NSX-T Data Center Installation Guide

vmware®

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

https://docs.vmware.com/

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

Contents

NSX-T Data Center Installation Guide

The *NSX-T Data Center Installation Guide* describes how to install the VMware NSX-T™ Data Center product. The information includes step-by-step configuration instructions, and suggested best practices.

Intended Audience

This information is intended for anyone who wants to install or use NSX-T Data Center. This information is written for experienced system administrators who are familiar with virtual machine technology and network virtualization concepts.

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-T Data Center

<div style="text-align: right">1</div>

In the same way that server virtualization programmatically creates and manages virtual machines, NSX-T Data Center network virtualization programmatically creates and managed software-based virtual networks.

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) 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.

NSX-T Data Center works by implementing three separate but integrated planes: management, control, and data. These planes are implemented as a set of processes, modules, and agents residing on two types of nodes: NSX Manager and transport nodes.

- Every node hosts a management plane agent.
- NSX Manager nodes host API services and the management plane cluster daemons.
- NSX Controller nodes host the central control plane cluster daemons.
- Transport nodes host local control plane daemons and forwarding engines.

NSX Manager provides a three-node clustering support which merges policy manager, management, and central control services on a cluster of nodes. NSX Manager clustering provides high availability of the user interface and API. The convergence of management and control plane nodes reduces the number of virtual appliances that must be deployed and managed by the NSX-T Data Center administrator.

The NSX Manager appliance is available in three different sizes for different deployment scenarios. A small appliance for lab or proof-of-concept deployments. A medium appliance for deployments up to 64 hosts and a large appliance for customers who deploy to a large-scale environment. See NSX Manager VM System Requirements and Configuration maximums tool.

This chapter includes the following topics:

- Key Concepts
- Overview of the NSX Manager

Key Concepts

The common NSX-T Data Center concepts that are used in the documentation and user interface.

Compute Manager

A compute manager is an application that manages resources such as hosts and VMs. One example is vCenter Server.

Control Plane

Computes runtime state based on configuration from the management plane. Control plane disseminates topology information reported by the data plane elements, and pushes stateless configuration to forwarding engines.

Data Plane

Performs stateless forwarding or transformation of packets based on tables populated by the control plane. Data plane reports topology information to the control plane and maintains packet level statistics.

External Network

A physical network or VLAN not managed by NSX-T Data Center. You can link your logical network or overlay network to an external network through an NSX Edge. For example, a physical network in a customer data center or a VLAN in a physical environment.

Fabric Node

Host that has been registered with the NSX-T Data Center management plane and has NSX-T Data Center modules installed. For a hypervisor host or NSX Edge to be part of the NSX-T Data Center overlay, it must be added to the NSX-T Data Center fabric.

Logical Port Egress

Outbound network traffic leaving the VM or logical network is called egress because traffic is leaving virtual network and entering the data center.

Logical Port Ingress

Inbound network traffic leaving the data center and entering the VM is called ingress traffic.

Logical Router

NSX-T Data Center routing entity.

Logical Router Port

Logical network port to which you can attach a logical switch port or an uplink port to a physical network.

Logical Switch

Entity that provides virtual Layer 2 switching for VM interfaces and Gateway interfaces. A logical switch gives tenant network administrators the logical equivalent of a physical Layer 2 switch, allowing them to connect a set of VMs to a common broadcast domain. A logical switch is a logical entity independent of the physical hypervisor infrastructure and spans many hypervisors, connecting VMs regardless of their physical location.

In a multi-tenant cloud, many logical switches might exist side-by-side on the same hypervisor hardware, with each Layer 2 segment isolated from the others. Logical switches can be connected using logical routers, and logical routers can provide uplink ports connected to the external physical network.

Logical Switch Port

Logical switch attachment point to establish a connection to a virtual machine network interface or a logical router interface. The logical switch port reports applied switching profile, port state, and link status.

Management Plane

Provides single API entry point to the system, persists user configuration, handles user queries, and performs operational tasks on all of the management, control, and data plane nodes in the system. Management plane is also responsible for querying, modifying, and persisting use configuration.

NSX Edge Cluster

Collection of NSX Edge node appliances that have the same settings as protocols involved in high-availability monitoring.

NSX Edge Node

Component with the functional goal is to provide computational power to deliver the IP routing and the IP services functions.

NSX Managed Virtual Distributed Switch or KVM Open vSwitch

The NSX managed virtual distributed switch (N-VDS, previously known as hostswitch)or OVS is used for shared NSX Edge and compute cluster. N-VDS is required for overlay traffic configuration.

An N-VDS has two modes: standard and enhanced datapath. An enhanced datapath N-VDS has the performance capabilities to support NFV (Network Functions Virtualization) workloads.

NSX Manager

Node that hosts the API services, the management plane, and the agent services. NSX Manager is an appliance included in the NSX-T Data Center installation package. You can deploy the appliance in the role of nsx-manager nsx-controller or nsx-cloud-service-manager. Currently, the appliance only supports one role at a time.

NSX Manager Cluster

A cluster of NSX Managers that can provide high availability.

Open vSwitch (OVS)

Open source software switch that acts as a virtual switch within XenServer, Xen, KVM, and other Linux-based hypervisors.

Overlay Logical Network

Logical network implemented using Layer 2-in-Layer 3 tunneling such that the topology seen by VMs is decoupled from that of the physical network.

Physical Interface (pNIC)

Network interface on a physical server that a hypervisor is installed on.

Segment

Entity that provides virtual Layer 2 switching for VM interfaces and Gateway interfaces. A segment gives tenant network administrators the logical equivalent of a physical Layer 2 switch, allowing them to connect a set of VMs to a common broadcast domain. A segment is a logical entity independent of the physical hypervisor infrastructure and spans many hypervisors, connecting VMs regardless of their physical location. A segment is also known as a logical switch.

In a multi-tenant cloud, many segments might exist side-by-side on the same hypervisor hardware, with each Layer 2 segment isolated from the others. Segments can be connected using gateways, which can provide connectivity to the external physical network.

Tier-0 Gateway or Tier-0 Logical Router

The Tier-0 Gateway is called the Tier-0 Logical Router in the **Advanced Networking & Security** tab. It interfaces with the physical network and can be realized as an active-active or active-standby cluster. The Tier-0 gateway runs BGP and peers with physical routers. In active-standby mode the gateway can also provide stateful services.

Tier-1 Gateway or Tier-1 Logical Router

The Tier-1 Gateway is called the Tier-1 Logical Router in the **Advanced Networking & Security** tab. It connects to one Tier-0 gateway for northbound connectivity and one or more overlay networks for southbound connectivity. A Tier-1 gateway can be an active-standby cluster that provides stateful services.

Transport Zone

Collection of transport nodes that defines the maximum span for logical switches. A transport zone represents a set of similarly provisioned hypervisors and the logical switches that connect VMs on those hypervisors.

Transport Node

A node capable of participating in an NSX-T Data Center overlay or NSX-T Data Center VLAN networking. For a KVM host, you can preconfigure the N-VDS, or you can have NSX Manager perform the configuration. For an ESXi host, NSX Manager always configures the N-VDS.

Uplink Profile

Defines policies for the links from hypervisor hosts to NSX-T Data Center logical switches or from NSX Edge nodes to top-of-rack switches. The settings defined by uplink profiles might include teaming policies, active/standby links, the transport VLAN ID, and the MTU setting. The transport VLAN set in the uplink profile tags overlay traffic only and the VLAN ID is used by the TEP endpoint.

VM Interface (vNIC)

Network interface on a virtual machine that provides connectivity between the virtual guest operating system and the standard vSwitch or vSphere distributed switch. The vNIC can be attached to a logical port. You can identify a vNIC based on its Unique ID (UUID).

Virtual Tunnel Endpoint

Each hypervisor has a Virtual Tunnel Endpoint (VTEP) responsible for encapsulating the VM traffic inside a VLAN header and routing the packet to a destination VTEP for further processing. Traffic can be routed to another VTEP on a different host or the NSX Edge gateway to access the physical network.

Overview of the NSX Manager

The NSX Manager provides a web-based user interface where you can manage the NSX-T environment. It also hosts the API server that processes API calls.

The NSX Manager web interface provides two methods of configuring resources.

- The Policy interface: the **Networking**, **Security**, **Inventory**, and **Plan & Troubleshoot** tabs.

- The Advanced interface: the **Advanced Networking & Security** tab.

When to Use Policy or Advanced Interfaces

Be consistent about which user interface you use. There are a few reasons to use one user interface over another.

- If you are deploying a new environment with NSX-T Data Center 2.4 or later, using the new policy-based user interface to create and manage your environment is the best choice in most situations.

  - Some features are not available in the policy-based user interface. If you need these features, use the Advanced user interface for all configurations.

- If you are upgrading to NSX-T Data Center 2.4 or later, continue to make configuration changes using the **Advanced Networking & Security** user interface.

Table 1-1. When to Use Policy or Advanced Interfaces

Policy Interface	Advanced Interface
Most new deployments should use the policy-based interface.	Deployments which were created using the advanced interface, for example, upgrades from versions before the policy-based interface was present.
NSX Cloud deployments	Deployments which integrate with other plugins. For example, NSX Container Plug-in, Openstack, and other cloud management platforms.
Networking features available in the Policy interface only: ■ DNS Services and DNS Zones ■ VPN ■ Forwarding policies for NSX Cloud	Networking features available in the Advanced interface only: ■ Layer 3 forwarding for IPv4 and IPv6 ■ Forwarding up timer ■ Change internal transit network IP ■ VIP HA support on Tier-0 ■ Standby relocation ■ Route advertisement filtering based on list of prefixes on Tier-1 ■ Loopback creation ■ BGP multihop ■ BGP source addresses ■ Static routes with BFD and interface as next-hop ■ Metadata proxy ■ DHCP server attached to an isolated segment and static binding

Security features available in the Policy interface only:
■ Endpoint Protection
■ Network Introspection (East-West Service Insertion)
■ Context Profiles
  ■ L7 applications
  ■ FQDN
■ New Distributed Firewall and Gateway Firewall Layout
  ■ Categories
  ■ Auto service rules
	Security features available in the Advanced interface only:
	■ Ability to enable or disable Distributed Firewall, Identity Firewall, and Gateway Firewall
	■ Distributed Firewall session timers
	■ Exclusion lists
	■ CPU and memory thresholds
	■ Sections for stateless rules
	■ Bridge Firewall
	■ Section Locking
	■ Distributed Firewall rule IDs
	■ Distributed Firewall rules based on IPs in source and destination

Using the Policy Interface

If you decide to use the policy interface, use it to create all objects. Do not use the advanced interface to create objects.

You can use the advanced interface to modify objects that have been created in the policy interface. The settings for a policy-created object might include a link for **Advanced Configuration**. This link takes you to the advanced interface where you can fine-tune the configuration. You can also view policy-created objects in the advanced interface directly. Settings that are managed by policy but are visible in the advanced interface have this icon next to them: ⊖. You cannot modify them from the advanced user interface.

Where to Find the Policy Interfaces and Advanced Interfaces

The policy-based and advanced interfaces appear in different parts of the NSX Manager user interface, and use different API URIs.

Table 1-2. Policy Interfaces and Advanced Interfaces

Policy Interface	Advanced Interface
■ **Networking** tab ■ **Security** tab ■ **Inventory** tab ■ **Plan & Troubleshoot** tab	**Advanced Networking & Security** tab
API URIs that begin with `/policy/api`	API URIs that begin with `/api`

Note The **System** tab is used for all environments. If you modify Edge nodes, Edge clusters, or transport zones, it can take up to 5 minutes for those changes to be visible on the policy-based user interface. You can synchronize immediately using `POST /policy/api/v1/infra/sites/default/enforcement-points/default?action=reload`.

For more information about using the policy API, see the NSX-T Policy API Getting Started Guide.

Names for Objects Created in the Policy and Advanced Interfaces

The objects you create have different names depending on which interface was used to create them.

Table 1-3. Object Names

Objects Created Using the Policy Interface	Objects Created Using the Advanced Interface
Segment	Logical switch
Tier-1 gateway	Tier-1 logical router
Tier-0 gateway	Tier-0 logical router
Group	NSGroup, IP Sets, MAC Sets
Security Policy	Firewall section
Rule	Firewall rule
Gateway firewall	Edge firewall

NSX-T Data Center Installation Workflows

2

You can install NSX-T Data Center on vSphere or KVM hosts. You can also configure a bare metal server to use NSX-T Data Center.

To install or configure any of the hypervisors or bare metal, follow the recommended tasks in the workflows.

This chapter includes the following topics:

- NSX-T Data Center Workflow for vSphere
- NSX-T Data Center Installation Workflow for KVM
- NSX-T Data Center Configuration Workflow for Bare Metal Server

NSX-T Data Center Workflow for vSphere

Use the checklist to track your installation progress on a vSphere host.

Follow the recommended order of procedures.

1 Review the NSX Manager installation requirements. See NSX Manager Installation.

2 Configure the necessary ports and protocols. See Ports and Protocols.

3 Install the NSX Manager. See Install NSX Manager and Available Appliances.

4 Log in to the newly created NSX Manager. See Log In to the Newly Created NSX Manager .

5 Configure a compute manger. See Add a Compute Manager.

6 Deploy additional NSX Manager nodes to form a cluster. See Deploy NSX Manager Nodes to Form a Cluster from UI.

7 Review the NSX Edge installation requirements. See NSX Edge Installation.

8 Install NSX Edges. See Install an NSX Edge on ESXi Using a vSphere GUI.

9 Create an NSX Edge cluster. See Create an NSX Edge Cluster.

10 Create transport zones. See Create Transport Zones.

11 Create host transport nodes. See Create a Standalone Host or Bare Metal Server Transport Node or Configure a Managed Host Transport Node.

A virtual switch is created on each host. The management plane sends the host certificates to the control plane, and the management plane pushes control plane information to the hosts. Each host connects to the control plane over SSL presenting its certificate. The control plane validates the certificate against the host certificate provided by the management plane. The controllers accept the connection upon successful validation.

Post-Installation

When the hosts are transport nodes, you can create transport zones, logical switches, logical routers, and other network components through the NSX Manager UI or API at any time. When NSX Edges and hosts join the management plane, the NSX-T Data Center logical entities and configuration state are pushed to the NSX Edges and hosts automatically.

For more information, see the *NSX-T Data Center Administration Guide*.

NSX-T Data Center Installation Workflow for KVM

Use the checklist to track your installation progress on a KVM host.

Follow the recommended order of procedures.

1   Prepare your KVM environment. See Set Up KVM.

2   Review the NSX Manager installation requirements. See NSX Manager Installation.

3   Configure the necessary ports and protocols. See Ports and Protocols.

4   Install the NSX Manager. See Install NSX Manager on KVM.

5   Log in to the newly created NSX Manager. See Log In to the Newly Created NSX Manager .

6   Configure third-party packages on the KVM host. See Install Third-Party Packages on a KVM Host.

7   Deploy additional NSX Manager nodes to form a cluster. See Deploy NSX Manager Nodes to Form a Cluster Using CLI .

8   Review the NSX Edge installation requirements. See NSX Edge Installation.

9   Install NSX Edges. See Install NSX Edge on Bare Metal .

10 Create an NSX Edge cluster. See Create an NSX Edge Cluster.

11 Create transport zones. See Create Transport Zones.

12 Create host transport nodes. See Create a Standalone Host or Bare Metal Server Transport Node.

A virtual switch is created on each host. The management plane sends the host certificates to the control plane, and the management plane pushes control plane information to the hosts. Each host connects to the control plane over SSL presenting its certificate. The control plane validates the certificate against the host certificate provided by the management plane. The controllers accept the connection upon successful validation.

Post-Installation

When the hosts are transport nodes, you can create transport zones, logical switches, logical routers, and other network components through the NSX Manager UI or API at any time. When NSX Edges and hosts join the management plane, the NSX-T Data Center logical entities and configuration state are pushed to the NSX Edges and hosts automatically.

For more information, see the *NSX-T Data Center Administration Guide*.

NSX-T Data Center Configuration Workflow for Bare Metal Server

Use the checklist to track your progress when configuring bare metal server to use NSX-T Data Center.

Follow the recommended order of procedures.

1   Review the bare metal requirements. See Bare Metal Server System Requirements.

2   Configure the necessary ports and protocols. See Ports and Protocols.

3   Install the NSX Manager. See Install NSX Manager on KVM.

4   Configure third-party packages on the bare metal server. See Install Third-Party Packages on a Bare Metal Server.

5   Create host transport nodes. See Create a Standalone Host or Bare Metal Server Transport Node.

A virtual switch is created on each host. The management plane sends the host certificates to the control plane, and the management plane pushes control plane information to the hosts. Each host connects to the control plane over SSL presenting its certificate. The control plane validates the certificate against the host certificate provided by the management plane. The controllers accept the connection upon successful validation.

6   Create an application interface for bare metal server workload. See Create Application Interface for Bare Metal Server Workloads.

Preparing for Installation

3

Before installing NSX-T Data Center, make sure your environment is prepared.

This chapter includes the following topics:

- System Requirements
- Ports and Protocols
- Installing NSX-T Data Center Components

System Requirements

Before you install NSX-T Data Center, your environment must meet specific hardware and resource requirements.

NSX Manager VM System Requirements

Before you install an NSX Manager, make sure that your environment meets the supported requirements.

Hypervisor Host Requirements for Transport Nodes

Hypervisor	Version	CPU Cores	Memory
vSphere	Supported vSphere version	4	16 GB
CentOS Linux KVM	7.4	4	16 GB
Red Hat Enterprise Linux (RHEL) KVM	7.6, 7.5, and 7.4	4	16 GB
SUSE Linux Enterprise Server KVM	12 SP3, SP4	4	16 GB
Ubuntu KVM	18.04 and 16.04.2 LTS	4	16 GB

Table 3-1. Supported Hosts for NSX Managers

Support Description	Hypervisor
ESXi	For supported hosts, see the VMware Product Interoperability Matrices.
KVM	RHEL 7.4 and Ubuntu 16.04 LTS

For ESXi hosts, NSX-T Data Center supports the Host Profiles and Auto Deploy features on vSphere 6.7 U1 or later. See *Understanding vSphere Auto Deploy* in the *VMware ESXi Installation and Setup* documentation for more information.

Caution On RHEL, the `yum update` command might update the kernel version and break the compatibility with NSX-T Data Center. Disable the automatic kernel update when you run `yum update`. Also, after running `yum install`, verify that NSX-T Data Center supports the kernel version.

Hypervisor Host Network Requirements

It is required that hypervisor hosts running NSX-T Data Center have a compatible NIC card. For supported NIC cards, see the VMware Compatibility Guide.

Tip To quickly identify compatible cards in the Compatibility Guide, apply the following criteria:

- Under **I/O Device Type**, select `Network`.

- Optionally, to use supported GENEVE encapsulation, under **Features**, select the GENEVE options.

- Optionally, to use Enhanced Data Path, select `N-VDS Enhanced Data Path`.

Enhanced Data Path NIC Drivers

Download the supported NIC drivers from the My VMware page.

NIC Card	NIC Driver
Intel 82599	ixgben 1.1.0.26-1OEM.670.0.0.7535516
Intel(R) Ethernet Controller X710 for 10GbE SFP+ Intel(R) Ethernet Controller XL710 for 40GbE QSFP+	i40en 1.2.0.0-1OEM.670.0.0.8169922

NSX Manager VM Resource Requirements

Thin virtual disk size is 3.8 GB and thick virtual disk size is 200 GB.

Appliance Size	Memory	vCPU	Disk Space	VM Hardware Version
NSX Manager Extra Small	8 GB	2	200 GB	10 or later
NSX Manager Small VM	16 GB	4	200 GB	10 or later

Appliance Size	Memory	vCPU	Disk Space	VM Hardware Version
NSX Manager Medium VM	24 GB	6	200 GB	10 or later
NSX Manager Large VM	48 GB	12	200 GB	10 or later

Note As of NSX-T 2.4, the NSX Manager provides multiple roles which previously required separate appliances. This includes the policy role, the management plane role, and the central control plane role. The central control plane role was previously provided by the NSX Controller appliance.

- The NSX Manager Extra Small VM resource requirements apply only to the Cloud Service Manager.

- The NSX Manager Small VM appliance size is suitable for lab and proof-of-concept deployments, and must not be used in production.

- The NSX Manager Medium VM appliance size is suitable for typical production environments and can support up to 64 hypervisors.

- The NSX Manager Large VM appliance size is suitable for large-scale deployments with more than 64 hypervisors.

For maximum scale using the NSX Manager Large VM appliance size, go to the VMware Configuration Maximums tool at https://configmax.vmware.com/guest and select NSX-T Data Center from the product list.

NSX Manager Browser Support

The following browsers are recommended for working with NSX Manager.

Browser	Windows 10	Mac OS X 10.13, 10.14	Ubuntu 18.04
Google Chrome 76	Yes	Yes	Yes
Mozilla Firefox 68	Yes	Yes	Yes
Microsoft Edge 44	Yes		
Apple Safari 12		Yes	

Note

- Internet Explorer is not supported.

- Supported Browser minimum resolution is 1280 x 800 px.

- Language support: NSX Manager has been localized into multiple languages: English, German, French, Japanese, Simplified Chinese, Korean, Traditional Chinese, and Spanish. However, because NSX Manager localization uses the browser language settings, ensure that your settings match the desired language. There is no language preference setting within the NSX Manager interface itself.

Network Latency Requirements

The maximum network latency between NSX Managers in a NSX Manager cluster is 10ms.

The maximum network latency between NSX Managers and Transport Nodes is 150ms.

Storage Requirements

- The maximum disk access latency is under 10ms.

- It is recommended that NSX Managers be placed on shared storage.

- Storage should be highly available to avoid a storage outage causing all NSX Manager file systems to be placed into read-only mode upon event of a storage failure.

  Consult the documentation for your storage technology on how to best design a highly available storage solution.

NSX Edge VM System Requirements

Before you install NSX Edge, make sure that your environment meets the supported requirements.

NSX Edge nodes are supported only on ESXi-based hosts with Intel-based chipsets. Otherwise, vSphere EVC mode may prevent NSX Edge nodes from starting, showing an error message in the console.

Note Only VMXNET 3 vNIC is supported for the NSX Edge VM.

NSX Cloud Note If using NSX Cloud, note that the NSX Public Cloud Gateway(PCG) is deployed in a single default size for each supported public cloud. See Deploy or Link NSX Public Cloud Gateways for details.

NSX Edge VM Resource Requirements

Appliance Size	Memory	vCPU	Disk Space	VM Hardware Version
NSX Edge Small	4 GB	2	200 GB	11 or later (vSphere 6.0 or later)
NSX Edge Medium	8 GB	4	200 GB	11 or later (vSphere 6.0 or later)
NSX Edge Large	32 GB	8	200 GB	11 or later (vSphere 6.0 or later)

Note

- The NSX Edge Small VM appliance size is suitable for lab and proof-of-concept deployments. L7 rules are not realized on a Tier-1 gateway if you deploy a small sized NSX Edge VM.

- The NSX Edge Medium appliance size is suitable for a typical production environments.

- The NSX Edge Large appliance size is suitable for environments with load balancing. See Scaling Load Balancer Resources in the *NSX-T Data Center Administration Guide*.

NSX Edge VM CPU Requirements

For the DPDK support, the underlaying platform needs to meet the following requirements:

■ CPU must have AESNI capability.

■ CPU must have 1 GB Huge Page support.

Hardware	Type
CPU	■ Intel Xeon E7-xxxx (Westmere-EX and later CPU generation) ■ Intel Xeon 56xx (Westmere-EP) ■ Intel Xeon E5-xxxx (Sandy Bridge and later CPU generation) ■ Intel Xeon Platinum (all generations) ■ Intel Xeon Gold (all generations) ■ Intel Xeon Silver (all generations) ■ Intel Xeon Bronze (all generations)

NSX Edge Bare Metal Requirements

Before you configure the NSX Edge bare metal, make sure that your environment meets the supported requirements.

NSX Edge nodes are supported only on ESXi-based hosts with Intel-based chipsets. Otherwise, vSphere EVC mode may prevent Edge nodes from starting, showing an error message in the console.

NSX Edge Bare Metal Memory, CPU, and Disk Requirements

Memory	CPU Cores	Disk Space
32 GB	8	200 GB

NSX Edge Bare Metal DPDK CPU Requirements

For the DPDK support, the underlaying platform needs to meet the following requirements:

■ CPU must have AES-NI capability.

■ CPU must have 1 GB Huge Page support.

Hardware	Type
CPU	■ Intel Xeon E7-xxxx (Westmere-EX and later CPU generation) ■ Intel Xeon 56xx (Westmere-EP) ■ Intel Xeon E5-xxxx (Sandy Bridge and later CPU generation) ■ Intel Xeon Platinum (all generations) ■ Intel Xeon Gold (all generations) ■ Intel Xeon Silver (all generations) ■ Intel Xeon Bronze (all generations)

NSX Edge Bare Metal Hardware Requirements

Verify that the bare metal NSX Edge hardware is listed in this URL https://certification.ubuntu.com/server/models/?release=18.04%20LTS&category=Server. If the hardware is not listed, the storage, video adapter, or motherboard components might not work on the NSX Edge appliance.

NSX Edge Bare Metal NIC Requirements

NIC Type	Description	PCI Device ID
Intel XXV710	I40E_DEV_ID_25G_B	0x158A
	I40E_DEV_ID_25G_SFP28	0x158B
Intel X520/Intel 82599	IXGBE_DEV_ID_82599_KX4	0x10F7
	IXGBE_DEV_ID_82599_KX4_MEZZ	0x1514
	IXGBE_DEV_ID_82599_KR	0x1517
	IXGBE_DEV_ID_82599_COMBO_BACKPLANE	0x10F8
		0x000C
	IXGBE_SUBDEV_ID_82599_KX4_KR_MEZZ	0x10F9
		0x10FB
	IXGBE_DEV_ID_82599_CX4	0x11A9
	IXGBE_DEV_ID_82599_SFP	0x1F72
	IXGBE_SUBDEV_ID_82599_SFP	0x17D0
	IXGBE_SUBDEV_ID_82599_RNDC	0x0470
	IXGBE_SUBDEV_ID_82599_560FLR	0x1507
	IXGBE_SUBDEV_ID_82599_ECNA_DP	0x154D
		0x154A
	IXGBE_DEV_ID_82599_SFP_EM	0x1558
	IXGBE_DEV_ID_82599_SFP_SF2	0x1557
	IXGBE_DEV_ID_82599_SFP_SF_QP	0x10FC
	IXGBE_DEV_ID_82599_QSFP_SF_QP	0x151C
	IXGBE_DEV_ID_82599EN_SFP	
	IXGBE_DEV_ID_82599_XAUI_LOM	
	IXGBE_DEV_ID_82599_T3_LOM	
Intel X540	IXGBE_DEV_ID_X540T	0x1528
	IXGBE_DEV_ID_X540T1	0x1560
Intel X550	IXGBE_DEV_ID_X550T	0x1563
	IXGBE_DEV_ID_X550T1	0x15D1
Intel X710	I40E_DEV_ID_SFP_X710	0x1572
	I40E_DEV_ID_KX_C	0x1581
	I40E_DEV_ID_10G_BASE_T	0x1586
Intel XL710	I40E_DEV_ID_KX_B	0x1580
	I40E_DEV_ID_QSFP_A	0x1583
	I40E_DEV_ID_QSFP_B	0x1584
	I40E_DEV_ID_QSFP_C	0x1585
Cisco VIC 1300 series	Cisco UCS Virtual Interface Card 1300	0x0043

Bare Metal Server System Requirements

Before you configure the bare metal server, make sure that your server meets the supported requirements.

Important   The user performing the installation may require `sudo` command permissions for some of the procedures. See Install Third-Party Packages on a Bare Metal Server.

Bare Metal Server Requirements

Operating System	Version	CPU Cores	Memory
CentOS Linux	7.4	4	16 GB
Red Hat Enterprise Linux (RHEL)	7.5 and 7.4	4	16 GB
SUSE Linux Enterprise Server	12 SP3	4	16 GB
Ubuntu	18.04 and 16.04.2 LTS	4	16 GB

Bare Metal Linux Container Requirements

For bare metal Linux container requirements, see the *NSX Container Plug-in for OpenShift - Installation and Administration Guide.*

Ports and Protocols

Ports and protocols allow node-to-node communication paths in NSX-T Data Center, the paths are secured and authenticated, and a storage location for the credentials are used to establish mutual authentication.

Note   The required ports and protocols must be open on both the physical and host hypervisor firewalls.

Figure 3-1. NSX-T Data Center Ports and Protocols
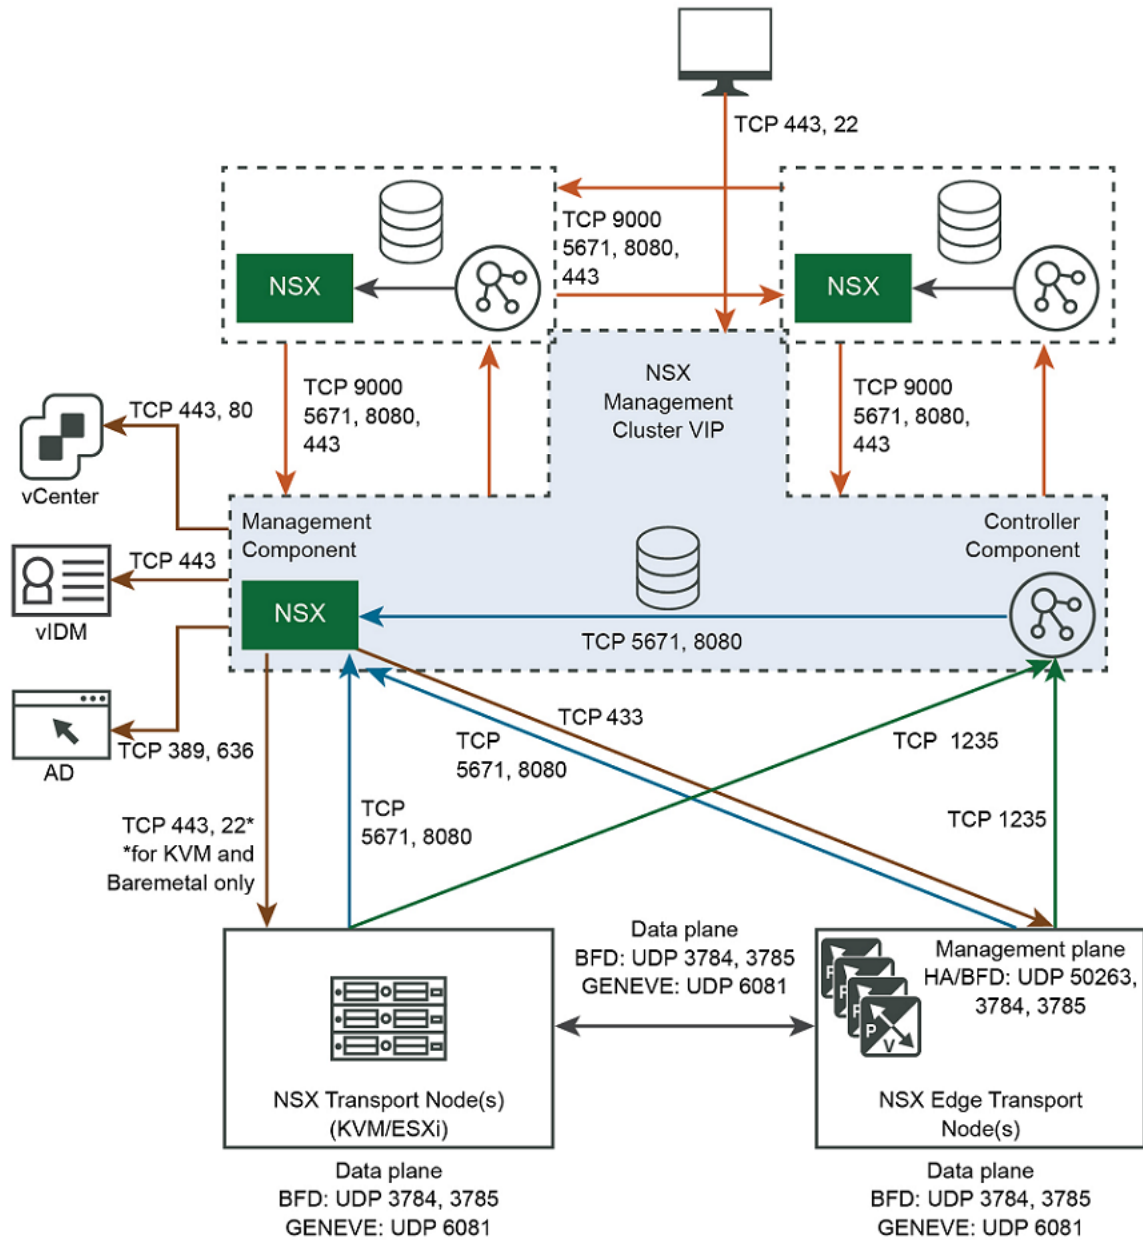
By default, all certificates are self-signed certificates. The northbound GUI and API certificates and private keys can be replaced by CA signed certificates.

There are internal daemons that communicate over the loopback or UNIX domain sockets:

■   KVM: MPA, netcpa, nsx-agent, OVS

- ESXi: netcpa, ESX-DP (in the kernel)

Note To get access to NSX-T Data Center nodes, you must enable SSH on these nodes.

NSX Cloud Note See Enable Access to ports and protocols on CSM for Hybrid Connectivity for a list of ports required for deploying NSX Cloud.

TCP and UDP Ports Used by NSX Manager

NSX Manager uses certain TCP and UDP ports to communicate with other components and products. These ports must be open in the firewall.

You can use an API call or CLI command to specify custom ports for transferring files (22 is the default) and for exporting Syslog data (514 and 6514 are the defaults). If you do, you will need to configure the firewall accordingly.

Table 3-2. TCP and UDP Ports Used by NSX Manager

Source	Target	Port	Protocol	Description
NSX Manager	Active Directory	389	TCP	Active Directory
NSX Controllers, NSX Edge nodes, Transport Nodes	NSX Manager	5671	TCP	NSX messaging
NSX Controllers, NSX Edge nodes, Transport Nodes, vCenter Server	NSX Manager	8080	TCP	Install-upgrade HTTP repository
NSX Manager	NSX Manager	9000	TCP	Internal datastore access
NSX Manager	DNS Servers	53	TCP	DNS
NSX Manager	DNS Servers	53	UDP	DNS
NSX Manager	NSX Edge	443	TCP	HTTPS
NSX Manager	Management SCP Servers	22	TCP	SSH (upload support bundle, backups, and so on)
NSX Manager	NTP Servers	123	UDP	NTP
NSX Manager	SNMP Servers	161, 162	TCP	SNMP
NSX Manager	SNMP Servers	161, 162	UDP	SNMP
NSX Manager	Syslog Servers	514	TCP	Syslog
NSX Manager	Syslog Servers	514	UDP	Syslog
NSX Manager	Syslog Servers	6514	TCP	Syslog
NSX Manager	Syslog Servers	6514	UDP	Syslog

Table 3-2. TCP and UDP Ports Used by NSX Manager (continued)

Source	Target	Port	Protocol	Description
NSX Manager	Traceroute Destination	33434 - 33523	UDP	Traceroute
NSX Manager	vCenter Server	80	TCP	NSX Manager to compute manager (vCenter Server) communication, when configured.
NSX Manager	vCenter Server	443	TCP	NSX Manager to compute manager (vCenter Server) communication, when configured.
NSX Manager	vIDM	443	TCP	vIDM
NSX Manager	NSX Manager	443	TCP	NSX Manager to NSX Manager communication
Management Clients	NSX Manager	22	TCP	SSH (Disabled by default)
Management Clients	NSX Manager	443	TCP	NSX API server
SNMP Servers	NSX Manager	161	UDP	SNMP

TCP and UDP Ports Used by NSX Edge

NSX Edge uses certain TCP and UDP ports to communicate with other components and products. These ports must be open in the firewall.

You can use an API call or CLI command to specify custom ports for transferring files (22 is the default) and for exporting Syslog data (514 and 6514 are the defaults). If you do, you will need to configure the firewall accordingly.

Table 3-3. TCP and UDP Ports Used by NSX Edge

Source	Target	Port	Protocol	Description
Management Clients	NSX Edge nodes	22	TCP	SSH (Disabled by default)
NSX Agent	NSX Edge nodes	5555	TCP	NSX Cloud - Agent on instance communicates to NSX Cloud Gateway.
NSX Edge nodes	DNS Servers	53	UDP	DNS
NSX Edge nodes	Management SCP or SSH Servers	22	TCP	SSH (upload support bundle, backups, and so on)
NSX Edge nodes	NSX Controller nodes	1235	TCP	netcpa
NSX Edge nodes	NSX Edge nodes	1167	TCP	DHCP backend

Table 3-3. TCP and UDP Ports Used by NSX Edge (continued)

Source	Target	Port	Protocol	Description
NSX Edge nodes	NSX Edge nodes	2480	TCP	Nestdb
NSX Edge nodes	NSX Edge nodes	6666	TCP	NSX Cloud - NSX Edge local communication.
NSX Edge nodes	NSX Edge nodes	50263	UDP	High-Availability
NSX Edge nodes	NSX Manager node	443	TCP	HTTPS
NSX Edge nodes	NSX Manager node	5671	TCP	NSX messaging
NSX Edge nodes	NSX Manager node	8080	TCP	NAPI, NSX-T Data Center upgrade
NSX Edge nodes	NTP Servers	123	UDP	NTP
NSX Edge nodes	OpenStack Nova API Server	3000 - 9000	TCP	Metadata proxy
NSX Edge nodes	SNMP Servers	161, 162	TCP	SNMP
NSX Edge nodes	SNMP Servers	161, 162	UDP	SNMP
NSX Edge nodes	Syslog Servers	514	TCP	Syslog
NSX Edge nodes	Syslog Servers	514	UDP	Syslog
NSX Edge nodes	Syslog Servers	6514	TCP	Syslog
NSX Edge nodes	Syslog Servers	6514	UDP	Syslog
NSX Edge nodes	Traceroute Destination	33434 - 33523	UDP	Traceroute
NSX Edge nodes, Transport Nodes	NSX Edge nodes	3784, 3785	UDP	BFD between the Transport Node TEP IP address in the data.
SNMP Servers	NSX Edge nodes	161	UDP	SNMP

TCP and UDP Ports Used by ESXi, KVM Hosts, and Bare Metal Server

ESXi, KVM hosts, and bare metal server when used as transport nodes need certain TCP and UDP ports available.

Table 3-4. TCP and UDP Ports Used by ESXi and KVM Hosts

Source	Target	Port	Protocol	Description
ESXi host	NSX Controller	1235	TCP	Control Plane - LCP to CCP communication
ESXi host	NSX Manager	5671	TCP	AMPQ Communication channel to NSX Manager
ESXi host	NSX Manager	8080	TCP	Install and upgrade HTTP repository
ESXi and KVM host	NSX Manager	443	TCP	Management and provisioning connection
ESXi and KVM host	NSX Manager	443	TCP	Install and upgrade HTTP repository
GENEVE Termination End Point (TEP)	GENEVE Termination End Point (TEP)	6081	UDP	Transport network
KVM host	NSX Manager	5671	TCP	AMPQ Communication channel to NSX Manager
KVM host	NSX Controller	1235	TCP	Control Plane - LCP to CCP communication
KVM host	NSX Manager	8080	TCP	Install and upgrade HTTP repository
NSX Manager	ESXi host	443	TCP	Management and provisioning connection
NSX Manager	KVM host	443	TCP	Management and provisioning connection
ESXi and KVM host	Syslog Servers	514	TCP	Syslog
ESXi and KVM host	Syslog Servers	514	UDP	Syslog
ESXi and KVM host	Syslog Servers	6514	TCP	Syslog
ESXi and KVM host	Syslog Servers	6514	UDP	Syslog
NSX-T Data Center transport node	NSX-T Data Center transport node	3784, 3785	UDP	BFD Session between TEPs, in the datapath using TEP interface

Installing NSX-T Data Center Components

You must install the NSX Manager and NSX Edge core components to use NSX-T Data Center.

NSX Manager Installation

NSX Manager provides a graphical user interface (GUI) and REST APIs for creating, configuring, and monitoring NSX-T Data Center components such as logical switches, logical routers, and firewalls.

NSX Manager provides a system view and is the management component of NSX-T Data Center.

For high availability, NSX-T Data Center supports a management cluster of three NSX Managers. For a production environment, deploying a management cluster is recommended. For a proof-of-concept environment, you can deploy a single NSX Manager.

NSX Manager Deployment, Platform, and Installation Requirements

The following table details the NSX Manager deployment, platform, and installation requirements

Requirements	Description
Supported deployment methods	■ OVA/OVF ■ QCOW2
Supported platforms	See NSX Manager VM System Requirements. On ESXi, it is recommended that the NSX Manager appliance be installed on shared storage.
IP address	An NSX Manager must have a static IP address. You cannot change the IP address after installation.
NSX-T Data Center appliance password	■ At least 12 characters ■ At least one lower-case letter ■ At least one upper-case letter ■ At least one digit ■ At least one special character ■ At least five different characters ■ No dictionary words ■ No palindromes ■ More than four monotonic character sequence is not allowed
Hostname	When installing NSX Manager, specify a hostname that does not contain invalid characters such as an underscore. If the hostname contains any invalid character, after deployment the hostname will be set to `nsx-manager`. For more information about hostname restrictions, see https://tools.ietf.org/html/rfc952 and https://tools.ietf.org/html/rfc1123.
VMware Tools	The NSX Manager VM running on ESXi has VMTools installed. Do not remove or upgrade VMTools.

Requirements	Description
System	- Verify that the system requirements are met. See System Requirements. - Verify that the required ports are open. See Ports and Protocols. - Verify that a datastore is configured and accessible on the ESXi host. - Verify that you have the IP address and gateway, DNS server IP addresses, domain search list, and the NTP server IP address for the NSX Manager to use. - If you do not already have one, create the target VM port group network. Place the NSX-T Data Center appliances on a management VM network. If you have multiple management networks, you can add static routes to the other networks from the NSX-T Data Center appliance. - Plan your NSX Manager IPv4 or IPv6 IP addressing scheme.
OVF Privileges	Verify that you have adequate privileges to deploy an OVF template on the ESXi host. A management tool that can deploy OVF templates, such as vCenter Server or the vSphere Client. The OVF deployment tool must support configuration options to allow for manual configuration. OVF tool version must be 4.0 or later.
Client Plug-in	The Client Integration Plug-in must be installed.

Note On an NSX Manager fresh install, reboot, or after an `admin` password change when prompted on first login, it might take several minutes for the NSX Manager to start.

NSX Manager Installation Scenarios

Important When you install NSX Manager from an OVA or OVF file, either from vSphere Client or the command line, OVA/OVF property values such as user names, passwords, or IP addresses are not validated before the VM is powered on.

- If you specify a user name for the `admin` or `audit` user, the name must be unique. If you specify the same name, it is ignored and the default names (`admin` and `audit`) is used.

- If the password for the `admin` user does not meet the complexity requirements, you must log in to NSX Manager through SSH or at the console as the `admin` user with the password `default`. You are prompted to change the password.

- If the password for the `audit` user does not meet the complexity requirements, the user account is disabled. To enable the account, log in to NSX Manager through SSH or at the console as the `admin` user and run the command `set user audit` to set the `audit` user's password (the current password is an empty string).

- If the password for the **root** user does not meet the complexity requirements, you must log in to NSX Manager through SSH or at the console as **root** with the password **vmware**. You are prompted to change the password.

Caution   Changes made to the NSX-T Data Center while logged in with the **root** user credentials might cause system failure and potentially impact your network. You can only make changes using the **root** user credentials with the guidance of VMware Support team.

Note   The core services on the appliance do not start until a password with sufficient complexity is set.

After you deploy NSX Manager from an OVA file, you cannot change the VM's IP settings by powering off the VM and modifying the OVA settings from vCenter Server.

Configuring NSX Manager for Access by the DNS Server

By default, transport nodes access NSX Managers based on their IP addresses. However, this can be based also on the DNS names of the NSX Managers.

By enabling FQDN usage (DNS) on NSX Managers, the IP address of the Managers can change without affecting the transport nodes.

You enable FQDN usage by publishing the FQDNs of the NSX Managers.

Note   Enabling FQDN usage (DNS) on NSX Managers is required for multisite Lite and NSXNSX Cloud and deployments. (It is optional for all other deployment types.) See *Multisite Deployment of NSX-T Data Center* in the *NSX-T Data Center Administration Guide* and Chapter 11 Installing NSX Cloud Components in this guide.

Publishing the FQDNs of the NSX Managers

After installing the NSX-T Data Center core components and CSM, to enable NAT using FQDN you would set up the entries for lookup and reverse lookup in the NSX-T DNS server in your deployment.

In addition, you must also enable publishing the NSX Manager FQDNs using the NSX-T API.

Example request: PUT `https://<nsx-mgr>/api/v1/configs/management`

```
{
  "publish_fqdns": true,
  "_revision": 0
}
```

Example response:

```
{
  "publish_fqdns": true,
```

```
    "_revision": 1
}
```

See the *NSX-T Data Center API Guide* for details.

Note   After publishing the FQDNs, validate access by the transport nodes as described in the next section.

Validating Access via FQDN by Transport Nodes

After publishing the FQDNs of the NSX Managers, verify that the transport nodes are successfully accessing the NSX Managers.

Using SSH, log into a transport node such as a hypervisor or Edge node, and run the `get controllers` CLI command.

Example response:

```
Controller IP    Port SSL     Status      Is Physical Master  Session State  Controller FQDN
192.168.60.5     1235 enabled connected   true                up             nsxmgr.corp.com
```

NSX Edge Installation

The NSX Edge provides routing services and connectivity to network NSX Edges that are external to the NSX-T Data Center deployment. An NSX Edge is required if you want to deploy a tier-0 router or a tier-1 router with stateful services such as network address translation (NAT), VPN, and so on.

Note   There can be only one tier-0 router per NSX Edge node. However, multiple tier-1 load routers can be hosted on one NSX Edge node. NSX Edge VMs of different sizes can be combined in the same cluster; however, it is not recommended.

Table 3-5. NSX Edge Deployment, Platforms, and Installation Requirements

Requirements	Description
Supported deployment methods	- OVA/OVF - ISO with PXE - ISO without PXE
Supported platforms	NSX Edge is supported only on ESXi or on bare metal. NSX Edge is not supported on KVM.
PXE installation	The Password string must be encrypted with sha-512 algorithm for the root and admin user password.

Table 3-5. NSX Edge Deployment, Platforms, and Installation Requirements (continued)

Requirements	Description
NSX-T Data Center appliance password	- At least 12 characters - At least one lower-case letter - At least one upper-case letter - At least one digit - At least one special character - At least five different characters - No dictionary words - No palindromes - More than four monotonic character sequence is not allowed
Hostname	When installing NSX Edge, specify a hostname that does not contain invalid characters such as an underscore. If the hostname contains any invalid character, after deployment the hostname will be set to **localhost**. For more information about hostname restrictions, see https://tools.ietf.org/html/rfc952 and https://tools.ietf.org/html/rfc1123.
VMware Tools	The NSX Edge VM running on ESXi has VMTools installed. Do not remove or upgrade VMTools.
System	Verify that the system requirements are met. See NSX Edge VM System Requirements.
Ports	Verify that the required ports are open. See Ports and Protocols.
IP Addresses	If you have multiple management networks, you can add static routes to the other networks from the NSX-T Data Center appliance. Plan your NSX Edge IPv4 or IPv6 IP addressing scheme.
OVF Template	- Verify that you have adequate privileges to deploy an OVF template on the ESXi host. - Verify that hostnames do not include underscores. Otherwise, the hostname is set to *nsx-manager*. - A management tool that can deploy OVF templates, such as vCenter Server or the vSphere Client. The OVF deployment tool must support configuration options to allow for a manual configuration. - The Client Integration Plug-in must be installed.
NTP Server	The same NTP server must be configured on all NSX Edge servers in an Edge cluster.

NSX Edge Installation Scenarios

Important   When you install NSX Edge from an OVA or OVF file, either from vSphere Web Client or the command line, OVA/OVF property values such as user names, passwords, or IP addresses are not validated before the VM is powered on.

- If you specify a user name for the `admin` or `audit` user, the name must be unique. If you specify the same name, it is ignored and the default names (`admin` and `audit`) is used.

- If the password for the `admin` user does not meet the complexity requirements, you must log in to NSX Edge through SSH or at the console as the `admin` user with the password `default`. You are prompted to change the password.

- If the password for the `audit` user does not meet the complexity requirements, the user account is disabled. To enable the account, log in to NSX Edge through SSH or at the console as the `admin` user and run the command `set user audit` to set the `audit` user's password (the current password is an empty string).

- If the password for the `root` user does not meet the complexity requirements, you must log in to NSX Edge through SSH or at the console as `root` with the password `vmware`. You are prompted to change the password.

Caution   Changes made to the NSX-T Data Center while logged in with the `root` user credentials might cause system failure and potentially impact your network. You can only make changes using the `root` user credentials with the guidance of VMware Support team.

Note   The core services on the appliance do not start until a password with sufficient complexity has been set.

After you deploy NSX Edge from an OVA file, you cannot change the VM's IP settings by powering off the VM and modifying the OVA settings from vCenter Server.

Join NSX Edge with the Management Plane

Joining NSX Edges with the management plane ensures that the NSX Manager and NSX Edges can communicate with each other.

Prerequisites

Verify that you have admin privileges to log in to the NSX Edges and NSX Manager appliance.

Procedure

1   Open an SSH session to the NSX Manager appliance.

2   Open an SSH session to the NSX Edge.

3   On the NSX Manager appliance, run the `get certificate api thumbprint` command.

   The command output is a string of alphanumeric numbers that is unique to this NSX Manager.

For example:

```
NSX-Manager1> get certificate api thumbprint
...
```

4 On the NSX Edge, run the **join management-plane** command.

Provide the following information:

- Hostname or IP address of the NSX Manager with an optional port number

- Username of the NSX Manager

- Certificate thumbprint of the NSX Manager

- Password of the NSX Manager

```
NSX-Edge1> join management-plane NSX-Manager1 username admin thumbprint <NSX-Manager1's-
thumbprint>
Password for API user: <NSX-Manager1's-password>
Node successfully registered and Edge restarted
```

Repeat this command on each NSX Edge node.

5 Verify the result by running the `get managers` command on your NSX Edges.

```
nsx-edge-1> get managers
- 192.168.110.47   Connected
```

6 In the NSX Manager UI, select **System > Fabric > Nodes > Edge Transport Nodes** page.

The NSX Manager connectivity should be Up. If NSX Manager connectivity is not Up, try refreshing the browser window.

What to do next

Add the NSX Edge as a transport node. See Create an NSX Edge Transport Node.

Installing NSX-T Data Center on vSphere

4

You can install the NSX-T Data Center components, NSX Manager and NSX Edge using the UI or CLI.

Make sure that you have the supported vSphere version. See vSphere support.

This chapter includes the following topics:

- Install NSX Manager and Available Appliances

- Install an NSX Edge on ESXi Using a vSphere GUI

Install NSX Manager and Available Appliances

You can use the vSphere Client to deploy NSX Manager or the Cloud Service Manager as a virtual appliance.

Cloud Service Manager is a virtual appliance that uses NSX-T Data Center components and integrates them with your public cloud.

Prerequisites

- Verify that the system requirements are met. See System Requirements.

- Verify that the required ports are open. See Ports and Protocols.

- Verify that a datastore is configured and accessible on the ESXi host.

- Verify that you have the IP address and gateway, DNS server IP addresses, domain search list, and the NTP server IP address for the NSX Manager to use.

- If you do not already have one, create the target VM port group network. Place the NSX-T Data Center appliances on a management VM network.

  If you have multiple management networks, you can add static routes to the other networks from the NSX-T Data Center appliance.

- Plan your NSX Manager IPv4 or IPv6 IP addressing scheme.

Procedure

1   Locate the NSX-T Data Center OVA file on the VMware download portal.

    Either copy the download URL or download the OVA file.

2   In the vSphere Client, select the host on which to install NSX-T Data Center.

3   Right-click and select **Deploy OVF template** to start the installation wizard.

4   Enter the download OVA URL or navigate to the OVA file.

5   Enter a name for the NSX Manager VM.

The name you enter appears in the vSphere inventory.

6   Select a compute resource for the NSX Manager appliance.

  ◆   To install on a ESXi host managed by vCenter, select a host on which to deploy the NSX Manager appliance.

  ◆   To install on a standalone ESXi host, select the host on which to deploy the NSX Manager appliance.

7   Verify the OVF template details.

8   For an optimal performance, reserve memory for the NSX Manager appliance.

Set the reservation to ensure that NSX Manager has sufficient memory to run efficiently. See NSX Manager VM System Requirements.

9   Select a datastore to store the NSX Manager appliance files.

10 Select a destination network for each source network.

11 Select the port group or destination network for the NSX Manager.

12 Enter the NSX Manager system root, CLI admin, and audit passwords.

Your passwords must comply with the password strength restrictions.

  ▪   At least 12 characters

  ▪   At least one lower-case letter

  ▪   At least one upper-case letter

  ▪   At least one digit

  ▪   At least one special character

  ▪   At least five different characters

  ▪   No dictionary words

  ▪   No palindromes

  ▪   More than four monotonic character sequence is not allowed

13 Enter the hostname of the NSX Manager.

Note   The host name must be a valid domain name. Ensure that each part of the host name (domain/subdomain) that is separated by dot must start with an alphabet character.

14 Accept the default **NSX Manager** role for VM.

Select the **nsx-cloud-service-manager** role from the drop-down menu to install the NSX Cloud appliance.

15 Enter the default gateway, management network IPv4, management network netmask, DNS, and NTP IP address.

16 Enable SSH and allow root SSH login to the NSX Manager command line.

By default, these options are disabled for security reasons.

17 Verify that all your custom OVF template specification is accurate and click **Finish** to initiate the installation.

The installation might take 7-8 minutes.

18 From the vSphere Client, open NSX Manager VM console to track the boot process.

19 After the NSX Manager boots, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.

20 Enter the `get services` command to verify that all the services are running.

If the services are not running, wait for all the services to start running.

Note The following services are not running by default: liagent, migration-coordinator, and snmp. You can start them as follows:

▪ `start service liagent`

▪ `start service migration-coordinator`

▪ For SNMP:

```
set snmp community <community-string>
start service snmp
```

21 Verify that your NSX Manager has the required connectivity.

Make sure that you can perform the following tasks.

▪ Ping your NSX Manager from another machine.

▪ The NSX Manager can ping its default gateway.

▪ The NSX Manager can ping the hypervisor hosts that are in the same network as the NSX Manager using the management interface.

▪ The NSX Manager can ping its DNS server and its NTP server.

▪ If you enabled SSH, make sure that you can SSH to your NSX Manager.

If connectivity is not established, make sure that the network adapter of the virtual appliance is in the proper network or VLAN.

What to do next

Log in to the NSX Manager from a supported web browser. See Log In to the Newly Created NSX Manager .

Install NSX Manager on ESXi Using the Command-Line OVF Tool

If you prefer to automate or use CLI for the NSX Manager installation, you can use the VMware OVF Tool, which is a command-line utility.

By default, `nsx_isSSHEnabled` and `nsx_allowSSHRootLogin` are both disabled for security reasons. When they are disabled, you cannot SSH or log in to the NSX Manager command line. If you enable `nsx_isSSHEnabled` but not `nsx_allowSSHRootLogin`, you can SSH to NSX Manager but you cannot log in as root.

Prerequisites

- Verify that the system requirements are met. See System Requirements.

- Verify that the required ports are open. See Ports and Protocols.

- Verify that a datastore is configured and accessible on the ESXi host.

- Verify that you have the IP address and gateway, DNS server IP addresses, domain search list, and the NTP server IP address for the NSX Manager to use.

- If you do not already have one, create the target VM port group network. Place the NSX-T Data Center appliances on a management VM network.

  If you have multiple management networks, you can add static routes to the other networks from the NSX-T Data Center appliance.

- Plan your NSX Manager IPv4 or IPv6 IP addressing scheme.

Procedure

1 Run the `ovftool` command with the appropriate parameters.

   The process depends on whether the host is standalone or managed by vCenter Server.

   - For a standalone host:

     - Windows example:

       ```
       C:\Program Files\VMware\VMware OVF Tool>ovftool \
       --sourceType=OVA \
       --name=nsx-manager \
       --X:injectOvfEnv \
       --X:logFile=<filepath>\nsxovftool.log \
       --allowExtraConfig \
       --datastore=<datastore name> \
       --network=<network name> \
       --acceptAllEulas \
       --noSSLVerify \
       --diskMode=thin \
       ```

```
--powerOn \
--prop:"nsx_role=nsx-manager nsx-controller" \
--prop:"nsx_ip_0=10.168.110.75" \
--prop:"nsx_netmask_0=255.255.255.0" \
--prop:"nsx_gateway_0=10.168.110.1" \
--prop:"nsx_dns1_0=10.168.110.10" \
--prop:"nsx_domain_0=corp.local" \
--prop:"nsx_ntp_0=10.168.110.10" \
--prop:"nsx_isSSHEnabled=<True|False>" \
--prop:"nsx_allowSSHRootLogin=<True|False>" \
--prop:"nsx_passwd_0=<password>" \
--prop:"nsx_cli_passwd_0=<password>" \
--prop:"nsx_cli_audit_passwd_0=<password>" \
--prop:"nsx_hostname=nsx-manager" \
<nsx-unified-appliance-release>.ova \
vi://root:<password>@10.168.110.51
```

Note The above Windows code block uses the backslash (\) to indicate the continuation of the command line. In actual use, omit the backslash and put the entire command in a single line.

Note In the above example, 10.168.110.51 is the IP address of the host machine where NSX Manager is to be deployed.

▪ Linux example:

```
mgrformfactor="small"
ipAllocationPolicy="fixedPolicy"
mgrdatastore="QNAP-Share-VMs"
mgrnetwork="Management-VLAN-210"

mgrname01="nsx-manager-01"
mgrhostname01="nsx-manager-01"
mgrip01="192.168.210.121"

mgrnetmask="255.255.255.0"
mgrgw="192.168.210.254"
mgrdns="192.168.110.10"
mgrntp="192.168.210.254"
mgrpasswd="<password>"
mgrssh="<True|False>"
mgrroot="<True|False>"
logLevel="trivia"

mgresxhost01="192.168.110.113"

ovftool --noSSLVerify --skipManifestCheck --powerOn \
--deploymentOption=$mgrformfactor \
--diskMode=thin \
--acceptAllEulas \
--allowExtraConfig \
--ipProtocol=IPv4 \
--ipAllocationPolicy=$ipAllocationPolicy \
```

```
--datastore=$mgrdatastore \
--network=$mgrnetwork \
--name=$mgrname01 \
--prop:nsx_hostname=$mgrhostname01 \
--prop:nsx_role="nsx-manager nsx-controller" \
--prop:nsx_ip_0=$mgrip01 \
--prop:nsx_netmask_0=$mgrnetmask \
--prop:nsx_gateway_0=$mgrgw \
--prop:nsx_dns1_0=$mgrdns \
--prop:nsx_ntp_0=$mgrntp \
--prop:nsx_passwd_0=$mgrpasswd \
--prop:nsx_cli_passwd_0=$mgrpasswd \
--prop:nsx_cli_audit_passwd_0=$mgrpasswd \
--prop:nsx_isSSHEnabled=$mgrssh \
--prop:nsx_allowSSHRootLogin=$mgrroot \
--X:logFile=nsxt-manager-ovf.log \
--X:logLevel=$logLevel \
/home/<user/nsxt-autodeploy/<nsx-unified-appliance-release>.ova \
vi://root:<password>@$mgresxhost01
```

The result should look something like this:

```
Opening OVA source: nsx-<component>.ova
The manifest validates
Source is signed and the certificate validates
Opening VI target: vi://root:<password>@10.168.110.51
Deploying to VI: vi://root:<password>@10.168.110.51
Transfer Completed
Powering on VM: nsx-manager nsx-controller
Task Completed
Completed successfully
```

- For a host managed by vCenter Server:

  - Windows example:

    ```
    C:\Users\Administrator\Downloads>ovftool
    --name=nsx-manager \
    --X:injectOvfEnv \
    --X:logFile=ovftool.log \
     --allowExtraConfig \
    --datastore=ds1 \
    --network="management" \
    --acceptAllEulas \
    --noSSLVerify \
    --diskMode=thin \
    --powerOn \
    --prop:"nsx_role=nsx-manager nsx-controller" \
    --prop:"nsx_ip_0=10.168.110.75" \
    --prop:"nsx_netmask_0=255.255.255.0" \
    --prop:"nsx_gateway_0=10.168.110.1" \
    --prop:"nsx_dns1_0=10.168.110.10" \
    --prop:"nsx_domain_0=corp.local" \
    ```

```
--prop:"nsx_ntp_0=10.168.110.10" \
--prop:"nsx_isSSHEnabled=<True|False>" \
--prop:"nsx_allowSSHRootLogin=<True|False>" \
--prop:"nsx_passwd_0=<password>" \
--prop:"nsx_cli_passwd_0=<password>" \
--prop:"nsx_hostname=nsx-manager" \
<nsx-unified-appliance-release>.ova \
vi://administrator@vsphere.local:<password>@10.168.110.24/?ip=10.168.110.51
```

Note The above Windows code block uses the backslash (\) to indicate the continuation of the command line. In actual use, omit the backslash and put the entire command in a single line.

- Linux example:

```
mgrformfactor="small"
ipAllocationPolicy="fixedPolicy"
mgrdatastore="QNAP-Share-VMs"
mgrnetwork="Management-VLAN-210"

mgrname01="nsx-manager-01"
mgrhostname01="nsx-manager-01"
mgrip01="192.168.210.121"

mgrnetmask="255.255.255.0"
mgrgw="192.168.210.254"
mgrdns="192.168.110.10"
mgrntp="192.168.210.254"
mgrpasswd="<password>"
mgrssh="<True|False>"
mgrroot="<True|False>"
logLevel="trivia"

vcadmin="administrator@vsphere.local"
vcpass="<password>"
vcip="192.168.110.151"
mgresxhost01="192.168.110.113"

ovftool --noSSLVerify --skipManifestCheck --powerOn \
--deploymentOption=$mgrformfactor \
--diskMode=thin \
--acceptAllEulas \
--allowExtraConfig \
--ipProtocol=IPv4 \
--ipAllocationPolicy=$ipAllocationPolicy \
--datastore=$mgrdatastore \
--network=$mgrnetwork \
--name=$mgrname01 \
--prop:nsx_hostname=$mgrhostname01 \
--prop:nsx_role="nsx-manager nsx-controller" \
--prop:nsx_ip_0=$mgrip01 \
--prop:nsx_netmask_0=$mgrnetmask \
--prop:nsx_gateway_0=$mgrgw \
--prop:nsx_dns1_0=$mgrdns \
```

```
--prop:nsx_ntp_0=$mgrntp \
--prop:nsx_passwd_0=$mgrpasswd \
--prop:nsx_cli_passwd_0=$mgrpasswd \
--prop:nsx_cli_audit_passwd_0=$mgrpasswd \
--prop:nsx_isSSHEnabled=$mgrssh \
--prop:nsx_allowSSHRootLogin=$mgrroot \
--X:logFile=nsxt-manager-ovf.log \
--X:logLevel=$logLevel \
/home/<user/nsxt-autodeploy/<nsx-unified-appliance-release>.ova \
vi://$vcadmin:$vcpass@$vcip/?ip=$mgresxhost01
```

The result should look something like this:

```
Opening OVA source: nsx-<component>.ova
The manifest validates
Source is signed and the certificate validates
Opening VI target: vi://administrator@vsphere.local@10.168.110.24:443/
Deploying to VI: vi://administrator@vsphere.local@10.168.110.24:443/
Transfer Completed
Powering on VM: nsx-manager nsx-controller
Task Completed
Completed successfully
```

2   For an optimal performance, reserve memory for the NSX Manager appliance.

Set the reservation to ensure that NSX Manager has sufficient memory to run efficiently. See NSX Manager VM System Requirements.

3   From the vSphere Client, open NSX Manager VM console to track the boot process.

4   After the NSX Manager boots, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.

5   Verify that your NSX Manager has the required connectivity.

Make sure that you can perform the following tasks.

- Ping your NSX Manager from another machine.

- The NSX Manager can ping its default gateway.

- The NSX Manager can ping the hypervisor hosts that are in the same network as the NSX Manager using the management interface.

- The NSX Manager can ping its DNS server and its NTP server.

- If you enabled SSH, make sure that you can SSH to your NSX Manager.

If connectivity is not established, make sure that the network adapter of the virtual appliance is in the proper network or VLAN.

What to do next

Log in to the NSX Manager from a supported web browser. See Log In to the Newly Created NSX Manager .

Configure NSX-T Data Center to Display the GRUB Menu at Boot Time

Configuring the NSX-T Data Center appliance to display the GRUB menu at boot time is required to reset the root password of the NSX-T Data Center appliance.

Important   If the configuration is not performed after deploying the appliance and you forget the root, admin, or audit password, resetting it is not possible.

Procedure

1   Log in to the VM as root.

2   Change the value for the parameter `GRUB_HIDDEN_TIMEOUT` in the `/etc/default/grub` file.

   `GRUB_HIDDEN_TIMEOUT=2`

3   (Optional) Change the GRUB password in the `/etc/grub.d/40_custom` file.

   The default password is `VMware1`.

4   Update the GRUB configuration.

   `update-grub`

Log In to the Newly Created NSX Manager

After you install NSX Manager, you can use the user interface to perform other installation tasks.

After you install NSX Manager, you can join the Customer Experience Improvement Program (CEIP) for NSX-T Data Center. See Customer Experience Improvement Program in the *NSX-T Data Center Administration Guide* for more information about the program, including how to join or leave the program later.

Prerequisites

Verify that NSX Manager is installed. See Install NSX Manager and Available Appliances.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

   The EULA appears.

2   Read and accept the EULA terms.

3   Select whether to join the VMware's Customer Experience Improvement Program (CEIP).

4   Click **Save**

Add a Compute Manager

A compute manager, for example, vCenter Server, is an application that manages resources such as hosts and VMs.

NSX-T Data Center polls compute managers to collect cluster information from vCenter Server.

When you add a vCenter Server compute manager, you must provide a vCenter Server user's credentials. You can provide the vCenter Server administrator's credentials, or create a role and a user specifically for NSX-T Data Center and provide this user's credentials. This role must have the following vCenter Server privileges:

`Extension.Register extension`
`Extension.Unregister extension`
`Extension.Update extension`
`Sessions.Message`
`Sessions.Validate session`
`Sessions.View and stop sessions`
`Host.Configuration.Maintenance`
`Host.Local Operations.Create virtual machine`
`Host.Local Operations.Delete virtual machine`
`Host.Local Operations.Reconfigure virtual machine`
`Host.Configuration.NetworkConfiguration`
`Tasks`
`Scheduled task`
`Global.Cancel task`
`Permissions.Reassign role permissions`
`Resource.Assign vApp to resource pool`
`Resource.Assign virtual machine to resource pool`
`Virtual Machine.Configuration`
`Virtual Machine.Guest Operations`
`Virtual Machine.Provisioning`
`Virtual Machine.Inventory`
`Network.Assign network`
`vApp`

For more information about vCenter Server roles and privileges, see the *vSphere Security* document.

Prerequisites

- Verify that you use the supported vSphere version. See Supported vSphere version.

- IPv6 and IPv4 communication with vCenter Server.

- Verify that you use the recommended number of compute managers. See https://configmax.vmware.com/home.

> **Note** NSX-T Data Center does not support the same vCenter Server to be registered with more than one NSX Manager.

Procedure

1 From your browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Compute Managers > Add**.

3 Complete the compute manager details.

Option	Description
Name and Description	Type the name to identify the vCenter Server.
	You can optionally describe any special details such as, the number of clusters in the vCenter Server.
Domain Name/IP Address	Type the IP address of the vCenter Server.
Type	Keep the default option.
Username and Password	Type the vCenter Server login credentials.
Thumbprint	Type the vCenter Server SHA-256 thumbprint algorithm value.

If you left the thumbprint value blank, you are prompted to accept the server provided thumbprint.

After you accept the thumbprint, it takes a few seconds for NSX-T Data Center to discover and register the vCenter Server resources.

4 If the progress icon changes from **In progress** to **Not registered**, perform the following steps to resolve the error.

    a   Select the error message and click **Resolve**. One possible error message is the following:

```
Extension already registered at CM <vCenter Server name> with id <extension ID>
```

    b   Enter the vCenter Server credentials and click **Resolve**.

        If an existing registration exists, it will be replaced.

Results

It takes some time to register the compute manager with vCenter Server and for the connection status to appear as UP.

You can click the compute manager's name to view the details, edit the compute manager, or to manage tags that apply to the compute manager.

After the vCenter Server is successfully registered, do not power off and delete the NSX Manager VM without deleting the compute manager first. Otherwise, when you deploy a new NSX Manager, you will not be able to register the same vCenter Server again. You will get the error that the vCenter Server is alredy registered with another NSX Manager.

Deploy NSX Manager Nodes to Form a Cluster from UI

You can deploy multiple NSX Manager nodes to provide high availability and reliability.

After the new nodes are deployed, these nodes connect to the NSX Manager node to form a cluster. The recommended number of clustered NSX Manager nodes is three.

Note Deploying multiple NSX Manager nodes using the UI is supported only on ESXi hosts managed by vCenter Server.

All the repository details and the password of the first deployed NSX Manager node are synchronized with the newly deployed nodes in the cluster.

Prerequisites

- Verify that an NSX Manager node is installed. See Install NSX Manager and Available Appliances.

- Verify that compute manager is configured. See Add a Compute Manager.

- Verify that the system requirements are met. See System Requirements.

- Verify that the required ports are open. See Ports and Protocols.

- Verify that a datastore is configured and accessible on the ESXi host.

- Verify that you have the IP address and gateway, DNS server IP addresses, domain search list, and the NTP server IP address for the NSX Manager to use.

- If you do not already have one, create the target VM port group network. Place the NSX-T Data Center appliances on a management VM network.

  If you have multiple management networks, you can add static routes to the other networks from the NSX-T Data Center appliance.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Appliances > Overview > Add Nodes**.

3 Enter the NSX Manager common attribute details.

Option	Description
Compute Manager	Registered resource compute manager is populated.
Enable SSH	Toggle the button to allow an SSH login to the new NSX Manager node.

Option	Description
Enable Root Access	Toggle the button to allow the root access to the new NSX Manager node.
CLI Username and Password Confirmation	Set the CLI password and password confirmation for the new node. Your password must comply with the password strength restrictions. ■ At least 12 characters ■ At least one lower-case letter ■ At least one upper-case letter ■ At least one digit ■ At least one special character ■ At least five different characters ■ No dictionary words ■ No palindromes ■ More than four monotonic character sequence is not allowed The CLI username is already set to admin.
Root Password and Password Confirmation	Set the root password and password confirmation for the new node. Your password must comply with the password strength restrictions. ■ At least 12 characters ■ At least one lower-case letter ■ At least one upper-case letter ■ At least one digit ■ At least one special character ■ At least five different characters ■ No dictionary words ■ No palindromes ■ More than four monotonic character sequence is not allowed
DNS Servers	Enter the DNS server IP address available in the vCenter Server.
NTP Servers	Enter the NTP server IP address.

4   Enter the NSX Manager node details.

Option	Description
Name	Enter a name for the NSX Manager node.
Cluster	Designate the cluster the node is going to join from the drop-down menu.
Resource Pool or Host	Assign either a resource pool or a host for the node from the drop-down menu.
Datastore	Select a datastore for the node files from the drop-down menu.
Network	Assign the network from the drop-down menu.
Management IP/Netmask	Enter the IP address and netmask.
Management Gateway	Enter the gateway IP address.

5   (Optional) Click **New Node** and configure another node.

Repeat steps 3-4.

6 Click **Finish**.

The new nodes are deployed. You can track the deployment process on the **System > Appliances > Overview** page or the vCenter Server.

7 Wait for 10-15 minutes for the deployment, cluster formation, and repository synchronization to complete.

All the repository details and the password of the first deployed NSX Manager node are synchronized with the newly deployed nodes in the cluster.

8 After the NSX Manager boots, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.

9 Enter the `get services` command to verify that all the services are running.

If the services are not running, wait for all the services to start running.

Note The following services are not running by default: liagent, migration-coordinator, and snmp. You can start them as follows:

- `start service liagent`

- `start service migration-coordinator`

- For SNMP:

```
set snmp community <community-string>
start service snmp
```

10 Log in to the first deployed NSX Manager node and enter the `get cluster status` command to verify that the nodes are successfully added to the cluster.

11 Verify that your NSX Manager has the required connectivity.

Make sure that you can perform the following tasks.

- Ping your NSX Manager from another machine.

- The NSX Manager can ping its default gateway.

- The NSX Manager can ping the hypervisor hosts that are in the same network as the NSX Manager using the management interface.

- The NSX Manager can ping its DNS server and its NTP server.

- If you enabled SSH, make sure that you can SSH to your NSX Manager.

If connectivity is not established, make sure that the network adapter of the virtual appliance is in the proper network or VLAN.

What to do next

Configure NSX Edge. See Install an NSX Edge on ESXi Using a vSphere GUI.

Deploy NSX Manager Nodes to Form a Cluster Using CLI

Joining the NSX Manager to form a cluster using CLI ensures that all the NSX Manager nodes in cluster can communicate with each other.

Prerequisites

The installation of NSX-T Data Center components must be complete.

Procedure

1 Open an SSH session to the first deployed NSX Manager node.

2 Log in with the administrator credentials.

3 On the NSX Manager node, run the `get certificate api thumbprint` command.

The command output is a string of numbers that is unique to this NSX Manager.

4 Run the `get cluster config` command to get the first deployed NSX Manager cluster ID.

5 Add a NSX Manager node to the cluster.

> **Note** You must run the join command on the newly deployed NSX Manager node.

Provide the followingNSX Manager information:

- Hostname or IP address node that you want to join

- Cluster ID

- User name

- Password

- Certificate thumbprint

You can use the CLI command or API call.

- CLI command

```
host> join <NSX-Manager-IP> cluster-id <cluster-id> username<NSX-Manager-username>
password<NSX-Manager-password> thumbprint <NSX-Manager1's-thumbprint>
```

- API call `POST https://<nsx-mgr>/api/v1/cluster?action=join_cluster`

The joining and cluster stabilizing process might 10-15 minutes.

6 Add the third NSX Manager node to the cluster.

Repeat step 5.

7 Verify the cluster status by running the `get cluster status` command on your hosts.

8 Select **System > Appliances > Overview** and verify the cluster connectivity.

What to do next

Create a transport zone. See Create a Standalone Host or Bare Metal Server Transport Node.

Configure a Virtual IP (VIP) Address for a Cluster

To provide fault tolerance and high availability to NSX Manager nodes, assign a virtual IP address (VIP) to a member of the NSX-T cluster.

NSX Managers of a cluster become part of an HTTPS group to service API and UI requests. The leader node of the cluster assumes ownership of the set VIP of the cluster to service any API and UI request. Any API and UI request coming in from clients is directed to the leader node.

Note   When assigning Virtual IP, all the NSX Manager VMs in the cluster must be configured in the same subnet.

If the leader node that owns VIP becomes unavailable, NSX-T elects a new leader. The new leader owns the VIP. It sends out a gratuitous ARP packet advertising the new VIP to MAC address mapping. After a new leader node is elected, new API and UI requests are sent to the new leader node.

Failover of VIP to a new leader node of the cluster might take a few minutes to become functional. If the VIP fails over to a new leader node because the previous leader node became unavailable, reauthenticate credentials so that API requests are directed to the new leader node.

Note   VIP is not designed to serve as a load-balancer and you cannot use it if you enable the vIDM **External Load Balancer Integration** from **System > Users > Configuration**. Do not set up a VIP if you want to use the External Load Balancer from vIDM. See Configure VMware Identity Manager Integration in the *NSX-T Data Center Administration Guide* for more details.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Go to **System** > **Overview**.

3 In the Virtual IP field, click **Edit**.

4 Enter the VIP for the cluster. Ensure that VIP is part of the same subnet as the other management nodes.

5 Click **Save**.

6 To verify the cluster status and the API leader of the HTTPS group, enter the NSX Manager CLI command `get cluster status verbose` in the NSX Manager console or over SSH.

The following is an example output with the leader marked in bold.

```
Group Type: HTTPS
Group Status: STABLE

Members:
    UUID                                    FQDN            IP
STATUS
    cdb93642-ccba-fdf4-8819-90bf018cd727    nsx-manager     192.196.197.84
```

```
    UP
        51a13642-929b-8dfc-3455-109e6cc2a7ae          nsx-manager          192.196.198.156
    UP
        d0de3642-d03f-c909-9cca-312fd22e486b          nsx-manager          192.196.198.54
    UP

    Leaders:
        SERVICE                                    LEADER                                LEASE
    VERSION
        api                                        cdb93642-ccba-fdf4-8819-90bf018cd727    8
```

7 To troubleshoot VIP, verify Reverse Proxy logs at `/var/log/proxy/reverse-proxy.log` and cluster manager logs at `/var/log/cbm/cbm.log` in the NSX Manager CLI.

Results

Any API requests to NSX-T is redirected to the virtual IP address of the cluster, which is owned by the leader node. The leader node then routes the request forward to the other components of the appliance.

Install an NSX Edge on ESXi Using a vSphere GUI

If you prefer an interactive NSX Edge installation, you can use the vSphere web client.

Important In NSX-T, the NSX Edge VM does not support vMotion.

Prerequisites

See NSX Edge network requirements in NSX Edge Installation.

Procedure

1 Locate the NSX Edge appliance OVA file on the VMware download portal.

Either copy the download URL or download the OVA file onto your computer.

2 In the vSphere Client, select the host on which to install NSX Edge appliance.

3 Right-click and select **Deploy OVF template** to start the installation wizard.

4 Enter the download OVA URL or navigate to the saved OVA file.

5 Enter a name for the NSX Edge VM.

The name you type appears in the inventory.

6 Select a compute resource for the NSX Edge appliance.

7 For an optimal performance, reserve memory for the NSX Edge appliance.

Set the reservation to ensure that NSX Edge has sufficient memory to run efficiently. See NSX Edge VM System Requirements.

8 Verify the OVF template details.

9 Select a datastore to store the NSX Edge appliance files.

10   Accept the default source and destination network interface.

You can accept the default network destination for the rest of the networks and change the network configuration after the NSX Edge is deployed.

11   Select the IP allocation from the drop-down menu.

12   Enter the NSX Edge system root, CLI admin, and audit passwords.

Your passwords must comply with the password strength restrictions.

- At least 12 characters

- At least one lower-case letter

- At least one upper-case letter

- At least one digit

- At least one special character

- At least five different characters

- No dictionary words

- No palindromes

- More than four monotonic character sequence is not allowed

13   Enter the default gateway, management network IPv4, management network netmask, DNS, and NTP IP address.

14   (Optional) Register the NSX Edge with the management plane, if you have a NSX Manager available.

a   Enter the parent NSX Manager node IP address and thumbprint.

b   Run the API call POST `https://<nsx-manager>/api/v1/aaa/registration-token` to retrieve the NSX Manager token.

15   Enter the hostname of the NSX Edge VM.

16   Enable SSH and allow root SSH login to the NSX Edge command line.

By default, these options are disabled for security reasons.

17   Verify that all your custom OVA template specification is accurate and click **Finish** to initiate the installation.

The installation might take 7-8 minutes.

18   Open the console of the NSX Edge to track the boot process.

If the console window does not open, make sure that pop-ups are allowed.

19   After the NSX Edge starts, log in to the CLI with admin credentials.

Note   After NSX Edge starts, if you do not log in with admin credentials for the first time, the data plane service does not automatically start on NSX Edge.

20 Run the `get interface eth0.<vlan_ID>` command to verify that the IP address was applied as expected

```
nsx-edge-1> get interface eth0.100

Interface: eth0.100
  Address: 192.168.110.37/24
  MAC address: 00:50:56:86:62:4d
  MTU: 1500
  Default gateway: 192.168.110.1
  Broadcast address: 192.168.110.255
  ...
```

Note When bringing up NSX Edge VMs on non-NSX managed host, verify that the MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

21 Run the `get managers` command to verify that the NSX Edge is registered.

```
- 10.29.14.136 Standby
- 10.29.14.135 Standby
- 10.29.14.134 Connected
```

22 Verify that the NSX Edge appliance has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge.

- You can ping your NSX Edge.

- NSX Edge can ping its default gateway.

- NSX Edge can ping the hypervisor hosts that are in the same network as the NSX Edge.

- NSX Edge can ping its DNS server and its NTP server.

23 Troubleshoot connectivity problems.

Note If connectivity is not established, make sure the VM network adapter is in the proper network or VLAN.

By default, the NSX Edge datapath claims all virtual machine NICs except the management NIC (the one that has an IP address and a default route). If you incorrectly assigned a NIC as the management interface, follow these steps to use DHCP to assign management IP address to the correct NIC.

a   Log in CLI and type the `stop service dataplane` command.

b   Type the `set interface` *interface* `dhcp plane mgmt` command.

c   Place *interface* into the DHCP network and wait for an IP address to be assigned to that *interface*.

d   Type the **start service dataplane** command.

The datapath fp-ethX ports used for the VLAN uplink and the tunnel overlay are shown in the **get interfaces** and **get physical-port** commands on the NSX Edge.

What to do next

Join the NSX Edge with the management plane. See Join NSX Edge with the Management Plane.

Install NSX Edge on ESXi Using the Command-Line OVF Tool

If you prefer to automate NSX Edge installation, you can use the VMware OVF Tool, which is a command-line utility.

Prerequisites

- Verify that the system requirements are met. See System Requirements.

- Verify that the required ports are open. See Ports and Protocols.

- Verify that a datastore is configured and accessible on the ESXi host.

- Verify that you have the IP address and gateway, DNS server IP addresses, domain search list, and the NTP server IP address for the NSX Manager to use.

- If you do not already have one, create the target VM port group network. Place the NSX-T Data Center appliances on a management VM network.

  If you have multiple management networks, you can add static routes to the other networks from the NSX-T Data Center appliance.

- Plan your NSX Manager IPv4 or IPv6 IP addressing scheme.

- See NSX Edge network requirements in NSX Edge Installation.

- Verify that you have adequate privileges to deploy an OVF template on the ESXi host.

- Verify that hostnames do not include underscores. Otherwise, the hostname is set to *localhost*.

- OVF Tool version 4.3 or later.

Procedure

◆ For a standalone host, run the `ovftool` command with the appropriate parameters.

```
C:\Users\Administrator\Downloads>ovftool
--name=nsx-edge-1
--deploymentOption=medium
--X:injectOvfEnv
--X:logFile=ovftool.log
--allowExtraConfig
```

```
--datastore=ds1
--net:"Network 0=Mgmt"
--net:"Network 1=nsx-tunnel"
--net:"Network 2=vlan-uplink"
--net:"Network 3=vlan-uplink"
--acceptAllEulas
--noSSLVerify
--diskMode=thin
--powerOn
--prop:nsx_ip_0=192.168.110.37
--prop:nsx_netmask_0=255.255.255.0
--prop:nsx_gateway_0=192.168.110.1
--prop:nsx_dns1_0=192.168.110.10
--prop:nsx_domain_0=corp.local
--prop:nsx_ntp_0=192.168.110.10
--prop:nsx_isSSHEnabled=True
--prop:nsx_allowSSHRootLogin=True
--prop:nsx_passwd_0=<password>
--prop:nsx_cli_passwd_0=<password>
--prop:nsx_hostname=nsx-edge
<path/url to nsx component ova>
vi://root:<password>@192.168.110.51
```

```
Opening OVA source: nsx-<component>.ova
The manifest validates
Source is signed and the certificate validates
Opening VI target: vi://root@192.168.110.24
Deploying to VI: vi://root@192.168.110.24
Transfer Completed
Powering on VM: nsx-edge-1
Task Completed
Completed successfully
```

◆ For a host managed by vCenter Server, run the `ovftool` command with the appropriate parameters.

```
C:\Users\Administrator\Downloads>ovftool
--name=nsx-edge-1
--deploymentOption=medium
--X:injectOvfEnv
--X:logFile=ovftool.log
--allowExtraConfig
--datastore=ds1
--net:"Network 0=Mgmt"
--net:"Network 1=nsx-tunnel"
--net:"Network 2=vlan-uplink"
--net:"Network 3=vlan-uplink"
--acceptAllEulas
--noSSLVerify
--diskMode=thin
--powerOn
```

```
--prop:nsx_ip_0=192.168.110.37
--prop:nsx_netmask_0=255.255.255.0
--prop:nsx_gateway_0=192.168.110.1
--prop:nsx_dns1_0=192.168.110.10
--prop:nsx_domain_0=corp.local
--prop:nsx_ntp_0=192.168.110.10
--prop:nsx_isSSHEnabled=True
--prop:nsx_allowSSHRootLogin=True
--prop:nsx_passwd_0=<password>
--prop:nsx_cli_passwd_0=<password>
--prop:nsx_hostname=nsx-edge
<path/url to nsx component ova>
vi://administrator@vsphere.local:<password>@192.168.110.24/?ip=192.168.210.53
```

```
Opening OVA source: nsx-<component>.ova
The manifest validates
Source is signed and the certificate validates
Opening VI target: vi://administrator@vsphere.local@192.168.110.24:443/
Deploying to VI: vi://administrator@vsphere.local@192.168.110.24:443/
Transfer Completed
Powering on VM: nsx-edge-1
Task Completed
Completed successfully
```

◆ For an optimal performance, reserve memory for the NSX Manager appliance.

Set the reservation to ensure that NSX Manager has sufficient memory to run efficiently. See NSX Manager VM System Requirements.

◆ Open the console of the NSX Edge to track the boot process.

◆ After the NSX Edge starts, log in to the CLI with admin credentials.

◆ Run the `get interface eth0.<vlan_ID>` command to verify that the IP address was applied as expected

```
nsx-edge-1> get interface eth0.100

Interface: eth0.100
  Address: 192.168.110.37/24
  MAC address: 00:50:56:86:62:4d
  MTU: 1500
  Default gateway: 192.168.110.1
  Broadcast address: 192.168.110.255
  ...
```

Note When bringing up NSX Edge VMs on non-NSX managed host, verify that the MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

◆ Verify that the NSX Edge appliance has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge.

   ▪ You can ping your NSX Edge.

   ▪ NSX Edge can ping its default gateway.

   ▪ NSX Edge can ping the hypervisor hosts that are in the same network as the NSX Edge.

   ▪ NSX Edge can ping its DNS server and its NTP server.

◆ Troubleshoot connectivity problems.

> **Note** If connectivity is not established, make sure the VM network adapter is in the proper network or VLAN.

By default, the NSX Edge datapath claims all virtual machine NICs except the management NIC (the one that has an IP address and a default route). If you incorrectly assigned a NIC as the management interface, follow these steps to use DHCP to assign management IP address to the correct NIC.

a   Log in CLI and type the **stop service dataplane** command.

b   Type the **set interface *interface* dhcp plane mgmt** command.

c   Place *interface* into the DHCP network and wait for an IP address to be assigned to that *interface*.

d   Type the **start service dataplane** command.

The datapath fp-ethX ports used for the VLAN uplink and the tunnel overlay are shown in the **get interfaces** and **get physical-port** commands on the NSX Edge.

What to do next

Join the NSX Edge with the management plane. See Join NSX Edge with the Management Plane.

Installing NSX-T Data Center on KVM

5

NSX-T Data Center supports KVM in two ways: as a host transport node and as a host for NSX Manager.

Make sure that you have the supported KVM versions. See NSX Manager VM System Requirements.

This chapter includes the following topics:

- Set Up KVM
- Manage Your Guest VMs in the KVM CLI
- Install NSX Manager on KVM
- Log In to the Newly Created NSX Manager
- Install Third-Party Packages on a KVM Host
- Verify Open vSwitch Version on RHEL KVM Hosts
- Deploy NSX Manager Nodes to Form a Cluster Using CLI
- Install NSX Edge on Bare Metal

Set Up KVM

If you plan to use KVM as a transport node or as a host for NSX Manager guest VM, but you do not already have a KVM setup, you can use the procedure described here.

Note The Geneve encapsulation protocol uses UDP port 6081. You must allow this port access in the firewall on the KVM host.

Procedure

1 (Only RHEL) Open the `/etc/yum.conf` file.

2 Search for the line `exclude`.

3 Add the line "`kernel* redhat-release*`" to configure YUM to avoid any unsupported RHEL upgrades.

`exclude=[existing list] kernel* redhat-release*`

If you plan to run NSX-T Data Center Container Plug-in, which has specific compatibility requirements, exclude the container-related modules as well.

`exclude=[existing list] kernel* redhat-release* kubelet-* kubeadm-* kubectl-* docker-*`

The supported RHEL versions are 7.4. and 7.5.

4 Install KVM and bridge utilities.

Linux Distribution	Commands
Ubuntu	`apt-get install -y qemu-kvm libvirt-bin ubuntu-vm-builder bridge-utils` `virtinst virt-manager virt-viewer libguestfs-tools`

RHEL or CentOS Linux	`yum groupinstall "Virtualization Hypervisor"`
	`yum groupinstall "Virtualization Client"`
	`yum groupinstall "Virtualization Platform"`
	`yum groupinstall "Virtualization Tools"`

SUSE Linux Enterprise Server	Start YaSt and select **Virtualization > Install Hypervisor and Tools**. YaSt allows you to automatically enable and configure the network bridge.

5 For NSX manager to automatically install NSX software packages on KVM host, prepare the network configuration of the uplink/data interface.

The KVM host can have multiple network interfaces. For the network interface that you plan to provide as an uplink interface (data interface) for NSX-T purposes, it is important to have network configuration files correctly populated. NSX-T looks at these network configuration files to create NSX-T specific network devices. On Ubuntu, populate `/etc/network/interfaces` file. On RHEL, CentOS, or SUSE, populate the `/etc/sysconfig/network-scripts/ifcfg-$uplinkdevice` file.

In the following examples, interface "ens32" is the uplink device (data interface). Depending on your deployment environment, this interface can use DHCP or static IP settings.

Note Interface names might vary in different environments.

Important For Ubuntu, all network configurations must be specified in `/etc/network/interfaces`. Do not create individual network configuration files such as `/etc/network/ifcfg-eth1`, which can lead to failure of transport node creation.

Linux Distribution	Network Configuration
Ubuntu	Edit /etc/network/interfaces: ``` auto eth0 iface eth0 inet manual auto ens32 iface ens32 inet manual ```

RHEL or CentOS Linux	Edit /etc/sysconfig/network-scripts/ifcfg-ens32:
	
	``
	DEVICE="ens32"
	  TYPE="Ethernet"
	  NAME="ens32"
	  UUID="<something>"
	  BOOTPROTO="none"
	  HWADDR="<something>"
	  ONBOOT="yes"
	  NM_CONTROLLED="no"
	``

SUSE Linux Enterprise Server	If a SLES host already exists, verify that data interfaces is already configured on the host. If you do not have a pre-configured SLES host, see the reference configuration for the management and data interface. Edit /etc/sysconfig/network/ifcfg-ens32: ``` DEVICE="ens32" NAME="ens32" UUID="<UUID>" BOOTPROTO="none" LLADDR="<HWADDR>" STARTMODE="yes" ```

6 Restart networking service `systemctl restart network` or reboot the Linux server for the networking changes take effect.

7 After the KVM host is configured as a transport node, the bridge interface 'nsx-vtep0.0' is automatically created by NSX-T.

In Ubuntu, the `/etc/network/interfaces` file has entries such as the following:

```
iface nsx-vtep0.0 inet static
pre-up ip addr flush dev nsx-vtep0.0
address <IP_pool_address>
netmask <subnet_mask>
mtu 1600
down ifconfig nsx-vtep0.0 down
up ifconfig nsx-vtep0.0 up
```

In RHEL, the host NSX agent (nsxa) creates a configuration file named `ifcfg-nsx-vtep0.0` that has entries such as the following:

```
DEVICE=nsx-vtep0.0
BOOTPROTO=static
NETMASK=<IP address>
IPADDR=<subnet mask>
MTU=1600
ONBOOT=yes
USERCTL=no
NM_CONTROLLED=no
```

In SUSE,

```
DEVICE=nsx-vtep0.0
BOOTPROTO=static
NETMASK=255.255.255.0
IPADDR=192.168.13.119
MACADDR=ae:9d:b7:ca:20:4a
MTU=1600
USERCTL=no
STARTMODE=auto
```

8   Configure the syslog rotation policy as time-based instead of size-based policy. With a size-based syslog rotation policy, the log files generated might be of very large sizes.

Manage Your Guest VMs in the KVM CLI

NSX Manager can be installed as KVM VMs. In addition, KVM can be used as the hypervisor for NSX-T Data Center transport nodes.

KVM guest VM management is beyond the scope of this guide. However, here are some simple KVM CLI commands to get you started.

To manage your guest VMs in the KVM CLI, use the `virsh` commands. Following are some common `virsh` commands. Refer to the KVM documentation for additional information.

```
# List running
virsh list

# List all
virsh list --all

# Control instances
virsh start <instance>
virsh shutdown <instance>
virsh destroy <instance>
virsh undefine <instance>
virsh suspend <instance>
```

```
virsh resume <instance>

# Access an instance's CLI
virsh console <instance>
```

In the Linux CLI, the `ifconfig` command shows the vnetX interface, which represents the interface created for the guest VM. If you add additional guest VMs, additional vnetX interfaces are added.

```
ifconfig
...

vnet0     Link encap:Ethernet  HWaddr fe:54:00:b0:a0:6d
          inet6 addr: fe80::fc54:ff:feb0:a06d/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:13183 errors:0 dropped:0 overruns:0 frame:0
          TX packets:181524 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:500
          RX bytes:4984832 (4.9 MB)  TX bytes:29498709 (29.4 MB)
```

Install NSX Manager on KVM

NSX Manager can be installed as a virtual appliance on a KVM host.

The QCOW2 installation procedure uses guestfish, a Linux command-line tool to write virtual machine settings into the QCOW2 file.

Prerequisites

- KVM set up. See Set Up KVM.

- Privileges to deploy a QCOW2 image on the KVM host.

- Verify that the password in the guestinfo adheres to the password complexity requirements so that you can log in after installation. See NSX Manager Installation.

- Familiarize yourself with the NSX Manager resource requirements. See NSX Manager VM System Requirements.

- If you plan to install Ubuntu OS, it is recommened to install Ubuntu version 18.04 before installing NSX Manager on the KVM host.

Procedure

1 Download the NSX Manager QCOW2 image from the **nsx-unified-appliance > exports > kvm** folder.

2 Copy it to the KVM machine that is going to run the NSX Manager using SCP or sync.

3 (Ubuntu only) Add the currently logged in user as a libvirtd user:

```
adduser $USER libvirtd
```

4 In the same directory where you saved the QCOW2 image, create a file called guestinfo.xml and populate it with the NSX Manager VM's properties.

Property	Description
■ **nsx_cli_passwd_0** ■ **nsx_cli_audit_passwd_0** ■ **nsx_passwd_0**	Your passwords must comply with the password strength restrictions. ■ At least 12 characters ■ At least one lower-case letter ■ At least one upper-case letter ■ At least one digit ■ At least one special character ■ At least five different characters ■ No dictionary words ■ No palindromes ■ More than four monotonic character sequence is not allowed
nsx_hostname	Enter the host name for the NSX Manager. The host name must be a valid domain name. Ensure that each part of the host name (domain/subdomain) that is separated by dot must start with an alphabet character.
nsx_role	■ *nsx-manager*: Required. This role-name installs the NSX Manager appliance. ■ *nsx-cloud-service-manager*: Optional. After installing NSX Manager, use this role-name to install the Cloud Service Manager appliance for NSX Cloud.
nsx_isSSHEnabled	You can enable or disable this property. If enabled, you can log in to the NSX Manager using SSH.
nsx_allowSSHRootLogin	You can enable or disable this property. If enabled, you can log in to the NSX Manager using SSH as the root user. To be able to use this property, `nsx_isSSHEnabled` must be enabled.
■ **nsx_dns1_0** ■ **nsx_ntp_0** ■ **nsx_domain_0** ■ **nsx_gateway_0** ■ **nsx_netmask_0** ■ **nsx_ip_0**	Enter IP addresses for the default gateway, management network IPv4, management network netmask, DNS, and NTP IP address.

For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<Environment
     xmlns="http://schemas.dmtf.org/ovf/environment/1"
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
     xmlns:oe="http://schemas.dmtf.org/ovf/environment/1">
   <PropertySection>
    <Property oe:key="nsx_cli_passwd_0" oe:value="<password>"/>
```

```
        <Property oe:key="nsx_cli_audit_passwd_0" oe:value="<password>"/>
        <Property oe:key="nsx_passwd_0" oe:value="<password>"/>
        <Property oe:key="nsx_hostname" oe:value="nsx-manager1"/>
        <Property oe:key="nsx_role" oe:value="nsx-manager"/>
        <Property oe:key="nsx_isSSHEnabled" oe:value="True"/>
        <Property oe:key="nsx_allowSSHRootLogin" oe:value="True"/>
        <Property oe:key="nsx_dns1_0" oe:value="10.168.110.10"/>
        <Property oe:key="nsx_ntp_0" oe:value="10.168.110.10"/>
        <Property oe:key="nsx_domain_0" oe:value="corp.local"/>
        <Property oe:key="nsx_gateway_0" oe:value="10.168.110.83"/>
        <Property oe:key="nsx_netmask_0" oe:value="255.255.252.0"/>
        <Property oe:key="nsx_ip_0" oe:value="10.168.110.19"/>
    </PropertySection>
</Environment>
```

Note In the example, `nsx_isSSHEnabled` and `nsx_allowSSHRootLogin` are both enabled. When they are disabled, you cannot SSH or log in to the NSX Manager command line. If you enable `nsx_isSSHEnabled` but not `nsx_allowSSHRootLogin`, you can SSH to NSX Manager but you cannot log in as root.

5   Use guestfish to write the `guestinfo.xml` file into the QCOW2 image.

Note After the `guestinfo` information is written into a QCOW2 image, the information cannot be overwritten.

```
sudo guestfish --rw -i -a nsx-unified-appliance-<BuildNumber>.qcow2 upload guestinfo /config/
guestinfo
```

6   Deploy the QCOW2 image with the `virt-install` command.

The vCPU and RAM values are suitable for a large VM. The network name and portgroup name are specific to your environment. The model must be `virtio`.

```
sudo virt-install \
--import \
--ram 48000 \
--vcpus 12 \
--name <manager-name> \
--disk path=<manager-qcow2-file-path>,bus=virtio,cache=none \
--network network=<network-name>,portgroup=<portgroup-name>,model=virtio \
--noautoconsole              \
--cpu mode=host-passthrough,cache.mode=passthrough

Starting install...
Domain installation still in progress. Waiting for installation to complete.
```

7 Verify that the NSX Manager is deployed.

```
virsh list --all

Id    Name            State
-------------------------------
18    nsx-manager1    running
```

8 Open the NSX Manager console and log in.

```
virsh console 18
Connected to domain nsx-manager1
Escape character is ^]

nsx-manager1 login: admin
Password:
```

9 After the NSX Manager boots, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.

10 Run `get services` to verify that the services are running.

11 Verify that your NSX Manager has the required connectivity.

Make sure that you can perform the following tasks.

- Ping your NSX Manager from another machine.

- The NSX Manager can ping its default gateway.

- The NSX Manager can ping the hypervisor hosts that are in the same network as the NSX Manager using the management interface.

- The NSX Manager can ping its DNS server and its NTP server.

- If you enabled SSH, make sure that you can SSH to your NSX Manager.

If connectivity is not established, make sure that the network adapter of the virtual appliance is in the proper network or VLAN.

12 Exit the KVM console.

```
control-]
```

13 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

Log In to the Newly Created NSX Manager

After you install NSX Manager, you can use the user interface to perform other installation tasks.

After you install NSX Manager, you can join the Customer Experience Improvement Program (CEIP) for NSX-T Data Center. See Customer Experience Improvement Program in the *NSX-T Data Center Administration Guide* for more information about the program, including how to join or leave the program later.

Prerequisites

Verify that NSX Manager is installed. See Install NSX Manager and Available Appliances.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

The EULA appears.

2 Read and accept the EULA terms.

3 Select whether to join the VMware's Customer Experience Improvement Program (CEIP).

4 Click **Save**

Install Third-Party Packages on a KVM Host

To prepare a KVM host to be a fabric node, you must install some third-party packages.

Prerequisites

- (RHEL and CentOS Linux) Before you install the third-party packages, run the following commands to install the virtualization packages.

```
yum groupinstall "Virtualization Hypervisor"
yum groupinstall "Virtualization Client"
yum groupinstall "Virtualization Platform"
yum groupinstall "Virtualization Tools"
```

If you are not able to install the packages, you can manually install them with the command `yum install glibc.i686 nspr` on a new installation.

- (Ubuntu) Before you install the third-party packages, run the following commands to install the virtualization packages.

```
apt install -y \
qemu-kvm \
libvirt-bin \
virtinst \
virt-manager \
virt-viewer \
ubuntu-vm-builder \
bridge-utils
```

- (SUSE Linux Enterprise Server) Before you install the third-party packages, run the following commands to install the virtualization packages.

```
libcap-progs
```

Procedure

◆ On Ubuntu, run `apt-get install <package_name>` to install the third-party packages manually.

Ubuntu 18.04 Packages	Ubuntu 16.04 Packages
traceroute python-mako python-netaddr python-simplejson python-unittest2 python-yaml python-openssl dkms make	libboost-chrono1.58.0 libboost-filesystem1.58.0 libgoogle-glog0v5 libgoogle-perftools4 libprotobuf9v5 traceroute python-mako python-netaddr python-simplejson python-unittest2 python-yaml python-openssl libboost-date-time1.58.0 libleveldb1v5 python-gevent python-protobuf libboost-program-options1.58.0 dkms

◆ On RHEL and CentOS Linux, run `yum install <package_name>` to install the third-party packages manually.

If you manually prepare the host that is already registered to RHEL or CentOS, you do not need to install third-party packages on the host.

RHEL 7.6, 7.5, and 7.4	CentOS Linux 7.5 and 7.4
wget PyYAML libunwind python-gevent python-mako python-netaddr redhat-lsb-core tcpdump	wget PyYAML libunwind python-gevent python-mako python-netaddr redhat-lsb-core tcpdump

◆ On SUSE, run `zypper install <package_name>` to install the third-party packages manually.

SUSE Linux Enterprise Server 12.0
python-simplejson python-PyYAML python-netaddr lsb-release

Verify Open vSwitch Version on RHEL KVM Hosts

If OVS packages exist on the RHEL host, you must remove the existing packages and install the supported packages.

The supported Open vSwitch version is 2.9.1.8614397-1.

Procedure

1 Verify that the current version of the Open vSwitch installed on the host.

`ovs-vswitchd --version`

If you have an Open vSwitch newer or older version, you must replace that Open vSwitch version with the supported one.

2 Open the Open vSwitch folder.

3 Delete the following Open vSwitch packages.

- kmod-openvswitch
- openvswitch
- openvswitch-selinux-policy

4 Alternatively, add the Open vSwitch packages required by NSX-T Data Center.

a Log in to the host as an administrator.

b Download and copy the nsx-lcp file into the /tmp directory.

c Untar the package.

`tar -zxvf nsx-lcp-<release>-rhel75_x86_64.tar.gz`

d Navigate to the package directory.

`cd nsx-lcp-rhel75_x86_64/`

e Replace existing Open vSwitch version with the supported one.

- For the newer Open vSwitch version, use the `--nodeps` command.

  For example, rpm -Uvh kmod-openvswitch-<new version>.e17.x86_64.rpm --nodeps

  rpm -Uvh openvswitch-*.rpm --nodeps

- For the older Open vSwitch version, use the `--force` command.

  For example, rpm -Uvh kmod-openvswitch-<new version>.e17.x86_64.rpm --nodeps --force

  rpm -Uvh openvswitch-*.rpm --nodeps --force

Deploy NSX Manager Nodes to Form a Cluster Using CLI

Joining the NSX Manager to form a cluster using CLI ensures that all the NSX Manager nodes in cluster can communicate with each other.

Prerequisites

The installation of NSX-T Data Center components must be complete.

Procedure

1   Open an SSH session to the first deployed NSX Manager node.

2   Log in with the administrator credentials.

3   On the NSX Manager node, run the `get certificate api thumbprint` command.

    The command output is a string of numbers that is unique to this NSX Manager.

4   Run the `get cluster config` command to get the first deployed NSX Manager cluster ID.

5   Add a NSX Manager node to the cluster.

    Note   You must run the join command on the newly deployed NSX Manager node.

    Provide the followingNSX Manager information:

    ■   Hostname or IP address node that you want to join

    ■   Cluster ID

    ■   User name

    ■   Password

    ■   Certificate thumbprint

    You can use the CLI command or API call.

    ■   CLI command

    ```
    host> join <NSX-Manager-IP> cluster-id <cluster-id> username<NSX-Manager-username>
    password<NSX-Manager-password> thumbprint <NSX-Manager1's-thumbprint>
    ```

    ■   API call POST `https://<nsx-mgr>/api/v1/cluster?action=join_cluster`

    The joining and cluster stabilizing process might 10-15 minutes.

6   Add the third NSX Manager node to the cluster.

    Repeat step 5.

7   Verify the cluster status by running the `get cluster status` command on your hosts.

8   Select **System > Appliances > Overview** and verify the cluster connectivity.

What to do next

Create a transport zone. See Create a Standalone Host or Bare Metal Server Transport Node.

Install NSX Edge on Bare Metal

Use PXE server to automate installation of NSX Edge on a bare metal server or use ISO file to install NSX Edge as a VM appliance or on a bare metal server.

PXE boot installation is not supported for NSX Manager. You also cannot configure networking settings, such as the IP address, gateway, network mask, NTP, and DNS.

Prerequisites

- If the NSX Edge Bare Metal server is running version 6.7u3 or earlier, do not upgrade the NSX Edge `virtualHW.version` to **14** or higher in vCenter Server. By default, the `virtualHW.version` is set to **13**.

- By default, NSX Edge Bare Metal bond devices aggregating Ethernet devices to form a LAG are optimized for load balancing. So, a bond device only uses network devices that are on a local NUMA node whose CPU transmits packets. If the devices forming the bond span multiple NUMA nodes but the CPUs allocated to packet processing belong to a subset of the NUMA nodes, then only some of the devices send out traffic. In short, not all devices are used for balancing traffic that is sent out of the bond device.You cannot disable the default optimization.

  However, if you want to use all Ethernet devices of the bond to load balance traffic, you must move all Ethernet devices to the NUMA nodes where the packet processing CPUs are attached.

  Note   Failover is exclusive of load balancing. If the Ethernet device attached to the local NUMA node is down, then the bond sends traffic to the other device even though it is not NUMA local. The load-balancing optimization does not impact failover functionality.

Install NSX Edge via ISO File as a Virtual Appliance

You can install NSX Edge VMs in a manual fashion using an ISO file.

Important   The NSX-T Data Center component virtual machine installations include VMware Tools. Removal or upgrade of VMware Tools is not supported for NSX-T Data Center appliances.

Prerequisites

- See NSX Edge network requirements in NSX Edge Installation.

Procedure

1   Go to your MyVMware account (myvmware.com) and navigate to **VMware NSX-T Data Center > Downloads**.

2   Locate and download the ISO file for NSX Edge.

3   In the vSphere Client, select the host datastore.

4   Select **Files > Upload Files > Upload a File to a Datastore**, browse to the ISO file, and upload.

   If you are using a self-signed certificate, open the IP address in a browser and accept the certificate and reupload the ISO file.

5   In the vSphere Client inventory, select the host you uploaded the ISO file. or in the vSphere Client,

6   Right-click and select **New Virtual Machine** .

7   Select a compute resource for the NSX Edge appliance.

8   Select a datastore to store the NSX Edge appliance files.

9   Accept the default compatibility for your NSX Edge VM.

10   Select the supported ESXi operating systems for your NSX Edge VM.

11   Configure the virtual hardware.

- New Hard Disk - `200 GB`

- New Network - `VM Network`

- New CD/DVD Drive - `Datastore ISO File`

  You must click **Connect** to bind the NSX Edge ISO file to the VM.

12   Power on the new NSX Edge VM.

13   During ISO boot, open the VM console and choose **Automated installation**.

There might be a pause of 10 seconds after you press Enter.

During installation, the installer prompts you to enter a VLAN ID for the management interface. Select **Yes** and enter a VLAN ID to create a VLAN subinterface for the network interface. Select **No** if you do not want to configure VLAN tagging on the packet.

During power-on, the VM requests a network configuration via DHCP. If DHCP is not available in your environment, the installer prompts you for IP settings.

By default, the root login password is `vmware`, and the admin login password is `default`.

When you log in for the first time, you are prompted to change the password. This password change method has strict complexity rules, including the following:

- At least 12 characters

- At least one lower-case letter

- At least one upper-case letter

- At least one digit

- At least one special character

- At least five different characters

- No dictionary words

- No palindromes

- More than four monotonic character sequence is not allowed

Important   The core services on the appliance do not start until a password with sufficient complexity has been set.

14 For an optimal performance, reserve memory for the NSX Edge appliance.

Set the reservation to ensure that NSX Edge has sufficient memory to run efficiently. See NSX Edge VM System Requirements.

15 After the NSX Edge starts, log in to the CLI with admin credentials.

Note After NSX Edge starts, if you do not log in with admin credentials for the first time, the data plane service does not automatically start on NSX Edge.

16 There are three ways to configure a management interface.

- Untagged interface. This interface type creates an out-of-band management interface.

  (DHCP) `set interface eth0 dhcp plane mgmt`

  (Static) `set interface eth0 ip <CIDR> gateway <gateway-ip> plane mgmt`

- Tagged interface.

  `set interface eth0 vlan <vlan_ID> plane mgmt`

  (DHCP) `set interface eth0.<vlan_ID> dhcp plane mgmt`

  (Static) `set interface eth0.<vlan_ID> ip <CIDR> gateway <gateway-ip> plane mgmt`

- In-band interface.

  `set interface mac <mac_address> vlan <vlan_ID> in-band plane mgmt`

  (DHCP) `set interface eth0.<vlan_ID> dhcp plane mgmt`

  (Static) `set interface eth0.<vlan_ID> ip <CIDR> gateway <gateway-ip> plane mgmt`

17 (Optional) Start SSH service. Run `start service ssh`.

18 Run the `get interface eth0.<vlan_ID>` command to verify that the IP address was applied as expected

```
nsx-edge-1> get interface eth0.100

Interface: eth0.100
  Address: 192.168.110.37/24
  MAC address: 00:50:56:86:62:4d
  MTU: 1500
  Default gateway: 192.168.110.1
  Broadcast address: 192.168.110.255
  ...
```

Note When bringing up NSX Edge VMs on non-NSX managed host, verify that the MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

19 (Tagged interface and In-band interface) Any existing VLAN management interface must be cleared before creating a new one.

`Clear interface eth0.<vlan_ID>`

To set a new interface, refer to step 15.

20 Verify that the NSX Edge appliance has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge.

- You can ping your NSX Edge.

- NSX Edge can ping its default gateway.

- NSX Edge can ping the hypervisor hosts that are in the same network as the NSX Edge.

- NSX Edge can ping its DNS server and its NTP server.

21 Troubleshoot connectivity problems.

Note If connectivity is not established, make sure the VM network adapter is in the proper network or VLAN.

By default, the NSX Edge datapath claims all virtual machine NICs except the management NIC (the one that has an IP address and a default route). If you incorrectly assigned a NIC as the management interface, follow these steps to use DHCP to assign management IP address to the correct NIC.

a   Log in CLI and type the **stop service dataplane** command.

b   Type the **set interface** *interface* **dhcp plane mgmt** command.

c   Place *interface* into the DHCP network and wait for an IP address to be assigned to that *interface*.

d   Type the **start service dataplane** command.

The datapath fp-ethX ports used for the VLAN uplink and the tunnel overlay are shown in the **get interfaces** and **get physical-port** commands on the NSX Edge.

What to do next

If you did not join the NSX Edge with the management plane, see Join NSX Edge with the Management Plane.

Install NSX Edge via ISO File on Bare Metal

You can install NSX Edge devices in a manual fashion on bare metal using an ISO file. This includes configuring networking settings, such as IP address, gateway, network mask, NTP, and DNS.

Prerequisites

- Verify that the system BIOS mode is set to Legacy BIOS.

- See NSX Edge network requirements in NSX Edge Installation.

Procedure

1   Locate the NSX Edge appliance ISO file under the **nsx-edgenode > publish > xenial_amd64** folder.

    Download the ISO file onto your computer.

2   Log in to the ILO of the bare metal.

3   Click **Launch** in the virtual console preview.

4   Select **Virtual Media > Connect Virtual Media**.

    Wait a few seconds for the virtual media to connect.

5   Select **Virtual Media > Map CD/DVD** and browse to the ISO file.

6   Select **Next Boot > Virtual CD/DVD/ISO**.

7   Select **Power > Reset System (warm boot)**.

    The installation duration depends on the bare metal environment.

8   Choose **Automated installation**.

    There might be a pause of 10 seconds after you press Enter.

9   Select the applicable primary network interface.

    During power-on, the installer requests a network configuration via DHCP. If DHCP is not available in your environment, the installer prompts you for IP settings.

    By default, the root login password is **vmware**, and the admin login password is **default**.

10 Open the console of the NSX Edge to track the boot process.

    If the console window does not open, make sure that pop-ups are allowed.

11 After the NSX Edge starts, log in to the CLI with admin credentials.

    Note   After NSX Edge starts, if you do not log in with admin credentials for the first time, the data plane service does not automatically start on NSX Edge.

12 After the reboot, you can log in with either admin or root credentials. The default root password is `vmware`.

13 There are three ways to configure a management interface.

    ■   Untagged interface. This interface type creates an out-of-band management interface.

        (DHCP) `set interface eth0 dhcp plane mgmt`

        (Static) `set interface eth0 ip <CIDR> gateway <gateway-ip> plane mgmt`

    ■   Tagged interface.

        `set interface eth0 vlan <vlan_ID> plane mgmt`

        (DHCP) `set interface eth0.<vlan_ID> dhcp plane mgmt`

(Static) `set interface eth0.<vlan_ID> ip <CIDR> gateway <gateway-ip> plane mgmt`

- In-band interface.

  `set interface mac <mac_address> vlan <vlan_ID> in-band plane mgmt`

  (DHCP) `set interface eth0.<vlan_ID> dhcp plane mgmt`

  (Static) `set interface eth0.<vlan_ID> ip <CIDR> gateway <gateway-ip> plane mgmt`

14 Run the `get interface eth0.<vlan_ID>` command to verify that the IP address was applied as expected

```
nsx-edge-1> get interface eth0.100

Interface: eth0.100
  Address: 192.168.110.37/24
  MAC address: 00:50:56:86:62:4d
  MTU: 1500
  Default gateway: 192.168.110.1
  Broadcast address: 192.168.110.255
  ...
```

Note When bringing up NSX Edge VMs on non-NSX managed host, verify that the MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

15 (Tagged interface and In-band interface) Any existing VLAN management interface must be cleared before creating a new one.

`clear interface eth0.<vlan_ID>`

To set a new interface, refer to step 13.

16 Verify that the NSX Edge appliance has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge.

- You can ping your NSX Edge.

- NSX Edge can ping its default gateway.

- NSX Edge can ping the hypervisor hosts that are in the same network as the NSX Edge.

- NSX Edge can ping its DNS server and its NTP server.

17 Troubleshoot connectivity problems.

Note If connectivity is not established, make sure the VM network adapter is in the proper network or VLAN.

By default, the NSX Edge datapath claims all virtual machine NICs except the management NIC (the one that has an IP address and a default route). If you incorrectly assigned a NIC as the management interface, follow these steps to use DHCP to assign management IP address to the correct NIC.

a   Log in CLI and type the **`stop service dataplane`** command.

b   Type the **`set interface`** *`interface`* **`dhcp plane mgmt`** command.

c   Place *interface* into the DHCP network and wait for an IP address to be assigned to that *interface*.

d   Type the **`start service dataplane`** command.

The datapath fp-ethX ports used for the VLAN uplink and the tunnel overlay are shown in the **get interfaces** and **get physical-port** commands on the NSX Edge.

What to do next

Join the NSX Edge with the management plane. See Join NSX Edge with the Management Plane.

Install NSX Edge on PXE Server

PXE is made up of several components: DHCP, HTTP, and TFTP. This procedure demonstrates how to set up a PXE server on Ubuntu.

DHCP dynamically distributes IP settings to NSX-T Data Center components, such as NSX Edge. In a PXE environment, the DHCP server allows NSX Edge to request and receive an IP address automatically.

TFTP is a file-transfer protocol. The TFTP server is always listening for PXE clients on the network. When it detects any network PXE client asking for PXE services, it provides the NSX-T Data Center component ISO file and the installation settings contained in a preseed file.

Prerequisites

■   A PXE server must be available in your deployment environment. The PXE server can be set up on any Linux distribution. The PXE server must have two interfaces, one for external communication and another for providing DHCP IP and TFTP services.

   If you have multiple management networks, you can add static routes to the other networks from the NSX-T Data Center appliance.

■   Verify that the preseeded configuration file has the parameters net.ifnames=0 and biosdevname=0 set after –– to persist after reboot.

■   See NSX Edge network requirements in NSX Edge Installation.

Procedure

1   (Optional) Use a kickstart file to set up a new TFTP or DHCP services on an Ubuntu server.

   A kickstart file is a text file that contains CLI commands that you run on the appliance after the first boot.

Name the kickstart file based on the PXE server it is pointing to. For example:

```
nsxcli.install
```

The file must be copied to your Web server, for example at /var/www/html/nsx-edge/nsxcli.install.

In the kickstart file, you can add CLI commands. For example, to configure the IP address of the management interface:

```
stop dataplane
set interface eth0 <ip-cidr-format> plane mgmt
start dataplane
```

To change the admin user password:

```
set user admin password <new_password> old-password <old-password>
```

If you specify a password in the preseed.cfg file, use the same password in the kickstart file. Otherwise, use the default password, which is "default".

To join the NSX Edge with the management plane:

```
join management-plane <manager-ip> thumbprint <manager-thumbprint> username <manager-username>
password <manager password>
```

2 Create two interfaces, one for management and another for DHCP and TFTP services.

Make sure that the DHCP/TFTP interface is in the same subnet that the NSX Edge resides in.

For example, if the NSX Edge management interfaces are going to be in the 192.168.210.0/24 subnet, place eth1 in that same subnet.

```
# The loopback network interface
auto lo
iface lo inet loopback

# PXE server's management interface
auto eth0
iface eth0 inet static
  address 192.168.110.81
  gateway 192.168.110.1
  netmask 255.255.255.0
  dns-nameservers 192.168.110.10

# PXE server's DHCP/TFTP interface
auto eth1
iface eth1 inet static
  address 192.168.210.82
  gateway 192.168.210.1
  netmask 255.255.255.0
  dns-nameservers 192.168.110.10
```

3 Install DHCP server software.

```
sudo apt-get install isc-dhcp-server -y
```

4 Edit the `/etc/default/isc-dhcp-server` file, and add the interface that provides the DHCP service.

```
INTERFACES="eth1"
```

5 (Optional) If you want this DHCP server to be the official DHCP server for the local network, uncomment the **authoritative;** line in the `/etc/dhcp/dhcpd.conf` file.

```
...
authoritative;
...
```

6 In the `/etc/dhcp/dhcpd.conf` file, define the DHCP settings for the PXE network.

For example:

```
subnet 192.168.210.0 netmask 255.255.255.0 {
    range 192.168.210.90 192.168.210.95;
    option subnet-mask 255.255.255.0;
    option domain-name-servers 192.168.110.10;
    option routers 192.168.210.1;
    option broadcast-address 192.168.210.255;
    default-lease-time 600;
    max-lease-time 7200;
}
```

7 Start the DHCP service.

```
sudo service isc-dhcp-server start
```

8 Verify that the DHCP service is running.

```
service --status-all | grep dhcp
```

9 Install Apache, TFTP, and other components that are required for PXE booting.

```
sudo apt-get install apache2 tftpd-hpa inetutils-inetd
```

10 Verify that TFTP and Apache are running.

```
service --status-all | grep tftpd-hpa
service --status-all | grep apache2
```

11 Add the following lines to the `/etc/default/tftpd-hpa` file.

```
RUN_DAEMON="yes"
OPTIONS="-l -s /var/lib/tftpboot"
```

12 Add the following line to the `/etc/inetd.conf` file.

```
tftp    dgram   udp    wait    root    /usr/sbin/in.tftpd /usr/sbin/in.tftpd -s /var/lib/tftpboot
```

13 Restart the TFTP service.

```
sudo /etc/init.d/tftpd-hpa restart
```

14 Copy or download the NSX Edge installer ISO file to a temporary folder.

15 Mount the ISO file and copy the install components to the TFTP server and the Apache server.

```
sudo mount -o loop ~/nsx-edge.<build>.iso /mnt
cd /mnt
sudo cp -fr install/netboot/* /var/lib/tftpboot/
sudo mkdir /var/www/html/nsx-edge
sudo cp -fr /mnt/* /var/www/html/nsx-edge/
```

16 (Optional) Edit the `/var/www/html/nsx-edge/preseed.cfg` file to modify the encrypted passwords.

You can use a Linux tool such as mkpasswd to create a password hash.

```
sudo apt-get install whois
sudo mkpasswd -m sha-512

Password:
$6$SUFGqs[...]FcoHLijOuFD
```

a   Modify the root password, edit `/var/www/html/nsx-edge/preseed.cfg` and search for the following line:

```
d-i passwd/root-password-crypted password $6$tgmLNLMp$9BuAHhN...
```

b   Replace the hash string.

You do not need to escape any special character such as $, ', ", or \.

c   Add the `usermod` command to `preseed.cfg` to set the password for root, admin, or both.

For example, search for the `echo 'VMware NSX Edge'` line and add the following command.

```
usermod --password '\$6\$VS3exId0aKmzW\$U3g0V7BF0DXlmRI.LR0v/VgloxVotEDp00bO2hUF8u/' root; \
usermod --password '\$6\$VS3exId0aKmzW\$U3g0V7BF0DXlmRI.LR0v/VgloxVotEDp00bO2hUF8u/' admin; \
```

The hash string is an example. You must escape all special characters. The root password in the first `usermod` command replaces the password that is set in `d-i passwd/root-password-crypted password $6$tgm....`

If you use the `usermod` command to set the password, the user is not prompted to change the password at the first login. Otherwise, the user must change the password at the first login.

17 Add the following lines to the `/var/lib/tftpboot/pxelinux.cfg/default` file.

Replace 192.168.210.82 with the IP address of your TFTP server.

```
label nsxedge
    kernel ubuntu-installer/amd64/linux
    ipappend 2
    append netcfg/dhcp_timeout=60 auto=true priority=critical vga=normal partman-lvm/
device_remove_lvm=true netcfg/choose_interface=auto debian-installer/allow_unauthenticated=true
preseed/url=http://192.168.210.82/nsx-edge/preseed.cfg mirror/country=manual mirror/http/
hostname=192.168.210.82 nsx-kickstart/url=http://192.168.210.82/nsx-edge/nsxcli.install mirror/
http/directory=/nsx-edge initrd=ubuntu-installer/amd64/initrd.gz mirror/suite=xenial --
```

18 Add the following lines to the `/etc/dhcp/dhcpd.conf` file.

Replace 192.168.210.82 with the IP address of your DHCP server.

```
allow booting;
allow bootp;

next-server 192.168.210.82; #Replace this IP address
filename "pxelinux.0";
```

19 Restart the DHCP service.

```
sudo service isc-dhcp-server restart
```

Note If an error is returned, for example: "stop: Unknown instance: start: Job failed to start", run `sudo /etc/init.d/isc-dhcp-server stop` and then `sudo /etc/init.d/isc-dhcp-server start`. The `sudo /etc/init.d/isc-dhcp-server start` command returns information about the source of the error.

What to do next

Install NSX Edge on bare metal using an ISO file. See Install NSX Edge via ISO File on Bare Metal or Install NSX Edge via ISO File as a Virtual Appliance.

Configuring Bare Metal Server to Use NSX-T Data Center

6

To use NSX-T Data Center on a bare metal server you must install supported third-party packages.

NSX-T Data Center supports the bare metal server in two ways: as a host transport node and as a host for NSX Manager.

Make sure that you have the supported bare metal server versions. See Bare Metal Server System Requirements.

Note If your NSX Edges are in VM form factor and you intend to use the NSX DHCP service (deployed on VLAN-based logical switch), you must set the forged transmits option to Accept on the baremetal hosts on which the NSX Edges are deployed. See Forged Transmits in the vSphere product documentation.

This chapter includes the following topics:

- Install Third-Party Packages on a Bare Metal Server

- Create Application Interface for Bare Metal Server Workloads

Install Third-Party Packages on a Bare Metal Server

To prepare a bare metal server to be a fabric node, you must install some third-party packages.

Prerequisites

- Verify that the user performing the installation has administrative permission to do the following actions, some of which may require `sudo` permissions:

    - Download and untar the bundle.

    - Run `dpkg` or `rpm` commands for installing/uninstalling NSX components.

    - Execute `nsxcli` command for executing join management plane commands.

- Verify that the virtualization packages are installed.

    - Redhat or CentOS - `yum install libvirt-libs`

    - Ubuntu - `apt-get install libvirt0`

    - SUSE - `zypper install libvirt-libs`

Procedure

◆ On Ubuntu, run `apt-get install <package_name>` to install the third-party packages.

Ubuntu18.04	Ubuntu16.04
traceroute python-mako python-netaddr python-simplejson python-unittest2 python-yaml python-openssl dkms libvirt0	libunwind8 libgflags2v5 libgoogle-perftools4 traceroute python-mako python-simplejson python-unittest2 python-yaml python-netaddr libboost-filesystem1.58.0 libboost-chrono1.58.0 libgoogle-glog0v5 dkms libboost-date-time1.58.0 python-protobuf python-gevent libsnappy1v5 libleveldb1v5 libboost-program-options1.58.0 libboost-thread1.58.0 libboost-iostreams1.58.0 libvirt0

◆ On RHEL or CentOS, run `yum install` to install the third-party packages.

RHEL 7.4, 7.5, and 7.6	CentOS 7.4, 7.5, and 7.6
tcpdump boost-filesystem PyYAML boost-iostreams boost-chrono python-mako python-netaddr python-six gperftools-libs libunwind snappy boost-date-time c-ares redhat-lsb-core wget net-tools yum-utils lsof python-gevent libev python-greenlet libvirt-libs	tcpdump boost-filesystem PyYAML boost-iostreams boost-chrono python-mako python-netaddr python-six gperftools-libs libunwind snappy boost-date-time c-ares redhat-lsb-core wget net-tools yum-utils lsof python-gevent libev python-greenlet libvirt-libs

◆ On SUSE, run `zypper install <package_name>` to install the third-party packages manually.

SUSE 12.0

```
net-tools
tcpdump
python-simplejson
python-netaddr
python-PyYAML
python-six
libunwind
wget
lsof
libcap-progs
libvirt-libs
```

Create Application Interface for Bare Metal Server Workloads

You must configure NSX-T Data Center install Linux third-party packages before you create or migrate an application interface for bare metal server workloads.

NSX-T Data Center does not support Linux OS interface bonding. You must use Open vSwitch (OVS) bonding for Bare Metal Server Transport Nodes. See Knowledge Base article 67835 Bare Metal Server supports OVS bonding for Transport Node configuration in NSX-T.

Procedure

1 Install the required third-party packages.

See Install Third-Party Packages on a Bare Metal Server.

2 Configure the TCP and UDP ports.

See TCP and UDP Ports Used by ESXi, KVM Hosts, and Bare Metal Server.

3 Add a bare metal server to the NSX-T Data Center fabric and create a transport node.

See Create a Standalone Host or Bare Metal Server Transport Node.

4 Use the Ansible playbook to create an application interface.

See https://github.com/vmware/bare-metal-server-integration-with-nsxt.

Configure the NSX Manager Cluster

The following subsections describe how to configure the NSX Manager cluster, details cluster requirements, and provides recommendations for specific site deployments. They also describe how you can use vSphere HA with NSX-T Data Center to enable quick recovery if the host running the NSX Manager node fails.

This chapter includes the following topics:

- NSX Manager Cluster Requirements
- NSX Manager Cluster Requirements for Single, Dual, and Multiple Sites

NSX Manager Cluster Requirements

The following requirements apply to NSX Manager cluster configuration:

- In a production environment, the NSX Manager cluster must have three members to avoid an outage to the management and control planes.

    Each cluster member should be placed on a unique hypervisor host with three physical hypervisor hosts in total. This is required to avoid a single physical hypervisor host failure impacting the NSX control plane. It is recommended you apply anti-affinity rules to ensure that all three cluster members are running on different hosts.

    The normal production operating state is a three-node NSX Manager cluster. However, you can add additional, temporary NSX Manager nodes to allow for IP address changes.

    Important   As of NSX-T Data Center 2.4, the NSX Manager contains the NSX Central Control Plane process. This service is critical for the operation of NSX. If there is a complete loss of NSX Managers, or if the cluster is reduced from three NSX Managers to one NSX Manager, you will not be able to make topology changes to your environment, and vMotion of machines depending on NSX will fail.

■ For lab and proof-of-concept deployments where there are no production workloads, you can run a single NSX Manager to save resources. NSX Manager nodes can be deployed on either ESXi or KVM. However, mixed deployments of managers on both ESXi and KVM are not supported.

Important The number of sites in an NSX-T Data Center deployment can affect requirements. See NSX Manager Cluster Requirements for Single, Dual, and Multiple Sites.

NSX Manager Cluster Requirements for Single, Dual, and Multiple Sites

Your NSX Manager cluster configuration will vary depending on whether your deployment is for single, dual, or multiple sites.

You can use vSphere HA with NSX-T Data Center to enable quick recovery if the host running the NSX Manager node fails.

Note See *Creating and Using vSphere HA Clusters* in the vSphere product documentation.

Single Site Requirements and Recommendations

The following recommendations apply to single site NSX-T Data Center deployments.

■ It is recommended that you place your NSX Managers on different hosts to avoid a single host failure impacting multiple managers.

■ Maximum latency between NSX Managers is 10ms.

■ You can place NSX Managers in different vSphere clusters or in a common vSphere cluster.

■ It is recommended that you place NSX Managers in different management subnets or a shared management subnet. When using vSphere HA it is recommended to use a shared management subnet soNSX Managers that are recovered by vSphere can preserve their IP address.

■ It is recommended that you place NSX Managers on shared storage also. For vSphere HA, please review the requirements for that solution.

You can also use vSphere HA with NSX-T to provide recovery of a lost NSX Manager when the host where the NSX Manager is running fails.

Scenario example:

■ A vSphere cluster in which all three NSX Managers are deployed.

■ The vSphere cluster consists of four or more hosts:

    ■ Host-01 with nsxmgr-01 deployed

    ■ Host-02 with nsxmgr-02 deployed

    ■ Host-03 with nsxmgr-03 deployed

- Host-04 with no NSX Manager deployed

- vSphere HA is configured to recover any lost NSX Manager (e.g., nsxmgr-01) from any host (e.g., Host-01) to Host-04.

Thus, upon the loss of any hosts where a NSX Manager is running, vSphere recovers the lost NSX Manager on Host-04.

Dual Site Requirements and Recommendations

The following recommendations apply to dual site (Site A/Site B) NSX-T Data Center deployments.

- It is not recommended to deploy NSX Managers in a dual-site scenario without vSphere HA. In this scenario, one site requires the deployment of twoNSX Managers and the loss of that site will impact the operation of NSX-T Data Center.

- Deployment of NSX Managers in a dual site scenario with vSphere HA can be done with the following considerations:

  - A single stretched vSphere cluster contains all the hosts for NSX Managers.

  - All three NSX Managers are deployed to a common management subnet/VLAN to allow IP address preservation upon recovery of a lost NSX Managers.

  - For latency between sites, see the storage product requirements.

Scenario example:

- A vSphere cluster in which all three NSX Managers are deployed.

- The vSphere cluster consists of six or more hosts, with three hosts in Site A and three hosts in Site B.

- The three NSX Managers are deployed to distinct hosts with additional hosts for placement of recovered NSX Managers:

  Site A:

  - Host-01 with nsxmgr-01 deployed

  - Host-02 with nsxmgr-02 deployed

  - Host-03 with nsxmgr-03 deployed

  Site B:

  - Host-04 with no NSX Manager deployed

  - Host-05 with no NSX Manager deployed

  - Host-06 with no NSX Manager deployed

- vSphere HA is configured to recover any lost NSX Manager (e.g., nsxmgr-01) from any host (e.g., Host-01) in Site A to one of the hosts in Site B.

Thus, upon failure of Site A, vSphere HA will recover all NSX Managers to hosts in site B.

Important   You must you properly configure anti-affinity rules to prevent NSX Managers from being recovered to the same common host.

Multiple (Three or More) Site Requirements and Recommendations

The following recommendations apply to multiple-site (Site A/Site B/Site C) NSX-T Data Center deployments.

In a scenario with three or more sites, you can deploy NSX Managers with or without vSphere HA.

If you deploy without vSphere HA:

- It is recommended that you use separate management subnets or VLANs per site.

- Maximum latency between NSX Managers is 10ms.

Scenario example (three sites):

- Three separate vSphere clusters, one per site.

- At least one host per site running NSX Manager:

    - Host-01 with nsxmgr-01 deployed

    - Host-02 with nsxmgr-02 deployed

    - Host-03 with nsxmgr-03 deployed

Failure scenarios:

- Single site failure: Two remaining NSX Managers in other sites continue to operate. NSX-T Data Center is in a degraded state but still operational. It is recommended you manually deploy a third NSX Manager to replace the lost cluster member.

- Two site failure: Loss of quorum and therefore impact to NSX-T Data Center operations.

Recovery of NSX Managers may take as long as 20 minutes depending on environmental conditions such as CPU speed, disk performance, and other deployment factors.

Transport Zones and Transport Nodes

8

Transport zones and transport nodes are important concepts in NSX-T Data Center.

This chapter includes the following topics:

- Create Transport Zones
- Create an IP Pool for Tunnel Endpoint IP Addresses
- Enhanced Data Path
- Configuring Profiles
- Create a Standalone Host or Bare Metal Server Transport Node
- Manual Installation of NSX-T Data Center Kernel Modules
- NSX Edge Networking Setup
- Create an NSX Edge Transport Node
- Create an NSX Edge Cluster

Create Transport Zones

Transport zones dictate which hosts and, therefore, which VMs can participate in the use of a particular network. A transport zone does this by limiting the hosts that can "see" a logical switch —and, therefore, which VMs can be attached to the logical switch. A transport zone can span one or more host clusters.

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

NSX-T Data Center does not allow connection of VMs that are in different transport zones in the Layer 2 network. 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.

The overlay transport zone is used by both host transport nodes and NSX Edges. When a host or NSX Edge transport node is added to an overlay transport zone, an N-VDS is installed on the host or NSX Edge.

The VLAN transport zone is used by the NSX Edge and host transport nodes for its VLAN uplinks. When an NSX Edge is added to a VLAN transport zone, a VLAN N-VDS is installed on the NSX Edge.

The N-VDS allows for virtual-to-physical packet flow by binding logical router uplinks and downlinks to physical NICs.

When you create a transport zone, you must provide a name for the N-VDS that will be installed on the transport nodes when they are later added to this transport zone. The N-VDS name can be whatever you want it to be.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Transport Zones > Add**.

3   Enter a name for the transport zone and optionally a description.

4   Enter a name for the N-VDS.

5   Select an N-VDS mode.

   ▪   **Standard** mode that applies to all the supported hosts.

   ▪   **Enhanced Datapath** is a networking stack mode that applies to only transport nodes of ESXi host version 6.7 and later type that can belong in a transport zone.

6   If the N-VDS mode is set to Standard, select a traffic type.

   The options are **Overlay** and **VLAN**.

7   If the N-VDS mode is set to Enhanced Datapath, select a traffic type.

   The options are **Overlay** and **VLAN**.

   Note   In the enhanced datapath mode, only specific NIC configurations are supported. Make sure that you configure the supported NICs.

8   Enter one or more uplink teaming policy names. These named teaming policies can be used by logical switches attached to the transport zone. If the logical switches do not find a matching named teaming policy, then the default uplink teaming policy is used.

9   View the new transport zone on the **Transport Zones** page.

10 (Optional) You can also view the new transport zone with the GET `https://<nsx-mgr>/api/v1/transport-zones` API call.

```
{
  "cursor": "00369b661aed-1eaa-4567-9408-ccbcfe50b416tz-vlan",
  "result_count": 2,
  "results": [
    {
```

```
          "resource_type": "TransportZone",
          "description": "comp overlay transport zone",
          "id": "efd7f38f-c5da-437d-af03-ac598f82a9ec",
          "display_name": "tz-overlay",
          "host_switch_name": "overlay-hostswitch",
          "transport_type": "OVERLAY",
          "transport_zone_profile_ids": [
            {
              "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a",
              "resource_type": "BfdHealthMonitoringProfile"
            }
          ],
          "_create_time": 1459547126454,
          "_last_modified_user": "admin",
          "_system_owned": false,
          "_last_modified_time": 1459547126454,
          "_create_user": "admin",
          "_revision": 0,
          "_schema": "/v1/schema/TransportZone"
        },
        {
          "resource_type": "TransportZone",
          "description": "comp vlan transport zone",
          "id": "9b661aed-1eaa-4567-9408-ccbcfe50b416",
          "display_name": "tz-vlan",
          "host_switch_name": "vlan-uplink-hostwitch",
          "transport_type": "VLAN",
          "transport_zone_profile_ids": [
            {
              "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a",
              "resource_type": "BfdHealthMonitoringProfile"
            }
          ],
          "_create_time": 1459547126505,
          "_last_modified_user": "admin",
          "_system_owned": false,
          "_last_modified_time": 1459547126505,
          "_create_user": "admin",
          "_revision": 0,
          "_schema": "/v1/schema/TransportZone"
        }
      ]
    }
```

What to do next

Optionally, create a custom transport-zone profile and bind it to the transport zone. You can create custom transport-zone profiles using the POST /api/v1/transportzone-profiles API. There is no UI workflow for creating a transport-zone profile. After the transport-zone profile is created, you can find it to the transport zone with the PUT /api/v1/transport-zones/<transport-zone-id> API.

Create a transport node. See Create a Standalone Host or Bare Metal Server Transport Node.

Create an IP Pool for Tunnel Endpoint IP Addresses

You can use an IP pool for the tunnel endpoints. Tunnel endpoints are the source and destination IP addresses used in the external IP header to identify the hypervisor hosts originating and end the NSX-T Data Center encapsulation of overlay frames. You can also use either DHCP or manually configured IP pools for tunnel endpoint IP addresses.

If you are using both ESXi and KVM hosts, one design option might be to use two different subnets for the ESXi tunnel endpoint IP pool (sub_a) and the KVM tunnel endpoint IP Pool (sub_b). In this case, on the KVM hosts a static route to sub_a must be added with a dedicated default gateway.

An example of the resulting routing table on an Ubuntu host where sub_a = 192.168.140.0 and sub_b = 192.168.150.0. (The management subnet, for example, might be 192.168.130.0).

Kernel IP routing table:

```
Destination         Gateway          Genmask          Iface
0.0.0.0             192.168.130.1    0.0.0.0          eth0
192.168.122.0       0.0.0.0          255.255.255.0    virbr0
192.168.130.0       0.0.0.0          255.255.255.0    eth0
192.168.140.0       192.168.150.1    255.255.255.0    nsx-vtep0.0
192.168.150.0       0.0.0.0          255.255.255.0    nsx-vtep0.0
```

The route can be added in at least two different ways. Of these two methods, the route persists after host reboot only if you add the route by editing the interface. Adding a route using the route add command does not persist after a host reboot.

```
route add -net 192.168.140.0 netmask 255.255.255.0 gw 192.168.150.1 dev nsx-vtep0.0
```

In /etc/network/interfaces before "up ifconfig nsx-vtep0.0 up" add this static route:

```
post-up route add -net 192.168.140.0 netmask 255.255.255.0 gw 192.168.150.1
```

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **Advanced Networking & Security > Inventory > Groups > IP Pools > Add**.

3 Enter the IP pool details.

Option	Parameter Example
Name and Description	Enter the IP pool and optional description.
IP Ranges	IP allocation ranges 192.168.200.100 - 192.168.200.115
Gateway	192.168.200.1

Option	Parameter Example
CIDR	Network address in a CIDR notation
	192.168.200.0/24
DNS Servers	Comma-separated list of DNS servers
	192.168.66.10
DNS Suffix	corp.local

Results

The IPv4 or IPv6 address pool is listed on the IP pool page.

You can also use the GET `https://<nsx-mgr>/api/v1/pools/ip-pools` API call to view the IP pool list.

What to do next

Create an uplink profile. See Create an Uplink Profile.

Enhanced Data Path

Enhanced data path is a networking stack mode, which when configured provides superior network performance. It is primarily targeted for NFV workloads, which requires the performance benefits provided by this mode.

The N-VDS switch can be configured in the enhanced data path mode only on an ESXi host. ENS also supports traffic flowing through Edge VMs.

In the enhanced data path mode, you can configure:

- Overlay traffic
- VLAN traffic

Supported VMkernel NICs

With NSX-T Data Center supporting multiple ENS host switches, the maximum number of VMkernel NICs supported per host is 32.

High-Level Process to Configure Enhanced Data Path

As a network administrator, before creating transport zones supporting N-VDS in the enhanced data path mode, you must prepare the network with the supported NIC cards and drivers. To improve network performance, you can enable the Load Balanced Source teaming policy to become NUMA node aware.

The high-level steps are as follows:

1   Use NIC cards that support the enhanced data path.

See VMware Compatibility Guide to know NIC cards that support enhanced data path.

On the VMware Compatibility Guide page, under the **IO devices** category, select **ESXi 6.7**, IO device Type as **Network**, and feature as **N-VDS Enhanced Datapath**.

2   Download and install the latest NIC drivers from the My VMware page.

a   Go to **Drivers & Tools** > **Driver CDs**.

b   Download NIC drivers:

VMware ESXi 6.7 ixgben-ens 1.1.3 NIC Driver for Intel Ethernet Controllers 82599, x520, x540, x550, and x552 family

VMware ESXi 6.7 i40en-ens 1.1.3 NIC Driver for Intel Ethernet Controllers X710, XL710, XXV710, and X722 family

3   Create an uplink policy.

See Create an Uplink Profile.

4   Create a transport zone with N-VDS in the enhanced data path mode.

See Create Transport Zones.

Note   ENS transport zones configured for overlay traffic: For a Microsoft Windows virtual machine running VMware tools version earlier to version 11.0.0 and vNIC type is VMXNET3, ensure MTU is set to 1500. For a Microsoft Windows virtual machine running vSphere 6.7 U1 and VMware tools version 11.0.0 and later, ensure MTU is set to a value less than 8900. For virtual machines running other supported OSes, ensure the virtual machine MTU is set to a value less than 8900.

5   Create a host transport node. Configure the enhanced data path N-VDS with logical cores and NUMA nodes.

See Create a Standalone Host or Bare Metal Server Transport Node.

Load Balanced Source Teaming Policy Mode Aware of NUMA

The Load Balanced Source teaming policy mode defined for an enhanced datapath N-VDS becomes aware of NUMA when the following conditions are met:

- The **Latency Sensitivity** on VMs is **High**.

- The network adapter type used is VMXNET3.

If the NUMA node location of either the VM or the physical NIC is not available, then the Load Balanced Source teaming policy does not consider NUMA awareness to align VMs and NICs.

The teaming policy functions without NUMA awareness in the following conditions:

- The LAG uplink is configured with physical links from multiple NUMA nodes.

- The VM has affinity to multiple NUMA nodes.

- The ESXi host failed to define NUMA information for either VM or physical links.

ENS Support for SCTP Applications

In SCTP environments, NFV workloads use multi-homing and redundancy features to increase resiliency and reliability to the traffic running on applications. Multi-homing is the ability to support redundant paths from a source VM to a destination VM.

Depending upon the number of physical NICs available to be used as an uplink for an overlay or VLAN network, those many redundant network paths are available for a VM to send traffic over to the target VM. The redundant paths are used when the pinned pNIC to a logical switch fails. So, traffic routed over SCTP protocol is provided redundant network paths by the Enhanced Data Path N-VDS.

Figure 8-1. ENS Traffic Running On SCTP Applications

Overlay Traffic on Logical Switch
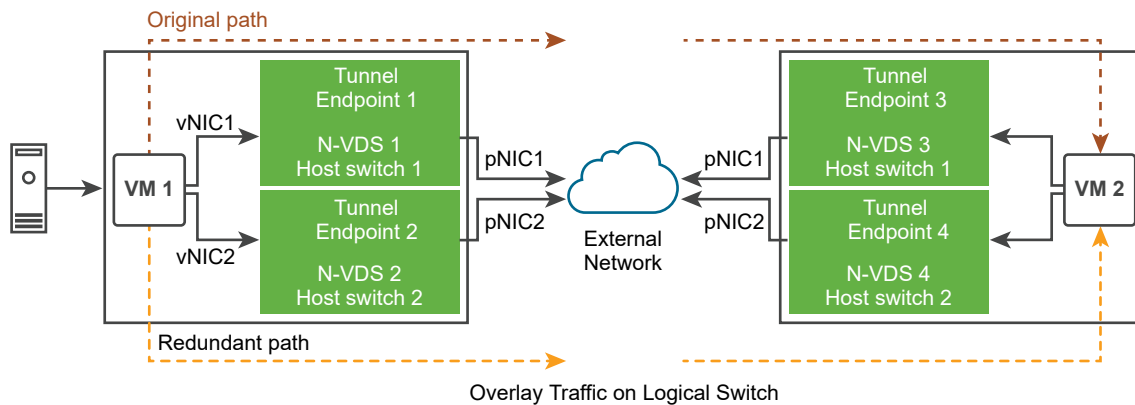
The high-level tasks are:

1   Prepare host as an NSX-T Data Center transport node.

2   Prepare VLAN or Overlay Transport Zone with two N-VDS switches in Enhanced Data Path mode.

3   On N-VDS 1, pin the first physical NIC to the switch.

4   On N-VDS 2, pin the second physical NIC to the switch.

The N-VDS in enhanced data path mode ensures that if pNIC1 becomes unavailable, then traffic from VM 1 is routed through the redundant path - vNIC 1 → tunnel endpoint 2 → pNIC 2 → VM 2. Note that vNIC1 of VM 1 and VM 2 are on one subnet. Similarly, vNIC2 of VM 1 and VM 2 are on another subnet.

Configuring Profiles

Profiles allow you to consistently configure identical capabilities for network adapters across multiple hosts or nodes.

Profiles are containers for the properties or capabilities that you want your network adapters to have. Instead of configuring individual properties or capabilities for each network adapter, you can specify the capabilities in the profiles, which you can then apply across multiple hosts or nodes.

Create an Uplink Profile

An uplink is a link from the NSX Edge nodes to the top-of-rack switches or NSX-T Data Center logical switches. A link is from a physical network interface on an NSX Edge node to a switch.

An uplink profile defines policies for the uplinks. The settings defined by uplink profiles can include teaming policies, active and standby links, transport VLAN ID, and MTU setting.

Configuring uplinks for VM appliance-based NSX Edge nodes and Host Transport nodes:

- If the Failover teaming policy is configured for an uplink profile, then you can only configure a single active uplink in the teaming policy. Standby uplinks are not supported and must not be configured in the failover teaming policy. When you install NSX Edge as a virtual appliance or host transport node, use the default uplink profile.

- If the Load Balanced Source teaming policy is configured for an uplink profile, then you can configure multiple active uplinks on the same N-VDS. Each uplink is associated with one physical NIC with a distinct name and IP address. The IP address assigned to an uplink endpoint is configurable using IP Assignment for the N-VDS.

You must use the **Load Balanced Source** teaming policy for traffic load balancing.

Prerequisites

- See NSX Edge network requirements in NSX Edge Installation.

- Each uplink in the uplink profile must correspond to an up and available physical link on your hypervisor host or on the NSX Edge node.

   For example, your hypervisor host has two physical links that are up: vmnic0 and vmnic1. Suppose vmnic0 is used for management and storage networks, while vmnic1 is unused. This might mean that vmnic1 can be used as an NSX-T Data Center uplink, but vmnic0 cannot. To do link teaming, you must have two unused physical links available, such as vmnic1 and vmnic2.

   For an NSX Edge, tunnel endpoint and VLAN uplinks can use the same physical link. For example, vmnic0/eth0/em0 might be used for your management network and vmnic1/eth1/em1 might be used for your fp-ethX links.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Profiles > Uplink Profiles > Add**.

3 Complete the uplink profile details.

Option	Description
Name and Description	Enter an uplink profile name. Add an optional uplink profile description.
LAGs	(Optional) In the LAGs section, click **Add** for Link aggregation groups (LAGs) using Link Aggregation Control Protocol (LACP) for the transport network. **Note** For LACP, multiple LAG is not supported on KVM hosts. The active and standby uplink names you create can be any text to represent physical links. These uplink names are referenced later when you create transport nodes. The transport node UI/API allows you to specify which physical link corresponds to each named uplink. Possible LAG hashing mechanism options: ■ Source MAC address ■ Destination MAC address ■ Source and destination MAC address ■ Source and destination IP address and VLAN ■ Source and destination MAC address, IP address, and TCP/UDP port
Teamings	In the Teaming section, you can either enter a default teaming policy or you can choose to enter a named teaming policy. Click **Add** to add a naming teaming policy. A teaming policy defines how N-VDS uses its uplink for redundancy and traffic load balancing. You can configure a teaming policy in the following modes: ■ **Failover Order**: An active uplink is specified along with an optional list of standby uplinks. If the active uplink fails, the next uplink in the standby list replaces the active uplink. No actual load balancing is performed with this option. ■ **Load Balance Source**: A list of active uplinks is specified, and each interface on the transport node is pinned to one active uplink. This configuration allows use of several active uplinks at the same time. **Note** ■ On KVM hosts: Only Failover Order teaming policy is supported, whereas Load Balance Source and Load Balance Source MAC teaming policies are not supported. ■ On NSX Edge: For default teaming policy, Load Balance Source and Failover Order teaming policies are supported. For named teaming policy, only Failover Order policy is supported. ■ On ESXi hosts: Load Balance Source MAC, Load Balance Source, and Failover Order teaming policies are supported. (ESXi hosts and NSX Edge) You can define the following policies for a transport zone: ■ A Named teaming policy for every VLAN-based logical switch or segment. ■ A Default teaming policy for the entire N-VDS.

Option	Description
	Named teaming policy: A named teaming policy means that for every VLAN-based logical switch or segment, you can define a specific teaming policy mode and uplinks names. This policy type gives you the flexibility to select specific uplinks depending on the traffic steering policy, for example, based on bandwidth requirement.
	■ If you define a named teaming policy, N-VDS uses that named teaming policy if it is attached to the VLAN-based transport zone and finally selected for specific VLAN-based logical switch or segment in the host.
	■ If you do not define any named teaming policies, N-VDS uses the default teaming policy.

4   Enter a Transport VLAN value. The transport VLAN set in the uplink profile tags overlay traffic only and the VLAN ID is used by the TEP endpoint.

5   Enter the MTU value.

The uplink profile MTU default value is 1600.

The global physical uplink MTU configures the MTU value for all the N-VDS instances in the NSX-T Data Center domain. If the global physical uplink MTU value is not specified, the MTU value is inferred from the uplink profile MTU if configured or the default 1600 is used. The uplink profile MTU value can override the global physical uplink MTU value on a specific host.

The global logical interface MTU configures the MTU value for all the logical router interfaces. If the global logical interface MTU value is not specified, the MTU value is inferred from the tier-0 logical router. The logical router uplink MTU value can override on a specific port the global logical interface MTU value.

Results

In addition to the UI, you can also view the uplink profiles with the API call `GET /api/v1/host-switch-profiles`.

What to do next

Create a transport zone. See Create Transport Zones.

Configuring Network I/O Control Profiles

Use the Network I/O Control (NIOC) profile to allocate the network bandwidth to business-critical applications and to resolve situations where several types of traffic compete for common resources.

NIOC profile introduces a mechanism to reserve bandwidth for the system traffic based on the capacity of the physical adapters on a host. Version 3 of the Network I/O Control feature offers improved network resource reservation and allocation across the entire switch.

Network I/O Control version 3 for NSX-T Data Center supports the resource management of the system traffic related to virtual machines and to infrastructure services, such as vSphere Fault Tolerance. System traffic is strictly associated with an ESXi host.

Bandwidth Guarantee to System Traffic

Network I/O Control version 3 provisions bandwidth to the network adapters of virtual machines by using constructs of shares, reservation, and limit. These constructs can be defined in the NSX-T Data Center Manager UI. The bandwidth reservation for virtual machine traffic is also used in the admission control. When you power on a virtual machine, admission control utility verifies that enough bandwidth is available before placing a VM on a host that can provide the resource capacity.

Bandwidth Allocation for System Traffic

You can configure Network I/O Control to allocate a certain amount of bandwidth for traffic generated by vSphere Fault Tolerance, vSphere vMotion, virtual machines, and so on.

- Management Traffic: is traffic for a host management

- Fault Tolerance (FT) traffic: is traffic for failover and recovery.

- NFS Traffic: is traffic related to a file transfer in the network file system.

- vSAN Traffic: is traffic generated by virtual storage area network.

- vMotion Traffic: is traffic for computing resource migration.

- vSphere Replication Traffic: is traffic for replication.

- vSphere Data Protection Backup Traffic: is traffic generated by backup of data.

- Virtual machine Traffic: is traffic generated by virtual machines.

- iSCSI Traffic: is traffic for Internet Small Computer System Interface.

vCenter Server propagates the allocation from the distributed switch to each physical adapter on the hosts that are connected to the switch.

Bandwidth Allocation Parameters for System Traffic

By using several configuration parameters, the Network I/O Control service allocates the bandwidth to traffic from basic vSphere system features. Allocation Parameters for System Traffic.

Allocation Parameters for System Traffic

- Shares: Shares, from 1 to 100, reflect the relative priority of a system traffic type against the other system traffic types that are active on the same physical adapter. The relative shares assigned to a system traffic type and the amount of data transmitted by other system features determine the available bandwidth for that system traffic type.

- Reservation: The minimum bandwidth, in Mbps, that must be guaranteed on a single physical adapter. The total bandwidth reserved among all system traffic types cannot exceed 75 percent of the bandwidth that the physical network adapter with the lowest capacity can provide. Reserved bandwidth that is unused becomes available to other types of system traffic. However, Network I/O Control does not redistribute the capacity that system traffic does not use to virtual machine placement.

- Limit: The maximum bandwidth, in Mbps or Gbps, that a system traffic type can consume on a single physical adapter.

Note You can reserve no more than 75 percent of the bandwidth of a physical network adapter.

For example, if the network adapters connected to an ESXi host are 10 GbE, you can only allocate 7.5 Gbps bandwidth to the various traffic types. You might leave more capacity unreserved. The host can allocate the unreserved bandwidth dynamically according to shares, limits, and use. The host reserves only the bandwidth that is enough for the operation of a system feature.

Configure Network I/O Control and Bandwidth Allocation for System Traffic on an N-VDS

To guarantee the minimum bandwidth to the system traffic running on NSX-T Data Center hosts, enable and configure a network resource management on an N-VDS.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Profiles > NIOC Profiles > Add**.

3   Enter the NIOC profile details.

Option	Description
Name and Description	Enter a NIOC profile name. You can optionally enter the profile details such as, the traffic types enabled.
Status	Toggle to enable the bandwidth allocations listed in the traffic resources.
Host Infra Traffic Resource	You can accept the default listed traffic resources. Click **Add** and enter your traffic resource to customize the NIOC profile. (Optional) Select an existing traffic type and click **Delete** to remove the resource from the NIOC profile.

The new NIOC profile is added to the NIOC profiles list.

Configure Network I/O Control and Bandwidth Allocation for System Traffic on an N-VDS Using APIs

You can use NSX-T Data Center APIs to configure the network and bandwidth for applications running on the host.

Procedure

1   Query the host to display both system-defined and user-defined host switch profiles.

2 GET `https://<nsx-mgr>/api/v1/host-switch-profiles?include_system_owned=true`.

The sample response displays the NIOC profile that is applied to the host.

```
{
  "description": "This profile is created for Network I/O Control (NIOC).",
  "extends": {
  "$ref": "BaseHostSwitchProfile"+
  },
  "id": "NiocProfile",
  "module_id": "NiocProfile",
  "polymorphic-type-descriptor": {
  "type-identifier": "NiocProfile"
  },
  "properties": {
  "_create_time": {
  "$ref": "EpochMsTimestamp"+,
  "can_sort": true,
  "description": "Timestamp of resource creation",
  "readonly": true
    },
  "_create_user": {
  "description": "ID of the user who created this resource",
  "readonly": true,
  "type": "string"
    },
  "_last_modified_time": {
  "$ref": "EpochMsTimestamp"+,
  "can_sort": true,
  "description": "Timestamp of last modification",
  "readonly": true
    },

  "_last_modified_user": {
  "description": "ID of the user who last modified this resource",
  "readonly": true,
  "type": "string"
    },

  "_links": {
  "description": "The server will populate this field when returning the resource. Ignored on PUT
and POST.",
  "items": {
  "$ref": "ResourceLink"+
    },

  "readonly": true,
  "title": "References related to this resource",
  "type": "array"
    },
  "_protection": {
  "description": "Protection status is one of the following:
    PROTECTED - the client who retrieved the entity is not allowed to modify it.
    NOT_PROTECTED - the client who retrieved the entity is allowed to modify it
    REQUIRE_OVERRIDE - the client who retrieved the entity is a super user and can modify it,
```

```
            but only when providing the request header X-Allow-Overwrite=true.
            UNKNOWN - the _protection field could not be determined for this entity.",
    "readonly": true,
    "title": "Indicates protection status of this resource",
    "type": "string"
       },

    "_revision": {
    "description": "The _revision property describes the current revision of the resource.
       To prevent clients from overwriting each other's changes, PUT operations must include the
                   current _revision of the resource,
       which clients should obtain by issuing a GET operation.
                        If the _revision provided in a PUT request is missing or stale, the operation
will be rejected.",
    "readonly": true,
    "title": "Generation of this resource config",
    "type": "int"
       },

    "_schema": {
    "readonly": true,
    "title": "Schema for this resource",
    "type": "string"
       },

    "_self": {
    "$ref": "SelfResourceLink"+,
    "readonly": true,
    "title": "Link to this resource"
       },

    "_system_owned": {
    "description": "Indicates system owned resource",
    "readonly": true,
    "type": "boolean"
       },

    "description": {
    "can_sort": true,
    "maxLength": 1024,
    "title": "Description of this resource",
    "type": "string"
       },

    "display_name": {
    "can_sort": true,
    "description": "Defaults to ID if not set",
    "maxLength": 255,
    "title": "Identifier to use when displaying entity in logs or GUI",
    "type": "string"
       },

    "enabled": {
    "default": true,
    "description": "The enabled property specifies the status of NIOC feature.
```

```
  When enabled is set to true, NIOC feature is turned on and the bandwidth allocations
    specified for the traffic resources are enforced.
  When enabled is set to false, NIOC feature is turned off and no bandwidth allocation is
guaranteed.

  By default, enabled will be set to true.",

  "nsx_feature": "Nioc",
  "required": false,
  "title": "Enabled status of NIOC feature",
  "type": "boolean"
    },

  "host_infra_traffic_res": {
  "description": "host_infra_traffic_res specifies bandwidth allocation for various traffic
resources.",
  "items": {
  "$ref": "ResourceAllocation"+
    },
  "nsx_feature": "Nioc",
  "required": false,
  "title": "Resource allocation associated with NiocProfile",
  "type": "array"
    },

  "id": {
  "can_sort": true,
  "readonly": true,
  "title": "Unique identifier of this resource",
  "type": "string"
    },

  "required_capabilities": {
  "help_summary":
                              "List of capabilities required on the fabric node if this profile is
used.
        The required capabilities is determined by whether specific features are enabled in the
profile.",
  "items": {
  "type": "string"
    },
  "readonly": true,
  "required": false,
  "type": "array"
    },

  "resource_type": {
  "$ref": "HostSwitchProfileType"+,
  "required": true
    },

  "tags": {
  "items": {
  "$ref": "Tag"+
```

```
     },

   "maxItems": 30,
   "title": "Opaque identifiers meaningful to the API user",
   "type": "array"
     }
   },
   "title": "Profile for Nioc",
   "type": "object"
   }
```

3   If a NIOC profile does not exist, create a NIOC profile.

POST https://<nsx-mgr>/api/v1/host-switch-profiles

```
 {
   "description": "Specify limit, shares and reservation for all kinds of traffic.
    Values for limit and reservation are expressed in percentage. And for shares,
    the value is expressed as a number between 1-100.\nThe overall reservation among all traffic
types should not exceed 75%.
    Otherwise, the API request will be rejected.",

   "id": "ResourceAllocation",
   "module_id": "NiocProfile",
   "nsx_feature": "Nioc",
   "properties": {
     "limit": {
       "default": -1.0,
       "description": "The limit property specifies the maximum bandwidth allocation for a given
       traffic type and is expressed in percentage. The default value for this
       field is set to -1 which means the traffic is unbounded for the traffic
       type. All other negative values for this property is not supported\nand will be rejected by
the API.",
       "maximum": 100,
       "minimum": -1,
       "required": true,
       "title": "Maximum bandwidth percentage",
       "type": "number"
     },

     "reservation": {
       "default": 0.0,
       "maximum": 75,
       "minimum": 0,
       "required": true,
       "title": "Minimum guaranteed bandwidth percentage",
       "type": "number"
     },

     "shares": {
       "default": 50,
       "maximum": 100,
       "minimum": 1,
       "required": true,
       "title": "Shares",
```

```
    "type": "int"
  },

  "traffic_type": {
    "$ref": "HostInfraTrafficType"+,
    "required": true,
    "title": "Resource allocation traffic type"
  }

},

"title": "Resource allocation information for a host infrastructure traffic type",
"type": "object"
```

4   Update the transport node configuration with the NIOC profile ID of the newly created NIOC
     profile.

PUT https://<nsx-mgr>/api/v1/transport-nodes/<TN-id>

```
{
    "resource_type": "TransportNode",
    "description": "Updated NSX configured Test Transport Node",
    "id": "77816de2-39c3-436c-b891-54d31f580961",
    "display_name": "NSX Configured TN",
    "host_switch_spec": {
      "resource_type": "StandardHostSwitchSpec",
      "host_switches": [
        {
        "host_switch_profile_ids": [
          {
            "value": "e331116d-f59e-4004-8cfd-c577aefe563a",
            "key": "UplinkHostSwitchProfile"
          },
          {
          "value": "9e0b4d2d-d155-4b4b-8947-fbfe5b79f7cb",
          "key": "LldpHostSwitchProfile"
          }
          {
           "value": "b0185099-8003-4678-b86f-edd47ca2c9ad",
           "key": "NiocProfile"
          }
          ],
          "host_switch_name": "nsxvswitch",
          "pnics": [
          {
           "device_name": "vmnic1",
           "uplink_name": "uplink1"
          }
          ],
          "ip_assignment_spec": {
          "resource_type": "StaticIpPoolSpec",
          "ip_pool_id": "ecddcdde-4dc5-4026-ad4f-8857995d4c92"
          }
          }
          ]
```

```
    },
     "transport_zone_endpoints": [
      {
      "transport_zone_id": "e14c6b8a-9edd-489f-b624-f9ef12afbd8f",
      "transport_zone_profile_ids": [
        {
          "profile_id": "52035bb3-ab02-4a08-9884-18631312e50a",
          "resource_type": "BfdHealthMonitoringProfile"
        }
      ]
      }
    ],

     "host_switches": [
      {
        "host_switch_profile_ids": [
          {
          "value": "e331116d-f59e-4004-8cfd-c577aefe563a",
          "key": "UplinkHostSwitchProfile"
         },
          {
          "value": "9e0b4d2d-d155-4b4b-8947-fbfe5b79f7cb",
          "key": "LldpHostSwitchProfile"
          }
        ],

        "host_switch_name": "nsxvswitch",
        "pnics": [
        {
          "device_name": "vmnic1",
          "uplink_name": "uplink1"
          }
        ],
        "static_ip_pool_id": "ecddcdde-4dc5-4026-ad4f-8857995d4c92"
      }
    ],
    "node_id": "41a4eebd-d6b9-11e6-b722-875041b9955d",
    "_revision": 0
  }
```

5 Verify that the NIOC profile parameters are updated in the `com.vmware.common.respools.cfg` file.

[root@ host:] `net-dvs -l`

```
        switch 1d 73 f5 58 99 7a 46 6a-9c cc d0 93 17 bb 2a 48 (vswitch)
  max ports: 2560
  global properties:

  com.vmware.common.opaqueDvs = true ,      propType = CONFIG
  com.vmware.nsx.kcp.enable = true ,      propType = CONFIG
  com.vmware.common.alias = nsxvswitch ,     propType = CONFIG
  com.vmware.common.uplinkPorts: uplink1    propType = CONFIG
  com.vmware.common.portset.mtu = 1600, propType = CONFIG
```

```
com.vmware.etherswitch.cdp = LLDP, listen  propType = CONFIG
com.vmware.common.respools.version = version3,   propType = CONFIG
com.vmware.common.respools.cfg:
netsched.pools.persist.ft:0:50:−1:255
netsched.pools.persist.hbr:0:50:−1:255
netsched.pools.persist.vmotion:0:50:−1:255
netsched.pools.persist.vm:0:100:−1:255
netsched.pools.persist.iscsi:0:50:−1:255
netsched.pools.persist.nfs:0:50:−1:255
netsched.pools.persist.mgmt:0:50:−1:255
netsched.pools.persist.vdp:0:50:−1:255
netsched.pools.persist.vsan:0:50:−1:255
propType = CONFIG
```

6 Verify NIOC profiles in the host kernel.

[root@ host:] /get /net/portsets/DvsPortset−1/ports/50335755/niocVnicInfo

```
Vnic NIOC Info
                   {
     Uplink reserved on:vmnic4
     Reservation in Mbps:200
     Shares:50
     Limit in Mbps:4294967295
     World ID:1001400726
     vNIC Index:0
     Respool Tag:0
     NIOC Version:3
     Active Uplink Bit Map:15
     Parent Respool ID:netsched.pools.persist.vm
   }
```

7 Verify the NIOC profile information.

[root@ host:] /get /net/portsets/DvsPortset−1/uplinks/vmnic4/niocInfo

```
Uplink NIOC Info
          {
   Uplink device:vmnic4
   Link Capacity in Mbps:750
   vm respool reservation:275
   link status:1
   NetSched Ready:1
   Infrastructure reservation:0
   Total VM reservation:200
   Total vnics on this uplink:1
   NIOC Version:3
   Uplink index in BitMap:0
   }
```

Results

NIOC profile is configured with a pre-defined bandwidth allocation for applications running on NSX-T Data Center hosts.

Add an NSX Edge Cluster Profile

The NSX Edge cluster profile defines the policies for the NSX Edge transport node.

Prerequisites

Verify that the NSX Edge cluster is available.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Profiles > Edge Cluster Profiles > Add**.

3 Enter the NSX Edge cluster profile details.

Option	Description
Name and Description	Enter a NSX Edge cluster profile name.
	You can optionally enter the profile details such as, the Bidirectional Forwarding Detection (BFD) setting.
BFD Probe Interval	Accept the default setting.
	BFD is detection protocol used to identify the forwarding path failures. You can set the interval timing for BFD to detect a forwarding path failure.
BFD Allowed Hops	Accept the default setting.
	You can set the number of multihop BFD sessions allowed for the profile.
BFD Declare Dead Multiple	Accept the default setting.
	You can set the number of number of times the BFD packet is not received before the session is flagged as down.
Stand By Relocation Threshold	Accept the default setting.

Add an NSX Edge Bridge Profile

The NSX Edge bridge profile defines the policies for the ESXi bridge cluster.

A bridge cluster is a collection of ESXi host transport nodes.

Prerequisites

- Verify that the NSX Edge cluster is available.
- Verify that the ESXi bridge cluster is available.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Profiles > Edge Bridge Profiles > Add**.

3 Enter the NSX Edge cluster profile details.

Option	Description
Name and Description	Enter a NSX Edge bridge cluster profile name. You can optionally enter the profile details such as, the primary and backup node details.
Edge Cluster	Select the NSX Edge cluster that you can to use.
Primary Node	Designate the preferred NSX Edge node from the cluster.
Backup Node	Designate the back up NSX Edge node if the primary node fails.
Failover Mode	Select either **Preemptive** or **Non-Preemptive** mode. The default HA mode is preemptive, which can slowdown traffic when the preferred NSX Edge node goes back online. The non-preemptive mode does not cause any traffic slowdown.

Add a Transport Node Profile

A transport node profile captures the configuration required to create a transport node. The transport node profile can be applied to an existing vCenter Server cluster to create transport nodes for the member hosts. Transport node profiles define transport zones, member hosts, N-VDS switch configuration including uplink profile, IP assignment, mapping of physical NICs to uplink virtual interfaces and so on.

Transport node creation begins when a transport node profile is applied to a vCenter Server cluster. NSX Manager prepares the hosts in the cluster and installs the NSX-T Data Center components on all the hosts. Transport nodes for the hosts are created based on the configuration specified in the transport node profile.

To delete a transport node profile, you must first detach the profile from the associated cluster. The existing transport nodes are not affected. New hosts added to the cluster are no longer automatically converted into transport nodes.

Considerations for Transport Node Profile Creation:

- You can add a maximum of four N-VDS switches for each configuration: enhanced N-VDS created for VLAN transport zone, standard N-VDS created for overlay transport zone, enhanced N-VDS created for overlay transport zone.

- There is no limit on the number of standard N-VDS switches created for VLAN transport zone.

- In a single host cluster topology running multiple standard overlay N-VDS switches and edge VM on the same host, NSX-T Data Center provides traffic isolation such that traffic going through the first N-VDS is isolated from traffic going through the second N-VDS and so on. The physical NICs on each N-VDS must be mapped to the edge VM on the host to allow the north-south traffic connectivity with the external world. Packets moving out of a VM on the first transport zone must be routed through an external router or an external VM to the VM on the second transport zone.

- Each N-VDS switch name must be unique. NSX-T Data Center does not allow use of duplicate switch names.

- Each transport zone ID must be unique. NSX-T Data Center does not allow use of duplicate IDs.

- You can add a maximum of 1000 transport zones in the transport node profile.

- To add a transport zone, it must be realized by any N-VDS present in the transport node profile.

Prerequisites

- Verify that the hosts are part of a vCenter Server cluster.

  vCenter Server must have at least one cluster.

- Verify that a transport zone is configured. See Create Transport Zones.

- Verify that a cluster is available. See Deploy NSX Manager Nodes to Form a Cluster from UI.

- Verify that an IP pool is configured, or DHCP must be available in the network deployment. See Create an IP Pool for Tunnel Endpoint IP Addresses.

- Verify that a compute manager is configured. See Add a Compute Manager.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Profiles > Transport Node Profiles > Add**.

3 Enter a name to identify the transport node profile.

  You can optionally add the description about the transport node profile.

4 Select the available transport zones and click the **>** button to include the transport zones in the transport node profile.

  Note You can add multiple transport zones.

5 Click the **N-VDS** tab and enter the switch details.

Option	Description
N-VDS Name	If the transport node is attached to a transport zone, then ensure the name entered for the N-VDS is the same as the N-VDS name specified in the transport zone. A transport node can be created without attaching it to a transport zone.
Associated Transport Zones	Shows the transport zones that are realized by the associated host switches. You cannot add a transport zone if it is not realized by any N-VDS in the transport node profile.

Option	Description
NIOC Profile	Select the NIOC profile from the drop-down menu.
	The bandwidth allocations specified in the profile for the traffic resources are enforced.
Uplink Profile	Select an existing uplink profile from the drop-down menu or create a custom uplink profile.
	Note The hosts in a cluster must have the same uplink profile.
	You can also use the default uplink profile.
LLDP Profile	By default, NSX-T only receives LLDP packets from a LLDP neighbor.
	However, NSX-T can be set to send LLDP packets to and receive LLDP packets from a LLDP neighbor.
IP Assignment	Select **Use DHCP**, **Use IP Pool**, or **Use Static IP List** to assign an IP address to virtual tunnel endpoints (VTEPs) of the transport node.
	If you select **Use Static IP List**, you must specify a list of comma-separated IP addresses, a gateway, and a subnet mask. All the VTEPs of the transport node must be in the same subnet otherwise bidirectional flow (BFD) session is not established.
IP Pool	If you selected **Use IP Pool** for an IP assignment, specify the IP pool name.
Physical NICs	Add physical NICs to the transport node. You can use the default uplink or assign an existing uplink from the drop-down menu.
	Click **Add PNIC** to configure additional physical NICs to the transport node.
	Note Migration of the physical NICs that you add in this field depends on how you configure **PNIC only Migration**, **Network Mappings for Install**, and **Network Mappings for Uninstall**.
	■ To migrate a used physical NIC (for example, by a standard vSwitch or a vSphere distributed switch) without an associated VMkernel mapping, ensure that **PNIC only Migration** is enabled. Otherwise, the transport node state remains in **partial success**, and the fabric node LCP connectivity fails to establish.
	■ To migrate a used physical NIC with an associated VMkernel network mapping, disable **PNIC only Migration** and configure the VMkernel network mapping.
	■ To migrate a free physical NIC, enable **PNIC only Migration**.

Option	Description

PNIC only Migration	Before setting this field, consider the following points:
	
	■ Know whether the physical NIC defined is a used NIC or a free NIC.
	
	■ Determine whether VMkernel interfaces of a host need to be migrated along with physical NICs.
	
	Set the field:
	
	■ Enable **PNIC only Migration** if you only want to migrate physical NICs from a VSS or DVS switch to an N-VDS switch.
	
	■ Disable **PNIC only Migration** if you want to migrate a used physical NIC and its associated VMkernel interface mapping. A free or available physical NIC is attached to the N-VDS switch when a VMkernel interface migration mapping is specified.
	
	On a host with multiple host switches:
	
	■ If all host switches are to migrate only PNICs, then you can migrate the PNICs in a single operation.
	
	■ If some hosts switches are to migrate VMkernel interfaces and the remaining host switches are to migrate only PNICs:
	
	  1 In the first operation, migrate only PNICs.
	
	  2 In the second operation, migrate VMkernel interfaces. Ensure that **PNIC only Migration** is disabled.
	
	Both PNIC only migration and VMkernel interface migration are not supported at the same time across multiple hosts.
	
	Note To migrate a management network NIC, configure its associated VMkernel network mapping and keep **PNIC only Migration** disabled. If you only migrate the management NIC, the host loses connectivity.
	
	For more information, see VMkernel Migration to an N-VDS Switch.

Option	Description
Network Mappings for Install	To migrate VMkernels to N-VDS switch during installation, map VMkernels to an existing logical switch. The NSX Manager migrates the VMkernel to the mapped logical switch on N-VDS.
	Caution Ensure that the management NIC and management VMkernel interface are migrated to a logical switch that is connected to the same VLAN that the management NIC was connected to before migration. If vmnic<*n*> and VMkernel<*n*> are migrated to a different VLAN, then connectivity to the host is lost.
	Caution For pinned physical NICs, ensure that the host switch mapping of physical NIC to a VMkernel interface matches the configuration specified in the transport node profile. As part of the validation procedure, NSX-T Data Center verifies the mapping and if the validation passes migration of VMkernel interfaces to an N-VDS switch is successful. It is also mandatory to configure the network mapping for uninstallation because NSX-T Data Center does not store the mapping configuration of the host switch after migrating the VMkernel interfaces to the N-VDS switch. If the mapping is not configured, connectivity to services, such as vSAN, can be lost after migrating back to the VSS or VDS switch.
	For more information, see VMkernel Migration to an N-VDS Switch .
Network Mappings for Uninstall	To revert the migration of VMkernels during uninstallation, map VMkernels to port groups on VSS or DVS, so that NSX Manager knows which port group the VMkernel must be migrated back to on the VSS or DVS. For a DVS switch, ensure that the port group is of the type `Ephemeral`.
	Caution For pinned physical NICs, ensure that the transport node profile mapping of physical NIC to VMkernel interface matches the configuration specified in the host switch. It is mandatory to configure the network mapping for uninstallation because NSX-T Data Center does not store the mapping configuration of the host switch after migrating the VMkernel interfaces to the N-VDS switch. If the mapping is not configured, connectivity to services, such as vSAN, can be lost after migrating back to the VSS or VDS switch.
	For more information, see VMkernel Migration to an N-VDS Switch.

6   To add another N-VDS switch, click **+ ADD N-VDS**.

7   Click **Save** to complete configuration.

What to do next

Apply the transport node profile to an existing vSphere cluster. See Configure a Managed Host Transport Node.

VMkernel Migration to an N-VDS Switch

To migrate VMkernel interfaces from a VSS or DVS switch to an N-VDS switch at a cluster-level, configure the transport node profile with network-mapping details required for migration (map VMkernel interfaces to logical switches). Similarly, to migrate VMkernel interfaces on a host node, configure the transport node configuration. To revert migrate VMkernel interfaces back to a VSS

or DVS switch, configure uninstall network-mapping (map logical ports to VMkernel interface) in the transport node profile to be realized during uninstallation.

During migration physical NICs currently in use are migrated to an N-VDS switch, while available or free physical NICs are attached to the N-VDS switch after migration.

Note Transport node profiles are applied to all member hosts of a cluster. But if you want to limit migration of VMkernel interfaces on specific hosts, you can directly configure the host. After migration, N-VDS handles traffic on the VLAN and overlay network for those interfaces attached to the N-VDS switch.

Important Configurations done to individual hosts are marked with the `Overridden` flag. Any further updates to the transport node profile are not applied to these overridden hosts. These hosts remain in overridden state until NSX-T Data Center is uninstalled.

In the following figure, if a host has only two physical NICs, you might want to assign both those NICs to the N-VDS for redundancy and their associated VMkernel interfaces so that the interfaces do not lose connectivity with the host.

Figure 8-2. Pre and Post Migration of Network Interfaces to an N-VDS
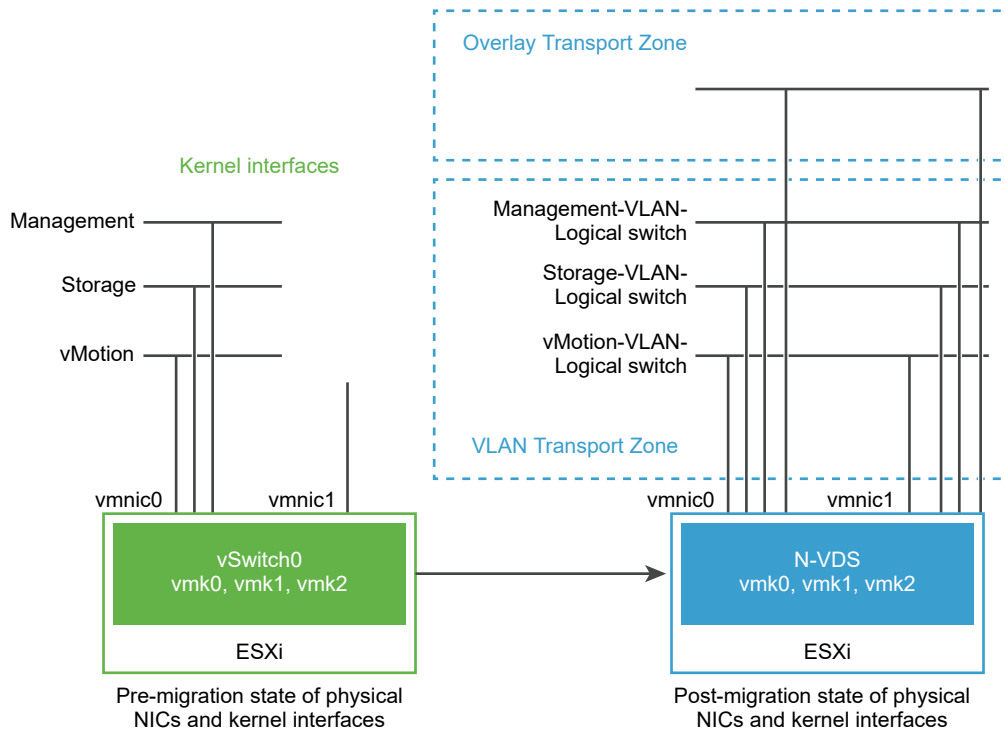
Before migration, the ESXi host has two uplinks derived from the two physical ports - vmnic0 and vmnic1. Here, vmnic0 is configured to be in an active state, attached to a VSS, whereas vmnic1 is unused. In addition, there are three VMkernel interfaces: vmk0, vmk1, and vmk2.

You can migrate VMkernel interfaces by using the NSX-T Data Center Manager UI or NSX-T Data Center APIs. See *NSX-T Data Center API Guide*.

Post migration, the vmnic0, vmnic1, and their VMkernel interfaces are migrated to the N-VDS switch. Both vmnic0 and vmnic1 are connected over VLAN and overlay transport zones.

Considerations for VMkernel Migration

- PNIC and VMkernel migration: Before you migrate pinned physical NICs and associated VMkernel interfaces to an N-VDS switch, make a note of the network-mapping (physical NICs to port group mapping) on the host switch.

- PNIC only migration: If you plan to only migrate PNICs, ensure that the management physical NIC connected to the management VMkernel interface is not migrated. It results in loss of connectivity with the host. For more details, see the **PNIC only Migration** field in Add a Transport Node Profile.

- Revert migration: Before you plan to revert migrate VMkernel interfaces to the VSS or DVS host switch for pinned physical NICs, ensure that you make a note of the network-mapping (physical NIC to port group mapping) on the host switch. It is mandatory to configure the transport node profile with the host switch mapping in the **Network Mapping for Uninstallation** field. Without this mapping, NSX-T Data Center does not know which port groups must the VMkernel interfaces be migrated back to. This situation can lead to loss of connectivity to the vSAN network.

- vCenter Server registration before migration: If you plan to migrate a VMkernel or PNIC connected to a DVS switch, ensure that a vCenter Server is registered with the NSX Manager.

- Match VLAN ID: After migration, the management NIC and management VMkernel interface must be on the same VLAN the management NIC was connected to before migration. If vmnic0 and vmk0 are connected to the management network and migrated to a different VLAN, then connectivity to the host is lost.

- Migration to VSS switch: Cannot migrate back two VMkernel interfaces to the same port group of a VSS witch.

- vMotion: Perform vMotion to move VM workloads to another host before VMkernel and/or PNIC migration. If migration fails, then workload VMs are not impacted.

- vSAN: If the vSAN traffic is running on the host, place the host in maintenance mode through vCenter Server and move VMs out of the host using vMotion functionality before VMkernel and/or PNIC migration.

- Migration: If a VMkernel is already connected to a target switch, it can still be selected to be migrated into the same switch. This property makes the VMK and/or PNIC migration operation idempotent. It helps when you want to migrate only PNICs into a target switch. As migration always requires at least one VMkernel and a PNIC, you select a VMkernel that is already migrated to a target switch when you migrate only PNICs into a target switch. If no VMkernel needs to be migrated, create a temp VMkernel through a vCenter Server in either the source switch or target switch. Then migrate it together with the PNICs, and delete the temp VMkernel through vCenter Server after the migration is finished.

- MAC sharing: If a VMkernel interface and a PNIC share the same MAC and they are in the same switch, they must be migrated together to the same target switch if they will be both used after migration. Always keep vmk0 and vmnic0 in the same switch.

  Check the MACs used by all VMKs and PNICs in the host by running the following commands:

  `esxcfg-vmknic -l`

  `esxcfg-nics -l`

- VIF logical ports created after migration: After you migrate VMkernel from a VSS or DVS switch to an N-VDS switch, a logical switch port of the type VIF is created on the NSX Manager. You must not create distributed firewall rules on these VIF logical switch ports.

Migrate VMkernel Interfaces to an N-VDS Switch

The high-level workflow to migrate VMkernel Interfaces to an N-VDS switch:

1 Create a logical switch if needed.

2 Power off VMs on the host from which VMkernel interfaces and PNICs are migrated to an N-VDS switch.

3 Configure a transport node profile with a network mapping that is used to migrate the VMkernel interfaces during the creation of transport nodes. Network mapping means mapping a VMkernel interface to a logical switch.

  For more details, see Add a Transport Node Profile.

4 Verify that the network adapter mappings in vCenter Server reflect a new association of the VMkernel switch with an N-VDS switch. In case of pinned physical NICs, verify the mapping in NSX-T Data Center reflects any VMkernels pinned to a physical NIC in the vCenter Server.

5 In NSX Manager, go to **Advanced Networking & Security** > **Networking** > **Switching**. On the **Switches** page, verify that the VMkernel interface is attached to the logical switch through a newly created logical port.

6 Go to **System** > **Nodes** > **Host Transport Node**. For each transport node, verify the status on the **Node Status** column is Success to confirm that the transport node configuration is successfully validated.

7 On the **Host Transport Node** page, verify the status on the **Configuration State** is Success to confirm that the host is successfully realized with the specified configuration.

After you migrate VMkernel interfaces and PNICs from a VDS to a N-VDS switch using NSX-T UI or transport node API, vCenter Server displays warnings for the VDS. If the host need be connected to the VDS, remove the host out of the VDS. The vCenter Server no longer displays any warning for VDS.

For details on errors that might encounter during migration, see VMkernel Migration Errors

Revert Migration of VMkernel Interfaces to a VSS or DVS Switch

The high-level workflow to revert migration of VMkernel Interfaces from an N-VDS switch to a VSS or DVS switch during NSX-T Data Center uninstallation:

1   On the ESXi host, power off VMs connected to the logical ports that hosts the VMkernel interface after migration.

2   Configure the transport node profile with network mapping that is used to migrate the VMkernel interfaces during the uninstallation process. Network mapping during uninstallation maps the VMkernel interfaces to a port group on VSS or DVS switch on the ESXi host.

   Note   Reverting migration of a VMkernel to a port group on a DVS switch, ensure that the port group type is set to `Ephemeral`.

   For more details, see Add a Transport Node Profile.

3   Verify the network adapter mappings in vCenter Server reflect a new association of the VMkernel switch with a port group of VSS or DVS switch.

4   In NSX Manager, go to **Advanced Networking & Security** > **Networking** > **Switching**. On the **Switches** page, verify that the logical switch containing VMkernel interfaces are deleted.

For details on errors that you might encounter during migration, see VMkernel Migration Errors

Update Host Switch Mapping

Important

■   Stateful hosts: Add and Update operations are supported. To update an existing mapping, you can add a new VMkernel interface entry to the network-mapping configuration. If you update the network mapping configuration of a VMkernel interface that is already migrated to the N-VDS switch, the updated network mapping is not realized on the host.

■   Stateless hosts: Add, Update, and Remove operations are supported. Any changes you make to the network-mapping configuration is realized after the host reboots.

   To update the VMkernel interfaces to a new logical switch, you can edit the transport node profile to apply the network mappings at a cluster level. If you only want the updates to be applied to a single host, configure the transport node using host-level APIs.

Note   After you update the transport node configuration for an individual host, then any new updates applied through the transport node profile are not applied to that host. That host state turns to `overriden`.

1   To update all hosts in a cluster, edit the **Network Mapping during Installation** field to update the VMkernel mapping to logical switches.

   For more details, see Add a Transport Node Profile.

2   Save the changes. Changes made to a transport node profile is automatically applied to all the member hosts of the cluster, except on hosts that are marked with the `overridden` state.

3   Similarly, to update an individual host, edit the VMkernel mapping in the transport node configuration.

Note   If you update the **Network Mapping during Installation** field with a new VMkernel mapping, then the same VMkernel interface must be added to the **Network Mapping during uninstallation** field.

For details on errors that you might encounter during migration, see VMkernel Migration Errors

Migrate VMkernel Interfaces on a Stateless Cluster

1   Prepare and configure a host as a reference host using transport node APIs.

2   Extract the host profile from the reference host.

3   In the vCenter Server, apply the host profile to the stateless cluster.

4   In NSX-T Data Center, apply the transport node profile to the stateless cluster.

5   Reboot each host of the cluster.

The cluster hosts might take several minutes to realize the updated states.

Migration Failure Scenarios

- If migration fails for some reason, the host attempts to migrate the physical NICs and VMkernel interfaces three times.

- If the migration still continues to fail, the host performs a rollback to the earlier configuration by retaining VMkernel connectivity with the management physical NIC, vmnic0.

- In case the rollback also fails such that the VMkernel configured to the management physical NIC was lost, you must reset the host.

Unsupported Migration Scenarios

The following scenarios are not supported:

- VMkernel interfaces from two different VSS or DVS switches are migrated at the same time.

- On stateful hosts, network mapping is updated to map VMkernel interface to another logical switch. For example, before migration the VMkernel is mapped to Logical Switch 1, and the VMkernel interface is mapped to Logical Switch 2.

VMkernel Migration Errors

You can encounter errors when migrating VMkernel interfaces and physical NICs from a VSS or DVS switch to an N-VDS switch or revert migrating interfaces to a VSS or DVS host switch.

Table 8-1. VMkernel Migration Errors

Error Code	Problem	Cause	Resolution
8224	Unable to find the host switch specified by the transport node configuration.	The host switch ID cannot be found.	■ Ensure that the transport zone is created with the host switch name and then create the transport node. ■ Ensure that a valid host switch is used in the transport node configuration.
8225	VMkernel migration is in progress.	Migration is in progress.	Wait for the migration to complete before performing another action.
8226	VMkernel migration is only supported on a ESXi host.	Migration is only valid for ESXi hosts.	Ensure that the host is a ESXi host before you initiate migration.
8227	VMkernel interface is not appended with the host switch name.	On a host with multiple host switches, NSX-T Data Center cannot identify association of each VMkernel interface with its host switch.	If the host has multiple N-VDS host switches, ensure the VMkernel interface is appended with the host switch name of the N-VDS the host is connected to. For example, the network mapping for uninstallation of a host with N-VDS host switch name nsxvswitch1 and VMkernel1 and another N-VDS host switch name nsxvswitch2 and VMkernel2 must be defined as follows: `device_name: VMkernel1@nsxvswitch1`, `destination_network: DPortGroup`.
8228	Host switch used in the `device_name` field not found on the host.	Incorrect host switch name.	Enter the correct host switch name.
8229	Transport node did not specify the transport zone of the logical switch.	Transport zone not added.	Add the transport zone to the transport node configuration.
8230	No physical NIC on the host switch.	There must be at least one physical NIC on the host switch.	Specify at least one physical NIC to an uplink profile and the VMkernel network mapping configuration to a logical switch.
8231	Host switch name does not match.	If the host switch name used in `vmk1@host_switch` does not match the host switch name used by the destination logical switch of the interface.	Ensure that the host switch name specified in the network mapping configuration matches the name used by the logical switch of the interface.

Table 8-1. VMkernel Migration Errors (continued)

Error Code	Problem	Cause	Resolution
8232	Logical switch not realized on the host.	Realization of logical switch on the host was unsuccessful.	Synchronize the host with the NSX Manager.
8233	Unexpected logical switch in the network interface mapping.	The network interface mapping for installation and uninstallation lists both logical switches and port groups.	Network mapping for installation must only contain logical switches as destination targets. Similarly, network mapping for uninstallation must only contain port groups as destination targets.
8294	Logical switch does not exist in the network interface mapping.	Logical switches not specified.	Ensure that logical switches are specified in the network interface mapping configuration.
8296	Host switch mismatch.	The network interface mapping for uninstallation is configured with the incorrect host switch name.	Ensure that the host switch name used in the mapping configuration matches the name entered on the host switch where the VMkernel interfaces reside on.
8297	Duplicate VMkernel.	Duplicate VMkernels are specified for migration.	Ensure that no duplicate VMkernel interfaces are specified in the installation or uninstallation mapping configuration.
8298	Mismatch of number of VMkernel interfaces and destinations.	Incorrect configuration.	Ensure that each VMkernel interface has a corresponding destination specified in the configuration.
8299	Cannot delete transport node as the VMkernel interface is using ports on N-VDS.	VMkernel interfaces are using ports from the N-VDS switch.	Revert the migration of all VMkernel interfaces from the N-VDS switch to a VSS/DVS switch. Then attempt to delete the transport node.
9412	VMkernel cannot be migrated from one N-VDS to another N-VDS.	Unsupported action.	Revert the migration of the VMkernel interface to a VSS or DVS switch. Then, you can migrate the VMkernel interface to another N-VDS switch.
9413	VMkernel interfaces cannot be migrated to a different logical switch.	On stateful hosts, a VMkernel connected to a logical switch cannot be migrated to another logical switch.	Revert the migration of the VMkernel from the logical switch to a VSS/DVS switch. Then, migrate the VMkernel to another logical switch on the N-VDS.

Table 8-1. VMkernel Migration Errors (continued)

Error Code	Problem	Cause	Resolution
9414	Duplicate VMkernel interfaces.	Duplicate VMkernel interfaces mapped in the installation and uninstallation mapping configuration.	Ensure that each VMkernel interface is unique in the installation and uninstallation mappings.
9415	Powered on VMs on the host.	With powered on VMs, migration does not proceed.	Power off the VMs on the host before you initiate migration of VMkernel interfaces.
9416	VMkernel cannot be found on the host.	Did not specify a VMkernel that exists on the host in the network mapping configuration.	Specify a VMkernel that exists in the network mapping configuration.
9417	Port group not found.	Did not specify a port group that exists on the host in the network mapping configuration.	Specify a port group that exists in the network mapping configuration.
9419	Logical switch not found during migration.	Did not find the logical switch defined in the network interface mapping configuration.	Specify a logical switch that exists in the network interface mapping configuration.
9420	Logical port not found during migration.	During migration, NSX-T Data Center does not find the ports created on the logical switch.	Ensure that no logical ports are deleted from the logical switch for migration to be successful.
9421	Host information missing to validate the migration process.	Unable to retrieve host information from inventory.	Retry the migration process.
9423	Pinned physical NICs to a VMkernel interface are not migrated to the correct host switch.	A pinned physical NIC was found in the environment but the VMkernel and physical NIC are not being migrated to the same host switch.	A physical NIC pinned with VMkernel interface must have a transport node configuration that maps the physical NIC with the VMkernel on the same host switch.
600	Object not found.	The specified transport zone used by the logical switch does not exist. The logical switch found in the VMK mapping destination cannot be found.	■ Specify a transport zone which exists in the environment. ■ Create the desired logical switch or use an existing VLAN logical switch.
8310	The logical switch type is incorrect.	The logical switch type is Overlay.	Create a VLAN logical switch.
9424	Cannot migrate if both PNIC only Migration and Network Mapping for install or uninstall settings are configured at the same time.	Migration progresses only when one of these settings is configured.	Ensure that either the PNIC only Migration or Network Mapping for install or uninstall setting is configured.

Create a Standalone Host or Bare Metal Server Transport Node

You must first add your ESXi host, KVM host, or bare metal server to the NSX-T Data Center fabric and then configure the transport node.

A fabric node is a node that has been registered with the NSX-T Data Center management plane and has NSX-T Data Center modules installed. For a host or bare metal server to be part of the NSX-T Data Center overlay, it must first be added to the NSX-T Data Center fabric.

A transport node is a node that participates in an NSX-T Data Center overlay or NSX-T Data Center VLAN networking.

For a KVM host or bare metal server, you can preconfigure the N-VDS, or you can have NSX Manager perform the configuration. For a ESXi host, NSX Manager always configures the N-VDS.

Note If you plan to create transport nodes from a template VM, make sure that there are no certificates on the host in `/etc/vmware/nsx/`. The netcpa agent does not create a certificate if a certificate exists.

Bare metal server supports an overlay and VLAN transport zone. You can use the management interface to manage the bare metal server. The application interface allows you to access the applications on the bare metal server.

Single physical NICs provide an IP address for both the management and application IP interfaces.

Dual physical NICs provide a physical NIC and a unique IP address for the management interface. Dual physical NICs also provide a physical NIC and a unique IP address for the application interface.

Multiple physical NICs in a bonded configuration provide dual physical NICs and a unique IP address for the management interface. Multiple physical NICs in a bonded configuration also provide dual physical NICs and an unique IP address for the application interface.

You can add a maximum of four N-VDS switches for each configuration: standard N-VDS created for VLAN transport zone, enhanced N-VDS created for VLAN transport zone, standard N-VDS created for overlay transport zone, enhanced N-VDS created for overlay transport zone.

In a single host cluster topology running multiple standard overlay N-VDS switches and edge VM on the same host, NSX-T Data Center provides traffic isolation such that traffic going through the first N-VDS is isolated from traffic going through the second N-VDS and so on. The physical NICs on each N-VDS must be mapped to the edge VM on the host to allow the north-south traffic connectivity with the external world. Packets moving out of a VM on the first transport zone must be routed through an external router or an external VM to the VM on the second transport zone.

Prerequisites

- The host must be joined with the management plane, and connectivity must be Up.

- A transport zone must be configured.

- An uplink profile must be configured, or you can use the default uplink profile.

- An IP pool must be configured, or DHCP must be available in the network deployment.

- At least one unused physical NIC must be available on the host node.

- Hostname

- Management IP address

- User name

- Password

- (Optional) (KVM) SHA-256 SSL thumbprint

- (Optional) (ESXi) SHA-256 SSL thumbprint

- Verify that the required third-party packages are installed. See Install Third-Party Packages on a KVM Host.

Procedure

1   (Optional) Retrieve the hypervisor thumbprint so that you can provide it when adding the host to the fabric.

   a   Gather the hypervisor thumbprint information.

      Use a Linux shell.

      ```
      # echo -n | openssl s_client -connect <esxi-ip-address>:443 2>/dev/null | openssl x509 -noout
      -fingerprint -sha256
      ```

      Use the ESXi CLI in the host.

      ```
      [root@host:~] openssl x509 -in /etc/vmware/ssl/rui.crt -fingerprint -sha256 -noout
      SHA256
      Fingerprint=49:73:F9:A6:0B:EA:51:2A:15:57:90:DE:C0:89:CA:7F:46:8E:30:15:CA:4D:5C:95:28:0A:9E:A
      2:4E:3C:C4:F4
      ```

   b   Retrieve the SHA-256 thumbprint from a KVM hypervisor, run the command in the KVM host.

      ```
      # awk '{print $2}' /etc/ssh/ssh_host_rsa_key.pub | base64 -d | sha256sum -b | sed 's/ .*$//'
      | xxd -r -p | base64
      ```

2   Select **System > Fabric > Nodes > Host Transport Nodes**.

3   From the Managed by field, select **Standalone Hosts** and click **+ Add**.

4 Enter the standalone host or bare metal server details to add to the fabric.

Option	Description
Name and Description	Enter the name to identify the standalone host or bare metal server. You can optionally add the description of the operating system used for the host or bare metal server.
IP Addresses	Enter the host or bare metal server IP address.
Operating System	Select the operating system from the drop-down menu. Depending on your host or bare metal server, you can select any of the supported operating systems. See System Requirements.
Username and Password	Enter the host user name and password.
SHA-256 Thumbprint	Enter the host thumbprint value for authentication. If you leave the thumbprint value empty, you are prompted to accept the server provided value. It takes a few seconds for NSX-T Data Center to discover and authenticate the host.

5 (Required) For a KVM host or bare metal server, select the N-VDS type.

Option	Description
NSX Created	NSX Manager creates the N-VDS. This option is selected by default.
Preconfigured	The N-VDS is already configured.

For a ESXi host, the N-VDS type is always set to **NSX Created**.

6 Enter the standard N-VDS details. Multiple N-VDS switches can be configured on a single host.

Option	Description
Transport Zone	Select the transport zone that this transport node belongs to from the drop-down menu.
N-VDS Name	Must be the same as the N-VDS name of the transport zone that this node belongs to.
NIOC Profile	For a ESXi host, select the NIOC profile from the drop-down menu.
Uplink Profile	Select an existing uplink profile from the drop-down menu or create a custom uplink profile. You can also use the default uplink profile.
LLDP Profile	By default, NSX-T only receives LLDP packets from a LLDP neighbor. However, NSX-T can be set to send LLDP packets to and receive LLDP packets from a LLDP neighbor.
IP Assignment	Select **Use DHCP**, **Use IP Pool**, or **Use Static IP List**. If you select **Use Static IP List**, you must specify a list of comma-separated IP addresses, a gateway, and a subnet mask.
IP Pool	If you selected **Use IP Pool** for IP assignment, specify the IP pool name.

Option	Description
Physical NICs	Add physical NICs to the transport node. You can use the default uplink or assign an existing uplink from the drop-down menu. Click **Add PNIC** to configure additional physical NICs to the transport node. **Note** Migration of the physical NICs that you add in this field depends on how you configure **PNIC only Migration**, **Network Mappings for Install**, and **Network Mappings for Uninstall**. ■ To migrate a used physical NIC (for example, by a standard vSwitch or a vSphere distributed switch) without an associated VMkernel mapping, ensure that **PNIC only Migration** is enabled. Otherwise, the transport node state remains in **partial success**, and the fabric node LCP connectivity fails to establish. ■ To migrate a used physical NIC with an associated VMkernel network mapping, disable **PNIC only Migration** and configure the VMkernel network mapping. ■ To migrate a free physical NIC, enable **PNIC only Migration**.
PNIC only Migration	Before setting this field, consider the following points: ■ Know whether the physical NIC defined is a used NIC or a free NIC. ■ Determine whether VMkernel interfaces of a host need to be migrated along with physical NICs. Set the field: ■ Enable **PNIC only Migration** if you only want to migrate physical NICs from a VSS or DVS switch to an N-VDS switch. ■ Disable **PNIC only Migration** if you want to migrate a used physical NIC and its associated VMkernel interface mapping. A free or available physical NIC is attached to the N-VDS switch when a VMkernel interface migration mapping is specified. On a host with multiple host switches: ■ If all host switches are to migrate only PNICs, then you can migrate the PNICs in a single operation. ■ If some hosts switches are to migrate VMkernel interfaces and the remaining host switches are to migrate only PNICs: 1 In the first operation, migrate only PNICs. 2 In the second operation, migrate VMkernel interfaces. Ensure that **PNIC only Migration** is disabled. Both PNIC only migration and VMkernel interface migration are not supported at the same time across multiple hosts. **Note** To migrate a management network NIC, configure its associated VMkernel network mapping and keep **PNIC only Migration** disabled. If you only migrate the management NIC, the host loses connectivity. For more information, see VMkernel Migration to an N-VDS Switch.

Option	Description
Network Mappings for Install	To migrate VMkernels to N-VDS switch during installation, map VMkernels to an existing logical switch. The NSX Manager migrates the VMkernel to the mapped logical switch on N-VDS.
	Caution Ensure that the management NIC and management VMkernel interface are migrated to a logical switch that is connected to the same VLAN that the management NIC was connected to before migration. If vmnic<*n*> and VMkernel<*n*> are migrated to a different VLAN, then connectivity to the host is lost.
	Caution For pinned physical NICs, ensure that the host switch mapping of physical NIC to a VMkernel interface matches the configuration specified in the transport node profile. As part of the validation procedure, NSX-T Data Center checks the mapping and if the validation passes migration of VMkernel interfaces to an N-VDS switch is successful. At the same time it is mandatory to configure the network mapping for uninstallation because NSX-T Data Center does not store the mapping configuration of the host switch after migrating the VMkernel interfaces to the N-VDS switch. If the mapping is not configured, connectivity to services, such as vSAN, can be lost after migrating back to the VSS or VDS switch.
	For more information, see VMkernel Migration to an N-VDS Switch.
Network Mappings for Uninstall	To revert the migration of VMkernels during uninstallation, map VMkernels to port groups on VSS or DVS, so that NSX Manager knows which port group the VMkernel must be migrated back to on the VSS or DVS. For a DVS switch, ensure the port group is of the type `Ephemeral`.
	Caution For pinned physical NICs, ensure that the transport node profile mapping of physical NIC to VMkernel interface matches the configuration specified in the host switch. It is mandatory to configure the network mapping for uninstallation because NSX-T Data Center does not store the mapping configuration of the host switch after migrating the VMkernel interfaces to the N-VDS switch. If the mapping is not configured, connectivity to services, such as vSAN, can be lost after migrating back to the VSS or VDS switch.
	For more information, see VMkernel Migration to an N-VDS Switch.

7 Enter the enhanced datapath N-VDS details. Multiple N-VDS switches can be configured on a single host.

Option	Description
N-VDS Name	Must be the same as the N-VDS name of the transport zone that this node belongs to.
IP Assignment	Select **Use DHCP**, **Use IP Pool**, or **Use Static IP List**. If you select **Use Static IP List**, you must specify a list of comma-separated IP addresses, a gateway, and a subnet mask.
IP Pool	If you selected **Use IP Pool** for an IP assignment, specify the IP pool name.

Option	Description
Physical NICs	Add physical NICs to the transport node. You can use the default uplink or assign an existing uplink from the drop-down menu. Click **Add PNIC** to configure additional physical NICs to the transport node. **Note** Migration of the physical NICs that you add in this field depends on how you configure **PNIC only Migration**, **Network Mappings for Install**, and **Network Mappings for Uninstall**. ■ To migrate a used physical NIC (for example, by a standard vSwitch or a vSphere distributed switch) without an associated VMkernel mapping, ensure that **PNIC only Migration** is enabled. Otherwise, the transport node state remains in **partial success**, and the fabric node LCP connectivity fails to establish. ■ To migrate a used physical NIC with an associated VMkernel network mapping, disable **PNIC only Migration** and configure the VMkernel network mapping. ■ To migrate a free physical NIC, enable **PNIC only Migration**.
Uplink	Select the uplink profile from the drop-down menu.
CPU Config	In the NUMA Node Index drop-down menu, select the NUMA node that you want to assign to an N-VDS switch. The first NUMA node present on the node is represented with the value 0. You can find out the number for NUMA nodes on your host by running the `esxcli hardware memory get` command. **Note** If you want to change the number of NUMA nodes that have affinity with an N-VDS switch, you can update the NUMA Node Index value.
	In the Lcore per NUMA node drop-down menu, select the number of logical cores that must be used by enhanced datapath. You can find out the maximum number of logical cores that can be created on the NUMA node by running the `esxcli network ens maxLcores get` command. **Note** If you exhaust the available NUMA nodes and logical cores, any new switch added to the transport node cannot be enabled for ENS traffic.

8 For a preconfigured N-VDS, provide the following details.

Option	Description
N-VDS External ID	Must be the same as the N-VDS name of the transport zone that this node belongs to.
VTEP	Virtual tunnel endpoint name.

9 View the connection status on the **Host Transport Nodes** page.

After adding the host or bare metal server as a transport node, the connection to NSX Manager changes to the Up status in 3-4 minutes.

10 Alternatively, view the connection status using CLI commands.

◆ For ESXi, type the `esxcli network ip connection list | grep 1234` command.

```
# esxcli network ip connection list | grep 1234
tcp   0   0  192.168.210.53:20514   192.168.110.34:1234   ESTABLISHED   1000144459   newreno
netcpa
```

◆ For KVM, type the command `netstat -anp --tcp | grep 1234`.

```
user@host:~$ netstat -anp --tcp | grep 1234
tcp  0   0 192.168.210.54:57794  192.168.110.34:1234   ESTABLISHED -
```

11 Verify that the NSX-T Data Center modules are installed on your host or bare metal server.

As a result of adding a host or bare metal server to the NSX-T Data Center fabric, a collection of NSX-T Data Center modules are installed on the host or bare metal server.

The modules on different hosts are packaged as follows:

- KVM on RHEL or CentOS - RPMs.

- KVM on Ubuntu - DEBs

- On ESXi, enter the command `esxcli software vib list | grep nsx`.

  The date is the day you performed the installation.

- On RHEL or CentOS, enter the command `yum list installed` or `rpm -qa`.

- On Ubuntu, enter the command `dpkg --get-selections`.

12 (Optional) Change the polling intervals of certain processes, if you have 500 hypervisors or more.

The NSX Manager might experience high CPU use and performance problems if there are more than 500 hypervisors.

a Use the NSX-T Data Center CLI command `copy file` or the API `POST /api/v1/node/file-store/<file-name>?action=copy_to_remote_file` to copy the `aggsvc_change_intervals.py` script to a host.

b Run the script, which is located in the NSX-T Data Center file store.

```
python aggsvc_change_intervals.py -m '<NSX ManagerIPAddress>' -u 'admin' -p '<password>' -i
900
```

c (Optional) Change the polling intervals back to their default values.

```
python aggsvc_change_intervals.py -m '<NSX ManagerIPAddress>' -u 'admin' -p '<password>' -r
```

Results

Note For an NSX-T Data Center created N-VDS, after the transport node is created, if you want to change the configuration, such as IP assignment to the tunnel endpoint, you must do it through the NSX Manager GUI and not through the CLI on the host.

What to do next

Migrate network interfaces from a vSphere Standard Switch to an N-VDS. See VMkernel Migration to an N-VDS Switch.

Configure a Managed Host Transport Node

If you have a vCenter Server, you can automate the installation and creation of transport nodes on all the NSX-T Data Center hosts instead of configuring manually.

If the transport node is already configured, then automated transport node creation is not applicable for that node.

Prerequisites

- Verify that all hosts in the vCenter Server are powered on.
- Verify that the system requirements are met. See System Requirements.
- Verify that a transport zone is available. See Create Transport Zones.
- Verify that a transport node profile is configured. See Add a Transport Node Profile.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Nodes > Host Transport Nodes**.

3 From the Managed By drop-down menu, select an existing vCenter Server.

The page lists the available vSphere clusters and/or ESXi hosts from the selected vCenter Server. You may need to expand a cluster to view the ESXi hosts.

4 Select a single host from the list and click **Configure NSX**.

The Configure NSX dialog box opens.

   a Verify the host name in the Host Details panel. Optionally, you can add a description.

   b Click **Next** to move to the **Configure NSX** panel.

   c Select the available transport zones and click the **>** button to include the transport zones in the transport node profile.

5 Verify the host name in the Host Details panel, and click **Next**.

Optionally, you can add a description.

6   In the **Configure NSX** panel, select the desired transport zones.

You can select more than one transport zone.

7   (Optional) View the ESXi connection status.

```
# esxcli network ip connection list | grep 1235
tcp   0   0  192.168.210.53:20514  192.168.110.34:1234   ESTABLISHED  1000144459  newreno  netcpa
```

8   From the Host Transport Node page, verify that the NSX Manager connectivity status of hosts in the cluster is Up and NSX-T Data Center configuration state is Success.

You can also see that the transport zone is applied to the hosts in the cluster.

9   (Optional) Remove an NSX-T Data Center installation and transport node from a host in the transport zone.

a   Select one or more hosts and click **Actions > Remove NSX**.

The uninstallation takes up to three minutes. Uninstallation of NSX-T Data Center removes the transport node configuration on hosts and the host is detached from the transport zone(s) and N-VDS switch. Any new host added to the vCenter Server cluster will not be automatically configured until the transport node profile is reapplied to the cluster.

10   (Optional) Remove a transport node from the transport zone.

a   Select a single transport node and click **Actions > Remove from Transport Zone**.

What to do next

Create a logical switch and assign logical ports. See the Advanced Switching section in the *NSX-T Data Center Administration Guide*.

Configure an ESXi Host Transport Node with Link Aggregation

This procedure describes how to create an uplink profile that has a link aggregation group configured, and how to configure an ESXi host transport node to use that uplink profile.

Prerequisites

- Familiarize yourself with the steps to create an uplink profile. See Create an Uplink Profile.

- Familiarize yourself with the steps to create a host transport node. See Create a Standalone Host or Bare Metal Server Transport Node.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Profiles > Uplink Profiles > Add**.

3   Enter a name and optionally a description.

For example, you enter the name `uplink-profile1`.

4   Under **LAGs**, click **Add** to add a link aggregation group.

For example, you add an LAG called `lag1` with 2 uplinks.

5   Under **Teamings**, select **Default Teaming**.

6   In the **Active Uplinks** field, enter the name of the LAG that you added in the step 4. In this example, the name is `lag1`.

7   Enter a value for the **Transport VLAN** and **MTU**.

8   Click **Add** at the bottom of the dialog box.

9   Under **Teamings**, click **Add** to add an entry for link aggregation.

10   Select **Fabric > Nodes > Host Transport Nodes > Add**.

11   In the **Host Details** tab, enter IP address, OS name, admin credentials, and SHA-256 thumbprint of the host.

12   In the **N-VDS** tab, select the uplink profile `uplink-profile1` that was created in step 3.

13   In the **Physical NICs** field, the physical NICs and uplinks dropdown list reflects the new NICs and uplink profile. Specifically, the uplinks `lag1-0` and `lag1-1`, corresponding to the LAG `lag1` that was created in step 4 are displayed. Select a physical NIC for `lag1-0` and a physical NIC for `lag1-1`.

14   Enter information for the other fields.

Fully Collapsed vSphere Cluster NSX-T Deployment

Configure NSX Manager, host transport nodes to run workload VMs, and NSX Edge VMs on a single cluster. Each host in the cluster provides two physical NICs that are configured for NSX-T.

Important   Deploy the fully collapsed single vSphere cluster topology starting with NSX-T 2.4.2 or 2.5 release.

The topology referenced in this procedure uses:

▪   vSAN configured with the hosts in the cluster.

▪   A minimum of two physical NICs per host.

▪   vMotion and Management VMkernel interfaces.

Figure 8-3. Topology: Single N-VDS Switch Managing Host Communication with NSX Edge and Guest VMs

Note Even if the host has four physical NICs, only two NICs can be used to deploy the fully collapsed topology. This procedure references physical NICs on the host as vmnic0 and vmnic1.

Prerequisites

- All the hosts must be part of a vSphere cluster.

- Each host has two physical NICs enabled.

- Register all hosts to a vCenter Server.

- Verify on the vCenter Server that shared storage is available to be used by the hosts.

- Ensure that the VLAN ID used for the TEP and the HOST TEP is NSX Edge different.

Procedure

1 Prepare four ESXi hosts with vmnic0 on vSS or vDS, vmnic1 is free.

2 On Host 1, install vCenter Server, configure a vSS/vDS port group, and install NSX Manager on the port group created on the host.

3 Prepare ESXi hosts 1, 2, 3 and 4 to be transport nodes.

    a Create VLAN transport zones with a named teaming policy. See Create Transport Zones.

    b Create an IP pool or DHCP for tunnel endpoint IP addresses for the hosts. See Create an IP Pool for Tunnel Endpoint IP Addresses.

    c Create an IP pool or DHCP for tunnel endpoint IP addresses for the Edge node. See Create an IP Pool for Tunnel Endpoint IP Addresses.

    d Create an uplink profile with a named teaming policy. See Create an Uplink Profile.

    e Configure hosts as a transport node by applying transport node profile. In this step, the transport node profile only migrates vmnic1, the unused physical NIC, to the N-VDS switch. After the transport node profile is applied to the cluster hosts, the N-VDS switch is created and vmnic1 is connected to the N-VDS switch. See Add a Transport Node Profile.

vmnic1 on all hosts are added to the N-VDS switch. So, out of the two physical NICs, one is migrated to the N-VDS switch. The vmnic0 interface is still connected to the vSS or vDS switch, which ensures connectivity to the host is available.

4   In the NSX Manager UI, create VLAN-backed segments for NSX Manager, vCenter Server, NSX Edge. Ensure to select the correct teaming policy for each of the VLAN-backed segments.

5 On Host 2, Host 3, and Host 4, you must migrate the vmk0 adapter and vmnic0 together from VSS/VDS to N-VDS switch. Update the NSX-T configuration on each host. While migrating ensure that vmnic0 is mapped to an active uplink.

6   In the vCenter Server, go to Host 2, Host 3, and Host 4, and verify that vmk0 adapter is connected to vmnic0 physical NIC on theN-VDS and must be reachable.

7   In the NSX Manager UI, go to Host 2, Host 3, and Host 4, and verify both pNICs are on the N-VDSswitch.

8   Create a logical segment and attach the NSX Manager to the logical segment. Wait for approximately 10 minutes for the cluster to form and verify that the cluster has formed.

9   On Host 2 and Host 3, from the NSX Manager UI, install NSX Manager.

10   Power off the first NSX Manager node. Wait for approximately 10 minutes.

11   Reattach the NSX Manager and vCenter Server to the previously created logical switch. On host 4, power on the NSX Manager. Wait for approximately 10 minutes to verify that the cluster is in a stable state. With the first NSX Manager powered off, perform cold vMotion to migrate the NSX Manager and vCenter Server from host 1 to host 4.

For vMotion limitations, see https://kb.vmware.com/s/article/56991.
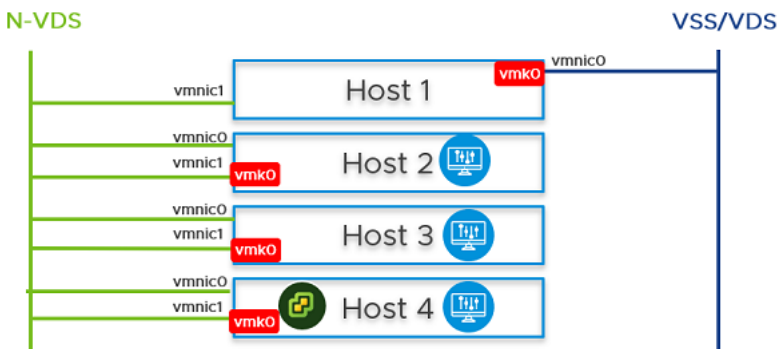
12   From the NSX Manager UI, go to Host 1, migrate vmk0 and vmnic0 together from VSS to N-VDS switch.

13 In the **Network Mapping for Install** field, ensure that the vmk0 adapter is mapped to the management logical segment on the N-VDS switch.

Network Mappings for Install

Host connectivity may be lost when vmnic0 and vmk0 are migrated.

Changing the Logical Switch for Stateful Host (Standalone or Clustered) will not affect and operation will fail.

+ ADD   🗑 DELETE

☐ VMkernel Adapter*	VLAN Segment/Logical Switch*
☐ vmk0	Seg-Vlan2200-ESXi-MGT

CANCEL   ADD
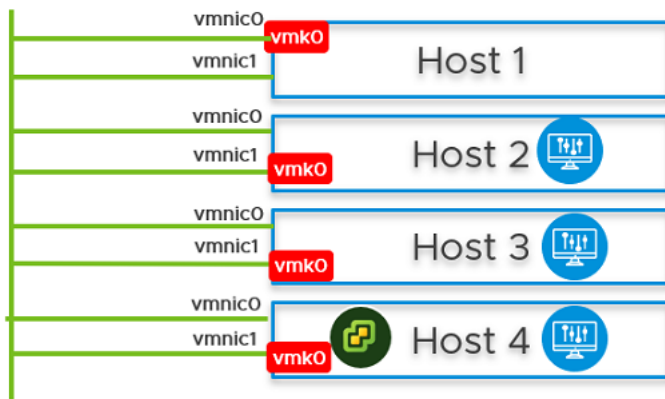
14 On Host 1, install the NSX Edge VM from the NSX Manager UI.

See Create an NSX Edge Transport Node.

15 Join the NSX Edge VM with the management plane.

See Join NSX Edge with the Management Plane.

16 To establish the north-south traffic connectivity, configure NSX Edge VM with an external router.

17 Verify that north-south traffic connectivity between the NSX Edge VM and the external router.

18 Set up and verify the BFD connectivity between NSX Manager and NSX Edge VM.

19 If there is a power failure scenario where the whole cluster is rebooted, the NSX-T management component might not come up and communicate with N-VDS. To avoid this scenario, perform the following steps:

Caution   Any API command that is incorrectly run results in a loss of connectivity with the NSX Manager.

Note   In a single cluster configuration, management components are hosted on an N-VDS switch as VMs. The N-VDS port to which the management component connects to by default is initialized as a blocked port due to security considerations. If there is a power failure requiring all the four hosts (minimum recommended) to reboot, the default reboot state the management VM port is in a blocked state. To avoid circular dependencies, it is recommended to create a port on N-VDS in the unblocked state. An unblocked port ensures that when the cluster is rebooted, the NSX-T management component can communicate with N-VDS to resume normal function.

At the end of the subtask, the migration command takes the :

■ UUID of the host node where the NSX Manager resides.

■ UUID of the NSX Manager VM and migrates it to the static logical port which is in an unblocked state.

If all the hosts are powered-off or powered-on or if an NSX Manager VM moves to another host, then after the NSX Manager comes back up it gets attached to the unblocked port, so preventing loss of connectivity with the management component of NSX-T.

a Go to **Advanced Networking & Security→ Switching**, select the MGMT-VLAN-Segment. In the **Overview** tab, find and copy the UUID. The UUID used in this example is, *c3fd8e1b-5b89-478e-abb5-d55603f04452*.

b To create logical ports that are initialized to be in `UNBLOCKED_VLAN` state, create four JSON files, three for NSX Managers and one for vCenter Server Appliance (VCSA). Replace the value for `logical_switch_id` with the UUID of the previously created MGMT-VLAN-Segment segment.

```
mgrhost.json
{
"admin_state": "UP",
"attachment": {
"attachment_type": "VIF",
"id": "nsxmgr-port-147"
},
"display_name": "NSX Manager Node 147 Port",
"init_state": "UNBLOCKED_VLAN",
"logical_switch_id": "c3fd8e1b-5b89-478e-abb5-d55603f04452"
}
```

c   Create the logical port for the Manager with an API client or using the curl command.

```
root@nsx-mgr-147:/var/CollapsedCluster# curl -X POST -k -u
'<username>:<password>' -H 'Content-Type:application/json' -d @mgr.json https://
localhost/api/v1/logical-ports
{
  "logical_switch_id" : "c3fd8e1b-5b89-478e-abb5-d55603f04452",
  "attachment" : {
    "attachment_type" : "VIF",
    "id" : "nsxmgr-port-147"
  },
  "admin_state" : "UP",
  "address_bindings" : [ ],
  "switching_profile_ids" : [ {
    "key" : "SwitchSecuritySwitchingProfile",
    "value" : "fbc4fb17-83d9-4b53-a286-ccdf04301888"
  }, {
    "key" : "SpoofGuardSwitchingProfile",
    "value" : "fad98876-d7ff-11e4-b9d6-1681e6b88ec1"
  }, {
    "key" : "IpDiscoverySwitchingProfile",
    "value" : "0c403bc9-7773-4680-a5cc-847ed0f9f52e"
  }, {
    "key" : "MacManagementSwitchingProfile",
    "value" : "1e7101c8-cfef-415a-9c8c-ce3d8dd078fb"
  }, {
    "key" : "PortMirroringSwitchingProfile",
    "value" : "93b4b7e8-f116-415d-a50c-3364611b5d09"
  }, {
    "key" : "QosSwitchingProfile",
    "value" : "f313290b-eba8-4262-bd93-fab5026e9495"
  } ],
  "init_state" : "UNBLOCKED_VLAN",
  "ignore_address_bindings" : [ ],
  "resource_type" : "LogicalPort",
  "id" : "02e0d76f-83fa-4839-a525-855b47ecb647",
  "display_name" : "NSX Manager Node 147 Port",
  "_create_user" : "admin",
  "_create_time" : 1574716624192,
  "_last_modified_user" : "admin",
  "_last_modified_time" : 1574716624192,
  "_system_owned" : false,
  "_protection" : "NOT_PROTECTED",
  "_revision" : 0
```


d   Move the NSX Manager to the statically created logical port.

e   To copy the NSX Manager VM Instance ID, go to Advanced Networking & Security →
     Inventory → Virtual Machines. Select the NSX Manager VM. In the **Overview** tab, find and
     copy the ID. The ID used in this example is *5028d756-d36f-719e-3db5-7ae24aa1d6f3*.

f   To find host ID where the NSX Manager is installed, go to **System** -> **Fabric** -> **Nodes** ->
     Host Transport Node. Select the host and click the **Overview** tab. Find and copy the host
     ID. The ID used in this example is *11161331-11f8-45c7-8747-34e7218b687f*.

g   Migrate the NSX Manager from the VM Network to the previously created logical-port on
     the MGMT-VLAN-Segment. The `vnic_migration_dest` value is the attachment ID of the
     ports created earlier for the NSX Manager.

```
root@nsx-mgr-147:/var/CollapsedCluster# curl -k -X PUT -u '<username>:<password>' -H
'Content-Type:application/json' -d @mgrhost.json
'https://localhost/api/v1/transport-nodes/11161331-11f8-45c7-8747-34e7218b687f?
vnic_migration_dest=nsxmgr-port-147'
```

h   In the NSX Manager UI, ensure that the statically created logical port is Up.

i   Repeat the preceding steps on every NSX Manager in the cluster.

Verify the Transport Node Status

Make sure that the transport node creation process is working correctly.

After creating a host transport node, the N-VDS gets installed on the host.

Procedure

1   Log in to the NSX-T Data Center.

2   Navigate to the Transport Node page and view the N-VDS status.

3   Alternatively, view the N-VDS on ESXi with the `esxcli network ip interface list`
     command.

     On ESXi, the command output should include a vmk interface (for example, vmk10) with a
     VDS name that matches the name you used when you configured the transport zone and the
     transport node.

```
# esxcli network ip interface list
...

vmk10
```

```
    Name: vmk10
    MAC Address: 00:50:56:64:63:4c
    Enabled: true
    Portset: DvsPortset-1
    Portgroup: N/A
    Netstack Instance: vxlan
    VDS Name: overlay-hostswitch
    VDS UUID: 18 ae 54 04 2c 6f 46 21-b8 ae ef ff 01 0c aa c2
    VDS Port: 10
    VDS Connection: 10
    Opaque Network ID: N/A
    Opaque Network Type: N/A
    External ID: N/A
    MTU: 1600
    TSO MSS: 65535
    Port ID: 67108895


 ...
```

If you are using the vSphere Client, you can view the installed N-VDS in the UI by selecting host **Configuration > Network Adapters**.

The KVM command to verify the N-VDS installation is `ovs-vsctl show`. Note that on KVM, the N-VDS name is nsx-switch.0. It does not match the name in the transport node configuration. This is by design.

```
# ovs-vsctl show
...
    Bridge "nsx-switch.0"
        Port "nsx-uplink.0"
            Interface "em2"
        Port "nsx-vtep0.0"
            tag: 0
            Interface "nsx-vtep0.0"
                type: internal
        Port "nsx-switch.0"
            Interface "nsx-switch.0"
                type: internal
    ovs_version: "2.4.1.3340774"
```

4   Check the transport node's assigned tunnel endpoint address.

The vmk10 interface receives an IP address from the NSX-T Data Center IP pool or DHCP, as shown here:

```
# esxcli network ip interface ipv4 get
Name   IPv4 Address    IPv4 Netmask    IPv4 Broadcast   Address Type  DHCP DNS
-----  -------------   -------------   --------------   -----------   --------
```

```
vmk0    192.168.210.53   255.255.255.0   192.168.210.255   STATIC          false
vmk1    10.20.20.53      255.255.255.0   10.20.20.255      STATIC          false
vmk10   192.168.250.3    255.255.255.0   192.168.250.255   STATIC          false
```

In KVM, you can verify the tunnel endpoint and IP allocation with the `ifconfig` command.

```
# ifconfig
...
nsx-vtep0.0 Link encap:Ethernet  HWaddr ba:30:ae:aa:26:53
          inet addr:192.168.250.4  Bcast:192.168.250.255  Mask:255.255.255.0
          ...
```

5 Check the API for transport node state information.

Use the GET `https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/state` API call. For example:

```
{
  "state": "success",
  "host_switch_states": [
    {
      "endpoints": [
        {
          "default_gateway": "192.168.250.1",
          "device_name": "vmk10",
          "ip": "192.168.250.104",
          "subnet_mask": "255.255.255.0",
          "label": 69633
        }
      ],
      "transport_zone_ids": [
        "efd7f38f-c5da-437d-af03-ac598f82a9ec"
      ],
      "host_switch_name": "overlay-hostswitch",
      "host_switch_id": "18 ae 54 04 2c 6f 46 21-b8 ae ef ff 01 0c aa c2"
    }
  ],
  "transport_node_id": "2d030569-5769-4a13-8918-0c309c63fdb9"
}
```

Visual Representation of N-VDS

You get a granular view of N-VDS at an individual host level. NSX-T Data Center provides a visual representation of the connectivity status between the uplink of the N-VDS and VMs associated to a transport zone. The objects represented visually include the teaming policy - uplink and physical NIC that provide connectivity to VMs. The other set of objects represented visually are

VMs, associated logical ports and switches, and status of VMs. The visual representation makes it easier to manage N-VDS.

Note Only ESXi hosts support visualization of N-VDS object.

Figure 8-4. N-VDS Visualization
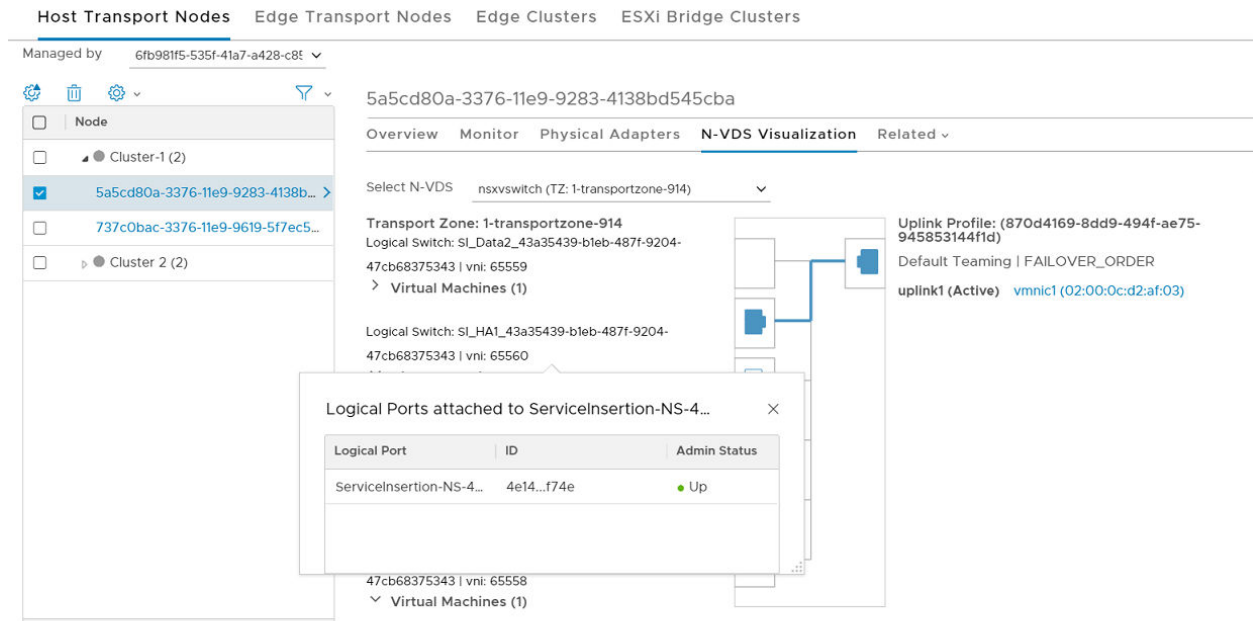
Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Nodes > Host Transport Nodes**.

3   From the Managed by field, select **Standalone Host** or a *compute manager*.

4   Select the host.

5   Click the **N-VDS Visualization** tab.

6   Select an N-VDS.

NSX-T visually represents uplink profiles connected to VMs, logical ports associated to VMs, logical switches connected to a transport zone.

7   To view uplink profiles connected to a VM and the logical port to which a VM is connected, select a VM.

NSX-T visually represents the connectivity between a VM and an uplink profile.

8   To view which VMs are connected to an uplink profile, select the uplink profile.

9   To view logical ports associated to a VM, expand the logical switch, click the VM.

The logical port details are displayed in a separate dialog box.

> **Note**   The admin status of a logical port is displayed on the dialog box. If the operational status is down it is not displayed on the dialog box.

Migrate ESXi VMkernel and Physical Adapters

After preparing a host as a transport node, you can make changes to the current migration configuration of VMkernel adapters and physical adapters.

Prerequisites

- Ensure that the host has at least one free physical adapter.

- Ensure that VMkernel adapters and port groups exist on the host.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Go to **System** > **Fabric** > **Host Transport Nodes**.

3   Select a transport node and click **Actions** > **Migrate ESX VMkernel and Physical Adapters**.

4   In the Migrate ESX VMkernel and Physical Adapters, enter the following details.

Field	Description
Direction	Make a selection: ■ **Migrate to Logical Switches**: To migrate VMkernel adapters from a VSS or VDS switch to an N-VDS switch in NSX-T Data Center. ■ **Migrate to Port Groups**: To migrate VMkernel adapters from an N-VDS switch to a VSS or VDS switch.
Select Switch	Select the switch from which you want to migrate the VMkernel adapters and physical adapters. You can select from the available switches.
Select VMkernel Adapters to Migrate	Click **Add** to enter the VMkernel adapter name and select destination as a logical switch or port group depending on where you want to migrate to.
Edit Physical Adapters in N-VDS	Click **Add** to enter the physical adapter name and map it to an uplink on the host switch.

5   Click **Save** to begin migration of VMkernel adapters and physical adapters.

Results

The updated VMkernel adapters and physical adapters are migrated to the N-VDS switch or revert migrated to the VSS or VDS switch in the ESXi host.

Manual Installation of NSX-T Data Center Kernel Modules

As an alternative to using the NSX-T Data Center **Fabric > Nodes > Hosts > Add** UI or the `POST /api/v1/fabric/nodes` API, you can install NSX-T Data Center kernel modules manually from the hypervisor command line.

Note You cannot manually install of NSX-T Data Center kernel modules on a bare metal server.

Manually Install NSX-T Data Center Kernel Modules on ESXi Hypervisors

To prepare hosts to participate in NSX-T Data Center, you must install NSX-T Data Center kernel modules on ESXi hosts. This allows you to build the NSX-T Data Center control-plane and management-plane fabric. NSX-T Data Center kernel modules packaged in VIB files run within the hypervisor kernel and provide services such as distributed routing, distributed firewall, and bridging capabilities.

You can download the NSX-T Data Center VIBs manually and make them part of the host image. The download paths can change for each release of NSX-T Data Center. Always check the NSX-T Data Center downloads page to get the appropriate VIBs.

Procedure

1 Log in to the host as root or as a user with administrative privileges

2 Navigate to the /tmp directory.

```
[root@host:~]: cd /tmp
```

3 Download and copy the nsx-lcp file into the /tmp directory.

4 Run the install command.

```
[root@host:/tmp]: esxcli software vib install -d /tmp/nsx-lcp-<release>.zip
Installation Result
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed: VMware_bootbank_nsx-aggservice_<release>, VMware_bootbank_nsx-da_<release>,
VMware_bootbank_nsx-esx-datapath_<release>, VMware_bootbank_nsx-exporter_<release>,
VMware_bootbank_nsx-host_<release>, VMware_bootbank_nsx-lldp_<release>, VMware_bootbank_nsx-
mpa_<release>, VMware_bootbank_nsx-netcpa_<release>, VMware_bootbank_nsx-python-
protobuf_<release>, VMware_bootbank_nsx-sfhc_<release>, VMware_bootbank_nsxa_<release>,
VMware_bootbank_nsxcli_<release>
   VIBs Removed:
   VIBs Skipped:
```

Depending on what was already installed on the host, some VIBs might be installed, some might be removed, and some might be skipped. A reboot is not required unless the command output says `Reboot Required: true`.

Results

As a result of adding an ESXi host to the NSX-T Data Center fabric, the following VIBs get installed on the host.

nsx-adf

(Automated Diagnostics Framework) Collects and analyzes performance data to produce both local (at host) and central (across datacenter) diagnoses of performance issues.

nsx-aggservice

Provides host-side libraries for NSX-T Data Center aggregation service. NSX-T Data Center aggregation service is a service that runs in the management-plane nodes and fetches runtime state from NSX-T Data Center components.

nsx-cli-libs

Provides the NSX-T Data Center CLI on hypervisor hosts.

nsx-common-libs

Provide some utilities classes such as AES, SHA-1, UUID, bitmap, and others.

nsx-context-mux

Provides NSX Guest Introspection relay functionality. Allows VMware Tools guest agents to relay guest context to inhouse and registered third-party partner appliances.

nsx-esx-datapath

Provides NSX-T Data Center data plane packet processing functionality.

nsx-exporter

Provides host agents that report runtime state to the aggregation service running in the management plane.

nsx-host

Provides metadata for the VIB bundle that is installed on the host.

nsx-metrics-libs

Provides metric utility classes for collecting daemon metrics.

nsx-mpa

Provides communication between NSX Manager and hypervisor hosts.

nsx-nestdb-libs

NestDB is a database that stores NSX configurations related to the host (desired/runtime state, etc).

nsx-netcpa

Provides communication between the central control plane and hypervisors. Receives logical networking state from the central control plane and programs this state in the data plane.

nsx-opsagent

Communicates operations agent executions (transport node realization, Link Layer Discovery Protocol - LLDP,traceflow, packet capture, etc.) with the management plane.

nsx-platform-client

Provides a common CLI execution agent, for centralized CLI and audit log collecting.

nsx-profiling-libs

Provides the functionality of profiling based on gpeftool which used for daemon process profiling.

nsx-proxy

Provides the only northbound contact point agent, which talks to the central control plane and management plane.

nsx-python-gevent

Contains Python Gevent.

nsx-python-greenlet

Contains Python Greenlet library (third party libraries).

nsx-python-logging

Contains the Python logs.

nsx-python-protobuf

Provides Python bindings for protocol buffers.

nsx-rpc-libs

This library provides nsx-rpc functionality.

nsx-sfhc

Service fabric host component (SFHC). Provides a host agent for managing the lifecycle of the hypervisor as a fabric host in the management plane's inventory. This provides a channel for operations such as NSX-T Data Center upgrade and uninstall and monitoring of NSX-T Data Center modules on hypervisors.

nsx-shared-libs

Contains the shared NSX libraries.

nsx-upm-libs

Provides unified profile management functionality for flattening client-side configuration and avoiding duplicate data transmission.

nsx-vdpi

Provides Deep Packet Inspection capabilities for NSX-T Data Center Distributed Firewall.

nsxcli

Provides the NSX-T Data Center CLI on hypervisor hosts.

vsipfwlib

Provides distributed firewall functionality.

To verify, you can run the `esxcli software vib list | grep nsx` and `esxcli software vib list | grep vsipfwlib` commands on the ESXi host. Or you can run the `esxcli software vib list | grep <yyyy–mm–dd>` command, where the date is the day that you performed the installation.

What to do next

Add the host to the NSX-T Data Center management plane. See Deploy NSX Manager Nodes to Form a Cluster Using CLI .

Manually Install NSX-T Data Center Kernel Modules on Ubuntu KVM Hypervisors

To prepare hosts to participate in NSX-T Data Center, you can manually install NSX-T Data Center kernel modules on Ubuntu KVM hosts. This allows you to build the NSX-T Data Center control-plane and management-plane fabric. NSX-T Data Center kernel modules packaged in DEB files run within the hypervisor kernel and provide services such as distributed routing, distributed firewall, and bridging capabilities.

You can download the NSX-T Data Center DEBs manually and make them part of the host image. Be aware that download paths can change for each release of NSX-T Data Center. Always check the NSX-T Data Center downloads page to get the appropriate DEBs.

Prerequisites

- Verify that the required third-party packages are installed. See Install Third-Party Packages on a KVM Host.

Procedure

1   Log in to the host as a user with administrative privileges.

2   (Optional) Navigate to the /tmp directory.

```
cd /tmp
```

3   Download and copy the nsx-lcp file into the /tmp directory.

4 Untar the package.

```
tar -xvf nsx-lcp-<release>-ubuntu-trusty_amd64.tar.gz
```

5 Navigate to the package directory.

```
cd nsx-lcp-trusty_amd64/
```

6 Install the packages.

```
sudo dpkg -i *.deb
```

7 Reload the OVS kernel module.

```
/etc/init.d/openvswitch-switch force-reload-kmod
```

If the hypervisor uses DHCP on OVS interfaces, restart the network interface on which DHCP is configured. You can manually stop the old dhclient process on the network interface and restart a new dhclient process on that interface.

8 To verify, you can run the `dpkg -l | grep nsx` command.

```
user@host:~$ dpkg -l | grep nsx

ii  nsx-agent                          <release>    amd64       NSX Agent
ii  nsx-aggservice                     <release>    all         NSX Aggregation Service Lib
ii  nsx-cli                            <release>    all         NSX CLI
ii  nsx-da                             <release>    amd64       NSX Inventory Discovery Agent
ii  nsx-host                           <release>    all         NSX host meta package
ii  nsx-host-node-status-reporter      <release>    amd64       NSX Host Status Reporter for
Aggregation Service
ii  nsx-lldp                           <release>    amd64       NSX LLDP Daemon
ii  nsx-logical-exporter               <release>    amd64       NSX Logical Exporter
ii  nsx-mpa                            <release>    amd64       NSX Management Plane Agent Core
ii  nsx-netcpa                         <release>    amd64       NSX Netcpa
ii  nsx-sfhc                           <release>    amd64       NSX Service Fabric Host
Component
ii  nsx-transport-node-status-reporter <release>    amd64       NSX Transport Node Status
Reporter
ii  nsxa                               <release>    amd64       NSX L2 Agent
```

Any errors are most likely caused by incomplete dependencies. The `apt-get install -f` command can attempt to resolve dependencies and re-run the NSX-T Data Center installation.

What to do next

Add the host to the NSX-T Data Center management plane. See Deploy NSX Manager Nodes to Form a Cluster Using CLI .

Manually Install NSX-T Data Center Kernel Modules on RHEL and CentOS KVM Hypervisors

To prepare hosts to participate in NSX-T Data Center, you can manually install NSX-T Data Center kernel modules on RHEL or CentOS KVM hosts.

This allows you to build the NSX-T Data Center control-plane and management-plane fabric. NSX-T Data Center kernel modules packaged in RPM files run within the hypervisor kernel and provide services such as distributed routing, distributed firewall, and bridging capabilities.

You can download the NSX-T Data Center RPMs manually and make them part of the host image. Be aware that download paths can change for each release of NSX-T Data Center. Always check the NSX-T Data Center downloads page to get the appropriate RPMs.

Prerequisites

Ability to reach a RHEL or CentOS repository.

Procedure

1   Log in to the host as an administrator.

2   Download and copy the nsx-lcp file into the /tmp directory.

3   Untar the package.

```
tar -zxvf nsx-lcp-<release>-rhel7.4_x86_64.tar.gz
```

4   Navigate to the package directory.

```
cd nsx-lcp-rhel74_x86_64/
```

5   Install the packages.

```
sudo yum install *.rpm
```

When you run the yum install command, any NSX-T Data Center dependencies are resolved, assuming the RHEL or CentOS hosts can reach their respective repositories.

6   Reload the OVS kernel module.

```
/etc/init.d/openvswitch force-reload-kmod
```

If the hypervisor uses DHCP on OVS interfaces, restart the network interface on which DHCP is configured. You can manually stop the old dhclient process on the network interface and restart a new dhclient process on that interface.

7   To verify, you can run the `rpm -qa | egrep 'nsx|openvswitch|nicira'` command.

The installed packages in the output must match the packages in the nsx-rhel74 or nsx-centos74 directory.

What to do next

Add the host to the NSX-T Data Center management plane. See Deploy NSX Manager Nodes to Form a Cluster Using CLI .

NSX Edge Networking Setup

NSX Edge can be installed using ISO, OVA/OVF, or PXE start. Regardless of the installation method, make sure that the host networking is prepared before you install NSX Edge.

High-Level View of NSX Edge Within a Transport Zone

The high-level view of NSX-T Data Center shows two transport nodes in a transport zone. One transport node is a host. The other is an NSX Edge.

Figure 8-5. High-Level Overview of NSX Edge
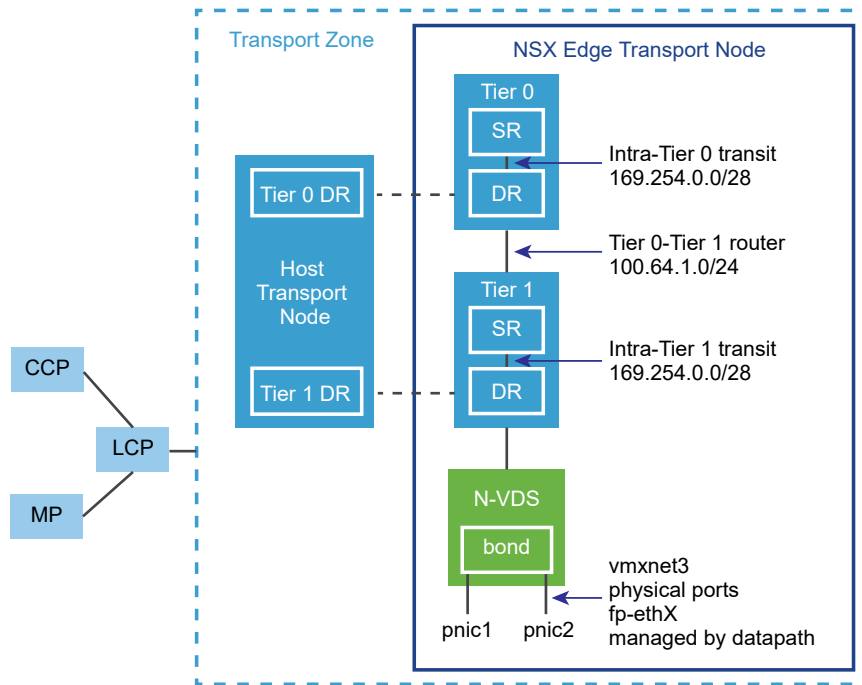
When you first deploy an NSX Edge, you can think of it as an empty container. The NSX Edge does not do anything until you create logical routers. The NSX Edge provides the compute backing for tier-0 and tier-1 logical routers. Each logical router contains a services router (SR) and a distributed router (DR). When we say that a router is distributed, we mean that it is replicated on all transport nodes that belong to the same transport zone. In the figure, the host transport node contains the same DRs contained on the tier-0 and tier-1 routers. A services router is required if the logical router is going to be configured to perform services, such as NAT. All tier-0 logical routers have a services router. A tier-1 router can have a services router if needed based on your design considerations.

By default, the links between the SR and the DR use the 169.254.0.0/28 subnet. These intra-router transit links are created automatically when you deploy a tier-0 or tier-1 logical router. You do not need to configure or modify the link configuration unless the 169.254.0.0/28 subnet is already in use in your deployment. On a tier-1 logical router, the SR is present only if you select an NSX Edge cluster when creating the tier-1 logical router.

The default address space assigned for the tier-0-to-tier-1 connections is 100.64.0.0/10. Each tier-0-to-tier-1 peer connection is provided a /31 subnet within the 100.64.0.0/10 address space. This link is created automatically when you create a tier-1 router and connect it to a tier-0 router. You do not need to configure or modify the interfaces on this link unless the 100.64.0.0/10 subnet is already in use in your deployment.

Each NSX-T Data Center deployment has a management plane cluster (MP) and a control plane cluster (CCP). The MP and the CCP push configurations to each transport zone's local control plane (LCP). When a host or NSX Edge joins the management plane, the management plane agent (MPA) establishes connectivity with the host or NSX Edge, and the host or NSX Edge becomes an NSX-T Data Center fabric node. When the fabric node is then added as a transport node, LCP connectivity is established with the host or NSX Edge.

Lastly, the figure shows an example of two physical NICs (pNIC1 and pNIC2) that are bonded to provide high availability. The datapath manages the physical NICs. They can serve as either VLAN uplinks to an external network or as tunnel endpoint links to internal NSX-T Data Center-managed VM networks.

It is a best practice to allocate at least two physical links to each NSX Edge that is deployed as a VM. Optionally, you can overlap the port groups on the same pNIC using different VLAN IDs. The first network link found is used for management. For example, on an NSX Edge VM, the first link found might be vnic1. On a bare-metal installation, the first link found might be eth0 or em0. The remaining links are used for the uplinks and tunnels. For example, one might be for a tunnel endpoint used by NSX-T Data Center-managed VMs. The other might be used for an NSX Edge-to-external TOR uplink.

You can view the physical link information of the NSX Edge, by logging in to the CLI as an administrator and running the `get interfaces` and `get physical-ports` commands. In the API, you can use the `GET fabric/nodes/<edge-node-id>/network/interfaces` API call. Physical links are discussed in more detail in the next section.

Whether you install NSX Edge as a VM appliance or on bare metal, you have multiple options for the network configuration, depending on your deployment.

Transport Zones and N-VDS

To understand NSX Edge networking, you must know something about transport zones and N-VDS. Transport zones control the reach of Layer 2 networks in NSX-T Data Center. N-VDS is a software switch that gets created on a transport node. The purpose of N-VDS is to bind logical router uplinks and downlinks to physical NICs. For each transport zone that an NSX Edge belongs to, a single N-VDS gets installed on the NSX Edge.

There are two types of transport zones:

- Overlay for internal NSX-T Data Center tunneling between transport nodes.

- VLAN for uplinks external to NSX-T Data Center.

An NSX Edge can belong to zero VLAN transport zones or many. For zero VLAN transport zones, the NSX Edge can still have uplinks because the NSX Edge uplinks can use the same N-VDS installed for the overlay transport zone. You might do this if you want each NSX Edge to have only one N-VDS. Another design option is for the NSX Edge to belong to multiple VLAN transport zones, one for each uplink.

The most common design choice is three transport zones: One overlay and two VLAN transport zones for redundant uplinks.

To use the same VLAN ID for a transport network for overlay traffic and other for VLAN traffic, such as a VLAN uplink, configure the ID on two different N-VDS, one for VLAN and the other for overlay.

Virtual-Appliance/VM NSX Edge Networking

When you install NSX Edge as a virtual appliance or VM, internal interfaces are created, called fp-ethX, where X is 0, 1, 2, and 3. These interfaces are allocated for uplinks to a top-of-rack (ToR) switches and for NSX-T Data Center overlay tunneling.

When you create the NSX Edge transport node, you can select fp-ethX interfaces to associate with the uplinks and the overlay tunnel. You can decide how to use the fp-ethX interfaces.

On the vSphere distributed switch or vSphere Standard switch, you must allocate at least two vmnics to the NSX Edge: One for NSX Edge management and one for uplinks and tunnels.

In the following sample physical topology, fp-eth0 is used for the NSX-T Data Center overlay tunnel. fp-eth1 is used for the VLAN uplink. fp-eth2 and fp-eth3 are not used. vNIC1 is assigned to the management network.

Figure 8-6. One Suggested Link Setup for NSX Edge VM Networking
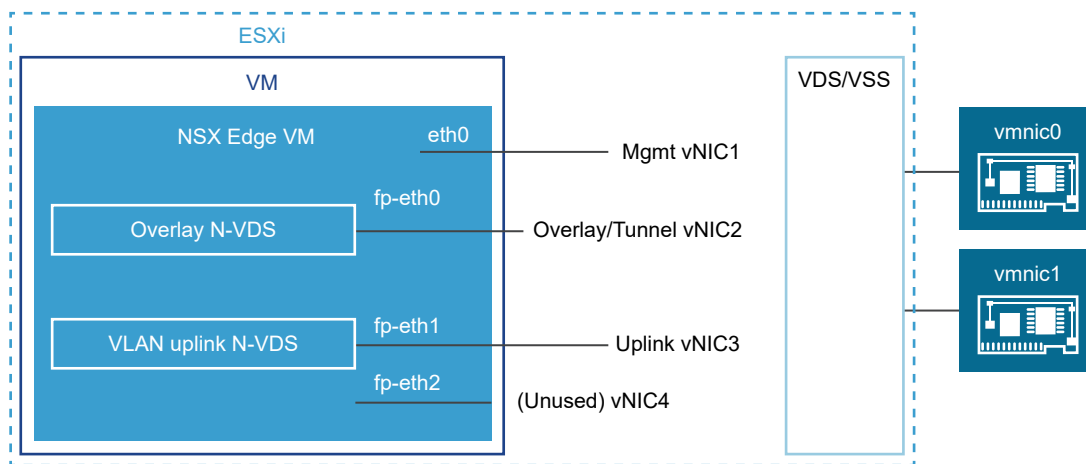
The NSX Edge shown in this example belongs to two transport zones (one overlay and one VLAN) and therefore has two N-VDS, one for tunnel and one for uplink traffic.

This screenshot shows the virtual machine port groups, nsx-tunnel, and vlan-uplink.
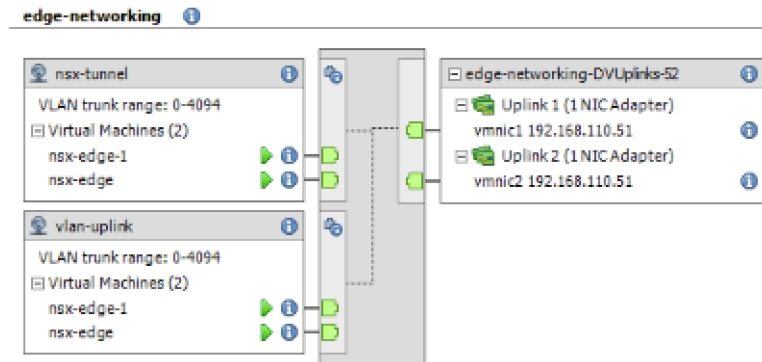
During deployment, you must specify the network names that match the names configured on your VM port groups. For example, to match the VM port groups in the example, your network ovftool settings can be as follows if you were using the ovftool to deploy NSX Edge:

```
--net:"Network 0-Mgmt" --net:"Network 1-nsx-tunnel" --net:"Network 2=vlan-uplink"
```

The example shown here uses the VM port group names Mgmt, nsx-tunnel, and vlan-uplink. You can use any names for your VM port groups.

The tunnel and uplink VM port groups configured for the NSX Edge do not need to be associated with VMkernel ports or given IP addresses. This is because they are used at Layer 2 only. If your deployment uses DHCP to provide an address to the management interface, make sure that only one NIC is assigned to the management network.

Notice that the VLAN and tunnel port groups are configured as trunk ports. This is required. For example, on a standard vSwitch, you configure trunk ports as follows: . **Host > Configuration > Networking > Add Networking > Virtual Machine > VLAN ID All (4095)**.

If you are using an appliance-based or VM NSX Edge, you can use standard vSwitches or vSphere distributed switches.

NSX Edge VM can be installed on an NSX-T Data Center prepared host and configured as a transport node. There are two types of deployment:

- NSX Edge VM can be deployed using VSS/VDS port groups where VSS/VDS consume separate pNIC(s) on the host. Host transport node consumes separate pNIC(s) for N-VDS installed on the host. N-VDS of the host transport node co-exists with a VSS or VDS, both consuming separate pNICs. Host TEP (Tunnel End Point) and NSX Edge TEP can be in the same or different subnets.

- NSX Edge VM can be deployed using VLAN-backed logical switches on the N-VDS of the host transport node. Host TEP and NSX Edge TEP must be in different subnets.

Optionally, you can install multiple NSX Edge appliances/VMs on a single host, and the same management, VLAN, and tunnel endpoint port groups can be used by all installed NSX Edges.

With the underlying physical links up and the VM port groups configured, you can install the NSX Edge.

Bare-Metal NSX Edge Networking

The bare-metal NSX Edge contains internal interfaces called fp-ethX, where X is 0, 1, 2, 3, or 4. The number of fp-ethX interfaces created depends on how many physical NICs your bare-metal NSX Edge has. Up to four of these interfaces can be allocated for uplinks to top-of-rack (ToR) switches and NSX-T Data Center overlay tunneling.

When you create the NSX Edge transport node, you can select fp-ethX interfaces to associate with the uplinks and the overlay tunnel.

You can decide how to use the fp-ethX interfaces. In the following sample physical topology, fp-eth0 and fp-eth1 are bonded and used for the NSX-T Data Center overlay tunnel. fp-eth2 and fp-eth3 are used as redundant VLAN uplinks to TORs.

Figure 8-7. One Suggested Link Setup for Bare-Metal NSX Edge Networking
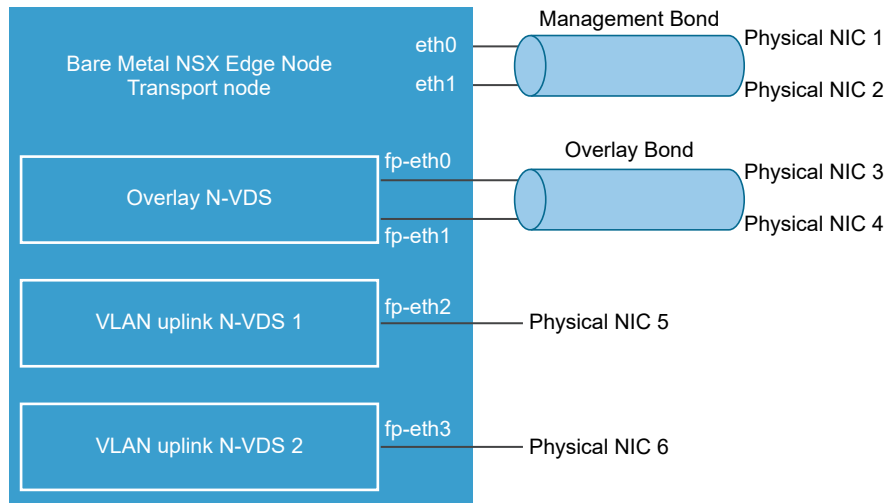
NSX Edge Uplink Redundancy

NSX Edge uplink redundancy allows two VLAN equal-cost multipath (ECMP) uplinks to be used on the NSX Edge-to-external TOR network connection.

When you have two ECMP VLAN uplinks, you must also have two TOR switches for high availability and fully meshed connectivity. Each VLAN logical switch has an associated VLAN ID.

When you add an NSX Edge to a VLAN transport zone, a new N-VDS is installed. For example, if you add an NSX Edge node to four VLAN transport zones, as shown in the figure, four N-VDS get installed on the NSX Edge.

Figure 8-8. One Suggested ECMP VLAN Setup for NSX Edges to TORs
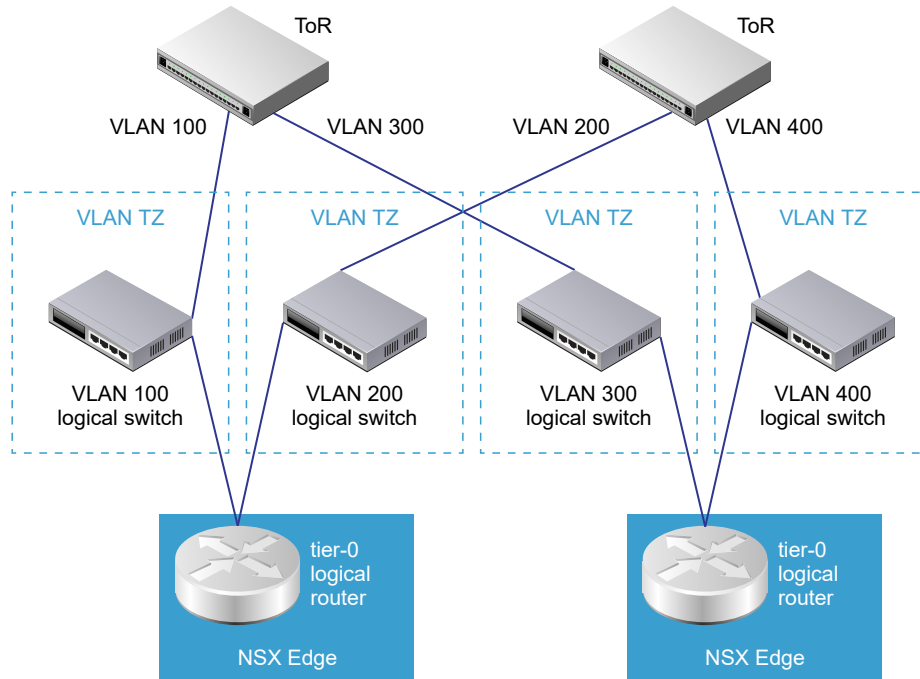
Note For an Edge VM deployed on an ESXi host that has the vSphere Distributed Switch (vDS) and not N-VDS, you must do the following:

- Enable forged transmit for DHCP to work.

- Enable promiscuous mode for the Edge VM to receive unknown unicast packets because MAC learning is disabled by default. This is not necessary for vDS 6.6 or later, which has MAC learning enabled by default.

Create an NSX Edge Transport Node

You can add an NSX Edge to the NSX-T Data Center fabric and proceed to configure the NSX Edge as a transport node.

A transport node is a node that is capable of participating in an NSX-T Data Center overlay or NSX-T Data Center VLAN networking. Any node can serve as a transport node if it contains an N-VDS. Such nodes include but are not limited to NSX Edges.

An NSX Edge can belong to one overlay transport zone and multiple VLAN transport zones. If a VM requires access to the outside world, the NSX Edge must belong to the same transport zone that the VM's logical switch belongs to. Generally, the NSX Edge belongs to at least one VLAN transport zone to provide the uplink access.

Note If you plan to create transport nodes from a template VM, make sure that there are no certificates on the host in `/etc/vmware/nsx/`. The netcpa agent does not create a certificate if a certificate already exists.

Prerequisites

- Transport zones must be configured.

- Verify that compute manager is configured. See Add a Compute Manager.

- An uplink profile must be configured or you can use the default uplink profile for bare-metal NSX Edge nodes.

- An IP pool must be configured or must be available in the network deployment.

- At least one unused physical NIC must be available on the host or NSX Edge node.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Nodes > Edge Transport Nodes > Add Edge VM**.

3 Type a name for the NSX Edge.

4 Type the Host name or FQDN from vCenter Server.

5 For an optimal performance, reserve memory for the NSX Edge appliance.

   Set the reservation to ensure that NSX Edge has sufficient memory to run efficiently. See NSX Edge VM System Requirements.

6 Specify the CLI and the root passwords for the NSX Edge.

   Your passwords must comply with the password strength restrictions.

   - At least 12 characters

   - At least one lower-case letter

   - At least one upper-case letter

   - At least one digit

   - At least one special character

   - At least five different characters

   - No dictionary words

   - No palindromes

   - More than four monotonic character sequence is not allowed

7 Enter the NSX Edge details.

Option	Description
Compute Manager	Select the compute manager from the drop-down menu.
	The compute manager is the vCenter Server registered in the Management Plane.
Cluster	Designate the cluster the NSX Edge is going to join from the drop-down menu.
Resource Pool or Host	Assign either a resource pool or a specific host for the NSX Edge from the drop-down menu.
Datastore	Select a datastore for the NSX Edge files from the drop-down menu.

8 Enter the NSX Edge interface details.

Option	Description
IP Assignment	Select **DHCP** or **Static** IP.
	If you select **Static**, you must specify a list of comma-separated IP addresses, a gateway, and a subnet mask.
Management Interface	Select the VM Network interface from the drop-down menu.

9 Select the transport zones that this transport node belongs to.

An NSX Edge transport node belongs to at least two transport zones, an overlay for NSX-T Data Center connectivity and a VLAN for uplink connectivity.

Note Multiple VTEPs in a transport zone must be configured to the same network segment. If VTEPs in a transport zone are configured to different network segments, BFD sessions cannot be established between the VTEPs.

10 Enter the N-VDS information.

Option	Description
Edge Switch Name	Select the overlay switch from the drop-down menu.
Uplink Profile	Select the uplink profile from the drop-down menu.
	The available uplinks depend on the configuration in the selected uplink profile.
IP Assignment	Select **Use IP Pool** or **Use Static IP List** for the overlay N-VDS.
	If you select **Use Static IP List**, you must specify a list of comma-separated IP addresses, a gateway, and a subnet mask.
IP Pool	If you selected **Use IP Pool** for IP assignment, specify the IP pool name.
Datapath Interfaces	Select the data path interface name for the uplink interface.

Note LLDP profile is not supported on an NSX Edge VM appliance.

11 View the connection status on the **Transport Nodes** page.

After adding the NSX Edge as a transport node, the connection status changes to Up in 10-12 minutes.

12 (Optional) View the transport node with the GET `https://<nsx-manager>/api/v1/transport-nodes/<transport-node-id>` API call.

13 (Optional) For status information, use the GET `https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/status` API call.

What to do next

Add the NSX Edge node to an NSX Edge cluster. See Create an NSX Edge Cluster.

Create an NSX Edge Cluster

Having a multi-node cluster of NSX Edges helps ensure that at least one NSX Edge is always available.

In order to create a tier-0 logical router or a tier-1 router with stateful services such as NAT, load balancer, and so on. You must associate it with an NSX Edge cluster. Therefore, even if you have only one NSX Edge, it must still belong to an NSX Edge cluster to be useful.

An NSX Edge transport node can be added to only one NSX Edge cluster.

An NSX Edge cluster can be used to back multiple logical routers.

After creating the NSX Edge cluster, you can later edit it to add additional NSX Edges.

Prerequisites

- Install at least one NSX Edge node.

- Join the NSX Edges with the management plane.

- Add the NSX Edges as transport nodes.

- Optionally, create an NSX Edge cluster profile for high availability (HA). You can also use the default NSX Edge cluster profile.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Nodes > Edge Clusters > Add**.

3 Enter the NSX Edge cluster a name.

4 Select an NSX Edge cluster profile from the drop-down menu.

5 Select either NSX Edge Node from the Member Type drop-down menu.

If the virtual machine is deployed in a public cloud environment, select Public Cloud Gateway otherwise select NSX Edge Node.

6 From the **Available** column, select NSX Edges and click the right-arrow to move them to the **Selected** column.

What to do next

You can now build logical network topologies and configure services. See the *NSX-T Data Center Administration Guide*.

Auto Deploy Stateless Cluster

9

Stateless hosts do not persist configuration, so they need an auto-deploy server to provide the required start files when hosts power on.

This section helps you to set up a stateless cluster using vSphere Auto Deploy and NSX-T Transport Node Profile to reprovision a host with a new image profile that contains a different version of ESXi and NSX-T. Hosts that are set up for vSphere Auto Deploy use an auto-deploy server and vSphere host profiles to customize hosts. These hosts can also be set up for NSX-T Transport Node Profile to configure NSX-T on the hosts.

So, a stateless host can be set up for vSphere Auto Deploy and NSX-T Transport Node Profile to reprovision a host with a custom ESXi and NSX-T version.

This chapter includes the following topics:

- High-Level Tasks to Auto Deploy Stateless Cluster
- Prerequisites and Supported Versions
- Create a Custom Image Profile for Stateless Hosts
- Associate the Custom Image with the Reference and Target Hosts
- Set Up Network Configuration on the Reference Host
- Configure the Reference Host as a Transport Node in NSX-T
- Extract and Verify the Host Profile
- Verify the Host Profile Association with Stateless Cluster
- Update Host Customization
- Trigger Auto Deployment on Target Hosts
- Troubleshoot Host Profile and Transport Node Profile

High-Level Tasks to Auto Deploy Stateless Cluster

High-level tasks to auto deploy a stateless cluster.

The high-level tasks to set up an auto deploy stateless cluster are:

1 Prerequisites and Supported Versions. See Prerequisites and Supported Versions.

2 (Reference host) Create a Custom Image Profile. See Create a Custom Image Profile for Stateless Hosts.

3 (Reference and Target hosts) Associate the Custom Image Profile. See Associate the Custom Image with the Reference and Target Hosts.

4 (Reference host) Set up Network Configuration in ESXi. See Set Up Network Configuration on the Reference Host.

5 (Reference host) Configure as a Transport Node in NSX. See Configure the Reference Host as a Transport Node in NSX-T.

6 (Reference host) Extract and Verify Host Profile. See Extract and Verify the Host Profile.

7 (Reference and Target hosts) Verify the Host Profile Association with Stateless Cluster. See Verify the Host Profile Association with Stateless Cluster.

8 (Reference host) Update Host Customization. See Update Host Customization.

9 (Target hosts) Trigger Auto Deployment. See Trigger Auto Deployment on Target Hosts.

    a Before applying Transport Node Profile. See Reboot Hosts Before Applying TNP.

    b Apply Transport Node Profile. See Apply TNP on Stateless Cluster.

    c After applying Transport Node Profile. See Reboot Hosts After Applying TNP.

10 Troubleshoot Host Profile and Transport Node Profile. See Troubleshoot Host Profile and Transport Node Profile.

Prerequisites and Supported Versions

Prerequisites and supported ESXi and NSX-T versions.

Supported Workflows

- With Image Profile and HostProfile

Prerequisites

- Only homogeneous clusters (all hosts within a cluster must be either stateless or stateful) are supported.

- Image builder service must be enabled.

- Auto deploy service must be enabled.

Supported NSX and ESXi Versions

Supported EXSi Version	ESXi 67ep6	ESXi 67u2	ESXi 67u3	ESXi 67ep7
NSX-T Data Center 2.4	Yes	Yes	No	No
NSX-T Data Center 2.4.1	Yes	Yes	No	No
NSX-T Data Center 2.4.2	Yes	Yes	No	No
NSX-T Data Center 2.4.3	Yes	Yes	No	No
NSX-T Data Center 2.5	Yes	Yes	Yes	Yes

Create a Custom Image Profile for Stateless Hosts

In your data center, identify a host to be prepared as the reference host.

The first time the reference host starts up, ESXi associates the default rule with the reference host. In this procedure, we are adding a custom image profile (ESXi and NSX VIBs) and associate the reference host with the new custom image. An image profile with the NSX-T image significantly reduces the installation time. The same custom image is associated with the target hosts in the stateless cluster.

Note Alternatively, you can add only an ESXi image profile to the reference and target stateless cluster. The NSX-T VIBs are downloaded when you apply the transport node profile on the stateless cluster. See Add a Software Depot.

Prerequisites

Ensure that the auto-deploy service and image builder service are enabled. See Using vSphere Auto Deploy to Reprovision Hosts.

Procedure

1   To import NSX-T packages, create a software depot.

2   Download the `nsx-lcp` packages.

   a   Log in to `https://my.vmware.com`.

   b   On the Download VMware NSX-T Data Center page, select the NSX-T version.

   c   In the Product Downloads page, search NSX-T Kernel Modules for a specific VMware ESXi version.

   d   Click **Download Now** to begin downloading the `nsx-lcp` package.

   e   Import `nsx-lcp` packages into the software depot.

3 Create another software depot to import ESXi packages.

The vSphere Web Client displays two depots created on the reference host.

4 Create a custom software depot to clone previously imported ESXi image and `nsx-lcp` packages.

    a   Select the ESXi Image profile from the ESXi software depot created in the earlier step.

    b   Click **Clone**.

    c   In the Clone Image Profile wizard, enter a name for the custom image to be created.

    d   Select the custom software depot where the cloned image (ESXi) must be available.

    e   In the Select software packages window, select the Acceptance level to **VMware Certified**. The ESXi VIBs are preselected.

    f   Identify and select the NSX-T packages manually from the list of packages and click **Next**.

    g   In the Ready to complete screen, verify the details and click **Finish** to create the cloned image containing ESXi and NSX-T packages into the custom software depot.

What to do next

Associate the custom image with the reference and target hosts. See Associate the Custom Image with the Reference and Target Hosts.

Associate the Custom Image with the Reference and Target Hosts

To start the reference host and target hosts with the new custom image containing ESXi and NSX packages, associate the custom image profile.

At this point in the procedure, the custom image is only being associated to the reference and target hosts but NSX installation does not happen.

Important   Perform this custom image association procedure on both reference and target hosts.

Prerequisites

Procedure

1   On the ESXi host, navigate to **Menu** > **Auto Deploy** > **Deployed Hosts**.

2   To associate the custom image profile with a host, select the custom image.

3   Click **Edit Image Profile Association**.

4   In the Edit Image Profile Association wizard, click **Browse** and select the custom depot and select the custom image profile.

5   Enable **Skip image profile signature check**.

6   Click **Ok**.

Software Depots	Deploy Rules	Deployed Hosts	Discovered Hosts	Script Bundles	Configure

ⓘ The image profile, host profile and location that Auto Deploy has associated with the hosts are listed below. The associations might differ from the actual state of the host.

CHECK HOST ASSOCIATIONS COMPLIANCE   REMEDIATE HOST ASSOCIATIONS   EDIT IMAGE PROFILE ASSOCIATION

☐	Host	▼	Associated Image Profile	▼	Associated Host Profile	▼	Associated Location	▼	Associated Script Bundle	▼
☐	🖥 10.144.139.147		⊙ CustomDepot(ESXi and NSX)				🗎 1-datacenter-1964			
☐	🖥 10.144.137.225		⊙ CustomDepot(ESXi and NSX)				🗐 Statless-Cluster			

Results

What to do next

Set up Network Configuration on the Reference Host. See Set Up Network Configuration on the Reference Host.

Set Up Network Configuration on the Reference Host

On the reference host, a standard switch with a VMkernel adapter is created to set up the network configuration on ESXi.

This network configuration is captured in the host profile which is extracted from the reference host. During a stateless deployment, the host profile replicates this network configuration setting on each target host.

Procedure

1   On the ESXi host, configure a vSphere Standard Switch (VSS) or Distributed Virtual switch (DVS) by adding a VMkernel adapter.

2   Verify that the newly added VSS/DVS switch is displayed in the VMkernel adapters page.
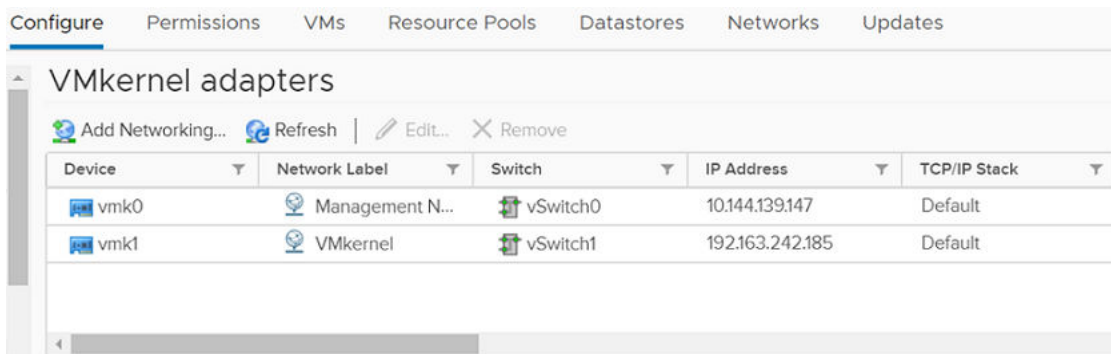
What to do next

Configure the Reference Host as a Transport Node in NSX-T. See Configure the Reference Host as a Transport Node in NSX-T.

Configure the Reference Host as a Transport Node in NSX-T

After the reference host is associated with the custom image profile and configured with a VSS switch, set up the reference host as a transport node in NSX-T.

Procedure

1   From a browser, log in to NSX-T at https://<NSXManager_IPaddress>.

2   To locate the reference host, navigate to **System** -> **Nodes** -> **Host Transport Node**.

3   Create a VLAN transport zone to define the span of the virtual network. The span is defined by attaching N-VDS switches to the transport zone. Based on this attachment, N-VDS can access segments defined within the transport zone. See Create a Transport Zone.

4   Create a VLAN segment on the transport zone. The created segment is displayed as a logical switch.

    a   Navigate to **Networking** -> **Segments**.

b   Select the transport zone to attach the segment.

c   Enter VLAN ID.

d   Click **Save**.

5   Create an uplink profile for the reference host that defines how an N-VDS connects to the physical network. See, Create an Uplink Profile.

6   Configure the reference host as a transport node. See Configure a Managed Host Transport Node.

a   In the Host Transport Node page, select the reference host.

b   Click Configure NSX and select the previously created transport zone, N-VDS, uplink profile.

7 In the Network Mappings to Install section, click **Add Mapping** to add the VMkernel to Segment/Logical switch mapping.

8 Click **Finish** to begin installation of NSX-T on the reference host.

During installation, VMkernel adapters and physical NICs are migrated from a VSS or DVS switch to an N-VDS switch. After installation, configuration status of the reference host is displayed as Success.

Note The reference host is listed under Other Hosts.

9 In vCenter Server, verify that the PNICs and VMkernels adapters on the VSS switch are migrated and connected to the N-VDS switch.

What to do next

Extract and Verify the Host Profile. See Extract and Verify the Host Profile.

Extract and Verify the Host Profile

After you extract the host profile from the reference host, verify the NSX-T configuration extracted in the host profile. It consists of ESXi and NSX-T configuration that is applied to target hosts.

Procedure

1 To extract the host profile, Extract and Configure Host Profile from the Reference Host.

2 Verify the NSX configuration in the extracted host profile.

Results

The host profile contains configuration related to ESXi and NSX as the host was prepared for both environments.

What to do next

Verify the Host Profile Association with Stateless Cluster. See Verify the Host Profile Association with Stateless Cluster.

Verify the Host Profile Association with Stateless Cluster

To prepare the target stateless cluster with ESXi and NSX configuration, associate the host profile extracted from the reference host to the target stateless cluster.

Without the host profile associated to the stateless cluster, new nodes joining the cluster cannot be auto deployed with ESXi and NSX VIBs.

Procedure

1   Attach or Detach Host Profile to Stateless Cluster. See Attach or Detach Entities from a Host Profile.

2   In the Deployed Hosts tab, verify that the existing stateless host is associated with the correct image and associated with the host profile.

3   If the host profile association is missing, select the target host and click Remediate Host Associations to force update the image and host profile to the target host.

Software Depots	Deploy Rules	Deployed Hosts	Discovered Hosts	Script Bundles	Configure

ⓘ The image profile, host profile and location that Auto Deploy has associated with the hosts are listed below. The associations might differ from the actual state of the host.

CHECK HOST ASSOCIATIONS COMPLIANCE   REMEDIATE HOST ASSOCIATIONS   EDIT IMAGE PROFILE ASSOCIATION

|	Host ▽	Associated Image Profile ▽	Associated Host Profile ▽	Associated Location ▽	Associated Script Bundle ▽
☐	🖥 10.144.139.147	◉ CustomDepot(ESXi and NSX)		📑 1-datacenter-1964	
☐	🖥 10.144.137.225	◉ CustomDepot(ESXi and NSX)	📇 Host Profile_ReferenceHost	🗔 Stateless-Cluster	

What to do next

Update Host Customization. See Update Host Customization.

Update Host Customization

After the attaching the host profile to the target cluster, additional custom entries might be required on the host to successfully auto deploy the ESXi and NSX-T packages on it.

Procedure

1 After attaching the host profile to the target cluster, if the hosts are not updated with custom values, the system displays the following message.

2 To update host customizations, navigate to the host profile, click **Actions** -> **Edit Host Customizations**.

3 For ESXi versions 67ep6, 67ep7, 67u2, enter the MUX user password.

4 Verify that all the required fields are updated with appropriate values.

What to do next

Trigger Auto Deployment on Target Hosts. See Trigger Auto Deployment on Target Hosts.

Trigger Auto Deployment on Target Hosts

When a new node is added to the cluster, it needs to be manually rebooted for the ESXi and NSX-T VIBs to be configured.

Note Only applies to stateless hosts.

There are two ways to prepare hosts to trigger auto-deployment of ESXi and NSX-T VIBs to be configured.

- Reboot hosts before applying TNP to the stateless cluster.

- Reboot hosts after applying TNP to the stateless cluster.

If you want to migrate VMkernel adapters when installing NSX-T on the hosts, see:

- Scenarios When the Stateless Host Is in the Target Cluster

- Scenarios When the Stateless Host Is Outside of the Target Cluster

What to do next

Reboot hosts before applying TNP to the stateless cluster. See Reboot Hosts Before Applying TNP.

Reboot Hosts Before Applying TNP

Only applies to stateless hosts. In this scenario, the transport node profile is not applied to the stateless cluster, which means that NSX-T is not installed and configured on the target host.

Procedure

1 Reboot hosts.

The target host starts with the ESXi image. After starting, the target host remains in maintenance mode until the TNP profile is applied to the target host and NSX-T installation is complete. Profiles are applied on hosts in the following order:

Profiles are applied on hosts in the following order.

- Image profile is applied to the host.

- Host profile configuration is applied to the host.

- NSX-T configuration is applied to the host.

2 On the ESXi host, the VMkernel adapter is attached to a temporary segment named <N-LogicalSegment> because the host is not yet a transport node. After NSX-T is installed the temporary switch is replaced with the actual N-VDS switch and logical segment.

Configure	Permissions	VMs	Datastores	Networks	Updates		

VMkernel adapters

🖧 Add Networking... 🔁 Refresh | ✏ Edit... ✕ Remove

Device	▼	Network Label	▼	Switch	▼	IP Address	▼	TCP/IP Stack	▼	\
vmk0		Management N...		vSwitch0		10.144.137.225		Default		
vmk1		N-Segement_a...		N-nsxvswitch		192.163.242.186		Default		

ESXi VIBs are applied to all the rebooted hosts. A temporary NSX switch in an ESXi host. When TNP is applied to the hosts, the temporary switch is replaced by the actual NSX-T switch.

What to do next

Apply TNP to the stateless cluster. See Apply TNP on Stateless Cluster.

Apply TNP on Stateless Cluster

NSX-T configuration and installation only happens on the target hosts when TNP is applied to the cluster.

Procedure

1 Note down the settings extracted in the Host Profile from the reference host. The corresponding entities in the TNP profile must have the same value. For example, the N-VDS name used in the Host Profile and TNP must be the same.

For more information on extracted host profile settings, see Extract and Verify the Host Profile.

2 Add a TNP. See Add a Transport Node Profile.

3 Ensure that values of the following parameters are the same on both the new TNP profile and the existing Host Profile.

- N-VDS Name: Ensure N-VDS name referenced in Host Profile and TNP is the same.

- Uplink Profile: Ensure uplink profile referenced in Host Profile and TNP is the same.

- PNIC: When mapping a physical NIC to an uplink profile, first verify the NIC used in the Host Profile and map that physical NIC to the uplink profile.

- Network mapping for install: When mapping network during installation, first verify the VMkernel to segment mapping on the Host Profile and add the same mapping in TNP.

- Network mapping for uninstall: When mapping network during uninstallation, first verify the VMkernel to VSS/DVS switch mapping on the Host Profile and add the same mapping in TNP.

4 Add a TNP by entering all required field. See Add a Transport Node Profile.

Ensure that values of the following parameters are the same on both the new TNP profile and the existing Host Profile.

- Transport Zone: Ensure transport zone referenced in Host Profile and TNP is the same.

- N-VDS Name: Ensure N-VDS name referenced in Host Profile and TNP is the same.

- Uplink Profile: Ensure uplink profile referenced in Host Profile and TNP is the same.

- PNIC: When mapping a physical NIC to an uplink profile, first verify the NIC used in the Host Profile and map that physical NIC to the uplink profile.

- Network mapping for install: When mapping network during installation, first verify the VMkernel to logical switch mapping on the Host Profile and add the same mapping in TNP.

- Network mapping for uninstall: When mapping network during uninstallation, first verify the VMkernel to VSS/DVS switch mapping on the Host Profile and add the same mapping in TNP.

After applying TNP on target nodes, if the TNP configuration does not match Host Profile configuration, the node might not come up because of compliance errors.

5 Verify that the TNP profile is successfully created.

6 Apply TNP profile to the target cluster and click **Save**.

7 Verify that the TNP profile is successfully applied to the target cluster. It means that NSX is successfully configured on all nodes of the cluster.

8 In vSphere, verify that the physical NICs or VMkernel adapters are attached to the N-VDS switch.

9 In NSX, verify that the ESXi host is configured successfully as a transport node.

What to do next

Alternatively, you can reboot a target host after applying TNP to the cluster. See Reboot Hosts After Applying TNP.

Reboot Hosts After Applying TNP

Only applies to stateless hosts. When a new node is added to the cluster, manually reboot the node for the ESXi and NSX-T packages to be configured on it.

Procedure

1 Apply TNP to the stateless cluster that is already prepared with host profile. See Create and Apply TNP on Stateless Cluster.

2 Reboot hosts.

After applying TNP profile to the stateless cluster, when you reboot any new node joining the cluster that node is automatically configured with NSX-T on the host.

What to do next

Ensure that you reboot any new node joining the cluster to automatically deploy and configure ESXi and NSX-T on the rebooted node.

To troubleshoot issues related to host profile and transport node profile when configuring auto-deployment, see Troubleshoot Host Profile and Transport Node Profile.

Scenarios When the Stateless Host Is in the Target Cluster

This section discusses use cases when a stateless host exists in the target cluster.

Important   On a stateless target host:

- Migration of vmk0 adapter from VSS/DVS to N-VDS is not supported on NSX-T 2.4 and NSX-T 2.4.1.

- Migration of vmk0 adapter from VSS/DVS to N-VDS is supported on NSX-T 2.5.

Target Host	Reference Host Configuration	Steps To Auto Deploy Target Hosts
Target host has vmk0 adapter configured.	The host profile extracted from the reference host has vmk0 configured on an N-VDS switch. In NSX-T, TNP has only vmk0 migration mapping configured.	1 Attach the host profile to the target host. The vmk0 adapter is attached to a vSwitch. 2 Update host customizations, if required. 3 Reboot the host. The host profile is applied to the host. vmk0 is attached to a temporary switch. 4 Apply TNP. The vmk0 adapter migrates to N-VDS. The target host is successfully deployed with ESXi and NSX-T VIBs.
Target host has vmk0 adapter configured.	The host profile extracted from the reference host has vmk0 on vSwitch and vmk1 is on an N-VDS switch. In NSX-T, TNP has only vmk1 migration mapping configured.	1 Attach the host profile to the target host. The vmk0 adapter is attached to a vSwitch, but vmk1 is not realized on any switch. 2 Update host customizations, if required. 3 Reboot the host. vmk0 is attached to a vSwitch and vmk1 is attached to a temporary NSX switch. 4 Apply TNP. The vmk1 adapter migrates to N-VDS. 5 (optional) If the host remains non-compliant with the host profile, reboot the host to make the host compliant. The target host is successfully deployed with ESXi and NSX-T VIBs.

Target Host	Reference Host Configuration	Steps To Auto Deploy Target Hosts
Target host has vmk0 adapter configured.	The host profile extracted from the reference host has vmk0 is configured on a vSwitch and vmk1 is configured on an N-VDS switch. In NSX-T, TNP has vmk0 and vmk1 migration mappings configured.	1 Attach the host profile to the target host. The vmk0 adapter is attached to a vSwitch, but vmk1 is not realized on any switch. 2 Update host customizations, if required. 3 Reboot the host. The vmk0 adapter is attached to a vSwitch and vmk1 is attached to a temporary NSX switch. 4 Apply TNP. 5 (optional) If the host remains non-compliant with the host profile, reboot the host to make the host compliant. The target host is successfully deployed with ESXi and NSX-T VIBs.
Target host has vmk0 and vmk1 adapters configured.	The host profile extracted from the reference host has vmk0 on vSwitch and vmk1 configured on an N-VDS switch. In NSX-T, TNP has a vmk1 migration mapping configured.	1 Attach the host profile to the target host. The vmk0 and vmk1 adapters are attached to a vSwitch. 2 Update host customizations, if required. 3 Reboot the host. 4 Apply TNP. The vmk0 adapter is attached to a vSwitch and vmk1 is attached to an N-VDS switch. 5 (optional) If the host remains non-compliant with the host profile, reboot the host to make the host compliant. The target host is successfully deployed with ESXi and NSX-T VIBs.
Target host has vmk0 and vmk1 adapters configured.	The host profile extracted from the reference host has vmk0 and vmk1 configured on an N-VDS switch. In NSX-T, TNP has vmk0 and vmk1 migration mappings configured.	1 Attach the host profile to the target host. The vmk0 and vmk1 adapters are attached to a vSwitch. 2 Update host customizations, if required. 3 Reboot the host. 4 Apply TNP. The vmk0 and vmk1 are migrated to an N-VDS switch. The target host is successfully deployed with ESXi and NSX-T VIBs.

Scenarios When the Stateless Host Is Outside of the Target Cluster

This section discusses use cases when a stateless host exists outside of the target cluster.

Important On stateless hosts:

- Migration of vmk0 adapter from VSS/DVS to N-VDS is not supported on NSX-T 2.4 and NSX-T 2.4.1.

- Migration of vmk0 adapter from VSS/DVS to N-VDS is supported on NSX-T 2.5.

.

Target Host State	Reference Host Configuration	Steps To Auto Deploy Target Hosts
Host is in powered-off state (first-time start). It is later added to the cluster. The default auto-deploy rule is configured for the target cluster and associated with the host profile. The TNP is applied on the cluster.	The host profile extracted from the reference host has VMkernel adapter 0 (vmk0) on vSwitch and VMkernel adapter 1 (vmk1) configured on an N-VDS switch. In NSX-T, TNP has only vmk1 migration mapping configured.	1 Power on the host. After the host is powered on. ■ The host gets added to cluster. ■ The host profile is applied on the target host. ■ The vmk0 adapter is on vSwitch and vmk1 adapter is on a temporary switch. ■ TNP is triggered. ■ After TNP is applied to cluster, the vmk0 adapter is on vSwitch and vmk1 is migrated to the N-VDS switch. 2 (Optional) If the host remains non-compliant with the host profile, reboot the host to make the host compliant. The host is successfully deployed with ESXi and NSX-T VIBs.
Host is in powered-off state (first-time start). It is later added to the cluster. The default auto-deploy rule is configured for the target cluster and associated to the host profile. The TNP is applied on the cluster.	The host profile extracted from the reference host has VMkernel adapter 0 (vmk0) and VMkernel adapter 1 (vmk1) configured on an N-VDS switch. In NSX-T, TNP has vmk0 and vmk1 migration configured.	1 Power on the host. After the host is powered on. ■ The host gets added to cluster. ■ The host profile is applied on the target host. ■ The vmk0 and vmk1 adapters are on a temporary switch. ■ TNP is triggered. ■ After TNP is applied to cluster, the vmk0 and vmk1 are migrated to the N-VDS switch. The host is successfully deployed with ESXi and NSX-T VIBs.

Target Host State	Reference Host Configuration	Steps To Auto Deploy Target Hosts
Host is in powered-on state. It is later added to the cluster. The default auto-deploy rule is configured for the target cluster and associated to the host profile. Target host only has a vmk0 adapter configured on it.	The host profile extracted from the reference host has VMkernel adapter 0 (vmk0) on vSwitch and VMkernel adapter 1 (vmk1) configured on an N-VDS switch. In NSX-T, TNP has a vmk1 migration mapping configured.	1 Move the host to be part of the cluster. 2 Reboot the host. After the host is rebooted the host profile is applied on the target host. - The vmk0 adapter is attached to a vSwitch, whereas the vmk1 adapter is attached to a temporary NSX switch. - TNP is triggered. - vmk1 is migrated to the N-VDS switch. 3 (Optional) If the host remains non-compliant with the host profile, reboot the host to make the host compliant. The host is successfully deployed with ESXi and NSX-T VIBs.
Host is in powered-on state. It is later added to the cluster. The default auto-deploy rule is configured for the target cluster and associated to the host profile. Target host only has a vmk0 adapter configured on it.	The host profile extracted from the reference host has VMkernel adapter 0 (vmk0) and VMkernel adapter 1 (vmk1) configured on N-VDS. In NSX-T, TNP has vmk0 and vmk1 migration configured.	1 Move the host to be part of the cluster. 2 Reboot the host. After the rebooting the host, the host profile is applied to the target host. - The vmk0 and vmk1 adapters are attached to a temporary NSX switch. - TNP is triggered. - vmk0 and vmk1 are attached to an N-VDS switch. The host is successfully deployed with ESXi and NSX-T VIBs.

Target Host State	Reference Host Configuration	Steps To Auto Deploy Target Hosts
Host is in powered-on state. It is later added to the cluster. The default auto-deploy rule is configured for the target cluster and associated to the host profile. Target host has vmk0 and vmk1 network mapping configured.	The host profile extracted from the reference host has VMkernel adapter 0 (vmk0) on vSwitch and VMkernel adapter 1 (vmk1) configured on an N-VDS switch. In NSX-T, TNP has a vmk1 migration configured.	1 Move the host to be part of the cluster. 2 Reboot the host. After the host is rebooted the host profile is applied on the target host. ■ The vmk0 adapter is attached to a vSwitch, whereas the vmk1 adapter is attached to a temporary NSX switch. ■ TNP is triggered. ■ vmk1 is migrated to the N-VDS switch. 3 (Optional) If the host remains non-compliant with the host profile, reboot the host to make the host compliant. The host is successfully deployed with ESXi and NSX-T VIBs.
Host is in powered-on state. It is later added to the cluster. The default auto-deploy rule is configured for the target cluster and associated to the host profile. Host has vmk0 and vmk1 network mapping configured.	In the reference host, the host profile has VMkernel adapter 0 (vmk0) and VMkernel adapter 1 (vmk1) configured on an N-VDS switch. In NSX-T, TNP has vmk0 and vmk1 migration configured.	1 Move the host to be part of the cluster. 2 Reboot the host. After the host is rebooted the host profile is applied on the target host. ■ The vmk0 and vmk1 adapters are attached to a temporary NSX switch. ■ TNP is triggered. ■ The vmk0 and vmk1 adapters are migrated to the N-VDS switch. The host is successfully deployed with ESXi and NSX-T VIBs.

Troubleshoot Host Profile and Transport Node Profile

Troubleshoot issues with host profiles and TNPs when they are used to auto deploy stateless clusters.

Scenario	Description
Host Profile is not portable.	Issue: None of the vCenter servers can use the host profile containing NSX-T configuration. Workaround: None.
Auto Deploy Rule Engine	Issue: Host profile cannot be used in auto deploy rules to deploy new clusters. If new clusters are deployed, the hosts get deployed with basic networking and remain in maintenance mode. Workaround: Prepare each cluster from NSX-T GUI. See Apply TNP on Stateless Cluster.

Scenario	Description
Check compliance errors.	Issue: Host profile remediation cannot fix the compliance errors related to the NSX-T configuration. ■ Physical NICs configured on Host Profile and TNP are different. ■ Mapping between vNIC to LS mapping. Host Profile finds a mismatch in the logical switch to vNIC mapping with the TNP profile. ■ VMkernel connected to N-VDS mismatch on Host Profile and TNP. ■ Opaque switch mismatch on Host Profile and TNP. Workaround: Ensure the NSX-T configuration matches on Host Profile and TNP. Reboot the host to realize the configuration changes. The host comes up.
Remediation	Issue: If there are any NSX-T specific compliance errors, host profile remediation on that cluster is blocked. Incorrect configuration: ■ Mapping between vNIC to LS mapping ■ Mapping of physical NICs Workaround: Ensure that the NSX-T configuration matches on Host Profile and TNP. Reboot the host to realize the configuration changes. The host comes up.
Attach	Issue: In a cluster configured with NSX-T, host profile cannot be attached at the host-level. Workaround: None.
Detach	Issue: Detaching and attaching a new host profile in a cluster configured with NSX-T does not remove the NSX-T configuration. Even though the cluster is compliant with newly attach the host profile, it still has the NSX-T configuration from a previous profile. Workaround: None.
Update	Issue: If the user has changed NSX-T configuration in the cluster, then extract a new host profile. Update the host profile manually for all the settings that were lost. Workaround: None.
Host-level transport node configuration	Issue: After anportsport node was auto-deployed, it acts as individual entity. Any update to that transport node might not match with the TNP. Workaround: Update the cluster. Any update in a standalone transport node cannot persist its migration specification. The migration might fail to post the reboot.
Cannot apply the host profile because mux_user password policy and password were not reset.	Issue: Only on hosts running versions earlier than vSphere 6.7 U3. Host remediation and host profile application on hosts might fail unless the `mux_user` password is reset. Workaround: Under Policies & Profiles, edit the host profile to modify the mux_user password policy and reset the `mux_user` password.

Scenario	Description
PeerDNS configuration is not supported on the VMkernel adapter selected for migration to the NVDS switch.	Issue: If a VMkernel adapter selected for migration to NVDS is peer-DNS enabled, then host profile application fails. Workaround: Edit the extracted host profile by disabling peer-DNS setting on the VMkernel adapter that must be migrated to an NVDS switch. Alternatively, ensure that you do not migrate peer-DNS enabled VMkernel adapters to an NVDS switch.
DHCP address of the VMkernel NIC address not retained	Issue: If the reference host is stateful, then any stateless hosts using profile extracted from the stateful reference host cannot retain their VMkernel management MAC address derived from PXE started MAC. It results in DHCP addressing issues. Workaround: Edit extracted host profile of stateful host and modify the '**Determine how MAC address for vmknic should be decided**' to '**Use the MAC address from which the system was PXE started**'.
Host Profile application failure in vCenter can lead to NSX configuration errors on the host.	Issue: If host profile application fails in vCenter, NSX configuration might also fail. Workaround: In vCenter, verify that host profile was successfully applied. Fix the errors and try again.
LAGS are not supported on stateless ESXi hosts.	Issue: The uplink profile configured as LAGs in NSX is not supported in a stateless ESXi host managed by a vCenter Server or in NSX. Workaround: None.

Uninstalling NSX-T Data Center from a Host Transport Node

10

The steps to uninstall NSX-T Data Center from a host transport node vary depending on the host type and how it is configured.

- **Verify Host Network Mappings for Uninstall**

  Before you uninstall NSX-T Data Center from an ESXi host, verify that you have appropriate network mappings for uninstall configured. The mappings are required if the ESXi host has VMkernel interfaces connected to N-VDS.

- **Uninstall NSX-T Data Center from a vSphere Cluster**

  If you have installed NSX-T Data Center on a vSphere Cluster using transport node profiles, you can follow these instructions to uninstall NSX-T Data Center from all hosts in the cluster.

- **Uninstall NSX-T Data Center from a Host in a vSphere Cluster**

  You can uninstall NSX-T Data Center from a single host that is managed by vCenter Server. The other hosts in the cluster are not affected.

- **Uninstall NSX-T Data Center from a Standalone Host**

  You can uninstall NSX-T Data Center from a standalone host. Standalone hosts can be ESXi or KVM.

Verify Host Network Mappings for Uninstall

Before you uninstall NSX-T Data Center from an ESXi host, verify that you have appropriate network mappings for uninstall configured. The mappings are required if the ESXi host has VMkernel interfaces connected to N-VDS.

The uninstall mapping determines where the interfaces are connected after the uninstall. There are uninstall mappings for physical interfaces (vmnicX) and VMkernel interfaces (vmkX). When you uninstall, VMkernel interfaces are moved from their current connections to the port groups specified in the uninstall mapping. If a physical interface is included in the uninstall mapping, the physical interface is connected to the appropriate vSphere Distributed Switch or vSphere Standard Switch based on the destination port group of the VMkernel interfaces.

Caution   Uninstalling NSX-T Data Center from an ESXi host is disruptive if the physical interfaces or VMkernel interfaces are connected to N-VDS. If the host or cluster is participating in other applications such as vSAN, those applications might be affected by the uninstall.

There are two places that you can configure network mappings for uninstall.

- In the transport node configuration, which applies to that host.

- In a transport node profile configuration, which can then be applied to a cluster.

  Note You must have a compute manager configured to apply a transport node profile to a cluster.

If a compute manager is configured, a host can have both a transport node configuration and a transport node profile configuration. If both exist, the transport node configuration is active. Verify that the network mappings for uninstall are correctly configured on the active configuration.

In this example, the cluster cluster-1 has transport node profile TNP-1 applied to it. The host tn-1 is displaying "Configuration Mismatch". This mismatch message indicates that a different configuration has been applied to tn-1. The mismatch persists until the transport node configuration matches the transport node profile configuration. Transport node tn-2 uses the network mappings from the transport node profile, and transport node tn-1 uses its own configuration.

Prerequisites

- Verify that you have appropriate port groups configured to use in the uninstall mapping. You must use vSphere Distributed Switch ephemeral port groups or vSphere Standard Switch port groups.

- Configure a compute manager if you want to use a vSphere Distributed Switch port group in the uninstall mappings for a standalone ESXi host. See Add a Compute Manager. If there is no compute manager configured, you must use a vSphere Standard Switch port group.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Select **System > Fabric > Nodes > Host Transport Nodes**.

3 For each host you want to uninstall, verify that the network mapping for uninstall includes a port group for each VMkernel interface that is on N-VDS. Add any missing mappings.

Important The port group in the network mapping for uninstall must be a vSphere Distributed Switch ephemeral port group or a vSphere Standard Switch port group.

a To view VMkernel interfaces, log in vCenter Server, select the host, and click **Configure > VMkernel Adapters**.

b If the transport node configuration is the active configuration, select the host and click **Edit** (for standalone hosts) or **Configure NSX** (for managed hosts). Click **Next**, then click **Network Mappings for Uninstall**. View the mappings in the **VMKNic Mappings** and **Physical NIC Mappings** tabs.

c If the transport node profile is the active configuration, click the name of the transport node profile for the cluster in the **NSX Configuration** column and click **Edit**. On the **N-VDS** tab, click **Network Mappings for Uninstall**. View the mappings in the **VMKNic Mappings** and **Physical NIC Mappings** tabs.

Uninstall NSX-T Data Center from a vSphere Cluster

If you have installed NSX-T Data Center on a vSphere Cluster using transport node profiles, you can follow these instructions to uninstall NSX-T Data Center from all hosts in the cluster.

For more information on transport node profiles, see Add a Transport Node Profile.

Caution Uninstalling NSX-T Data Center from an ESXi host is disruptive if the physical interfaces or VMkernel interfaces are connected to N-VDS. If the host or cluster is participating in other applications such as vSAN, those applications might be affected by the uninstall.

If you have not used a transport node profile to install NSX-T Data Center, or if you want to remove NSX-T Data Center from a subset of the hosts in the cluster, see Uninstall NSX-T Data Center from a Host in a vSphere Cluster.

Note Removing a host from a cluster does not uninstall NSX-T Data Center. Follow these instructions to uninstall NSX-T Data Center from a host in a cluster: Uninstall NSX-T Data Center from a Host in a vSphere Cluster.

Prerequisites

- Verify that the hosts you want to uninstall have network uninstall mappings configured. See Verify Host Network Mappings for Uninstall.

- Verify that the hosts you want to uninstall are in maintenance mode in vSphere.

Procedure

1 From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Nodes > Host Transport Nodes**.

3   From the **Managed by** drop-down menu, select the vCenter Server.

4   Select the cluster you want to uninstall, and click **Remove NSX**.

5   Verify that the NSX-T Data Center software is removed from the host.

    a   Log into the host's command-line interface as root.

    b   Run this command to check for NSX-T Data Center VIBs

```
esxcli software vib list | grep –E 'nsx|vsipfwlib'
```

    If the NSX-T Data Center software is successfully removed, no VIBs are listed. If any NSX VIBs remain on the host, contact VMware Support.

Uninstall NSX-T Data Center from a Host in a vSphere Cluster

You can uninstall NSX-T Data Center from a single host that is managed by vCenter Server. The other hosts in the cluster are not affected.

Caution   Uninstalling NSX-T Data Center from an ESXi host is disruptive if the physical interfaces or VMkernel interfaces are connected to N-VDS. If the host or cluster is participating in other applications such as vSAN, those applications might be affected by the uninstall.

Prerequisites

- Verify that the hosts you want to uninstall have network uninstall mappings configured. See Verify Host Network Mappings for Uninstall.

- Verify that the hosts you want to uninstall are in maintenance mode in vSphere.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Nodes > Host Transport Nodes**.

3   From the **Managed by** drop-down menu, select the vCenter Server.

4   If the cluster has a transport node profile applied, select the cluster, and click **Actions > Detach TN Profile**.

    If the cluster has a transport node profile applied, the **NSX Configuration** column for the cluster displays the profile name.

5   Select the host and click **Remove NSX**.

6   Verify that the NSX-T Data Center software is removed from the host.

    a   Log into the host's command-line interface as root.

    b   Run this command to check for NSX-T Data Center VIBs

```
esxcli software vib list | grep —E 'nsx|vsipfwlib'
```

    If the NSX-T Data Center software is successfully removed, no VIBs are listed. If any NSX VIBs remain on the host, contact VMware Support.

7   If the cluster had a Transport Node Profile applied, and you want to reapply it, select the cluster, click **Configure NSX**, and select the profile from the **Select Deployment Profile** drop-down menu.

Uninstall NSX-T Data Center from a Standalone Host

You can uninstall NSX-T Data Center from a standalone host. Standalone hosts can be ESXi or KVM.

Caution   Uninstalling NSX-T Data Center from an ESXi host is disruptive if the physical interfaces or VMkernel interfaces are connected to N-VDS. If the host or cluster is participating in other applications such as vSAN, those applications might be affected by the uninstall.

Prerequisites

If you are uninstalling NSX-T Data Center from a standalone ESXi host, verify the following settings:

- Verify that the hosts you want to uninstall have network uninstall mappings configured. See Verify Host Network Mappings for Uninstall.

- Verify that the hosts you want to uninstall are in maintenance mode in vSphere.

Procedure

1   From a browser, log in with admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2   Select **System > Fabric > Nodes > Host Transport Nodes**.

3   From the **Managed by** drop-down menu, select **None: Standalone Hosts**.

4   Select the host and click **Delete**. In the confirmation dialog box that is displayed, make sure **Uninstall NSX Components** is selected, and **Force Delete** is deselected. Click **Delete**.

    The NSX-T Data Center software is removed from the host. It might take up to 5 minutes for all NSX-T Data Center software to be removed.

5 If the uninstall fails, select the host and click **Delete** again. In the confirmation dialog box, deselect **Uninstall NSX Components** and select **Force Delete**.

The host transport node is deleted from the management plane, but the host might still have NSX-T Data Center software installed.

6 Verify that the NSX-T Data Center software is removed from the host.

a Log into the host's command-line interface as root.

b Run the appropriate command to check for NSX-T Data Center software packages.

Table 10-1. Package List Commands

Host Operating System	Command
ESXi	`esxcli software vib list	grep -E 'nsx|vsipfwlib'`
Red Hat Enterprise Linux and CentOS Linux	`rpm -qa	grep -E 'nsx|vsipfwlib'`
Ubuntu	`dpkg -l	grep -E 'nsx|vsipfwlib'`
SUSE Linux Enterprise Server	`zypper packages --installed-only	grep -E 'nsx|vsipfwlib'`

If the NSX-T Data Center software is successfully removed, no packages are listed. If any NSX software packages remain on the host, contact VMware Support.

Installing NSX Cloud Components 11

NSX Cloud provides a single pane of glass for managing your public cloud networks.

NSX Cloud is agnostic of provider-specific networking that does not require hypervisor access in a public cloud.

It offers several benefits:

- You can develop and test applications using the same network and security profiles used in the production environment.

- Developers can manage their applications until they are ready for deployment.

- With disaster recovery, you can recover from an unplanned outage or a security threat to your public cloud.

- If you migrate your workloads between public clouds, NSX Cloud ensures that similar security policies are applied to workload VMs regardless of their new location.

This chapter includes the following topics:

- NSX Cloud Architecture and Components
- Overview of Installing and Configuring NSX Cloud Components for your Public Cloud
- Install CSM and Connect with NSX Manager
- Connect Public Cloud with On-prem Deployment
- Add your Public Cloud Account
- Deploy or Link NSX Public Cloud Gateways
- Undeploy PCG

NSX Cloud Architecture and Components

NSX Cloud integrates the NSX-T Data Center core components with your public cloud to provide network and security across your implementations.

Figure 11-1. NSX Cloud Architecture
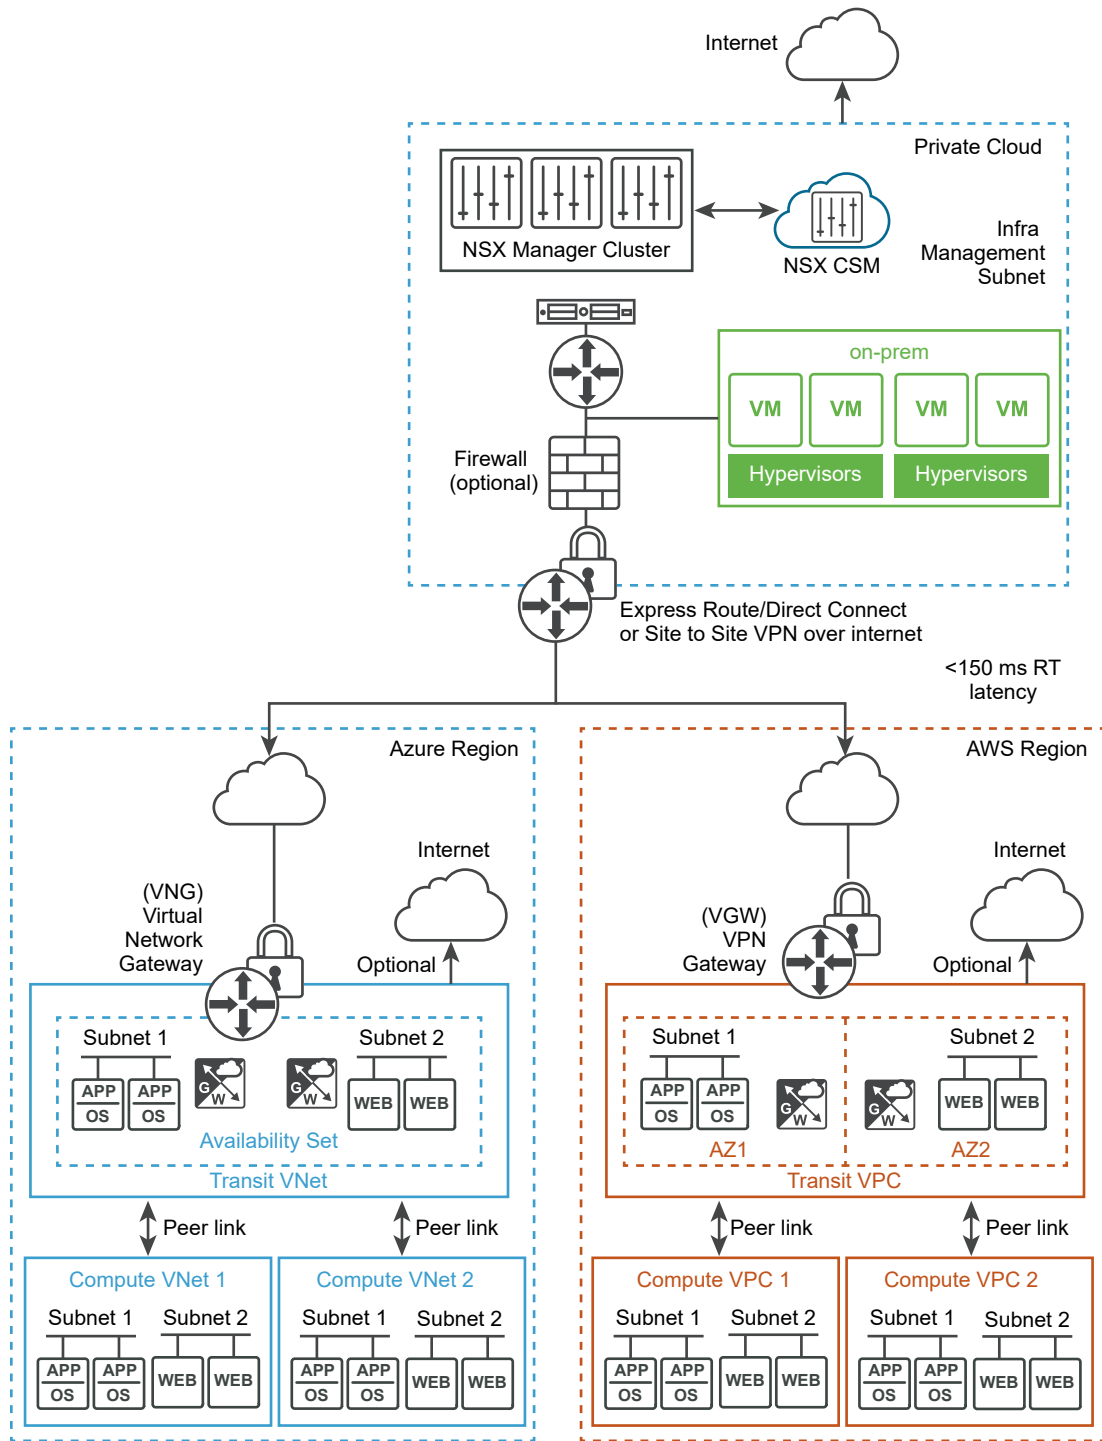
Core Components

The core NSX Cloud components are:

- *NSX Manager* for the management plane with policy-based routing, role-based access control (RBAC), control plane and runtime states defined.

- *Cloud Service Manager (CSM)* for integration with NSX Manager to provide public cloud-specific information to the management plane.

- *NSX Public Cloud Gateway (PCG)* for connectivity to the NSX management and control planes, NSX Edge gateway services, and for API-based communications with the public cloud entities. See Deploy or Link NSX Public Cloud Gateways for details.

- *NSX Agent* functionality that provides NSX-managed datapath for workload VMs.

Modes of Deployment

The NSX Public Cloud Gateway can either be a standalone gateway appliance or a shared between your public cloud VPCs or VNets to achieve a hub and spoke topology.

Self-managed VPC or VNet serves as a Transit VPC: When you deploy the PCG in a VPC or VNet, it qualifies the VPC or VNet as self-managed, that is, you can bring VMs hosted in this VPC or VNet under NSX management. This VPC or VNet also qualifies as a Transit VPC or VNet because you can use the PCG deployed on it to onboard VMs hosted in other VPCs or VNets.

Compute VPC or VNet links to Transit VPC or VNet: VPCs or VNets that do not have the PCG deployed on them but link to a Transit VPC or VNet are called *Compute* VPCs or VNets.

Overview of Installing and Configuring NSX Cloud Components for your Public Cloud

Refer to the checklist for an overview of the steps involved in enabling NSX-T Data Center to manage your workload VMs in the public cloud.

Day-0 Workflow for Connecting NSX Cloud with your Public Cloud

This workflow provides an overview of the steps required to get started with NSX Cloud for your public cloud.

Note While planning your deployment, ensure that the on-prem NSX-T Data Center appliances have good connectivity with the PCG deployed in the public cloud. Also, Transit VPCs/VNets must be in the same region as the Compute VPCs/VNets.

Table 11-1. Day-0 Workflow for connecting NSX Cloud with your public cloud

Task	Instructions
☐ Install CSM and connect with NSX Manager.	See Install CSM and Connect with NSX Manager.
☐ Add one or more of your public cloud accounts in CSM.	See Add your Public Cloud Account.

Table 11-1. Day-0 Workflow for connecting NSX Cloud with your public cloud (continued)

Task	Instructions
❏ Deploy PCG in your Transit VPCs or VNets and link to your Compute VPCs or VNets.	See Deploy or Link NSX Public Cloud Gateways.
❏ Onboard workload VMs by tagging in your public cloud and installing the NSX Agent on them.	Follow instructions at Onboard Workload VMs in the *NSX-T Data Center Administration Guide.*

Install CSM and Connect with NSX Manager

Use the Setup Wizard to connect CSM with NSX Manager and set up proxy servers, if any.

Install CSM

The Cloud Service Manager (CSM) is an essential component of NSX Cloud.

Install CSM after installing the core NSX-T Data Center components.

See Install NSX Manager and Available Appliances for detailed instructions.

Note Installing NSX Cloud requires you to enable FQDN usage (DNS) on NSX Managers. See Publishing the FQDNs of the NSX Managers.

Join CSM with NSX Manager

You must connect the CSM appliance with NSX Manager to allow these components to communicate with each other.

Prerequisites

- NSX Manager must be installed and you must have the username and password for the admin account to log in to NSX Manager

- CSM must be installed and you must have the Enterprise Administrator role assigned in CSM.

Procedure

1 From a browser, log in to CSM.

2 When prompted in the setup wizard, click **Begin Setup**.

3 Enter the following details in the NSX Manager Credentials screen:

Option	Description
NSX Manager Host Name	Enter the fully qualified domain name (FQDN) of the NSX Manager, if available. You may also enter the IP address of NSX Manager.
Admin Credentials	Enter an Enterprise Administrator username and password for NSX Manager.
Manager Thumbprint	Optionally, enter the NSX Manager's thumbrpint value. If you leave this field blank, the system identifies the thumbprint and displays it in the next screen.

4 (Optional) If you did not provide a thumbprint value for NSX Manager, or if the value was incorrect, the **Verify Thumbprint** screen appears. Select the checkbox to accept the thumbprint discovered by the system.

5 Click **Connect**.

> **Note** If you missed this setting in the setup wizard or if you want to change the associated NSX Manager, log in to CSM, click **System > Settings**, and click **Configure** on the panel titled **Associated NSX Node**.

CSM verifies the NSX Manager thumbprint and establishes connection.

6 (Optional) Set up the Proxy server. See instructions in (Optional) Configure Proxy Servers.

(Optional) Configure Proxy Servers

If you want to route and monitor all internet-bound HTTP/HTTPS traffic through a reliable HTTP Proxy, you can configure up to five proxy servers in CSM.

All public cloud communication from PCG and CSM is routed through the selected proxy server.

Proxy settings for PCG are independent of proxy settings for CSM. You can choose to have none or a different proxy server for PCG.

You can choose the following levels of authentication:

- Credentials-based authentication.

- Certificate-based authentication for HTTPS interception.

- No authentication.

Procedure

1 Click **System > Settings**. Then click **Configure** on the panel titled **Proxy Servers**.

> **Note** You can also provide these details when using the CSM Setup Wizard that is available when you first install CSM.

2   In the Configure Proxy Servers screen, enter the following details:

Option	Description
Default	Use this radio button to indicate the default proxy server.
Profile Name	Provide a proxy server profile name. This is mandatory.
Proxy Server	Enter the proxy server's IP address. This is mandatory.
Port	Enter the proxy server's port. This is mandatory.
Authentication	Optional. If you want to set up additional authentication, select this check box and provide valid username and password.
Username	This is required if you select the Authentication checkbox.
Password	This is required if you select the Authentication checkbox.
Certificate	Optional. If you want to provide an authentication certificate for HTTPS interception, select this checkbox and copy-paste the certificate in the text box that appears.
No Proxy	Select this option if you do not want to use any of the proxy servers configured.

(Optional) Set Up vIDM for Cloud Service Manager

If you use VMware Identity Manager, you can set it up to access CSM from within NSX Manager.

Procedure

1   Configure vIDM for NSX Manager and CSM. See instructions at Configure VMware Identity Manager Integration in the *NSX-T Data Center Administration Guide*.

2   Assign the same role to the vIDM user for NSX Manager and CSM, for example, **Enterprise Admin** role assigned to the user named `vIDM_admin`. You must log in to NSX Manager and CSM each and assign the same role to the same username. See Add a Role Assignment or Principal Identity in the *NSX-T Data Center Administration Guide* for detailed instructions.

3   Log in to NSX Manager. You are redirected to the vIDM login.

4   Enter the vIDM user's credentails. Once you log in, you can switch between NSX Manager and CSM by clicking the Applications icon.

Connect Public Cloud with On-prem Deployment

You must use suitable connectivity options to connect your on-prem deployment with your public cloud accounts or subscriptions.

Enable Access to ports and protocols on CSM for Hybrid Connectivity

Open up necessary network ports and allow the required protocols on NSX Manager to enable public cloud connectivity.

Allow access to NSX Manager from the Public Cloud

Open up the following network ports and protocols to allow connectivity with your on-prem NSX Manager deployment:

Table 11-2.

From	To	Protocol/Port	Description
PCG	NSX Manager	TCP/5671	Inbound traffic from public cloud to on-prem NSX-T Data Center for Management Plane Communication.
PCG	NSX Manager	TCP/8080	Inbound traffic from public cloud to on-prem NSX-T Data Center for access to an HTTP repository for upgrading NSX Cloudcomponents.
PCG	NSX Controller	TCP/1234, TCP/1235	Inbound traffic from public cloud to on-prem NSX-T Data Center for Control Plane Communication.
PCG	DNS	UDP/53	Inbound traffic from public cloud to on-prem NSX-T Data Center DNS, (if you are using the on-prem DNS Server).
CSM	PCG	TCP/7442	CSM Config Push
Any	NSX Manager	TCP/443	NSX Manager UI
Any	CSM	TCP/443	CSM UI.

Important   All NSX-T Data Center infrastructure communication leverages SSL-based encryption. Ensure your firewall allows SSL traffic over non-standard ports.

Connect your Microsoft Azure Network with your On-prem NSX-T Data Center Deployment

A connection must be established between your Microsoft Azure network and your on-prem NSX-T Data Center appliances.

Note You must have already installed and connected NSX Manager with CSM in your on-prem deployment.

Overview

- Connect your Microsoft Azure subscription with on-prem NSX-T Data Center.

- Configure your VNets with the necessary CIDR blocks and subnets required by NSX Cloud.

- Synchronize time on the CSM appliance with the Microsoft Azure Storage server or NTP.

Connect your Microsoft Azure subscription with on-prem NSX-T Data Center

Every public cloud provides options to connect with an on-premises deployment. You can choose any of the available connectivity options that suit your requirements. See Microsoft Azure reference documentation for details.

Note You must review and implement the applicable security considerations and best practices by Microsoft Azure, for example, all privileged user accounts accessing the Microsoft Azure portal or API should have Multi Factor Authentication (MFA) enabled. MFA ensures only a legitimate user can access the portal and reduces the likelihood of access even if credentials are stolen or leaked. For more information and recommendations, refer to the Azure Security Center Documentation.

Configure your VNet

In Microsoft Azure, create routable CIDR blocks and set up the required subnets.

- One management subnet with a recommended range of at least /28, to handle:

  - control traffic to on-prem appliances

  - API traffic to cloud-provider API endpoints

- One downlink subnet with a recommended range of /24, for the workload VMs.

- One, or two for HA, uplink subnets with a recommended range of /24, for routing of north-south traffic leaving from or entering the VNet.

See Deploy or Link NSX Public Cloud Gateways for details on how these subnets are used.

Connect your Amazon Web Services (AWS) Network with your On-prem NSX-T Data Center Deployment

A connection must be established between your Amazon Web Services (AWS) network and your on-prem NSX-T Data Center appliances.

Note You must have already installed and connected NSX Manager with CSM in your on-prem deployment.

Overview

- Connect your AWS account with on-prem NSX Manager appliances using any of the available options that best suit your requirements.

- Configure your VPC with subnets and other requirements for NSX Cloud.

Connect your AWS account with your on-prem NSX-T Data Center deployment

Every public cloud provides options to connect with an on-premises deployment. You can choose any of the available connectivity options that suit your requirements. See AWS reference documentation for details.

Note You must review and implement the applicable security considerations and best practices by AWS; see AWS Security Best Practices.

Configure your VPC

You need the following configurations:

- six subnets for supporting PCG with High Availability

- an Internet gateway (IGW)

- a private and a public route table

- subnet association with route tables

- DNS resolution and DNS hostnames enabled

Follow these guidelines to configure your VPC:

1 Assuming your VPC uses a /16 network, for each gateway that needs to be deployed, set up three subnets.

   Important If using High Availability, set up three additional subnets in a different Availability Zone.

   - **Management subnet**: This subnet is used for management traffic between on-prem NSX-T Data Center and PCG. The recommended range is /28.

- **Uplink subnet**: This subnet is used for north-south internet traffic. The recommended range is /24.

- **Downlink subnet**: This subnet encompasses the workload VM's IP address range, and should be sized accordingly. Bear in mind that you may need to incorporate additional interfaces on the workload VMs for debugging purposes.

Note Label the subnets appropriately, for example, `management-subnet`, `uplink-subnet`, `downlink-subnet`,because you will need to select the subnets when deploying PCG on this VPC.

See Deploy or Link NSX Public Cloud Gateways for details.

2   Ensure you have an Internet gateway (IGW) that is attached to this VPC.

3   Ensure the routing table for the VPC has the **Destination** set to `0.0.0.0/0` and the **Target** is the IGW attached to the VPC.

4   Ensure you have DNS resolution and DNS hostnames enabled for this VPC.

Add your Public Cloud Account

To add your public cloud inventory, you need to create roles in your public cloud to allow access to NSX Cloud and then add the required information in CSM.

Set Up Secure Access to Your Microsoft Azure Inventory

For NSX Cloud to operate in your subscription, create a Service Principal to grant the required permissions, and roles for CSM and PCG based on the Microsoft Azure feature for managing identities for Azure Resources.

Note If you already added an AWS account to CSM, update the MTU in **NSX Manager > Fabric > Profiles > Uplink Profiles > PCG-Uplink-HostSwitch-Profile** to 1500 before adding the Microsoft Azure account. This can also be done using the NSX Manager REST APIs.

Overview:

- Your Microsoft Azure subscription contains one or more VNets that you want to bring under NSX-T Data Center management. The VNet might be in Transit mode or Compute mode. Transit VNet is one in which you deploy the PCG. You can link other VNets to the Transit VNet and onboard workload VMs hosted in them. The VNets linked to the Transit VNet are called Compute VNets.

- NSX Cloud provides a PowerShell script to generate the Service Principal and roles that use the managed identity feature of Microsoft Azure to manage authentication while keeping your Microsoft Azure credentials secure. You can also include multiple subscriptions under one Service Principal using this script.

- You have the option of reusing the Service Principal for all your subscriptions, or to create new Service Principals as required. There is an additional script if you want to create separate Service Principals for additional subscriptions.

- For multiple subscriptions, whether you are using a single Service Principal for all, or multiple Service Principals, you must update the JSON files for the CSM and PCG roles to add each additional subscription name under the section *AssignableScopes*.

- If you already have an NSX Cloud Service Principal in your VNet, you can update it by running the scripts again and leaving out the Service Principal name from the parameters.

- The Service Principal name must be unique for your Microsoft Azure Active Directory. You may use the same Service Principal in different subscriptions under the same Active Directory domain, or different Service Principals per subscription. But you cannot create two Service Principals with the same name.

- You must either be the owner of or have permissions to create and assign roles in all the Microsoft Azure subscriptions.

- The following scenarios are supported:

  - **Scenario 1:** You have a single Microsoft Azure Subscription that you want to enable with NSX Cloud.

  - **Scenario 2:** You have multiple Microsoft Azure Subscriptions under the same Microsoft Azure Directory, that you want to enable with NSX Cloud, but want to use one NSX Cloud Service Principal across all your subscriptions.

  - **Scenario 3:** You have multiple Microsoft Azure Subscriptions under the same Microsoft Azure Directory, that you want to enable with NSX Cloud, but want to use different NSX Cloud Service Principal names for different subscriptions.

Here is an outline of the process:

1 Use the NSX Cloud PowerShell script to:

  - Create a Service Principal account for NSX Cloud.

  - Create a role for CSM.

  - Create a role for PCG.

2 (Optional) Create Service Principals for other subscriptions you want to link.

3 Add the Microsoft Azure subscription in CSM.

   Note If using multiple subscriptions, whether using the same or different Service Principals, you must add each subscription separately in CSM.

Generate the Service Principal and Roles

NSX Cloud provides PowerShell scripts that help you generate the required service principal and roles for one or multiple subscriptions.

Prerequisites

■ You must have PowerShell 5.0+ with the AzureRM Module installed.

■ You must either be the owner of or have permissions to create and assign roles in all the Microsoft Azure subscriptions.

Note The response time from Microsoft Azure can cause the script to fail when you run it the first time. If the script fails, try running it again.

Procedure

1 On a Windows desktop or server, download the ZIP file named `CreateNSXCloudCredentials.zip` from the NSX-T Data Center **Download page > Drivers & Tools > NSX Cloud Scripts > Microsoft Azure**.

2 Extract the following contents of the ZIP file in your Windows system:

Script/File	Description
CreateNSXRoles.ps1	The PowerShell script to generate the NSX Cloud Service Principal and managed identity roles for CSM and PCG. This script takes the following parameters: ■ –subscriptionId <the Transit_VNet's_Azure_subscription_ID> ■ (optional) –servicePrincipalName <Service_Principal_Name> ■ (optional) –useOneServicePrincipal
AddServicePrincipal.ps1	An optional script required if you want to add multiple subscriptions and assign different Service Principals to each subscription. See **Scenario 3** in the following steps. This script takes the following parameters: ■ –computeSubscriptionId <the_Compute_VNet's_Azure_subscription_ID> ■ –transitSubscriptionId <the Transit_VNet's_Azure_Subscription_ID> ■ –csmRoleName <CSM_Role_Name> ■ –servicePrincipalName <Service_Principal_Name>
nsx_csm_role.json	A JSON template for the CSM role name and permissions. This file is required as an input to the PowerShell script and must be in the same folder as the script.
nsx_pcg_role.json	A JSON template for the PCG role name and permissions. This file is required as an input to the PowerShell script and must be in the same folder as the script. **Note** The default PCG (Gateway) Role Name is `nsx-pcg-role`. You need to provide this value when adding your subscription in CSM.

3 **Scenario 1:** You have a single Microsoft Azure Subscription that you want to enable with NSX Cloud.

   a   From a PowerShell instance, go to the directory where you downloaded the Microsoft Azure scripts and JSON files.

   b   Run the script named CreateNSXRoles.ps1 with the parameter -SubscriptionId, as follows:

```
.\CreateNSXRoles.ps1 —subscriptionId <the_single_Azure_subscription_ID>
```

   Note If you want to override the default Service Principal name of `nsx—service—admin`, you can also use the parameter `-servicePrincipalName`. The Service Principal name must be unique in your Microsoft Azure Active Directory.

4 **Scenario 2:** You have multiple Microsoft Azure Subscriptions under the same Microsoft Azure Directory, that you want to enable with NSX Cloud, but want to use one NSX Cloud Service Principal across all your subscriptions.

   a   From a PowerShell instance, go to the directory where you downloaded the Microsoft Azure scripts and JSON files.

   b   Edit each of the JSON files to add a list of other subscription IDs under the section titled *"AssignableScopes"*, for example:

```
"AssignableScopes": [

"/subscriptions/aaaaaaaa—bbbb—cccc—dddd—eeeeeeeeeeee",

"/subscriptions/aaaaaaaa—bbbb—cccc—dddd—ffffffffffff",

"/subscriptions/aaaaaaaa—bbbb—cccc—dddd—000000000000"
```

   Note You must use the format shown in the example to add subscription IDs: `"/subscriptions/<Subscription_ID>"`

   c   Run the script named `CreateNSXRoles.ps1` with the parameters `—subscriptionID` and `—useOneServicePrincipal`:

```
.\CreateNSXRoles.ps1 —subscriptionId <the_Transit_VNet's_Azure_subscription_ID> —
useOneServicePrincipal
```

   Note Omit the Service Principal name here if you want to use the default name: `nsx—service—admin`. If that Service Principal name already exists in your Microsoft Azure Active Directory, running this script without a Service Principal name updates that Service Principal.

5   **Scenario 3:** You have multiple Microsoft Azure Subscriptions under the same Microsoft Azure Directory, that you want to enable with NSX Cloud, but want to use different NSX Cloud Service Principal names for different subscriptions.

   a   From a PowerShell instance, go to the directory where you downloaded the Microsoft Azure scripts and JSON files.

   b   Follow steps **b** and **c** from the second scenario to add multiple subscriptions to the *AssignableScopes* section in each of the JSON files.

   c   Run the script named `CreateNSXRoles.ps1` with the parameters -subscriptionID:

   ```
   .\CreateNSXRoles.ps1 -subscriptionId <One of the subscription_IDs>
   ```

   Note   Omit the Service Principal name here if you want to use the default name: `nsx-service-admin`. If that Service Principal name exists in your Microsoft Azure Active Directory, running this script without a Service Principal name updates that Service Principal.

   d   Run the script named `AddServicePrincipal.ps1` with the following parameters:

Parameter	Value
`-computeSubscriptionId`	The Compute_VNet's Azure Subscription ID
`-transitSubscriptionId`	The Transit VNet's Azure Subscription ID
`-csmRoleName`	Get this value from the file `nsx_csm_role.JSON`
`-servicePrincipalName`	New Service Principal name

   ```
   ./AddServicePrincipal.ps1 -computeSubscriptionId <the_Compute_VNet's_Azure_subscription_ID>
    -transitSubscriptionId <the_Tranist_VNet's_Azure_Subscription_ID>
   -csmRoleName <CSM_Role_Name>
   -servicePrincipalName <new_Service_Principal_Name>"
   ```

6   Look for a file in the same directory where you ran the PowerShell script. It is named like: `NSXCloud_ServicePrincipal_<your_subscription_ID>_<NSX_Cloud_Service_Principal_name>`. This file contains the information required to add your Microsoft Azure subscription in CSM.

   ■   Client ID

   ■   Client Key

   ■   Tenant ID

   ■   Subscription ID

Results

The following constructs are created:

■   an Azure AD application for NSX Cloud.

- an Azure Resource Manager Service Principal for the NSX Cloud application.

- a role for CSM attached to the Service Principal account.

- a role for PCG to enable it to work on your public cloud inventory.

- a file named like
  `NSXCloud_ServicePrincipal_<your_subscription_ID>_<NSX_Cloud_Service_Principal_name>` is created in the same directory where you ran the PowerShell script. This file contains the information required to add your Microsoft Azure subscription in CSM.

Note Refer to the JSON files that are used to create the CSM and PCG roles for a list of permissions available to them after the roles are created.

What to do next

Add your Microsoft Azure Subscription in CSM

Note When enabling NSX Cloud for multiple subscriptions, you must add each separate subscription to CSM individually, for example, if you have five total subscriptions you must add five Microsoft Azure accounts in CSM with all other values the same but different subscription IDs.

Add your Microsoft Azure Subscription in CSM

Once you have the details of the NSX Cloud Service Principal and the CSM and PCG roles, you are ready to add your Microsoft Azure subscription in CSM.

Prerequisites

- You must have the Enterprise Administrator role in NSX-T Data Center.

- You must have the output of the PowerShell script with details of the NSX Cloud Service Principal.

- You must have the value of the PCG role you provided when running the PowerShell script to create the roles and the Service Principal. The default value is `nsx-pcg-role`.

Procedure

1 Log in to CSM using an account with the Enterprise Administrator role.

2 Go to **CSM > Clouds > Azure**.

3 Click **+Add** and enter the following details:

Option	Description
Name	Provide a suitable name to identify this account in CSM. You may have multiple Microsoft Azure subscriptions that are associated with the same Microsoft Azure tenant ID. Name your account account and you can name them appropriately in CSM, for example, Azure-DevOps-Account, Azure-Finance-Account, etc.
Client ID	Copy paste this value from the output of the PowerShell script.

Option	Description
Key	Copy paste this value from the output of the PowerShell script.
Subscription ID	Copy paste this value from the output of the PowerShell script.
Tenant ID	Copy paste this value from the output of the PowerShell script.
Gateway Role Name	The default value is `nsx-pcg-role`. This value is available from the *nsx_pcg_role.json* file if you changed the default.
Cloud Tags	By default this option is enabled and allows your Microsoft Azure tags to be visible in NSX Manager

4 Click **Save**.

CSM adds the account and you can see it in the **Accounts** section within three minutes.

What to do next

Set Up Secure Access to Your AWS Inventory

You might have one or more AWS accounts with VPCs and workload VMs that you want to bring under NSX-T Data Center management.

Overview:

- You can use the Transit/Compute VPC topology where you deploy the PCG in one VPC, making it the Transit VPC, and link other VPCs to it, which are called Compute VPCs.

- NSX Cloud provides a shell script that you can run from the AWS CLI of your AWS account to create the IAM profile and role, and create a trust relationship for Transit and Compute VPCs .

- The following scenarios are supported:

  - **Scenario 1:** You want to use a single AWS account with NSX Cloud.

  - **Scenario 2:** You want to use multiple sub-accounts in AWS that are managed by a master AWS account.

  - **Scenario 3:** You want to use multiple AWS accounts with NSX Cloud.

Here is an outline of the process:

1 Use the NSX Cloud shell script, that requires AWS CLI, to do the following:

- Create an IAM profile.

- Create a role for PCG.

- (Optional) Create a trust relationship between the AWS account hosting the Transit VPC and the AWS account hosting the Compute VPC.

2 Add the AWS account in CSM.

Generate the IAM Profile and PCG Role

NSX Cloud provides a SHELL script to help set up one or more of your AWS accounts by generating an IAM profile and a role for PCG attached to the profile that provides necessary permissions to your AWS account.

If you plan to host a Transit VPC linked to multiple Compute VPCs in two different AWS accounts, you can use the script to create a trust relationship between these accounts.

Note The PCG (Gateway) role name is `nsx_pcg_service` by default. If you want a different value for the Gateway Role Name, you can change it in the script, but make a note of this value because it is required for adding the AWS account in CSM.

Prerequisites

You must have the following installed and configured on your Linux or compatible system before you run the script:

- AWS CLI

- jq (A JSON parser)

- `openssl`

Note If using multiple AWS accounts, the accounts must be peered using a suitable method.

Procedure

1 On a Linux or compatible desktop or server, download the SHELL script named `nsx_csm_iam_script.sh` from the NSX-T Data Center **Download page > Drivers & Tools > NSX Cloud Scripts > AWS**.

2 **Scenario 1:** You want to use a single AWS account with NSX Cloud.

  a Run the script, for example:

```
bash nsx_csm_iam_script.sh
```

  b Enter yes when prompted with the question `Do you want to create an IAM user for CSM and an IAM role for PCG? [yes/no]`

  c Enter a name for the IAM user when asked `What do you want to name the IAM User?`

   Note The IAM user name must be unique in your AWS account.

  d Enter `no` when asked `Do you want to add trust relationship for any Transit VPC account? [yes/no]`

When the script runs successfully, the IAM profile and a role for PCG is created in your AWS account. The values are saved in the output file named `aws_details.txt` in the same directory where you ran the script. Next, follow instructions at Add your AWS Account in CSM and then Deploy PCG in a Self-Managed or Transit VPC to finish the process of setting up a Transit or Self-Managed VPC.

3 **Scenario 2:** You want to use multiple sub-accounts in AWS that are managed by one master AWS account.

   a Run the script from your AWS master account.

   ```
   bash nsx_csm_iam_script.sh
   ```

   b Enter yes when prompted with the question `Do you want to create an IAM user for CSM and an IAM role for PCG? [yes/no]`

   c Enter a name for the IAM user when asked `What do you want to name the IAM User?`

      Note The IAM user name must be unique in your AWS account.

   d Enter `no` when asked `Do you want to add trust relationship for any Transit VPC account? [yes/no]`

      Note With a master AWS account, if your Transit VPC has permission to view Compute VPCs in the sub-accounts, you do not need to establish a trust relationship with your sub-accounts. If not, follow the steps for **Scenario 3** to set up multiple accounts.

   When the script runs successfully, the IAM profile and a role for PCG is created in your AWS master account. The values are saved in the output file in the same directory where you ran the script. The filename is `aws_details.txt`. Next, follow instructions at Add your AWS Account in CSM and then Deploy PCG in a Self-Managed or Transit VPC to finish the process of setting up a Transit or Self-Managed VPC.

4 **Scenario 3:** You want to use multiple AWS accounts with NSX Cloud.

   Note Verify that the AWS accounts are peered before you proceed.

   a Make a note of the 12-digit AWS account number where you want to host the Transit VPC.

   b Set up the Transit VPC in the AWS account by following steps *a* through *d* for *Scenario 1* and finish the process of adding the account in CSM and deploying a PCG in it.

   c Download and run the NSX Cloud script from a Linux or compatible system in your other AWS account where you want to host the Compute VPCs.

      Note Alternatively, you can use AWS profiles with different account credentials to use the same system to run the script again for your other AWS account.

d   Enter yes when asked `Do you want to create an IAM user for CSM and an IAM role for PCG? [yes/no]`

> **Note** If you already added this AWS account into CSM and want to reuse the script to connect to a different AWS account, you can enter `no` and skip the creation of the IAM user.

e   Enter a name for the IAM user when asked `What do you want to name the IAM User?`

> **Note** The IAM user name must be unique in your AWS account.

f   Enter yes when asked `Do you want to add trust relationship for any Transit VPC account? [yes/no]`

g   Enter or copy-paste the 12-digit AWS account number that you noted in step 1 when asked `What is the Transit VPC account number?`

An IAM Trust Relationship is established between the two AWS accounts and an ExternalID is generated by the script.

When the script runs successfully, the IAM profile and a role for PCG is created in your AWS master account. The values are saved in the output file in the same directory where you ran the script. The filename is `aws_details.txt`. Next, follow instructions at Add your AWS Account in CSM and then Link to a Transit VPC or VNet to finish the process of linking to a Transit VPC.

Add your AWS Account in CSM

Add your AWS account using values generated by the script.

Procedure

1   Log in to CSM using the Enterprise Administrator role.

2   Go to **CSM > Clouds > AWS**.

3   Click **+Add** and enter the following details using the output file aws_details.txt generated from the NSX Cloud script:

Option	Description
Name	Enter a descriptive name for this AWS Account
Access Key	Enter your account's Access Key
Secret Key	Enter your account's Secret Key
Cloud Tags	By default this option is enabled and allows your AWS tags to be visible in NSX Manager
Gateway Role Name	The default value is `nsx_pcg_service`. You can find this value in the output of the script in the file *aws_details.txt*.

Results

The AWS account gets added in CSM.

In the VPCs tab of CSM, you can view all the VPCs in your AWS account.

In the Instances tab of CSM, you can view the EC2 Instances in this VPC.

What to do next

Deploy or Link NSX Public Cloud Gateways

The NSX Public Cloud Gateway (PCG) provides north-south connectivity between the public cloud and the on-prem management components of NSX-T Data Center.

The PCG can either be a standalone gateway appliance or shared between your public cloud VPCs or VNets to achieve a hub and spoke topology.

Note The PCG is deployed in a single default size for each supported public cloud:

Public Cloud	PCG instance type
AWS	C4.xlarge **Note** Some regions may not support the C4.xlarge instance type. Refer to AWS documentation for details.
Microsoft Azure	Standard DS3 v.2

Figure 11-2. NSX Public Cloud Gateway Architecture
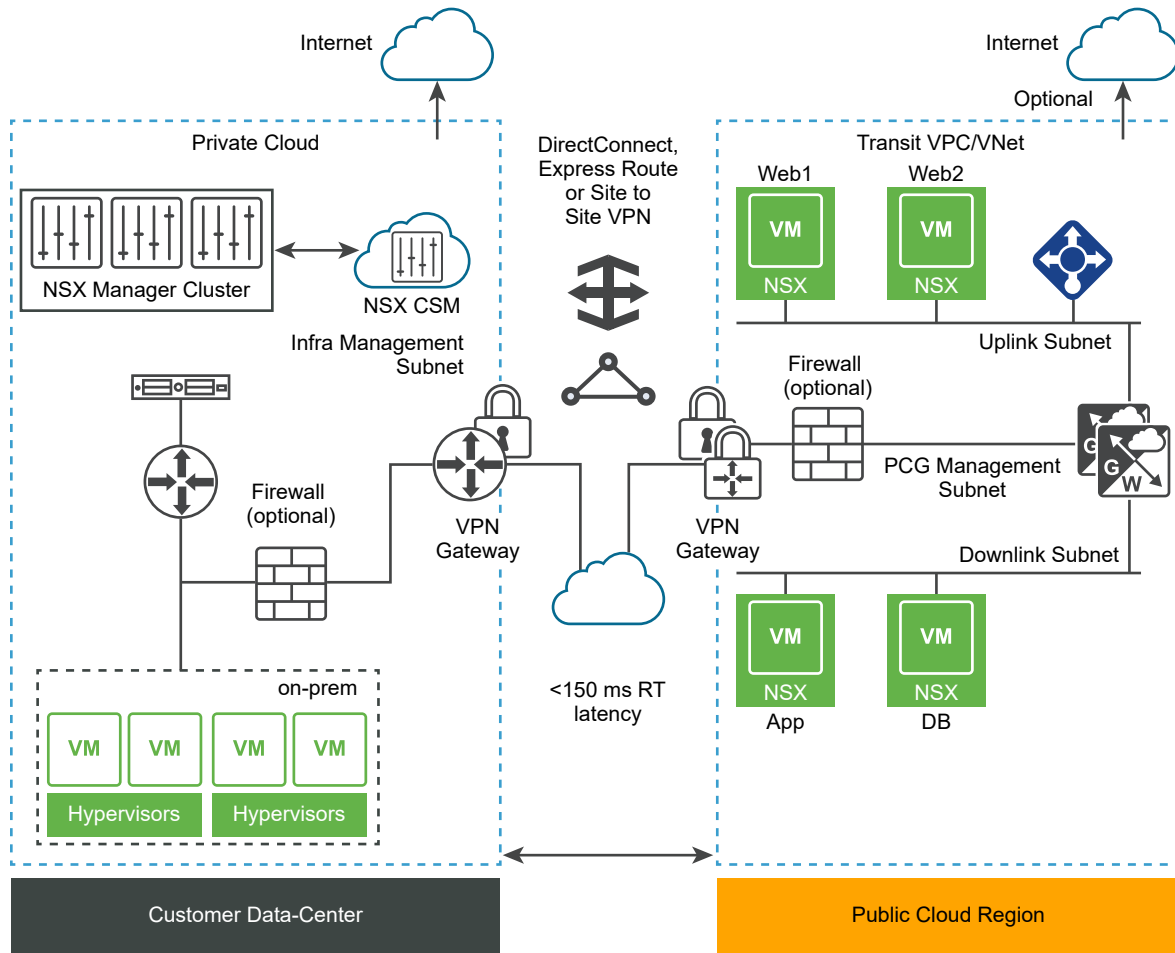
Transit or Self-managed VPC or VNet: When you deploy the PCG in a VPC or VNet, it qualifies the VPC or VNet as *self-managed*, that is, you can bring VMs hosted in this VPC or VNet under NSX management. This VPC or VNet also qualifies as a *Transit* VPC or VNet because you can use the PCG deployed on it to onboard VMs hosted in other VPCs or VNets. The PCG utilizes the following subnets that you set up in your VPC/VNet. See Connect your Microsoft Azure Network with your On-prem NSX-T Data Center Deployment or Connect your Amazon Web Services (AWS) Network with your On-prem NSX-T Data Center Deployment.

- **Management subnet**: This subnet is used for management traffic between on-prem NSX-T Data Center and PCG. The recommended range is /28.

- **Uplink subnet**: This subnet is used for north-south internet traffic. The recommended range is /24.

- **Downlink subnet**: This subnet encompasses the workload VM's IP address range, and should be sized accordingly. Bear in mind that you may need to incorporate additional interfaces on the workload VMs for debugging purposes.

Compute VPC or VNet: VPCs or VNets that do not have the PCG deployed on them but link to a Transit VPC or VNet are called *Compute* VPCs or VNets.

PCG deployment aligns with your network addressing plan with FQDNs for the NSX-T Data Center components and a DNS server that can resolve these FQDNs.

Note It is not recommended to use IP addresses for connecting the public cloud with NSX-T Data Center using PCG, but if you choose that option, do not change your IP addresses.

Deploy PCG in a Self-Managed or Transit VNet

Follow these instructions to deploy PCG in your Microsoft Azure VNet.

The VNet in which you deploy a PCG can act as a Transit VNet to which other VNets can connect (known as Compute VNets). This VNet can also manage VMs and act as a self-managed VNet.

Follow these instructions to deploy a PCG. If you want to link to an existing Transit VNet, see Link to a Transit VPC or VNet .

Prerequisites

- Your public cloud accounts must be already added into CSM.

- The VNet on which you are deploying PCG must have the required subnets appropriately adjusted for High Availability: *uplink*, *downlink*, and *management*.

Procedure

1 Log in to CSM using an account with the Enterprise Administrator role.

2 Click **Clouds > Azure** and go to the **VNets** tab.

3 Click a VNet where you want to deploy PCG.

4 Click **Deploy Gateways**. The **Deploy Primary Gateway** wizard opens.

5 For General Properties, use the following guidelines:

Option	Description
SSH Public Key	Provide an SSH public key that can be validated while deploying PCG. This is required for each PCG deployment.
Quarantine Policy on the Associated VNet	Leave this in the default **disabled** mode when you first deploy PCG. You can change this value after onboarding VMs. See **Manage Quarantine Policy** in the *NSX-T Data Center Administration Guide* for details.
Local Storage Account	When you add a Microsoft Azure subscription to CSM, a list of your Microsoft Azure Storage Accounts is available to CSM. Select the Storage Account from the drop-down menu. When proceeding with deploying PCG, CSM copies the publicly available VHD of the PCG into this Storage Account of the selected region.
	Note If the VHD image has been copied to this storage account in the region already for a previous PCG deployment, then the image is used from this location for subsequent deployments to reduce the overall deployment time.

Option	Description
VHD URL	If you want to use a different PCG image that is not available from the public VMware repository, you can enter the URL of the PCG's VHD here. The VHD must be present in the same account and region where this VNet is created.
	Note The VHD must be in the correct URL format. We recommend that you use the **Click to copy** option in Microsoft Azure.
Proxy Server	Select a proxy server to use for internet-bound traffic from this PCG. The proxy servers are configured in CSM. You can select the same proxy server as CSM if one, or select a different proxy server from CSM, or select **No Proxy Server**.
	See (Optional) Configure Proxy Servers for details on how to configure proxy servers in CSM.
Advanced	The advanced DNS settings provide flexibility in selecting DNS servers for resolving NSX-T Data Center management components.
Obtain via Public Cloud Provider's DHCP	Select this option if you want to use Microsoft Azure DNS settings. This is the default DNS setting if you do not pick either of the options to override it.
Override Public Cloud Provider's DNS Server	Select this option if you want to manually provide the IP address of one or more DNS servers to resolve NSX-T Data Center appliances as well as the workload VMs in this VNet.
Use Public Cloud Provider's DNS server only for NSX-T Data Center Appliances	Select this option if you want to use the Microsoft Azure DNS server for resolving the NSX-T Data Center management components. With this setting, you can use two DNS servers: one for PCG that resolves NSX-T Data Center appliances; the other for the VNet that resolves your workload VMs in this VNet.

6 Click **Next**.

7 For **Subnets**, use the following guidelines:

Option	Description
Enable HA for NSX Cloud Gateway	Select this option to enable High Availability.
Subnets	Select this option to enable High Availability.
Public IP on Mgmt NIC	Select **Allocate New IP address** to provide a public IP address to the management NIC. You can manually provide the public IP address if you want to reuse a free public IP address.
Public IP on Uplink NIC	Select **Allocate New IP address** to provide a public IP address to the uplink NIC. You can manually provide the public IP address if you want to reuse a free public IP address.

What to do next

Onboard your workload VMs. See **Onboarding and Managing Workload VMs** in the *NSX-T Data Center Administration Guide* for the Day-N workflow.

Deploy PCG in a Self-Managed or Transit VPC

Follow these instructions to deploy PCG in your AWS VPC.

The VPC in which you deploy a PCG can act as a Transit VPC to which other VPCs can connect (known as Compute VPCs). This VPC can also manage VMs and act as a self-managed VPC.

Follow these instructions to deploy a PCG. If you want to link to an existing Transit VPC, see Link to a Transit VPC or VNet .

Prerequisites

- Your public cloud accounts must be already added into CSM.

- The VPC on which you are deploying PCG must have the required subnets appropriately adjusted for High Availability: *uplink*, *downlink*, and *management*.

- The configuration for your VPC's network ACL must include an ALLOW inbound rule.

Procedure

1 Log in to CSM using an account with the Enterprise Administrator role.

2 Click **Clouds > AWS > `<AWS_account_name>`** and go to the **VPCs** tab.

3 In the **VPCs** tab, select an AWS region name, for example, `us-west`. The AWS region must be the same where you created the compute VPC.

4 Select a compute VPC configured for NSX Cloud.

5 Click Deploy Gateways.

6 Complete the general gateway details:

Option	Description
PEM File	Select one of your PEM files from the drop-down menu. This file must be in the same region where NSX Cloud was deployed and where you created your compute VPC. This uniquely identifies your AWS account.
Quarantine Policy on the Associated VPC	Leave this in the default **disabled** mode when you first deploy PCG. You can change this value after onboarding VMs. See **Manage Quarantine Policy** in the *NSX-T Data Center Administration Guide* for details.
Proxy Server	Select a proxy server to use for internet-bound traffic from this PCG. The proxy servers are configured in CSM. You can select the same proxy server as CSM if one, or select a different proxy server from CSM, or select **No Proxy Server**. See (Optional) Configure Proxy Servers for details on how to configure proxy servers in CSM.
Advanced	The advanced settings provide extra options if required.
Override AMI ID	Use this advanced feature to provide a different AMI ID for the PCG from the one that is available in your AWS account.
Obtain via Public Cloud Provider's DHCP	Select this option if you want to use AWS settings. This is the default DNS setting if you do not pick either of the options to override it.

Option	Description
Override Public Cloud Provider's DNS Server	Select this option if you want to manually provide the IP address of one or more DNS servers to resolve NSX-T Data Center appliances as well as the workload VMs in this VPC.
Use Public Cloud Provider's DNS server only for NSX-T Data Center Appliances	Select this option if you want to use the AWS DNS server for resolving the NSX-T Data Center management components. With this setting, you can use two DNS servers: one for PCG that resolves NSX-T Data Center appliances; the other for the VPC that resolves your workload VMs in this VPC.

7 Click Next.

8 Complete the Subnet details.

Option	Description
Enable HA for Public Cloud Gateway	The recommended setting is Enable, that sets up a High Availability Active/Standby pair to avoid an unscheduled downtime.
Primary gateway settings	Select an Availability Zone such as `us-west-1a`, from the drop-down menu as the primary gateway for HA. Assign the uplink, downlink, and management subnets from the drop-down menu.
Secondary gateway settings	Select another Availability Zone such as `us-west-1b`, from the drop-down menu as the secondary gateway for HA. The secondary gateway is used when the primary gateway fails. Assign the uplink, downlink, and management subnets from the drop-down menu.
Public IP on Mgmt NIC	Select **Allocate New IP address** to provide a public IP address to the management NIC. You can manually provide the public IP address if you want to reuse a free public IP address.
Public IP on Uplink NIC	Select **Allocate New IP address** to provide a public IP address to the uplink NIC. You can manually provide the public IP address if you want to reuse a free public IP address.

Click Deploy.

9 Monitor the status of the primary (and secondary, if you selected it) PCG deployment. This process can take 10-12 minutes.

10 Click Finish when PCG is successfully deployed.

What to do next

Onboard your workload VMs. See **Onboarding and Managing Workload VMs** in the *NSX-T Data Center Administration Guide* for the Day-N workflow.

Link to a Transit VPC or VNet

You can link one or more compute VPCs or VNets to a Transit VPC or VNet.

Prerequisites

- Verify that you have a Transit VPC or VNet with a PCG in the **Up** state.

- Verify that the VPC/VNet you want to link is connected to the Transit VPC or VNet through VPN or peering.

- Verify that the Transit VPC/VNet in the same region as the Compute VPC/VNet.

Note In route-based IPSec VPN configuration, you must specify the IP address for the virtual tunnel interface (VTI) port. This IP must be in a different subnet than workload VMs. This prevents workload VM inbound traffic from being directed to the VTI port, from which it will be dropped.

Note In the public cloud, a default limit exists for the number of inbound/outbound rules per security group and NSX Cloud creates default security groups. This affects how many Compute VPCs/VNets can be linked to a Transit VPC/VNet. Assuming 1 CIDR block per VPC/VNet, NSX Cloud supports 10 Compute VPCs/VNets per Transit VPC/VNet. If you have more than 1 CIDR in any Compute VPC/VNet, the number of supported Compute VPCs/VNets per Transit VPC/VNet reduces. You can adjust the default limits by reaching out to your public cloud provider.

Procedure

1 Log in to CSM using an account with the Enterprise Administrator role.

2 Click **Clouds > AWS / Azure > `<public cloud_account_name>`** and go to the **VPCs / VNets** tab.

3 In the **VPCs** or **VNets** tab, select a region name where you are hosting one or more compute VPCs or VNets.

4 Select a compute VPC or VNet configured for NSX Cloud.

5 Click **LINK TO TRANSIT VPC** or **LINK TO TRANSIT VNET**

6 Complete the options in the **Link Transit VPC or VNet** window:

Option	Description
Transit VPC or VNet	Select a Transit VPC or VNet from the dropdown menu. The Transit VPC or VNet you select must be already linked with this VPC by way of VPN or peering.
	Note If connecting to Transit VNet, you must have a DNS forwarder configured in that VNet. See Microsoft Azure documentation for more information.
Default Quarantine Policy	Leave this in the default **disabled** mode when you first deploy PCG. You can change this value after onboarding VMs. See **Manage Quarantine Policy** in the *NSX-T Data Center Administration Guide* for details.

What to do next

Onboard your workload VMs. See **Onboarding and Managing Workload VMs** in the *NSX-T Data Center Administration Guide* for the Day-N workflow.

Auto-Created Logical Entities and Cloud-native Security Groups

The deployment of PCG in a Transit VPC/VNet and linking a compute VPC/VNet to it triggers necessary configurations in NSX-T Data Center and the public cloud.

Auto-created NSX-T Logical Entities

In the NSX-T Data Center, a set of logical entities are created.

Important Do not delete any of these auto-created entities.

System Entities

You can see the following entities under **System**:

Table 11-3. Auto-Created System Entities

Logical System Entity	How many are created?	Nomenclature	Scope
Transport Zones	Two Transport Zones are created for each Transit VPC/VNet	■ TZ-\<VPC/VNet-ID>-OVERLAY ■ TZ-\<VPC/VNet-ID>-VLAN	Scope: Global
Edge Transport Nodes	One Edge Transport Node is created for each deployed PCG, two if deployed in high availability mode.	■ PublicCloudGateway TN-\<VPC/VNET-ID> ■ PublicCloudGateway TN-\<VPC/VNET-ID>-preferred	Scope: Global
Edge Cluster	One Edge Cluster is created per deployed PCG, whether one or in a high availability pair.	PCG-cluster-\<VPC/VNet-ID>	Scope: Global

Inventory Entities

The following entities are created under **Inventory**:

Table 11-4. Auto-Created Inventory Entities

Logical Inventory Entity	How many are created?	Nomenclature	Scope
Domain **Note** The **Domain** object is an experimental feature in NSX-T Data Center 2.4 and the auto-created Domains are visible in the user interface. However, Domains are no longer visible in the NSX-T Data Center 2.4.1 user interface.	One per Transit VPC/VNet	cloud-<Transit VPC/VNet-ID>	Scope: shared across all PCGs.
Groups	Two Groups under the **default** Domain **Note** In NSX-T Data Center you can see the default Domain. However, in NSX-T Data Center 2.4.1, the Domain object is not visible.	■ cloud-default-route ■ cloud-metadata services	Scope: Shared across all PCGs
Groups	One Group created at Transit VPC/VNet level as a parent group for individual segments created at the Compute VPC/VNet level.	cloud-<Transit VPC/VNet ID>-all-segments	Scope: shared across all Compute VPCs/VNets
Groups	Two Groups: ■ Network CIDR Group for all CIDRs of the Compute VPC/VNet ■ Local Segment Group for all managed segments within the Compute VPC/VNet	■ cloud-<Compute VPC/VNet ID>-cidr ■ cloud-<Compute VPC/VNet ID>-local-segments	Scope: shared across all Compute VPC/VNets

Security Entities

Table 11-5. Auto-Created Security Entities

Logical Security Entity	How many are created?	Nomenclature	Scope
Distributed Firewall (East-West)	Two per Transit VPC/VNet: ■ Stateless ■ Stateful	■ cloud-stateless-\<VPC/VNet ID> ■ cloud-stateful-\<VPC/VNet ID>	■ Stateful rule to allow traffic within local managed segments ■ Stateful rule to reject traffic from unmanaged VMs
Gateway Firewall (North-South)	One per Transit VPC/VNet	cloud-\<Transit VPC/VNet ID>	

Networking Entities

The following entities are created at different stages of onboarding:

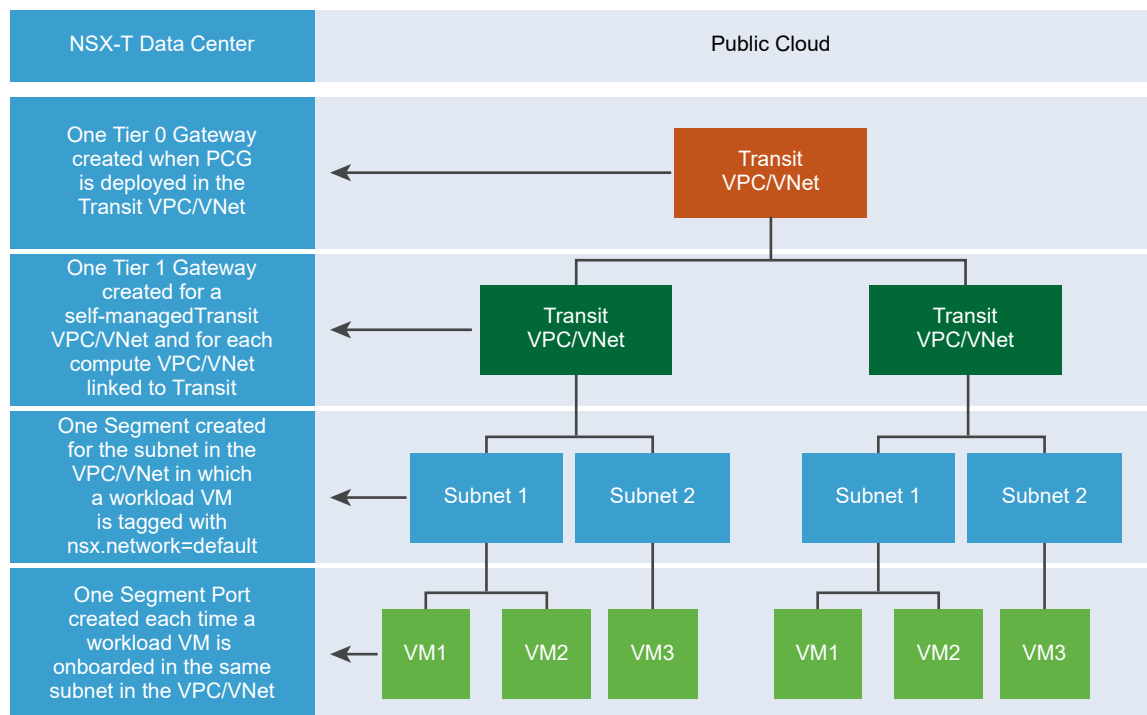Figure 11-3. Auto-created NSX-T Data Center Networking Entities After PCG is Deployed

Table 11-6. Auto-Created Networking Entities

Onboarding Task	Logical Entities Created in NSX-T Data Center
PCG deployed on Transit VPC/VNet	■ Tier-0 Gateway ■ Infra Segment (Default VLAN switch) ■ Tier-1 router
Compute VPC or VNet linked to the Transit VPC/VNet	■ Tier-1 router

Table 11-6. Auto-Created Networking Entities (continued)

Onboarding Task	Logical Entities Created in NSX-T Data Center
A workload VM with the NSX agent installed on it is tagged with the "nsx.network:default" key:value in a subnet of a compute or self-managed VPC/VNet	■ A Segment is created for this specific subnet of the compute or self-managed VPC or VNet ■ Hybrid ports are created for each tagged workload VM that has the NSX agent installed on it
More workload VMs are tagged in the same subnet of the Compute or self-managed VPC/VNet	■ Hybrid ports are created for each tagged workload VM that has the NSX agent installed on it

Forwarding Policies

The following three forwarding rules are set up for a Compute VPC/VNet, including Self-managed Transit VPC/VNet:

■ Access any CIDR of the same Compute VPC over the public cloud's network (underlay)

■ Route traffic pertaining to public cloud metadata services over the public cloud's network (underlay)

■ Route everything not in the Compute VPC/VNet's CIDR block, or a known service, through the NSX-T Data Center network (overlay)

Auto-created cloud-native SGs

In your public clouds, cloud-native security groups are created.

Public Cloud Configurations

In AWS:

■ In the AWS VPC, a new Type A Record Set gets added with the name `nsx-gw.vmware.local` into a private hosted zone in Amazon Route 53. The IP address mapped to this record matches the Management IP address of the PCG which is assigned by AWS using DHCP and will differ for each VPC. This DNS entry in the private hosted zone in Amazon Route 53 is used by NSX Cloud to resolve the PCG's IP address.

> **Note** When you use custom DNS domain names defined in a private hosted zone in Amazon Route 53, the **DNS Resolution** and **DNS Hostnames** attributes must be set to **Yes** for your VPC settings in AWS.

■ A secondary IP for the uplink interface for PCG is created. An AWS Elastic IP is associated with this secondary IP address. This configuration is for SNAT.

In AWS and Microsoft Azure:

The **gw** security groups are applied to the respective PCG interfaces.

Table 11-7. Public Cloud Security Groups created by NSX Cloud for PCG interfaces

Security Group name	Available in Microsoft Azure?	Available in AWS?	Full Name
gw-mgmt-sg	Yes	Yes	Gateway Management Security Group
gw-uplink-sg	Yes	Yes	Gateway Uplink Security Group
gw-vtep-sg	Yes	Yes	Gateway Downlink Security Group

Table 11-8. Public Cloud Security Groups created by NSX Cloud for Workload VMs

| Security Group name	Available in Microsoft Azure?	Available in AWS?	Descriptiom
quarantine	Yes	No	Quarantine security group for Microsoft Azure
default	No	Yes	Quarantine security group for AWS
vm-underlay-sg	Yes	Yes	VM Non-Overlay security group
vm-override-sg	Yes	Yes	VM Override Security Group
vm-overlay-sg	Yes	Yes	VM Overlay security group (this is not used in the current release)
vm-outbound-bypass-sg	Yes	Yes	VM Outbound Bypass Security Group (this is not used in the current release)
vm-inbound-bypass-sg	Yes	Yes	VM Inbound Bypass Security Group (this is not used in the current release)

Undeploy PCG

Refer to this flowchart for the steps involved in undeploying PCG.

Before undeploying PCG, you must do the following:

- Make sure that no workload VMs in the VPC or VNet are NSX-managed.

- Disable the Quarantine Policy.

- Delete all user-created logical entities associated with the PCG.

Figure 11-4. Undeploying PCG
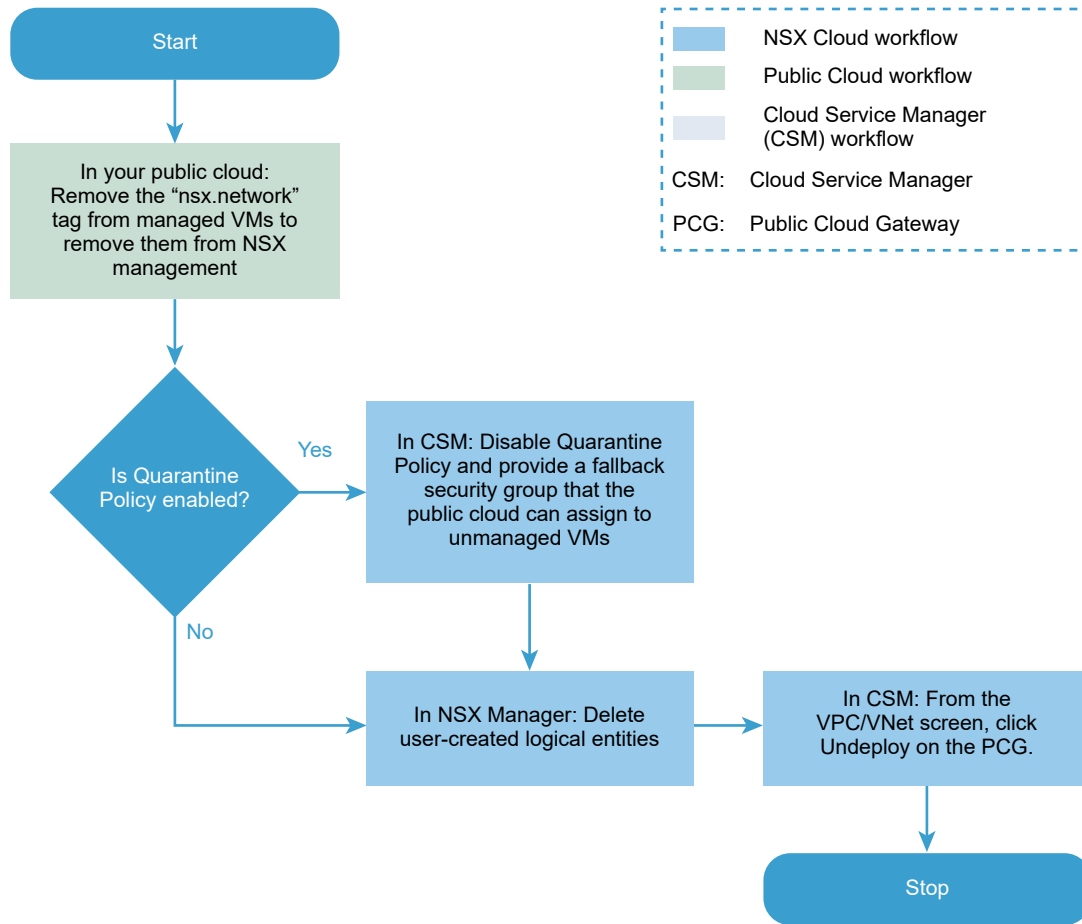
Procedure

1 Untag VMs in the Public Cloud

Before you can undeploy PCG, all VMs must be unmanaged.

2 Disable Quarantine Policy, if Enabled

If previously enabled, Quarantine Policy must be disabled to undeploy PCG.

3 Delete User-created Logical Entities

All user-created logical entities associated with the PCG must be deleted.

4 **Undeploy** from CSM

To undeploy PCG after completing the prerequisites, click Undeploy Gateway from **Clouds > <Public_Cloud> > <VNet/VPC>** in CSM.

Untag VMs in the Public Cloud

Before you can undeploy PCG, all VMs must be unmanaged.

Go to the VPC or VNet in your public cloud and remove the `nsx.network` tag from the managed VMs.

# Disable Quarantine Policy, if Enabled

If previously enabled, Quarantine Policy must be disabled to undeploy PCG.

With Quarantine Policy enabled, your VMs are assigned security groups defined by NSX Cloud. When you undeploy PCG, you need to disable Quarantine Policy and specify a fallback security group that the VMs can be assigned to when they are removed from the NSX Cloud security groups.

Note The fallback security group must be an existing user-defined security group in your public cloud. You cannot use any of the NSX Cloud security groups as a fallback security group. See Auto-Created Logical Entities and Cloud-native Security Groups for a list of NSX Cloud security groups.

Disable Quarantine Policy for the VPC or VNet from which you are undeploying PCG:

- Go to the VPC or VNet in CSM.

- From **Actions > Edit Configurations** >, turn off the setting for **Default Quarantine** .

- Enter a value for a fallback security group that VMs will be assigned.
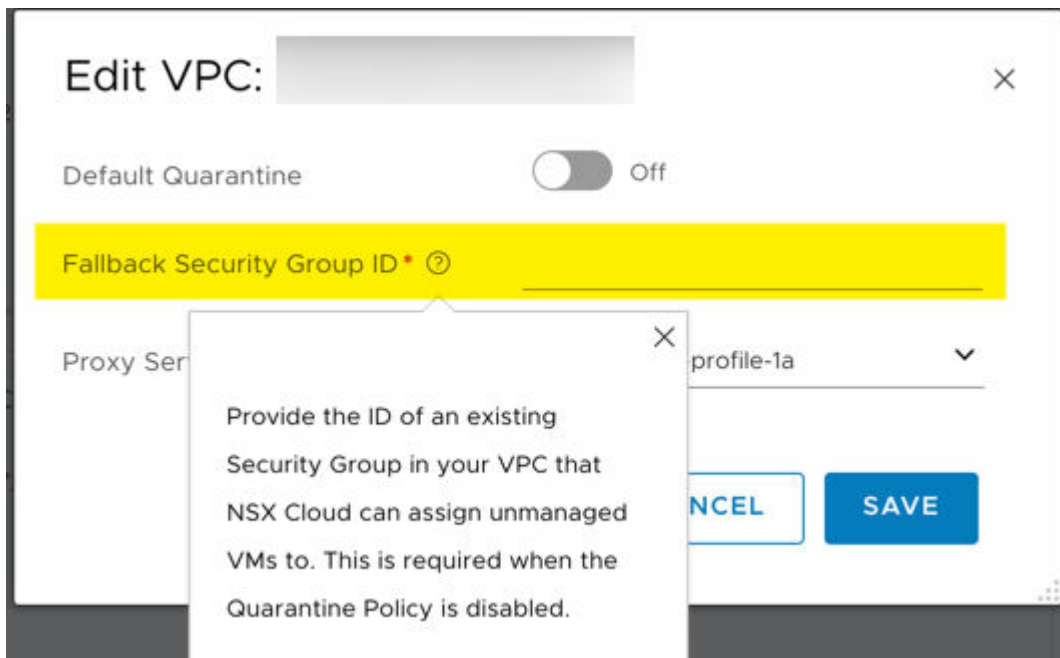
- All VMs that are unmanaged or quarantined in this VPC or VNet will get the fallback security group assigned to them.

- If all VMs are unmanaged, they get assigned to the fallback security group.

- If there are managed VMs while disabling Quarantine Policy, they retain their NSX Cloud-assigned security groups. The first time you remove the `nsx.network` tag from such VMs to take them out from NSX management, they are also assigned the fallback security group.

Note See **Managing Quarantine Policy** in the *NSX-T Data Center Administration Guide* for instructions and more information on the effects of enabling and disabling the Quarantine Policy.

Delete User-created Logical Entities

All user-created logical entities associated with the PCG must be deleted.

Identify entities which are associated with the PCG and delete them.

Note Do not delete the auto-created logical entities. These are deleted automatically after you click **Undeploy Gateway** from CSM. See Auto-Created Logical Entities and Cloud-native Security Groups for the list of auto-created logical entities.

Undeploy from CSM

To undeploy PCG after completing the prerequisites, click Undeploy Gateway from **Clouds > <Public_Cloud> > <VNet/VPC>** in CSM.

1   Log in to CSM and go to your public cloud:

- If using AWS, go to **Clouds > AWS > VPCs**. Click on the VPC on which one or a pair of PCGs is deployed and running.

- If using Microsoft Azure, go to **Clouds > Azure > VNets**. Click on the VNet on which one or a pair of PCGs is deployed and running.

2   Click Undeploy Gateway.

The default entities created by NSX Cloud are removed automatically when a PCG is undeployed.