

NSX Installation Guide

Modified on 09 SEPT 2024
VMware NSX 4.1

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

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

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

Copyright © 2017-2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to <https://www.broadcom.com>. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contents

NSX Installation Guide 10

1 Overview of NSX 11

Key Concepts 12

NSX Manager 17

Configure the User Interface Settings 20

2 NSX Installation Workflows 22

NSX Workflow for vSphere 22

NSX Configuration Workflow for Bare Metal Server 23

3 Preparing for Installation 24

System Requirements 24

NSX Manager VM and Host Transport Node System Requirements 24

Supported Hypervisor Version 27

NSX Edge VM System Requirements 29

NSX Edge Bare Metal Requirements 31

Bare Metal Server System Requirements 35

Bare Metal Linux Container Requirements 37

Ports and Protocols 37

Configuring a vSphere Distributed Switch 37

Checklist Before Deploying Infrastructure 40

4 NSX Manager Installation Requirements 43

Configure NSX Manager for Access by DNS Server 47

Publish FQDN of the NSX Managers 48

Modifying the Default Admin Password Expiration 49

5 NSX Manager Cluster Requirements 50

Cluster Requirements for an Individual Site 51

Cluster Requirements for Multiple Sites 52

6 Installing NSX Manager Cluster on vSphere 55

Install NSX Manager and Available Appliances 55

Log In to the Newly Created NSX Manager 61

Add a Compute Manager 61

Deploy NSX Manager Nodes to Form a Cluster from the UI 66

Install NSX Manager on ESXi Using the Command-Line OVF Tool 72

Form an NSX Manager Cluster Using the CLI	79
Verify NSX Manager Clustering	81
Configure an Appliance to Display the GRUB Menu at Boot Time	81
Configure a Virtual IP Address for a Cluster	83
Configuring an External Load Balancer	85
Verify Appliance Proxy Hub on all NSX Manager Nodes are Connected	87
7 Installing and Configuring NSX Embedded using VMware vCenter Plugin	89
Install NSX Manager from vSphere Client	89
Install Additional NSX Manager Nodes to Form a Cluster from VMware vCenter Plugin	93
Configure NSX for Virtual Networking from vSphere Client	97
Configuring NSX-T for Security from vSphere Client	103
Prepare Clusters for NSX Security	103
Create Groups	104
Define and Publish Communication Strategies for Groups	107
Viewing NSX Alarms in vSphere Web Client UI	110
8 Transport Zones and Profiles	111
Create Transport Zones	111
Create an IP Pool for Tunnel Endpoint IP Addresses	114
Enhanced Data Path	117
Enhanced Datapath Supported Features	120
Guidance to Set Maximum Transmission Unit	121
Configuring Profiles	123
Create an Uplink Profile	123
Configure Named Teaming Policy	128
Add and attach NSX Edge Bridge Profile to a Segment	130
Add an NSX Edge Cluster Profile	134
Prepare a vSphere Distributed Switch for NSX	134
Add a Transport Node Profile	135
9 Host Transport Nodes	141
Preparing ESXi Hosts as Transport Nodes	141
Transport Node Profiles	141
Sub-TNPs and Sub-clusters	143
Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP	145
Change Sub-cluster	150
Detach Cluster TNP	151
Migrate VMkernels and Physical NICs to a vSphere Distributed Switch	151
Prepare ESXi Individual Hosts as Transport Nodes	153
Verify the Transport Node Status	159

Manual Installation of NSX Kernel Modules	161
Manually Install NSX Kernel Modules on ESXi Hypervisors	161
Preparing Physical Servers as NSX Transport Nodes	165
Install Third-Party Packages on a Linux Physical Server	165
Configure a Physical Server as a Transport Node from GUI	168
Ansible Server Configuration for Physical Server	176
Create Application Interface for Physical Server Workloads	176
Windows Physical Server Supported Teaming Topologies	177
Secure Workloads on Windows Server 2016/2019 Bare Metal Servers	184
Configure an ESXi Host Transport Node with Link Aggregation Group	186
Quick Start Wizard to Prepare ESXi Cluster hosts for Security-only or Networking and Security	187
Prepare ESXi cluster Hosts for Networking and Security	187
Install Distributed Security for vSphere Distributed Switch	189
Deploy a Fully Collapsed vSphere Cluster NSX on Hosts Running N-VDS Switches	190
Multiple NSX Managers Managing a Single VMware vCenter	201
Troubleshoot Multi-NSX Issues	205
Cannot Enable or Disable Multiple NSX on a Single VMware vCenter	205
Override NSX Ownership Constraints	205
Preparing Standalone Host Cluster Results in Failure	207
Traffic Performance Issues Related to NSX Edge VMs in Multiple NSX Environment	208
Distributed Virtual Switch is Not Available for Selection in NSX	209
Managing Transport Nodes	211
Switch Visualization	211
NSX Maintenance Mode	211
Migrate ESXi VMkernel and Physical Adapters	212
View Bidirectional Forwarding Detection Status	213
IPv6 Unsupported Features	215
10 Installing NSX Edge	216
NSX Edge Installation Requirements	216
NSX Edge Networking Setup	219
NSX Edge Installation Methods	225
Create an NSX Edge Transport Node	226
Configure NSX Edge DPDK Interfaces	235
Manually Deploying NSX Edge Node Outside of NSX	238
Install an NSX Edge on ESXi Using the vSphere GUI	238
Install NSX Edge on ESXi Using the Command-Line OVF Tool	242
Install NSX Edge via ISO File as a Virtual Appliance	247
Install NSX Edge on Bare Metal	251
Join NSX Edge with the Management Plane	264
Edit NSX Edge Transport Node Configuration	266

- Create an NSX Edge Cluster 271
- Remove NSX Edge Nodes from an Edge Cluster 272
- Relocate and Remove an NSX Edge Node from an NSX Edge Cluster 274

11 NSX IPv6 Configuration 277

- IPv6 Support in the NSX Platform Infrastructure 277
- IPv6 Limitations Between Management Plane/Control Plane and Transport Nodes 279
- Supported Topologies for IPv6 279
- IPv6 Configuration Workflow for NSX Manager and Transport Node Communication 281
 - Configure a Brownfield NSX Manager Node with IPv6 282
 - Configure a Brownfield NSX Edge Transport Node with IPv6 283
 - Change the IP Address for an IPv6 NSX Edge Transport Node 285
- IPv6 Tunnel Endpoint Deployment for ESXi Host and NSX Edge on New Transport Zones 286
- IPv6 Tunnel Endpoint Configuration for ESXi Host and NSX Edge on Existing Transport Zones 287
- Troubleshooting for IPv6 Configuration 288

12 vSphere Lifecycle Manager with NSX 291

- Prepare an NSX Cluster with vSphere Lifecycle Manager 292
- Enable vSphere Lifecycle Manager on an NSX Cluster 295
- NSX on vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine 297
 - Configure NSX host transport node on DPU-based vSphere Lifecycle Manager-enabled cluster 299
 - Displaying DPU-related information on NSX Manager Interface 301
- NSX with vSphere Lifecycle Manager Scenarios 301
- vSphere Lifecycle Manager Failed to Prepare a Host for NSX Networking 303
- vSphere Lifecycle Manager Failed to Prepare NSX Cluster 304
- Delete a NSX Depot on vCenter Server 305

13 Host Profile integration with NSX 306

- Auto Deploy Stateless Cluster 306
 - High-Level Tasks to Auto Deploy Stateless Cluster 306
 - Prerequisites and Supported Versions 307
 - Create a Custom Image Profile for Stateless Hosts 308
 - Associate the Custom Image with the Reference and Target Hosts 309
 - Set Up Network Configuration on the Reference Host 310
 - Configure the Reference Host as a Transport Node in NSX 310
 - Extract and Verify the Host Profile 311
 - Verify the Host Profile Association with Stateless Cluster 313
 - Update Host Customization 313
 - Trigger Auto Deployment on Target Hosts 314
 - Troubleshoot Host Profile and Transport Node Profile 317

Stateful Servers	319
Supported NSX and ESXi versions	320
Prepare a Target Stateful Cluster	321
14 Getting Started with NSX Federation	323
NSX Federation Key Concepts	323
NSX Federation Requirements	324
Configuring the Global Manager and Local Managers	325
Install the Active and Standby Global Manager	326
Make the Global Manager Active and Add Standby Global Manager	327
Add a Location	328
Remove a Location	335
15 Install NSX Advanced Load Balancer Appliance Cluster	337
Troubleshooting NSX Advanced Load Balancer Controller Issues	341
NSX Advanced Load Balancer does not register with NSX Manager	341
The Second NSX Advanced Load Balancer Controller Remains in Queued State	341
NSX Advanced Load Balancer Controller Password Change Caused Cluster Failure	342
Unable to Delete NSX Advanced Load Balancer Controller	343
NSX Advanced Load Balancer Cluster HA Status is Compromised	343
Credential Mismatch After Changing NSX Advanced Load Balancer Controller Password	344
Deployment of NSX Advanced Load Balancer Controller Failed	345
Cluster Unstable After Two Controllers Are Down	345
16 Getting Started with NSX Cloud	347
NSX Cloud Architecture and Components	348
Overview of Deploying NSX Cloud	349
Deploy NSX On-Prem Components	349
Install CSM	349
Join CSM with NSX Manager	350
Specify CSM IPs for Access by PCG	351
(Optional) Configure Proxy Servers	351
(Optional) Set Up vIDM for Cloud Service Manager	352
Obtain Thumbprint of NSX Manager	353
Connect your Public Cloud with On-prem NSX	354
Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image	357
Deploy NSX Cloud Components in Microsoft Azure using Terraform scripts	359
Deploy NSX Cloud Components in Microsoft Azure without using Terraform scripts	364
Add your Public Cloud Account	368
Adding your Microsoft Azure Subscription	369

Adding your AWS Account	375
Managing Regions in CSM	379
NSX Public Cloud Gateway: Architecture and Modes of Deployment	380
Deploy PCG or Link to a PCG	384
Deploy PCG in a VNet	384
Deploy PCG in a VPC	387
Link to a Transit VPC or VNet	390
Auto-Configurations after PCG Deployment or Linking	391
Auto-created NSX Logical Entities	391
Auto-created Public Cloud Configurations	395
Integrate Horizon Cloud Service with NSX Cloud	397
Auto-Created Entities After Horizon Cloud Integration	400
(Optional) Install NSX Tools on your Workload VMs	403
Un-deploy NSX Cloud	403
Undeploy or Unlink PCG	404
17 Uninstalling NSX from a Host Transport Node	407
Uninstall NSX from a vSphere Cluster	407
Uninstall NSX from a Managed Host in a vSphere Cluster	410
Uninstall NSX from a Physical Host	412
Triggering Uninstallation from the vSphere Web Client	416
Uninstall NSX from a vSphere Lifecycle Manager cluster through NSX Manager	418
18 Troubleshooting Installation Issues	420
Troubleshooting Installation to check for Basic Infrastructure Services	420
Troubleshoot OVA Deployment and Appliance Bringup	422
Troubleshoot NSX Manager Cluster	423
Manually deployed NSX Manager failed to join the NSX Manager Cluster	423
NSX Manager Cluster status Degraded As Datastore-related Components Are Down	424
Manager and HTTPS Services Are Down Frequently Due to Incorrect NAT Configuration	425
NSX Manager is Slow To Load And Tasks Fail	425
NSX Manager UI is not loading even when NSX Clustering is up	426
NSX Manager cluster is DOWN or UNAVAILABLE if all nodes part of the the NSX Manager cluster is down or majority nodes are down	426
Troubelshoot NSX Appliance or NSX Clustering issues Using APIs	427
Troubleshooting Host Transport Nodes	427
Host Fails to Install as Transport Node	427
Accessing the NSX CLI Terminal	428
Transport Node Installation Failure Due To Pending Reboot	428
Unable to Reach Transport Node Due to Incorrect Credentials As Host is in Orphaned State	428

Transport node profile Fails to Prepare Transport Nodes Due to Stale Objects	429
Transport Node Creation is in Partial Success State Due to Logical Switch or Segment Full-sync Realization Error	430
Transport Node status is degraded when its interface is down	430
Transport Node status is Disconnected or Unknown	431
Transport Node is Down as Agent Service is Down	432
Transport Node Connectivity to Controller is Down	432
Unable to power on VMs or vMotion VMs on Transport Node	433
Transport Node Tunnels Down	434
Installation Fails Due to Insufficient Space in Bootbank on ESXi Host	435
NSX Agent on ESXi Transport Nodes Times Out Communicating with NSX Manager	435
Troubleshooting NSX Edge Nodes	436
Accessing the NSX Edge CLI Terminal	437
NSX Edge MPA Connectivity Down	437
NSX Edge Status DOWN or DEGRADED As BFD Tunnel(s) are Down	437
NSX Edge Router High Availability Status is Down	439
NSX Edge Node Status Down Due to PNIC Bond Status Down	440
NSX Edge Transport Node Connectivity to Controller is Down	440
NSX Edge node status is Down As Controller Is Unavailable	441
Transport Node Failed as Named Teaming Defined with No Active Uplink	442
NSX Edge Transport Node goes into NSX Maintenance Mode On HA Failover	442
VMotion of NSX Edge VM Fails Due To ESXi Running Out Of Resources	443
State Not Consistent of NSX BFD BGP or HA Functionality	444
NSX Edge Node Status Down and BFD Tunnel Data is Missing	444
NSX Edge status DOWN or DEGRADED due to BFD tunnels between Edge and ESXi down	444

NSX Installation Guide

The *NSX Installation Guide* describes how to install the VMware NSX® 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. 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 <https://www.vmware.com/topics/glossary>.

Related Documentation

You can find the VMware NSX® Intelligence™ documentation at <https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html>. The NSX Intelligence 1.0 content was initially included and released with the NSX 2.5 documentation set.

Overview of NSX

1

In the same way that server virtualization programmatically creates and manages virtual machines, NSX network virtualization programmatically creates and manages 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 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 RESTful API services, management plane (MP) cluster daemons and central control plane (CCP) cluster daemons.
- Transport nodes host local control plane (LCP) daemons and forwarding engines implementing the NSX data plane.

NSX supports a NSX Manager cluster with three nodes, which merges policy manager, management, and central control services on a cluster of nodes. Starting NSX release version 2.4, NSX controller elements reside in NSX Manager appliance. The convergence of management and control plane nodes reduces the number of virtual appliances that must be deployed and managed by the NSX administrator. NSX Manager clustering runs on top of a distributed data platform called Corfu and provides high availability of the user interface and the API .

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.
- A large appliance for customers who deploy to a large-scale environment.

See [NSX Manager VM and Host Transport Node System Requirements and Configuration maximums](#) tool.

Read the following topics next:

- [Key Concepts](#)
- [NSX Manager](#)

Key Concepts

The common NSX 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. NSX supports VMware vCenter as a compute manager.

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 (transport nodes). NSX control plane is split into two components - the Central Control Plane (CCP) and the Local Control Plane (LCP). The CCP is implemented on NSX Manager Cluster, while the LCP on all of the NSX transport nodes.

Corfu services

Run on each NSX Manager node to create the highly available distributed datastore Corfu.

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. Data plane is implemented by NSX transport nodes.

Data Processing Unit (DPU)

A DPU device is a SmartNIC device, or a high-performance network interface card, with added embedded CPU cores, memory, and a hypervisor running on the device independently from the ESXi hypervisor installed on the server.

Note We will refer to SmartNIC as DPU across our user guides.

External Network

A physical network or VLAN not managed by NSX. You can link your logical network or overlay network to an external network through a Tier-0 Gateway, Tier-1 Gateway or L2 bridge.

External Interface

Tier-0 Gateway Interface connecting to the physical infrastructure or router. Static routing and BGP are supported on this interface. This interface was referred to as uplink interface in previous releases.

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 entering the VM is called ingress traffic.

Gateway

NSX routing entity that provides connectivity between different L2 networks. Configuring a gateway through NSX Manager instantiates a gateway (Tier-0 or Tier-1) on transport nodes and provides optimized distributed routing as well as centralized routing and services like NAT, Load balancer, DHCP and other supported services on each hypervisor.

Gateway Port

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

Segment Port

Logical switch attachment point to establish a connection to a virtual machine network interface, container, physical appliances or a gateway interface. The segment 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.

NSX Edge Cluster

Is a collection of NSX Edge node appliances that have the same settings and provide high availability if one of the NSX Edge node fails.

NSX Edge Node

Edge nodes are service appliances (Bare Metal or VM form factor) with pools of capacity, dedicated to running network and security services that cannot be distributed to the hypervisors.

NSX Managed Virtual Distributed Switch (N-VDS, host-switch)

The NSX managed virtual distributed switch forwards traffic between logical and physical ports of the device. On ESXi hosts, the N-VDS implementation is derived from VMware vSphere® Distributed Switch™ (VDS) and it shows up as an opaque network in vCenter. With

any other kind of transport node (KVM hypervisors, Edges, Bare Metal servers, cloud VMs and so on) the N-VDS implementation is derived from the Open vSwitch (OVS).

Note VMware has removed support of the NSX N-VDS virtual switch on ESXi hosts starting release 4.0.0.1 because it is recommended to deploy NSX on top of vCenter VDS. N-VDS will remain the supported virtual switch on NSX Edge nodes, native public cloud NSX agents, and Bare Metal workloads.

An N-VDS has two modes: Standard and Enhanced Datapath. The Enhanced Data Path N-VDS is optimized for the Network Function Virtualization, where the workloads perform networking functions under demanding latency and packet rate requirements.

vSphere Distributed Switch (VDS)

Starting with NSX 3.0, NSX can run directly on top of a vSphere Distributed Switch version 7 or later. It is recommended that you use the VDS switch for deployment of NSX on ESXi hosts. You can create Overlay or VLAN backed segments on VDS switches, similar to the N-VDS, VDS switches that can be configured in Standard or Enhanced Datapath mode.

vSphere Distributed Services Engine

Sphere 8.0 introduces VMware vSphere Distributed Services Engine, which leverages data processing units (DPUs) as the new hardware technology to overcome the limits of core CPU performance while delivering zero-trust security and simplified operations to vSphere environments. With NSX 4.0.1.1, vSphere Distributed Services Engine provides the ability to offload some of the network operations from your server CPU to a DPU.

NSX Manager

Node that hosts the API services, the management plane, the control plane and the agent services. It is accessible through CLI, Web UI, or API. NSX Manager is an appliance included in the NSX installation package. You can deploy the appliance in the role of `NSX Manager` or `nsx-cloud-service-manager`. Currently, the appliance only supports one role at a time.

NSX Manager Cluster

A cluster of NSX Manager virtual machine appliances providing high availability of the user interface and the API.

Open vSwitch (OVS)

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

Opaque Network

An opaque network is a network created and managed by a separate entity outside of vSphere. For example, logical networks that are created and managed by N-VDS switch running on NSX appear in vCenter Server as opaque networks of the type `nsx.LogicalSwitch`. You can choose an opaque network as the backing for a VM network adapter. To manage an opaque network, use the management tools associated with the opaque network, such as NSX Manager or the NSX API management tools.

Overlay Logical Network

Logical network implemented using GENEVE encapsulation protocol as mentioned in <https://www.rfc-editor.org/rfc/rfc8926.txt>. 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

Previously known as logical switch. It is an 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.

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.

Service Interface

Tier-0 Interface connecting VLAN segments to provide connectivity and services to VLAN backed physical or virtual workloads. A service interface can also be connected to overlay segments for Tier-1 standalone load balancer use cases. Starting with NSX release version 3.0, service interface supports static and dynamic routing.

Tier-0 Gateway

A Tier-0 Gateway provides north-south connectivity and connects to the physical routers. It can be configured as an active-active or active-standby cluster. The Tier-0 Gateway runs BGP and peers with physical routers.

Tier-0 Gateways consists of two components:

- distributed routing component (DR) that runs on all transport nodes. DR of the Tier-0 gateway is instantiated on the hypervisors and Edge transport nodes upon its creation.

- centralized services routing component (SR) runs on edge cluster nodes. SR gets instantiated on the edge nodes upon association of gateway with edge cluster and creation of external interfaces.

Tier-1 Gateway

A Tier-1 Gateway connects to one Tier-0 Gateway for northbound connectivity of the subnetworks attached to it (multi-tier routing model). It connects to one or more overlay networks for southbound connectivity to its subnetworks. A Tier-1 Gateway can be configured as an active-standby cluster. Like Tier-0 gateway, when a Tier-1 gateway is created, a distributed component (DR) of the Tier-1 gateway is instantiated on the hypervisors and Edge transport nodes but the service component (SR) will only be created if gateways is associated with edge cluster and external interfaces created.

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. It also has been registered with the NSX management plane and has NSX modules installed. For a hypervisor host or NSX Edge to be part of the NSX overlay, it must be added to the NSX transport zone.

Transport Node

A fabric node is prepared as a transport node so that it becomes capable of participating in an NSX overlay or NSX VLAN networking. For an ESXi host, you must configure a VDS switch.

Uplink Profile (host-switch-profile)

Defines policies for the links from transport nodes to NSX segments or from NSX Edge nodes to top-of-rack switches. The settings defined by uplink profiles might include teaming policies, 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 NSX segment. The vNIC can be attached to a logical port. You can identify a vNIC based on its Unique ID (UUID).

Tunnel Endpoint (TEP)

Each transport node has a Tunnel Endpoint (TEP) responsible for encapsulating the overlay VM traffic inside a VLAN header and routing the packet to a destination TEP for further processing. TEPs are the source and destination IP addresses used in the external IP header to identify the ESXi hosts that originate and end the NSX encapsulation of overlay frames. Traffic can be routed to another TEP on a different host or the NSX Edge gateway to access the physical network. TEPs create a GENEVE tunnel between the source and destination endpoints.

NSX Manager

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

The NSX Manager interface provides two modes for configuring resources:


- Policy mode
- Manager mode

The Policy Mode is the default and recommended mode. The Manager mode will be deprecated over time, as all features or functionality is transitioned to Policy.

Note: For details on promoting Manager objects to Policy objects, see the Promote Manager Objects to Policy Objects topic in the *NSX Administration Guide*.

Accessing Policy Mode and Manager Mode

If present, you can use the **Policy** and **Manager** buttons to switch between the Policy and Manager modes. Switching modes controls which menus items are available to you.



- By default, if your environment contains only objects created through Policy mode, your user interface is in Policy mode and you do not see the **Policy** and **Manager** buttons.
- By default, if your environment contains any objects created through Manager mode, you see the **Policy** and **Manager** buttons in the top-right corner.

These defaults can be changed by modifying the user interface settings. See [Configure the User Interface Settings](#) for more information.

The same **System** tab is used in the Policy and Manager interfaces. If you modify Edge nodes, Edge clusters, or transport zones, it can take up to 5 minutes for those changes to be visible in Policy mode. You can synchronize immediately using `POST /policy/api/v1/infra/sites/default/enforcement-points/default?action=reload`.

When to Use Policy Mode or Manager Mode

VMware recommends to use NSX Policy UI as all the new features are implemented only on Policy UI/API.

Be consistent about which mode you use. There are a few reasons to use one mode over the other.

- If you are deploying a new NSX environment, using **Policy** mode to create and manage your environment is the best choice in most situations.
 - Some features are not available in Policy mode. If you need these features, use **Manager** mode for all configurations.
- If you plan to use NSX Federation, use **Policy** mode to create all objects. Global Manager supports only Policy mode.
- If you are upgrading from an earlier version of NSX and your configurations were created using the Advanced Networking & Security tab, use **Manager** mode.

The menu items and configurations that were found under the Advanced Networking & Security tab are available in NSX 3.0 in **Manager** mode.

Important If you decide to use Policy mode, use it to create all objects. Do not use Manager mode to create objects.

Similarly, if you need to use Manager mode, use it to create all objects. Do not use Policy mode to create objects.

Table 1-1. When to Use Policy Mode or Manager Mode

Policy Mode	Manager Mode
Most new deployments should use Policy mode. NSX Federation supports only Policy mode. If you want to use NSX Federation, or might use it in future, use Policy mode.	Deployments which were created using the advanced interface, for example, upgrades from versions before Policy mode was available.
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 Policy mode only: <ul style="list-style-type: none"> ■ DNS Services and DNS Zones ■ VPN ■ Forwarding policies for NSX Cloud 	Forwarding up timer
Security features available in Policy mode only: <ul style="list-style-type: none"> ■ Endpoint Protection ■ Network Introspection (East-West Service Insertion) ■ Context Profiles <ul style="list-style-type: none"> ■ L7 applications ■ FQDN ■ New Distributed Firewall and Gateway Firewall Layout <ul style="list-style-type: none"> ■ Categories ■ Auto service rules ■ Drafts 	Security features available in Manager mode only: <ul style="list-style-type: none"> ■ Bridge Firewall

Names for Objects Created in Policy Mode and Manager Mode

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

Table 1-2. Object Names

Objects Created Using Policy Mode	Objects Created Using Manager Mode
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
Gateway firewall	Edge firewall

Policy and Manager APIs

The NSX Manager provides two APIs: Policy and Manager.

- The Policy API contains URIs that begin with `/policy/api`.
- The Manager API contains URIs that begin with `/api`.

Policy APIs support partial patching of objects. This feature needs to be explicitly enabled. If enabled, you can provide the partial payload for updating the existing object using PATCH APIs.

To enable the feature, use the Partial Patch Config API

```
PATCH /policy/api/v1/system-config/nsx-partial-patch-config
```

```
{ "enable_partial_patch": "true" }
```

The default is 'false'.

For more information about using the Policy API, see the [NSX Policy API: Getting Started Guide](#).

Security

NSX Manager has the following security features:

- NSX Manager has a built-in user account called **admin**, which has access rights to all resources, but does not have rights to the operating system to install software. NSX upgrade files are the only files allowed for installation.
- NSX Manager supports session timeout and automatic user logout. NSX Manager does not support session lock. Initiating a session lock can be a function of the workstation operating system being used to access NSX Manager. Upon session termination or user logout, users are redirected to the login page.

- Authentication mechanisms implemented on NSX follow security best practices and are resistant to replay attacks. The secure practices are deployed systematically. For example, sessions IDs and tokens on NSX Manager for each session are unique and expire after the user logs out or after a period of inactivity. Also, every session has a time record and the session communications are encrypted to prevent session hijacking.

You can view and change the session timeout value with the following CLI commands:

- The command `get service http` displays a list of values including session timeout.
- To change the session timeout value, run the following commands:


```
set service http session-timeout <timeout-value-in-seconds>
restart service ui-service
```

Configure the User Interface Settings

You can configure how your users view the NSX user interface. These settings are valid for NSX Manager, and in NSX Federation for Global Managers and Local Managers.

Prior to NSX 3.2, users could access two possible modes in the user interface: Policy and Manager. You can control which mode is default, and whether users can switch between them using the user interface mode buttons. The Policy mode is the default. New users of release 4.0 will not see the **Manager** button.

If present, you can use the **Policy** and **Manager** buttons to switch between the Policy and Manager modes. Switching modes controls which menus items are available to you.



- By default, if your environment contains only objects created through Policy mode, your user interface is in Policy mode and you do not see the **Policy** and **Manager** buttons.
- By default, if your environment contains any objects created through Manager mode, you see the **Policy** and **Manager** buttons in the top-right corner.

You can use the User Interface settings to modify these defaults.

See [NSX Manager](#) for more information about the modes.

Procedure

- 1 With admin privileges, log in to NSX Manager.
- 2 Navigate to **System > General System Settings**.
- 3 Select **User Interface** and click **Edit** in the User Interface Mode Toggle pane.

4 Modify the user interface settings: **Toggle Visibility** and **Default Mode**.

Toggle Visibility	Description
Visible to All Users	If Manager mode objects are present, the mode buttons are visible to all users.
Visible to Users with the Enterprise Admin Role	If Manager mode objects are present, the mode buttons are visible to users with the Enterprise Admin role.
Hidden from All Users	Even if Manager mode objects are present, the mode buttons are hidden from all users. This displays Policy mode UI only, even if Manager mode objects are present.
Default Mode	Can be set to Policy or Manager, if available.

NSX Installation Workflows

2

You can install NSX on vSphere hosts. You can also configure a bare metal server to use NSX.

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

Read the following topics next:

- [NSX Workflow for vSphere](#)
- [NSX Configuration Workflow for Bare Metal Server](#)

NSX 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 [Chapter 4 NSX Manager Installation Requirements](#).
- 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 manager. 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 the UI](#).
- 7 Review the NSX Edge installation requirements. See [NSX Edge Installation Requirements](#).
- 8 Install NSX Edges. See [Install an NSX Edge on ESXi Using the 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 [Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP](#).
- 12 Create edge transport nodes. See [Create an NSX Edge Transport Node](#).
- 13 Create an edge cluster. See [Create an NSX Edge Cluster](#).

- 14 Create a Tier-0 gateway (if deploying single-tier or multi-tier routing topology in NSX). See *NSX Administration Guide*.

NSX Configuration Workflow for Bare Metal Server

Use the checklist to track your progress when configuring bare metal server to use NSX.

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 and Available Appliances](#).
- 4 Configure third-party packages on the bare metal server. See [Install Third-Party Packages on a Linux Physical Server](#).
- 5 Create host transport nodes.

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 Physical Server Workloads](#).

Preparing for Installation

3

Before installing NSX, make sure your environment is prepared.

Read the following topics next:

- [System Requirements](#)
- [Ports and Protocols](#)
- [Configuring a vSphere Distributed Switch](#)
- [Checklist Before Deploying Infrastructure](#)

System Requirements

Before you install NSX, your environment must meet specific hardware and resource requirements.

Before you configure Gateway Firewall features, make sure that the NSX Edge form factor supports the features. See *Supported Gateway Firewall Features on NSX Edge* topic in the *NSX Administration Guide*.

NSX Manager VM and Host Transport Node System Requirements

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

Hypervisor Host Network Requirements

The NIC card used must be compatible with the ESXi version that is running NSX. For supported NIC card, see the [VMware Compatibility Guide](#). To find the supported driver for the NIC card vendor, see [Firmware Versions of Supported Drivers](#).

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

- Under **I/O Devices**, 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 [Broadcom Support](#) 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
Cisco VIC 1400 series	nenic_ens

NSX Manager VM Resource Requirements

Thin virtual disk size is 3.8 GB and thick virtual disk size is 300 GB.

Appliance Size	Memory	vCPU	Shares	Reservations	Disk Space	VM Hardware Version
NSX Manager Extra Small (NSX 3.0 onwards)	8 GB	2	81920, Normal	8192 MB	300 GB	10 or later
NSX Manager Small VM (NSX 2.5.1 onwards)	16 GB	4	163840, Normal	16384 MB	300 GB	10 or later

Appliance Size	Memory	vCPU	Shares	Reservations	Disk Space	VM Hardware Version
NSX Manager Medium VM	24 GB	6	24576 0, Normal	24576 MB	300 GB	10 or later
NSX Manager Large VM	48 GB	12	49152 0, Normal	49152 MB	300 GB	10 or later

Note 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 provide by the NSX Controller appliance.

- You can use the Extra Small VM resource size only for the Cloud Service Manager appliance (CSM). Deploy CSM in the Extra Small VM size or higher, as required. See [Overview of Deploying NSX Cloud](#) for more information.
- 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 the autoselected appliance size during deployment and is suitable for typical production environments. An NSX management cluster formed using this appliance size can support up to 128 hypervisors. Starting with NSX 3.1, a single NSX Manager cluster is supported.
- The NSX Manager Large VM appliance size is suitable for large-scale deployments. An NSX management cluster formed using this appliance size can support more than 128 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 from the product list.

Language Support

NSX Manager has been localized into multiple languages: English, German, French, Japanese, Simplified Chinese, Korean, Traditional Chinese, and Spanish.

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 80	Yes	Yes	Yes
Mozilla Firefox 72	Yes	Yes	Yes

Browser	Windows 10	Mac OS X 10.13, 10.14	Ubuntu 18.04
Microsoft Edge 80	Yes		
Apple Safari 13		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 utilizes 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.

Important NSX Appliance VMs backed by vSAN clusters may see intermittent disk write latency spikes of 10+ms. This is expected due to the way VSAN handles data (burst of incoming IOs resulting in queuing of data and delay). As long as the average disk access latency continues to be less than 10ms, intermittent latency spike should not have impact on NSX Appliance VMs.

Storage Requirements

- NSX appliance VMs that are backed by VSAN clusters may see intermittent disk write latency spikes of 10+ms. This is expected due to the way VSAN handles data (burst of incoming IOs resulting in queuing of data and delay). As long as the average disk access latency continues to be less than 10ms, intermittent latency spike should not have impact on NSX Appliance VMs.
- It is recommended that NSX Managers be placed on shared storage.
- Storage must 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.

Refer to the product documentation of the storage technology you are using to know how to optimally design a highly available storage solution.

Supported Hypervisor Version

Supported Hypervisor for Host Transport Nodes

To find interoperability between NSX and vCenter Server/ESXi host, do the following:

- 1 In the **Solution 1** section, from the **Select a Solution** drop-down menu, select **VMware NSX** and from the **Select a Version** drop-down menu, select a version. For example, 4.x.x.
- 2 In the **Platform/Solution 1** section, from the **Select a Solution** drop-down menu, select **VMware vCenter Server and/or VMware vSphere Hypervisor (ESXi)** and select a version. For example, 7.x.
- 3 Click **Check Interoperability**.

Hypervisor	Version	CPU Cores	Memory
vSphere	Supported vSphere version	4	16 GB

Note To avoid memory errors on a hypervisor host running vSphere ESXi version 7.x.x, ensure that at least 16 GB is available before deploying NSX Manager.

Table 3-1. Supported Hosts for NSX Managers

Support Description	Hypervisor
ESXi	For supported hosts, see the VMware Product Interoperability Matrices .

For ESXi hosts, NSX supports the Host Profiles and Auto Deploy features on vSphere 6.7 EP6 or higher. See *Understanding vSphere Auto Deploy* in the *VMware ESXi Installation and Setup* documentation for more information.

Caution On RHEL and Ubuntu, the `yum update` command might update the kernel version, which must not be greater than 4.19.x, and break the compatibility with NSX. Disable the automatic kernel update when you run `yum update`. Also, after running `yum install`, verify that NSX supports the kernel version.

NSX Edge VM System Requirements

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

Note The following conditions apply to the hosts for the NSX Edge nodes:

- NSX Edge nodes are supported only on ESXi-based hosts with Intel-based and AMD-based chipsets.

Otherwise, vSphere EVC mode may prevent NSX Edge nodes from starting, showing an error message in the console.

- If vSphere EVC mode is enabled for the host for the NSX Edge VM, the CPU must be Haswell or later generation.
- Only VMXNET3 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 [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#) for details.

NSX Edge VM Resource Requirements

Appliance Size	Memory	vCPU	Disk Space	VM Hardware Version	Notes
NSX Edge Small	4 GB	2	200 GB	11 or later (vSphere 7.0 or later)	<p>Proof-of-concept deployments only.</p> <p>Note L7 rules for firewall, load balancing and so on are not realized on a Tier-1 gateway if you deploy a small sized NSX Edge VM.</p>
NSX Edge Medium	8 GB	4	200 GB	11 or later (vSphere 7.0 or later)	<p>Suggested for L2 through L4 features and when the total throughput requirement is less than 2 Gbps:</p> <ul style="list-style-type: none"> ■ NAT ■ Routing ■ L4 firewall ■ L4 load balancer

Appliance Size	Memory	vCPU	Disk Space	VM Hardware Version	Notes
NSX Edge Large	32 GB	8	200 GB	11 or later (vSphere 7.0 or later)	<p>Suggested for L2 through L4 features and when the total throughput requirement is between 2 ~ 10 Gbps.</p> <ul style="list-style-type: none"> ■ NAT ■ Routing ■ L4 Firewall ■ L4 Load Balancer ■ L7 Load Balancer (for example, when SSL offload is required) ■ TLS Inspection <p>See Scaling Load Balancer Resources in the <i>NSX Administration Guide</i>. For more information about what the different load balance sizes and NSX Edge form factors can support, see https://configmax.vmware.com.</p>
NSX Edge Extra Large	64 GB	16	200 GB	11 or later (vSphere 7.0 or later)	<p>Suggested for features that have a higher total throughput requirement than the NSX Edge Large form factor:</p> <ul style="list-style-type: none"> ■ L7 load balancer ■ VPN ■ TLS Inspection ■ L7 Access Profile (URL filtering) ■ IDS/IPS ■ Malware prevention (Advanced Threat Protection) <p>See Scaling Load Balancer Resources in the <i>NSX Administration Guide</i>. For more information about what the different load balance sizes and NSX Edge form factors can support, see https://configmax.vmware.com.</p>

NSX Edge VM CPU Requirements

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

- CPU must have AESNI capability.

- CPU must have 1 GB Huge Page support.

Hardware	Type
CPU	<ul style="list-style-type: none"> ■ 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) <hr/> <ul style="list-style-type: none"> ■ AMD EPYC Series processors

NSX Edge Bare Metal Requirements

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

NSX Edge Bare Metal Memory, CPU, and Disk Requirements

Minimum Requirements

Memory	CPU Cores	Disk Space
32 GB	8	200 GB

Recommended Requirements

Memory	CPU Cores	Disk Space
256 GB	24	200 GB

NSX Edge Bare Metal DPDK CPU Requirements

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

- CPU must have AES-NI capability.
- CPU must have 1 GB Huge Page support.
- (NSX 4.1.0) NSX Edge Bare Metal supports up to 64 cores for the entire system. This means that on a server with a single socket, its CPU can have up to 64 cores. On a server with 2 sockets, each socket cannot have more than 32 cores.
- (NSX 4.1.1) NSX Edge Bare Metal supports up to 80 cores for the entire system. This means that on a server with a single socket, its CPU can have up to 80 cores. On a server with 2 sockets, each socket cannot have more than 40 cores.

Hardware	Type
CPU	<ul style="list-style-type: none"> ■ 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) <hr/> <ul style="list-style-type: none"> ■ AMD EPYC Series processors

NSX Edge Bare Metal Hardware Requirements

Verify that the bare metal NSX Edge hardware is listed in this URL <https://ubuntu.com/certified?category=Server&release=20.04%20LTS&category=Server>. If the hardware is not listed, the storage, video adapter, or motherboard components might not work on the NSX Edge appliance. Bare Metal supports both UEFI and legacy BIOS modes.

NSX Edge Bare Metal NIC Requirements

NIC Type	Description	Vendor ID	PCI Device ID	Firmware Version
Mellanox ConnectX-4 Lx EN	PCI_DEVICE_ID_MELLANOX_CONNECTX4	15b3	0x1013	14.24.1000 and later versions
Mellanox ConnectX-4 Lx	PCI_DEVICE_ID_MELLANOX_CONNECTX4LX	15b3	0x1015	14.31.22.50 and later versions
Mellanox ConnectX-5	PCI_DEVICE_ID_MELLANOX_CONNECTX5	15b3	0x1017	16.21.1000 and later versions
Mellanox ConnectX-5 EX	PCI_DEVICE_ID_MELLANOX_CONNECTX5EX	15b3	0x1019	16.21.1000 and later versions
Mellanox ConnectX-6	PCI_DEVICE_ID_MELLANOX_CONNECTX6	15b3	0x101B	20.27.0090 and later versions
Mellanox ConnectX-6 Dx	PCI_DEVICE_ID_MELLANOX_CONNECTX6DX	15b3	0x101D	22.27.6008 and later versions
Intel X520/Intel 82599	IXGBE_DEV_ID_82599_KX4	8086	0x10F7	19.5.12 and later versions
	IXGBE_DEV_ID_82599_KX4_MEZZ	8086	0x1514	19.5.12 and later versions
	IXGBE_DEV_ID_82599_KR	8086	0x1517	19.5.12 and later versions
	IXGBE_DEV_ID_82599_COMBO_BACKPLANE	8086	0x10F8	19.5.12 and later versions
	IXGBE_DEV_ID_82599_CX4	8086	0x10F9	19.5.12 and later versions

NIC Type	Description	Vendor ID	PCI Device ID	Firmware Version
	IXGBE_DEV_ID_82599_SFP	8086	0x10FB	19.5.12 and later versions
	IXGBE_SUBDEV_ID_82599_SFP	8086	0x11A9	19.5.12 and later versions
	IXGBE_SUBDEV_ID_82599_RNDC	8086	0x1F72	19.5.12 and later versions
	IXGBE_SUBDEV_ID_82599_560FLR	8086	0x17D0	19.5.12 and later versions
	IXGBE_SUBDEV_ID_82599_ECNA_DP	8086	0x0470	19.5.12 and later versions
	IXGBE_DEV_ID_82599_SFP_EM	8086	0x1507	19.5.12 and later versions
	IXGBE_DEV_ID_82599_SFP_SF2	8086	0x154D	19.5.12 and later versions
	IXGBE_DEV_ID_82599_SFP_SF_QP	8086	0x154A	19.5.12 and later versions
	IXGBE_DEV_ID_82599_QSFP_SF_QP	8086	0x1558	19.5.12 and later versions
	IXGBE_DEV_ID_82599_EN_SFP	8086	0x1557	19.5.12 and later versions
	IXGBE_DEV_ID_82599_XAUI_LOM	8086	0x10FC	19.5.12 and later versions
	IXGBE_DEV_ID_82599_T3_LOM	8086	0x151C	19.5.12 and later versions
Intel X540	IXGBE_DEV_ID_X540T	8086	0x1528	n/a
	IXGBE_DEV_ID_X540T1	8086	0x1560	n/a
Intel X550	IXGBE_DEV_ID_X550T	8086	0x1563	21.5.9 and later versions
	IXGBE_DEV_ID_X550T1	8086	0x15D1	21.5.9 and later versions
	IXGBE_DEV_ID_X550E_M_A_10G_T	8086	0x15C8	21.5.9 and later versions
	IXGBE_DEV_ID_X550E_M_A_QSFP	8086	0x15CA	21.5.9 and later versions
	IXGBE_DEV_ID_X550E_M_A_QSFP_N	8086	0x15CC	21.5.9 and later versions
	IXGBE_DEV_ID_X550E_M_X_10G_T	8086	0x15AD	21.5.9 and later versions
	IXGBE_DEV_ID_X550E_M_X_XFI	8086	0x15B0	21.5.9 and later versions

NIC Type	Description	Vendor ID	PCI Device ID	Firmware Version
Intel X710	I40E_DEV_ID_SFP_X710	8086	0x1572	NSX 4.1.0: 7.00 and later versions NSX 4.1.1: 8.30
	I40E_DEV_ID_KX_C	8086	0x1581	NSX 4.1.0: 7.00 and later versions NSX 4.1.1: 8.30
	I40E_DEV_ID_10G_BA SE_T	8086	0x1586	NSX 4.1.0: 7.00 and later versions NSX 4.1.1: 8.30
	I40E_DEV_ID_10G_BA SE_T4	8086	0x1589	NSX 4.1.0: 7.00 and later versions NSX 4.1.1: 8.30
	I40E_DEV_ID_10G_SF P	8086	0x104E	n/a
	I40E_DEV_ID_10G_B	8086	0x104F	n/a
Intel XL710	I40E_DEV_ID_KX_B	8086	0x1580	NSX 4.1.x: 9.0 and later versions
	I40E_DEV_ID_QSFP_A	8086	0x1583	NSX 4.1.x: 9.0 and later versions
	I40E_DEV_ID_QSFP_B	8086	0x1584	NSX 4.1.x: 9.0 and later versions
	I40E_DEV_ID_QSFP_C	8086	0x1585	NSX 4.1.x: 9.0 and later versions
	I40E_DEV_ID_20G_KR 2	8086	0x1587	NSX 4.1.x: 9.0 and later versions
	I40E_DEV_ID_20G_KR 2_A	8086	0x1588	NSX 4.1.x: 9.0 and later versions
	I40E_DEV_ID_10G_BA SE_T_BC	8086	0x15FF	NSX 4.1.x: 9.0 and later versions
Intel XXV710	I40E_DEV_ID_25G_B	8086	0x158A	NSX 4.1.x: 9.0 and later versions
	I40E_DEV_ID_25G_SF P28	8086	0x158B	NSX 4.1.x: 9.0 and later versions
Intel E810-C (100G)	ICE_DEV_ID_E810C_Q SFP	8086	0x1592	NSX 4.1.1: 4.0.0 and later versions
Intel E810-XXV (25G)	ICE_DEV_ID_E810_XX V_SFP	8086	0x159B	NSX 4.1.1: 4.0.0 and later versions
	ICE_DEV_ID_E810_XX V_BACKPLANE	8086	0x1599	NSX 4.1.1: 4.0.0 and later versions
	ICE_DEV_ID_E810_XX V_QSFP	8086	0x159A	NSX 4.1.1: 4.0.0 and later versions

NIC Type	Description	Vendor ID	PCI Device ID	Firmware Version
Intel E810-C	ICE_DEV_ID_E810C_B ACKPLANE	8086	0x1591	NSX 4.1.1: 4.0.0 and later versions
	ICE_DEV_ID_E810C_SF P	8086	0x1593	NSX 4.1.1: 4.0.0 and later versions
Intel E822-C	ICE_DEV_ID_C822N_B ACKPLANE	8086	0x1890	NSX 4.1.1: 3.1 and later versions
	ICE_DEV_ID_C822N_Q SFP	8086	0x1891	NSX 4.1.1: 3.1 and later versions
	ICE_DEV_ID_C822N_S FP	8086	0x1892	NSX 4.1.1: 3.1 and later versions
Cisco VIC 1300 series	Cisco UCS Virtual Interface Card 1300	1137	0x0043	n/a
Cisco VIC 1400 series	Cisco UCS Virtual Interface Card 1400	1137	0x0043	n/a

Note For all the supported NICs listed above, verify that the media adapters and cables you use follow the vendor's supported media types. Any media adapter or cables not supported by the vendor can result in unpredictable behavior, including the inability to boot up due to an unrecognized media adapter. See the NIC vendor documentation for information about supported media adapters and cables.

Important Cisco VICs: To successfully claim Cisco VICs for the NSX Edge datapath, configure multiple RX and TX queues from the Cisco UCS Manager. The number of queues configured must be sufficient for datapath to have one queue per core. For configuration details, refer to the Cisco documentation.

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 Linux Physical Server](#).

Bare Metal Server Requirements

Operating System	Version	CPU Cores	Memory
CentOS Linux	7.6 (kernel: 3.10.0-957) 7.7 7.9 8.2 8.4	4	16 GB
Red Hat Enterprise Linux (RHEL)	7.6 (kernel: 3.10.0-957) 7.7 7.9 8.2 8.4 8.6 (starting with NSX 4.0.1)	4	16 GB
Oracle Linux	7.6 (kernel: 3.10.0-957) 7.7 7.8 7.9 8.6 (starting with NSX 4.0.1)	4	16 GB
SUSE Linux Enterprise Server	12 SP3 12 SP4 12 SP5 (starting with NSX 3.2.1)	4	16 GB
Ubuntu	16.04.2 LTS (kernel: 4.4.0-*) 18.04 20.04	4	16 GB
Windows Server	2012 R2 Datacenter and Standard 2016 (minor version 14393.2248 and later) 2019	4	16 GB

- Ensure MTU is set to 1600 for Jumbo frame support by NIC or OS drivers.
- Hosts running Ubuntu 18.04.2 LTS must be upgraded from 16.04 or freshly installed.

Supported Topologies

To find the complete list of supported topologies, see the [NSX Physical Server Encyclopedia](#) slide deck.

Physical NICs

For physical servers running Linux: There is no restriction on the physical NIC other than being supported by the operating system.

For physical servers running Windows on Segment-VLAN with MTU at 1500, there is also no restrictions on the physical NIC other than being supported by the operating system.

For physical servers running Windows on Segment-Overlay or with Segment-VLAN with large MTU (> 1500), validate its associated driver support jumbo packet. To verify whether jumbo packet is supported, run the following command:

```
$ (Get-NetAdapterAdvancedProperty -Name "<Ethernet>").DisplayName -Contains "Jumbo Packet"
```

Where, <Ethernet> must be replaced with the real adapter name of each physical NIC."

Table 3-2. Virtual NICs

NIC Type	Description	PCI BUS ID	Firmware Version
e1000e	Intel(R) 82574L Gigabit Network	0000:1b:00	12.15.22.6 and later version
vmxnet3	vmxnet3 Ethernet Adapter	0000:0b:00	1.9.2.0 and later version

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, the paths are secured and authenticated, and a storage location for the credentials are used to establish mutual authentication.

Configure the ports and protocols required to be open on both the physical and the host hypervisor firewalls in NSX. Refer to <https://ports.vmware.com/home/NSX-T-Data-Center> for more details.

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:

- ESXi: nsx-cfgagent, ESX-DP (in the kernel)

Note To get access to NSX nodes, you must enable SSH on these nodes.

Configuring a vSphere Distributed Switch

When a transport node is configured on a VDS host switch, some network parameters can only be configured in VMware vCenter.

The following requirements must be met to install NSX on a VDS host switch:

- VMware vCenter 7.0 or a later version
- ESXi 7.0 or a later version

The created VDS switch can be configured to centrally manage networking for NSX hosts.

Configuring a VDS switch for NSX networking requires objects to be configured on NSX and in vCenter Server.

- In vSphere:
 - Create a VDS switch.
 - Set MTU to at least 1600
 - Add ESXi hosts to the switch. These hosts are later prepared as NSX transport nodes.
 - Assign uplinks to physical NICs.
- In NSX:
 - When configuring a transport node, map uplinks created in NSX uplink profile with uplinks in VDS.

For more details on preparing a host transport node on a VDS switch, see the *NSX Installation Guide*.

The following parameters can only be configured in a VMware vCenter on a VDS backed host switch:

Configuration	VDS	NSX	Description
MTU	<p>In VMware vCenter, set an MTU value on the switch.</p> <hr/> <p>Note A VDS switch must have an MTU of 1600 or higher.</p> <p>In VMware vCenter, select VDS, click Actions → Settings → Edit Settings.</p>	Any MTU value set in an NSX uplink profile is overridden.	As a host transport node that is prepared using VDS as the host switch, the MTU value needs to be set on the VDS switch in vCenter Server.
Uplinks/LAGs	<p>In VMware vCenter, configure Uplinks/LAGs on a VDS switch.</p> <p>In VMware vCenter, select VDS, click Actions → Settings → Edit Settings.</p>	When a transport node is prepared, the teaming policy on NSX is mapped to uplinks/LAGs configured on a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the uplink or LAG are configured on the VDS switch. During configuration, NSX requires teaming policy be configured for the transport node. This teaming policy is mapped to the uplinks/LAGs configured on the VDS switch.
NIOC	<p>Configure in VMware vCenter.</p> <p>In VMware vCenter, select VDS, click Actions → Settings → Edit Settings.</p>	NIOC configuration is not available when a host transport node is prepared using a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the NIOC profile can only be configured in vCenter Server.

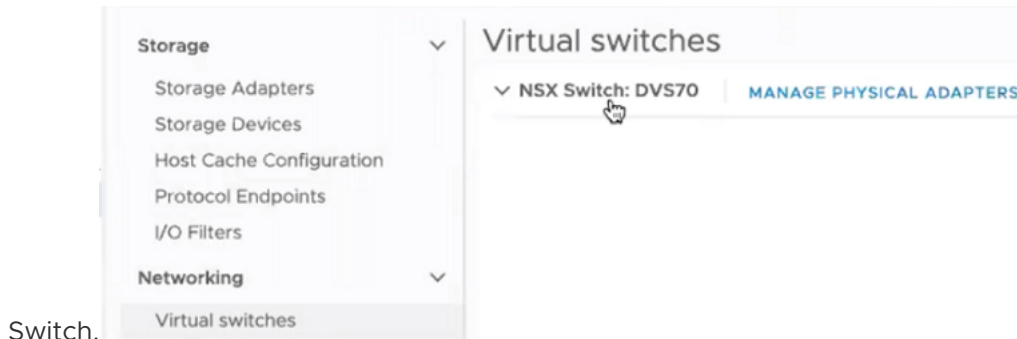
Configuration	VDS	NSX	Description
Link Layer Discovery Protocol (LLDP)	Configure in VMware vCenter. In VMware vCenter, select VDS, click Actions → Settings → Edit Settings .	LLDP configuration is not available when a host transport node is prepared using a VDS switch.	As a host transport node that is prepared using VDS as the host switch, the LLDP profile can only be configured in vCenter Server.
Add or Manage Hosts	Manage in VMware vCenter. In VMware vCenter, go to Networking → VDS Switch → Add and Manage Host..	Prepared as transport nodes in NSX.	Before preparing a transport node using a VDS switch, that node must be added to the VDS switch in vCenter Server.

Note NIOC profiles, Link Layer Discovery Protocol (LLDP) profile, and Link Aggregation Group (LAG) for these virtual machines are managed by VDS switches and not by NSX. As a vSphere administrator, configure these parameters from VMware vCenter UI or by calling VDS API commands.

After preparing a host transport node with VDS as a host switch, the host switch type displays VDS as the host switch. It displays the configured uplink profile in NSX and the associated transport zones.



In VMware vCenter, the VDS switch used to prepare NSX hosts is created as an NSX



Switch.

Checklist Before Deploying Infrastructure

Note Ensure NSX components (controllers, transport nodes, and so on, including management interfaces on hosts as well as on NSX Edge nodes (bare metal and VM form factors) and compute hosts have IP Connectivity between them.

VMware vCenter Infrastructure Checklist

- 1 At least four vSphere hosts must be available on management and compute host clusters in vCenter. This requirement is a best practice for vSphere HA and vSphere Dynamic Resource Scheduling (DRS) features.

Consider these points related to workloads:

- To be resilient to rack failures, spread clustered elements between racks.
 - Use DRS anti-affinity rules to spread the workloads across racks.
- 2 Enable vSphere DRS, HA, vSAN capabilities on existing host clusters.
 - 3 Do one of the following:
 - Configure management VDS to span the management cluster and Configure NSX Edge cluster and compute VDS to span the compute clusters.
 - Configure a single VDS to span both management and compute clusters.
 - 1 VDS portgroups for NSX Appliances: NSX Manager appliances are deployed on a hypervisor backed by a standard VLAN backed portgroup.
 - 2 VDS portgroups for ESXi Hypervisors:
 - A vmk0 portgroup for ESXi host mgmt.
 - A vmotion portgroup for vmotion traffic.
 - A storage portgroup for storage traffic.

Note All ESXi vmk interfaces can reside on the same portgroup or different management portgroups. Optionally, the vmks can be configured with a VLAN ID to provide logical isolation between traffic types.

- 3 VDS trunk portgroups for NSX Edge nodes (when using VM form factor for Edges):
 - A management VLAN trunk portgroup for NSX Edge management.
 - NSX Edge VLAN trunk portgroups to connect to DPDK fast-path interfaces. In NSX v3.2 and later, you can enable up to four datapath interfaces on NSX Edge VMs.
 - Configure one NSX Edge trunk portgroup to connect to all NSX Edge fast-path interfaces.

- Alternatively, configure multiple VLAN trunk portgroups pinned to specific VLAN IDs. And configure one active uplink (with a standby uplink) to steer VLAN traffic to a specific TOR through the specified uplink. Then connect each NSX Edge fast-path interface to separate NSX Edge trunk portgroups in vCenter.
- 4 Change VDS MTU from default value 1500 bytes to 1700 or 9000 bytes (recommended value). See [Guidance to Set Maximum Transmission Unit](#).
 - 5 Create VDS portgroups for Tunnel Endpoints (TEPs): Hypervisor TEP VLAN and NSX Edge TEP VLAN can share same VLAN ID if TEPs are connected to NSX VLAN segments. Otherwise, use different VLANs for the TEPs.

Note Management hypervisors do not need TEP VLAN as these do not need to be configured as transport nodes.

NSX Infrastructure Checklist

- 1 For each of the three NSX Manager nodes that you will deploy in the management cluster, set up reserve IP addresses and FQDN.
- 2 Ensure IP address and FQDN used as VIP of NSX Manager are in same subnet as the NSX appliances IP address. If you are using an external Load Balancer for VIP, then use a VIP from a different subnet.
- 3 Configure LDAP server.
- 4 Configure Backup server.
- 5 Configure Syslog server.
- 6 Configure DHCP server or DHCP relay server.
- 7 (Optional) Configure CA-signed certificates to apply to each NSX Manager node and NSX Manager cluster VIP. The default value is self-signed.

Note After you configure CA-signed certificates, you must configure FQDN for NSX Manager appliances, VIP and vCenter.

- 8 (optional) If NSX Edge appliances are going to connect to NSX VLAN segments (instead of a VDS on vCenter), then reserve VLAN IDs to be used when you create VLAN segments for NSX Edge management and NSX Edge DPDK fast-path interfaces.
- 9 Configure IP Pools that can be used when you configure TEPs for compute hypervisors.
- 10 Configure IP Pools that can be used when you configure TEPs for NSX Edge bare metal (physical server) or NSX Edge VMs.
- 11 (Optional) Global MTU.
- 12 (Optional) Partial Patch support.
- 13 (Optional) Configure NSX failure domain for the NSX Edge clusters.
- 14 (Optional) Cluster FQDN.

15 If configuring single tier or multi-tier topology in NSX, then following additional configuration required.:

- Configure Tier-0 Gateway external interfaces (uplink interfaces): Reserve at least two VLANs that will be used as NSX Edge uplinks to TOR and configured as Tier-0 gateway interfaces.
- The Service Router component of the Tier-0 Gateways gets instantiated on Edge TN Cluster node upon creation of external interface. Example: If a Tier-0 gateway resides on NSX Edge cluster with four NSX Edge nodes, create at least two VLAN segments that will act as two uplinks to each of the four edge nodes.

Example:

Create two VLAN segments:

segment-vlan2005

segment-vlan2006

Create the first external interfaces on the Tier-0 gateway using these segments.

Name: **EXT-INT-VLAN-2005**

Type: **External**

IP Address Mask: **VLAN-X CIDR address**

Connected To(Segment): **segment-vlan2005**

NSX Edge Node: **Edge-1, Edge-2, Edge-3, Edge-4**

Create the second external interfaces on the Tier-0 gateway using these segments.

Name: **EXT-INT-VLAN-2006**

Type: **External**

IP Address Mask: **VLAN-Y CIDR address**

Connected To (Segment): **segment-vlan2006**

NSX Edge Node: **Edge-1, Edge-2, Edge-3, Edge-4**

16 Configure dedicated BGP peers (Remote/Local ASN no) or static routing on Tier0 gateway interfaces.

17 Enable BFD per BGP neighbor for faster failover. Ensure BGP and BFD configuration is between ToR switches and NSX Edge nodes.

18 Set route redistribution for TO and T1 routes.

19 Reserve VLAN IDs and IP addresses to be used to configure a Service Interface. It connects VLAN-backed segments or logical switches to VLAN-backed physical or virtual workloads. This interface acts as a gateway for these VLAN backed workloads and is supported both on Tier-0 and Tier-1 Gateways configured in active/standby HA configuration mode.

NSX Manager Installation Requirements

4

NSX Manager provides a graphical user interface (GUI) and REST APIs for creating, configuring, and monitoring NSX components such as logical switches (segments), logical routers (Tier-0, Tier-1 Gateways) and firewalls.

NSX Manager provides a system view and is the management component of NSX.

For high availability, NSX supports a management cluster of three NSX Managers. For a production environment, configuring a management cluster is recommended. Starting with NSX 3.1, a single NSX Manager cluster deployment is supported.

In a vSphere environment, the following functions are supported by NSX Manager:

- vCenter Server can use the vMotion function to live migrate NSX Manager across hosts and clusters.
- vCenter Server can use the Storage vMotion function to live migrate file system of an NSX Manager across hosts and clusters.
- vCenter Server can use the Distributed Resource Scheduler function to rebalance NSX Manager across hosts and clusters.
- vCenter Server can use the Anti-affinity function to manage NSX Manager across hosts and clusters.

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
Supported platforms	See NSX Manager VM and Host Transport Node System Requirements . On ESXi, it is recommended that the NSX Manager appliance be installed on shared storage.

Requirements	Description
IP address	<p>An NSX Manager must have a static IP address. You can change the IP address after installation. Both IPv4 and IPv6 are supported. You can choose IPv4 only or use dual stack (both IPv4 and IPv6).</p> <hr/> <p>Note If you choose to use one IPv4 only, then the NSX Manager services (for example, SNMP, NTP, vIDM, etc.) must have IPv4 addresses.</p>
NSX appliance password	<ul style="list-style-type: none"> ■ 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 ■ Default password complexity rules are enforced by the following Linux PAM module arguments: <ul style="list-style-type: none"> ■ <code>retry=3</code>: The maximum number of times a new password can be entered, for this argument at the most 3 times, before returning with an error. ■ <code>minlen=12</code>: The minimum acceptable size for the new password. In addition to the number of characters in the new password, credit (of +1 in length) is given for each different kind of character (other, upper, lower and digit). ■ <code>difok=0</code>: The minimum number of bytes that must be different in the new password. Indicates similarity between the old and new password. With a value 0 assigned to <code>difok</code>, there is no requirement for any byte of the old and new password to be different. An exact match is allowed. ■ <code>lcredit=1</code>: The maximum credit for having lower case letters in the new password. If you have less than or 1 lower case letter, each letter will count +1 towards meeting the current <code>minlen</code> value. ■ <code>ucredit=1</code>: The maximum credit for having upper case letters in the new password. If you have less than or 1 upper case letter each letter will count +1 towards meeting the current <code>minlen</code> value. ■ <code>dcredit=1</code>: The maximum credit for having digits in the new password. If you have less than or 1 digit, each digit will count +1 towards meeting the current <code>minlen</code> value. ■ <code>ocredit=1</code>: The maximum credit for having other characters in the new password. If you have less than or 1 other characters, each character will count +1 towards meeting the current <code>minlen</code> value. ■ <code>enforce_for_root</code>: The password is set for the root user. <hr/> <p>Note For more details on Linux PAM module to check the password against dictionary words, refer to the man page.</p> <p>For example, avoid simple and systematic passwords such as <code>VMware123!123</code> or <code>VMware12345</code>. Passwords that meet complexity standards are not simple and systematic but are a combination of letters, alphabets, special characters, and numbers, such as <code>VMware123!45</code>, <code>VMware 1!2345</code> or <code>VMware@1az23x</code>.</p>

Requirements	Description
Hostname	<p>When installing NSX Manager, specify a hostname that does not contain invalid characters such as an underscore or special characters such as dot ".". If the hostname contains any invalid character or special characters, after deployment the hostname will be set to nsx-manager.</p> <hr/> <p>Important When you install NSX installation as a dual stack (IPv4 and IPv6) and/or you plan to configure CA-signed certificates, configure a hostname with valid domain name in NSX Manager VMs and Cluster VIP (if configured).</p> <hr/> <p>For more information about hostname restrictions, see https://tools.ietf.org/html/rfc952 and https://tools.ietf.org/html/rfc1123.</p>
VMware Tools	<p>The NSX Manager VM running on ESXi has VMTTools installed. Do not remove or upgrade VMTTools.</p>
System	<ul style="list-style-type: none"> ■ 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 or FQDN for the NSX Manager to use. ■ Create a management VDS and target VM port group in vCenter. Place the NSX appliances onto this management VDS port group network. See Prepare a vSphere Distributed Switch for NSX. <p>Multiple management networks can be used as long as the NSX Manager nodes has consistent connectivity and recommended latency between them.</p> <hr/> <p>Note If you plan to use Cluster VIP, all NSX Manager appliances should belong to same subnet.</p> <hr/> <ul style="list-style-type: none"> ■ Plan your NSX Manager IP and NSX Manager Cluster VIP addressing scheme. <hr/> <p>Note Verify that you have the hostname for NSX Manager to use. The Hostname format must be <code>nsx-manager-fqdn@domain-name.com</code>. This format is required if NSX installation is dual stack (IPv4, IPv6) and/or if planning to configure CA-signed certificates.</p>
OVF Privileges	<p>Verify that you have adequate privileges to deploy an OVF template on the ESXi host.</p> <p>A management tool that can deploy OVF templates, such as VMware vCenter or the vSphere Client. The OVF deployment tool must support configuration options to allow for manual configuration.</p> <p>OVF tool version must be 4.0 or later.</p>
Client Plug-in	<p>The Client Integration Plug-in must be installed.</p>
Certificates	<p>If you plan to configure internal VIP on a NSX Manager cluster, you can apply a different certificate to each NSX Manager node of the cluster. See Configure a Virtual IP Address for a Cluster.</p> <p>If you plan to configure an external load balancer, ensure only a single certificate is applied to all NSX Manager cluster nodes. See Configuring an External Load Balancer.</p>

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 as a standalone ESXi host, OVA/OVF required fields as well as user names and password criteria are not validated before the VM is powered on. However, the static IP address field is a mandatory field to install NSX Manager. When you install NSX Manager as a managed host in VMware vCenter, OVA/OVF required fields as well as user names and password criteria are validated before the VM is powered on.

- If you specify a user name for any local user, the name must be unique. If you specify the same name, it is ignored and the default names (for example, **admin** and **audit**) are used.
 - If the password for the **root** or **admin** user does not meet the complexity requirements, you must log in to NSX Manager through SSH or at the console as **root** with password **vmware** and **admin** with password **default**. You are prompted to change the password.
 - If the password for other local users (for example, **audit**) 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 local_user_name** to set the local user's password (the current password is an empty string). You can also reset passwords in the UI using **System > User Management > Local Users**.
-

Caution Changes made to the NSX 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 VMware vCenter.

Read the following topics next:

- [Configure NSX Manager for Access by DNS Server](#)
- [Publish FQDN of the NSX Managers](#)
- [Modifying the Default Admin Password Expiration](#)

Configure NSX Manager for Access by DNS Server

In a few scenarios, NSX Managers require DNS to be configured so that the Manager can perform DNS lookups. This is important to know so that you can configure your DNS server before you deploy the NSX Manager.

NSX requires forward and reverse DNS entries for any of the following scenarios:

- Starting with NSX 4.1 and later versions, in a dual stack environment (that is, both IPv4 and IPv6 have been configured).
- Use cases where `"publish_fqdns": true`.
- NSX Manager that uses CA-signed certificates.
- NSX Manager with Multisite deployments. (It is optional for all other deployment types.) See *Multisite Deployment of NSX* in the *NSX Administration Guide*.

Note If you did not provide a fully-qualified hostname (FQDN) while deploying NSX Manager in a dual stack environment, then you may be required to replace the REST API certificates because they may not have generated correctly during first boot. You will notice this problem if your browser does not trust the NSX Manager certificate in which case the browser will ask you if you want to ignore the problem. You can either continue to ignore the problem or replace the Manager's REST API certificate.

To ensure a valid FQDN is configured for both the IPv4 and IPv6 addresses used to deploy the NSX Manager and that both address types point to the same FQDN, use the following workflow.

Prerequisites

Understand the scenarios in which you plan to use the NSX Manager. To avoid any problems, ensure you always configure the NSX Manager hostname to be fully qualified. If the NSX Manager hostname is always fully qualified, then the initial certificates will be generated correctly and will match the DNS server records. If any of the following scenarios are present that require NSX to use forward and reverse DNS entries and you have already deployed NSX, make sure you complete step one in the following procedure.

Procedure

- 1 For NSX Manager hostname queries to work successfully for the scenarios mentioned in this topic, configure the DNS server with a DNS A record, a DNS AAAA record, and PTR records for both IPv4 and IPv6 addresses.

Note that the AAAA record is only required if you have configured IPv6 addresses.

Most customers currently do not use IPv6 addresses.

- 2 Deploy a NSX Manager with a hostname that is fully qualified (that is, FQDN) so that initial certificates are generated correctly and match the DNS server records (configured in Step 1).

For details, see the step that covers the setting the NSX Manager hostname of the OVF template. For example, [Install NSX Manager and Available Appliances](#) or [Install NSX Manager on ESXi Using the Command-Line OVF Tool](#).

What to do next

Run `get hostname` CLI to confirm if FQDN is set.

Publish FQDN of the NSX Managers

Configure FQDN on NSX Manager in a dual stack mode.

When the Publish FQDN functionality is enabled in NSX Manager, transport nodes can access an NSX Manager using its DNS name instead of IP address. By default, transport nodes access NSX Manager based on their IP addresses. However, this can be based also on the DNS names of the NSX Manager. This gives ability to change IP addresses for disaster recovery.

Note Enabling the FQDN publish functionality is not a requirement for single-site deployment. However, you must enable it for NSX Multisite deployments as this gives transport nodes ability to communicate with the NSX Manager during disaster recovery, when NSX Manager IP address changes. This is because the failover process requires a manual restore of the NSX Managers to the DR site where they will use a new IP address. If the NSX Managers are not registered to use FQDN then all the transport nodes will not know that the IP has changed and will fail to reconnect.

Prerequisites

- Add unique hostname records of IPv4 and IPv6 address to the respective DNS server.
For example, (IPv4 address) Name: nsx-mgr-01.eng.vmware.com, Address:10.176.132.45
(IPv6) Name: nsx-mgr-01.eng.vmware.com, Address: 2620:124:6020:1045::b

Procedure

- ◆ To enable FQDN usage of NSX Managers by Transport Node, using the following NSX API.

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

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


Example response

```
{ "publish_fqdns": true,
  "_revision": 1
}
```

Important When using this feature, configure both the forward and reverse lookup entries for the NSX Managers FQDN on the DNS servers and with a short TTL, for example, 600 seconds.

What to do next

After publishing FQDNs, you must verify that transport nodes are accessing the NSX Managers using their FQDN.

- Log in to a transport node such as a hypervisor or NSX Edge node using SSH, and run the `get controllers` CLI command to verify that 'Controller FQDN' field is populated correctly with NSX Manager FQDN value.

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

Modifying the Default Admin Password Expiration

By default, the administrative password for the NSX Manager and NSX Edge appliances expires after 90 days. However, you can reset the expiration period after initial installation and configuration.

If the password expires, you will be unable to log in and manage components. Additionally, any task or API call that requires the administrative password to execute will fail. If your password expires, see Knowledge Base article 70691 [NSX-T admin password expired](#).

Procedure

- 1 Use a secure program to connect to the NSX CLI console.
- 2 Reset the expiration period.

You can set the expiration period for between 1 and 9999 days.

```
nsxcli> set user admin password-expiration <1 - 9999>
```

Note Alternatively, you can use API commands to set the admin password expiration period.

- 3 (Optional) You can disable password expiry so the password never expires.

```
nsxcli> clear user admin password-expiration
```

NSX Manager Cluster Requirements

5

The following subsections describe cluster requirements for NSX appliances and provides recommendations for specific site deployments.

Cluster Requirements

- In a production environment, the NSX Manager (Local Manager in an NSX Federation environment) or Global Manager cluster must have three members for scaling out and for redundancy of its 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 cluster of the NSX Manager (Local Manager in an NSX Federation environment) or Global Manager. However, you can add additional, temporary nodes to allow for IP address changes.

Important NSX Manager cluster requires a quorum to be operational. This means at least two out of three of its members must be up and running. If a quorum is lost, NSX management and control planes do not work. This results in no access to the NSX Manager's UI and API for configuration updates using the management plane as well as no access to add or vMotion VMs on NSX segments. The dynamic routing on Tier-0 remains operational.

- Starting with NSX 3.1.1, NSX supports the deployment of a single NSX Manager or Global Manager in production deployments. To recover from a failed NSX Manager, configure vSphere HA for the single NSX Manager. The time required to recover a single NSX Manager that uses backup/restore or vSphere HA might be much longer than the availability provided by a cluster of NSX Manager nodes. Similarly, for lab and proof-of-concept deployments where there are no production workloads, you can run a single NSX Manager or Global Manager to save resources. NSX Manager or Global Manager nodes can be deployed on either ESXi or KVM. However, mixed deployments of managers on both ESXi and KVM are not supported.

- If you deploy a single NSX Manager or Global Manager in a production environment, you can also enable vSphere HA on the host cluster for the manager nodes to get automatic recovery of the manager VM in case of ESXi failure. For more information on vSphere HA, see the *vSphere Availability* guide in the vSphere Documentation Center.
- If one of the NSX Manager node goes down, the three node NSX Manager cluster provides redundancy. However, if one of the NSX Manager node loses its disk, it goes into read only state. It might still continue to participate in the cluster if its network connectivity is Up. This case can result in the management plane and control plane to become unavailable. To resolve the issue, power off the impacted NSX Manager.

Configure Anti-affinity rules in VMware vCenter

Set up DRS Anti-Affinity rules to prevent, whenever possible, two NSX Manager VMs from running on the same host. Follow the vSphere guide to configure VM to VM affinity rules. See [Create a VM-VM Affinity Rule](#).

VM to Host affinity rules. See [Create a VM-Host Affinity Rule](#).

Recovery with vSphere HA

You can use vSphere HA (High Availability) with NSX to enable quick recovery if the host running the NSX Manager node fails. See *Creating and Using vSphere HA Clusters* in the vSphere product documentation.

Read the following topics next:

- [Cluster Requirements for an Individual Site](#)
- [Cluster Requirements for Multiple Sites](#)

Cluster Requirements for an Individual Site

Each NSX appliance cluster – Global Manager, Local Manager or NSX Manager – must contain three VMs.

Those three VMs can be physically all deployed in the same data center or in different data centers, as long as latency between VMs in the cluster is below 10ms. In a 3-node Manager cluster, only one of the nodes can fail at any given time for the cluster to self-heal.

Single Site Requirements and Recommendations

The following recommendations apply to single site NSX deployments and cover cluster recommendations for Global Manager, Local Manager and NSX Manager appliances:

- Place your NSX appliances on different hosts (from anti-affinity rules in VMware vCenter) to avoid a single host failure impacting multiple managers.
- Maximum latency between NSX appliances is 10ms.

- You can place NSX appliances in different vSphere clusters or in a common vSphere cluster.
- It is recommended that you place NSX appliances in different management subnets or a shared management subnet. When using vSphere HA it is recommended to use a shared management subnet so NSX appliances that are recovered by vSphere can preserve their IP address.
- It is recommended that you place NSX appliances on shared storage also. For vSphere HA, please review the requirements for that solution.

You can also use vSphere HA with NSX to provide recovery of a lost NSX appliance when the host where the NSX appliance 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.

Cluster Requirements for Multiple Sites

The following recommendations apply to dual site (Site A/Site B) and multiple-site (Site A/Site B/Site C) NSX deployments.

Dual Site Requirements and Recommendations

Note Starting from NSX v3.0.2 onwards, VMware Site Recovery Manager (SRM) is supported.

- Do not deploy NSX Managers in a dual-site scenario without vSphere HA or VMware SRM. In this scenario, one site requires the deployment of two NSX Managers and the loss of that site will impact the operation of NSX.
- Deployment of NSX Managers in a dual site scenario with vSphere HA or VMware SRM 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 or VMware SRM 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 or VMware SRM will recover all NSX Managers to hosts in site B.

Important You must 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

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

If you deploy without vSphere HA or VMware SRM:

- 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 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 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.

Installing NSX Manager Cluster on vSphere

6

As an administrator, your first task will be to install NSX in setting up your NSX environment.

NSX Manager is the application that you use to administer your NSX environment. You can use the NSX Manager UI, API or CLI to manage workloads and NSX Edge nodes. In a production environment, for fault tolerance, deploy a cluster of three NSX Manager nodes, each running on a separate ESXi host.

NSX Manager can be deployed either on an ESXi host or a KVM host and this section covers the use of vSphere Client to deploy the NSX Manager virtual appliances OVA/OVF on an ESXi host.

As an admin, your first task is to install NSX Manager as part of setting up your NSX environment.

Make sure that you have the supported vSphere version. See [vSphere support](#).

If you deploy the NSX appliances using OVAs, then there is no need to disable snapshots because the OVA template is preconfigured to handle snapshots.

Read the following topics next:

- [Install NSX Manager and Available Appliances](#)
- [Configure a Virtual IP Address for a Cluster](#)
- [Configuring an External Load Balancer](#)
- [Verify Appliance Proxy Hub on all NSX Manager Nodes are Connected](#)

Install NSX Manager and Available Appliances

You can use the vSphere Client to deploy NSX Manager virtual appliances. The same OVF file can be used to deploy three different types of appliances: NSX Manager, NSX Cloud Service Manager for NSX Cloud, and Global Manager for NSX Federation.

Cloud Service Manager is a virtual appliance that uses NSX 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 or FQDN for the NSX Manager to use.
- Create a management VDS and target VM port group in vCenter. Place the NSX appliances onto this management VDS port group network. See [Prepare a vSphere Distributed Switch for NSX](#).

Multiple management networks can be used as long as the NSX Manager nodes has consistent connectivity and recommended latency between them.

Note If you plan to use Cluster VIP, all NSX Manager appliances should belong to same subnet.

- Plan your NSX Manager IP and NSX Manager Cluster VIP addressing scheme.

Note Verify that you have the hostname for NSX Manager to use. The Hostname format must be `nsx-manager-fqdn@domain-name.com`. This format is required if NSX installation is dual stack (IPv4, IPv6) and/or if planning to configure CA-signed certificates.

Procedure

- 1 Locate and download the NSX OVA file from the **My Downloads** panel on the [Broadcom Support](#) page. You can either copy the download URL or download the OVA file.
- 2 In the vSphere Client, select the ESXi host or ESXi host cluster on which to install NSX.
- 3 Right-click Host and select **Deploy OVF template** to start the installation wizard.
- 4 Enter the download OVA URL or navigate to the OVA file, and click **Next**.
- 5 Enter a name and a location for the NSX Manager VM, and click **Next**.

The name you enter appears in the vSphere and VMware vCenter inventory.

- 6 Select a compute resource for the NSX Manager appliance, and click **Next**.
 - ◆ 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 Review and verify the OVF template details, and click **Next**.
- 8 Specify the deployment configuration size, and click **Next**.

The Description panel on the right side of the wizard shows details of the selected configuration.

- 9 Specify storage for the configuration and disk files.
 - a Select the virtual disk format.
 - b Select the VM storage policy.

- c Specify the datastore to store the NSX Manager appliance files.
 - d Click **Next**.
- 10 Select a destination network for each source network.
 - 11 Select the port group or destination network for the NSX Manager.
 - 12 Configure IP Allocation settings.
 - a For IP allocation, specify **Static - Manual**.
 - b For IP protocol, select **IPv4** or **IPv6**.

Note You can ignore the IP Allocation settings. You can select either IPv4 or IPv6. It would not impact ingress or egress network traffic of NSX Manager.

- 13 Click **Next**.

The following steps are all located in the Customize Template section of the Deploy OVF Template wizard.

- 14 In the Application section, enter the system root, CLI admin, and audit passwords for the NSX Manager. The **root** and **admin** credentials are mandatory fields.

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
- Default password complexity rules are enforced by the following Linux PAM module arguments:
 - `retry=3`: The maximum number of times a new password can be entered, for this argument at the most 3 times, before returning with an error.
 - `minlen=12`: The minimum acceptable size for the new password. In addition to the number of characters in the new password, credit (of +1 in length) is given for each different kind of character (other, upper, lower and digit).
 - `difok=0`: The minimum number of bytes that must be different in the new password. Indicates similarity between the old and new password. With a value 0 assigned to `difok`, there is no requirement for any byte of the old and new password to be different. An exact match is allowed.

- `lcredit=1`: The maximum credit for having lower case letters in the new password. If you have less than or 1 lower case letter, each letter will count +1 towards meeting the current `minlen` value.
- `ucredit=1`: The maximum credit for having upper case letters in the new password. If you have less than or 1 upper case letter each letter will count +1 towards meeting the current `minlen` value.
- `dcredit=1`: The maximum credit for having digits in the new password. If you have less than or 1 digit, each digit will count +1 towards meeting the current `minlen` value.
- `ocredit=1`: The maximum credit for having other characters in the new password. If you have less than or 1 other characters, each character will count +1 towards meeting the current `minlen` value.
- `enforce_for_root`: The password is set for the root user.

Note For more details on Linux PAM module to check the password against dictionary words, refer to the man page.

For example, avoid simple and systematic passwords such as `VMware123!123` or `VMware12345`. Passwords that meet complexity standards are not simple and systematic but are a combination of letters, alphabets, special characters, and numbers, such as `VMware123!45`, `VMware 1!2345` or `VMware@1az23x`.

15 In the Optional parameters section, leave the password fields blank. It avoids the risk of compromising passwords set for VMC roles by a user who has access to the VMware vCenter. When deploying VMC for NSX, this field is used internally to set passwords for the Cloud Admin and Cloud Operator roles.

16 In the Network Properties section, 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 starts with an alphabet character. Also, NSX accepts only latin alphabets that do not have an accent mark, as in í, ó, ú, ý.

Important If you plan to install NSX in dual stack (IPv4 and IPv6) and/or if you plan to configure CA-signed certificates, then enter a Hostname with valid domain name to NSX Manager VMs and Cluster VIP (if configured).

17 Select a **Rolename** for the appliance. The default role is **NSX Manager**.

- To install an NSX Manager appliance, select the **NSX Manager** role.
- To install a Global Manager appliance for a NSX Federation deployment, select the **NSX Global Manager** role.

See [Chapter 14 Getting Started with NSX Federation](#) for details.

- To install a Cloud Service Manager (CSM) appliance for an NSX Cloud deployment, select the **nsx-cloud-service-manager** role.

See [Overview of Deploying NSX Cloud](#) for details.

- 18 Enter a default gateway, management network IP address (required), and management network netmask (required).

Note Entering a default gateway is optional. However, you cannot configure it after deploying NSX Manager.

- 19 In the DNS section, enter DNS Server list and Domain Search list.

- 20 In the Services Configuration section, enter NTP Server IP or FQDN.

Optionally, you can enable SSH service and allow root SSH login. But, it is not recommended to allow root access to SSH service.

- 21 Verify that all your custom OVF template specification is accurate and click **Finish** to begin installation.

The installation might take 7-8 minutes.

- 22 From the vSphere Client, verify that the VM is powered on. Open the VM console to track the boot process of the node.

- 23 After the VM node boots a second time, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.

- 24 Enter the `get services` command after waiting for about 5 minutes to verify that all default services are running.

The following services are not required by default and do not start automatically.

- `liagent`
- `migration-coordinator`: This service is used only when running migration coordinator. See the *NSX Migration Guide* before starting this service.
- `snmp`: For information on starting SNMP see *Simple Network Management Protocol* in the *NSX Administration Guide*.
- `nsx-message-bus`: This service is not used in NSX 3.0 and later releases.

- 25 After deployment, verify that the NSX Manager UI comes up by accessing the following URL, `https://nsx-manager-ip` OR `https://nsx-manager-fqdn`.

- 26 Verify that your NSX Manager, or Global Manager node has the required connectivity.

Perform the following tasks:

- Ping your node from another machine.
- The node can ping its default gateway.
- The node can ping the hypervisor hosts that are in the same network using the management interface.
- The node can ping its DNS server and its NTP Server or FQDN.

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

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

27 Troubleshooting OVA failures:

Note During deployment, if you entered incorrect configuration details, delete the appliance and redeploy with correct configuration.

- Verify that the datastore chosen for deployment is mounted on all the hosts that are members of a cluster. Redeploy and choose ESXi host instead of VMware vCenter to bypass VMware vCenter cluster related checks.
- If proxy enabled on VMware vCenter, edit file `/etc/sysconfig/proxy` and add line `.*.domainname` to bypass proxy for ESXi hosts. See, <https://kb.vmware.com/s/article/81565>.
- If deployment of appliance through OVF tool gives error `ovf descriptor not found`, view the file contents in terminal `cat -A <filepath/filename>` and remove hidden formatting characters. Then try again.

28 Troubleshooting issues related to bringing up the appliance. SSH to NSX Manager CLI as admin and run followign commands to troubleshoot:

- Run `get configuration` and verify `hostname/name-server/search-domain/ntp` settings are correct.
- Run `get services` and verify all required services are running (other than `nsx-message-bus`, `snmp`, `migration-coordinator`). If these services are not running, try restarting the service by running `restart service <service-name>`.
- Run `get cluster status` and verify all manager cluster components are up. If any component is down, try restarting the service associated to the component by running `restart service <associated-component-service-name>`.
- Run `get core-dumps` to verify no core dumps generated in `/var/log/core` or `/image/core`. If you find any core dumps, contact VMware Support.
- Run `get filesystem-status` to verify that no disk partition is full, especially those partitions that are consumed by NSX.
- Alternatively, you can run API commands to know the node and service status.

```
GET api/v1/node/status
```

```
GET api/v1/node/services
```

What to do next

- Log in to the NSX Manager from a supported web browser. See [Log In to the Newly Created NSX Manager](#) .

- [Deploy NSX Manager Nodes to Form a Cluster from the UI.](#)

Log In to the Newly Created NSX Manager

After you install NSX Manager, log in to the NSX Manager UI to perform additional tasks.

After you install NSX Manager, you can join the Customer Experience Improvement Program (CEIP) for NSX. See [Customer Experience Improvement Program](#) in the *NSX Administration Guide* for more information about the program, including how to join or leave the program 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>> or <https://<nsx-manager-fqdn>>.
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**.
- 5 Go to **System** → **Licenses** and click **Add** to add your NSX license.

Add a Compute Manager

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

NSX polls compute managers to collect cluster information from VMware vCenter.

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

Prerequisites

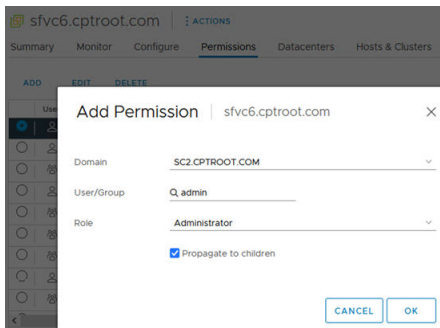
- Verify that you use the supported vSphere version. See [Supported vSphere version](#).
- IPv4 communication with VMware vCenter.
- Verify that you use the recommended number of compute managers. See <https://configmax.vmware.com/home>.

- Decide the hashing algorithm type you want to use for stamping NSX Manager thumbprint in compute manager extension. SHA1 and SHA256 algorithm types are supported. The default is SHA1. If you use SHA256 there might be communication issues between WCP component in VC and NSX Manager.

- To set the hashing algorithm, run API PUT `https://<nsx-mgr>/api/v1/fabric/compute-managers/thumbprint-hashing-algorithm`

```
{
  "hashing_algorithm_type": "SHA1"
}
```

- Provide credentials of a VMware vCenter user. You can provide the credentials of VMware vCenter administrator, or create a role and a user specifically for NSX and provide this user's credentials. Go to the **Administration > Global Permissions** tab. Add global permissions to the newly created user and role and select **Propagate to Children**.



Create an admin role with the following VMware vCenter privileges:

Global	Cancel task
Extension	Register extension
Extension	Unregister extension
Extension	Update extension
Host	Configuration.Maintenance
Host	Configuration.NetworkConfiguration
Host	Local Operations.Create virtual machine
Host	Local Operations.Delete virtual machine
Host	Local Operations.Reconfigure virtual machine
Network	Assign network
Permissions	Reassign role permissions
Resource	Assign vApp to resource pool
Resource	Assign virtual machine to resource pool
Sessions	Message
Sessions	Validate session

Sessions	View and stop sessions
Scheduled task	Select all privileges
Tasks	Select all privileges
vApp	Select all privileges
Virtual Machine.	Configuration
Virtual Machine	Guest Operations
Virtual Machine	Provisioning
Virtual Machine	Inventory

To use the NSX license for the vSphere Distributed Switch 7.0 feature, the VMware vCenter user must either be an administrator, or the user must have *Global.Licenses* privileges and be a member of the *LicenseService.Administrators* group.

- Before you create a service account for the compute manager, add these additional VMware vCenter privileges to the admin user role:

Permissions	Modify permission
Permissions	Modify role
Service Account Management	Administer
VMware vSphere Lifecycle Manager	ESXi Health Perspectives.Read
VMware vSphere Lifecycle Manager	Lifecycle Manager: General Privileges.Read
VMware vSphere Lifecycle Manager	Lifecycle Manager: Image Privileges.Read
VMware vSphere Lifecycle Manager	Lifecycle Manager: Image Privileges.Write
VMware vSphere Lifecycle Manager	Lifecycle Manager: Image Remediation Privileges.Write
VMware vSphere Lifecycle Manager	Lifecycle Manager: Settings Privileges.Read
VMware vSphere Lifecycle Manager	Lifecycle Manager: Settings Privileges.Write
VMware vSphere Lifecycle Manager	Lifecycle Manager: General Privileges.Write

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System > Fabric > Compute Managers > Add Compute Manager**.
- 3 Complete the compute manager details.

Option	Description
Name and Description	Type the name to identify the VMware vCenter. You can optionally describe any special details such as, the number of clusters in the VMware vCenter.
Type	The default compute manager type is set to VMware vCenter.

Option	Description
Multi NSX	<p>Starting with NSX 3.2.2, you can register the same vCenter Server with multiple NSX Managers.</p> <p>Enable this field if you want to allow multiple NSX instances to manage a single VMware vCenter. This functionality is supported from VMware vCenter 7.0 or later versions.</p> <hr/> <p>Note Cannot be enabled on a Workload Control Plane (WCP) cluster or vSphere Lifecycle Manager (vLCM) cluster.</p> <hr/> <p>See Multiple NSX Managers Managing a Single VMware vCenter.</p>
FQDN or IP Address	<p>Type the FQDN or IP address of the VMware vCenter.</p> <hr/> <p>Note If you plan to deploy NSX Manager in dual stack mode (IPv4 and IPv6) and if you plan to configure NSX Manager with CA signed certificates, you must set a FQDN with valid domain name.</p>
HTTPS Port of Reverse Proxy	<p>The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances.</p> <p>Set the reverse proxy port to register the compute manager in NSX.</p>
Username and Password	<p>Type the VMware vCenter login credentials.</p>
SHA-256 Thumbprint	<p>(Optional) Type the VMware vCenter SHA-256 thumbprint algorithm value. If you configured the VMware vCenter WCP (Workload Control Plane) feature, using the SHA256 setting results in communication issues between the WCP component in VMware vCenter and NSX Manager. In such cases, use the SHA1 algorithm instead.</p>
Create Service Account	<p>(Optional) Enable this field for features such as vSphere Lifecycle Manager that need to authenticate with NSX APIs. Log in with the administrator@vsphere.local credential to register a compute manager. After registration, the compute manager creates a service account.</p> <hr/> <p>Note Service account creation is not supported on a global NSX Manager.</p> <p>If service account creation fails, the compute manager's registration status is set to <code>Registered with errors</code>. The compute manager is successfully registered. However, vSphere Lifecycle Manager cannot be enabled on NSX clusters.</p> <p>If a VMware vCenter admin deletes the service account after it was successfully created, vSphere Lifecycle Manager tries to authenticate the NSX APIs and the compute manager's registration status is set to <code>Registered with errors</code>.</p>

Option	Description
Enable Trust	(Optional) Enable this field to establish trust between NSX and compute manager, so that services running in vCenter Server can establish trusted communication with NSX. For vSphere Lifecycle Manager to be enabled on NSX clusters, you must enable the Enable Trust field. Supported only on VMware vCenter 7.0 and later versions.
Access Level	Enable one of the options based on your requirement: <ul style="list-style-type: none"> ■ Full Access to NSX: Is selected by default. This access level gives the compute manager complete access to NSX. Full access ensures vSphere for Kubernetes and vSphere Lifecycle Manager can communicate with NSX. The VMware vCenter user's role must be set to an Enterprise Admin. ■ Limited Access to NSX: This access level ensures vSphere Lifecycle Manager can communicate with NSX. The VMware vCenter user's role must be set to Limited vSphere Admin.

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 to discover and register the VMware vCenter resources.

Note If the FQDN, IP, or thumbprint of the compute manager changes after registration, edit the computer manager and enter the new values.

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 VMware vCenter credentials and click **Resolve**.

If an existing registration exists, it will be replaced.

Results

It takes some time to register the compute manager with VMware vCenter 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 VMware vCenter 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 VMware vCenter again. You will get the error that the VMware vCenter is already registered with another NSX Manager.

Note After a vCenter Server (VC) compute manager is successfully added, it cannot be removed if you successfully performed any of the following actions:

- Transport nodes are prepared using VDS that is dependent on the VC.
- Service VMs deployed on a host or a cluster in the VC using NSX service insertion.
- You use the NSX Manager UI to deploy Edge VMs or NSX Manager nodes on a host or a cluster in the VC.

If you try to perform any of these actions and you encounter an error (for example, installation failed), you can remove the VC if you have not successfully performed any of the actions listed above.

If you have successfully prepared any transport node using VDS that is dependent on the VC or deployed any VM, you can remove the VC after you have done the following:

- Unprepare all transport nodes. If uninstalling a transport node fails, you must force delete the transport node.
- Undeploy all service VMs, all NSX Edge VMs, and all NSX Manager nodes. The undeployment must be successful or in a failed state.
- If an NSX Manager cluster consists of nodes deployed from the VC (manual method) and nodes deployed from the NSX Manager UI, and you had to undeploy the manually deployed nodes, then you cannot remove the VC. To successfully remove the VC, ensure that you re-deploy an NSX Manager node from the VC.

This restriction applies to a fresh installation of NSX as well as an upgrade.

Deploy NSX Manager Nodes to Form a Cluster from the UI

Forming an NSX Manager or Global Manager cluster provides high availability and reliability. Deploying nodes using the UI is supported only on ESXi hosts managed by VMware vCenter that is added as a compute manager.

For deploying additional NSX Manager nodes on a VMware vCenter that is not added as a compute manager, see [Install NSX Manager on ESXi Using the Command-Line OVF Tool and Form an NSX Manager Cluster Using the CLI](#).

When you deploy a new node from the UI, the node connects to the first deployed node to form a cluster. All the repository details and the password of the first deployed node are synchronized with the newly deployed node. The first node is known as the orchestrator node because it contains the original copy of the VIBs and installation files required to prepare hosts of the cluster. The orchestrator node also help identify the node on which the Upgrade-Coordinator is running. When new nodes are added to the cluster, NSX uses the repository IP to synchronize the repository of VIBs and installation files on the new nodes of the cluster.

To create an NSX Manager cluster, deploy two additional nodes to form a cluster of three nodes total.

Note Data is replicated to all the active NSX Manager nodes of the cluster. So, when the NSX Manager cluster is stable, every NSX Manager node contains the same data.

To create a Global Manager cluster, deploy two additional nodes to form a cluster of three nodes total. However, if your Global Manager has NSX 3.0.0 installed, deploy only one node, and do not form a cluster. See [Install the Active and Standby Global Manager](#) .

Prerequisites

- Verify that an NSX Manager or Global 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 or FQDN for the NSX Manager to use.
- Create a management VDS and target VM port group in vCenter. Place the NSX appliances onto this management VDS port group network. See [Prepare a vSphere Distributed Switch for NSX](#).

Multiple management networks can be used as long as the NSX Manager nodes has consistent connectivity and recommended latency between them.

Note If you plan to use Cluster VIP, all NSX Manager appliances should belong to same subnet.

Procedure

- 1 From a browser, log in with admin privileges to the NSX Manager or Global Manager at `https://<manager-ip-address>`.
- 2 Deploy an appliance.
 - From NSX Manager, select **System > Appliances > NSX Manager > Add NSX Appliance**.

- From Global Manager, select **System > Global Manager Appliances > Add NSX Appliance**.

3 Enter the appliance information details.

Option	Description
Host Name or FQDN	Enter a name for the node. Note If you plan to deploy NSX Manager in dual stack mode (IPv4 and IPv6) and if you plan to configure NSX Manager with CA signed certificates, you must set a FQDN with valid domain name.
IP Type	Select the IP type. The appliance can have IPv4 address only or both IPv4 and IPv6 addresses.
Management IPv4/Netmask	Enter an IPv4 address to be assigned to the node.
Management Gateway IPv4	Enter a gateway IPv4 address to be used by the node.
Management IPv6/Netmask	Enter an IPv6 address to be assigned to the node. This option appears when IP Type is set to Both IPv4 and IPv6 .
Management Gateway IPv6	Enter a gateway IPv4 address to be used by the node. This option appears when IP Type is set to Both IPv4 and IPv6 .
DNS Servers	Enter DNS server IP addresses to be used by the node.
NTP Server	Enter an NTP server IP address to be used by the node.
Node Size	Select the form factor to deploy the node from the following options: <ul style="list-style-type: none"> ■ Small (4 vCPU, 16 GB RAM, 300 GB storage) ■ Medium (6 vCPU, 24 GB RAM, 300 GB storage) ■ Large (12 vCPU, 48 GB RAM, 300 GB storage) For Global Manager select size: <ul style="list-style-type: none"> ■ Medium GM appliance for deployments up to four locations and 128 hypervisors across all locations ■ Large GM appliance for deployments with higher scale Do not use Small GM appliance for scale deployment.

4 Enter the configuration details.

Option	Description
Compute Manager	Select the VMware vCenter to provision compute resources for deploying the node.
Compute Cluster	Select the cluster the node is going to join.
(Optional) Resource Pool	Select a resource pool for the node from the drop-down menu.
(Optional) Host	Select a host for the node from the drop-down menu.
Datastore	Select a datastore for the node files from the drop-down menu.

Option	Description
Virtual Disk Format	<ul style="list-style-type: none"> ■ For NFS datastores, select a virtual disk format from the available provisioned policies on the underlying datastore. <ul style="list-style-type: none"> ■ With hardware acceleration, Thin Provision, Thick Provision Lazy Zeroed, and Thick Provision Eager Zeroed formats are supported. ■ Without hardware acceleration, only Thin Provision format is supported. ■ For VMFS datastores, Thin Provision, Thick Provision Lazy Zeroed, and Thick Provision Eager Zeroed formats are supported. ■ For vSAN datastores, you cannot select a virtual disk format because the VM storage policy defines the format. <ul style="list-style-type: none"> ■ The vSAN storage policies determine the disk format. The default virtual disk format for vSAN is Thin Provision. You can change the vSAN storage policies to set a percentage of the virtual disk that must be thick-provisioned. <p>By default, the virtual disk for an NSX Manager or Global Manager node is prepared in the Thin Provision format.</p> <p>Note You can provision each node with a different disk format based on which policies are provisioned on the datastore.</p>
Network	Click Select Network to select the management network for the node.

5 Enter the access and credentials details.

Option	Description
Enable SSH	Toggle the button to allow an SSH login to the new node.
Enable Root Access	Toggle the button to allow root access to the new node.

Option	Description
System Root Credentials	<p>Set the root password and confirm the password for the new node. Your password must comply with the password strength restrictions.</p> <ul style="list-style-type: none"> ■ 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 ■ Default password complexity rules are enforced by the following Linux PAM module arguments: <ul style="list-style-type: none"> ■ <code>retry=3</code>: The maximum number of times a new password can be entered, for this argument at the most 3 times, before returning with an error. ■ <code>minlen=12</code>: The minimum acceptable size for the new password. In addition to the number of characters in the new password, credit (of +1 in length) is given for each different kind of character (other, upper, lower and digit). ■ <code>difok=0</code>: The minimum number of bytes that must be different in the new password. Indicates similarity between the old and new password. With a value 0 assigned to <code>difok</code>, there is no requirement for any byte of the old and new password to be different. An exact match is allowed. ■ <code>lcredit=1</code>: The maximum credit for having lower case letters in the new password. If you have less than or 1 lower case letter, each letter will count +1 towards meeting the current <code>minlen</code> value. ■ <code>ucredit=1</code>: The maximum credit for having upper case letters in the new password. If you have less than or 1 upper case letter each letter will count +1 towards meeting the current <code>minlen</code> value. ■ <code>dcredit=1</code>: The maximum credit for having digits in the new password. If you have less than or 1 digit, each digit will count +1 towards meeting the current <code>minlen</code> value. ■ <code>ocredit=1</code>: The maximum credit for having other characters in the new password. If you have less than or 1 other characters, each character will count +1 towards meeting the current <code>minlen</code> value. ■ <code>enforce_for_root</code>: The password is set for the root user. <p>Note For more details on Linux PAM module to check the password against dictionary words, refer to the man page.</p> <p>For example, avoid simple and systematic passwords such as <code>VMware123!123</code> or <code>VMware12345</code>. Passwords that meet complexity standards are not simple and systematic but are a combination of letters, alphabets, special characters, and numbers, such as <code>VMware123!45, VMware 1!2345</code> or <code>VMware@1az23x</code>.</p>
Admin CLI Credentials and Audit CLI Credentials	<p>Select the Same as root password check box to use the same password that you configured for root, or deselect the check box and set a different password.</p>

6 Click **Install Appliance**.

The new node is deployed. You can track the deployment process in the **System > Appliances** page for NSX Manager, the **System > Global Manager Appliances** for Global Manager, or the VMware vCenter for either. Do not add additional nodes until the installation is finished and the cluster is stable.

7 Wait for the deployment, cluster formation, and repository synchronization to finish.

The joining and cluster stabilizing process might take from 10 to 15 minutes. After the node boots (as part of appliance bringup) and is back up, log in to the first deployed NSX Manager CLI as an admin. Run `get cluster status` to view the status. Verify that the first node and second node are members of the cluster and the status for every cluster service group is `UP` and cluster status `STABLE` before making any other cluster changes.

Note

- If you reboot the first deployed NSX Manager node, when the deployment of a new node is in progress, the new node might fail to register with the cluster. It displays the `Failed to Register` message on the new node's thumbnail. To redeploy the node manually on the cluster, delete and redeploy the failed node again.
 - If a node deployment fails, you cannot reuse the same IP address to deploy another node until the failed node is deleted.
 - NSX Manager nodes that are removed from the cluster should be powered off or deleted. Do not reuse it in your deployments.
-

8 After the VM node boots a second time, 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 Verify that your NSX Manager, or Global Manager node has the required connectivity.

Perform the following tasks:

- Ping your node from another machine.
- The node can ping its default gateway.
- The node can ping the hypervisor hosts that are in the same network using the management interface.
- The node can ping its DNS server and its NTP Server or FQDN.
- If you enabled SSH, make sure that you can SSH to your node.

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

10 If your cluster has only two nodes, add another appliance.

- From NSX Manager, select **System > Appliances > NSX Manager > Add NSX Appliance** and repeat the configuration steps.

- From Global Manager, select **System > Global Manager Appliances > Add NSX Appliance** and repeat the configuration steps.

Important If the orchestrator node (first NSX Manager node deployed) goes down or becomes unreachable while additional nodes in the cluster have not finished replicating with first node fully, further operations on the cluster will fail. To avoid this, bring up the first node up successfully and then delete and redeploy the additional nodes that had failed to finish replication of data with the first node.

Note NSX Manager nodes that are removed from the cluster should be powered off or deleted. Do not reuse it in your deployments.

- 11 Before adding or removing an existing node from the cluster, place all NSX Intelligence and NDR SaaS agents in maintenance mode.

```
PUT https://nsx-manager-ip/policy/api/v1/infra/sites/agents/intelligence/
maintenance
```

```
{
  "enable": true
}
```

- 12 After adding or removing nodes from the cluster, take all the NSX Intelligence and NDR SaaS agents out of maintenance mode.

```
PUT https://nsx-manager-ip/policy/api/v1/infra/sites/agents/intelligence/
maintenance
```

```
{
  "enable": false
}
```

What to do next

Configure NSX Edge. See [Install an NSX Edge on ESXi Using the vSphere GUI](#).

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

- You can download the latest OVF tool from the [Broadcom Support](#) page.
- 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 or FQDN for the NSX Manager to use.
- Create a management VDS and target VM port group in vCenter. Place the NSX appliances onto this management VDS port group network. See [Prepare a vSphere Distributed Switch for NSX](#).

Multiple management networks can be used as long as the NSX Manager nodes has consistent connectivity and recommended latency between them.

Note If you plan to use Cluster VIP, all NSX Manager appliances should belong to same subnet.

- Plan your NSX Manager IP and NSX Manager Cluster VIP addressing scheme.

Note Verify that you have the hostname for NSX Manager to use. The Hostname format must be `nsx-manager-fqdn@domain-name.com`. This format is required if NSX installation is dual stack (IPv4, IPv6) and/or if planning to configure CA-signed certificates.

Procedure

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

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

- For a standalone host:

Note On a standalone host, if you enter an incorrect role in the `nsx_role` property, then the appliance is deployed in the NSX Manager role.

- Windows example:

```
C:\Program Files\VMware\VMware OVF Tool>ovftool \
--name=<nsxmanager>
--X:injectOvfEnv
--X:logFile=ovftool.log
--sourceType=OVA
--vmFolder='Folder-in-VC'
--allowExtraConfig
--datastore=<datastore>
--net:"<network-name-of-OVF>=<network-name>"
--acceptAllEulas
--skipManifestCheck
--noSSLVerify
--diskMode=thin
--quiet
--hideEula
--powerOn
```

```

--prop:nsx_ip_0=10.196.176.81
--prop:nsx_netmask_0=255.255.252.0
--prop:nsx_gateway_0=10.196.179.253
--prop:nsx_dns1_0=10.142.7.1
--prop:nsx_domain_0=eng.vmware.com
--prop:nsx_ntp_0=10.128.243.14
--prop:nsx_isSSHEnabled=True
--prop:"nsx_passwd_0=<password>"
--prop:"nsx_cli_passwd_0=<password-cli>"
--prop:"nsx_cli_audit_passwd_0=<password-cli-audit>"
--prop:nsx_hostname=<hostname>
--prop:mgrhostname01="nsx-manager-02@vmware.com"
--prop:nsx_allowSSHRootLogin=True
--prop:nsx_role="NSX Manager"
--X:logFile=/root/ovftool/<ovf-folder>.log
--X:logLevel=trivia
--ipProtocol=IPv4
--ipAllocationPolicy="fixedPolicy"
<nsx-unified-appliance>.ova
'vi://admin@vSphere.local:password@vc-or-ip.eng.vmware.com/<datacenter>/host/
Install/10.196.6.78/'

```

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.

Note In the above example, --deploymentOption is set to the default size Medium. To know the other supported sizes, see [NSX Manager VM and Host Transport Node System Requirements](#).

- 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
--name=<nsxmanager>
--X:injectOvfEnv
--X:logFile=ovftool.log
--sourceType=OVA
--vmFolder='Folder-in-VC'
--allowExtraConfig
--datastore=<datastore>
--net:"<network-name-of-OVF>=<network-name>"
--acceptAllEulas
--skipManifestCheck
--noSSLVerify
--diskMode=thin
--quiet
--hideEula
--powerOn
--prop:nsx_ip_0=10.196.176.81
--prop:nsx_netmask_0=255.255.252.0
--prop:nsx_gateway_0=10.196.179.253
--prop:nsx_dns1_0=10.142.7.1
--prop:nsx_domain_0=eng.vmware.com
--prop:nsx_ntp_0=10.128.243.14
--prop:nsx_isSSHEnabled=True
--prop:"nsx_passwd_0=<password>"
--prop:"nsx_cli_passwd_0=<password-cli>"
--prop:"nsx_cli_audit_passwd_0=<password-cli-audit>"
--prop:nsx_hostname=<hostname>
--prop:mgrhostname01="nsx-manager-02@vmware.com"
--prop:nsx_allowSSHRootLogin=True
--prop:nsx_role="NSX Manager"
--X:logFile=/root/ovftool/<ovf-folder>.log
--X:logLevel=trivia
--ipProtocol=IPv4
--ipAllocationPolicy="fixedPolicy"
<nsx-unified-appliance>.ova
'vi://admin@vsphere.local:password@vc-or-ip.eng.vmware.com/<datacenter>/host/
Install/10.196.6.78/'

```

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>@<esxi-IP-address>
Deploying to VI: vi://root:<password>@<esxi-IP-address>
Transfer Completed
Powering on VM: NSX Manager
Task Completed
Completed successfully

```

- For a host managed by VMware vCenter:
 - Windows example:

```
C:\Users\Administrator\Downloads>ovftool
--name=<nsxmanager>
--X:injectOvfEnv
--X:logFile=ovftool.log
--sourceType=OVA
--vmFolder='Folder-in-VC'
--allowExtraConfig
--datastore=<datastore>
--net:"<network-name-of-OVF>=<network-name>"
--acceptAllEulas
--skipManifestCheck
--noSSLVerify
--diskMode=thin
--quiet
--hideEula
--powerOn
--prop:nsx_ip_0=10.196.176.81
--prop:nsx_netmask_0=255.255.252.0
--prop:nsx_gateway_0=10.196.179.253
--prop:nsx_dns1_0=10.142.7.1
--prop:nsx_domain_0=eng.vmware.com
--prop:nsx_ntp_0=10.128.243.14
--prop:nsx_isSSHEnabled=True
--prop:"nsx_passwd_0=<password>"
--prop:"nsx_cli_passwd_0=<password-cli>"
--prop:"nsx_cli_audit_passwd_0=<password-cli-audit>"
--prop:nsx_hostname=<hostname>
--prop:mgrhostname01="nsx-manager-02@vmware.com"
--prop:nsx_allowSSHRootLogin=True
--prop:nsx_role="NSX Manager"
--X:logFile=/root/ovftool/<ovf-folder>.log
--X:logLevel=trivia
--ipProtocol=IPv4
--ipAllocationPolicy="fixedPolicy" <nsx-unified-appliance>.ova
'vi://admin@vsphere.local:password@vc-or-ip.eng.vmware.com/<datacenter>/host/
Install/10.196.6.78/'
```

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, --deploymentOption is set to the default size Medium. To know the other supported sizes, see [NSX Manager VM and Host Transport Node System Requirements](#).

- 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"

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

ovftool
--name=<nsxmanager>
--X:injectOvfEnv
--X:logFile=ovftool.log
--sourceType=OVA
--vmFolder='Folder-in-VC'
--allowExtraConfig
--datastore=<datastore>
--net:"<network-name-of-OVF>=<network-name>"
--acceptAllEulas
--skipManifestCheck
--noSSLVerify
--diskMode=thin
--quiet
--hideEula
--powerOn
--prop:nsx_ip_0=10.196.176.81
--prop:nsx_netmask_0=255.255.252.0
--prop:nsx_gateway_0=10.196.179.253
--prop:nsx_dns1_0=10.142.7.1
--prop:nsx_domain_0=eng.vmware.com
--prop:nsx_ntp_0=10.128.243.14
--prop:nsx_isSSHEnabled=True
--prop:"nsx_passwd_0=<password>"
--prop:"nsx_cli_passwd_0=<password-cli>"
--prop:"nsx_cli_audit_passwd_0=<password-cli-audit>"
--prop:nsx_hostname=<hostname>
--prop:mgrhostname01="nsx-manager-02@vmware.com"
--prop:nsx_allowSSHRootLogin=True
--prop:nsx_role="NSX Manager"

```

```
--X:logfile=/root/ovftool/<ovf-folder>.log
--X:LogLevel=trivia
--ipProtocol=IPv4
--ipAllocationPolicy="fixedPolicy" <nsx-unified-appliance>.ova
'vi://admin@vsphere.local:password@vc-or-ip.eng.vmware.com/<datacenter>/host/
Install/10.196.6.78/'
```

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@<esxi-IP-address:port>/
Deploying to VI: vi://administrator@vsphere.local@<esxi-IP-address:port>/
Transfer Completed
Powering on VM: NSX Manager
Task Completed
Completed successfully
```

- 2 You can also run the OVF tool in Probe mode to view contents of a source. OVA and OVF packages can be probed among a list of other supported source types. You can use the information returned by the Probe mode to configure deployments.

```
$> \ovftool --allowExtraConfig <OVA path or URL>
```

Where, `--allowExtraConfig` is the supported appliance type for Cloud Service Manager (CSM).

- 3 For an optimal performance, reserve memory for the appliance.
Set the reservation to ensure that NSX Manager has sufficient memory to run efficiently. See [NSX Manager VM and Host Transport Node System Requirements](#).
- 4 From the vSphere Client, verify that the VM is powered on. Open the VM console to track the boot process of the node.
- 5 After the VM node boots a second time, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.
- 6 Enter the `get services` command after waiting for about 5 minutes to verify that all default services are running.

The following services are not required by default and do not start automatically.

- `liagent`
- `migration-coordinator`: This service is used only when running migration coordinator. See the *NSX Migration Guide* before starting this service.
- `snmp`: For information on starting SNMP see *Simple Network Management Protocol* in the *NSX Administration Guide*.
- `nsx-message-bus`: This service is not used in NSX 3.0 and later releases.

- 7 After deployment, verify that the NSX Manager NSX Manager UI comes up by accessing the following URL, <https://nsx-manager-ip> OR <https://nsx-manager-fqdn>.
- 8 Verify that your NSX Manager, or Global Manager node has the required connectivity.

Perform the following tasks:

- Ping your node from another machine.
- The node can ping its default gateway.
- The node can ping the hypervisor hosts that are in the same network using the management interface.
- The node can ping its DNS server and its NTP Server or FQDN.
- If you enabled SSH, make sure that you can SSH to your node.

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](#) .

Note NSX Manager nodes that are removed from the cluster should be powered off or deleted. Do not reuse the same NSX Manager again in your environment.

- If deploying second and third NSX Manager nodes as OVA/OVF, join the manager nodes to first deployed manager node to create NSX Manager Cluster. see [Form an NSX Manager Cluster Using the CLI](#).

Form an NSX Manager Cluster Using the CLI

Forming an NSX Manager or Global Manager cluster provides high availability and reliability. If all or at least one of the three NSX Manager appliances is deployed as an OVA/OVF using vCenter, use the `join` command to join the NSX Manager nodes and create a cluster.

Prerequisites

- To create an NSX Manager cluster, deploy three NSX Manager nodes from the OVF tool CLI or VMware vCenter UI.
- To create a Global Manager cluster, deploy three nodes to create the cluster. However, if your Global Manager has NSX 3.0.0 installed, deploy only one node, and do not form a cluster. See [Install the Active and Standby Global Manager](#) .

Procedure

- 1 Open an SSH or console session to the first deployed NSX Manager or Global Manager node and log in with the administrator credentials.

- 2 On the first deployed node, run the following commands.
 - a Run the `get certificate api thumbprint` command.
The command output is a string that is unique to this node.
 - b Run the `get cluster config` command to get the cluster ID of the first deployed node.

```
mgr-first> get cluster config
Cluster Id: 7b50abb9-0402-4ed5-afec-363587c3c705
Cluster Configuration Version: 0
Number of nodes in the cluster: 1
...
```

- 3 Open an SSH or console session to the new node and log in with the administrator credentials.
- 4 On the new node that is joining the cluster, run the `join` command.
Provide the following information about the first deployed node in the `join` command:
 - IP address
 - Cluster ID
 - User name
 - Password
 - Certificate thumbprint

```
mgr-new> join <Manager-IP> cluster-id <cluster-id> username <Manager-username> password
<Manager-password> thumbprint <Manager-thumbprint>
```

The joining and cluster stabilizing process might take from 10 to 15 minutes. Run `get cluster status` to view the status. Verify that the status for every cluster service group is `UP` before making any other cluster changes.

- 5 Add the third node to the cluster.
Repeat step 4 on the third node.
- 6 Verify the cluster status on the web interface.
 - On NSX Manager, log in to the NSX Manager web interface and select **System > Appliances**.
 - On Global Manager, log in to the Global Manager web interface and select **System > Global Manager Appliances**.

What to do next

Create a transport zone. See [Create Transport Zones](#).

Verify NSX Manager Clustering

Verify whether individual nodes are listed as members of the NSX Manager cluster.

Log in to any one of the nodes and perform one of these steps to verify whether all nodes are listed as members of the NSX Manager cluster.

Prerequisites

Procedure

- 1 To verify NSX Manager cluster status from CLI, perform these steps:
 - a Log in as an admin to any one of the NSX Manager nodes.
 - b run 'get cluster status' to view the correct number of manager nodes are listed as members of the cluster and cluster status is STABLE.
- 2 To verify NSX Manager cluster status from UI, perform these steps:
 - a Go to **System** → **Appliances**. Verify the Cluster status is STABLE and a correct number of nodes are shown to be part of the cluster.
- 3 To verify NSX Manager cluster status from API, run these commands:
 - `GET /api/v1/cluster/status`
 - `GET api/v1/cluster/nodes` or `GET api/vi/cluster/<node-id>`

Configure an Appliance to Display the GRUB Menu at Boot Time

You must configure an NSX appliance to display the GRUB menu at boot time if you want to reset the root password of the appliance.

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

Procedure

- 1 Log in to the VM as root.
- 2 In the `/etc/default/grub` file, set the `GRUB_TIMEOUT_STYLE` to **menu** or **countdown**.
 - If this option set to **menu**, then GRUB will display the menu and then wait for the timeout set by `GRUB_TIMEOUT` to expire before booting the default entry. Pressing a key interrupts the timeout.
 - If this option is set to **countdown**, then before displaying the menu, GRUB will wait for the timeout set by `GRUB_TIMEOUT` to expire. If ESC or F4 are pressed, or SHIFT is held down during that time, it will display the menu and wait for input. It will show a one-line indication of the remaining time.

- 3 In the `/etc/default/grub` file, change the value for the parameter `GRUB_TIMEOUT`.

```
GRUB_TIMEOUT=4
```

- 4 (Optional) Generate a new password by running the following command:

```
grub-mkpasswd-pbkdf2
```

- 5 (Optional) In the `/etc/grub.d/40_custom` file, replace the existing GRUB password.

The default password is **NSX@VM!WaR10**.

- 6 Update the GRUB configuration.

```
update-grub
```

Configure GRUB Menu Using CLI or API

You must configure an NSX appliance to display the GRUB menu at boot time if you want to reset the root password of the appliance.

Starting with NSX 4.0.1.1, you can use CLI or API commands to set the GRUB timeout value and password. You can follow these commands post deployment of NSX.

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

Procedure

- 1 Using CLI to set GRUB menu:

- a Log in to the NSX command line interface.
- b Run `set grub menu timeout <value>`.

Where `<value>` is time in seconds. The default timeout value is 4.

- c Run `set grub user root password <newpassword>`.

OR

- d Run `set grub user root password`

```
Enter password:<newpassword>
```

```
Confirm password:<newpassword>
```

2 Using API to set GRUB menu:

- a Use the GET API to retrieve GRUB menu values.

```
GET https://<nsx-mgr>/api/v1/node/grub
```

Example Response:

```
{
  "timeout": 4,
  "users": [
    {
      "username": "root"
    }
  ]
}
```

- b Set the GRUB timeout value.

```
PUT https://<nsx-mgr>/api/v1/node/grub { "timeout": 4 }
```

Example Response:

```
{
  "timeout": 4
}
```

- c Set the GRUB menu password.

```
PUT https://<nsx-mgr>/api/v1/node/grub/root { "password": "Str0ng_Pwd!Wins$" }
```

Example Response:

```
{
  "username": "root"
}
```

3 Get GRUB timeout value.

```
get grub menu timeout
```

```
GRUB Menu Timeout = 4
```

Configure a Virtual IP Address for a Cluster

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

NSX Manager nodes 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. But if external load balancer is used to configure cluster VIP, then NSX Managers and VIP can belong to a different subnet.

If the leader node that owns VIP becomes unavailable, NSX 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 the NSX Manager 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 Administration Guide* for more details.

Important If you reset the cluster VIP, then vIDM configurations that are using the VIP is cleared. You will need to reconfigure vIDM configuration with the new VIP.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Go to **System > Appliances**.
- 3 In the Virtual IP field, click **Set Virtual IP**.
- 4 Enter the IPv4 and/or IPv6 address to use as VIP for the cluster.

Ensure that VIP is part of the same subnet as the other management nodes. When you deploy a NSX Manager in a dual stack environment (IPv4 and IPv6), use a valid FQDN address and use the same FQDN address for both IPv4 and IPv6 addresses.

- 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:

```
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 Verify that the VIP is working correctly.

From a browser, log in to the NSX Manager using the virtual IP address assigned to the cluster at `https://<VIP-address>`.

Results

Any API requests to NSX are 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.

Note If you deploy NSX Manager in dual stack mode (IPv4, IPv6) and/or if you plan to configure NSX Manager deployment with CA-signed certificates, associate the VIP IP address with DNS name. Also, configure reverse and forward proxy for the VIP IP address on the DNS server. Then access VIP from `https://<VIP-dns-name>`.

Configuring an External Load Balancer

You can configure an external load balancer to distribute traffic to the NSX Managers in a manager cluster.

An NSX Manager cluster does not require an external load balancer. The NSX Manager virtual IP (VIP) provides resiliency in the event of a Manager node failure but has the following limitations:

- VIP does not perform load balancing across the NSX Managers.
- VIP requires all the NSX Managers to be in the same subnet.
- VIP recovery takes about 1 - 3 minutes in the event of a Manager node failure.

An external load balancer can provide the following benefits:

- Load balance across the NSX Managers.
- The NSX Managers can be in different subnets.
- Fast recovery time in the event of a Manager node failure.

An external load balancer will not work with the NSX Manager VIP. Do not configure an NSX Manager VIP if you use an external load balancer.

Authentication Methods When Accessing NSX Manager

The following authentication methods are supported by NSX Manager. For more information about the authentication methods, see the *NSX API Guide*.

- HTTP Basic Authentication

- Session-Based Authentication
- Authentication using an X.509 certificate and a Principal Identity
- Authentication in VMware Cloud on AWS

The session-based authentication method (used when you access NSX Manager from a browser) requires source-IP persistence (all requests from the client must go to the same NSX Manager). The other methods do not require source-IP persistence (requests from the client can go to different NSX Managers).

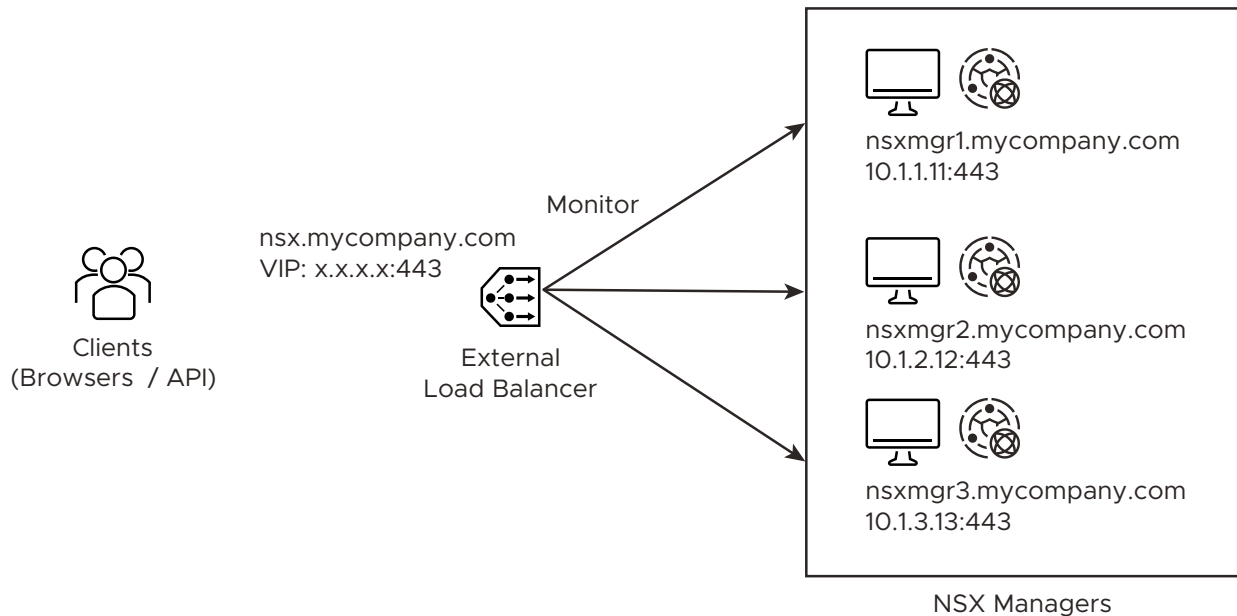
Recommendations

- Create a single VIP on the load balancer with source-IP persistence configured to handle all the authentication methods.
- If you have applications or scripts that might generate a lot of requests to NSX Manager, create a second VIP without source-IP persistence for these applications or scripts. Use the first VIP for browser access to NSX Manager only.

The VIP must have the following configurations:

- Type: Layer4-TCP
- Port: 443
- Pool: NSX Manager Pool
- Persistence: Source-IP persistence for the first VIP. None for the second VIP (if present).

Example of an external load balancer configuration:



NSX Manager's Certificate

The clients access NSX Manager using FQDN name (for example, nsx.mycompany.com). This FQDN is resolved to the load balancer's VIP. To avoid any certificate mismatch, each NSX Manager must have a certificate that is valid for the VIP's FQDN name. Therefore, you must configure each NSX Manager with a SAN certificate that is valid for its own name (for example, nsxmgr1.mycompany.com) and the VIP's FQDN.

Monitoring the Health of NSX Managers

The load balancer can check that each NSX Manager is running with the following API:

```
GET /api/v1/reverse-proxy/node/health
```

The request headers are:

- Header 1
 - Name: Content-Type
 - Value: application/json
- Header 2
 - Name: Accept
 - Value: application/json

A response indicating that the NSX Manager is running will be:

```
"healthy" : true
```

Note that the format of the response is "healthy"<space>:<space>true.

Verify Appliance Proxy Hub on all NSX Manager Nodes are Connected

Prerequisites

Appliance Proxy Hub (APH) acts as a communication channel between NSX Manager and a transport node. It runs as a service on NSX Manager and provides secure connection between a transport node and NSX Manager.

Before you set a virtual IP (VIP) address for the cluster nodes, ensure the APH running on each NSX Manager cluster nodes are connected to each other. If APH on all cluster nodes are not connected, you might face issues when configuring the VIP for the cluster.

Procedure

- 1 To verify APH on all the cluster nodes are connected, go to one of the NSX Manager nodes and run the `GET /api/v1/messaging/cluster-connection/status` call.

```
{
  "results": [
    {
      "address": "ssl-tcp://192.161.1.52:38522",
      "conn_status": "Connected",
      "node_id": "e85ceb93-df30-43fa-84d7-32d88f68a2ba",
      "node_type": "APPLIANCE_PROXY_HUB"
    },
    {
      "address": "ssl://192.161.1.51:1234",
      "conn_status": "Connected",
      "node_id": "46cc9a95-6194-4d84-94bd-0744fb46e225",
      "node_type": "APPLIANCE_PROXY_HUB"
    }
  ]
}
```

- 2 Repeat Step 1 on the remaining cluster nodes.
- 3 Wait for 60 seconds before you set the cluster VIP for the cluster.

Results

When APH on all cluster nodes are connected to each other, you can proceed to configure VIP for the cluster.

What to do next

Set up the cluster VIP address.

Installing and Configuring NSX Embedded using VMware vCenter Plugin

7

NSX embedded (NSXe) integrates NSX with VMware vCenter to enable the consumption of NSX from VMware vCenter. As a VI admin, you can install NSX Manager and NSX for virtual networking or security-only use case by installing and configuring the NSX plugin in VMware vCenter.

There are two workflows - Virtual Networking and Security Only - allowed from the NSX page on the vSphere Client. The Virtual Networking deployment workflow includes both networking and security use cases. In contrast, if you choose to configure Security Only type of deployment, then you cannot configure virtual networking on the selected cluster hosts.

Note

- You must not enable **Multi NSX** on the VMware vCenter on which you plan to install and configure the NSX plugin. It might result in unexpected results when using the NSX plugin from the VMware vCenter.
 - VMware does not support registration of multiple VMware vCenter servers to the same NSX Manager.
-

Supported Browser Resolution

The minimum supported browser resolution is 1280 x 800 px.

Read the following topics next:

- [Install NSX Manager from vSphere Client](#)
- [Install Additional NSX Manager Nodes to Form a Cluster from VMware vCenter Plugin](#)
- [Configure NSX for Virtual Networking from vSphere Client](#)
- [Configuring NSX-T for Security from vSphere Client](#)
- [Viewing NSX Alarms in vSphere Web Client UI](#)

Install NSX Manager from vSphere Client

As a VI admin working in the vSphere environment, you can completely install NSX Manager appliance from the vSphere Client. You do not need to perform any installation operations from

the NSX Manager UI. After NSX Manager is installed, NSX appears as a plug-in in VMware vCenter that is ready to install NSX for Virtual Networking or Security-only use cases.

Important In NSX 3.2, only a single NSX Manager cluster is supported.

Prerequisites

- Ensure that ESXi host version is compatible with VMware vCenter version v7.0.3.
- Ensure that VMware vCenter version is v7.0.3 or later.
- To provision a thick disk, ensure the disk size on host has at least 300GB free space.
- Configure a vSphere Distributed Switch (VDS) switch on hosts. Only VDS 6.6 or later is supported.
- Ensure VMware vCenter points to an FQDN address and the DNS server must be able to resolve the address.
- To ensure time is synchronized, configure NTP server on NSX Manager and ESXi hosts. See the Time Synchronization between NSX Manager, vIDM, and Related Components topic in the *NSX Administration Guide*.

Procedure

- 1 From a browser, log in with admin privileges to an VMware vCenter at `https://<vcenter-server-ip-address>`.
- 2 On the vSphere Client UI, select **vSphere Client** menu and click **NSX**.
- 3 On the screen, click **Install NSX**.
- 4 Enter the download OVF URL or navigate to the OVF file, and click **Next**.

Important If you enter a URL to download the OVF file, ensure the URL points to a secure HTTPS server. For example, `https://<OVF-URL>`. There is a separate OVF file available for NSX Manager deployed from vSphere Client. You must select a OVF file name using the following convention: `nsx-embedded-unified-appliance-<releaseversion.buildversion>.ova`. Do not use the `nsx-unified-appliance-<releaseversion.buildversion>.ova` file.

- 5 To verify the thumbprint of the SSL certificate of the HTTPS server, click **Yes**.
- 6 Enter a name and a location for the NSX Manager VM, and click **Next**.

The selected location also indicates the VMware vCenter where the NSX Manager is deployed and which VMware vCenter is managed by the NSX instance.

The name you enter appears in the vSphere and VMware vCenter inventory.

- 7 Select a compute resource for the NSX Manager appliance, and click **Next**.
- 8 Review and verify the OVF template details, and click **Next**.

- 9 Select a form factor to deploy the NSX appliance. You must deploy the NSX Manager in either **Medium** or **Large** form factor. If you select any other form factor, then installation fails and NSX appliance is not registered to VMware vCenter.
- 10 Specify storage for the configuration and disk files.
 - a Select the virtual disk format.
 - b Select the VM storage policy.
 - c Specify the datastore to store the NSX Manager appliance files.
 - d Click **Next**.
- 11 Select a destination network for each source network.
- 12 Select the port group or destination network for the NSX Manager.
- 13 Configure IP Allocation settings.
 - a For IP allocation, specify **Static - Manual**.
 - b For IP protocol, select **IPv4** or **IPv6**.

Note You can ignore the IP Allocation settings. You can select either IPv4 or IPv6. It would not impact ingress or egress network traffic of NSX Manager.

- 14 Click **Next**.
- 15 In the Application section, enter the **System Root User 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
 - Default password complexity rules are enforced by the following Linux PAM module arguments:
 - `retry=3`: The maximum number of times a new password can be entered, for this argument at the most 3 times, before returning with an error.
 - `minlen=12`: The minimum acceptable size for the new password. In addition to the number of characters in the new password, credit (of +1 in length) is given for each different kind of character (other, upper, lower and digit).
 - `difok=0`: The minimum number of bytes that must be different in the new password. Indicates similarity between the old and new password. With a value 0 assigned to `difok`, there is no requirement for any byte of the old and new password to be different. An exact match is allowed.

- `lcredit=1`: The maximum credit for having lower case letters in the new password. If you have less than or 1 lower case letter, each letter will count +1 towards meeting the current `minlen` value.
- `ucredit=1`: The maximum credit for having upper case letters in the new password. If you have less than or 1 upper case letter each letter will count +1 towards meeting the current `minlen` value.
- `dcredit=1`: The maximum credit for having digits in the new password. If you have less than or 1 digit, each digit will count +1 towards meeting the current `minlen` value.
- `ocredit=1`: The maximum credit for having other characters in the new password. If you have less than or 1 other characters, each character will count +1 towards meeting the current `minlen` value.
- `enforce_for_root`: The password is set for the root user.

Note For more details on Linux PAM module to check the password against dictionary words, refer to the man page.

For example, avoid simple and systematic passwords such as `VMware123!123` or `VMware12345`. Passwords that meet complexity standards are not simple and systematic but are a combination of letters, alphabets, special characters, and numbers, such as `VMware123!45`, `VMware 1!2345` or `VMware@1az23x`.

Important If the password you set does not meet the password complexity requirements (additionally, the password length must not exceed 128 characters), installation of the NSX Manager fails. If installation fails, you need to redeploy the NSX Manager again.

- 16 In the Network Properties section, 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 starts with an alphabet character. Also, NSX accepts only latin alphabets that do not have an accent mark, as in í, ó, ú, ý.

Important If you plan to install NSX in dual stack (IPv4 and IPv6) and/or if you plan to configure CA-signed certificates, then enter a Hostname with valid domain name to NSX Manager VMs and Cluster VIP (if configured).

- 17 Enter a default gateway, management network IP address (required), and management network netmask (required).
- 18 In the DNS section, enter DNS Server list and Domain Search list.
- 19 In the Services Configuration section, enter NTP Server IP or FQDN.

Optionally, you can enable SSH service and allow root SSH login. But, it is not recommended to allow root access to SSH service.

20 Verify that all your custom OVF template specification is accurate and click **Finish** to begin installation.

See the installation progress in the **Recent Tasks** tab.

21 On the NSX page, you can either click **Start NSX Onboarding** to load the plugin or skip the onboarding workflow and access the NSX Manager UI from the vSphere Client.

What to do next

Apply NSX license.

- 1 Click **Go To NSX Getting Started**.
- 2 In the **NSX License Key** section, enter the NSX license key and click **Apply**.

After you successfully apply the NSX license, configure NSX for Virtual Networking or Security use case on the vSphere platform. See [Configure NSX for Virtual Networking from vSphere Client](#).

Install Additional NSX Manager Nodes to Form a Cluster from VMware vCenter Plugin

Forming an NSX Manager cluster provides high availability and reliability of NSX management function if one of the NSX Manager goes down.

For other environments, see [Form an NSX Manager Cluster Using the CLI](#).

To create an NSX Manager cluster, deploy two additional nodes to form a cluster of three nodes total.

Note Data is replicated to all the active NSX Manager nodes of the cluster. So, when the NSX Manager cluster is stable, every NSX Manager node contains the same data.

Prerequisites

- Verify that an NSX Manager node is installed. See [Install NSX Manager from vSphere Client](#).
- 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 or FQDN for the NSX Manager to use.
- Create a management VDS and target VM port group in vCenter. Place the NSX appliances onto this management VDS port group network. See [Prepare a vSphere Distributed Switch for NSX](#).

Multiple management networks can be used as long as the NSX Manager nodes has consistent connectivity and recommended latency between them.

Note If you plan to use Cluster VIP, all NSX Manager appliances should belong to same subnet.

Procedure

- 1 From a browser, log in with admin privileges to an vCenter Server at `https://<vcenter-server-ip-address>`.
- 2 On the vSphere Web Client UI, select vSphere Web Client menu and click NSX.
- 3 Deploy an appliance. Go to **System > Appliances > NSX Manager > Add NSX Appliance**.
- 4 Enter the appliance information details.

Option	Description
Host Name or FQDN	Enter a name for the node.
IP Type	Select the IP type. The appliance can have IPv4 address only or both IPv4 and IPv6 addresses.
Management IPv4/Netmask	Enter an IPv4 address to be assigned to the node.
Management Gateway IPv4	Enter a gateway IPv4 address to be used by the node.
Management IPv6/Netmask	Enter an IPv6 address to be assigned to the node. This option appears when IP Type is set to Both IPv4 and IPv6 .
Management Gateway IPv6	Enter a gateway IPv4 address to be used by the node. This option appears when IP Type is set to Both IPv4 and IPv6 .
DNS Servers	Enter DNS server IP addresses to be used by the node.
NTP Server	Enter an NTP server IP address to be used by the node.
Node Size	Select the form factor to deploy the node from the following options: <ul style="list-style-type: none"> ■ Small (4 vCPU, 16 GB RAM, 300 GB storage) ■ Medium (6 vCPU, 24 GB RAM, 300 GB storage) ■ Large (12 vCPU, 48 GB RAM, 300 GB storage)

- 5 Enter the configuration details.

Option	Description
Compute Manager	Select the VMware vCenter to provision compute resources for deploying the node.
Compute Cluster	Select the cluster the node is going to join.
Resource Pool	Select either a resource pool or a host for the node from the drop-down menu.
Host	If you did not select a resource pool, select a host for the node.
Datastore	Select a datastore for the node files from the drop-down menu.

Option	Description
Virtual Disk Format	<ul style="list-style-type: none"> ■ For NFS datastores, select a virtual disk format from the available provisioned policies on the underlying datastore. <ul style="list-style-type: none"> ■ With hardware acceleration, Thin Provision, Thick Provision Lazy Zeroed, and Thick Provision Eager Zeroed formats are supported. ■ Without hardware acceleration, only Thin Provision format is supported. ■ For VMFS datastores, Thin Provision, Thick Provision Lazy Zeroed, and Thick Provision Eager Zeroed formats are supported. ■ ■ For vSAN datastores, you cannot select a virtual disk format because the VM storage policy defines the format. <ul style="list-style-type: none"> ■ The vSAN storage policies determine the disk format. The default virtual disk format for vSAN is Thin Provision. You can change the vSAN storage policies to set a percentage of the virtual disk that must be thick-provisioned. <p>By default, the virtual disk for an NSX Manager node is prepared in the Thin Provision format.</p> <hr/> <p>Note You can provision each node with a different disk format based on which policies are provisioned on the datastore.</p> <hr/>
Network	Click Select Network to select the management network for the node.

6 Enter the access and credentials details.

Option	Description
Enable SSH	Toggle the button to allow an SSH login to the new node.
Enable Root Access	Toggle the button to allow root access to the new node.

Option	Description
System Root Credentials	<p>Set the root password and confirm the password for the new node. Your password must comply with the password strength restrictions.</p> <ul style="list-style-type: none"> ■ 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 ■ Default password complexity rules are enforced by the following Linux PAM module arguments: <ul style="list-style-type: none"> ■ <code>retry=3</code>: The maximum number of times a new password can be entered, for this argument at the most 3 times, before returning with an error. ■ <code>minlen=12</code>: The minimum acceptable size for the new password. In addition to the number of characters in the new password, credit (of +1 in length) is given for each different kind of character (other, upper, lower and digit). ■ <code>difok=0</code>: The minimum number of bytes that must be different in the new password. Indicates similarity between the old and new password. With a value 0 assigned to <code>difok</code>, there is no requirement for any byte of the old and new password to be different. An exact match is allowed. ■ <code>lcredit=1</code>: The maximum credit for having lower case letters in the new password. If you have less than or 1 lower case letter, each letter will count +1 towards meeting the current <code>minlen</code> value. ■ <code>ucredit=1</code>: The maximum credit for having upper case letters in the new password. If you have less than or 1 upper case letter each letter will count +1 towards meeting the current <code>minlen</code> value. ■ <code>dcredit=1</code>: The maximum credit for having digits in the new password. If you have less than or 1 digit, each digit will count +1 towards meeting the current <code>minlen</code> value. ■ <code>ocredit=1</code>: The maximum credit for having other characters in the new password. If you have less than or 1 other characters, each character will count +1 towards meeting the current <code>minlen</code> value. ■ <code>enforce_for_root</code>: The password is set for the root user. <p>Note For more details on Linux PAM module to check the password against dictionary words, refer to the man page.</p> <p>For example, avoid simple and systematic passwords such as <code>VMware123!123</code> or <code>VMware12345</code>. Passwords that meet complexity standards are not simple and systematic but are a combination of letters, alphabets, special characters, and numbers, such as <code>VMware123!45, VMware 1!2345</code> or <code>VMware@1az23x</code>.</p>
Admin CLI Credentials and Audit CLI Credentials	<p>Select the Same as root password check box to use the same password that you configured for root, or deselect the check box and set a different password.</p>

7 Click Install Appliance.

The new node is deployed. You can track the NSX Manager deployment progress in the **System > Appliances** page (NSX UI) in VMware vCenter. Do not add additional nodes until the installation is finished and the cluster is stable.

8 Wait for the deployment, cluster formation, and repository synchronization to finish.**9 Verify that installed NSX Manager node has the required connectivity.**

Make sure that you can perform the following tasks.

- Ping your node from another machine.
- The node can ping its default gateway.
- The node can ping the hypervisor hosts that are in the same network using the management interface.
- The node can ping its DNS server and its NTP Server IP or FQDN list.
- If you enabled SSH, make sure that you can SSH to your node.

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

10 If your cluster has only two nodes, add another appliance.

- From NSX Manager, select **System > Appliances > NSX Manager > Add NSX Appliance** and repeat the configuration steps.

Results

After the cluster is formed, VMware vCenter displays the IP address of all the three nodes on the NSX UI page.

What to do next

After the cluster is formed, you can choose to set a virtual IP address (VIP) for the cluster or you can choose to not have a VIP for the cluster. See [Configure a Virtual IP Address for a Cluster](#). Even after you configure a VIP for the cluster, the NSX plugin in VMware vCenter continues to access NSX UI using the primary IP address of the current HTTPS leader node. During failover, the NSX plugin in VMware vCenter automatically starts using the primary IP address of the new leader node.

Configure NSX for Virtual Networking from vSphere Client

As a VI administrator working in the vSphere environment, you can configure NSX for virtual networking. The workflow involves configuring logical segments to establish connectivity between hosts even in different subnets, configuring NSX Edge nodes, Tier-0 gateways, Tier-1 gateways and segments. Finally, workload VMs connected to these segments can pass north-south and east-west traffic.

Prerequisites

- Ensure that ESXi hosts are compatible with VMware vCenter version v7.0.3 or later.
- Ensure that VMware vCenter version is v7.0.3 or later.
- Configure a vSphere Distributed Switch (VDS) switch on hosts. Only VDS 6.6 or later is supported.
- On a vSphere Lifecycle Manager enabled cluster, edit the VMware vCenter from the NSX Manager UI to:
 - Create a service account and enable trust between NSX and VMware vCenter. See [Add a Compute Manager](#).

Procedure

- 1 From a browser, log in with admin privileges to an VMware vCenter at `https://<vcenter-server-ip-address>`.
- 2 On the vSphere Client UI, select **vSphere Client** menu and click **NSX**.
- 3 On the **Welcome to NSX** screen, on the **Virtual Networking** card, click **Getting Started**.
- 4 In the **Host Cluster Preparation** tab, perform the following tasks.
- 5 Expand the **Host Cluster** section, select the clusters that you want to prepare for virtual networking and click **Next**.

Note Any cluster with an incompatible ESXi host is not allowed for host preparation.

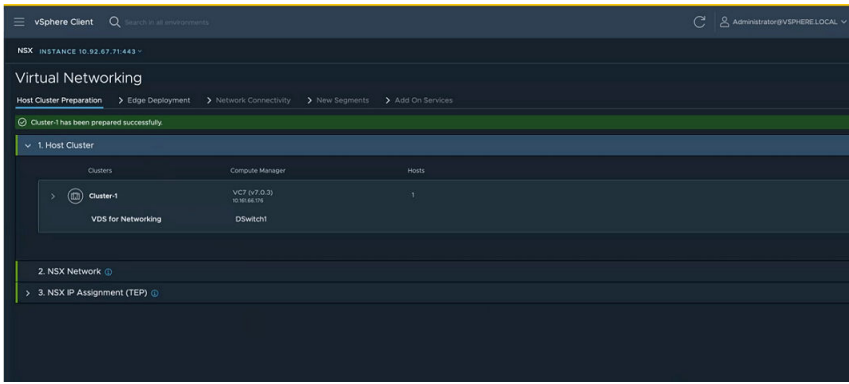
- 6 Expand the **NSX Network** section, enter a VLAN ID that will tag the overlay traffic and click **Next**.
- 7 Expand the **NSX IP Assignment (TEP)** section, and enter IP details:

Field	Description
IP Assignment	Select the mode of IP assignment, from between static and DHCP. If you select IP pool, enter a name for the pool, IP range, subnet along with prefix (subnet/prefix) and default gateway. NSXe does not support the IPv6 address type to be assigned as a TEP IP address of a host.

- 8 Click **Prepare Cluster** to begin installation of NSX.

Cluster preparation begins. View installation progress at each host.

Alternatively, you can also verify the progress in NSX Manager UI. NSX creates a new transport node profile using the configuration that you defined in the installation section. The switch is set to VDS. The transport node profile is applied to the cluster to prepare hosts of the cluster as transport nodes.



- In the **Edge Deployment** tab, expand the **Management Network for Edge Connectivity** and enter the following details:

Field	Description
Management VDS	Select a vSphere Distributed Switch for management traffic on NSX Edge nodes.
Management Network	Select a network for management traffic of NSX Edge nodes.
Management Gateway	Select the gateway to route management traffic. Enter a static IP address.

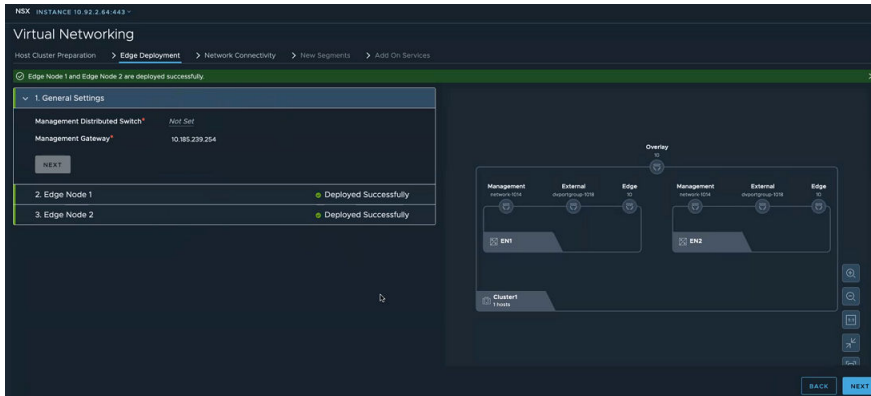
- Click **Next**.

- In the **Edge Deployment** tab, expand the **Edge Node 1** and enter the following details

Field	Description
Name	Enter a name for the Edge node.
Fully Qualified Domain Name	Enter a Fully Qualified Domain Name that resolves to the IP address of the NSX Edge node.
Management IP address	Enter an IP address for management traffic of the NSX Edge node.
External Network Connectivity	Select a distributed port group to be used as a data path interface. This distributed port group manages the ingress and egress traffic of workload VMs that are processed by the NSX Edge node. Note Even though there are three data path interfaces on an NSX Edge node, this workflow uses only one interface for a distributed port group.
Edge Node Settings	Select Apply same settings for all Edges if you want to replicate the settings across all NSX Edge nodes.
Password	Enter a password that conforms to the required password complexity. Additionally, the password length must not exceed 128 characters. Enter the same password in the next field.
Virtual Machine Size	Select a form factor to deploy the Edge node.
Storage Location	Select the datastore as storage location for installation and configuration files and data generated by the Edge node.

- On the **Edge Deployment** tab, verify that the visualization is updated with management network, external network and other details related to NSX Edge node.

- 13 Enter details for **Edge Node 2**.
- 14 Click **Deploy Edge**.
- 15 On the confirmation window, click **Deploy**.
- 16 Observe the topology created based on the configuration details entered on the **Edge Deployment** tab. After NSX Edge nodes are realized, the dotted line turns to a solid line, indicating NSX Edge node is realized.



- 17 Click **Next** to configure network connectivity.
- 18 In the **Network Connectivity** tab, expand the **Physical Router** section and enter the following details:

Field	Description
<p>Do you want to peer with a physical router now?</p>	<p>After deploying the NSX Edge node, it can establish a peer connection with a physical router.</p> <ul style="list-style-type: none"> ■ Select Yes if you want to set up Border Gateway Protocol (BGP) or static routing to your physical router. <ul style="list-style-type: none"> ■ BGP Local AS: Enter the local autonomous system number for use in BGP. ■ Select No if you do not want to set up BGP or static routing to your router. However, you will need to set up NAT to connect to workloads to external networks. <ul style="list-style-type: none"> ■ In the Physical Routing IP address field, enter a static IP address.
<p>How many physical routers do you want to peer with?</p>	<p>Based on your selection, enter the following details for one or two physical routers:</p> <p>If you want to allow other routers to peer with your router, then enter the following details:</p> <p>For each peer router, enter these details:</p> <ul style="list-style-type: none"> ■ BGP Local AS: Enter the local autonomous system number used by the BGP neighbor. ■ BGP Neighbors: <ul style="list-style-type: none"> ■ IP Address: Enter the IP address of the physical router, which is the BGP neighbor. ■ Remote AS: Remote autonomous system number used by BGP neighbors.

19 Click **Next**.

20 In the **Network Connectivity** tab, expand the **NSX Gateway** section and enter the following details:

Field	Description
Gateway Name Prefix	Enter a prefix for the gateway. Every object, such as Tier-0, Tier-1 gateways, that are created for the gateway is prefixed with this value. You can search for objects with a specific prefix to get a list of objects related to a particular gateway.
Uplink VLAN for Router 1	Enter the VLAN ID to tag VLAN traffic going from NSX Edge node to physical router 1.
Uplink VLAN for Router 2	Enter the VLAN ID to tag VLAN traffic going from NSX Edge node to physical router 2.

21 Click **Next**.

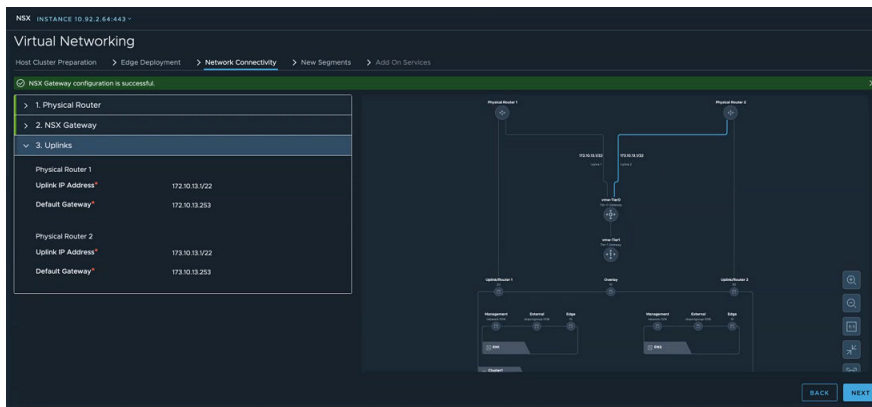
22 In the **Network Connectivity** tab, expand the **Uplinks** section and enter the following details:

Field	Description
IP Address for Uplink 1	Enter the IP address for uplink from NSX gateway or Tier-0 gateway to physical router 1.
IP Address for Uplink 2	Enter the IP address for uplink from NSX gateway or Tier-0 gateway to physical router 2.
Physical Router 1	Enter subnet mask and default gateway for physical router 1.
Physical Router 2	Enter subnet mask and default gateway for physical router 2.

23 Verify the visualization created based on the network details you entered.

24 Click **Create Gateways**.

25 On the confirmation window, click **Create Gateways**.



The NSX Gateway is successfully created.

26 In the **New Segments** tab, create a segment where workloads VMs will be running. For example, create a segment for a web group. Enter the following details:

Field	Description
Name	Enter the name of the segment.
Subnet/Prefix Length	Enter the subnet network for the segment.
Default Gateway	Enter the default gateway that segments should forward traffic to.

27 To create additional segments, click **Add Segment** and enter the required details.

28 Click **Create Segment**.

29 After segments are created, add them to NSX distributed virtual port group.

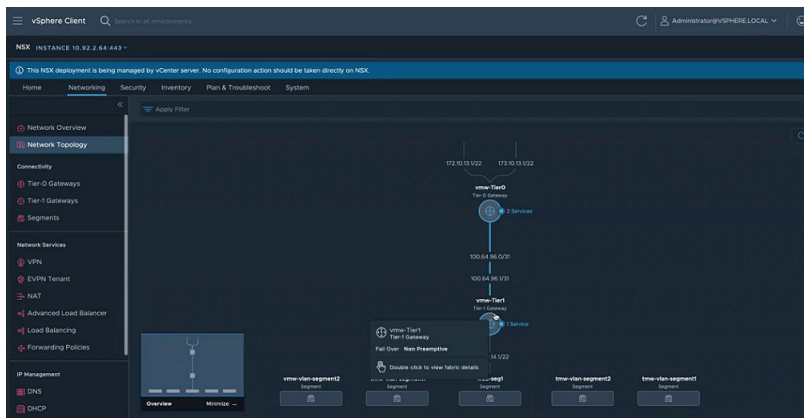
30 Click **Next**.

31 (Optional) In the **Add-on Services** tab, enter Network Address Translation (NAT) details. On the **NAT Only** window, enter the following details:

Field	Description
Name	Enter a name for NAT service.
Source	Select a segment so that IP addresses of local hosts connected to this segment are translated and protected and a single translated IP address is presented to an external network.
Translated	The IP address that is presented to external network thus protecting local hosts from exposing their IP addresses to an external network.

32 Click **Next**.

33 View the created topology created in NSX.



Results

NSX is configured for virtual networking.

Example:


What to do next

- 1 From a browser log in to NSX Manager with `https://<NSX Manager-IP-Address>/`.
- 2 After logging in to NSX Manager, verify networking configuration is successfully created in NSX.

Configuring NSX-T for Security from vSphere Client

As a VI admin working in the vSphere environment, you can use the simplified workflow to prepare ESXi clusters for NSX security.

Use the vSphere Client to prepare ESXi clusters for NSX security. On such clusters, you can enable micro-segmentation, URL filtering and distributed IDS on application workloads. These clusters are not prepared for NSX virtual networking.



High-level tasks:

- Prepare Host Cluster.
- Create Firewall Rules
 - Create Groups for infrastructure services (Active Directory, DNS, and so on), environment groups (production or testing), and application groups (web, database, application).
 - Define communication strategy. Some of the actions you can take are:
 - Define communication between any workload and infrastructure services.
 - Define communication so that no environment can talk to each other.
 - Limit communication to a specific port or protocol.
 - Specify source workloads.
 - Set up exceptions after setting up communication strategies for workloads.
 - Define Action for Default Firewall Rule (to process traffic that does not match firewall rules defined in Communication section).
 - Review and publish firewall rules.

Prepare Clusters for NSX Security

Select a host cluster to prepare it for NSX security.

The Getting Started section gives you the option to select between **Security Only** or **Virtual Networking**. When you choose to enable clusters only for security, the wizard asks you to define security rules and uses those rules to automatically configure NSX security on the distributed virtual port groups of the selected clusters.

Prerequisites

- Ensure that ESXi hosts are compatible with VMware vCenter version v7.0.3 or later.
- Ensure that VMware vCenter version is v7.0.3 or later.
- Configure a vSphere Distributed Switch (VDS) switch on hosts. Only VDS 6.6 or later is supported.
- On a vSphere Lifecycle Manager enabled cluster, edit the VMware vCenter from the NSX Manager UI to:
 - Create a service account and enable trust between NSX and VMware vCenter. See [Add a Compute Manager](#).

Procedure

- 1 From a browser, log in with admin privileges to a VMware vCenter at `https://<vcenter-server-ip-address>`.
- 2 On the vSphere Client UI, select the vSphere Client menu and click **NSX**.
- 3 On the Welcome to NSX screen, on the **Security Only** card, click **Getting Started**.
- 4 On the **Host Cluster Preparation** section, select the clusters that you want to prepare for security only and click **Install NSX**.
- 5 On the Install Security pop-up window, confirm you want to process by clicking **Install**.

Note Any cluster with an incompatible ESXi host is not allowed for host preparation.

- 6 Click **Next** to define firewall rules.

Results

NSX is installed on the host cluster.

What to do next

To avoid any loss of connectivity, add VMware vCenter and NSX Manager to the DFW Exclusion list.

Create Groups

As part of firewall creation, define infrastructure group that run selected services, such as DHCP, define environment groups, such as production, testing, or so on, comprising of selected group members and define application groups with selected group members.

Prerequisites

- Install NSX on the host cluster.

Procedure

- 1 In the **Create Firewalls Rules** tab, select **Create Groups**.
- 2 In the **Create Groups** page, expand **Create Infrastructure Groups**.
- 3 Click **Add Group**.
- 4 From the **Infrastructure Service** drop-down menu, select a service, such as Active Directory. In the next step, you assign this service to a group comprising of members that form the infrastructure group. You can create an infrastructure service only once in a workflow. It cannot be edited once you create it.
- 5 To define an infrastructure group, click [**Define Group**].

An infrastructure can be a combination of VMs, IP address range, or distributed virtual port groups.

- a (Optional) In the **Group Name** field, modify the default group name.
- b (Optional) In the **NSX Tag** field, modify the default tag name. The defined tag is applied to all VMs and distributed virtual port groups selected for the group. You can edit the default tag name.
- c Expand the **Select VMs to add NSX Tag** section and select VMs that must be part of the infrastructure group.
- d Expand the **IP Address** section and enter an IP address, IP addresses in CIDR format, or an IP range. Both IPv4 and IPv6 formats are supported.
- e Expand the **Select DVPGs to add NSX Tag** section and select the distributed virtual port groups that must be part of the infrastructure group.
- f Click **Save**.

The wizard automatically creates the group and applies the NSX tag on all the selected members of the group. For example, if the defined group includes one VM, one distributed virtual port group, and 1 IP address, and DHCP is the selected infrastructure service, then wizard tags all group members with the defined tag.

- 6 Click **Next**.
- 7 In the **Create Groups** page, expand **Create Environment Group**.
- 8 Click **Add Group**.
- 9 From the **Environment** drop-down menu, select the environment for the group. For example, an environment can be a production, testing, partner or a custom environment that you want to define in your topology.

- 10 To define an environment group, click [**Define Group**].
 - a (Optional) In the **Group Name** field, modify the default group name.
 - b (Optional) In the **NSX Tag** field, modify the default NSX tag name. This tag name is applied to all VMs and distributed virtual port group selected for the environment group.
 - c Expand the **Select VMs to add NSX Tag** section and select VMs that must be part of the environment group.
 - d Expand the **IP Address** section and enter an IP address, IP addresses in CIDR format, or an IP range. Both IPv4 and IPv6 formats are supported.
 - e Expand the **Select DVPGs to add NSX Tag** section and select the distributed virtual port groups that must be part of the environment group.
 - f Click **Save**.
- 11 Click **Next**.
- 12 In the **Create Groups** page, expand **Create Application Group**.
- 13 Click **Add Group**.
- 14 From the **Application Group Name** drop-down menu, select the type of application group you want to create.
- 15 To define an application group, click [**Define Group**].
 - a (Optional) In the **Group Name** field, modify the default group name for the application group.
 - b (Optional) In the **NSX Tag** field, modify the default tag name. This tag name is applied to all VMs and distributed virtual port group selected for the application group, enter a NSX tag.
 - c Expand the **Select VMs to add NSX Tag** section and select VMs that must be part of the application group.
 - d Expand the **IP Address** section and enter an IP address, IP addresses in CIDR format, or an IP range. Both IPv4 and IPv6 formats are supported.
 - e Expand the **Select DVPGs to add NSX Tag** section and select the distributed virtual port groups that must be part of the application group.
 - f Click **Save**.
- 16 Click **Next**.

Results

You created infrastructure groups, environment groups and application groups.

What to do next

After creating groups, define firewall rules that govern communication among workloads and these different groups.

Define and Publish Communication Strategies for Groups

After creating groups, define firewall rules to govern communication between groups, define exceptions and ports or protocols for communication.

Prerequisites

- Install NSX on the host cluster.
- Create Infrastructure groups, Environment groups, and Application groups.

Procedure

- 1 Expand the **Access to infrastructure services** section and define specific workloads that can access shared infrastructure services.

Field	Description
Source	In the Source column, select the workloads that can access the target infrastructure service.
Target	Is the defined infrastructure service that is accessed by source workloads.
(NSX3.2.2) Service Entry	<p>Click the Edit icon to add or edit service entries.</p> <p>In the Service Entry window, select a service type and properties for the service type.</p> <p>Note In NSX 3.2.1 and previous versions, the field name was L4.</p>

- 2 Click **Next**.

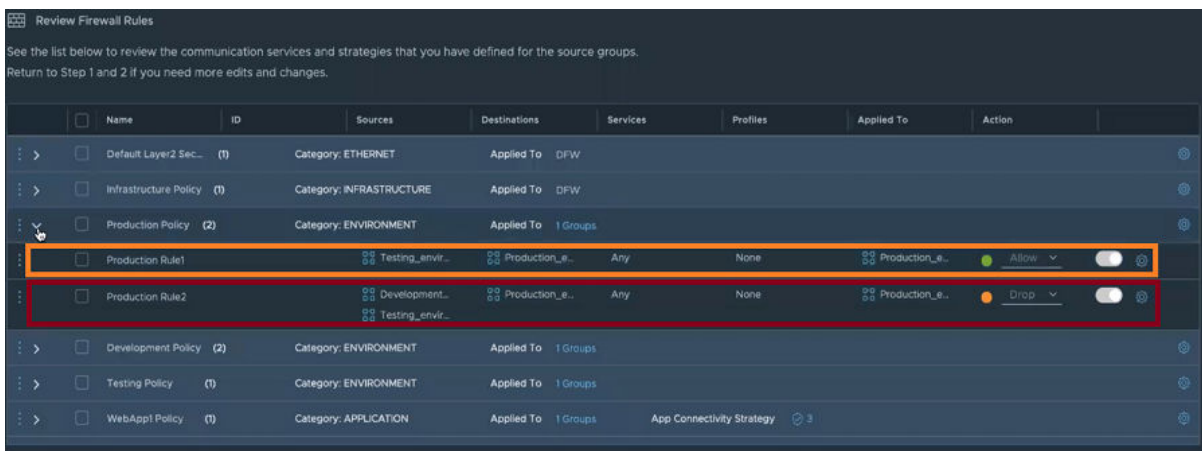
- 3 Expand the **Define communication between environments (Optional)** section and define communication between groups.

Field	Description
Source	<p>Expand the section to define which source environment must communicate with a target environment.</p> <p>(NSX 3.2.2): For each source group listed, select a communication method: Unprotected, Allowed or Blocked.</p> <hr/> <p>Note To allow all communication between all source groups and the target group, select Allow All Communication.</p> <hr/> <p>(NSX 3.2.1 and previous versions): To allow communication between a Development environment and a Production environment, click the red dotted line between Development and Production. The enabled state is displayed when a green line is established between groups.</p>
Environment	Is the target environment selected by the system.
(NSX3.2.2) Service Entry	<p>Select the service type, ports and properties over which the workloads in source and target environments communicate with each other.</p> <p>Click Apply.</p> <hr/> <p>Note In NSX 3.2.1 and previous versions, the field name was L4.</p>

- 4 Click **Next**.
- 5 Expand the **Define communication strategies for applications (Optional)** section and define communication for application groups.

Field	Description
Source	Select an application group for which you can select communication rules to manage incoming or outgoing traffic.
Strategy	<p>Select a firewall strategy to apply to an application group.</p> <p>Supported firewall rules are:</p> <ul style="list-style-type: none"> ■ Allow all external traffic. ■ Deny incoming and allow outgoing traffic. ■ Allow incoming and deny outgoing traffic. ■ Deny all external traffic. <hr/> <p>Note If you want to apply one firewall rule to all application groups, click Select Strategy, select the rule and click Apply.</p>
Exception	<p>Based on how you want to configure firewall rule, you might want to add exceptions.</p> <p>By default, no exceptions are added. To add an exception, click the No Exceptions link. Edit these fields to add exceptions:</p> <ul style="list-style-type: none"> ■ Source: Select the source. ■ Service Entry: Select the service, port and properties. ■ L7 App ID: Select the App ID. ■ FQDN: Select FQDN of the application. <p>.Click Apply.</p>

- 6 Click **Next**.
- 7 Expand the **Define Action for Default Firewall Rules (Optional)** section and define an action that is applied to traffic that does not match the defined criteria.
- 8 In the Default rule action, select from one of the following:
 - **Allow**: Is the default rule set. Allows all traffic that does not match the defined criteria.
 - **Drop** or **Reject**: To enforce firewall rules insider your network, you might choose to drop traffic that does not match the defined criteria.
- 9 Click **Next**.
- 10 In the Review and Publish page, review the communication strategies and firewall rules that you applied to the groups.



In the screenshot, Production Rule 1 is a user-defined rule and Production Rule 2 is system-defined default rule, where the default action is set to **Drop**.

- 11 Click **Publish Policies**.

Results

The wizard ends and firewall policies you defined are applied to the groups. The NSX UI is available in VMware vCenter.

What to do next

To verify the firewall rules published from vSphere Client are realized on NSX Manager UI.


- 1 In the NSX Manager UI, go to **Inventory** → **Groups**.
- 2 On the Groups page, verify whether the workload groups you defined in vSphere Client are realized in NSX Manager.
- 3 Go to **Security** → **Distributed Firewall** page.
- 4 On the Distributed Firewall page, verify whether the firewall rules you applied in vSphere Client are realized in NSX Manager.

Viewing NSX Alarms in vSphere Web Client UI

As a VI admin you can view NSX generated alarms in VMware vCenter.

Note The NSX Alarms forwarding functionality is supported starting with vSphere version 8.0 onwards.

After you install and configure NSX from the NSX page in VMware vCenter environment, NSX automatically registers its alarm definitions with VMware vCenter. Any NSX alarms generated are forwarded and displayed on the vSphere Web Client UI. Alarm forwarding is only applicable with NSX deployments in VMware vCenter environment.



View NSX Alarms on vSphere Web Client UI at:

- **Monitor page** → **Issues and Alarms** → **Triggered Alarms**
- **Alarms** section

To identify a NSX alarm, go to the **Triggered Alarms** page and check the user field that triggered the alarm. The user `com.vmware.nsx.management.nsxxt` identifies that the event was generated by NSX.

Note If there are multiple alarms generated for a certain NSX event, in vSphere Web Client a summary event is displayed. The summary event takes you to NSX Alarms UI page where all alarms of the same type are displayed. For example, if NSX generates five alarms related to certificate expiry, the **Alarms** tab in vSphere Web Client displays only a single summary event for all the five certificate expiry alarms.

Caution Do not delete any NSX alarms from VMware vCenter or reset any alarms (alarm turns to green indicating it is resolved) . If you delete an **Alarm Definition** in VMware vCenter, any alarm generated for that event may not be forwarded to VMware vCenter.

If one of the NSX Manager goes down, another NSX Manager in the cluster takes over as the active manager. VMware vCenter synchronizes existing alarm definitions and triggered alarms with the current state of alarms on the new NSX Manager.

Transport Zones and Profiles



Transport zones and profiles are building blocks to prepare hosts for NSX networking.

Read the following topics next:

- [Create Transport Zones](#)
- [Create an IP Pool for Tunnel Endpoint IP Addresses](#)
- [Enhanced Data Path](#)
- [Guidance to Set Maximum Transmission Unit](#)
- [Configuring Profiles](#)
- [Prepare a vSphere Distributed Switch for NSX](#)
- [Add a Transport Node Profile](#)

Create Transport Zones

Transport zones dictate which hosts transport nodes 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 segment—and, therefore, which VMs can be attached to the segment. A transport zone can span one or more host clusters. Also, a host transport node can be associated to multiple transport zones.

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

NSX does not allow connection of VMs that are in different transport zones in the Layer 2 network. The span of a segment is limited to a transport zone.

Both host transport nodes and NSX Edge nodes use Overlay and VLAN transport zones. Host transport nodes connect to VDS switches while N-VDS switch is configured on NSX Edge transport nodes.

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.

Note vMotion is not supported between two segments or logical switches on different VLAN transport zones.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System > Fabric > Transport Zones > Add Zone**.
- 3 Enter a name for the transport zone and optionally a description.
- 4 For the **Traffic Type** drop-down menu, specify whether the transport zone is for overlay or VLAN traffic.

The options are **Overlay** and **VLAN**.

For a transport zone to forward IPv6 traffic, it must be set for an overlay traffic type.

- 5 For overlay transport zones, specify the forwarding mode which determines the underlay transport protocol for the traffic encapsulation.

The forwarding mode can be for **IPv4** or **IPv6**.

Note Dual stack (both IPv4 and IPv6) is not supported.

- 6 (Optional) In the **VLANs** field, specify the VLAN IDs or range of VLANs IDs that are allowable for use.

For example, entering **1-1000**, **2000-3000**, **4094** would allow VLAN IDs 1 through 1000, 2000 through 3000, and 4094.

Segments of the transport zone must have VLANs within the specified ranges or match the specified values.

- 7 (Optional) For VLAN transport zones, you can take one of the following actions based on the NSX release version:
 - NSX 4.1.1 and earlier: Enter the uplink teaming policy names. These named teaming policies can be used by segments attached to the VLAN transport zone which in turn use the named teaming policy specified in uplink profiles to direct traffic. For more information, see [Configure Named Teaming Policy](#).

Note If you define named teaming policies, ensure that you enter the exact named teaming policy name in associated VLAN segments and uplink profiles as well. If segments do not find a matching named teaming policy, then NSX uses the default uplink teaming policy.

- NSX 4.1.2 and later: Select the uplink teaming policy from the drop-down list. These named teaming policies can be used by segments attached to the VLAN transport zone which in turn use the named teaming policy specified in uplink profiles to direct traffic. For more information, see [Configure Named Teaming Policy](#).

- 8 After you add the transport zone, go to the **Transport Zones** page and view the newly added transport zone either from the UI or by running the following API command.

```
GET /policy/api/v1/global-infra/sites/<site-id>/enforcement-points/
<enforcementpoint-id>/transport-zones
```

```
{
  "sort_ascending": true,
  "sort_by": "display_name",
  "result_count": 1,
  "results": [
    {
      "tz_type": "OVERLAY_BACKED",
      "is_default": true,
      "transport_zone_profile_paths": [
        "/infra/transport-zone-profiles/tzp"
      ],
    },
    "nested_nsx": false,
    "resource_type": "PolicyTransportZone",
    "id": "tz",
    "display_name": "tz",
    "path": "/infra/sites/default/enforcement-points/default/transport-zones/tz",
    "relative_path": "tz",
    "parent_path": "/infra/sites/default/enforcement-points/default",
    "unique_id": "8f4a026d-e3f5-4f23-a3ef-46309d573dc1",
    "marked_for_delete": false,
    "overridden": false,
    "_create_user": "admin",
    "_create_time": 1607501697823,
    "_last_modified_user": "admin",
    "_last_modified_time": 1607582307987,
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 5
  ]
}
```

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 (deprecated) `PATCH /api/v1/infra/transport-zone-profiles` API. There is no UI workflow for creating a transport-zone profile. After the transport-zone profile is created, you can attach it to the transport zone with the `PATCH https://<policy-mgr>/policy/api/v1/infra/sites/default/enforcement-points/nsxt-ep/transport-zones/<transport-zone-id>` API.

```
{
  "tz_type": "OVERLAY_BACKED",
  "is_default": true,
  "nested_nsx": false,
  "transport_zone_profile_paths": [
    "/infra/transport-zone-profiles/tzp"
  ]
}
```

Create an IP Pool for Tunnel Endpoint IP Addresses

NSX Transport Nodes are configured with tunnel endpoint (TEP pool) IP addresses. 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 encapsulation of overlay frames. You can also use either DHCP or manually configured IP pools for tunnel endpoint IP addresses.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Navigate to **Networking > IP Address Pools**.
- 3 Click **Add IP Address Pool**.
- 4 Enter a name and optionally a description.
- 5 From the **Subnets** column, click **Set** to add the subnets.
- 6 Click the **Add Subnet** drop-down menu and choose either **IP Block** or **IP Ranges**:
 - Choose **IP Block** to use an IP block of the allowable IP addresses:

Setting	Description
Select IP Blocks	Choose an IP address block.
Size	Specify the size or the number of IP addresses in the IP address block.

Setting	Description
Allocation Range	Specify the allocation range for the IP address block. Note For an IPv6 address block, you must specify an allocation range. The allocation range value cannot be more than 1048576.
Auto Assign Gateway	Click the toggle to enable or disable automatic gateway IP assignment.

- Choose **IP Ranges** to specify the range of IP addresses that can be used:

Setting	Description
Enter IPv4 or IPv6 Ranges	Enter the range of IPv4 or IPv6 addresses. Note Transport nodes and tunnel endpoints (TEPs) do not allow a mixed IP address pool to be used (an IP pool with both IPv4 and IPv6 ranges). Ensure that the IP range consists of only IPv4 or only IPv6 addresses.
CIDR	Enter the IP address range in CIDR format.
Gateway IP	Enter the gateway address.
DNS Server	Enter the list of the DNS (domain name system) server IP addresses separated by commas.
DNS Suffix	Enter the DNS suffix. For example, <code>corp.local</code> .

7 Click **Add** and then **Apply**.

8 Click **Save**.

Note IP Pools are configured for only those hosts that participate in overlay traffic (hosts that are members of an overlay transport zone). So, TEP tunnels are not required for host clusters that are member of VLAN transport zones only and/or management host clusters (hosting NSX Appliances / NSX Edge VMs) if management appliances are connected to VLAN backed VDS portgroups.

Results

The IPv4 or IPv6 address pool is listed on the IP pool page.

If you are using API to view the list of IP addresses allocated from an IP pool, run the NSX Manager API. Do not run the Policy API to view IP pool allocations.

For example, the Policy API call does not retrieve IP addresses allocated from an IP pool.

GET https://<nsxmanager-IP>/policy/api/v1/infra/ip-pools/TEP-IP-Pool/ip-allocations

```
{
  "results": [],
  "result_count": 0,
  "sort_by": "display_name",
  "sort_ascending": true
}
```

The NSX Manager API retrieves IP addresses allocated from an IP pool.

GET https://<nsxmanager-IP>/api/v1/pools/ip-pools/<ip-pool-UUID>/allocations

```
{
  "results": [
    {
      "allocation_id": "192.85.85.12",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.13",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.14",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.15",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.16",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.17",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.18",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.19",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.20",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.21",
      "_protection": "NOT_PROTECTED"
    }
  ]
}
```

```

    },
    {
      "allocation_id": "192.85.85.11",
      "_protection": "NOT_PROTECTED"
    },
    {
      "allocation_id": "192.85.85.22",
      "_protection": "NOT_PROTECTED"
    }
  ],
  "result_count": 12
}

```

What to do next

Create an uplink profile. See [Create an Uplink Profile](#).

Enhanced Data Path

Enhanced Data Path (EDP) is a networking stack mode, which when configured provides superior network performance. It is primarily targeted for NFV workloads, which offer performance benefits leveraging DPDK capability.

The VDS switch can be configured in the enhanced data path mode only on an ESXi host. Enhanced Data Path also supports traffic flowing through Edge VMs.

In the enhanced data path mode, both traffic modes are supported:

- Overlay traffic
- VLAN traffic

Supported VMkernel NICs

With NSX supporting multiple Enhanced Data Path 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 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.
- 2 On the VMware Compatibility Guide page, from the **Systems/Servers** drop-down menu, select **IO Devices**.

- 3 On the **IO devices** page, in the **Product Release Version** section, select **ESXi <version>**.
- 4 In the **IO device Type** section, select **Network**.
- 5 In the **Features** section, **Enhanced Datapath - Interrupt Mode** or **Enhanced Datapath - Poll Mode**.
- 6 Click **Update and View Results**.
- 7 In the search results, you will find the supported NIC cards compatible with the **ESXi** version you selected.
- 8 Identify the brand for which you want to download the driver and click the **Model** URL to view and download the driver.
- 9 Download and install the latest NIC drivers from the [Broadcom Support](#) page.
 - a Select the VMware vSphere version.
 - b Go to **Drivers & Tools > Driver CDs**.
 - c Download the NIC drivers.
 - d To use the host as an Enhanced Data Path host, at least one Enhanced Data Path capable NIC must be available on the system. If there are no Enhanced Data Path capable NICs, the management plane will not allow hosts to be added to Enhanced Data Path transport zones.
 - e List the Enhanced Data Path driver.


```
esxcli software vib list | grep -E "i40|ixgben"
```
 - f Verify whether the NIC is capable to process Enhanced Data Path traffic.

```
esxconfig-nics -e
```

Name	Driver	ENS Capable	ENS Driven	MAC Address	Description
vmnic0	ixgben	True	False	e4:43:4b:7b:d2:e0	Intel (R) Ethernet Controller X550
vmnic1	ixgben	True	False	e4:43:4b:7b:d2:e1	Intel (R) Ethernet Controller X550
vmnic2	ixgben	True	False	e4:43:4b:7b:d2:e2	Intel (R) Ethernet Controller X550
vmnic3	ixgben	True	False	e4:43:4b:7b:d2:e3	Intel (R) Ethernet Controller X550
vmnic4	i40en	True	False	3c:fd:fe:7c:47:40	Intel (R) Ethernet Controller X710/X557-AT 10GBASE-T
vmnic5	i40en	True	False	3c:fd:fe:7c:47:41	Intel (R) Ethernet Controller X710/X557-AT 10GBASE-T
vmnic6	i40en	True	False	3c:fd:fe:7c:47:42	Intel (R) Ethernet Controller X710/X557-AT 10GBASE-T
vmnic7	i40en	True	False	3c:fd:fe:7c:47:43	Intel (R) Ethernet Controller X710/X557-AT 10GBASE-T

- g Install the Enhanced Data Path driver.

```
esxcli software vib install -v file:///<DriverInstallerURL> --no-sig-check
```

- h Alternately, download the driver to the system and install it.

```
wget <DriverInstallerURL>
```

```
esxcli software vib install -v file:///<DriverInstallerURL> --no-sig-check
```

- i Reboot the host to load the driver. Proceed to the next step.
- j To unload the driver, follow these steps:

```
vmkload_mod -u i40en
```

```
ps | grep vmkdevmgr
```

```
kill -HUP "$(ps | grep vmkdevmgr | awk {'print $1'})"
```

```
ps | grep vmkdevmgr
```

```
kill -HUP <vmkdevmgrProcessID>
```

```
kill -HUP "$(ps | grep vmkdevmgr | awk {'print $1'})"
```

- k To uninstall the Enhanced Data Path driver, `esxcli software vib remove --vibName=i40en-ens --force --no-live-install`.

Note Enhanced Data Path transport zones configured for overlay traffic: For a Microsoft Windows virtual machine running VMware tools version prior 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 that the virtual machine MTU is set to a value less than **8900**.

- 10 Create a host transport node. Configure mode in Enhanced Datapath on a VDS switch with logical cores and NUMA nodes.

Load Balanced Source Teaming Policy Mode Aware of NUMA

The Load Balanced Source teaming policy mode defined for an enhanced datapath 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.

Enhanced Data Path Support for Applications Requiring Traffic Reliability

NFV workloads might use multi-homing and redundancy features provided by Stream Control Transmission Protocol (SCTP) 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 a 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. The enhanced data path switch provides redundant network paths between the hosts.

Figure 8-1. Multi-homing and Redundancy of Traffic over Enhanced Data Path



The high-level tasks are:

- 1 Prepare host as an NSX transport node.
- 2 Prepare VLAN or Overlay Transport Zone with two VDS switches in Enhanced Data Path mode.
- 3 On VDS 1, pin the first physical NIC to the switch.
- 4 On VDS 2, pin the second physical NIC to the switch.

The 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.

Enhanced Datapath Supported Features

Note In NSX 3.2.X version and earlier, the term used is ENS (Enhanced Network Stack). In NSX 4.0.X version and later, the term used is Enhanced Datapath.

Supported Features

Name	Notes
Segments (Logical Switches) – VLAN based	
Segments (Logical Switches) – Overlay based	

Name	Notes
Differentiated Services Code Point (DSCP) trust/re-marking	
Ingress and Egress rate limiting	
Distributed Firewall – L2	
Distributed Firewall – L3 / L4	
IDFW	Traffic will be flowing. However, the performance will be equal to regular non-Enhanced Datapath performance.
IDS/IPS	Traffic will be flowing. However, the performance will be equal to regular non-Enhanced Datapath performance.
East/West Service Chaining	Traffic will be flowing. However, the performance will be equal to regular non-Enhanced Datapath performance.
Distributed Load Balancing (DLB)	DLB on Enhanced Datapath works at full performance.
IP Discovery (DHCP/DHCPv6 Snooping, ARP snooping, ND snooping, etc.)	
Spoofguard (IP, MAC, etc.)	
Segment Security Filter (Bridge Protocol Data Units (BPDU), DHCP and so on.)	
DHCP Relay & DHCP Local and Gateway Server	
L2 Bridging through Edge	
L2 Bridging through L2VPN	
Distributed Routing (Overlay)	
EVPN in route server mode	
EVPN through Edge	
L2 Multicast	Traffic will be flowing. However, the performance will be equal to regular non-Enhanced Datapath performance.
IPFIX	
Port Mirroring	Support for ERSPAN with Enhanced Datapath performance, other port mirroring methods will be equal to regular non-Enhanced Datapath performance.

Unsupported Features

Name	Notes
Network I/O Control (NIOC)	Disable NIOC on VDS switches used while configuring NSX.

Guidance to Set Maximum Transmission Unit

Get guidance on how to set the Maximum Transmission Unit (MTU) value in the different objects or profiles in NSX.

Jumbo Frame Support

The minimum required MTU is 1600 bytes. However, MTU of 1700 bytes is recommended to address the whole possibility of a variety of functions and future proof the environment for an expanding Geneve header. In order to get better performance for applications generating large packets, and for optimal throughput, increase MTU to atleast 9,000 bytes as long as underlay physical infrastructure supports it and is also set to use jumbo frame MTU of 9000 bytes.

VM MTU

In most of the deployments, the guest VM MTU is set to 1500 bytes. So no change is required for the VM MTU if the physical fabric has an MTU of 1700 bytes or higher. To improve the throughput, one can increase the MTU up to 8800 (a estimated number to accommodate bridging and future header expansion) only if underlay physical infrastructure is set to use 9000 bytes. VM MTU should be set 100 bytes or more (200 preferred) lower than the MTU of the physical fabric.

MTU Configuration

- **Global Tunnel Endpoint MTU:** To configure the MTU value, go to **System** → **Settings** → **Global Fabric Settings**. The default value for MTU is 1700 bytes. When you set this MTU value, NSX configures the MTU value for all the N-VDS instances used in NSX Transport Nodes.
- **Global Logical Interface MTU:** To configure the MTU value, go to **Networking** → **Global Networking Config**. The default value for MTU is 1500. When you set this MTU value, NSX configures the MTU value for all the logical router interfaces. If the Global Logical interface MTU value is not specified, the MTU value is taken from the Tier-0 logical router (T-O Gateway). However, on a specific port, the logical router uplink MTU value can override the Global Logical interface MTU value.
- **Uplink Profile MTU:** To configure the MTU value, go to **System** → **Profiles** → **UplinkProfiles**. When you set this MTU value, NSX configures the MTU value for NSX Transport nodes that use N-VDS switch. This MTU field is optional in the uplink profile. If you do not configure it, NSX takes the value set in the global Tunnel Endpoint MTU.
- **(vSphere) VDS MTU:** To configure the MTU value, go to the VMware vCenter and modify the VDS directly. When you set this MTU value, NSX configures the MTU value for NSX Transport Nodes that use vSphere VDS. In this case, MTU value set on the attached uplink profile is not used.

Design Guidance

For optimal throughput, set the Global Tunnel Endpoint MTU , Uplink Profile MTU, and vSphere VDS MTU to at least 9000 bytes as long as:

- The underlying infrastructure supports 9000 bytes.
- The underlying infrastructure is set to use jumbo frame MTU of 9000 bytes.

Otherwise, configure Global TEP MTU, Uplink Profile MTU, and vSphere VDS MTU to minimum 1600 bytes or minimum recommended 1700 bytes.

The Gateway Interface MTU can continue to have default value. If you modified the Gateway Interface MTU, the modified value must be at least 200 bytes less than the Fabric MTU (which refers to Global Tunnel Endpoint MTU or VDS MTU or Uplink Profile MTU).

Important When adjusting the fabric MTU packet size, you must also configure the entire network path (VMkernel ports, virtual switches, physical switches and routers) to support the same MTU packet size. If a device along the path does not support the required frame size and receives a frame larger than its MTU, the device will drop the frame.

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 or hypervisor nodes to the top-of-rack switches or NSX logical switches. A link is from a physical network interface on an NSX Edge node or hypervisor nodes 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.

Consider the following points when configuring Failover Teaming Policy for VM appliance-based NSX Edge nodes and Bare Metal NSX Edge:

- For uplinks used by a teaming policy, you cannot use the same uplinks in a different uplink profile for a given NSX Edge transport node. Standby uplinks are not supported and must not be configured in the failover teaming policy. If the teaming policy uses more than one uplink (active/standby list), you cannot use the same uplinks in the same or a different uplink profile for a given NSX Edge transport node.
- Supported scenarios:
 - Bare Metal NSX Edge supports a single active uplink and a standby uplink. They do not support multiple standby uplinks.
 - NSX Edge VMs do not support any standby uplinks - single or multiple standby uplinks.

Consider the following points when configuring Load Balance Source for VM appliance-based NSX Edge nodes:

- Supports multiple active uplinks.

- You cannot use LAG to configure the teaming policy.
- In the **Active Uplinks** field, enter uplink labels that will be associated to physical NICs when you prepare transport nodes. For example, uplink1, uplink2. When you prepare transport nodes, you will associate uplink1 to pnic1 and uplink2 to pnic2.
- You must use the **Load Balanced Source** teaming policy for traffic load balancing.

Consider the following points when configuring Load Balance Source for Bare Metal NSX Edge:

- Supports multiple active uplinks.
- In the **Active Uplinks** field, you can use LAGs or enter individual uplink labels. For example, LAG1 or uplink1, uplink2.
- A LAG must have two physical NICs on the same N-VDS.
- The number of LAGs that you can actually use depends on the capabilities of the underlying physical environment and the topology of the virtual network. For example, if the physical switch supports up to four ports in an LACP port channel, you can connect up to four physical NICs per host to a LAG.
- In the LACP section, Bare Metal NSX Edge only supports **Source and destination MAC address, IP address and TCP/UDP port**.
- If multiple LAG uplinks are configured on a Bare Metal NSX Edge, enter a unique LAG name for each LAG uplink profile.
- If multi-vtep uplink profile is used for Bare Metal NSX Edge or edge VMs, NSX only supports **Load Balance Source** teaming policy.
- You must use the **Load Balanced Source** teaming policy for traffic load balancing.

Prerequisites

- See NSX Edge network requirements in [NSX Edge Installation Requirements](#).
- 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 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>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System > Fabric > Profiles > Uplink Profiles > Add Profile**.

3 Complete the uplink profile details.

Option	Description
Name and Description	Enter an uplink profile name. Add an optional uplink profile description.
LAGs	<p>(Optional) In the LAGs section, click Add for Link aggregation groups (LAGs) using Link Aggregation Control Protocol (LACP) for the transport network. 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.</p> <p>Possible LAG hashing mechanism options:</p> <ul style="list-style-type: none"> ■ 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 <p>Supported LAG hashing mechanisms on hosts types:</p> <ul style="list-style-type: none"> ■ NSX Edge nodes: Source and destination MAC address, IP address, and TCP/UDP port. ■ ESXi hosts with VDS in Enhanced Networking Stack (ENS) mode: Source MAC address, Destination MAC address, and Source and destination MAC address. ■ ESXi hosts with VDS in Standard mode: Source MAC address, Destination MAC address, Source and destination MAC address, and Source and destination IP address and VLAN. ■ ESXi hosts with vSphere Distributed Switch (v 7.0 and later that supports NSX): LACP is not configured in NSX. You need to configure it in VMware vCenter. ■ Physical server hosts: Source MAC address.
Teamings	<p>In the Teaming section, you can either enter a default teaming policy or you can choose to enter a named teaming policy that is only applicable to VLAN networks. Click Add to add a naming teaming policy. A teaming policy defines how VDS uses its uplink for redundancy and traffic load balancing. You can configure a teaming policy in the following modes:</p> <ul style="list-style-type: none"> ■ Failover Order: Specify an active uplink 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. Standby uplinks and multiple active uplinks are not supported for NSX Edge transport nodes. Also, for an NSX Edge transport node, active uplink used in one profile should must not be used in another profile.

Option	Description
	<ul style="list-style-type: none"> <li data-bbox="635 226 1425 443">■ Load Balance Source: Maps a virtual interface of a VM to an uplink. Traffic sent by this virtual interface will leave the host through this uplink only, and traffic destined to this virtual interface will necessarily enter the virtual switch through this uplink. Select a list of active uplinks. When you configure a transport node, you can pin each interface of the transport node to one active uplink. This configuration allows use of several active uplinks at the same time. No standby uplink is configured in this case. <hr/> <p data-bbox="671 468 1425 554">Important To manage VLAN traffic, if you configure a default teaming policy in Load Balance Source mode, then on failure of the first uplink, traffic will not fail over to the second uplink interface.</p> <hr/> <ul style="list-style-type: none"> <li data-bbox="635 569 1425 655">■ Load Balance Source MAC Address: Select an uplink based on a hash of the source Ethernet. NSX Edge transport nodes do not support this teaming policy.
	<p data-bbox="635 684 687 705">Note</p> <ul style="list-style-type: none"> <li data-bbox="635 720 1425 1108">■ On hypervisor hosts: <ul style="list-style-type: none"> <li data-bbox="671 751 1425 842">■ ESXi hosts: Default teaming policies - Load Balance Source MAC, Load Balance Source, and Failover Order teaming policies are supported. <li data-bbox="671 856 1425 947">■ Physical server hosts (Linux): Only Failover Order teaming policy is supported. Load Balance Source and Load Balance Source MAC teaming policies are not supported. <li data-bbox="671 961 1425 1108">■ Physical server hosts (Windows): Supports Load Balance Source and Load Balance Source Mac teaming policies. Load Balance Source teaming policy on NSX is mapped to Address Hash on Windows. Load Balance Source Mac Address on NSX is mapped to Mac Addresses on Windows. <li data-bbox="635 1123 1425 1213">■ 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.
	<p data-bbox="635 1234 1425 1297">(ESXi hosts and NSX Edge) You can define the following policies for a transport zone:</p> <ul style="list-style-type: none"> <li data-bbox="635 1308 1425 1371">■ A Named teaming policy for every VLAN-based logical switch or segment. <li data-bbox="635 1381 1425 1402">■ A Default teaming policy for the entire VDS. <p data-bbox="635 1413 1425 1560">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.</p> <ul style="list-style-type: none"> <li data-bbox="635 1581 1425 1665">■ If you define a named teaming policy, 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. <li data-bbox="635 1686 1425 1738">■ If you do not define any named teaming policies, VDS uses the default teaming policy. <p data-bbox="635 1749 1425 1770">For more details, see Configure Named Teaming Policy.</p>

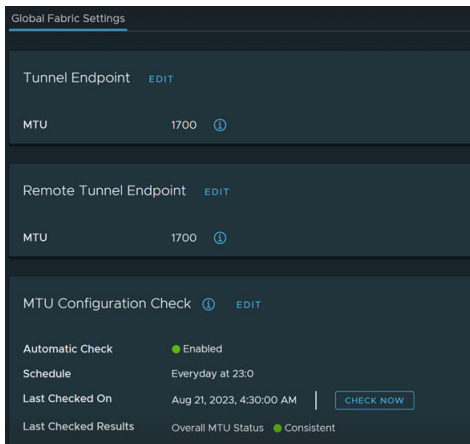
- 4 Enter a Transport VLAN ID value. The transport VLAN set in the uplink profile tags overlay traffic only and the VLAN ID is used by the Tunnel Endpoint Pools (TEP IP Pools).

Important While you can choose any of the available pre-created default uplink profiles, note that you can only edit and configure the transport VLAN ID field to a value of your choice. You cannot edit any other field of a pre-created default uplink profile.

- 5 Enter the MTU value.

For hosts that use vSphere VDS, configure MTU on the VDS from VMware vCenter. The uplink profile MTU default value is 1700 bytes and is applicable to transport-nodes that use N-VDS.

Note The MTU field is optional. If you do not configure it, NSX takes the value set in the Tunnel Endpoint MTU field. If both MTU fields are set, uplink profile MTU value overrides the tunnel endpoint MTU value.

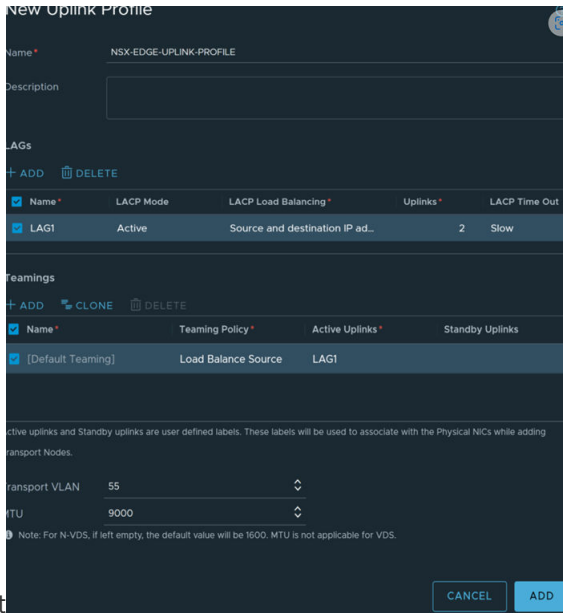


For more information on MTU guidance, see [Guidance to Set Maximum Transmission Unit](#).

An example of Active/Active uplink profile for ESXi host with named-teaming (optional)

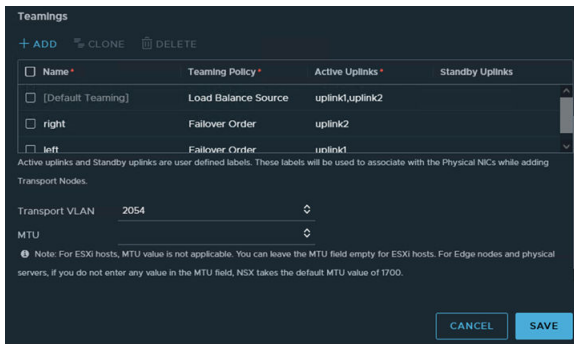


An example of uplink profile using LAG for ESXi



host

An example of uplink profile for NSX Edge with named-teaming policy




6 Configure global tunnel endpoint MTU.

Results

In addition to the UI, you can also view the uplink profiles with the API call `GET /policy/api/v1/infra/host-switch-profiles`.

Configure Named Teaming Policy

On a host running NSX with an N-VDS, you can use the named teaming policies to override the default teaming policy that was configured for some specific VLAN backed segments. This capability is used to steer VLAN traffic to specific uplinks.



In the figure, VMs (VM1, VM2) connected to overlay networks follow default teaming policy and use uplinks u1 and u2. However, an additional teaming policy named **teaming policy VLAN-traffic** is configured to steer vlan traffic (for VMs 'VM3, VM4' connected to VLAN segment) to uplinks u3 and u4 to separate overlay and VLAN traffic.


Prerequisites

Procedure

- 1 Configure an uplink profile, which will be consumed by transport node members of a VLAN Transport zone and define following:
 - a Set a default teaming policy with active uplinks (standby uplinks not supported for NSX Edge Transport nodes).
 - b Set a named teaming policy <p1> with mode 'failover order' and active uplink1.
 - c Set a named teaming policy <p2> with mode 'failover order' and active uplink2 and so on.

Note For overlay segments, the default teaming is used by default. For VLAN segments, when a named teaming policy is defined, NSX ensures that the named teaming policy overrides the default teaming policy.


- d Set the Transport VLAN and MTU.
- e Click **Save**.



- 2 Configure VLAN Transport zone with teaming policy names specified in the uplink profile.



- 3 Create VLAN **segment1** on VLANX and associate with teaming policy name <p1>.



- 4 Create a VLAN **segment2** on VLAN Y and associate it with teaming policy name <p2>.

- Configure a transport node to use uplink profile configured with the named teaming policy.

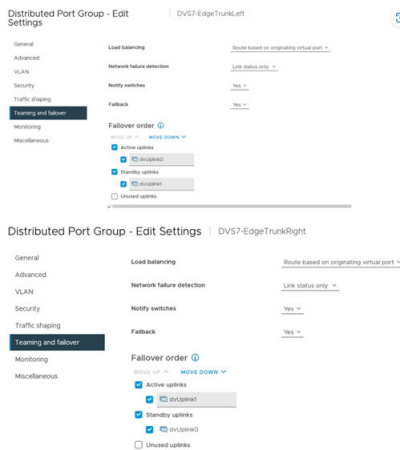
Use a named teaming policy to pin traffic that egresses from a specific uplink and ingresses into a ToR switch .

Note A named teaming policy is only applicable to VLAN segments/logical switches. A named uplink profile only supports the Failover teaming policy. If a VLAN segment does not specify which teaming policy to use, the default teaming policy in the uplink profile will be used.

If you use (Failover) Active/Standby for a VLAN-based segment, the traffic will be always forwarded through the active uplink. Traffic will not failover to the standby uplink. If you want to configure failover, define it in the distributed port group. This is valid only for Edge VM.

NSX teaming policies apply to traffic entering and existing N-VDS but not the traffic exiting the hypervisor configured with PNICs on VDS. Hosts configured with PNICs on VDS, will follow the teaming policy defined in their respective VDS DVPGs in VMware vCenter.

An example of a VDS portgroup teaming that is defined in VMware vCenter for NSX Edge VM.



Add and attach NSX Edge Bridge Profile to a Segment

The NSX Edge bridge profile specifies the primary NSX Edge node that will be the preferred node for the active bridge and backup node that will be preferred for the backup bridge.

NSX Edge Bridge connects NSX overlay logical segment with a traditional VLAN at layer 2. The edge bridge leverages DPDK for high performance forwarding. The traffic bridged in or out of the NSX domain is subject to an edge bridge firewall instance. The NSX Edge bridge functionality is mainly for migration scenarios (physical to virtual or virtual to virtual) or for integration of physical, non-virtualized appliances to the virtualized environment.

Note Starting with NSX 2.5, the same segment can be attached to several bridges on different Edge Clusters and VLANs.

At the time of the creation of the Bridge Profile, no Bridge is instantiated yet. The Bridge Profile is just a template for the creation of one or several Bridge pairs. Once a Bridge Profile is created, you can attach a segment to it. By doing so, an active Bridge instance is created on the primary Edge, while a standby Bridge is provisioned on the backup Edge. NSX creates a Bridge Endpoint object, which represents this pair of Bridges. The attachment of the segment to the Bridge Endpoint is represented by a dedicated logical port.

Prerequisites

- Verify that the NSX Edge cluster is available with minimum of two edge nodes (BareMetal or VM form factor).
- For the VM form factor NSX Edge, verify NSX Edge uplink VDS Trunk port group in VMware vCenter has the following configuration:
 - Forged transmit
 - MAC learning (recommended) or promiscuous mode/sink port configured
 - Active/Standby teaming policy.

Note Consider dedicating an NSX Edge uplink (vNIC) to bridged traffic so that other kinds of traffic to and from the NSX Edge do not suffer from the performance impact related to promiscuous mode.

- Verify NSX Edge TNs member of NSX Edge Cluster are attached to overlay as well as VLAN Transport Zone.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **Networking** → **Segments** → **Profiles** → **Edge Bridge Profiles** → **Add Edge Bridge Profile**.
- 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.

Option	Description
Backup Node	Designate the back up NSX Edge node if the primary node fails.
Failover Mode	<p>Select either Preemptive or Non-Preemptive mode.</p> <p>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.</p> <p>In the preemptive mode, the Bridge on the primary Edge will always become the active bridge forwarding traffic between overlay and VLAN as soon as it is available. In the non-preemptive mode, the Bridge on the primary Edge will remain standby if it becomes available when the Bridge on the backup Edge is already active.</p>

- 4 After you create a Bridge Profile, associate it to a segment.
- 5 Select **Networking > Segments > NSX > Add Segment**.
- 6 Enter the required details, connect to overlay transport zone and click **Save**.
- 7 Edit the segment to which you want to add the Bridge Profile.
- 8 In the Additional Settings section, in the Edge Bridges field, select **Set**.
- 9 Click **Add Edge Bridge**.
- 10 Select the Edge Bridge Profile.
- 11 Select the Transport Zone where the bridged traffic is sent to the N-VDS selected by the transport zone.
- 12 Select the VLAN ID for the VLAN traffic as well as the physical port you select on the NSX Edge for sending or receiving this VLAN traffic.
- 13 (Optional) Select the teaming policy to decide how N-VDS balances traffic across its uplinks.
- 14 Click **Add**.
- 15 Click **Save**.

Results

The newly created NSX Edge Bridge Profile is associated to a segment to balance VLAN traffic.

What to do next

- Verify the configuration and state of L2 bridges on the NSX Edge.
 - a SSH to the NSX Edge as an admin.
 - b Run `get bridge`.
 - c Verify that the Device State is Up.

NSX Edge bridge state is down

NSX Edge bridge state is down.

Problem

NSX Edge bridge state has gone down.

Cause

Failure of VLAN uplink or failure of BFD tunnel between two NSX Edge nodes that are part of the bridge.

Solution

- 1 SSH to NSX Edge as admin.
- 2 Run `get bridge summary` CLI to view existing bridges
- 3 Run `cli get bridge name <name> or get bridge <uuid> or get bridge vlan <vlan-id>` to view edge device state status.
- 4 Run `cli get bridge mac-sync-table` to verify overlay segment mac is learned successfully.
- 5 run `cli get edge-cluster-status` to verify edge status, admin-state, health-check for all uplink interfaces is up with no issues.
- 6 Verify NSX Edge bridge state is now active.

What to do next

If redundant bridges exist in environment with one active (down) and other standby, do the following:

- Run following cli on standby to put bridge state from standby(blocked) to active: `set bridge <uuid> state active` OR
- Run `set bridge vlan <vlan-id> state active`

NSX Edge bridge does not get created

NSX Edge bridge does not get created.

Problem

From the NSX Edge node, when you run `get bridge summary`, the output does not show the bridge is created.

Cause

VLAN ID or uplink conflict.

Solution

- 1 Verify that VLAN ID used for bridging on a specific uplink is not being used by another VLAN segment for configuration of any other feature.
- 2 Verify all NSX Edge transport node member of NSX Edge cluster configured with bridging belong to same overlay and VLAN transport zone.

- 3 Verify the bridge is not configured with multiple VLAN uplinks in a failover teaming policy.

Add an NSX Edge Cluster Profile

The NSX Edge cluster profile defines policies for NSX Edge transport nodes that are part of an NSX Edge cluster.

Prerequisites

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System > Fabric > Profiles > Edge Cluster Profiles > Add Profile**.
- 3 Enter the NSX Edge cluster profile details.

Option	Description
Name and Description	Enter an NSX Edge cluster profile name. Optionally, describe the profile you want to create. For example, Bidirectional Forwarding Detection (BFD) setting for the profile .
BFD Probe Interval	Accept the default setting. BFD is a detection protocol used to identify the forwarding path failures. To detect a forwarding path failure, you can set the interval timing for BFD .
BFD Allowed Hops	Accept the default setting. You can set the maximum number of hops multi-hop BFD sessions can transit.
BFD Declare Dead Multiple	Accept the default setting. You can set the number of times the BFD packet is not received before the session is flagged as down.
Stand By Relocation Threshold	Accept the default setting. The minimum threshold value must be 10 min. The recommended threshold value is 30 min.

Prepare a vSphere Distributed Switch for NSX

Before you configure an NSX transport node using vSphere Distributed Switch (VDS) as a host switch, ensure that the VDS created on a VMware vCenter 7.0 or a later version is configured to manage NSX traffic.

High-level tasks to configure a cluster or a standalone managed host using a VDS switch.

Important To create a VDS switch supporting NSX networking, the following conditions must be met:

- VMware vCenter 7.0 or a later version
- ESXi 7.0 or a later version

Prerequisites

- Verify that ESXi hosts have the required number of physical NICs to meet networking requirements. For example, if you plan to configure teaming policies and remote span port mirroring, ensure that a free physical NIC is available to avoid uplink conflicts.
- Ensure that the MTU value of the physical switch port or LACP port is set to 1700 bytes.

Procedure

- 1 In a VMware vCenter, create a VDS. For more information about creating a VDS, see [Prepare a vSphere Distributed Switch for NSX](#).

- Set the MTU value for the VDS to at least 1600.

Note [Prepare a vSphere Distributed Switch for NSX](#) On VDS 7.0 or later, the default MTU size is 1500. To prepare a VDS for NSX overlay networking, the MTU size of the VDS must be at least 1600. Starting in NSX 3.2.1, if the MTU size of the VDS is below 1600, NSX Manager notifies you that the MTU size will be automatically increased to 1600.

- Add hosts, that you want to prepare for NSX networking, to VDS created in VMware vCenter.
 - Assign Physical NICS of each of the host as uplinks on the VDS.
- 2 In NSX, add an uplink profile that defines a teaming policy mapping NSX uplinks with VDS uplinks.
 - 3 In NSX, prepare an ESXi host using VDS as the host switch.

At the end of the configuration, the host is prepared as NSX transport node with VDS as the host switch.

What to do next

Configure the host as a transport node. See [Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP](#).

Add a Transport Node Profile

A transport node profile is a template to define configuration that is applied to a group of hosts that are part of a VMware vCenter cluster. It is not applied to prepare standalone hosts. Prepare VMware vCenter cluster hosts as transport nodes by applying a transport node profile. Transport node profiles define transport zones, member hosts, switch configuration including uplink profile, IP assignment, mapping of physical NICs to uplink virtual interfaces and so on.

Note Transport node profiles are only applicable to ESXi hosts member of VMware vCenter cluster. It cannot be applied to NSX Edge transport nodes.

Transport node creation begins when a transport node profile is applied to a VMware vCenter cluster. NSX Manager prepares the hosts in the cluster and installs the NSX components on all the hosts. Transport nodes for the hosts are created based on the configuration specified in the transport node profile.

On a cluster prepared with a transport node profile, these outcomes are true:

- When you move an unprepared host into a cluster applied with a transport node profile, NSX automatically prepares the host as a transport node using the transport node profile.
- When you move a transport node from the cluster to an unprepared cluster or directly as a standalone host under the data center, first the transport node configuration applied to the node is removed and then NSX VIBs are removed from the host. See [Triggering Uninstallation from the vSphere Web Client](#).

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.

Points to note when you create a Transport Node Profile:

- You can add a maximum of four VDS switches for each configuration: enhanced VDS created for VLAN transport zone, standard VDS created for overlay transport zone, enhanced VDS created for overlay transport zone.
- There is no limit on the number of standard VDS switches created for VLAN transport zone.
- In a single host cluster topology running multiple standard overlay VDS switches and edge VM on the same host, NSX provides traffic isolation such that traffic going through the first VDS is isolated from traffic going through the second VDS and so on. The physical NICs on each 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 VDS switch name must be unique. NSX does not allow use of duplicate switch names.
- Each transport zone ID associated with each VDS host in a transport node configuration or transport node profile configuration must be unique.

Prerequisites

- Verify that the hosts are part of a VMware vCenter cluster.
- Verify that cluster hosts are member of VDS version 7.0 or later with at least one uplink on VDS port group.
- Verify that a transport zone is configured. See [Create Transport Zones](#).
- Verify that NSX Manager cluster nodes are up and available. To verify cluster status, go to **System** → **Appliances** → **Cluster**. See [Deploy NSX Manager Nodes to Form a Cluster from the 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](#).
- Verify uplink profile to use for Host configuration is configured. See [Create an Uplink Profile](#).

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System > Fabric > Hosts > Transport Node Profile > Add Transport Node Profile**.
- 3 Enter a name to identify the transport node profile.
You can optionally add the description about the transport node profile.
- 4 Click **Set** under Host Switch to add details of the new switch.
- 5 Before you proceed, decide which type of host switch you want to configure on nodes of a cluster.
- 6 Configure the following fields:

Option	Description
Name	(Hosts managed by a vSphere cluster) Select the VMware vCenter that manages the host switch. Select the VDS that is created in VMware vCenter and attached to your ESXi hosts.
Transport Zones	In the Show section, select Overlay, VLAN or All to view and select the type of transport zones you want for the host switch. These transport zones are realized by associated host switches. Supported transport zone configurations: <ul style="list-style-type: none"> ■ You can add multiple VLAN transport zones per host switch. ■ You must add only one overlay transport zone per host switch. NSX Manager UI does not allow adding multiple overlay transport zones.
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. If you keep the MTU value empty, the NSX takes the global default MTU value 1700. If you enter a MTU value in NSX uplink profile, that MTU value will override the global default MTU value. Note Link Aggregation Groups defined in an uplink profile cannot be mapped to VDS uplinks.
IP Address Type (TEP)	Select between IPv4 and IPv6 to specify the IP version for the tunnel endpoints (TEPs) of the transport node.
IPv4 Assignment	Choose how IPv4 addresses are assigned to the TEPs. The options are: <ul style="list-style-type: none"> ■ Use DHCP: IPv4 addresses are assigned from a DHCP server. ■ Use IPv4 Pool: IPv4 addresses are assigned from an IP pool. Specify the IPv4 pool name to be used for TEPs.


Option	Description
IPv6 Assignment	<p>Choose how IPv6 addresses are assigned to the TEPs. The options are:</p> <ul style="list-style-type: none"> ■ Use DHCPv6: IPv6 addresses are assigned from a DHCP server. ■ Use IPv6 Pool: IPv6 addresses are assigned from an IP pool. Specify the IPv6 pool name to be used for TEPs. ■ Use AutoConf: IPv6 addresses are assigned from Router Advertisement (RA).
Advanced Configuration Mode (NSX 4.0.0.1 only)	<p>Choose between the following mode options:</p> <ul style="list-style-type: none"> ■ Standard: Is the standard mode that is available to all supported hypervisors by NSX. ■ ENS Interrupt: Is a variant of the Enhanced Datapath mode. ■ Enhanced Datapath: Is the mode that provides accelerated networking performance. This mode requires nodes to use VMXNET3 vNIC enabled network cards. It is not supported on NSX Edge nodes and Public Gateways. The supported hypervisor is ESXi. It is recommended to run ESXi v6.7 U2 and later versions.
Mode (Starting with NSX 4.0.1.1)	<p>Choose between the following mode options:</p> <ul style="list-style-type: none"> ■ Standard: This mode applies to all transport nodes. The data-plane in the transport node automatically selects the host switch mode as per the uplink capabilities. ■ Enhanced Datapath - Standard: This mode is a variant of the Enhanced Data Path mode. It is available only on ESXi hypervisor 7.0 and later versions. Please consult your account representative for applicability. ■ Enhanced Datapath - Performance: This is the Enhanced Data Path switch mode for ESXi host transport node. This mode provides accelerated networking performance. It requires nodes to use VMXNET3 vNIC enabled network cards. It is not supported on NSX Edge nodes and Public Gateways. The supported hypervisor is ESXi. It is recommended to run ESXi v6.7 U2 and later versions. ■ Legacy: This mode was formerly called Standard. You can select this mode only through API, since the Legacy field is read-only in NSX Manager UI. It applies to all transport nodes. When the host switch mode is set to Legacy, the packet handler stack is enabled. On NSX Manager UI, you will see this mode set to Standard and the Legacy field to 'Yes'. <p>You can run the following Host Transport Node or Transport Node Profile policy API to set the host switch mode to Legacy:</p> <ul style="list-style-type: none"> ■ Create or update Host Transport Node: <div data-bbox="710 1493 1430 1629" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>PUT https://<NSX-Manager-IP-ADDRESS>/POST/policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id></pre> </div> ■ Create or update policy Host Transport Node Profile: <div data-bbox="710 1682 1430 1795" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>PUT https://<NSX-Manager-IP-ADDRESS>/POST/policy/api/v1/infra/host-transport-node-profiles/<transport-node-profile-id></pre> </div>

Option	Description
<p>CPU Config</p>	<p>You can configure the CPU Config field only when the Mode is set to Enhanced Datapath.</p> <ol style="list-style-type: none"> Click Set. In the CPU Config window, click Add. Enter values for the NUMA Node Index and LCores per NUMA Node fields. To save the values, click Add and Save.
<p>Teaming Policy Uplink Mapping</p>	<p>Before you map uplinks profiles in NSX with uplinks in VDS, ensure uplinks are configured on the VDS switch. To configure or view the VDS switch uplinks, go to VMware vCenter → <i>vSphere Distributed Switch</i>. Click Actions → Settings → Edit Settings.</p> <p>Map uplinks defined in the selected NSX uplink profile with VDS uplinks. The number of NSX uplinks that are presented for mapping depends on the uplink profile configuration.</p> <p>For example, in the uplink-1 (active) row, go to the Physical NICs column, click the edit icon, and type in the name of VDS uplink to complete mapping it with uplink-1 (active). Likewise, complete mapping for the other uplinks.</p>

Note Uplinks/LAGs, NIOC profile, LLDP profile are defined in VMware vCenter. These configurations are not available in NSX Manager. To manage VMkernel adapters on a VDS switch, go to VMware vCenter to attach VMkernel adapters to Distributed Virtual port groups or NSX port groups.

- If you have selected multiple transport zones, you can add them to the same switch. To add a new switch, click **Add Switch** again to configure a new switch for the other transport zones.

NSX switches can attach to a single overlay transport zone and multiple VLAN transport zones at the same time.



- Click **Add** to complete the configuration.

What to do next

Apply the transport node profile to an existing vSphere cluster. See [Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP](#).

Host Transport Nodes

9

You can prepare ESXi hosts and physical servers as NSX transport nodes. Before you prepare physical servers for NSX networking, ensure the required third-party packages are installed on hosts.

Read the following topics next:

- [Preparing ESXi Hosts as Transport Nodes](#)
- [Manual Installation of NSX Kernel Modules](#)
- [Preparing Physical Servers as NSX Transport Nodes](#)
- [Secure Workloads on Windows Server 2016/2019 Bare Metal Servers](#)
- [Configure an ESXi Host Transport Node with Link Aggregation Group](#)
- [Quick Start Wizard to Prepare ESXi Cluster hosts for Security-only or Networking and Security](#)
- [Deploy a Fully Collapsed vSphere Cluster NSX on Hosts Running N-VDS Switches](#)
- [Multiple NSX Managers Managing a Single VMware vCenter](#)
- [Troubleshoot Multi-NSX Issues](#)
- [Managing Transport Nodes](#)

Preparing ESXi Hosts as Transport Nodes

After you create transport zones, IP pools, uplink profiles for transport nodes, prepare a host as a transport node. An ESXi or Bare Metal (physical server) can be prepared as a transport node. ESXi hosts only support vSphere Distributed Switch (VDS), while Physical servers only support N-VDS host switch type.

Transport Node Profiles

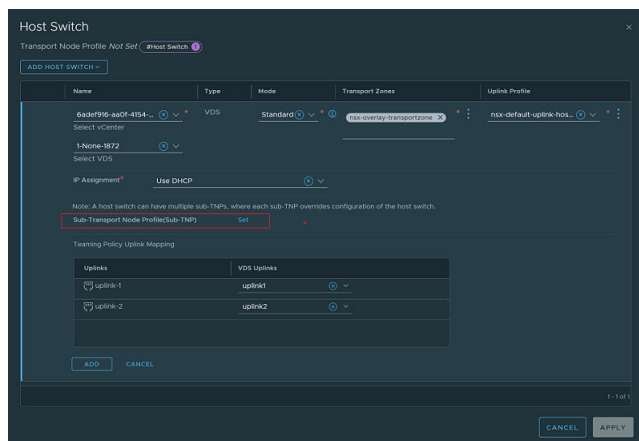
A Transport Node Profile (TNP) is a template to define networking configuration that is applied to a cluster.

Transport Node Profile

Before you create a Transport Node Profile consider these points:

- You can add a maximum of four N-VDS or VDS switches for each configuration: enhanced N-VDS or VDS created for VLAN transport zone, standard N-VDS or VDS created for overlay transport zone, enhanced N-VDS or VDS created for overlay transport zone.
- There is no limit on the number of standard N-VDS switches created for VLAN transport zone.
- Each N-VDS switch name must be unique. NSX does not allow use of duplicate switch names.
- Each transport zone ID associated with each N-VDS or VDS host in a transport node configuration or transport node profile configuration must be unique.

To create a TNP, you configure Host Switches with transport zones, uplink profiles, mapping uplinks to VDS uplinks (for a VDS switch) and other configurations. TNP can be created on VDS or N-VDS Host Switches.



A stretched cluster is a cluster that extends across multiple TEP subnets. A non-stretched cluster is a cluster that is confined to a single TEP subnet. In a non-stretched cluster, a single TNP is sufficient to be applied to a cluster. However, when a cluster is stretched across multiple subnets or L3 domains, you can create sub-clusters consisting of hosts that need the same configuration.

While a TNP represents the global configuration applied to a Host Switch, a sub-TNP represents the local configuration applied to a sub-cluster. When you apply a sub-TNP to a sub-cluster, all the configuration from sub-TNP takes precedence over the host switch configuration.

Note A Sub-TNP can only be created for vSphere Distributed Switch (VDS) switches.

Transport node creation begins when a transport node profile is applied to a VMware vCenter cluster. NSX Manager prepares the hosts in the cluster and installs the NSX components on all the hosts. Transport nodes for the hosts are created based on the configuration specified in the transport node profile.

Note TNP is not used to prepare standalone hosts.

On a cluster prepared with a transport node profile, these outcomes are true:

- When you move an unprepared host into a cluster applied with a transport node profile, NSX automatically prepares the host as a transport node using the transport node profile.
- When you move a transport node from the cluster to an unprepared cluster or directly as a standalone host under the data center, first the transport node configuration applied to the node is removed and then NSX VIBs are removed from the host. See [Triggering Uninstallation from the vSphere Web Client](#).

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.

Sub-TNPs and Sub-clusters

A Sub-TNP is a template to define configuration that is applied to hosts that are part of a Sub-cluster. To accommodate different cluster configurations required for sub-clusters, create sub-TNPs.


Sub-Transport Node Profile

Starting with NSX 3.2.2, you can configure Sub-TNPs and Sub-clusters.

Note The maximum number of Sub-TNPs under a hostswitch in a TNP is 16.

While a TNP represents the global configuration applied to a Host Switch, a sub-TNP represents the local configuration applied to a sub-cluster. When you apply a sub-TNP to a sub-cluster, the host switch is overridden on the sub-cluster.

You can only apply a sub-TNP to a VDS switch. A sub-TNP can only override the following fields of a host switch: VDS Host Switches ID, uplink profiles and IP assignment.



NSX Manager prepares the hosts in sub-clusters where the sub-TNP is applied and installs the NSX 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 Sub-TNP, you must first detach the global TNP 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.

Sub-clusters

Each cluster can have up to 16 sub-clusters.

To manage a cluster that is stretched across different subnets or L3 domains, some hosts will need different configuration, such as some hosts will need to be in the TEP subnet-1, whereas some other hosts will need to be in TEP subnet-2. So, hosts with similar configuration requirements can be placed in one sub-cluster. And that sub-cluster can be applied with a sub-TNP that provides configuration.

In the following image, the cluster has three hosts: one in each sub-cluster and the last host is not part of any sub-cluster.



Moving Hosts Between Sub-clusters

- 1 Move host from sub-cluster1 to sub-cluster2.
- 2 Sub-TNP config of sub-cluster1 is removed from the host.
- 3 Sub-TNP config of sub-cluster2 is applied to the host.

Limitations of using Sub-clusters

- Sub-clusters can have hosts only from the cluster under which it is created.
- Sub-cluster cannot have hosts which are part of other sub-clusters.
- Maximum number of sub-clusters under a cluster is 16.
- Sub-cluster cannot be deleted if it has some hosts in it. After removing the hosts, sub-cluster can be deleted.
- Minimum supported ESXi version for this feature is 7.0.0.
- Stateless hosts cannot be added to a sub-cluster.

Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP

If a cluster of ESXi hosts is registered to a VMware vCenter, you can apply transport node profiles on the VMware vCenter cluster to prepare all hosts part of the cluster as NSX transport nodes or you can prepare each host individually.

Note (Host in lockdown mode) If your exception list for vSphere lockdown mode includes expired user accounts, NSX installation on vSphere fails. If your host is part of the vLCM-enabled cluster, several users such as `lldp-vim-user`, `nsx-user`, `mux-user`, and `da-user`, are created automatically and added to the exception users list on an ESXi host when NSX VIBS are installed. Ensure that you delete all expired user accounts before you begin installation. For more information on accounts with access privileges in lockdown mode, refer to *Specifying Accounts with Access Privileges in Lockdown Mode* in the *vSphere Security Guide*. For more details on these NSX user accounts on the ESXi host, refer to the KB article, <https://kb.vmware.com/s/article/87795>.

Prerequisites


- Verify all hosts that you want to configure as transport nodes are powered on in VMware vCenter.
- Verify all hosts are members of VDS in VMware vCenter with correct uplinks.
- Verify that the system requirements are met. See [System Requirements](#).
- Verify vCenter is added as compute manager to NSX Manager.
- Verify NSX Manager cluster is up and stable.
 - UI: Go to **System** → **Appliances** → **NSX Appliances**.
 - CLI: SSH to one of the NSX Manager nodes as an admin and run `get cluster status`.
- The reverse proxy service on all nodes of the NSX Manager cluster must be `Up` and running. To verify, run `get service http`. If the service is down, restart the service by running `restart service http` on each NSX Manager node. If the service is still down, contact VMware support.
- If you deployed VMware vCenter on a custom port or a non-default port, apply these rules to NSX Manager:
 - IPv4 rules must be applied on NSX Manager manually before starting the host preparation.


```
iptables -A INPUT -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT
```

- `iptables -A OUTPUT -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT`
- IPv6 table rules must be applied on NSX Manager manually before starting the host preparation.
 - `ip6tables -A OUTPUT -o eth0 -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT`
 - `ip6tables -A INPUT -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT`
- 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>` or `https://<nsx-manager-fqdn>`.
- 2 Select **System > Fabric > Hosts**.
- 3 On the **Cluster** tab, select a cluster and click **Configure NSX**.
- 4 In the **NSX Installation** pop-up window, from the **Transport Node Profile** drop-down menu, select the transport node profile to apply to the cluster. If a transport node is not created, click **Create New Transport Node Profile** to create a new one.



- 5 Click **Apply** to begin the process of transport node creation of all hosts in the cluster. See [Add a Transport Node Profile](#).
- 6 If you only want to prepare individual hosts as transport nodes, click the menu icon (dots) for the host and select **Configure NSX**.
- 7 Verify the host name in the Host Details panel, and click **Next**. Optionally, you can add a description.
- 8 In the **Configure NSX** panel, click **Add Host Switch**.

9 Configure the following fields:

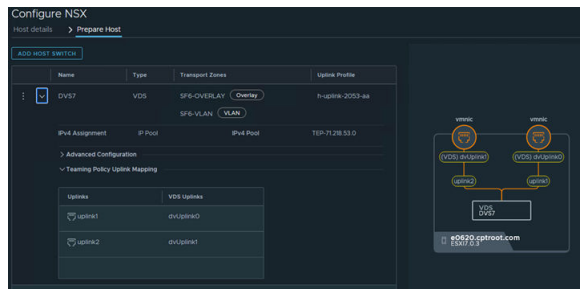
Option	Description
Name	<p>(Hosts managed by a vSphere cluster) Select the VMware vCenter that manages the host switch.</p> <p>Select the VDS that is created in VMware vCenter and attached to your ESXi hosts.</p>
Transport Zones	<p>Shows transport zones that are realized by associated host switches.</p> <p>Supported transport zone configurations:</p> <ul style="list-style-type: none"> ■ You can add multiple VLAN transport zones per host switch. ■ You must add only one overlay transport zone per host switch. NSX Manager UI does not allow adding multiple overlay transport zones.
Uplink Profile	<p>Select an existing uplink profile from the drop-down menu or create a custom uplink profile. You can also use the default uplink profile.</p> <p>If you keep the MTU value empty, the NSX takes the global default MTU value 1700. If you enter a MTU value in NSX uplink profile, that MTU value will override the global default MTU value. See Guidance to Set Maximum Transmission Unit.</p> <p>Note Link Aggregation Groups defined in an uplink profile cannot be mapped to VDS uplinks.</p>
IPv4 Assignment	<p>This field appears when the forwarding mode of the selected transport zones are set to IPv4. For details on configuring the forwarding mode of transport zones, see Create Transport Zones.</p> <p>Choose how IPv4 addresses are assigned to the TEPs. The options are:</p> <ul style="list-style-type: none"> ■ Use DHCP: IPv4 addresses are assigned from a DHCP server. ■ Use IPv4 Pool: IPv4 addresses are assigned from an IP pool. Specify the IPv4 pool name to be used for TEPs. ■ Use Static List: IPv4 addresses are assigned from a static list. Specify the static list, IPv4 gateway, and the subnet mask.
IPv6 Assignment	<p>Important For ESXi host TEPs to use IPv6, it is required that the ESXi version be 8.0 Update 1 or later.</p> <p>This field appears when the forwarding mode of the selected transport zones are set to IPv6. For details on configuring the forwarding mode of transport zones, see Create Transport Zones.</p> <p>Choose how IPv6 addresses are assigned to the TEPs. The options are:</p> <ul style="list-style-type: none"> ■ Use DHCPv6: IPv6 addresses are assigned from a DHCP server. ■ Use IPv6 Pool: IPv6 addresses are assigned from an IP pool. Specify the IPv6 pool name to be used for TEPs. ■ Use AutoConf: IPv6 addresses are assigned from Router Advertisement (RA).

Option	Description
Advanced Configuration Mode (NSX 4.0.0.1 only)	<p>Choose between the following mode options:</p> <ul style="list-style-type: none"> ■ Standard: Is the standard mode that is available to all supported hypervisors by NSX. ■ ENS Interrupt: Is a variant of the Enhanced Datapath mode. ■ Enhanced Datapath: Is the mode that provides accelerated networking performance. This mode requires nodes to use VMXNET3 vNIC enabled network cards. It is not supported on NSX Edge nodes and Public Gateways. The supported hypervisor is ESXi. It is recommended to run ESXi v6.7 U2 and later versions.
Mode (Starting with NSX 4.0.1.1)	<p>Choose between the following mode options:</p> <ul style="list-style-type: none"> ■ Standard: This mode applies to all transport nodes. The data-plane in the transport node automatically selects the host switch mode as per the uplink capabilities. ■ Enhanced Datapath - Standard: This mode is a variant of the Enhanced Data Path mode. It is available only on ESXi hypervisor 7.0 and later versions. Please consult your account representative for applicability. ■ Enhanced Datapath - Performance: This is the Enhanced Data Path switch mode for ESXi host transport node. This mode provides accelerated networking performance. It requires nodes to use VMXNET3 vNIC enabled network cards. It is not supported on NSX Edge nodes and Public Gateways. The supported hypervisor is ESXi. It is recommended to run ESXi v6.7 U2 and later versions. ■ Legacy: This mode was formerly called Standard. You can select this mode only through API, since the Legacy field is read-only in NSX Manager UI. It applies to all transport nodes. When the host switch mode is set to Legacy, the packet handler stack is enabled. On NSX Manager UI, you will see this mode set to Standard and the Legacy field to 'Yes'. <p>You can run the following Host Transport Node or Transport Node Profile policy API to set the host switch mode to Legacy:</p> <ul style="list-style-type: none"> ■ Create or update Host Transport Node: <div data-bbox="710 1270 1430 1402" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>PUT https://<NSX-Manager-IP-ADDRESS>/POST/policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id></pre> </div> ■ Create or update policy Host Transport Node Profile: <div data-bbox="710 1459 1430 1566" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>PUT https://<NSX-Manager-IP-ADDRESS>/POST/policy/api/v1/infra/host-transport-node-profiles/<transport-node-profile-id></pre> </div>

Option	Description
CPU Config	<p>You can configure the CPU Config field only when the Mode is set to Enhanced Datapath.</p> <ol style="list-style-type: none"> Click Set. In the CPU Config window, click Add. Enter values for the NUMA Node Index and LCores per NUMA Node fields. To save the values, click Add and Save.
Teaming Policy Uplink Mapping	<p>Before you map uplinks profiles in NSX with uplinks in VDS, ensure uplinks are configured on the VDS switch. To configure or view the VDS switch uplinks, go to VMware vCenter → <i>vSphere Distributed Switch</i>. Click Actions → Settings → Edit Settings.</p> <p>Map uplinks defined in the selected NSX uplink profile with VDS uplinks. The number of NSX uplinks that are presented for mapping depends on the uplink profile configuration.</p> <p>For example, in the uplink-1 (active) row, go to the Physical NICs column, click the edit icon, and type in the name of VDS uplink to complete mapping it with uplink-1 (active). Likewise, complete mapping for the other uplinks.</p>

- If you have selected multiple transport zones, click **+ Add Switch** again to configure the switch for the other transport zones.
- Click **Add** to complete the configuration.
- (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
nsx-proxy
```



- 13 On the **Cluster** tab, verify that the NSX Manager connectivity status of hosts in the cluster is Up and NSX configuration state is Success. During the configuration process, each transport node displays the percentage of progress of the installation process. If installation fails at any stage, you can restart the process by clicking the **Resolve** link that is available against the failed stage of the process.

You can also see that the transport zone is applied to the hosts in the cluster.

Note

- If you individually prepare a host that is already prepared using TNP, then the configuration state of the node shows a `Configuration Mismatch` since configuration does not match the applied TNP.
 - The Host Transport Node page displays TEP addresses of the host in addition to management IP addresses. TEP address is the address assigned to the VMkernel NIC of the host.
-

- 14 (Optional) Remove an NSX VIBs on the host.

- a Select one or more hosts or cluster with applied TNP and click **Actions > Remove NSX**.

The uninstallation takes up to three minutes. Uninstallation of NSX removes the transport node configuration on hosts and the host is detached from the transport zone(s) and switch. Similar to the installation process, you can follow the percentage of the uninstallation process completed on each transport node. If uninstallation fails at any stage, you can restart the process by clicking the **Resolve** link that is available against the failed stage of the process.

- 15 (Optional) Remove a transport node from the transport zone.

- a Select a single transport node and click **Actions > Manage Transport Zone** and choose the transport zone to remove.

What to do next

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 Edge nodes and hosts join the management plane, the NSX logical entities and configuration state are pushed to the NSX Edge nodes and hosts automatically. You can create transport zones, logical switches, logical routers, and other network components through the NSX Manager UI or API at any time. When the hosts are transport nodes, these entities gets realized on the host.

Create a logical switch and assign logical ports. See the Advanced Switching section in the *NSX Administration Guide*.

Change Sub-cluster

When you want to apply a different networking configuration to a host in a sub-cluster, such as a different uplink profile, move the host to a different Sub-cluster.

To change Sub-cluster a host belongs to, perform these steps:

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System** → **Fabric** → **Hosts**.
- 3 Select the **Clusters** tab.
- 4 Select a node and click **Actions** → **Change Sub-cluster**.
- 5 In the Change Sub-cluster window, select a Sub-cluster that the host must be moved to.
- 6 Click **Save**.

Detach Cluster TNP

When you want to apply a different cluster TNP, detach the existing one before applying a new TNP.

To detach cluster TNP, perform these steps:

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System** → **Fabric** → **Hosts**.
- 3 Select the **Clusters** tab.
- 4 Select a cluster and click **Actions** → **Detach Transport Node Profile**.

NSX detached the existing cluster TNP. At the end of the action, the cluster will not have any TNP applied to it.

Migrate VMkernels and Physical NICs to a vSphere Distributed Switch

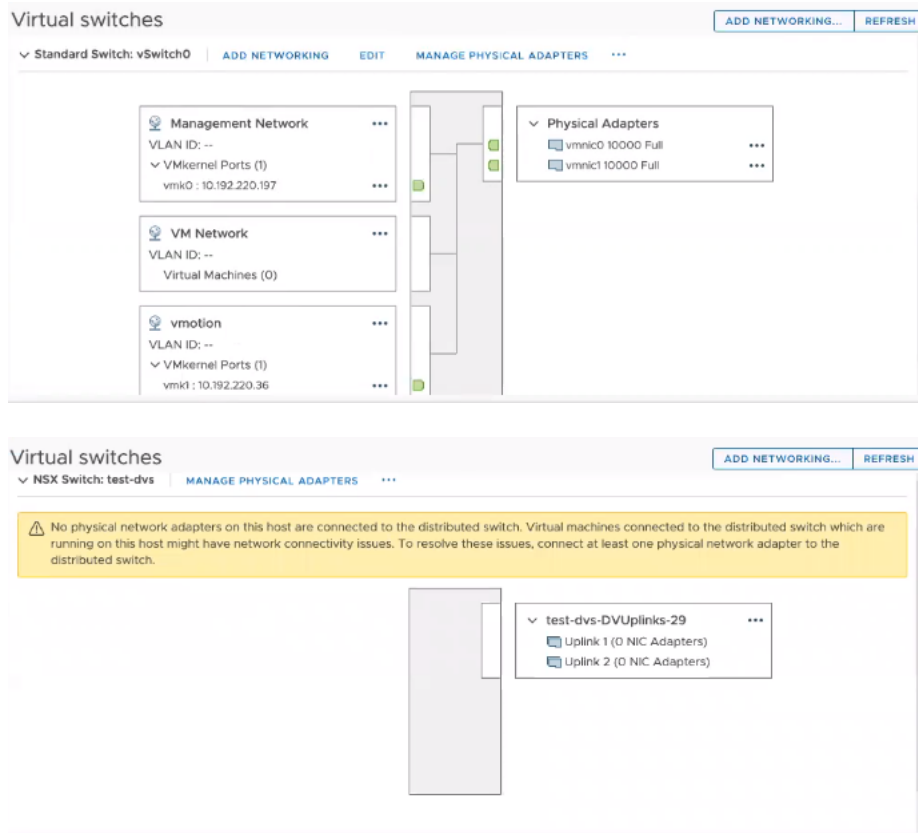
Manually migrate VMkernel adapters from a vSphere Standard Switch or from an N-VDS switch to a vSphere Distributed Switch.

Note Starting with NSX 3.0, transport nodes can be created using vSphere Distributed Switch.

After preparing the transport node with vSphere Distributed Switch host switch type (referred to as an NSX Switch in VMware vCenter), manually migrate VMkernel adapters (vmks) and physical NICs (vmnics) to an NSX Switch on the ESXi host.

In the procedure below, consider this switch configuration:

- vmk0, vmk1 are connected to vSwitch0, and vmnic0, vmnic1 are configured as uplink 1 and 2 respectively on the vSwitch0.
- NSX Switch does not have any vmnic or VMkernel adapter configured.



At the end of the procedure, vmnic0, vmnic1 and vmk0, vmk1 are migrated to vSphere Distributed Switch (referred to as an NSX Switch in VMware vCenter).

Prerequisites

- ESXi hosts are prepared as transport nodes using vSphere Distributed Switch.

Procedure

- 1 From a browser, log in with admin privileges to a VMware vCenter at <https://<vCenterServer-ip-address>>.
- 2 Navigate to **Host** → **Configure** → **Virtual Switches**.
- 3 View existing vmknics configured on vSwitch0.
- 4 Make a note of the vmknics to be migrated to the distributed virtual port group of the NSX Switch.
- 5 Navigate to **Home** → **Networking**, to view all switches configured in the data center.
- 6 In the Switch page, click **Actions** → **Add and Manage Hosts**.
- 7 Select **Manage Host Networking**.
- 8 Click **Next**.
- 9 In the Select Member Hosts window, select **hosts**.

- 10 Click **Ok**.
- 11 In the Manage physical adapters window, claim unassigned adapters, as there are available vmnics that can be attached to a switch.
 - a Select an unclaimed uplink and click **Assign uplink**.
 - b Map a vmnic to an uplink on the NSX Switch.
 - c Click **Ok**.
- 12 In the Manage VMkernel adapters window, assign port groups to NSX Switch.
 - a Select a vmk on vSwitch0 and click **Assign port group**.
 - b Select a NSX port group to assign a vmk to an NSX segment.
 - c Perform steps a and b for the remaining hosts that are managed by the switch.
- 13 Finish the Add and Manage Hosts wizard.
- 14 To verify vmk0 and pnics are migrated from vSwitch0 to NSX Switch on the ESXi host, navigate to **Host** → **Configure** → **Virtual Switches**. View the updated switch configuration.
- 15 Alternatively, run the API command, `https://<NSXManager-IP-address>/api/v1/logical-ports`, to verify migration of VMkernel adapters is successful.

Note All vmk0 ports are set to `Unblocked VLAN` state because management traffic and services are managed by vmk0 ports. These vmk0 ports in `Unblocked VLAN` state allows admins to connect to the vmk0 port if hosts lose connectivity.

What to do next

Navigate to NSX Manager. In the **System > Fabric > Hosts > Clusters** tab, verify configuration status changed from `Degraded` to `Success`, as vmnics and vmks are migrated to the NSX Switch, the vSphere Distributed Switch.

Prepare ESXi Individual Hosts as Transport Nodes

You can configure NSX on individual ESXi hosts

Prerequisites

- Verify that the individual host you want to prepare is powered on.
- Verify that the system requirements are met. See [System Requirements](#).
- The reverse proxy service on all nodes of the NSX Manager cluster must be `Up` and running.

To verify, run `get service http`. If the service is down, restart the service by running `restart service http` on each NSX Manager node. If the service is still down, contact VMware support.

- If you deployed VMware vCenter on a custom port or a non-default port, apply these rules to NSX Manager:
 - IPv4 rules must be applied on NSX Manager manually before starting the host preparation.
 - `iptables -A INPUT -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT`
 - `iptables -A OUTPUT -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT`
 - IPv6 table rules must be applied on NSX Manager manually before starting the host preparation.
 - `ip6tables -A OUTPUT -o eth0 -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT`
 - `ip6tables -A INPUT -p tcp -m tcp --dport <CUSTOM_PORT> --tcp-flags FIN,SYN,RST,ACK SYN -j ACCEPT`
- (Host in lockdown mode) If your exception list for vSphere lockdown mode includes expired user accounts such as lldpvim-user, NSX installation on vSphere fails. This user automatically gets created on ESXi to talk to hostd to get the LLDP neighbor information and then gets deleted. Ensure that you delete all expired user accounts before you begin installation. For more information on accounts with access privileges in lockdown mode, see *Specifying Accounts with Access Privileges in Lockdown Mode* in the *vSphere Security Guide*.

Procedure

- 1 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:A2:4E:3C:C4:F4
```

- 2 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 3 Select **System > Fabric > Hosts**.

- 4 Select **Other Nodes** and select a host.
- 5 On the **Host Details** page, enter details for the following fields.

Option	Description
Name and	Enter the name to identify the standalone host.
IP Addresses	Enter the host IP address.
Description	You can optionally add the description of the operating system used for the host.
Tags	Enter a tag that you want to associate with the host. A tag can be used when you want to group all hosts having a certain OS version, ESXi version, and so on.

- 6 Click **Next**.
- 7 On the **Prepare Host** tab, click **Add Host Switch**.
- 8 From the **Select VDS** drop-down menu, select a VDS switch.
- 9 Configure the following fields:

Option	Description
Name	(Hosts managed by a vSphere cluster) Select the VMware vCenter that manages the host switch. Select the VDS that is created in VMware vCenter and attached to your ESXi hosts.
Transport Zones	In the Show section, select Overlay , VLAN or All to view and select the type of transport zones you want for the host switch. These transport zones are realized by associated host switches. Supported transport zone configurations: <ul style="list-style-type: none"> ■ You can add multiple VLAN transport zones per host switch. ■ You must add only one overlay transport zone per host switch. NSX Manager UI does not allow adding multiple overlay transport zones.
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. If you keep the MTU value empty, the NSX takes the global default MTU value 1700. If you enter a MTU value in NSX uplink profile, that MTU value will override the global default MTU value. Note Link Aggregation Groups defined in an uplink profile cannot be mapped to VDS uplinks.
IP Address Type (TEP)	Select between IPv4 and IPv6 to specify the IP version for the tunnel endpoints (TEPs) of the transport node.
IPv4 Assignment	Choose how IPv4 addresses are assigned to the TEPs. The options are: <ul style="list-style-type: none"> ■ Use DHCP: IPv4 addresses are assigned from a DHCP server. ■ Use IPv4 Pool: IPv4 addresses are assigned from an IP pool. Specify the IPv4 pool name to be used for TEPs.

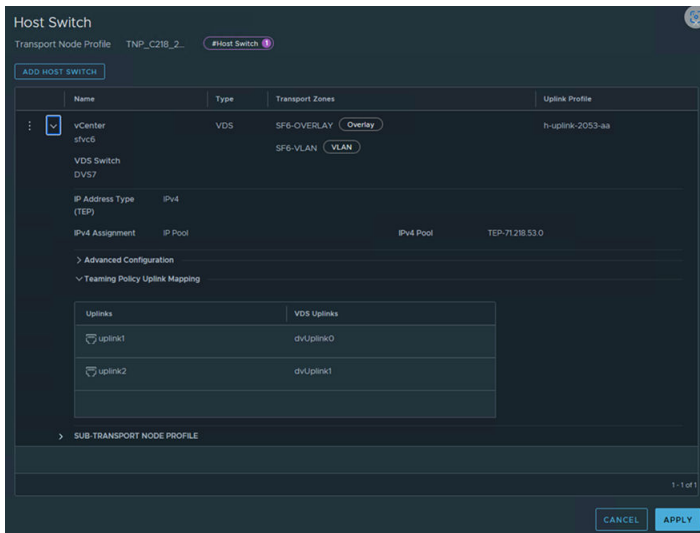
Option	Description
IPv6 Assignment	Choose how IPv6 addresses are assigned to the TEPs. The options are: <ul style="list-style-type: none"> ■ Use DHCPv6: IPv6 addresses are assigned from a DHCP server. ■ Use IPv6 Pool: IPv6 addresses are assigned from an IP pool. Specify the IPv6 pool name to be used for TEPs. ■ Use AutoConf: IPv6 addresses are assigned from Router Advertisement (RA).
Advanced Configuration Mode (NSX 4.0.0.1 only)	Choose between the following mode options: <ul style="list-style-type: none"> ■ Standard: Is the standard mode that is available to all supported hypervisors by NSX. ■ ENS Interrupt: Is a variant of the Enhanced Datapath mode. ■ Enhanced Datapath: Is the mode that provides accelerated networking performance. This mode requires nodes to use VMXNET3 vNIC enabled network cards. It is not supported on NSX Edge nodes and Public Gateways. The supported hypervisor is ESXi. It is recommended to run ESXi v6.7 U2 and later versions.
Mode (Starting with NSX 4.0.1.1)	Choose between the following mode options: <ul style="list-style-type: none"> ■ Standard: This mode applies to all transport nodes. The data-plane in the transport node automatically selects the host switch mode as per the uplink capabilities. ■ Enhanced Datapath - Standard: This mode is a variant of the Enhanced Data Path mode. It is available only on ESXi hypervisor 7.0 and later versions. Please consult your account representative for applicability. ■ Enhanced Datapath - Performance: This is the Enhanced Data Path switch mode for ESXi host transport node. This mode provides accelerated networking performance. It requires nodes to use VMXNET3 vNIC enabled network cards. It is not supported on NSX Edge nodes and Public Gateways. The supported hypervisor is ESXi. It is recommended to run ESXi v6.7 U2 and later versions. ■ Legacy: This mode was formerly called Standard. You can select this mode only through API, since the Legacy field is read-only in NSX Manager UI. It applies to all transport nodes. When the host switch mode is set to Legacy, the packet handler stack is enabled. On NSX Manager UI, you will see this mode set to Standard and the Legacy field to 'Yes'. <p>You can run the following Host Transport Node or Transport Node Profile policy API to set the host switch mode to Legacy:</p> <ul style="list-style-type: none"> ■ Create or update Host Transport Node: <pre data-bbox="715 1493 1430 1629">PUT https://<NSX-Manager-IP-ADDRESS>/POST/policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id></pre> ■ Create or update policy Host Transport Node Profile: <pre data-bbox="715 1682 1430 1787">PUT https://<NSX-Manager-IP-ADDRESS>/POST/policy/api/v1/infra/host-transport-node-profiles/<transport-node-profile-id></pre>

Option	Description
CPU Config	<p>You can configure the CPU Config field only when the Mode is set to Enhanced Datapath.</p> <ol style="list-style-type: none"> Click Set. In the CPU Config window, click Add. Enter values for the NUMA Node Index and LCores per NUMA Node fields. To save the values, click Add and Save.
Teaming Policy Uplink Mapping	<p>Before you map uplinks profiles in NSX with uplinks in VDS, ensure uplinks are configured on the VDS switch. To configure or view the VDS switch uplinks, go to VMware vCenter → <i>vSphere Distributed Switch</i>. Click Actions → Settings → Edit Settings.</p> <p>Map uplinks defined in the selected NSX uplink profile with VDS uplinks. The number of NSX uplinks that are presented for mapping depends on the uplink profile configuration.</p> <p>For example, in the uplink-1 (active) row, go to the Physical NICs column, click the edit icon, and type in the name of VDS uplink to complete mapping it with uplink-1 (active). Likewise, complete mapping for the other uplinks.</p>

Note Uplinks/LAGs, NIOC profile, LLDP profile are defined in VMware vCenter. These configurations are not available in NSX Manager. To manage VMkernel adapters on a VDS switch, go to VMware vCenter to attach VMkernel adapters to Distributed Virtual port groups or NSX port groups.

- If you have selected multiple transport zones, you can add them to the same switch. To add a new switch, click **Add Switch** again to configure a new switch for the other transport zones.

NSX switches can attach to a single overlay transport zone and multiple VLAN transport zones at the same time.



- Click **Add** to complete the configuration.

12 (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
nsx-proxy
```

- 13 On the **Other Nodes** tab, verify that the NSX Manager connectivity status of hosts in the cluster is Up and NSX configuration state is Success. During the configuration process, each transport node displays the percentage of progress of the installation process. If installation fails at any stage, you can restart the process by clicking the **Resolve** link that is available against the failed stage of the process.

You can also see that the transport zone is applied to the hosts in the cluster.

Note If you again configure a host that is part of a cluster that is already prepared by a transport node profile, the configuration state of a node is in *Configuration Mismatch* state.

Note The **Other Nodes** tab displays TEP addresses of the host in addition to IP addresses. TEP address is the address assigned to the VMkernel NIC of the host, whereas IP address is the management IP address.

14 (Optional) Remove an NSX VIBs on the host.

- a Select one or more hosts and click **Actions > Remove NSX**.

The uninstallation takes up to three minutes. Uninstallation of NSX removes the transport node configuration on hosts and the host is detached from the transport zone(s) and switch. Similar to the installation process, you can follow the percentage of the uninstallation process completed on each transport node. If uninstallation fails at any stage, you can restart the process by clicking the **Resolve** link that is available against the failed stage of the process.

15 (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

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 Edge nodes and hosts join the management plane, the NSX logical entities and configuration state are pushed to the NSX Edge nodes and hosts automatically. You can create transport zones, logical switches, logical routers, and other network components through the NSX Manager UI or API at any time. When the hosts are transport nodes, these entities gets realized on the host.

Create a logical switch and assign logical ports. See the Advanced Switching section in the *NSX Administration Guide*.

Static IP List Assigned to VTEPs During Transport Node Configuration

Static IP addresses are assigned to VTEPs during transport node configuration in the order they are entered in the static IP address list. Know the limitations and workaround when working with static IP list.

How NSX assigns static IP addresses during transport node configuration

- During Transport Node creation, NSX will honor the static IP list order when assigning these IP addresses to VTEPs. For example, a Static IPv4 List ordered as 10.10.10.1, 10.10.10.2, 10.10.10.3, 10.10.10.4, will be assigned to VTEPs vmk1, vmk2, vmk3, vmk4 in the same order, 10.10.10.1 will be assigned to vmk1, 10.10.10.2 will be assigned to vmk2 and so on.
- During a Transport Node update, NSX will not honor static IP address order if the same list was already in use. Re-shuffling of the list elements will not lead to update in IP Addresses of existing VTEPs.
- During NSX or Transport Node upgrade, NSX will not honor static IP list order and existing VTEPs will remain as is.

Scenario: Assign IP addresses to VTEPs in a specific order when a transport node is already created

If you want Transport Nodes to have VTEPs being assigned IP addresses in a specific order, when the Transport Node is already created, you can do any one of the following:

- Delete and Configure NSX again with the desired order of IP addresses in Static IPv4 List.
- Reconfigure Transport Node to change Uplink Profile assigned to Transport Node:
 - a Reconfigure the Transport Node to use an Uplink Profile with Teaming Policy as 'Failover Order', with the Static IPv4 List containing a new single IP address. Alternatively, use any Uplink Profile with Teaming Policy Load Balance Source, if it contains only one Active Uplink.
 - b Reconfigure Transport Node again to use the desired Uplink Profile with desired order of IP-addresses in Static IPv4 List. Use a new or a different Ip address that will not be used in the static IP list. This is to avoid re-use of IP-addresses and VTEPs from the intended Static IPv4 List.
- Reconfigure Transport Node to change IP Assignment Type to other options (DHCP, IP Pool v4 Pool), keeping all other parameters same. Then, reconfigure Transport Node again to use the desired Uplink Profile with desired order of IP-addresses in Static IPv4 List.

Verify the Transport Node Status

Verify the transport node is created and VDS configured successfully.

Procedure

- 1 Log in to the NSX Manager UI.
- 2 Navigate to **System** → **Fabric** → **Hosts** page.

- 3 To view the host and VDS status, navigate to **Host Details** → **Overview** tab.
- 4 Verify host controller connectivity and manager connectivity is UP.
- 5 Navigate to the Transport Node page and view the VDS status.
- 6 To view active manager node for this transport node, run `nsxcli -c get managers *`.
- 7 To view the master controller with 'connected' status and session state up for this transport node, run `nsxcli -c get controllers`.
- 8 Alternatively, view the VDS on ESXi with the `esxcli network ip interface list` command.

On ESXi, verify correct number of vmk interfaces are configured 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 VDS in the UI by selecting host **Configuration > Network Adapters**.

- 9 Check the transport node's assigned tunnel endpoint address.

The vmk10 interface receives an IP address from the NSX 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

10 Check the API for transport node state information.

Call the (deprecated API) GET `https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/state` API call.

Or call the GET `/policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcement-point-id>/transport-node-status-report`, where the default value for `<site-id>` and `<enforcement-point-id>` is default.

For example:

```
{ "transport_node_id": "55120a1a-51c6-4c20-b4a3-6f59662c9f6a", "host_switch_states":
  [ { "host_switch_id": "50 21 0c 52 94 22 aa 20-b7 f0 0b da 1c 7c 29 ea",
    "host_switch_name": "dvs1",
    "endpoints": [ { "device_name": "vmk10", "ip": "172.16.223.175", "default_gateway": "",
    "subnet_mask": "255.255.224.0", "label": 53249 } ],
    "transport_zone_ids": [ "1b3a2f36-bfd1-443e-a0f6-4de01abc963e" ], "host_switch_type":
    "VDS" } ],
  "maintenance_mode_state": "DISABLED",
  "node_deployment_state": { "state": "success", "details": [] },
  "deployment_progress_state": { "progress": 100, "current_step_title": "Configuration
  complete" },
  "state": "success",
  "details": [ { "sub_system_id": "55120a1a-51c6-4c20-b4a3-6f59662c9f6a",
    "sub_system_type": "HostConfig", "state": "success" },
    { "sub_system_id": "55120a1a-51c6-4c20-b4a3-6f59662c9f6a", "sub_system_type": "AppInit",
    "state": "success" },
    { "sub_system_id": "55120a1a-51c6-4c20-b4a3-6f59662c9f6a", "sub_system_type":
    "LogicalSwitchFullSync",
    "state": "success" } ] }
```

Manual Installation of NSX Kernel Modules

As an alternative to using the NSX UI or the POST `/api/v1/transport-nodes` API (deprecated) API or PUT `/api/v1/infra/sites/site-id/enforcement-points/enforcement-point-id/host-transport-nodes/` API, you can install NSX kernel modules manually from the hypervisor command line.

Note You cannot manually install of NSX kernel modules on a bare metal server.

Manually Install NSX Kernel Modules on ESXi Hypervisors

To prepare hosts to participate in NSX, you must install NSX kernel modules on ESXi hosts. This allows you to build the NSX control-plane and management-plane fabric. NSX kernel modules packaged in VIB files run within the hypervisor kernel and provide services such as distributed routing, distributed firewall, and bridging capabilities.

You can download the NSX VIBs manually and make them part of the host image. The download paths can change for each release of NSX. Always check the NSX 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: The update completed successfully, but the system needs to be rebooted for the
  changes to be effective.
  Reboot Required: true
  VIBs Installed: VMware_bootbank_nsx-adf-<release>, VMware_bootbank_nsx-
  aggservice-<release>, VMware_bootbank_nsx-cli-libs-<release>, VMware_bootbank_nsx-
  common-libs-<release>, VMware_bootbank_nsx-context-mux-<release>, VMware_bootbank_nsx-
  esx-datapath-<release>, VMware_bootbank_nsx-exporter-<release>, VMware_bootbank_nsx-
  host-<release>, VMware_bootbank_nsx-metrics-libs-<release>, VMware_bootbank_nsx-
  mpa-<release>, VMware_bootbank_nsx-nestdb-libs-<release>, VMware_bootbank_nsx-
  nestdb-<release>, VMware_bootbank_nsx-netcpa-<release>, VMware_bootbank_nsx-
  netopa-<release>, VMware_bootbank_nsx-opsagent-<release>, VMware_bootbank_nsx-platform-
  client-<release>, VMware_bootbank_nsx-profiling-libs-<release>, VMware_bootbank_nsx-
  proxy-<release>, VMware_bootbank_nsx-python-gevent-<release>, VMware_bootbank_nsx-python-
  greenlet-<release>, VMware_bootbank_nsx-python-logging-<release>, VMware_bootbank_nsx-
  python-protobuf-<release>, VMware_bootbank_nsx-rpc-libs-<release>, VMware_bootbank_nsx-
  sfhc-<release>, VMware_bootbank_nsx-shared-libs-<release>, VMware_bootbank_nsx-upm-
  libs-<release>, VMware_bootbank_nsx-vdpi-<release>, VMware_bootbank_nsxcli-<release>,
  VMware_bootbank_vsipfwlib-<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 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-aggsservice

Provides host-side libraries for NSX aggregation service. NSX aggregation service is a service that runs in the management-plane nodes and fetches runtime state from NSX components.

nsx-cfgagent

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-cli-libs

Provides the NSX 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 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

NestDB is a database that stores NSX configurations related to the host (desired/runtime state, etc).

nsx-opsagent

Communicates operations agent executions (transport node realization, Link Layer Discovery Protocol - LLDP, traceflow, packet capture, etc.) with the management plane.

nsx-netcpa

Provides communication required by the different components.

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 upgrade and uninstall and monitoring of NSX 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 Distributed Firewall.

vsipfwlib

Provides distributed firewall functionality.

nsxcli

Provides the NSX CLI on hypervisor hosts.

To verify, you can run the **esxcli software vib list | grep -E 'nsx|vsip'** or **esxcli software vib list | grep <yyyy-mm-dd>** command on the ESXi host, where the date is the day that you performed the installation.

What to do next

Add the host to the NSX management plane. See [Form an NSX Manager Cluster Using the CLI](#).

Preparing Physical Servers as NSX Transport Nodes

To use NSX on a physical server (also known as Bare Metal server), you must install supported third-party packages.

Physical Server Concepts:

- Application - represents the actual application running on the physical server server, such as a web server or a data base server.
- Application Interface - represents the network interface card (NIC) which the application uses for sending and receiving traffic. One application interface per physical server server is supported.
- Management Interface - represents the NIC which manages the physical server server.
- VIF - the peer of the application interface which is attached to the logical switch. This is similar to a VM vNIC.

NSX supports the physical server server in two ways: as a host transport node and as a host for NSX Manager.

Make sure that you have the supported physical server 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 physical server hosts on which the NSX Edges are deployed. See *seciton on Forged Transmits in the [vSphere product documentation](#)*.

Install Third-Party Packages on a Linux Physical Server

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

Prerequisites

- Disable SELinux on all Linux physical hosts.

- Disable secure boot on all Linux physical hosts.
- 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, CentOS or Oracle Linux- `yum install libvirt-libs`
 - Ubuntu - `apt-get install libvirt0`
 - Oracle Linux - `rpm-qa | grep xxx`
 - SUSE - `zypper install libvirt-libs`

Procedure

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

Ubuntu 20.04, 18.04.2	Ubuntu 16.04
traceroute python-mako python-netaddr python-simplejson python-unittest2 python-yaml python-openssl dkms libvirt0 libelf-dev python3-netifaces	libunwind8 libgflags2v5 libgoogle-perftools4 traceroute python-mako python-simplejson python-unittest2 python-yaml python-netaddr python-openssl 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 libelf-dev python3-netifaces

- ◆ On RHEL, and CentOS 8.4 and 8.2, run `yum install` to install the third-party packages.

RHEL 8.4 and 8.2	CentOS 8.4 and 8.2
tcpdump	tcpdump
boost-filesystem	boost-filesystem
python3-pyyaml	python3-pyyaml
boost-iostreams	boost-iostreams
boost-chrono	boost-chrono
python3-mako	python3-mako
python3-netaddr	python3-netaddr
python3-six	python3-six
snappy	snappy
boost-date-time	boost-date-time
c-ares	c-ares
redhat-lsb-core	redhat-lsb-core
wget	wget
net-tools	net-tools
yum-utils	yum-utils
lsof	lsof
libvirt-libs	libvirt-libs
python3-gevent	python3-gevent
libev	libev
python3-greenlet	python3-greenlet
python3	python3
libbbpf	libbbpf
python3-netifaces	python3-netifaces
python3-pyOpenSSL	python3-pyOpenSSL
network-scripts	network-scripts

- ◆ On Oracle Linux 8.6 run `yum install` to install the third-party packages.

Oracle 8.6
tcpdump
boost-filesystem
python3-pyyaml
boost-iostreams
boost-chrono
python3-mako
python3-netaddr
python3-six
snappy
boost-date-time
c-ares
redhat-lsb-core
wget
net-tools
yum-utils
lsof
libvirt-libs
python3-gevent
libev
python3-greenlet
python3
libbbpf
python3-netifaces
python3-pyOpenSSL

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

RHEL 7.9, 7.7, and 7.6	CentOS 7.9, 7.7, and 7.6	Oracle Linux 7.9, 7.8, 7.7 and 7.6
tcpdump	tcpdump	tcpdump
boost-filesystem	boost-filesystem	boost-filesystem
PyYAML	PyYAML	PyYAML
boost-iostreams	boost-iostreams	boost-iostreams
boost-chrono	boost-chrono	boost-chrono
python-mako	python-mako	python-mako
python-netaddr	python-netaddr	python-netaddr
python-six	python-six	python-six
gperftools-libs	gperftools-libs	gperftools-libs
libunwind	libunwind	libunwind
elfutils-libelf-devel	elfutils-libelf-devel	snappy
snappy	snappy	boost-date-time
boost-date-time	boost-date-time	c-ares
c-ares	c-ares	redhat-lsb-core
redhat-lsb-core	redhat-lsb-core	wget
wget	wget	net-tools
net-tools	net-tools	yum-utils
yum-utils	yum-utils	lsof
lsof	lsof	libvirt-libs
python-gevent	python-gevent	python3-netifaces
libev	libev	python-greenlet
python-greenlet	python-greenlet	libev
libvirt-libs	libvirt-libs	python-gevent
python3-netifaces	python3-netifaces	python3
python3	python3	python3-netifaces
wget	wget	
redhat-lsb-core	redhat-lsb-core	
	python3-netifaces	

- ◆ On SUSE 12 SP3, 12 SP4, and 12 SP5 (starting in NSX 3.2.1), run `zypper install <package_name>` to install the third-party packages manually.

```
net-tools
tcpdump
python-simplejson
python-netaddr
python-PyYAML
python-six
libunwind
wget
lsof
libcap-progs
libvirt-libs
python3-netifaces
```

Configure a Physical Server as a Transport Node from GUI

As an admin, you can configure a physical server as a standalone transport node through the NSX Manager GUI.

Alternatively, you can run the Ansible script to achieve the same goal. See [Ansible Server Configuration for Physical Server](#) for configuring Windows physical servers using Ansible. However, it is recommended to use the NSX Manager UI to prepare physical servers for NSX networking.

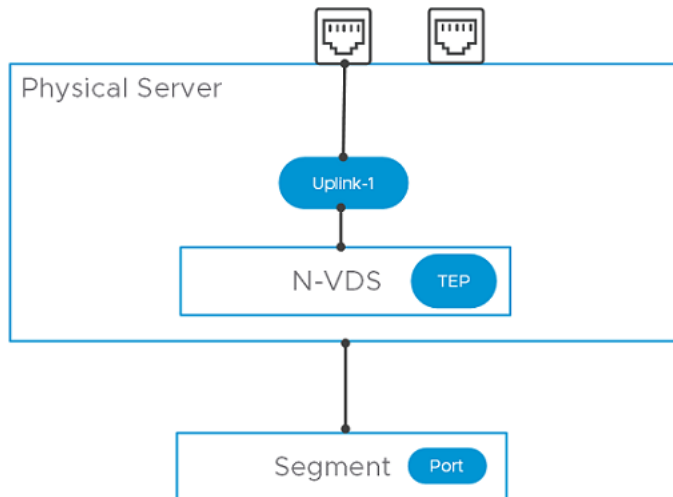
Physical servers supports an overlay and VLAN transport zone. Management interface is used to manage the physical server while application interface allows you to access the applications on the physical server. The application interface allows you to access the applications on the physical server. These NIC configurations are supported on a physical server:

- Single physical NIC cards provide an IP address for both the management and the application IP interfaces.
- Dual physical NIC cards provide a physical NIC and a unique IP address for the management interface. Dual physical NIC cards also provide a physical NIC, and a unique IP address for the application interface.
- Windows servers: Multiple physical NIC cards in a bonded configuration provide dual physical NIC cards - providing a unique IP address for both the management interface and the application interface. Such physical NIC bonds are supported through bonds created in the OS. Bond must be configured in the Switch Independent mode. Traffic running on management network is not supported on a bonded teaming interface.
- Linux servers: Bond interface only supports underlay mode (VLAN 0). CentOS 7.9, RHEL 7.9 are supported. Physical NIC bonds are supported in Active/Active and Active/Standby mode through OVS switch.

Unlike preparation of a standalone or a managed ESXi host that ends when it becomes a transport node, for a physical server, complete server preparation extends to attaching the application interface of the physical server to an NSX segment.

After preparing the host as a transport node, you must complete the following tasks to finish configuring a physical server.

- 1 Create a segment port on an NSX segment.
- 2 Attach application interface of the physical server to the segment port.



Prerequisites

- 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 physical NIC must be available on the host node.
- Hostname
- Management IP address
- User name
- Password
- A segment (VLAN or Overlay), depending upon your requirement, must be available to attach to the application interface of the physical server.
- Verify that the required third-party packages are installed. Third party packages must be installed on the physical server so that its physical NICs are available during transport node configuration. See [Install Third-Party Packages on a Linux Physical Server](#).
- On Linux physical servers, you can update the `sudoers` file to add custom users with minimal privileges. The custom users allows you to install NSX without root permissions.

After configuring visudo, run the following command to access the `/etc/sudoers` file.

```
$ sudo visudo
```

RHEL/CentOS/OEL/SLES:

```
tester ALL=(ALL) /usr/bin/rpm, /usr/bin/nsxcli, /usr/bin/systemctl restart openvswitch
```

Ubuntu:

```
tester ALL=(ALL) /bin/ls, /usr/bin/sudo, /usr/bin/dpkg, /bin/nsxcli
```

Procedure

- 1 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:A2:4E:3C:C4:F4
```

- 2 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 3 Go to **System > Fabric > Hosts**.
- 4 Select **Standalone** and **+ Add Host Node**.
- 5 On the **Host Details** window, enter the following details.

Option	Description
Name and Description	Enter the name to identify the physical server. You can optionally add the description of the operating system used for the host or physical server server.
IP Addresses	Enter the host or physical server server IP address.
Operating System	Select an operating system that mentions physical server. For example, if the operating system on the physical server is CentOS, select CentOS Physical Server. NSX identifies bare metal servers as physical servers. Depending on your physical server, you can select any of the supported operating systems. See Bare Metal Server System Requirements . Important Among the different flavors of Linux supported, you must know the distinction between a physical server server running a Linux distribution versus using a Linux distribution as a hypervisor host. For example, selecting Ubuntu Server as the operating system means setting up a physical server server running a Linux server.
Username and Password	Enter the host user name and password.
SHA-256 Thumbprint	This is an optional step. 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 to discover and authenticate the host.

6 Click **Next**.

7 On the **Prepare Host** window, enter the following details. You can only configure a single N-VDS switch for a single physical server.

Option	Description
Name	Enter a name for the N-VDS host switch.
Transport Zone	From the drop-down menu, select a transport zone that this transport node.
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 only receives LLDP packets from a LLDP neighbor. However, NSX can be set to send LLDP packets to and receive LLDP packets from a LLDP neighbor.
Uplinks-Physical NICs Mapping	<p>To map an uplink in NSX with a physical NIC or a bonded interface, enter the name of the physical NIC or bonded interface as configured on the physical server. For example, if teaming1 is the name of the interface you configured on the Windows server, then enter teaming1 in the Physical NICs field.</p> <p>Important</p> <ul style="list-style-type: none"> ■ You cannot map one uplink to a physical NIC and another uplink to a bonded interface. ■ If you are using a bonded interface, both NICs must be configured to function at the same packet transfer speed. <p>On Windows servers, you can configure teaming interfaces (bonded interfaces). The supported load balancing algorithms for teaming interfaces on Windows servers are:</p> <ul style="list-style-type: none"> ■ TransportNodes load balancing algorithm ■ MacAddresses load balancing algorithm ■ IPAddresses load balancing algorithm <p>In the teaming interface configuration, set Teaming Mode to Switch Independent mode. For more details, see Windows documentation.</p> <p>On Linux servers, you can configure a bonded interface by updating the <code>network-scripts</code> files. For more information, see Linux documentation.</p>

8 Click **Next**.

9 As the host is configured, the physical server progress is displayed.

- 10 On the **Configure NSX** window, verify status of host preparation. Based on whether you want to proceed with further configuration, these choices are available:

	Description
Click Select Segment	If the physical server preparation was successful, click Select Segment . In the next part of the procedure, you select a segment to attach the physical server's application interface through the NSX agent. Proceed to the next step.
Click Continue Later	If you click Continue Later button, then preparation ends without the application interface configured. You can later attach the segment port to the application interface. Go to Networking → Segments . Configure application interface for the BMS.
Preparation Failed	If preparation failed, go to System > Fabric > Hosts > Standalone . Identify the physical server, check if the Configuration State is in Failed state. Click Resolve to retry host preparation.

- 11 If you proceed to select a segment for the physical server, perform the following steps:
- From the list of segments connected to the transport zone you configured for the physical server, select the one to configure for the server.
 - Click the vertical ellipses and click **Edit** to customize segment properties.

Note Only properties related to a segment can be edited. Admin can modify: Segment Name, Connected Gateway, Subnet, Uplink Teaming Policy, IP Address Pool.

- 12 To add a new segment port on an NSX segment, go to the Select Segment window, click Add Segment Port. The segment port page is auto-populated.

Option	Description
Name	Enter the Segment Port name.
ID	The virtual interface UUID is auto-populated.
Type	Static is auto-populated as the node is of the type, physical server.
Context ID	Transport node UUID is auto-populated.

Note Alternatively, you can also run the API command, `https://<NSX-Manager-IP-address>/PATCH /policy/api/v1/infra/segments/<segment-id>/ports/<port-id>`.

Where, <port-id> is the virtual interface UUID, which is displayed on NSX Manager.

- 13 To attach application interface of physical server to a segment port, go to the **Set Segment Port** window, expand the **Attach Application Interface** section and enter these details:

Note The **Attach Application Interface** section is only applicable for physical servers.

Option	Description
Name	You can change the system-generated application interface name. On a Linux physical server, run <code>ovs-vsctl show</code> to verify the application interface name.
Context ID	To enable the application interface configuration, enter the host node ID.
Assign Existing IP	Use an existing IP so that it can be used for migration of the application interface.
Assign New IP	Used when configuring an overlay network. Select an IP assignment method on the segment - IP pool, DHCP, or Static. When you assign a new IP address for the application interface, complete the configuration by providing the IP address , Routing Table and Default Gateway details.

- 14 Click **Save**.
- 15 View the summary of the network configuration represented by topology diagram.
- 16 On the **Standalone** tab, select the physical server, and click **Switch Visualization** for the server. It must represent the network you configured on the physical server.
- 17 Verify that the NSX modules are installed on your host.

As a result of adding a host to the NSX fabric, a collection of NSX modules are installed on the host.

The modules on different hosts are packaged as follows:

- RHEL, CentOS, Oracle Linux, or SUSE - RPMs.
- Ubuntu - DEBs
- On RHEL, CentOS, or Oracle Linux, enter the command `yum list installed or rpm -qa`.
- On Ubuntu, enter the command `dpkg --get-selections`.
- On SUSE, enter the command `rpm -qa | grep nsx`.
- On Windows, open Task Manager. Or, from the command line enter `tasklist /V | grep nsx findstr "nsx ovs`

Results

The physical server is configured for NSX networking.

Attach Segment Port to Application Interface of Physical Server

When configuring a physical server as a transport node, if you did not attach a segment port to application interface of the server, complete the task to ensure NSX is configured on the physical server.

Make changes to the segment or segment port properties in the following scenarios:

- If the application interface of the physical server was not attached to an NSX segment, revisit the physical server transport node to complete the configuration.
- If you want to change segment port parameters, such as assigning a different IP address to the application interface and so on to application interface of the physical server.

Prerequisites

- Ensure the physical server is prepared as an NSX transport node. See [Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP](#).

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 If you did not configure the application interface of physical server, on the **System > Fabric > Hosts > Standalone** tab the node status is Success. The node indicates by showing a warning icon that segment port is not attached to application interface of physical server.
- 3 To configure the segment with the application interface, select the physical server, click **Actions → Manage Segment**.

The **Manage Segment** window displays the segment that is already attached to the application interface of the physical server.

- 4 For the selected segment, click **Edit Segment Port**, configure the **Application Interface** section and click **Save**. The segment port status displays as **Success**, if everything is functional.
- 5 To add a new port on the segment, on the **Manage Segment** window, click **Add Segment Port** and add the required details.

Alternatively, if a port already exists, click **Edit Segment** to proceed with configuration.

Option	Description
Context ID	To enable the application interface configuration, enter the host node ID.
Assign Existing IP	Use an existing IP so that it can be used for migration of the application interface.
Assign New IP	Used when configuring an overlay network. Select an IP assignment method on the segment - IP pool, DHCP, or Static. When you assign a new IP address for the application interface, complete the configuration by providing the IP address , Routing Table and Default Gateway details.

Option	Description
IP Address	Enter IP address for the application interface of the physical server.
Routing table	Enter routing table details.
Default Gateway	Enter the IP address of the gateway.

- 6 On the **System > Fabric > Hosts > Standalone** tab, select the physical server, and click **Switch Visualization** for the server. It must represent the network you configured on the physical server.

Results

A new segment port is attached to application interface of the physical server.

Ansible Server Configuration for Physical Server

When virtual interfaces (VIFs) are being configured, unique IDs of the VIFs have to be configured to be used as the segment port.

It is recommended to configure segment ports for application interface of physical server through UI. See [Configure a Physical Server as a Transport Node from GUI](#).

Ansible support modes are a set of automated scripts that set up the application interface for physical servers.

- Static Mode - Application interface IP Address is configured manually.
- DHCP Mode - Application interface IP Address is configured dynamically.
- Migrate Mode - This mode supports management and application sharing the same IP address. Also called underlay mode or VLAN-0.

For all Linux or Windows VM and physical servers:

- 1 Install Ansible based on the operating system: –https://docs.ansible.com/ansible/latest/installation_guide/intro_installation.html
- 2 Run the command `ansible-version`, and check that Ansible is version is 2.4.3 or later.
- 3 Download and extract the physical server integration with NSXfrom Github: –<https://github.com/vmware/bare-metal-server-integration-with-nsxt>

For Windows physical servers only:

- 1 Install pip for pywinrm.
- 2 Install pywinrm, and run `pip install pywinrm`.

Create Application Interface for Physical Server Workloads

Before you create or migrate an application interface for physical server workloads, you must configure NSX and install Linux third-party packages.

NSX does not support Linux OS interface bonding. You must use Open vSwitch (OVS) bonding for Physical Server Transport Nodes. See Knowledge Base article 67835 [Bare Metal Server supports OVS bonding for Transport Node configuration in NSX-T](#). The supported bond configuration for Linux is Active/Standby.

Procedure


- 1 Install the required third-party packages.
See [Install Third-Party Packages on a Linux Physical Server](#).
- 2 Configure the TCP and UDP ports.
See <https://ports.vmware.com/home/NSX-T-Data-Center>.
- 3 Add a physical server to the NSX fabric and create a transport node.
See [Configure a Physical Server as a Transport Node from GUI](#).
- 4 Use the Ansible playbook to create an application interface.
See [Ansible Server Configuration for Physical Server](#).

Windows Physical Server Supported Teaming Topologies

In some of the supported Windows physical server topologies you can pre-configure teaming in the Windows OS before deploying NSX, while in other supported topologies you can configure teaming in NSX and then deploy NSX.

Supported Topologies: Pre-configured NIC Teaming on Windows Physical Server

Figure 9-1. Topology 1, 2 and 3: NIC Teaming Configured on Windows OS



Topology 1:

- Configured pNIC1 and pNIC2 as bonding interface in Windows OS.
- Use IP1 on NSX for VLAN traffic.

Topology 2:

- Configured pNIC2 and pNIC3 as bonding interface in Windows OS.
- Use IP1 on pNIC1 for management traffic on the host.
- Use IP2 on NSX for VLAN traffic.

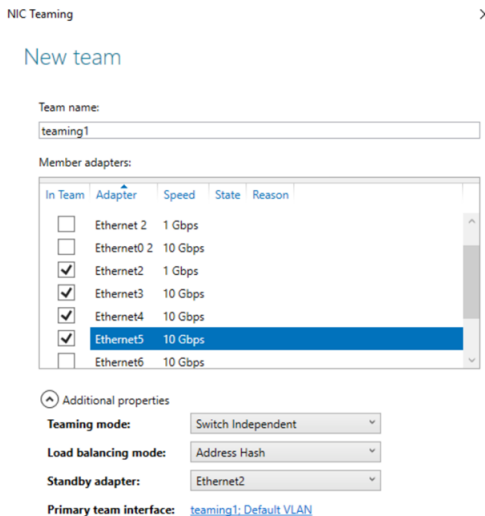
Topology 3:

- Configured pNIC2 and pNIC3 as bonding interface in Windows OS.

- Use IP1 on pNIC1 for management traffic on the host.
- Use IP2 as TEP on NSX for overlay traffic.

Windows NIC Teaming Configured with Active-Standby Adapters

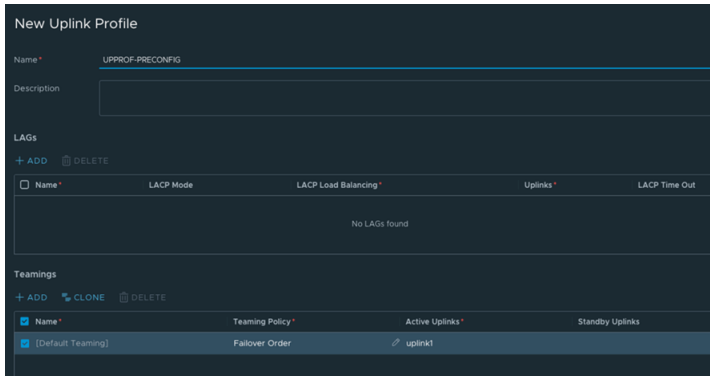
- 1 In the Windows OS interface, create a NIC team teaming1 with **Ethernet2**, **Ethernet3**, **Ethernet4**, **Ethernet5** as the NIC members and set the Teaming Mode to **Switch Independent**.



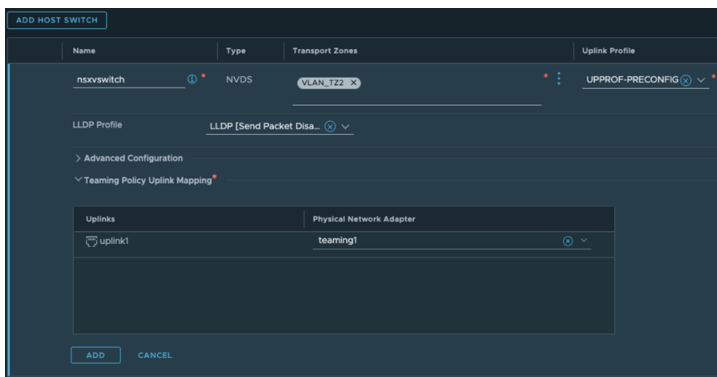
- 2 On the NIC teaming window, set the **Load balancing mode** to **Address Hash**. You can also set the mode to **IP Addresses** or **MAC Addresses**.

Note To set IP Addresses and Mac Addresses as the Load balancing mode, open a PowerShell session and run `Set-NetLbfoTeam`.

- 3 Set one of the ethernet members as the standby adapter. For example, set **Ethernet2** as the standby member.
- 4 From a browser window, open the NSX Manager UI interface.
- 5 Go to **System > Fabric > Profiles > Uplink Profiles**.
- 6 Create a new uplink profile, such as, *UPPROF-PRECONFIG* and use the default teaming policy with Active Uplink set to the Teaming interface, *uplink1*.



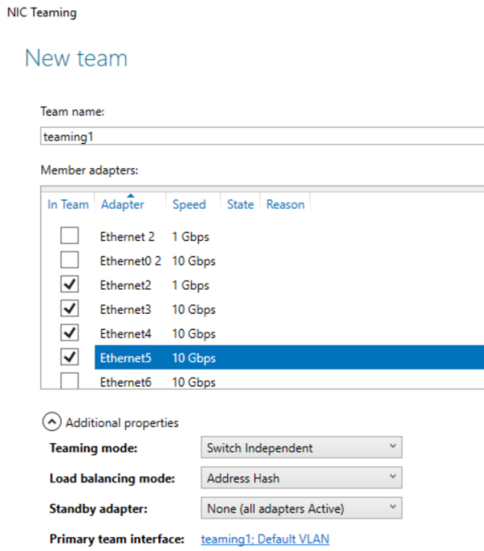
- 7 To prepare the Windows Host as a transport node, go to **System > Fabric > Host > Standalone**.
- 8 Select the Windows physical server and click **Configure NSX**.
- 9 In the Add Host Switch configuration window, select the uplink profile to *UPPROF-PRECONFIG*.



- 10 In the Teaming Policy Uplink Mapping section, set the uplink1 teaming interface to the physical network adapter, *teaming1*. You are using the NIC teaming, *teaming1*, set in the Windows OS that you configured at the beginning of the procedure. The NIC teaming comprises of the bond of Ethernet2, Ethernet3, Ethernet4, and Ethernet5.
- 11 Proceed with deployment of the Windows physical server host as a transport node.

Windows NIC Teaming Configured with Active/Active Adapters

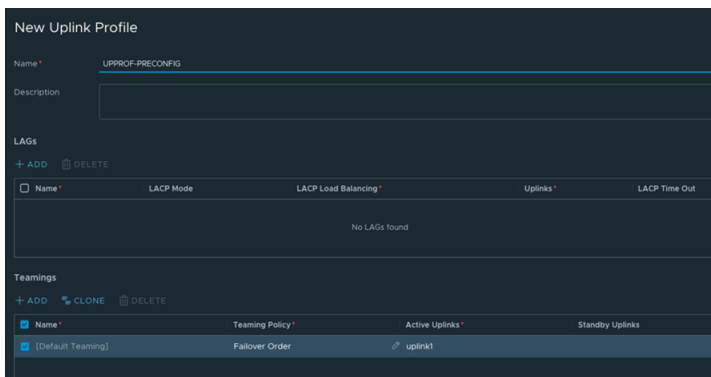
- 1 In the Windows OS interface, create a NIC team teaming1 with Ethernet2, Ethernet3, Ethernet4, and Ethernet5 as the NIC members and set the Teaming Mode to **Switch Independent**.



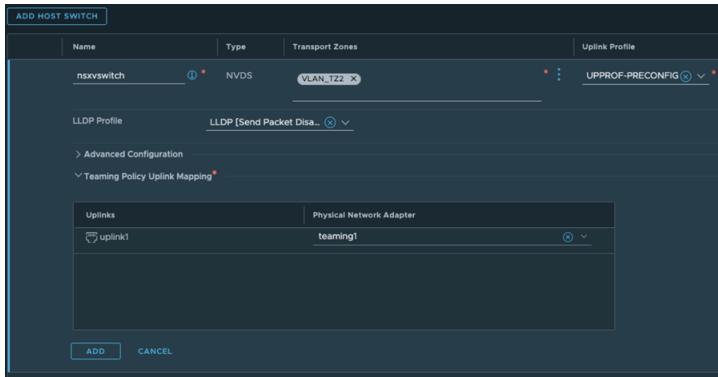
- On the NIC teaming window, set the **Load balancing mode** to **Address Hash**. You can also set the mode to **IP Addresses** or **MAC Addresses**.

Note To set IP Addresses and Mac Addresses as the Load balancing mode, open a PowerShell session and run `Set-NetLbfoTeam`.

- On the NIC teaming window, set the standby adapter to **None**. All NIC members are active in this configuration.
- From a browser window, open the NSX Manager UI interface.
- Go to **System > Fabric > Profiles > Uplink Profiles**.
- Create a new uplink profile, *UPPROF-PRECONFIG* and use the default teaming policy with Active Uplink set to the Teaming interface, *uplink1*.



- To prepare the Windows Host as a transport node, go to **System > Fabric > Host > Standalone**.
- Select the Windows physical server and click **Configure NSX**.
- In the Add Host Switch configuration window, select the uplink profile to *UPPROF-PRECONFIG*.




- 10 In the Teaming Policy Uplink Mapping section, set the uplink1 teaming interface to the physical network adapter, *teaming1*. You are using the NIC teaming, *teaming1*, set in the Windows OS that you configured at the beginning of the procedure. The NIC teaming comprises of the bond of Ethernet2, Ethernet3, Ethernet4, and Ethernet5.
- 11 Proceed with deployment of the Windows physical server host as a transport node.

Supported Topologies: NSX Configured NIC Teaming on Windows Physical Server

In the following topologies you configure NIC teaming bond interface in the NSX UI and then deploy NSX on the Windows physical server.

Figure 9-2. Topology 1 and 2: NIC Teaming Configured in NSX



Topology 1:

- Configured pNIC2 and pNIC3 as bonding interface in NSX.
- Use IP2 on NSX for VLAN traffic.
- Use IP1 for management traffic on the host.

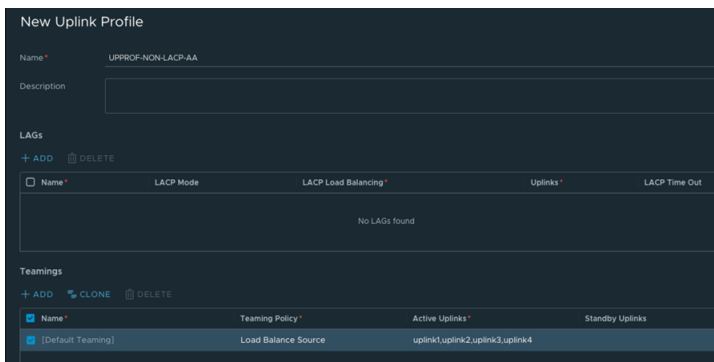
Topology 2:

- Configured pNIC2 and pNIC3 as bonding interface in NSX.
- Use IP2 for TEP on NSX for overlay traffic.
- Use IP1 for management traffic on the host.

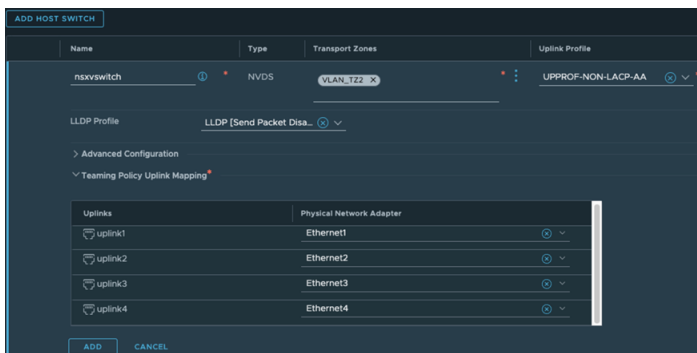
NSX Configured NIC Teaming Without LACP

In this configuration, you create an uplink profile using the default teaming policy as Active-Standby or Active-Active without configuring LACP profile along with the teaming policy.

- 1 From a browser window, open the NSX Manager UI interface.
- 2 Go to **System > Fabric > Profiles > Uplink Profiles**.
- 3 Create a new uplink profile, *UPPROF-NON-LACP-AA* and you can use the Teaming Policy as Failover or Load Balance Source.
- 4 The Active Uplinks field can be set to multiple uplinks, for example, uplink1, uplink2, uplink3, uplink4. These uplinks will be mapped later during transport node configuration to the Ethernet adapters on the Windows physical server.



- 5 To prepare the Windows Host as a transport node, go to **System > Fabric > Host > Standalone**.
- 6 Select the Windows physical server and click **Configure NSX**.
- 7 In the Add Host Switch configuration window, select the uplink profile to *UPPROF-NON-LACP-AA*.

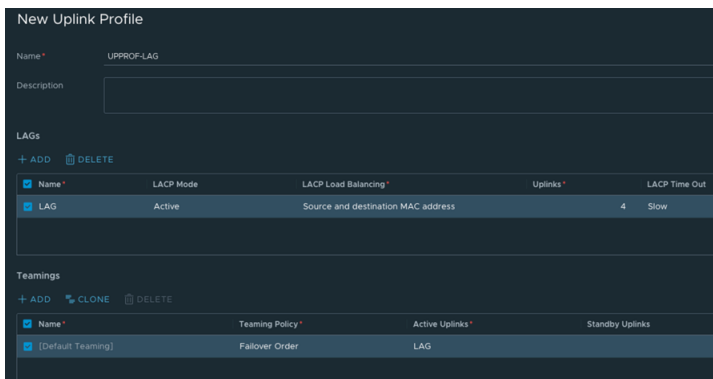


- 8 In the Teaming Policy Uplink Mapping section, map uplink1 to **Ethernet1**, uplink2 to **Ethernet2**, uplink3 to **Ethernet3**, uplink4 to **Ethernet4**.
- 9 Deploy the Windows physical server as a transport node. The deployed transport node will be configured to carry traffic using the ethernet adapters that are mapped the uplinks in the host switch.

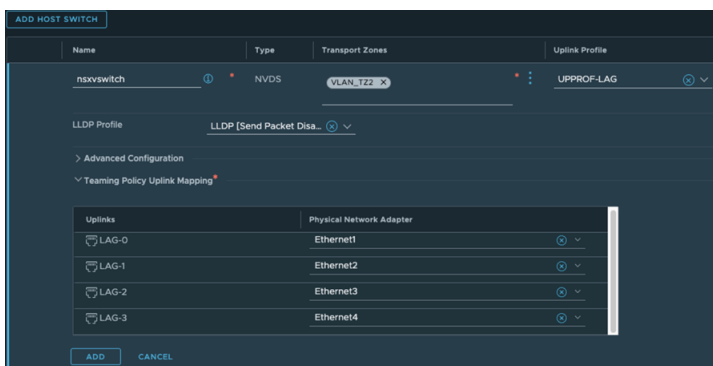
NSX Configured NIC Teaming With LACP

In this configuration, you create an uplink profile using the default teaming policy as Active-Active with LACP configuration.

- 1 From a browser window, open the NSX Manager UI interface.
- 2 Go to **System > Fabric > Profiles > Uplink Profiles**.
- 3 Create a new uplink profile, such as, *UPPROF-LAG*.
- 4 Create a LAG profile, such as *LAG*, LACP mode to **Active** and the LACP Load Balancing mode to **Source and Destination MAC address**.
- 5 Use the Teaming Policy as Failover and set the Active Uplink to *LAG*.
- 6 These uplinks will be mapped later during transport node configuration to the Ethernet adapters on the Windows physical server.



- 7 To prepare the Windows Host as a transport node, go to **System > Fabric > Host > Standalone**.
- 8 Select the Windows physical server and click **Configure NSX**.
- 9 In the Add Host Switch configuration window, select the uplink profile to *UPPROF-LAG*.

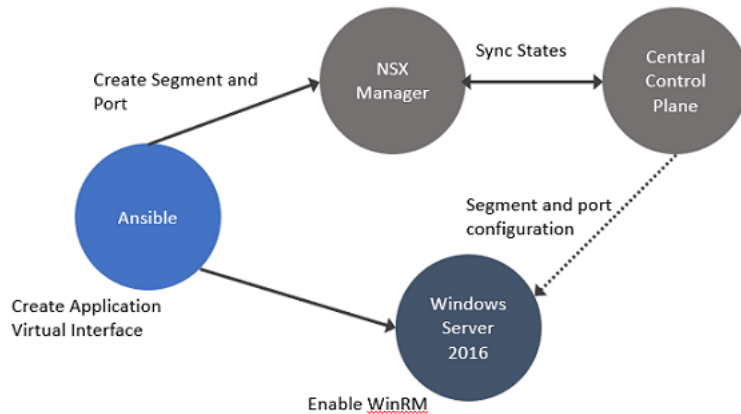


- 10 In the Teaming Policy Uplink Mapping section, map LAG-0 to **Ethernet1**, LAG-1 to **Ethernet2**, LAG-2 to **Ethernet3**, LAG-3 to **Ethernet4**.

- 11 Deploy the Windows physical server as a transport node. The deployed transport node will be configured to carry traffic using the ethernet adapters that are mapped the uplinks in the host switch.

Secure Workloads on Windows Server 2016/2019 Bare Metal Servers

The NSX agent installed on the servers provides connectivity and security to the bare metal workloads.



In this procedure, establish connectivity between the workloads and NSX Manager. Then, configure DFW rules to secure ingress and egress traffic running between virtual or physical and Windows Server 2016 or 2019 bare metal workloads.

Prerequisites

- Configure your own proxy settings on the physical server.

Procedure

- 1 Enable Windows Remote Management (WinRM) on Windows Server 2016 to allow the Windows server to interoperate with third-party software and hardware. To enable the WinRM service with a self-signed certificate.
 - a Run `PS$ wget -o ConfigureWinRMService.ps1 https://raw.githubusercontent.com/vmware/bare-metal-server-integration-with-nsxt/master/bms-ansible-nsx/windows/ConfigureWinRMService.ps1`.
 - b Run `PS$ powershell.exe -ExecutionPolicy ByPass -File ConfigureWinRMService.ps1`.
- 2 Configure WinRM to use HTTPS. The default port used for HTTPS is 5986.
 - a Run PowerShell as an administrator.
 - b Run `winrm quickconfig`.
 - c Run `winrm set winrm/config/service/auth '@{Basic="true"}'`.

- d Run `winrm set winrm/config/service @{AllowUnencrypted="true"}`.
- e Run `winrm create winrm/config/Listener?Address=*&Transport=HTTPS @{Hostname="win16-colib-001";CertificateThumbprint="[output of the 2nd command]"}`.
- f Verify configuration of WinRM. Run `winrm e winrm/config/listener`.

- 3 Add the bare metal server as a standalone transport node. See [Configure a Physical Server as a Transport Node from GUI](#).
- 4 Verify whether OVS bridges are created on the Windows server. The OVS bridge connects the application virtual interface to the NSX switch on the transport node.

```
ovs-vsctl show
```

The output must show the bridges created from `nsxswitch` and `nsx managed host` component. The `nsxswitch` bridge is for the transport node that was created. The `nsx managed` bridge is created for the application virtual interface on the Windows host. These bridge entries indicate that communication channel is established between the NSX switch and Windows remote listener.

- 5 On the overlay-backed transport node, verify:
 - The static IP address is reflected as the IP address of the overlay segment to which the Windows Server workload is connected.
 - The GENEVE tunnels are created between the NSX switch and the NSX managed host component on the Windows host.

Note Likewise, on a VLAN-backed transport node, verify that the static IP address is reflected as the IP address of the VLAN segment to which the Windows Server workload is connected.

- 6 In Windows, customize OVSIM driver for the Windows server to create two new network adapters - application virtual interfaces and virtual tunnel endpoint (VTEP) for overlay-backed workload.

```
$:> Get-NetAdapter
```

```
vEthernet1-VTEP: Used for overlay-backed VTEP interface. Not needed for a VLAN-backed workload.
```

```
vEthernet1-VIF1: Used for virtual interface or application interface of the bare metal Windows server.
```

- 7 To verify network adapters, go to the Windows server and run `Get-NetAdapter`.
- 8 Verify connectivity between the application, Windows bare metal server, and NSX Manager .
- 9 Add and publish L2 or L3 DFW rules for the overlay or VLAN-backed bare metal workload.
- 10 Verify ingress and egress traffic between virtual or physical workloads and bare metal workloads is flowing as per the DFW rules published.

Configure an ESXi Host Transport Node with Link Aggregation Group

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.
- For ESXi hosts (connecting to VDS), the LAG configuration is only in vCenter, on the VDS. Therefore configure LAG on the VDS switch in vCenter.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Select **System > Fabric > Profiles > Uplink Profiles > Add Profile**.
- 3 Enter a name and optionally a description.
For example, you enter the name `uplink-lag1`.
- 4 Under **Teamings**, select **Default Teaming**.
- 5 In the **Active Uplinks** field, do one of the following:
 - a Enter the name of the LAG that you configured on the VDS switch in VMware vCenter.
- 6 Enter the VLAN ID for **Transport VLAN**.
- 7 Click **Add** at the bottom of the dialog box.
- 8 Select **System > Fabric > Hosts**.
- 9 Select the **Cluster** tab and select an ESXi host.
- 10 Select **Configure NSX**.
- 11 In the **Host Details** tab, enter IP address, OS name, admin credentials, and SHA-256 thumbprint of the host.
- 12 During the switch configuration, depending on the VDS switch, select the uplink profile `uplink-lag1` 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 uplink LAG that is configured on the VDS switch is displayed in the uplink profile drop-down list. Select the VDS uplink for LAG.

- 14 Enter information for the other fields and complete host preparation.

The ESXi host is prepared as a transport node using the LAG profile.

Note You can only use LAGs that you define in NSX uplink profile for transport nodes that are prepared using N-VDS. However, for ESXi transport nodes prepared using VDS, you can only use LAGs configured on vSphere and it is referred to in the uplink profile.

- 15 SSH in to ESXi host as root and run the following commands:

- a Run `nsxcli`.
- b To verify LACP bond is up and active, run `get bonds`.

Quick Start Wizard to Prepare ESXi Cluster hosts for Security-only or Networking and Security

Using the Quick Start wizard, you can configure ESXi clusters for either Networking and Security or Security-only.

Note If you configure a cluster for Security-only then you cannot configure networking for that cluster.

Note Ensure that ESXi host version is v7.0.3 or later if you want to use the same VDS to prepare clusters for both the use cases - Security-only or Networking and Security. To avoid issues post cluster nodes preparation, upgrade ESXi host to v7.0.3 or later before preparing clusters using the same VDS.

Prepare ESXi cluster Hosts for Networking and Security

Use the Quick Start wizard to prepare ESXi clusters for networking and security using NSX recommended host configurations.

The Quick Start wizard gives you two options to prepare ESXi clusters: Networking and Security or Security Only. With either of these options, the wizard helps you finish installation with minimum user input, thus, simplifying the installation process. By default, VLAN networking is the default selection in the wizard.

Based on the type of host, the quick start wizard considers the following default configurations:

- ESXi hosts running 7.0 and later are prepared on the VDS switch. Configure the desired number of uplinks on the VDS switch in vCenter Server and set the MTU to 1700.

Each host switch is assigned an auto-created transport zone, uplink profile and transport node profile.


Prerequisites

- Register VMware vCenter as a compute manager in NSX.
- ESXi cluster hosts member of VDS in VMware vCenter with correct uplinks.

- (Optional) If you are using VLAN for VTEP pool, configure IP Pool for Host VTEP assignment and Uplink Profile using the TEP VLAN.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Navigate to **System > Quick Start**.
- 3 On the **Prepare Clusters for Networking and Security** card, click **Get Started**.
- 4 Select the clusters you want to prepare for NSX networking.
- 5 Click **Install NSX** and then select **Networking and Security**.
- 6 Depending on your requirement, you can prepare the same cluster for both VLAN and Overlay networking or for one type of networking. With Overlay networking, each host switch is added with a TEP IP address, which is required for overlay networking.



- 7 View the NSX recommended Host Switch configuration.

However, you can customize the settings for the cluster, even though it is an optional step.

Note A dotted line originating from a switch to a physical NIC indicates that it is an existing configuration on the host switch, which will be replaced by a firm line going to the same physical NIC.

- 8 Even though NSX provides recommendations, you can still customize the configuration. To customize the host switch, select the switch and change the recommended configuration.
 - a **IP Assignment:** Is applicable if overlay is selected for the host switch. Choose IP assignment type to be DHCP or a pre-created IP Pool for the overlay VTEP Pool.
 - b **VDS:** Select the VDS switch as the host switch.
 - c **Transport Zone:** Select a different transport zone that you want the host to be associated with.

- d **Uplink Profile:** If needed, select a different uplink profile in place of the recommended uplink profile.

Note If you configure two VDS switches with the same configuration, the wizard recommends the same uplink profile for both the switches.

- e **Uplink to Physical NIC mapping:** On a VDS switch, all uplinks configured on the VDS switch are mapped to the uplinks in NSX.

A change to host switch type or uplink to vmnic mapping is reflected in the Host Switch Configuration network representation.

9 Click **Install**.

View the progress of installation on the **Prepare Clusters for Networking and Security** card. If installation on any of the host fails, retry installation by resolving the error.

10 To view successfully prepared hosts, go to **System > Fabric > Hosts > Clusters**.

Results

The transport nodes are ready for VLAN and Overlay networking.

Install Distributed Security for vSphere Distributed Switch

NSX allows you to install Distributed Security for vSphere Distributed Switch (VDS). As the host switch is of the type VDS, DFW capabilities can be enabled on workload VMs..

Distributed Security provides security-related functionality to your VDS such as:

- Distributed Firewall (DFW)
- Distributed IDS/IPS
- Identity Firewall
- L7 App ID
- Fully Qualified Domain Name (FQDN) Filtering
- NSX Intelligence
- NSX Malware Prevention
- NSX Guest Introspection

Prerequisites

The following are the requirements for installing Distributed Security for VDS:

- vSphere 7.0 or later.
- The vSphere cluster should have at least one VDS with distributed switch version 6.6 or later configured and ESXi cluster hosts must be members of a VDS with uplinks configured.
- A compute manager must be registered in NSX. See [Add a Compute Manager](#).

- Before you deploy and configure Distributed Security on hosts, ensure that NSX is not deployed on such hosts.

Procedure

- 1 From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
- 2 Navigate to **System > Quick Start**.
- 3 On the **Prepare Clusters for Networking and Security** card, click **Get Started**.
- 4 Select the clusters that you want to install Distributed Security.
- 5 Click **Install NSX** and then select **Security Only**.
- 6 In the dialog box, click **Install**.

Note If the VDS spans across multiple clusters, Distributed Security installs only to the clusters that you selected.

The installation process for Distributed Security starts.

- 7 To view VDS with Distributed Security installed, navigate to **System > Fabric > Hosts**.

Note vSphere clusters prepared for Distributed Security are identified by the **Security** label.

Results

Distributed Security is installed and you can begin using security capabilities such as creating DFW policies and rules for the VDS.

Deploy a Fully Collapsed vSphere Cluster NSX on Hosts Running N-VDS Switches

You can configure NSX Manager, host transport nodes, and NSX Edge VMs on a single cluster. Each host in the cluster provides two physical NICs that are configured for NSX.

Prerequisites

- All the hosts must be part of a vSphere cluster.
- Each host has two physical NICs enabled.
- Register all hosts to a VMware vCenter.
- Verify on the VMware vCenter that shared storage is available to be used by the hosts.

- Starting with NSX 3.1, the TEP of an NSX Edge VM can communicate directly with the TEP of the host on which the NSX Edge VM is running. Before NSX 3.1, the NSX Edge VM TEP had to be in a different VLAN than the host TEP.

Note Alternatively, you can deploy the configuration described in this topic by using vSphere Distributed Switches. With vSphere Distributed Switches configured on hosts, the procedure is simple. It does not involve vSphere Distributed Switch to N-VDS migration and NSX Manager deployment on NSX distributed virtual port groups. But, the NSX Edge VM must connect to a VLAN segment to be able to use the same VLAN as the ESXi host TEP.

- Starting with NSX 3.2.2, you can enable multiple NSX Managers to manage a single VMware vCenter. See [Multiple NSX Managers Managing a Single VMware vCenter](#).

After you configure a collapsed cluster, if you enable the VMware vCenter to work with multiple NSX Managers (multiple NSX Managers are registered to the same VMware vCenter), you cannot deploy new NSX Manager nodes to the cluster. The workaround is to create a new cluster and deploy NSX Manager nodes.

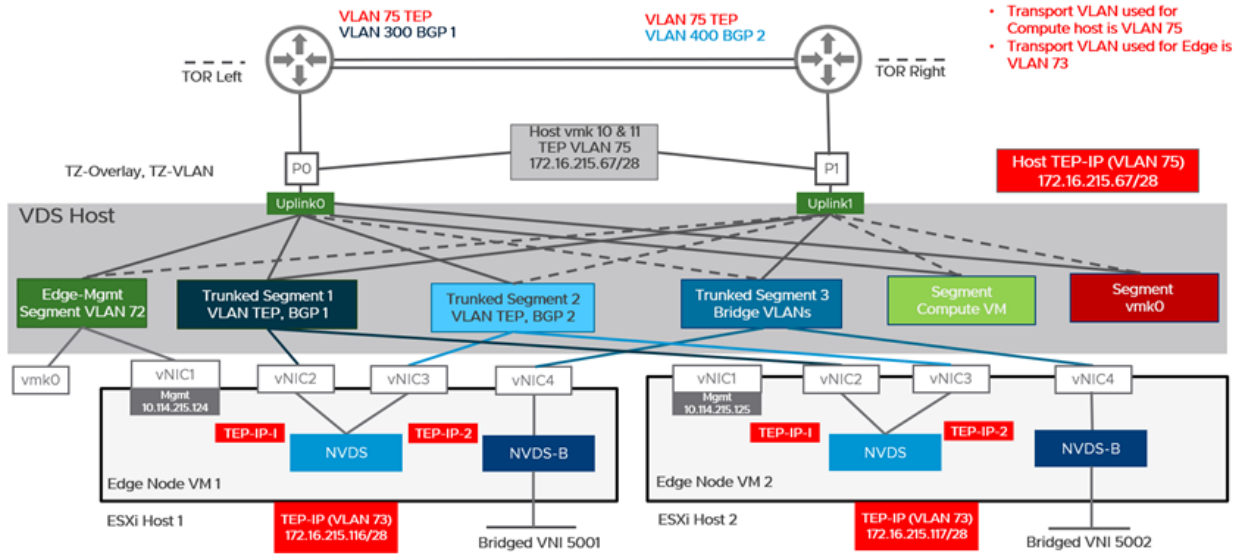
Note Deploy the fully collapsed single vSphere cluster topology starting with NSX 2.4.2 or 2.5 release.

The topology referenced in this procedure has:

- vSAN configured with the hosts in the cluster.
 - A minimum of two physical NICs per host.
 - vMotion and Management VMkernel interfaces.
- Starting with NSX 4.0, transport node hosts supports only VDS switch. However, NSX Edge is configured using an NVDS switch.

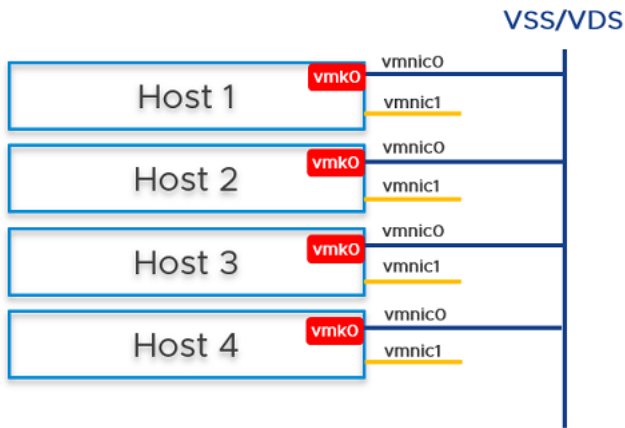
Figure 9-3. Topology: Single N-VDS Switch Managing Host Communication with NSX Edge and Guest VMs

Topology: Single VDS Host Switch Managing Host Communication with NSX Edge and Guest VMs
 Edge node VMs on NVDS




Procedure

- 1 Prepare four ESXi hosts with vmnic0 on vSS or vDS, vmnic1 is free.



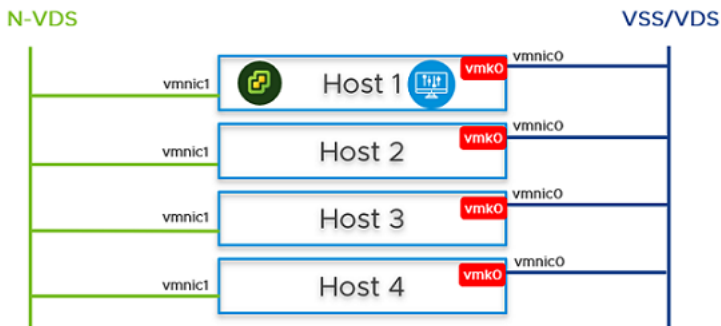
- 2 On Host 1, install VMware vCenter, configure a vSS/vDS port group, and install NSX Manager on the port group created on the host.



- 3 Prepare ESXi hosts 1, 2, 3 and 4 to be transport nodes.
 - a Create VLAN transport zone and overlay transport zone 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 transport nodes by applying a transport node profile. In this step, the transport node profile only migrates vmnic1 (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](#).

Edit Transport Node Profile - TNP-host ? ×


N-VDS Name*	nvds-host	▼
Associated Transport Zones	TZ-Overlay, TZ-Vlan	
NIOC Profile*	nsx-default-nioc-hostswitch-profile	▼
	OR Create New NIOC Profile	
Uplink Profile*	nsx-default-uplink-hostswitch-profile	▼
	OR Create New Uplink Profile	
LLDP Profile*	LLDP [Send Packet Enabled]	▼
IP Assignment*	Use IP Pool	▼
IP Pool*	TEP-Pool	▼
	OR Create and Use a new IP Pool	
Physical NICs	vmnic1	uplink-2 ▼
	Add PNIC	
PNIC only Migration	<input checked="" type="checkbox"/>	Yes
Enable this option if no vmks exist on PNIC selected for migration		
Network Mappings for Install	Add Mapping	
Network Mappings for Uninstall	Add Mapping	



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, VMware vCenter, and NSX Edge. Ensure to select the correct teaming policy for each of the VLAN-backed segments. Do not use VLAN trunk logical switch as the target. When creating the target segments in NSX Manager UI, in the **Enter List of VLANs** field, enter only one VLAN value.
- 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
 - vmk0 is mapped to **Edge Management Segment** .


- vmnic0 is mapped to an active uplink, **uplink-1**.



- 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 the N-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-VDS switch.

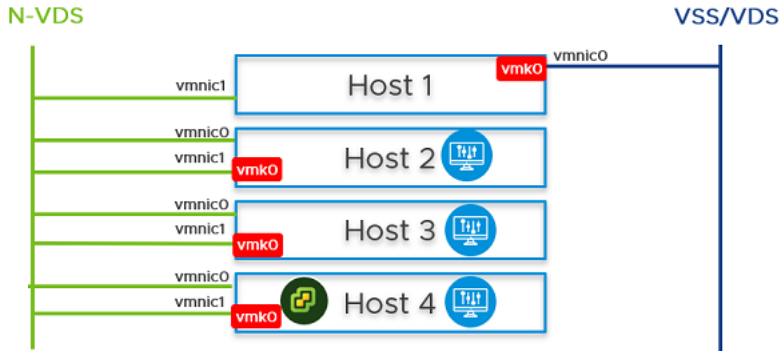


- 8 On Host 2 and Host 3, from the NSX Manager UI, install NSX Manager and attach NSX Manager to the segment. Wait for approximately 10 minutes for the cluster to form and verify that the cluster has formed.





- 9 Power off the first NSX Manager node. Wait for approximately 10 minutes.
- 10 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 VMware vCenter from host 1 to host 4.

For vMotion limitations, see <https://kb.vmware.com/s/article/56991>.




- 11 From the NSX Manager UI, go to Host 1, migrate vmk0 and vmnic0 together from VSS to N-VDS switch.
- 12 In the **Network Mapping for Install** field, ensure that the vmk0 adapter is mapped to the **Edge Management Segment** on the N-VDS switch.



- 13 On Host 1, install the NSX Edge VM from the NSX Manager UI.
See [Create an NSX Edge Transport Node](#).

N-VDS



- 14 Join the NSX Edge VM with the management plane.
See [Join NSX Edge with the Management Plane](#).
- 15 To establish the north-south traffic connectivity, configure NSX Edge VM with an external router.
- 16 Verify that north-south traffic connectivity between the NSX Edge VM and the external router.
- 17 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 to reboot, the management VM port will be initialized 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 In the NSX Manager UI, go to **Manager Mode > Networking > Logical Switches** tab (3.0 and later releases). Search for the **Segment Compute VM** segment. Select the **Overview** tab, find and copy the UUID. The UUID used in this example is, *c3fd8e1b-5b89-478e-abb5-d55603f04452*.
- b Create a JSON payload for each NSX Manager.
 - In the JSON payload, create logical ports with initialization status in **UNBLOCKED_VLAN** state by replacing the value for `logical_switch_id` with the UUID of the previously created **Edge Management Segment**.
 - In the payload for each NSX Manager, the `attachment_type_id` and `display_name` values will be different.

Important Repeat this step to create a total of four JSON files - three for NSX Managers and one for vCenter Server Appliance (VCSA).

```
port1.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"
}
```

Where,

- `admin_state`: This is state of the port. It must UP.
- `attachment_type`: Must be set to VIF. All VMs are connected to NSX-T switch ports using a VIF ID.
- `id`: This is the VIF ID. It must be unique for each NSX Manager. If you have three NSX Managers, there will be three payloads, and each one of them must have a different VIF ID. To generate a unique UUID, log into the root shell of the NSX Manager and run `/usr/bin/uuidgen` to generate a unique UUID.
- `display_name`: It must be unique to help NSX admin identify it from other NSX Manager display names.
- `init_state`: With the value set to `UNBLOCKED_VLAN`, NSX unblocks the port for NSX Manager, even if the NSX Manager is not available.
- `logical_switch_id`: This is the logical switch ID of the **Edge Management Segment**.

- c If there are three NSX Managers deployed, you need to create three payloads, one for each logical port of a NSX Manager. For example, port1.json, port2.json, port3.json.

Run the following commands to create payloads.

```
curl -X POST -k -u '<username>:<password>' -H 'Content-Type:application/json'
-d @port1.json https://nsxmgr/api/v1/logical-ports
```

```
curl -X POST -k -u '<username>:<password>' -H 'Content-Type:application/json'
-d @port2.json https://nsxmgr/api/v1/logical-ports
```

```
curl -X POST -k -u '<username>:<password>' -H 'Content-Type:application/json'
-d @port3.json https://nsxmgr/api/v1/logical-ports
```

An example of API execution to create a logical port.

```
root@nsx-mgr-147:/var/CollapsedCluster# curl -X POST -k -u
'<username>:<password>' -H 'Content-Type:application/json' -d @port1.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 Verify that the logical port is created.

Logical Port	ID	Admin Status	Operational Status	Switching Profiles	Attachment	Logical Switch
NSX Manager Node 147 Port	02e0_b647	Up	Down	nsx-default-switch-se...	VIF:nsxm...147	Seg-ESXi-MGT-VlanL...
nsx-mgr-147/nsx-mgr-147.vmx@1161331-11...	1ab1_bbd4	Up	Down	nsx-default-switch-se...	VIF:5028_f0c0	Seg-ESXi-MGT-VlanL...
nsx-mgr-147/nsx-mgr-147.vmx@1161331-11...	4c9f_379d	Up	Down	nsx-default-switch-se...	VIF:5028_8123	Seg-ESXi-MGT-VlanL...
nsx-mgr-147/nsx-mgr-147.vmx@1161331-11...	c7d0_8cfe	Up	Up	nsx-default-switch-se...	VM:nsx-mgr-147	Seg-ESXi-MGT-VlanL...
nsx-mgr-157/nsx-mgr-157.vmx@1161331-11...	a6f3_536f	Up	Up	nsx-default-switch-se...	VM:nsx-mgr-157	Seg-ESXi-MGT-VlanL...
vmk0@en-vds-1@1161331-116-45c7-8747-3...	0fa3_9685	Up	Up	nsx-default-switch-se...	VIF:8793_8168	Seg-ESXi-MGT-VlanL...

- e Find out the VM instance ID for each of the NSX Manager. You can retrieve the instance ID from the **Inventory** → **Virtual Machines**, select the NSX Manager VM, select the **Overview** tab and copy the instance ID. Alternatively, search the instance ID from the managed object browser (MOB) of VMware vCenter. Add **:4000** to the ID to get the VNIC hardware index of an NSX Manager VM.

For example, if the instance UUID of the VM is 503c9e2b-0abf-a91c-319c-1d2487245c08, then its vnic index is 503c9e2b-0abf-a91c-319c-1d2487245c08:4000. The three NSX Manager vnic indices are:

```
mgr1 vnic: 503c9e2b-0abf-a91c-319c-1d2487245c08:4000
```

```
mgr2 vnic: 503c76d4-3f7f-ed5e-2878-cffc24df5a88:4000
```

```
mgr3 vnic: 503cafd5-692e-d054-6463-230662590758:4000
```

- f Find out the transport node ID that hosts NSX Managers. If you have three NSX Manager, each hosted on a different transport node, note down the transport node IDs. For example, the three transport node IDs are:

```
tn1: 12d19875-90ed-4c78-a6bb-a3b1dfe0d5ea
```

```
tn2: 4b6e182e-0ee3-403f-926a-fb7c8408a9b7
```

```
tn3: d7cec2c9-b776-4829-beea-1258d8b8d59b
```

- g Retrieve the transport node configuration that is to be used as payloads when migrating the NSX Manager to the newly created port.

For example,

```
curl -k -u '<user>:<password' https://nsxmgr/api/v1/transport-nodes/12d19875-90ed-4c78-a6bb-a3b1dfe0d5ea > tn1.json
```

```
curl -k -u '<user>:<password' https://nsxmgr/api/v1/transport-nodes/4b6e182e-0ee3-403f-926a-fb7c8408a9b7 > tn2.json
```

```
curl -k -u '<user>:<password' https://nsxmgr/api/v1/transport-nodes/d7cec2c9-b776-4829-beea-1258d8b8d59b > tn3.json
```


- h Migrate the NSX Manager from the previous port to the newly created unblocked logical port on the **Edge Management Segment**. The VIF-ID value is the attachment ID of the port created previously for the NSX Manager.

The following parameters are needed to migrate NSX Manager:

- Transport node ID
- Transport node configuration
- NSX Manager VNIC hardware index
- NSX Manager VIF ID

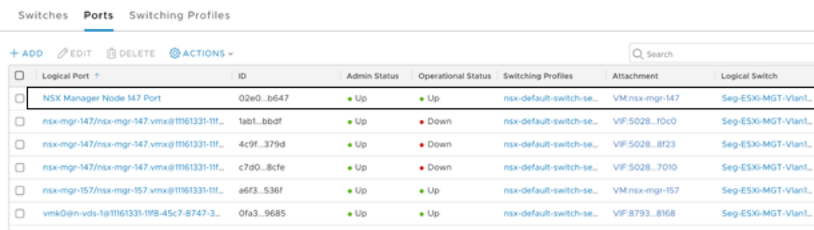
The API command to migrate NSX Manager to the newly created unblocked port is:

```
/api/v1/transport-nodes/<TN-ID>?vnic=<VNIC-ID>&vnic_migration_dest=<VIF-ID>
```

For example,

```
root@nsx-mgr-147:/var/CollapsedCluster# curl -k -X
PUT -u 'admin:VMware1!VMware1!' -H 'Content-
Type:application/json' -d @<tn1>.json 'https://localhost/api/v1/
transport-nodes/11161331-11f8-45c7-8747-34e7218b687f?vnic=5028d756-
d36f-719e-3db5-7ae24aa1d6f3:4000&vnic_migration_dest=nsxmgr-port-147'
```

- i Ensure that the statically created logical port is Up.



Logical Port	ID	Admin Status	Operational Status	Switching Profiles	Attachment	Logical Switch
NSX Manager Node 147 Port	02e0_b647	Up	Up	nsx-default-switch-se...	VM nsx-mgr-147	Seg-ESXi-MGT-Vlan...
nsx-mgr-147/nsx-mgr-147.vmx@1161331-11f...	1ab1_bbd8f	Up	Down	nsx-default-switch-se...	VIF-5028_10c0	Seg-ESXi-MGT-Vlan...
nsx-mgr-147/nsx-mgr-147.vmx@1161331-11f...	4c9f_379d	Up	Down	nsx-default-switch-se...	VIF-5028_8f23	Seg-ESXi-MGT-Vlan...
nsx-mgr-147/nsx-mgr-147.vmx@1161331-11f...	c7d0_8cfe	Up	Down	nsx-default-switch-se...	VIF-5028_7010	Seg-ESXi-MGT-Vlan...
nsx-mgr-157/nsx-mgr-157.vmx@1161331-11f...	a6f3_536f	Up	Up	nsx-default-switch-se...	VM nsx-mgr-157	Seg-ESXi-MGT-Vlan...
vmk0@in-vds-1@1161331-11f8-45c7-8747-3...	0fa3_9685	Up	Up	nsx-default-switch-se...	VIF-8793_8168	Seg-ESXi-MGT-Vlan...

- j Repeat the preceding steps on every NSX Manager in the cluster.

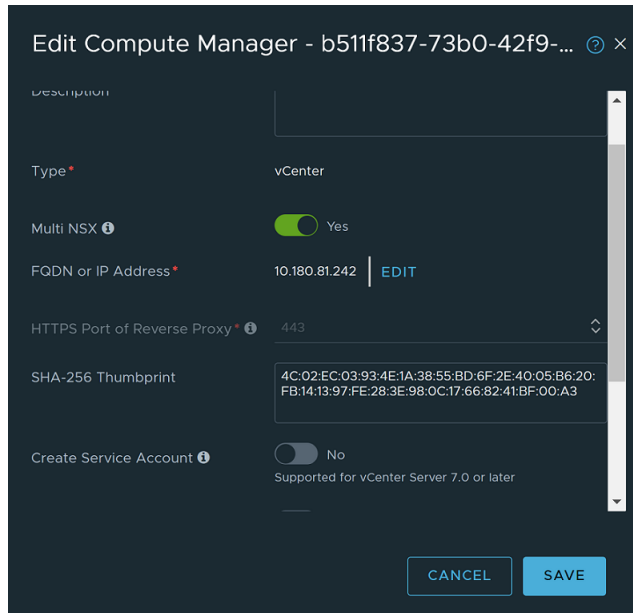
Multiple NSX Managers Managing a Single VMware vCenter

Improve the operational efficiency with multiple NSX Managers managing a single VMware vCenter. Admins can manage different clusters in the same VMware vCenter by using different NSX Managers.

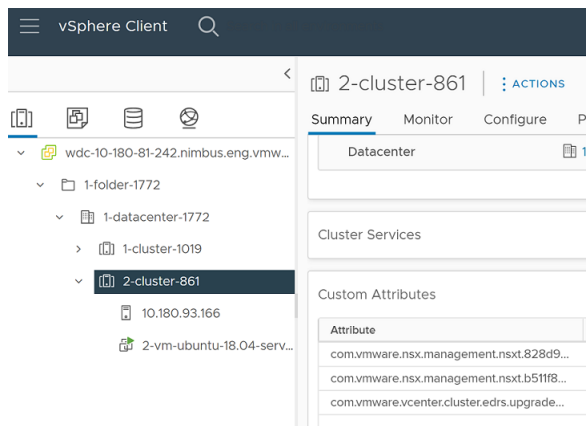
Starting with NSX 3.2.2, you can enable support of multiple NSX Managers managing a single VMware vCenter.

Important You can only enable the multiple NSX Managers managing a single VMware vCenter feature (configured by enabling the **Multi NSX** flag on the NSX UI) in VMware vCenter 7.0 and later versions.

For more details on how the **Multi NSX** flag is enabled in VMware vCenter, see [Add a Compute Manager](#).

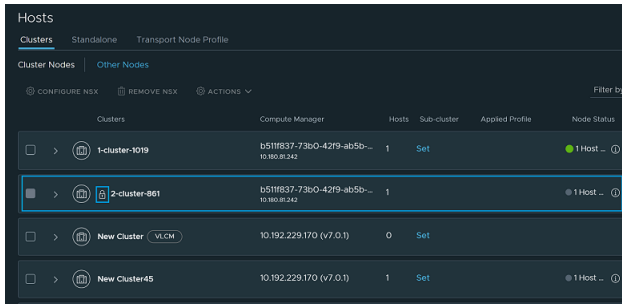


After you prepare a cluster or a host as a transport node, NSX appends the cluster, hosts and DVS extensions with a key that indicates these objects are managed by a certain VMware vCenter.



NSX changes the old extension (`com.vmware.nsx.management.nsx`) to a custom extension key (`com.vmware.nsx.management.nsx.<computemanager-id>`), where the `<computemanager-id>` is the VMware vCenter ID in NSX.

On the **Hosts** → **Clusters** page, NSX discovers all clusters managed by the same VMware vCenter, where a different NSX Manager can manage its own clusters. If another NSX Manager owns a cluster, you cannot prepare or edit it. These clusters are in read-only mode.



Switching between Single NSX Mode to Multiple NSX Mode

When you enable **Multi NSX** flag on the VMware vCenter, NSX appends its managed objects of NSX (cluster, host, Distributed Virtual Switch switch) with the custom extension (`com.vmware.nsx.management.nsx.<computemanager-id>`).

In Multiple NSX Mode, all NSX Managers registered to the same VMware vCenter must have **Multi NSX** flag enabled. You cannot configure **Multi NSX** enabled for NSX Manager-1 and deactivate on NSX Manager-2.

Host Movement Scenarios

Scenario	Action/Result
<ul style="list-style-type: none"> ■ Prepare Cluster-1 using TNP of NSX-1. ■ Prepare Cluster-2 using TNP of NSX-2. ■ In VMware vCenter UI, move one host from Cluster-1 to Cluster-2. <ul style="list-style-type: none"> ■ On the host that is moved to cluster-2, NSX-1 uninstalls NSX vib from the host. NSX Manager-1 removes its ownership from the host. Only after NSX Manager-1 removes its ownership and the Lock icon disappears from the System > Fabric > Hosts > Clusters does NSX Manager-2 start installation on the host. 	<ul style="list-style-type: none"> ■ NSX is uninstalled on the host that is moved to Cluster-2. After the host is moved, the host will be prepared by the TNP of NSX-2, which is attached to Cluster-2. ■ If there are any uninstallation-related errors, check the Host → Clusters page. Click Resolve to fix issues and proceed.
<ul style="list-style-type: none"> ■ Host-1 is prepared individually as a transport node in any of the following environments where: <ul style="list-style-type: none"> ■ It is not part of a VMware vCenter. ■ It is part of a VMware vCenter but TNP is detached from the cluster. ■ It is part of the datacenter in a VMware vCenter. 	<ul style="list-style-type: none"> ■ Before moving Host-1 to a NSX-managed cluster do the following: <ul style="list-style-type: none"> ■ Uninstall NSX from Host-1 transport node. ■ Add Host-1 to a VMware vCenter-managed cluster. ■ TNP is automatically applied to Host-1 and NSX is installed.
<ul style="list-style-type: none"> ■ Prepare Cluster-1 comprising of Host-1 Transport Node using TNP in NSX Manager-1. ■ Prepare Cluster-2 using TNP in NSX Manager-2. ■ Host-1 is a static member of NSGroup in NSX Manager-1. ■ From VMware vCenter, move Host-1 Transport Node to Cluster-2. 	<ul style="list-style-type: none"> ■ NSX cannot be removed from Host-1 Transport Node because it is part of NSGroup and the same host cannot be prepared in NSX Manager-2. You can find more details in log files. <p>Note This issue can occur even if Multi NSX functionality is not enabled. It can happen when you try to move a host between clusters.</p>

Limitations of Multiple NSX Managers Managing a Single VMware vCenter Setup

- If NSX Manager-1 has stamped its ownership on managed objects (cluster, host, or Distributed Virtual Switch (DVS)), these objects cannot be owned by NSX Manager-2 until the first manager gives up the ownership or ownership is forcefully passed on to another manager.
- Even though you can enable **Multi NSX** on a VMware vCenter, where the version is NSX 3.2.2, do not register the same VMware vCenter with NSX 3.2.1 or any previous release.
- Ensure the desired user roles have permissions to update `Global.ManageCustomFields` in VMware vCenter. The NSX Custom Attribute must not be appended on any of the managed objects. It can lead to disruption of the setup.
- With **Multi NSX** enabled on a VMware vCenter, you cannot enable Kubernetes cluster or vLCM cluster to work on the same VMware vCenter.
- If you deactivate **Multi NSX** on a VMware vCenter, you cannot use the same VMware vCenter to register with another NSX instance.
- If any custom or legacy VMware vCenter extension is not deleted from VMware vCenter for reason such as failure of NSX to come up, you will have to manually delete extension from VMware vCenter.
- Does not support collapsed cluster environments (where management and workloads are deployed on the same transport node). With Multi NSX flag enabled in a collapsed cluster environment, you cannot deploy new NSX Manager nodes. The workaround is to create a new cluster and deploy NSX Manager nodes.

Interoperability Matrix

The following table lists the solutions that are interoperable with the Multi NSX feature.

Feature/Solution	Supported
NSX Guest Introspection (GI) Platform	No
NSX Service Insertion (SI)	Yes
VMware vSphere with Tanzu	No
vSphere Lifecycle Manager (vLCM)	No
NSX Virtual Distributed Switch (N-VDS)	Yes
Note N-VDS is supported in versions before NSX 4.0.	
NSX Federation	Yes
VMware vSphere Distributed Resource Scheduler (DRS), VMware vSphere High Availability (HA) , VMware vMotion	Yes

Troubleshoot Multi-NSX Issues

Troubleshoot issues related to Multi-NSX environments. In a Multi-NSX environment, more than one NSX Manager is registered to a single VMware vCenter.

Cannot Enable or Disable Multiple NSX on a Single VMware vCenter

If an old NSX extension exists in VMware vCenter, you cannot enable or disable multiple NSX on the VMware vCenter as NSX is already registered.

Problem

Old extension still exists in VMware vCenter.

Cause

You registered NSX without enabling or disabling **Multi NSX** in VMware vCenter and deleted NSX Manager without unregistering the VMware vCenter.

Solution

- 1 Log in to <https://<vCenter Server hostname or IP address>/mob/?moid=ExtensionManager>.
- 2 On the ExtensionManager page, go to method "unregisterExtension".
- 3 Add old extension "com.vmware.nsx.management.nuxt" and click **Invoke** method.

Override NSX Ownership Constraints

When a transport node profile (TNP) is applied to a cluster if any validations (VMs are running on hosts) fail then Transport Node is not created. After ensuring validations pass, you can reapply cluster configuration using API. However, if you want the node to give up ownership of its managed objects you can call the `override_nsx_ownership` parameter in the API call.

You must only override ownership of managed objects from a NSX instance if any one of the following condition is true:

- NSX instance is not responding or unusable
- NSX is no longer actively managing its objects

Caution If the managed objects affected by this operation are actively used by the NSX that owns these objects, then it could corrupt host switch configurations that are pushed down by the NSX instance.

Prerequisites

Procedure

- ◆ Call the following API to override NSX ownership constraints.

```
POST https://<nsx-
mgr>/api/v1/fabric/discovered-nodes/5c669dc6-47a8-4508-3077-6a48f26c5a4g?
action=reapply_cluster_config&override_nsx_ownership=true
```

Where, `override_nsx_ownership=true` when set to True overrides NSX ownership of managed objects.

However, it is not recommended to pass this parameter. If you use this parameter it indicates you want to own certain managed objects owned by another NSX instance.

```
{
  "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"
          }
        ],
        "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"
        },
        "vmknics": [
          {
            "device_name": "vmk1",
            "logical_switch_id": "849e339e-64b7-47cb-9480-33068f70dc5a"
          }
        ]
      }
    ]
  },
  "transport_zone_endpoints": [
    {
```

```

        "transport_zone_id": "e14c6b8a-9edd-489f-b624-f9ef12afbd8f",
        "transport_zone_profile_ids": []
    }
],
"node_id": "41a4eebd-d6b9-11e6-b722-875041b9955d",
"node_deployment_info": {
    "resource_type": "HostNode",
    "id": "41a4eebd-d6b9-11e6-b722-875041b9955d",
    "display_name": "FN1",
    "fqdn": "sc2-rdops-vm03-dhcp-110-133.eng.vmware.com",
    "ip_addresses": [
        "10.192.156.199"
    ],
    "external_id": "a5113680-6a56-4386-9017-adffbe56c99e",
    "discovered_ip_addresses": [],
    "os_type": "ESXI",
    "os_version": "",
    "managed_by_server": "",
    "_create_time": 1538632396987,
    "_last_modified_user": "admin",
    "_protection": "NOT_PROTECTED",
    "_last_modified_time": 1538632396987,
    "_create_user": "admin",
    "_revision": 0
},
"_create_time": 1485299990773,
"_last_modified_user": "admin",
"_last_modified_time": 1485301913130,
"_create_user": "admin",
"_revision": 0
}

```

Note For more information on Override APIs, refer to the *NSX API guide*.

NSX no longer owns managed objects. These objects can be now owned by another NSX instance.

Preparing Standalone Host Cluster Results in Failure

Preparing a standalone host cluster for NSX 3.2.2 results in failure.

Problem

In this scenario, there are two NSX deployments connected to same VMware vCenter in NSX 3.1.x. The VMware vCenter is registered with NSX-1 as compute manager and NSX-2 is using standalone cluster registered to the same VMware vCenter.

You upgrade NSX-1 to v3.2.2, then enable multi-NSX flag on the VMware vCenter. However, you do not upgrade the NSX-2 which remains on 3.1.x. Since NSX-1 is enabled with multi-NSX, the NSX Manager UI lists the standalone cluster of NSX-2 as an unprepared cluster.

If you try to prepare the standalone NSX-2 cluster from NSX-1, the standalone host goes into failure state.

Cause

NSX-2 hosts go into failure state because it is running NSX v3.1.x.

Solution

- 1 Log in to `https://<vCenter-Server-IP>`.
- 2 Migrate all VMs from that host to the other hosts using VMware vCenter.
- 3 Remove NSX from both managers using below API:

```
DELETE https://<NSX Manager-IP>/api/v1/transport-nodes/<transportnode-id>?
force=true&unprepare_host=false
```

- 4 Configure NSX on the desired NSX Manager.

Traffic Performance Issues Related to NSX Edge VMs in Multiple NSX Environment

In a multiple NSX environment, only the NSX Manager which deployed NSX Edge VM can use it for routing and Inter TEP communication. None of the other NSX Managers registered to the same VMware vCenter can use it for routing and Inter TEP communication. The other NSX Manager instances consider the NSX Edge VM as a regular VM. This scenario can cause traffic performance issues on the NSX Edge VM.

Problem

In a multiple NSX scenario, you have the following configuration:

- NSX Manager-1 and NSX Manager-2 are registered to the same VMware vCenter (compute manager).
- NSX Manager-1 deployed the NSX Edge VM.
- NSX Manager-2 prepared the ESXi host.
- From vSphere Web Client, you perform vMotion of NSX Edge VM to an ESXi host prepared by NSX Manager-2. NSX Manager-2 did not deploy the NSX Edge VM.
- NSX Manager-1 does not recognize NSX Edge as an inventory VM. So, NSX Manager-1 does not apply any DFW rules on it.

After moving the NSX Edge VM to the new ESXi host:

- NSX Manager-2 categorizes NSX Edge as a regular VM and not as an NSX Edge VM. If there are any DFW rules configured, NSX Manager-2 applies any DFW rules on the NSX Edge VM.

See a sample output,

```
https://<NSX Manager-2>/api/v1/fabric/virtual-machines
{
  "host_id": "59ac4c38-56b1-4b82-a131-dd9ad119f53d",
  "source": {
    "target_id": "59ac4c38-56b1-4b82-a131-dd9ad119f53d",
    "target_display_name": "10.172.17.133",
```



```

        "target_type": "HostNode",
        "is_valid": true
    },
    ...
    "type": "REGULAR",
    "guest_info": {
        "os_name": "Ubuntu Linux (64-bit)",
        "computer_name": "vm"
    },
    "resource_type": "VirtualMachine",
    "display_name": "mgr2_edge1",
    "_last_sync_time": 1663802733277
},

```

Cause

As NSX Manager-2 exclude list does not filter out NSX Edge VM, it is considered as a regular VM and not as a NSX Edge VM. So, DFW rules or any third-party firewall rules configured for workloads are applied to the NSX Edge VM too. This scenario might cause traffic disruption.

Solution

- 1 Log in to the VMware vCenter, <https://vCenter-Server-IP>.
- 2 As NSX Manager-2 considers the NSX Edge VM as a regular VM, create an NS Group "Edge-VMs-From-Other-Managers" and add the Edge VMs to the NS Group.
- 3 To identify Edge VMs from NSX Manager-1 that must be added to Exclude lists on NSX Manager-2, follow these steps:
 - a Use `display_name` that you get after call the following API, https://<NSX Manager-1>/api/v1/transport-nodes?node_types=EdgeNode.
 - b Match the name with the `display_name` in API response from NSX Manager-2, <https://<NSX Manager-2>/api/v1/fabric/virtual-machines>.
- 4 Add NS Group "Edge-VMs-From-Other-Managers" to the DFW Exclusion List and SI Exclusion Lists on NSX Manager-2.
- 5 Verify and exclude NSX Edge VM from third party firewalls.
- 6 If Edge VM Id changes, update the NS Group.
- 7 Before you delete the Edge VM, remove the entry from the Exclude lists.

Note Inter TEP communication on the NSX Edge VM is not supported on NSX Manager-2.

Distributed Virtual Switch is Not Available for Selection in NSX

In a Multi NSX configuration, if one of the NSX instances registered with a VMware vCenter is removed, DVS switch associated with the removed NSX instance is not available for selection in the NSX UI.

Problem

To know whether the DVS associated to the NSX instance you removed is ready to be associated to another NSX instance, verify the status of the `owner_nsx` parameter. In the response, the `owner_nsx` parameter of the DVS did not reset. The value is `Other`. So, the DVS instance does not show up as an option when preparing NSX.

Run `GET https://<nsx-manager-ip>/api/v1/fabric/compute-collections/<compute-collection-id>:domain-<id>`

```
{
  "external_id" : "c39f2dea-fccd-4023-ab85-7e243a5df3e3:domain-c3633",
  "origin_type" : "VC_Cluster",
  "origin_id" : "c39f2dea-fccd-4023-ab85-7e243a5df3e3",
  "cm_local_id" : "domain-c3633",
  "owner_id" : "",
  "origin_properties" : [ {
    "key" : "lifecycleManaged",
    "value" : "false"
  }, {
    "key" : "dasConfig.enabled",
    "value" : "false"
  }, {
    "key" : "drsConfig.enabled",
    "value" : "false"
  }, {
    "key" : "drsConfig.defaultVmBehavior",
    "value" : "fullyAutomated"
  }, {
    "key" : "configManagerEnabled",
    "value" : "false"
  }, {
    "key" : "configurationEx.vsanConfigInfo.enabled",
    "value" : "false"
  } ],
  "owner_nsx" : "OTHER",
  "resource_type" : "ComputeCollection",
  "display_name" : "a1",
  "description" : "",
  "_last_sync_time" : 1686185879004
}
```

The `owner_nsx` field is set to `Other` indicating it is still used by the NSX instance. And the DVS is not available for selection in the NSX UI.

Solution

- 1 Go to the VMware vCenter, clean up DVS entries.
- 2 Using VMware vCenter MOB UI interface, remove the custom field definition for the DVS managed object.

For example, to delete **dvs-1897** from the VMware vCenter MOB UI interface, you need to find the data center and the network folder where the DVS object exists.

In the VMware vCenter MOB UI interface, go to **content** → **group-d1 (datacenters)** → **datacenter-1695 (Data center)** → **group-n1699 (network folder)** → **more**. Find the DVS folder, **dvs-1897(DVS7N-EXT)** → **value [109]** or **customerValue[109]**.

- 3 Navigate to **CustomFieldsManager** → **RemoveCustomFieldDef** and provide custom key integer value.
- 4 Delete the DVS entry from the VMware vCenter and then retry to assign the same DVS from the NSX UI. After cleaning up the DVS entries in VMware vCenter, new NSX instances can use the DVS object that is deleted. Similarly, you can use the same DVS within a Transport Node Profile to configure clusters.

Managing Transport Nodes

After preparing hosts as transport nodes, you can view status of hosts, switch visualization, and other configuration settings related to transport nodes. You can use it to debug issues if the host transport node is in a degraded or failed state.

Switch Visualization

You get a granular view of a vSphere Distributed Switch (VDS) at an individual host level.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System > Fabric > Hosts**.
- 3 Select **Clusters** or **Standalone Hosts** tab.
- 4 For any host, click **View Details**.
- 5 In the **Host Details** window, select the **Overview** tab and see visualization of VDS switch.
The VDS switch visualization shows the uplinks and the NICs the switch is connected to.

NSX Maintenance Mode

If you want to avoid vMotion of VMs to a transport node that is not functional, place that transport node in NSX Maintenance Mode.

To put a transport node in NSX Maintenance Mode, select the node, click **Actions** → **NSX Maintenance Mode**.

When you put a host in NSX Maintenance Mode, the transport node cannot participate in networking. Therefore, you must vMotion all VMs to another host before initiating NSX Maintenance Mode. Also, VMs running on other transport nodes that have VDS as the host switch cannot be vMotioned to this transport node. In addition, logical network cannot be configured on ESXi hosts as the status of host switch on the ESXi host would be shown as **Down**.

Scenarios to put the transport node in NSX Maintenance Mode:

- A transport node is not functional.
- If a host has hardware or software issues that are unrelated to NSX, but you want to retain the node and its configurations in NSX, place the host in NSX Maintenance Mode.
- A transport node is automatically put in NSX Maintenance Mode when an upgrade on that transport node fails.

Any transport node put in the NSX Maintenance Mode is not upgraded.

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>> or <https://<nsx-manager-fqdn>>.
- 2 Go to **System > Fabric > Hosts** and select the **Cluster** tab.
- 3 Expand a cluster and select a host.
- 4 Click **Actions > Migrate ESX VMkernel and Physical Adapters**.
- 5 In the Migrate ESX VMkernel and Physical Adapters, enter the following details.

Field	Description
Direction	Make a selection: <ul style="list-style-type: none"> ■ Migrate to Logical Switches: To migrate VMkernel adapters from a VSS or VDS switch to an N-VDS switch in NSX. ■ 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.

- 6 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.

View Bidirectional Forwarding Detection Status

To monitor the health of NSX overlay fabric, view Bidirectional Forwarding Detection (BFD) status between transport nodes.

Each transport node creates a full mesh of BFD sessions to all the Tunnel Endpoints (TEPs) that are active on one or more logical spans. NSX displays the BFD status among other details related to the transport node.

Both Host Transport nodes (standalone and hosts registered to a vCenter) and Edge nodes display the tunnel status. BFD packets support both GENEVE and STT encapsulation. GENEVE is the default encapsulation.

Note For compute transport nodes such as an ESXi host, BFD tunnels are formed if ESXi has an active port attached to an NSX segment. It means a powered-on VM with a vNIC is connected to an NSX segment.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select **System > Fabric > Hosts > Clusters**.
- 3 Select a host and click **View Details**.
- 4 On the **Host Details** window, select **Monitor** and expand **Transport Node Status**.
- 5 Select **Tunnels**.
- 6 On the Tunnel Endpoint, filter tunnels based on the encapsulation protocol they are using. Choose between **GENEVE** or **VXLAN**.
- 7 In the **Filter by BFD Status** drop-down menu, select ALL to view all BFD statuses or a specific status.

The Monitor page displays the status of tunnel, BFD diagnostic code, remote node UUID, encapsulation on BFD packets, and tunnel name.

The tunnel BFD diagnostic code indicates the reason for the change in the session state.

Code	Description	Action
0	No Diagnostic	Code 0: Default diagnostic code seen when the tunnel is Up.
1	Control Detection Time Expired	Code 1: BFD timer has expired. It is seen when the local interface has not received a BFD packet from the remote system resulting in BFD timer expiring. Check if BFD timer is too aggressive with respect to system load and path traffic load. Default BFD timer is 1-sec, up to 3 misses. Change in BFD parameters are not disruptive.
2	Echo Function Failed	Code 2: BFD echo packet loop verification failed. Verify health of transport-node.
3	Neighbor Signalled Session Down	Code 3: Peer node voluntarily brings the session down. Check if peer transport node is in maintenance mode or not healthy. un ICMP ping to verify connectivity to TEPs.
4	Forwarding Plane Reset	Code 4: When Forwarding-Plane is reset and peer does not reply on BFD, so session is marked down
5	Path Down	Code 5: The path to the remote node is down. Validate IP connectivity between the TEPs using ICMP ping. Note that the TEP interfaces on ESXi are instantiated on the vxlan netstack and on the tunnel VRF on edges. Be sure to initiate the ping from within the vxlan netstack on ESXi or from the tunnel VRF on edges. If you have more than one TEP, be sure to specify the source IP address or interface used for the ping. On ESXi hosts: <code>ping ++netstack=vxlan -I <vmk adapter> <remote address></code> On Edge nodes: <code>get logical-routers vrf 0 ping <dst-vtep> source <src-vtep> repeat 3</code>
6	Concatenated Path Down	Concat Path Down represents that the Edge transport node has lost all the BGP/OSPF (northbound) sessions to the northbound router. This error is safe to ignore when not using the Edge cluster for Tier-0 BGP routing purposes, while just the Tier-1 services are used.

Code	Description	Action
7	Administratively Down	Code 7: Session is marked down by the administrator. Verify if local Transport Node is in maintenance mode. Admin CLI to run on TN: <code>get maintenance-mode</code>
8	Reverse Concatenated Path Down	Code 8: The path from the remote node to local is down. Test IP connectivity from remote node to local node.

Results

The fabric health BFD sessions are created between the TEP addresses. The tunnel status is a true reflection of the IP connectivity and capability of the network to forward Geneve packets therefore status for all BFD sessions should be `Up`. If the BFD status is down, use the diagnostic code to troubleshoot the issue.

To know the status of BFD sessions on fabric nodes, run the following CLI commands:

- For ESXi, run `nsxcli bfd sessions list`.
- For Edge TN, run `get bfd-sessions`.

To verify the fabric health of transport node, call the following API:

```
GET /policy/api/v1/infra/sites/<site-id>/enforcement-points/<enforcement-point-id>/transport-node-status-report
```

where, `<site-id>` and `<enforcement-point-id>` can use the value `default`.

IPv6 Unsupported Features

NSX does not support a few features if you select IPv6 as the addressing scheme while configuring any of these objects or components:

- Host transport node
- NSX Edge transport node
- Transport node profile
- QuickStart card to configure NSX

The unsupported features are:

- EVPN
- Multicast Routing
- NSX Federation

Installing NSX Edge

10

Install NSX Edge on ESXi using the NSX UI, the vSphere web client, or the command-line OVF tool.

Read the following topics next:

- [NSX Edge Installation Requirements](#)
- [NSX Edge Networking Setup](#)
- [NSX Edge Installation Methods](#)
- [Create an NSX Edge Transport Node](#)
- [Configure NSX Edge DPDK Interfaces](#)
- [Manually Deploying NSX Edge Node Outside of NSX](#)
- [Create an NSX Edge Cluster](#)
- [Remove NSX Edge Nodes from an Edge Cluster](#)
- [Relocate and Remove an NSX Edge Node from an NSX Edge Cluster](#)

NSX Edge Installation Requirements

The NSX Edge provides routing services and connectivity to networks that are external to the NSX deployment. NSX Edge nodes provide a pool of capacity for running centralized services and 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 logical 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 10-1. NSX Edge Deployment, Platforms, and Installation Requirements

Requirements	Description
Supported deployment methods	<ul style="list-style-type: none"> ■ OVA/OVF ■ ISO with PXE ■ ISO without PXE
Supported platforms	NSX Edge is supported as VM (only on ESXi) or as physical server (bare metal). Both leverage the data plane development kit (DPDK) for faster packet processing and high performance.
PXE installation	The Password string must be encrypted with sha-512 algorithm for the root and admin user password.
NSX appliance password	<ul style="list-style-type: none"> ■ 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 VMTTools installed. Do not remove or upgrade VMTTools.
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	Plan your NSX Edge IPv4 and IPv6 IP addressing scheme. However, a NSX Edge configured for IPv4 and IPv6 stack is not supported on an ESXi node that is configured only for IPv6 addressing scheme. Also, NSX Edge tunnel endpoints on a dual IPv4 and IPv6 scheme cannot communicate with an ESXi node using only IPv6 addressing scheme.

Table 10-1. NSX Edge Deployment, Platforms, and Installation Requirements (continued)

Requirements	Description
OVF Template	<ul style="list-style-type: none"> ■ 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 <i>localhost</i>. ■ 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 VMs or Bare Metal Edges in an Edge cluster.

Intel-based Chipsets

NSX Edge nodes are supported on ESXi-based hosts with Intel chipsets. If an unsupported chipset type is used, vSphere EVC mode may prevent Edge nodes from starting, showing an error message in the console. See [NSX Edge VM System Requirements](#).

AMD EPYC

NSX Edge nodes are also supported on AMD-based chipsets. NSX Edge nodes can now be deployed on AMD EPYC series chipsets. See [NSX Edge VM System Requirements](#).

NSX Edge Support of vSphere Business Continuity Features

Starting in NSX 2.5.1, vMotion, DRS, and vSphere HA are supported for NSX Edge nodes. Use vMotion and DRS with caution to avoid traffic disruption, especially, when the BFD timers are set to less than one second.

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 any of the local users, the name must be unique. If you specify the same name, it is ignored and the default names (for example, **admin** or **audit**) are used.
- If the password for the **root** or **admin** user does not meet the complexity requirements, you must log in to NSX Edge through SSH or at the console as **root** with password **vmware** and **admin** with password **default**. You are prompted to change the password.

- If the password for other local users (for example, **audit**) 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 local_user_name** to set the local user's password (the current password is an empty string). You can also reset passwords in the UI using System > User Management > Local Users.

Caution Changes made to the NSX 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 VMware vCenter.

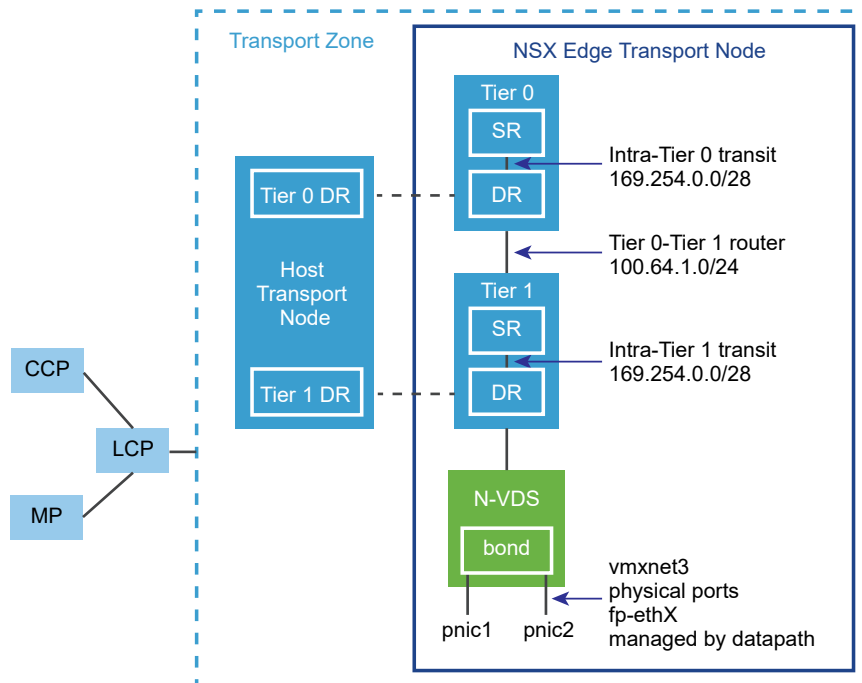
NSX Edge Networking Setup

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 shows two transport nodes in a transport zone. One transport node is a host. The other is an NSX Edge.

Figure 10-1. 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 100.64.0.0/16. Each tier-0-to-tier-1 peer connection is provided a /31 subnet within the 100.64.0.0/16 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 100.64.0.0/16 subnet is already in use in your deployment.

Each NSX 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 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-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-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 /api/v1/transport-nodes/{transport-node-id}/node/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.

Figure 10-2. NSX Edge Transport Node in VM Form Factor within ESXi


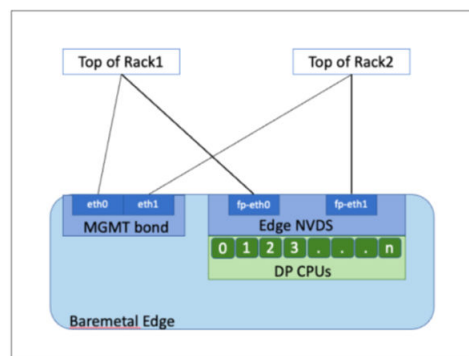


Figure 10-3. Bare Metal Edge Transport Node



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. 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 tunneling between transport nodes.
- VLAN for uplinks external to NSX.

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

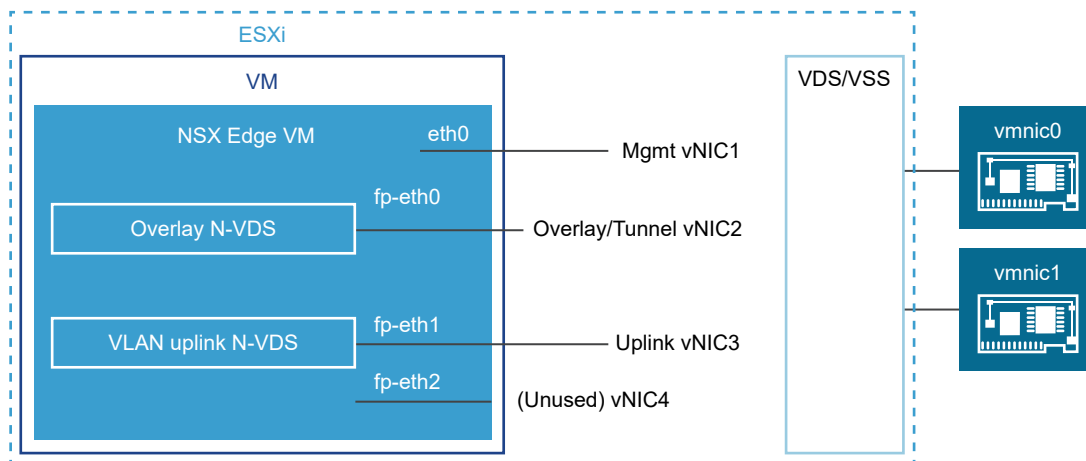
When you install NSX Edge as a virtual appliance or VM, internal interfaces are created, called fp-ethX, where X is 0, 1, 2, and 3. These interfaces are allocated for uplinks to a top-of-rack (ToR) switches and for NSX overlay tunneling.

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 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.

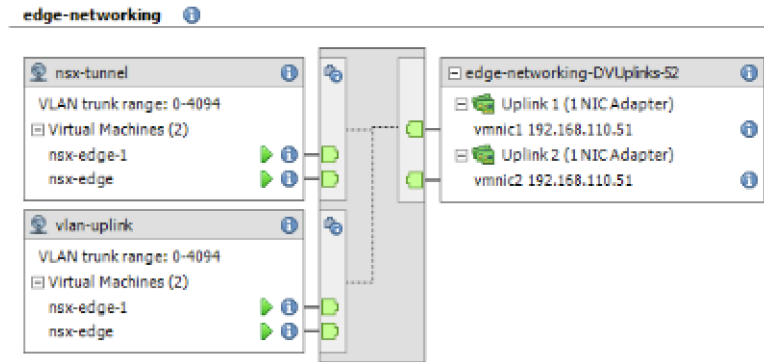
Figure 10-4. One Suggested Link Setup for NSX Edge VM Networking



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.

Figure 10-5. Edge Networking in vSphere



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 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 can be in the same VLAN or subnet.

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.

Note If NSX Edge VM uplinks (management, fast-path interfaces) will be connecting to VLAN based portgroups then underlying ESXi host does not need to be an NSX transport node. But if NSX Edge VM uplinks will be connecting to NSX VLAN segments (logical switches), then you must configure ESXi hosts hosting the edges as transport nodes.

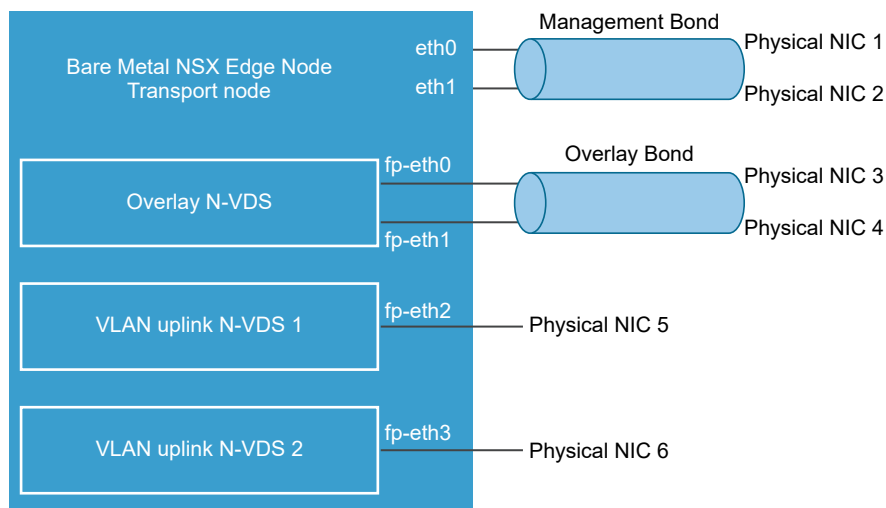
Bare-Metal NSX Edge Networking

The bare-metal NSX Edge contains internal interfaces called fp-ethX, where X is up to 16 interfaces. 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 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 overlay tunnel. fp-eth2 and fp-eth3 are used as redundant VLAN uplinks to TORs.

Figure 10-6. 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 10-7. 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:

- On a vDS switch running version prior (<) to 6.6, enable promiscuous mode for the port connected to NSX Edge VM virtual NIC that provides VLAN connectivity. These settings are needed to support bridging or L2VPN and DHCP functionality for VLAN networks.
- On a vDS switch running version equal to or greater than (>=) 6.6, enable mac learning and disable promiscuous mode. These settings ensure that packets are received at the destination where destination mac address does not match the virtual NIC effective MAC address. These settings also ensure packets are received at destinations that are on an NSX Segment. These settings are needed to support bridging or L2VPN and DHCP functionality for VLAN networks.
- Enable forged transmit on the vDS switch. Forged transmit enables sending packets with source mac address not matching the virtual NIC effective MAC addresses. These settings are needed to support bridging or L2VPN and DHCP functionality for VLAN networks.

NSX Edge Installation Methods

Install NSX Edge on an ESXi host using NSX Manager UI (recommended method), vSphere Web Client (from UI or vSphere command-line OVF tool) or as physical servers.

NSX Edge Installation Methods

Installation Method	Instructions
NSX Manager (recommended method to install an NSX Edge VM appliance only)	<ul style="list-style-type: none"> ■ Ensure NSX Edge network requirements are met. See NSX Edge Installation Requirements. ■ Create an NSX Edge transport node. See Create an NSX Edge Transport Node. ■ Create an NSX Edge cluster. See Create an NSX Edge Cluster.
vSphere web client or vSphere command-line OVF tool	<ul style="list-style-type: none"> ■ Ensure NSX Edge network requirements are met. See NSX Edge Installation Requirements. ■ Choose vSphere web client or vSphere command-line OVF tool to install NSX Edge. <ul style="list-style-type: none"> ■ (Web Client) Install NSX Edge on ESXi. See Install an NSX Edge on ESXi Using the vSphere GUI. ■ (Command-line OVF tool) Install NSX Edge on ESXi. See Install NSX Edge on ESXi Using the Command-Line OVF Tool. ■ Join NSX Edge with the Management Plane. See Join NSX Edge with the Management Plane. ■ Configure an NSX Edge as a transport node. See Edit NSX Edge Transport Node Configuration. ■ Create an NSX Edge cluster. See Create an NSX Edge Cluster.
Physical server (Automated or Interactive mode using ISO file) or NSX Edge VM appliance	<p>Install NSX Edge using ISO file using PXE on physical servers. Note that PXE boot installation procedure is not supported on NSX Manager.</p> <ul style="list-style-type: none"> ■ Ensure NSX Edge network requirements are met. See NSX Edge Installation Requirements. ■ Prepare PXE server. See Prepare a PXE Server for Bare Metal NSX Edge Installation. Choose from one of the supported installation methods: <ul style="list-style-type: none"> ■ (Automated installation) Install NSX Edge using ISO File on physical servers. See Install Bare Metal NSX Edge Automatically using ISO File. ■ (Automated installation) Install NSX Edge using ISO File as a Virtual Appliance. See Install NSX Edge via ISO File as a Virtual Appliance. ■ (Manual installation) Manually Install NSX Edge using ISO File. See Install Bare Metal NSX Edge Interactively using ISO File. ■ Join NSX Edge with the Management Plane. See Join NSX Edge with the Management Plane. ■ Configure an NSX Edge as a transport node. See Edit NSX Edge Transport Node Configuration. ■ Create an NSX Edge cluster. See Create an NSX Edge Cluster.

Create an NSX Edge Transport Node

NSX Edge nodes are service appliances with pools of capacity, dedicated to running network and security services.

NSX Edge nodes when configured as transport nodes host Tier-0 and Tier-1 gateways. They can be instantiated as a bare metal appliance or in virtual machine form factor. They are grouped in one or several clusters. Each cluster is representing a pool of capacity.

An NSX Edge can belong to one overlay transport zone and multiple VLAN transport zones. An 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/`. nsx-proxy does not create a certificate if a certificate already exists.

Important When you deploy an Edge Node through NSX Manager, the system records the node's MO-REF. This MO-REF is required to make requests to VMware vCenter for any subsequent operations that needs to be performed on the node, such as redeploy and delete. However, through customer inventory operations at VMware vCenter the MO-REF could change. If MO-REF changes, the NSX operations for that edge node will fail. For example, an edge node redeploy will fail to get rid of the node and the new node will get created with the same IP as the old one. To help you mitigate this issue, the system generates some alarms. For more information about these alarms, see the *NSX Administration Guide*.

Prerequisites

- Transport zones must be configured. See [Create Transport Zones](#).
- 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 NSX Edge nodes. See [Create an Uplink Profile](#).
- An IP pool must be configured or must be available in the network deployment. See [Create an IP Pool for Tunnel Endpoint IP Addresses](#).
- Prepare uplinks. For example, distributed port groups as trunk in vCenter Server or NSX Segments in NSX.
 - Create distributed trunk port groups in VMware vCenter for management, TEP and overlay networks if you plan to connect NSX Edge network interfaces to a VDS in VMware vCenter.
 - Create VLAN trunk segments in NSX if you plan to connect NSX Edge network interfaces to NSX VLAN segments or logical switches.
- Before you can use NSX Edge VM datapath interfaces in Uniform Passthrough (UPT) mode, meet the following conditions:

Note UPT mode is not supported on NSX Edge Bare Metal hosts.

- NSX Edge hardware version is 20 (vmx-20) or later. Previous NSX Edge hardware versions do not support UPT mode.
- Verify that the memory reservation on the configured NSX Edge is set to 100%.
- From the vSphere Web Client, enable UPT on the NSX Edge VM network adapter. See the *Change the Virtual Machine Network Adapter Configuration* topic in *vSphere Virtual Machine Administration* guide.

- At least one of the NSX Edge VM datapath interface must be backed by an ESXi host that hosts a Data Processing Unit-based SmartNIC. A SmartNIC is a NIC card that provides network traffic processing using a Data Processing Unit (DPU), a programmable processor on the NIC card, in addition to the traditional functions of a NIC card. For more information related to DPU, see [NSX on vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine](#).
- Starting with NSX 4.0.1.1, NSX Edge VM hardware version will no longer default to `virtualHW.version 13`. NSX Edge VM hardware will depend on the underlying version of the ESXi host. VM hardware versions compatible with ESXi hosts are listed in KB article [2007240](#).

■

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Select **System > Fabric > Nodes > Edge Transport Nodes > Add Edge Node**.
- 3 Type a name for the NSX Edge.
- 4 Type the Host name or FQDN in the format *subdomain.example.com*.
- 5 Select the form factor for the NSX Edge VM appliance.

- 6 To customize CPU and memory allocated to an NSX Edge VM appliance, tune the following parameters. However, for maximum performance NSX Edge VM appliance must be assigned 100% of the available resources.

Caution If you customize resources allocated to the NSX Edge VM, turn back the reservation later on to 100% to get maximum performance.

Option	Description
Memory Reservation (%)	<p>Reservation percentage is relative to the pre-defined value in the form factor.</p> <p>100 indicates 100% of memory is reserved for the NSX Edge VM.</p> <p>If you enter 50, it indicates that 50% of the allocated memory is reserved for the Edge transport node.</p> <p>Note If you want to use NSX Edge VM datapath interfaces in UPT mode, reserve 100% of the allocated memory for the NSX Edge transport node.</p>
CPU Reservation Priority	<p>Select the number of shares to be allocated to an NSX Edge VM relative to other VMs that are contending for shared resources.</p> <p>The following shares are for an NSX Edge VM in Medium form factor:</p> <ul style="list-style-type: none"> ■ Low - 2000 shares ■ Normal - 4000 shares ■ High - 8000 shares ■ Extra High - 10000 shares
CPU Reservation (MHz)	<p>Caution Unless you need fine grained control over CPU reservations, do not use this field. Instead, change CPU reservations from the CPU Reservation Priority field.</p> <p>The maximum CPU reservation value must not exceed the number of vCPUs multiplied by the normal CPU operation rate of the physical CPU core.</p> <p>If the MHz value entered exceeds the maximum CPU capacity of the physical CPU cores, the NSX Edge VM might fail to start even though the allocation was accepted.</p> <p>For example, consider a system with two Intel Xeon E5-2630 CPUs. Each CPU contains ten cores running at 2.20 GHz. The maximum CPU allocation for a VM configured with two vCPUs is 2 x 2200 MHz = 4400 MHz. If CPU reservation is specified as 8000 MHz, the reconfiguration of the VM completes successfully. However, the VM fails to power on.</p>

- 7 In the Credentials window, enter the following details.
- 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
- To enable SSH for an administrator, toggle the **Allow SSH Login** button.
- To enable SSH for a root user, toggle the **Allow Root SSH Login** button.
- Enter credentials for the Audit role. If you do not enter credentials in the **Audit Credentials** section, the audit role remains disabled.

Note After deploying the NSX Edge node, you cannot change the SSH setting for a root user that you set during deployment. For example, you cannot enable SSH for a root user if you disabled it during deployment.

8 Enter the NSX Edge details.

Option	Description
Compute Manager	Select the compute manager from the drop-down menu. The compute manager is the VMware vCenter 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.

9 Enter the NSX Edge management interface details.

Option	Description
Management IP Assignment	<p>This specifies the IP version used for the IP address assigned to the NSX Edge node which is required to communicate with NSX Manager and NSX Controller.</p> <p>Select IPv4 Only or IPv4 & IPv6.</p> <ul style="list-style-type: none"> ■ If you select IPv4 Only, select DHCP or Static IP. <ul style="list-style-type: none"> If you select Static, enter the values for: <ul style="list-style-type: none"> ■ Management IP: Enter the IP address of NSX Edge in the CIDR notation. ■ Default gateway: Enter the gateway IP address of NSX Edge. ■ If you select IPv4 & IPv6, enter the values for: <ul style="list-style-type: none"> ■ Management IP: Enter the IP address of NSX Edge in the CIDR notation. ■ Default gateway: Enter the gateway IP address of NSX Edge.
Management Interface	<p>From the drop-down menu, select the interface that connects to the NSX Edge management network. This interface must either be reachable from NSX Manager or must be in the same management interface as NSX Manager and NSX Controller.</p> <p>The NSX Edge management interface establishes communication with the NSX Manager management interface.</p> <p>The NSX Edge management interface is connected to distributed port groups or segments.</p>
Search Domain Names	Enter domain names in the format 'example.com' or enter an IP address.
DNS Servers	Enter the IP address of the DNS server.
NTP Servers	Enter the IP address or FQDN of the NTP server.
Enable UPT mode for datapath interface	<p>Enable Uniform Passthrough (UPT) mode on NSX Edge datapath interfaces to have direct I/O access or passthrough to the virtual network adapter. It improves overall performance of the NSX Edge node.</p> <p>Before you enable this field, ensure:</p> <ul style="list-style-type: none"> ■ NSX Edge hardware version is 20 or vmx-20 or later. Earlier hardware version do not support UPT mode. ■ ESXi host version must be 8.0 or later. <p>Caution To make UPT settings effective on NSX Edge VM virtual network adapters, NSX Manager puts NSX Edge VM into maintenance mode, powers it off and powers it back on again.</p>

10 Enter the N-VDS information.

Consider these points before you configure vNICs of NSX Edge nodes:

An N-VDS switch is hosted inside the Edge node VM with four fast path vNICs and one management vNIC.

- One vNIC is dedicated to management traffic.
- One vNIC is dedicated to overlay traffic (fp-eth0 DPDK fastpath interface).

- Two vNICs are dedicated to external traffic (fp-eth1, fp-eth2 DPDK fastpath interfaces).

Option	Description
Edge Switch Name	Enter a name for the switch or keep the default name.
Transport Zone	<p>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 connectivity and a VLAN for uplink connectivity.</p> <p>Note NSX Edge nodes support multiple overlay tunnels (multi-TEP) when the following prerequisites are met:</p> <ul style="list-style-type: none"> ■ TEP configuration must be done on one N-VDS only. ■ All TEPs must use the same transport VLAN for overlay traffic. ■ All TEP IPs must be in the same subnet and use the same default gateway.
Uplink Profile	<p>Select the uplink profile from the drop-down menu. The available uplinks depend on the configuration in the selected uplink profile.</p> <p>Note NSX Edge nodes support uplink profiles with Failover teaming policy (with single active uplink and no standby) and Loadbalancer Source teaming policy (with multiple active uplinks) only.</p>
IP Address Type (TEP)	<p>Select the IP version to be used for the tunnel endpoint (TEP). The options are IPv4 and IPv6.</p> <p>Important Ensure that the transport node forwarding mode and TEP IP address type are the same. For example, if the transport node forwarding mode is set to IPv6, set the TEP IP address type to IPv6. If they are different, a loss of traffic may result.</p>
IPv4 Assignment (TEP)	<p>This field appears when IP Address Type (TEP) is set to IPv4.</p> <p>Choose how IPv4 addresses are assigned to the NSX Edge switch that is configured. It is used as the tunnel endpoint of the NSX Edge. The options are:</p> <ul style="list-style-type: none"> ■ Use IP Pool: Select the IPv4 pool. ■ Use Static IPv4 List: Specify the following fields: <ul style="list-style-type: none"> ■ Static IP List: Enter a list of comma-separated IPv4 addresses to be used by the NSX Edge. ■ IPv4 Gateway: Enter the default gateway of the TEP, which is used to route packets another TEP in another network. For example, ESXi TEP is in 20.20.20.0/24 and NSX Edge TEPs are in 10.10.10.0/24 then we use the default gateway to route packets between these networks. ■ IPv4 Subnet Mask: Enter the subnet mask of the TEP network used on the NSX Edge.

Option	Description
IPv6 Assignment (TEP)	<p>This field appears when IP Address Type (TEP) is set to IPv6.</p> <p>Choose how IPv6 addresses are assigned to the NSX Edge switch that is configured. It is used as the tunnel endpoint of the NSX Edge. The options are:</p> <ul style="list-style-type: none"> ■ Use IP Pool: Select the IPv4 pool. ■ Use Static IPv6 List: Specify the following fields: <ul style="list-style-type: none"> ■ Static IP List: Enter a list of comma-separated IPv4 addresses to be used by the NSX Edge. ■ IPv6 Gateway: Enter the default gateway of the TEP, which is used to route packets another TEP in another network. ■ IPv6 Subnet Mask: Enter the subnet mask of the TEP network used on the NSX Edge.
DPDK Fastpath Interfaces / Virtual NICs	<p>Map uplinks to DPDK fastpath interfaces.</p> <p>Starting with NSX release 2.5, single N-VDS deployment mode is recommended for both bare metal and NSX Edge VM. See Configure NSX Edge DPDK Interfaces.</p> <p>Starting with NSX 4.0.1, you can map uplinks to DPDK fastpath interfaces that are backed by smartNIC-enabled DVPGs, VLAN logical switches or segments. The prerequisite is to enable UPT mode on NSX Edge VM virtual network adapters. The UPT mode requires at least one DPDK interface to be backed by smartNIC-enabled hardware also known as Data Processing Unit (DPU)-backed networks.</p> <p>Note If the uplink profile applied to the NSX Edge node is using a Named Teaming policy, ensure the following condition is met:</p> <ul style="list-style-type: none"> ■ All uplinks in the Default Teaming policy must be mapped to the corresponding physical network interfaces on the Edge VM for traffic to flow through a logical switch that uses the Named Teaming policies. See Configure Named Teaming Policy.
	<p>You can configure a maximum of four unique data path interfaces as uplinks on a NSX Edge VM.</p> <p>When mapping uplinks to DPDK Fastpath Interfaces, if NSX Edge does not display all the available interfaces (four in total), it means that either the additional interface is not yet added to the NSX Edge VM or the uplink profile has fewer number of uplinks.</p> <p>For NSX Edge VMs upgraded from an earlier version of NSX to 3.2.1 or later, invoke the redeploy API call to redeploy the NSX Edge VM. Invoking the redeploy API ensures the NSX Edge VM deployed recognizes all the available datapath interfaces in NSX Manager UI. Make sure the Uplink profile is correctly configured to use additional datapath NIC.</p> <p>For more information on configuring NSX Edge DPDK fastpath interfaces, see Configure NSX Edge DPDK Interfaces.</p> <ul style="list-style-type: none"> ■ For autodeployed NSX Edges (edge nodes deployed from the NSX Manager UI or API), call the redeploy API. The following API is deprecated. <pre style="background-color: #f0f0f0; padding: 10px;">POST api/v1/transport-nodes/<transport-node-id>? action=redeploy</pre>

Option	Description
	<ul style="list-style-type: none"> ■ For manually deployed edges (edges deployed using OVA/OVF file from the VMware vCenter UI or API), deploy a new NSX Edge VM. Ensure all the vmx customizations of the old NSX Edge VM are also done for the new NSX Edge VM. <p>Performing vMotion on an NSX Edge VM can result in ESXi running out of resources from a shared buffer pool if you create large VMs with multiple vNICs that use large sized ring buffers. To increase the depth of the shared buffer, modify the <code>ShareCOSBufSize</code> parameter in ESXi. To configure buffer size, see https://kb.vmware.com/s/article/76387.</p>

Note

- LLDP profile is not supported on an NSX Edge VM appliance.
- Uplink interfaces are displayed as **DPDK Fastpath Interfaces** if the NSX Edge is installed using NSX Manager or on a Bare Metal server.
- Uplink interfaces are displayed as **Virtual NICs** if the NSX Edge is installed manually using vCenter Server.

- 11 View the connection status on the **Transport Nodes** page.

After adding the NSX Edge as a transport node, the **Edge Transport Nodes** page will show the Configuration status as `Success` and Node Status as `Up` in about 10-12 mins.

- 12 Verify the transport node status by running the `Get edge-cluster-status | get managers | get controllers | get host-switch` CLI command.
- 13 (Optional) View the transport node by calling the `GET /api/v1/transport-nodes/{transport-node-id}/status | state` (deprecated) API call.

```
GET api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/
host-transport-nodes/<host-transport-node-id>/state | status
```

The default values for `enforcementpoint-id` and `site-id` is **default**.

Note After an NSX Edge node is migrated to a new host using vCenter Server, you might find NSX Manager UI reporting stale configuration details (Compute, Datastore, Network, SSH, NTP, DNS, Search Domains) of the NSX Edge. To refresh latest NSX Edge configuration details on NSX Manager, run the API command. `POST api/v1/transport-nodes/<transport-node-id>?action=refresh_node_configuration&resource_type=EdgeNode`

Important You can change the IP address of the NSX Edge node from the command line interface. At the CLI terminal, run `set interface eth0 ip <Gateway_IPAddress> gateway <NSXEdge_IPAddress> plane mgmt`. For example, `set interface eth0 ip <edge-new-ip-address/cidr> gateway <gateway-ip-address> plane mgmt`.

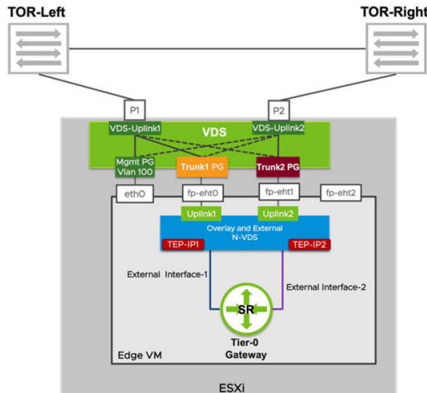
What to do next

Add the NSX Edge node to an NSX Edge cluster. See [Create an NSX Edge Cluster](#).

Configure NSX Edge DPDK Interfaces


Configure an NSX Edge VM with an N-VDS switch to carry both overlay and external traffic.

In this topology, NSX Edge VM is configured with an N-VDS switch to carry overlay and external traffic.



In this topology, the uplink profile attached is configured to use Multi-TEP to provide load balancing for overlay traffic by selecting default teaming to be Load Balancing Source teaming policy with active 'uplink1' and 'uplink2' on transport VLAN 200.

(optional) The uplink profile attached is also configured to use named teaming policy, where 'vlan300-policy' is mapped to uplink1 and 'vlan400-policy' is mapped to uplink2.

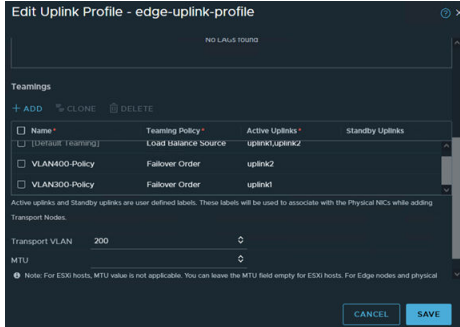


To create the topology, follow these steps.


Prerequisites

Procedure

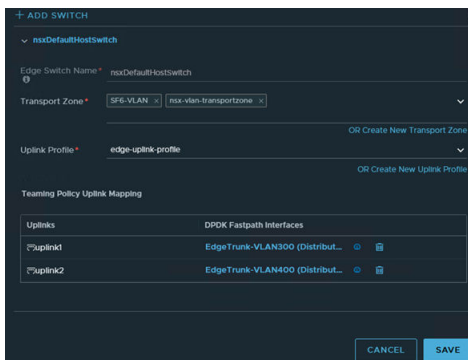
- 1 Uplink Profile creation with named-teaming policy for VLAN networks and default teaming for overlay networks. If named-teaming does not exist in uplink profile then default teaming is used for all networks.



- 2 VLAN Transport Zone is created or modified to use named teaming policy VLAN300-Policy and VLAN400-Policy (if using named-teaming policy).



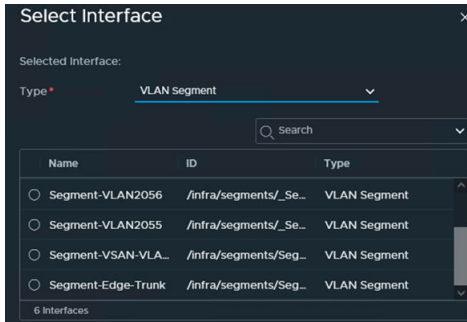
- 3 In the NSX Edge configuration, “Uplink1” (fp-eth0) is vNIC2 on Edge VM and is mapped to use VLAN 300 Trunk PG and “Uplink2”(fp-eth1) is vNIC3 on Edge VM and is mapped to use VLAN 400 Trunk portgroup.



- 4 NSX Edge VM interfaces can connect to VDS Portgroups in vCenter or NSX VLAN Segments.

Note If Edge interfaces will be connecting to NSX VLAN Segments, ESXi Hosts (hosting the Edge VMs) should be configured as Host Transport Nodes and member of VLAN transport zone.

An example of connecting VM interfaces to NSX VLAN segments.




An example of an NSX Edge VM with interfaces on VMware vCenter VDS Edge Trunk Portgroups

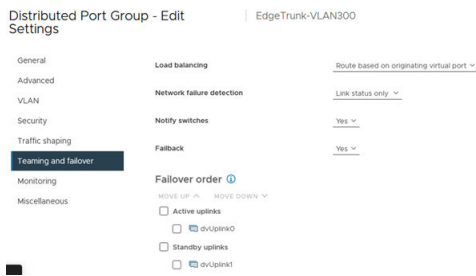
VM Hardware

> CPU	8 CPU(s)
> Memory	32 GB, 0.96 GB memory active
> Hard disk 1	196 GB
> Network adapter 1	DVS7-mgmt (connected)
> Network adapter 2	EdgeTrunk-VLAN300 (connected)
> Network adapter 3	EdgeTrunk-VLAN400 (connected)
> Network adapter 4	DVS7-EdgeTrunk (connected)

- a If VM interfaces connect to NSX VLAN segments, named-teaming is enabled on the segments. The diagram shows that External VLAN segment 300 is configured to use a named teaming policy “Vlan300-Policy” that sends traffic from this VLAN on “Uplink1” (vNIC2 of Edge VM). “External VLAN segment 400” is configured to use a named teaming policy “Vlan400-Policy” that sends traffic from this VLAN on “Uplink2” (vNIC3 of Edge VM).



- b If VM interfaces connect to VDS Portgroups in VMware vCenter, then “Trunk1 PG” is configured to use active uplink as “VDS-Uplink1” and standby uplink as “VDS-Uplink2”. “Trunk2 PG” is configured to use active uplink as “VDS-Uplink2” and standby uplink as “VDS-Uplink1”. This configuration ensures that the traffic sent on “External VLAN Segment 300” uses vNIC2 of Edge VM to exit the Edge VM and then “VDS-Uplink1” and is sent to the left TOR switch. Similarly, traffic sent on VLAN 400 uses “VDS-Uplink2” and is sent to the TOR switch on the right.



Important

- ESXi TEP and Edge TEP share the same VLAN: Use this configuration only if NSX Edge Trunk 1 portgroup and NSX Edge Trunk 2 portgroup are created from NSX as VLAN segments. Because any traffic between the Edge TEP to its own hypervisor's TEP does not need to exit the hypervisor.
- ESXi TEP and Edge TEP use different VLANs: Use this configuration if NSX EdgeTrunk1 portgroup and NSX EdgeTrunk2 portgroup are created as VDS portgroups from VMware vCenter. Any traffic between the NSX Edge TEP and its hypervisor TEP must exit the ESXi. The top of rack switch must then route it back toward the ESXi.

Manually Deploying NSX Edge Node Outside of NSX

In addition to deploying and configuring NSX Edge from the NSX Manager, you can manually deploy NSX Edge either as a VM in VMware vCenter or a Bare Metal server.

Install an NSX Edge on ESXi Using the vSphere GUI

You can use the vSphere Web Client or vSphere Client to interactively install an NSX Edge on ESXi.

Note Starting in NSX 2.5.1, the NSX Edge VM supports vMotion.

Prerequisites

See NSX Edge network requirements in [NSX Edge Installation Requirements](#).

Procedure

- 1 Locate the NSX Edge node appliance OVA file on the [Broadcom Support](#) page. 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 node 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, and click **Next**.
- 5 Enter a name and location for the NSX Edge node , and click **Next**.

The name you type appears in the VMware vCenter and vSphere inventory.

- 6 Select a compute resource for the NSX Edge node appliance, and click **Next**.
- 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 Manager VM and Host Transport Node System Requirements](#).
- 8 Review and verify the OVF template details, and click **Next**.
- 9 Select a deployment configuration, **Small**, **Medium**, **Large**, or **XLarge** and click **Next**.
The Description panel on the right side of the wizard shows the details of selected configuration.
- 10 Select storage for the configuration and disk files, and click **Next**.
 - a Select the virtual disk format.
 - b Select the VM storage policy.
 - c Specify the datastore to store the NSX Edge node appliance files.
- 11 Select a destination network for each source network.
 - a For network 0, select the VDS management portgroup.
 - b For networks 1, 2, and 3, select the previously configured VDS trunk portgroups.
- 12 Configure IP Allocation settings.
 - a For IP allocation, specify **Static - Manual**.
 - b For IP protocol, select **IPv4**.
- 13 Click **Next**.
The following steps are all located in the Customize Template section of the Deploy OVF Template wizard.
- 14 Enter the NSX Edge node system root, CLI admin, and audit passwords.

Note In the Customize Template window, ignore the message `All properties have valid values` that is displayed even before you have entered values in any of the fields. This message is displayed because all parameters are optional. The validation passes as you have not entered values in any of the fields.

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.

- 15 (Optional) If you have an available NSX Manager and want to register the NSX Edge node with the management plane during the OVA deployment, complete the Manager IP, Username, Password, and Thumbprint.

- Manager IP: Enter the NSX Manager node IP address.

Note Do not register the NSX Edge node with the virtual IP (VIP) address of the management plane during the OVA deployment.

- Manager Username: Enter the NSX Manager username.
- Manager Password: Enter the NSX Manager password.
- Manager Thumbprint: Enter the NSX Manager thumbprint.

Note An NSX Manager thumbprint is required to join an NSX Edge node to an NSX Manager. To retrieve thumbprint on an NSX Manager node, run `get certificate api thumbprint`.

- Node ID: Leave the field blank. The Node UUID field is only for internal use.

- 16 (Optional) If you want to deploy the NSX Edge node as an autonomous edge in a L2 VPN topology, enable the option **External and HA** sections. An autonomous edge is not managed by NSX. Do not enable the option if you want to deploy an NSX Edge node that provides centralized edge services to host transport nodes in an NSX topology.

Note The fields in the External and HA sections are required only when you configure an autonomous NSX Edge node.

- 17 Enter the hostname of the NSX Edge.

- 18 Enter the default gateway, management network IPv4, and management network netmask address.

Skip any VMC network settings.

- 19 Enter the DNS Server list, the Domain Search list, and the NTP Server IP or FQDN list.

- 20 (Optional) Do not enable SSH if you prefer to access NSX Edge using the console. However, if you want root SSH login and CLI login to the NSX Edge command line, enable the SSH option.

By default, SSH access is disabled for security reasons.

- 21 (Optional) In the **Internal Use Only** section, if you want to enable NSX Edge in uniform passthrough (UPT) or direct access mode to IO devices for improved performance, enable **Datapath UPT Mode Enabled** field.

Note Meet these prerequisites before you enable UPT on NSX Edge:

- NSX Edge hardware version is 20 or vmx-20 or later. Earlier hardware version do not support UPT mode.
 - ESXi host version must be 8.0 or later.
-

Caution Enabling UPT requires restart of the NSX Edge node.

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

The installation might take 7-8 minutes.

- 23 Power on the appliance. Open the console of the NSX Edge node to track the boot process. If the console window does not open, make sure that pop-ups are allowed.

- 24 After the NSX Edge node starts, log in to the Edge node using the console or SSH (provided SSH is enabled at the time of install) with admin credentials.

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

- 25 Run the `get interface eth0` (without VLAN) to verify that the IP address was applied as expected.

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

Interface: eth0
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 nodes on non-NSX managed host, verify that the minimum MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

- 26 Run the `get managers` command to verify that the NSX Edge node is registered.

```
- 10.173.161.17 Connected (NSX-RPC)
- 10.173.161.140 Connected (NSX-RPC)
- 10.173.160.204 Connected (NSX-RPC)*
```

The NSX Manager with * next to it is the active manager for the NSX Edge transport node VM.

27 If NSX Edge is not registered with the management plane, see [Join NSX Edge with the Management Plane](#).

28 Verify that the NSX Edge node has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge node and verify the following:

- You can ping your NSX Edge node management interface.
- From the NSX Edge node, you can ping the node's default gateway.
- From the NSX Edge node, you can ping the hypervisor hosts that are either in the same network or a network reachable through routing.
- From the NSX Edge node, you can ping the DNS server and NTP Server.

29 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 node 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 assign management IP address to the correct NIC.

- a Log in to the NSX Edge CLI and type the **stop service dataplane** command.
- b (Static IP) Run the `set interface <interface-name> ip <x.x.x.x/24> gateway <x.x.x.x> plane mgmt` command.
- c (DHCP) Run the `set interface interface-name dhcp plane mgmt` command.
- 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 node.

What to do next

Configure NSX Edge as a transport node. See [Edit NSX Edge Transport Node Configuration](#).

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 or FQDN for the NSX Manager to use.
- Create a management VDS and target VM port group in vCenter. Place the NSX appliances onto this management VDS port group network. See [Prepare a vSphere Distributed Switch for NSX](#).

Multiple management networks can be used as long as the NSX Manager nodes has consistent connectivity and recommended latency between them.

Note If you plan to use Cluster VIP, all NSX Manager appliances should belong to same subnet.

- Plan your NSX Manager IP and NSX Manager Cluster VIP addressing scheme.

Note Verify that you have the hostname for NSX Manager to use. The Hostname format must be `nsx-manager-fqdn@domain-name.com`. This format is required if NSX installation is dual stack (IPv4, IPv6) and/or if planning to configure CA-signed certificates.

- See NSX Edge network requirements in [NSX Edge Installation Requirements](#).
- 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.
- Know parameters that you can use to deploy a NSX Edge VM and join it to the management plane.

Field Name	OVF Parameter	Field Type
System root password	<code>nsx_passwd_0</code>	Required to install. NSX Edge
CLI admin password	<code>nsx_cli_passwd_0</code>	Required to install NSX Edge.
CLI audit password	<code>nsx_cli_audit_passwd_0</code>	Optional
CLI admin username	<code>nsx_cli_username</code>	Optional
CLI audit username	<code>nsx_cli_audit_username</code>	Optional
NSX Manager IP	<code>mpIp</code>	Required to join NSX Edge VM to NSX Manager.

Field Name	OVF Parameter	Field Type
NSX Manager token	mpToken	Required to join NSX Edge VM to NSX Manager. This token must be unique for every NSX Edge node. You must not use it to deploy any other NSX Edge node. To retrieve token, on the NSX Manager, run POST <code>https://<nsx-manager>/api/v1/aaa/registration-token</code> .
NSX Manager thumbprint	mpThumbprint	Required to join NSX Edge VM to NSX Manager. To retrieve thumbprint, on the NSX Manager node, run <code>get certificate api thumbprint</code> .
Node Id	mpNodeId	Only for internal use.
Hostname	nsx_hostname	Optional
Default IPv4 gateway	nsx_gateway_0	Optional
Management network IP address	nsx_ip_0	Optional
Management network netmask	nsx_netmask_0	Optional
DNS servers	nsx_dns1_0	Optional
Domain Search suffixes	nsx_domain_0	Optional
NTP Servers	nsx_ntp_0	Optional
Is SSH service enabled	nsx_isSSHEnabled	Optional
Is SSH enabled for root login	nsx_allowSSHRootLogin	Optional
Is autonomous Edge	is_autonomous_edge	Optional. Valid values: True, False (default)

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
--prop:mpIp=<NSXManager-IP>
--prop:mpToken=<NSXManager-Token>
--prop:mpThumbprint=<NSXManager-Thumbprint>
--prop:is_autonomous_edge=False
<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 VMware vCenter, 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
--prop:mpIp=<NSXManager-IP>
--prop:mpToken=<NSXManager-Token>
--prop:mpThumbprint=<NSXManager-Thumbprint>
--prop:is_autonomous_edge=False
<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 appliance.
Set the reservation to ensure that NSX Manager has sufficient memory to run efficiently. See [NSX Manager VM and Host Transport Node System Requirements](#).
- ◆ Power on the appliance. Open the console of the NSX Edge node to track the boot process.
- ◆ After the NSX Edge node starts, log in to the Edge node using the console or SSH (provided SSH is enabled at the time of install) with admin credentials.
- ◆ Run the `get interface eth0` (without VLAN) to verify that the IP address was applied as expected.

```

nsx-edge-1> get interface eth0

Interface: eth0
  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 nodes on non-NSX managed host, verify that the minimum MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

- ◆ Verify that the NSX Edge node has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge node and verify the following:

- You can ping your NSX Edge node management interface.
 - From the NSX Edge node, you can ping the node's default gateway.
 - From the NSX Edge node, you can ping the hypervisor hosts that are either in the same network or a network reachable through routing.
 - From the NSX Edge node, you can ping the DNS server and 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 node 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 assign management IP address to the correct NIC.

- Log in to the NSX Edge CLI and type the **stop service dataplane** command.
- (Static IP) Run the `set interface <interface-name> ip <x.x.x.x/24> gateway <x.x.x.x> plane mgmt` command.
- (DHCP) Run the `set interface interface-name dhcp plane mgmt` command.
- 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 node.

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 as a Virtual Appliance

You can manually install NSX Edge using an ISO file.

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

Prerequisites

- See NSX Edge network requirements in [NSX Edge Installation Requirements](#).

Procedure

- 1 Go to the [Broadcom Support](#) page. Select the **VMware Cloud Foundation** division on the top panel and go to the **My Downloads** panel.
- 2 Search VMware NSX and select the appropriate product version.
- 3 Locate and download the ISO file for NSX Edge.
- 4 In the vSphere Client, select the host datastore.
- 5 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.

- 6 In the vSphere Client inventory, select the host you uploaded the ISO file. or in the vSphere Client,
- 7 Right-click and select **New Virtual Machine** .
- 8 Select a compute resource for the NSX Edge appliance.
- 9 Select a datastore to store the NSX Edge appliance files.
- 10 Accept the default compatibility for your NSX Edge VM.
- 11 Select the supported ESXi operating systems for your NSX Edge VM.
- 12 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.

- 13 Power on the new NSX Edge VM.
- 14 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.

- 15 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 Manager VM and Host Transport Node System Requirements](#).

- 16 After the NSX Edge node starts, log in to the Edge node using the console or SSH (provided SSH is enabled at the time of install) with admin credentials.

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

- 17 There are three ways to configure a management interface.

Note If the server uses Mellanox NIC cards, do not configure the Edge in In-band 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
```

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

19 Run the `get interface eth0` (without VLAN) to verify that the IP address was applied as expected.

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

Interface: eth0
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 nodes on non-NSX managed host, verify that the minimum MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

20 (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.

21 Verify that the NSX Edge node has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge node and verify the following:

- You can ping your NSX Edge node management interface.
- From the NSX Edge node, you can ping the node's default gateway.
- From the NSX Edge node, you can ping the hypervisor hosts that are either in the same network or a network reachable through routing.
- From the NSX Edge node, you can ping the DNS server and NTP Server.

22 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 node 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 assign management IP address to the correct NIC.

- a Log in to the NSX Edge CLI and type the **stop service dataplane** command.
- b (Static IP) Run the `set interface <interface-name> ip <x.x.x.x/24> gateway <x.x.x.x> plane mgmt` command.
- c (DHCP) Run the `set interface interface-name dhcp plane mgmt` command.
- 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 node.

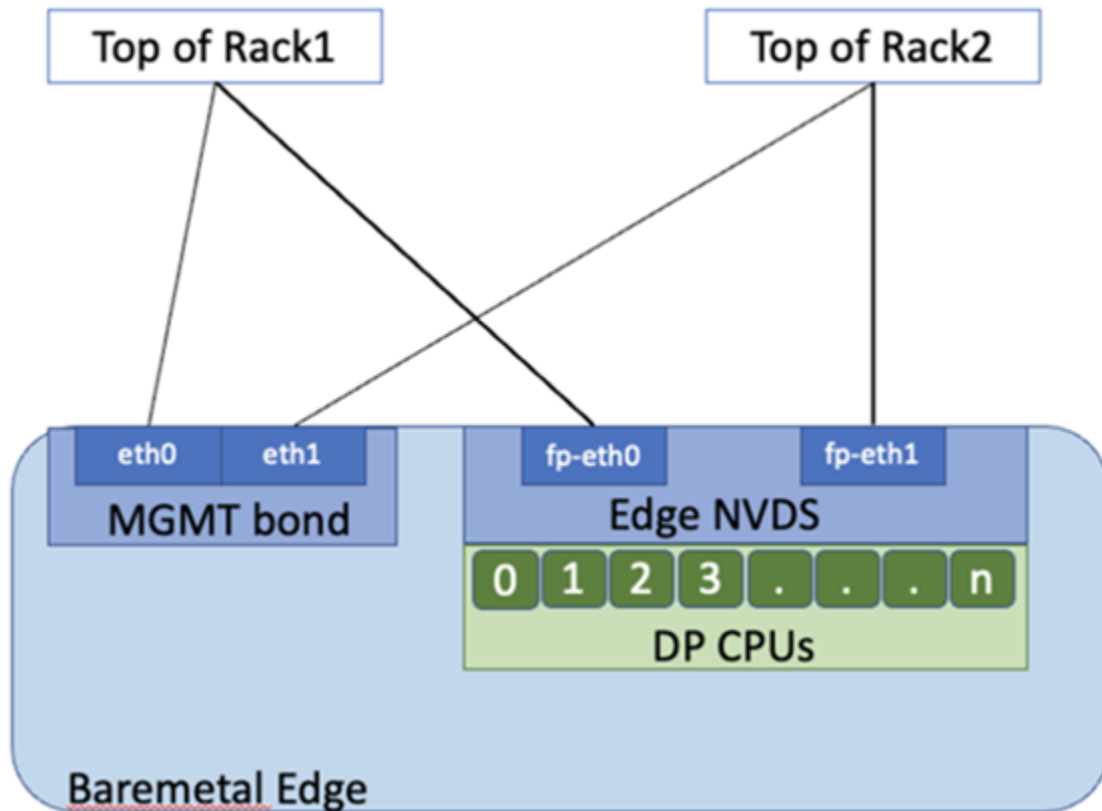
What to do next

If you did not join the NSX Edge with the management plane, see [Join NSX Edge with the Management Plane](#).

Install NSX Edge 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 on a bare metal server.

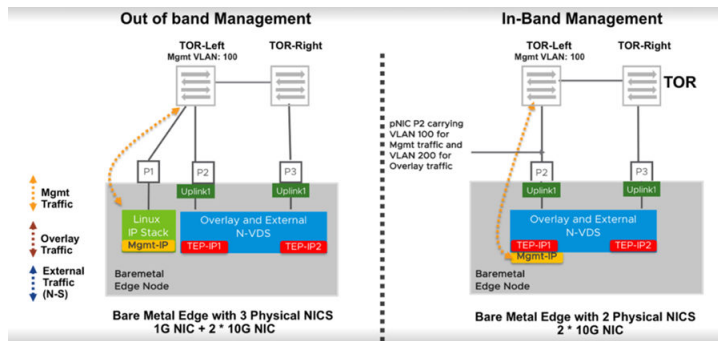
The NSX Edge bare metal node is a dedicated physical server that runs a special version NSX Edge software. The bare metal NSX Edge node requires a NIC supporting the Data Plane Development Kit (DPDK). VMware maintains a list of the compatibility with various vendor NICs. See the [Bare Metal Server System Requirements](#).



The NSX Edge nodes in the diagram are illustrated with a single N-VDS switch, configured with two datapath interfaces and two dedicated NICs for high availability of management plane.

Bare Metal NSX Edge nodes can be configured with more than 2 datapath interfaces depending on the number of NICs available on the server. Bare Metal NSX Edge nodes use pNICs as uplinks that directly connect to the top of rack switches. To provide high availability for bare metal edge management, configure two pNICs on the server as an active/standby Linux bond.

The CPUs on the edge are assigned as either datapath (DP) CPUs that provide routing and stateful services or services CPUs that provide load balancing and VPN services.



When a bare metal Edge node is installed, a dedicated interface is retained for management. This configuration is called out-of-band management. If redundancy is desired, two NICs can be used for management plane high availability. Bare metal Edge also supports in-band management where management traffic can leverage an interface being used for overlay or external (N-S) traffic as in the diagram.

For bare metal edge nodes, all the cores from the first node of a multi-NUMA node server will be assigned to NSX datapath. If the bare metal has only one NUMA node, then 50% of the cores will be assigned for the NSX datapath.

Bare Metal NSX Edge nodes support a maximum of two NUMA nodes.

Important VMware does not support the sub-NUMA clustering, which is a server hardware feature, on a Bare Metal NSX Edge is not supported as it may result in a shortage of heap memory.

The sub-NUMA clustering functionality changes the socket of heap memory from two NUMA domains to four NUMA domains. This change limits the size of heap memory allocated to each socket and causes a shortage of heap memory for socket 0 that the datapath requires. You must disabled sub-NUMA functionality in the BIOS. Any changes made to the BIOS will require a reboot.

To check if the sub-NUMA functionality is enabled, log in to the bare metal NSX Edge as root and run `lscpu`. The output is also captured in the support bundle. If there are more than two NUMA nodes, it implies that the sub-NUMA functionality is enabled and it must be disabled and must be disabled via the BIOS.

Note

- 1 When configuring LACP LAG bonds on Bare Metal NSX Edge nodes, datapath cores (backing NICs) should belong to same NUMA node for load balancing to occur on both devices. If devices forming the bond span multiple NUMA nodes, then bond only uses network device CPU that is local NUMA node (0) to transmit packets. So, not all devices do not get used for balancing traffic that is sent out of the bond device.

In this case, failover still works as 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.

Run `get dataplane` command to view the NUMA node associated with each datapath interface. To move nics associated with datapath to single NUMA node, physical reconfiguration of the server is required via the BIOS.

Prerequisites

- Disable sub-NUMA clustering by editing the BIOS settings. NSX does not support sub-NUMA clustering. For more details, refer to the KB article, <https://kb.vmware.com/s/article/91790>.

- Starting with NSX v3.1.3, on Bare Metal NSX Edge nodes, you do not need to disable hyperthreading. Hyperthreading is automatically disabled.

Prepare a PXE Server for Bare Metal NSX Edge Installation

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 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 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.
- Verify that the preseeded configuration file has the parameters `net.ifnames=0` and `biosdevname=0` set after `--` to persist after reboot.
- See [Bare Metal Server System Requirements](#).
- See NSX Edge network requirements in [NSX Edge Installation Requirements](#).

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 service dataplane
set interface eth0 <ip-cidr-format> plane mgmt
start service 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 and configure required settings to set up the PXE server. For more details, see Linux documentation.
- 4 Install the Apache server and TFTP and other components required to configure the PXE server.
- 5 Copy or download the NSX Edge installer ISO file to a temporary folder.
- 6 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/
```

- 7 (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:
$6$SUFQqs[...]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-rypted 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, add the following command.

```
usermod --password '\$6\$VS3exId0aKzW\$U3g0V7BF0DX1mRI.LR0v/Vg1oxVotEDp00b02hUF8u/'
root; \
usermod --password '\$6\$VS3exId0aKzW\$U3g0V7BF0DX1mRI.LR0v/Vg1oxVotEDp00b02hUF8u/'
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-rypted password 6tgm...`

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.

- 8 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=bionic netcfg/do_not_use_netplan=true --
```


- 9 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";
```

- 10 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 Bare Metal NSX Edge Automatically using ISO File](#).

Install Bare Metal NSX Edge Automatically using ISO File

Use the automated installation of NSX Edge on the bare metal servers because it automatically configures the NSX Edge node according to the hardware configuration and the most frequently used NSX Edge configuration. After the installation is complete, configure the management interface of the nodes as out-of-band or as in-band and provide the final IP configuration of the management interface statically or over DHCP.

Prerequisites

- Both BIOS and UEFI boot modes are supported.
- See NSX Edge network requirements in [NSX Edge Installation Requirements](#).
- See [Bare Metal Server System Requirements](#).
- If the bare metal servers for NSX Edge nodes have integrated SD card devices, disable those SD cards. For more information, see <https://kb.vmware.com/s/article/67363>.
- To prepare a bare metal server as an NSX Edge node, ensure there is a minimum of 200 Gb disk space.

For more details about the supported NICs, see [VMware Compatibility Guide](#).

Procedure

- 1 Go to the [Broadcom Support](#) page. Select the **VMware Cloud Foundation** division on the top panel and go to the **My Downloads** panel.

- 2 Search VMware NSX and select the appropriate product version.
- 3 Locate and download the ISO file for NSX Edge for Bare Metal.
- 4 Log in to the Integrated Lights-Out (ILO) interface of the bare metal server.
- 5 Click **Launch** in the virtual console preview.
- 6 Select **Virtual Media > Connect Virtual Media**.

Wait a few seconds for the virtual media to connect.

- 7 Select **Virtual Media > Map CD/DVD** and browse to the ISO file.
- 8 Select **Next Boot > Virtual CD/DVD/ISO**.
- 9 Select **Power > Reset System (warm boot)**.

The installation duration depends on the bare metal environment.

- 10 Choose **Automated installation**.

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

- 11 Select the applicable primary network interface. This is for the management network interface.

During power-on, the installer requests a network configuration. Select static IP settings. If the static IP is not available, use DHCP.

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

Note During deployment, the installer automatically selects the largest disk to install NSX on the NSX Edge node.

- 12 Power on the appliance. Open the console of the NSX Edge node to track the boot process. If the console window does not open, make sure that pop-ups are allowed.
- 13 After the NSX Edge node starts, log in to the Edge node using the console or SSH (provided SSH is enabled at the time of install) with admin credentials.

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

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

- 15 There are three ways to configure a management interface.

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

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

```
(DHCP) set interface eth0 dhcp plane mgmt
```

- Tagged interface.

```
set interface eth0 vlan <vlan_ID> plane mgmt
```

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

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

- In-band interface.

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

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

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

- Tagged interface and In-band interface.

Any existing VLAN management interface must be cleared before creating a new one.

```
clear interface eth0.<vlan_ID>
```

- (Optional) Create a **bond0** interface for management HA interface with multiple interfaces.

You can configure a bond management interface on an NSX Edge using the following CLI command. Use console to clear existing management IP before you create a bond and add an interface to it.

Note Only active-backup mode is allowed on a bond interface. You can configure VLANs on a bond0 interface.

Create a bond interface

```
set interface bond0 ip x.x.x.x/mask gateway x.x.x.x plane mgmt mode active-backup members eth0,eth1 primary eth0
```

Create vlan interface on bond0

```
set interface bond0 vlan Y plane mgmt
```

Assign IP address to bond0.yyy

```
set interface bond0.yyy ip x.x.x.x/24 gateway z.z.z.z plane mgmt
```

- 16 Run the `get interface eth0` (without VLAN) to verify that the IP address was applied as expected.

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

```
Interface: eth0
```

```
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 nodes on non-NSX managed host, verify that the minimum MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

17 Set physical NICs to be used by NSX dataplane from the list of available PCI devices.

- a `get dataplane device list`
- b `reset dataplane device list`
- c `restart service dataplane`
- d `get physical-port`

After selecting physical NICs, restart NSX dataplane services for changes to take effect.

Note Starting in NSX 3.1.2, you can claim up to 16 physical NICs.

Note To configure custom NICs for dataplane, run the `set dataplane device list <NIC1>, <NIC2>, <NIC3>` command.

18 To avoid network configuration errors, verify that the physical NICs selected match the NICs configured in the uplink profiles.

19 Verify that the NSX Edge node has the required connectivity.

If you enabled SSH, make sure that you can SSH to your NSX Edge node and verify the following:

- You can ping your NSX Edge node management interface.
- From the NSX Edge node, you can ping the node's default gateway.
- From the NSX Edge node, you can ping the hypervisor hosts that are either in the same network or a network reachable through routing.
- From the NSX Edge node, you can ping the DNS server and NTP Server.

20 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 node 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 assign management IP address to the correct NIC.

- a Log in to the NSX Edge CLI and type the **stop service dataplane** command.
- b (Static IP) Run the `set interface <interface-name> ip <x.x.x.x/24> gateway <x.x.x.x> plane mgmt` command.
- c (DHCP) Run the `set interface interface-name dhcp plane mgmt` command.
- 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 node.

Install Bare Metal NSX Edge Interactively using ISO File

Install NSX Edge devices on bare metal using an ISO file in the interactive mode.

Prerequisites

- Starting in NSX 3.2 both UEFI or Legacy BIOS modes are supported.
- See NSX Edge network requirements in [NSX Edge Installation Requirements](#).

Procedure

- 1 Go to the [Broadcom Support](#) page. Select the **VMware Cloud Foundation** division on the top panel and go to the **My Downloads** panel.
 - 2 Search VMware NSX and select the appropriate product version.
 - 3 Locate and download the ISO file for NSX Edge for Bare Metal.
 - 4 Log in to the ILO of the bare metal.
 - 5 Click **Launch** in the virtual console preview.
 - 6 Select **Virtual Media > Connect Virtual Media**.
- Wait a few seconds for the virtual media to connect.
- 7 Select **Virtual Media > Map CD/DVD** and browse to the ISO file.
 - 8 Select **Next Boot > Virtual CD/DVD/ISO**.
 - 9 Select **Power > Reset System (warm boot)**.

The installation duration depends on the bare metal environment.

- 10 Choose **Interactive Install**.

This operation takes several seconds to start.

- 11 In the Configure the keyboard window, select **Yes** if the installer must auto-detect the keyboard or select **No** if the keyboard must not be detected by the console.
- 12 Select English US as the language.
- 13 In the Configure the network window, select the applicable primary network interface.
- 14 Enter the host name that connects to the selected primary interface and click **Ok**.

During power-on, the installer requests a network configuration mode. Select static IP address and provide one. If static IP address is not available, select DHCP.

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

- 15 In the Configure NSX appliance using kickstart window:
 - Enter the URL of the NSX kickstart config file if you want to automate NSX configuration on the bare metal server.
 - Leave the field blank if you want to manually configure NSX on the bare metal server.
- 16 Select a disk size that meets the system requirements.
- 17 In the Partition disks window, choose one of the following options:
 - Select **Yes** if you want to unmount existing partitions so that new partitions can be created on disks.
 - Select **No** if you want to use existing partitions.
- 18 After the NSX Edge node starts, log in to the Edge node using the console or SSH (provided SSH is enabled at the time of install) with admin credentials.

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

- 19 Run the `get interface eth0` (without VLAN) to verify that the IP address was applied as expected.

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

Interface: eth0
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 nodes on non-NSX managed host, verify that the minimum MTU setting is set to 1600 (instead of 1500) on the physical host switch for the data NIC.

20 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 node 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 assign management IP address to the correct NIC.

- a Log in to the NSX Edge CLI and type the **stop service dataplane** command.
- b (Static IP) Run the `set interface <interface-name> ip <x.x.x.x/24> gateway <x.x.x.x> plane mgmt` command.
- c (DHCP) Run the `set interface interface-name dhcp plane mgmt` command.
- 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 node.

Intel QAT Support for IPsec VPN Bulk Cryptography

Beginning with the NSX 3.0 release, support for the Intel QuickAssist Technology (QAT) is provided on bare metal servers. Intel QAT provides the hardware acceleration for various cryptography operations.

The QAT feature is enabled by default if the NSX Edge is deployed on a bare metal server with an Intel QuickAssist PCIe card that is based on the installed C62x chipset (Intel QuickAssist Adapter 8960 or 8970). The single root I/O virtualization (SR-IOV) interface must be enabled in the BIOS firmware.

To check the status of the QAT feature, enter the following command on the NSX Edge bare metal server CLI.

```
get dataplane qat
```

The possible responses you might receive are listed in the following table.

Status of QAT Feature	Definition
QAT present, enabled, running	The QAT feature is enabled and running.
QAT present, enabled, not running	The QAT feature has been enabled, but the dataplane service must be restarted for the status change to take effect.
QAT present, disabled, not running	The QAT feature is disabled.
QAT present, disabled, running	The QAT feature has been disabled, but the dataplane service must be restarted for the status change to take effect.

Status of QAT Feature	Definition
QAT not present	The bare metal server on which you ran the CLI command does not have a QAT device installed.
QAT not supported in VM	You ran the CLI command on a VM edge.

To disable or enable the use of an installed QAT device, use the following CLI commands. The expected responses are also shown.

```
set dataplane qat disabled
QAT disabled. Please restart service dataplane to take effect.
```

```
set dataplane qat enabled
QAT enabled. Please restart service dataplane to take effect.
```

Important You must enter the `restart service dataplane` command at the CLI prompt for the QAT feature status change to take effect.

Join NSX Edge with the Management Plane

To establish communication between NSX Edges and NSX Manager or NSX Manager cluster, join NSX Edges with NSX Manager. You only need to register NSX Edges with one NSX Manager to ensure communication with the management plane.

Prerequisites

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

Procedure

- 1 Open an SSH session or console session to one of the NSX Manager appliances.
- 2 Open an SSH session or console session to the NSX Edge node VM.
- 3 To retrieve the thumbprint of the NSX Manager appliance, at the NSX Manager appliance console, 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
659442c1435350edbbc0e87ed5a6980d892b9118f851c17a13ec76a8b985f57
```

- 4 Alternatively, to retrieve the thumbprint of the cluster, at the NSX Manager appliance console, run `get certificate cluster thumbprint`.

The output is a string of alphanumeric numbers that is unique to the cluster.

- 5 To join the NSX Edge node (VM or Bare Metal) to the NSX Manager appliance, run the **join management-plane** command.

Provide the following information:

- Hostname or IP address of the NSX Manager with an optional port number
- User name of the NSX Manager
- Certificate thumbprint of the NSX Manager
- Password of the NSX Manager

```
NSX-Edge1> join management-plane <Manager-IP> thumbprint <Manager-thumbprint> username
admin
```

Repeat this command on each NSX Edge node VM.

- 6 If cluster VIP is configured for the NSX Manager, then join NSX Edge Node using the cluster thumbprint. Run the **join management-plane** command.

Provide the following information:

- Virtual IP address of the NSX Manager cluster with an optional port number
- User name of the NSX Manager
- Certificate thumbprint of the NSX Manager cluster
- Password of the NSX Manager

```
NSX-Edge1> join management-plane <Cluster-VIP> username <Manager-username> password
<Manager-password> thumbprint <Cluster-tumbprint>
```

- 7 Verify the result by running the `get managers` command on your NSX Edge node VMs.

```
nsx-edge-1> get managers
- 10.173.161.17 Connected (NSX-RPC)
- 10.173.161.140 Connected (NSX-RPC)
- 10.173.160.204 Connected (NSX-RPC)
```

- 8 In the NSX Manager UI, navigate to **System > Fabric > Nodes > Edge Transport Nodes**.

On the NSX Edge Transport Node page:

- The **Configuration State** column displays `Configure NSX`. Click `Configure NSX` to begin configuration on the node. If the **NSX Version** column does not display the version number installed on the node, try refreshing the browser window.
- Before you configure NSX on the NSX Edge node, the **Node Status** and **Tunnel Status** columns display state `Not Available`. The **Transport Zones** and **N-VDS** switches columns display value 0, indicating there are no transport zones attached or N-VDS switches configured on the NSX Edge node.

What to do next

When installing NSX Edge using NSX Manager see [Create an NSX Edge Transport Node](#).

When installing NSX Edge manually, see [Edit NSX Edge Transport Node Configuration](#).

Edit NSX Edge Transport Node Configuration

After manually installing NSX Edge VM on an ESXi host or as a Bare Metal server, you can edit a NSX Edge configuration.

A transport node is a node that is capable of participating in an NSX overlay or NSX 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.

Prerequisites

- VLAN and Overlay transport zones must be configured.
- Verify that compute manager is configured. See [Add a Compute Manager](#).
- An IP pool (to be used as NSX Edge TEP pool) must be configured or must be available in the network deployment.
- (NSX 4.0.1.1) Before you can use NSX Edge VM datapath interfaces in Uniform Passthrough (UPT) mode, meet the following conditions:

Note UPT mode is not supported on NSX Edge Bare Metal hosts.

- NSX Edge hardware version is 20 (vmx-20) or later. Previous NSX Edge hardware versions do not support UPT mode.
- Verify that the memory reservation on the configured NSX Edge is set to 100%.
- From the vSphere Web Client, enable UPT on the NSX Edge VM network adapter. See the *Change the Virtual Machine Network Adapter Configuration* topic in *vSphere Virtual Machine Administration* guide.
- At least one of the NSX Edge VM datapath interface must be backed by an ESXi host that hosts a Data Processing Unit-based SmartNIC. A SmartNIC is a NIC card that provides network traffic processing using a Data Processing Unit (DPU), a programmable processor on the NIC card, in addition to the traditional functions of a NIC card. For more information related to DPU, see [NSX on vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine](#).

- Starting with NSX 4.0.1.1, NSX Edge VM hardware version will no longer default to `virtualHW.version 13`. NSX Edge VM hardware will depend on the underlying version of the ESXi host. VM hardware versions compatible with ESXi hosts are listed in KB article [2007240](#).

Procedure

- From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- (NSX 4.0.1.1) To enable UPT mode on the NSX Edge node:
 - Select **System** → **Fabric** → **Nodes** → **Edge Transport Nodes**.
 - Select the NSX Edge node to enable UPT, click **Actions** and **Change Node Settings**.
 - In the **Change Node Settings** window, ensure the **Enable UPT mode for datapath interface** field is enabled. This setting enables UPT on all datapath interfaces that support UPT mode or support network offloads.
 - Click **Save**.
- To prepare the NSX Edge node as a transport node, select **System** > **Fabric** > **Nodes** > **Edge Transport Nodes** > **Edit Edge**. Configure the following fields to complete preparation of a NSX Edge node as a transport node.
- Enter the N-VDS information.

Consider these points before you configure vNICs of NSX Edge nodes:

An N-VDS switch is hosted inside the Edge node VM with four fast path vNICs and one management vNIC.

- One vNIC is dedicated to management traffic.
- One vNIC is dedicated to overlay traffic (fp-eth0 DPDK fastpath interface).
- Two vNICs are dedicated to external traffic (fp-eth1, fp-eth2 DPDK fastpath interfaces).

Option	Description
Edge Switch Name	Enter a name for the switch or keep the default name.
Transport Zone	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 connectivity and a VLAN for uplink connectivity. Note NSX Edge nodes support multiple overlay tunnels (multi-TEP) when the following prerequisites are met: <ul style="list-style-type: none"> TEP configuration must be done on one N-VDS only. All TEPs must use the same transport VLAN for overlay traffic. All TEP IPs must be in the same subnet and use the same default gateway.

Option	Description
Uplink Profile	<p>Select the uplink profile from the drop-down menu. The available uplinks depend on the configuration in the selected uplink profile.</p> <p>Note NSX Edge nodes support uplink profiles with Failover teaming policy (with single active uplink and no standby) and Loadbalancer Source teaming policy (with multiple active uplinks) only.</p>
IP Address Type (TEP)	<p>Select the IP version to be used for the tunnel endpoint (TEP). The options are IPv4 and IPv6.</p> <p>Important Ensure that the transport node forwarding mode and TEP IP address type are the same. For example, if the transport node forwarding mode is set to IPv6, set the TEP IP address type to IPv6. If they are different, a loss of traffic may result.</p>
IPv4 Assignment (TEP)	<p>This field appears when IP Address Type (TEP) is set to IPv4.</p> <p>Choose how IPv4 addresses are assigned to the NSX Edge switch that is configured. It is used as the tunnel endpoint of the NSX Edge. The options are:</p> <ul style="list-style-type: none"> ■ Use IP Pool: Select the IPv4 pool. ■ Use Static IPv4 List: Specify the following fields: <ul style="list-style-type: none"> ■ Static IP List: Enter a list of comma-separated IPv4 addresses to be used by the NSX Edge. ■ IPv4 Gateway: Enter the default gateway of the TEP, which is used to route packets another TEP in another network. For example, ESXi TEP is in 20.20.20.0/24 and NSX Edge TEPs are in 10.10.10.0/24 then we use the default gateway to route packets between these networks. ■ IPv4 Subnet Mask: Enter the subnet mask of the TEP network used on the NSX Edge.

Option	Description
IPv6 Assignment (TEP)	<p>This field appears when IP Address Type (TEP) is set to IPv6.</p> <p>Choose how IPv6 addresses are assigned to the NSX Edge switch that is configured. It is used as the tunnel endpoint of the NSX Edge. The options are:</p> <ul style="list-style-type: none"> ■ Use IP Pool: Select the IPv4 pool. ■ Use Static IPv6 List: Specify the following fields: <ul style="list-style-type: none"> ■ Static IP List: Enter a list of comma-separated IPv4 addresses to be used by the NSX Edge. ■ IPv6 Gateway: Enter the default gateway of the TEP, which is used to route packets another TEP in another network. ■ IPv6 Subnet Mask: Enter the subnet mask of the TEP network used on the NSX Edge.
DPDK Fastpath Interfaces / Virtual NICs	<p>Map uplinks to DPDK fastpath interfaces.</p> <p>Starting with NSX release 2.5, single N-VDS deployment mode is recommended for both bare metal and NSX Edge VM. See Configure NSX Edge DPDK Interfaces.</p> <p>Starting with NSX 4.0.1, you can map uplinks to DPDK fastpath interfaces that are backed by smartNIC-enabled DVPGs, VLAN logical switches or segments. The prerequisite is to enable UPT mode on NSX Edge VM virtual network adapters. The UPT mode requires at least one DPDK interface to be backed by smartNIC-enabled hardware also known as Data Processing Unit (DPU)-backed networks.</p> <p>Note If the uplink profile applied to the NSX Edge node is using a Named Teaming policy, ensure the following condition is met:</p> <ul style="list-style-type: none"> ■ All uplinks in the Default Teaming policy must be mapped to the corresponding physical network interfaces on the Edge VM for traffic to flow through a logical switch that uses the Named Teaming policies. See Configure Named Teaming Policy.
	<p>You can configure a maximum of four unique data path interfaces as uplinks on a NSX Edge VM.</p> <p>When mapping uplinks to DPDK Fastpath Interfaces, if NSX Edge does not display all the available interfaces (four in total), it means that either the additional interface is not yet added to the NSX Edge VM or the uplink profile has fewer number of uplinks.</p> <p>For NSX Edge VMs upgraded from an earlier version of NSX to 3.2.1 or later, invoke the redeploy API call to redeploy the NSX Edge VM. Invoking the redeploy API ensures the NSX Edge VM deployed recognizes all the available datapath interfaces in NSX Manager UI. Make sure the Uplink profile is correctly configured to use additional datapath NIC.</p> <p>For more information on configuring NSX Edge DPDK fastpath interfaces, see Configure NSX Edge DPDK Interfaces.</p> <ul style="list-style-type: none"> ■ For autodeployed NSX Edges (edge nodes deployed from the NSX Manager UI or API), call the redeploy API. The following API is deprecated. <pre style="background-color: #f0f0f0; padding: 10px;">POST api/v1/transport-nodes/<transport-node-id>? action=redeploy</pre>

Option	Description
	<ul style="list-style-type: none"> ■ For manually deployed edges (edges deployed using OVA/OVF file from the VMware vCenter UI or API), deploy a new NSX Edge VM. Ensure all the vmx customizations of the old NSX Edge VM are also done for the new NSX Edge VM. <p>Performing vMotion on an NSX Edge VM can result in ESXi running out of resources from a shared buffer pool if you create large VMs with multiple vNICs that use large sized ring buffers. To increase the depth of the shared buffer, modify the <code>ShareCOSBufSize</code> parameter in ESXi. To configure buffer size, see https://kb.vmware.com/s/article/76387.</p>

Note

- LLDP profile is not supported on an NSX Edge VM appliance.
- Uplink interfaces are displayed as **DPDK Fastpath Interfaces** if the NSX Edge is installed using NSX Manager or on a Bare Metal server.
- Uplink interfaces are displayed as **Virtual NICs** if the NSX Edge is installed manually using vCenter Server.

5 Click **Save**.

6 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.

Note (NSX 4.0.1.1) When you enable the **Actions > Change Node Settings > Enable UPT mode for datapath interface** field, the NSX Manager puts the NSX Edge VM into maintenance mode, applies configuration, and removes NSX Edge from maintenance mode which makes the UPT configuration effective on the NSX Edge transport node.

7 To successfully configure firewall rules on the NSX Edge node, enable service core on the transport node.

```
set debug
set dataplane service-core enabled
restart service dataplane
```

8 (Optional) View the transport node with the `GET https://<nsx-manager>/api/v1/transport-nodes/<transport-node-id>` API call.

9 (Optional) For status information, use the `GET https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-id>/status` API call.

- 10 After an NSX Edge node is migrated to a new host using vCenter Server, you might find NSX Manager UI reporting stale configuration details (Compute, Datastore, Network, SSH, NTP, DNS, Search Domains) of the NSX Edge. To get the latest configuration details of NSX Edge on the new host, run the API command.

```
POST api/v1/transport-nodes/<transport-node-id>?
action=refresh_node_configuration&resource_type=EdgeNode
```

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.

Note Multiple NSX Edge clusters can be deployed within a single NSX Manager, allowing for the creation of pool of capacity that can be dedicated to specific services (for example, NAT at Tier-0 gateways or NAT at Tier-1 gateways). Within a single NSX Edge cluster, all NSX Edge nodes must be the same type – either physical servers (bare metal) or VMs. However, you can have NSX Edge node VMs of different sizes within the same NSX Edge cluster.

Prerequisites

- Ensure there is at least one NSX Edge node with Node status `Up` and it must not be part of any existing cluster.
- 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>` or `https://<nsx-manager-fqdn>`.
- 2 Select **System > Fabric > Nodes > Edge Clusters > Add Edge Clusters**.
- 3 Enter the NSX Edge cluster name.
- 4 Select an NSX Edge cluster profile from the drop-down menu.

- 5 In Member Type drop-down menu, select either **Edge Node** if the virtual machine is deployed on-premises or **Public Cloud Gateway** if the virtual machine is deployed in a public cloud.
- 6 From the **Available** column, select NSX Edges and click the right-arrow to move them to the **Selected** column.
- 7 Click **Add**.

What to do next

To know the current status of the NSX Edge nodes within a cluster, run the `get edge-cluster status` CLI command.

Understand these states of NSX Edge nodes:

- `Up (routing down)`: The NSX Edge node state is `Up` but Tier-0 SR routing daemon is not running on this NSX Edge node as no services are enabled on the it.
- `Admin Down`: The NSX Edge is in NSX maintenance mode and all services or traffic forwarding is disabled on this edge.
- `Down`: The datapath process is not running, link is down or VTEP tunnels are down.
- `Unreachable`: Between two NSX Edge nodes, one BFD session is run on the management interface, and at least one BFD session is run on each of the VTEP interfaces. An NSX Edge considers its peer as unreachable only when all BFD sessions to that NSX Edge (management one and all VTEP ones) are down.

You can now build logical network topologies and configure services. See the *NSX Administration Guide*.

Remove NSX Edge Nodes from an Edge Cluster

Remove NSX Edge nodes with Tier-1 Gateways or Tier-0 Gateway configured with service router (SR), DHCP and metadata proxy.

Before you remove NSX Edge nodes, relocate the gateway configurations to a new standby node:

- To remove Tier-0 gateway configurations on an NSX Edge node, you must manually relocate Tier-0 configurations such as Tier-0 SR, DHCP and metadata proxy configurations to a standby NSX Edge node.
- To remove Tier-1 gateway configurations on an NSX Edge node, do the following in these scenarios:
 - If Tier-1 SR, DHCP and metadata proxy configurations are auto allocated to the NSX Edge node, you can enable the standby relocation functionality to relocate Tier-1 configurations to a new standby NSX Edge node. The procedure describes how to use the standby relocation functionality to relocate the configurations to a new standby node.

- If Tier-1 SR, DHCP and metadata proxy configurations are manually allocated to the NSX Edge node, you need to manually relocate Tier-1 configurations to a new standby NSX Edge node.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 On NSX Edge nodes configured with Tier-0 configurations, manually move configurations to some other NSX Edgenode.

Note The standby relocation functionality does not relocate Tier-0 configurations such as Tier-0 SRs, DHCP and metadata proxy configurations to a standby NSX Edge node.

- a Select **Networking** → **Tier-0 Gateways**.
- b To edit a Tier-0 Gateway, select the gateway, click vertical ellipses and click **Edit**.
- c Navigate to **Interfaces** section and click the **External and Service Interfaces**.
- d In the **Set Interfaces** window, edit the interface configured for NSX Edge node.
- e Remove the existing NSX Edge node associated with the interface and select a new NSX Edge node that is configured with the same VLAN connectivity required for the interface and click **Save**.

The above procedure deletes Tier-0 SRs, DHCP and metadata proxy on Tier-0 Gateways of the existing NSX Edge and moves them to the new NSX Edge node.

- f Delete the NSX Edge node from the edge cluster.
- 3 On NSX Edge nodes that are auto allocated with Tier-1 SR, DHCP and metadata proxy configurations, follow these steps to trigger the standby relocation functionality to relocate those configurations and remove the NSX Edge node from the NSX Edge cluster:
 - a (Optional) For faster failover, change the BFD timers for NSX Edge cluster by setting it to **500 ms**.
 - b Apply the new NSX Edge cluster profile to the transport node profile. It ensures faster failover when the NSX Edge node is powered off.
 - c Select **Networking** → **Tier-1 Gateways**.
 - d To edit a Tier-1 Gateway, select the gateway, click vertical ellipses and click **Edit**.
 - e In the Edit view, select the NSX Edge cluster and enable the **Enable Standby Relocation** field.

Important For standby relocation to function successfully, there must be an additional healthy NSX Edge node in the edge cluster. During the process of removing an NSX Edge node, Tier-1, DHCP or metadata proxy configurations are relocated from an existing NSX Edge to a new standby node.

- f Select **System** → **Fabric** → **Profiles** → **Edge Cluster Profiles**.
- g Select the edge cluster profile and click **Edit**.
- h Set **Standby Relocation Threshold (mins)** that is applied to the edge cluster. The default recommended value is 30 mins and the minimum value is 10 mins.

Note Only auto allocated Tier-1 SR, DHCP and metadata proxy configurations are relocated to a standby NSX Edge. If the NSX Edge node to be removed contains any manually allocated configurations, such configurations will not be relocated out from existing NSX Edge node to a standby node. You need to manually change the allocation for those Tier-1 configurations.

- i Power off the NSX Edge without taking down the node into maintenance mode. If the NSX Edge is running any active service, then all such active service configurations will failover to another NSX Edge because of the HA failover trigger when the node is powered off.
- j Wait for the duration of standby relocation threshold timeout. After the threshold limit is reached, all Tier-1 service configurations that have standby relocation enabled will be removed from the edge node being powered-off and relocated to some other NSX Edge in the cluster. There can be minor delays in standby relocation to perform the relocation task.
- k After Tier-1 configurations are relocated to a standby node, remove the NSX Edge that was powered off from the cluster. Select the Edge Cluster, click **Edit** and remove the NSX Edge node from the cluster and click **Save**.

Relocate and Remove an NSX Edge Node from an NSX Edge Cluster

Starting with NSX 4.0.1.1, you can use the NSX relocate and remove API to relocate the service configurations of an NSX Edge node to another standby NSX Edge node in the same NSX Edge cluster and then remove the Edge node from the Edge cluster.

The relocate and remove API relocates the following service configurations:

- Logical routers
- DHCP server
- Metadata proxy
- L2 forwarder

Prerequisites

To relocate and remove an Edge node from an Edge cluster, the following conditions are required:

- The Edge node must not have any manually allocated service configurations. Only auto allocated service configurations can be relocated.
- To be available for relocation, standby Edge nodes must not be configured with Layer 2 bridging.
- The Edge cluster must have at least two healthy Edge nodes where the auto allocated service configurations can be relocated to.
- For HA (high availability), the Edge cluster must have more than two Edge nodes that are possible for relocation.

Procedure

- 1 Run the API command to get the `member_index` value of the Edge node that you want to relocate and remove from an Edge cluster:

```
GET https://<nsx-manager-IP>/policy/api/v1/edge-clusters/<edge-cluster-id>

{
  "deployment_type": "VIRTUAL_MACHINE",
  "members": [
    {
      "member_index": 11,
      "transport_node_id": "21a19cbf-eaba-4a59-b18d-ff71fe5d76aa",
      "display_name": "edgeVm1New"
    },
    {
      "member_index": 13,
      "transport_node_id": "740cf97d-892b-47bb-97e7-889d92252e80",
      "display_name": "edgeVm2New"
    },
    {
      "member_index": 14,
      "transport_node_id": "cd5ab447-a36a-4bc3-94ff-0a4eea9fb2ad",
      "display_name": "edgeVm3New"
    }
  ],
}
```

The `member_index` value is used to specify the Edge node to relocate and remove. Assume that you want to relocate the service configurations for the Edge node named `edgeVm1New`, then its `member_index` value is 11.

- 2 Enter the relocate and remove API command and the `member_value` value of the Edge node to relocate and remove:

```
POST https://<nsx-manager-IP>/api/v1/edge-clusters/<edge-cluster-id>?action=relocate_remove

{
  "member_index": 11
}
```

- 3 Run the API command.

The Edge node enters maintenance mode and its service configurations are transferred to one of the standby Edge nodes in the cluster. After the service configurations are transferred, the Edge node is removed from the Edge cluster and exits maintenance mode.

Note The API command will not work if:

- The Edge node has any manually allocated service configurations.
- The Edge cluster does not have at least two healthy standby Edge nodes.

Caution It is possible for the API command to give a success response, but in the background, the relocation operation fails. If this scenario occurs, then an alarm with the **Event Type** of `Edge Cluster Member Relocate Failure` is raised.

The recommended action for this scenario is to review the available capacity of the Edge cluster. If more capacity is required, scale your Edge cluster and then retry the API command.

NSX IPv6 Configuration

11

Configure IPv6 communication for NSX.

Read the following topics next:

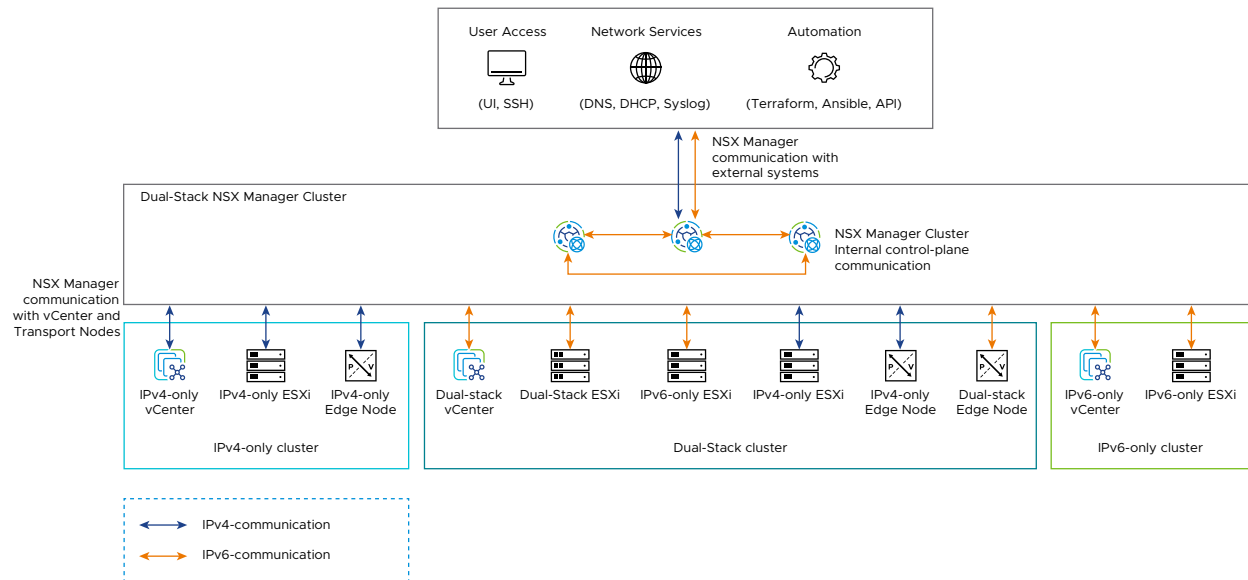
- [IPv6 Support in the NSX Platform Infrastructure](#)
- [IPv6 Limitations Between Management Plane/Control Plane and Transport Nodes](#)
- [Supported Topologies for IPv6](#)
- [IPv6 Configuration Workflow for NSX Manager and Transport Node Communication](#)
- [IPv6 Tunnel Endpoint Deployment for ESXi Host and NSX Edge on New Transport Zones](#)
- [IPv6 Tunnel Endpoint Configuration for ESXi Host and NSX Edge on Existing Transport Zones](#)
- [Troubleshooting for IPv6 Configuration](#)

IPv6 Support in the NSX Platform Infrastructure

NSX offers IPv6 addressing capability in the platform infrastructure for the management plane and control plane communication between NSX Managers and transport nodes.

The following communication and interface types provide support for IPv6:

- Communication between NSX Manager and external systems.
- Communication between NSX Manager and VMware vCenter.
- Communication within the NSX Manager cluster internal communication.
- Communication between NSX Manager and transport nodes in the control plane.



NSX Manager and External Systems Communication

The NSX Manager management interface can be configured with IPv4 only or with dual stack (both IPv4 and IPv6) addresses. The address type (IPv4 or IPv6) that is used to establish the connection is always used for the duration of the communication. For example:

- Users connecting to the NSX Manager UI using IPv6 or a FQDN which is resolved to an IPv6 address by a DNS server.
- API calls from users or automation tools using IPv6 or a FQDN which is resolved to an IPv6 address by a DNS server.
- Communication with VMware vCenter configured using IPv6 or a FQDN which is resolved to an IPv6 address by a DNS server.

For details on NSX Manager IPv6 configuration, see [IPv6 Configuration Workflow for NSX Manager and Transport Node Communication](#).

The NSX Manager cluster virtual IP address (VIP) can be configured as IPv4 only, IPv6 only, or dual stack. For details, see [Configure a Virtual IP Address for a Cluster](#).

NSX Manager and VMware vCenter Communication

Communication with VMware vCenter can be configured using IPv6 or a FQDN which is resolved to an IPv6 address by a DNS server.

For details on NSX Manager IPv6 configuration, see [Add a Compute Manager](#).

NSX Manager and Transport Node Communication

These communications refer to the following:

- NSX Manager control/management internal communication.

- NSX Edge transport node control/management internal communication.
- ESX control/management internal communication.

IPv6 Tunnel Endpoint (TEP) Support

IPv6 TEP with Geneve encapsulation is supported for the transport node types of NSX Edge and ESXi hosts. This capability allows overlay transport zones to use IPv6 as the underlay transport protocol.

IPv6 Limitations Between Management Plane/Control Plane and Transport Nodes

There are some limitations for the IPv6 communication between the management plane/control plane and transport nodes.

The following limitations in IPv6 support are:

- IPv6 only for NSX Manager and NSX Edge is not supported.
- NSX Global Managers (NSX Federation) are not supported for IPv6 in the management plane and control plane interfaces.
- IPv6 TEP with Geneve encapsulation is supported for the transport node types of NSX Edge and ESXi hosts only.
- For the IPv6 communication between the management plane/control plane and transport nodes, after transport nodes are registered to NSX Manager, the IP version configured for the transport node cannot be switched.

For example, if a transport node is configured as IPv4 and added to NSX Manager, you cannot switch its IP version to IPv6 only or vice versa.

- IPv6 communication is prioritized between dual stack NSX Managers and dual stack transport nodes. IPv6 communication to IPv4 communication fallback mechanism is not supported.
- Autonomous Edge with IPv6 APIs is not supported.
- Windows, Linux physical servers do not support IPv6 communication with NSX Manager.

Supported Topologies for IPv6

The following table provides the supported topologies for IP version configuration.

Topology #	NSX Manager			NSX Edge Node - Management Interface			VMware vCenter			ESX - Management Interface			Supported Topology?
	IPv4 only	Dual stack ^	IPv6 only*	IPv4 only	Dual stack ^	IPv6 only*	IPv4 only	Dual stack ^	IPv6 only	IPv4 only	Dual stack ^	IPv6 only	
1	✓	—	✗	✓	—	✗	✓	—	—	✓	✓	—	Supported
2	✓	—	✗	✓	—	✗	—	—	✓	—	✓	✓	Not supported
3	✓	—	✗	✓	—	✗	—	✓	—	✓	✓	—	Supported
4	—	✓	✗	✓	—	✗	✓	—	—	✓	✓	—	Supported
5	—	✓	✗	✓	—	✗	—	—	✓	—	✓	✓	Supported
6	—	✓	✗	✓	—	✗	—	✓	—	✓	✓	✓	Supported
7	✓	—	✗	—	✓	✗	✓	—	—	✓	✓	—	Supported
8	✓	—	✗	—	✓	✗	—	—	✓	—	✓	✓	Not supported
9	✓	—	✗	—	✓	✗	—	✓	—	✓	✓	—	Supported
10	—	✓	✗	—	✓	✗	✓	—	—	✓	✓	—	Supported
11	—	✓	✗	—	✓	✗	—	—	✓	—	✓	✓	Supported
12	—	✓	✗	—	✓	✗	—	✓	—	✓	✓	✓	Supported

Footnote:

- ^ = Dual stack means both IPv4 and IPv6.
- * = IPv6 only is unsupported.

Limitations and known issues:

- For topology #3, adding an NSX Manager with only an IPv4 address will fail when using the **Add NSX Appliance** wizard on a dual stack ESXi which is onboarded to a dual stack VMware vCenter cluster using its IPv6 address.
The workaround is to onboard the ESXi host to the VMware vCenter cluster with its IPv4 address.
- For topology #9, when there is an IPv4 only NSX Manager and a dual stack ESX host which are registered in VMware vCenter using IPv6, there is a limitation where the ESX host cannot communicate with an IPv4 only NSX Manager.
- For topologies #10 and #12, IPv6 to IPv4 fallback between dual stack NSX Managers and dual stack transport nodes is not supported.

IPv6 Configuration Workflow for NSX Manager and Transport Node Communication

To enable IPv6 communication between NSX Manager and transport nodes, you can use the following workflow.

Procedure

1 Configure the IPv6 address on the management interface of the transport nodes:

- For ESX transport nodes, you can configure the static IP address or with DHCP.

Note It is recommended to use the static IP address.

- For NSX Edge transport nodes, you can only manually configure the IPv6 address.

For greenfield deployments, you can:

- Create a new NSX Edge node with static IPv4 and static IPv6.
- Create a new NSX Edge node with DHCPv4 and static IPv6.

For brownfield deployments, you can:

- Configure a static IPv6 NSX Edge node to become dual stack (static IPv4 and IPv6).
See [Configure a Brownfield NSX Edge Transport Node with IPv6](#).
- Configure a DHCP IPv6 NSX Edge node to become dual stack (DHCP IPv4 and static IPv6).
See [Configure a Brownfield NSX Edge Transport Node with IPv6](#).
- In the NSX CLI, edit and remove the NSX Edge node IPv6 address.

Note You cannot configure NSX Edge transport nodes as IPv6 only.

2 For greenfield NSX Managers, configure the IPv6 address on the management interface of the NSX Manager by creating a new NSX Manager node with static IPv4 and IPv6.

3 For brownfield NSX Managers, configure the IPv6 address depending on the brownfield scenario:

- For brownfield NSX Managers configured as IPv4 only from an NSX version prior to NSX 4.1, configure the IPv6 address through the CLI.

See [Configure a Brownfield NSX Manager Node with IPv6](#).

- For brownfield NSX Managers configured as dual stack from NSX 4.0.0.1, upgrade the NSX Manager cluster to NSX 4.1.

See the *NSX Upgrade Guide*.

Important It is required for all NSX Managers in the cluster to be upgraded to NSX 4.1 or later. If any of the NSX Manager nodes in the cluster are not upgraded to NSX 4.1 or later, then the communication between NSX Manager and transport nodes will use IPv4 and not IPv6.

Configure a Brownfield NSX Manager Node with IPv6

For existing NSX Managers with only an IPv4 address, you can configure an IPv6 address so that the NSX Manager becomes dual stack (both IPv4 and IPv6).

You can configure an IPv6 address for NSX Manager with the CLI. It is not possible to perform this task through the NSX Manager UI or by API.

Note NSX Managers do not support an IPv6 only configuration.

After configuring NSX Manager with an IPv6 address, the communication between NSX Manager and hosts automatically switches from IPv4 to IPv6 provided that the hosts are configured with IPv6 or dual stack.

Note For communication between dual stack NSX Managers and dual stack hosts, IPv6 is prioritized.

Prerequisites

- Ensure that you have NSX Manager configured with an IPv4 address. See [Chapter 4 NSX Manager Installation Requirements](#).
- Ensure that the IPv6 address to be used for the NSX Manager has the same hostname on the DNS server as the existing IPv4 address configured to NSX Manager.

Procedure

- 1 SSH or log into the NSX CLI.
- 2 Run the command `set interface eth0 ipv6 <NSX Manager's IPv6 address> gateway <NSX Manager's gateway IP address>` to configure the IPv6 address for the NSX Manager node.

Example:

```
NSX Manager> set interface eth0 ipv6 2620:124:6020:1045::ad/64 gateway
2620:124:6020:1045::253
```

If the IPv6 address is configured successfully, the following confirmation appears:

```
IPv6 address successfully updated. Node may take some time for IPv6 migration.
Node restart is required for IPv6 functionality to work properly.
```

Attention After running the command, it may take a few minutes for the communication between the NSX Manager and hosts to switch to IPv6.

- Restart the NSX Manager by running the command:

```
nsx-manager> reboot
Are you sure you want to reboot (yes/no): y
```

- (Optional) Run the command `get interface eth0` to verify that the IPv6 address has been configured for the NSX Manager node.
- Run the command `set name-servers <DNS IPv6 address>` to configure the IPv6 address for the DNS server.

Configure a Brownfield NSX Edge Transport Node with IPv6

For existing NSX Edge transport nodes with only an IPv4 address, you can configure an IPv6 address so the NSX Edge transport node becomes dual stack (both IPv4 and IPv6).

After configuring NSX Edge transport node with an IPv6 address, the communication between the transport node and NSX Manager automatically switches from IPv4 to IPv6 provided that the NSX Manager are configured as dual stack.

Note

- NSX Edge transport nodes do not support an IPv6 only configuration.
- In NSX version 4.1, NSX Edge only supports static IPv6 address configuration. DHCPv6 and SLAACv6 are unsupported.
- For communication between dual stack NSX Managers and dual stack hosts, IPv6 is prioritized.

The NSX Edge transport establishes a connection with the Appliance Proxy Hub (APH) through port 1234 and with the central control plane (CCP) through port 1235.

Prerequisites

There are existing NSX Edge transport nodes configured as IPv4 only.

Procedure

- 1 Configure the IPv6 address for the NSX Edge transport node by one of the following methods:

Method	Procedure
UI	<ol style="list-style-type: none"> 1 With admin privileges, log in to NSX Manager. 2 In the NSX Manager UI, navigate to System > Fabric > Nodes > Edge Transport Nodes. 3 Select the NSX Edge transport node that you want to configure an IPv6 address. 4 Select Actions > Change Node Settings. 5 In the Change Node Settings pop-window, enter the corresponding IPv6 addresses in the Management IP and Default Gateway fields. 6 Click Save.
API	<ol style="list-style-type: none"> 1 Enter the API command <code>PUT https://<nsx-manager>/api/v1/transport-nodes/<transport-node-id></code>. 2 For the body parameter <code>ipv6_assignment_type</code>, enter the value of STATIC. <p>Note In NSX version 4.1, only the IPv6 assignment type of static is supported. The IPv6 assignment types DHCPv6 and SLAAC are not supported.</p> 3 Create new body parameters for <code>ip_addresses</code> and <code>prefix_length</code> and enter the corresponding values for the IPv6 address. <p>Example:</p> <pre>"management_port_subnets": [{ "ip_addresses": ["2620:124:6020:1045::14"], "prefix_length": 64 }]</pre> 4 For the body parameter <code>default_gateway_addresses</code>, enter a new value for the default gateway address of the IPv6 address. <p>Example:</p> <pre>], "default_gateway_addresses": ["10.176.135.253", "2620:124:6020:1045::253"],</pre> 5 Send the API command. <p>Note It may take a few minutes for the IPv6 address to appear for the transport node.</p>
CLI	<ol style="list-style-type: none"> 1 SSH into the NSX CLI. 2 Run the CLI command <code>set interface</code> command: <pre>set interface <interface name> ip <IPv6 address> gateway <gateway IP address> plane mgmt</pre>

Method	Procedure
	<p>The following is an example:</p> <pre>set interface eth0 ip 245:124:6020:202::10/64 gateway 245:124:6020:202::253 plane mgmt</pre>

- (Optional) You can verify that the IPv6 address has been added for the NSX Edge transport node by navigating to **System > Fabric > Nodes > Edge Transport Nodes** and checking the **Management IP** column for the IPv6 address.

What to do next

After an NSX Edge transport node is configured with an IPv6 address, for any further IP address changes, see [Change the IP Address for an IPv6 NSX Edge Transport Node](#).

Change the IP Address for an IPv6 NSX Edge Transport Node

For NSX Edge transport nodes configured with an IPv6 address, you can change the IP address.

To change the IP address of an NSX Edge transport node configured with an IPv6 address, you must first use the NSX CLI to delete the configured IP addresses before reconfiguring the new IP addresses.

Prerequisites

This procedure is for NSX Edge transport nodes with an existing IPv6 address.

Procedure

- SSH into the NSX CLI.
- Run the CLI command `clear interface <interface name> ip` to delete all the IP addresses configured on the NSX Edge node.
- Configure the new IP addresses with one of the following methods:

Method	Procedure
API	<p>Note The API does not allow configuring IPv4 addresses.</p> <ol style="list-style-type: none"> Enter the API PUT <code>https://<nsx-manager>/api/v1/node/interfaces<interface-id></code>. For the <code>ip_addresses</code> and <code>prefix_length</code> body parameters, change the corresponding values. If necessary, change the existing value of the body parameter <code>default_gateway_addresses</code>. Send the API command.
CLI	<ol style="list-style-type: none"> SSH into the NSX CLI. Run the CLI command <code>set interface</code> command: <pre>nsx edge> set interface <interface name> ip <IPv4 or IPv6 address> gateway <gateway IP address> plane mgmt</pre>

- In the NSX CLI, run the command `service nsx-proxy restart` to restart `nsx-proxy`.

IPv6 Tunnel Endpoint Deployment for ESXi Host and NSX Edge on New Transport Zones

For IPv4 address exhaustion issues in new transport zones, you can use this workflow to create tunnel endpoints (TEP) that have IPv6 addresses for ESXi hosts and NSX Edge transport nodes.

With IPv6 TEPs, IPv6 communication is enabled between transport nodes. IPv6 communication between transport nodes supports bidirectional forwarding detection (BFD).

Prerequisites

- For the physical infrastructure, the MTU should be increased by 20 bytes to support additional overhead for IPv6 header.
- For ESXi hosts, the required software version is ESXi 8.0 Update 1 or later.

Procedure

- 1 Create a new transport zone and configure it for IPv6:
 - a From **System > Fabric > Transport Zones** page, click **Add Transport Zone**.
 - b Set the traffic type as overlay.
 - c Set the forwarding mode to IPv6.See [Create Transport Zones](#).
- 2 Configure the ESXi host for IPv6 TEPs:
 - a From the **System > Fabric > Hosts > Clusters** page, expand the cluster and then click the menu icon (3 dots) for the ESXi host and choose **Configure NSX**.
 - b From the **Configure NSX > Host details** page, click **Next**.
 - c From the **Configure NSX > Prepare Host** page, click the menu icon (3 dots) for the host switch and choose **Edit**.
 - d Add the transport zone for IPv6 to the host switch and configure the host switch IPv6 assignment details.See [Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP](#).
- 3 Create a new NSX Edge transport node and configure it for IPv6:
 - a Include the transport zone created from the previous step.
 - b Set the TEP IP address type to IPv6 and configure the IPv6 assignment details.See [Create an NSX Edge Transport Node](#).
- 4 Add the NSX Edge transport node to an edge cluster.
See [Create an NSX Edge Cluster](#).

- 5 Create and configure a tier-0 gateway:
 - a Attach the NSX Edge cluster.
 - b Create segments for the tier-0 gateway.
 - c Create and connect the tier-0 uplink and downlink interfaces.

See *Add a Tier-0 Gateway* in the *NSX Administration Guide*.

IPv6 Tunnel Endpoint Configuration for ESXi Host and NSX Edge on Existing Transport Zones

For IPv4 address exhaustion issues on existing transport zones, you can use this workflow to change existing tunnel endpoints (TEP) that have IPv4 addresses to use IPv6 addresses instead for ESXi hosts and NSX Edge transport nodes.

With IPv6 TEPs, IPv6 communication is enabled between transport nodes. IPv6 communication between transport nodes supports bidirectional forwarding detection (BFD).

Prerequisites

- For the physical infrastructure, the MTU should be increased by 20 bytes to support additional overhead for IPv6 header.
- For ESXi hosts, the required software version is ESXi 8.0 Update 1 or later.

Procedure

- 1 Turn the admin state off for the segments of the transport zones to be used for IPv6:
 - a Navigate to **Networking > Segments**.
 - b Click the menu icon (3 dots) for the segment and choose **Edit**.
 - c Set the **Admin State** toggle to off.
- 2 Configure the existing NSX Edge transport nodes for IPv6 TEPs:
 - a Edit the NSX Edge transport node and change the TEP IP address type to IPv6 and configure the IPv6 assignment details.
See [Edit NSX Edge Transport Node Configuration](#).
- 3 Configure the ESXi host for IPv6 TEPs:
 - a From the **System > Fabric > Hosts > Clusters** page, expand the cluster and then click the menu icon (3 dots) for the ESXi host and choose **Configure NSX**.
 - b From the **Configure NSX > Host details** page, click **Next**.
 - c From the **Configure NSX > Prepare Host** page, click the menu icon (3 dots) for the host switch and choose **Edit**.
 - d Configure the host switch IPv6 assignment details.

See [Prepare ESXi Cluster Hosts as Transport Nodes by Using TNP](#) for details.

4 Edit the transport zones and configure it for IPv6:

Note Before changing the transport zone to IPv6, ensure that all transport nodes using the transport zone are configured to IPv6.

- a From **System > Fabric > Transport Zones** page, click the menu icon (3 dots) for a transport zone and choose **Edit**.
- b Set the traffic type as overlay.
- c Set the forwarding mode to IPv6.

See [Create an NSX Edge Transport Node](#).

5 (Optional) Verify that transport nodes have IPv6 address:

- a Navigate to **System > Fabric > Nodes > Edge Transport Nodes**.
- b Click the NSX Edge cluster and then click the **Tunnel** tab.
- c Verify that the tunnels have IPv6 addresses.

Troubleshooting for IPv6 Configuration

Refer to this topic for possible issues and limitations that may occur when configuring IPv6 for NSX.

Table 11-1. VMware vCenter Fails to be Added as a Compute Manager When Using its IPv6 Address

Problem	Workflow that results in the problem: <ol style="list-style-type: none"> 1 Deploy a dual stack (both IPv4 and IPv6) NSX Manager cluster. 2 Deploy a dual stack VMware vCenter. 3 Add the VMware vCenter as a compute manager with its IPv6 address, and it fails to be added as a compute manager.
Cause	It is a requirement for a dual stack NSX Manager to have both its IPv4 and IPv6 addresses to point to the same FQDN that is used to configure the NSX Manager.
Solution	Use the same FQDN for both IPv4 and IPv6 addresses that are used to deploy NSX Manager. See Chapter 4 NSX Manager Installation Requirements .

Table 11-2. ESXi Host Fails to Deploy in a IPv4 Only NSX Manager

Problem	<p>Workflow that results in the problem:</p> <ol style="list-style-type: none"> 1 Deploy a topology with the following: <ul style="list-style-type: none"> ■ Dual stack (both IPv4 and IPv6) VMware vCenter ■ Dual stack ESXi ■ IPv4 only NSX Manager 2 Deploy NSX Managers as IPv4 only. 3 Add the VMware vCenter as a compute manager. 4 Start the Add NSX Appliance wizard and select a dual stack ESXi host from the cluster and deploy with the NSX Manager IPv4 address. 5 The installation fails with the following error: <pre style="background-color: #f0f0f0; padding: 10px;">Error occurred during vmdk transfer. java.net.SocketException Protocol family unavailable</pre>
Cause	<p>For an IPv4 only NSX Manager and a dual stack ESXi host which are registered in VMware vCenter using IPv6, there is a limitation where the ESXi host cannot communicate with an IPv4 only NSX Manager.</p>
Solution	<p>The workaround is to onboard the ESXi host to the VMware vCenter cluster with its IPv4 address.</p>

Table 11-3. Deploying NSX Manager with Virtual Disk Format of Thick Provision Eager Zeroed Fails to Deploy

Problem	<p>Workflow that results in the problem:</p> <ol style="list-style-type: none"> 1 Deploy a supported VMware vCenter and ESXi topology. See Supported Topologies for IPv6. 2 Add the VMware vCenter as the compute manager. 3 Start the Add NSX Appliance wizard and deploy the third NSX Manager with a Thick Provision Eager Zeroed virtual disk format which results in the deployment progress to be stuck at 1%.
Cause	<p>The time for a thick provision eager zeroed disk creation exceeds the timeout length.</p>
Solution	<p>In VMware vCenter, for the vpxa configuration setting <code>task.completedLifetime</code>, increase the default value of 600 seconds:</p> <pre style="background-color: #f0f0f0; padding: 10px;">[root@sc1-10-78-185-35:/] configstorecli config default get -c esx -g services -k vpxa { ... "task": { ... "completed_lifetime": 600, ... }, ... }</pre>

Table 11-4. Adding an IPv6 Only ESX Host to a vLCM Cluster Fails

Problem	<p>Workflow that results in the problem:</p> <ol style="list-style-type: none"> 1 Deploy a topology with the following: <ul style="list-style-type: none"> ■ Dual stack (both IPv4 and IPv6) NSX Manager ■ IPv6 only VMware vCenter ■ IPv6 only ESX 2 Add the VMware vCenter as a compute manager. 3 Create a vSphere Lifecycle Manager (vLCM) cluster. 4 Create an NSX transport node profile and apply the profile to the vLCM cluster. 5 Add an ESX host to the vLCM cluster which causes an error and fails.
Cause	For IPv6 only ESX and IPv6 only VMware vCenter, the vLCM workflow is not supported in software versions 7.x or earlier for ESX and vCenter.
Solution	None. This is a known limitation.

Table 11-5. Uninstalling Transport Node Clears the IPv6 DNS Configuration

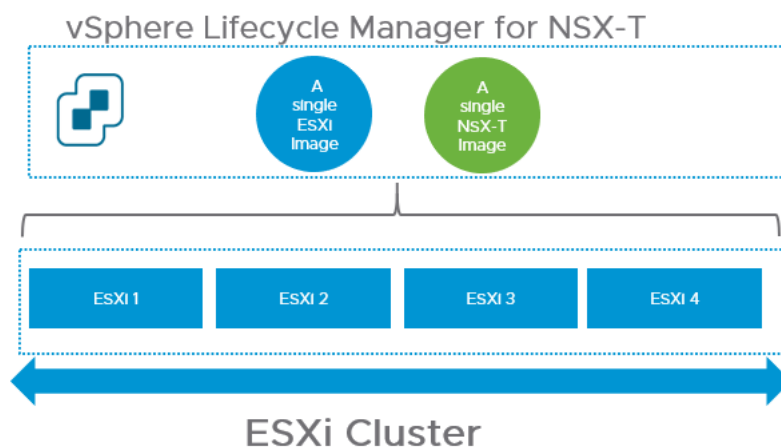
Problem	<p>Workflow that results in the problem:</p> <ol style="list-style-type: none"> 1 Deploy a topology with the following: <ul style="list-style-type: none"> ■ Dual stack (both IPv4 and IPv6) NSX Manager ■ Dual stack VMware vCenter ■ IPv6 only and dual stack ESXi 2 Install NSX on the cluster, add a transport node profile, and VLAN transport zone with DHCP. 3 Uninstall NSX from the cluster which clears the IPV6 DNS configuration for IPv6 only and dual stack ESXi hosts.
Cause	This problem occurs in vCenter and ESX software versions 7.0.3 and earlier.
Solution	None for ESXi 7.0.3 or earlier. Workaround is to reconfigure the DNS server. You can upgrade ESXi to 8.0.0.1 for the issue fix.

vSphere Lifecycle Manager with NSX

12

By enabling VMware vSphere® vSphere Lifecycle Manager on a cluster, you can ensure that all ESXi hosts participating in a cluster are prepared using a single ESXi image and a single NSX image. The vSphere Lifecycle Manager functionality minimizes errors and lowers cluster and host maintenance cycles.

Starting from vCenter Server 7.0 U1, ESXi 7.0 U1, and NSX 3.1.0 onwards, a vSphere Lifecycle Manager-enabled cluster can manage installation of ESXi and NSX VIBs.



vSphere Lifecycle Manager requires two images: one for ESXi and another one for NSX. It retrieves the ESXi image from the image directory in VMware vCenter. Ensure the ESXi image is uploaded to VMware vCenter. vSphere Lifecycle Manager gets the NSX image only when a cluster is prepared for NSX networking, which is possible from the NSX Manager user interface. The NSX image is automatically uploaded to VMware vCenter when NSX cluster preparation begins. For other clusters in the VMware vCenter, vSphere Lifecycle Manager references the already uploaded NSX image. vSphere Lifecycle Manager refers to NSX as a solution, as it does with other solutions such as HA, DRS and so on.

For more information on the usage of the terminology, such as images and solutions in vSphere Lifecycle Manager, refer to the *Managing Host and Cluster Lifecycle guide* in the VMware vSphere® Documentation center.

The following clusters can be enabled as vSphere Lifecycle Manager clusters:

- Clusters with ESXi hosts that are prepared for NSX networking using a transport node profile.

- Clusters with ESXi hosts that are not prepared for NSX networking.

Unsupported Scenarios

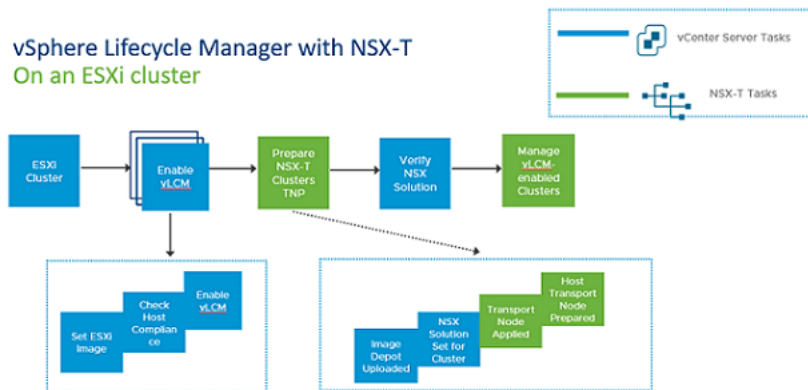
If vSphere Lifecycle Manager is enabled on an ESXi cluster, then you cannot apply an NSX transport node profile to prepare the cluster hosts as transport nodes.

Read the following topics next:

- [Prepare an NSX Cluster with vSphere Lifecycle Manager](#)
- [Enable vSphere Lifecycle Manager on an NSX Cluster](#)
- [NSX on vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine](#)
- [NSX with vSphere Lifecycle Manager Scenarios](#)
- [vSphere Lifecycle Manager Failed to Prepare a Host for NSX Networking](#)
- [vSphere Lifecycle Manager Failed to Prepare NSX Cluster](#)
- [Delete a NSX Depot on vCenter Server](#)

Prepare an NSX Cluster with vSphere Lifecycle Manager

You can prepare NSX on vSphere Lifecycle Manager-enabled clusters.



For vSphere Lifecycle Manager to get access to the NSX image, you must configure the cluster with a transport node profile. When you begin configuring the cluster, NSX local control plane bundle (in the format - nsx-lcp-bundle-*<release_version.build_version>*) is uploaded to the image repository in VMware vCenter.

During host preparation, vSphere Lifecycle Manager accesses the depot and sets NSX as a solution for that cluster. It applies the NSX solution to the cluster, which begins with the process of remediating hosts. Every host is remediated by vSphere Lifecycle Manager before the NSX switch is configured on the host. vSphere Lifecycle Manager remediation happens when a new ESXi host is added to a vSphere Lifecycle Manager cluster.

vSphere Lifecycle Manager remediates hosts so that the image on each host is the same as the ESXi version set for the cluster. Any drift must be resolved before host preparation can progress in NSX. During cluster preparation, if the cluster fails, NSX sets the cluster state to Failed. As an admin, you must retrigger host remediation by taking appropriate actions either from the NSX Manager user interface or from the vSphere Client.

Prerequisites

- Ensure all hosts in a cluster are running at least ESXi 7.0 U1 version or higher.
- Ensure Lockdown mode is not enabled on any of the hosts. vSphere Lifecycle Manager might fail to prepare hosts that are enabled to function in Lockdown mode.
- Ensure there is not drift in images between hosts and cluster. Otherwise, you cannot enable vSphere Lifecycle Manager on the cluster. Remediate hosts in VMware vCenter to ensure base image matches on host and cluster.
- Ensure vSphere Lifecycle Manager is enabled on the cluster. See VMware vSphere® documentation.
- Register Compute Manager with the following settings:
 - Enable **Trust** and set access level to vSphere Lifecycle Manager. Trust is mandatory to establish communication between NSX and vSphere Lifecycle Manager.
 - Enable **Create Service Account**.
- Create a transport node profile using a vSphere Distributed Switch host switch. N-VDS switch is not supported on a vSphere Lifecycle Manager-enabled.
- If you configure a Web Proxy on a VMware vCenter Appliance, add all NSX Manager IP addresses to the NO_PROXY list, otherwise vLCM cannot connect to NSX Managers.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Go to **System > Fabric > Hosts > Clusters**.
- 3 Select the cluster and click **Configure NSX**.

Note Identify vSphere Lifecycle Manager-enabled cluster when a cluster is accompanied with vLCM text.

- 4 Select a transport node profile that uses vSphere Distributed Switch as the host switch.
- 5 Click **Apply TNP**.

Important If vLCM Config Manager is enabled on this cluster, you will not be able to apply a TNP to this cluster. To remediate this issue, move the host to a cluster where vLCM Config Manager is not enabled.

If this is the first cluster that is enabled for vSphere Lifecycle Manager, NSX uploads the NSX LCP bundle to the image repository in VMware vCenter. vSphere Lifecycle Manager sets NSX as a solution on the cluster. It sets the desired state to the NSX image uploaded to VMware vCenter. Then, vSphere Lifecycle Manager begins installation of NSX VIBs on each host, followed by configuration of NSX switch on each transport node.

As part of host preparation, vSphere Lifecycle Manager remediates the host, registers the host with NSX Manager, configures NSX switch on the host and completes the configuration.

Note Installing NSX on a vSphere Lifecycle Manager-enabled cluster might take a little more time than when installing on a non-vSphere Lifecycle Manager-enabled cluster. This difference is due to the additional health checks that are included in this combination of products

6 Troubleshooting issues:

If vSphere Lifecycle Manager could not apply NSX as a solution to the cluster, the NSX cluster in NSX Manager goes into Failed state. To remediate the hosts in the cluster, do one of the following:

a Go to the VMware vCenter, verify the following conditions are met:

- Hosts are compliant.
- Hosts are not powered off or in maintenance mode.

b Verify cluster status through UI or API. Even if a host in the cluster is in Failed state, the cluster status remains in unrealized state.

Run the following API to verify the cluster state, GET /<NSX-Manager-IP>/api/v1/transport-node-collections/<transport-node-collection-id>.

c If any one of the host fails, the remaining hosts in the cluster go into Install Skipped state. To remediate, read the error message and take any necessary action. Then, click **Resolve** to retry remediation of the host and NSX preparation. Note that remediation happens serially, one host at a time.

d If the cluster is still in Install Failed state, click **Resolve** for the cluster in UI or run the API to realize the transport node profile on the cluster. Along with remediating the cluster, the following API also tries to prepare those hosts that are in the Install Skipped state. It retries remediation on the entire cluster. It tries to prepare the hosts where installation is skipped.

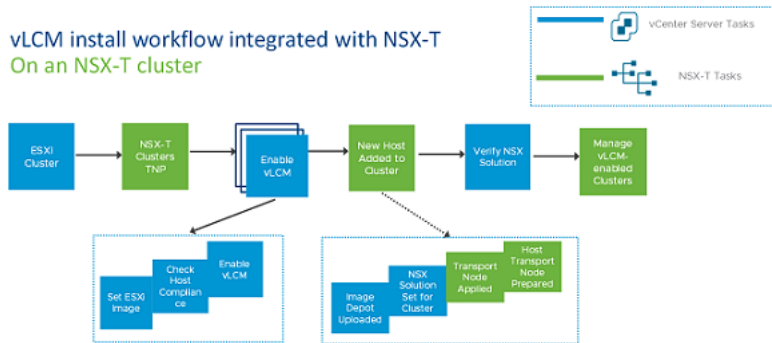
POST /api/v1/transport-node-collections/<transport-node-collection-id>?action=retry_profile_realization

Results

Sphere Lifecycle Manager prepares all hosts in the cluster as NSX transport nodes.

Enable vSphere Lifecycle Manager on an NSX Cluster

As you can prepare an NSX cluster on a cluster that is already enabled with vSphere Lifecycle Manager, similarly, you can enable vSphere Lifecycle Manager on an existing NSX prepared cluster.



Prerequisites

- Ensure all hosts in a cluster are running ESXi 7.0 U1 version.
- Register Compute Manager with the following settings:
 - Enable **Trust** and set access level to vSphere Lifecycle Manager. Trust is mandatory to establish communication between NSX and vSphere Lifecycle Manager.
 - Enable **Create Service Account**.
- Prepare the cluster by applying a Transport Node Profile (using VDS as the host switch type) to the cluster.

Note N-VDS host switch is not supported on a vSphere Lifecycle Manager-enabled cluster.

Procedure

- 1 From a browser, log in with admin privileges to a VMware vCenter at <https://<vcenter-server-ip-address>>.
- 2 Select the cluster on which the vSphere Lifecycle Manager functionality must be enabled.
- 3 On the Images page, confirm that all hosts are compliant. If any of the host is in non-compliant state, remediate the host to be compliant with the ESXi image set for the cluster.

- 4 Verify that vSphere Lifecycle Manager sets the solution for the cluster to NSX. To verify that the NSX solution is set on the vSphere Lifecycle Manager cluster, you can do one of the following:

- a In VMware vCenter, on the Images page, click **Check Compliance** and check the Components section for an NSX entry.
- b Alternatively, run the API command and verify that component and version are correctly set to NSX.

```
GET https://{server}/api/esx/settings/clusters/{cluster}/software/solutions/
com.vmware.nsxt?vmw-task=true
  components" : [
    {
      "component" : "nsx-lcp-bundle"
    }
  ],
  "version" : "3.1-0"
```

- 5 When a new host is added to the vSphere Lifecycle Manager-enabled cluster, NSX calls vSphere Lifecycle Manager to check host compliance with the ESXi image set for the cluster. If there is no drift in host and cluster image, then transport node profile is applied to the host. NSX VIBs on the host. The final part of the installation is followed by registration with NSX Manager and NSX switch configuration.

- 6 Troubleshooting issues:

If vSphere Lifecycle Manager could not apply NSX as a solution to the cluster, the NSX cluster in NSX Manager goes into **Failed** state. To remediate the hosts in the cluster, do one of the following:

- a Go to the vCenter Server, verify the following conditions are met:
 - Hosts are compliant.
 - Hosts are not powered off or in maintenance mode.
- b Verify cluster status through UI or API. Even if a host in the cluster is in **Failed** state, the cluster status remains in unrealized state.

Run the following API to verify the cluster state, GET /<NSX-Manager-IP>/api/v1/transport-node-collections/<transport-node-collection-id>.

- c If any one of the host fails, the remaining hosts in the cluster go into `Install Skipped` state. To remediate, read the error message and take any necessary action. Click **Resolve** to retry remediation of the host and NSX preparation. Note that remediation happens serially, one host at a time.
- d If the cluster is still in `Failed` state, click **Resolve** for the cluster in UI or run the API to realize the transport node profile on the cluster. Along with remediating the cluster, the following API also tries to prepare those hosts that are in the `Install Skipped` state. It retries remediation on the entire cluster. It tries to prepare the hosts where installation is skipped.

```
POST /api/v1/transport-node-collections/<transport-node-collection-id>?
action=retry_profile_realization
```

Results

vSphere Lifecycle Manager is enabled on a NSX cluster.

What to do next

After enabling vSphere Lifecycle Manager on the cluster, you can remediate drifts between hosts in VMware vCenter with the image set for the cluster.

NSX on vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine

Starting with NSX 4.0.1.1, vSphere Distributed Services Engine provides the ability to offload some of the network operations from your server CPU to a data processing unit (DPU also known as SmartNIC).

Using DPU devices for network acceleration frees up the CPU capacity for business-critical workloads. Besides accelerating networking performance, DPU devices improve network visibility and security acceleration.

Note DPUs are supported on hosts in a vSphere Lifecycle Manager-enabled cluster that are running at least ESXi 8.0 version or higher.

NSX supports NVIDIA BlueField-2 (25G) and AMD Pensando (25G and 100G) DPUs only. However, starting with NSX 4.1.1, NSX also supports NVIDIA BlueField-2 (100G).

License


To utilize the NSX DPU-based acceleration or NSX offload capabilities, you need to purchase an NSX Enterprise Plus Term license (Per Core basis) or an NSX Enterprise Plus with Threat Prevention Term license (Per Core basis).

Note For vSphere offload capabilities, you do not need to purchase a separate NSX license. You just need to have the vSphere ENT + Term license and be on vSphere 8. NSX Manager is available as a part of vSphere ENT+.

Once the license key is applied, you should be able to offload the routing and DFW capabilities to the DPU.

Enable Network Offloads

Workflow to Enable Network Offloads:



Reference Topics

For more information, refer to the following topics:

VMware vSphere Distributed Services Engine	See the topic, <i>Introducing VMware vSphere® Distributed Services Engine™ and Networking Acceleration by Using DPUs</i> , in the <i>VMware ESXi Installation and Setup</i> guide under VMware vSphere Product Documentation.
vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine	See the topic, <i>Using vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine</i> , in the <i>Managing Host and Cluster Lifecycle</i> guide under VMware vSphere Product Documentation.
Network Offloads and vCenter Server specific steps	See the topic, <i>What is Network Offloads Compatibility</i> , in the <i>vSphere Networking</i> guide under VMware vSphere Product Documentation.
ESXi on DPU	See the topic, <i>Introducing ESXi on Data Processing Units (DPUs)</i> , in the <i>VMware ESXi Installation and Setup</i> guide under VMware vSphere Product Documentation.
Prepare an NSX Cluster with vSphere Lifecycle Manager and troubleshooting issues	See Prepare an NSX-T Cluster with vSphere Lifecycle Manager

Configure NSX host transport node on DPU-based vSphere Lifecycle Manager-enabled cluster

You need to configure NSX and enable Enhanced Datapath for vSphere Distributed Services Engine to offload some of the network operations from your server CPU to the DPU.

Configuring NSX host transport node on a DPU-based vSphere Lifecycle Manager-enabled cluster is similar to [Prepare an NSX Cluster with vSphere Lifecycle Manager](#).

vSphere Distributed Switch (VDS) backed by the DPU on ESXi supports the offloading mode after NSX is enabled. Traffic forwarding logic is offloaded from the ESXi host to the VDS backed by DPU.

Note A DPU is supported only on vSphere Lifecycle Managed clusters. DPU requires a minimum version combination of vSphere 8.0, NSX 4.0.1.1, and Edge hardware version 20.

To learn more about Network Offloads Capability, see *What is Network Offloads Capability* in the VMware vSphere Documentation.

Prerequisites

- vSphere offloads: DPUs are supported on hosts in a vSphere Lifecycle Manager-enabled cluster that are running on ESXi 8.0 version or higher.

Note For vSphere offload capabilities, you do not need to purchase a separate NSX license. You just need to have the vSphere ENT + Term license and be on vSphere 8. NSX Manager is available as a part of vSphere ENT+.

- NSX offloads: To utilize the NSX DPU-based acceleration capabilities, you need to purchase an NSX Enterprise Plus Term license (Per Core basis) or an NSX Enterprise Plus with Threat Prevention Term license (Per Core basis).

Note You do not need any additional NSX licenses if you have an NSX Enterprise Plus Term license (Per Core) or an NSX Enterprise Plus with Threat Prevention Term license (Per Core).

For more information, see *NSX Feature and Edition Guide*.

- Lockdown mode is not enabled on any of the hosts. vSphere Lifecycle Manager might fail to prepare hosts that are enabled to function in Lockdown mode.
- There is not drift in images between hosts and cluster. Otherwise, you cannot enable vSphere Lifecycle Manager on the cluster. Remediate hosts in vCenter Server to ensure base image matches on host and cluster.
- vSphere Lifecycle Manager is enabled on the cluster. See VMware vSphere® documentation.
- Register Compute Manager with the following settings:
 - Enable **Trust** and set access level to vSphere Lifecycle Manager. Trust is mandatory to establish communication between NSX and vSphere Lifecycle Manager.
 - Enable **Create Service Account**.

- Create a transport node profile using a vSphere Distributed Switch host switch. NSX Virtual Distributed Switch (N-VDS) is not supported on vSphere Lifecycle Managed clusters.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Go to **System** → **Fabric** → **Hosts** → **Clusters**.
- 3 Select the cluster and click **Configure NSX**.

Note Identify vSphere Lifecycle Manager-enabled cluster when a cluster is accompanied with vLCM text.

- 4 To create a transport node profile, click **Add Profile**.
- 5 Enter Transport Node Profile (TNP) details as required.

Note

- For VDS backed by the DPU, select the Host TN and TNP Host Switch mode as **Enhanced Datapath - Standard** (Recommended) or **Enhanced Datapath - Performance**.
 - In **Add Switch**, select VDS (Distributed Switch version 8.0 created in vSphere Client) with offload compatibility. If there is a mismatch, the host will not be compatible.
 - One cluster allows only one type of host. For example, when you create two distributed switches on vSphere client: one with NVIDIA network offload compatibility and one with AMD Pensando network offload compatibility, both options will appear on the UI. Select the one specific to your requirement. Clusters of different types cannot use the same TNP. Therefore, you will need two separate TNPs: one for NVIDIA DPU and one for AMD Pensando DPU.
-
- 6 If TNP is already created, select a transport node profile that uses vSphere Distributed Switch as the host switch and datapath mode as Enhanced Datapath.
 - 7 Click **Apply TNP**.

If this is the first cluster that is enabled for vSphere Lifecycle Manager, NSX uploads the NSX LCP bundle to the image repository in vCenter Server. vSphere Lifecycle Manager sets NSX as a solution on the cluster. It sets the desired state to the NSX image uploaded to vCenter Server. Then, vSphere Lifecycle Manager begins installation of NSX VIBs on each host, followed by configuration of NSX switch on each transport node.

As part of host preparation, vSphere Lifecycle Manager remediates the host, registers the host with NSX Manager, configures NSX switch on the host and completes the configuration.


Note

- vSphere Lifecycle Manager puts the ESXi host backed by AMD Pensando DPU in maintenance mode and reboots it as part of host remediation. If vSphere Lifecycle Manager fails to place the host in maintenance mode, you need to manually power off all VMs and then retry NSX installation
 - Installing NSX on a vSphere Lifecycle Manager-enabled cluster might take a little more time than when installing on a non-vSphere Lifecycle Manager-enabled cluster. This difference is due to the additional health checks that are included in this combination of products.
-

Displaying DPU-related information on NSX Manager Interface

After you install and configure NSX on a host, you can monitor inventory objects of Host Transport Nodes and view the DPU-related information on the NSX Manager interface.

To view the DPU-related information for the host transport node on a DPU-based vSphere Lifecycle Manager-enabled cluster, go to **System > Fabric > Hosts**. Select the desired host and then click **View Details** to display the information:

- **DPU:** On the **Monitor** tab, the DPU chart displays the CPU cores allocated and the system memory used by a host on DPU. Click on the info  icon beside the DPU field to view the DPU-related information, such as the device number, vendor or model name, firmware version, and OS version.
- **DPU Backed:** On the **Physical Adapters** tab, the **DPU Backed** field displays if the hypervisor host is backed by the DPU or not. This is to know the DPU presence on the ESXi host:
 - If DPU Backed is 'Yes', the interface is backed by the DPU and is in the 'MANAGED' state.
 - If DPU Backed is 'No', it is the standard hypervisor host.

Note ESXi on DPU is used as a traditional NIC until NSX transport node is configured. The VDS on vCenter Server indicates if network offloading is permitted when NSX is enabled.

Hosts backed by DPU are associated with VDS. You get a granular view of VDS backed by the DPU at an individual host level. Click the **Switch Visualization** tab to view the uplinks configured on the VDS backed by the DPU that is connected to uplinks.

NSX with vSphere Lifecycle Manager Scenarios

Installation and uninstallation scenarios to consider when you work with vSphere Lifecycle Manager (vLCM) for NSX clusters.

Scenario	Result
You try to enable vLCM on a cluster where transport node profile is not applied, but some hosts are individually prepared as host transport nodes.	vLCM cannot be enabled on the cluster because a transport node profile was not applied to the cluster.
You try to enable vLCM on a cluster using a transport node profile configured to apply N-VDS host switch.	vCenter Server checks for cluster eligibility so that the cluster can be converted to a vLCM cluster. As N-VDS host switch type is not supported, apply a transport node profile that is configured to use a VDS host switch.
You move an unprepared host from a non-vLCM cluster to a vLCM cluster.	If the vLCM cluster is prepared with a transport node profile, the unprepared host is prepared as an NSX transport node by vLCM. If the vLCM cluster is not prepared with a transport node profile, the host remains in unprepared state.
You move a transport node from a vLCM cluster to a non-vLCM cluster that is not prepared for NSX.	The NSX VIBs are deleted from the host, but the NSX Solution-related data (set by vLCM) is not deleted. Now, if you try to enable vLCM cluster on the cluster, NSX Manager notifies that NSX Solution will be removed from the host. This notification is misleading because NSX VIBs were already removed on the host.
After you perform the Remove NSX operation on a vSphere Lifecycle Manager cluster, if vLCM is unable to delete NSX from the desired state, all nodes go into <code>Uninstall Failed</code> state. Now, you try to remove NSX on individual transport nodes.	If you remove NSX on each individual transport node, then even though NSX VIBs are removed on the host, the cluster continues to have NSX as the desired state in vLCM. This state shows up as drift in host compliance in vCenter Server. So, you must perform Remove NSX on the cluster to remove NSX from the vLCM configuration.
You prepare a vLCM cluster consisting of a host by applying a TNP. The VDS switch type is configured in the TNP. You move the host into maintenance mode and move the vLCM cluster out of the vLCM cluster into datacenter. And finally, move the host back to the vLCM cluster.	NSX installation fails with the following message: <pre>Failed to install software on host. Solution apply failed on host: '192.196.178.156' Deployment status of host bfeedeb69-48d3-4f3b-9ebc-ce4eb177a968 is INSTALL_IN_PROGRESS with 0 errors. Expected status is INSTALL_SUCCESSFUL with no errors.Deployment status of host bfeedeb69-48d3-4f3b-9ebc-ce4eb177a968 is INSTALL_IN_PROGRESS with 0 errors. Expected status is INSTALL_SUCCESSFUL with no errors.Solution apply failed on host: '192.196.178.156'</pre> Workaround: On the host, click Resolve and reapply TNP.
You move a host that failed installation to a non-vLCM cluster with or without a TNP applied.	NSX does not perform any operation.
You move a host that failed installation to a vLCM cluster with TNP applied.	NSX installation begins automatically.
You move a host that failed installation to a vLCM cluster without TNP applied.	NSX does not perform any operation.
You move a host that failed installation to a datacenter.	NSX does not perform any operation.

Scenario	Result
VMware vCenter is added as a compute manager with Multi NSX flag enabled. Apply TNP on another existing vLCM cluster.	NSX allows preparation of the existing vLCM cluster using the TNP.
VMware vCenter is added as a compute manager with Multi NSX flag enabled. Then try to change the already prepared cluster to a vLCM cluster.	NSX does not allow preparation of the existing vLCM cluster.
VMware vCenter is added as a compute manager with Multi NSX flag enabled. Then try creating a new vLCM cluster.	NSX allows preparation of the existing vLCM cluster.
VMware vCenter already contains a vLCM cluster and you try to add the VMware vCenter as a compute manager with Multi NSX flag enabled.	NSX fails this operation because the VMware vCenter already contains a vLCM cluster.
You move a DPU-enabled host from a TNP applied cluster to non-TNP applied cluster.	Vibs are not deleted from ESXi host and DPU. You need to remediate the host from vSphere Lifecycle Manager. vSphere Lifecycle Manager deletes NSX vibs from ESXi host and DPU, and reboots the host.
You remove NSX VIBs from a DPU-enabled host by using 'del nsx' nsxcli.	After running the 'del nsx' command, you need to reboot the ESXi host to complete the NSX VIBs removal process (VIBs are removed from ESXi and DPU).
NSX VIBs are not deleted from DPU-enabled host even if it is shown as Not Configured on the NSX UI.	<p>When a DPU-enabled host is disconnected during the Remove NSX operation, the operation fails. After a while, the TN deletion continues and the host is displayed as Not Configured but the VIBs are not removed.</p> <p>Workaround: Go to the vLCM UI and remediate the cluster.</p> <hr/> <p>Note The DPU-enabled host reboots as part of the remediation.</p>
When you enable vLCM on a cluster that is configured using a transport node profile the transition is successful. But the vLCM remediation (Apply NSX task on vLCM) fails.	Host transport nodes will continue to have their status before enabling vLCM. Go to the vLCM UI to check whether the Apply NSX operation has failed. Also, verify cluster details to find more details about the drift between the desired state and host.
You cannot prepare a host that is already in Uninstall Failed state. The host is in dirty state, which means that the transport node record and files are not completely removed from the host.	Before you prepare a failed host, delete transport node entry by using the Force Delete option. This operation deletes the transport node record. Then, using <code>del nsx</code> command remove NSX from the host. After this try to prepare host or move host into a TNP applied cluster.

vSphere Lifecycle Manager Failed to Prepare a Host for NSX Networking

vSphere Lifecycle Manager failed to prepare some hosts in the cluster for NSX Networking.

Problem

In a cluster containing many hosts, vSphere Lifecycle Manager successfully prepared some hosts, whereas vSphere Lifecycle Manager failed to realize NSX on one of the host.

Cause

Hosts can take different states when vSphere Lifecycle Manager triggers installation of NSX.

- Cluster goes into `Install Failed` if vSphere Lifecycle Manager fails to remediate the entire cluster.
- If one or more individual hosts fail, failed hosts go into `Install Failed` state. If there are other hosts in the cluster yet to be prepared, those hosts go into `Install Skipped` state. Both cluster and individual hosts display failure states.

Solution

- 1 On the NSX Manager, go to **System > Fabric > Hosts > Clusters**.
- 2 Identify the failed cluster to view the error state. Click the error link to open a popup window.
- 3 If the cluster is in `Install Failed` state, click **Resolve** at the to initiate transport node profile realization on the cluster.

Important With the cluster in `Install Failed` state, first try to resolve the remediation issues at the cluster and then try to remediate individual hosts. If you overlook cluster-level errors and directly try to remediate host-level errors, the UI does not allow you to perform any remediation action at the host-level.

- 4 If one or more hosts failed but the cluster remediation status is `Success`, then navigate to the failed host and click **Resolve** to remediate hosts.
- 5 You can also try to realize the transport node profile on the cluster by executing the following API command, `POST /api/v1/transport-node-collections/{tnc id}?action=retry_profile_realization`.

This API command re-triggers the transport node profile on the cluster.

vSphere Lifecycle Manager Failed to Prepare NSX Cluster

vSphere Lifecycle Manager failed to prepare an NSX cluster.

Problem

Transport node profile could not be applied to a vSphere Lifecycle Manager cluster, which caused the cluster to go into `Failed` state.

Cause

As vSphere Lifecycle Manager failed to set and apply NSX as a solution on the cluster, none of the hosts in the cluster were prepared as transport nodes. View the error state on the cluster in NSX Manager UI.

Solution

- 1 On the NSX Manager, go to **System > Fabric > Hosts > Clusters**.
- 2 Identify the failed cluster to view the error state. Click the error link to open a popup window.
- 3 Click **Resolve** to initiate transport node profile realization on the cluster.
- 4 Alternatively, run the API command, `POST /api/v1/transport-node-collections/<tnc id>?action=retry_profile_realization`. This command initiates vLCM to realize NSX on the cluster.

Delete a NSX Depot on vCenter Server

Delete a NSX Depot on vCenter Server.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Unregister the compute manager from NSX.

When the compute manager is unregistered, NSX invokes API to delete the depot in VMware vCenter. Depot is deleted from the image depot of vSphere Lifecycle Manager in VMware vCenter

- 3 Troubleshooting:
 - a If NSX is unable to delete the depot after unregistering the compute manager, run the following API.
 - b Remove the depot entry from the payload.
 - c To verify that the depot is successfully deleted, run the API command.

```
DELETE https://{server}/api/esx/settings/depots/offline/{depot}
```

```
GET https://{server}/api/esx/settings/depots/offline/{depot}
```

Host Profile integration with NSX

13

Integrate host profiles extracted from an ESXi host with NSX to deploy ESXi and NSX VIBs on stateful and stateless servers.

Read the following topics next:

- [Auto Deploy Stateless Cluster](#)
- [Stateful Servers](#)

Auto Deploy Stateless Cluster

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 Transport Node Profile to reprovision a host with a new image profile that contains a different version of ESXi and NSX. 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 Transport Node Profile to configure NSX on the hosts.

So, a stateless host can be set up for vSphere Auto Deploy and NSX Transport Node Profile to reprovision a host with a custom ESXi and NSX version.

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](#).
- 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 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 ESXi Version	ESXi 67ep6	ESXi 67u2	ESXi 67u3	ESXi 67ep7	ESXi 67ep15	ESXi 7.0
NSX 2.4	Yes	Yes	No	No	No	No
NSX 2.4.1	Yes	Yes	No	No	No	No
NSX 2.4.2	Yes	Yes	No	No	No	No
NSX 2.4.3	Yes	Yes	No	No	No	No
NSX 2.5	Yes	Yes	Yes	Yes	No	No
NSX 2.5.1	Yes	Yes	Yes	Yes	Yes	No
NSX3.0	Yes	Yes	Yes	Yes	Yes	Yes
NSX3.1 and later	Yes	Yes	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 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 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 packages, create a software depot.
- 2 Download the `nsx-1cp` packages.
 - a Go to the [Broadcom Support](#) page. Select the **VMware Cloud Foundation** division on the top panel and go to the **My Downloads** panel.
 - b In the **My Downloads** panel, search NSX Kernel Modules for a specific VMware ESXi version.
 - c Click **Download Now** to begin downloading the `nsx-1cp` package.
 - d Import `nsx-1cp` 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-1cp` 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 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 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**.

The screenshot shows the 'Deployed Hosts' tab in the vSphere Client. At the top, there are navigation tabs: Software Depots, Deploy Rules, **Deployed Hosts** (selected), Discovered Hosts, Script Bundles, and Configure. Below the tabs, a message states: '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.' Below this message are three action links: CHECK HOST ASSOCIATIONS COMPLIANCE, REMEDIATE HOST ASSOCIATIONS, and EDIT IMAGE PROFILE ASSOCIATION. The main content is a table with the following columns: Host, Associated Image Profile, Associated Host Profile, Associated Location, and Associated Script Bundle. There are two rows of data in the table.

<input type="checkbox"/>	Host	Associated Image Profile	Associated Host Profile	Associated Location	Associated Script Bundle
<input type="checkbox"/>	10.144.139.147	CustomDepot(ESXi and NSX)		1-datacenter-1964	
<input type="checkbox"/>	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](#).

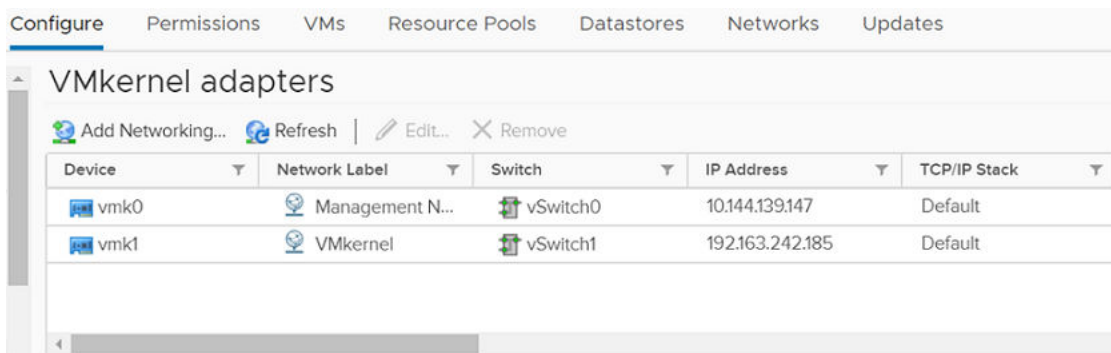
Set Up Network Configuration on the Reference Host

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. See [Configure the Reference Host as a Transport Node in NSX](#).


Configure the Reference Host as a Transport Node in NSX

After the reference host is associated with the custom image profile and configured with a VSS or DVS switch, deploy the reference host as a NSX transport node with NSX enabled DVS switch.

Procedure

- 1 From a browser, log in to NSX at https://<NSXManager_IPaddress>.
- 2 To locate the reference host, navigate to **System > Fabric > Hosts > Clusters**.
- 3 You also need to create a VLAN transport zone to define span of a virtual network. The span is defined by attaching VDS switches to the transport zone. Based on this attachment, 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 VDS switch 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 On the **Clusters** tab, select the reference host.
 - b (On a VDS switch) Click **Configure NSX** and select the previously created transport zone, VDS, uplink profile.
- 7 Click **Finish** to begin installation of NSX on the reference host.

(On a VDS switch) After installation, configuration status of the reference host is displayed as `Success`. In the vCenter Server, the VDS switch is displayed as NSX switch.

Note The reference host is listed under Other Hosts.

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 configuration extracted in the host profile. It consists of ESXi and NSX configuration that is applied to target hosts.

Procedure

- 1 To extract the host profile, see [Extract and Configure Host Profile from the Reference Host](#).

2 In the extracted host profile verify NSX configuration.

The screenshot shows the configuration for 'NSX Host vNIC : Segement_autodeploy'. The left sidebar lists various configuration categories, with 'NSX Host vNIC' expanded to show 'NSX Host vNIC : Seg...' selected. The main panel contains the following settings:

- Determine the LogicSwitch this virtual NIC should be connected to**
 - Choose a LogicSwitch connect to: *LogicSwitch Name (Segement_autodeploy)
 - Determine when the virtual NIC in LogicSwitch will be created: **Always create the object**
 - Stateless boot properties for virtual NIC in LogicSwitch: **Stateless boot config parameters (See doc before changing)**
 - Standby Uplinks used (See doc before changing):
 - *VLAN (See doc before changing): 0
 - Active Uplinks used (See doc before changing): vmnic1
 - *Teaming policy (See doc before changing): first uplink
- Determine how MAC address for vmknic should be decided**
 - Prompt the user for the MAC Address if no default is available: [dropdown]
 - VMkernel Network Adapter Name Policy: **Interface Name assigned**
 - VMkernel Network Adapter: vmk1
 - MTU policy: **Assign the specified MTU**
 - *MTU: 1500
 - TCP/IP stack: **Netstack Instance to which vmknic is connected**
 - *Name: defaultTcpipStack

3 To verify DVS switch is enabled on NSX, select **Policies and Profiles** → **Host Profiles** → **Configure** → **vSphere Distributed Switch**.

The screenshot shows the 'Policies and Profiles' configuration page for host profile 'c2306'. The 'Host Profiles' section is expanded to show 'DVS7' selected. The configuration for 'DVS7' is as follows:

- Determine whether NSX-T should be enabled on the DVS**
 - Specify NSX-T enabled on DVS: true
 - *Flag indicating if NSX-T should be enabled on DVS: true

4 Select the DVS switch and determine whether NSX is enabled on DVS.

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.

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

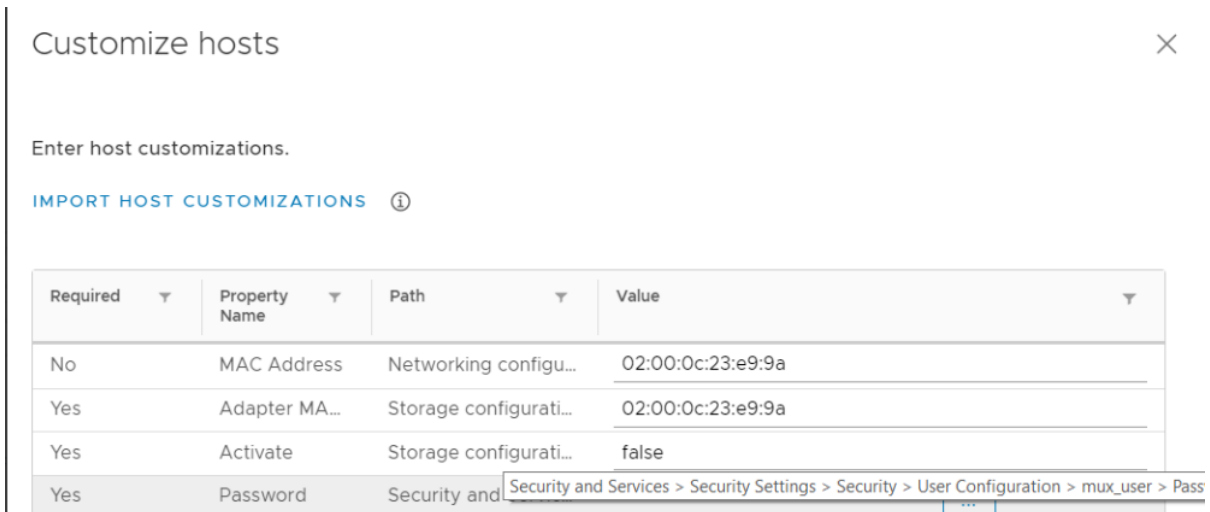
After the attaching the host profile to the target cluster, additional custom entries might be required on the host to successfully auto deploy the ESXi and NSX packages on it.

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 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 VIBs to be configured.

- Reboot hosts before applying TNP to the stateless cluster.
- Reboot hosts after applying TNP to the stateless 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 is not installed and configured on the target host.

Procedure

- ◆ 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 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 configuration is applied to the host.

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 switch.

What to do next

Apply TNP to the stateless cluster. See [Apply TNP on Stateless Cluster](#).

Apply TNP on Stateless Cluster

NSX 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 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.

3 Add a TNP by entering all required field.

Ensure that values of the following parameters are the same on both the new TNP profile and the existing Host Profile.

Note On a VDS switch, migration of VMkernel adapters and physical NIC migration is not supported.

- Transport Zone: Ensure transport zone referenced in Host Profile and TNP is the same.
- VDS Name: Ensure 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.
- Teaming Policy:
 - (On a VDS switch) In vCenter Server, when creating VDS uplinks, verify the NIC used in the Host Profile and map that physical NIC to the VDS uplink. In NSX-T, you map NSX uplinks to VDS uplinks. So, verify the configuration on the VDS switch in vCenter Server.

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.

- 4 Verify that the TNP profile is successfully created.
- 5 Apply TNP profile to the target cluster and click **Save**.
- 6 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.
- 7 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 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 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 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](#).

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
When multiple VMkernel adapters enabled to support Management, vMotion and other traffic are migrated to the same logical switch, VMkernel adapters get migrated to logical switch after reboot. But the service on one VMkernel adapter is enabled on a different adapter.	<p>For example, before migration, vmk0 is enabled to support Management traffic and vmk1 is enabled for vMotion traffic. After host reboot, vmk0 supports vMotion traffic and vmk1 supports Management traffic. This results in non-compliant error after reboot.</p> <p>Workaround: None. There is no impact as both VMkernel adapters are on the same logical switch.</p>
Host preparation progress is stuck at 60% while the node status displays <code>UP</code> .	<p>Issue: When a TNP is applied on a cluster, NSX is successfully installed on the host and node status displays <code>UP</code>, but GUI still shows 60% progress.</p> <p>Workaround: Reapply the TNP or TN configuration without any change in the config. This will fix the status to 100% on the GUI.</p>
Even though VMkernel migration is successful there was a validation error on the TN before host switches are removed.	<p>Issue: When you migrate vmk0 the management interface from vSwitch to a logical switch, NSX is successfully installed on the host. VMkernel migration is successful, but TN status shows Partial Success with error.</p> <pre>Validation before host switches removal failed: [error: No management vmk will have PNIC after ['vmk1'] in ['9a bb eb c1 04 81 40 e2-bc 3f 3e aa bd 14 62 1e'] lose all PNICs.]; LogicalSwitch full- sync: LogicalSwitch full-sync realization query skipped.</pre> <p>Workaround: None. Ignore the error message as VMkernel migration is successful.</p>
Reapplying a TNP where the Network Mapping for Install lists vmk0 results in host losing connectivity.	<p>Issue: When a TNP configuration consists of vmk0 in the Networking Mapping for Install, the hosts loses connectivity.</p> <p>Workaround: Instead of reapplying the TNP, reboot the host with necessary configurations in TNP.</p>
Cannot apply the host profile because MUX user password policy and password were not reset.	<p>Issue: Only on hosts running versions earlier than vSphere 6.7 U3. Host remediation and host profile application on hosts might fail unless the <code>mux_user</code> password is reset.</p> <p>Workaround: Under Policies & Profiles, edit the host profile to modify the <code>mux_user</code> password policy and reset the <code>mux_user</code> password.</p>

Scenario	Description
Host Profile is not portable.	<p>Issue: None of the vCenter servers can use the host profile containing NSX configuration.</p> <p>Workaround: None.</p>
Auto Deploy Rule Engine	<p>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.</p> <p>Workaround: Prepare each cluster from NSX GUI. See Apply TNP on Stateless Cluster.</p>
Check compliance errors.	<p>Issue: Host profile remediation cannot fix the compliance errors related to the NSX configuration.</p> <ul style="list-style-type: none"> ■ 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. <p>Workaround: Ensure the NSX configuration matches on Host Profile and TNP. Reboot the host to realize the configuration changes. The host comes up.</p>
Remediation	<p>Issue: If there are any NSX specific compliance errors, host profile remediation on that cluster is blocked.</p> <p>Incorrect configuration:</p> <ul style="list-style-type: none"> ■ Mapping between vNIC to LS mapping ■ Mapping of physical NICs <p>Workaround: Ensure that the NSX configuration matches on Host Profile and TNP. Reboot the host to realize the configuration changes. The host comes up.</p>
Attach	<p>Issue: In a cluster configured with NSX, host profile cannot be attached at the host-level.</p> <p>Workaround: None.</p>
Detach	<p>Issue: Detaching and attaching a new host profile in a cluster configured with NSX does not remove the NSX configuration. Even though the cluster is compliant with newly attach the host profile, it still has the NSX configuration from a previous profile.</p> <p>Workaround: None.</p>
Update	<p>Issue: If the user has changed NSX configuration in the cluster, then extract a new host profile. Update the host profile manually for all the settings that were lost.</p> <p>Workaround: None.</p>

Scenario	Description
Host-level transport node configuration	<p>Issue: After anportsport node was auto-deployed, it acts as individual entity. Any update to that transport node might not match with the TNP.</p> <p>Workaround: Update the cluster. Any update in a standalone transport node cannot persist its migration specification. The migration might fail to post the reboot.</p>
PeerDNS configuration is not supported on the VMkernel adapter selected for migration to the NVDS switch.	<p>Issue: If a VMkernel adapter selected for migration to NVDS is peer-DNS enabled, then host profile application fails.</p> <p>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.</p>
DHCP address of the VMkernel NIC address not retained	<p>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.</p> <p>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'.</p>
Host Profile application failure in vCenter can lead to NSX configuration errors on the host.	<p>Issue: If host profile application fails in vCenter, NSX configuration might also fail.</p> <p>Workaround: In vCenter, verify that host profile was successfully applied. Fix the errors and try again.</p>
LAGS are not supported on stateless ESXi hosts.	<p>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.</p> <p>Workaround: None.</p>
A stateless host does not boot up with MAC address of PXE NIC when it is applied with a host profile extracted from a stateful host.	<p>Issue: If a stateless host is attached with a host profile extracted from a statelful host, then the VMkernel adapter (vmknic) of the stateless host does not boot up with the MAC address of PXE NIC of the host because a stateful host does not boot up as a PXE-enabled system.</p> <p>Workaround: When you are setting up autodeployment of stateless hosts, ensure that the host profile extracted is from a from a host that boots up as a PXE-enabled system.</p>

Stateful Servers

Integrate host profiles of an ESXi host with NSX on stateful servers.

A stateful host is a host that retains all configurations and the installed VIBs even after it is rebooted. While an auto-deploy server is needed for stateless hosts because the boot up files required to bring up a stateless hosts are stored on the auto-deploy server, a stateful host does not need a similar infrastructure. Because the boot up files required to bring up a stateful host is stored on its hard drive.

In this procedure, the reference host is outside of the stateful cluster and the target hosts in the cluster. A target host can be within a cluster or a standalone host outside of the cluster. Prepare a cluster by applying host profile and transport node profile (TN profile) , so that any new target hosts joining the cluster is automatically prepared with NSX VIBs. Configure the target host as a transport node. Similarly, for a standalone host, apply the host profile and configure NSX to install NSX VIBs and when NSX configuration is complete, it becomes a transport node.

Note NSX VIBs are installed from TN profile and ESXi host configurations are applied by the Host Profiles.

Supported NSX and ESXi versions

Supported NSX and ESXi versions on stateful servers.

Version Name	67 ep 6	67 U 2	67 U3	67e p7	67 U2 C	6.5 U3	6.5 p0 3	7. 0. 1	7 0 P 3	7 0 U 2	7. 0. 1	6. 7 U1	6. 5 P 9	6. 5 P 27	6. 7P 8	6. 7E 24	7 0 P 6	7. 0 U 3 P 8	7 0 U 3c	8.0 .0.1
NSX 2.4	Yes	No	No	No	No	No	Yes	No	No	No	No	No	No	No	No	No	No	No	No	No
NSX 2.4.1	Yes	Yes	No	No	No	No	Yes	No	No	No	No	No	No	No	No	No	No	No	No	No
NSX 2.4.2	Yes	Yes	No	No	No	No	Yes	No	No	No	No	No	No	No	No	No	No	No	No	No
NSX 2.4.3	Yes	Yes	No	No	No	No	Yes	No	No	No	No	No	No	No	No	No	No	No	No	No
NSX 2.5	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No	No	No
NSX 2.5.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No	No	No
NSX 3.0	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
NSX 3.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
NSX 3.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No

Version Name	67ep6	67U2	67U3	67ep7	67U2C	6.5U3	6.5p03	7.0.01	7.0.032	7.0.01	6.7.0U1	6.5.0P9	6.5.0P27	6.7.0P8	6.7.0P24	7.0.0P6	7.0.0P8	7.0.0P3c	8.0.01
NSX 4.0	No	No	No	No	No	No	No	Yes	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes
NSX 4.1	No	No	No	No	No	No	No	Yes	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes

Prepare a Target Stateful Cluster

Prepare a target stateful cluster so that any new host joining the cluster is automatically deployed with ESXi and NSX VIBs.

You can select a host either within the cluster or outside of the cluster to be the reference host. You need to create a reference host because the host profile from the reference host is extracted and applied to a target host. VDS switch type supports migration of VMkernel adapters.

In this procedure, as an example the instructions are to migrate vmk0 (management traffic) and vmk1 (vMotion traffic) to an VDS switch.

Prerequisites

Procedure

- 1 On the Reference host, deploy a supported ESXi build.
 - a In vSphere, add vmk1 adapter. vmk0 is already present to serve management traffic.
- 2 Configure the reference node as a transport node.
 - a Using vSphere Web Client, ensure a logical switch is created in NSX.
 - b Using vSphere Web Client, ensure that vmk0 and vmk1 are connected to a logical switch on VDS switch.
- 3 Extract the host profile from the reference host.
- 4 On a target host that is a standalone host:
 - a Attach the host profile to the target host.
 - b Manually configure NSX on the host. When configuring the host as a transport node because the host profile on the ESXi, ensure the following conditions are met.
 - c Host must belong to the same transport zone.

- d Target host must use the same IP pool that is used by the reference host.
 - e Uplink profile, LLDP, Networkork mapping for install, VDS configured on the target host must be the same as configured on the reference host.
- 5 On a target host that is part of a cluster:
- a Attach the host profile to the stateful target cluster.
 - b Create and apply the TN profile on the cluster.
 - c To apply TN profile on the cluster.

What to do next

Scenarios when VMkernel adapters are migrated with and without host profiles applied to NSX.

Getting Started with NSX Federation

14

To get started with NSX Federation, you install the Global Manager, configure the Global Manager as active, and add locations.

Task	Details
Check the requirements for Federation.	See NSX Federation Requirements .
Install the Global Manager.	See Install the Active and Standby Global Manager .
Make the Global Manager cluster active.	See Make the Global Manager Active and Add Standby Global Manager .
Add Locations to the active Global Manager.	See Add a Location .

For further configuration tasks, such as preparing Edge clusters for stretched networking, and creating objects from the Global Manager, see *Federation* in the *NSX Administration Guide*.

Procedure

1 [NSX Federation Key Concepts](#)

NSX Federation introduces some new terms and concepts, such as remote tunnel endpoint (RTEP), span, and region.

2 [NSX Federation Requirements](#)

To support NSX Federation, your environment must meet various requirements, including round-trip time, software versions, and ports.

3 [Configuring the Global Manager and Local Managers](#)

An NSX Federation environment contains an active and a standby Global Manager cluster and one or more Local Manager clusters.

NSX Federation Key Concepts

NSX Federation introduces some new terms and concepts, such as remote tunnel endpoint (RTEP), span, and region.

NSX Federation Systems: Global Manager and Local Manager

An NSX Federation environment includes two types of management systems:

- Global Manager: A system similar to NSX Manager that federates multiple Local Managers.
- Local Manager: An NSX Manager system in charge of network and security services for a location.

NSX Federation Span: Local and Stretched

When you create a networking object from Global Manager, it can span one or more locations.

- Local: The object spans only one location.
- Stretched: The object spans more than one location.

You do not directly configure the span of a segment. A segment has the same span as the gateway it is attached to.

NSX Federation Regions

Security objects have a region. The region can be one of the following:

- Location: Each location automatically creates a region. This region has the span of that location.
- Global: A region that has the span of all available locations.
- Custom Region: You can create regions that include a subset of the available locations.

NSX Federation Tunnel Endpoints

In an NSX Federation environment, there are two types of tunnel endpoints.

- Tunnel End Point (TEP): The IP address of a transport node (Edge node or Host) used for Geneve encapsulation within a location.
- Remote Tunnel End Points (RTEP): The IP address of a transport node (Edge node only) used for Geneve encapsulation across locations.

NSX Federation Requirements

To support NSX Federation, your environment must meet various requirements, including round-trip time, software versions, and ports.

- There must be a maximum round-trip time of 500 ms between the following nodes:
 - Active Global Manager and standby Global Manager.
 - Global Manager and Local Manager.
 - Local Manager and remote Local Manager if you have cross-location security configuration only and VMware NSX® Edge™ Nodes RTEP and remote Edge Nodes RTEP if you have cross-network configurations.

- The Global Manager and Local Manager appliances must all have NSX 3.1.0 or later installed. All appliances in an NSX Federation environment must have the same version installed.
- The required ports must be open to allow communication between the Global Manager and Local Manager. See VMware Ports and Protocols at <https://ports.vmware.com/home/NSX>.
- There must be connectivity without NAT between the following nodes:
 - Global Manager and Local Manager.
 - Local Manager and remote Local Manager.
 - Edge node RTEP and remote Edge node RTEP.
- Verify that each location has a default overlay transport zone configured. From each Local Manager, select **System > Fabric > Transport Zones**. Select an overlay transport zone, and click **Actions > Set as Default Transport Zone**.
- Global Manager supports only Policy Mode. NSX Federation does not support Manager Mode. See [NSX Manager](#) for more information.

An NSX Federation environment has the following configuration maximums:

- For most configurations, the Local Manager cluster has the same configuration maximums as an NSX Manager cluster. Go to [VMware Configuration Maximums tool](#) and select NSX.
Select the NSX Federation category for NSX in the [VMware Configuration Maximums tool](#) for exceptions and other NSX Federation-specific values.
- For a given location, the following configurations contribute to the configuration maximum:
 - Objects that were created on the Local Manager.
 - Objects that were created on the Global Manager and include the location in its span.You can view the capacity and usage on each Local Manager. See *View the Usage and Capacity of Categories of Objects* in the *NSX Administration Guide*.

Configuring the Global Manager and Local Managers

An NSX Federation environment contains an active and a standby Global Manager cluster and one or more Local Manager clusters.

Active
Global Manager Cluster



Standby
Global Manager Cluster



Local Manager Cluster

Location 1



Local Manager Cluster

Location 2



Local Manager Cluster

Location 3

Install the Active and Standby Global Manager

To use NSX Federation, you must install the Global Manager.

Installing a Global Manager appliance is similar to installing an NSX Manager appliance. The only difference is that when you deploy the appliance, you select *NSX Global Manager* for the role.

Install a standby Global Manager appliance for high availability and disaster recovery. The standby Global Manager appliance must be installed in a different location with a latency of 500ms or less.

Prerequisites

- Verify that your environment meets the requirements for NSX Manager. See [NSX Manager VM and Host Transport Node System Requirements](#).

- Decide which locations will contain the active and standby Global Manager appliances.
- Verify that you are installing the Global Manager appliance with NSX 3.1.0 or later.

Important All Global Manager and Local Manager appliances in an NSX Federation environment must have the same version of NSX installed.

Procedure

- 1 To install the first Global Manager appliance on vSphere: [Install NSX Manager and Available Appliances](#).
 - Select `Medium` or `Large` for the deployment configuration size. Do not use `Small`.
 - Select `NSXGlobal Manager` for the **RoleName**.
- 2 Log in to the NSX Manager appliance.
See [Log In to the Newly Created NSX Manager](#) .
- 3 Configure a compute manager.
See [Add a Compute Manager](#).

Note If you are at this step while installing the standby Global Manager, you must configure a separate compute manager. Do not use the same compute manager that you configured for the active Global Manager.

- 4 Create a Global Manager cluster. See [Cluster Requirements for an Individual Site](#) for design recommendations.
 - On vSphere with a compute manager configured: See [Deploy NSX Manager Nodes to Form a Cluster from the UI](#).
 - On vSphere without a compute manager configured: Repeat the NSX Manager install on vSphere steps to install three appliances, then form the cluster. See [Form an NSX Manager Cluster Using the CLI](#).
- 5 Configure a VIP for the Global Manager cluster.
See [Configure a Virtual IP Address for a Cluster](#).
- 6 In a different location, install a standby Global Manager appliance and form a cluster by repeating these steps.

What to do next

Select the designated Global Manager appliance as active and connect it with the standby Global Manager.

Make the Global Manager Active and Add Standby Global Manager

After you have deployed a Global Manager appliance, you can make the Global Manager active.

Adding a standby Global Manager is optional but recommended for high availability of the Global Manager.

Procedure

- 1 Log in to the appliance at `https://global-manager-ip-or-fqdn/`.
- 2 Select **System > Location Manager**. In the **Global Manager** tile, click **Make Active**. Provide a descriptive name for the active Global Manager and click **Save**.
- 3 (Optional) Add a standby Global Manager cluster.
 - a Install a new Global Manager appliance in a secondary location and start it. Follow the same instructions as for installing the primary Global Manager, see [Install the Active and Standby Global Manager](#).
 - b From the active Global Manager, add this newly installed Global Manager appliance as standby.

Navigate back to your active Global Manager and click **Add Standby** and provide the following information:

Option	Description
Global Manager Name	Provide a name for the standby Global Manager.
FQDN/IP	Enter the FQDN or IP address of the Global Manager cluster VIP at the secondary location. Do not enter an individual Global Manager FQDN or IP.
Username and Password	Provide the admin user's credentials for the Global Manager at the secondary location.
SHA-256 Thumbprint	<p>Log in to any Global Manager node at the secondary location and run this command:</p> <pre>get certificate cluster thumbprint</pre> <p>The result is the cluster VIP certificate: bfae1a0a...</p> <p>If the GM-Standby is a single Manager VM, use the same command.</p>
Check Compatibility	Click Check Compatibility to ensure that the Global Manager can be added as standby. This checks that the NSX version is compatible.

- c Click **Save**.
- d Click **Make Standby**.

Add a Location

After you add a location to Global Manager, you can create objects from Global Manager that span that location.

You can find the number of supported locations in the [VMware Configuration Maximums tool](#). Select the appropriate version of NSX, select the NSX Federation category, and click **View Limits**.

Only use the admin account credentials to register the Local Manager with the Global Manager. After you add a location to the Global Manager, the NSX Manager is called a Local Manager (LM).

Prerequisites

- Verify that the NSX environment you are adding has the latest version of NSX installed.
This new NSX location can be a new NSX environment or an NSX environment with an existing Network and Security configuration.
- The NSX environment in the new location must have three NSX Manager nodes deployed and a cluster VIP configured. See [Configure a Virtual IP Address for a Cluster](#).
For a proof-of-concept environment, you can add a location that has only one NSX Manager node, but you must still configure a cluster VIP.
- Verify that the latency between the Global Manager and the location is 500 ms or less for non-stretched networks or 150 ms or less for stretched networks.

Procedure

- 1 Log in to the Global Manager at <https://global-manager-ip-or-fqdn/>.
- 2 Select **System > Location Manager** and click **Add On-Prem Location**.
- 3 In the **Add New Location** dialog box, enter the Location details.

Option	Description
Location Name	Provide a name for the location.
FQDN/IP	Enter the FQDN or IP address of the NSX Manager cluster VIP. Do not enter an individual NSX Manager FQDN or IP.
Username and Password	Provide the admin user's credentials for the NSX Manager at the location. Do not use any other account to register the Local Manager with the Global Manager.
SHA-256 Thumbprint	Log in to any NSX Manager node in the cluster and run this command: <pre>get certificate cluster thumbprint</pre> The result is the cluster VIP certificate: bfae1a0a...
Check Compatibility	Click Check Compatibility to ensure that the location can be added. This checks that the NSX version is compatible.

What to do next

If you want to create gateways and segments that span more than one location, you must configure a remote tunnel endpoint (RTEP) on Edge nodes in each location to handle the cross-

location traffic. See [Configure Edge Nodes for Stretched Networking](#). After you add a location to your Global Manager, you can import your configurations from that location's Local Manager appliance into the Global Manager. See [Importing Configurations from Local Manager](#).

Importing Configurations from Local Manager

After you successfully add a Local Manager location to the Global Manager, you can import all network and security Local Manager configurations to the Global Manager.

You can only import the entire Local Manager configuration into the Global Manager. There is no support for partial configuration import. You can only import the configurations once.

Local Manager Configurations Supported for Importing into Global Manager

- Context Profiles
- DHCP
- DNS
- Firewall Security Policies
- Gateway Profiles
- Groups
- NAT
- Security Profiles
- Services
- T0 Gateway
- T1 Gateway
- Time-based firewall (import/onboard now supported)

Local Manager Configurations Not Supported for Importing into Global Manager

The following features are not supported with NSX Federation. Import of configurations into the Global Manager is blocked if you have any of these configurations in your Local Manager. You must remove unsupported configurations to proceed with importing. After your supported Local Manager configurations are successfully imported into Global Manager, you can add the configurations for any of the unsupported features back into your Local Manager.

- DHCP dynamic binding
- Distributed IDS
- Distributed security for vCenter VDS Port Group only (Global Manager does not see the vCenter VDS port groups to assign them in security groups. However, Global Manager can use dynamic membership in groups based on vCenter VDS port groups tags added by Local Managers.)
- Endpoint protection

- Forwarding policies
- Guest introspection
- Identity firewall
- IDS/IPS
- L2 Bridge
- Load balancer
- Malware prevention
- Metadata proxy
- Multicast
- Network detection and response
- Network introspection
- Routing protocols (OSPF)
- Routing VPN and EVPN
- Service insertion
- TO VRF
- TLS inspection
- URL filtering

Note If you use a load balancer in the Local Manager, you cannot import the load balancer, but you can still import other configurations if you meet certain conditions as described in the section "Importing Configurations if you have a Load Balancer service on the Local Manager".

Importing Configurations if you have a Load Balancer service on the Local Manager

If your Local Manager has a load balancer service, you can import configurations except the load balancer, if you meet the following conditions:

- The load balancer service must be in one-arm mode on a standalone tier-1 gateway.
- The standalone tier-1 gateway that the one-arm load balancer is attached to:
 - must have only the load balancer service and no other services
 - must not have any downlink segments
 - must not share Gateway Firewall rules with any other tier-0 or tier-1 gateways.
- Groups used in load balancer service must not be used in any firewall rules. If you have groups common to both load balancer and firewall rules, you must either remove the group from the firewall rule or create an identical group to use with the load balancer.

Configurations Created by a Principal Identity User in Local Manager

If you have configurations in the Local Manager that are created by the Principal Identity user and the same Principal Identity user is not present in the Global Manager, import is blocked.

You have the following options for importing these entities:

- The system displays a list of Principal Identity usernames that are being used on the Local Manager to create configurations. Create each of these Principal Identity users in the Global Manager before proceeding to import.
- If you do not want to create Principal Identity usernames in Global Manager, remove all configurations in the Local Manager that are created using the Principal Identity username. You can then proceed with importing other configurations from the Local Manager.

Prerequisites


- The Local Manager appliance must register with the Global Manager.
- The Local Manager appliance must have a backup that you can restore in case the importing procedure fails.
- You must remove configurations for unsupported features from your Local Manager appliance. You are provided guidance in the NSX UI on how to resolve any importing conflicts.

Procedure

- 1 Log in to the Global Manager and navigate to **System > Location Manager**.
- 2 A system message appears for each location that has been successfully added into the Global Manager and has objects that can be imported.
- 3 Click **Import Now** from the system message. You can also import objects by clicking **Actions > Import** from the location tile.

- 4 You see a list of objects that can be imported into the Global Manager.
- a If there are naming conflicts, you can provide a prefix or suffix for configurations. The total length of the object including the prefix and suffix must not exceed 256 characters.
- The prefix or suffix gets applied to the following objects being imported:
- Tier-0 gateway
 - Tier-1 gateway
 - Segments
 - DNS zones
 - DHCP profiles
 - Switching profiles: IPv6, VNI Pool, Gateway QoS, BFD, IPFIX
 - Security profiles: IPFIX, Flood-Protection, DNS Security, Session Timer, Context Profiles
 - L4-L7 services (all services listed under **Inventory > Services**).
- The prefix or suffix does not get applied to the inventory and the firewall objects, that is: groups, firewall policies and firewall rules, and to system-created profiles and services.
- b For other conflicts, follow the guidance provided in the UI.

Results

Global Manager owns any Local Manager objects imported into the Global Manager. These objects appear in the Local Manager with this icon: . You can only modify these objects from the Global Manager now.

Configure Edge Nodes for Stretched Networking

If you want to create gateways and segments that span more than one location, you must configure a remote tunnel endpoint (RTEP) on Edge nodes in each location.

When you configure an RTEP, do it on an Edge cluster basis. All Edge nodes in the cluster must have an RTEP configured. You do not need to configure all Edge clusters with RTEP. RTEPs are required only if the Edge cluster is used to configure a gateway that spans more than one location.

You can configure the TEP and RTEP to use the same physical NIC on the Edge node or use separate physical NICs.

This procedure describes this task starting from your Local Manager. You can also configure RTEPs from your Global Manager by using the Location Manager site selection drop-down to choose the Local Manager.

Prerequisites

- Verify that each location participating in the stretched network has at least one Edge cluster.

- For RTEP networks, determine which layer 3 networks and VLANs to use.
 - Intra-location tunnel endpoints (TEP) and inter-location tunnel endpoints (RTEP) must use separate VLANs and layer 3 subnets.
- Verify that all RTEP networks used in a given NSX Federation environment have IP connectivity to each other.
- Verify that external firewalls allow cross-location RTEP tunnels. See VMware Ports and Protocols at <https://ports.vmware.com/home/NSX>.
- Configure the MTU for RTEP on each Local Manager. The default is 1500. Set the RTEP MTU to be as high as your physical network supports. On each Local Manager, select **System > Fabric > Settings**. Click **Edit** next to **Remote Tunnel Endpoint**.
- Optionally, if you do not use DHCP for RTEP, configure the RTEP IP pool for your site to configure the RTEPs for the Edge cluster. For details, go to "Add an IP Address Pool" in the *NSX Administration Guide*.

Procedure

- 1 From your browser, log in with admin privileges to the Local Manager at <https://<local-manager-ip-address>>.
- 2 To configure a new RTEP, select **System > Quick Start**.
- 3 Click **Configure Remote Tunnel Endpoint > Getting Started**.
- 4 You can select all Edge Nodes in this cluster or one node at a time. Provide the following details for the RTEP configuration:

Option	Description
Edge Switch	Select an edge switch from the drop-down menu.
Teaming Policy Name	(Optional) Select a teaming policy if you have one configured.
RTEP VLAN	Enter the VLAN ID for the RTEP network. Valid values are between 1 and 4094.
IP Assignment	Select an option from the drop-down. For example, select Use IP Pool and choose an option from the drop-down list.
IP Pool for all nodes	Select an IP pool for all nodes in this Edge Cluster. If you want to assign an IP address to an individual node, you can edit the RTEP configuration later.
Inter Location MTU	The default is 1500.

- 5 Click **Save**.
The green notification banner shows that all the Edge nodes in this edge cluster are configured successfully.
- 6 To add the RTEPs to the edge cluster for your other site locations, repeat steps 2 through 5.

- 7 To view or edit an existing Edge transport node:
 - a Select **System > Fabric > Nodes > Edge Transport Nodes**.
 - b Select an Edge node, then click **Tunnels**. If an RTEP is configured, it is displayed in the **Remote Tunnel Endpoint** section.
 - c Click **Edit** to modify the RTEP configuration.

The Configure Edge Nodes for Stretched Networking screen opens in the Local Manager with that Edge cluster selected.

Remove a Location

Removing a location from the Global Manager removes all objects created from the Global Manager that have a Global scope and are not specific to this location.

Removing a location is disruptive.

All configurations created by the Global Manager that are not specific to this location, such as groups and firewall sections with a Global scope, will be removed from the NSX Manager at this location.

Prerequisites

Before you remove a location, you must delete all objects created from the Global Manager that are specific to this location, such as tier-0 gateway, or firewall policies.

Procedure

- 1 Log in to the Global Manager at <https://global-manager-ip-or-fqdn/>.
- 2 Navigate to **System > Location Manager**.
- 3 From the location that you want to remove, click **Actions > Remove**.

You see a message describing the effect of removing a location. If you have not previously removed objects created from the Global Manager that are specific to this location, such as tier-0 gateways or firewall policies, you cannot remove the location. The system automatically removes all other configurations from the NSX Manager at this location, that have a Global scope in your NSX Federation deployment.

Remove a Location When the Global Manager is Unavailable

This procedure is used only if the Global Manager has been deleted **BEFORE** the Local Managers (LM) are aware of its deletion.

This is a possible use case at the completion of a Federation test. In this case, the LM will constantly try to connect to its configured Global Manager (which has been deleted), and share with other Local Managers its stretch GM group members. The API calls will remove the Global Manager constructs on each LM, even when the GM has been deleted.

The below API calls on the Local Manager, will remove the Local Manager, its registration to the Global Manager, and its registration to other Local Managers.

Note The following API calls done on the Local Manager are disruptive.

All configurations created by the Global Manager that are not specific to this location, such as groups and firewall sections with a Global scope, will be removed from the NSX Manager at this location.

Procedure

- 1 To remove a Global Manager that is active, a Global Manager that is standby, and other Local Managers registrations from the Local Manager use the site manager API (at the Local Manager) `POST https://<LM>/api/v1/sites?action=offboard_local..`
- 2 To remove Global Manager objects from a Local Manager, run the Local Manager API (at the Local Manager) `POST https://<LM>/policy/api/v1/infra/site?action=offboard.`
- 3 (Optional) The offboarding progress can be monitored by using the API call `GET https://<LM>/policy/api/v1/infra/site/offboarding-status.`

Results

Install NSX Advanced Load Balancer Appliance Cluster

15

The NSX Advanced Load Balancer is also known as VMware® Avi™ Load Balancer. From NSX Manager you can form an Avi controller cluster comprising of three NSX Advanced Load Balancer appliances. Objects, such as virtual services, profiles, pools and pool groups, that you later create in the Avi Load Balancer UI will need access to a management network. Use the controller cluster to provide these objects access to a management network.

The VMware NSX Advanced Load Balancer is a distributed and highly scalable cloud-native application distribution solution. Starting with NSX version 3.2, you can deploy and configure the NSX Advanced Load Balancer (AVI) using NSX Manager. The existing NSX Load balancer will be deprecated. The NSX Manager UI provides a single UI to install and manage all NSX components.

Important NSX Advanced Load Balancer(Avi Load Balancer) does not support deploying an appliance using an IPv6 address through the NSX Manager UI. However, you can deploy the controller appliance using an IPv6 address directly from a vCenter Server.

Prerequisites

- Supported Avi controller versions: 20.1.7, 21.1.2 or later versions
- Reserve four IP addresses (same subnet) in the management network to be assigned to the three controller appliances and one to the Virtual IP of NSX Advanced Load Balancer appliance cluster.
- Cluster VIP and all controllers management network must be in same subnet.
- Download the controller OVA from [Broadcom Support](#) page. To know more about downloading the controller OVA image, see <https://kb.vmware.com/s/article/82049>.

Procedure

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

Note You must log in with enterprise privileges. You cannot install **NSX Advanced Load Balancer** controller nodes with only load balancer privileges.

- 2 Select **System > Appliances > NSX Advanced Load Balancer**.

- 3 Click **Set Virtual IP** and enter the VIP for the cluster. It is mandatory to set a VIP for the cluster.

Note Verify that the virtual IP address you set is correct. If you set an incorrect cluster virtual IP address, then NSX Manager, API clients and end users cannot access the NSX Advanced Load Balancer controller. The only workaround is to delete all appliances and reconfigure the cluster with correct virtual IP address before proceeding with deployment.

- 4 Click **Save**.
- 5 Select the **Add NSX Advanced Load Balancer** card.
- 6 Choose the **Upload OVA File** or **Remote OVA Link** option.
- 7 Enter the URL and click **Upload**. Wait for the upload to finish.

Note

- If instead of the **Upload OVA File** option, the active bundle version is seen with a message to log in to the appliance, then directly log in to the appliance to upload the load balancer.
 - OVA upload can fail if the OVA file version being uploaded is different from the already deployed OVA files. For example, the second or third OVA deployment version is different from the first OVA deployment.
-

- 8 In the **Add Appliance** wizard, enter the deployment parameters for the first NSX Advanced Load Balancer appliance.
- 9 Click **Upload**.
- 10 On the Add Appliance window, configure these fields:

Field	Description
Hostname	Enter a valid hostname or FQDN (preferred) for the appliance. To enter a hostname that resolves to a FQDN, contact the DNS owner.
Management IP/Netmask	Enter a static IP address for the management IP address and netmask. For example, 192.168.1.2/22
Management Gateway	Enter a static IP address for the management gateway. The management gateway is used by NSX Advanced Load Balancer controller to communicate with NSX Manager and other NSX objects.
DNS Server	Enter the IP address of the DNS server.
NTP Server	Enter the IP address of the NTP server.
Node Size	Select the node size you want to deploy based on the requirements of your network. Supported node sizes are: <ul style="list-style-type: none"> ■ Small: 8 vCPU, 24 GB RAM, 128 GB storage ■ Medium: 16 vCPU, 32 GB RAM, 256 GB storage ■ Large: 24 vCPU, 48 GB RAM, 512 GB storage

- 11 Click **Next**.
- 12 On the Configuration window, configure these fields:

Field	Description
Compute Manager	Select a compute manager that registers the appliance.
Compute Cluster	Select a compute cluster where appliance will be deployed.
Resource Pool	(Optional) Select a resource pool that will be used during appliance deployment.
Host	Select a host where appliance will be deployed. Note Select either a host or a resource pool as storage location for deployment.
Datastore	Select a datastore that will be provide storage capacity for appliance.
Virtual Disk Format	By default, the Thin Provision format is selected. However, you can select a format that is feasible in your environment.
Network	Click Select Network to select the port group that will provide network connectivity to the appliance.

Note If incorrect compute manager details are provided, deployment fails. As a workaround, you must force delete the deployment and redeploy the appliance by providing the correct compute manager details.

- 13 Click **Next**.
- 14 On the Access & Credentials window, enter an admin password that complies with the required complexity.

Important Enter the same password when deploying all the controllers.

- 15 (Optional) In the **SSH Key** field, enter the private key of the SSH key pair to access controller using SSH key.
- 16 Click **Install Appliance**.

Do not try to delete the controller when NSX is registering the controller.

- 17 Follow steps 1-14 to deploy the second and third appliance.

Note Cluster formation only happens after the third appliance is deployed.

- 18 If clustering fails on the deployed controller nodes, the **NSX Advanced Load Balancer** displays an error message. Click **Start Clustering** to retrigger clustering of the deployed controller nodes. If clustering still fails, force delete the controller and reinstall it again.

NSX forms a cluster of the deployed controller nodes.

Results

NSX Advanced Load Balancer appliance controller cluster is deployed successfully and UI shows cluster status as `Stable`. Verify that the NSX Advanced Load Balancer controller cluster UI is accessible using its VIP, `https://<vip-fqdn>`.

What to do next

(Optional) Install portal certificate for ALB controllers.

- Run the following API to create a portal Certificate Signing Request (CSR) for ALB controller.

POST `/alb/controller-nodes/certificate/csr`

```
Payload:
{
  "common_name": "avi",
  "email": "xyz@vmware.com",
  "organization": "vm",
  "organization_unit": "VM",
  "locality": "BLR",
  "country": "IN",
  "state_name": "KA",
  "subject_alt_names": [
    "10.50.50.28"
  ],
  "algorithm": "SSL_KEY_ALGORITHM_RSA",
  "key_size": "SSL_KEY_2048_BITS"
}

Response:
{
  "name": "System-Portal-Cert-e8abab64",
  "csr": "-----BEGIN CERTIFICATE REQUEST ----- END CERTIFICATE REQUEST-----"
}
```

- (Optional) Run the following API to install and update portal certificate in ALB Controller.

POST `/alb/controller-nodes/certificate/install`

```
Payload:
{
  "name": "System-Portal-Cert-e8abab64",
  "cert": "-----BEGIN CERTIFICATE ----- END CERTIFICATE-----"
}

Response:
{
  "name": "System-Portal-Cert-14:58:30",
  "cert": "-----BEGIN CERTIFICATE ----- END CERTIFICATE-----"
}
```

After successfully deploying NSX Advanced Load Balancer appliance cluster, configure a NSX Cloud Connector in the AVI UI and then configure virtual services that will load balance traffic across servers.

For troubleshooting installation issues related to NSX Advanced Load Balancer appliance cluster, see [Troubleshooting NSX Advanced Load Balancer Controller Issues](#).

To know the best practices to install and run NSX Advanced Load Balancer, see the following link, <https://communities.vmware.com/t5/VMware-NSX-Documents/NSX-Advanced-Load-Balancer-by-Avi-Networks-NSX-T-Integration/ta-p/2890567>.

Note Starting with NSX 3.2.2, you can deploy NSX Advanced Load Balancer controller from the NSX Manager. However, you need to log in to the AVI controller to configure and consume load balancer services.

In upgraded environments such as 3.2.0 or 3.2.1 to 3.2.2 or higher with advanced load balancer activated, deactivate the NSX Advanced Load Balancer by clicking **Deactivate NSX-T ALB** in the banner message on the UI. For more details, see the *NSX Administration Guide*.

If the environment is running NSX Load Balancer, use NSX to Avi Migration Tool to migrate from NSX LB to NSX Advanced Load Balancer. See the AVI documentation.

Troubleshooting NSX Advanced Load Balancer Controller Issues

Troubleshoot issues when installing NSX Advanced Load Balancer controller.

NSX Advanced Load Balancer does not register with NSX Manager

Problem

Registration with compute manager failed.

Cause

If the status of deployment is 'VM registration in progress' and the appliance does not register even after 45 min, registration has failed.

Solution

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 To delete the controller, go to
System > Appliances > NSX Advanced Load Balancer
- 3 Click **Actions** and click **Delete**.
- 4 Re-deploy the appliance.

The Second NSX Advanced Load Balancer Controller Remains in Queued State

Problem

After initiating deployment of second controller, its status shows as `Controller VM Deployment Queued`

Cause

The second NSX Advanced Load Balancer controller remains in queued state till the third controller is deployed.

Solution

- 1 To retrieve controller deployments that are in pending state, run the following API command:

```
https://<NSX-Manager-IP-Address>/api/v1/alb/controller-nodes/deployments?
state=PENDING
```

Deployment of the second appliance begins only when you initiate deployment of the third appliance. Till then, the deployment status of the second appliance remains in queued state.

- 2 To retrieve status of deployments, run the following API command:

```
https://<NSX-Manager-IP-Address>/api/v1/alb/controller-nodes/deployments
```

- 3 To retrieve controller deployments that are in deployed state, run the following API command:

```
https://<NSX-Manager-IP-Address>/api/v1/alb/controller-nodes/deployments?
state=DEPLOYED
```

NSX Advanced Load Balancer Controller Password Change Caused Cluster Failure

Problem

NSX Advanced Load Balancer cluster failed because controller password was changed.

Cause

If you changed password of the leader controller outside of NSX, then NSX Advanced Load Balancer cluster goes into failed state. During the password change, if another node was made the leader, then the original leader node loses all its configured objects.

Solution

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Update the password in NSX.
- 3 On the failed controller nodes, click **Start Clustering**.

- 4 If clustering still fails, delete the controller. Go to
System > Appliances > NSX Advanced Load Balancer
- 5 Click **Actions** and click **Delete**.
- 6 Re-deploy the appliance.

Unable to Delete NSX Advanced Load Balancer Controller

Problem

Unable to delete NSX Advanced Load Balancer controller.

Cause

If NSX Advanced Load Balancer objects exist and there is only one node left, NSX does not allow you to delete all the deployed NSX Advanced Load Balancer controller nodes. This issue occurred because NSX cannot access the node or you manually deleted the node from compute manager.

Solution

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Delete all existing NSX Advanced Load Balancer objects.
- 3 To delete the controller, go to
System > Appliances > NSX Advanced Load Balancer
- 4 Click **Actions** and click **Delete** or **Force Delete**.
- 5 If load balancer objects are present, you cannot delete the controller node remaining in the cluster. To delete the last controller node, run the following API command.

Note Pass the `inaccessible` flag only if NSX cannot access the node and it is the last node in the cluster.

```
/policy/api/v1/alb/controller-nodes/deployments/{{node_id}}?
action=delete&force_delete=true&inaccessible
```

The API command deletes the controller node but the load balancer objects still exist in the system.

NSX Advanced Load Balancer Cluster HA Status is Compromised

Problem

NSX Advanced Load Balancer cluster status shows `Cluster Up HA Compromised`.

Cause

If you deleted one or two of the three NSX Advanced Load Balancer controllers from a stable cluster or if a controller is down, then the cluster status changes from `Cluster UP HA Active (Stable)` to `Cluster Up HA Compromised (Degraded)`.

Solution

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Verify all controllers are Up.
- 3 If you delete two nodes out of three nodes, SSH to the controller node with admin privileges and run `/opt/avi/scripts/recover_cluster.py`.
- 4 If the cluster still remains unstable, delete the controller and reinstall again.

Credential Mismatch After Changing NSX Advanced Load Balancer Controller Password

Problem

Credential mismatch after changing NSX Advanced Load Balancer controller password.

Cause

If you change controller password for admin user from outside of NSX from the Avi Vantage Platform UI, the new password does not refresh in NSX. Any change of password outside of NSX is not reflected in NSX Manager. NSX Manager does not reflect the state of the cluster.

Solution

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 If you changed the password outside of NSX, run the following API to change password.

```
PUT https://<NSX-Manager-IPaddress>/api/v1/alb/controller-nodes/
deployments/<node_id>
```

```
{
  "form_factor": "SMALL",
  "user_settings": {
    "admin_password": "Tilak@123456"
  },
  "deployment_config": {
    "vc_id": "755bd5cb-3700-456c-b74e-25f5140f4a50",
    "compute_id": "domain-c201",
    "host_id": null,
    "storage_id": "datastore-206",
    "management_network_id": "network-207",
    "hostname": "controller-AA",
```



```

"placement_type": "AlbControllerVsphereClusterNodeVmDeploymentConfig",
"disk_provisioning": "THIN",
"dns_servers": [
  "8.8.8.8"
]
}
}

```

- 3 If you changed the password from NSX, run the following API to change password, DNS and NTP servers in NSX and on the NSX Advanced Load Balancer controller.

PUT https://<NSX-Manager-IPAddress>/api/v1/alb/controller-nodes/deployments/<node-ID>?running_config=true

Deployment of NSX Advanced Load Balancer Controller Failed

Problem

Deployment of NSX Advanced Load Balancer controller failed.

Cause

Deployment might fail because of a number of reasons. Some of the reasons are:

- Deployment of controller failed
- Controller failed to power on
- Controller failed to register with compute manager
- Controller failed to power off
- Controller could not be undeployed
- Connection between NSX and compute manager is broken
- Compute manager registered to controller is deleted
- Controller is deployed with a hostname that already exist in the cluster

Solution

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 To delete the controller, go to
System > Appliances > NSX Advanced Load Balancer
- 3 Click **Actions** and click **Delete**.

Cluster Unstable After Two Controllers Are Down

Problem

A cluster becomes unstable after two NSX Advanced Load Balancer controllers go down.

Cause

When a three node cluster loses two controllers, the cluster quorum is lost. It becomes unstable.

Solution

- 1 Open a SSH session and log in to the controller that is up and running.
- 2 At the terminal, run the `/opt/avi/scripts/recover_cluster.py` script.
- 3 Verify the VIP of cluster is back up and the cluster status is stable.

Getting Started with NSX Cloud

16

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.

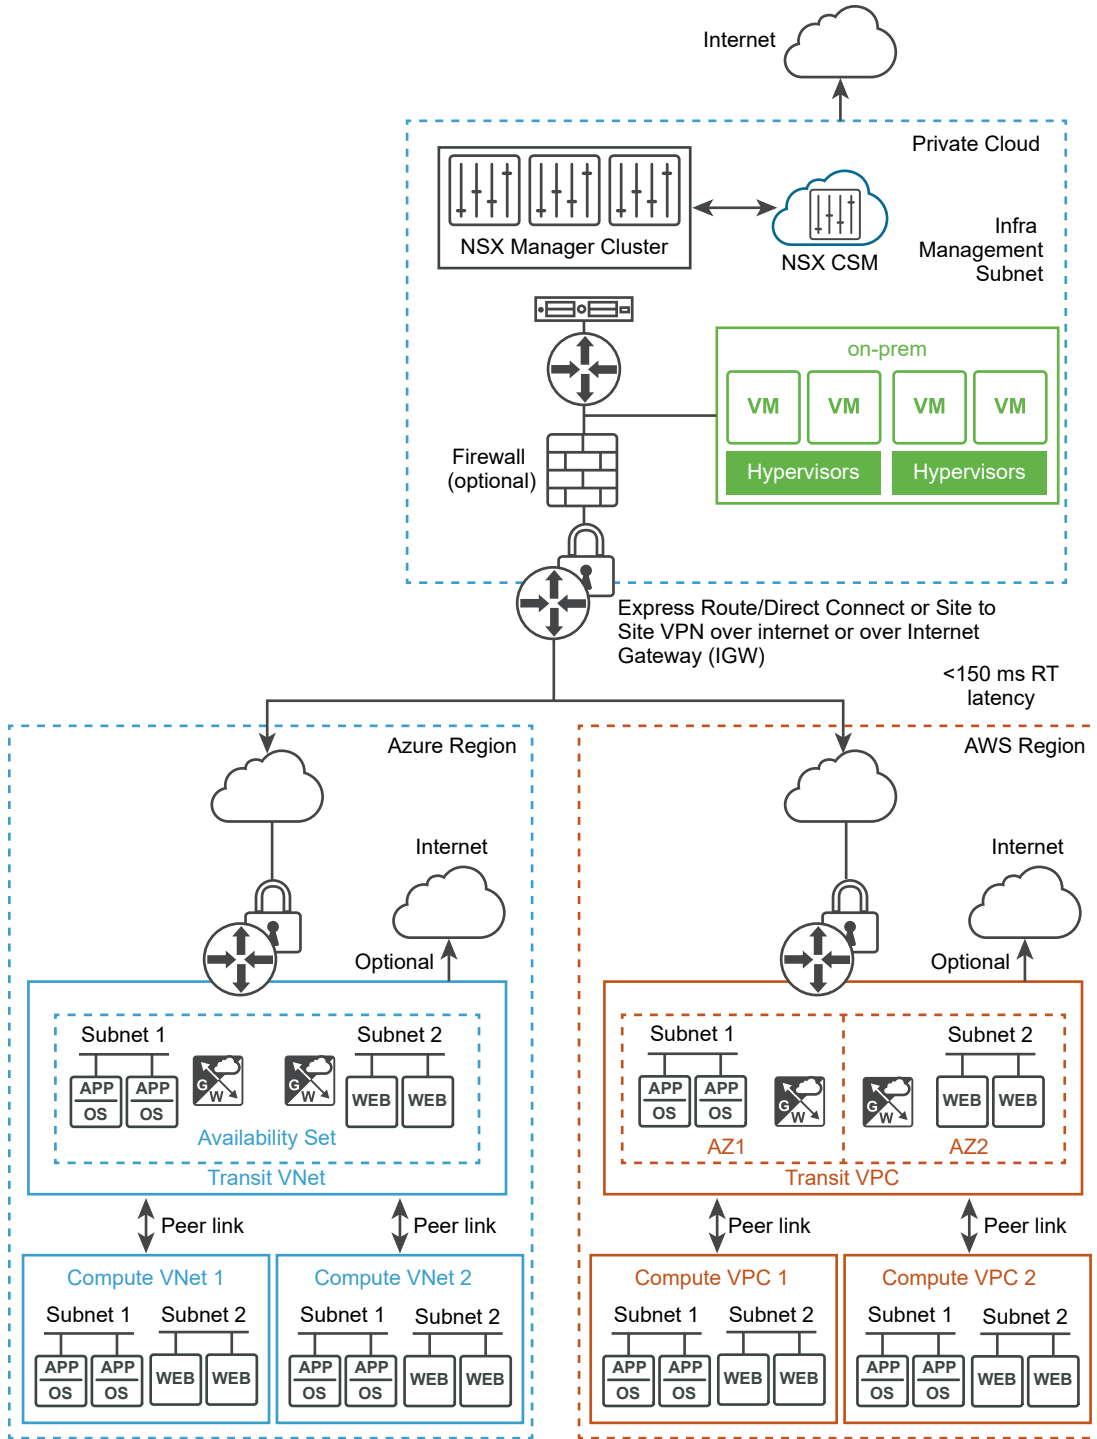
Read the following topics next:

- [NSX Cloud Architecture and Components](#)
- [Overview of Deploying NSX Cloud](#)
- [Deploy NSX On-Prem Components](#)
- [Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image](#)
- [Add your Public Cloud Account](#)
- [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#)
- [Deploy PCG or Link to a PCG](#)
- [Auto-Configurations after PCG Deployment or Linking](#)
- [Integrate Horizon Cloud Service with NSX Cloud](#)
- [\(Optional\) Install NSX Tools on your Workload VMs](#)
- [Un-deploy NSX Cloud](#)

NSX Cloud Architecture and Components

NSX Cloud integrates the NSX core components with your public cloud to provide network and security across your implementations.

Figure 16-1. 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.
- **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.
- **NSX Tools** functionality that provides NSX-managed datapath for workload VMs.

Overview of Deploying NSX Cloud

Refer to this overview to understand the overall process of installing and configuring NSX Cloud components to enable NSX to manage your public cloud workload VMs.

Note While planning your deployment, ensure that NSX appliances have good connectivity with the PCG deployed in the public cloud and Transit VPCs/VNets are in the same region as the Compute VPCs/VNets.

Table 16-1. Workflow for deploying NSX Cloud

Task	Instructions
<input type="checkbox"/> Install CSM and connect with NSX Manager. Note Starting in NSX 3.1.1, if you are using Microsoft Azure, you can deploy NSX Cloud components in your Microsoft Azure subscription. See Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image .	See Deploy NSX On-Prem Components .
<input type="checkbox"/> Add one or more of your public cloud accounts in CSM.	See Add your Public Cloud Account .
<input type="checkbox"/> Deploy PCG in your Transit VPCs or VNets and link to your Compute VPCs or VNets.	See NSX Public Cloud Gateway: Architecture and Modes of Deployment .
What to do next?	Follow instructions at "Using NSX Cloud" in the <i>NSX Administration Guide</i> .

Deploy NSX On-Prem Components

You must have already installed NSX Manager to proceed with installing CSM.

Install CSM

The Cloud Service Manager (CSM) is a core component of NSX Cloud.

After installing NSX Manager, install CSM by following the same steps as for installing NSX Manager and selecting **nsx-cloud-service-manager** as the VM role. See [Install NSX Manager and Available Appliances](#) for instructions.

You can deploy CSM in the Extra Small VM size or higher, as required. See [NSX Manager VM and Host Transport Node System Requirements](#) for details.

Ports and Protocols

For a list of ports and protocols required for inbound and outbound access for CSM, see <https://ports.esp.vmware.com/home/NSX>.

NTP Server Configuration

Many features require that CSM has a valid NTP server entry. You can configure the NTP server at the time of installing CSM or later. Also see *Configuring NTP on Appliances and Transport Nodes* in the *NSX Administration Guide* for other options for configuring NTP.

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 user name 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**.
You can also go to **System > Settings**, and click **Configure** on the **Associated NSX Node** tile.
- 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	Enter the NSX Manager's thumbprint value. For details, see the Obtain Thumbprint of NSX Manager topic in the <i>NSX Installation Guide</i> .

Important Thumbprint value is mandatory to proceed further. Not providing the thumbprint value for the NSX Manager will result in errors.

4 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.

5 (Optional) Set up the Proxy server. See instructions in [\(Optional\) Configure Proxy Servers](#).

Specify CSM IPs for Access by PCG

After CSM is deployed, run the following API to use an IP/subnet pool for CSM visible to PCG.

Whenever you run this API, CSM updates `gw-mgmt-sg` associated with the PCG in your public cloud, to append these IP addresses to allow inbound communication on PCG from CSM over these IP addresses or IP address ranges. See [Auto-created Public Cloud Configurations](#) for a list of constructs created in the public cloud after PCG is deployed.

```
PUT https://<csm-ip>/api/v1/csm/configs/system-config
```

Example Request Body where `10.1.1.1/24` is the IP address of CSM as seen by PCG.

```
{
  "mgmt_ip_config": [
    "10.1.1.1/24",
    "192.168.0.0/24"
  ]
}
```

(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 > General Settings > Internet Proxy Server**.

Note You can also provide these details when using the CSM Setup Wizard that is available when you first install CSM.

- 2 On the Internet Proxy Server 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 checkbox 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 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 Administration Guide* for detailed instructions.
- 3 Log in to NSX Manager. You are redirected to the vIDM login.
- 4 Enter the vIDM user's credentials. Once you log in, you can switch between NSX Manager and



CSM by clicking the Applications icon.

Obtain Thumbprint of NSX Manager

While configuring CSM, you must enter thumbprint value of NSX Manager.

Prerequisites

You can get thumbprint value in two ways; NSX Manager's User Interface (UI) or CLI.

Procedure

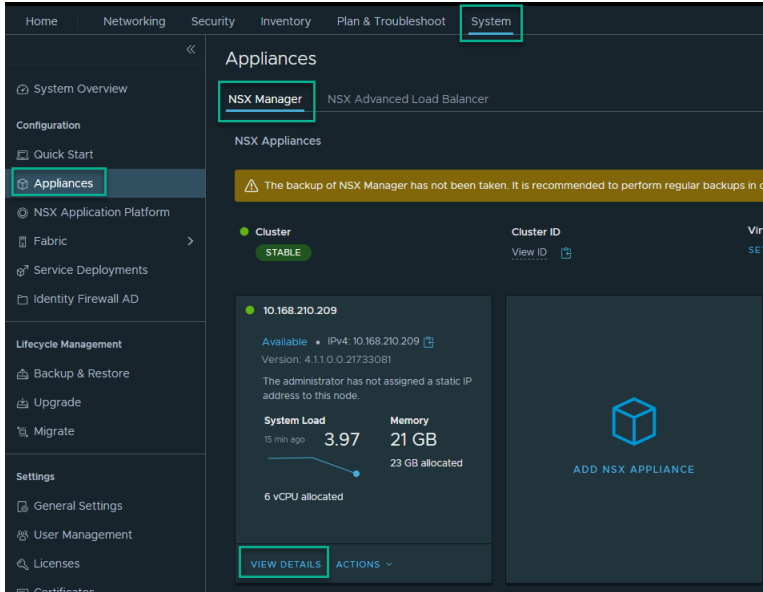
- 1 Obtain the thumbprint value from the NSX Manager CLI as follows:
 - a Open an SSH session or console session to one of the NSX Manager appliances.
 - b To retrieve the thumbprint of the NSX Manager appliance, at the NSX Manager appliance console, run the `get certificate api thumbprint` command. The command output is a string of alphanumeric numbers that is unique to this NSX Manager. For example:

```
NSX-Manager1> get certificate api thumbprint
659442c1435350edbbc0e87ed5a6980d892b9118f851c17a13ec76a8b985f57
```

Alternatively, to retrieve the thumbprint of the cluster, at the NSX Manager appliance console, run the `get certificate cluster thumbprint` command. The output is a string of alphanumeric numbers that is unique to the cluster.

See the *NSX Command-Line Interface Reference* for details on CLI commands.

- 2 Obtain the thumbprint value from the NSX Manager UI as follows:
 - a From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
 - b Navigate to **System > Appliances** and in the **NSX Manager** tab, click **View Details**.



- c Under the **Appliance details** section, use the copy icon next to **Cert Thumbprint** to copy the thumbprint value.



Results

The thumbprint value of the NSX Manager is available.

Connect your Public Cloud with On-prem NSX

Connect your public cloud account with the on-prem deployment of NSX,

Connect VPCs/VNets with on-prem using suitable methods, such as Direct Connect for AWS or Express Route for Microsoft Azure. You can also use site-to-site VPN with any VPN endpoint on-prem and PCG acting as the VPN endpoint in your public cloud.

To use the Transit/Compute topology, you must have peering connections established between the Transit and Compute VPCs/VNets. You can have a single PCG manage multiple compute VPCs/VNets. You can also have a flat compute VPC/VNet architecture with a PCG pair installed in each VPC/VNet. See [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#) for details on PCG deployment options.

Connect Microsoft Azure with On-prem NSX

A connection must be established between your Microsoft Azure network and your on-prem NSX 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.
- 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

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. Refer to 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 Microsoft 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 [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#) for details on how these subnets are used.

Connect AWS with On-prem NSX

A connection must be established between your Amazon Web Services (AWS) network and your on-prem NSX 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 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. Refer to AWS Reference Documentation for details.

Note You must review and implement the applicable security considerations and best practices by AWS; refer to AWS Security Best Practices for details.

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 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 [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#) 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.


Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image

Starting in version 3.1.1, you can use the Terraform scripts to deploy management components – NSX Manager and CSM – within your Microsoft Azure subscription.

provides Terraform scripts that deploy components in your Microsoft subscription. See [Deploy NSX Cloud Components in Microsoft Azure using Terraform scripts](#).

If you cannot use Terraform scripts to deploy components, you can also manually deploy them in your Microsoft Azure subscription. See [Deploy NSX Cloud Components in Microsoft Azure without using Terraform scripts](#).

Figure 16-2. NSX Cloud Components Deployed in Microsoft Azure



Accepting Azure Marketplace Terms

As the NSX Cloud images are published at the Azure Marketplace, you must first accept the legal terms for the images.

To accept the terms, you can use the PowerShell command. For details, see <https://docs.microsoft.com/en-us/powershell/module/az.marketplaceordering/set-azmarketplaceterms?view=azps-6.2.0>.

Use the following values for the PowerShell command. Note that the product value is based on the marketplace image.

Parameter	Description
Publisher	vmware-inc
Name	byol_release-3-1
Product For Public Cloud Gateway (PCG)	nsx-public-cloud-gateway

Parameter	Description
Product For Management Plane (MP)	nsx-policy-manager
Product For Cloud Service Manager (CSM)	nsx-cloud-service-manager

Following is an example for accepting terms for CSM using the PowerShell command:

```
Set-AzMarketplaceTerms -Publisher "vmware-inc" -Product "nsx-cloud-service-manager" -Name
"byol_release-3-1" -Terms $agreementTerms -Accept
```

Deploy NSX Cloud Components in Microsoft Azure using Terraform scripts

Follow these steps to deploy NSX Cloud using the NSX Cloud Marketplace image in Microsoft Azure using the Terraform scripts provided by NSX Cloud.

Prerequisites

- Verify that you have access to the NSX Cloud Marketplace image in your Microsoft subscription.
- Verify that you have accepted Microsoft Azure's Marketplace legal terms in the subscription where you are deploying NSX cloud appliances.
- You must have Microsoft Azure CLI installed and configured on the system. This is required for authenticating and running Azure APIs that are used in the Terraform scripts.

If possible, use the same system to run the Terraform scripts that you use to access your Microsoft subscription from. This ensure that your Microsoft Azure credentials can be used from within the system and you do not have share this information with a different system.

Also, as a security recommendation, run these scripts on a Linux/Unix or macOS system that supports the Python crypt module.

- Verify that you have binaries of Terraform 0.13 or higher on the system where you plan to run the Terraform scripts.
- You must have Python 3.0 or higher installed on this system.

Procedure

- 1 Download the Terraform scripts by logging in to your account and navigating to: **NSX > Drivers and Tools > VMware NSX Terraform Provider > Go To Downloads > Download Now**. For example, after you log in to your account, go to the [Download page for Drivers and Tools](#).
- 2 Extract the contents of the file named `NSXCloudScriptsforAddingPublicCloudAccounts.tar.gz`. The Terraform scripts and related files are in the folder `NSXCloudScripts/cloud-native-deployment/azure/igw`.

3 Update the Terraform configuration files.

- a In `config.auto.vars`, add the following information:

Parameter	Description
<code>subscription_id</code>	Provide the subscription ID for your Microsoft Azure account.
<code>location</code>	Specify the Microsoft Azure location that the NSX Cloud Management VNet will be deployed in.
<code>deployment_prefix</code>	This is the deployment name that will be prefixed to all auto-created entities. Ensure that this is unique for each Microsoft <code>subscription_id</code> and <code>location</code> .

- b In `credentials_nsx.auto.tfvars`, add the following information:

Parameter	Description
<code>mgr_public_key_path</code>	This is the path to the public key to be applied to the NSX Manager appliance.
<code>csm_public_key_path</code>	This is the path to the public key to be applied to the CSM appliance.
<code>license_key</code>	This is the license key for NSX Manager. You must have the NSX Enterprise Plus license.

- c Verify advanced configuration information, and update as necessary, in the file `advanced_config.auto.tfvars`:

Parameter	Description
<code>mgmt_vnet_address_space</code>	This is the address space for the newly deployed NSX Cloud Management VNet.
<code>mgmt_subnet_address_prefix</code>	This is the subnet for the NSX Cloud management appliances deployed in the NSX Cloud Management VNet.

4 Run the following commands in the specified order:

<code>~/terraform init</code>	This command collects all the modules required for deployment.
<code>~/terraform plan</code>	This command displays the list of steps or a blueprint of the procedure involved in the deployment.
<code>~/terraform apply</code>	This command executes the script. If something goes wrong during execution, you are shown the corresponding error messages. After you fix the errors, you can resume the deployment from where it stopped.

5 Follow these steps to change the passwords generated for NSX Manager and CSM by the Terraform scripts.

- a After the scripts run successfully, make a note of the following passwords for NSX Manager and CSM:

- admin_password
- root_password

These passwords are displayed on the screen at the end of the deployment. You can also find these passwords in the file `NSXCloudScripts/cloud-native-deployment/azure/igw/terraform.tfstate`, under the section "outputs", for example:

```
"outputs": {
  "csm": {
    "value": {
      "admin_password": "<pwd>",
      "admin_username": "nsxadmin",
      "private_ip": "<private IP>",
      "public_ip": "<public IP>",
      "root_password": "<pwd>"
    },
    "mgrs": {
      "value": [
        {
          "admin_password": "<pwd>",
          "admin_username": "nsxadmin",
          "private_ip": "<private IP>",
          "public_ip": "<public IP>",
          "root_password": "<pwd>"
        }
      ],
    },
  },
}
```

- b In Microsoft Azure, navigate to the Network Security Groups created for NSX Manager and CSM, named `<deployment_prefix>-nsx-mgr-sg` and `<deployment_prefix>-nsx-csm-sg`, and add the following temporary inbound "allow" rule for SSH:

Priority	Name	Port	Protocol	Source	Destination	Action
1010	AllowInboundRuleSSH	22	TCP	Any	Any	Allow

- c Log in to the NSX Manager appliance using your private key and change the password generated by the Terraform scripts:

```
$ ssh -i <nsx_mgr_key> nsxadmin@<NSX Manager public IP address>
WARNING: Your password has expired.
You must change your password now and login again!
Changing password for nsxadmin.
(current) UNIX password: <Enter mgr_admin_pwd from the Terraform scripts>
New password: <Enter new password conforming to NSX password complexity>
Retype new password:
passwd: password updated successfully
```

- d Log in to CSM using your private key and change the password generated by Terraform scripts:

```
$ ssh -i <nsx_csm_key> nsxadmin@<CSM public IP address>
WARNING: Your password has expired.
You must change your password now and login again!
Changing password for nsxadmin.
(current) UNIX password: <Enter csm_admin_pwd from the Terraform scripts>
New password: <Enter new password conforming to NSX password complexity>
Retype new password:
passwd: password updated successfully
```

- 6 Log in to the CSM appliance using the new password you have set and run the following NSX CLI command to join CSM with the NSX Manager cluster:

```
join <nsx-manager-ip-address & port(optional)> cluster-id <nsx-manager-cluster-id>
username <username> password <password> thumbprint <nsx-manager-api-thumbprint>
```

It takes a few minutes for CSM UI to appear. Run the `get cluster status` command on the CSM appliance CLI. If the status is `stable`, then continue to the next step.

You can run the NSX CLI command `get cluster status` from any NSX Manager node to get the `cluster-id`. You can get the NSX Manager thumbprint by running the `get certificate api thumbprint` command on the specified NSX Manager. See the *NSX Command-Line Interface Reference* for details on CLI commands and [Obtain Thumbprint of NSX Manager](#) .

Note If the NSX Manager node that you joined the CSM appliance to is lost, you can either run this NSX CLI command to join CSM with one of the other healthy NSX Manager nodes, or you can redeploy the lost NSX Manager node using its image file named, `<deployment_prefix>nsx-mgr-image` and CSM will automatically rejoin this node when this node is back online. See *Redeploying NSX Manager from nsx_mgr_image in Microsoft Azure* in the *NSX Administration Guide* for details.

- 7 To connect CSM with NSX Manager, add details in the NSX Manager Credentials screen as described in [Join CSM with NSX Manager](#) .

Results

The scripts deploy the following in your Microsoft Azure subscription:

- A VNet to host the NSX Cloud management appliances. This VNet is named `<deployment_prefix>-nsx-mgmt-vnet`.
- An Availability Set in which the three nodes of the NSX Manager cluster are deployed. This Availability Set is named `<deployment_prefix>-nsx-aset`.
- Microsoft Azure Resource Group named `<deployment_prefix>nsx-mgmt-rg`.
- The following resources for the each of the NSX Manager nodes and for the CSM appliance:
 - a VMs named `<deployment_prefix>nsx-csm` for CSM, and `<deployment_prefix>nsx-mgr0`, `<deployment_prefix>nsx-mgr1` and `<deployment_prefix>nsx-mgr2` for the NSX Manager cluster.
 - b OS Disk for each VM.
 - c Network interface (NIC) for each VM.
 - d Public IP address for each VM.
 - e Data disk for each VM.
- Network Security Groups for NSX Cloud management components that allow connectivity for these appliances.
 - `<deployment_prefix>-nsx-mgr-sg`:

Table 16-2. Inbound Rules for NSX Manager deployed using the Terraform scripts

Priority	Name	Port	Protocol	Source	Destination	Action
1000	AllowInboundRuleAPI	443	TCP	Any	Any	Allow

Table 16-3. Outbound Rules for NSX Manager deployed using the Terraform scripts

Priority	Name	Port	Protocol	Source	Destination	Action
100	AllowOutboundRuleAPI	Any	TCP	Any	Any	Allow

- `<deployment_prefix>-nsx-csm-sg`:

Table 16-4. Inbound Rules for CSM deployed using the Terraform scripts

Priority	Name	Port	Protocol	Source	Destination	Action
1000	AllowInboundRuleAPI	443	TCP	Any	Any	Allow

Table 16-5. Outbound Rules for CSM deployed using the Terraform scripts

Priority	Name	Port	Protocol	Source	Destination	Action
100	AllowOutboundRuleAPI	80,443	TCP	Any	Any	Allow

Note Consider updating the `Source` field of these auto-created network security groups to a restricted set of CIDRs from which you want to access NSX Manager and CSM. The default `Any` is not safe.

- A Microsoft Azure Recovery Service Vault with a vault policy to perform a recurring backup of all three NSX Manager nodes and the CSM appliance. The vault policy is named `<deployment_prefix>-nsx-vault` and the default backup schedule is set to: daily recurring at 11PM UTC.

See *Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure* in the *NSX Administration Guide* for details on restore options.

What to do next

[Deploy PCG in a VNet](#)

Deploy NSX Cloud Components in Microsoft Azure without using Terraform scripts

Follow these steps to manually deploy NSX Cloud components in Microsoft Azure using the Microsoft Azure marketplace image, without using Terraform scripts provided by NSX Cloud.

The following steps are performed in your Microsoft Azure subscription:

- 1 Create a resource group for NSX Cloud management resources with a descriptive name, for example, `nsx-mgmt-rg`.
- 2 In this resource group, create an availability set in which you will deploy three NSX Manager nodes.
- 3 In this resource group, create a VNet where you will deploy NSX Cloud management components.
- 4 In this VNet, create a subnet for NSX Cloud management components.

5 Create Security groups for NSX Manager and CSM appliances.

- Security groups for NSX Manager named like **nsx-mgr-sg**:

Table 16-6. Inbound Rules for NSX Manager

Priority	Name	Port	Protocol	Source	Destination	Action
1000	AllowInboundRuleAPI	443	TCP	Any	Any	Allow

Table 16-7. Outbound Rules for NSX Manager

Priority	Name	Port	Protocol	Source	Destination	Action
100	AllowOutboundRuleAPI	Any	TCP	Any	Any	Allow

- Security groups for CSM named like **nsx-csm-sg**:

Table 16-8. Inbound Rules for CSM

Priority	Name	Port	Protocol	Source	Destination	Action
1000	AllowInboundRuleAPI	443	TCP	Any	Any	Allow

Table 16-9. Outbound Rules for CSM

Priority	Name	Port	Protocol	Source	Destination	Action
100	AllowOutboundRuleAPI	80,443	TCP	Any	Any	Allow

6 Deploy one CSM VM using the CSM marketplace image URN with a public IP address. Use the following configurations as specified. For all other configurations you can select the default values or the best options for your requirements.

Parameter	Value
	Basic
Virtual machine name	Any descriptive name.
Size	The minimum requirement is: Standard_D4s_v3-4vcpus, 16 GB memory.
Authentication type	SSH
Username	Enter the default NSX Manager username: <code>nsxadmin</code> .
SSH Public Key Source	Provide the public key of the SSH key-pair you will use to log in to the appliance over SSH.
	Networking
Public IP	Click Create new and select Static for the Assignment option.

Parameter	Value
NIC network security group	Select Advanced
Configure network security group	Select the network security group created for CSM, for example, nsx-csm-sg as described in an earlier step.
	Advanced
Custom data	<p>Copy-paste the following, ensuring that you use your deployment's username and password:</p> <pre>#cloud-config hostname: <hostname> chpasswd: expire: false list: - nsxadmin:<admin_password> - root:<root_password></pre> <p>For example:</p> <pre>#cloud-config hostname: nsx-datacenter1-csm chpasswd: expire: false list: - nsxadmin:MySecretNsxAdminPassword - root:MySecretNsxRootPassword</pre>

- 7 Deploy three NSX Manager VMs using the NSX Manager marketplace image URN with a public IP address. Use the following configurations as specified. For all other configurations you can select the default values or the best options for your requirements.

Parameter	Value
	Basic
Virtual machine name	Any descriptive name.
Size	The minimum requirement is: Standard_D4s_v3-4vcpus, 16 GB memory.
Authentication type	SSH
Username	Enter the default NSX Manager username: <code>nsxadmin</code> .
SSH Public Key Source	Provide the public key of the SSH key-pair you will use to log in to the appliance over SSH.
	Disks
OS Disk type	Standard HDD
Data disks	<p>Click Create and attach a new disk and select Standard HDD. for Disk SKU with a custom size of 100 GiB.</p> <p>Note Ensure that the data disk host caching is set to read/write.</p>
	Networking

Parameter	Value
Public IP	Click Create new and select Static for the Assignment option.
NIC network security group	Select Advanced
Configure network security group	Select the network security group created in a previous step, from the example in this topic: <code>nsx-mgr-sg</code>
Advanced	
Custom data	Copy-paste the following, ensuring that you use your deployment's username and password: <pre>#cloud-config hostname: <hostname> bootcmd: - [cloud-init-per, instance, lvmdiskscan, lvmdiskscan] - [cloud-init-per, instance, secondary_partition, /opt/vmware/nsx-node-api/bin/set_secondary_partition.sh] chpasswd: expire: false list: - nsxadmin:<admin_password> - root:<root_password></pre>

- Configure a Microsoft Azure Recovery Service Vault with a vault policy to perform a recurring backup of all three NSX Manager nodes and the CSM appliance. For example, you could use this policy named **nsx-vault** and the default backup schedule set to daily recurring at 11PM UTC.

See *Managing Backup and Restore of NSX Manager and CSM in Microsoft Azure* in the *NSX Administration Guide* for details on restore options.

- Add a temporary network security group to allow SSH access for NSX Manager and CSM.

Table 16-10. Temporary rule for both NSX Manager and CSM to allow SSH access

Priority	Name	Port	Protocol	Source	Destination	Action
1010	AllowInboundRuleSSH	22	TCP	Any	Any	Allow

- Log in to the NSX Manager and CSM appliances using your private key and the passwords you provided in user data when launching the VMs.
- Create an NSX Manager cluster with the three NSX Manager nodes deployed. See [Form an NSX Manager Cluster Using the CLI](#).
- Add an NSX license:
 - From your browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.
 - Select **System > Licenses > Add License**.
 - Enter a license key. You must have the NSX Enterprise Plus license.

- 13 Log in to the CSM appliance and run the following NSX CLI command to join CSM with the NSX Manager cluster:

```
join <nsx-manager-ip-address & port(optional)> cluster-id <nsx-manager-ip-address>
username <username> password <password> thumbprint <nsx-manager-api-thumbprint>
```

It takes a few minutes for CSM UI to appear. Run the `get cluster status` command on the CSM appliance CLI. If the status is `stable`, then continue to the next step.

You can run the NSX CLI command `get cluster status` from any NSX Manager node to get the `cluster-id`. You can get the NSX Manager thumbprint by running the `get certificate api thumbprint` command on the specified NSX Manager. See the *NSX Command-Line Interface Reference* for details on CLI commands and [Obtain Thumbprint of NSX Manager](#).

- 14 To connect CSM with NSX Manager, add details in the NSX Manager Credentials screen as described in [Join CSM with NSX Manager](#).

Add your Public Cloud Account

Create roles with appropriate permissions in your public cloud account to add the account into CSM.

Overview

For each public cloud account that you want to bring under the control of NSX, you have the option of creating appropriate roles with appropriate permissions. NSX Cloud provides scripts that you can use to generate these roles.

If you want to restrict public clouds that can be added into CSM, run the following CSM API:

```
PUT /api/v1/csm/desired-clouds
```

Example Request:

```
PUT https://<nsx-csm>/api/v1/csm/desired-clouds
```

```
{
  "cloud_types": [
    {
      "cloud_type": "aws"
      "enabled": true,
    }
    {
      "cloud_type": "azure",
      "enabled": true,
    }
    {
      "cloud_type": "aws-gov-us-east"
      "enabled": false,
    }
    {
      "cloud_type": "aws-gov-us-west",
```



```

    "enabled": false,
  }
  {
    "cloud_type": "azure-gov-us",
    "enabled": false,
  }
]
}

```

See the latest version of the *NSX API Guide* at <https://code.vmware.com/> for API details.

Adding your Microsoft Azure Subscription

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.

Overview:

- 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. If you have the new Azure Powershell Az module, you must run the `Enable-AzureRmAlias` command to ensure that the AzureRM cmdlets for NSX Cloud run successfully .
- 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 **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	<p>The PowerShell script to generate the NSX Cloud Service Principal and managed identity roles for CSM and PCG. This script takes the following parameters:</p> <ul style="list-style-type: none"> ■ -subscriptionId <the Transit_VNet's_Azure_subscription_ID> ■ (optional) -servicePrincipalName <Service_Principal_Name> ■ (optional) -useOneServicePrincipal
AddServicePrincipal.ps1	<p>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:</p> <ul style="list-style-type: none"> ■ -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	<p>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.</p>
nsx_pcg_role.json	<p>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.</p> <p>Note The default PCG (Gateway) Role Name is <code>nsx-pcg-role</code>. You need to provide this value when adding your subscription in CSM.</p>

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
<code>-computeSubscriptionId</code>	The Compute_VNet's Azure Subscription ID
<code>-transitSubscriptionId</code>	The Transit VNet's Azure Subscription ID
<code>-csmRoleName</code>	Get this value from the file <code>nsx_csm_role.JSON</code>
<code>-servicePrincipalName</code>	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.
- 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 in CSM, for example, Azure-DevOps-Account , Azure-Finance-Account , etc.
Client ID	Copy paste this value from the output of the PowerShell script.
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.

Option	Description
Gateway Role Name	The default value is <code>nsx-pcg-role</code> . This value is available from the <code>nsx_pcg_role.json</code> 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.

5 (Optional) If you have a brownfield deployment, mark all the VMs as **User Managed** in the VNet where you want VMs managed to prevent automatic security group assignment under the Quarantine Policy.

6 (Optional) Manage access to regions. See [Managing Regions in CSM](#).

What to do next

[Deploy PCG in a VNet](#)

Adding your AWS Account

You might have one or more AWS accounts with VPCs and workload VMs that you want to bring under NSX management.

Overview:

- 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 primary AWS account.
 - **Scenario 3:** You want to use multiple AWS accounts with NSX Cloud, designating one account where you will install the PCG, that is a Transit VPC, and other accounts that will link to this PCG, that is, Compute VPCs. See [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#) for details on PCG deployment options.

Here is an outline of the process:

- 1 Use the NSX Cloud shell script to do the following. This step requires AWS CLI configured with the account you want to add.
 - 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 configured for the account and the default region.
- `jq` (a JSON parser).
- `openssl` (network security requirement).

Note If using AWS GovCloud (US) accounts, ensure that your AWS CLI is configured for the GovCloud (US) account and the default region is specified in the AWS CLI configuration file.

Procedure

- ◆ On a Linux or compatible desktop or server, download the SHELL script named `nsx_csm_iam_script.sh` from the NSX **Download page > Drivers & Tools > NSX Cloud Scripts > AWS**.
- ◆ **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 VPC](#) to finish the process of setting up a Transit or Self-Managed VPC.

- ◆ **Scenario 2:** You want to use multiple sub-accounts in AWS that are managed by one primary AWS account.

- a Run the script from your AWS primary 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 primary 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 primary 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 VPC](#) to finish the process of setting up a Transit or Self-Managed VPC.

- ◆ **Scenario 3:** You want to use multiple AWS accounts with NSX Cloud, designating one account for Transit VPC and other accounts for Compute VPCs. See [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#) for details on PCG deployment options.

- 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.
- 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. 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 The script poses the question: Do you want to create an IAM user for CSM and an IAM role for PCG? [yes/no]. Use the following guidance for the appropriate response:

This AWS account was already added to CSM.	Enter no in response to Do you want to create an IAM user for CSM and an IAM role for PCG? [yes/no]
This account has not been added to CSM before.	Enter yes in response to Do you want to create an IAM user for CSM and an IAM role for PCG? [yes/no]

- e (Optional) If you answered **yes** to creating an IAM user for CSM and PCG in the previous question, enter a name for the IAM user when asked What do you want to name the IAM User?. 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 primary 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
Discover 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 <i>nsx_pcg_service</i> . You can find this value in the output of the script in the file <i>aws_details.txt</i> .

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.

- 4 (Optional) If you have a brownfield deployment, mark all the VMs as **User Managed** in the VPC where you want VMs managed to prevent automatic security group assignment under the Quarantine Policy.
- 5 (Optional) Manage access to regions. See [Managing Regions in CSM](#).

What to do next

[Deploy PCG in a VPC](#)

Managing Regions in CSM

For your public cloud account added in CSM, you can get a list of supported regions and restrict access to specific regions.

- To get a list of specific regions, run the following API:

```
GET https://<csm-IP>/csmapi/api/v1/csm/supported-regions
```

Note If you do not see all the available regions in your AWS account, check whether you have enabled regions in your AWS account. See AWS documentations for details on enabling regions.

- To restrict regions in Microsoft Azure, run the following API:

```
PUT https://<csm-IP>/api/v1/csm/azure/accounts/<account_id>/desired-regions
```

Example Request:

```
PUT https://<nsx-csm>/api/v1/csm/azure/accounts/9174ffd1-41b1-42d6-a28d-05c61a0698e2/
desired-regions
{
  "regions": [
    {
      "id": "westus",
      "display_name": "westus",
      "enabled": true,
    },
    {
      "id": "eastus2",
      "display_name": "eastus2",
      "enabled": false,
    }
  ]
}
```

- To restrict regions in AWS, run the following API:

```
PUT https://<csms-IP>/api/v1/csm/aws/accounts/<account_id>/desired-regions
```

Example Request:

```
PUT https://<nsx-csm>/api/v1/csm/aws/accounts/9174ffd1-41b1-42d6-a28d-05c61a0698e2/desired-regions
{
  "default_client_region": "us-east-1",
  "regions": [
    {
      "id": "us-west-2",
      "display_name": "Oregon",
      "enabled": false,
    },
    {
      "id": "us-east-1",
      "display_name": "N. Virginia",
      "enabled": true,
    }
  ]
}
```

See the latest version of the *NSX API Guide* at <https://code.vmware.com/> for API details.

NSX Public Cloud Gateway: Architecture and Modes of Deployment

The NSX Public Cloud Gateway (PCG) provides north-south connectivity between the public cloud and the on-prem management components of NSX.

Familiarize yourself with the following terminology explaining the PCG's architecture and deployment modes for workload VM-management.

Note The PCG is deployed in a single default size for each supported public cloud:

Public Cloud	PCG instance type
AWS	<p>c5.xlarge.</p> <p>Some regions might not support this instance type. Refer to AWS documentation for details.</p> <p>Note If you see high CPU usage alerts with this instance type, resize PCG instances to c5.2xlarge.</p> <p>If you have a high availability pair of PCG instances, resize the standby PCG first by stopping, resizing and restarting it. Stop the currently active PCG next and wait until the standby PCG becomes active. Resize and restart this PCG and it should become active. See AWS documentation for details on resizing instance types.</p>
Microsoft Azure	Standard DS3 v.2

Architecture

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.

Modes of Deployment

Self-managed VPC/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.

Transit VPC/VNet: A self-managed VPC/VNet becomes a *Transit* VPC/VNet when you link Compute VPCs/VNets to it.

Compute VPC/VNet: VPCs/VNets that do not have the PCG deployed on them but link to a Transit VPC/VNet are called *Compute* VPCs/VNets.

AWS Transit Gateway with PCG: Starting in NSX 3.1.1, you can use AWS Transit Gateway to connect a Transit VPC with Compute VPCs. See [Using PCG with AWS Transit Gateway](#) for details.

Subnets Required in Your VPC/VNet to deploy PCG

The PCG uses the following subnets that you set up in your VPC/VNet. See [Connect Microsoft Azure with On-prem NSX](#) or [Connect AWS with On-prem NSX](#).

- **Management subnet:** This subnet is used for management traffic between on-prem NSX and PCG. Example range: /28.
- **Uplink subnet:** This subnet is used for north-south internet traffic. Example range: /24.
- **Downlink subnet:** This subnet encompasses the workload VM's IP address range. Size this subnet bearing in mind that you might need additional interfaces on the workload VMs for debugging.

PCG deployment aligns with your network addressing plan with FQDNs for the NSX components and a DNS server that can resolve these FQDNs.

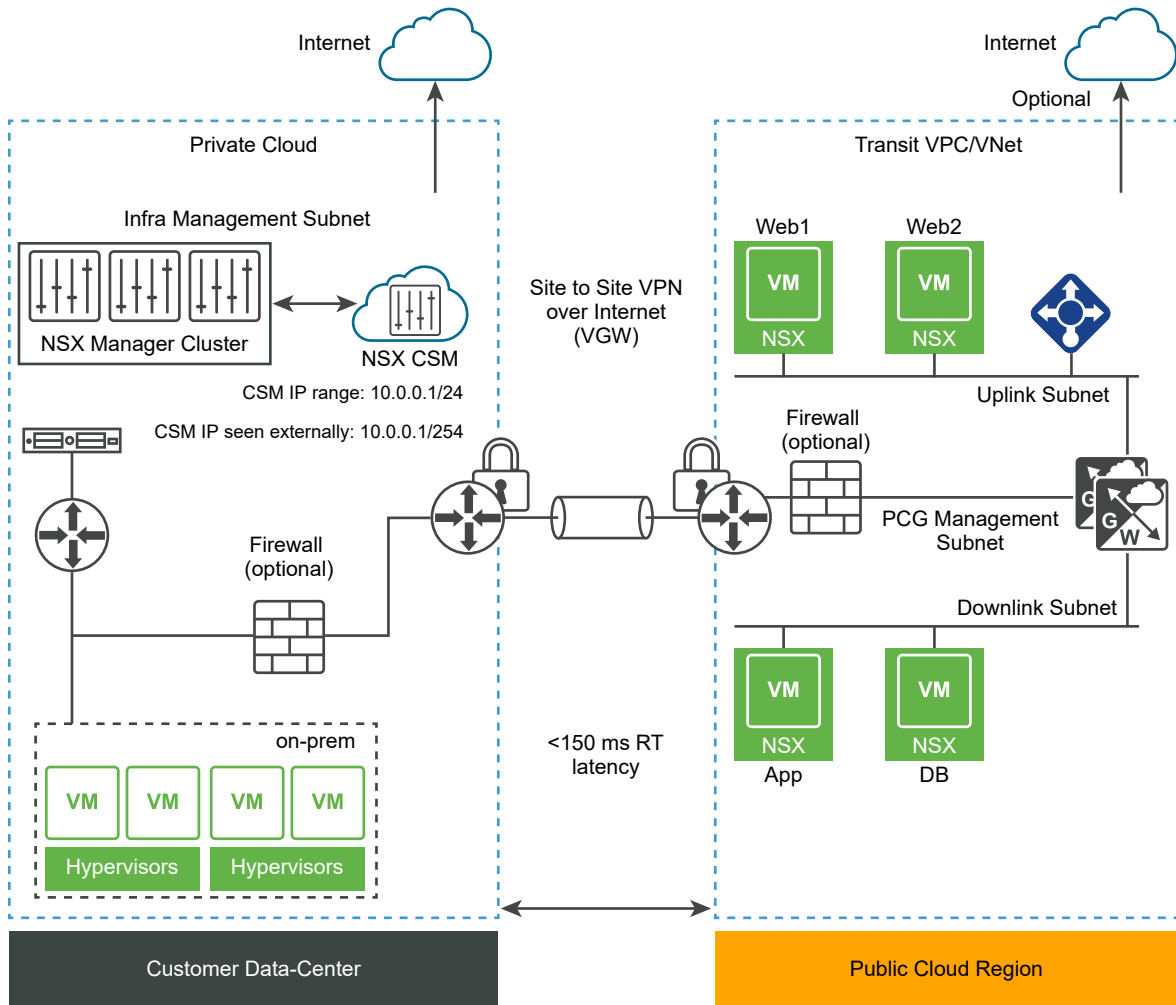
Note It is not recommended to use IP addresses for connecting the public cloud with NSX using PCG, but if you do, you must not change your IP addresses.

Impact of on-prem and public cloud connectivity mode on PCG's discovery of CSM

After PCG is deployed in your public cloud, it must interact with CSM as the management interface for your public cloud inventory. To ensure that PCG can reach the IP address of CSM, follow these guidelines:

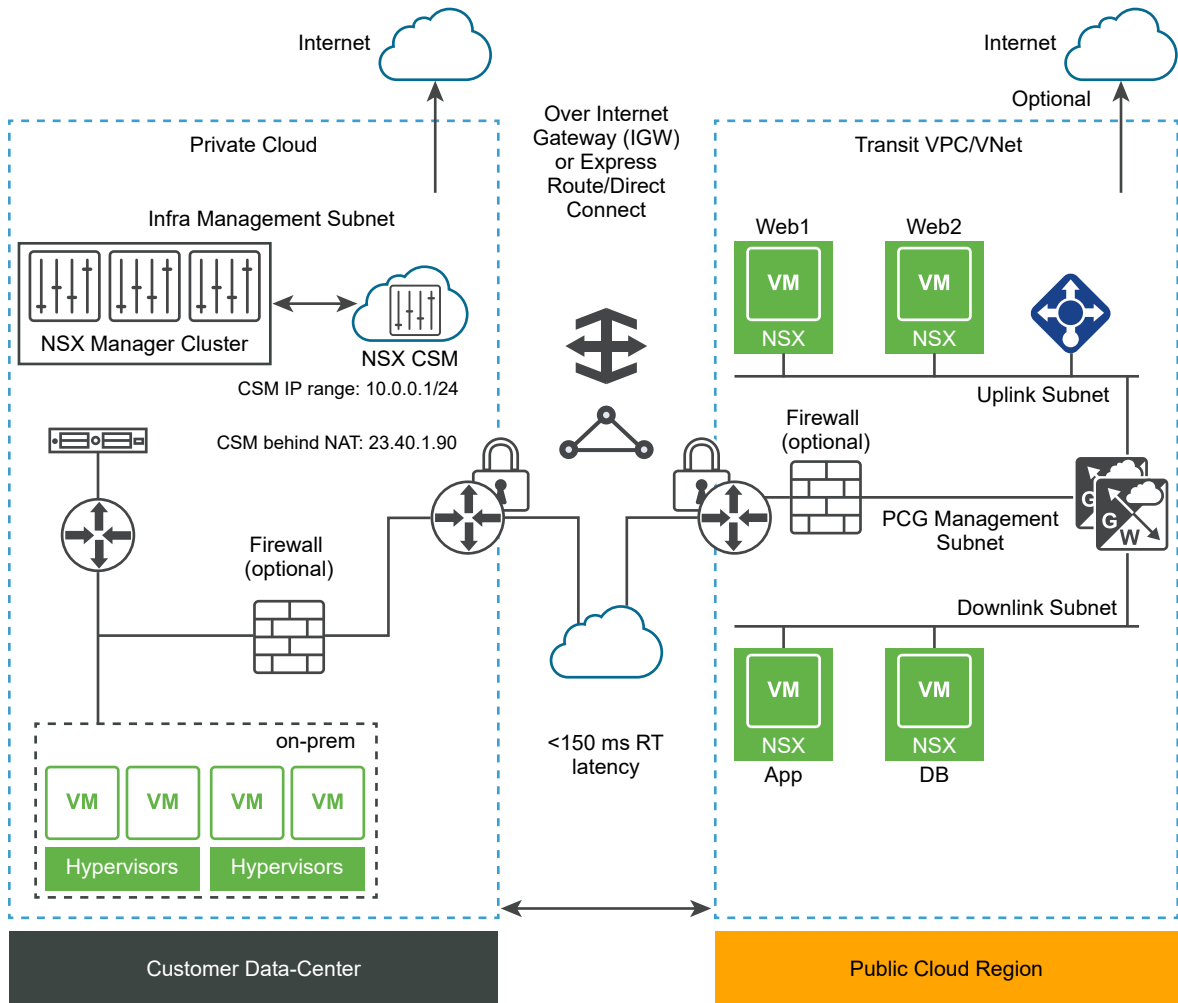
- For PCG deployed in Private IP mode (VGW): PCG discovers CSM with either the actual CSM IP address or with one of the IP addresses in the subnet range provided.

Figure 16-3. NSX Public Cloud Gateway Architecture with VGW Connectivity



- For PCG deployed in Public IP mode (IGW): PCG can discover CSM using CSM's NAT-translated IP address that allows access to the real IP address or subnet range for CSM.

Figure 16-4. NSX Public Cloud Gateway Architecture with IGW Connectivity



Modes of VM-management

NSX Enforced Mode: In this mode, workload VMs are brought under NSX management by installing NSX Tools on each workload VM to which you apply the tag `nsx.network=default` in your public cloud.

Native Cloud Enforced Mode: In this mode, your workload VMs can be brought under NSX management without the use of NSX Tools.

Quarantine Policy

Quarantine Policy: NSX Cloud's threat detection feature that works with your public cloud security groups.

- In the NSX Enforced Mode, you can enable or disable Quarantine Policy. As a best practice, disable Quarantine Policy and add all your VMs to the **User Managed** list when onboarding workload VMs.

- In the Native Cloud Enforced Mode Quarantine Policy is always enabled and cannot be disabled.

Design Options

Regardless of the mode you deploy the PCG in, you can link a Compute VPC/VNet to it in either mode.

Table 16-11. Design Options with PCG Deployment Modes

PCG Deployment Mode in Transit VPC/VNet	Supported Modes when linking Compute VPCs/VNets to this Transit VPC/VNet
NSX Enforced Mode	<ul style="list-style-type: none"> ■ NSX Enforced Mode ■ Native Cloud Enforced Mode
Native Cloud Enforced Mode	<ul style="list-style-type: none"> ■ NSX Enforced Mode ■ Native Cloud Enforced Mode

Note Once a mode is selected for a Transit or Compute VPC/VNet, you cannot change the mode. If you want to switch modes, you must undeploy the PCG and redeploy it in the desired mode.

Deploy PCG or Link to a PCG

Follow these instructions for deploying PCG or linking to a PCG.

NSX Cloud creates networking and security constructs in NSX Manager and your public cloud after PCG is deployed. See [Auto-Configurations after PCG Deployment or Linking](#) .

If you are using AWS Transit Gateway, see [Using PCG with AWS Transit Gateway](#).

Note For Native Cloud PCG deployment, the default *admin* and *root* password is same and is set as *NSX-`<instance-ID>`*. The instance-ID of the AWS or Azure VM is prefixed. For example, if the instance ID is *i-0349f21e34dc27b08*, the initial password is set as *NSX-i-0349f21e34dc27b08*

The *root* password gets expired by default. If you want to access root shell of the PCG appliance, you must reset the password using the initial password which is *NSX-`<instance-ID>`*.

Deploy PCG in a 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

- If you have deployed NSX Cloud components on-prem, ensure the VNet is connected with your on-prem NSX.

If you have deployed NSX Cloud components in Microsoft Azure, ensure that the VNet is peered with the NSX Cloud Management VNet. See deployment architecture details at [Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image](#).

- Verify that your Microsoft Azure subscription is added into CSM.
- Verify that you have the required subnets in the VNet where you are deploying PCG: *uplink*, *downlink*, and *management*. For High Availability, you must have an uplink subnet for the secondary PCG that is different from the primary PCG.

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 the PCG.
- 4 Click **Deploy Gateways**. The **Deploy 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.
Manage with NSX Tools	Leave in the default disabled state to onboard workload VMs in the Native Cloud Enforced Mode. If you want to install NSX Tools on your workload VMs to use the NSX Enforced Mode, enable this option.
Quarantine Policy on the Associated VNet	You can only change the Quarantine Policy setting if you choose to manage workload VMs using NSX Tools (NSX Enforced Mode). Quarantine Policy is always enabled in the Native Cloud Enforced Mode. 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 <i>NSX Administration Guide</i> for details.
Auto-install NSX Tools	This is only available when you enable Manage with NSX Tools. If selected, NSX Tools are auto-installed on all workload VMs in the Transit/Self-managed/linked Compute VNet if the tag <code>nsx.network=default</code> is applied to them.
Gateway Connectivity Mode	The PCG can be accessed from CSM using a public IP address or a private IP address depending on the connectivity mode between your public cloud and your on-prem NSX installation. If you select Auto Detect , they system attempts to connect with CSM over VGW first, and if that fails, over IGW. If the system cannot connect with CSM, the deployment fails. See Impact of on-prem and public cloud connectivity mode on PCG's discovery of CSM for details.

Option	Description
Use Marketplace Image	<p>This option is only available in NSX 3.1.1.</p> <p>It is enabled by default when a compatible marketplace image is available to deploy in Microsoft Azure. See Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image for details.</p>
Azure Marketplace Terms	<p>If you are using the marketplace image to deploy PCG, you must accept Microsoft Azure terms of use. NSX Cloud provides the terms for you to download and read. Select the checkbox to accept the terms to proceed.</p>
Local Storage Account	<p>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.</p> <p>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.</p>
VHD URL	<p>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.</p> <p>Note The VHD must be in the correct URL format. We recommend that you use the Click to copy option in Microsoft Azure.</p>
Proxy Server	<p>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.</p> <p>See (Optional) Configure Proxy Servers for details on how to configure proxy servers in CSM.</p>

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

Follow instructions at *Using NSX Cloud* in the *NSX Administration Guide*.

Deploy PCG in a 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](#).

If you are using AWS Transit Gateway, see [Using PCG with AWS Transit Gateway](#).

Prerequisites

- Ensure the VPC is connected with your on-prem NSX.
- Verify that your AWS account is already added into CSM.
- Verify that the VPC on which you are deploying PCG has the required subnets appropriately adjusted for High Availability: *uplink*, *downlink*, and *management*.
- Verify that the configuration for your VPC's network ACL includes 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 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.
Manage with NSX Tools	Leave in the default disabled state to onboard workload VMs in the Native Cloud Enforced Mode. If you want to install NSX Tools on your workload VMs to use the NSX Enforced Mode, enable this option.
Quarantine Policy on the Associated VPC	You can only change the Quarantine Policy setting if you choose to manage workload VMs using NSX Tools (NSX Enforced Mode). Quarantine Policy is always enabled in the Native Cloud Enforced Mode 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 <i>NSX Administration Guide</i> for details.

Option	Description
Gateway Connectivity Mode	<p>The PCG can be accessed from CSM using a public IP address or a private IP address depending on the connectivity mode between your public cloud and your on-premises NSX installation. If you select Auto Detect, the system attempts to connect with CSM over VGW first, and if that fails, over IGW. If the system cannot connect with CSM, the deployment fails.</p> <p>See Impact of on-prem and public cloud connectivity mode on PCG's discovery of CSM for details.</p>
InstanceType	<p>Select any one of the sizes from the drop-down menu list based on your requirement. There are four Instance Type sizes available:</p> <ul style="list-style-type: none"> ■ Small ■ Medium ■ Large ■ Extra Large <p>See NSX Public Cloud Gateway: Architecture and Modes of Deployment for more information on PCG instance type.</p> <p>Note You can enable Firewall features like Application ID, IDPS, and URL Enforcement only on Large and Extra Large size PCG deployment. However, in NSX 3.2, Firewall features are available in Tech Preview mode. Use these features only for experimental purposes and VMware does not officially provide support for these features.</p>
Proxy Server	<p>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.</p> <p>See (Optional) Configure Proxy Servers for details on how to configure proxy servers in CSM.</p>
Override AMI ID	<p>Use this advanced feature to provide a different AMI ID for the PCG from the one that is available in your AWS account.</p>

7 Click Next.

8 Complete the Subnet details.

Option	Description
Enable HA for Public Cloud Gateway	<p>The recommended setting is Enable, that sets up a High Availability Active/Standby pair to avoid an unscheduled downtime.</p>
Primary gateway settings	<p>Select an Availability Zone such as <code>us-west-1a</code>, from the drop-down menu as the primary gateway for HA.</p> <p>Assign the uplink, downlink, and management subnets from the drop-down menu.</p>
Secondary gateway settings	<p>Select another Availability Zone such as <code>us-west-1b</code>, from the drop-down menu as the secondary gateway for HA.</p> <p>The secondary gateway is used when the primary gateway fails.</p> <p>Assign the uplink, downlink, and management subnets from the drop-down menu.</p>

Option	Description
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

Follow instructions at "Using NSX Cloud" in the *NSX Administration Guide*.

Using PCG with AWS Transit Gateway

If you use AWS Transit Gateway, you can deploy the PCG in any VPC and connect this VPC with the Transit Gateway.

Follow instructions at [Deploy PCG in a VPC](#).

Any other VPCs connected to the Transit Gateway can have their workload VMs managed by NSX for micro-segmentation.

NSX Cloud does not manage networking between the Transit and Compute VPCs or the workload VMs. All NSX networking constructs are created upon PCG deployment but only the Security constructs are valid if you are working with AWS Transit Gateway. See [Security Entities](#) for a list of auto-created security policies after PCG deployment.

- Currently only NSX Enforced Mode is supported. You must install NSX Tools in your workload VMs. See *NSX Enforced Mode* in the *NSX Administration Guide* for instructions.
- The VPC where you deploy PCG – Transit VPC – must have the same subnets as required by a Transit VPC that is not using the AWS Transit Gateway. See [Subnets Required in Your VPC/VNet to deploy PCG](#) for details.
- You must link compute VPCs to the Transit VPC. See [Link to a Transit VPC or VNet](#) for instructions.
- You must ensure that workload VMs with NSX Tools installed on them have connectivity with the management subnet of the Transit VPC.
- To utilize micro-segmentation, you must add a Forwarding Policy with the following values:

Option	Value
Sources	A Group in NSX Manager that contains all NSX-Managed VMs from your Transit and Compute VPCs
Destinations	All (0.0.0.0/0)

Option	Value
Services	Any
Action	Route to Underlay

See *Add or Edit Forwarding Policies* in the *NSX Administration Guide* for details about Forwarding Policies.

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.
- Verify that the VPC/VNet you want to link is connected to the Transit VPC or VNet through VPN or peering.
- Verify that the Compute VPC/VNet is in the same region as the Transit 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/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.
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 <i>NSX Administration Guide</i> for details.
Manage with NSX Tools	Leave in the default disabled state to onboard workload VMs in the Native Cloud Enforced Mode. If you want to install NSX Tools on your workload VMs to use the NSX Enforced Mode, enable this option.
Auto-install NSX Tools	This is only available when you choose to manage with NSX Tools and only for Microsoft Azure VNets. If selected, NSX Tools are auto-installed on all workload VMs in the Transit/Self-managed/linked Compute VNet if the tag <code>nsx.network=default</code> is applied to them.

What to do next

Follow instructions at [Using NSX Cloud](#) in the *NSX Administration Guide*.

Auto-Configurations after PCG Deployment or Linking

The deployment of PCG in a Transit VPC/VNet and linking a compute VPC/VNet to it triggers necessary configurations in NSX and the public cloud.

Auto-created NSX Logical Entities

A set of logical entities are auto-created in NSX Manager.

Log in to NSX Manager to view the auto-created logical entities.

Important Do not delete any of these auto-created entities except if you are manually undeploying PCG. See [Troubleshooting PCG Undeployment](#) for details.

System Entities

You can see the following entities under the **System** tab:

Table 16-12. 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	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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 available under the **Inventory** tab:

Table 16-13. Groups

Groups	Scope
Two Groups named: <ul style="list-style-type: none"> ■ cloud-default-route ■ cloud-metadata services 	Scope: Shared across all PCGs
One Group created at the 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

Table 16-13. Groups (continued)

Groups	Scope
Two Groups for each Compute VPC/VNet: <ul style="list-style-type: none"> ■ Network CIDR Group for all CIDRs of the Compute VPC/VNet: <code>cloud-<Compute VPC/VNet ID>-cidr</code> ■ Local Segment Group for all managed segments within the Compute VPC/VNet: <code>cloud-<Compute VPC/VNet ID>-local-segments</code> 	Scope: shared across all Compute VPC/VNets
The following Groups are created for the currently supported public cloud services: <ul style="list-style-type: none"> ■ <code>aws-dynamo-db-service-endpoint</code> ■ <code>aws-elb-service-endpoint</code> ■ <code>aws-rds-service-endpoint</code> ■ <code>aws-s3-service-endpoint</code> ■ <code>azure-cosmos-db-service-endpoint</code> ■ <code>azure-load-balancer-service-endpoint</code> ■ <code>azure-sql-service-endpoint</code> ■ <code>azure-storage-service-endpoint</code> 	Scope: Shared across all PCGs

Note For PCGs deployed or linked to in the Native Cloud Enforced Mode, all the workload VMs in the VPC/VNet become available under Virtual Machines in NSX Manager.

Networking Entities

The following entities are created at different stages of onboarding and can be found under the **Networking** tab:

Figure 16-5. Auto-created NSX Networking Entities After PCG is Deployed

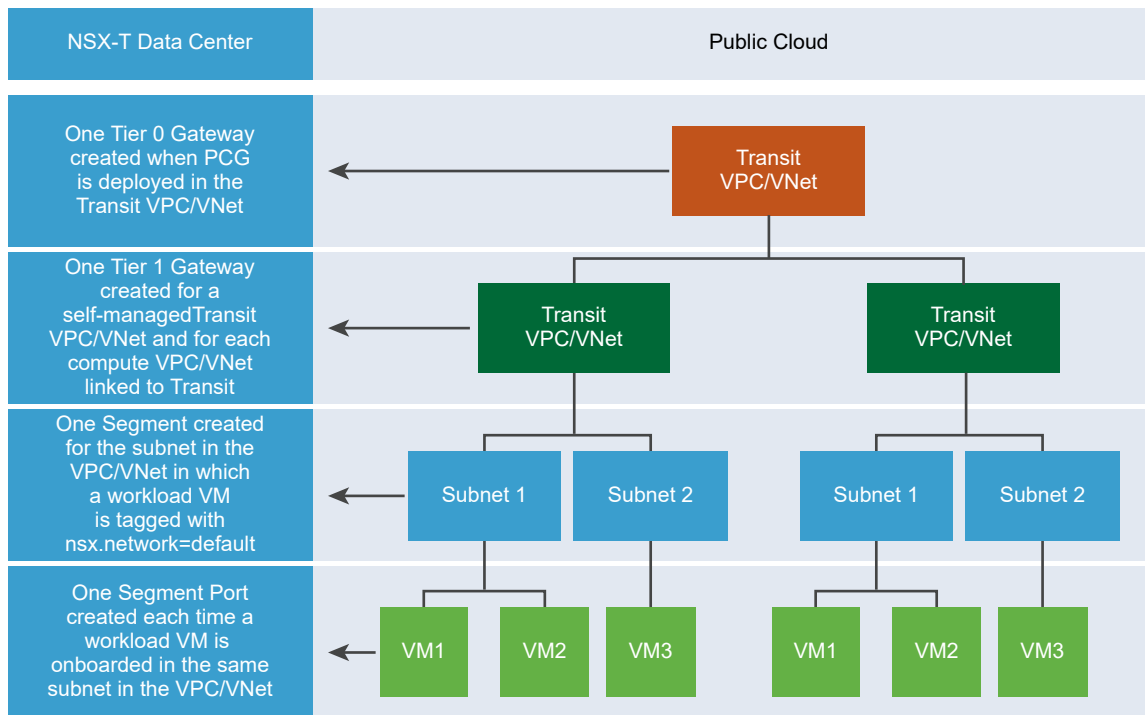


Table 16-14. Auto-Created Networking Entities

Onboarding Task	Logical Entities Created in NSX
PCG deployed on Transit VPC/VNet	<ul style="list-style-type: none"> ■ Tier-0 Gateway ■ Infra Segment (Default VLAN switch) ■ Tier-1 router
Compute VPC or VNet linked to the Transit VPC/VNet	<ul style="list-style-type: none"> ■ Tier-1 router
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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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 network (overlay)

Security Entities

The following entities are available under the **Security** tab:

Table 16-15. Auto-Created Security Entities

Logical Security Entity	How many are created?	Nomenclature	Scope
Distributed Firewall (East-West)	Two per Transit VPC/VNet: <ul style="list-style-type: none"> ■ Stateless ■ Stateful 	<ul style="list-style-type: none"> ■ cloud-stateless-<VPC/VNet ID> ■ cloud-stateful-<VPC/VNet ID> 	<ul style="list-style-type: none"> ■ 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>	

Auto-created Public Cloud Configurations

In your public clouds, some configurations are set up automatically after you deploy PCG.

Some auto configurations are common to all public clouds and both NSX management modes. Other configurations are specific to either the public cloud or the NSX management mode.

Specific to AWS

The following are specific to 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.

Specific to Microsoft Azure

The following are specific to Microsoft Azure:

- A common Resource Group is created per region, per subscription. It is named like: `nsx-default-<region-name>-rg`, for example: `nsx-default-westus-rg`. All VNets in this region share this Resource Group. This Resource Group and all the NSX-created security groups named like `default-<vnet-ID>-sg` are not deleted from the Microsoft Azure region after you off-board a VNet in this region from NSX Cloud.

Common to both modes and all public clouds

The following are created in all public clouds and for both NSX-management modes: NSX Enforced Mode and Native Cloud Enforced Mode:

- The **gw** security groups are applied to the respective PCG interfaces in VPCs or VNets.

Table 16-16. Public Cloud Security Groups created by NSX Cloud for PCG interfaces

Security Group name	Description
gw-mgmt-sg	Gateway Management Security Group
gw-uplink-sg	Gateway Uplink Security Group
gw-vtep-sg	Gateway Downlink Security Group

Specific to Native Cloud Enforced Mode

The following security groups are created when the PCG is deployed in the Native Cloud Enforced Mode.

After workload VMs are matched with groups and corresponding security policies in NSX Manager, security groups named like `nsx-<GUID>` are created in the public cloud for each matching security policy.

Note In AWS, Security Groups are created. In Microsoft Azure, Application Security Groups are created corresponding to Groups in NSX Manager and Network Security Groups are created corresponding to Security Policies in NSX Manager.

Security Group name	Available in Microsoft Azure?	Available in AWS?	Description
default-vnet-<vnet-id>-sg	Yes	No	NSX Cloud-created security group in the common Microsoft Azure Resource Group for assigning to VMs that are not matched with a security policy in NSX.
default	No	Yes	An existing security group in AWS used by NSX Cloud for assigning to VMs that are not matched with a security policy in NSX.
vm-overlay-sg	Yes	Yes	VM overlay security group (this is not used in the current release)

Specific to NSX Enforced Mode

The following security groups are created for workload VMs when you deploy PCG in the NSX Enforced Mode.

Table 16-17. Public Cloud Security Groups created by NSX Cloud for Workload VMs in the NSX Enforced Mode

Security Group name	Available in Microsoft Azure?	Available in AWS?	Description
default-vnet-<vnet-id>-sg	Yes	No	NSX Cloud-created security group in Microsoft Azure for threat-detection workflows in the NSX Enforced Mode
default	No	Yes	An existing security group in AWS used by NSX Cloud for threat-detection workflows in the NSX Enforced Mode
vm-underlay-sg	Yes	Yes	VM underlay security group
vm-overlay-sg	Yes	Yes	VM overlay security group (this is not used in the current release)

Integrate Horizon Cloud Service with NSX Cloud

Starting in NSX 3.1.1, you can integrate your Horizon Cloud deployment with NSX Cloud with fewer manual steps than in earlier releases.

Horizon Cloud integration is supported with NSX Cloud management components – NSX Manager and Cloud Service Manager (CSM) – deployed either on-premise, or natively in Microsoft Azure, starting in NSX 3.1.1.

The diagrams depict the options available to deploy NSX Cloud management components and shared subnets between PCG and Horizon Cloud Management components as possible scenarios for this integration.

Figure 16-6. Horizon Cloud Integration with NSX Cloud Components deployed On-prem





Figure 16-7. Horizon Cloud Integration with NSX Cloud with NSX Cloud Components deployed in Microsoft Azure



You must deploy your Horizon Cloud pod in the same Self-managed/Transit VNet where you have deployed the PCG. After deploying the Horizon Cloud pod in a Self-managed/Transit VNet, the system auto-creates the necessary NSX policies to enable communication between Horizon Cloud management components and VDIs deployed in Microsoft Azure. You can create security policies for VDIs as necessary.

Note Auto-creation of entities is a feature of NSX version 3.1.1 and later only.

Prerequisites

- Verify that NSX Cloud management components are already deployed and active. See [Deploy NSX On-Prem Components](#) or [Deploy NSX Cloud Components in Microsoft Azure using the NSX Cloud Marketplace Image](#) for relevant instructions for your deployment model.
- Verify that a PCG is already deployed in the same VNet where you are deploying the Horizon Cloud pod. A VNet that has the PCG deployed in it is called a Self-managed or Transit VNet in NSX Cloud terminology. See [NSX Public Cloud Gateway: Architecture and Modes of Deployment](#).

- If NSX Cloud management components are deployed on-prem, verify that you have a Self-managed or Transit VNet connected with the on-prem components using VGW or IGW connectivity.
- If you are deploying NSX Cloud management components in the public cloud, verify that you have a peering connection between the VNet where NSX Cloud components are deployed and the Self-managed/ Transit VNet.

Procedure

- 1 From your Microsoft Azure subscription, deploy the Horizon Cloud pod in a Self-managed/ Transit VNet. See [Horizon Pods on Microsoft Azure - Using the Horizon Cloud Administration Console Capacity Page to Add More Pods to Your Pod Fleet](#) in the *Horizon Cloud Service Product Documentation* for detailed steps. Note the following configurations in the **Add Microsoft Azure Capacity** wizard that are specific to integration with NSX Cloud:
 - a In the **Subscription** tab, select the Microsoft Azure subscription where PCG is deployed.
 - b In the **Pod Setup** tab, under the section **Networking** for **Virtual Network**, select the Microsoft Azure VNet in which PCG is deployed.

Use subnets that allow communication between the PCG and Horizon Cloud management components.
 - c Click **Validate and Proceed**
- 2 Create a base image with NSX Tools installed.

If you are using pod manifest 2632.0 or later, and your customer record has access to the new Image Management feature that auto-installs NSX Tools into the base image, enable **Install NSX Tools**.

If you are using pod manifests prior to 2632.0, follow the steps described at [Install the NSX Agent in the Horizon Cloud Imported Image VM](#) to manually install NSX Tools in the base image.
- 3 Use the base image to create the VDI desktop assignment, enabling **NSX Cloud Managed**. See [Horizon Cloud Workflows and NSX Cloud](#) for details.

What to do next

Verify the [Auto-Created Entities After Horizon Cloud Integration](#) .

Auto-Created Entities After Horizon Cloud Integration

After you deploy your Horizon Cloud pod in a Self-managed/Transit VNet, the following entities are auto-created in CSM and NSX Manager.

Note Auto-creation of entities is a feature of NSX version 3.1.1 and later only.

Horizon Cloud VMs in CSM

- Horizon Cloud VMs can be management VMs or VDIs for end users.
- CSM distinguishes Horizon Cloud management VMs from end-user consumable VDIs as follows:
 - The three types of Horizon Cloud management VMs – UAG, Base, and Node are labeled as **Horizon Management VMs** in **CSM > Clouds > Azure > Instances**. The Horizon Cloud administrator has complete control over the security groups assigned to these VMs in Microsoft Azure.
 - All VDIs launched in the Horizon Cloud pod, using any image with NSX Cloud enabled, are NSX-managed if they enable NSX Cloud when launched. NSX Tools are installed on such VDIs and they are managed like other managed VMs in the NSX Enforced mode. In **CSM > Clouds > Azure > Instances**, you can see these VDIs with the label **Horizon VDI**.

See "Managing VMs in the NSX Enforced Mode" in the *NSX Administration Guide* for details.

Also see "VMware NSX Cloud and Horizon Cloud Pods in Microsoft Azure" in the [Horizon Cloud Service Product Documentation](#).

Horizon Cloud Entities Created in NSX Manager

NSX Manager Component	Auto-created Entities	Details
Inventory > Services	HorizonUAGPolicyService	This service allows communication between the Horizon Cloud UAG and VDIs. See this table for details: Table 16-18. DFW Policy Auto-created for Horizon Cloud Integration under the Infrastructure category
Inventory > Services	HorizonNodeVMPolicyService	This service is used to allow communication from the VDIs to Horizon Cloud Management Node VMs. See this table for details: Table 16-18. DFW Policy Auto-created for Horizon Cloud Integration under the Infrastructure category

NSX Manager Component	Auto-created Entities	Details
Inventory > Groups	<ul style="list-style-type: none"> ■ vmw-hcs-<pod-id>-base ■ vmw-hcs-<pod-id>-node ■ vmw-hcs-<pod-id>-uag ■ vmw-hcs-<pod-id>-vdi 	<p>The group definition for these groups is as follows:</p> <ul style="list-style-type: none"> ■ instance-type label that Horizon Cloud applies to these VMs in Microsoft Azure. ■ Microsoft Azure ID of the Self-managed/Transit VNet that also hosts the Horizon Cloud pod. <p>You manage the VDIs that are included in the <code>vmw-hcs-<id>-vdi</code> group. The other groups are managed by Horizon Cloud.</p> <p>The Horizon Cloud jumpbox VMs are grouped under <code>vmw-hcs-<id>-node</code></p>
Inventory > Virtual Machines	Horizon Cloud VDIs with names provided by Horizon Cloud	<p>These are the VDIs in Horizon Cloud that are categorized as Virtual Machines in NSX Manager. All security policies and other configurations in NSX Manager are targeted towards these Virtual Machines.</p>
Inventory > Tags	<ul style="list-style-type: none"> ■ Tag Scope: <code>azure:instance_type</code> ■ Tag Values: <ul style="list-style-type: none"> ■ HORIZON_MGMT ■ HORIZON_BASE ■ HORIZON_UAG ■ HORIZON_VDI 	<p>These system tags are used to create groups for security policies.</p>

Security Policy

Under **Security > Distributed Firewall > Infrastructure** a DFW policy is created with the name: `vmw-hcs-<pod_id>-security-policy`. This policy has the following **Allow** rules.

Table 16-18. DFW Policy Auto-created for Horizon Cloud Integration under the Infrastructure category

DFW Rule Name	Source	Destination	Service/Ports	Protocols
AllowHCSUAGToVDI	Unified Access Gateway	VDI	HorizonUAGPolicyService TCP (Source: Any; Destination: 22443,32111,4172,443,8443,9427) UDP (Source: Any Destination: 22443,4172)	TCP and UDP
AllowVDIToHCSNode	VDI	Node VMs	HorizonNodeVMPolicyService (Source: Any; Destination: 3099,4001,4002)	TCP

Note All networking for NSX-managed VDIs within the VNet is through Microsoft Azure. NSX only manages traffic going out of the VNet.

See "Group VMs using NSX and Public Cloud Tags" in the *NSX Administration Guide* for details on discovered tags: these are tags that you apply in Microsoft Azure to your VDIs and they are visible in NSX Manager to enable tag-based grouping.

(Optional) Install NSX Tools on your Workload VMs

If you are using the NSX Enforced Mode, proceed to installing NSX Tools in your workload VMs.

See instructions and further details at "Onboarding VMs in the NSX Enforced Mode" in the *NSX Administration Guide*.

Un-deploy NSX Cloud

You must un-deploy NSX Cloud configurations before decommissioning the CSM appliance.

To un-deploy NSX Cloud, perform the following steps:

1 [Undeploy or Unlink PCG](#).

For any issues, see [Troubleshooting PCG Undeployment](#).

2 Decommission the CSM appliance.

- a Log in to the vSphere Client using your administrator credentials.
- b Power off the CSM appliance VM.
- c To delete the CSM appliance VM from the datastore, right-click the appliance VM.
- d Select **Delete from Disk**, and then click **OK**. For details, refer to the vSphere documentation.

Note In case of Azure deployment, you must delete the entire Resource Group that contains three MPs and one CSM.

Undeploy or Unlink PCG

You can undeploy or unlink PCG after you have removed some NSX Cloud configurations.

In the NSX Enforced Mode

- Remove the `nsx.network=default` tag from NSX-managed workload VMs.
- Disable the Quarantine Policy if it is enabled.
- Delete all user-created logical entities associated with the PCG.

In the Native Cloud Enforced Mode

- Delete all user-created logical entities associated with the PCG.

Follow the steps relevant to the NSX management mode your PCG is deployed in.

What to read next

Procedure

- 1 [Remove the `nsx.network` tag in the Public Cloud](#)
Before you can undeploy PCG, all VMs must be unmanaged.
- 2 [Disable Quarantine Policy in the NSX Enforced Mode](#)
If using the NSX Enforced Mode you must disable Quarantine Policy if previously enabled. .
- 3 [Delete User-created Logical Entities](#)
All user-created logical entities associated with the PCG must be deleted.
- 4 [Undeploy or Unlink from CSM](#)
Follow these instructions to undeploy or unlink PCG after completing the prerequisites.
- 5 [Troubleshooting PCG Undeployment](#)
If PCG undeployment fails, you have to manually delete all the NSX Cloud-created entities in NSX Manager as well as in the public cloud.

Remove the `nsx.network` tag in the Public Cloud

Before you can undeploy PCG, all VMs must be unmanaged.

Note This is only applicable in the NSX Enforced Mode.

Go to the VPC or VNet in your public cloud and remove the `nsx.network=default` tag from the managed VMs.

Disable Quarantine Policy in the NSX Enforced Mode

If using the NSX Enforced Mode you must disable Quarantine Policy if previously enabled. .

This step is only applicable to the NSX Enforced Mode.

With Quarantine Policy enabled, your VMs are assigned security groups in your public cloud that are defined by NSX Cloud.

When you undeploy PCG, you must disable Quarantine Policy. Follow these steps: :

- 1 Go to the VPC or VNet in CSM.
- 2 From **Actions > Edit Configurations >**, turn off the setting for **Default Quarantine** .
- 3 All VMs that are unmanaged or quarantined in this VPC or VNet will be assigned to the `default` security group in AWS and the `default-vnet-<vnet-id>-sg` security group in Microsoft Azure.
- 4 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=default` tag from such VMs to take them out from NSX management, they are also assigned to the `default` security group in AWS and the `default-vnet-<vnet-id>-sg` security group in Microsoft Azure.

Note The common Resource Group created in Microsoft Azure, that is named like: `nsx-default-<region-name>-rg`, for example: `nsx-default-westus-rg`, is not removed when you undeploy PCG. This Resource Group and all the NSX-created security groups named like `default-<vnet-ID>-sg` are not deleted from the Microsoft Azure region. You can remove the NSX Cloud-specific security group any time after the VNet is off-boarded.

See [Auto-Configurations after PCG Deployment or Linking](#) for a list of NSX Cloud security groups.

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** or **Unlink from Transit VPC/VNet** from CSM. See [Auto-Configurations after PCG Deployment or Linking](#) for details.

Undeploy or Unlink from CSM

Follow these instructions to undeploy or unlink PCG after completing the prerequisites.

- 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** or **Unlink from Transit VPC/VNet**.

The default entities created by NSX Cloud are removed automatically when the PCG is undeployed or unlinked.

Troubleshooting PCG Undeployment

If PCG undeployment fails, you have to manually delete all the NSX Cloud-created entities in NSX Manager as well as in the public cloud.

- In your public cloud:
 - Terminate all PCGs in the Transit VPC/VNet.
 - Move all your workload VMs to a security group not created by NSX Cloud.
 - For Microsoft Azure, also delete the NSX Cloud-created Resource Group named like **nsx-gw-`<vnet ID>`-rg**.
- Delete the auto-created entities with the VPC/VNet ID in NSX Manager as listed here: [Auto-created NSX Logical Entities](#).

Note Do not delete the global entities that are auto-created. Only delete the ones that have the VPC/VNet ID in their name.

- Restart the CSM Service. Log in to the CSM appliance CLI and run the `restart service cloud-service-manager` command.

Important If the PCG un-deployment fails even after performing the steps mentioned in this topic or you need to redeploy in same VPC/VNet, a database cleanup for PCG may be required. The database cleanup needs engineering assistance. If you want to clean-up the PCG database, contact the VMware support team.

Uninstalling NSX from a Host Transport Node

17

The steps to uninstall NSX from a host transport node vary depending on the host type and how it is configured.

- [Uninstall NSX from a vSphere Cluster](#)

If you have installed NSX on a vSphere Cluster using transport node profiles, you can follow these instructions to uninstall NSX from all hosts in the cluster.

- [Uninstall NSX from a Managed Host in a vSphere Cluster](#)

You can uninstall NSX from a single host that is managed by VMware vCenter. The other hosts in the cluster are not affected.

- [Uninstall NSX from a Physical Host](#)

You can uninstall NSX from a physical host.

- [Triggering Uninstallation from the vSphere Web Client](#)

In the vSphere Web Client, if you move a host from a cluster prepared with a transport node profile to either another cluster, outside of the cluster as a standalone host, or outside of the data center, then NSX is uninstalled on the host that is moved. Such an uninstallation is not triggered when a host that is individually prepared with a transport node configuration is moved.

- [Uninstall NSX from a vSphere Lifecycle Manager cluster through NSX Manager](#)

You can trigger uninstallation of NSX on a host that is part of a vSphere Lifecycle Manager cluster through NSX Manager.

Uninstall NSX from a vSphere Cluster

If you have installed NSX on a vSphere Cluster using transport node profiles, you can follow these instructions to uninstall NSX from all hosts in the cluster.

For more information on transport node profiles, see [Add a Transport Node Profile](#).

If you have not used a transport node profile to install NSX, or if you want to remove NSX from a subset of the hosts in the cluster, see [Uninstall NSX from a Managed Host in a vSphere Cluster](#).

Note Follow these instructions to uninstall NSX from a host in a cluster: [Uninstall NSX from a Managed Host in a vSphere Cluster](#).

Prerequisites

- Ensure there are no VIF ports associated with hosts.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Select **System > Fabric > Hosts**.
- 3 On the **Cluster** tab, select a cluster, click **Actions** menu and select **Detach Transport Node Profile**.
- 4 Select all cluster host nodes and select **Remove NSX**.

Note If NSX Intelligence is also deployed on the host, NSX uninstallation fails because all transport nodes become part of a default network security group. To successfully uninstall NSX, you also need to select the **Force Delete** option before proceeding with uninstallation.

- 5 On the Remove NSX window, click **Remove**.
- 6 Verify that the NSX 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 VIBs

```
esxcli software vib list | grep -E 'nsx|vsipfwlib'
```

- 7 (IPv4 only or IPv4 and IPv6 stack) If the host goes into failed state and NSX VIBs cannot be removed, then run the `nsxcli -c del nsx` command to remove NSX from the host.
 - a Before running the `del nsx` command, perform the following steps:
 - If there are VMkernel adapters on NSX port groups on the VDS switch, you must manually migrate or remove vmks from NSX port group to DV port groups on the VDS switch. If there are any vmks available on the NSX port groups, `del nsx` command execution fails.
 - Put the ESXi host in maintenance mode. The VMware vCenter does not allow the host to be put in maintenance mode unless all running VMs on the host are in powered off state or moved to a different host.
 - Permanently disconnect the ESXi host transport node from NSX Manager by stopping `nsx-proxy` service running on the ESX host transport node. Log in to the ESXi CLI terminal and run `/etc/init.d/nsx-proxy stop`.
 - Refresh the NSX Manager UI.
 - Verify that the state of the ESXi host transport node is `Disconnected` from NSX Manager.
 - b Log in to the ESXi CLI terminal.

- c Run `nsxcli -c del nsx`.
- d Read the warning message. Enter **Yes** if you want to go ahead with NSX uninstallation.

```
Carefully read the requirements and limitations of this command:
1. Read NSX documentation for 'Remove a Host from NSX or Uninstall NSX Completely'.
2. Deletion of this Transport Node from the NSX UI or API failed, and this is the last resort.
3. If this is an ESXi host:
    a. The host must be in maintenance mode.
    b. All resources attached to NSXPGs must be moved out.
    c. If this is a SmartNIC-enabled host, the host must be rebooted after vib removal is completed. Verify this by checking /var/run/log/esxupdate.log for the thread which shows nsx-lcp component removal and confirm it completed without error. It will look something like:

        In(14) esxupdate[2150621]: Starting runnable component remove -n nsx-lcp-bundle:4.1.1.0.0-8.0.21958016 with 6e3446d0-8393-5869-8873-076a95930f56
        ...
        Db(15) esxupdate[2150621]: Finished execution of command =
component.remove
If the above conditions for ESXi hosts are not met, the command WILL fail.
4. If this is a Linux host:
    a. If KVM is managing VM tenants then shut them down before running this command.
    b. This command should be run from the host console and may fail if run from an SSH client
        or any other network based shell client.
    c. The 'nsxcli -c del nsx' form of this command is not supported
5. If this is a Windows host:
    NOTE: This will completely remove all NSX instances (image and config) from the host.
6. For command progress check /scratch/log/nsxcli.log on ESXi host or /var/log/nsxcli.log on Linux host or 'c:/Programdata/VMware/NSX/Logs/nsxcli.log' on Windows host.
Are you sure you want to remove NSX on this host? (yes/no) yes
```

Important After running the `del nsx` command, do not use the **Resolve** functionality in the NSX Manager UI to reprepare the host that is in **Disconnected** state . If you use the **Resolve** functionality, the host might go into **Degraded** state.

- e On the ESXi host, verify that the system message displayed is `Terminated`. This message indicates that NSX is completely removed from the host.
- f On a SmartNIC-enabled host, reboot the host once the command removes all required vibs.
- g Go to the ESXi host, select **Force Delete** and begin uninstallation. All existing host switches are removed, transport node is detached from NSX Manager, and NSX VIBs are removed.

- h To verify whether any NSX VIBs still remain on the host, run `esxcli software vib list | grep -E 'nsx|vsipfwlib'`. If you find any VIBs on the host, it means that `del nsx` has failed. When you executed the command, the host connectivity with NSX might have come up.
- i As uninstallation has failed, try to gracefully delete NSX either from the NSX Manager UI or call API.
- j If uninstallation fails again due to host disconnectivity from NSX Manager, repeat the procedure to remove NSX using the `del nsx` command.
- k If uninstallation is still unsuccessful, contact VMware support.

Results

NSX objects and all related services are completely removed from the host. However, if you applied the predefined or a custom high-performance switch profile to a cluster, NSX retains these profile properties on cluster hosts after uninstallation. For more information, see the *Configure High-Performance Host Switch Profiles* topic in the *NSX Administration Guide*.

Uninstall NSX from a Managed Host in a vSphere Cluster

You can uninstall NSX from a single host that is managed by VMware vCenter. The other hosts in the cluster are not affected.

Prerequisites

- On an ESXi host that is put into `Locked` state, ensure that the root user is added to the exception list, so that an SSH session can be established with the host.
- Ensure there are no VIF ports associated with hosts.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Select **System > Fabric > Hosts**.
- 3 On the **Cluster** tab, select a cluster.
- 4 If the cluster has a transport node profile applied, select the cluster, and click **Actions > Detach Transport Node 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 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 VIBs

```
esxcli software vib list | grep -E 'nsx|vsipfwlib'
```

7 If a Transport Node Profile is applied to the cluster, and you want to reapply it, select the cluster, click **Configure NSX**, and select the profile from the **Select Deployment Profile** drop-down menu.**8** (Host on a VDS 7.0 switch) If the host goes into failed state and NSX VIBs cannot be removed, then run the `nsxcli -c del nsx` command to remove NSX from the host.

- a Before running the `del nsx` command, do the following:
 - If there are VMkernel adapters on NSX port groups on the VDS switch, you must manually migrate or remove vmks from NSX port group to DV port groups on the VDS switch. If there are any vmks available on the NSX port groups, `del nsx` command execution fails.
 - Put the ESXi host in maintenance mode. The VMware vCenter does not allow the host to be put in maintenance mode unless all running VMs on the host are in powered off state or moved to a different host.
 - Permanently disconnect the ESXi host transport node from NSX Manager by stopping `nsx-proxy` service running on the ESX host transport node. Log in to the ESXi SSH terminal and run `/etc/init.d/nsx-proxy stop`.
 - Refresh the NSX Manager UI.
 - Verify that the state of the ESXi host transport node is `Disconnected` from NSX Manager.
- b Disable SNMP on the ESXi host.

```
esxcli system snmp set --enable false
```
- c Log in to the ESXi CLI terminal.
- d Run `nsxcli -c del nsx`.

- e Read the warning message. Enter **Yes** if you want to go ahead with NSX uninstallation.

```
Carefully read the requirements and limitations of this command:
1. Read NSX documentation for 'Remove a Host from NSX or Uninstall NSX Completely'.
2. Deletion of this Transport Node from the NSX UI or API failed, and this is the last resort.
3. If this is an ESXi host:
  a. The host must be in maintenance mode.
  b. All resources attached to NSXPGs must be moved out.
   If the above conditions for ESXi hosts are not met, the command WILL fail.
4. For command progress check /scratch/log/nsxcli.log on ESXi host or /var/log/nsxcli.log on non-ESXi host.
Are you sure you want to remove NSX on this host? (yes/no)
```

Important After running the `del nsx` command, do not use the **Resolve** functionality in the NSX Manager UI to re-prepare the host that is in **Disconnected** state . If you use the **Resolve** functionality, the host might go into **Degraded** state.

- f Select each host and click **Remove NSX**.
- g In the pop-up window, select **Force Delete** and begin uninstallation.
- h On the ESXi host, verify that system message displayed is `Terminated`. This message indicates that NSX is completely removed from the host.
- All existing host switches are removed, transport node is detached from NSX Manager, and NSX VIBs are removed. If any NSX VIBs remain on the host, contact VMware support.
 - On a host part of a vSphere Lifecycle Manager, after you perform `del nsx` and **Remove NSX** from NSX Manager, the host status in vCenter Server is compliant with the cluster image. The system displays, `All hosts in the cluster are compliant`.

Results

NSX objects and all related services are completely removed from the host. However, if you applied the predefined or a custom high-performance switch profile to individual transport nodes, NSX retains these profile properties on hosts after uninstallation. For more information, see the *Configure High-Performance Host Switch Profiles* topic in the *NSX Administration Guide*.

Uninstall NSX from a Physical Host

You can uninstall NSX from a physical host.

You can uninstall NSX from a physical host either from the NSX Manager or from the Windows Powershell terminal on Windows hosts or the CLI terminal on Linux hosts.

Prerequisites

If you are uninstalling NSX from a standalone physical host, verify the following settings:

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>` or `https://<nsx-manager-fqdn>`.
- 2 Select **System > Fabric > Hosts**.
- 3 Select the **Standalone** tab.
- 4 Select the physical host, click **Delete**. In the confirmation dialog box, by default **Uninstall NSX Components** is selected. Deselect **Force Delete**, click **Delete**.

The NSX software is removed from the host.

- 5 If the uninstall fails, select the host and click **Delete** again. In the confirmation dialog box, check **Force Delete** and click **Delete**.

The system deletes the host transport node from the management plane, but the host might still have NSX software installed.

Go to the next step only if NSX cannot be uninstalled from the NSX Manager. In the following steps, you will remove NSX from the CLI terminal of physical hosts.

- 6 Before deleting NSX from the CLI terminal, verify whether NSX packages are removed from the host.

On Windows Powershell, run `Get-ItemProperty HKLM:\Software\Microsoft\Windows\CurrentVersion\Uninstall* | Select-Object DisplayName, DisplayVersion, Publisher, InstallDate | Format-Table -AutoSize | findstr NSX`

On Ubuntu CLI terminal, run `Ubuntu: apt list --installed | grep nsx`

On RHEL/SLES/CentOS CLI terminal, run `rpm -qa | grep nsx`

- 7 (Linux physical hosts) If the host goes into a failed state and NSX install bundles cannot be removed, then to forcefully remove NSX from the host, run the `del nsx` command .
 - a Log into the host's command-line interface as root.
 - b Run `nsxcli -c del nsx`.

- c Read the warning message. Enter **Yes** if you want to go ahead with NSX uninstallation.

Carefully read the requirements and limitations of this command:

1. Read NSX documentation for 'Remove a Host from NSX or Uninstall NSX Completely'.

2. Deletion of this Transport Node from the NSX UI or API failed, and this is the last resort.

3. If this is an ESXi host:

a. The host must be in maintenance mode.

b. All resources attached to NSXPGs must be moved out.

If the above conditions for ESXi hosts are not met, the command WILL fail.

4. If this is a Linux host:

a. If KVM is managing VM tenants then shut them down before running this command.

b. This command should be run from the host console and may fail if run from an SSH client or any other network based shell client.

c. The 'nsxcli -c del nsx' form of this command is not supported.

5. If this is a Windows host:

Note: This will completely remove all NSX-T instances (image and config) from the host.

6. For command progress check /scratch/log/nsxcli.log on ESXi host or /var/log/nsxcli.log on non-ESXi host.

Are you sure you want to remove NSX on this host? (yes/no)

Important After running the `del nsx` command, do not use the **Resolve** functionality in the NSX Manager UI to reprepare the host that is in **Disconnected** state. If you use the **Resolve** functionality, the host might go into **Degraded** state.

- d On the physical host, verify that system message displayed is `Terminated`. This message indicates that NSX is completely removed from the host including the application interface that was created for the physical host.

After running `del nsx`, NSX packages and the application interface are removed from the host.

- 8 On Windows physical hosts, if the host goes into a failed state and NSX install bundles cannot be removed, then to forcefully remove NSX from the host, follow these steps.

- a Log into Windows powershell interface as one of the administrators.
- b Go to the NSX directory.

```
PS C:\program files\VMware\nsx\nsx-cli> .\nsxclibms.bat -c del nsx
```

- c Read the warning message. Enter **Yes** if you want to go ahead with NSX uninstallation.

Carefully read the requirements and limitations of this command:

1. Read NSX documentation for 'Remove a Host from NSX or Uninstall NSX Completely'.

2. Deletion of this Transport Node from the NSX UI or API failed, and this is the last resort.

3. If this is an ESXi host:

- a. The host must be in maintenance mode.

- b. All resources attached to NSXPGs must be moved out.

If the above conditions for ESXi hosts are not met, the command WILL fail.

4. If this is a Linux host:

- a. If KVM is managing VM tenants then shut them down before running this command.

- b. This command should be run from the host console and may fail if run from an SSH client or any other network based shell client.

- c. The 'nsxcli -c del nsx' form of this command is not supported.

5. If this is a Windows host:

Note: This will completely remove all NSX-T instances (image and config) from the host.

6. For command progress check /scratch/log/nsxcli.log on ESXi host or /var/log/nsxcli.log on non-ESXi host.

Are you sure you want to remove NSX on this host? (yes/no)

Important After running the `del nsx` command, do not use the **Resolve** functionality in the NSX Manager UI to reprepare the host that is in **Disconnected** state. If you use the **Resolve** functionality, the host might go into **Degraded** state.

After running `del nsx`, the following actions are performed:

- Application Interface on Windows server is uninstalled.
- Transport Node configuration is deleted.
- NSX packages are deleted.

Results

When the NSX software is successfully removed, no packages are listed. If any NSX software packages remain on the host, contact VMware support.

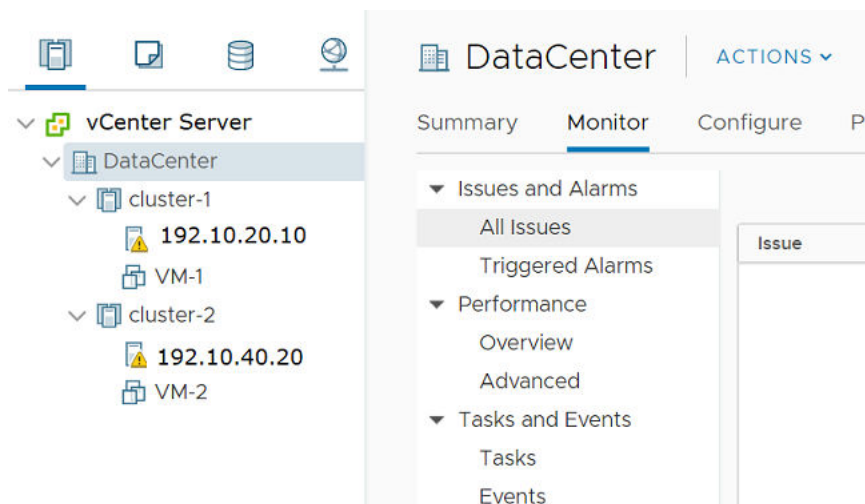
What to do next

Remove the segment port that was connected to the application interface of the physical host.

Triggering Uninstallation from the vSphere Web Client

In the vSphere Web Client, if you move a host from a cluster prepared with a transport node profile to either another cluster, outside of the cluster as a standalone host, or outside of the data center, then NSX is uninstalled on the host that is moved. Such an uninstallation is not triggered when a host that is individually prepared with a transport node configuration is moved.

NSX Uninstallation Scenarios from the vSphere Web Client



Action	Steps/Description	Result
In the VMware vCenter, move an ESXi host in cluster-1 (prepared by applying transport node profile) to the data center as a standalone host (not to another cluster).	<ol style="list-style-type: none"> 1 From the vSphere Web Client, log in to the VMware vCenter. 2 Move the host in maintenance mode. 3 Move the host from cluster-1 that is prepared with a transport node profile out of the cluster as a standalone managed host. NSX triggers uninstallation of the configuration and NSX VIBs. 4 During uninstallation, the transport node is deleted, NSX VIBs are uninstalled. 5 In the NSX UI, the uninstalled host is displayed under Other Hosts on the same VMware vCenter. 	<p>The host is turned into a standalone managed host, which is displayed under 'Other Hosts'. NSX is uninstalled on the host.</p> <p>If the host is under Configuration Mismatch state, then the host remains in that state after it is moved.</p>
In the VMware vCenter, move a prepared host from cluster-1 with transport node profile-1 to cluster-2 with transport node profile-2.	<ol style="list-style-type: none"> 1 From the vSphere Web Client, log in to the VMware vCenter. 2 Move the host in maintenance mode. 3 As cluster-1 is prepared with transport node profile-1, when a host from cluster-1 is moved to cluster-2, then transport node profile-2 is applied to the host. Only the new transport node profile-2 configuration is applied to the host, whereas NSX VIBs are not uninstalled from the host. 4 If the NSX host is in a failed configuration state, then it is not configured after it moves to cluster-2. The host remains in the failed state. 	<p>The host is moved from cluster-1 to cluster-2. A successfully configured host is applied with transport node profile-2.</p> <p>If the host is in the failed state, ensure that the host is successfully configured in NSX.</p>
In the VMware vCenter, move a host that is in the Configuration Mismatch state (NSX state) from cluster-1 with transport node profile-1 to cluster-2 with transport node profile-2.	<ol style="list-style-type: none"> 1 From the vSphere Web Client, log in to the VMware vCenter. 2 Move the host in maintenance mode. 3 As the host is in the Configuration Mismatch state, even though it is moved to cluster-2, transport node profile-2 is not applied to it. 	The host remains in Configuration Mismatch state.
In the VMware vCenter, move a host from cluster-1 with transport node profile-1 to cluster-3 not applied with any transport node profile.	<ol style="list-style-type: none"> 1 From the vSphere Web Client, log in to the VMware vCenter. 2 Move the host in maintenance mode. 3 Move the host from cluster-1 to cluster-3. 	<p>If the NSX host is successfully configured, then NSX uninstallation begins on the host.</p> <p>If the NSX host is in the failed configuration state, then after it moves to cluster-3 the node remains in the failed state.</p>

Action	Steps/Description	Result
In the VMware vCenter, delete a host that is in the Configuration Mismatch state (NSX state) because the host has two different configurations applied to it - transport node configuration and transport node profile configurations.	<ol style="list-style-type: none"> 1 From the vSphere Web Client, log in to the VMware vCenter. 2 Move the host in maintenance mode. 3 Remove the host from the vCenter Server inventory. 	<p>Uninstallation of NSX does not begin because the node configuration was in the Configuration Mismatch state.</p> <p>To ensure that uninstallation begins, ensure that the transport node is configured with a single configuration, either at the host-level or at the cluster-level.</p> <p>After uninstallation, go to the NSX Manager UI and verify that the managed host is moved out of the cluster to become a standalone unmanaged host.</p>
<p>In the VMware vCenter, move a prepared host from cluster-1 with transport node profile-1 applied to:</p> <ul style="list-style-type: none"> ■ Another cluster without any transport node profile applied ■ Data center ■ Outside of the data center 	<ol style="list-style-type: none"> 1 From the vSphere Web Client, log in to the VMware vCenter. 2 Move the host in maintenance mode. 3 Perform one of the actions: <ul style="list-style-type: none"> ■ Move the host to another cluster without any transport node profile applied. ■ Move the host as a standalone host in the data center. ■ Move the host to outside of the data center 	NSX is uninstalled from the host.

Uninstall NSX from a vSphere Lifecycle Manager cluster through NSX Manager

You can trigger uninstallation of NSX on a host that is part of a vSphere Lifecycle Manager cluster through NSX Manager.

Procedure

- 1 From a browser, log in with admin privileges to an NSX Manager at <https://<nsx-manager-ip-address>> or <https://<nsx-manager-fqdn>>.
- 2 Select the cluster and click **Remove NSX**.

vSphere Lifecycle Manager removes the NSX solution applied to the cluster in VMware vCenter.

Important You cannot detach a transport node profile on a vSphere Lifecycle Manager-enabled cluster. You must remove NSX on the cluster.

- 3 If vSphere Lifecycle Manager fails to remove the NSX Solution applied to the cluster in vCenter Server or remove NSX on one or more transport nodes, then vSphere Lifecycle Manager marks the cluster and or transport nodes in `Uninstall Failed` state.
 - a Select the cluster and click **Remove NSX** on the cluster to retry uninstallation.

Note vSphere Lifecycle Manager puts the DPU-backed hosts in maintenance mode and reboots it as part of host remediation. If vSphere Lifecycle Manager fails to place the host in maintenance mode, you need to manually power off all VMs and then retry NSX installation.

Troubleshooting Installation Issues

18

A list of issues related to NSX installation and configuration

Issue	Solution
Installation Fails Due to Insufficient Space in Bootbank on ESXi host.	https://kb.vmware.com/s/article/74864

Read the following topics next:

- [Troubleshooting Installation to check for Basic Infrastructure Services](#)
- [Troubleshoot OVA Deployment and Appliance Bringup](#)
- [Troubleshoot NSX Manager Cluster](#)
- [Troubleshooting Host Transport Nodes](#)
- [Troubleshooting NSX Edge Nodes](#)

Troubleshooting Installation to check for Basic Infrastructure Services

This section provides information about troubleshooting installation issues.

Basic Infrastructure Services

The following services must be running on the appliances and hypervisors, also on vCenter Server if it is used as a compute manager.

- NTP
- DNS

Make sure that firewall is not blocking traffic between NSX components and hypervisors. Make sure that the required ports are open between the components.

To flush the DNS cache on the NSX Manager, SSH as root to the manager and run the following command:

```
root@nsx-mgr-01:~# /etc/init.d/resolvconf restart
[ ok ] Restarting resolvconf (via systemctl): resolvconf.service.
```

You can then check the DNS configuration file.

```
root@nsx-mgr-01:~# cat /etc/resolv.conf
# Dynamic resolv.conf(5) file for glibc resolver(3) generated by resolvconf(8)
#     DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
nameserver 192.168.253.1
search mgt.sg.lab
```

Log in as root user and run `su admin` to launch `nsxcli` on NSX manager. As an admin user, `nsxcli` is default prompt.

Check DNS servers using following `nsxcli` command:

```
get name-servers
```

```
198.10.10.1
198.10.10.2
198.10.10.3
```

Checking Communication from Host to Controller and Manager

On an ESXi host using NSX CLI commands:

```
esxi-01.corp.local> get managers
- 192.168.110.19 Connected

esxi-01.corp.local> get controllers
Controller IP      Port      SSL      Status      Is Physical Master  Session State
Controller FQDN
192.168.110.16    1235     enabled  connected   true
up                NA
```

On an ESXi host using host CLI commands:

```
[root@esxi-01:~] esxcli network ip connection list | grep 1235
tcp          0      0 192.168.110.53:42271          192.168.110.16:1235
ESTABLISHED 67702 newreno nsx-proxy
[root@esxi-01:~]
[root@esxi-01:~] esxcli network ip connection list | grep 5671
tcp          0      0 192.168.110.253:11721        192.168.110.19:5671 ESTABLISHED
2103688 newreno mpa
tcp          0      0 192.168.110.253:30977        192.168.110.19:5671 ESTABLISHED
2103688 newreno mpa
```

Host Registration Failure

If NSX uses the wrong IP address, host registration will fail. This can happen when a host has multiple IP addresses. Trying to delete the transport node leaves it in the Orphaned state. To resolve the issue:

- Go to **Fabric > Nodes > Hosts**, edit the host and remove all IP addresses except the management one.

- Click on the errors and select **Resolve**.

Configuration Error when Deploying an Edge VM

After deploying an Edge VM, NSX Manager shows the VM's status as **configuration error**. The manager log has a message similar to the following:

```
nsx-manager NSX - FABRIC [nsx@6876 comp="nsx-manager" errorCode="MP16027" subcomp="manager"]
Edge 758ad396-0754-11e8-877e-005056abf715 is not ready for configuration error occurred,
error detail is NSX Edge configuration has failed. The host does not support required cpu
features: ['aes'].
```

Restarting the edge datapath service and then the VM should resolve the issue.

Force Removing a Transport Node

You can remove a transport node that is stuck in the Orphaned state by making the following API call:

```
DELETE https://<NSX Manager>/api/v1/transport-nodes/<TN ID>?force=true
```

NSX Manager will not do any validations as to whether you have any active VMs running on the host. You are responsible for deleting the N-VDS and VIBs. If you have the node added through Compute Manager, delete the Compute Manager first and then delete the node. The transport node will be deleted as well.

Troubleshoot OVA Deployment and Appliance Bringup

Troubleshooting OVA Failures

Note During deployment, if you entered incorrect configuration details, delete the appliance and redeploy with correct configuration.

- Verify that the datastore chosen for deployment is mounted on all the hosts that are members of a cluster. Redeploy and choose ESXi host instead of VMware vCenter to bypass VMware vCenter cluster related checks.
- If proxy enabled on VMware vCenter, edit file `/etc/sysconfig/proxy` and add line `.*.domainname` to bypass proxy for ESXi hosts. See, <https://kb.vmware.com/s/article/81565>.
- If deployment of appliance through OVF tool gives error `ovf descriptor not found`, view the file contents in terminal `cat -A <filepath/filename>` and remove hidden formatting characters. Then try again.

Troubleshooting Issues Related to Bringing Up the Appliance

SSH to NSX Manager CLI as admin and run following commands to troubleshoot.

- Run `get configuration` and verify `hostname/name-server/search-domain/ntp` settings are correct.
- Run `get services` and verify all required services are running (other than `nsx-message-bus`, `snmp`, `migration-coordinator`). If these services are not running, try restarting the service by running `restart service <service-name>`.
- Run `get cluster status` and verify all manager cluster components are up. If any component is down, try restarting the service associated to the component by running `restart service <associated-component-service-name>`.
- Run `get core-dumps` to verify no core dumps generated in `/var/log/core` or `/image/core`. If you find any core dumps, contact VMware Support.
- Run `get filesystem-stats` to verify that no disk partition is full, especially those partitions that are consumed by NSX.
- Alternatively, you can run API commands to know the node and service status.

```
GET api/v1/node/status
```

```
GET api/v1/node/services
```

Troubleshoot NSX Manager Cluster

You must configure a three-node NSX Manager cluster and only one of the nodes can fail at any given time for the cluster to self-heal.

For an NSX Manager cluster to self-recover from a node failure, the majority of nodes must not have failed (the number of active nodes must be greater than the failed nodes) otherwise the entire cluster becomes unavailable. It means that all write operations are blocked, all clustering related API/CLIs fail. However, local API/CLI commands continue to work.

NSX Manager logs are written into `/var/log/syslog`, directory.

Admin CLI

To activate Admin CLI, log in as an admin to a NSX Manager. But if you login as root, you can run singleton admin CLI directly from the root shell using the `su admin -c <cmd-to-run>` command. You can switch on interactive mode of admin CLI using the `su admin` command and then run admin commands.

Manually deployed NSX Manager failed to join the NSX Manager Cluster

A NSX Manager that is manually deployed in VMware vCenter using OVF or OVA file failed to join the NSX Manager cluster.

Cause

You used an incorrect thumbprint value to join the new NSX Manager node to the existing NSX Manager cluster. If NSX VIP is configured, then use the cluster thumbprint command. If external VIP or no VIP is configured, use the API thumbprint command. Try registering the NSX Manager node with the cluster using correct thumbprint.

Solution

- ◆ You can manually deploy NSX Manager by running one of the following commands.
 - Run the `get certificate cluster thumbprint` command and then run `Join management-plane <cluster-vip> username <manager-username> password <manager-pwd> thumbprint <cluster-thumbprint>` command. The procedure to join the management node with the cluster is complete. Alternatively, you can use the NSX Manager thumbprint, as shown in the following steps to join the NSX Manager with the cluster.
 - Or run the `get certificate api thumbprint` command and then run the `Join management-plane <manager-ip> username <manager-username> password <manager-pwd> thumbprint <manager-thumbprint>` command.

NSX Manager Cluster status Degraded As Datastore-related Components Are Down

NSX Manager status is degraded due to DATASTORE component and/or CORFU_NONCONFIG component down on one or more manager nodes member of NSX Cluster.

Solution

- 1 SSH to NSX Manager CLI terminal as an admin.
- 2 To identify the manager node with components down, run `get cluster status`.
- 3 Verify underlying datastore is available with recommended disk access latency. To get and fix the disk access latency numbers, see https://kb.vmware.com/s/article/87075?lang=en_US.

Note Datastore outage can cause NSX Manager appliance VMs to go into read-only mode. Linux does not provide a usable utility to recover from this error. Follow the KB, Recover NSX Manager upon storage outage, if appliance reboot does not fix the issue. Verify backend datastore has fully recovered before rebooting the affected NSX Manager VM.

- 4 To verify no disk partition consumed by NSX is full or close-to-full, run 'get filesystem-stats'.

Note Datastore specific logs can be found at `var/log/corfu/corfu.9000.log` and `/var/log/corfu/tanuki.log`.

- 5 Clean up disk space and / or reboot all NSX Manager nodes (after storage issues are resolved) to remove the read-only mode.

- 6 If components continue to be down, contact VMware Support.

Important With the datastore down, the NSX Manager continues to participate in the cluster if its network connectivity is up. This state might result in the management plane and control plane to become unavailable. If datastore issue cannot be resolved, replace the impacted NSX Manager by powering off the impacted NSX Manager (as long as cluster quorum is maintained by having majority number of nodes up).

- 7 Remove NSX Manager node from existing NSX Manager Cluster. On the NSX Manager, run `detach node <node-id>`.
- 8 Verify the problematic node is not a member of the cluster anymore by running the `get cluster status`.
- 9 Deploy a new NSX Manager node and join it to the existing Manager Cluster.

Manager and HTTPS Services Are Down Frequently Due to Incorrect NAT Configuration

If you incorrectly configured NAT, the Manager and HTTPS service might go down frequently.

Solution

- 1 Do not use 0.0.0.0/0, x.0.0.0/8, x.x.0.0/16 subnet ranges in the NAT rules for SNAT Translated IP section (SNAT rule) and for Destination Network section (DNAT rule).
- 2 Fix the incorrect NAT configuration.
- 3 Configure NAT rule where IP address range and count does not exceed 256 for SNAT Translated IP section (SNAT rule) and for Destination Network section (DNAT rule).

NSX Manager is Slow To Load And Tasks Fail

NSX Manager is slow to load and tasks fail with message `server is overloaded or too many requests`.

Cause

This issue occurs upon hitting the nsx manager incoming API rate limit. The solution is to either limit the number of incoming API requests to NSX Manager or modify the default values for Client API rate limit or Client API concurrency limit or Global API concurrently limit from API or CLI.

Solution

- 1 To view current default values, run `get cluster api-service` as admin or `GET API /api/v1/cluster/api-service`.
- 2 To set a new API rate value, run `set cluster api-service client-api-rate-limit` or `PUT /api/v1/cluster/api-service` with updated values.

- 3 To view logs of incoming API requests, go to `/var/log/proxy` directory, reverse-`proxy.log`.

Note Try to avoid configuring the API rate limits. Instead, design your API client to gracefully deal with situations where limits are exceeded.

NSX Manager UI is not loading even when NSX Clustering is up

NSX Manager is slow to load and tasks fail with message `server is overloaded or too many requests`.

Solution

- 1 SSH to one of the NSX Manager that is a member of the cluster.
- 2 Run admin CLI `get cluster status`.
- 3 If cluster status is stable, run `get services` or `get service <service-name>` and verify that the following services are running: `http`, `manager`, `search`, `ui-service`.
- 4 To restart or start the stopped service, run `restart | start service <service-name>`.

NSX Manager cluster is DOWN or UNAVAILABLE if all nodes part of the the NSX Manager cluster is down or majority nodes are down

NSX Manager is down or unavailable if majority of the nodes in the cluster are down.

Problem

NSX Manager UI will fail to load with following error `Some appliance components are not functioning properly. Component health: POLICY:UNKNOWN, MANAGER:UNKNOWN, SEARCH:UNKNOWN, NO` and clustering related commands will fail using the CLI and API.

Solution

- 1 SSH to each of the affected NSX Manager nodes and run following commands:
 - a Run `get filesystem-stats` and verify `/config` and `/image` is not 100% full.
 - b Run `get core-dumps` to verify no cores have gotten generated in NSX Manager.
 - c Verify there was no datastore outage. See [NSX Manager Cluster status Degraded As Datastore-related Components Are Down](#).
 - d Check logs for out-of-memory errors. See `/var/log/proton/proton-tomcat-wrapper.log`

- 2 To restore clustering and UI, any two nodes in a three node cluster must be up. If you are not able to bring any failed node back up, but if there is a healthy node available, then do one of the following steps to restore clustering:
 - Deploy a new manager node (as 4th member node), join the existing cluster and then detach one of the failed nodes using CLI cmd `detach node <node-uuid>` or API `POST /api/v1/cluster/<node-uuid>?action=remove_node`. The commands should be executed from one of the healthy nodes. Alternatively, you can follow the next bulleted point to deactivate the cluster.
 - (Optional) Run the `deactivate cluster` command on active node such that you end up with single node cluster. Now continue to add the new additional nodes to make a 3-member NSX Manager cluster.

Note NSX Manager nodes that are removed from the cluster should be powered off and deleted.

Troubelshoot NSX Appliance or NSX Clustering issues Using APIs

Use APIs to troubleshoot NSX appliance and clustering issues.

Verify Cluster Status

```
GET https://<nsx-manager>/api/v1/cluster/status
```

```
GET https://<nsx-manager>/api/v1/cluster/<node-id>/status
```

```
GET/api/v1/reverse-proxy/node/health
```

Verify CPU Usage

```
GET /api/v1/systemhealth/appliances/process/status
```

```
GET https://<nsx-manager-ip>/api/vi/systemhealth/appliances/<appliance-id>/process/status
```

Verify Network Latency

```
GET https://<nsx-manager>/api/v1/systemhealth/appliances/latency/status
```

```
GET https://<nsx-manager-ip>/api/v1/systemhealth/appliances/<appliance-id>/latency/status
```

Troubleshooting Host Transport Nodes

Host Fails to Install as Transport Node

Problem

When many hosts are prepared simultaneously (in non vLCM way) some hosts may fail in the internal step of downloading nsx lcp bundle to the host. Error seen is "Failed to download NSX components on host."

Solution

- 1 Check the host connectivity and ensure that the "/tmp" folder has enough space.
- 2 If it is fine then check if there are many hosts getting prepared at this time. Wait for host prep of other hosts to complete.
- 3 Follow the usual error resolution steps from the UI or API to re trigger the operation on the failed host.

Accessing the NSX CLI Terminal

After logging into the root shell of an ESXi host, you can do one of the following:

- Use a singleton NSX CLI directly from the root shell by running the `nsxcli -c <cmd-to-run>` command.
- Go into interactive mode by running the `nsxcli` command.

The NSX Syslog directory on the ESXi host is `/var/log/nsx-syslog`.

Transport Node Installation Failure Due To Pending Reboot**Problem**

Transport Node install task fails due to pending reboot.

Cause

This issue occurs if host is pending reboot from previous installation or upgrade of vib on the host.

Solution

- ◆ Run following CLI on the host as root to verify.

```
vim-cmd hostsvc/hostsummary|grep -i reboot requireSecureBoot = <unset>,
rebootRequired = true
```

If *rebootRequired* is set to **true**, reboot the host and then try installation of NSX modules again on the host.

Unable to Reach Transport Node Due to Incorrect Credentials As Host is in Orphaned State

Problem

- Host is not reachable. Cannot complete login due to an incorrect username or password.
- The Transport Node Apply Task fails with error `Node already exists`.
- Host is in orphaned state.

Cause

This issue occurs due to a race condition under heavy traffic load. Run the API (deprecated) `GET /api/v1/transport-nodes/<TN-UUID>/status` or `GET api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id>/state | status`, where default values for *enforcementpoint-id* and *site-id* is default to review the transport node shows transport nodes status unknown and node deployment status shows as `Failed`.

All these cases occur for host TN that did not get cleaned correctly upon initiation of NSX removal task therefore is still registered with NSX Manager.

In this case, the `GET transport node api` and `GET transport node status api` will fail but `GET transport node state api` works and will show failure message `Failed to uninstall the software on host....`

Solution

- ◆ To fix the existence of stale entry, you must forcefully remove NSX from the host and also run the following api to delete stale host entries in the setup.
 - a (NSX Manager UI) On **Hosts** page, select the **Force Delete** option and click **Remove NSX**.
 - b (API) To forcefully delete NSX, run the API, `https://{{MPIP}}/api/v1/transport-nodes/<Transport-Node-UUID>?force=true&unprepare_host=false`.
 - c (API) To remove stale entries, run the API, `https://{{nsx-mgr-ip}}/api/v1/transport-nodes?action=clean_stale_entries`.

Transport node profile Fails to Prepare Transport Nodes Due to Stale Objects

Problem

Transport node profile (TNP) fails to prepare transport nodes due to stale objects in TNP.

Cause

NSX displays the following error when TNP fails to prepare hosts into transport nodes: `The requested object <uuid> could not be found. Object identifiers are case sensitive. Please make the appropriate changes and reapply the configuration.`

This issue can occur if any configuration specified in the TNP does not exist in NSX Manager or VMware vCenter.

Solution

- 1 Search the UUID in NSX Manager to find out the object being referred to in the error.
- 2 Edit the Transport Node Profile to remove the referenced stale object (could be TNP itself, IP Pool, Profile etc) and then retrigger host configuration.

Transport Node Creation is in Partial Success State Due to Logical Switch or Segment Full-sync Realization Error

Problem

Transport Node has `partial success` status with error `LogicalSwitch full-sync realization query skipped`.

Cause

This issue can occur if uplink profile associated with host has teaming defined with no active uplinks or if host switch uplinks is down or if failure in creation of VTEP.

Run transport-node state API to view more info on failure details using one of the following API command:

- (deprecated) `GET api/v1/transport-nodes/<uuid>/state`
- `GET api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id>/state`, where default values for `enforcementpoint-id` and `site-id` is `default`.

Solution

- 1 Edit the uplink profile associated with the transport node and ensure active uplinks are selected for each teaming profile defined.
- 2 Verify PNIC link status is up by running CLI on the host: `esxcli network nic | esxcli network nic get -n <vmnic-name>`.
- 3 If uplink profile is active-active, verify correct number of VTEPS got created (rather than just one).

Run `net-vd12 -l |more` to verify VTEP count is correct, is assigned a valid IP Address, Gateway and state of each VTEP interface is `UP`.

Transport Node status is degraded when its interface is down

Problem

Transport Node status is degraded when its interface is down.

Cause

The NSX Manager declares a transport node to be in a degraded state if any of its interface that is used by host switch is `Down`.

Solution

- 1 To check for the status of interfaces, run `Esxcfg-vswitch -l`.
- 2 To check status of uplinks, run `esxcli network nic list`.
- 3 Fix the PNIC that is down and in use by the VDS.
- 4 Alternatively, remove or replace the PNIC in VDS uplink configuration.

Solution

Transport Node status is Disconnected or Unknown

Problem

ESXi host prepared as a Transport Node status goes into `Unknown` or `Disconnected` state due to lost connection to NSX Manager. NSX displays the following error: `Heart beating between NSX management node and host <uuid> is down`.

Cause

Host infrastructure services being down due to ESXi disks being full or memory leak may result in this condition. If ESXi.

If ESXi disks are full or if there is a memory leak, it can cause certain processes to crash and cause the transport node to go into `Disconnected` state. When you run `admin cli get managers`, NSX might return active manager node if the crash occurred post successful manager registration. When you run `admin cli get controllers`, NSX gives error `Failed to get controller list`.

Solution

- 1 Run `admin cli get core-dumps` to see if any cores got generated (in `/var/core` or `/image/core`) due to service crash.
- 2 If core-dump happens, run `cmd esxtop` to see which NSX process is consuming too much memory and `df -h` to verify disk partitions used by nsx is not full or close to full.
- 3 Run `/etc/init.d/nsx-proxy | nsx-nestdb status` to get status of infrastructure services on the host.
- 4 Clean up the disk space, then start any stopped infrastructure services on the host by issuing command `/etc/init.d/<service-name> start` (as a temporary workaround).
- 5 Open a support case with VMware if you see any cores.

Transport Node is Down as Agent Service is Down

Problem

The agent service is Down.

Cause

Solution

1 To view the host agent status in UI, go to **Host > Transport Node > View Details > Monitor** tab.

2 To view the status of three host agent services (`nsx-cfgagent`, `nsx-opsagent`, `nsx-nestdb`) in CLI, run `/etc/init.d/<service-name> status`.

3 To view the status of agents in API, call these APIs:

(deprecated) `GET api/v1/transport-nodes/<uuid>/status`

`GET api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id>/status`

The default value for `enforcementpoint-id` and `site-id` is **default**.

4 Restart service that is down by running the following CLI command: `/etc/init.d/<service-name> start`.

Transport Node Connectivity to Controller is Down

Problem

The issue is seen when a transport node's connectivity to NSX Manager is Up but its controller is Down. When you run `get managers`, NSX returns active manager node, while `get controllers` does not return any active controller for this transport node, which is in Connected state and its session state is also Up.

Cause

Solution

1 Verify transport node is not in NSX Maintenance Mode via `admin cli get maintenance-mode`.

2 Call one of the following API:

a (Deprecated) `GET API/v1/transport-nodes/<uuid>/status`

b `GET api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id>/status`, where default values for `<site-id>` and `<enforcementpoint-id>` is default.

- 3 Verify FQDN property (used by transport nodes to talk with NSX Manager or controller) is set by running API: `GET /api/v1/configs/management` and view value for `publish_fqdns`.
If FQDN is set, verify the controller FQDN is reachable and FQDN value is being used to by the transport node to talk to controller by first running ICMP ping to controller FQDN followed by `admin cli get controllers` to verify controller FQDN value is getting populated correctly.
- 4 Verify host agent services are running by following host agent troubleshooting step mentioned above.
- 5 Verify `controller.xml` file exists and contains Host transport node as its member: `/etc/vmware/nsx/controller-info.xml`.
- 6 If host is in NSX Maintenance Mode, run `admin cli set maintenance-mode false` or API to take host out of NSX maintenance mode: `POST /api/v1/transport-nodes/<node-id>?action=exit_maintenance_mode`.
- 7 If FQDN set and ICMP ping works for controller FQDN then try unsetting and setting the FQDN property again by running API `PUT /api/v1/configs/management` with value for `publish_fqdns` as **false** and then run the API again with value **true**.
- 8 Start agent services on the host (if any stopped) by running the CLI command `etc/init.d/<service-name> start`.
- 9 If `controller.xml` file has incorrect data, restart `nsx-proxy` service on the host to trigger re-creation of file.

Unable to power on VMs or vMotion VMs on Transport Node

Problem

The Virtual Distributed L2 (VDL2) component is down. This component must be up for NSX to successfully complete VM operations that are attached to a segment on the transport node.

Cause

NSX displays the following error message: `Currently connected network interface "Network adapter 1" uses network 'VM_NETWORK:vd12 down)`

- 1 SSH to host and run following command to verify status of vdl2 component, `net-dvs | grep "component.vdl2"`

```
com.vmware.common.opaqueDvs.status.component.vdl2 = down , propType = RUNTIME
```

- 2 Run `net-vdl2 -l` to verify VTEP interface is assigned valid IP Address, Gateway and status of each interface is UP.
- 3 Run `esxcfg-vswitch -l` to verify minimum MTU of minimum 1600 bytes is setup on the VDS switch used by NSX and uplinks assigned to VTEP interfaces are UP.

- 4 To view host switch information, run one of the following transport node state API:
 - a (deprecated) `GET api/v1/transport-nodes/<uuid>/state`
 - b `GET api/v1/infra/sites/<site-id>/enforcement-points/<enforcementpoint-id>/host-transport-nodes/<host-transport-node-id>/state`, where default values for enforcementpoint-id and site-id is 'default' or `GET api/v1/transport-nodes/<uuid>/state` (deprecated).

Solution

- 1 Ensure that the configuration details entered in these fields are correct:
 - VTEP IP Pool
 - VTEP VLAN
 - VDS MTU
 - Status of assigned PNIC (must be Up)
- 2 If VTEP pool is configured using DHCP, verify that DHCP server is assigning valid IP addresses to the VTEP pool.

Transport Node Tunnels Down

Problem

ESXi hosts that are prepared as transport nodes have their tunnel status as Down.

Cause

Solution

- 1 If all tunnels are down, verify existence of valid VTEP on the host by running the `net-vd12 -l | more` command.
- 2 Verify these configuration details are correct: VTEP state is UP, VTEP count is correct, and a valid VTEP addresses is assigned.
- 3 To find VTEP BFD tunnel sessions that are down on the ESXi transport node, run the CLI command `nsxdp-cli bfd sessions list | grep down`.
- 4 For sessions that are down, validate IP connectivity between the TEPs using ICMP ping.

Note The TEP interfaces on ESXi are instantiated on the vxlan netstack and on the tunnel VRF for NSX Edge nodes. Therefore, initiate the ping from within the vxlan netstack on ESXi and if you have more than one TEP, be sure to specify the source IP address or interface used for the ping.

- 5 Run `ping ++netstack=vxlan <remote address> -I vmk10 .`
- 6 If ping fails, verify underlay connectivity of the network.

- 7 If ping is successful, verify Fabric MTU is configured correctly on the underlay network and within NSX. Use ICMP with the don't fragment bit to test proper delivery for large packets.
- 8 Run `ping ++netstack=vxlan <remote-vtep-ip=-address> -I vmk10 -d -s 1600 .`

Installation Fails Due to Insufficient Space in Bootbank on ESXi Host

NSX installation might fail if there is insufficient space in the bootbank or in the alt-bootbank on an ESXi host.

Problem

On the ESXi host, you might see a similar log (`esxupdate.log`) message:

```
20**-**-**T13:37:50Z esxupdate: 5557508: BootBankInstaller.pyc:
ERROR: The pending transaction requires 245 MB free space,
however the maximum supported size is 239 MB.^@
```

Cause

Unused VIBs on the ESXi host can be relatively large in size. These unused VIBs can result in insufficient space in the bootbank or in the alt-bootbank when installing the required VIBs.

Solution

- Uninstall the VIBs that are no longer required and free up additional disk space.

For more information on deleting the unused VIBs, see the VMware knowledge base article at <https://kb.vmware.com/s/article/74864>.

NSX Agent on ESXi Transport Nodes Times Out Communicating with NSX Manager

In a large-scale environment with many transport nodes and VMs on ESXi hosts, NSX agents, which run on ESXi hosts, might time out when communicating with NSX Manager.

Problem

Some operations, such as when a VM vnic tries to attach to a logical switch, fail.

The `/var/run/log/nsx-opsagent.log` has messages such as:

```
level="ERROR" errorCode="MPA41542"] [MP_AddVnicAttachment] RPC call [0e316296-13-14] to NSX
management plane timeout
2017-05-15T05:32:13Z nsxa: [nsx@6876 comp="nsx-esx" subcomp="NSXA[VifHandlerThread:-2282640]"
tid="1000017079" level="ERROR" errorCode="MPA42003"] [DoMpVifAttachRpc]
MP_AddVnicAttachment() failed: RPC call to NSX management plane timeout
```

Cause

In a large-scale environment, some operations might take longer than usual and fail because the default timeout values are exceeded.

Solution**1** Increase the NSX agent timeout (seconds) value.

- a On the ESXi host, stop the NSX ops agent with the following command:

On NSX 2.3 or later releases:

```
/etc/init.d/nsx-opsagent stop
```

On NSX 2.1 or previous releases:

```
/etc/init.d/nsxa stop
```

- b Edit the file `/etc/vmware/nsx-opsagent/nsxa.json` and change the `vifOperationTimeout` value from 25 seconds to, for example, 55 seconds.

```
"mp" : {
  /* timeout for VIF operation */
  "vifOperationTimeout" : 25,
```

Note This timeout value must be less than the `hostd` timeout value that you set in step 2.

- c Start the NSX ops agent with the following command:

```
/etc/init.d/nsx-opsagent start
```

2 Increase the `hostd` timeout (seconds) value.

- a On the ESXi host, stop the `hostd` agent with the following command:

```
/etc/init.d/hostd stop
```

- b Edit the file `/etc/vmware/hostd/config.xml`. Under `<opaqueNetwork>`, uncomment the entry for `<taskTimeout>` and change the value from 30 seconds to, for example, 60 seconds.

```
<opaqueNetwork>
  <!-- maximum message size allowed in opaque network manager IPC, in bytes. -->
  <!-- <maxMsgSize> 65536 </maxMsgSize> -->
  <!-- maximum wait time for opaque network response -->
  <!-- <taskTimeout> 30 </taskTimeout> -->
```

- c Start the `hostd` agent with the following command:

```
/etc/init.d/hostd start
```

Troubleshooting NSX Edge Nodes

Admin CLI is activated when you log in to an NSX Edge using admin credential. As a root user, you can run singleton admin CLI directly from the root shell `su admin -c <cmd-to-run>` or go into interactive mode of admin CLI `su admin` and then run admin commands.

Accessing the NSX Edge CLI Terminal

Admin CLI is activated if you login to NSX Edge as admin.

As a root user, you can:

- Use a singleton NSX Edge CLI directly from the root shell by running the `su admin -c <cmd-to-run>` command.
- Go into interactive mode by running the `su admin` command.

NSX Edge MPA Connectivity Down

Problem

The NSX Edge MPA connectivity Down due to infrastructure service crash.

Cause

NSX Edge node disks being full or memory leak can cause certain processes to crash and lead to this failure. Admin cli `get managers` may return active manager node (if crash occurred post successful manager registration) and admin cli `get controller` will give error `Failed to get controller list`.

Solution

- 1 Run admin cli `get diagnosis config` OR `GET API /api/v1/transport-nodes/{transport-node-id}/node/diagnosis` to diagnose failures related to health of NSX Edge nodes that are caused when services go down.
- 2 Run admin cli `get cores-dumps` to see if any cores got generated (in `/var/core` or `/image/core`) due to service crash. If core dump is seen, run cmd `top -o %MEM` as root to see which nsx process is consuming too much memory and admin cli `get filesystem-status` to verify if partitions used by nsx is not full or close to full.
- 3 Run root cli `/etc/init.d/nsx-proxy | nsx-nestdb status` to get status of infrastructure services running on the NSX Edge node.
- 4 Clean up the disk space, then start any stopped infrastructure services on the host by issuing command `/etc/init.d/<service-name> start` (as a temporary workaround). Open support case with VMware if any cores are seen.

NSX Edge Status DOWN or DEGRADED As BFD Tunnel(s) are Down

NSX Edge

Problem

NSX Edge status DOWN or DEGRADED due to BFD tunnel(s) to remote NSX Edge down.

Cause

Between two NSX Edge, one BFD session is run on management interface, and one or more BFD session on each VTEP interface. An NSX Edge considers its peer as unreachable only when all BFD sessions to that edges (management one and all VTEP ones) are down.

Solution

- 1 To get information about NSX Edge VTEP devices, run `admin cli get host-switch`.
- 2 To verify physical port status, run `get physical-port <vtep device>`. Then on edge-1> `get phy fp-eth0`.

```
Physical Port
ADMIN_STATUS : up <----- should be "up"
DRIVER       : net_vmxnet3
DUPLEX       : full
ID           : 0
LINK         : up <----- should be "up"
```

- 3 Run `admin cli get diagnosis topology` and `get edge-cluster status` to verify Edge is healthy with edge cluster High Availability State Up, edge node status Up, admin status Up. Then verify if VTEP State Up and status of BFD healthcheck sessions.

```
Interface      : nsx-edge-vtep
  Device       : fp-eth0
  Session      : 71.23.54.3:71.23.54.1
  Status       : Unreachable
  Interface    : nsx-edge-vtep.1
  Device       : fp-eth1
  Session      : 71.23.54.4:71.23.54.2
  Status       : Unreachable
```

If status is unreachable, or Neighbor Signal Down, validate IP connectivity using ICMP Ping.

For all other status, check the BFD error code explanation in the guide. See [View Bidirectional Forwarding Detection Status](#).

- 4 Since TEP interfaces reside on the tunnel VRF for Edges therefore initiate ping from the tunnel VRF 0 on edges and if you have more than one TEP, specify the source IP address or interface used for the ping.
- 5 Run `admin cli 'get logical-routers'` to get tunnel vrf followed by ping.

```
vrf 0
ping 71.23.47.8 source 71.23.46.1 repeat 3
```

- 6 Run `admin cli get neighbor` to check if ARP is getting resolved for the BFD session.
- 7 Run `admin cli get interface` to check status of interface with BFD tunnel down.

- 8 If any of the status is unreachable, verify underlay wiring is correct.
- 9 If ICMP ping is working yet VTEP status is unreachable, verify the VTEP IP addresses are not already in use.

NSX Edge Router High Availability Status is Down

NSX Edge

Problem

Edge Router High Availability Status is Down. NSX Edge status in UI may also show as DOWN.

- After you run the CLI `get edge-cluster status`, the Edge Node Status shows as UP but Routing Status shows as Down.

```
Edge Node Id           : 9e60f8c7-c0ac-42ef-8854-5466cd0cc7eb
Edge Node Status      : Up (Routing Down)
Admin State           : Up
Service Status        :
Datapath Config Channel : Up
Datapath Status Channel : Up
Routing Status Channel : Up
Routing Status        : Down
```

- In the NSX Manager UI, navigate to **Networking** → **Tier-0 Logical Routers** → **Logical Router** → **Overview** → **High Availability State** shows as Down.
- In the Edge CLI, run `get logical-routers → vrf (Tier0 SR) → get high-availability status` of SR shows as Down.

Cause

This issue occurs if all BFP and/or BGP sessions for the router goes down.

Solution

- 1 To troubleshoot BFD sessions, see the previous described troubleshooting cases related to BFD sessions. see [NSX Edge status DOWN or DEGRADED due to BFD tunnels between Edge and ESXi down](#) and [NSX Edge Status DOWN or DEGRADED As BFD Tunnel\(s\) are Down](#).
- 2 To troubleshoot BGP sessions, follow these steps:
 - a Within Tier0 SR vrf (get logical-routers; vrf x), run `admin cli, get bgp neighbor summary`.
 - b If there is no connection established, ping bgp neighbor addresses and source addresses (TO upinks/interfaces) from inside as well as outside SR vrf to verify the interfaces are correctly setup up as BGP peers on the TOR and BGP neighbors are up and accessible.

- c If connection status is established, then run `cli get bgp neighbor <neighbor-ip> advertised-routes` followed by `get route <ip-address> | get route connected | get route bgp` to view if BGP routes are getting advertised.
- d Run `get logical-routers` to view logical routers ID of Tier-0 Service Router followed by `get logical-router <logical-router-id> interfaces stats` to view if TX or RX drops are seen on the Service Router interfaces.
- e Run `admin cli get diagnosis topology` to view status of edge uplink interfaces, bgp peers.

NSX Edge Node Status Down Due to PNIC Bond Status Down

NSX Edge node status is Down because the PNIC bond status is down.

Problem

Cause

This issue can occur if one or more Edge interfaces used by host-switch are down.

Solution

- 1 Verify if uplinks used by fast path interfaces are not down on the physical server hosting the uplinks.
- 2 If fast-path components are down, verify if dataplane service is running by running `admin cli get service dataplane`.
- 3 Verify status of Edge VTEP Device via `admin cli get host-switch` to find out vtep device name.
- 4 Run `get physical-port <vtep device>`.

```
edge-1> get phy fp-eth0
Physical Port
ADMIN_STATUS : up <----- should be "up"
DRIVER : net_vmxnet3
DUPLEX : full
ID : 0
LINK : up <----- should be "up"
```

- 5 If VTEP device `admin_status` or link is down, fix this infrastructure issue.
- 6 If dataplane service is stopped, start by running `start | restart service dataplane`.
- 7 After resolving the issue, run `admin cli get logical-routers` on the edge transport node to make sure edge is functional again.

NSX Edge Transport Node Connectivity to Controller is Down

NSX Edge Transport Node connectivity to controller is down.

Problem**Cause**

This issue is seen when connectivity to manager is up but connectivity to controller is down. Admin cmd 'get managers' returns active manager node while cmd 'get controllers' does not return any active controller for this transport node with status connected and/or session-state UP.

Solution

- 1 Verify transport node is not in NSX Maintenance Mode using admin cli `get maintenance-mode` or run the API, `GET api/v1/transport-nodes/<tn-uuid>| state | status`.
- 2 Verify if FQDN property (used by transport-nodes to talk with NSX Manager/Controller) is set by running API, `GET /api/v1/configs/management` and view value for `publish_FQDNS`.
- 3 If FQDN set, verify the controller FQDN is reachable and FQDN value is being used to by TN to talk to controller by first running ICMP ping to controller FQDN followed by admin cli `get controllers` to verify controller FQDN value is getting populated correctly.
- 4 Verify node agent services are running by following node agent troubleshooting step outlined before.
- 5 If edge transport node is in NSX Maintenance Mode, run admin cli `set maintenance-mode false` or API `POST /api/v1/transport-nodes/<node-id>?action=exit_maintenance_mode` to take node out of NSX Maintenance Mode
- 6 If FQDN is set and ICMP ping works for controller FQDN then try unsetting and setting the FQDN property again by running API `PUT /api/v1/configs/management` with value for `publish_fqdns false` followed by `true`.
- 7 Verify that agent services are running on the NSX Edge node by running `get edge diagnosis config`. If any service shows as failed, restart by running admin cli `start service <service-name>` or root cli `etc/init.d/<service-name> start`.

NSX Edge node status is Down As Controller Is Unavailable

NSX Edge node status is Down due to controller unavailability.

Problem

NSX Edge node status is Down with controller connectivity status 'unavailable' because agent service down.

Cause

The `get manager` and `get controller` commands will return correct values but `GET api/v1/transport-nodes/<tn-uuid>/status` shows agent-status is Down. The error is due to the service `nsx-opsagent` is not running.

Solution

- 1 Run `admin cli get status nsx-opsagent` to verify the service status.
- 2 Start or restart the service by running `admin cli, restart | start service nsx-opsagent`

Transport Node Failed as Named Teaming Defined with No Active Uplink

Transport Node `configuration has failed` due to named teaming defined with no active uplink.

Problem**Cause**

This issue can occur if uplink profile associated with transport-node has named teaming defined with no active uplinks.

Solution

- 1 Make sure active uplinks are selected for any teaming policies defined.
- 2 Edit the uplink profile associated with the transport node by either removing the teaming with no active uplink or assigning a valid uplink to teaming policy defined.

NSX Edge Transport Node goes into NSX Maintenance Mode On HA Failover

NSX Edge Transport Node goes into NSX Maintenance Mode automatically upon HA failover.

Problem

NSX Edge Transport Nodes may go automatically into NSX Maintenance Mode if there are issues with its datapath or with the usage of heap memory.

Cause

To view edge node high availability status, status changes and reasons for it, run `admin cli get edge-cluster status` and `get edge-cluster history state`. If the Edge STATE would be DOWN, implying either datapath process is not running or physical link is down or vtep tunnels are down.

Solution

- 1 Run `admin cli get diagnosis config` and `get service dataplane` to verify core services are up.
- 2 Run `admin cli get diagnosis topology` to view detailed edge config state.
- 3 Run `admin cli get host-switch` to get vtep-device name and physical-port-name.

- 4 Run `admin cli get physical-port <port-name>` to view status of host switch port followed by `get physical-port <interface-name> stats` and look for `rx_misses` (ingress buffer) or `tx_drops` (egress buffer) counter to determine occurrence of packet loss. Packet loss may be seen if edge is flooded with traffic rate higher than the datapath CPUs can process. Packets are first held in input/egress buffer and gets dropped if buffers are full. To check the current buffer size configuration, use the CLI `get dataplane | find ring`.
- 5 If dataplane service is stopped, start by issuing `cmd start service dataplane` (as a temporary workaround).
- 6 If host switch port is down, start by issuing `cmd set physical-port fp-eth0 state up` (as a temporary workaround).
- 7 If packet loss is seen or issue with host switch status/state, file a ticket with VMware Support Desk.

You can also try to change the rx/tx buffer configuration (to enhance edge interface traffic management capacity) using the CLI `set dataplane ring-size <rx/tx> <size>`. The supported buffer size range is **128-4096** bytes and dataplane service needs to be restarted in order to make the new configuration effective resulting in a downtime of about **60** seconds.

- a For example, `set dataplane ring-size rx 2048`. Restart dataplane service for the change to take effect.
- b `set dataplane ring-size tx 2048`. Restart dataplane service for the change to take effect.

```
restart service dataplane
get dataplane | find ring
Bfd_ring_size      : 512
Lacp_ring_size     : 512
Learning_ring_size : 512
Livetrace_ring_size : 512
Rx_ring_size       : 2048
Slowpath_ring_size : 512
Tx_ring_size       : 2048
```

VMotion of NSX Edge VM Fails Due To ESXi Running Out Of Resources

VMotion of NSX Edge VM may fail in vCenter due to ESXi running out of resources from a shared buffer pool.

Cause

This issue can occur for an NSX Edge of XL VM form factor of 16 vCPU cores with error `reservation failed`.

Solution

- ◆ To fix the issue, increase the P2M buffer slots for all the virtual machines on the host. For more information, see the Kb, <https://kb.vmware.com/s/article/76387>.

State Not Consistent of NSX BFD BGP or HA Functionality

State Not Consistent of NSX BFD BGP or HA Functionality

Problem

Pings are not consistent. BFD status changes occur. Edge HA failovers occur consistently.

Cause**Solution**

- 1 Run `admin cli Get edge-cluster history state`.

The CLI showcases history of edge cluster HA transition and reason. The only acceptable state is Active. Anything else means edge node is down and no service is able to function. For the purpose of identifying events leading to such HA state, focus on the timestamp and "Reason" line.

- 2 Run `admin cli get diagnosis topology` to view status of all attached interfaces and config.

NSX Edge Node Status Down and BFD Tunnel Data is Missing**Cause**

This issue can occur if NSX Edge dataplane service goes down due to memory leak in certain NSX Edge processes. In this case, certain NSX Edge CLIs may fail with `error encountered`.

Solution

- 1 Run `admin cli get service dataplane` to view the status of service.
- 2 If status is down, run `cli start service dataplane`.
- 3 Once service has started, run the following CLIs to view the CPU and Memory consumption by NSX Edge dataplane module.
- 4 Get dataplane cpu stats from `/var/log/vmware/top-cpu.log` and `top-mem.log` to watch for values going above 60% usage.

NSX Edge status DOWN or DEGRADED due to BFD tunnels between Edge and ESXi down

NSX Edge

Problem**Cause****Solution**

- 1 Run `admin cli get bfd-sessions` to find VTEP BFD tunnel sessions that are down on the edge transport node. As a guard against VTEP misconfiguration, when edge loses all its BFD to hypervisor, it brings down itself.
- 2 For down sessions, validate IP connectivity between the TEPs using ICMP ping.

Note The TEP interfaces reside on the tunnel VRF for Edges and vxlan netstack for ESXi. Therefore initiate ping from the tunnel VRF on edges and if you have more than one TEP, be sure to specify the source IP address or interface used for the ping.

- 3 Run `admin cli 'get logical-routers'` to get tunnel vrf.

```
vrf 0
  ping 48.13.47.8 source 48.13.46.1 repeat 3
```

- 4 ESXi: `ping ++netstack=vxlan <remote-vtep-ip==address> -I vmk10 -d -s 1600`
- 5 As shown in the CLI, the ping should be initiated by specifying both the remote and local IP address relevant to the BFD session. While not required to test simple connectivity, specify the payload size that's 100 bytes less than the TEP's configured MTU, and the `dfbit` set to enable to prevent the underlay network from fragmenting the packet. Testing with bigger payloads will validate that your underlay network has been properly setup to support your NSX Geneve overlay configuration.
- 6 Validate if ARP is getting resolved for neighbor VTEP address.

```
edge-1(vrf)> get neighbor
Logical Router
UUID      : 736a80e3-23f6-5a2d-81d6-bbefb2786666
VRF       : 0
LR-ID    : 0
Name     :
Type     : TUNNEL
Neighbor
  Interface : 4d9091fe-b971-5d3c-9201-4cb9c7f455fe
  IP        : 202.1.1.2 <----- peer TN VTEP IP
  MAC      : 00:50:56:a6:7d:9b <---- resolved
  State    : reach <----- ARP reachable state
  Timeout  : 37
```

- 7 Run `get interface cmd` followed by `get logical-router interface <uuid> status` to get VTEP interface status.

```
Interface   : ac80718b-72d3-5028-bb07-8f3c4ea2231a
Ifuid      : 258
Name       :
```

```
Fwd-mode      : IPV4_AND_IPV6
Internal name : uplink-258
Mode          : lif
Port-type     : uplink
IP/Mask       : 71.23.46.1/24
MAC           : 00:50:56:b8:2c:c4
VLAN          : 2046
Access-VLAN   : untagged
LS port       : d31578e5-bc91-5466-97c1-8e4a6aa1b2e8
Urpf-mode     : PORT_CHECK
DAD-mode      : LOOSE
RA-mode       : RA_INVALID
Admin         : up
Op_state      : up
Enable-mcast  : True
MTU           : 8800
arp_proxy
```

- 8 Run `get bfd-session stats` to look for RX drops and TX misses counter values.
- 9 If ICMP ping fails or ARP is not reachable, verify your underlay connectivity and peer host TEP interface address is valid. If large packet MTU ping fails, fix the NSX Fabric and/or underly infrastructure MTU to correct values.
- 10 Validate Edge TEP address is not in use by another transport node and validate Edge TEP VLAN and Host TEP VLAN are not using the same VLAN and uplink.