

Cross-vCenter NSX Installation Guide

Update 3

Modified on 20 NOV 2017

VMware NSX Data Center for vSphere 6.2



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2010 – 2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1 Cross-vCenter Installation Guide 5**
- 2 Overview of NSX 6**
 - NSX Components 7
 - NSX Edge 10
 - NSX Services 13
- 3 Overview of Cross-vCenter Networking and Security 15**
 - Benefits of Cross-vCenter NSX 15
 - How Cross-vCenter NSX Works 16
 - Support Matrix for NSX Services in Cross-vCenter NSX 17
 - Universal Controller Cluster 18
 - Universal Transport Zone 19
 - Universal Logical Switches 19
 - Universal Logical (Distributed) Routers 19
 - Universal Firewall Rules 20
 - Universal Network and Security Objects 20
 - Cross-vCenter NSX Topologies 21
 - Modifying NSX Manager Roles 24
- 4 Preparing for Installation 26**
 - System Requirements for NSX 26
 - Ports and Protocols Required by NSX 28
 - NSX and vSphere Distributed Switches 30
 - Example: Working with a vSphere Distributed Switch 32
 - NSX Installation Workflow and Sample Topology 40
 - Cross-vCenter NSX and Enhanced Linked Mode 42
- 5 Tasks for Primary and Secondary NSX Managers 43**
 - Install the NSX Manager Virtual Appliance 43
 - Configure Single Sign On 48
 - Register vCenter Server with NSX Manager 50
 - Specify a Syslog Server 54
 - Install and Assign NSX for vSphere License 55
 - Exclude Virtual Machines from Firewall Protection 57
- 6 Configuring the Primary NSX Manager 59**
 - Deploy NSX Controllers on the Primary NSX Manager 59

- Prepare Hosts on the Primary NSX Manager 64
- Configure VXLAN from the Primary NSX Manager 68
- Assign a Segment ID Pool to the Primary NSX Manager 72
- Assign Primary Role to NSX Manager 75
- Assign a Universal Segment ID Pool on the Primary NSX Manager 76
- Add a Universal Transport Zone on the Primary NSX Manager 77
- Add a Universal Logical Switch on the Primary NSX Manager 78
- Connect Virtual Machines to a Logical Switch 81
- Add a Universal Logical (Distributed) Router on the Primary NSX Manager 81

7 Configuring Secondary NSX Managers 93

- Add a Secondary NSX Manager 93
- Prepare Hosts on a Secondary NSX Manager 95
- Configure VXLAN from the Secondary NSX Manager 96
- Assign a Segment ID Pool for a Secondary NSX Manager 97
- Add Clusters to the Universal Transport Zone 98

8 After Configuring Primary and Secondary NSX Managers 100

- Add a Transport Zone 100
- Add a Logical Switch 105
- Add a Logical (Distributed) Router 112
- Add an Edge Services Gateway 125
- Add a Host to a Prepared Cluster 135
- Remove a Host from an NSX Prepared Cluster 135

9 Uninstalling NSX Components 137

- Uninstall an NSX Edge Services Gateway or a Distributed Logical Router 137
- Uninstall a Logical Switch 137
- Safely Remove an NSX Installation 138

Cross-vCenter Installation Guide

This manual, describes how to install the cross-vCenter NSX environment. The information includes step-by-step configuration instructions, and suggested best practices.

Intended Audience

This manual is intended for anyone who wants to install cross-vCenter NSX. The information in this manual is written for experienced system administrators who are familiar with virtual machine technology and virtual datacenter operations. This manual assumes familiarity with VMware Infrastructure 4.x, including VMware ESX, vCenter Server, and the vSphere Client.

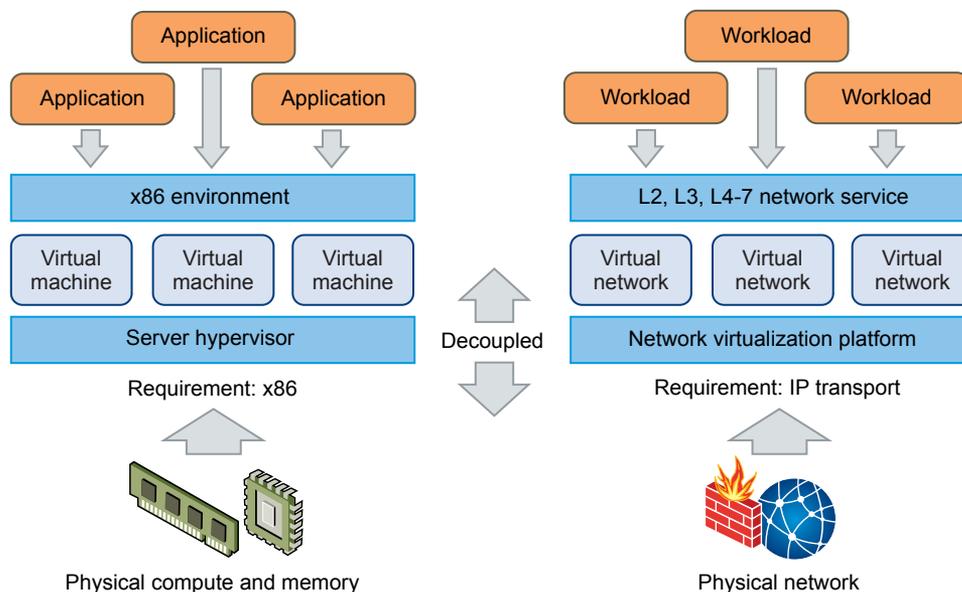
VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Overview of NSX

IT organizations have gained significant benefits as a direct result of server virtualization. Server consolidation reduced physical complexity, increased operational efficiency and the ability to dynamically re-purpose underlying resources to quickly and optimally meet the needs of increasingly dynamic business applications.

VMware's Software Defined Data Center (SDDC) architecture is now extending virtualization technologies across the entire physical data center infrastructure. VMware NSX[®], the network virtualization platform, is a key product in the SDDC architecture. With NSX, virtualization delivers for networking what it has already delivered for compute and storage. In much the same way that server virtualization programmatically creates, snapshots, deletes and restores software-based virtual machines (VMs), NSX network virtualization programmatically creates, snapshots, deletes, and restores software-based virtual networks. The result is a completely transformative approach to networking that not only enables data center managers to achieve orders of magnitude better agility and economics, but also allows for a vastly simplified operational model for the underlying physical network. With the ability to be deployed on any IP network, including both existing traditional networking models and next-generation fabric architectures from any vendor, NSX is a completely non-disruptive solution. In fact, with NSX, the physical network infrastructure you already have is all you need to deploy a software-defined data center.



The figure above draws an analogy between compute and network virtualization. With server virtualization, a software abstraction layer (server hypervisor) reproduces the familiar attributes of an x86 physical server (for example, CPU, RAM, Disk, NIC) in software, allowing them to be programmatically assembled in any arbitrary combination to produce a unique VM in a matter of seconds.

With network virtualization, the functional equivalent of a network hypervisor reproduces the complete set of Layer 2 through Layer 7 networking services (for example, switching, routing, access control, firewalling, QoS, and load balancing) in software. As a result, these services can be programmatically assembled in any arbitrary combination, to produce unique, isolated virtual networks in a matter of seconds.

With network virtualization, benefits similar to server virtualization are derived. For example, just as VMs are independent of the underlying x86 platform and allow IT to treat physical hosts as a pool of compute capacity, virtual networks are independent of the underlying IP network hardware and allow IT to treat the physical network as a pool of transport capacity that can be consumed and repurposed on demand. Unlike legacy architectures, virtual networks can be provisioned, changed, stored, deleted, and restored programmatically without reconfiguring the underlying physical hardware or topology. By matching the capabilities and benefits derived from familiar server and storage virtualization solutions, this transformative approach to networking unleashes the full potential of the software-defined data center.

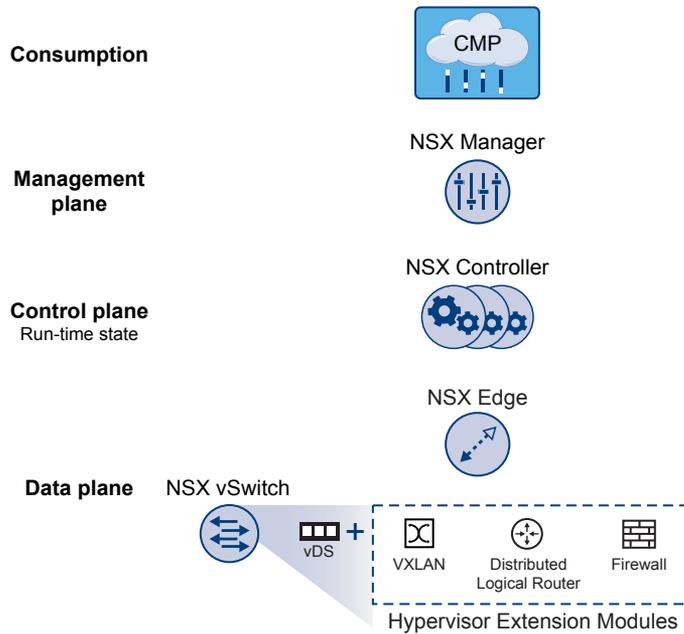
NSX can be configured through the vSphere Web Client, a command-line interface (CLI), and a REST API.

This chapter includes the following topics:

- [NSX Components](#)
- [NSX Edge](#)
- [NSX Services](#)

NSX Components

This section describes the components of the NSX solution.



Note that a cloud management platform (CMP) is not a component of NSX, but NSX provides integration into virtually any CMP via the REST API and out-of-the-box integration with VMware CMPs.

Data Plane

The NSX data plane consists of the NSX vSwitch, which is based on the vSphere Distributed Switch (VDS) with additional components to enable services. NSX kernel modules, userspace agents, configuration files, and install scripts are packaged in VIBs and run within the hypervisor kernel to provide services such as distributed routing and logical firewall and to enable VXLAN bridging capabilities.

The NSX vSwitch (vDS-based) abstracts the physical network and provides access-level switching in the hypervisor. It is central to network virtualization because it enables logical networks that are independent of physical constructs, such as VLANs. Some of the benefits of the vSwitch are:

- Support for overlay networking with protocols (such as VXLAN) and centralized network configuration. Overlay networking enables the following capabilities:
 - Reduced use of VLAN IDs in the physical network.
 - Creation of a flexible logical Layer 2 (L2) overlay over existing IP networks on existing physical infrastructure without the need to re-architect any of the data center networks
 - Provision of communication (east–west and north–south), while maintaining isolation between tenants
 - Application workloads and virtual machines that are agnostic of the overlay network and operate as if they were connected to a physical L2 network
- Facilitates massive scale of hypervisors
- Multiple features—such as Port Mirroring, NetFlow/IPFIX, Configuration Backup and Restore, Network Health Check, QoS, and LACP—provide a comprehensive toolkit for traffic management, monitoring, and troubleshooting within a virtual network

The logical routers can provide L2 bridging from the logical networking space (VXLAN) to the physical network (VLAN).

The gateway device is typically an NSX Edge virtual appliance. NSX Edge offers L2, L3, perimeter firewall, load balancing, and other services such as SSL VPN and DHCP.

Control Plane

The NSX control plane runs in the NSX Controller cluster. NSX Controller is an advanced distributed state management system that provides control plane functions for NSX logical switching and routing functions. It is the central control point for all logical switches within a network and maintains information about all hosts, logical switches (VXLANs), and distributed logical routers.

The controller cluster is responsible for managing the distributed switching and routing modules in the hypervisors. The controller does not have any dataplane traffic passing through it. Controller nodes are deployed in a cluster of three members to enable high-availability and scale. Any failure of the controller nodes does not impact any data-plane traffic.

NSX Controllers work by distributing network information to hosts. To achieve a high level of resiliency the NSX Controller is clustered for scale out and HA. NSX Controllers must be deployed in a three-node cluster. The three virtual appliances provide, maintain, and update the state of all network functioning within the NSX domain. NSX Manager is used to deploy NSX Controller nodes.

The three NSX Controller nodes form a control cluster. The controller cluster requires a quorum (also called a majority) in order to avoid a "split-brain scenario." In a split-brain scenario, data inconsistencies originate from the maintenance of two separate data sets that overlap. The inconsistencies can be caused by failure conditions and data synchronization issues. Having three controller nodes ensures data redundancy in case of failure of one NSX Controller node.

A controller cluster has several roles, including:

- API provider
- Persistence server
- Switch manager
- Logical manager
- Directory server

Each role has a master controller node. If a master controller node for a role fails, the cluster elects a new master for that role from the available NSX Controller nodes. The new master NSX Controller node for that role reallocates the lost portions of work among the remaining NSX Controller nodes.

NSX supports three logical switch control plane modes: multicast, unicast and hybrid. Using a controller cluster to manage VXLAN-based logical switches eliminates the need for multicast support from the physical network infrastructure. You don't have to provision multicast group IP addresses, and you also don't need to enable PIM routing or IGMP snooping features on physical switches or routers. Thus, the unicast and hybrid modes decouple NSX from the physical network. VXLANs in unicast control-plane mode do not require the physical network to support multicast in order to handle the broadcast, unknown

unicast, and multicast (BUM) traffic within a logical switch. The unicast mode replicates all the BUM traffic locally on the host and requires no physical network configuration. In the hybrid mode, some of the BUM traffic replication is offloaded to the first hop physical switch to achieve better performance. Hybrid mode requires IGMP snooping on the first-hop switch and access to an IGMP querier in each VTEP subnet.

Management Plane

The NSX management plane is built by the NSX Manager, the centralized network management component of NSX. It provides the single point of configuration and REST API entry-points.

The NSX Manager is installed as a virtual appliance on any ESX™ host in your vCenter Server environment. NSX Manager and vCenter have a one-to-one relationship. For every instance of NSX Manager, there is one vCenter Server. This is true even in a cross-vCenter NSX environment.

In a cross-vCenter NSX environment, there is both a primary NSX Manager and one or more secondary NSX Managers. The primary NSX Manager allows you to create and manage universal logical switches, universal logical (distributed) routers and universal firewall rules. Secondary NSX Managers are used to manage networking services that are local to that specific NSX Manager. There can be up to seven secondary NSX Managers associated with the primary NSX Manager in a cross-vCenter NSX environment.

Consumption Platform

The consumption of NSX can be driven directly through the NSX Manager user interface, which is available in the vSphere Web Client. Typically end users tie network virtualization to their cloud management platform for deploying applications. NSX provides rich integration into virtually any CMP through REST APIs. Out-of-the-box integration is also available through VMware vCloud Automation Center, vCloud Director, and OpenStack with the Neutron plug-in for NSX.

NSX Edge

You can install NSX Edge as an edge services gateway (ESG) or as a distributed logical router (DLR). The number of edge appliances including ESGs and DLRs is limited to 250 on a host.

Edge Services Gateway

The ESG gives you access to all NSX Edge services such as firewall, NAT, DHCP, VPN, load balancing, and high availability. You can install multiple ESG virtual appliances in a datacenter. Each ESG virtual appliance can have a total of ten uplink and internal network interfaces. With a trunk, an ESG can have up to 200 subinterfaces. The internal interfaces connect to secured port groups and act as the gateway for all protected virtual machines in the port group. The subnet assigned to the internal interface can be a publicly routed IP space or a NATed/routed RFC 1918 private space. Firewall rules and other NSX Edge services are enforced on traffic between network interfaces.

Uplink interfaces of ESGs connect to uplink port groups that have access to a shared corporate network or a service that provides access layer networking. Multiple external IP addresses can be configured for load balancer, site-to-site VPN, and NAT services.

Distributed Logical Router

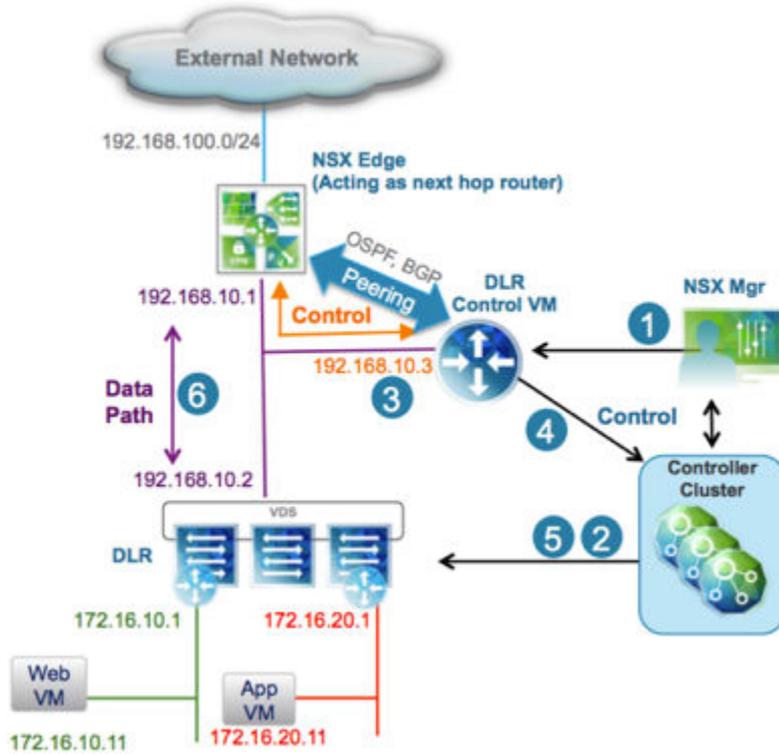
The DLR provides East-West distributed routing with tenant IP address space and data path isolation. Virtual machines or workloads that reside on the same host on different subnets can communicate with one another without having to traverse a traditional routing interface.

A logical router can have eight uplink interfaces and up to a thousand internal interfaces. An uplink interface on a DLR generally peers with an ESG, with an intervening Layer 2 logical transit switch between the DLR and the ESG. An internal interface on a DLR peers with a virtual machine hosted on an ESX hypervisor with an intervening logical switch between the virtual machine and the DLR.

The DLR has two main components:

- The DLR control plane is provided by the DLR virtual appliance (also called a control VM). This VM supports dynamic routing protocols (BGP and OSPF), exchanges routing updates with the next Layer 3 hop device (usually the edge services gateway) and communicates with the NSX Manager and the NSX Controller cluster. High-availability for the DLR virtual appliance is supported through active-standby configuration: a pair of virtual machines functioning in active/standby modes are provided when you create the DLR with HA enabled.
- At the data-plane level, there are DLR kernel modules (VIBs) that are installed on the ESXi hosts that are part of the NSX domain. The kernel modules are similar to the line cards in a modular chassis supporting Layer 3 routing. The kernel modules have a routing information base (RIB) (also known as a routing table) that is pushed from the controller cluster. The data plane functions of route lookup and ARP entry lookup are performed by the kernel modules. The kernel modules are equipped with logical interfaces (called LIFs) connecting to the different logical switches and to any VLAN-backed port-groups. Each LIF has assigned an IP address representing the default IP gateway for the logical L2 segment it connects to and a vMAC address. The IP address is unique for each LIF, whereas the same vMAC is assigned to all the defined LIFs.

Figure 2-1. Logical Routing Components



- 1 A DLR instance is created from the NSX Manager UI (or with API calls), and routing is enabled, leveraging either OSPF or BGP.
- 2 The NSX Controller leverages the control plane with the ESXi hosts to push the new DLR configuration including LIFs and their associated IP and vMAC addresses.
- 3 Assuming a routing protocol is also enabled on the next-hop device (an NSX Edge [ESG] in this example), OSPF or BGP peering is established between the ESG and the DLR control VM. The ESG and the DLR can then exchange routing information:
 - The DLR control VM can be configured to redistribute into OSPF the IP prefixes for all the connected logical networks (172.16.10.0/24 and 172.16.20.0/24 in this example). As a consequence, it then pushes those route advertisements to the NSX Edge. Notice that the next hop for those prefixes is not the IP address assigned to the control VM (192.168.10.3) but the IP address identifying the data-plane component of the DLR (192.168.10.2). The former is called the DLR "protocol address," whereas the latter is the "forwarding address."
 - The NSX Edge pushes to the control VM the prefixes to reach IP networks in the external network. In most scenarios, a single default route is likely to be sent by the NSX Edge, because it represents the single point of exit toward the physical network infrastructure.
- 4 The DLR control VM pushes the IP routes learned from the NSX Edge to the controller cluster.

- 5 The controller cluster is responsible for distributing routes learned from the DLR control VM to the hypervisors. Each controller node in the cluster takes responsibility of distributing the information for a particular logical router instance. In a deployment where there are multiple logical router instances deployed, the load is distributed across the controller nodes. A separate logical router instance is usually associated with each deployed tenant.
- 6 The DLR routing kernel modules on the hosts handle the data-path traffic for communication to the external network by way of the NSX Edge.

NSX Services

The NSX components work together to provide the following functional services.

Logical Switches

A cloud deployment or a virtual data center has a variety of applications across multiple tenants. These applications and tenants require isolation from each other for security, fault isolation, and non-overlapping IP addresses. NSX allows the creation of multiple logical switches, each of which is a single logical broadcast domain. An application or tenant virtual machine can be logically wired to a logical switch. This allows for flexibility and speed of deployment while still providing all the characteristics of a physical network's broadcast domains (VLANs) without physical Layer 2 sprawl or spanning tree issues.

A logical switch is distributed and can span across all hosts in vCenter (or across all hosts in a cross-vCenter NSX environment). This allows for virtual machine mobility (vMotion) within the data center without limitations of the physical Layer 2 (VLAN) boundary. The physical infrastructure is not constrained by MAC/FIB table limits, because the logical switch contains the broadcast domain in software.

Logical Routers

Routing provides the necessary forwarding information between Layer 2 broadcast domains, thereby allowing you to decrease the size of Layer 2 broadcast domains and improve network efficiency and scale. NSX extends this intelligence to where the workloads reside for East-West routing. This allows more direct VM-to-VM communication without the costly or timely need to extend hops. At the same time, NSX logical routers provide North-South connectivity, thereby enabling tenants to access public networks.

Logical Firewall

Logical Firewall provides security mechanisms for dynamic virtual data centers. The Distributed Firewall component of Logical Firewall allows you to segment virtual datacenter entities like virtual machines based on VM names and attributes, user identity, vCenter objects like datacenters, and hosts, as well as traditional networking attributes like IP addresses, VLANs, and so on. The Edge Firewall component helps you meet key perimeter security requirements, such as building DMZs based on IP/VLAN constructs, and tenant-to-tenant isolation in multi-tenant virtual data centers.

The Flow Monitoring feature displays network activity between virtual machines at the application protocol level. You can use this information to audit network traffic, define and refine firewall policies, and identify threats to your network.

Logical Virtual Private Networks (VPNs)

SSL VPN-Plus allows remote users to access private corporate applications. IPsec VPN offers site-to-site connectivity between an NSX Edge instance and remote sites with NSX or with hardware routers/VPN gateways from 3rd-party vendors. L2 VPN allows you to extend your datacenter by allowing virtual machines to retain network connectivity while retaining the same IP address across geographical boundaries.

Logical Load Balancer

The NSX Edge load balancer distributes client connections directed at a single virtual IP address (VIP) across multiple destinations configured as members of a load balancing pool. It distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing thus helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload.

Service Composer

Service Composer helps you provision and assign network and security services to applications in a virtual infrastructure. You map these services to a security group, and the services are applied to the virtual machines in the security group using a Security Policy.

Data Security provides visibility into sensitive data stored within your organization's virtualized and cloud environments and reports any data security violations.

NSX Extensibility

3rd-party solution providers can integrate their solutions with the NSX platform, thus enabling customers to have an integrated experience across VMware products and partner solutions. Data center operators can provision complex, multi-tier virtual networks in seconds, independent of the underlying network topology or components.

Overview of Cross-vCenter Networking and Security

3

NSX 6.2 allows you to manage multiple vCenter NSX environments from a single primary NSX Manager.

This chapter includes the following topics:

- [Benefits of Cross-vCenter NSX](#)
- [How Cross-vCenter NSX Works](#)
- [Support Matrix for NSX Services in Cross-vCenter NSX](#)
- [Universal Controller Cluster](#)
- [Universal Transport Zone](#)
- [Universal Logical Switches](#)
- [Universal Logical \(Distributed\) Routers](#)
- [Universal Firewall Rules](#)
- [Universal Network and Security Objects](#)
- [Cross-vCenter NSX Topologies](#)
- [Modifying NSX Manager Roles](#)

Benefits of Cross-vCenter NSX

NSX environments containing more than one vCenter Server system can be managed centrally.

There are many reasons multiple vCenter Server systems may be required, for example:

- To overcome scale limits of vCenter Server
- To accommodate products that require dedicated or multiple vCenter Server systems, such as Horizon View or Site Recovery Manager
- To separate environments, for example by business unit, tenant, organization, or environment type

In NSX 6.1 and earlier, if multiple vCenter NSX environments are deployed, they must be managed separately. In NSX 6.2 you can create universal objects on the primary NSX Manager, which are synchronized across all vCenter Servers systems in the environment.

Cross-vCenter NSX includes these features:

- Increased span of NSX logical networks. The same logical networks are available across the vCenter NSX environment, so it's possible for VMs on any cluster on any vCenter Server system to be connected to the same logical network.
- Centralized security policy management. Firewall rules are managed from one centralized location, and apply to the VM regardless of location or vCenter Server system.
- Support of new mobility boundaries in vSphere 6, including cross vCenter and long distance vMotion across logical switches.
- Enhanced support for multi-site environments, from metro distance to 150ms RTT. This includes both active-active and active-passive datacenters.

Cross-vCenter NSX environments have many benefits:

- Centralized management of universal objects, reducing administration effort.
- Increased mobility of workloads - VMs can be vMotioned across vCenter Servers without having to reconfigure the VM or change firewall rules.
- Enhanced NSX multi-site and disaster recovery capabilities.

Note Cross-vCenter NSX functionality is supported only with vSphere 6.0.

How Cross-vCenter NSX Works

In a cross-vCenter NSX environment, you can have multiple vCenter Servers, each of which must be paired with its own NSX Manager. One NSX Manager is assigned the role of primary NSX Manager, and the others are assigned the role of secondary NSX Manager.

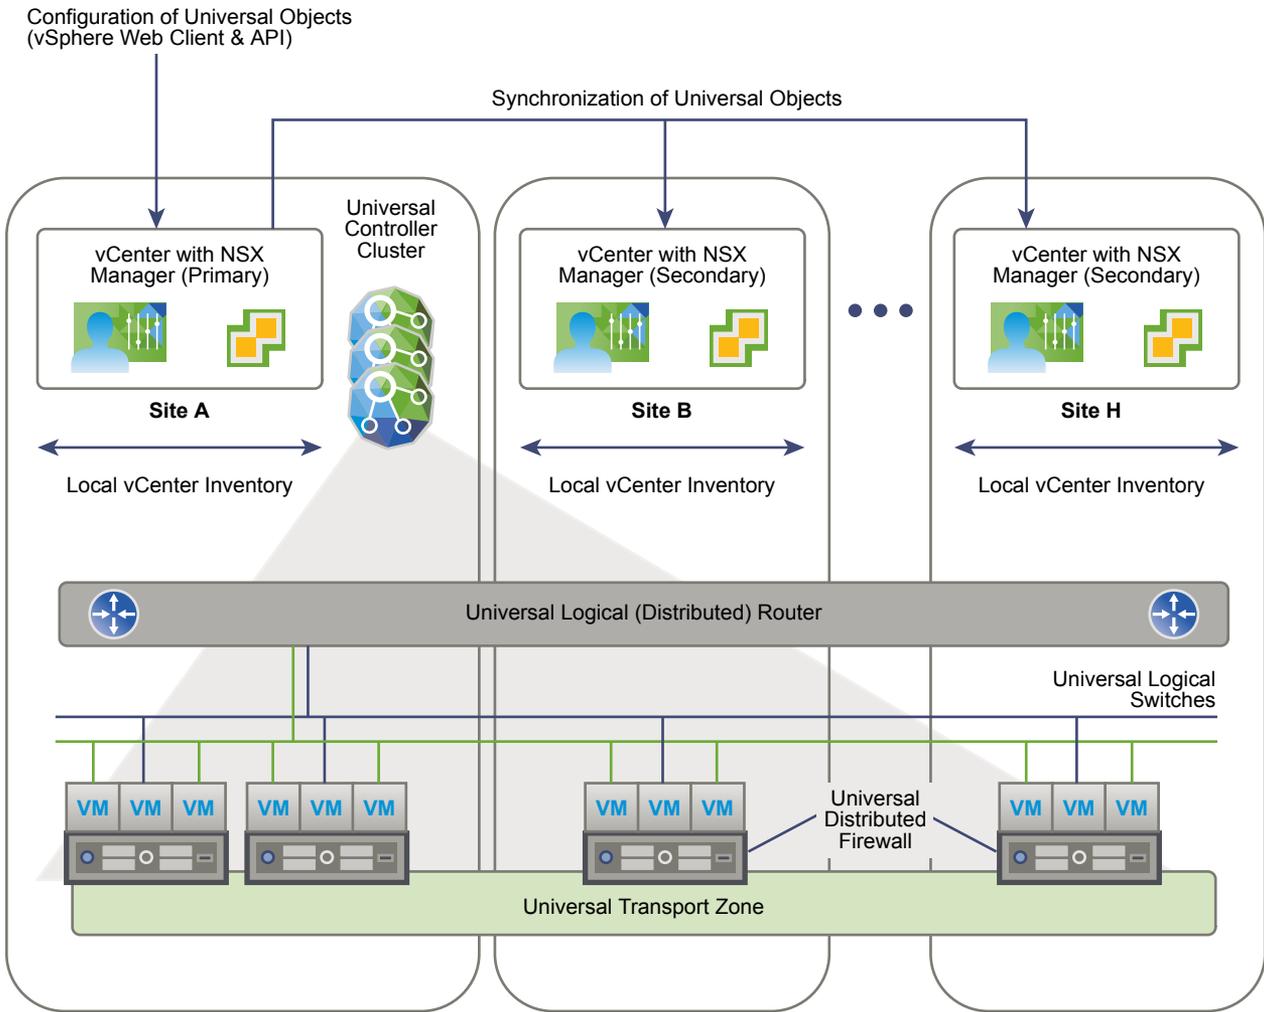
The primary NSX Manager is used to deploy a universal controller cluster that provides the control plane for the cross-vCenter NSX environment. The secondary NSX Managers do not have their own controller clusters.

The primary NSX Manager can create universal objects, such as universal logical switches. These objects are synchronized to the secondary NSX Managers by the NSX Universal Synchronization Service. You can view these objects from the secondary NSX Managers, but you cannot edit them there. You must use the primary NSX Manager to manage universal objects. The primary NSX Manager can be used to configure any of the secondary NSX Managers in the environment.

On both primary and secondary NSX Managers, you can create objects that are local to that specific vCenter NSX environment, such as logical switches, and logical (distributed) routers. They will exist only within the vCenter NSX environment in which they were created. They will not be visible on the other NSX Managers in the cross-vCenter NSX environment.

NSX Managers can be assigned the standalone role. This is equivalent to pre-NSX 6.2 environments with a single NSX Manager and single vCenter. A standalone NSX Manager cannot create universal objects.

Note If you change the role of a primary NSX Manager to standalone and any universal objects exist in the NSX environment, the NSX Manager will be assigned the transit role. The universal objects remain, but they cannot be changed, and no other universal objects can be created. You can delete universal objects from the transit role. The transit role should only be used temporarily, for example, when changing which NSX Manager is the primary.



Support Matrix for NSX Services in Cross-vCenter NSX

A subset of NSX Services are available for universal synchronization in cross-vCenter NSX.

Table 3-1. Support matrix for NSX Services in cross-vCenter NSX

NSX Service	Details	Supports cross-vCenter NSX synchronization in NSX 6.2?
Logical switch	Transport zone	Yes
	Logical switch	Yes

Table 3-1. Support matrix for NSX Services in cross-vCenter NSX (Continued)

NSX Service	Details	Supports cross-vCenter NSX synchronization in NSX 6.2?
L2 bridges		No
Routing	Logical (distributed) router	Yes
	Logical (distributed) router appliance	No by design. Appliances must be created on each NSX Manager if multiple appliances are required per universal logical router. This allows for different configurations per appliance, which may be required in an environment with local egress configured.
	NSX Edge services gateway	No
Logical firewall	Distributed firewall	Yes
	Exclude list	No
	SpoofGuard	No
	Flow monitoring for aggregate flows	No
	Network service insertion	No
	Edge firewall	No
VPN		No
Logical load balancer		No
Other edge services		No
Service composer		No
Data security		No
Network extensibility		No
Network and security objects	IP address groups (IP sets)	Yes
	MAC address groups (MAC sets)	Yes
	IP pools	No
	Security groups	Yes, but universal security groups can contain only included objects, no dynamic membership or excluded objects
	Services	Yes
	Service groups	Yes

Universal Controller Cluster

Each cross-vCenter NSX environment has one universal controller cluster associated with the primary NSX Manager. Secondary NSX Managers do not have a controller cluster.

As the universal controller cluster is the only controller cluster for the cross-vCenter NSX environment, it maintains information about universal logical switches and universal logical routers as well as logical switches and logical routers that are local to a vCenter NSX pair.

In order to avoid any overlap in object IDs, separate ID pools are maintained for universal objects and local objects.

Universal Transport Zone

In a cross-vCenter NSX environment, there can be only one universal transport zone.

The universal transport zone is created on the primary NSX Manager, and is synchronized to the secondary NSX Managers. Clusters that need to participate in universal logical networks must be added to the universal transport zone from their NSX Managers.

Universal Logical Switches

Universal logical switches allow layer 2 networks to span multiple sites.

When you create a logical switch in a universal transport zone, you create a universal logical switch. This switch is available on all clusters in the universal transport zone. The universal transport zone can include clusters in any vCenter in the cross-vCenter NSX environment.

The segment ID pool is used to assign VNIs to logical switches, and the universal segment ID pool is used to assign VNIs to universal logical switches. These pools must not overlap.

You must use a universal logical router to route between universal logical switches. If you need to route between a universal logical switch and a logical switch, you must use an Edge Services Gateway.

Universal Logical (Distributed) Routers

Universal Logical (Distributed) Routers offer centralized administration and a routing configuration that can be customized at the universal logical router, cluster, or host level.

When you create a universal logical router you must choose whether to enable local egress, as this cannot be changed after creation. Local egress allows you to control what routes are provided to ESXi hosts based on an identifier, the locale ID.

Each NSX Manager is assigned a locale ID, which is set to the NSX Manager UUID by default. You can override the locale ID at the following levels:

- Universal logical router
- Cluster
- ESXi host

If you do not enable local egress the locale ID is ignored and all ESXi hosts connected to the universal logical router will receive the same routes. Whether or not to enable local egress in a cross-vCenter NSX environment is a design consideration, but it is not required for all cross-vCenter NSX configurations.

Universal Firewall Rules

Distributed Firewall in a cross-vCenter NSX environment allows centralized management of rules that apply to all vCenter Servers in your environment. It supports cross-vCenter vMotion which enables you to move workloads or virtual machines from one vCenter Server to another and seamlessly extends your software defined datacenter security.

As your datacenter needs scale out, the existing vCenter Server may not scale to the same level. This may require you to move a set of applications to newer hosts that are managed by a different vCenter Server. Or you may need to move applications from staging to production in an environment where staging servers are managed by one vCenter Server and production servers are managed by a different vCenter Server. Distributed Firewall supports these cross-vCenter vMotion scenarios by replicating firewall policies that you define for the primary NSX Manager on up to seven secondary NSX Managers.

From the primary NSX Manager you can create a distributed firewall rule section that is marked for universal synchronization. You can create one universal L2 rule section and one universal L3 rule section. These sections and their rules are synchronized to all secondary NSX Managers in your environment. Rules in other sections remain local to the appropriate NSX Manager.

The following Distributed Firewall features are not supported in a cross-vCenter NSX environment:

- Exclude list
- SpoofGuard
- Flow monitoring for aggregate flows
- Network service insertion
- Edge Firewall

Service Composer does not support universal synchronization, so you cannot use it to create distributed firewall rules in the universal section.

Universal Network and Security Objects

You can create custom network and security objects to use in Distributed Firewall rules in the universal section.

- Universal IP Sets
- Universal MAC Sets
- Universal Security Groups
- Universal Services
- Universal Service Groups

Universal network and security objects can be created only from the primary NSX Manager.

Universal security groups can contain only universal IP sets, universal MAC sets, and universal security groups. Membership is defined by included objects only, you cannot use dynamic membership or excluded objects.

Universal security groups cannot be created from Service Composer. Security groups created from Service Composer will be local to that NSX Manager.

Cross-vCenter NSX Topologies

You can deploy cross-vCenter NSX in a single physical site, or across multiple sites.

Multi-Site and Single Site Cross-vCenter NSX

A cross-vCenter NSX environment allows you to use the same logical switches and other network objects across multiple vCenter NSX setups. The vCenters can be located in the same site, or in different sites.

Whether the cross-vCenter NSX environment is contained within a single site or crosses multiple sites, a similar configuration can be used. These two example topologies consist of the following:

- A universal transport zone that includes all clusters in the site or sites.
- Universal logical switches attached to the universal transport zone. Two universal logical switches are used to connect VMs and one is used as a transit network for the router uplink.
- VMs added to the universal logical switches
- A universal logical router with an NSX Edge appliance to enable dynamic routing. The universal logical router appliance has internal interfaces on the VM universal logical switches and an uplink interface on the transit network universal logical switch.
- Edge Services Gateways (ESGs) connected to the transit network and the physical egress router network.

Figure 3-1. Cross-vCenter NSX in a single site

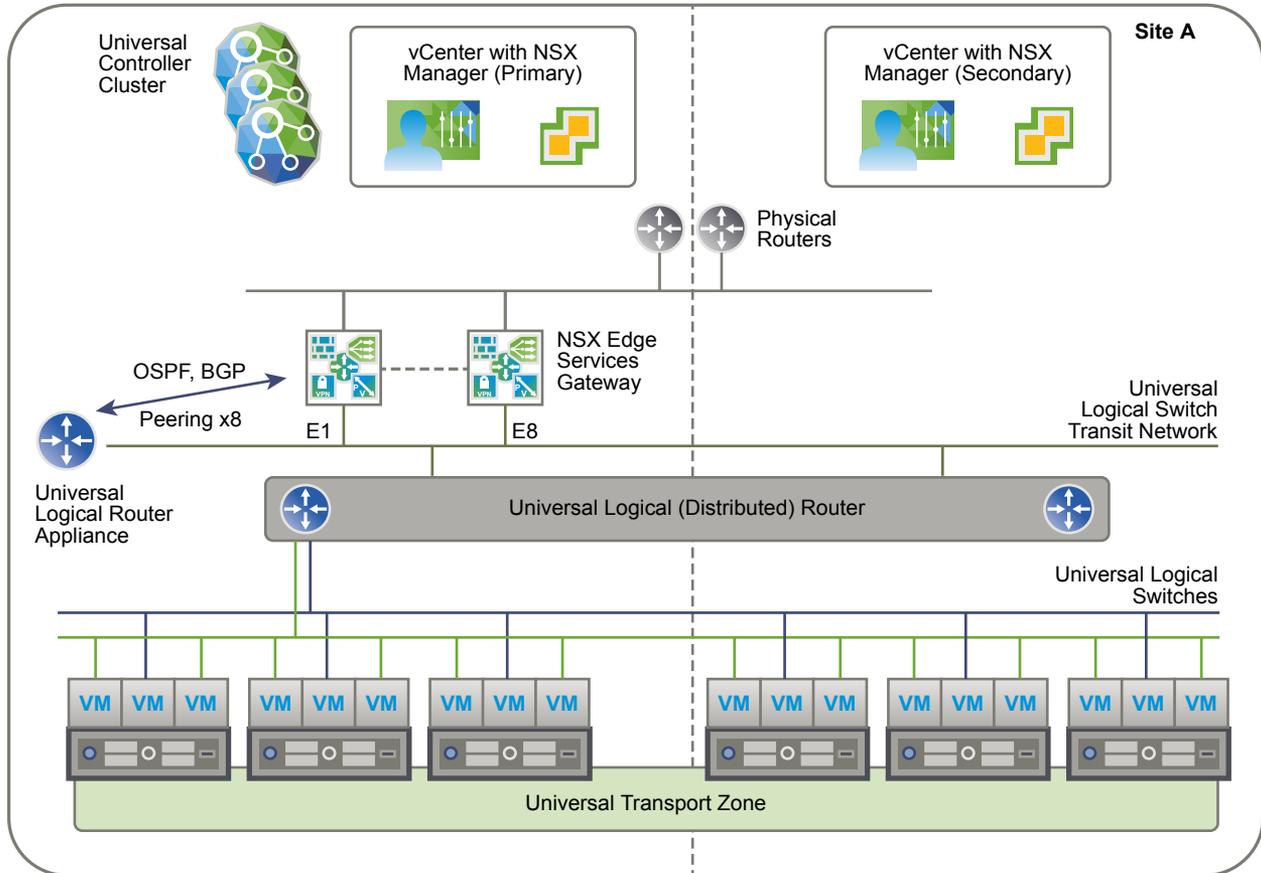
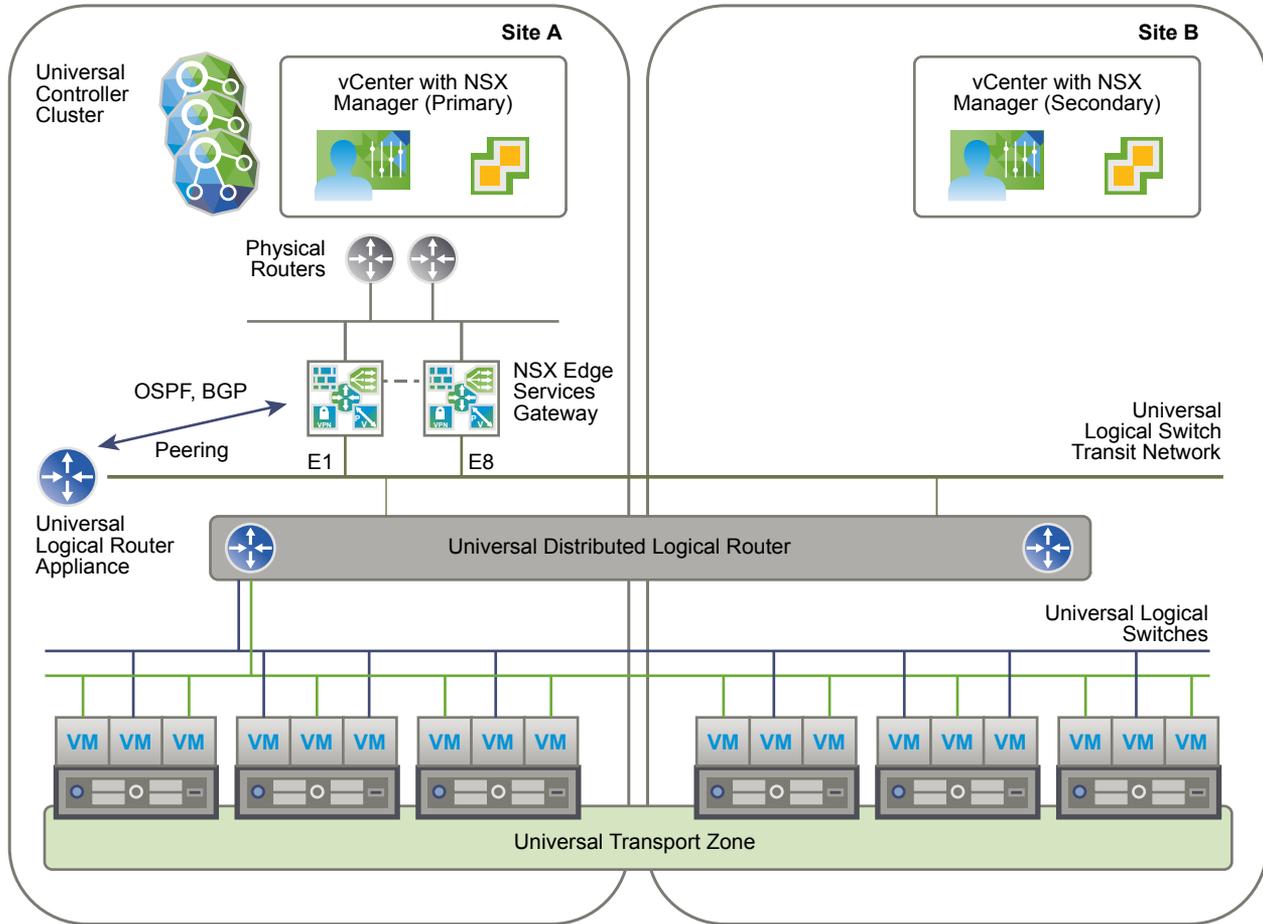


Figure 3-2. Cross-vCenter NSX spanning two sites

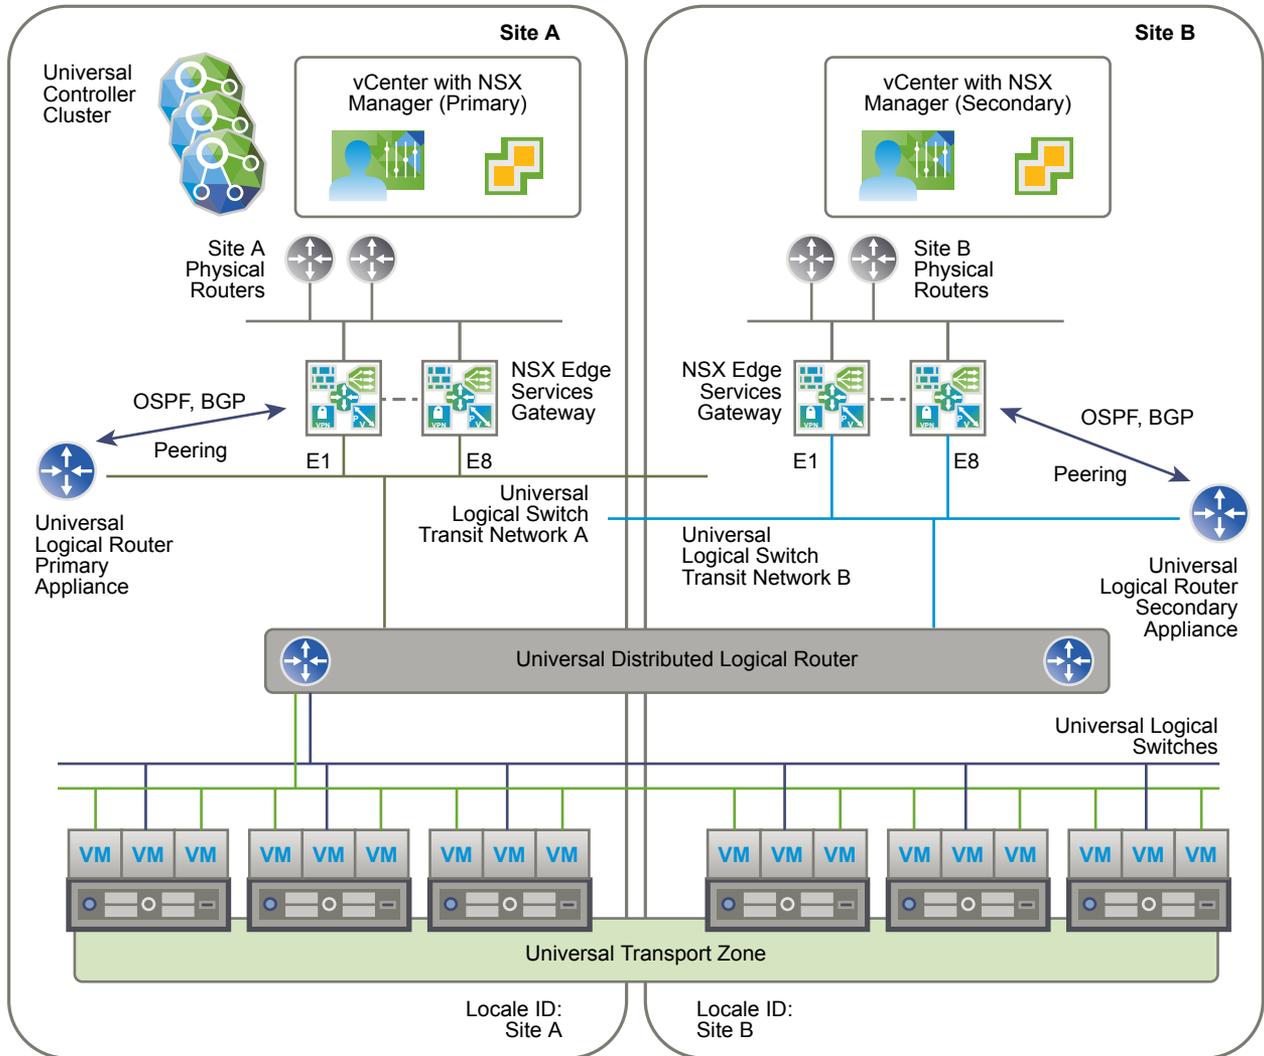


Local Egress

All sites in a multi-site cross-vCenter NSX environment can use the same physical routers for egress traffic. However, if egress routes need to be customized, the local egress feature must be enabled when the universal logical router is created. This allows you to customize routes at the universal logical router, cluster, or host level.

This example of a cross-vCenter NSX environment in multiple sites has local egress enabled. The edge services gateways (ESGs) in each site have a default route that sends traffic out through that site's physical routers. The universal logical router is configured with two appliances, one in each site. The appliances learn routes from their site's ESGs. The learned routes are sent to the universal controller

cluster. Because local egress is enabled, the locale ID for that site is associated with those routes. The universal controller cluster sends routes with matching locale IDs to the hosts. Routes learned on the site A appliance are sent to the hosts in site A, and routes learned on the site B appliance are sent to the hosts in site B.



Modifying NSX Manager Roles

An NSX Manager can have the primary role, the secondary role, or the standalone role. Special synchronization software runs on the primary NSX Manager, synchronizing all universal objects to secondary NSX Managers.

It is important to understand what happens when you change an NSX Manager's role.

Set as primary	This operation sets the role of an NSX Manager to primary and starts the synchronization software. This operation fails if the NSX Manager is already the primary or already a secondary.
Set as standalone (from secondary)	This operation sets the role of NSX Manager to standalone or transit mode. This operation might fail if the NSX Manager already has the standalone role.
Set as standalone (from primary)	This operation resets the primary NSX Manager to standalone or transit mode, stops the synchronization software, and unregisters all secondary NSX Managers. This operation might fail if the NSX Manager is already standalone or if any of the secondary NSX Managers are unreachable.
Disconnect from primary	When you run this operation on a secondary NSX Manager, the secondary NSX Manager is unilaterally disconnected from the primary NSX Manager. This operation should be used when the primary NSX Manager has experienced an unrecoverable failure, and you want to register the secondary NSX Manager to a new primary. If the original primary NSX Manager does come up again, its database continues to list the secondary NSX Manager as registered. To resolve this issue, include the force option when you disconnect or unregister the secondary from the original primary. The force option removes the secondary NSX Manager from the original primary NSX Manager's database.

Preparing for Installation

This section describes the system requirements for NSX as well as the ports that must be open.

This chapter includes the following topics:

- [System Requirements for NSX](#)
- [Ports and Protocols Required by NSX](#)
- [NSX and vSphere Distributed Switches](#)
- [Example: Working with a vSphere Distributed Switch](#)
- [NSX Installation Workflow and Sample Topology](#)
- [Cross-vCenter NSX and Enhanced Linked Mode](#)

System Requirements for NSX

Before you install or upgrade NSX, consider your network configuration and resources. You can install one NSX Manager per vCenter Server, one instance of Guest Introspection and Data Security per ESXi™ host, and multiple NSX Edge instances per datacenter.

Hardware

Table 4-1. Hardware Requirements

Appliance	Memory	vCPU	Disk Space
NSX Manager	16 GB (24 GB with certain NSX deployment sizes*)	4 (8 with certain NSX deployment sizes*)	60 GB
NSX Controller	4 GB	4	20 GB
NSX Edge	<ul style="list-style-type: none"> ▪ Compact: 512 MB ▪ Large: 1 GB ▪ Quad Large: 1 GB ▪ X-Large: 8 GB 	<ul style="list-style-type: none"> ▪ Compact: 1 ▪ Large: 2 ▪ Quad Large: 4 ▪ X-Large: 6 	<ul style="list-style-type: none"> ▪ Compact: 1 disk 500MB ▪ Large: 1 disk 500MB + 1 disk 512MB ▪ Quad-Large: 1 disk 500MB + 1 disk 512MB ▪ X-Large: 1 disk 500MB + 1 disk 2GB

Table 4-1. Hardware Requirements (Continued)

Appliance	Memory	vCPU	Disk Space
Guest Introspection	1 GB	2	4 GB
NSX Data Security	512 MB	1	6 GB per ESXi host

As a general guideline, you should increase NSX Manager resources to 8 vCPU and 24 GB of RAM if your NSX managed environment contains more than 256 hypervisors or more than 2000 VMs.

For specific sizing details contact VMware support.

For information about increasing the memory and vCPU allocation for your virtual appliances, see *Allocate Memory Resources*, and *Change the Number of Virtual CPUs in vSphere Virtual Machine Administration*.

Software

For the latest interoperability information, see the Product Interoperability Matrixes at http://partnerweb.vmware.com/comp_guide/sim/interop_matrix.php.

For recommended versions of NSX, vCenter Server, and ESXi, see the release notes at <https://docs.vmware.com/en/VMware-NSX-for-vSphere/index.html>.

Note that for an NSX Manager to participate in a cross-vCenter NSX deployment the following conditions are required:

Component	Version
NSX Manager	6.2 or later
NSX Controller	6.2 or later
vCenter Server	6.0 or later
ESXi	<ul style="list-style-type: none"> ■ ESXi 6.0 or later ■ Host clusters prepared with NSX 6.2 or later VIBs

To manage all NSX Managers in a cross-vCenter NSX deployment from a single vSphere Web Client, you must connect your vCenter Servers in Enhanced Linked Mode. See *Using Enhanced Linked Mode in vCenter Server and Host Management*.

To check the compatibility of partner solutions with NSX, see the VMware Compatibility Guide for Networking and Security at

<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=security>.

Client and User Access

- If you added ESXi hosts by name to the vSphere inventory, ensure that forward and reverse name resolution is working. Otherwise, NSX Manager cannot resolve the IP addresses.
- Permissions to add and power on virtual machines

- Access to the datastore where you store virtual machine files, and the account permissions to copy files to that datastore
- Cookies enabled on your Web browser, to access the NSX Manager user interface
- From NSX Manager, ensure port 443 is accessible from the ESXi host, the vCenter Server, and the NSX appliances to be deployed. This port is required to download the OVF file on the ESXi host for deployment.
- A Web browser that is supported for the version of vSphere Web Client you are using. See Using the vSphere Web Client in the *vCenter Server and Host Management* documentation for details.

Ports and Protocols Required by NSX

The following ports must be open for NSX to operate properly.

Table 4-2. Ports and Protocols required by NSX

Source	Target	Port	Protocol	Purpose	Sensitive	TLS	Authentication
Client PC	NSX Manager	443	TCP	NSX Manager Administrative Interface	No	Yes	PAM Authentication
Client PC	NSX Manager	80	TCP	NSX Manager VIB Access	No	No	PAM Authentication
ESXi Host	vCenter Server	443	TCP	ESXi Host Preparation	No	No	
vCenter Server	ESXi Host	443	TCP	ESXi Host Preparation	No	No	
ESXi Host	NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password
ESXi Host	NSX Controller	1234	TCP	User World Agent Connection	No	Yes	
NSX Controller	NSX Controller	2878, 2888, 3888	TCP	Controller Cluster - State Sync	No	Yes	IPsec
NSX Controller	NSX Controller	7777	TCP	Inter-Controller RPC Port	No	Yes	IPsec
NSX Controller	NSX Controller	30865	TCP	Controller Cluster - State Sync	No	Yes	IPsec
NSX Manager	NSX Controller	443	TCP	Controller to Manager Communication	No	Yes	User/Password
NSX Manager	vCenter Server	443	TCP	vSphere Web Access	No	Yes	
NSX Manager	vCenter Server	902	TCP	vSphere Web Access	No	Yes	

Table 4-2. Ports and Protocols required by NSX (Continued)

Source	Target	Port	Protocol	Purpose	Sensitive	TLS	Authentication
NSX Manager	ESXi Host	443	TCP	Management and provisioning connection	No	Yes	
NSX Manager	ESXi Host	902	TCP	Management and provisioning connection	No	Yes	
NSX Manager	DNS Server	53	TCP	DNS client connection	No	No	
NSX Manager	DNS Server	53	UDP	DNS client connection	No	No	
NSX Manager	Syslog Server	514	TCP	Syslog connection	No	No	
NSX Manager	Syslog Server	514	UDP	Syslog connection	No	No	
NSX Manager	NTP Time Server	123	TCP	NTP client connection	No	Yes	
NSX Manager	NTP Time Server	123	UDP	NTP client connection	No	Yes	
vCenter Server	NSX Manager	80	TCP	Host Preparation	No	Yes	
REST Client	NSX Manager	443	TCP	NSX Manager REST API	No	Yes	User/Password
VXLAN Tunnel End Point (VTEP)	VXLAN Tunnel End Point (VTEP)	8472 (default before NSX 6.2.3) or 4789 (default in new installs of NSX 6.2.3 and later)	UDP	Transport network encapsulation between VTEPs	No	Yes	
ESXi Host	ESXi Host	6999	UDP	ARP on VLAN LIFs	No	Yes	
ESXi Host	NSX Manager	8301, 8302	UDP	DVS Sync	No	Yes	
NSX Manager	ESXi Host	8301, 8302	UDP	DVS Sync	No	Yes	
Guest Introspection VM	NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password

Table 4-2. Ports and Protocols required by NSX (Continued)

Source	Target	Port	Protocol	Purpose	Sensitive	TLS	Authentication
Primary NSX Manager	Secondary NSX Manager	443	TCP	Cross-vCenter NSX Universal Sync Service	No	Yes	
Primary NSX Manager	vCenter Server	443	TCP	vSphere API	No	Yes	
Secondary NSX Manager	vCenter Server	443	TCP	vSphere API	No	Yes	
Primary NSX Manager	NSX Universal Controller Cluster	443	TCP	NSX Controller REST API	No	Yes	User/Password
Secondary NSX Manager	NSX Universal Controller Cluster	443	TCP	NSX Controller REST API	No	Yes	User/Password
ESXi Host	NSX Universal Controller Cluster	1234	TCP	NSX Control Plane Protocol	No	Yes	
ESXi Host	Primary NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password
ESXi Host	Secondary NSX Manager	5671	TCP	RabbitMQ	No	Yes	RabbitMQ User/Password

Ports for Cross-vCenter NSX and Enhanced Linked Mode

If you have a cross-vCenter NSX environment and your vCenter Server systems are in Enhanced Linked Mode, in order to manage any NSX Manager from any vCenter Server system each NSX Manager appliance must have the required connectivity to each vCenter Server system in the environment.

NSX and vSphere Distributed Switches

In an NSX domain, NSX vSwitch is the software that operates in server hypervisors to form a software abstraction layer between servers and the physical network.

NSX vSwitch is based on vSphere distributed switches (VDSs), which provide uplinks for host connectivity to the top-of-rack (ToR) physical switches. As a best practice, VMware recommends that you plan and prepare your vSphere distributed switches before installing NSX for vSphere.

A single host can be attached to multiple VDSs. A single VDS can span multiple hosts across multiple clusters. For each host cluster that will participate in NSX, all hosts within the cluster must be attached to a common VDS.

For instance, say you have a cluster with Host1 and Host2. Host1 is attached to VDS1 and VDS2. Host2 is attached to VDS1 and VDS3. When you prepare a cluster for NSX, you can only associate NSX with VDS1 on the cluster. If you add another host (Host3) to the cluster and Host3 is not attached to VDS1, it is an invalid configuration, and Host3 will not be ready for NSX functionality.

Often, to simplify a deployment, each cluster of hosts is associated with only one VDS, even though some of the VDSs span multiple clusters. For example, suppose your vCenter contains the following host clusters:

- Compute cluster A for app tier hosts
- Compute cluster B for web tier hosts
- Management and edge cluster for management and edge hosts

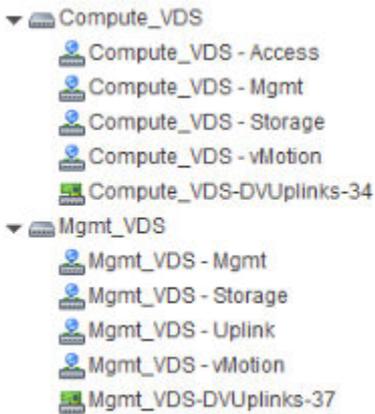
The following screen shows how these clusters appear in vCenter.



For a cluster design such as this, you might have two VDSs called Compute_VDS and Mgmt_VDS. Compute_VDS spans both of the compute clusters, and Mgmt_VDS is associated with only the management and edge cluster.

Each VDS contains distributed port groups for the different types of traffic that need to be carried. Typical traffic types include management, storage, and vMotion. Uplink and access ports are generally required as well. Normally, one port group for each traffic type is created on each VDS.

For example, the following screen shows how these distributed switches and ports appear in vCenter.

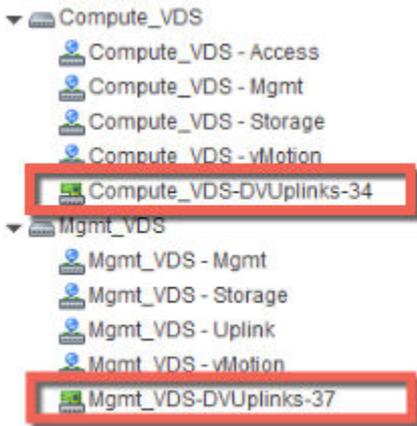


Each port group can, optionally, be configured with a VLAN ID. The following list shows an example of how VLANs can be associated with the distributed port groups to provide logical isolation between different traffic types:

- Compute_VDS - Access---VLAN 130
- Compute_VDS - Mgmt---VLAN 210

- Compute_VDS - Storage---VLAN 520
- Compute_VDS - vMotion---VLAN 530
- Mgmt_VDS - Uplink---VLAN 100
- Mgmt_VDS - Mgmt---VLAN 110
- Mgmt_VDS - Storage---VLAN 420
- Mgmt_VDS - vMotion---VLAN 430

The DVUplinks port group is a VLAN trunk that is created automatically when you create a VDS. As a trunk port, it sends and receives tagged frames. By default, it carries all VLAN IDs (0-4094). This means that traffic with any VLAN ID can be passed through the vmnic network adapters associated with the DVUplink slot and filtered by the hypervisor hosts as the distributed switch determines which port group should receive the traffic.



If your existing vCenter environment contains standard vSwitches instead of distributed switches, you can migrate your hosts to distributed switches.

Example: Working with a vSphere Distributed Switch

This example shows how to create a new vSphere distributed switch (VDS); add port groups for management, storage, and vMotion traffic types; and migrate hosts on a standard vSwitch to the new distributed switch.

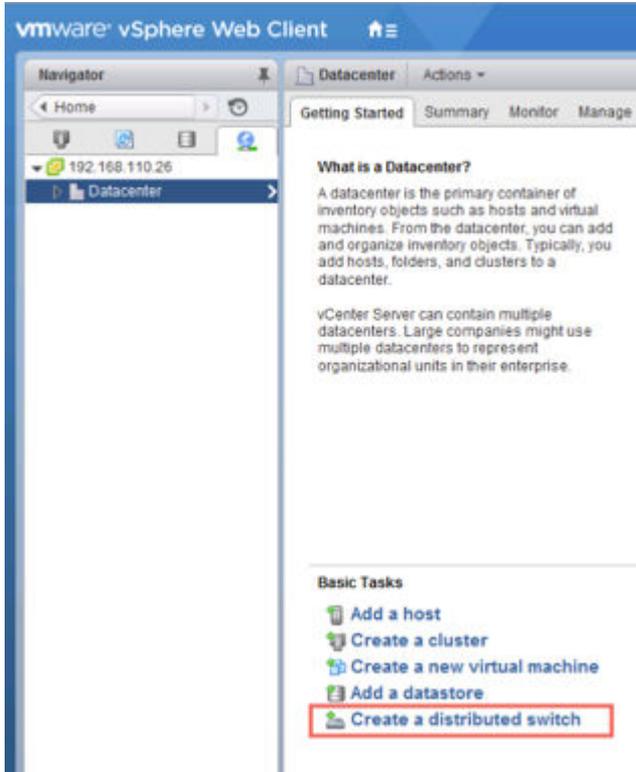
Note that this is just one example used to show the procedure. For detailed VDS physical and logical uplink considerations, see the *VMware NSX for vSphere Network Virtualization Design Guide* at <https://communities.vmware.com/docs/DOC-27683>.

Prerequisites

This example assumes that each ESX host to be connected to the vSphere distributed switch has at least one connection to a physical switch (one vmnic uplink). This uplink can be used for the distributed switch and NSX VXLAN traffic.

Procedure

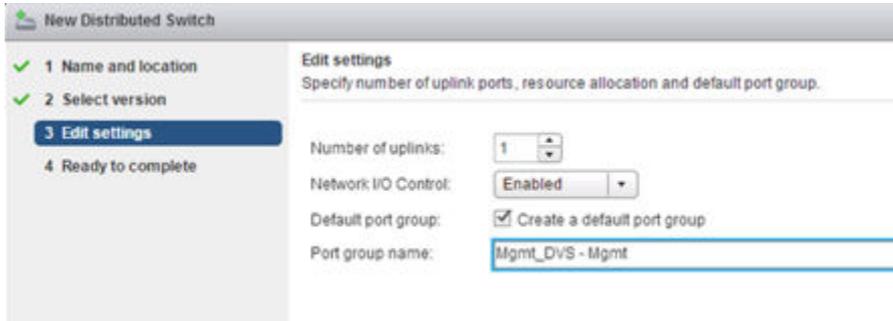
- 1 In the vSphere Web Client, navigate to a datacenter.
- 2 Click **Create a Distributed Switch**.



- 3 Give the switch a meaningful name based on the host cluster that will be associated with this switch. For example, if a distributed switch will be associated with a cluster of datacenter management hosts, you could name the switch VDS_Mgmt.
- 4 Provide at least one uplink for the distributed switch, keep IO control enabled, and provide a meaningful name for the default port group. Note that it is not mandatory to create the default port group. The port group can be manually created later.

By default, four uplinks are created. Adjust the number of uplinks to reflect your VDS design. The number of uplinks required is normally equal to the number of physical NICs you allocate to the VDS.

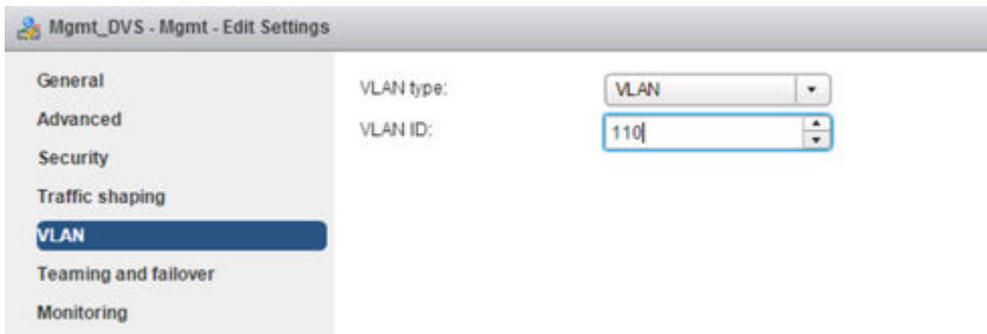
The following screen shows example settings for management traffic on the management host cluster.



The default port group is just one of the port groups that this switch will contain. You will have an opportunity after the switch is created to add port groups for different traffic types. Optionally, you can untick **Create a default port group** option when creating a new VDS. This may in fact be the best practice; it's best to be explicit when creating port groups.

- 5 (Optional) Upon completion of the New Distributed Switch wizard, edit the settings of the default port group to place it in the correct VLAN for management traffic.

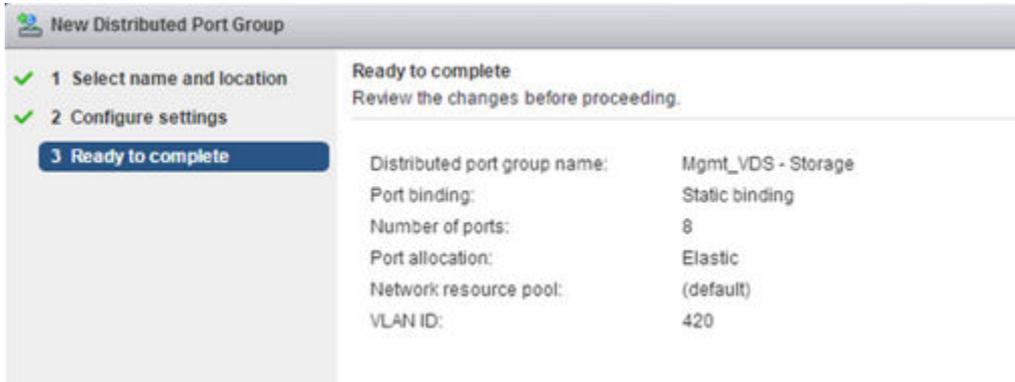
For example, if your host management interfaces are in VLAN 110, place the default port group in VLAN 110. If your host management interfaces are not in a VLAN, skip this step.



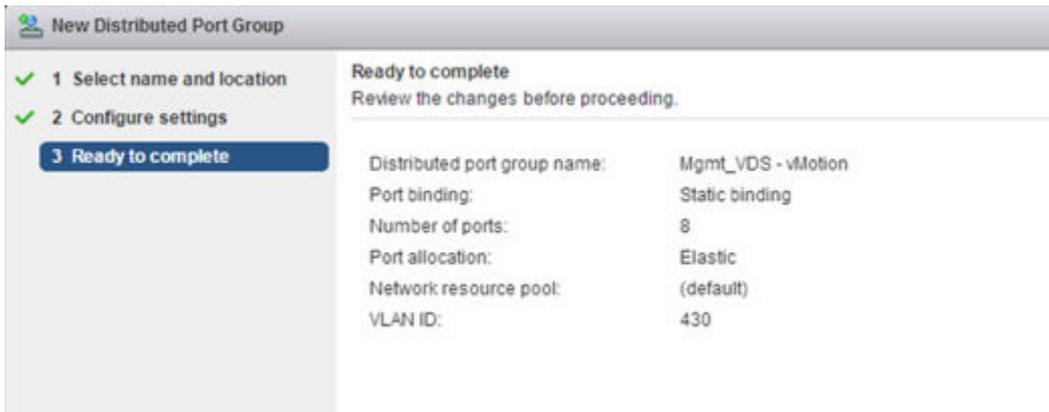
- 6 Upon completion of the New Distributed Switch wizard, right-click the distributed switch and select **New Distributed Port Group**.

Repeat this step for each traffic type, making sure to provide a meaningful name for each port group and making sure to configure the proper VLAN ID based on the traffic separation requirements of your deployment.

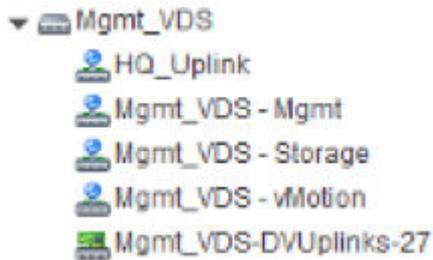
Example group settings for storage.



Example group settings for vMotion traffic.

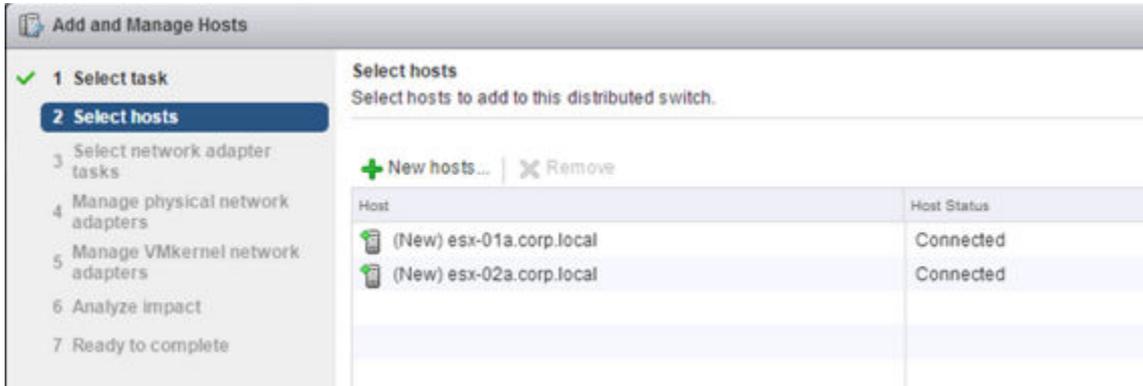


The completed distributed switch and port groups looks like this.

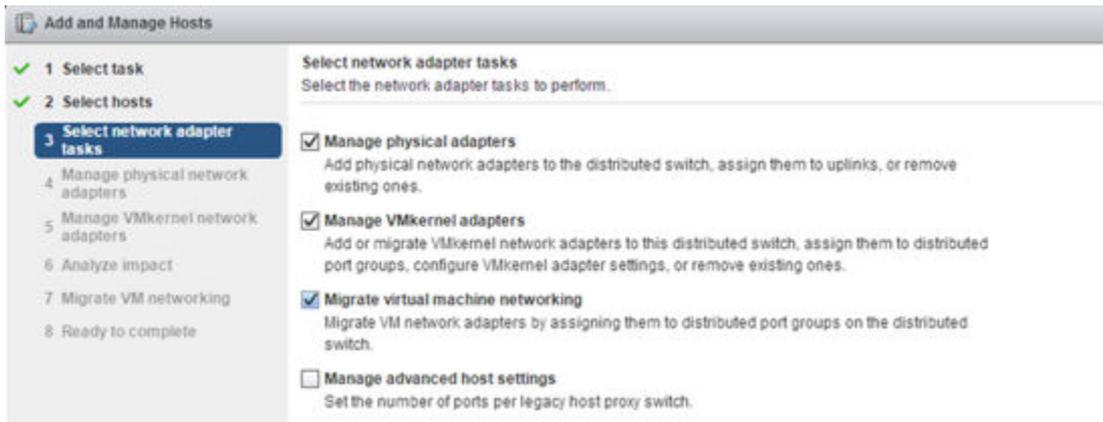


- Right-click the distributed switch, select **Add and Manage Hosts**, and select **Add Hosts**.

Attach all hosts that are in the associated cluster. For example, if the switch is for management hosts, select all of the hosts that are in the management cluster.

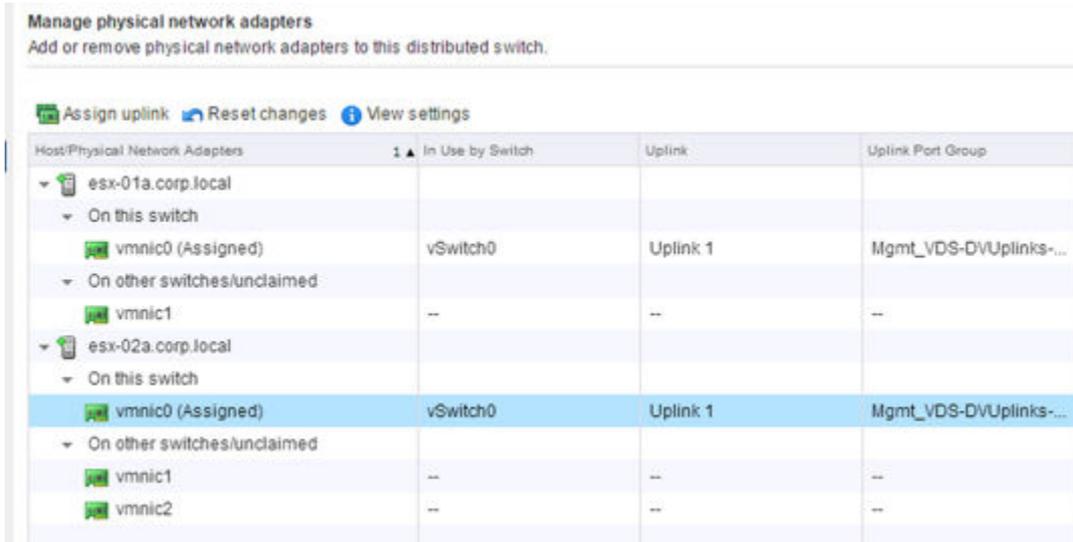


- Select the options to migrate physical adapters, VMkernel adapters, and virtual machine networking.



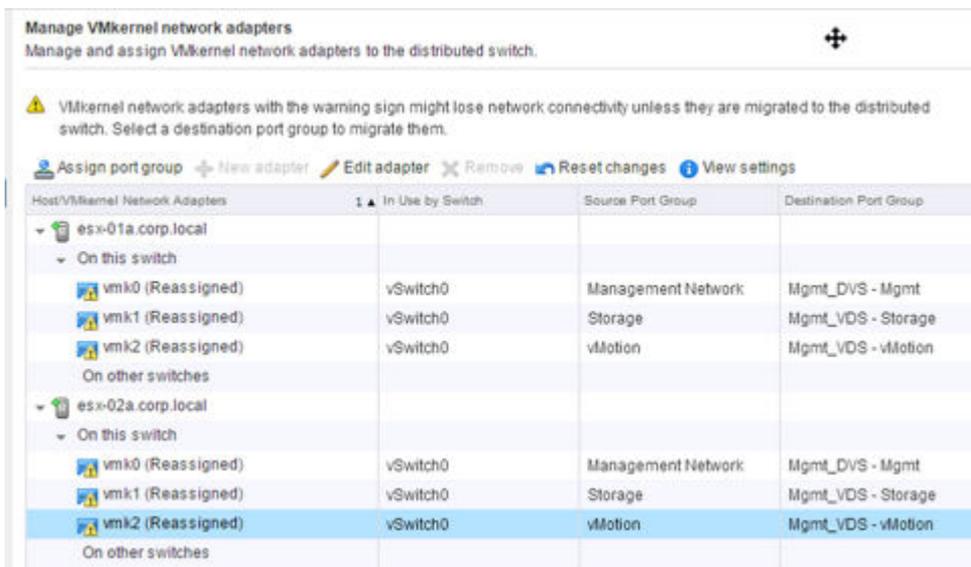
- 9 Select a vmnic and click **Assign uplink** to migrate the vmnic from the standard vSwitch to the distributed switch. Repeat this step for each host that you are attaching to the distributed vSwitch.

For example, this screen shows two hosts with their vmnic0 uplinks configured to migrate from their respective standard vSwitch to the distributed Mgmt_VDS-DVUplinks port group, which is a trunk port that can carry any VLAN ID.



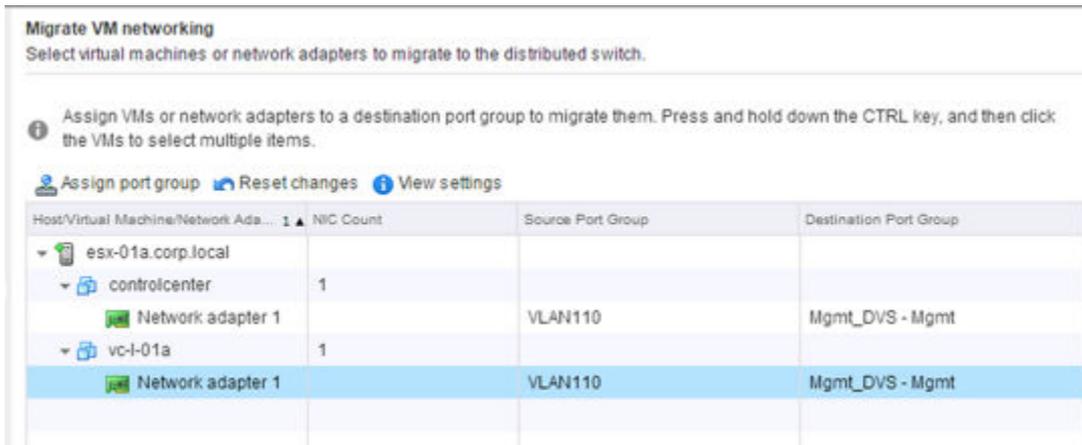
- 10 Select a VMKernel network adapter and click **Assign port group**. Repeat this step for all network adapters on all hosts that you are attaching to the distributed vSwitch.

For example, this screen shows three vmk network adapters on two hosts configured to be migrated from the standard port groups to the new distributed port groups.



11 Move any VMs that are on the hosts to a distributed port group.

For example, this screen shows two VMs on a single host configured to be migrated from the standard port group to the new distributed port group.



After the procedure is complete, in the host CLI you can verify the results by running the following commands:

```

~ # esxcli network vswitch dvs vmware list
Mgmt_VDS
  Name: Mgmt_VDS
  VDS ID: 89 78 26 50 98 bb f5 1e-a5 07 b5 29 ff 86 e2 ac
  Class: etherswitch
  Num Ports: 1862
  Used Ports: 5
  Configured Ports: 512
  MTU: 1600
  CDP Status: listen
  Beacon Timeout: -1
  Uplinks: vmnic0
  VMware Branded: true
  DVPort:
    Client: vmnic0
    DVPortgroup ID: dvportgroup-306
    In Use: true
    Port ID: 24

    Client: vmk0
    DVPortgroup ID: dvportgroup-307
    In Use: true
    Port ID: 0

    Client: vmk2
    DVPortgroup ID: dvportgroup-309
    In Use: true
    Port ID: 17

```

```

Client: vmk1
DVPortgroup ID: dvportgroup-308
In Use: true
Port ID: 9

```

```

■ ~ # esxcli network ip interface list
vmk2
  Name: vmk2
  MAC Address: 00:50:56:6f:2f:26
  Enabled: true
  Portset: DvsPortset-0
  Portgroup: N/A
  Netstack Instance: defaultTcpipStack
  VDS Name: Mgmt_VDS
  VDS UUID: 89 78 26 50 98 bb f5 1e-a5 07 b5 29 ff 86 e2 ac
  VDS Port: 16
  VDS Connection: 1235399406
  MTU: 1500
  TSO MSS: 65535
  Port ID: 50331650

vmk0
  Name: vmk0
  MAC Address: 54:9f:35:0b:dd:1a
  Enabled: true
  Portset: DvsPortset-0
  Portgroup: N/A
  Netstack Instance: defaultTcpipStack
  VDS Name: Mgmt_VDS
  VDS UUID: 89 78 26 50 98 bb f5 1e-a5 07 b5 29 ff 86 e2 ac
  VDS Port: 2
  VDS Connection: 1235725173
  MTU: 1500
  TSO MSS: 65535
  Port ID: 50331651

vmk1
  Name: vmk1
  MAC Address: 00:50:56:6e:a4:53
  Enabled: true
  Portset: DvsPortset-0
  Portgroup: N/A
  Netstack Instance: defaultTcpipStack
  VDS Name: Mgmt_VDS
  VDS UUID: 89 78 26 50 98 bb f5 1e-a5 07 b5 29 ff 86 e2 ac
  VDS Port: 8
  VDS Connection: 1236595869
  MTU: 1500
  TSO MSS: 65535
  Port ID: 50331652

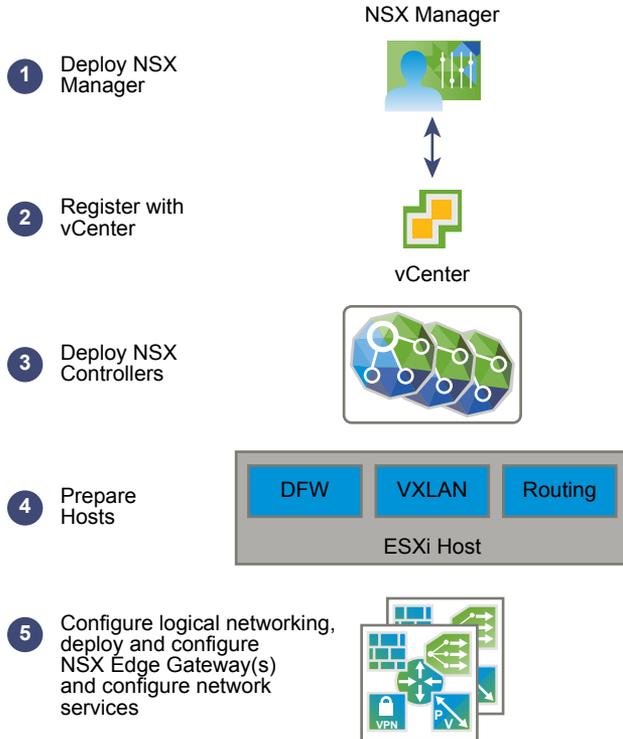
```

What to do next

Repeat the migration process for all vSphere distributed switches.

NSX Installation Workflow and Sample Topology

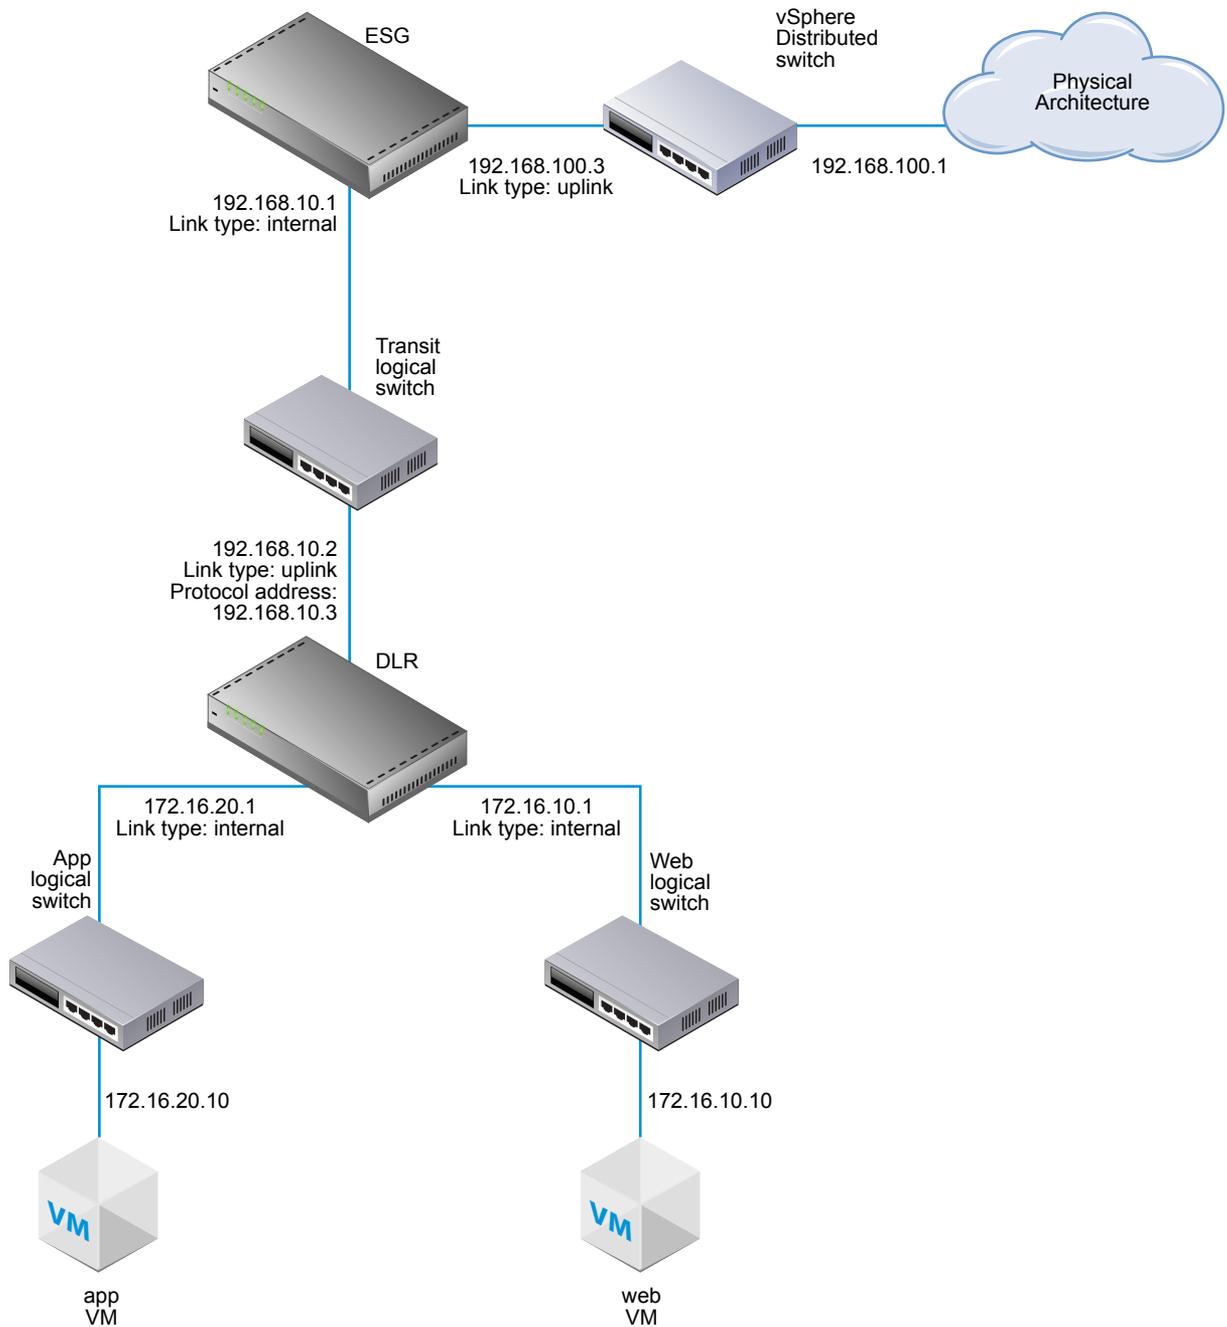
NSX installation involves the deployment of several virtual appliances, some ESX host preparation, and some configuration to allow communication across all of the physical and virtual devices.



The process begins by deploying an NSX Manager OVF/OVA template and ensuring that the NSX Manager has full connectivity to the management interfaces of the ESX hosts that it will manage. After that, the NSX Manager and a vCenter instance need to be linked with each other through a registration process. This then allows a cluster of NSX controllers to be deployed. NSX controllers, like the NSX Manager, run as virtual appliances on ESX hosts. The next step is to prepare the ESX hosts for NSX by installing several VIBs on the hosts. These VIBs enable the Layer 2 VXLAN functionality, distributed routing, and distributed firewall. After configuring VXLANs, specifying virtual network interface (VNI) ranges, and creating transport zones, you can build out your NSX overlay topology.

This installation guide describes in detail each step in the process.

While being applicable to any NSX deployment, this guide also leads you through the creation of a sample NSX overlay topology that you can use for practice, guidance, and reference purposes. The sample overlay has a single NSX logical distributed router (sometimes called a DLR), an edge services gateway (ESG), and an NSX logical transit switch connecting the two NSX routing devices. The sample topology includes elements of an underlay as well, including two sample virtual machines. These virtual machines are each connected to a separate NSX logical switch that allow connectivity through the NSX logical router (DLR).



Cross-vCenter NSX and Enhanced Linked Mode

vSphere 6.0 introduces Enhanced Linked Mode, which links multiple vCenter Server systems by using one or more Platform Services Controllers. This allows you to view and search the inventories of all linked vCenter Server systems within the vSphere Web Client. In a cross-vCenter NSX environment, Enhanced Linked Mode allows you to manage all NSX Managers from a single vSphere Web Client.

In large deployments where there are multiple vCenter Servers, it might make sense for you to use Cross-vCenter NSX with Enhanced Linked Mode for vCenter. These two features are complementary but separate from each other.

Combining Cross-vCenter NSX with Enhanced Linked Mode

In cross-vCenter NSX, you have a primary NSX Manager and multiple secondary NSX Managers. Each of these NSX Managers is linked to a separate vCenter Server. On the primary NSX Manager, you can create universal NSX components (such as switches and routers) that are viewable from the secondary NSX Managers.

When the individual vCenter Servers are deployed with Enhanced Linked Mode, all of the vCenter Servers can be viewed and managed from a single vCenter Server (sometimes called a single pane of glass).

Thus, when cross-vCenter NSX is combined with Enhanced Linked Mode for vCenter, you can view and manage any of the NSX Managers and all of the universal NSX components from any of the linked vCenter Servers.

Using Cross-vCenter NSX Without Enhanced Linked Mode

Enhanced Linked Mode is not a prerequisite or requirement for cross-vCenter NSX. Without Enhanced Linked Mode, you can still create cross-vCenter universal transport zones, universal switches, universal routers, and universal firewall rules. However, without Enhanced Linked Mode in place, you must log in to the individual vCenter Servers to access each NSX Manager instance.

Further Information About vSphere 6.0 and Enhanced Linked Mode

For more information about vSphere 6.0 and enhanced linked mode, see <https://blogs.vmware.com/consulting/2015/03/vsphere-datacenter-design-vcenter-architecture-changes-vsphere-6-0-part-1.html>.

If you decide to use Enhanced Linked Mode see the *vSphere 6.0 Installation and Setup Guide* <http://pubs.vmware.com/vsphere-60/topic/com.vmware.vsphere.install.doc/GUID-7C9A1E23-7FCD-4295-9CB1-C932F2423C63.html> or the *vSphere 6.0 Upgrade Guide* <http://pubs.vmware.com/vsphere-60/topic/com.vmware.vsphere.upgrade.doc/GUID-18B7B4BB-C24A-49CD-AE76-13285157B29F.html> for the latest requirements for vSphere 6.0 and Enhanced Linked Mode.

Tasks for Primary and Secondary NSX Managers

5

In a cross-vCenter environment there can be one primary NSX Manager and up to seven secondary NSX Managers. Some setup tasks are performed on every NSX Manager, whether it will become a primary NSX Manager, or a secondary NSX Manager.

This chapter includes the following topics:

- [Install the NSX Manager Virtual Appliance](#)
- [Configure Single Sign On](#)
- [Register vCenter Server with NSX Manager](#)
- [Specify a Syslog Server](#)
- [Install and Assign NSX for vSphere License](#)
- [Exclude Virtual Machines from Firewall Protection](#)

Install the NSX Manager Virtual Appliance

The NSX Manager provides the graphical user interface (GUI) and the REST APIs for creating, configuring, and monitoring NSX components, such as controllers, logical switches, and edge services gateways. The NSX Manager provides an aggregated system view and is the centralized network management component of NSX. NSX Manager is installed as a virtual appliance on any ESX host in your vCenter environment.

The NSX Manager virtual machine is packaged as an OVA file, which allows you to use the vSphere Web Client to import the NSX Manager into the datastore and virtual machine inventory.

For high availability, VMware recommends that you deploy NSX Manager in a cluster configured with HA and DRS. Optionally, you can install the NSX Manager in a different vCenter than the one that the NSX Manager will be interoperating with. A single NSX Manager serves a single vCenter Server environment.

In cross-vCenter NSX installations, make sure that each NSX Manager has a unique UUID. NSX Manager instances deployed from OVA files have unique UUIDs. An NSX Manager deployed from a template (as in when you convert a virtual machine to a template) will have the same UUID as the original NSX Manager used to create the template, and these two NSX Managers cannot be used in the same cross-vCenter NSX installation. In other words, for each NSX Manager, you should install a new appliance from scratch as outlined in this procedure.

The NSX Manager virtual machine installation includes VMware Tools. Do not attempt to upgrade or install VMware Tools on the NSX Manager.

During the installation, you can choose to join the Customer Experience Improvement Program (CEIP) for NSX. See Customer Experience Improvement Program in the *NSX Administration Guide* for more information about the program, including how to join or leave the program.

Prerequisites

- Before installing NSX Manager, make sure that the required ports are open. See [Ports and Protocols Required by NSX](#).
- Make sure that a datastore is configured and accessible on the target ESX host. Shared storage is recommended. HA requires shared storage, so that the NSX Manager appliance can be restarted on another host if the original host fails.
- Make sure that you know the IP address and gateway, DNS server IP addresses, domain search list, and the NTP server IP address that the NSX Manager will use.
- Decide whether NSX Manager will have IPv4 addressing only, IPv6 addressing only, or dual-stack network configuration. The host name of the NSX Manager will be used by other entities. Hence, the NSX Manager host name must be mapped to the right IP address in the DNS servers used in that network.
- Prepare a management traffic distributed port group on which NSX Manager will communicate. See [Example: Working with a vSphere Distributed Switch](#). The NSX Manager management interface, vCenter Server, and ESXi host management interfaces must be reachable by NSX Guest Introspection instances.
- The Client Integration Plug-in must be installed. The Deploy OVF template wizard works best in the Firefox web browser. Sometimes in the Chrome web browser, an error message about installing the Client Integration Plug-in is displayed even though the plug-in is already successfully installed. To install the Client Integration Plug-in:
 - a Open a Web browser and type the URL for the vSphere Web Client.
 - b At the bottom of the vSphere Web Client login page, click Download Client Integration Plug-in.

If the Client Integration Plug-In is already installed on your system, you will not see the link to download the plug-in. If you uninstall the Client Integration Plug-In, the link to download it will display on the vSphere Web Client login page.

Procedure

- 1 Locate the NSX Manager Open Virtualization Appliance (OVA) file.
Either copy the download URL or download the OVA file onto your computer.
- 2 In Firefox, open vCenter.
- 3 Select **VMs and Templates**, right-click your datacenter, and select **Deploy OVF Template**.
- 4 Paste the download URL or click **Browse** to select the file on your computer.

5 Tick the checkbox **Accept extra configuration options**.

This allows you to set IPv4 and IPv6 addresses, default gateway, DNS, NTP, and SSH properties during the installation, rather than configuring these settings manually after the installation.

6 Accept the VMware license agreements.

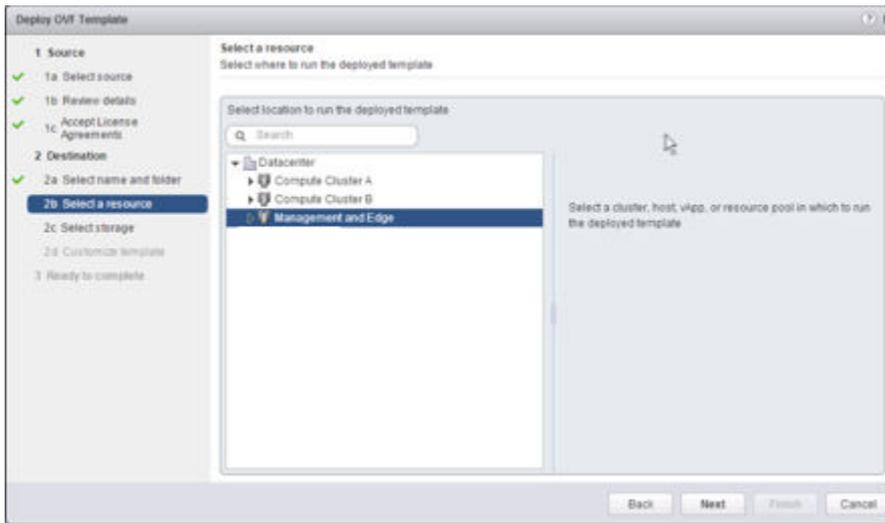
7 Edit the NSX Manager name (if required).select the location for the deployed NSX Manager

The name you type will appear in the vCenter inventory.

The folder you select will be used to apply permissions to the NSX Manager.

8 Select a host or cluster on which to deploy the NSX Manager appliance.

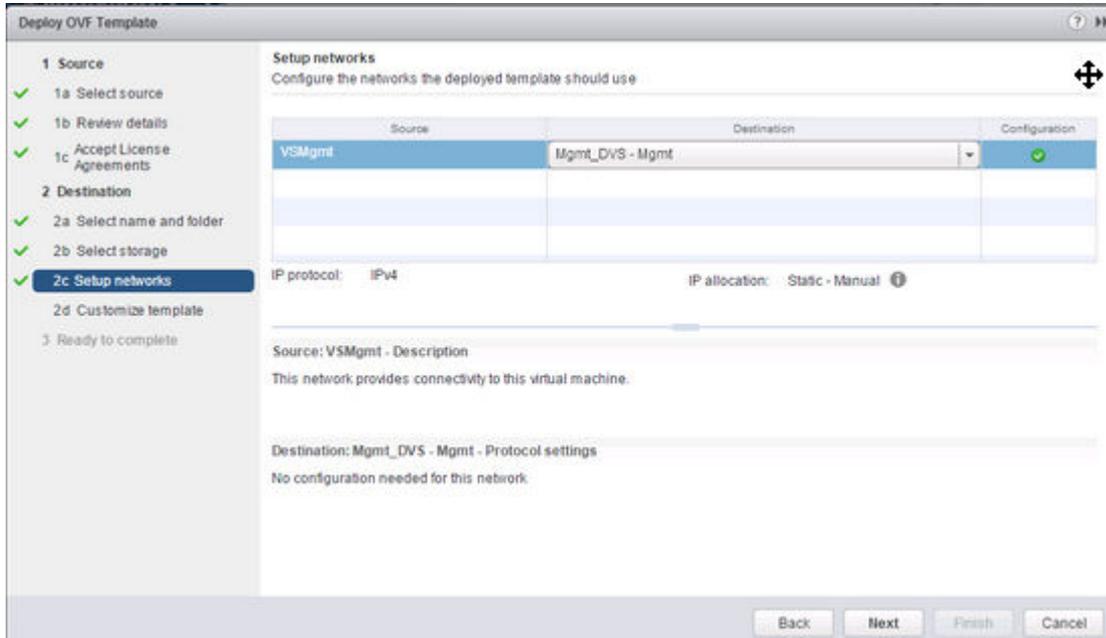
For example:



9 Change the virtual disk format to **Thick Provision**, and select the destination datastore for the virtual machine configuration files and the virtual disks.

10 Select the port group for the NSX Manager.

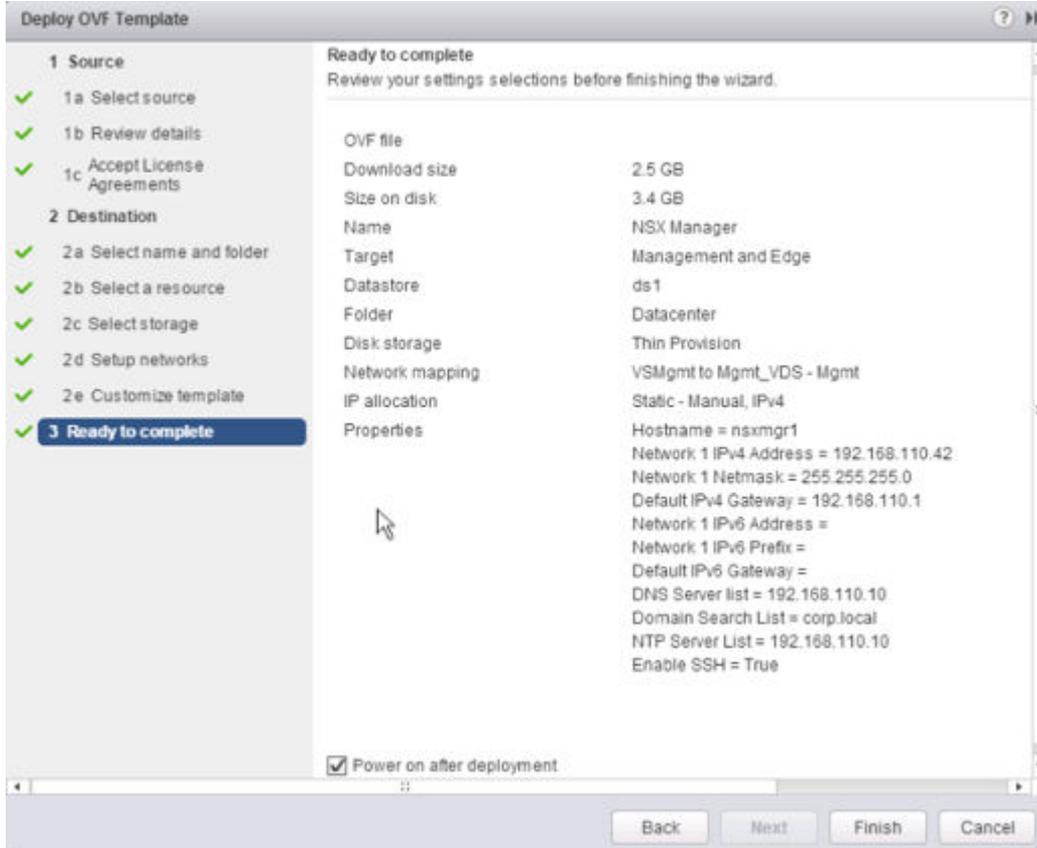
For example, this screen shot shows the selection of the Mgmt_DVS - Mgmt port group.



11 (Optional) Select the **Join the Customer Experience Improvement Program** checkbox.

12 Set the NSX Manager extra configuration options.

For example, this screen shows the final review screen after all the options are configured in an IPv4-only deployment.



Open the console of the NSX Manager to track the boot process.

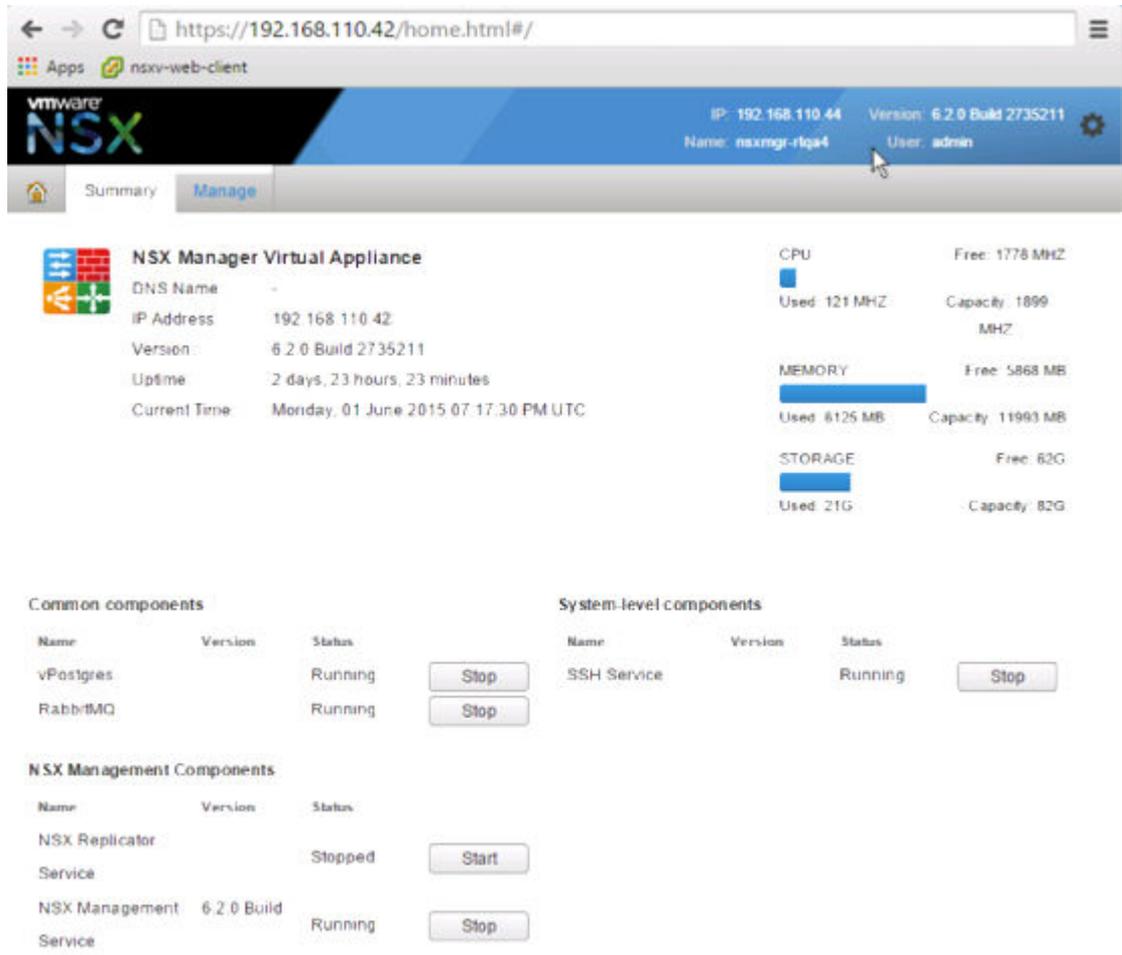
After the NSX Manager is completely booted, log in to the CLI and run the `show interface` command to verify that the IP address was applied as expected.

```
nsxmgr1> show interface
Interface mgmt is up, line protocol is up
  index 3 metric 1 mtu 1500 <UP,BROADCAST,RUNNING,MULTICAST>
  HWaddr: 00:50:56:8e:c7:fa
  inet 192.168.110.42/24 broadcast 192.168.110.255
  inet6 fe80::250:56ff:fe8e:c7fa/64
  Full-duplex, 0Mb/s
  input packets 1370858, bytes 389455808, dropped 50, multicast packets 0
  input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 0, missed 0
  output packets 1309779, bytes 2205704550, dropped 0
  output errors 0, aborted 0, carrier 0, fifo 0, heartbeat 0, window 0
  collisions 0
```

Make sure that the NSX Manager can ping its default gateway, its NTP server, the vCenter Server, and the IP address of the management interface on all hypervisor hosts that it will manage.

Connect to the NSX Manager appliance GUI by opening a web browser and navigating to the NSX Manager IP address or hostname.

After logging in as **admin** with the password you set during installation, click **View Summary** and make sure that the following services are running: vPostgres, RabbitMQ, and NSX Management Services.



For optimal performance, VMware recommends that you reserve memory for the NSX Manager virtual appliance. A memory reservation is a guaranteed lower bound on the amount of physical memory that the host reserves for a virtual machine, even when memory is overcommitted. Set the reservation to a level that ensures NSX Manager has sufficient memory to run efficiently.

What to do next

Register the vCenter Server with the NSX Manager.

Configure Single Sign On

SSO makes vSphere and NSX more secure by allowing the various components to communicate with each other through a secure token exchange mechanism, instead of requiring each component to authenticate a user separately. You can configure lookup service on the NSX Manager and provide the SSO administrator credentials to register NSX Management Service as an SSO user. Integrating the single sign on (SSO) service with NSX improves the security of user authentication for vCenter users and enables NSX to authenticate users from other identity services such as AD, NIS, and LDAP.

With SSO, NSX supports authentication using authenticated Security Assertion Markup Language (SAML) tokens from a trusted source via REST API calls. NSX Manager can also acquire authentication SAML tokens for use with other VMware solutions.

NSX caches group information for SSO users. Changes to group memberships will take up to 60 minutes to propagate from the identity provider (for example, active directory) to NSX.

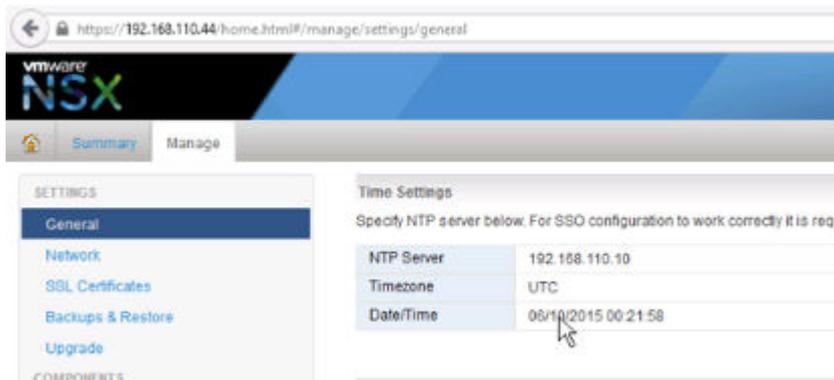
Prerequisites

- To use SSO on NSX Manager, you must have vCenter Server 5.5 or later, and single sign on (SSO) authentication service must be installed on the vCenter Server. Note that this is for embedded SSO. Instead, your deployment might use an external centralized SSO server.

For information about SSO services provided by vSphere, see <http://kb.vmware.com/kb/2072435> and <http://kb.vmware.com/kb/2113115>.

- NTP server must be specified so that the SSO server time and NSX Manager time is in sync.

For example:



Procedure

- 1 Log in to the NSX Manager virtual appliance.

In a Web browser, navigate to the NSX Manager appliance GUI at <https://<nsx-manager-ip>> or <https://<nsx-manager-hostname>>, and log in as admin with the password that you configured during NSX Manager installation.

- 2 Click the **Manage** tab, then click **NSX Management Service**.
- 3 Type the name or IP address of the host that has the lookup service.

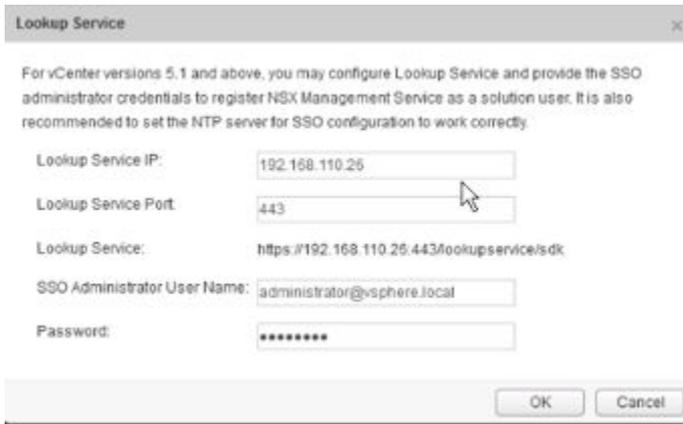
If you are using vCenter to perform the lookup service, enter the vCenter Server's IP address or hostname, and enter the vCenter Server user name and password.

- 4 Type the port number.

Enter port 443 if you are using vSphere 6.0. For vSphere 5.5, use port number 7444.

The Lookup Service URL is displayed based on the specified host and port.

For example:

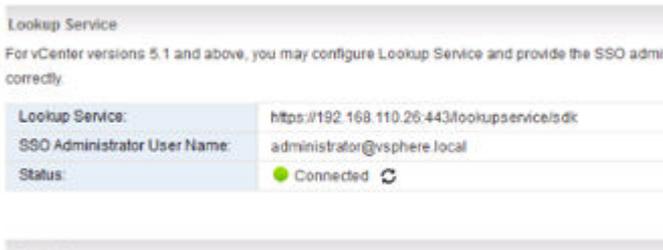


- 5 Check that the certificate thumb print matches the certificate of the vCenter Server.

If you installed a CA-signed certificate on the CA server, you are presented with the thumbprint of the CA-signed certificate. Otherwise, you are presented with a self-signed certificate.

- 6 Confirm that the Lookup Service status is **Connected**.

For example:



What to do next

Assign a role to the SSO user.

Register vCenter Server with NSX Manager

NSX Manager and vCenter have a one-to-one relationship. For every instance of NSX Manager, there is one vCenter Server. This is true even if you are using the NSX cross-vCenter feature. After you install NSX Manager and ensure that the NSX management service is running, the next step is to register a vCenter Server with the NSX Manager.

Only one NSX Manager can be registered with a vCenter. Changing a vCenter registration can lead to an issue in which the change does not get communicated correctly to all affected vCenters and NSX Managers.

For example, suppose you have the following initial configuration of NSX Managers and vCenters:

- NSX1 ----> VC1
- NSX2 ----> VC2

If you change the configuration on NSX1 so that its vCenter is VC2, the following occurs:

- NSX1 correctly reports that its vCenter is VC2.
- VC2 correctly reports that its NSX Manager is NSX1.
- VC1 incorrectly reports that its NSX Manager is NSX1.
- NSX2 incorrectly reports that its vCenter is VC2.

In other words, if a vCenter is already registered with one NSX Manager, and then a different NSX Manager registers the same vCenter, the vCenter automatically removes its connection with the first NSX Manager and connects to the new NSX Manager. However, when you log in to the first NSX Manager, it continues to report a connection to the vCenter.

To prevent this issue, remove the NSX Manager plug-in from VC1 before registering it with VC2. For instructions, see [Safely Remove an NSX Installation](#).

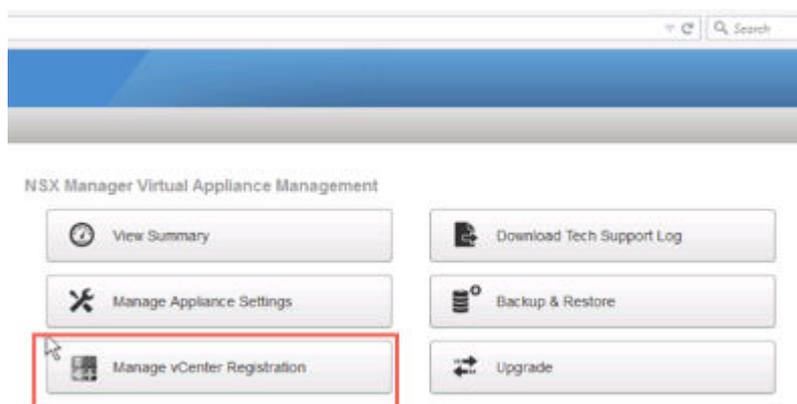
Prerequisites

- The NSX management service must be running. You can verify this by using a Web browser to open the NSX Manager appliance GUI at <https://<nsx-manager-ip>> and looking at the **Summary** tab.
- You must have a vCenter Server user account with the Administrator role to synchronize NSX Manager with the vCenter Server. If your vCenter password has non-ASCII characters, you must change it before synchronizing the NSX Manager with the vCenter Server.

Procedure

- 1 In a Web browser, navigate to the NSX Manager appliance GUI at <https://<nsx-manager-ip>> or <https://<nsx-manager-hostname>>, and log in as admin with the password that you configured during NSX Manager installation.
- 2 Under Appliance Management, click **Manage vCenter Registration**.

For example:



- 3 Edit the vCenter Server element to point to the vCenter Server's IP address or hostname, and enter the vCenter Server user name and password.

For the user name, the best practice is to enter administrator@vsphere.local or an alternative account that you have created. Do not use the root account.

- 4 Check that the certificate thumbprint matches the certificate of the vCenter Server.

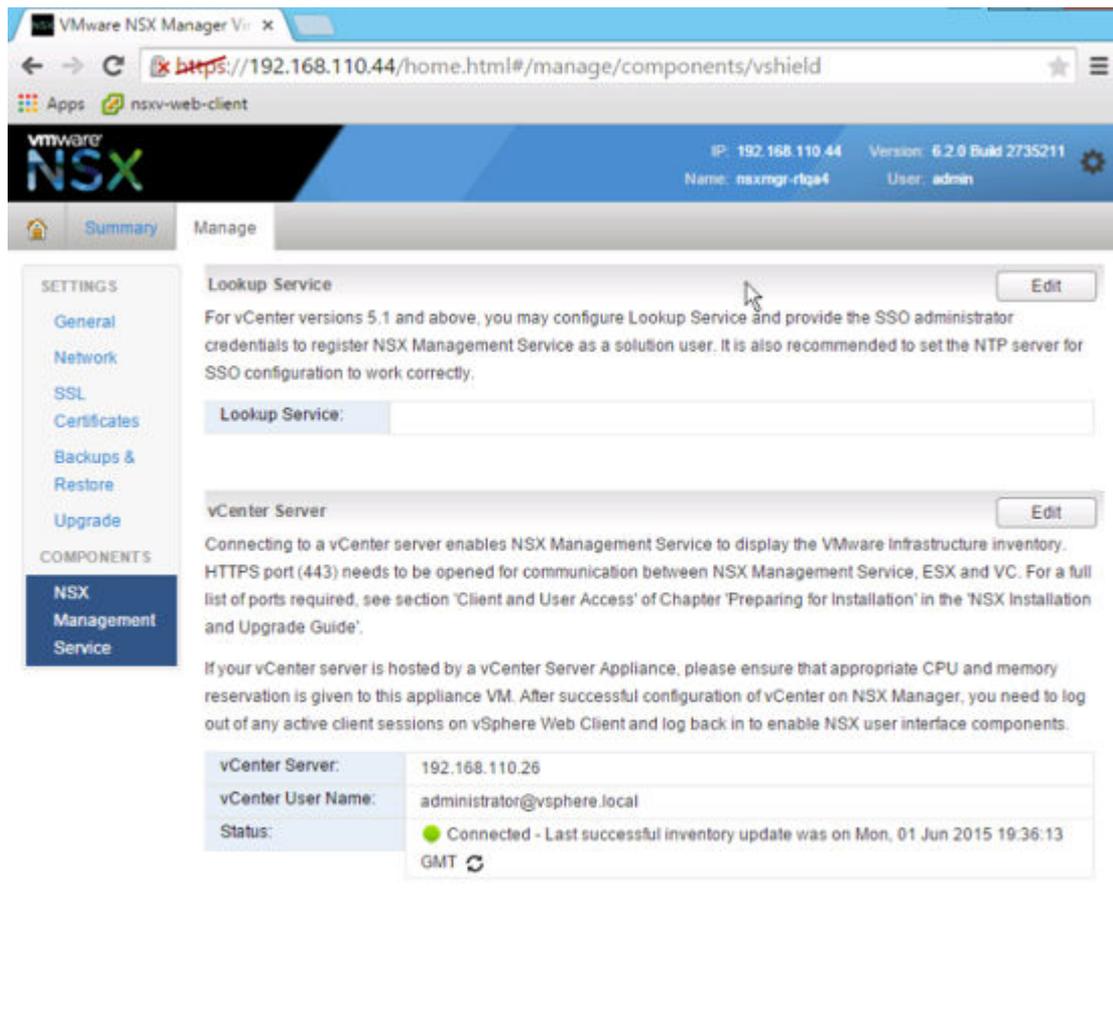
If you installed a CA-signed certificate on the CA server, you are presented with the thumbprint of the CA-signed certificate. Otherwise, you are presented with a self-signed certificate.

- 5 Do not tick **Modify plugin script download location**, unless the NSX Manager is behind a firewall type of masking device.

This option allows you to enter an alternate IP address for NSX Manager. Note that putting NSX Manager behind a firewall of this type is not recommended.

- 6 Confirm that the vCenter Server status is **Connected**.

For example:



- 7 If vCenter Web Client is already open, log out of vCenter and log back in with the same Administrator role used to register NSX Manager with vCenter.

If you do not do this, vCenter Web Client will not display the **Networking & Security** icon on the **Home** tab.

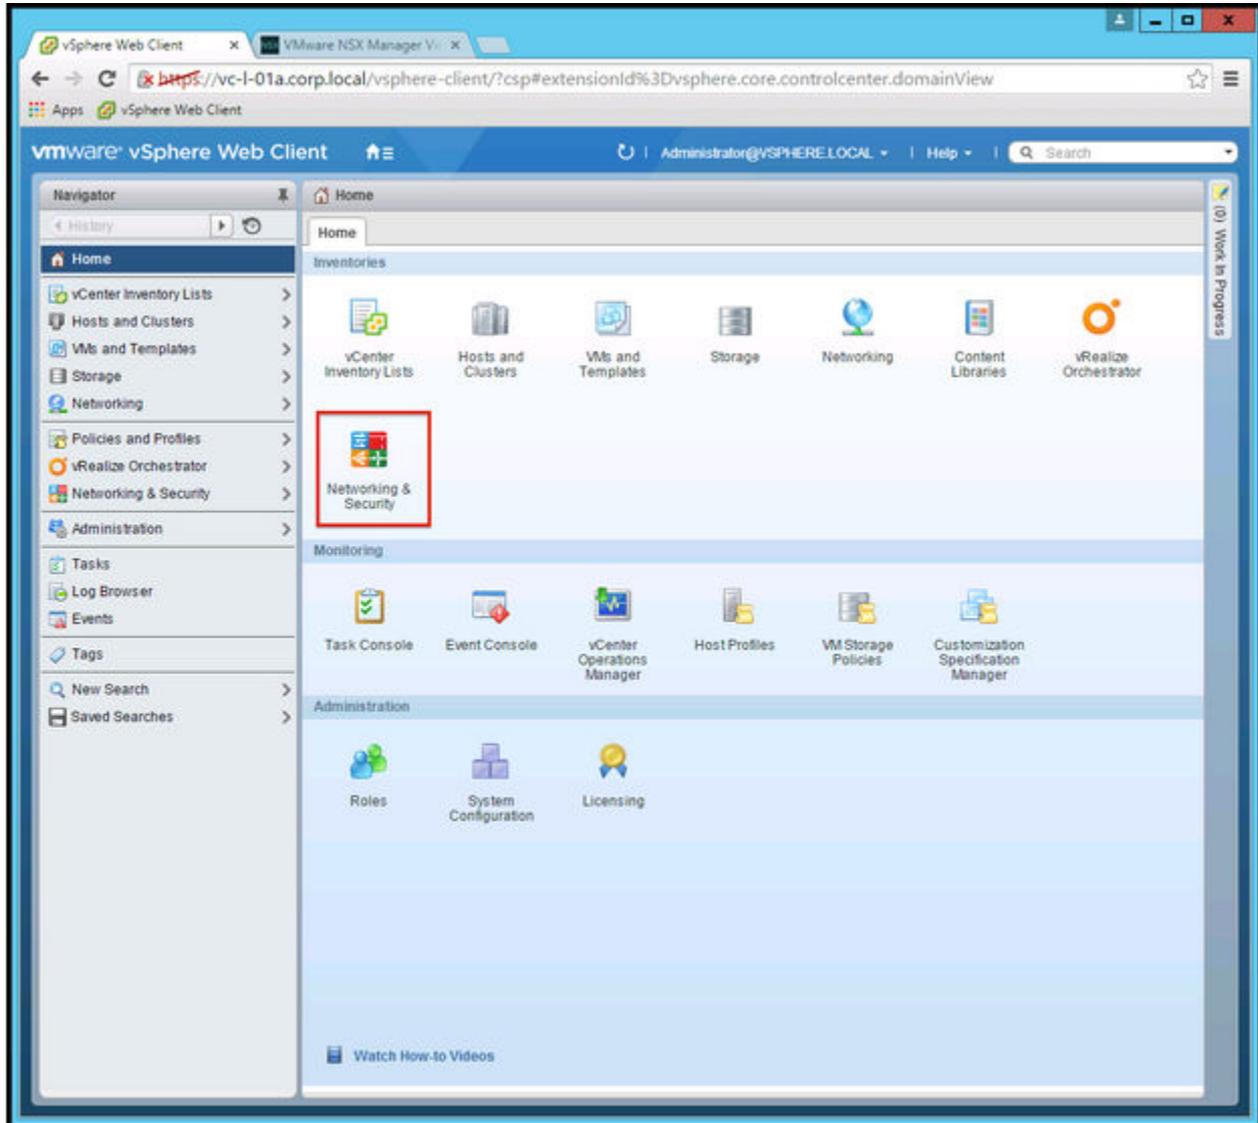
Click the **Networking & Security** icon and confirm that you can see the newly deployed NSX Manager.

What to do next

VMware recommends that you schedule a backup of NSX Manager data right after installing NSX Manager.

If you have an NSX partner solution, refer to partner documentation for information on registering the partner console with NSX Manager.

Login to the vSphere Web Client and make sure that the **Networking & Security** icon appears on the **Home** tab. If you are already logged in, the icon will not appear. Re-login to the vSphere Web Client to view the new icon.



You can now install and configure NSX components.

Specify a Syslog Server

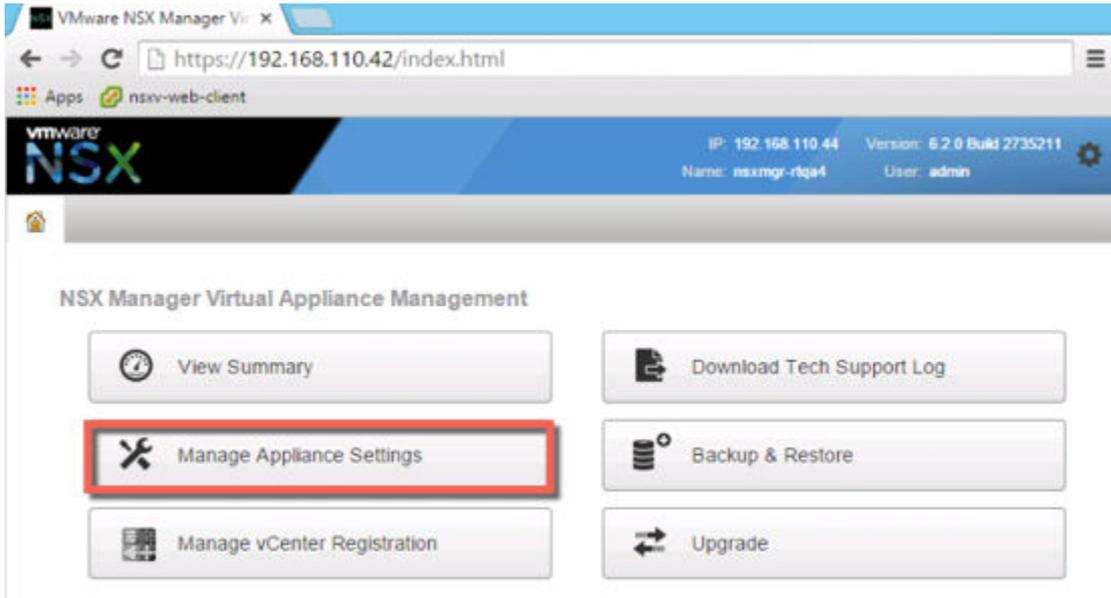
If you specify a syslog server, NSX Manager sends all audit logs and system events to the syslog server. Syslog data is useful for troubleshooting and reviewing data logged during installation and configuration. NSX Edge supports two syslog servers. NSX Manager and NSX Controllers support one syslog server.

Procedure

- 1 In a Web browser, navigate to the NSX Manager appliance GUI at <https://<nsx-manager-ip>> or <https://<nsx-manager-hostname>>.
- 2 Log in as admin with the password that you configured during NSX Manager installation.

3 Click **Manage Appliance Settings**.

For example:



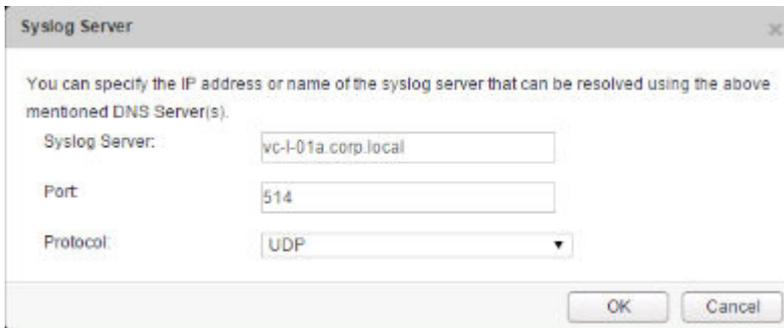
4 From the Settings panel, click **General**.

5 Click **Edit** next to **Syslog Server**.

6 Type the IP address or hostname, port, and protocol of the syslog server.

If you do not specify a port, the default UDP port for the IP address/host name of the syslog server is used.

For example:



7 Click **OK**.

vCenter Server remote logging is enabled, and logs are stored in your standalone syslog server.

Install and Assign NSX for vSphere License

You can install and assign an NSX for vSphere license after NSX Manager installation is complete by using the vSphere Web Client.

Starting in NSX 6.2.3, the default license upon install will be NSX for vShield Endpoint. This license enables use of NSX for deploying and managing vShield Endpoint for anti-virus offload capability only, and has hard enforcement to restrict usage of VXLAN, firewall, and Edge services, by blocking host preparation and creation of NSX Edges.

If you need other NSX features, including logical switches, logical routers, Distributed Firewall, or NSX Edge, you must either purchase an NSX license to use these features, or request an evaluation license for short-term evaluation of the features.

See the NSX License FAQ <https://www.vmware.com/files/pdf/products/nsx/vmware-nsx-editions-faq.pdf>

Procedure

- In vSphere 5.5, complete the following steps to add a license for NSX.
 - a Log in to the vSphere Web Client.
 - b Click **Administration** and then click **Licenses**.
 - c Click the **Solutions** tab.
 - d Select NSX for vSphere in the Solutions list. Click **Assign a license key**.
 - e Select **Assign a new license key** from the drop-down menu.
 - f Type the license key and an optional label for the new key.
 - g Click **Decode**.

Decode the license key to verify that it is in the correct format, and that it has enough capacity to license the assets.
 - h Click **OK**.
- In vSphere 6.0, complete the following steps to add a license for NSX.
 - a Log in to the vSphere Web Client.
 - b Click **Administration** and then click **Licenses**.
 - c Click the **Assets** tab, then the **Solutions** tab.
 - d Select NSX for vSphere in the Solutions list. From the **All Actions** drop-down menu, select **Assign license....**
 - e Click the **Add (+)** icon. Enter a license key and click **Next**. Add a name for the license, and click **Next**. Click **Finish** to add the license.
 - f Select the new license.
 - g (Optional) Click the **View Features** icon to view what features are enabled with this license. View the **Capacity** column to view the capacity of the license.
 - h Click **OK** to assign the new license to NSX.

What to do next

For more information about NSX licensing, see <http://www.vmware.com/files/pdf/vmware-product-guide.pdf>.

Exclude Virtual Machines from Firewall Protection

You can exclude a set of virtual machines from NSX distributed firewall protection.

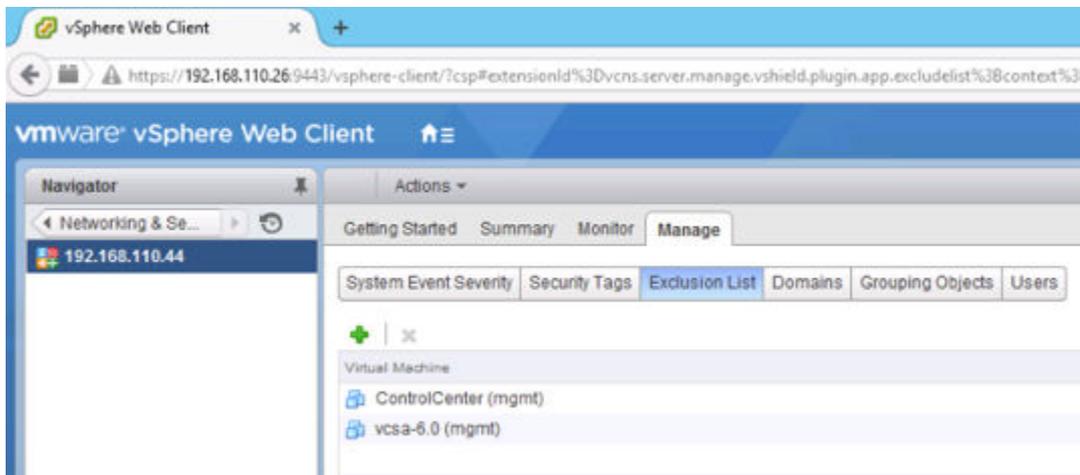
NSX Manager, NSX Controllers, and NSX Edge virtual machines are automatically excluded from NSX distributed firewall protection. In addition, VMware recommends that you place the following service virtual machines in the Exclusion List to allow traffic to flow freely.

- vCenter Server. It can be moved into a cluster that is protected by Firewall, but it must already exist in the exclusion list to avoid connectivity issues.
- Partner service virtual machines.
- Virtual machines that require promiscuous mode. If these virtual machines are protected by NSX distributed firewall, their performance may be adversely affected.
- The SQL server that your Windows-based vCenter uses.
- vCenter Web server, if you are running it separately.

Procedure

- 1 In the vSphere Web Client, click **Networking & Security**.
- 2 In **Networking & Security Inventory**, click **NSX Managers**.
- 3 In the **Name** column, click an NSX Manager.
- 4 Click the **Manage** tab and then click the **Exclusion List** tab.
- 5 Click the **Add (+)** icon.
- 6 Type the name of the virtual machine that you want to exclude and click **Add**.

For example:



7 Click **OK**.

If a virtual machine has multiple vNICs, all of them are excluded from protection. If you add vNICs to a virtual machine after it has been added to the Exclusion List, Firewall is automatically deployed on the newly added vNICs. In order to exclude these vNICs from firewall protection, you must remove the virtual machine from the Exclusion List and then add it back to the Exclusion List. An alternative workaround is to power cycle (power off and then power on) the virtual machine, but the first option is less disruptive.

Configuring the Primary NSX Manager

6

There is only one primary NSX Manager in a cross-vCenter NSX environment. Select which NSX Manager will be your primary NSX Manager, and complete the configuration tasks to complete NSX installation, assign the primary role to the NSX Manager, and create universal objects.

The primary NSX Manager is used to deploy a universal controller cluster that provides the control plane for the cross-vCenter NSX environment. The secondary NSX Managers do not have their own controller clusters.

This chapter includes the following topics:

- [Deploy NSX Controllers on the Primary NSX Manager](#)
- [Prepare Hosts on the Primary NSX Manager](#)
- [Configure VXLAN from the Primary NSX Manager](#)
- [Assign a Segment ID Pool to the Primary NSX Manager](#)
- [Assign Primary Role to NSX Manager](#)
- [Assign a Universal Segment ID Pool on the Primary NSX Manager](#)
- [Add a Universal Transport Zone on the Primary NSX Manager](#)
- [Add a Universal Logical Switch on the Primary NSX Manager](#)
- [Connect Virtual Machines to a Logical Switch](#)
- [Add a Universal Logical \(Distributed\) Router on the Primary NSX Manager](#)

Deploy NSX Controllers on the Primary NSX Manager

NSX controller is an advanced distributed state management system that provides control plane functions for NSX logical switching and routing functions. It serves as the central control point for all logical switches within a network and maintains information about all hosts, logical switches (VXLANs), and distributed logical routers. Controllers are required if you are planning to deploy 1) distributed logical routers or 2) VXLAN in unicast or hybrid mode. In cross-vCenter NSX, once the NSX Manager is assigned the primary role, its controller cluster becomes the universal controller cluster for the whole cross-vCenter NSX environment.

No matter the size of the NSX deployment, VMware requires that each NSX Controller cluster contain three controller nodes. Having a different number of controller nodes is not supported.

The cluster requires that each controller's disk storage system has a peak write latency of less than 300ms, and a mean write latency of less than 100ms. If the storage system does not meet these requirements, the cluster can become unstable and cause system downtime.

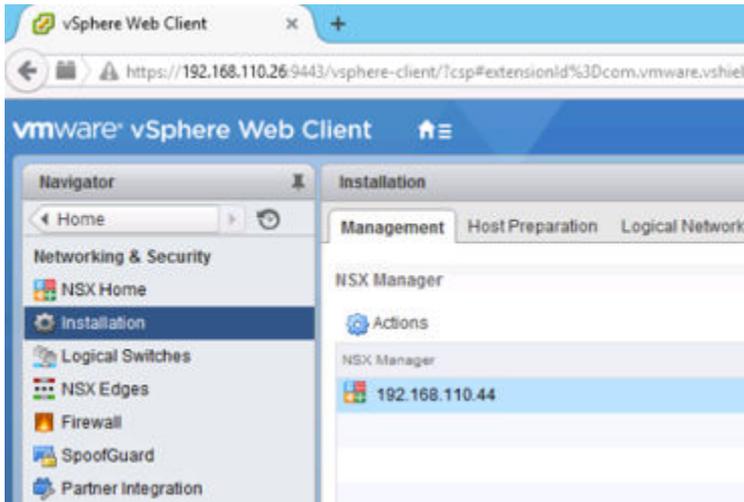
Prerequisites

- Before deploying NSX Controllers, you must deploy an NSX Manager appliance and register vCenter with NSX Manager.
- Determine the IP pool settings for your controller cluster, including the gateway and IP address range. DNS settings are optional. The NSX Controller IP network must have connectivity to the NSX Manager and to the management interfaces on the ESXi hosts.

Procedure

- 1 In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Management** tab.

For example:

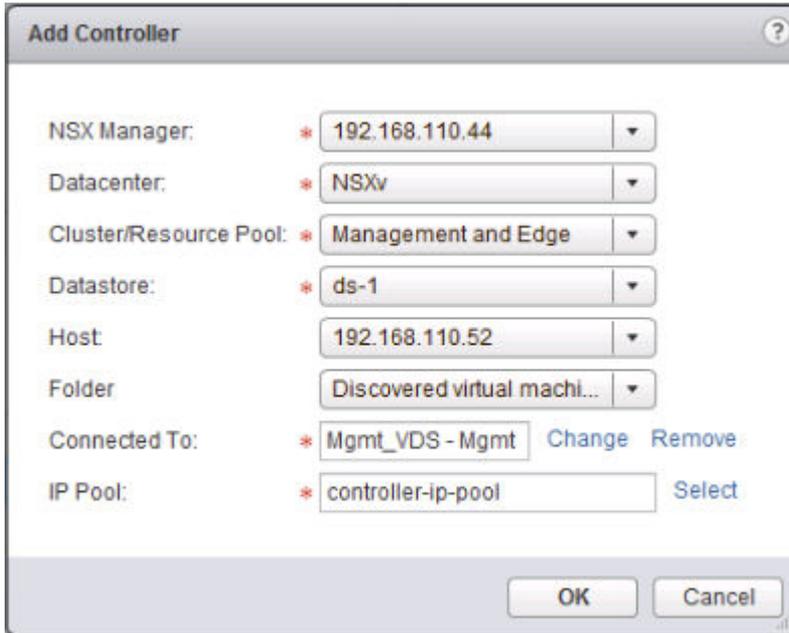


- 2 In the NSX Controller nodes section, click the **Add Node** (+) icon.

- 3 Enter the NSX Controller settings appropriate to your environment.

NSX Controllers should be deployed to a vSphere Standard Switch or vSphere Distributed Switch port group which is not VXLAN based and has connectivity to the NSX Manager, other controllers, and to hosts via IPv4.

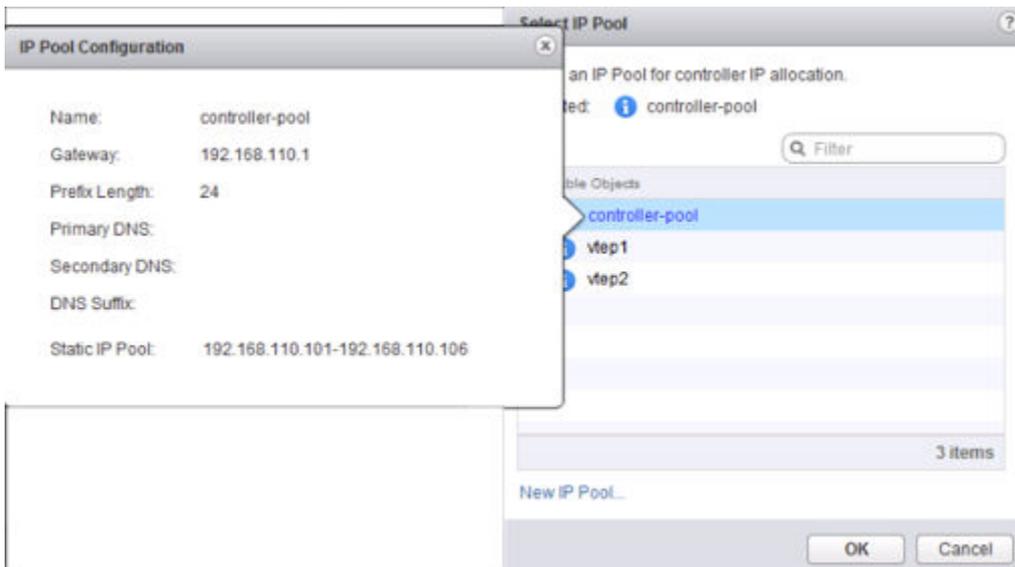
For example:



- 4 If you have not already configured an IP pool for your controller cluster, configure one now by clicking **New IP Pool**.

Individual controllers can be in separate IP subnets, if necessary.

For example:



5 Type and re-type a password for the controller.

Note Password must not contain the username as a substring. Any character must not consecutively repeat 3 or more times.

The password must be at least 12 characters and must follow 3 of the following 4 rules:

- At least one upper case letter
- At least one lower case letter
- At least one number
- At least one special character

6 After the first controller is completely deployed, deploy two additional controllers.

Having three controllers is mandatory. We recommend configuring a DRS anti-affinity rule to prevent the controllers from residing on the same host.

When successfully deployed, the controllers have a **Normal** status and display a green check mark.

SSH to each controller and make sure they can ping the host management interface IP addresses. If the ping fails, make sure all controllers have the correct default gateway. To view a controller routing table, run the **show network routes** command. To change a controller's default gateway run the **clear network routes** command followed by the **add network default-route <IP-address>** command.

Run the following commands to verify the control cluster is behaving as expected.

- `show control-cluster status`

Type	Status	Since
Join status:	Join complete	05/04 02:36:03
Majority status:	Connected to cluster majority	05/19 23:57:23
Restart status:	This controller can be safely restarted	05/19 23:57:12
Cluster ID:	ff3ebaeb-de68-4455-a3ca-4824e31863a8	
Node UUID:	ff3ebaeb-de68-4455-a3ca-4824e31863a8	
Role	Configured status	Active status
api_provider	enabled	activated
persistence_server	enabled	activated
switch_manager	enabled	activated
logical_manager	enabled	activated
directory_server	enabled	activated

For Join status, verify the controller node is reporting Join Complete.

For Majority status, verify the controller is connected to the cluster majority.

For Cluster ID, all the controller nodes in a cluster should have the same cluster ID.

For Configured status and Active status, verify that the all the controller roles are enabled and activated.

- `show control-cluster roles`

	Listen-IP	Master?	Last-Changed	Count
<code>api_provider</code>	Not configured	Yes	06/02 08:49:31	4
<code>persistence_server</code>	N/A	Yes	06/02 08:49:31	4
<code>switch_manager</code>	127.0.0.1	Yes	06/02 08:49:31	4
<code>logical_manager</code>	N/A	Yes	06/02 08:49:31	4
<code>directory_server</code>	N/A	Yes	06/02 08:49:31	4

One controller node will be the master for each role. In this example, a single node is the master for all roles.

If a master NSX Controller instance for a role fails, the cluster elects a new master for that role from the available NSX Controller instances.

NSX Controller instances are on the control plane, so an NSX Controller failure does not affect data plane traffic.

- `show control-cluster connections`

role	port	listening	open conns
<code>api_provider</code>	<code>api/443</code>	Y	2
<code>persistence_server</code>	<code>server/2878</code>	Y	2
	<code>client/2888</code>	Y	1
	<code>election/3888</code>	Y	0
<code>switch_manager</code>	<code>ovsmgmt/6632</code>	Y	0
	<code>openflow/6633</code>	Y	0
<code>system</code>	<code>cluster/7777</code>	Y	0

This command shows the intra-cluster communication status.

The controller cluster majority leader listens on port 2878 (as shown by the “Y” in the “listening” column). The other controller nodes will have a dash (-) in the “listening” column for Port 2878.

All other ports should be listening on all three controller nodes.

The open conns column shows the number of open connections that the controller node has with other controller nodes. In a 3-node controller cluster, the controller node should have no more than two open connections.

What to do next

Caution While a controller status is **Deploying**, do not add or modify logical switches or distributed routing in your environment. Also, do not continue to the host preparation procedure. After a new controller is added to the controller cluster, all controllers are inactive for a short while (no more than 5 minutes). During this downtime, any operation related to controllers—for example, host preparation—might have unexpected results. Even though host preparation might seem to complete successfully, the SSL certification might not establish correctly, thus causing issues in the VXLAN network.

If you need to delete a deployed controller, see Recover from an NSX Controller Failure in the *NSX Administration Guide*.

On the hosts where the NSX Controller nodes are first deployed, NSX enables automatic VM startup/shutdown. If the controller node VMs are later migrated to other hosts, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, VMware recommends that you check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html.

Prepare Hosts on the Primary NSX Manager

Host preparation is the process in which the NSX Manager 1) installs NSX kernel modules on ESXi hosts that are members of vCenter clusters and 2) builds the NSX control-plane and management-plane fabric. NSX kernel modules packaged in VIB files run within the hypervisor kernel and provide services such as distributed routing, distributed firewall, and VXLAN bridging capabilities.

To prepare your environment for network virtualization, you must install network infrastructure components on a per-cluster level for each vCenter server where needed. This deploys the required software on all hosts in the cluster. When a new host is added to this cluster, the required software is automatically installed on the newly added host.

If you are using ESXi in stateless mode (meaning that ESXi does not actively persist its state across reboots), you must download the NSX VIBs manually and make them part of the host image. You can find the download paths for the NSX VIBs from the page:

`https://<NSX_MANAGER_IP>/bin/vdn/nwfabric.properties`. Be aware that download paths can change for each release of NSX. Always check the `https://<NSX_MANAGER_IP>/bin/vdn/nwfabric.properties` page to get the appropriate VIBs. See Deploying VXLAN through Auto Deploy <https://kb.vmware.com/kb/2041972> for more information.

Prerequisites

- Register vCenter with NSX Manager and deploy NSX controllers.

- Verify that DNS reverse lookup returns a fully qualified domain name when queried with the IP address of NSX Manager. For example:

```
C:\Users\Administrator>nslookup 192.168.110.42
Server: localhost
Address: 127.0.0.1

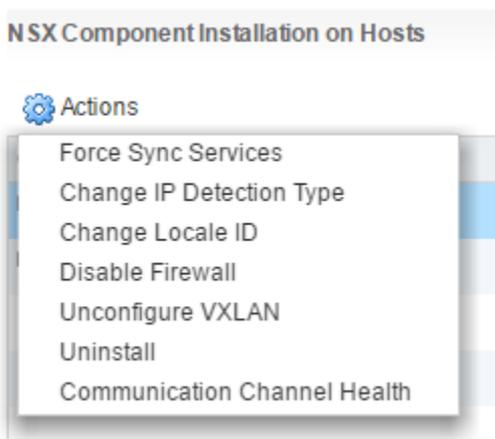
Name: nsxmgr-1-01a.corp.local
Address: 192.168.110.42
```

- Verify that hosts can resolve the DNS name of vCenter server.
- Verify that hosts can connect to vCenter Server on port 80.
- Verify that the network time on vCenter Server and ESXi hosts is synchronized.
- For each host cluster that will participate in NSX, verify that hosts within the cluster are attached to a common VDS.

For instance, say you have a cluster with Host1 and Host2. Host1 is attached to VDS1 and VDS2. Host2 is attached to VDS1 and VDS3. When you prepare a cluster for NSX, you can only associate NSX with VDS1 on the cluster. If you add another host (Host3) to the cluster and Host3 is not attached to VDS1, it is an invalid configuration, and Host3 will not be ready for NSX functionality.

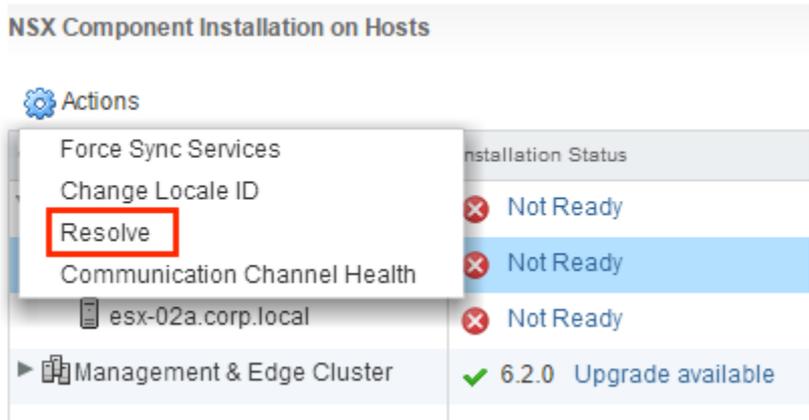
- If you have vSphere Update Manager (VUM) in your environment, you must disable it before preparing clusters for network virtualization. For information on how to check if VUM is enabled and how to disable it if necessary, see <http://kb.vmware.com/kb/2053782>.
- Before beginning the NSX host preparation process, always make sure that the cluster is in the resolved state---meaning that the **Resolve** option does not appear in the cluster's **Actions** list.

For example:

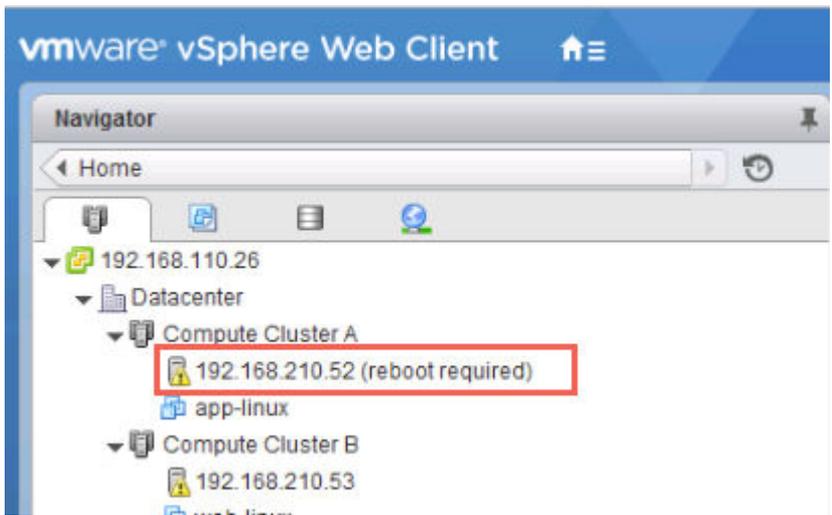


The **Resolve** option sometimes appears because one or more hosts in the cluster need to be rebooted.

Other times the **Resolve** option appears because there is an error condition that needs to be resolved. Click the **Not Ready** link to view the error. If you can, clear the error condition. If you cannot clear an error condition on a cluster, one workaround is to move the hosts to a new or different cluster and delete the old cluster.



Any process that removes one or more VIBs from a host requires a host reboot. Processes that remove VIBs include NSX upgrades, removing the NSX Manager plug-in from vCenter, removing a host from a cluster that has been prepared for NSX, and manually removing NSX VIBs from a host. When a host needs to be rebooted, the **reboot required** tag appears in the Hosts and Clusters view. For example:

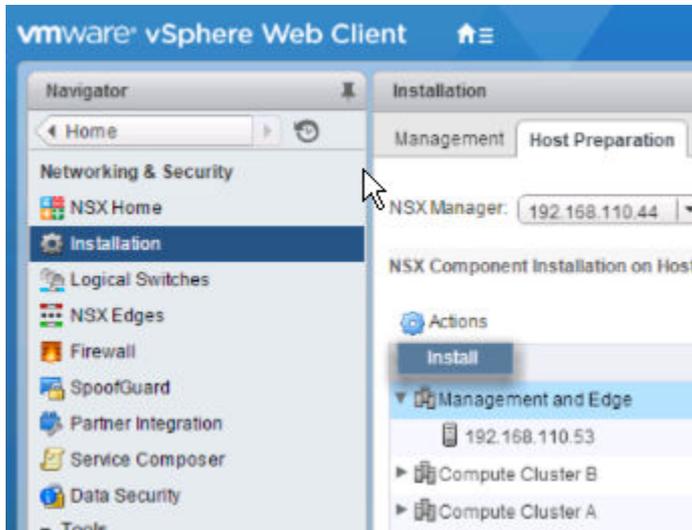


The required reboot does not happen automatically. Before rebooting a host, power off or move its VMs (or allow DRS to move them). Then, on the Host Preparation tab, click the **Resolve** option in the **Actions** list. The **Resolve** action places the hosts into maintenance mode, reboots the hosts, and then removes the hosts from maintenance mode. If you powered off a host's VMs, you must manually power them back on.

Procedure

- 1 In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Host Preparation** tab.

For example:



- 2 For all clusters that will require NSX logical switching, routing, and firewalls, click the gear icon and click **Install**.

A compute cluster (also known as a payload cluster) is a cluster with application VMs (web, database, and so on). If a compute cluster will have NSX switching, routing, or firewalls, you must click **Install** for the compute cluster.

In a shared "Management and Edge" cluster (as shown in the example), NSX Manager and controller VMs share a cluster with edge devices, such as distributed logical routers (DLRs) and edge services gateways (ESGs). In this case, it is important to click **Install** for the shared cluster.

Conversely, if Management and Edge each has a dedicated, non-shared cluster---as is recommended in a production environment---click **Install** for the Edge cluster but for the Management cluster.

Note While the installation is in progress, do not deploy, upgrade, or uninstall any service or component.

- 3 Monitor the installation until the **Installation Status** column displays a green check mark.

If the **Installation Status** column displays a red warning icon and says **Not Ready**, click **Resolve**. Clicking **Resolve** might result in a reboot of the host. If the installation is still not successful, click the warning icon. All errors are displayed. Take the required action and click **Resolve** again.

When the installation is complete, the **Installation Status** column displays **6.2 Uninstall** and the **Firewall** column displays **Enabled**. Both columns have a green check mark. If you see **Resolve** in the **Installation Status** column, click **Resolve** and then refresh your browser window.

VIBs are installed and registered with all hosts within the prepared cluster:

- esx-vsip
- esx-vxlan

To verify, SSH to each host and run the `esxcli software vib list | grep esx` command. In addition to displaying the VIBs, this command shows the version of the VIBs installed.

```
[root@host:~] esxcli software vib list | grep esx
...
esx-vsip      6.0.0-0.0.2732470   VMware VMwareCertified  2015-05-29
esx-vxlan    6.0.0-0.0.2732470   VMware VMwareCertified  2015-05-29
...
```

After host preparation, a host reboot is not required.

If you add a host to a prepared cluster, the NSX VIBs automatically get installed on the host.

If you move a host to an unprepared cluster, the NSX VIBs automatically get uninstalled from the host. In this case, a host reboot is required to complete the uninstall process.

Configure VXLAN from the Primary NSX Manager

The VXLAN network is used for Layer 2 logical switching across hosts, potentially spanning multiple underlying Layer 3 domains. You configure VXLAN on a per-cluster basis, where you map each cluster that is to participate in NSX to a vSphere distributed switch (VDS). When you map a cluster to a distributed switch, each host in that cluster is enabled for logical switches. The settings chosen here will be used in creating the VMkernel interface.

If you need logical routing and switching, all clusters that have NSX VIBs installed on the hosts should also have VXLAN transport parameters configured. If you plan to deploy distributed firewall only, you do not need to configure VXLAN transport parameters.

Prerequisites

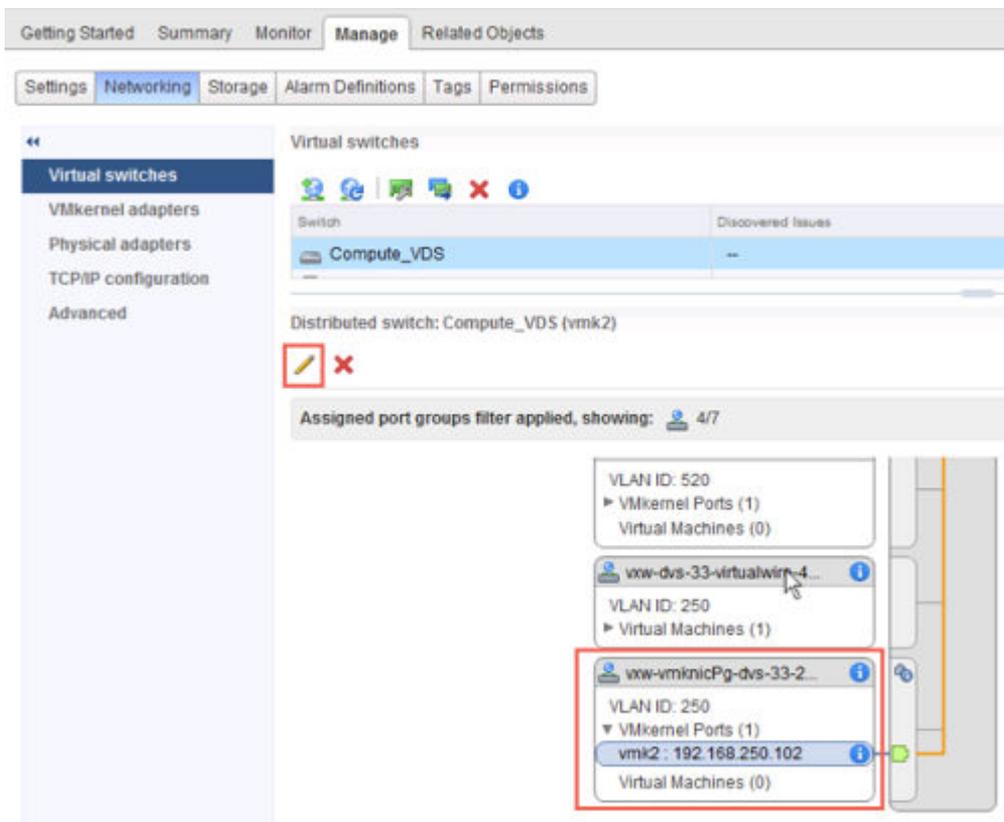
- All hosts within the cluster must be attached to a common VDS.
- NSX Manager must be installed.
- NSX controllers must be installed, unless you are using multicast replication mode for the control plane.
- Plan your NIC teaming policy. Your NIC teaming policy determines the load balancing and failover settings of the VDS.

Do not mix different teaming policies for different portgroups on a VDS where some use Etherchannel or LACPv1 or LACPv2 and others use a different teaming policy. If uplinks are shared in these different teaming policies, traffic will be interrupted. If logical routers are present, there will be routing problems. Such a configuration is not supported and should be avoided.

The best practice for IP hash-based teaming (EtherChannel, LACPv1 or LACPv2) is to use all uplinks on the VDS in the team, and do not have portgroups on that VDS with different teaming policies. For more information and further guidance, see the VMware® NSX for vSphere Network Virtualization Design Guide at <https://communities.vmware.com/docs/DOC-27683>.

- Plan the IP addressing scheme for the VXLAN tunnel end points (VTEPs). VTEPs are the source and destination IP addresses used in the external IP header to uniquely identify the ESX hosts originating and terminating the VXLAN encapsulation of frames. You can use either DHCP or manually configured IP pools for VTEP IP addresses.

If you want a specific IP address to be assigned to a VTEP, you can either 1) use a DHCP fixed address or reservation that maps a MAC address to a specific IP address in the DHCP server or 2) use an IP pool and then manually edit the VTEP IP address assigned to the vmknic in **Manage > Networking > Virtual Switches**. For example:



VTEPs have an associated VLAN ID. You can, however, specify VLAN ID = 0 for VTEPs, meaning frames will be untagged.

- For clusters that are members of the same VDS, the VLAN ID for the VTEPs and the NIC teaming must be the same.
- As a best practice, export the VDS configuration before preparing the cluster for VXLAN. See <http://kb.vmware.com/kb/2034602>.

Procedure

- 1 In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Host Preparation** tab.
- 2 From the NSX Manager dropdown menu, select the Primary NSX Manager.
- 3 Click **Not Configured** in the **VXLAN** column.
- 4 Set up logical networking.

This involves selecting a VDS, a VLAN ID, an MTU size, an IP addressing mechanism, and a NIC teaming policy.

The MTU for each switch must be set to 1550 or higher. By default, it is set to 1600. If the vSphere distributed switch (VDS) MTU size is larger than the VXLAN MTU, the VDS MTU will not be adjusted down. If it is set to a lower value, it will be adjusted to match the VXLAN MTU. For example, if the VDS MTU is set to 2000 and you accept the default VXLAN MTU of 1600, no changes to the VDS MTU will be made. If the VDS MTU is 1500 and the VXLAN MTU is 1600, the VDS MTU will be changed to 1600.

These example screens show a configuration for a management cluster with IP pool address range of 182.168.150.1-192.168.150.100, backed by VLAN 150, and with a failover NIC teaming policy.

The screenshot shows a dialog box titled "Configure VXLAN networking" with a help icon and a right arrow. Below the title bar, it says "Configuring all hosts in cluster 'Management and Edge' for VXLAN networking." The configuration fields are as follows:

- Switch: * Mgmt_VDS (dropdown menu)
- VLAN: * 150 (text input)
- MTU: * 1600 (text input)
- VMKNic IP Addressing: * Use DHCP, Use IP Pool
- IP Pool: New IP Pool... (dropdown menu)
- VMKNic Teaming Policy: * Fail Over (dropdown menu)
- VTEP: * 1 (text input)

At the bottom right, there are "OK" and "Cancel" buttons.

The number of VTEPs is not editable in the UI. The VTEP number is set to match the number of dvUplinks on the vSphere distributed switch being prepared.

Add Static IP Pool

Name: * mgmt-edge-ip-pool

Gateway: * 192.168.150.1
A gateway can be any IPv4 or IPv6 address.

Prefix Length: * 24

Primary DNS: 192.168.110.10

Secondary DNS:

DNS Suffix: corp.local

Static IP Pool: * 192.168.150.1-192.168.150.100

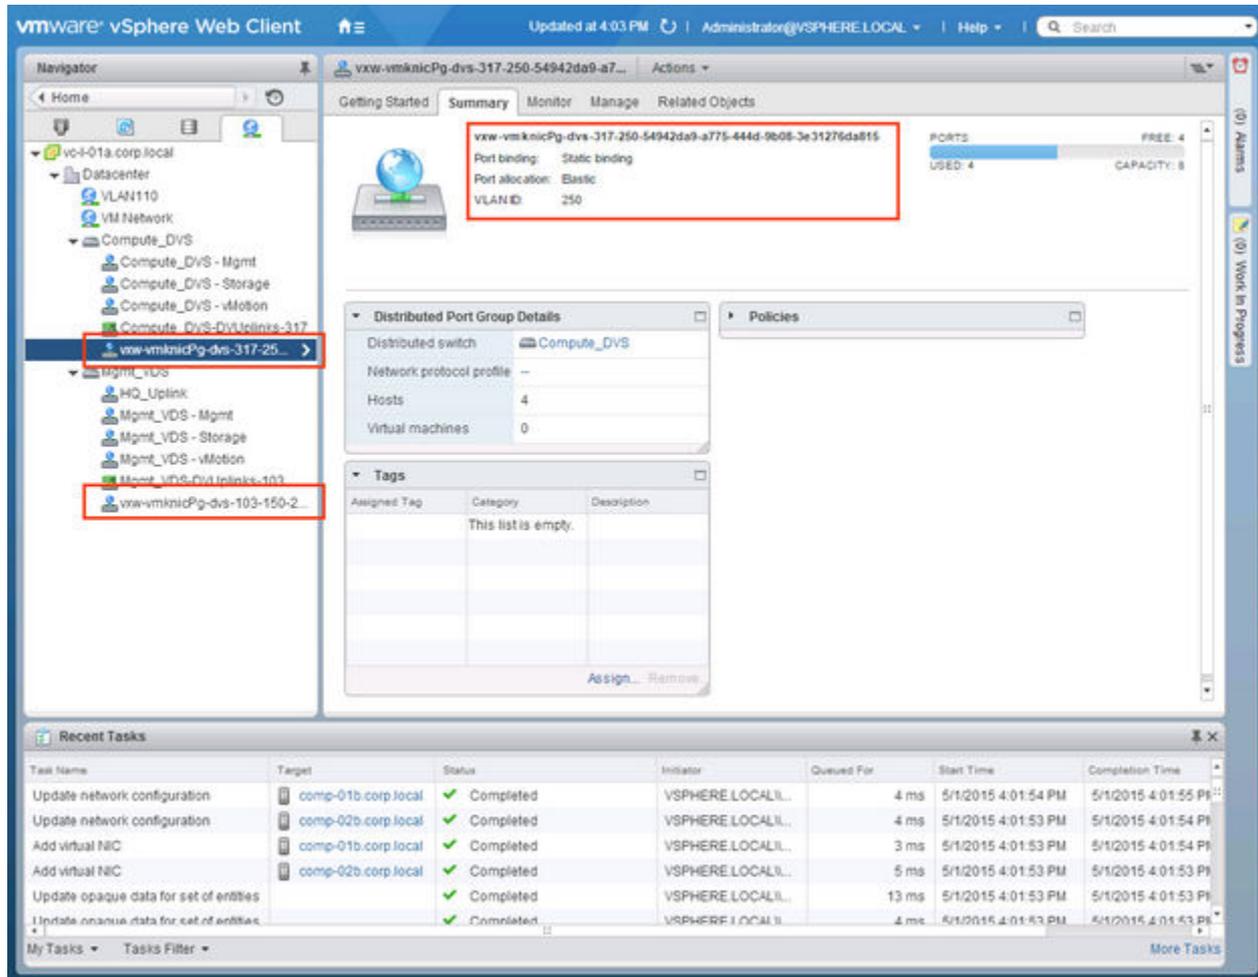
for example 192.168.1.2-192.168.1.100 or
abcd:87:87::10-abcd:87:87::20

OK Cancel

For compute clusters, you may want to use different IP address settings (for example, 192.168.250.0/24 with VLAN 250). This would depend on how the physical network is designed, and likely won't be the case in small deployments.

Configuring VXLAN results in the creation of new distributed port groups.

For example:



Assign a Segment ID Pool to the Primary NSX Manager

VXLAN segments are built between VXLAN tunnel end points (VTEPs). Each VXLAN tunnel has a segment ID. You must specify a segment ID pool for the primary NSX Manager to isolate your network traffic. If an NSX controller is not deployed in your environment, you must also add a multicast address range to spread traffic across your network and avoid overloading a single multicast address.

If you wish to configure multiple segment ID ranges in a single vCenter---for example, 5000-5999, 7000-7999---this is not currently supported in the vSphere Web Client UI, but you can do this using the NSX API.

```
POST https://<nsxmgr-ip>/api/2.0/vdn/config/segments

<segmentRange>
<name>Segment ID Pool 1</name>
<begin>5000</begin>
<end>5999</end>
</segmentRange>
```

```
POST https://<nsxmgr-ip>/api/2.0/vdn/config/segments
```

```
<segmentRange>
<name>Segment ID Pool 2</name>
<begin>7000</begin>
<end>7999</end>
</segmentRange>
```

Prerequisites

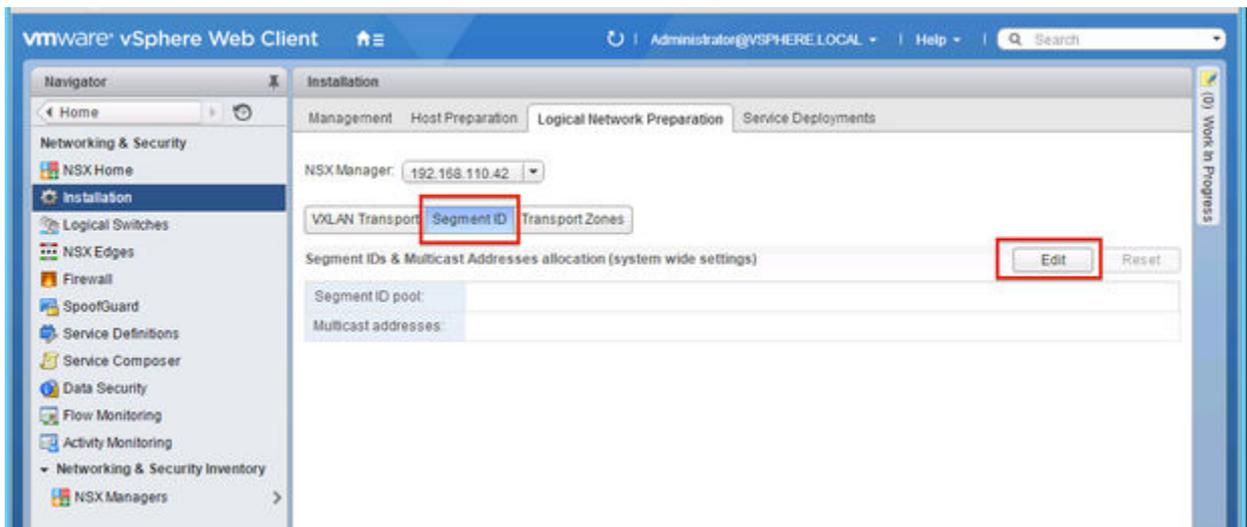
When determining the size of each segment ID pool, keep in mind that the segment ID range controls the number of logical switches that can be created. Choose a small subset of the 16 million potential VNIs. You should not configure more than 10,000 VNIs in a single vCenter because vCenter limits the number of dvPortgroups to 10,000.

If VXLAN is in place in another NSX deployment, consider which VNIs are already in use and avoid overlapping VNIs. Non-overlapping VNIs is automatically enforced within a single NSX Manager and vCenter environment. Local VNI ranges can't be overlapping. However, it's important for you make sure that VNIs do not overlap in your separate NSX deployments. Non-overlapping VNIs is useful for tracking purposes and helps to ensure that your deployments are ready for a cross-vCenter environment.

Procedure

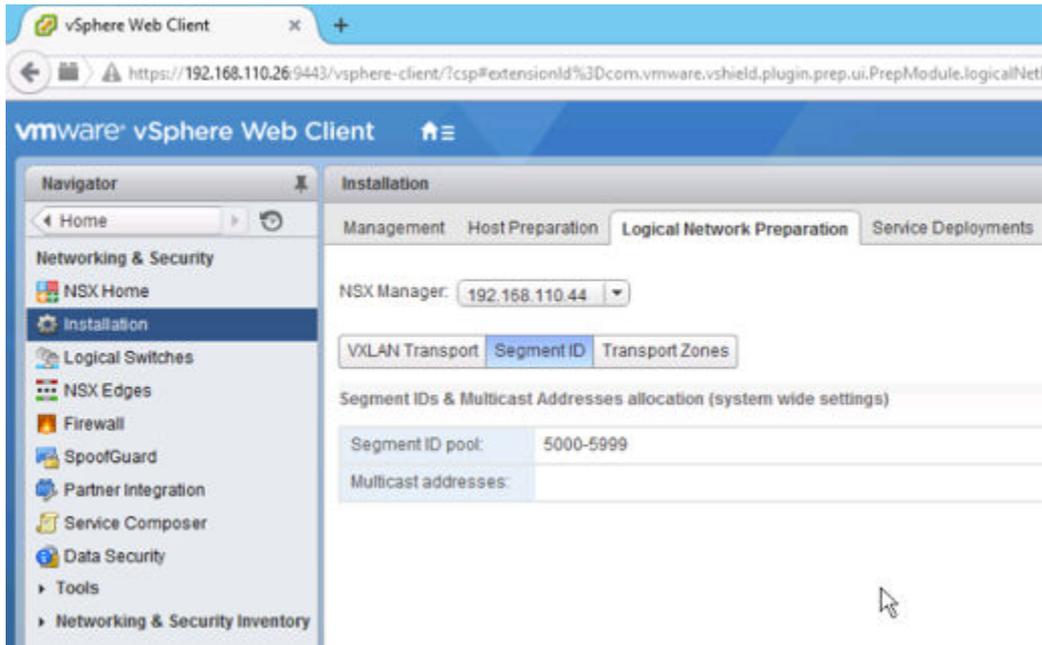
- 1 In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Logical Network Preparation** tab.
- 2 Click **Segment ID > Edit**.

For example:



- 3 Type a range for segment IDs, such as **5000–5999**.

For example:



- 4 If any of your transport zones will use multicast or hybrid replication mode, add a multicast address or a range of multicast addresses.

Having a range of multicast addresses spreads traffic across your network, prevents the overloading of a single multicast address, and better contains BUM replication.

When VXLAN multicast and hybrid replication modes are configured and working correctly, a copy of multicast traffic is delivered only to hosts that have sent IGMP join messages. Otherwise, the physical network floods all multicast traffic to all hosts within the same broadcast domain. To avoid such flooding, you must do the following:

- Make sure that the underlying physical switch is configured with an MTU larger than or equal to 1600.
- Make sure that the underlying physical switch is correctly configured with IGMP snooping and an IGMP querier in network segments that carry VTEP traffic.
- Make sure that the transport zone is configured with the recommended multicast address range. The recommended multicast address range starts at 239.0.1.0/24 and excludes 239.128.0.0/24.

Do not use 239.0.0.0/24 or 239.128.0.0/24 as the multicast address range, because these networks are used for local subnet control, meaning that the physical switches flood all traffic that uses these addresses. For more information about unusable multicast addresses, see <https://tools.ietf.org/html/draft-ietf-mboned-ipv4-mcast-unusable-01>.

When you configure logical switches, each logical switch receives a segment ID from the pool.

Assign Primary Role to NSX Manager

The primary NSX Manager runs the controller cluster. Additional NSX Managers are secondary. The controller cluster that is deployed by the primary NSX Manager is a shared object and is referred to as the universal controller cluster. Secondary NSX Managers automatically import the universal controller cluster. There can be one primary NSX Manager and up to seven secondary NSX Managers in a cross-vCenter NSX environment.

NSX Managers can have one of four roles:

- Primary
- Secondary
- Standalone
- Transit

To view the role of an NSX Manager, log in to the vCenter linked to the NSX Manager, and navigate to **Home > Networking & Security > Installation** and select the **Management** tab. The role is displayed in the Role column in the NSX Managers section. If there is no Role column shown, the NSX Manager has the standalone role.

Prerequisites

- The version of the NSX Managers (the primary NSX Manager and NSX Managers that will be assigned the secondary role) must match.
- The node IDs of the primary NSX Manager and the NSX Managers that will be assigned the secondary role must be present and different. NSX Manager instances deployed from OVA files have unique node IDs. An NSX Manager deployed from a template (as in when you convert a virtual machine to a template) will have the same node ID as the original NSX Manager used to create the template, and these two NSX Managers cannot be used in the same cross-vCenter NSX installation.
- Each NSX Manager must be registered with a separate and unique vCenter Server.
- The UDP ports used for VXLAN must be the same for all NSX Managers.
- When assigning the secondary role to a NSX Manager, the vCenter linked to it must not have any deployed NSX Controllers.
- The segment ID pool of the NSX Manager being assigned the secondary role must not overlap with the segment ID pools of the primary NSX Manager or the segment ID pool of any other secondary NSX Manager .

- The NSX Manager being assigned the secondary role must have the standalone or transit role.

Note You can view the NSX Manager node ID with the following REST API call:

```
GET https://NSX-Manager-IP-Address/api/2.0/services/vsmconfig
```

You can view the UDP ports used for VXLAN with the following REST API call:

```
GET https://NSX-Manager-IP-Address/api/2.0/vdn/config/vxlan/udp/port
```

Procedure

- 1 Log into the vCenter linked to the primary NSX Manager using the vSphere Web Client.
- 2 Navigate to **Home > Networking & Security > Installation** and select the **Management** tab.
- 3 Select the NSX Manager you want to assign as the primary and click **Actions** and then **Assign Primary Role**.

The selected NSX Manager is assigned the primary role.

Assign a Universal Segment ID Pool on the Primary NSX Manager

The universal segment ID pool specifies a range for use when building logical network segments. Cross-vCenter NSX deployments use a unique universal segment ID Pool to ensure that the universal logical switches VXLAN network identifiers (VNIs) are consistent across all secondary NSX Managers.

The universal segment ID is defined once on the primary NSX Manager and then synced to all of the secondary NSX Managers. The universal segment ID range controls the number of universal logical switches that can be created. Note that the segment ID range must be unique across any NSX Manager that you plan use in a cross-vCenter NSX deployment. This example uses a high range to provide future scalability.

Procedure

- 1 In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Logical Network Preparation** tab.
- 2 Click **Segment ID > Edit**.
- 3 Enter a range for universal segment IDs, such as 900000-909999.

Note This will be the range of IDs used by Universal Logical Switches. The Universal Segment ID pool must be unique across any NSX Managers that will be used in a cross-vCenter NSX deployment.

- 4 If any of your transport zones will use multicast or hybrid replication mode, check **Enable Universal multicast addressing** and type a universal multicast address or a range of universal multicast addresses.

Caution In a cross-vCenter NSX environment, if you select hybrid mode, you must ensure that the multicast address specified does not conflict with any other multicast addresses assigned on any NSX Manager in the environment.

Having a range of multicast addresses spreads traffic across your network, prevents the overloading of a single multicast address, and better contains BUM replication. The recommended multicast address range starts at 239.0.1.0/24 and excludes 239.128.0.0/24.

Do not use 239.0.0.0/24 or 239.128.0.0/24 as the multicast address range, because these networks are used for local subnet control, meaning that the physical switches flood all traffic that uses these addresses. For more information about unusable multicast addresses, see <https://tools.ietf.org/html/draft-ietf-mboned-ipv4-mcast-unusable-01>.

Later, after you configure universal logical switches, each universal logical switch receives a universal segment ID from the pool.

Add a Universal Transport Zone on the Primary NSX Manager

Universal transport zones control the hosts that a universal logical switch can reach. A universal transport zone is created by the primary NSX Manager and is replicated to the secondary NSX Managers. Universal transport zones can span one or more vSphere clusters across the cross-vCenter NSX environment.

After creation, a universal transport zone is available on all secondary NSX Managers in the cross-vCenter NSX environment. There can be only one universal transport zone.

Prerequisites

Configure a universal transport zone after a primary NSX Manager has been created.

Procedure

- 1 In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Logical Network Preparation** tab.
- 2 Click **Transport Zones** and click the **New Transport Zone (+)** icon.
- 3 Select **Mark this object for universal synchronization**.

These transport zones will be synchronized with the secondary NSX Managers.

4 Select the control plane mode:

- **Multicast:** Multicast IP addresses in the physical network are used for the control plane. This mode is recommended only when you are upgrading from older VXLAN deployments. Requires PIM/IGMP in the physical network.
- **Unicast:** The control plane is handled by an NSX controller. All unicast traffic leverages optimized headend replication. No multicast IP addresses or special network configuration is required.
- **Hybrid:** Offloads local traffic replication to the physical network (L2 multicast). This requires IGMP snooping on the first-hop switch and access to an IGMP querier in each VTEP subnet, but does not require PIM. The first-hop switch handles traffic replication for the subnet.

5 Select the clusters to add to the transport zone.

The universal transport zone is available on all NSX Managers in the cross-vCenter NSX environment.

Name	Description	Control Plane Mode	Logical Switches
Transport-Zone		Unicast	1
Universal-Transport-Zone		Unicast	4

What to do next

Next, create a universal logical switch.

Add a Universal Logical Switch on the Primary NSX Manager

In a cross-vCenter NSX deployment, you can create universal logical switches, which can span all vCenters. The transport zone type determines whether the new switch is a logical switch or a universal logical switch. When you add a logical switch to a universal transport zone, the logical switch is universal.

Prerequisites

Table 6-1. Prerequisites for creating a logical switch or universal logical switch

Logical Switch	Universal Logical Switch
<ul style="list-style-type: none"> ■ vSphere distributed switches must be configured. ■ NSX Manager must be installed. ■ Controllers must be deployed. ■ Host clusters must be prepared for NSX. ■ VXLAN must be configured. ■ A transport zone must be created. ■ A segment ID pool must be configured. 	<ul style="list-style-type: none"> ■ vSphere distributed switches must be configured. ■ NSX Manager must be installed. ■ Controllers must be deployed. ■ Host clusters must be prepared for NSX. ■ VXLAN must be configured. ■ A primary NSX Manager must be assigned. ■ A universal transport zone must be created. ■ A universal segment ID pool must be configured.

Procedure

- 1 In the vSphere Web Client, navigate to **Home > Networking & Security > Logical Switches**.
- 2 Select the primary NSX Manager.

- 3 Click the **New Logical Switch** (+) icon.
- 4 Type a name and optional description for the logical switch.
- 5 In the Transport Zone section, click **Change** to select a Transport Zone. Select the universal transport zone to create a universal logical switch.

By default, the logical switch inherits the control plane replication mode from the transport zone. You can change it to one of the other available modes. The available modes are unicast, hybrid, and multicast.

The case in which you might want to override the inherited transport zone's control plane replication mode for an individual logical switch is when the logical switch you are creating has significantly different characteristics in terms of the amount of BUM traffic it will to carry. In this case, you might create a transport zone that uses as unicast mode, and use hybrid or multicast mode for the individual logical switch.

Important If you create a universal logical switch and select hybrid as the replication mode, you must ensure that the multicast address used does not conflict with any other multicast addresses assigned on any NSX Manager in the environment.

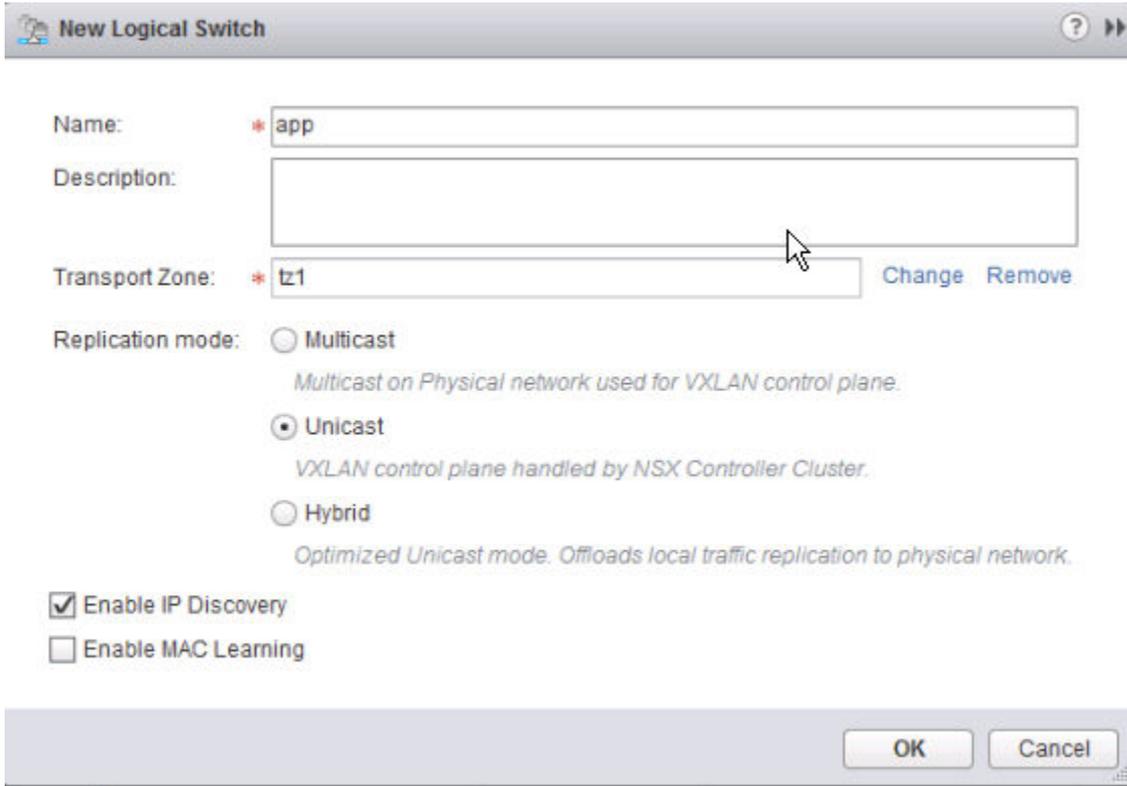
- 6 (Optional) Click **Enable IP Discovery** to enable ARP suppression.

This setting minimizes ARP traffic flooding within individual VXLAN segments---in other words, between VMs connected to the same logical switch. IP discovery is enabled by default.

- (Optional) Click **Enable MAC learning** if your VMs have multiple MAC addresses or are using virtual NICs that are trunking VLANs.

Enabling MAC learning builds a VLAN/MAC pair learning table on each vNIC. This table is stored as part of the dvfilter data. During vMotion, dvfilter saves and restores the table at the new location. The switch then issues RARPs for all the VLAN/MAC entries in the table.

This example shows the app logical switch with default settings.



DB-Tier-00 is logical switch connected to a transport zone. It is available only on the NSX Manager on which it was created.

DB-Tier-01 is a universal logical switch connected to a universal transport zone. It is available on any of the NSX Managers in the cross-vCenter NSX environment.

The logical switch and the universal logical switch have segment IDs from different segment ID pools.

DB-Tier-00	✓ Normal	Transport-Zone	10000	Unicast
DB-Tier-01	✓ Normal	Universal-Transport-Zone	900003	Unicast

What to do next

Add VMs to a universal logical switch.

Optionally, create a universal logical router and attach it to your universal logical switch to enable connectivity between VMs that are connected to different universal logical switches.

Connect Virtual Machines to a Logical Switch

You can connect virtual machines to a logical switch or universal logical switch.

Procedure

- 1 In **Logical Switches**, select the logical switch to which you want to add virtual machines.
- 2 Click the **Add Virtual Machine** () icon.
- 3 Select the virtual machines you want to add to the logical switch.
- 4 Select the vNICs that you want to connect.
- 5 Click **Next**.
- 6 Review the vNICs you selected.
- 7 Click **Finish**.

Add a Universal Logical (Distributed) Router on the Primary NSX Manager

Logical router kernel modules in the host perform routing between VXLAN networks, and between virtual and physical networks. An NSX Edge appliance provides dynamic routing ability if needed. A universal logical router provides east-west routing between universal logical switches.

Prerequisites

- You must have been assigned the Enterprise Administrator or NSX Administrator role.
- You must have an operational controller cluster in your environment before installing a logical router.
- You must create a local segment ID pool, even if you have no plans to create NSX logical switches.
- A logical router cannot distribute routing information to hosts without the help of NSX controllers. A logical router relies on NSX controllers to function, while edge services gateways (ESGs) do not. Make sure the controller cluster is up and available before creating or changing a logical router configuration.
- If a logical router is to be connected to VLAN dvPortgroups, ensure that all hypervisor hosts with a logical router appliance installed can reach each other on UDP port 6999 for logical router VLAN-based ARP proxy to work.
- Logical router interfaces and bridging interfaces cannot be connected to a dvPortgroup with the VLAN ID set to 0.
- A given logical router instance cannot be connected to logical switches that exist in different transport zones. This is to ensure that all logical switches and logical router instances are aligned.

- A logical router cannot be connected to VLAN-backed portgroups if that logical router is connected to logical switches spanning more than one vSphere distributed switch (VDS). This is to ensure correct alignment of logical router instances with logical switch dvPortgroups across hosts.
- Logical router interfaces should not be created on two different distributed portgroups (dvPortgroups) with the same VLAN ID if the two networks are in the same vSphere distributed switch.
- Logical router interfaces should not be created on two different dvPortgroups with the same VLAN ID if two networks are in different vSphere distributed switches, but the two vSphere distributed switches share the same hosts. In other words, logical router interfaces can be created on two different networks with the same VLAN ID if the two dvPortgroups are in two different vSphere distributed switches, as long as the vSphere distributed switches do not share a host.
- Unlike NSX version 6.0 and 6.1, NSX version 6.2 allows logical router-routed logical interfaces (LIFs) to be connected to a VXLAN that is bridged to a VLAN.
- When selecting placement of a logical router virtual appliance, avoid placing it on the same host as one or more of its upstream ESGs if you use ESGs in an ECMP setup. You can use DRS anti-affinity rules to enforce this, thus reducing the impact of host failure on logical router forwarding. This guideline does not apply if you have one upstream ESG by itself or in HA mode. For more information, see the *VMware NSX for vSphere Network Virtualization Design Guide* at <https://communities.vmware.com/docs/DOC-27683>.
- Determine if you need to enable local egress. Local egress allows you to selectively send routes to hosts. You may want this ability if your NSX deployment spans multiple sites. See [Cross-vCenter NSX Topologies](#) for more information. You cannot enable local egress after the universal logical router has been created.

Procedure

- 1 In the vSphere Web Client, navigate to **Home > Networking & Security > NSX Edges**.
- 2 Select the Primary NSX Manager to add a universal logical router.
- 3 Click the **Add** () icon.
- 4 Select **Universal Logical (Distributed) Router**.
- 5 (Optional) Enable local egress.
- 6 Type a name for the device.

This name appears in your vCenter inventory. The name should be unique across all logical routers within a single tenant.

Optionally, you can also enter a hostname. This name appears in the CLI. If you do not specify the host name, the Edge ID, which gets created automatically, is displayed in the CLI.

Optionally, you can enter a description and tenant.

7 (Optional) Deploy an edge appliance.

Deploy Edge Appliance is selected by default. An edge appliance (also called a logical router virtual appliance) is required for dynamic routing and the logical router appliance's firewall, which applies to logical router pings, SSH access, and dynamic routing traffic.

You can deselect the edge appliance option if you require only static routes, and do not want to deploy an Edge appliance. You cannot add an Edge appliance to the logical router after the logical router has been created.

8 (Optional) Enable High Availability.

Enable High Availability is not selected by default. Select the Enable High Availability check box to enable and configure high availability. High availability is required if you are planning to do dynamic routing.

9 Type and re-type a password for the logical router.

The password must be 12-255 characters and must contain the following:

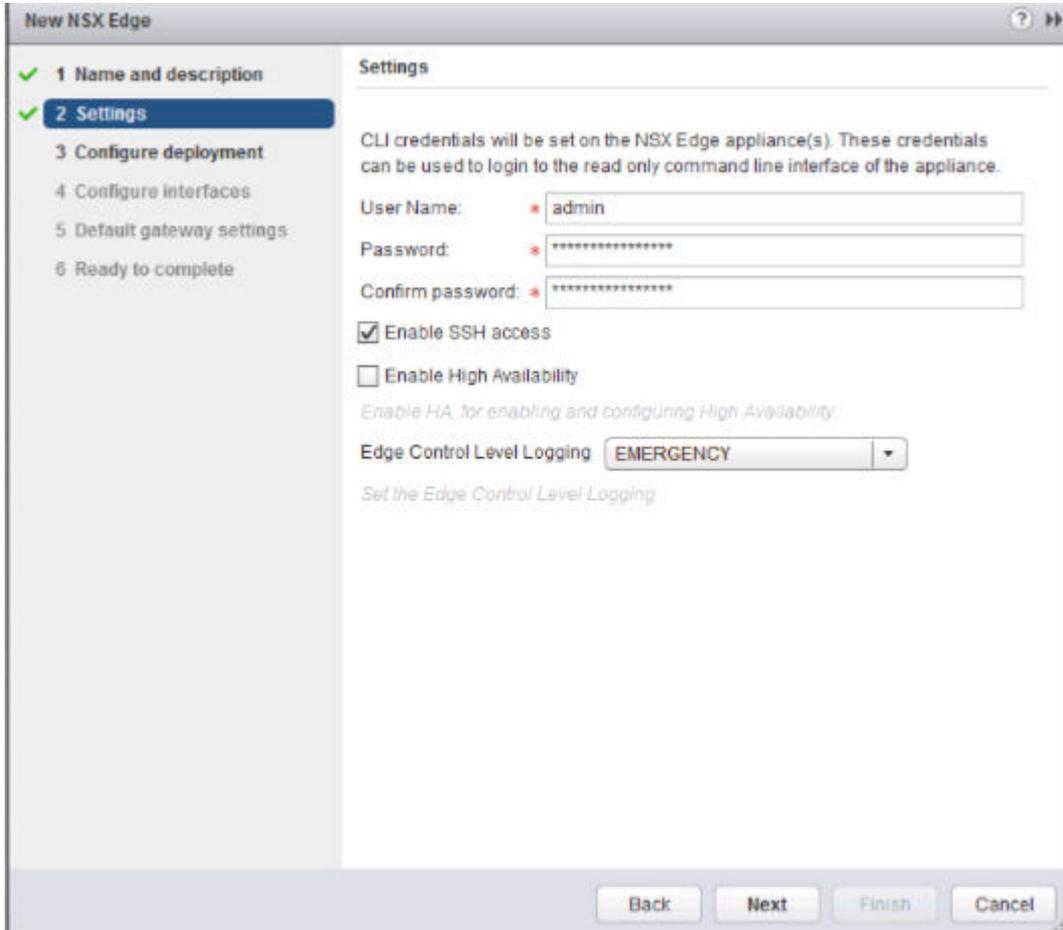
- At least one upper case letter
- At least one lower case letter
- At least one number
- At least one special character

10 (Optional) Enable SSH and set the log level.

By default, SSH is disabled. If you do not enable SSH, you can still access the logical router by opening the virtual appliance console. Enabling SSH here causes the SSH process to run on the logical router virtual appliance, but you will also need to adjust the logical router firewall configuration manually to allow SSH access to the logical router's protocol address. The protocol address is configured when you configure dynamic routing on the logical router.

By default, the log level is emergency.

For example:



11 Configure deployment.

- ◆ If you did not select **Deploy NSX Edge**, the **Add (+)** icon is grayed out. Click **Next** to continue with configuration.
- ◆ If you selected **Deploy NSX Edge**, enter the settings for the logical router virtual appliance that will be added to your vCenter inventory.

For example:

The screenshot shows a dialog box titled "Edit NSX Edge Appliance" with a help icon. The instruction reads: "Specify placement parameters for the NSX Edge appliance." The configuration fields are as follows:

Cluster/Resource Pool:	Management and Edg...
Datastore:	ds-1
Host:	192.168.110.52
Folder:	Discovered virtual mac...

Buttons: OK, Cancel

12 Configure interfaces.

On logical routers, only IPv4 addressing is supported.

In the HA Interface Configuration, if you selected **Deploy NSX Edge** you must connect the interface to a distributed port group. It is recommended to use a VXLAN logical switch for the HA interface. An IP address for each of the two NSX Edge appliances is chosen from the link local address space, 169.250.0.0/16. No further configuration is necessary to configure the HA service.

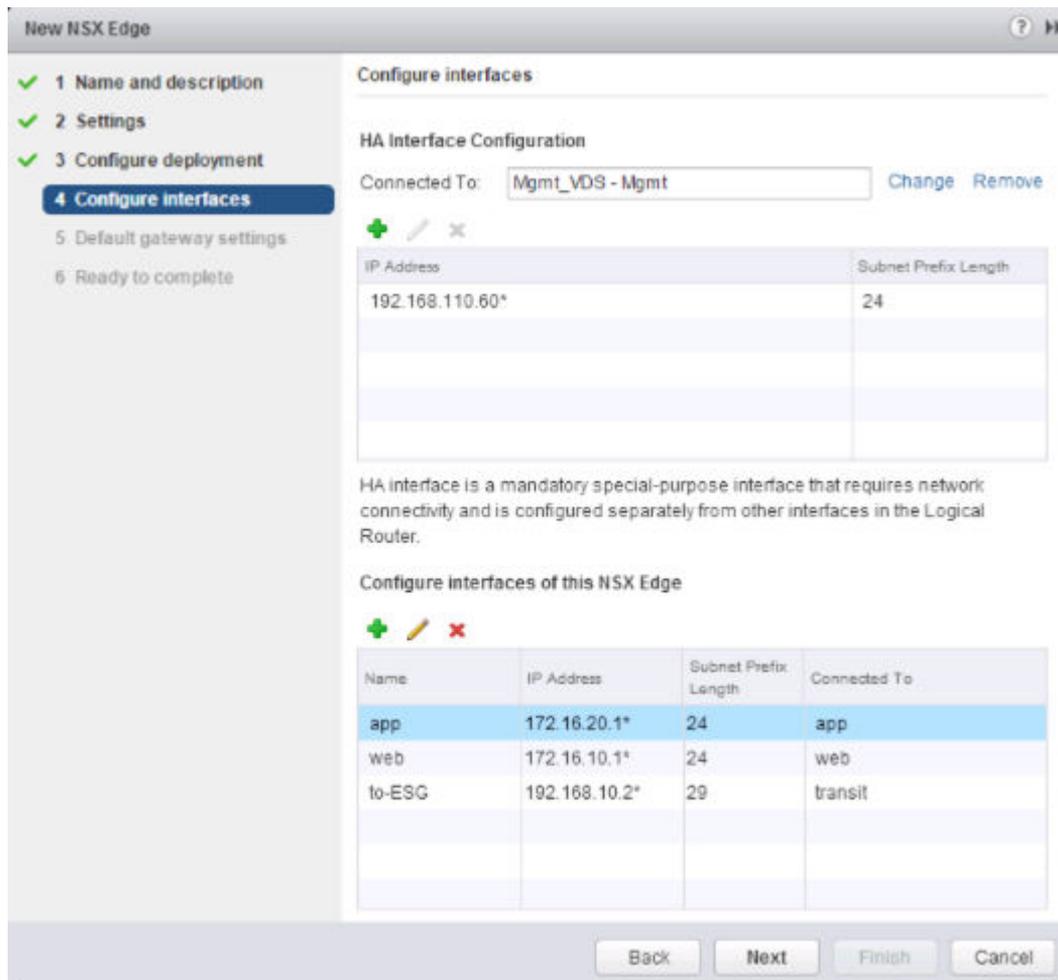
Note In prior releases of NSX, the HA interface was called the management interface. The HA interface is not supported for remote access to the logical router. You cannot SSH into the HA interface from anywhere that isn't on the same IP subnet as the HA interface. You cannot configure static routes that point out of the HA interface, which means that RPF will drop incoming traffic. You could, in theory, disable RPF, but this would be counterproductive to high availability. For SSH, use the logical router's protocol address, which is configured later when you configure dynamic routing.

In NSX 6.2, the HA interface of a logical router is automatically excluded from route redistribution.

In **Configure interfaces of this NSX Edge** the internal interfaces are for connections to switches that allow VM-to-VM (sometimes called East-West) communication. Internal interfaces are created as pseudo vNICs on the logical router virtual appliance. Uplink interfaces are for North-South communication. A logical router uplink interface might connect to an NSX edge services gateway, a third-party router VM for that, or a VLAN-backed dvPortgroup to make the logical router connect to a physical router directly. You must have at least one uplink interface for dynamic routing to work. Uplink interfaces are created as vNICs on the logical router virtual appliance.

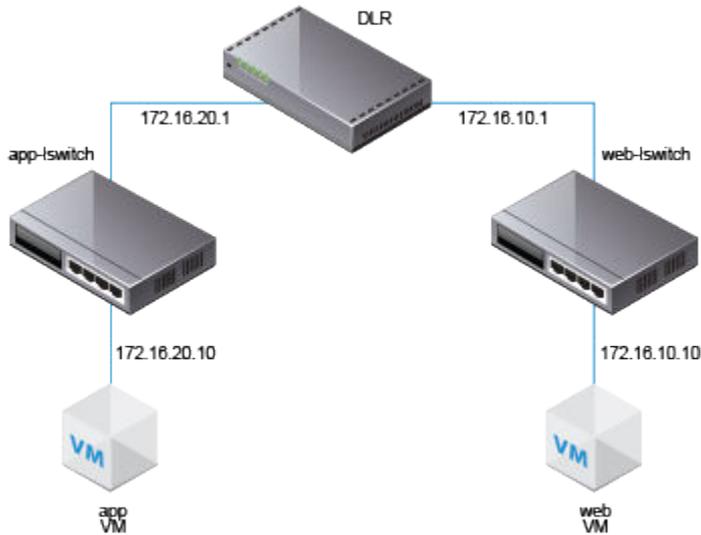
The interface configuration that you enter here is modifiable later. You can add, remove, and modify interfaces after a logical router is deployed.

The following example shows an HA interface connected to the management distributed port group. The example also shows two internal interfaces (app and web) and an uplink interface (to-ESG).



- 13 Make sure any VMs attached to the logical switches have their default gateways set properly to the logical router interface IP addresses.

In the following example topology, the default gateway of app VM should be 172.16.20.1. The default gateway of web VM should be 172.16.10.1. Make sure the VMs can ping their default gateways and each other.



Log in via SSH to the NSX Manager, and run the following commands:

- List all logical router instance information.

```
nsxmgr-l-01a> show logical-router list all
Edge-id          Vdr Name          Vdr id           #Lifs
-----          -
edge-1          default+edge-1    0x00001388      3
```

- List the hosts that have received routing information for the logical router from the controller cluster.

```
nsxmgr-l-01a> show logical-router list dlr edge-1 host
ID              HostName
-----
host-25        192.168.210.52
host-26        192.168.210.53
host-24        192.168.110.53
```

The output includes all hosts from all host clusters that are configured as members of the transport zone that owns the logical switch that is connected to the specified logical router (edge-1 in this example).

- List the routing table information that is communicated to the hosts by the logical router. Routing table entries should be consistent across all of the hosts.

```
nsx-mgr-l-01a> show logical-router host host-25 dlr edge-1 route

VDR default+edge-1 Route Table
Legend: [U: Up], [G: Gateway], [C: Connected], [I: Interface]
Legend: [H: Host], [F: Soft Flush] [!: Reject] [E: ECMP]

Destination      GenMask          Gateway          Flags  Ref Origin  UpTime  Interface
-----
0.0.0.0          0.0.0.0         192.168.10.1    UG     1  AUTO      4101    138800000002
```

172.16.10.0	255.255.255.0	0.0.0.0	UCI	1	MANUAL	10195	13880000000b
172.16.20.0	255.255.255.0	0.0.0.0	UCI	1	MANUAL	10196	13880000000a
192.168.10.0	255.255.255.248	0.0.0.0	UCI	1	MANUAL	10196	138800000002
192.168.100.0	255.255.255.0	192.168.10.1	UG	1	AUTO	3802	138800000002

- List additional information about the router from the point of view of one of the hosts. This is helpful to learn which controller is communicating with the host.

```
nsx-mgr-l-01a> show logical-router host host-25 dlr edge-1 verbose
```

```
VDR Instance Information :
```

```
-----
Vdr Name:                default+edge-1
Vdr Id:                  0x00001388
Number of Lifs:          3
Number of Routes:        5
State:                   Enabled
Controller IP:           192.168.110.203
Control Plane IP:        192.168.210.52
Control Plane Active:    Yes
Num unique nexthops:     1
Generation Number:       0
Edge Active:             No
```

Check the Controller IP field in the output of the `show logical-router host host-25 dlr edge-1 verbose` command.

SSH to a controller, and run the following commands to display the controller's learned VNI, VTEP, MAC, and ARP table state information.

```
192.168.110.202 # show control-cluster logical-switches vni 5000
VNI      Controller      BUM-Replication ARP-Proxy Connections
5000     192.168.110.201 Enabled          Enabled          0
```

The output for VNI 5000 shows zero connections and lists controller 192.168.110.201 as the owner for VNI 5000. Log in to that controller to gather further information for VNI 5000.

```
192.168.110.201 # show control-cluster logical-switches vni 5000
VNI      Controller      BUM-Replication ARP-Proxy Connections
5000     192.168.110.201 Enabled          Enabled          3
```

The output on 192.168.110.201 shows three connections. Check additional VNIs.

```
192.168.110.201 # show control-cluster logical-switches vni 5001
VNI      Controller      BUM-Replication ARP-Proxy Connections
5001     192.168.110.201 Enabled          Enabled      3
```

```
192.168.110.201 # show control-cluster logical-switches vni 5002
VNI      Controller      BUM-Replication ARP-Proxy Connections
5002     192.168.110.201 Enabled          Enabled      3
```

Because 192.168.110.201 owns all three VNI connections, we would expect to see zero connections on the other controller, 192.168.110.203.

```
192.168.110.203 # show control-cluster logical-switches vni 5000
VNI      Controller      BUM-Replication ARP-Proxy Connections
5000     192.168.110.201 Enabled          Enabled      0
```

- Before checking the MAC and ARP tables, start pinging from one VM to the other VM.

From app VM to web VM:

```
vmware@vmware-virtual-machine:~$ ping 172.16.10.10
PING 172.16.10.10 (172.16.10.10) 56(84) bytes of data.
64 bytes from 172.16.10.10: icmp_seq=1 ttl=63 time=2.15 ms
64 bytes from 172.16.10.10: icmp_seq=2 ttl=63 time=1.60 ms
```

Check the MAC tables.

```
192.168.110.201 # show control-cluster logical-switches mac-table 5000
VNI      MAC              VTEP-IP          Connection-ID
5000     00:50:56:a6:23:ae 192.168.250.52  7
```

```
192.168.110.201 # show control-cluster logical-switches mac-table 5001
VNI      MAC              VTEP-IP          Connection-ID
5001     00:50:56:a6:8d:72 192.168.250.51  23
```

Check the ARP tables.

```
192.168.110.201 # show control-cluster logical-switches arp-table 5000
VNI      IP              MAC              Connection-ID
5000     172.16.20.10   00:50:56:a6:23:ae 7
```

```
192.168.110.201 # show control-cluster logical-switches arp-table 5001
VNI      IP              MAC              Connection-ID
5001     172.16.10.10   00:50:56:a6:8d:72 23
```

Check the logical router information. Each logical router Instance is served by one of the controller nodes.

The instance sub-command of `show control-cluster logical-routers` command displays a list of logical routers that are connected to this controller.

The interface-summary sub-command displays the LIFs that the controller learned from the NSX Manager. This information is sent to the hosts that are in the host clusters managed under the transport zone.

The routes sub-command shows the routing table that is sent to this controller by the logical router's virtual appliance (also known as the control VM). Note that unlike on the ESXi hosts, this routing table does not include directly connected subnets because this information is provided by the LIF configuration. Route information on the ESXi hosts includes directly connected subnets because in that case it is a forwarding table used by ESXi host's datapath.

```

controller # show control-cluster logical-routers instance all
LR-Id      LR-Name      Universal Service-Controller Egress-Locale
0x1388     default+edge-1  false      192.168.110.201  local
    
```

Note the LR-Id and use it in the following command.

```

controller # show control-cluster logical-routers interface-summary 0x1388
Interface      Type  Id      IP[]
13880000000b   vxlan 0x1389  172.16.10.1/24
13880000000a   vxlan 0x1388  172.16.20.1/24
138800000002   vxlan 0x138a  192.168.10.2/29
    
```

```

controller # show control-cluster logical-routers routes 0x1388
Destination      Next-Hop[]      Preference Locale-Id      Source
192.168.100.0/24  192.168.10.1   110          00000000-0000-0000-0000-000000000000 CONTROL_VM
0.0.0.0/0         192.168.10.1   0            00000000-0000-0000-0000-000000000000 CONTROL_VM
    
```

```

[root@comp02a:~] esxcfg-route -l
VMkernel Routes:
Network      Netmask      Gateway      Interface
10.20.20.0   255.255.255.0  Local Subnet  vmk1
192.168.210.0 255.255.255.0  Local Subnet  vmk0
default      0.0.0.0       192.168.210.1  vmk0
    
```

- Display the controller connections to the specific VNI.

```
192.168.110.203 # show control-cluster logical-switches connection-table 5000
Host-IP      Port ID
192.168.110.53 26167 4
192.168.210.52 27645 5
192.168.210.53 40895 6
```

```
192.168.110.202 # show control-cluster logical-switches connection-table 5001
Host-IP      Port ID
192.168.110.53 26167 4
192.168.210.52 27645 5
192.168.210.53 40895 6
```

These Host-IP addresses are vmk0 interfaces, not VTEPs. Connections between ESXi hosts and controllers are created on the management network. The port numbers here are ephemeral TCP ports that are allocated by the ESXi host IP stack when the host establishes a connection with the controller.

- On the host, you can view the controller network connection matched to the port number.

```
[root@192.168.110.53:~] #esxcli network ip connection list | grep 26167
tcp      0      0 192.168.110.53:26167      192.168.110.101:1234  ESTABLISHED
96416    newreno  netcpa-worker
```

- Display active VNIs on the host. Observe how the output is different across hosts. Not all VNIs are active on all hosts. A VNI is active on a host if the host has a VM that is connected to the logical switch.

```
[root@192.168.210.52:~] # esxcli network vswitch dvs vmware vxlan network list --vds-name
Compute_VDS
VXLAN ID  Multicast IP          Control Plane          Controller Connection
Port Count  MAC Entry Count  ARP Entry Count  VTEP Count
-----
5000      N/A (headend replication)  Enabled (multicast proxy,ARP proxy)  192.168.110.203
(up)      1                  0                  0          0
5001      N/A (headend replication)  Enabled (multicast proxy,ARP proxy)  192.168.110.202
(up)      1                  0                  0          0
```

Note To enable the vxlan namespace in vSphere 6 and later, run the `/etc/init.d/hostd restart` command.

For logical switches in hybrid or unicast mode, the `esxcli network vswitch dvs vmware vxlan network list --vds-name <vds-name>` command should contain the following output:

- Control Plane is enabled.
- Multicast proxy and ARP proxy are listed. AARP proxy is listed even if you disabled IP discovery.
- A valid controller IP address is listed and the connection is up.

- If a logical router is connected to the ESXi host, the port Count is at least 1, even if there are no VMs on the host connected to the logical switch. This one port is the vdrPort, which is a special dvPort connected to the logical router kernel module on the ESXi host.
- First ping from VM to another VM on a different subnet and then display the MAC table. Note the Inner MAC is the VM entry while the Outer MAC and Outer IP refer to the VTEP.

```
~ # esxcli network vswitch dvs vmware vxlan network mac list --vds-name=Compute_VDS --vxlan-id=5000
Inner MAC          Outer MAC          Outer IP          Flags
-----
00:50:56:a6:23:ae  00:50:56:6a:65:c2  192.168.250.52   00000111
```

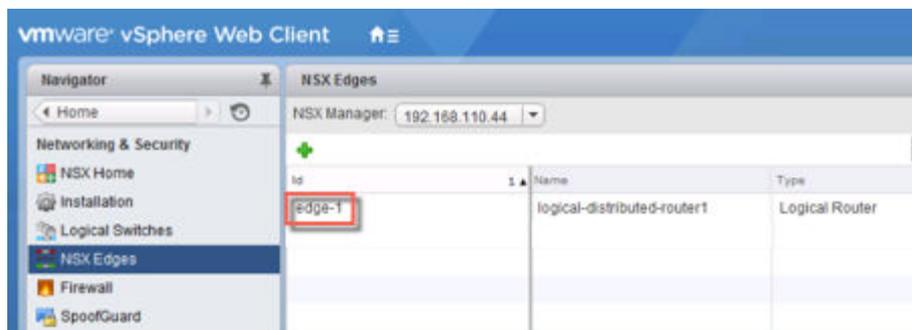
```
~ # esxcli network vswitch dvs vmware vxlan network mac list --vds-name=Compute_VDS --vxlan-id=5001
Inner MAC          Outer MAC          Outer IP          Flags
-----
02:50:56:56:44:52  00:50:56:6a:65:c2  192.168.250.52   00000101
00:50:56:f0:d7:e4  00:50:56:6a:65:c2  192.168.250.52   00000111
```

What to do next

On the hosts where NSX edge appliances are first deployed, NSX enables automatic VM startup/shutdown. If the appliance VMs are later migrated to other hosts, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, VMware recommends that you check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html.

After the logical router is deployed, double-click the logical router ID to configure additional settings, such as interfaces, routing, firewall, bridging, and DHCP relay.

For example:



Configuring Secondary NSX Managers

7

After a primary cross-vCenter NSX Manager has been configured, you can configure your secondary NSX Managers. Secondary NSX Managers use the same universal control cluster that was deployed by the primary NSX Manager. There can be up to seven secondary NSX Managers in a cross-vCenter NSX environment. Once an NSX Manager is assigned the secondary role it can use universal objects such as universal logical switches.

Complete the configuration tasks for each secondary NSX Manager in the cross-vCenter NSX environment.

This chapter includes the following topics:

- [Add a Secondary NSX Manager](#)
- [Prepare Hosts on a Secondary NSX Manager](#)
- [Configure VXLAN from the Secondary NSX Manager](#)
- [Assign a Segment ID Pool for a Secondary NSX Manager](#)
- [Add Clusters to the Universal Transport Zone](#)

Add a Secondary NSX Manager

You can add up to seven secondary NSX Managers in a cross-vCenter NSX environment. Universal objects configured on the primary NSX Manager are synchronized to the secondary NSX Managers.

NSX Managers can have one of four roles:

- Primary
- Secondary
- Standalone
- Transit

To view the role of an NSX Manager, log in to the vCenter linked to the NSX Manager, and navigate to **Home > Networking & Security > Installation** and select the **Management** tab. The role is displayed in the Role column in the NSX Managers section. If there is no Role column shown, the NSX Manager has the standalone role.

Prerequisites

- There should be at least two NSX Managers, one with the primary role and one with the standalone or transit role.
- The version of the NSX Managers (the primary NSX Manager and NSX Managers that will be assigned the secondary role) must match.
- The node IDs of the primary NSX Manager and the NSX Managers that will be assigned the secondary role must be present and different. NSX Manager instances deployed from OVA files have unique node IDs. An NSX Manager deployed from a template (as in when you convert a virtual machine to a template) will have the same node ID as the original NSX Manager used to create the template, and these two NSX Managers cannot be used in the same cross-vCenter NSX installation.
- Each NSX Manager must be registered with a separate and unique vCenter Server.
- The UDP ports used for VXLAN must be the same for all NSX Managers.
- When assigning the secondary role to a NSX Manager, the vCenter linked to it must not have any deployed NSX Controllers.
- The segment ID pool of the NSX Manager being assigned the secondary role must not overlap with the segment ID pools of the primary NSX Manager or the segment ID pool of any other secondary NSX Manager .
- The NSX Manager being assigned the secondary role must have the standalone or transit role.

Note You can view the NSX Manager node ID with the following REST API call:

```
GET https://NSX-Manager-IP-Address/api/2.0/services/vsmconfig
```

You can view the UDP ports used for VXLAN with the following REST API call:

```
GET https://NSX-Manager-IP-Address/api/2.0/vdn/config/vxlan/udp/port
```

Procedure

- 1 Log in to the vCenter linked to the primary NSX Manager.
- 2 Navigate to **Home > Networking & Security > Installation** and select the **Management** tab.
- 3 Click the primary NSX Manager. Then select **Actions > Add Secondary NSX Manager**.
- 4 Enter the IP address, user name, and password of the secondary NSX Manager.
- 5 Click **OK**.
- 6 Check that the certificate thumbprint matches the certificate of the vCenter Server.

If you installed a CA-signed certificate on the CA server, you are presented with the thumbprint of the CA-signed certificate. Otherwise, you are presented with a self-signed certificate.

- After successful registration the role changes from Standalone to Secondary.

If your vCenter Server systems are in Enhanced Linked Mode, you can see the roles of all NSX Managers associated with those vCenter Server systems from the **Home > Networking & Security > Installation** tab.

NSX Manager	Role	IP Address	vCenter	Version
192.168.110.42	Primary	192.168.110.42	vc-l-01a.corp.local	6.2.0.2654989
192.168.110.43	Secondary	192.168.110.43	vc-l-01b.corp.local	6.2.0.2654989

If your environment does not employ Enhanced Linked Mode, log in to the vCenter linked to the secondary NSX Manager to view the NSX Manager's role.

Note Initially, the controller status will show disconnected. Wait a few seconds and then refresh the vSphere Web Client and the status will change to Normal.

- Log out of the vSphere Web Client and back in to ensure that the vSphere Web Client displays the new NSX Manager roles.

Prepare Hosts on a Secondary NSX Manager

During host preparation the secondary NSX Manager installs NSX kernel modules on ESXi hosts that are members of vCenter clusters, and builds the NSX control-plane and management-plane fabric. NSX kernel modules packaged in VIB files run within the hypervisor kernel and provide services such as distributed routing, distributed firewall, and VXLAN bridging capabilities.

Prerequisites

For details on prerequisites for host preparation, see [Prepare Hosts on the Primary NSX Manager](#)

Procedure

- Log in to the vCenter linked to the secondary NSX Manager and navigate to **Home > Networking & Security > Installation** and select the **Host Preparation** tab..

If your vCenters are in Enhanced Linked Mode, you can configure the secondary NSX Manager from any linked vCenter. Navigate to the **Host Preparation** tab, and select the secondary NSX Manager from the drop-down menu.

- For each cluster, click **Install**.

Note While the installation is in progress, do not deploy, upgrade, or uninstall any service or component.

- 3 Monitor the installation until the **Installation Status** column displays a green check mark.

If the **Installation Status** column displays a red warning icon and says **Not Ready**, click **Resolve**. Clicking **Resolve** might result in a reboot of the host. If the installation is still not successful, click the warning icon. All errors are displayed. Take the required action and click **Resolve** again.

Configure VXLAN from the Secondary NSX Manager

The VXLAN network is used for Layer 2 logical switching across hosts, potentially spanning multiple underlying Layer 3 domains. You configure VXLAN on a per-cluster basis, where you map each cluster that is to participate in NSX to a vSphere distributed switch (VDS). When you map a cluster to a distributed switch, each host in that cluster is enabled for logical switches. The settings chosen here will be used in creating the VMkernel interface.

Prerequisites

For details on prerequisites see [Configure VXLAN from the Primary NSX Manager](#)

Procedure

- 1 Log in to the vCenter linked to the secondary NSX Manager and navigate to **Home > Networking & Security > Installation** and select the **Host Preparation** tab.

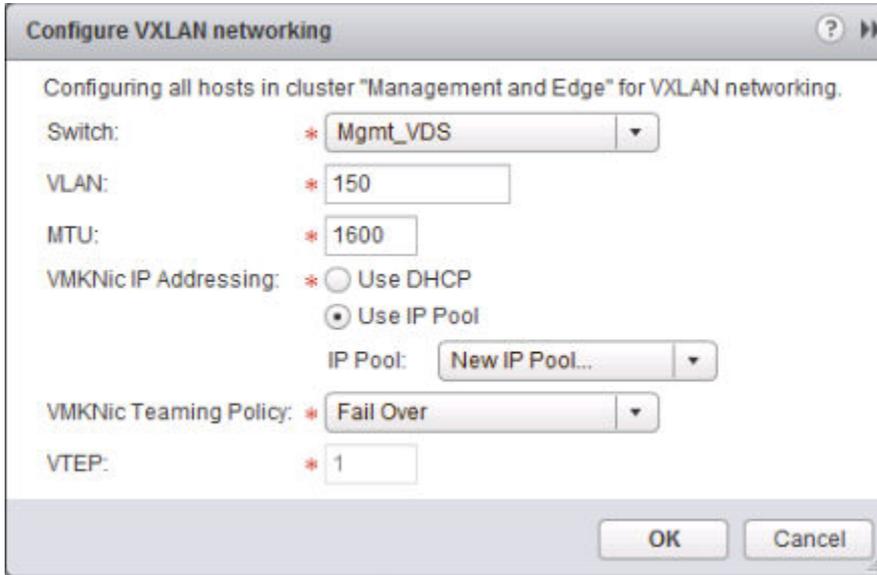
If your vCenters are in Enhanced Linked Mode, you can configure the secondary NSX Manager from any linked vCenter. Navigate to the **Host Preparation** tab, and select the secondary NSX Manager from the drop-down menu.

- 2 Click **Not Configured** in the **VXLAN** column.
- 3 Set up logical networking.

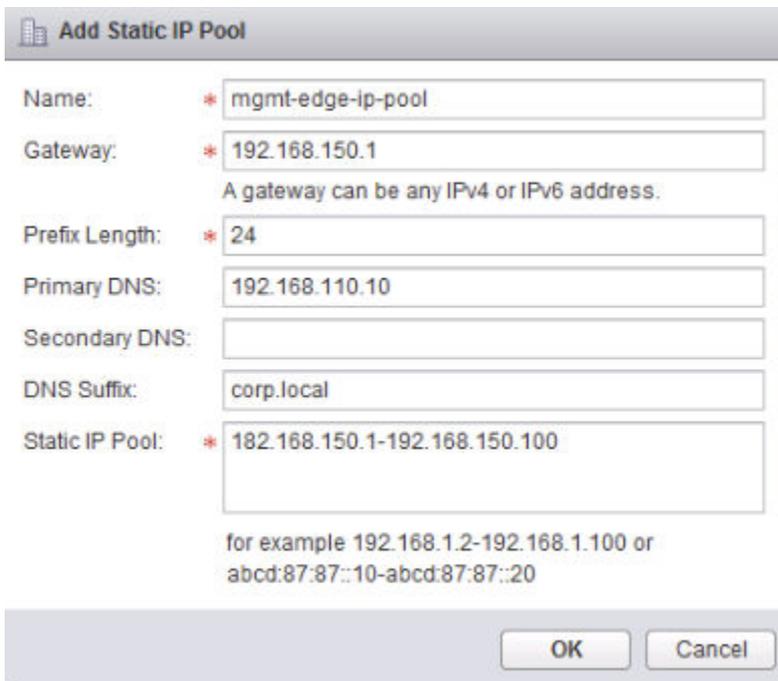
This involves selecting a VDS, a VLAN ID, an MTU size, an IP addressing mechanism, and a NIC teaming policy.

The MTU for each switch must be set to 1550 or higher. By default, it is set to 1600. If the vSphere distributed switch (VDS) MTU size is larger than the VXLAN MTU, the VDS MTU will not be adjusted down. If it is set to a lower value, it will be adjusted to match the VXLAN MTU. For example, if the VDS MTU is set to 2000 and you accept the default VXLAN MTU of 1600, no changes to the VDS MTU will be made. If the VDS MTU is 1500 and the VXLAN MTU is 1600, the VDS MTU will be changed to 1600.

These example screens show a configuration for a management cluster with IP pool address range of 182.168.150.1-192.168.150.100, backed by VLAN 150, and with a failover NIC teaming policy.



The number of VTEPs is not editable in the UI. The VTEP number is set to match the number of dvUplinks on the vSphere distributed switch being prepared.



For compute clusters, you may want to use different IP address settings (for example, 192.168.250.0/24 with VLAN 250). This would depend on how the physical network is designed, and likely won't be the case in small deployments.

Assign a Segment ID Pool for a Secondary NSX Manager

The secondary NSX Manager will display the universal segment ID pool, which is synced from the primary NSX Manager. Additionally, you can create a segment ID pool that is local to the secondary NSX Manager, which will be used to create logical switches that are local to that NSX Manager.

Prerequisites

For details on prerequisites and guidance for planning segment ID pools see [Assign a Segment ID Pool to the Primary NSX Manager](#).

Procedure

- 1 Log in to the vCenter linked to the secondary NSX Manager and navigate to **Home > Networking & Security > Installation > Logical Network Preparation** and select the **Segment ID** tab.

If your vCenters are in Enhanced Linked Mode, you can configure the secondary NSX Manager from any linked vCenter. Navigate to the **Segment ID** tab, and select the secondary NSX Manager from the drop-down menu.

- 2 Enter a range for local segment IDs, such as **20000–29999**.

Caution The specified ranges for local and universal segment IDs must be unique and not overlap.

- 3 If any of your transport zones will use multicast or hybrid replication mode, check **Enable multicast addressing** and type a multicast address or a range of multicast addresses.

Caution In a cross-vCenter NSX environment, if you select hybrid mode, you must ensure that the multicast address specified does not conflict with any other multicast addresses assigned on any NSX Manager in the environment.

Having a range of multicast addresses spreads traffic across your network, prevents the overloading of a single multicast address, and better contains BUM replication. The recommended multicast address range starts at 239.0.1.0/24 and excludes 239.128.0.0/24.

Do not use 239.0.0.0/24 or 239.128.0.0/24 as the multicast address range, because these networks are used for local subnet control, meaning that the physical switches flood all traffic that uses these addresses. For more information about unusable multicast addresses, see <https://tools.ietf.org/html/draft-ietf-mboned-ipv4-mcast-unusable-01>.

The secondary NSX Manager now has both imported universal segment IDs provided by the primary NSX Manager and local segment IDs.

Add Clusters to the Universal Transport Zone

Add the clusters on the vCenter that is linked to the secondary NSX Manager to the universal transport zone to allow those clusters to use universal logical switches.

Procedure

- 1 Log in to the vCenter linked to the secondary NSX Manager and navigate to **Home > Networking & Security > Installation > Logical Network Preparation** and select the **Transport Zones** tab.

If your vCenters are in Enhanced Linked Mode, you can configure the secondary NSX Manager from any linked vCenter. Navigate to the **Transport Zones** tab, and select the secondary NSX Manager from the drop-down menu.

- 2 Select the universal transport zone and click **Actions** () > **Connect Clusters**. Select the clusters to add to the universal transport zone and click OK.

After Configuring Primary and Secondary NSX Managers

8

You now have a configured primary NSX Manager and at least one secondary NSX Manager. In addition to creating universal objects from the primary NSX Manager, you can also create objects that are local to that specific vCenter NSX environment, such as logical switches, and logical (distributed) routers. These can be created on primary or secondary NSX Managers. They will exist only within the vCenter NSX environment in which they were created. They will not be visible on the other NSX Managers in the Cross-vCenter NSX environment. Additionally, you can add hosts to or remove hosts from clusters.

See the *NSX Administration Guide* for details on additional administrative tasks you may want to complete.

This chapter includes the following topics:

- [Add a Transport Zone](#)
- [Add a Logical Switch](#)
- [Add a Logical \(Distributed\) Router](#)
- [Add an Edge Services Gateway](#)
- [Add a Host to a Prepared Cluster](#)
- [Remove a Host from an NSX Prepared Cluster](#)

Add a Transport Zone

A transport zone controls which hosts a logical switch can reach and can span one or more vSphere clusters. Transport zones dictate which clusters VMs can participate in the use of a particular network. Transport Zones can be created on both primary and secondary NSX Managers in a cross-vCenter NSX environment. Both primary and secondary NSX Managers can have local transport zones, but only primary NSX Managers can have a universal transport zone.

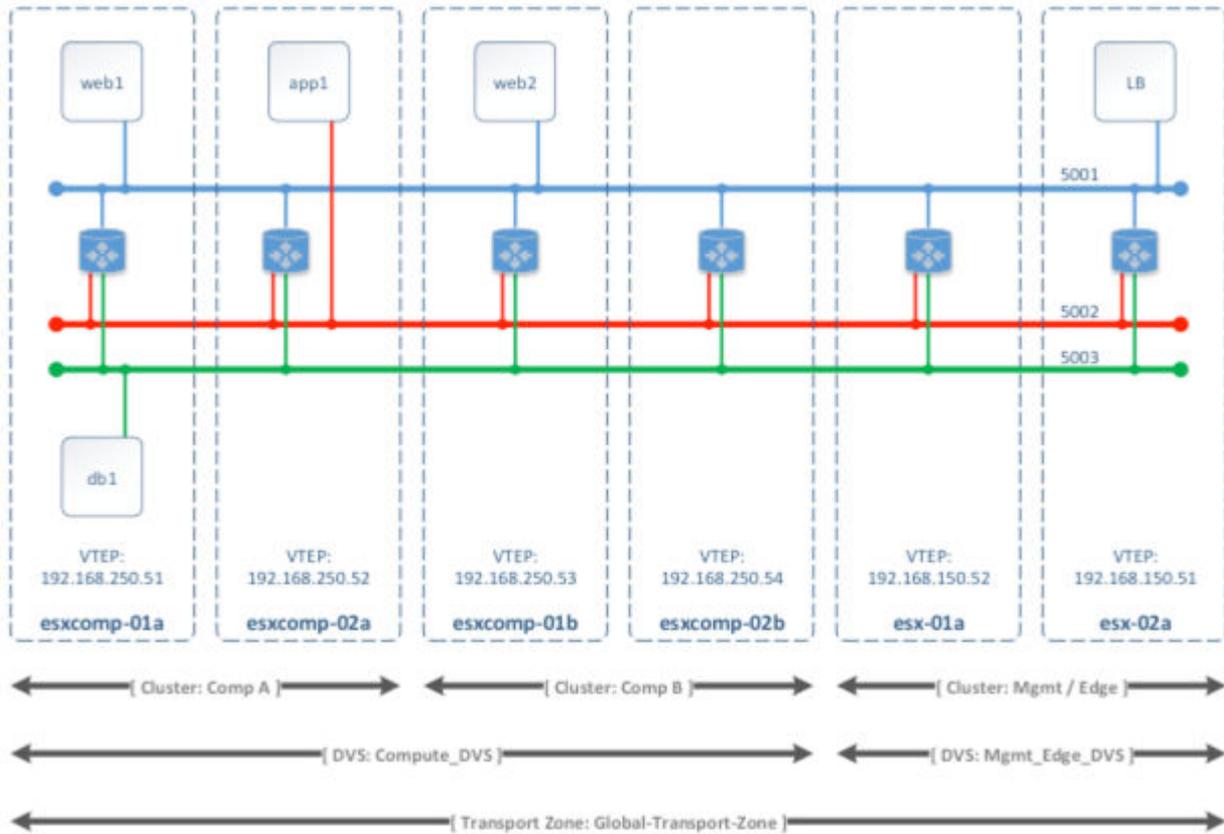
An NSX environment can contain one or more transport zones based on your requirements. A host cluster can belong to multiple transport zones. A logical switch can belong to only one transport zone.

NSX does not allow connection of VMs that are in different transport zones. The span of a logical switch is limited to a transport zone, so virtual machines in different transport zones cannot be on the same Layer 2 network. A distributed logical router cannot connect to logical switches that are in different transport zones. After you connect the first logical switch, the selection of further logical switches is limited to those that are in the same transport zone. Similarly, an edge services gateway (ESG) has access to logical switches from only one transport zone.

The following guidelines are meant to help you design your transport zones:

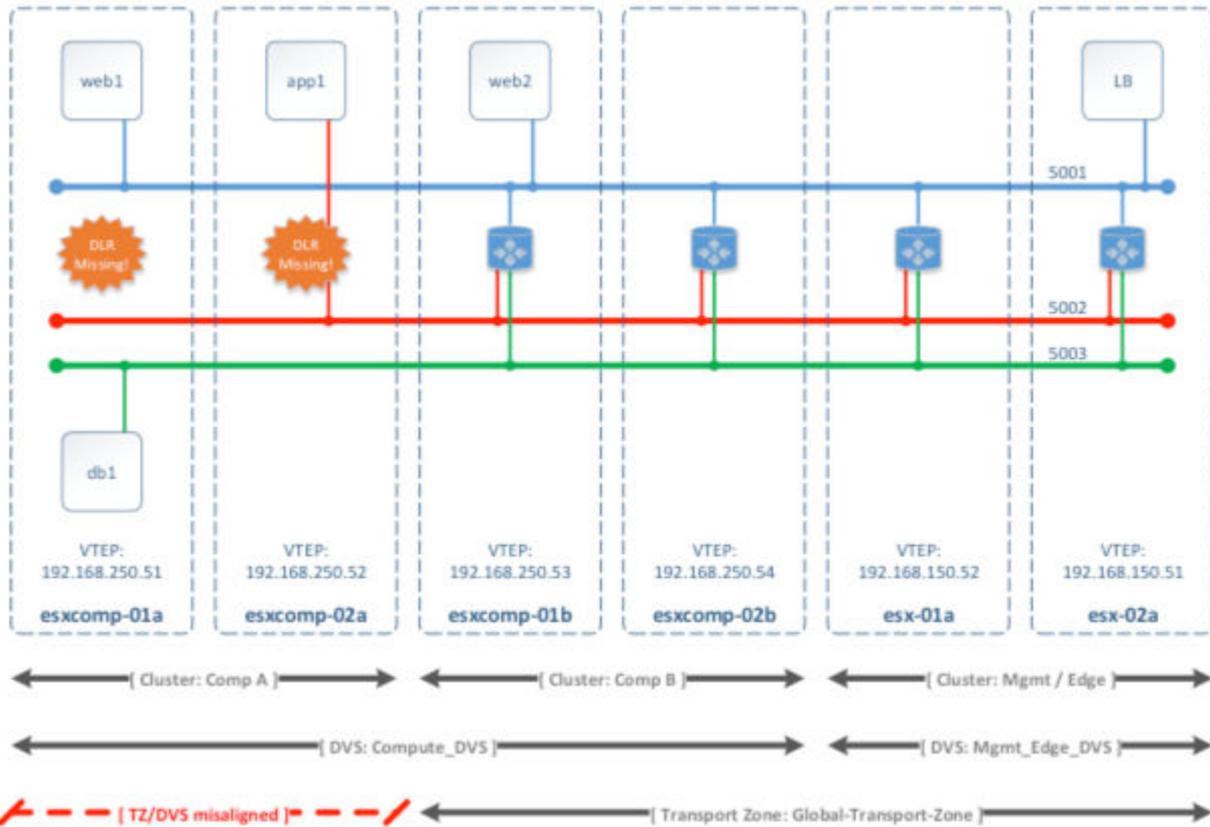
- If a cluster requires Layer 3 connectivity, the cluster must be in a transport zone that also contains an edge cluster, meaning a cluster that has Layer 3 edge devices (distributed logical routers and edge services gateways).
- Suppose you have two clusters, one for web services and another for application services. To have VXLAN connectivity between the VMs in these two clusters, both of the clusters must be included in the transport zone.
- Keep in mind that all logical switches included in the transport zone will be available and visible to all VMs within the clusters that are included in the transport zone. If a cluster includes secured environments, you might not want to make it available to VMs in other clusters. Instead, you can place your secure cluster in a more isolated transport zone.
- The span of the vSphere distributed switch (VDS or DVS) should match the transport zone span. When creating transport zones in multi-cluster VDS configurations, make sure all clusters in the selected VDS are included in the transport zone. This is to ensure that the DLR is available on all clusters where VDS dvPortgroups are available.

The following diagram shows a transport zone correctly aligned to the VDS boundary.



If you do not follow this best practice, keep in mind that if a VDS spans more than one host cluster and the transport zone includes only one (or a subset) of these clusters, any logical switch included within this transport zone can access VMs within all clusters spanned by the VDS. In other words, the transport zone will not be able to constrain the logical switch span to a subset of the clusters. If this logical switch is later connected to a DLR, you must ensure that the router instances are created only in the cluster included in the transport zone to avoid any Layer 3 issues.

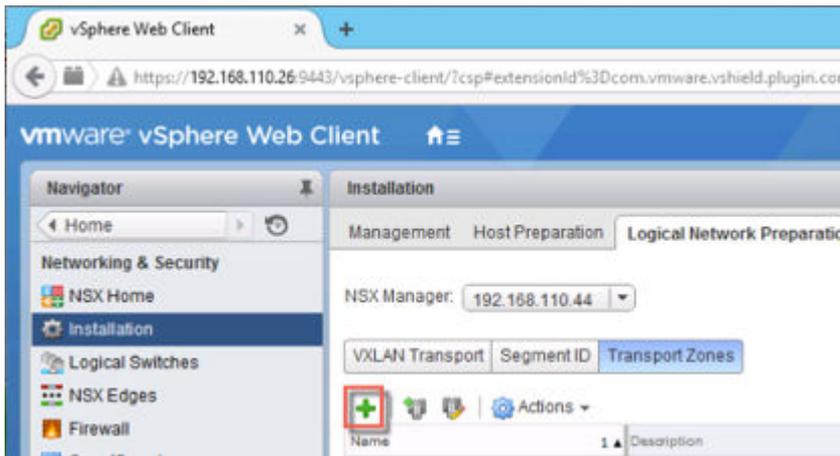
For example, when a transport zone is not aligned to the VDS boundary, the scope of the logical switches (5001, 5002 and 5003) and the DLR instances that these logical switches are connected to becomes disjointed, causing VMs in cluster Comp A to have no access to the DLR logical interfaces (LIFs).



Procedure

- 1 In vCenter, navigate to **Home > Networking & Security > Installation** and select the **Logical Network Preparation** tab.
- 2 Click **Transport Zones** and click the **New Transport Zone (+)** icon.

For example:



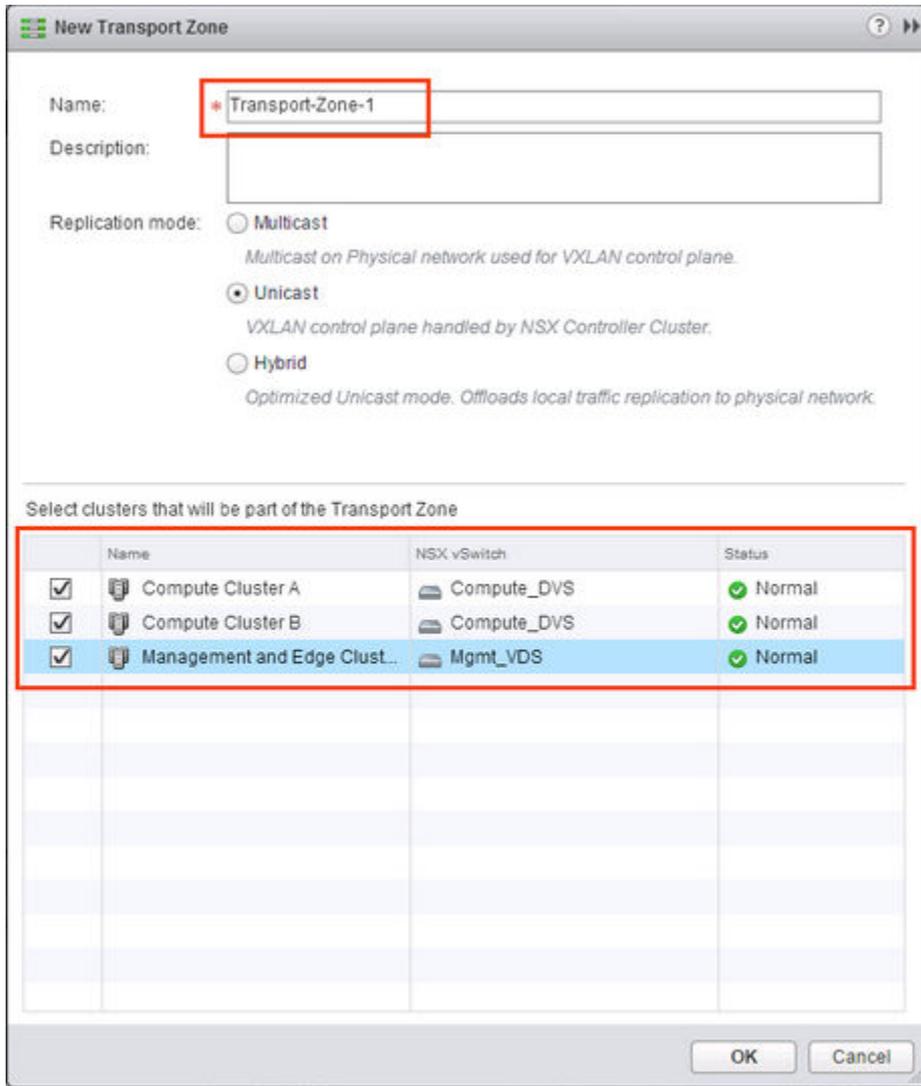
- 3 In the New Transport Zone dialog box, type a name and an optional description for the transport zone.

4 Depending on whether you have a controller node in your environment, or you want to use multicast addresses, select the control plane mode.

- **Multicast:** Multicast IP addresses in the physical network are used for the control plane. This mode is recommended only when you are upgrading from older VXLAN deployments. Requires PIM/IGMP in the physical network.
- **Unicast:** The control plane is handled by an NSX controller. All unicast traffic leverages optimized headend replication. No multicast IP addresses or special network configuration is required.
- **Hybrid:** Offloads local traffic replication to the physical network (L2 multicast). This requires IGMP snooping on the first-hop switch and access to an IGMP querier in each VTEP subnet, but does not require PIM. The first-hop switch handles traffic replication for the subnet.

5 Select the clusters to be added to the transport zone.

For example:



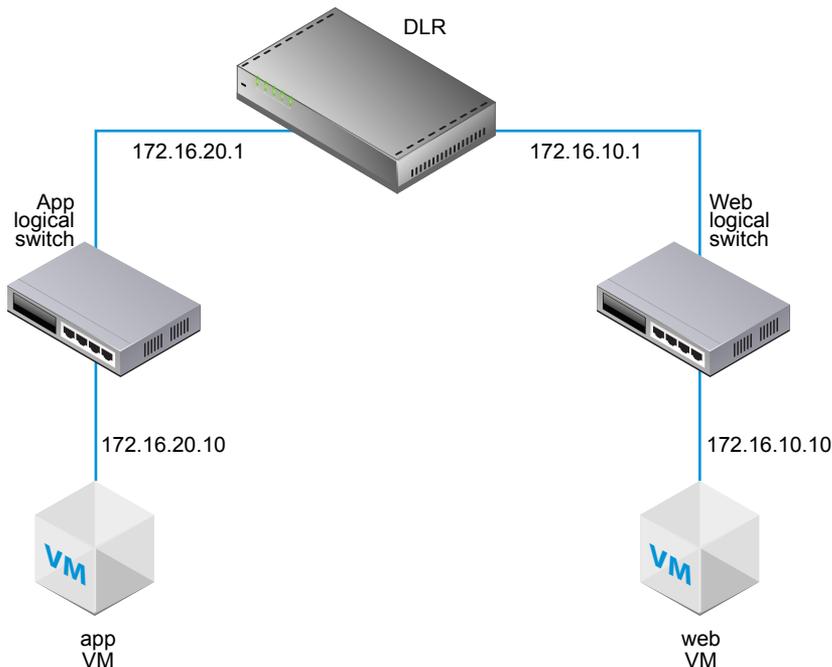
What to do next

Now that you have a transport zone, you can add logical switches.

Add a Logical Switch

Both primary and secondary NSX Managers can have logical switches that are local to the environment in which they are created. Universal switches are created in universal transport zones, and logical switches are created in transport zones that are local to that NSX Manager.

When you are adding logical switches, it is important to have in mind a particular topology that you are building. For example, the following simple topology shows two logical switches connected to a single distributed logical router (DLR). In this diagram, each logical switch is connected to a single VM. The two VMs can be on different hosts or the same host, in different host clusters or in the same host cluster. If a DLR does not separate the VMs, the underlying IP addresses configured on the VMs can be in the same subnet. If a DLR does separate them, the IP addresses on the VMs must be in different subnets (as shown in the example).



Prerequisites

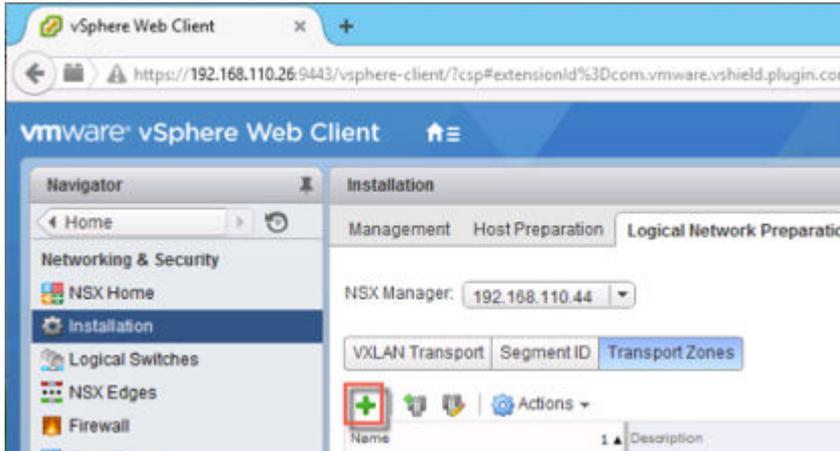
- vSphere distributed switches must be configured.
- NSX Manager must be installed.
- Controllers must be deployed.
- Host clusters must be prepared for NSX.
- VXLAN must be configured.
- A segment ID pool must be configured.

- A transport zone must be created.

Procedure

- 1 In the vSphere Web Client, navigate to **Home > Networking & Security > Logical Switches**.
- 2 Click the **New Logical Switch (+)** icon.

For example:



- 3 Type a name and optional description for the logical switch.
- 4 Select the transport zone in which you want to create the logical switch.

By default, the logical switch inherits the control plane replication mode from the transport zone. You can change it to one of the other available modes. The available modes are unicast, hybrid, and multicast.

The case in which you might want to override the inherited transport zone's control plane replication mode for an individual logical switch is when the logical switch you are creating has significantly different characteristics in terms of the amount of BUM traffic it will carry. In this case, you might create a transport zone that uses as unicast mode, and use hybrid or multicast mode for the individual logical switch.

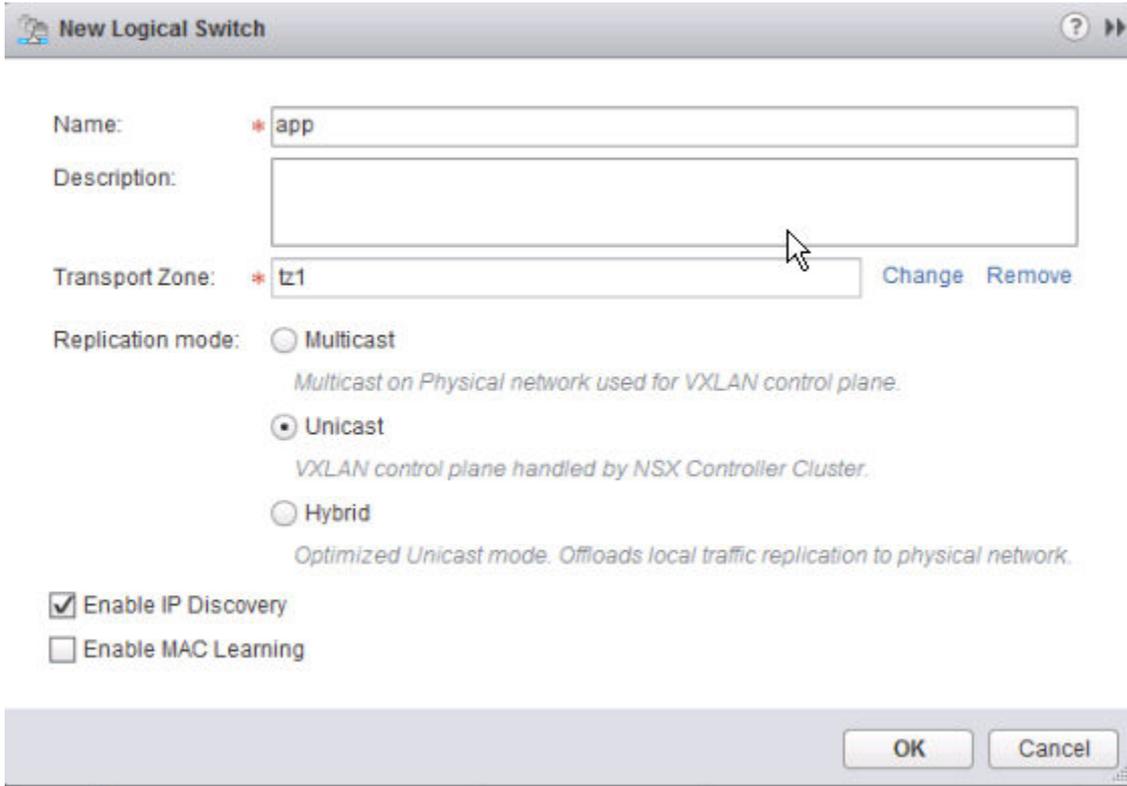
- 5 (Optional) Click **Enable IP Discovery** to enable ARP suppression.

This setting minimizes ARP traffic flooding within individual VXLAN segments---in other words, between VMs connected to the same logical switch. IP discovery is enabled by default.

- (Optional) Click **Enable MAC learning** if your VMs have multiple MAC addresses or are using virtual NICs that are trunking VLANs.

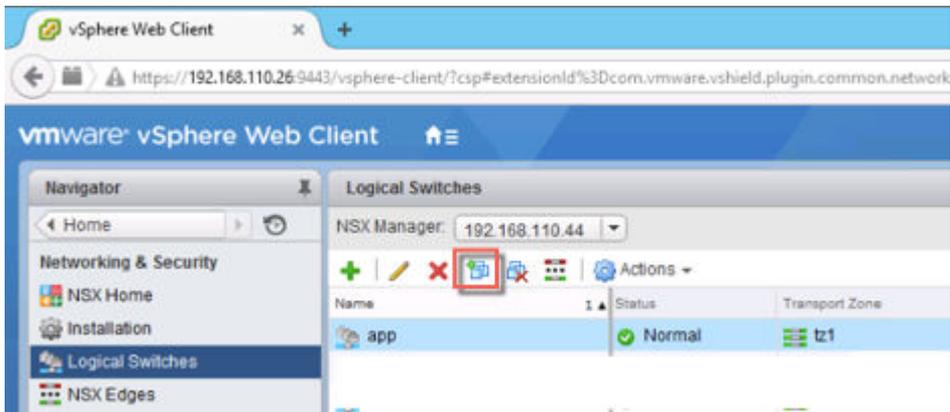
Enabling MAC learning builds a VLAN/MAC pair learning table on each vNIC. This table is stored as part of the dvfilter data. During vMotion, dvfilter saves and restores the table at the new location. The switch then issues RARPs for all the VLAN/MAC entries in the table.

This example shows the app logical switch with default settings.



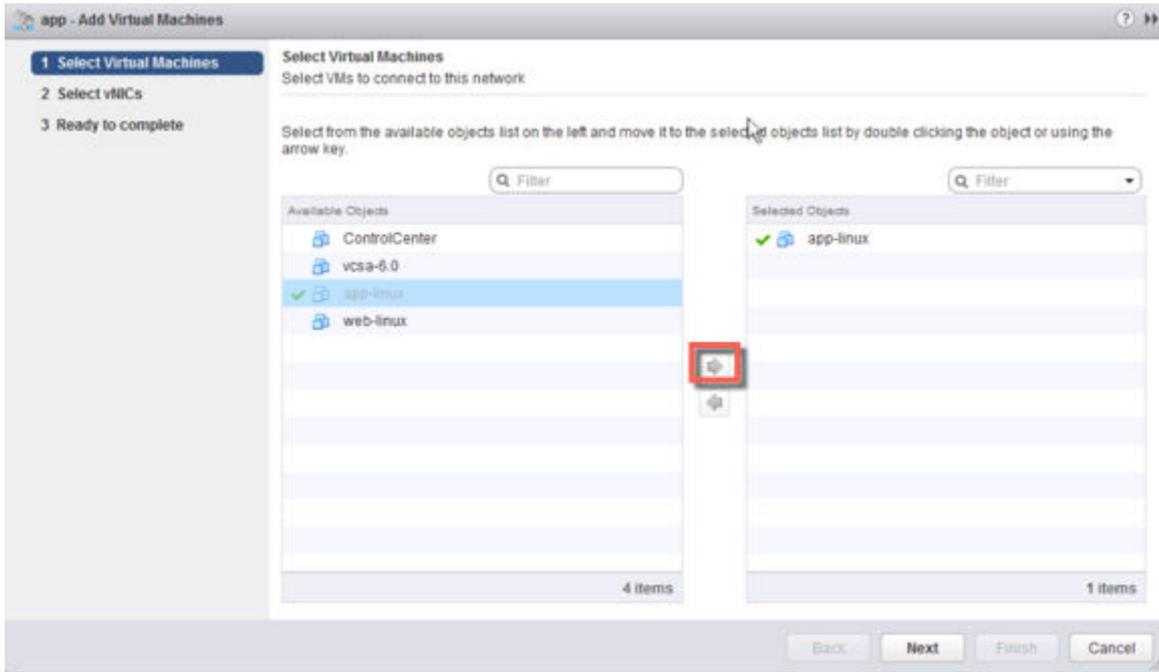
- Attach a VM to the logical switch by selecting the switch and clicking the **Add Virtual Machine** () icon.

For example:



- 8 Select the VM and click the right-arrow button.

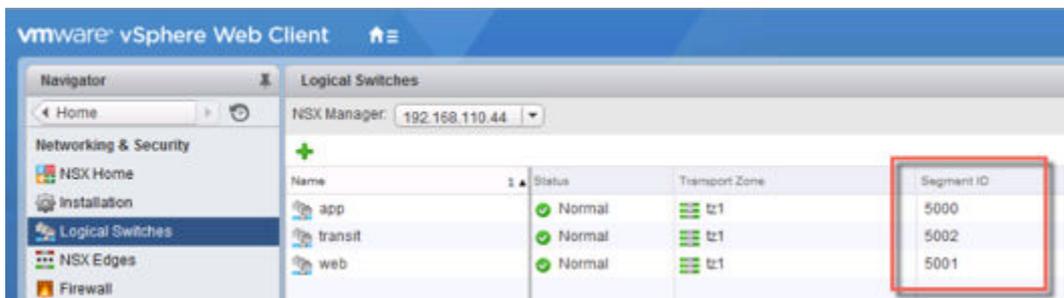
For example:



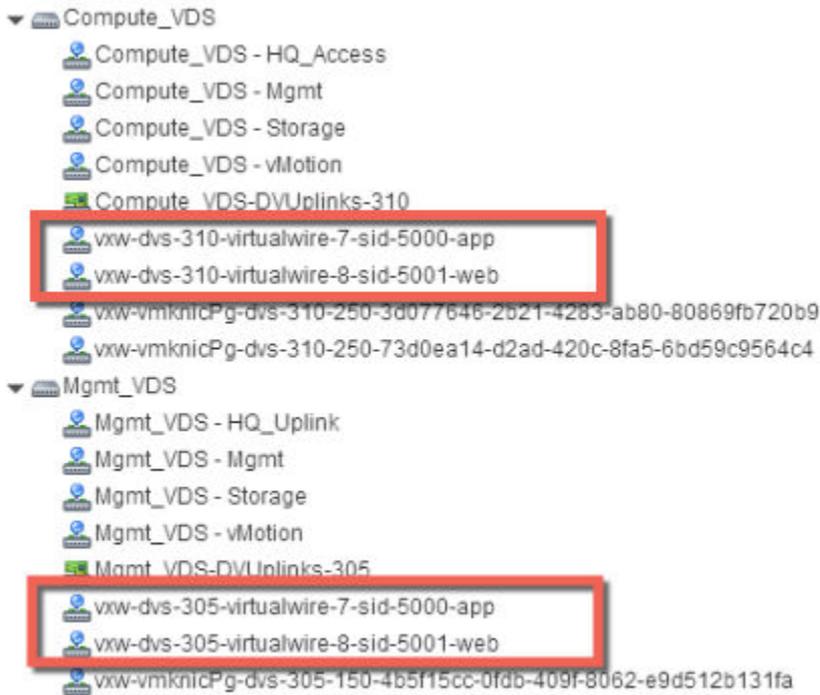
- 9 Select a vNIC.

Each logical switch that you create receives an ID from the segment ID pool, and a virtual wire is created. A virtual wire is a dvPortgroup that is created on each vSphere distributed switch. The virtual wire descriptor contains the name of the logical switch and the logical switch's segment ID. Assigned segment IDs appear in multiple places, as shown in the following examples.

In **Home > Networking & Security > Logical Switches**:

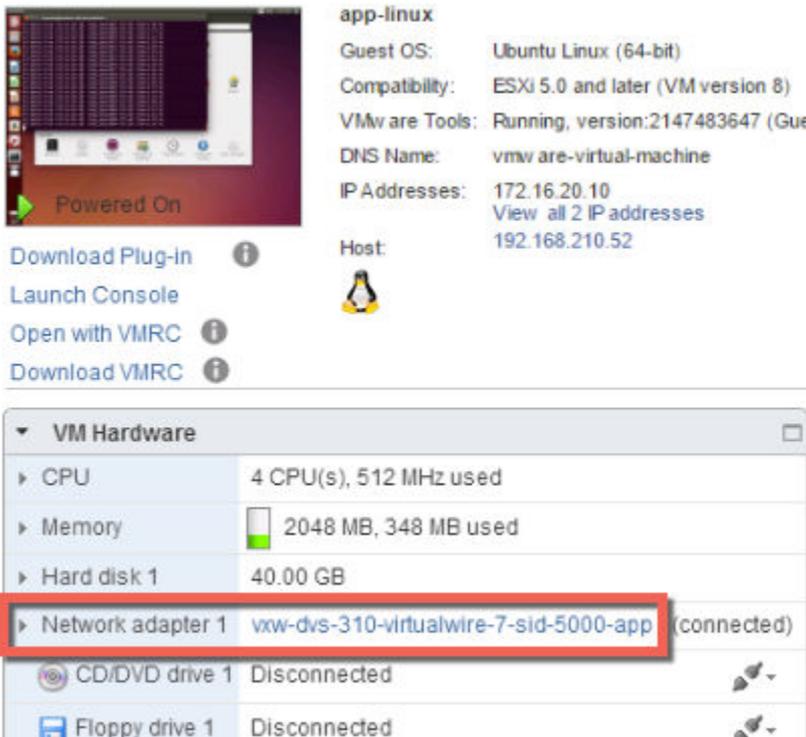


In **Home > Networking**:



Notice that the virtual wires are created on both of the vSphere distributed switches, Compute_VDS and Mgmt_VDS. This is because both of these vSphere distributed switches are members of the transport zone that is associated with the web and app logical switches.

In Home > Hosts and Clusters > VM > Summary:



On the hosts that are running the VMs that are attached to the logical switch, log in and execute the following commands to view local VXLAN configuration and state information.

- Displays host-specific VXLAN details.

```

~ # esxcli network vswitch dvs vmware vxlan list
VDS ID                               VDS Name      MTU  Segment ID  Gateway IP
Gateway MAC      Network Count  Vmknics Count
-----
88 eb 0e 50 96 af 1d f1-36 fe c1 ef a1 51 51 49 Compute_VDS  1600  192.168.250.0  192.168.250.1
ff:ff:ff:ff:ff:ff           0           1

```

Note If the `esxcli network vswitch dvs vmware vxlan` command produces the "Unknown command or namespace" error message, run the `/etc/init.d/hostd restart` command on the host and then try again.

VDS Name displays the vSphere distributed switch to which the host is attached.

The Segment ID is the IP network used by VXLAN.

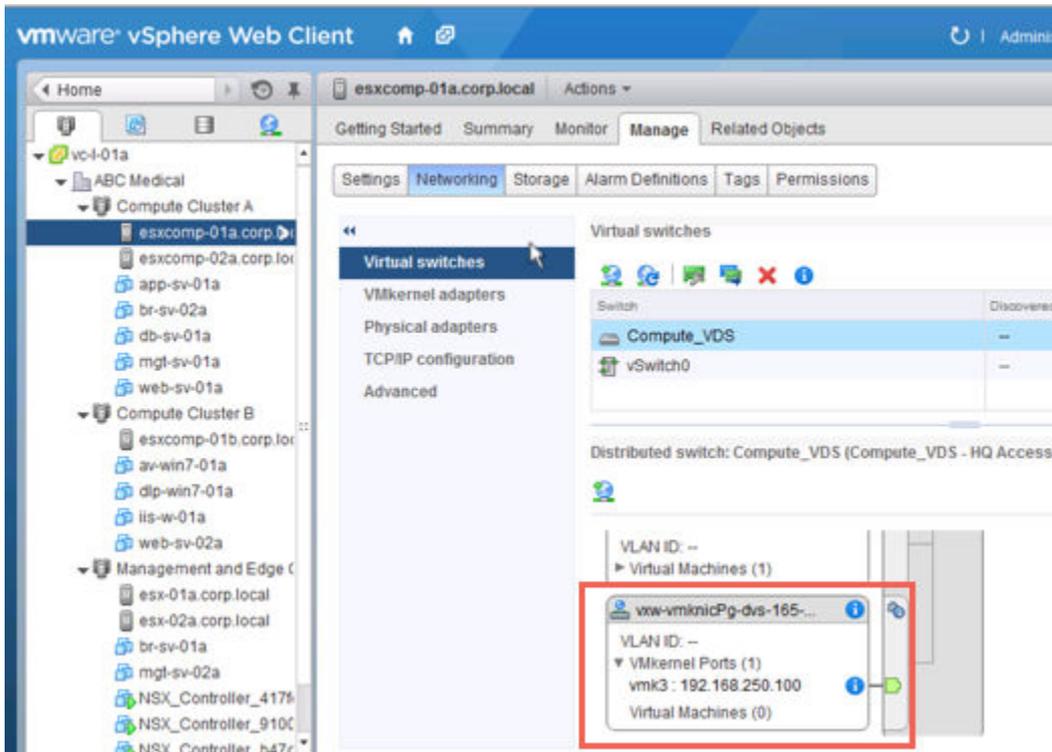
The Gateway IP is the gateway IP address used by VXLAN.

The Gateway MAC address remains `ff:ff:ff:ff:ff:ff`.

The Network Count remains 0 unless a DLR is attached to the logical switch.

The Vmknics count should match the number of VMs attached to the logical switch.

- Test IP VTEP interface connectivity, and verify the MTU has been increased to support VXLAN encapsulation. Ping the vmknic interface IP address, which can be found on the host's **Manage > Networking > Virtual switches** page in the vCenter Web Client.



The -d flag sets the don't-fragment (DF) bit on IPv4 packets. The -s flag sets the packet size.

```
root@esxcomp-02a ~ # vmkping ++netstack=vxlan -d -s 1570 192.168.250.100
PING 192.168.250.100 (192.168.250.100): 1570 data bytes
1578 bytes from 192.168.250.100: icmp_seq=0 ttl=64 time=1.294 ms
1578 bytes from 192.168.250.100: icmp_seq=1 ttl=64 time=0.686 ms
1578 bytes from 192.168.250.100: icmp_seq=2 ttl=64 time=0.758 ms

--- 192.168.250.100 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.686/0.913/1.294 ms
~ #
```

```
root@esxcomp-01a ~ # vmkping ++netstack=vxlan -d -s 1570 192.168.250.101
PING 192.168.250.101 (192.168.250.101): 1570 data bytes
1578 bytes from 192.168.250.101: icmp_seq=0 ttl=64 time=0.065 ms
1578 bytes from 192.168.250.101: icmp_seq=1 ttl=64 time=0.118 ms

--- 192.168.250.101 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0.065/0.091/0.118 ms
```

What to do next

Create a DLR and attach it to your logical switches to enable connectivity between VMs that are connected to different logical switches.

Add a Logical (Distributed) Router

Logical router kernel modules in the host perform routing between VXLAN networks, and between virtual and physical networks. An NSX Edge appliance provides dynamic routing ability if needed. Logical routers can be created on both primary and secondary NSX Managers in a cross-vCenter NSX environment, but universal logical routers can be created only on the primary NSX Manager.

The following list describes feature support by interface type (uplink and internal) on the logical router:

- Dynamic routing protocols (BGP and OSPF) are supported only on uplink interfaces.
- Firewall rules are applicable only on uplink interfaces and are limited to control and management traffic that is destined to the edge virtual appliance.
- For more information about the DLR Management Interface see the Knowledge Base Article, Considerations for Management Interface of Distributed Logical Router Control VM, <http://kb.vmware.com/kb/2122060>.

Prerequisites

- You must have been assigned the Enterprise Administrator or NSX Administrator role.
- You must have an operational controller cluster in your environment before installing a logical router.
- You must create a local segment ID pool, even if you have no plans to create NSX logical switches.
- A logical router cannot distribute routing information to hosts without the help of NSX controllers. A logical router relies on NSX controllers to function, while edge services gateways (ESGs) do not. Make sure the controller cluster is up and available before creating or changing a logical router configuration.
- If a logical router is to be connected to VLAN dvPortgroups, ensure that all hypervisor hosts with a logical router appliance installed can reach each other on UDP port 6999 for logical router VLAN-based ARP proxy to work.
- Logical router interfaces and bridging interfaces cannot be connected to a dvPortgroup with the VLAN ID set to 0.
- A given logical router instance cannot be connected to logical switches that exist in different transport zones. This is to ensure that all logical switches and logical router instances are aligned.
- A logical router cannot be connected to VLAN-backed portgroups if that logical router is connected to logical switches spanning more than one vSphere distributed switch (VDS). This is to ensure correct alignment of logical router instances with logical switch dvPortgroups across hosts.
- Logical router interfaces should not be created on two different distributed portgroups (dvPortgroups) with the same VLAN ID if the two networks are in the same vSphere distributed switch.

- Logical router interfaces should not be created on two different dvPortgroups with the same VLAN ID if two networks are in different vSphere distributed switches, but the two vSphere distributed switches share the same hosts. In other words, logical router interfaces can be created on two different networks with the same VLAN ID if the two dvPortgroups are in two different vSphere distributed switches, as long as the vSphere distributed switches do not share a host.
- Unlike NSX version 6.0 and 6.1, NSX version 6.2 allows logical router-routed logical interfaces (LIFs) to be connected to a VXLAN that is bridged to a VLAN.
- When selecting placement of a logical router virtual appliance, avoid placing it on the same host as one or more of its upstream ESGs if you use ESGs in an ECMP setup. You can use DRS anti-affinity rules to enforce this, thus reducing the impact of host failure on logical router forwarding. This guideline does not apply if you have one upstream ESG by itself or in HA mode. For more information, see the *VMware NSX for vSphere Network Virtualization Design Guide* at <https://communities.vmware.com/docs/DOC-27683>.

Procedure

- 1 In the vSphere Web Client, navigate to **Home > Networking & Security > NSX Edges**.
- 2 Click the **Add** () icon.
- 3 Select **Logical (Distributed) Router** and type a name for the device.

This name appears in your vCenter inventory. The name should be unique across all logical routers within a single tenant.

Optionally, you can also enter a hostname. This name appears in the CLI. If you do not specify the host name, the Edge ID, which gets created automatically, is displayed in the CLI.

Optionally, you can enter a description and tenant.

For example:

The screenshot shows the 'New NSX Edge' configuration window. On the left, a navigation pane lists steps: 1 Name and description (selected), 2 Settings, 3 Configure deployment, 4 Configure interfaces, 5 Default gateway settings, and 6 Ready to complete. The main area is titled 'Name and description'. Under 'Install Type', there are two radio buttons: 'Edge Services Gateway' (unselected) and 'Logical (Distributed) Router' (selected). Below the radio buttons are input fields for 'Name' (containing 'logical-distributed-router1'), 'Hostname', 'Description', and 'Tenant'. At the bottom, there are two checkboxes: 'Deploy Edge Appliance' (checked) and 'Enable High Availability' (unchecked). At the very bottom of the window are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

4 (Optional) Deploy an edge appliance.

Deploy Edge Appliance is selected by default. An edge appliance (also called a logical router virtual appliance) is required for dynamic routing and the logical router appliance's firewall, which applies to logical router pings, SSH access, and dynamic routing traffic.

You can deselect the edge appliance option if you require only static routes, and do not want to deploy an Edge appliance. You cannot add an Edge appliance to the logical router after the logical router has been created.

5 (Optional) Enable High Availability.

Enable High Availability is not selected by default. Select the Enable High Availability check box to enable and configure high availability. High availability is required if you are planning to do dynamic routing.

6 Type and re-type a password for the logical router.

The password must be 12-255 characters and must contain the following:

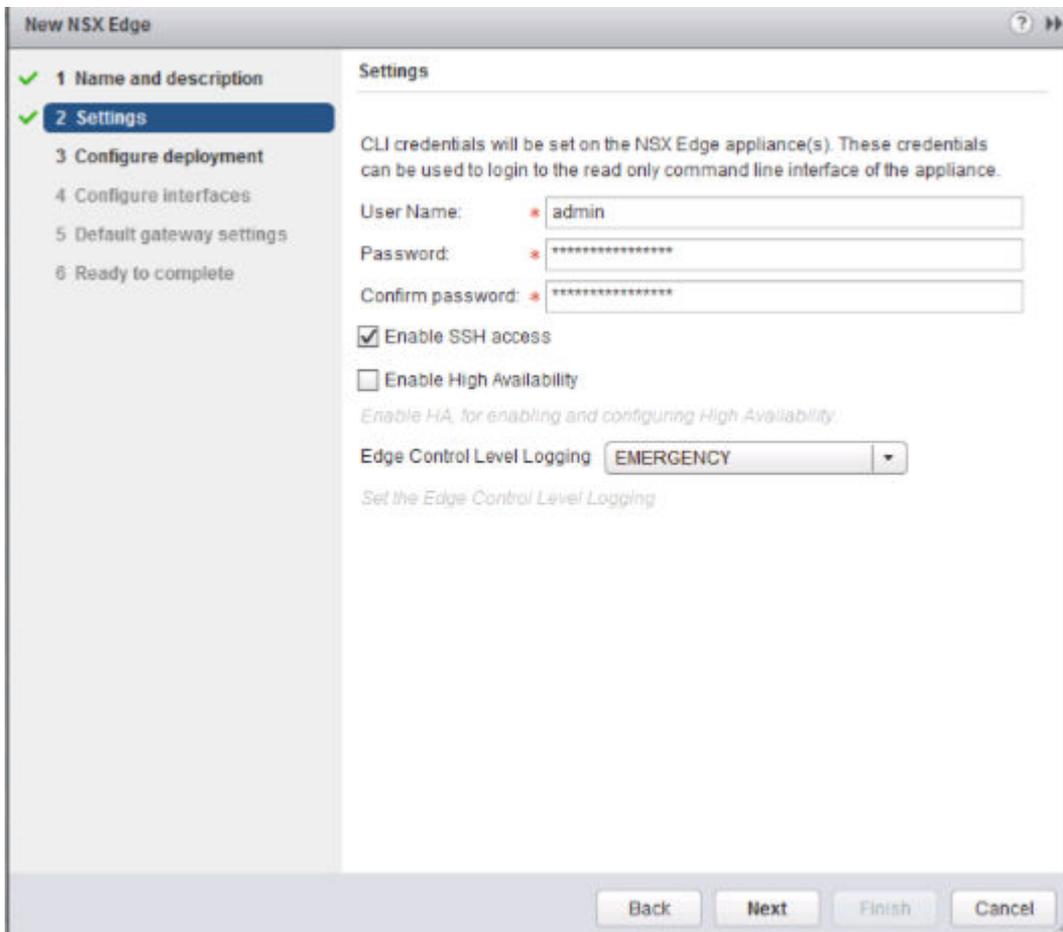
- At least one upper case letter
- At least one lower case letter
- At last one number
- At least one special character

7 (Optional) Enable SSH and set the log level.

By default, SSH is disabled. If you do not enable SSH, you can still access the logical router by opening the virtual appliance console. Enabling SSH here causes the SSH process to run on the logical router virtual appliance, but you will also need to adjust the logical router firewall configuration manually to allow SSH access to the logical router's protocol address. The protocol address is configured when you configure dynamic routing on the logical router.

By default, the log level is emergency.

For example:



8 Configure deployment.

- ◆ If you did not select **Deploy NSX Edge**, the **Add (+)** icon is grayed out. Click **Next** to continue with configuration.
- ◆ If you selected **Deploy NSX Edge**, enter the settings for the logical router virtual appliance that will be added to your vCenter inventory.

For example:

The screenshot shows a dialog box titled "Edit NSX Edge Appliance" with a help icon. The instruction reads: "Specify placement parameters for the NSX Edge appliance." The configuration fields are as follows:

Field	Value
Cluster/Resource Pool	Management and Edg...
Datastore	ds-1
Host	192.168.110.52
Folder	Discovered virtual mac...

Buttons: OK, Cancel

9 Configure interfaces.

On logical routers, only IPv4 addressing is supported.

In the HA Interface Configuration, if you selected **Deploy NSX Edge** you must connect the interface to a distributed port group. It is recommended to use a VXLAN logical switch for the HA interface. An IP address for each of the two NSX Edge appliances is chosen from the link local address space, 169.250.0.0/16. No further configuration is necessary to configure the HA service.

Note In prior releases of NSX, the HA interface was called the management interface. The HA interface is not supported for remote access to the logical router. You cannot SSH into the HA interface from anywhere that isn't on the same IP subnet as the HA interface. You cannot configure static routes that point out of the HA interface, which means that RPF will drop incoming traffic. You could, in theory, disable RPF, but this would be counterproductive to high availability. For SSH, use the logical router's protocol address, which is configured later when you configure dynamic routing.

In NSX 6.2, the HA interface of a logical router is automatically excluded from route redistribution.

In **Configure interfaces of this NSX Edge** the internal interfaces are for connections to switches that allow VM-to-VM (sometimes called East-West) communication. Internal interfaces are created as pseudo vNICs on the logical router virtual appliance. Uplink interfaces are for North-South communication. A logical router uplink interface might connect to an NSX edge services gateway, a third-party router VM for that, or a VLAN-backed dvPortgroup to make the logical router connect to a physical router directly. You must have at least one uplink interface for dynamic routing to work. Uplink interfaces are created as vNICs on the logical router virtual appliance.

The interface configuration that you enter here is modifiable later. You can add, remove, and modify interfaces after a logical router is deployed.

The following example shows an HA interface connected to the management distributed port group. The example also shows two internal interfaces (app and web) and an uplink interface (to-ESG).

New NSX Edge ? >>

- ✓ 1 Name and description
- ✓ 2 Settings
- ✓ 3 Configure deployment
- 4 Configure interfaces**
- 5 Default gateway settings
- 6 Ready to complete

Configure interfaces

HA Interface Configuration

Connected To: [Change](#) [Remove](#)

+ ✎ ✕

IP Address	Subnet Prefix Length
192.168.110.60*	24

HA interface is a mandatory special-purpose interface that requires network connectivity and is configured separately from other interfaces in the Logical Router.

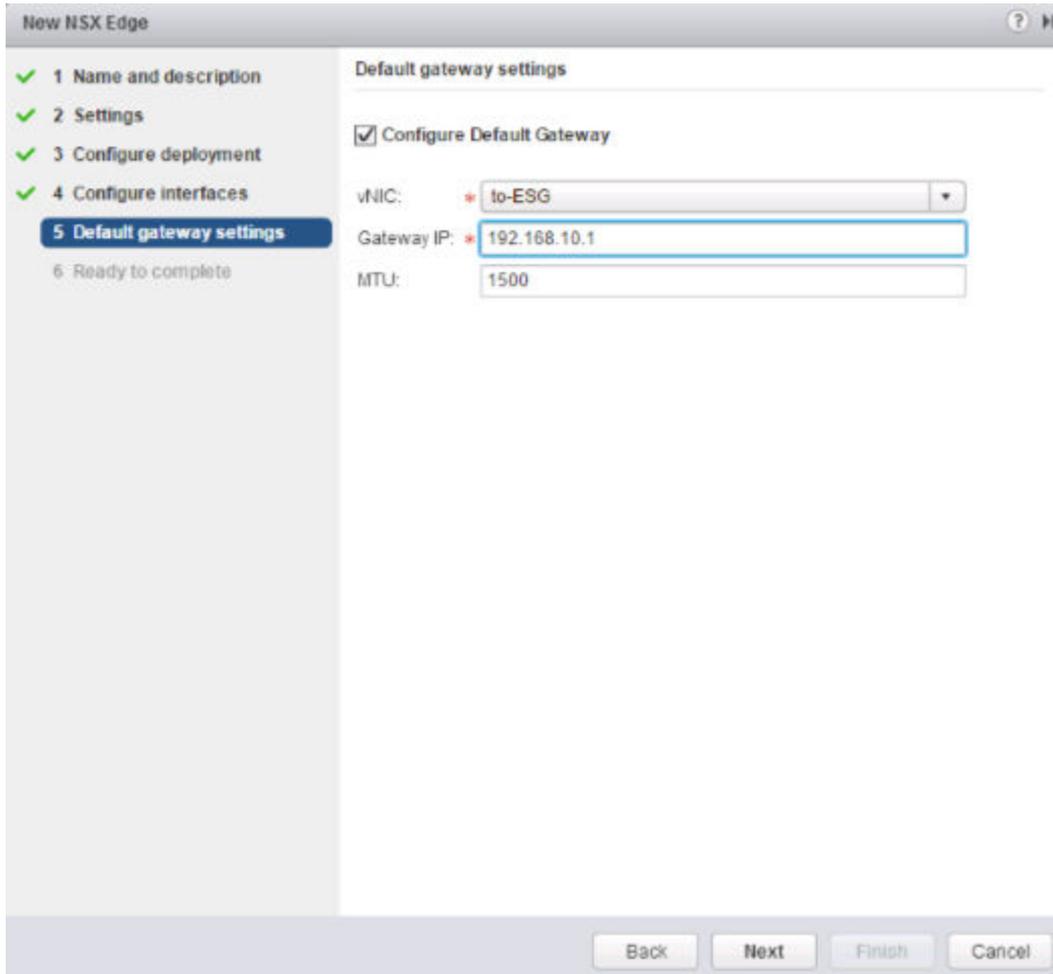
Configure interfaces of this NSX Edge

+ ✎ ✕

Name	IP Address	Subnet Prefix Length	Connected To
app	172.16.20.1*	24	app
web	172.16.10.1*	24	web
to-ESG	192.168.10.2*	29	transit

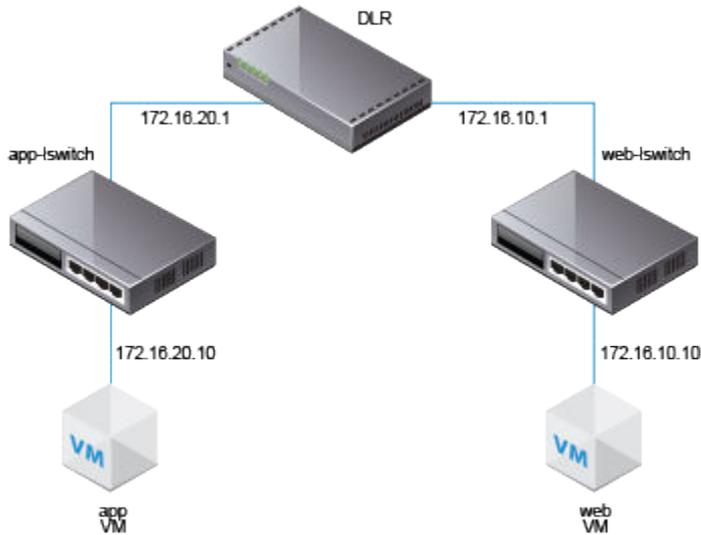
10 Configure a default gateway.

For example:



11 Make sure any VMs attached to the logical switches have their default gateways set properly to the logical router interface IP addresses.

In the following example topology, the default gateway of app VM should be 172.16.20.1. The default gateway of web VM should be 172.16.10.1. Make sure the VMs can ping their default gateways and each other.



Log in via SSH to the NSX Manager, and run the following commands:

- List all logical router instance information.

```
nsxmgr-l-01a> show logical-router list all
Edge-id          Vdr Name          Vdr id           #Lifs
-----          -
edge-1          default+edge-1    0x00001388      3
```

- List the hosts that have received routing information for the logical router from the controller cluster.

```
nsxmgr-l-01a> show logical-router list dlr edge-1 host
ID              HostName
-----
host-25        192.168.210.52
host-26        192.168.210.53
host-24        192.168.110.53
```

The output includes all hosts from all host clusters that are configured as members of the transport zone that owns the logical switch that is connected to the specified logical router (edge-1 in this example).

- List the routing table information that is communicated to the hosts by the logical router. Routing table entries should be consistent across all of the hosts.

```
nsx-mgr-l-01a> show logical-router host host-25 dlr edge-1 route

VDR default+edge-1 Route Table
Legend: [U: Up], [G: Gateway], [C: Connected], [I: Interface]
Legend: [H: Host], [F: Soft Flush] [!: Reject] [E: ECMP]

Destination      GenMask          Gateway           Flags  Ref Origin  UpTime  Interface
-----
0.0.0.0          0.0.0.0         192.168.10.1    UG     1  AUTO      4101  138800000002
```

172.16.10.0	255.255.255.0	0.0.0.0	UCI	1	MANUAL	10195	13880000000b
172.16.20.0	255.255.255.0	0.0.0.0	UCI	1	MANUAL	10196	13880000000a
192.168.10.0	255.255.255.248	0.0.0.0	UCI	1	MANUAL	10196	138800000002
192.168.100.0	255.255.255.0	192.168.10.1	UG	1	AUTO	3802	138800000002

- List additional information about the router from the point of view of one of the hosts. This is helpful to learn which controller is communicating with the host.

```

nsx-mgr-l-01a> show logical-router host host-25 dlr edge-1 verbose

VDR Instance Information :
-----

Vdr Name:                default+edge-1
Vdr Id:                  0x00001388
Number of Lifs:          3
Number of Routes:        5
State:                   Enabled
Controller IP:           192.168.110.203
Control Plane IP:        192.168.210.52
Control Plane Active:    Yes
Num unique nexthops:     1
Generation Number:       0
Edge Active:             No
    
```

Check the Controller IP field in the output of the `show logical-router host host-25 dlr edge-1 verbose` command.

SSH to a controller, and run the following commands to display the controller's learned VNI, VTEP, MAC, and ARP table state information.

- ```

192.168.110.202 # show control-cluster logical-switches vni 5000
VNI Controller BUM-Replication ARP-Proxy Connections
5000 192.168.110.201 Enabled Enabled 0

```

The output for VNI 5000 shows zero connections and lists controller 192.168.110.201 as the owner for VNI 5000. Log in to that controller to gather further information for VNI 5000.

```

192.168.110.201 # show control-cluster logical-switches vni 5000
VNI Controller BUM-Replication ARP-Proxy Connections
5000 192.168.110.201 Enabled Enabled 3

```

The output on 192.168.110.201 shows three connections. Check additional VNIs.

```
192.168.110.201 # show control-cluster logical-switches vni 5001
VNI Controller BUM-Replication ARP-Proxy Connections
5001 192.168.110.201 Enabled Enabled 3
```

```
192.168.110.201 # show control-cluster logical-switches vni 5002
VNI Controller BUM-Replication ARP-Proxy Connections
5002 192.168.110.201 Enabled Enabled 3
```

Because 192.168.110.201 owns all three VNI connections, we would expect to see zero connections on the other controller, 192.168.110.203.

```
192.168.110.203 # show control-cluster logical-switches vni 5000
VNI Controller BUM-Replication ARP-Proxy Connections
5000 192.168.110.201 Enabled Enabled 0
```

- Before checking the MAC and ARP tables, start pinging from one VM to the other VM.

From app VM to web VM:

```
vmware@vmware-virtual-machine:~$ ping 172.16.10.10
PING 172.16.10.10 (172.16.10.10) 56(84) bytes of data.
64 bytes from 172.16.10.10: icmp_seq=1 ttl=63 time=2.15 ms
64 bytes from 172.16.10.10: icmp_seq=2 ttl=63 time=1.60 ms
```

Check the MAC tables.

```
192.168.110.201 # show control-cluster logical-switches mac-table 5000
VNI MAC VTEP-IP Connection-ID
5000 00:50:56:a6:23:ae 192.168.250.52 7
```

```
192.168.110.201 # show control-cluster logical-switches mac-table 5001
VNI MAC VTEP-IP Connection-ID
5001 00:50:56:a6:8d:72 192.168.250.51 23
```

Check the ARP tables.

```
192.168.110.201 # show control-cluster logical-switches arp-table 5000
VNI IP MAC Connection-ID
5000 172.16.20.10 00:50:56:a6:23:ae 7
```

```
192.168.110.201 # show control-cluster logical-switches arp-table 5001
VNI IP MAC Connection-ID
5001 172.16.10.10 00:50:56:a6:8d:72 23
```

Check the logical router information. Each logical router Instance is served by one of the controller nodes.

The instance sub-command of `show control-cluster logical-routers` command displays a list of logical routers that are connected to this controller.

The interface-summary sub-command displays the LIFs that the controller learned from the NSX Manager. This information is sent to the hosts that are in the host clusters managed under the transport zone.

The routes sub-command shows the routing table that is sent to this controller by the logical router's virtual appliance (also known as the control VM). Note that unlike on the ESXi hosts, this routing table does not include directly connected subnets because this information is provided by the LIF configuration. Route information on the ESXi hosts includes directly connected subnets because in that case it is a forwarding table used by ESXi host's datapath.

```

controller # show control-cluster logical-routers instance all
LR-Id LR-Name Universal Service-Controller Egress-Locale
0x1388 default+edge-1 false 192.168.110.201 local

```

Note the LR-Id and use it in the following command.

```

controller # show control-cluster logical-routers interface-summary 0x1388
Interface Type Id IP[]
13880000000b vxlan 0x1389 172.16.10.1/24
13880000000a vxlan 0x1388 172.16.20.1/24
138800000002 vxlan 0x138a 192.168.10.2/29

```

```

controller # show control-cluster logical-routers routes 0x1388
Destination Next-Hop[] Preference Locale-Id Source
192.168.100.0/24 192.168.10.1 110 00000000-0000-0000-0000-000000000000 CONTROL_VM
0.0.0.0/0 192.168.10.1 0 00000000-0000-0000-0000-000000000000 CONTROL_VM

```

```

[root@comp02a:~] esxcfg-route -l
VMkernel Routes:
Network Netmask Gateway Interface
10.20.20.0 255.255.255.0 Local Subnet vmk1
192.168.210.0 255.255.255.0 Local Subnet vmk0
default 0.0.0.0 192.168.210.1 vmk0

```

- Display the controller connections to the specific VNI.

```
192.168.110.203 # show control-cluster logical-switches connection-table 5000
Host-IP Port ID
192.168.110.53 26167 4
192.168.210.52 27645 5
192.168.210.53 40895 6
```

```
192.168.110.202 # show control-cluster logical-switches connection-table 5001
Host-IP Port ID
192.168.110.53 26167 4
192.168.210.52 27645 5
192.168.210.53 40895 6
```

These Host-IP addresses are vmk0 interfaces, not VTEPs. Connections between ESXi hosts and controllers are created on the management network. The port numbers here are ephemeral TCP ports that are allocated by the ESXi host IP stack when the host establishes a connection with the controller.

- On the host, you can view the controller network connection matched to the port number.

```
[root@192.168.110.53:~] #esxcli network ip connection list | grep 26167
tcp 0 0 192.168.110.53:26167 192.168.110.101:1234 ESTABLISHED
96416 newreno netcpa-worker
```

- Display active VNIs on the host. Observe how the output is different across hosts. Not all VNIs are active on all hosts. A VNI is active on a host if the host has a VM that is connected to the logical switch.

```
[root@192.168.210.52:~] # esxcli network vswitch dvs vmware vxlan network list --vds-name
Compute_VDS
VXLAN ID Multicast IP Control Plane Controller Connection
Port Count MAC Entry Count ARP Entry Count VTEP Count

5000 N/A (headend replication) Enabled (multicast proxy,ARP proxy) 192.168.110.203
(up) 1 0 0 0
5001 N/A (headend replication) Enabled (multicast proxy,ARP proxy) 192.168.110.202
(up) 1 0 0 0
```

**Note** To enable the vxlan namespace in vSphere 6 and later, run the `/etc/init.d/hostd restart` command.

For logical switches in hybrid or unicast mode, the `esxcli network vswitch dvs vmware vxlan network list --vds-name <vds-name>` command should contain the following output:

- Control Plane is enabled.
- Multicast proxy and ARP proxy are listed. AARP proxy is listed even if you disabled IP discovery.
- A valid controller IP address is listed and the connection is up.

- If a logical router is connected to the ESXi host, the port Count is at least 1, even if there are no VMs on the host connected to the logical switch. This one port is the vdrPort, which is a special dvPort connected to the logical router kernel module on the ESXi host.
- First ping from VM to another VM on a different subnet and then display the MAC table. Note the Inner MAC is the VM entry while the Outer MAC and Outer IP refer to the VTEP.

```
~ # esxcli network vswitch dvs vmware vxlan network mac list --vds-name=Compute_VDS --vxlan-id=5000
Inner MAC Outer MAC Outer IP Flags

00:50:56:a6:23:ae 00:50:56:6a:65:c2 192.168.250.52 00000111
```

```
~ # esxcli network vswitch dvs vmware vxlan network mac list --vds-name=Compute_VDS --vxlan-id=5001
Inner MAC Outer MAC Outer IP Flags

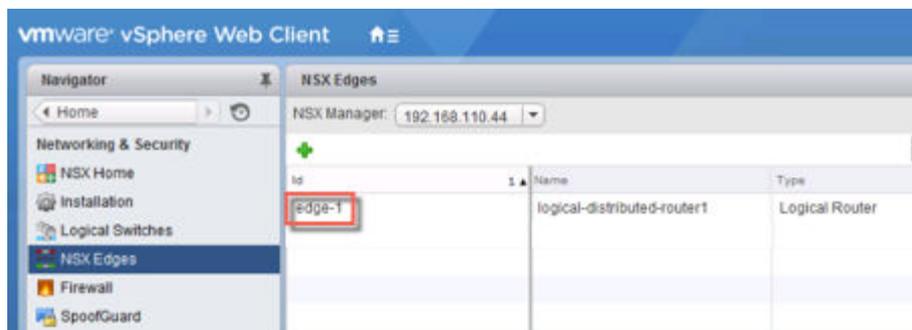
02:50:56:56:44:52 00:50:56:6a:65:c2 192.168.250.52 00000101
00:50:56:f0:d7:e4 00:50:56:6a:65:c2 192.168.250.52 00000111
```

**What to do next**

On the hosts where NSX edge appliances are first deployed, NSX enables automatic VM startup/shutdown. If the appliance VMs are later migrated to other hosts, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, VMware recommends that you check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See [http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm\\_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html](http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html).

After the logical router is deployed, double-click the logical router ID to configure additional settings, such as interfaces, routing, firewall, bridging, and DHCP relay.

For example:



## Add an Edge Services Gateway

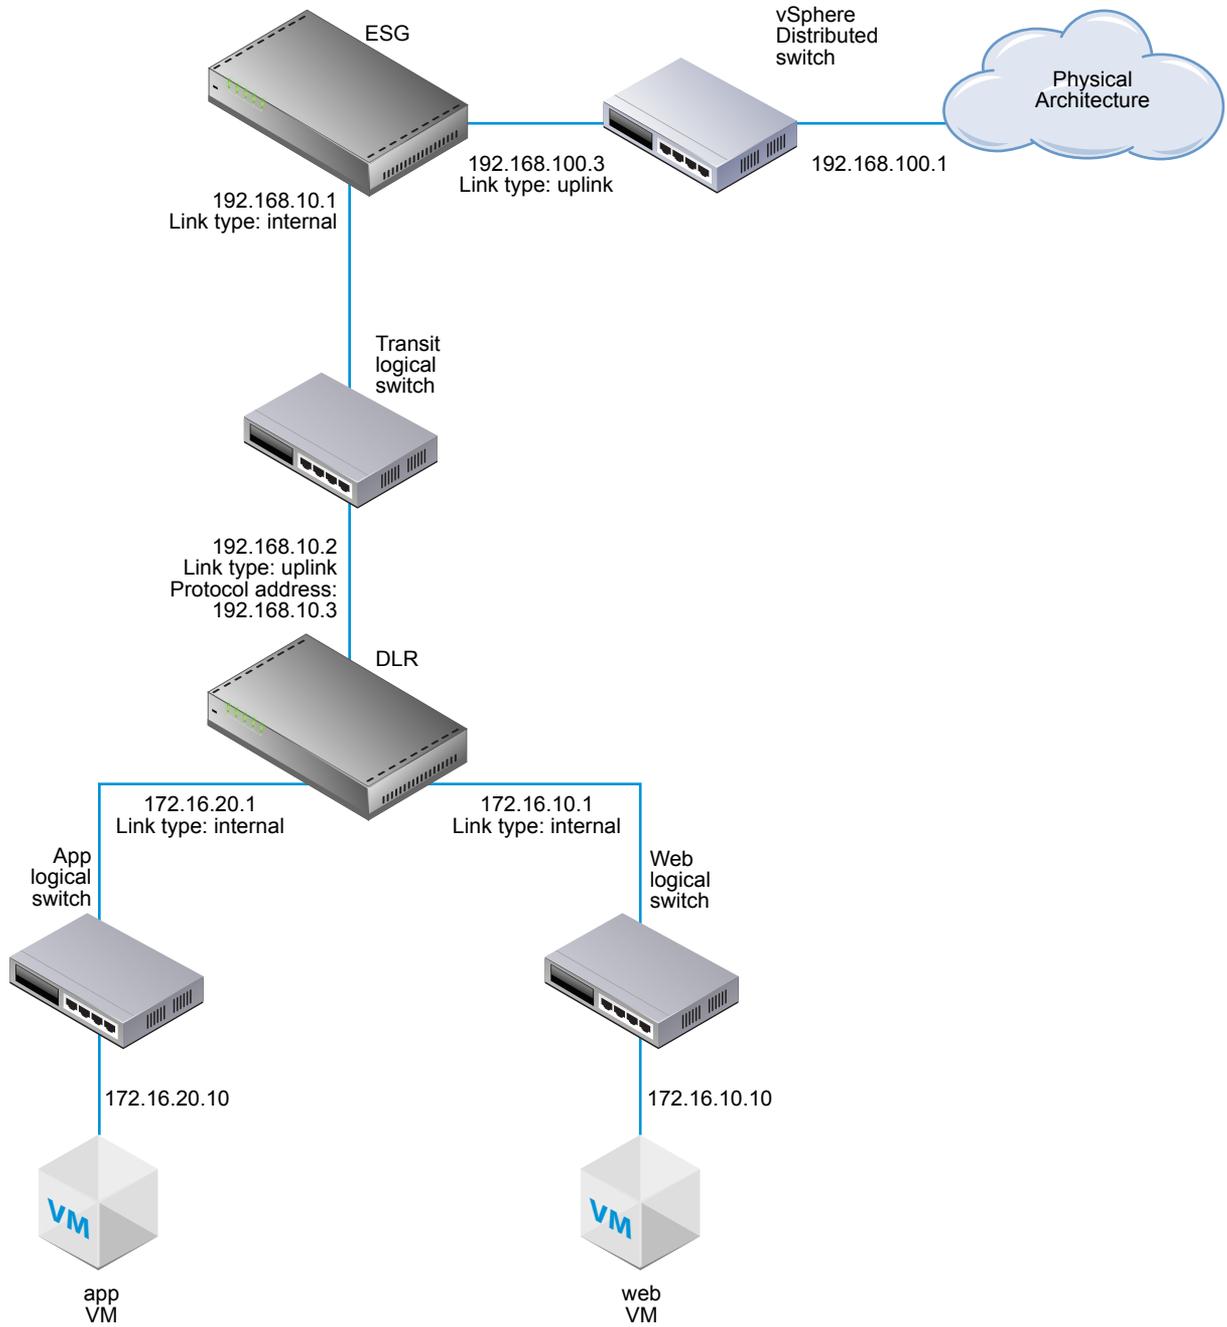
You can install multiple NSX Edge services gateway virtual appliances in a data center. Each NSX Edge virtual appliance can have a total of ten uplink and internal network interfaces. The internal interfaces connect to secured port groups and act as the gateway for all protected virtual machines in the port group. The subnet assigned to the internal interface can be a publicly routed IP address space or a NATed/routed RFC 1918 private space. Firewall rules and other NSX Edge services are enforced on traffic between interfaces.

Uplink interfaces of an ESG connect to uplink port groups that have access to a shared corporate network or a service that provides access layer networking.

The following list describes feature support by interface type (internal and uplink) on an ESG.

- DHCP: Not supported on uplink interface.
- DNS Forwarder: Not supported on uplink interface.
- HA: Not supported on uplink interface, requires at least one internal interface.
- SSL VPN: Listener IP must belong to uplink interface.
- IPsec VPN: Local site IP must belong to uplink interface.
- L2 VPN: Only internal networks can be stretched.

The following figure shows a sample topology with an ESG's uplink interface connected to physical infrastructure through the vSphere distributed switch and the ESG's internal interface connect to an NSX logical router through an NSX logical transit switch.



Multiple external IP addresses can be configured for load balancing, site-to-site VPN, and NAT services.

**Prerequisites**

You must have been assigned the Enterprise Administrator or NSX Administrator role.

Verify that the resource pool has enough capacity for the edge services gateway (ESG) virtual appliance to be deployed. See [System Requirements for NSX](#).

## Procedure

- 1 In vCenter, navigate to **Home > Networking & Security > NSX Edges** and click the **Add (+)** icon.
- 2 Select **Edge Services Gateway** and type a name for the device.

This name appears in your vCenter inventory. The name should be unique across all ESGs within a single tenant.

Optionally, you can also enter a hostname. This name appears in the CLI. If you do not specify the host name, the Edge ID, which gets created automatically, is displayed in the CLI.

Optionally, you can enter a description and tenant and enable high availability.

For example:

The screenshot shows the 'New NSX Edge' configuration wizard. The left sidebar lists the steps: 1 Name and description (selected), 2 Settings, 3 Configure deployment, 4 Configure interfaces, 5 Default gateway settings, 6 Firewall and HA, and 7 Ready to complete. The main area is titled 'Name and description' and contains the following fields and options:

- Install Type:** Radio buttons for 'Edge Services Gateway' (selected) and 'Logical (Distributed) Router'.
- Name:** Text box containing 'ESG-1'.
- Hostname:** Empty text box.
- Description:** Empty text area.
- Tenant:** Empty text box.
- Deploy NSX Edge:** Checked checkbox.
- Enable High Availability:** Unchecked checkbox.

At the bottom of the wizard are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

- 3 Type and re-type a password for the ESG.

The password must be at least 12 characters and must follow 3 of the following 4 rules:

- At least one upper case letter
- At least one lower case letter
- At last one number

- At least one special character

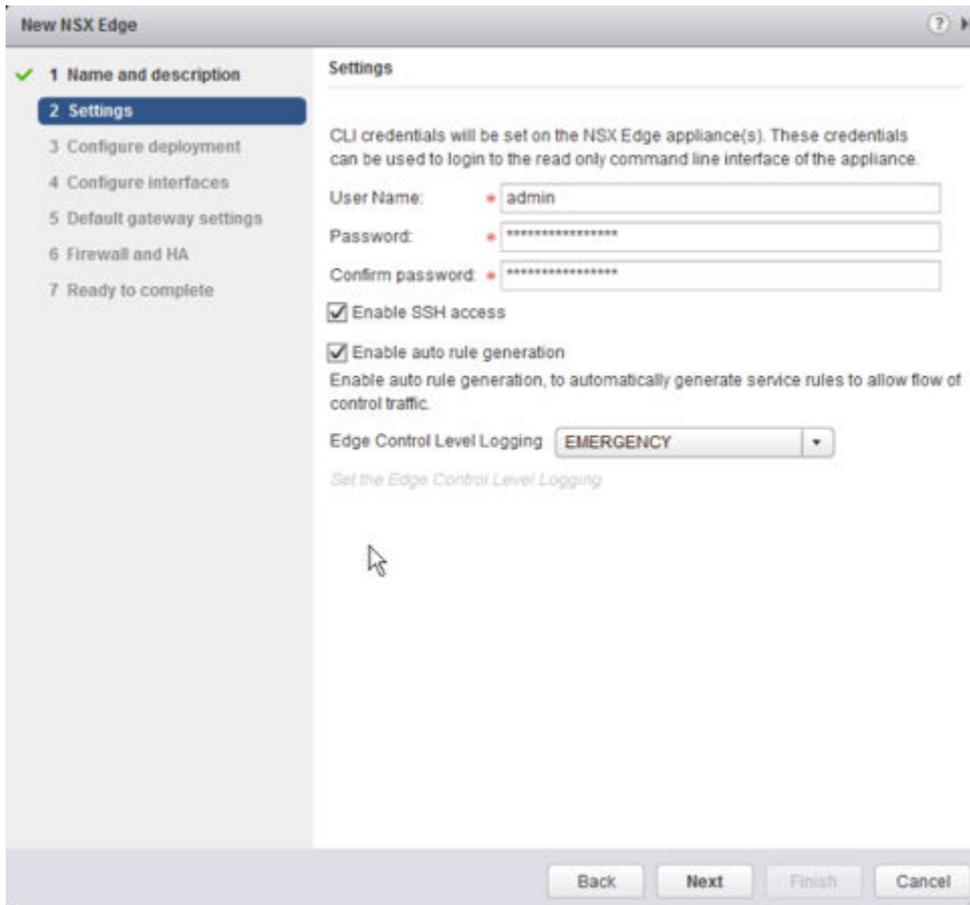
4 (Optional) Enable SSH, high availability, and automatic rule generation, and set the log level.

If you do not enable automatic rule generation, you must manually add firewall, NAT, and routing configuration to allow control traffic for certain, NSX Edge services, including as load balancing and VPN. Auto rule generation does not create rules for data-channel traffic.

By default, SSH and high availability are disabled, and automatic rule generation is enabled. By default, the log level is emergency.

By default, logging is enabled on all new NSX Edge appliances. The default logging level is NOTICE.

For example:



5 Select the size of the NSX Edge instance based on your system resources.

The **Large** NSX Edge has more CPU, memory, and disk space than the **Compact** NSX Edge, and supports a larger number of concurrent SSL VPN-Plus users. The **X-Large** NSX Edge is suited for environments that have a load balancer with millions of concurrent sessions. The Quad Large NSX Edge is recommended for high throughput and requires a high connection rate.

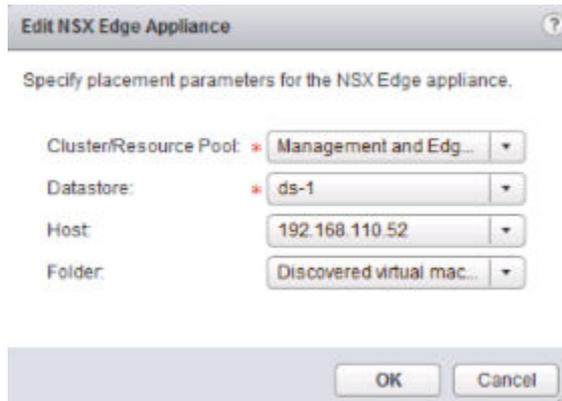
See [System Requirements for NSX](#).

## 6 Create an edge appliance.

Enter the settings for the ESG virtual appliance that will be added to your vCenter inventory. If you do not add an appliance when you install NSX Edge, NSX Edge remains in an offline mode until you add an appliance.

If you enabled HA you can add two appliances. If you add a single appliance, NSX Edge replicates its configuration for the standby appliance and ensures that the two HA NSX Edge virtual machines are not on the same ESX host even after you use DRS and vMotion (unless you manually vMotion them to the same host). For HA to work correctly, you must deploy both appliances on a shared datastore.

For example:



- 7 Select **Deploy NSX Edge** to add the Edge in a deployed mode. You must configure appliances and interfaces for the Edge before it can be deployed.

## 8 Configure interfaces.

On ESGs, both IPv4 and IPv6 addresses are supported.

You must add at least one internal interface for HA to work.

An interface can have multiple non-overlapping subnets.

If you enter more than one IP address for an interface, you can select the primary IP address. An interface can have one primary and multiple secondary IP addresses. NSX Edge considers the primary IP address as the source address for locally generated traffic, for example remote syslog and operator-initiated pings.

You must add an IP address to an interface before using it on any feature configuration.

Optionally, you can enter the MAC address for the interface.

If HA is enabled, you can optionally enter two management IP addresses in CIDR format. Heartbeats of the two NSX Edge HA virtual machines are communicated through these management IP addresses. The management IP addresses must be in the same L2/subnet and be able to communicate with each other.

Optionally, you can modify the MTU.

Enable proxy ARP if you want to allow the ESG to answer ARP requests intended for other machines. This is useful, for example, when you have the same subnet on both sides of a WAN connection.

Enable ICMP redirect to convey routing information to hosts.

Enable reverse path filtering to verify the reachability of the source address in packets being forwarded. In enabled mode, the packet must be received on the interface that the router would use to forward the return packet. In loose mode, the source address must appear in the routing table.

Configure fence parameters if you want to reuse IP and MAC addresses across different fenced environments. For example, in a cloud management platform (CMP), fencing allow you to run several cloud instances simultaneous with the same IP and MAC addresses completely isolated or “fenced.”

For example:

**Edit NSX Edge Interface**

vNIC#: 1

Name: \* Internal

Type:  Internal  Uplink

Connected To: transit-switch Change Remove

Connectivity Status:  Connected  Disconnected

Configure subnets

| IP Address    | Subnet Prefix Length |
|---------------|----------------------|
| 192.168.10.1* | 29                   |
|               |                      |
|               |                      |

MAC Addresses:

You can specify a MAC address or leave it blank for auto generation. In case of HA, two different MAC addresses are required.

MTU:

Options:  Enable Proxy ARP  Send ICMP Redirect Reverse Path Filter Disable ▾

Fence Parameters:

Example: ethernet0.filter1.param1=1

OK Cancel

The following example shows two interfaces, one attaching the ESG to the outside world through an uplink portgroup on a vSphere distributed switch and the other attaching the ESG to a logical transit switch to which a distributed logical router is also attached.

**New NSX Edge**

- ✓ 1 Name and description
- ✓ 2 Settings
- ✓ 3 Configure deployment
- ✓ 4 Configure interfaces**
- 5 Default gateway settings
- 6 Firewall and HA
- 7 Ready to complete

**Configure interfaces**

Configure interfaces of this NSX Edge

+ ✎ ✖

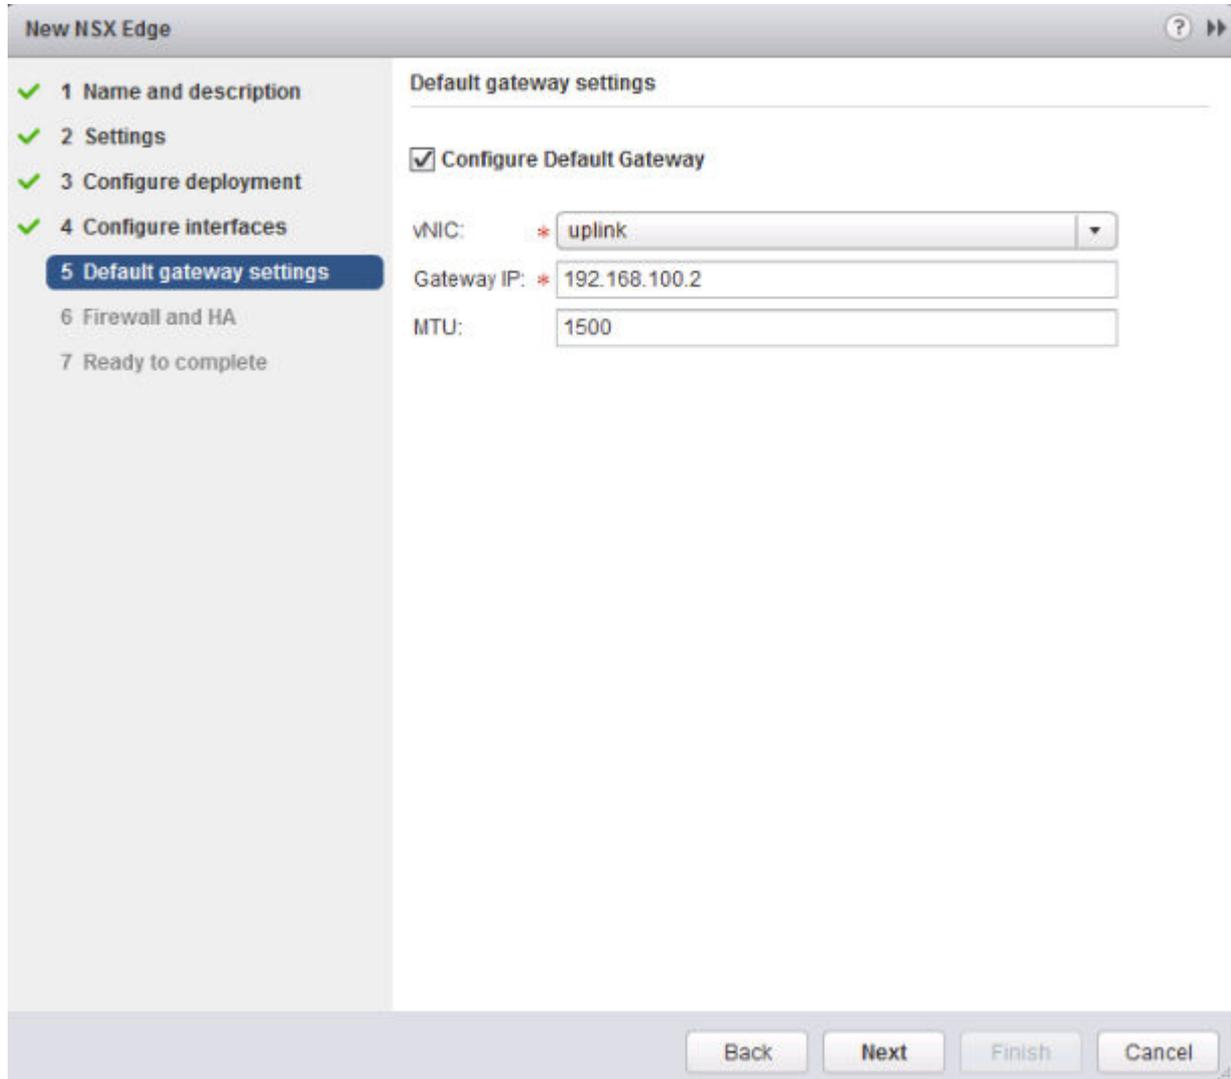
| vNIC# | Name     | IP Address    | Subnet Prefix Length | Connected To         |
|-------|----------|---------------|----------------------|----------------------|
| 0     | uplink   | 192.168.100.3 | 24                   | Mgmt_VDS - HQ_Uplink |
| 1     | internal | 192.168.10.1* | 29                   | transit-switch       |
|       |          |               |                      |                      |
|       |          |               |                      |                      |
|       |          |               |                      |                      |
|       |          |               |                      |                      |
|       |          |               |                      |                      |

Back Next Finish Cancel

9 Configure a default gateway.

You can edit the MTU value, but it cannot be more than the configured MTU on the interface.

For example:



10 Configure the firewall policy, logging, and HA parameters.

**Caution** If you do not configure the firewall policy, the default policy is set to deny all traffic.

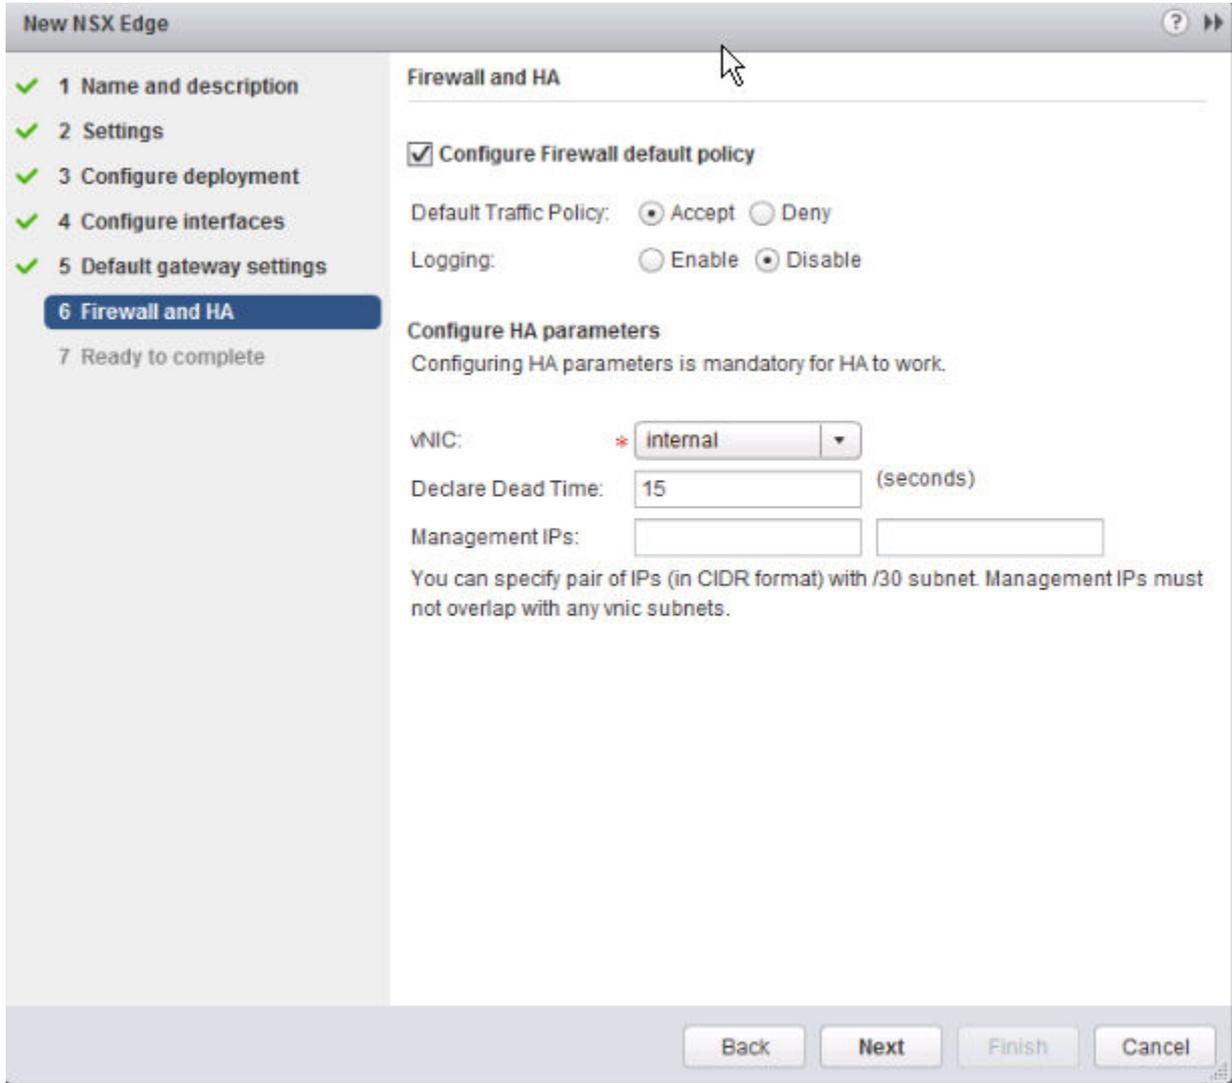
By default, logs are enabled on all new NSX Edge appliances. The default logging level is NOTICE. If logs are stored locally on the ESG, logging may generate too many logs and affect the performance of your NSX Edge. For this reason, it is recommended that you configure remote syslog servers, and forward all logs to a centralized collector for analysis and monitoring.

If you enabled high availability, complete the HA section. By default, HA automatically chooses an internal interface and automatically assigns link-local IP addresses. NSX Edge supports two virtual machines for high availability, both of which are kept up to date with user configurations. If a heartbeat failure occurs on the primary virtual machine, the secondary virtual machine state is

changed to active. Thus, one NSX Edge virtual machine is always active on the network. NSX Edge replicates the configuration of the primary appliance for the standby appliance and ensures that the two HA NSX Edge virtual machines are not on the same ESX host even after you use DRS and vMotion. Two virtual machines are deployed on vCenter in the same resource pool and datastore as the appliance you configured. Local link IP addresses are assigned to HA virtual machines in the NSX Edge HA so that they can communicate with each other. Select the internal interface for which to configure HA parameters. If you select ANY for interface but there are no internal interfaces configured, the UI does not display an error. Two Edge appliances are created but since there is no internal interface configured, the new Edge remains in standby and HA is disabled. Once an internal interface is configured, HA will get enabled on the Edge appliance. Type the period in seconds within which, if the backup appliance does not receive a heartbeat signal from the primary appliance, the primary appliance is considered inactive and the backup appliance takes over. The default interval is 15 seconds. Optionally, you can enter two management IP addresses in CIDR format to override the

local link IP addresses assigned to the HA virtual machines. Ensure that the management IP addresses do not overlap with the IP addresses used for any other interface and do not interfere with traffic routing. You should not use an IP address that exists somewhere else on your network, even if that network is not directly attached to the NSX Edge.

For example:



After the ESG is deployed, go to the Hosts and Clusters view and open the console of the edge virtual appliance. From the console, make sure you can ping the connected interfaces.

## What to do next

On the hosts where NSX edge appliances are first deployed, NSX enables automatic VM startup/shutdown. If the appliance VMs are later migrated to other hosts, the new hosts might not have automatic VM startup/shutdown enabled. For this reason, VMware recommends that you check all hosts in the cluster to make sure that automatic VM startup/shutdown is enabled. See

[http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm\\_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html](http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc%2FGUID-5FE08AC7-4486-438E-AF88-80D6C7928810.html).

Now you can configure routing to allow connectivity from external devices to your VMs.

## Add a Host to a Prepared Cluster

This section describes how to add a host to a cluster prepared for network virtualization.

### Procedure

- 1 Add the host to vCenter Server as a standalone host.

See *ESXi and vCenter Server 5.5 Documentation*.

- 2 Add the host to the vSphere Distributed Switch mapped to the cluster where you want to add the host.

All hosts in the cluster must be in the vSphere Distributed Switch being leveraged by NSX.

- 3 Place the host into maintenance mode.

- 4 Add the host to the cluster.

Since this is a prepared cluster, the required software is automatically installed on the newly added host.

- 5 Remove the host from maintenance mode.

DRS balances virtual machines onto the host.

## Remove a Host from an NSX Prepared Cluster

This section describes how to remove a host from a cluster prepared for network virtualization. You might want to do this if, for example, you decide that the host is not going to participate in NSX.

### Procedure

- 1 Place the host into maintenance mode and wait for DRS to evacuate the host, or manually vMotion running VMs from the host.

- 2 Remove host from the prepared cluster by either moving it to an unprepared cluster or making it a standalone host outside of any cluster.

NSX uninstalls the network virtualization components and service virtual machines from the host.

- 3 For the change to take effect, move or stop all the VMs, put the host into maintenance mode, and reboot the host.

Optionally, you can defer the reboot to a maintenance window.

The NSX VIBs are removed from the host. To verify, SSH to the host and run the `esxcli software vib list | grep esx` command. Make sure the following VIBs are not present on the host:

- `esx-vsip`
- `esx-vxlan`

If the VIBs remain on the host, you might want to view the logs to determine why automated VIB removal did not work.

You can remove the VIBs manually by running the following commands:

- `esxcli software vib remove --vibname=esx-vxlan`
- `esxcli software vib remove --vibname=esx-vsip`

---

**Important** After running these commands, you must reboot the host for the change to take effect.

---

# Uninstalling NSX Components

This chapter details the steps required to uninstall NSX components from your vCenter inventory.

---

**Note** Do not remove NSX appliances from vCenter directly. Always manage and remove NSX appliances using the Networking and Security tab of the vSphere Web Client.

---

This chapter includes the following topics:

- [Uninstall an NSX Edge Services Gateway or a Distributed Logical Router](#)
- [Uninstall a Logical Switch](#)
- [Safely Remove an NSX Installation](#)

## Uninstall an NSX Edge Services Gateway or a Distributed Logical Router

You can uninstall an NSX Edge by using the vSphere Web Client.

### Prerequisites

You must have been assigned the Enterprise Administrator or NSX Administrator role.

### Procedure

- 1 Log in to the vSphere Web Client.
- 2 Click **Networking & Security** and then click **NSX Edges**.
- 3 Select an NSX Edge and click the **Delete** (✖) icon.

## Uninstall a Logical Switch

You can uninstall a logical switch by using the vSphere Web Client.

### Prerequisites

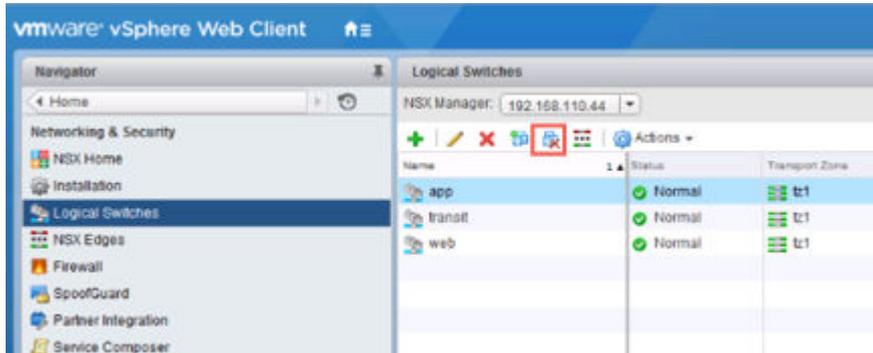
You must have been assigned the Enterprise Administrator or NSX Administrator role.

### Procedure

- 1 In the vSphere Web Client, navigate to **Home > Networking & Security > Logical Switches**.

- 2 Select a logical switch and remove any attached virtual machines by clicking the Remove Virtual Machine icon.

For example:



- 3 With the logical switch selected, click the **Delete** (✖) icon.

## Safely Remove an NSX Installation

A complete uninstall of NSX removes host VIBs, the NSX Manager, controllers, all VXLAN configuration, logical switches, logical routers, NSX firewall, and the vCenter NSX plug in. Make sure to follow the steps for all hosts in the cluster. VMware recommends that you uninstall the network virtualization components from a cluster before removing the NSX plug-in from vCenter Server.

Removing NSX completely requires two host reboots. The first reboot is required after uninstalling the NSX VIBs. The second reboot is required after removing the host VTEPs and dvPortgroups used for VTEPs.

---

**Note** Do not remove NSX appliances from vCenter directly. Always manage and remove NSX appliances using the Networking and Security tab of the vSphere Web Client.

---

If you want to remove NSX from individual hosts (instead of from the entire the entire cluster), see [Remove a Host from an NSX Prepared Cluster](#).

### Prerequisites

- You must have been assigned the Enterprise Administrator or NSX Administrator role.
- Remove any registered partner solutions, as well as endpoint services before reversing host preparation so that service VMs in the cluster are removed gracefully.
- Delete all NSX Edges. See [Uninstall an NSX Edge Services Gateway or a Distributed Logical Router](#).
- Detach virtual machines in the transport zone from the logical switches, and delete the logical switches. See [Uninstall a Logical Switch](#).

## Procedure

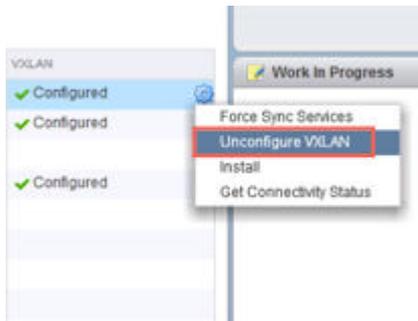
- 1 Remove the cluster from its transport zone.

Go to **Logical Network Preparation > Transport Zones** and disconnect the cluster from the transport zone.

If the cluster is grayed out and you cannot disconnect it from the transport zone, this might be because 1) a host in the cluster is disconnected or is not powered on, or 2) the cluster might contain one or more virtual machines or appliances that are attached to the transport zone. For example, if the host is in a management cluster and has NSX controllers installed on it, first remove or move the controllers.

- 2 Delete the transport zone.
- 3 Unconfigure VXLAN on the cluster.

For example:

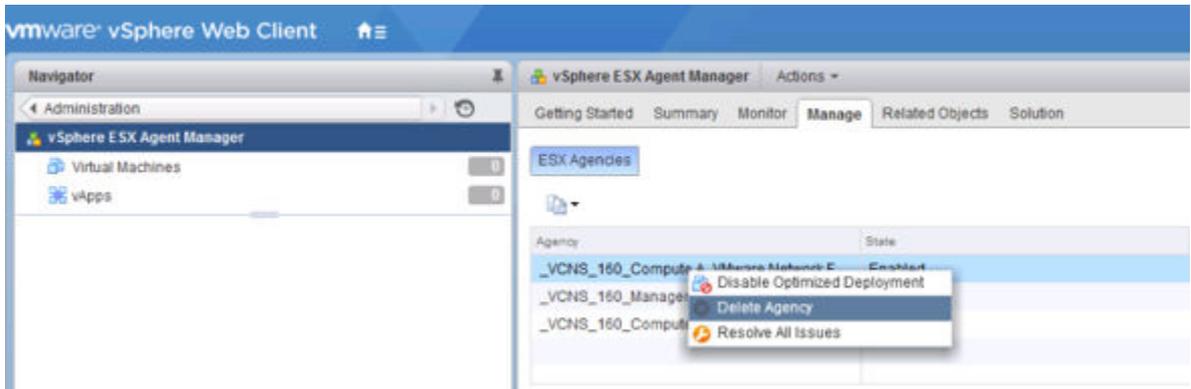


- 4 Reverse host preparation by uninstalling the NSX VIBs.

Select one of the following methods for uninstalling NSX VIBs. The first two options uninstall NSX VIBs on all hosts in a cluster. The last two options uninstall the VIBs on one host at a time.

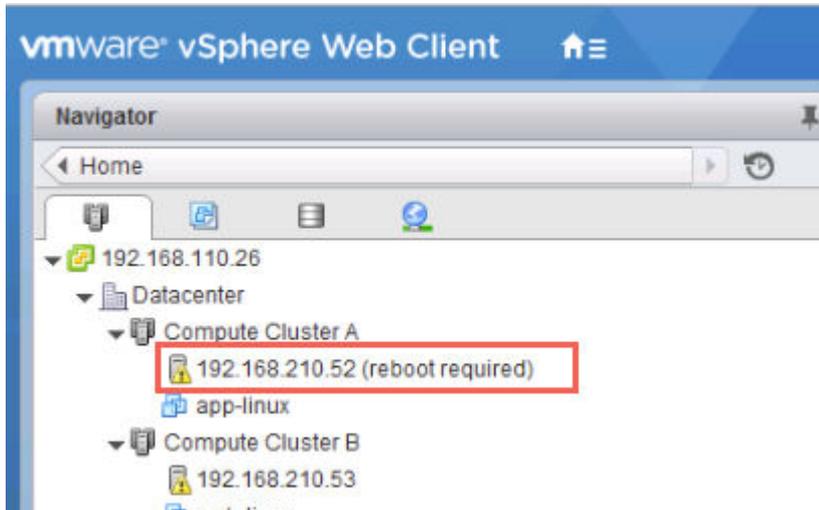
- In the vCenter Web Client, go to **Networking & Security > Installation > Host Preparation**, and click **Uninstall**.
- In the vCenter Web Client, go to **Administration > vCenter Server Extensions > vSphere ESX Agent Manager**. On the **Management** tab, right-click the VCNS agency and select **Delete Agency**.

For example:



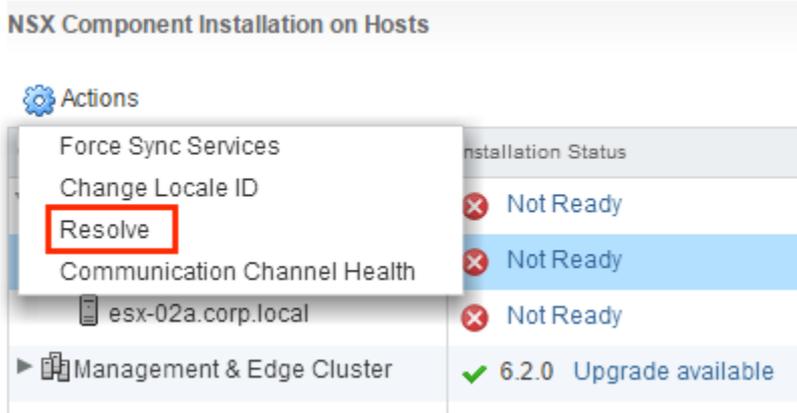
- Move the host from the prepared cluster to an unprepared cluster.
  - On the host, run the following commands:
    - `esxcli software vib remove --vibName=esx-vxlan`
    - `esxcli software vib remove --vibName=esx-vsip`
- 5 Reboot the hosts.

Removing VIBs from a host requires a host reboot. The required reboot does not happen automatically. When a host needs to be rebooted, the **(reboot required)** tag appears in the Hosts and Clusters view. For example:



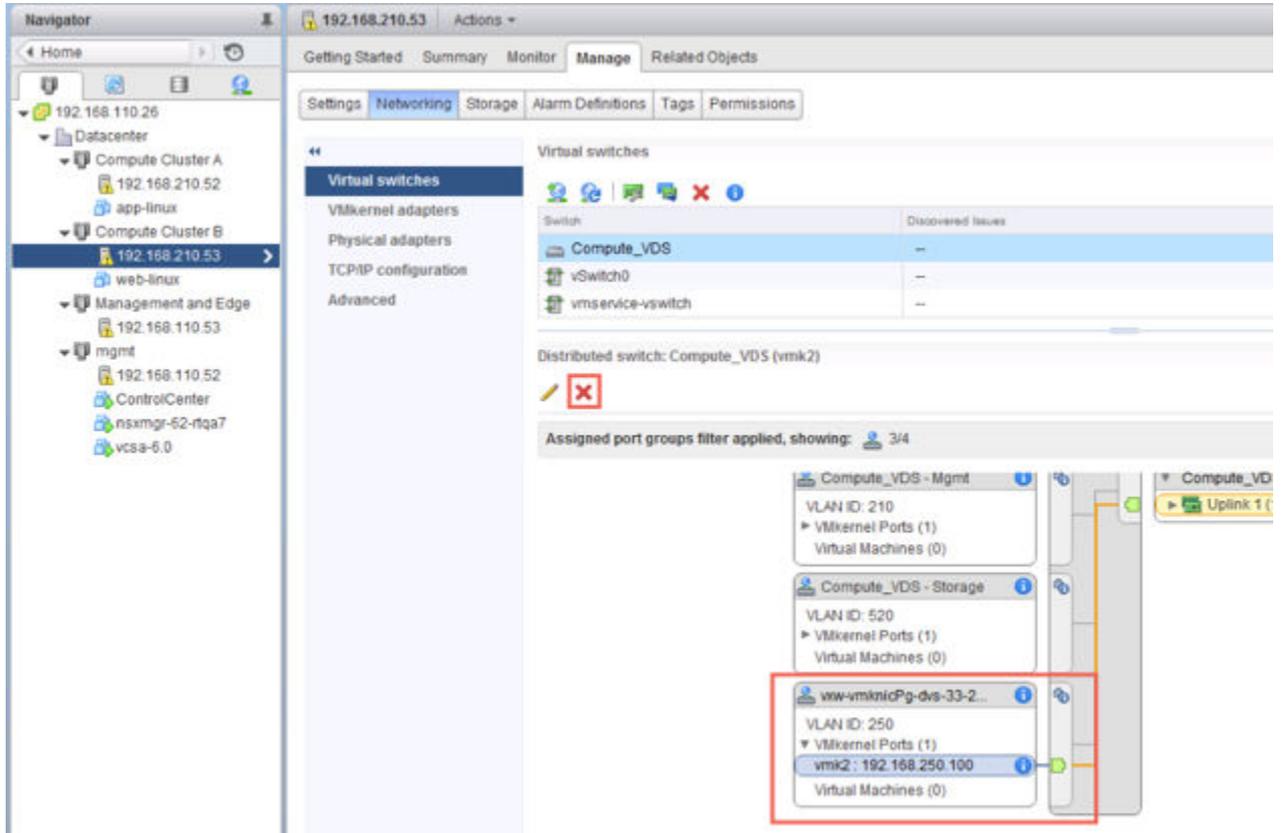
Use one of the following procedures to reboot the hosts.

- Reboot the host manually.
- Select the cluster and click the **Resolve** action. This action reboots all hosts in the cluster. If the cluster has DRS enabled, DRS will attempt to reboot the hosts in a controlled fashion that allows the VMs to continue running. If DRS fails for any reason, the **Resolve** action halts. In this case, you may need to move the VMs manually and then retry the **Resolve** action or reboot the hosts manually.



- 6 Delete the NSX Manager appliance and all NSX controller appliance VMs from the disk.
- 7 Remove any leftover VTEP vmkernel interfaces.

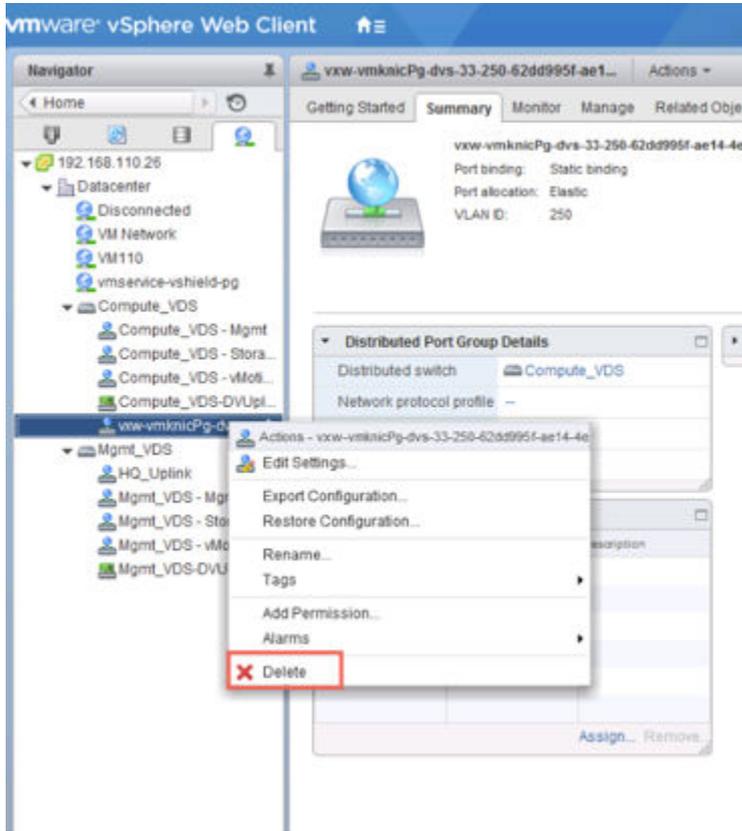
For example:



Generally, the VTEP vmkernel interfaces are already deleted as a result of earlier uninstall operations.

- 8 Remove any leftover dvPortgroups used for VTEPs.

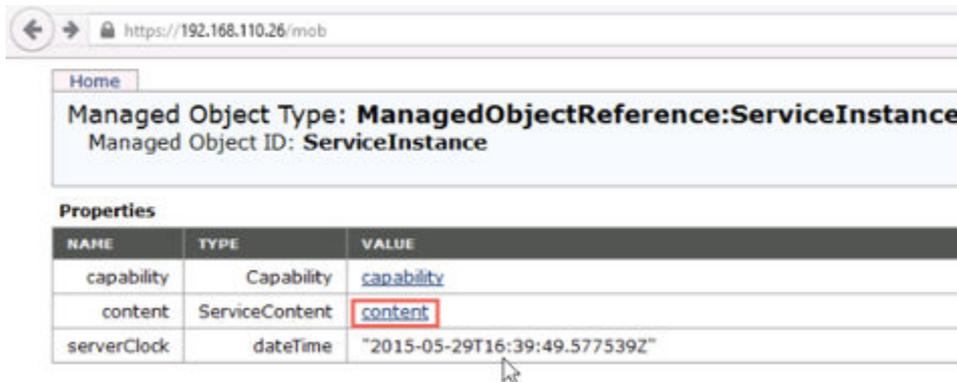
For example:



Generally, the dvPortgroups used for VTEPs are already deleted as a result of earlier uninstall operations.

- 9 Reboot the hosts.
- 10 For the vCenter on which you want to remove the NSX Manager plug-in, log in to the managed object browser at [https://your\\_vc\\_server/mob](https://your_vc_server/mob).
- 11 Click **Content**.

For example:



12 Click **ExtensionManager**.

Home **Data Object Type: ServiceContent**  
 Parent Managed Object ID: **ServiceInstance**  
 Property Path: **content**

**Properties**

| NAME                      | TYPE                                                   | VALUE                                     |
|---------------------------|--------------------------------------------------------|-------------------------------------------|
| about                     | AboutInfo                                              | <a href="#">about</a>                     |
| accountManager            | ManagedObjectReference:HostLocalAccountManager         | Unset                                     |
| alarmManager              | ManagedObjectReference:AlarmManager                    | <a href="#">AlarmManager</a>              |
| authorizationManager      | ManagedObjectReference:AuthorizationManager            | <a href="#">AuthorizationManager</a>      |
| certificateManager        | ManagedObjectReference:CertificateManager              | <a href="#">certificateManager</a>        |
| clusterProfileManager     | ManagedObjectReference:ClusterProfileManager           | <a href="#">ClusterProfileManager</a>     |
| complianceManager         | ManagedObjectReference:ProfileComplianceManager        | <a href="#">MoComplianceManager</a>       |
| customFieldsManager       | ManagedObjectReference:CustomFieldsManager             | <a href="#">CustomFieldsManager</a>       |
| customizationSpecManager  | ManagedObjectReference:CustomizationSpecManager        | <a href="#">CustomizationSpecManager</a>  |
| datastoreNamespaceManager | ManagedObjectReference:DatastoreNamespaceManager       | <a href="#">DatastoreNamespaceManager</a> |
| diagnosticManager         | ManagedObjectReference:DiagnosticManager               | <a href="#">DiagMgr</a>                   |
| dvSwitchManager           | ManagedObjectReference:DistributedVirtualSwitchManager | <a href="#">DVSwitchManager</a>           |
| eventManager              | ManagedObjectReference:EventManager                    | <a href="#">EventManager</a>              |
| extensionManager          | ManagedObjectReference:ExtensionManager                | <a href="#">ExtensionManager</a>          |
| fileManager               | ManagedObjectReference:FileManager                     | <a href="#">FileManager</a>               |
| guestOperationsManager    | ManagedObjectReference:GuestOperationsManager          | <a href="#">guestOperationsManager</a>    |
| hostProfileManager        | ManagedObjectReference:HostProfileManager              | <a href="#">HostProfileManager</a>        |

13 Click **UnregisterExtension**.

**Methods**

| RETURN TYPE                            | NAME                                            |
|----------------------------------------|-------------------------------------------------|
| Extension                              | <a href="#">FindExtension</a>                   |
| string                                 | <a href="#">GetPublicKey</a>                    |
| ExtensionManagerIpAllocationUsage[]    | <a href="#">QueryExtensionIpAllocationUsage</a> |
| ManagedObjectReference:ManagedEntity[] | <a href="#">QueryManagedBy</a>                  |
| void                                   | <a href="#">RegisterExtension</a>               |
| void                                   | <a href="#">SetExtensionCertificate</a>         |
| void                                   | <a href="#">SetPublicKey</a>                    |
| void                                   | <a href="#">UnregisterExtension</a>             |
| void                                   | <a href="#">UpdateExtension</a>                 |

- Enter the string `com.vmware.vShieldManager` and click on **Invoke Method**.

Managed Object Type:  
**ManagedObjectReference:ExtensionManager**  
 Managed Object ID: **ExtensionManager**  
 Method: **UnregisterExtension**

**void UnregisterExtension**

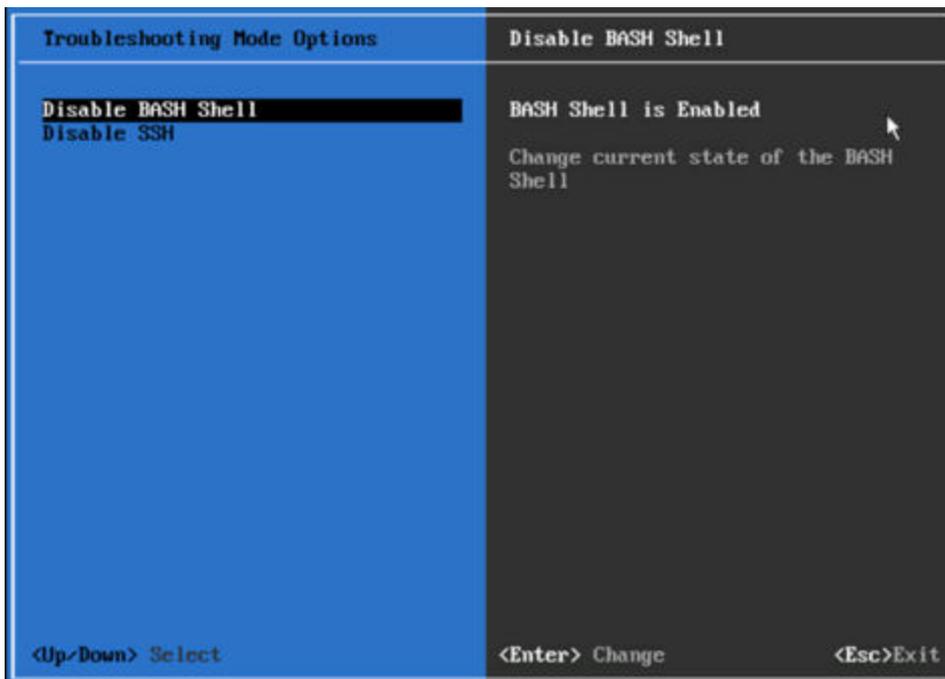
---

**Parameters**

| NAME                           | TYPE   | VALUE                                                  |
|--------------------------------|--------|--------------------------------------------------------|
| <b>extensionKey</b> (required) | string | <input type="text" value="com.vmware.vShieldManager"/> |

Invoke Method

- If you are running the vSphere 6 vCenter Appliance, launch the console and enable the BASH shell under **Troubleshooting Mode Options**.



Another way to enable the BASH shell is to log in as root and run the `shell.set - -enabled true` command.

- Delete the vSphere Web Client directories for NSX and then restart the Web Client service.  
 The vSphere Web Client directories for NSX are called `com.vmware.vShieldManager.**` and are located as follows:

- vCenter Server 5.x
  - Windows 2003 – %ALLUSERSPROFILE%\Application Data\VMware\vSphere Web Client\vc-packages\vsphere-client-serenity\

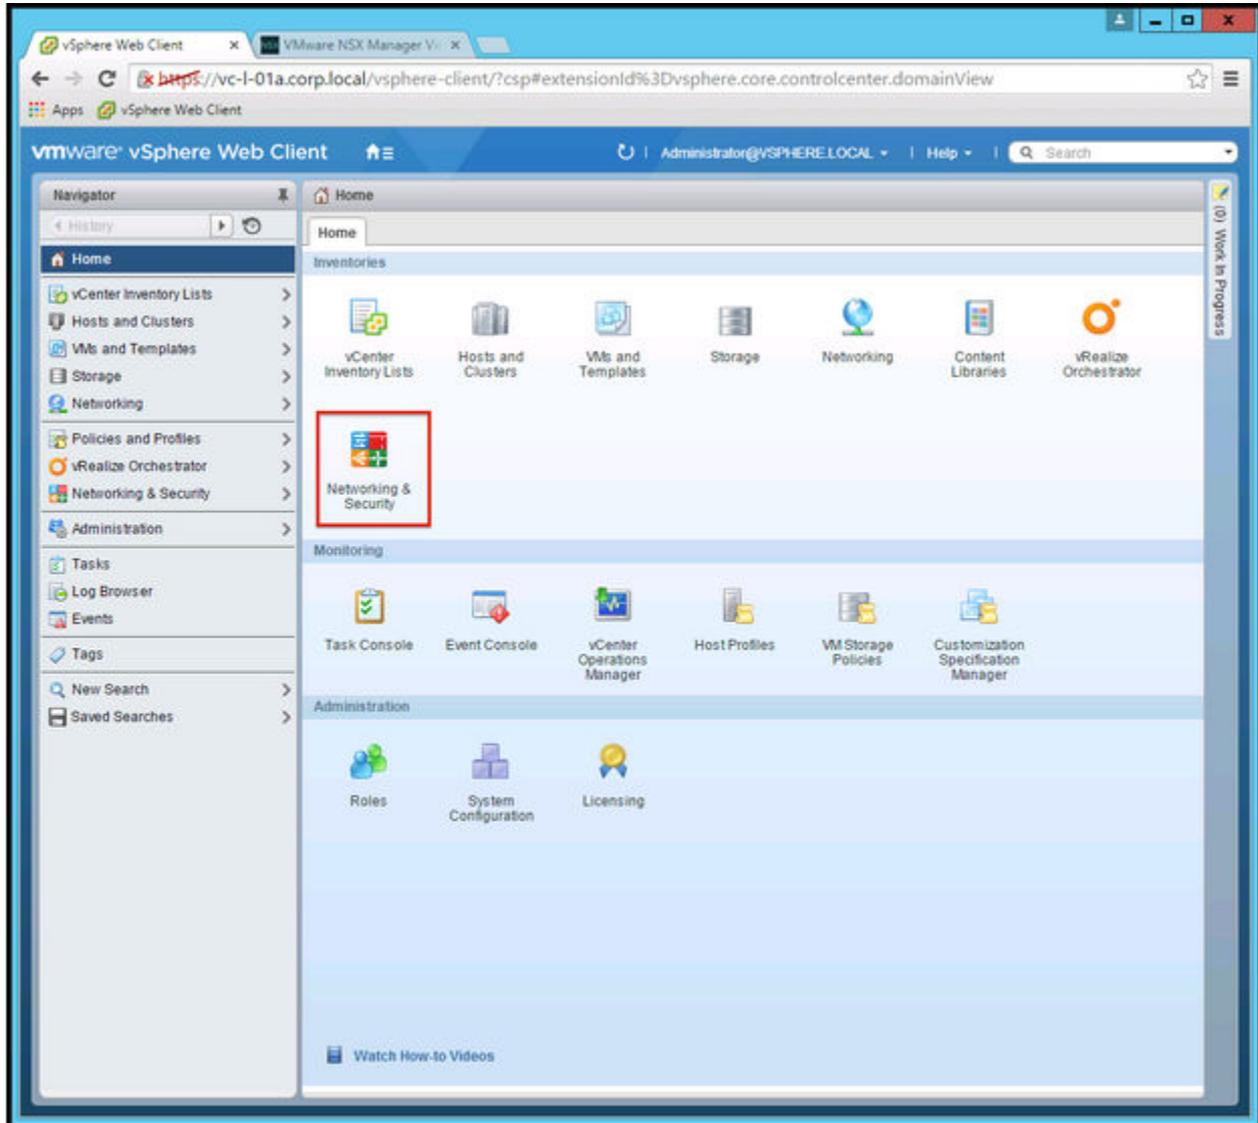
- Windows 2008/2012 – %ALLUSERSPROFILE%\VMware\vSphere Web Client\vc-packages\vsphere-client-serenity\
  - VMware vCenter Server Appliance – /var/lib/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
- vCenter Server 6.0.x
  - Windows 2008/2012 – C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity\
    - VMware vCenter Server Appliance – /etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/

For vCenter Server Appliance, run the `service vsphere-client restart` command in the appliance shell.

For Windows-based vCenter, run `services.msc`, right-click **vSphere Web Client**, and click **Start**.

The NSX Manager plug-in is removed from vCenter. To confirm, log out of vCenter and log back in.

The NSX Manager plug-in **Networking & Security** icon no longer appears on the Home screen in the vCenter Web Client.



Go to **Administration > Client Plug-Ins** and verify that the list of plug-ins does not include **NSX User Interface** plugin.

