration point between Cisco Nexus 1000V and vCloud Director via Cisco Nexus 1000V Series.

VMware vShield Edge

VMware vShield Edge is an edge gateway firewall providing policy enforcement, VPN and NAT capabilities for multitenant hosting services

Cisco Nexus 1000V Series Switches

Cisco Nexus 1000V Series VSM The Cisco Nexus 1000V Series VSM controls multiple VEMs as one logical modular switch. Instead of physical line-card modules, the VSM supports multiple VEMs running in software inside the physical servers

Cisco Nexus 1000V Series Virtual Ethernet Module

The Cisco Nexus 1000V Series VEM runs as part of the VMware ESX or ESXi kernel and replaces the VMware virtual switch feature

For More Information

For more information about the Cisco Nexus 1000V Series, please refer to the following URLs:

- Cisco Nexus 1000V Series product information: http://www.cisco.com/go/1000v
- Cisco Nexus 1000V Series technical documentation: http://www.cisco.com/go/1000vdocs
- Cisco Nexus 1000V community: http://www.cisco.com/go/1000vcommunity
- Cisco VSG: http://www.cisco.com/en/US/partner/products/ps11208/index.html
- VMware vCloud Director: http://www.vmware.com/products/vcloud-director
- VMware vSphere: http://www.vmware.com/go/vsphere

- Deployment guide for Cisco Nexus 1000V Series Switches: http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/guide_c07-556626.html
- Enable Cisco Virtual Security Gateway Service on a Virtual Extensible LAN Network in VMware vCloud Director: http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c11-715721.html

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 C07-703595-01 02/13