NSX-T Data Center Migration Guide

Modified on 13 AUG 2024
VMware NSX-T Data Center 3.2

vmware by Broadcom

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

Contents

NSX-T Data Center Migration Guide

The *NSX-T Data Center Migration Guide* provides information about migrating a VMware NSX for vSphere™ environment to a VMware NSX-T Data Center™ environment.

For the NSX 4.0 version of this guide, see NSX Migration Guide.

Intended Audience

This manual is intended for anyone who wants to migrate an NSX for vSphere environment or vSphere networking to an NSX-T Data Center environment. The information is written for experienced network and system administrators who are familiar with virtual machine technology and data center operations.

Migration Modes

1

There are multiple ways you can migrate NSX-V to NSX-T.

The following standard migration modes are available:

- Migrate NSX for vSphere
  - Fixed Topology

    Fixed topology migration is limited to the topologies supported in Fixed Topologies Supported for End-to-End Migration.
  - User-defined Topology

    User-defined topologies allow migrating any topology from NSX-V. You create the layer-3 topology on NSX-T and maps the NSX-v Edge Services Gateways (ESGs) and Distributed Logical Routers (DLRs) to NSX-T tier-0 or tier-1 gateways.

    There are two modes available when you choose user-defined topology:
    - Complete Migration
    - Configuration Migration
    - Configuration and Edge Migration (available in NSX-T 3.2.2 as a tech preview feature)

  The folowing screen shows the migration modes that are available when you select NSX for vSphere:

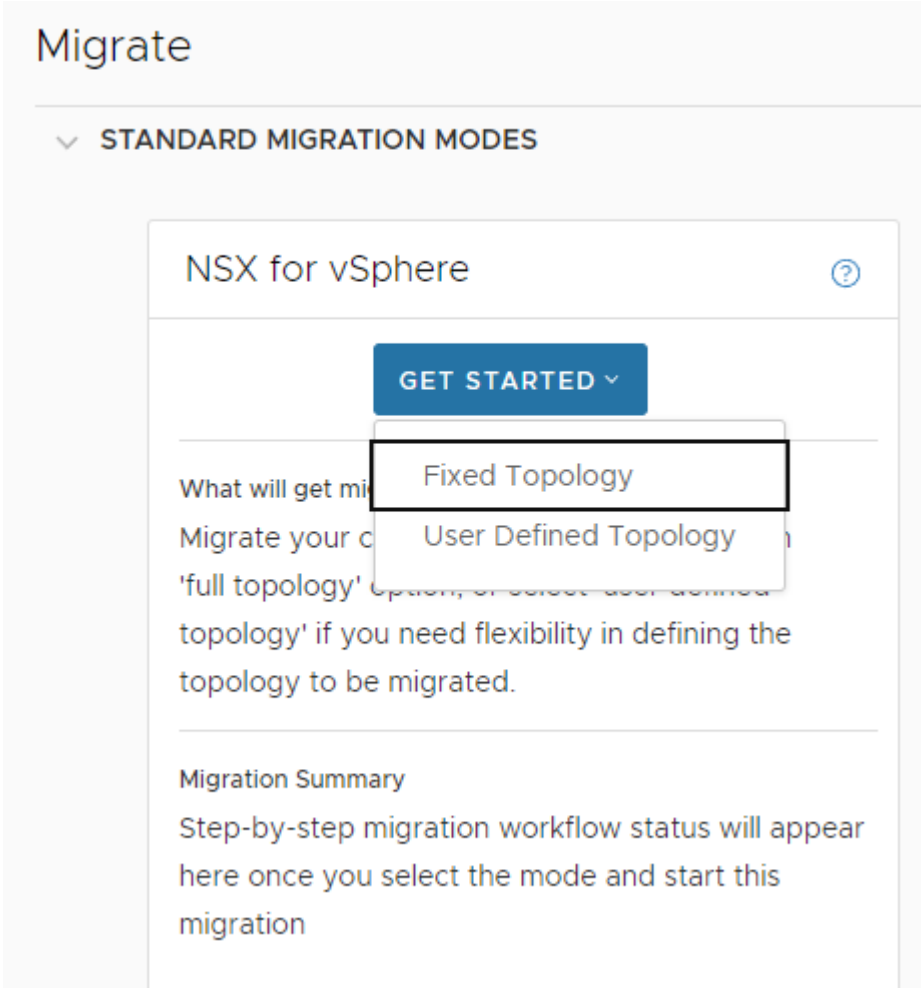

The folowing screen shows the migration modes that are available when you select User-defined Topology:

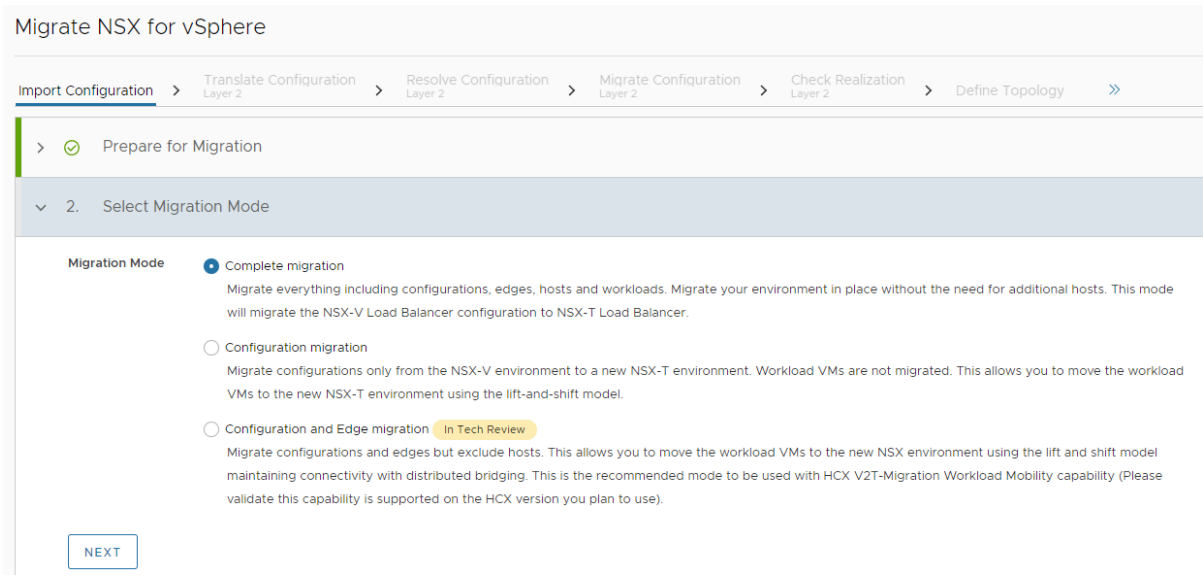

The following table shows the type of migration for each mode:

Mode	Migration Type
Fixed Topology	End-to-end
User-Defined Topology + Complete Migration	End-to-end
User-Defined Topology + Configuration Migration	Lift-and-shift
User-defined Topology + Configuration and Edge Migration	Lift-and-shift. Supports HCX for workload migration.

An end-to-end migration will migrate the whole NSX-V environment. A lift-and-shift migration will migrate configurations only, such as firewall or load balancer. After the configurations are migrated, you migrate workload VMs using vMotion and a layer-2 bridge between the NSX-V and NSX-T environments. You can also do a lift-and-shift migration using the advanced migration mode "Migrate Distributed Firewall (DFW)".

- Migrate vSphere Networking

  In this mode, you migrate vSphere Distributed Switch (VDS) 6.5.0 and 6.6.0 to NSX Virtual Distributed Switch (N-VDS). For more information, see Chapter 9 Migrating vSphere Networking.

- Migrate NSX for vSphere with vRealize Automation

  Similar to the "Migrate NSX for vSphere" mode, you can select Fixed Topology or User Defined Topology. Before migrating a user-defined topology, you must check the vRealize Automation documentation to ensure that your version of vRealize Automation supports the migration of a user-defined topology. For more information about this migration mode, see Chapter 10 Migrating NSX-V with vRealize Automation - Fixed Topology or Chapter 11 Migrating NSX-V with vRealize Automation - User-Defined Topology.

The following advanced migration modes are available:

- Migrate Edge Cutover

  In this mode, north-south traffic is migrated from NSX-V to NSX-T. For more information, see Migrating North-South Traffic to NSX-T Edges Using Edge Cutover.

- Migrate Distributed Firewall (DFW)

  In this mode, you do a lift-and-shift migration of the distributed firewall. For more information, see Chapter 7 Migrating Distributed Firewall Configuration.

- Migrate Distributed Firewall, Host and Workload

In this mode, you do an in-place migration of the distributed firewall, hosts, and workload VMs. For more information, see Chapter 8 In-Place Migration of Specific Parts of NSX-V.

Note If ESGs are present in the NSX-V environment, it is recommended that you choose the "Migrate NSX for vSphere" mode (or the "Migrate NSX for vSphere with vRealize Automation" mode if appropriate). This is the optimal way to migrate topologies with ESGs. When you choose the "NSX for vSphere" migration mode, you can still do a lift-and-shift migration if you choose "User-Defined Topology + Configuration Migration" or "User-defined Topology + Configuration and Edge Migration."

Important If the migration mode involves the migration of ESXi hosts, note the following:

- During the migration or before removing NSX-V, do not perform any lifecycle operation, or restart NSX-V Manager or vCenter server. Lifecycle operations include, but are not limited to, upgrading or patching vCenter server, NSX-V Manager or NSX-T Data Center, renewing certificates, or changing passwords.
- NSX-V should be removed as soon as possible after the migration has completed successfully.
- If you need to perform lifecycle operations, you must do so before the start of the migration.

Important For all migration modes, you must run the migration from a single NSX Manager node. If you have an NSX Manager cluster, start the `migration-coordinator` service on only one NSX Manager and always access that NSX Manager's UI using its IP address or host name. Do not use the manager cluster's virtual IP (VIP).

Migrating a User-Defined Topology

In the **Migrate NSX for vSphere** mode, if you choose **User Defined Topology**, you have the following options:

- **Complete Migration** - This will migrate everything (configurations, Edge Services Gateways, Distributed Logical Routers, hosts and workloads) in place without the need for additional hosts. The NSX-V load balancer will be migrated to an NSX-T load balancer.
- **Configuration Migration** - This will migrate configurations only. After the migration, you can migrate the workload VMs using vMotion. The NSX-V load balancer will be migrated to NSX-T Advanced Load Balancer (ALB). This is the only way to migrate the NSX-V load balancer to ALB.

- **Configuration and Edge Migration** (available in NSX-T 3.2.2 as a tech preview feature) - This will migrate configurations, bridge the NSX-V logical switches to their corresponding NSX-T segments, and migrate Edge nodes for north-south traffic cutover. You can migrate workload VMs after the Edge nodes are migrated. This mode supports the HCX V2T Migration Workload Mobility capability.

Note In NSX-T 3.2.0 and 3.2.1, when migrating a single-site environment, the NSX-V load balancer will be migrated to NSX-T Advanced Load Balancer (ALB). Starting with NSX-T 3.2.2, The NSX-V load balancer will be migrated to an NSX-T load balancer.

Migrating a Cross-vCenter Environment to NSX Federation

Starting with NSX-T 3.2.1, you can migrate an NSX-V cross-vCenter environment to an NSX Federation environment in NSX-T. You must perform the migration from the Global Manager, choose the **Migrate NSX for vSphere** mode and select **User Defined Topology**. You can then choose either **Complete Migration** or **Configuration Migration**. Migrating a cross-vCenter environment to NSX Federation is not supported in any other migration mode. Also, in NSX-T 3.2.1, migrating a cross-vCenter environment to NSX Federation does not support migrating the NSX-V load balancer. Starting with NSX-T 3.2.2, The NSX-V load balancer will be migrated to an NSX-T load balancer.

Migrating Identity Firewall (End-to-End and Lift-and-Shift)

2

If you plan to migrate Identity Firewall (IDFW), some preparations are required.

Before the migration, make sure that the following requirements are met:

- The Active Directory (AD) domains registered in NSX-V are registered in NSX-T.
- The LDAP servers registered in NSX-V are registered in NSX-T.
- The event log servers registered in NSX-V are registered in NSX-T.
- A successful full sync for each newly registered AD domain is completed in NSX-T.
- The IDFW environment in NSX-V is supported by NSX-T. For more information, see the topic Identity Firewall Supported Configurations in the *NSX-T Data Center Administration Guide*.

Note the following:

- During the migration, do not allow new users to log in.
- Some IDFW rules in NSX-V are not supported in NSX-T. Those rules cannot be migrated to NSX-T. You must skip or change them to continue the migration.
- For IP-based IDFW connections, users must re-login after the migration for IDFW to work. If you want IDFW connections for these users to be maintained during the migration, you must manually create shadow firewall rules for these users.
- For SID-based IDFW connections, users do not need to re-login for IDFW to work.
- In NSX-T, IDFW can be configured on a global level and on a cluster level. Because NSX-V does not support IDFW on a cluster level, after the migration, IDFW will be enabled for all clusters in NSX-T.
- You must manually undeploy Guest Introspection (GI) in NSX-V after the migration if GI is not undeployed by other migration operations.

Creating and deleting a shadow firewall rule

To create a shadow firewall rule, after the configuration is imported, do the following in NSX-T:

1 Create an IP set for the directory group.

2 Add the IP set to the same NSGroup that the directory group belongs to.

3 Find the IP addresses of the VMs that users are logged in to.

4 Add the IP addresses to the IP set.

After the VMs are migrated and the users are logged out of the VMs, do the following:

1 Remove the IP addresses from the IP set.

2 After all the IP addresses are removed from the IP-Set, remove the IP set from the NSGroup and delete the IP-set.

Preparing the NSX-V and NSX-T Environments for a User-Defined Topology Migration

3

Before migrating a user-defined topology, you must prepare the NSX-V and NSX-T environments.

If you are migrating an NSX-V environment without vRealize Automation, see Preparing for a User-Defined Topology Migration.

If you are migrating an NSX-V environment with vRealize Automation, see Preparing for an NSX-V with vRA User-Defined Topology Migration.

Performing Post-Migration Tasks

4

For some migration modes, you must perform post-migration tasks after the migration is completed and the NSX-T environment is working as expected.

Migration Mode	Post-Migration Tasks
Migrate NSX for vSphere, Fixed Topology	See Post-Migration Tasks.
Migrate NSX for vSphere, User-Defined Topology, Complete Migration	See Post-Migration Tasks in End-to-End Migration
Migrate NSX for vSphere with vRealize Automation, Fixed Topology	See Post-Migration Tasks
Migrate NSX for vSphere with vRealize Automation, User-defined Topology	See Post-Migration Tasks

Migrating a Fixed Topology

5

If your NSX-V environment has a fixed topology that is supported for migration, you can do an end-to-end migration that is simpler than migrating a user-defined topology.

This migration uses the existing hardware of your NSX-V environment. You do not need extra servers for the migration. The end-to-end migration will migrate the following:

- Logical network topology
- Logical networking and security configurations
- Edges
- Hosts
- Workload VMs

To minimize disruption during migration, ensure that:

- NSX-V and NSX-T edges are on different ESXi hosts.
- Workload VMs directly connected to an Edge Services Gateway (ESG) are on a different ESXi host than the ESG.

Read the following topics next:

- Finish the Migration
- Post-Migration Tasks

Overview - Migrating a Fixed Topology

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.
- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.
- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 5-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.

Table 5-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.

NSX-V Configuration	Supported	Details
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. Migration Coordinator will only migrate from an NSX-V Manager with the role of Primary or Standalone. You can modify the NSX-V environment by changing the status of the secondary managers in order to migrate each NSX-V environment independently.
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	
Network I/O Control (NIOC) having vNIC with reservation	No	
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	
Host profiles	No	
ESXi lockdown mode	No	Not supported in NSX-T.

NSX-V Configuration	Supported	Details
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.
VXLAN port number other than 4789	No	
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	

NSX-V Configuration	Supported	Details
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration

Details	Supported	Notes
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes for in-place migration No for lift-and-shift migration	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.

NSX-V Configuration	Supported	Details
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	

NSX-V Configuration	Supported	Details
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled

NSX-V Configuration	Supported	Details
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network

NSX-V Configuration	Supported	Details
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.

NSX-V Configuration	Supported	Details
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	

NSX-V Configuration	Supported	Details
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.

NSX-V Configuration	Supported	Details
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 5-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 5-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.

Table 5-3. DHCP Features (continued)

NSX-V Configuration	Supported	Details

DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration.
		For example, dhcp option '80' is not migrated.
		<pre><dhcpOptions>
		<other>
		<code>80</code>
		<value>2f766172</value>
		</other>
		</dhcpOptions></pre>

Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 5-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.

NSX-V Configuration	Supported	Details
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ VMs not on NSX-V hosts	No	NSX-T does not support referencing objects not connected to NSX-T portgroups or networks. Those objects will be lost from the source or destination when migration is completed. To avoid this issue use IP addresses in NSX-V to reference those objects before migration.
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group

NSX-V Configuration	Supported	Details
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.

NSX-V Configuration	Supported	Details
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	

NSX-V Configuration	Supported	Details
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 5-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of ‘Groups’ after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 5-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.

Table 5-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.

Table 5-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 5-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 5-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 5-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 5-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000

Table 5-10. Single-Site Limits (continued)

Feature	Limit
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 5-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Fixed Topologies Supported for End-to-End Migration

If your NSX-V environment has a topology that is the same as one of those described below, you can migrate it end to end by choosing the Fixed Topology option.

When you migrate a fixed topology, the NSX-V load balancer is migrated to NSX-T load balancer. To migrate to NSX-T Advanced Load Balancer (ALB), you must migrate using a user-defined topology. For more information, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

Support for firewall is independent of the topology. Every topology listed below supports the following:

- NSX Manager
- Distributed Firewall
- Service Composer
- Grouping Objects

Unsupported Features

In all topologies, the following features are not supported:

- IP Multicast.
- IPv6.
- SSL VPN

For detailed information about which features and configurations are supported, see Detailed Feature Support for Migration.

ESG with High Availability and L4-L7 Services (Topology 1)

This topology contains the following configurations:

- A Distributed Logical Router (DLR) peering with Edge Services Gateway (ESG).
- ECMP is not configured.

- The ESGs are in a high availability configuration.
- BGP, OSPF or static routing is configured between the ESG and top-of-rack (ToR) northbound routers. If BGP is configured, all ESGs must be configured with the same global BGP settings.
- The ESGs can be running L4-L7 services:
  - VPN, NAT, DHCP server, DHCP relay, DNS forwarding, Edge Firewall are supported services.
  - Load balancer is not supported in this topology.

About migrating DHCP relay:

- Although DHCP relay can be configured on either ESG or DLR, only DHCP relay on DLR will be migrated.
- In this topology, if DHCP relay is running on the DLR, and DHCP server is running on the ESG, both DHCP relay and DHCP server will be migrated to the same NSX-T gateway. They will not be migrated separately.

After migration, this configuration is replaced with a tier-0 gateway.

- The tier-0 gateway service router is in active/standby mode.
- The IP addresses of the DLR interfaces are configured as downlinks on the tier-0 gateway.
- The BGP, OSPF or static routing configuration of the ESG is translated to a BGP, OSPF or static routing configuration on the tier-0 gateway.

  Note When static routing is used, the NSX-T HA Virtual IP (VIP) address is not configured automatically. You must add the NSX-T HA VIP address manually after the migration.

- Supported services are migrated to the tier-0 gateway.

Note Depending on your configuration, you might need to provide new IP addresses for the tier-0 gateway uplinks. For example, on an ESG, you can use the same IP address for the router uplink and for the VPN service. On a tier-0 gateway, you must use the different IP address for VPN and uplinks. See Example Configuration Issues for more information.

Figure 5-1. Topology 1: Before and After Migration

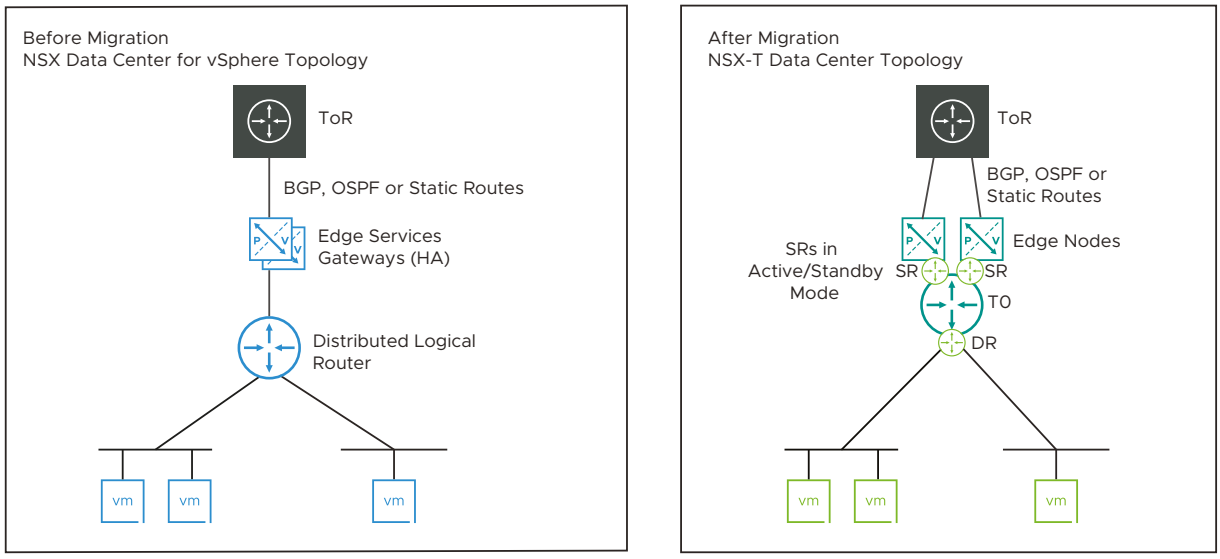

ESG with No L4-L7 Services (Topology 2)

This topology contains the following configurations:

- The DLR has ECMP enabled and peers with multiple ESGs.
- BGP, OSPF or static routing is configured between the ESG and top-of-rack (ToR) northbound routers. If BGP is configured, all ESGs must be configured with the same global BGP settings.
- If BGP is configured between the DLR and ESG, all BGP neighbors on the DLR must have the same weight.
- The ESGs must not be running L4-L7 services.

After migration, this configuration is replaced with a tier-0 gateway.

- The tier-0 gateway service router is in active/active mode.
- The IPs of the DLR interfaces are configured as downlinks on the tier-0 Gateway.
- The BGP or OSPF configuration of the ESGs is translated to a BGP or OSPF configuration, respectively, on the tier-0 gateway. Route redistribution configuration is translated.
- Static routes from ESGs and DLRs are translated to static routes on the tier-0 gateway.

Figure 5-2. Topology 2: Before and After Migration

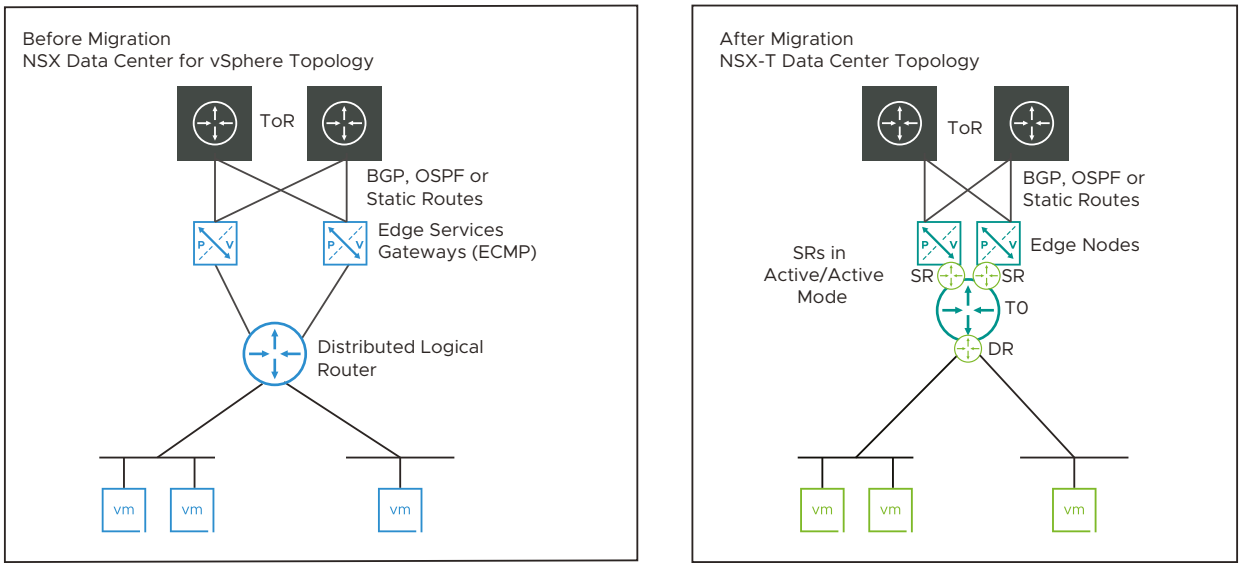

Two Levels of ESG with L4-L7 Services on Second-Level ESG (Topology 3)

The topology contains the following configurations:

- Two levels of ESGs with DLR.
- The first-level (ToR-facing) ESGs must not be running L4-L7 services.
- BGP, OSPF or static routing is configured between the first-level ESGs and top-of-rack (ToR) northbound routers. If BGP is configured, all ESGs must be configured with the same global BGP settings.
- The first-level ESGs have ECMP enabled and peer with the second-level ESGs.
- The second-level ESGs can run L4-L7 services:
  - NAT, DHCP server, DHCP relay, DNS forwarding, inline load balancer, and Edge firewall are supported.
  - VPN is not supported.

About migrating DHCP relay:

- Although DHCP relay can be configured on either ESG or DLR, only DHCP relay on DLR will be migrated.
- In this topology, if DHCP relay is running on the DLR, and DHCP server is running on the ESG, both DHCP relay and DHCP server will be migrated to the same NSX-T gateway. They will not be migrated separately.

After migration, this configuration is replaced with a tier-0 gateway and a tier-1 gateway.

- The first-level ESGs are replaced with a tier-0 gateway. The service router is in active/active mode.
- The IPs of the first-level ESG uplinks are used for the tier-0 gateway uplinks.

- The tier-0 gateway peers with northbound routers (ToR) using BGP or OSPF.
- The second-level ESGs are translated to a tier-1 gateway, which is linked to the tier-0 gateway.
- The IPs of the DLR interfaces are configured as downlinks on the tier-1 Gateway.
- Any services running on the second-level ESG are migrated to the tier-1 gateway. The active/passive Service Routers on the tier-1 gateway use the same Edge nodes that are used for the tier-0 gateway.
- The BGP or OSPFconfiguration on the first-level ESGs is translated to a BGP or OSPF configuration, respectively, on the tier-0 gateway. Route redistribution configuration is translated.
- Static routes from ESGs and DLRs are translated to static routes on the tier-0 gateway. Static routes between the DLR and second-level ESGs are not needed, and so are not translated.

Figure 5-3. Topology 3: Before and After Migration

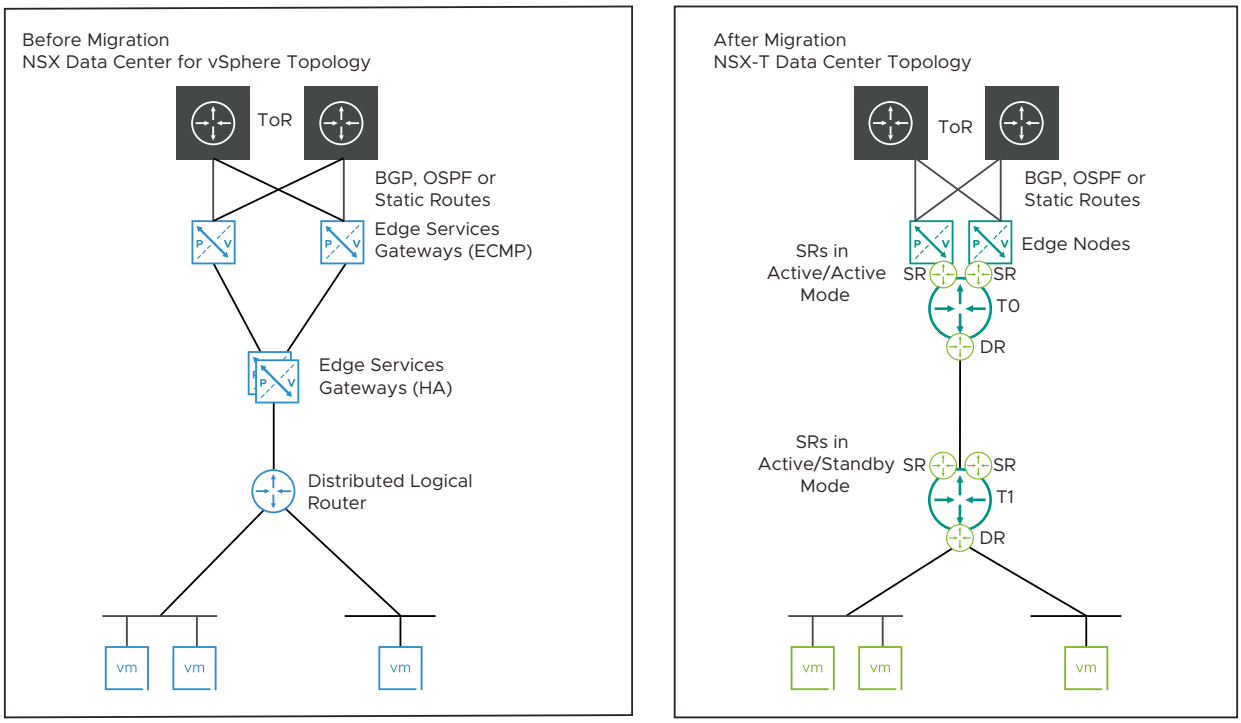

One-Armed Load Balancer (Topology 4)

This topology contains the following configurations:

- The DLR has ECMP enabled and peers with multiple ESGs.
- BGP, OSPF or static routing is configured between the ESG and top-of-rack (ToR) northbound routers. If BGP is configured, all ESGs must be configured with the same global BGP settings.
- If BGP is configured between the DLR and ESG, all BGP neighbors on the DLR must have the same weight.

- The ToR-facing ESGs must not be running L4-L7 services.
- An ESG is a single-arm load balancer attached to a Logical Switch, which is connected to a DLR. This ESG can also run Edge firewall and DHCP.

After migration, the top-level (ToR-facing) Edge Services Gateways and the DLR are replaced with a tier-0 gateway. The ESG performing load balancing service is replaced with a tier-1 gateway.

- The tier-0 gateway service router is in active/active mode.
- The IPs of the DLR interfaces are configured as downlinks on the tier-0 Gateway.
- The BGP or OSPF configuration of the top-level ESGs is translated to a BGP or OSPF configuration, respectively, on the tier-0 gateway. Route redistribution configuration is translated.
- Static routes from the top-level ESGs and DLRs are translated to static routes on the tier-0 gateway.
- The load balancing configuration on the ESG is translated to a one-arm load balancer using Service Interface (SI) configuration on the tier-1 Service Router.

Figure 5-4. Topology 4: Before and After Migration

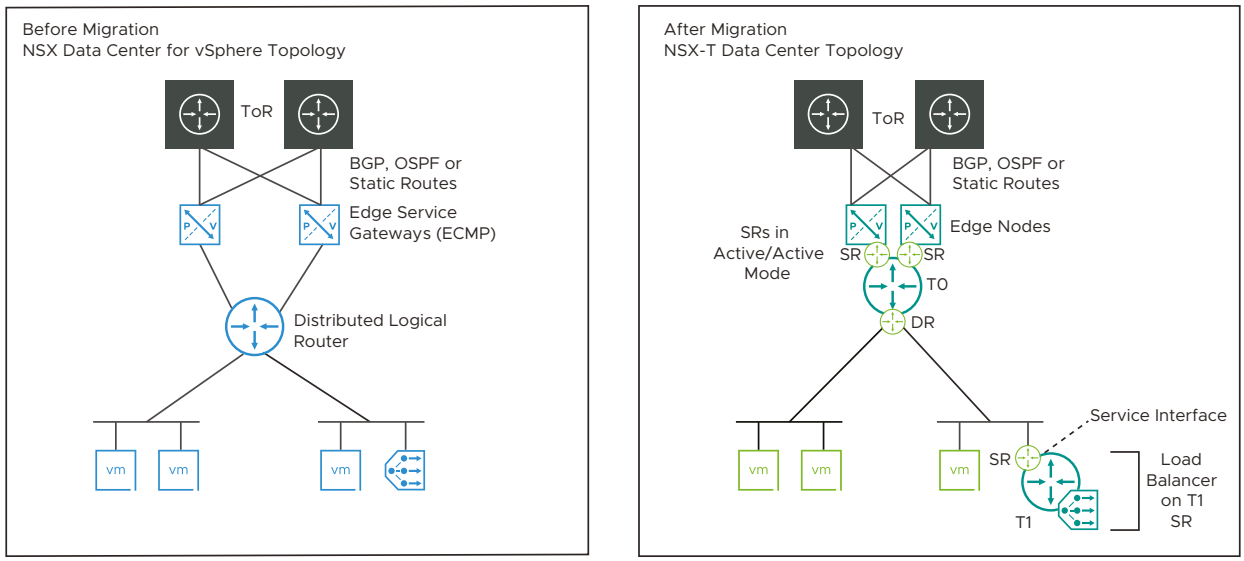

VLAN-Backed Micro-Segmentation (Topology 5)

This topology uses Distributed Firewall to provide firewall protection to workloads connected to VLAN-backed distributed port groups.

This topology uses the following NSX-V features:

- NSX Manager
- Host Preparation (Distributed Firewall only)
- Distributed Firewall

- Service Composer
- Grouping Objects

This topology must not contain the following features:

- Transport Zone
- VXLAN
- Logical Switch
- Edge Services Gateway
- Distributed Logical Router

Changes Made During Host Migration in an End-to-End Migration

During the host migration step in an end-to-end migration, changes are made to migrate NSX-V hosts to NSX-T hosts.

- NSX-V software is uninstalled.
- NSX-T software is installed.
- For vSphere Distributed Switch versions 6.5.0 and 6.6.0:

  Hosts are configured with N-VDS to replace vSphere Distributed Switches:

  - Each N-VDS is created with a name that references the distributed switch name. For example, distributed switch `ComputeSwitchA` is created as N-VDS `nvds.ComputeSwitchA`.
  - If different clusters use different distributed switches to back logical switches, an N-VDS is created with a name that combines all the distributed switch names. For example, if `ComputeCluster1` and `ComputeCluster2` use distributed switch `ComputeSwitchA` to back logical switches and `ComputeCluster3` uses `ComputeSwitchB` to back logical switches, the N-VDS is created as `nvds.ComputeSwitchA.ComputeSwitchB`.
  - PNICs and vmks in the vSphere Distributed Switch are migrated to N-VDS.
  - NSX-V VTEPs are migrated to NSX-T TEPs.
- For vSphere Distributed Switches version 7.0:

  Hosts configured for vSphere Distributed Switch version 7.0 continue using the same switch after migration.

  - PNICs and vmks in the vSphere Distributed Switch remain connected on the same vSphere Distributed Portgroups.
  - NSX-V VTEPs are migrated to NSX-T TEPs and connected to standalone ports on the same vSphere Distributed Switch.

Note: If NSX-V has multiple VTEPs and a single LAG configuration, after migration the LAG will have a single TEP and in failover mode.

Virtual Machine Deployment During an End-to-End Migration

After you start an end-to-end migration, do not change the NSX-V environment. If you want to deploy VMs during the migration, wait until some of the NSX-V hosts have migrated to NSX-T and deploy the VMs on NSX-T hosts. Connect the VMs to NSX-T segments and install VMware Tools on the VMs.

Deploying on NSX-T with VMware Tools installed ensures that the VMs are populated into security groups and receive the intended Distributed Firewall policies.

Caution VMs deployed without VMware Tools installed, or deployed on NSX-V do not receive the intended Distributed Firewall policies.

If you use vSphere templates to deploy VMs, update the templates to use NSX-T segments for the VM network configuration. Specifying NSX-T segments ensures that any VMs deployed using the templates are deployed on NSX-T hosts.

If you use automation tools to deploy VMs on vSphere, but do not use vSphere templates, you might need to change your automation tool configuration to ensure that the VMs are deployed on NSX-T.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1
	Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Preparing the NSX-V Environment for a Fixed Topology Migration

Check the Configurations of the NSX-V Environment

Check your NSX-V environment before starting an end-to-end migration.

System State

Check the following system states:

- If your environment is vSphere 7.0 or later, upgrade the VDS to 7.0 or later.
- Verify that the NSX-V components are in a green state on the NSX Dashboard.
- Verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering maintenance mode.
- Verify that no NSX-V upgrades are in progress.

- Verify the publish status of Distributed Firewall and Service Composer to make sure that there are no unpublished changes.
- You can have vSphere High Availability (HA) enabled if the NSX-V environment has VDS 7.0 or later.

  Note: HA is not supported for previous versions of VDS This is because if the NSX-V environment has VDS 6.5 or 6.7, and the vmkernel ports (vmk's) are attached to VDSes, during an in-place migration, the hosts and VMs may lose network connectivity for a period of time long enough to trigger HA. The HA mechanism will try to power off, migrate and restart VMs. This might fail because the NSX-V environment is being migrated to NSX-T. As a result, after the migration, VMs might remain in a powered-off state or have no network connectivity if powered on. To avoid this situation, disable HA or attach the management vmk to a VSS before starting the migration.

General Configuration

- Back up the NSX-V and vSphere environments. See "NSX Backup and Restore" in the *NSX Administration Guide*.
- The VXLAN port must be set to 4789. If your NSX-V environment uses a different port, you must change it before you can migrate. See "Change VXLAN Port" in the NSX-V *NSX Administration Guide*.

Controller Configuration

- NSX-V transport zones using multicast or hybrid replication mode are not supported for migration. An NSX Controller cluster is required if VXLAN is in use. VLAN-backed micro-segmentation topologies do not use VXLAN and so do not require an NSX Controller cluster.

Host Configuration

- On all host clusters in the NSX-V environment, check these settings and update if needed:
  - Set vSphere DRS accordingly.

    Disable vSphere DRS if one of the following apply:

    - **In-Place** migration mode will be used. In this mode hosts are not put in maintenance mode during migration and VMs will experience a network outage and network storage outage during the migration. This mode is only available if the environment is vSphere 6.x (VDS will be migrated to N-VDS).
    - **Manual Maintenance** migration mode will be used. If you decide to use vMotion for migrating VMs, you can disable vSphere DRS, or set the vSphere DRS automation level to Manual, Partially Automated, or Fully Automated.
    - **Automated Maintenance** migration mode will be used and the VDS version is 6.5 or 6.7.

    Set vSphere DRS mode to Fully Automated if:

    - **Automated Maintenance** migration mode will be used and the VDS version is 7.0.

Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

- To migrate Network Introspection service rules, use the **Maintenance** host migration mode. **In-Place** migration mode is not supported.
- If you have hosts that have NSX-V installed, but are not added to a vSphere Distributed Switch, you must add them to distributed switches if you want to migrate them to NSX-T. See Configure Hosts Not Attached to vSphere Distributed Switches for more information.
- On each cluster that has NSX-V installed, check whether Distributed Firewall is enabled. You can view the enabled status at **Installation & Upgrade > Host Preparation**.

  If Distributed Firewall is enabled on any NSX-V clusters before migration, Distributed Firewall is enabled on all clusters when they migrate to NSX-T. Determine the impact of enabling Distributed Firewall on all clusters and change the Distributed Firewall configuration if needed.

Edge Services Gateway Configuration

- You might need to make changes to your NSX-V route redistribution configuration before migration starts.
  - Redistribution filters are not migrated. For BGP, filters can be moved to the BGP neighbor level.
  - After migration, dynamically learned routes between Distributed Logical Router and Edge Services Gateway are converted to static routes and all static routes are redistributed in BGP or OSPF. If you need to filter any of these routes, you can configure them at the BGP neighbor level or manually configure the redistribution rules on NSX-T after the configuration migration is completed and before cutover. Note that if you roll back, the manual configuration of redistribution rules will also be removed.
  - The default MTU setting is 1500 on NSX-T. If you have non-default MTU setting requirements, you can change the setting. See Change the Global MTU Setting.
- NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the **Migrate Configuration** step fails.
- If you have an Edge Services gateway performing one-armed load balancer function, you must change the following configurations if present before you import the configuration:
  - If the Edge Services Gateway has an interface configured for management, you must delete it before migration. You can have only one connected interface on an Edge Services Gateway providing one-arm load balancer function. If it has more than one interface, the **Migrate Configuration** step fails.

- If the Edge Services Gateway firewall is disabled, and the default rule is set to deny, you must enable the firewall and change the default rule to accept. After migration the firewall is enabled on the tier-1 gateway, and the default rule accept takes effect. Changing the default rule to accept before migration prevents incoming traffic to the load balancer from being blocked.

- Verify that Edge Services Gateways are all connected correctly to the topology being migrated. If Edge Services Gateways are part of the NSX-V environment, but are not correctly attached to the rest of the environment, they are not migrated.

  For example, if an Edge Services Gateway is configured as a one-armed load balancer, but has one of the following configurations, it is not migrated:

  - The Edge Services Gateway does not have an uplink interface connected to a logical switch.
  - The Edge Services Gateway has an uplink interface connected to a logical switch, but the uplink IP address not does match the subnet associated with the distributed logical router that connects to the logical switch.

Security Configuration

- If you plan to use vMotion to move VMs during the migration, disable all SpoofGuard policies in NSX-V to prevent packet loss.
  - Automated Maintenance mode uses DRS and vMotion to move VMs during migration.
  - In Manual Maintenance mode, you can optionally use vMotion to move VMs during migration.
  - In-Place migration mode does not use vMotion.

Security Group Configuration

If existing Security Policies contain Guest Introspection service rules that are applied to Security Groups with static VM members or dynamic members other than VMs, do these steps:

1. Create new Security Groups with VMs only in the dynamic membership criteria. Make sure that the dynamic membership criteria produces the same effective VM members as your original Security Groups.
2. Before starting the migration, update the existing Security Policies to apply the new Security Groups to the Guest Introspection service rules.

   If you prefer not to update your existing Security Policies before the migration, you can still keep the new Security Groups ready with the correct dynamic membership criteria in your NSX-V environment. In the **Resolve Configuration** step of the migration process, you will be prompted to provide alternative Security Groups.

Service Composer Synchronization

Ensure that the Service Composer is in sync with Distributed Firewall before you start the migration. A manual synchronization ensures that if you make any last-minute changes in the policy configuration before starting the migration, these changes are applied to the Security Policies that are created using Security Composer too. For example, you edit the name of the Security Group that is used in a firewall rule before starting the migration.

To verify whether the Service Composer status is in sync, do these steps:

1 In the vSphere Client, navigate to **Networking and Security > Security > Service Composer**.

2 Click the **Security Policies** tab.

3 Verify that the Sync Status is **In Sync**. If it is not in sync, click **Synchronize**.

As a best practice, always click the **Synchronize** button before starting the migration even when the sync status is green. Do this manual synchronization regardless of whether you performed any last-minute changes in the policy configuration.

During migration, if the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of the Security Policies created using the Service Composer by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get the Service Composer in sync with Distributed Firewall, and restart the migration.

Configure Hosts Not Attached to vSphere Distributed Switches

An NSX-V environment can contain hosts that have NSX-V installed, but are not added to a vSphere Distributed Switch. You must add the hosts to a vSphere Distributed Switch before you can migrate them.

You can use a distributed switch you already have in your environment, or create a new distributed switch for this purpose. Right click the distributed switch and select **Add and Manage Hosts** to add the hosts to the distributed switch. You do not need to assign physical uplinks or VMkernel network adapters to the distributed switch.

See "Add Hosts to a vSphere Distributed Switch" in the *vSphere Networking Guide* for more information.

If you import the configuration before you make this change, you must restart the migration to import the updated configuration. See Make Changes to the NSX-V Environment.

After the migration has finished, the hosts are no longer required to be attached to the distributed switch.

- If you added the hosts to an existing distributed switch, you can remove them from the distributed switch.
- If you added the hosts to a new distributed switch that you are not using for another purpose, you can delete the distributed switch.

Tag Management VMs in a Collapsed Cluster Environment

You can migrate an environment that uses a collapsed cluster.

In a collapsed cluster design, all management VMs, workload VMs, and optionally edges run on the same vSphere cluster that is prepared for NSX-V. The management VMs of the NSX-T must be initially attached to dvPortgroups. After migration, the management VMs of NSX-T will be attached to NSX-T VLAN segments.

The management VMs in the NSX-T include appliances such as NSX Manager, vCenter Server, VMware Identity Manager, and so on. The NSX-T VLAN segment ports to which these management VMs connect are blocked in two cases: when these management VMs are rebooted after they were migrated in-place by the Migration Coordinator, or when they are moved from NSX-V hosts to NSX-T hosts by vMotion in maintenance migration mode. Therefore, the management VMs might lose connectivity in such cases.

To prevent this problem, create a "management_vms" tag category, and add tags in this category. Assign a tag from this category to all the management VMs in the NSX-T environment. These VMs will be attached to unblocked VLAN segment ports.

Procedure

1. Log in to the vSphere Client.
2. Click **Menu > Tags & Custom Attributes**.
3. Click **Categories**, and then click **New** to add a category.

   Create a category with name **`management_vms`**.
4. Click the **Tags** tab and add a tag in the management_vms category.
5. Navigate to **Menu > Hosts and Clusters**.
6. Expand the collapsed cluster from the left Navigator view, right-click the name of the NSX Manager VM, and select **Tags & Custom Attributes > Assign Tag**.
7. Assign a tag from the management_vms category to the NSX Manager VM.
8. Repeat steps 6 and 7 for all the management VMs in the cluster.

   For a detailed information about tag categories and tags, see the *vCenter Server and Host Management* documentation.
9. Log in to NSX Manager.
10. Navigate to **Inventory > Groups** and create a group.
11. Click **Set** to set members.
12. Set **Group Type** to **IP Addresses Only** and include the IP addresses of the management VMs that will be migrating from the dvPortgroups to the NSX-T segments.
13. Navigate to **Security > Distributed Firewall > Actions > Exclusion List** and add this group.

Delete Partner Service Deployments

If your NSX-V environment uses a partner service for Guest Introspection, or both Guest Introspection and Network Introspection, delete the partner service deployment before migration.

You must also delete the Guest Introspection instance (GI-SVM) so that the Guest Introspection module is uninstalled from the clusters.

If your NSX-V environment uses a partner service only for Network Introspection, you have the flexibility to decide whether to delete the partner service deployment before or after the migration. When a partner service deployment is deleted, the partner service virtual machines (SVMs) are removed from the NSX-V-prepared host cluster, and security protection is lost.

Important Consult the VMware partner before migrating their service that is running on the NSX-V workloads. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their services to NSX-T.

Procedure

1. Log in to the vSphere Client and navigate to **Networking and Security > Installation and Upgrade > Service Deployment**.
2. Select the deployed service and click **Delete**.

Preparing the NSX-T Environment for a Fixed Topology Migration

Deploy an NSX Manager Appliance

You must deploy a new NSX Manager appliance to run the migration coordinator. Do not deploy an NSX Global Manager.

In other words, you cannot merge your NSX-V environment into an existing NSX-T environment, which has NSX-T already installed on the vSphere host clusters.

For details on deploying a licensed version of the NSX Manager appliance, see *Install NSX Manager and Available Appliances* in the *NSX-T Data Center Installation Guide*.

Install one appliance to perform the migration. Deploy additional appliances to form a cluster after the migration is finished. See Finish Deploying the NSX Manager Cluster.

If you install the NSX Manager appliance on an ESXi host that is a part of the NSX-V environment that is migrating, do not attach the appliance interfaces to an NSX-V logical switch. To prevent the management VMs in NSX-T from losing connectivity after the VMs are rebooted post migration, tag the management VMs. For more information, see Tag Management VMs in a Collapsed Cluster Environment.

Add a Compute Manager

Before you can start the migration process, you must add the vCenter Server that is associated with NSX-V as a compute manager in NSX-T.

Prerequisites

Log into the NSX-V NSX Manager web interface to retrieve the settings used for vCenter Server registration. You must use the same settings. For example, if an IP address is specified, use the IP address and not the FQDN. Note that the FQDN of the vCenter Server is case-sensitive. If you enter the FQDN in the procedure below, be sure that it matches the vCenter Server's FQDN exactly.

Procedure

1. From a browser, log in with admin privileges to NSX Manager at `https://<nsx-manager-ip-address>`.
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 vCenter Server. You can optionally describe any special details such as, the number of clusters in the vCenter Server.
FQDN or IP Address	Type the FQDN or IP address of the vCenter Server.
Type	The default compute manager type is set to vCenter Server.
HTTPS Port of Reverse Proxy	The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances. Set the reverse proxy port to register the compute manager in NSX-T Data Center.
Username and Password	Type the vCenter Server login credentials.
SHA-256 Thumbprint	Type the vCenter Server SHA-256 thumbprint algorithm value.
Create Service Account	Enable this field for features such as vSphere Lifecycle Manager that need to authenticate with NSX-T Data Center APIs. Log in with the administrator@vsphere.local credential to register a compute manager. After registration, the compute manager creates a service account. **Note** Service account creation is not supported on a global NSX Manager. If service account creation fails, the compute manager's registration status is set to `Registered with errors`. The compute manager is successfully registered. However, vSphere Lifecycle Manager cannot be enabled on NSX-T Data Center clusters. If a vCenter Server admin deletes the service account after it was successfully created, vSphere Lifecycle Manager tries to authenticate the NSX-T Data Center APIs and the compute manager's registration status is set to `Registered with errors`.

Option	Description
Enable Trust	Enable this field to establish trust between NSX-T Data Center and compute manager, so that services running in vCenter Server can establish trusted communication with NSX-T Data Center. For example, for vSphere Lifecycle Manager to be enabled on NSX-T Data Center clusters, you must enable this field. Supported only on vCenter Server 7.0 and later versions.
Access Level	Enable one of the options based on your requirement: ■ Full Access to NSX: Is selected by default. This access level gives the compute manager complete access to NSX-T Data Center. Full access ensures vSphere for Kubernetes and vSphere Lifecycle Manager can communicate with NSX-T Data Center. The vCenter Server 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-T Data Center. The vCenter Server 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-T Data Center to discover and register the vCenter Server 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 vCenter Server credentials and click **Resolve**.

If an existing registration exists, it will be replaced.

Results

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

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

After the vCenter Server is successfully registered, do not power off and delete the NSX Manager VM without deleting the compute manager first. Otherwise, when you deploy a new NSX Manager, you will not be able to register the same vCenter Server again. You will get the error that the vCenter Server is 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, NSX Intelligence VM, 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, any NSX Intelligence VM, 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 sucessfully remove the VC, ensure that you re-deploy an NSX Manager node from the VC.

This restriction applies to a fresh installation of NSX-T Data Center as well as an upgrade.

Change the Global MTU Setting

When an Edge Services Gateway (ESG) is migrated, the MTU setting of the interfaces is not migrated. A default value of 1500 is used. You can change the default value using the API.

You can also modify the MTU setting for the interfaces after the migration.

Procedure

1 Make the following API call to retrieve the current configuration.

```
GET /api/v1/global-configs/RoutingGlobalConfig
```

2 Change the value for `logical_uplink_mtu` and make the following call.

```
PUT /api/v1/global-configs/RoutingGlobalConfig
```

Create an IP Pool for Edge Tunnel End Points

If your NSX-V environment uses Edge Services Gateways, you must create an IP pool in the NSX-T environment for the Edge Tunnel End Points (TEP) before you start the migration.

Prerequisites

- Identify existing IP pools or DHCP ranges for NSX-V VTEPs.
- Determine which IP addresses to use to create an IP pool for Edge TEPs.

  The IP range and VLAN must not already be in use in the NSX-V environment.
- Verify that the NSX-T TEP IP addresses have network connectivity to the NSX-V VTEP IP addresses.

Procedure

1. From a browser, log in with admin privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Select **Networking > IP Management > IP Address Pools**.
3. Click **Add IP Address Pool**.
4. Enter a name for the new IP pool.
5. (Optional) Enter a description.
6. In the **Subnets** column, click **Set** to add subnets.
7. Specify the IP ranges.
   a. Select **Add Subnets > IP Ranges**.
   b. Enter IPv4 or IPv6 ranges.
   c. Enter the subnet address in a CIDR format.
   d. Enter the Gateway IP address for this subnet.
   e. (Optional) Enter DNS servers.
   f. (Optional) Enter DNS suffix.
   g. Click **Add**, and then click **Apply**.
8. Click **Save**.

Deploy NSX Edge Nodes

You must deploy NSX Edge nodes as a virtual machine on ESXi using an OVA or OVF file.

Do not deploy on bare metal. Do not deploy from the NSX Manager user interface.

Snapshots of NSX appliances (including Edge node VMs) are not supported and must be disabled. For information on how to disable snapshots, see the topic Disable Snapshots on an NSX Appliance in the *NSX-T Data Center Installation Guide*.

NSX Edge nodes must be connected to trunk portgroups. To learn more about NSX Edge networking, see "NSX Edge Networking Setup" in the *NSX-T Data Center Installation Guide*.

Caution If you deploy the NSX Edge node VM on an NSX-V-prepared host, connectivity of the Edge node might be affected by a Distributed Firewall deny rule in NSX-V. To avoid this issue, add the Edge node VM to the Distributed Firewall's exclusion list.

Prerequisites

- You must have sufficient ESXi hosts with appropriate resources available to accommodate the NSX Edge appliances.

Procedure

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

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

2 In the vSphere Client, select the host on which to install NSX Edge 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 vCenter Server and vSphere inventory.

6 Select a compute resource for the NSX Edge node appliance, and click **Next**.

7 Review and verify the OVF template details, and click **Next**.

8 Select a deployment configuration and click **Next**.

   See the **Import Configuration** step for details on the size of Edge nodes you must deploy.

9 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.

10 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.

   Post-migration, the NSX Edge node is connected to one of these three trunk networks using only a single fastpath interface. The network settings can be adjusted or verified after the NSX Edge node is deployed.

11 Configure IP Allocation settings.

   a For IP allocation, specify **`Static - Manual`**.

   b For IP protocol, select **`IPv4`**.

12 Click **Next**.

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

13 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.

14 Enter the hostname of the NSX Edge.

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

   Skip any VMC network settings.

16 Enter the DNS Server list, the Domain Search list, and the NTP Server IP or FQDN list.

17 (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.

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

   The installation might take 7-8 minutes.

19 Start the NSX Edge node VM manually.

20 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.

21 After the NSX Edge node starts, log in to the CLI 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.

22 Run the `get interface eth0` (without VLAN) or `get interface eth0.<vlan_ID>` (with a VLAN) command to verify that the IP address was applied as expected.

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

Interface: eth0.100
  Address: 192.168.110.37/24
  MAC address: 00:50:56:86:62:4d
```

```
MTU: 1500
Default gateway: 192.168.110.1
Broadcast address: 192.168.110.255
...
```

23 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 IP or FQDN list.

24 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 use DHCP 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 Type the **`set interface`** ***`interface`*** **`dhcp plane mgmt`** command.

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

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

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

Join NSX Edge Node VM with the Management Plane

You must join the NSX Edge node VM you created to the management plane.

Do not join the NSX Edge node VM to the management plane using any other method. Do not create transport nodes from the NSX Edge node VM.

Procedure

1 Open an SSH session or console session to the NSX Manager appliance.

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 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.

5 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)
```

6 In the NSX Manager UI, you can navigate to **System > Fabric > Nodes > Edge Transport Nodes** and see the NSX Edge node. The **Configuration State** column will display **Pending**. If you click the name of the Edge node, you will be prompted to configure the node. Do not configure the node. The configuration will occur during the migration.

Register Third-Party Guest Introspection Service with NSX-T

If Security Policies in your NSX-V environment use third-party Guest Introspection service provided by a partner, register the partner service with NSX-T before you start the migration.

You might need to upgrade the Partner Console to register the service with the version of NSX-T that is used for this migration. For more information, see the partner documentation.

Complete the following procedure to register a partner service for endpoint protection with the NSX Manager in your NSX-T environment.

Procedure

1 Log in to the Partner Console with **Admin** privileges.

2 Update the NSX endpoint in the Partner Console. Specify the following details:

   - IP address of NSX-T NSX Manager
   - Port (default is 443)
   - User name and password of the NSX-T NSX Manager

   Make sure to test the connection before proceeding to the next step. If you need help with using the Partner Console, see the partner documentation.

   The partner service and the vendor templates that are associated with this partner service are now created in NSX-T.

3 Verify that the partner service is registered with NSX-T.

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

   b Navigate to **System > Service Deployments > Deployment**.

   c Click the **Partner Service** drop-down menu, and check that the partner service is listed.

Register Third-Party Network Introspection Services with NSX-T

If Security Policies in your NSX-V environment use third-party Network Introspection services provided by partners, partner services must be registered with NSX-T before you start the migration.

You might need to upgrade the Partner Console to ensure that the partner service is registered with the version of NSX-T that is used for this migration. For more information, see the partner documentation.

The following types of east-west network introspection services are supported for migration:

- Intrusion detection services (IDS)
- Intrusion protection services (IPS)
- Network monitoring services
- Next-generation firewall services

A partner registers the service, vendor template, and the Partner Management Console/Partner Service Manager. Then, either you or the partner can create the service profile. It can vary from one partner to another. See the partner documentation.

In the following procedure, step 2 is required when your NSX-V environment uses only Network Introspection service.

If your environment uses a combination of both Guest Introspection and Network Introspection services from a single partner (partner A), partner does step 1. Step 2 is not required.

If your environment uses Guest Introspection service from one partner (partner A) and Network Introspection service from another partner (partner B), then:

- Use the Partner Console of partner A to register the Guest Introspection service. See the partner documentation for help on registering the service.
- Partner B registers the Network Introspection service (step 1 of the procedure). Either you or the partner can create the service profile, as explained in step 2.

Procedure

1 Partner registers the partner service, vendor template, and the partner Service Manager using NSX-T APIs.

2 Create a service profile to specify attributes of a vendor template for a given partner service.

   For a network introspection service, multiple service profiles can be associated with a single vendor template.

   You can create a service profile either by using the NSX-T API or the NSX Manager UI. For detailed steps on creating the service profile by using the NSX Manager UI, see the *NSX-T Data Center Administration Guide*.

   When you use the NSX Manager UI to create a service profile, the service reference is internally created, if it is not already present.

   If you decide to use the NSX-T APIs to create a service profile, do the following steps:

   a Create a service reference.

   ```
   PATCH https://{policy-mgr-ip}/policy/api/v1/infra/service-references/
   {service-reference-id}
   ```

   b Use the *service-reference-id* from the previous step to create the service profile.

   ```
   PATCH https://{policy-mgr-ip}/policy/api/v1/infra/service-references/
   {service-reference-id}/service-profiles/{service-profile-id}
   ```

   For a detailed information about these APIs, see the *NSX-T Data Center API Guide*.

Import the NSX-V Configuration

The first step of the migration process is to import the NSX-V configuration.

Note During this step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.

Procedure

1 From a browser, log in to NSX Manager as **admin**.

2 Navigate to **System > Migrate**.

3 In the **NSX for vSphere** pane, click **Get Started** and select **Fixed Topology**.

4 From the **Import Configuration** page, click **Select NSX** and provide the credentials for vCenter Server and NSX-V.

Note The drop-down menu for vCenter Server displays all vCenter Server systems that are registered as compute managers. Click **Add New** if you need to add a compute manager.

5 Click **Start** to import the configuration.

6 When the import has finished, click **Continue** to proceed to the **Resolve Configuration** page.

If the import fails due to incorrect edge node configuration translation, click the **Failed** flag to view information about the number and size of the required NSX Edge resources. After you deploy the correct number and size of edge nodes, click **Rollback** to roll back this migration attempt and restart the configuration import.

Results

When the NSX-V topology is imported successfully, the **View Imported Topology** link is enabled. Click this link to view a graph of the imported topology. However, the topology viewer does not work for a scaled NSX-V environment.

Roll Back a Migration

After you have started the migration process, you can roll back the migration to undo some or all your progress.

You can roll back or undo the migration from some of the migration steps. After the migration has started, you can click **Rollback** on the furthest step completed. The button is disabled on all other pages.

Table 5-12. Details about Rolling Back a Migration

Migration Step	Rollback Details
Import Configuration	Click **Rollback** on this page to roll back the **Import Configuration** step.
Resolve Configuration	Rollback is not available here. Click **Rollback** from the **Import Configuration** page.

Table 5-12. Details about Rolling Back a Migration (continued)

Migration Step	Rollback Details
Migrate Configuration	Click **Rollback** on this page to roll back the migration of the configuration to NSX-T and the input provided on the **Resolve Configuration** page. Before rolling back the migration from the **Migrate Configuration** page, it is recommended to collect the Support Bundle. For more information, see *Collect Support Bundles* in the *NSX-T Data Center Administration Guide*. Verify that the rollback was successful before you start a new migration. Log into the NSX Manager web interface and switch to **Manager** mode. Verify that all configurations have been removed. For more information about Manager mode, see *Overview of the NSX Manager* in the *NSX-T Data Center Administration Guide*. **Note** If you experience problems rolling back the **Migrate Configuration** step, you can start a new migration instead. 1 Remove the vCenter Server that you added as a compute manager. 2 Delete the current NSX Manager appliance. 3 Deploy a new NSX-T environment with NSX Manager and NSX Edge appliances. 4 Start a new migration. If you forget to remove the compute manager, see the topic "Remove NSX-T Data Center Extension from vCenter Server" in the *NSX-T Data Center Administration Guide* on how to remove the NSX-T Data Center extension from the vCenter Server.
Migrate Edges	Click **Rollback** on this page to roll back the migration of Edge routing and services to NSX-T. **Caution** If you roll back the **Migrate Edges** step, verify that the traffic is going back through the NSX-V Edge Services Gateways. You might need to take manual action to assist the rollback.
Migrate Hosts	Rollback is not available in this step. However, you can still do a manual rollback to remove NSX-T from the migrated hosts and reinstall NSX-V on the hosts. If you are migrating from NSX-V 6.4.8 or later, run the following REST API on the NSX-V NSX Manager before doing the manual rollback, and then reinstall NSX-V on the hosts. `POST api/2.0/nwfabric/blockEamEvents?action=unblock` This API enables the vSphere ESX Agent Manager (EAM) on the hosts so that the NSX-V VIBs can be installed correctly. If you are migrating from NSX-V 6.4.4, 6.4.5, or 6.4.6, this API is not needed. **Note** It is better to do a manual rollback of a failed host after the cluster, which contains the failed host, has stopped. If you are doing an In-Place host migration, and chosen parallel migration order within the cluster, then wait until the migration of all hosts in the cluster stops, regardless of failure or success. If you have chosen serial migration order of hosts within the cluster, migration stops when a host migration fails.

When you roll back a migrated configuration that contains Network Introspection redirection rules, you might see the following error message:

```
Service Insertion failed with '400: The object path=[/infra/segments/service-segments/
vdnscope-1]
cannot be deleted as either it has children or it is being referenced by other objects path=
[/infra/service-chains/Service-Chain-serviceprofile-1].
```

This error occurs because a service segment in NSX-T depends on a service chain. A service chain is not deleted until all the redirection rules referenced by it are deleted. Wait for approximately five minutes and try rolling back the migrated configuration again.

Resolve Configuration Issues

After you have imported the configuration from your environment, you must review and resolve the reported configuration issues before you can continue with the migration.

Review Migration Information

The **Resolve Configuration** page contains information about the features and configurations that are not supported for migration, and the issues that must be fixed in the NSX-V environment before you can migrate.

After reviewing the blocking issues and warnings, you might need to change configurations in your NSX-V environment before you can migrate to NSX-T. If you change the NSX-V environment, you must restart the migration to pick up the new configuration. Review all migration feedback before you provide input to avoid duplication of work.

Note For some NSX-V features, there might be automatic configurations such as certificates present. If these configurations are for features that are not supported for the specific topology, these automatic configurations are flagged as issues that need to be skipped from migration. For example, in topologies that don't support L4-L7 services on Edge Services Gateways, the certificates present for VPN and DNS will raise issues to skip these configurations from migration.

Procedure

1 From the **Resolve Configuration** page, review the reported issues in the **Blocking** category to identify blocking issues that require changes to your NSX-V environment.

Figure 5-5. Blocking Issues on the Resolve Configuration Page

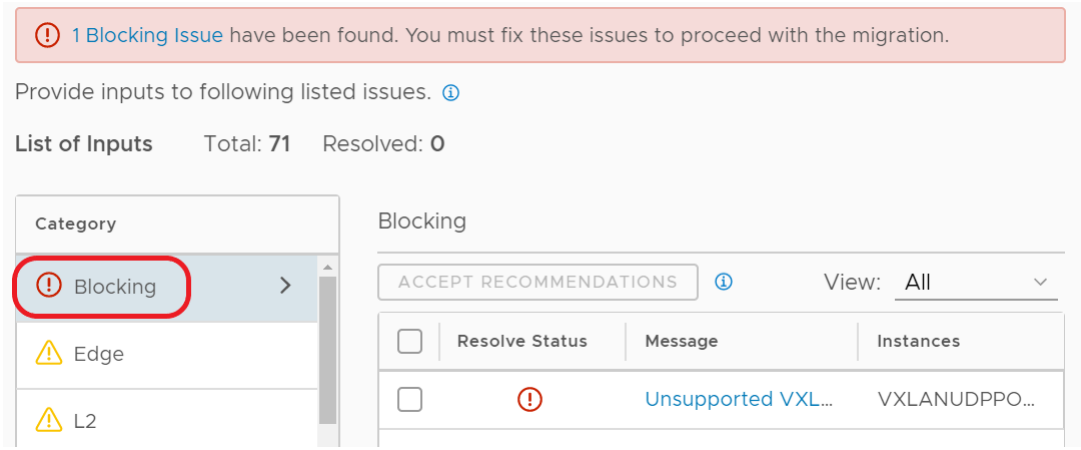

Some examples of blocking issues are:

- Incorrect DRS configuration of Maintenance mode migration.
- vMotion vmknics not configured on host for Maintenance mode migration.
- Unsupported VXLAN UDP port.

2 Review the warnings and issues reported in each category.

Figure 5-6. Warnings and Categories of Issues on the Resolve Configuration Page

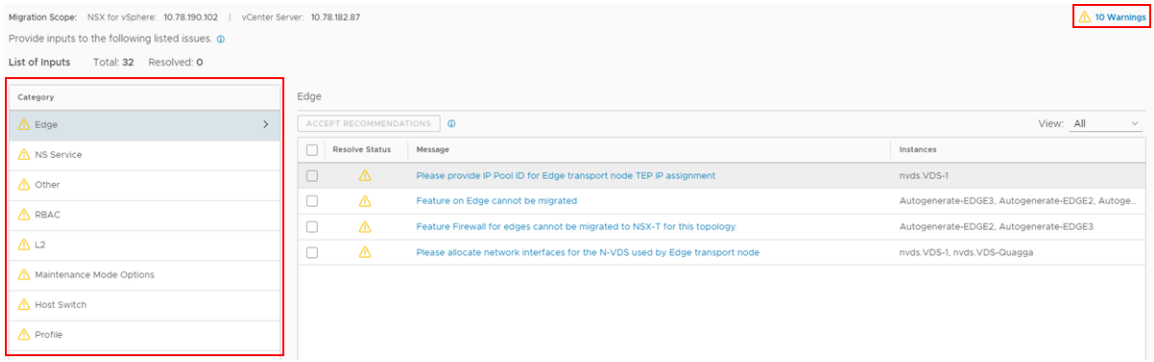

a Click **Warnings** and review the information there.

b Review the reported issues in all categories.

What to do next

If you find blocking issues, fix them in the NSX-V environment before you can proceed with the migration. See Make Changes to the NSX-V Environment.

If you did not find any blocking issues or other configurations that require a change in the NSX-V environment, you can proceed with the migration. See Provide Input for Configuration Issues.

Make Changes to the NSX-V Environment

If you find blocking issues or other configuration issues that must be fixed in your NSX-V environment, fix those issues before you can proceed with the migration. After you make the configuration changes, you must import the configuration again.

Prerequisites

Verify that Host or Edge migration has not started. See Roll Back the vSphere Networking Migration for more information about restarting the migration.

Procedure

1. Make the required changes in the NSX-V environment.
2. Navigate to the **Import Configuration** page and click **Rollback**.
3. Click **Start** to import the updated NSX-V configuration.

Results

The migration starts with the new NSX-V configuration.

What to do next

Continue the migration process. See Resolve Configuration Issues.

Provide Input for Configuration Issues

After you have reviewed the migration information and are ready to proceed with the migration, you can provide inputs for the reported configuration issues. The input you provide determines how the NSX-T environment is configured.

Multiple people can provide the input over multiple sessions. You can return to a submitted input and modify it. Depending on your configuration, you might run through the **Resolve Issues** process multiple times, update your NSX-V environment as needed, and restart the migration.

Important If you have changed the NSX-V environment for any reason since you last imported the configuration, you must restart the migration. For example, if you have connected a new VM to a logical switch, made a firewall rule change, or installed NSX-V on new hosts. See Make Changes to the NSX-V Environment for information on restarting the migration.

Note For some NSX-V features, there might be automatic configurations such as certificates present. If these configurations are for features that are not supported for the specific topology, these automatic configurations are flagged as issues that need to be skipped from migration. For example, in topologies that don't support L4-L7 services on Edge Services Gateways, the certificates present for VPN and DNS will raise issues to skip these configurations from migration.

Prerequisites

- Verify that you have reviewed all issues and warning messages and are ready to continue with the migration.
- Verify that you have addressed all blocking issues and other issues requiring a change to the NSX-V.

Procedure

1. Navigate to **System > Migrate**.
2. Go to the **Migrate NSX for vSphere** pane, and click **Resolve Configuration**.
3. Click each issue and provide input.

   Each issue can cover multiple configuration items. For each item there might be one or more possible resolutions to the issue, for example, skip, configure, or select a specific value.

   For issues that apply to multiple configuration items, you can provide input for each item individually, or select all and provide one answer for all items.
4. After the input is provided, a **Submit** button is displayed on the **Resolve Configuration** page. Click **Submit** to save your progress.
5. When you have provided input for all configuration issues, click **Submit**.

   The input is validated. You are prompted to update any invalid input. Additional input might be required for some configuration items.
6. After you have submitted all requested input, click **Continue** to proceed to the **Migrate Configuration** page.

Example: Configuration Issues

For some examples of configuration issues and their required inputs, see Example Configuration Issues.

Example Configuration Issues

You must provide inputs on various configuration issues, including Maintenance Mode migration options and configuration details for the new NSX-T Edge nodes.

Migrating Hosts in vCenter Server 7.0 Using Automated Maintenance Migration Mode

Consider the following scenario:

- The NSX-V environment uses vSphere Distributed Switch 7.0.
- On the **Resolve Configuration** page, Host Maintenance mode is set to **Automated**.
- vSphere DRS is not enabled on the clusters that are being migrated.

In this scenario, the following blocking issue messages are displayed on the **Resolve Configuration** page:

```
Incorrect DRS Configuration for Maintenance Mode migration.
```

```
Vmotion vmknics not configured on host for Maintenance mode migration.
```

To resolve the DRS configuration issue, go to the vSphere Client, and enable DRS on each cluster that is being migrated. Ensure that the DRS Automation Level is set to **Fully Automated**.

To resolve the second blocking issue, go to the vSphere Client, and enable vMotion on the VMkernel adapter of each host in the cluster. For detailed steps about enabling vMotion on the VMkernel adapter, see the *vSphere 7.0* product documentation.

After fixing the blocking configuration issues in the NSX-V environment, roll back the current migration, and import the configuration again.

Edge Node Networking Configuration

During **Resolve Configuration**, you provide information about the NSX Edge nodes that you have created to replace your NSX-V Edge Services Gateways. The configuration might have to change to work correctly on NSX-T. You might need to use a different IP address and VLAN than you used in NSX-V.

Migrating Edge Services Gateway with L4-L7 Services

Using the same interface for the router uplink and services such as VPN is supported in NSX-V. This configuration is not supported in NSX-T. You can assign new IP addresses for the NSX Edge node uplinks so that you do not need to change the IP address for the services running on the NSX Edge node.

Migrating Edge Services Gateway in a High Availability Configuration

The NSX-V topology that contains Edge Services Gateways in a high availability configuration can contain an Edge Services Gateway with two uplinks connected to two different distributed port groups on different networks.

In NSX-T, this configuration is replaced by two NSX Edge nodes, both of which must have their uplinks on the same network.

For example, an Edge Services Gateway with HA might have this configuration:

- vnic1 has IP address 192.178.14.2/24 and is attached to port group Public-DVPG which uses VLAN 11.
- vnic4 has IP address 192.178.44.2/24 and is attached to port group Public-DVPG-2 which uses VLAN 15.

To work after migration, at least one of these IP addresses has to change, as they both must be on the same network.

Here is an example of the information that might be provided during Resolve Configuration.

For the first NSX Edge node:

- ID is fa3346d8-2502-11e9-8013-000c2936d594.
- IP address is 192.178.14.2/24.
- VLAN is 11.

For the second NSX Edge node:

- ID is fa2de198-2502-11e9-9d7a-000c295cffc6.
- IP address is 192.178.14.4/24.
- You do not need to provide the VLAN because the same VLAN configured for the first NSX Edge node is assumed for the second node.

Both NSX Edge nodes must have connectivity to this network.

Add Additional Uplinks on NSX-T Edge Nodes

In the **Resolve Configuration** step of the migration, you can optionally add additional uplink interfaces on the NSX-T Edge nodes.

A maximum of 15 additional uplinks are supported on each NSX-T Edge node. Therefore, each Edge node can have a maximum of 16 uplinks.

- One required uplink that corresponds to an existing uplink on the NSX-V Edge Services Gateway (ESG).
- Up to15 additional uplinks.

On the **Resolve Configuration** page, you must first submit inputs to resolve feedback messages that are displayed for adding the required uplink on each NSX-T Edge node corresponding to an existing uplink on each NSX-V Edge Services Gateway. After the feedback messages for the required uplinks are resolved for each NSX-T Edge, follow the steps in this procedure to add additional uplinks (if necessary) on the desired NSX-T Edge nodes.

The default recommendation is to skip the creation of additional uplinks on the NSX-T Edge nodes. If you accept the default recommendation, no additional uplinks are created during the migration.

Example

Assume that your NSX-V topology has two north-facing ESGs configured in ECMP and connected to upstream physical ToR switches. The existing uplink interfaces on the two ESGs have the following configuration:

- ESG-1: The uplink of this ESG is attached to a Public-DVPG1 port group, which uses VLAN 13. The uplink IP is 40.40.40.2/24.
- ESG-2: The uplink of this ESG is attached to a Public-DVPG2 port group, which uses VLAN 16. The uplink IP is 20.20.20.2/24.

The migrated NSX-T topology contains two NSX-T Edge nodes (Edge-1, Edge-2) in an active/active configuration. You want to add one additional uplink on both Edge-1 and Edge-2.

The procedure in this topic explains the workflow to add two additional uplinks for this example.

Note If the migrated NSX-T Edge nodes in an active-passive configuration, the feedback messages are slightly different. However, the overall workflow remains the same. You must follow the messages in the UI and submit appropriate inputs to resolve the feedback messages.

Procedure

1 On the **Resolve Configuration** page, click the following feedback message:

`Enter the number of additional uplinks needed.`

Note If the Edge nodes are in an active-passive configuration, enter the number of additional uplinks you want to add on an Edge node. The same number of additional uplinks will automatically be added on the other Edge node.

For the example above, the NSX-T Edge nodes are in an active-active configuration. You must enter the total number of uplinks that are required for all the active Edge nodes. That is, enter **2** to create two additional uplinks.

2 Click the following feedback message:

`Select an Edge Transport Node for additional uplinks. See System > Fabric > Nodes > Edge Transport Nodes for available Edge Transport Nodes in NSX-T.`

Note This feedback message is displayed only when the Edge nodes are in an active-active configuration. If the Edge nodes are in an active-passive configuration, the active Edge node is automatically selected for creating the additional uplinks.

Because you entered **2** in the earlier step, two additional uplinks named `Additional Uplink-1` and `Additional Uplink-2` are displayed. Select an Edge node for each additional uplink.

For the example above, select **Edge-1** for `Additional Uplink-1`, and **Edge-2** for `Additional Uplink-2`.

3 Click the following feedback message:

`Enter the IP address for additional uplink.`

The additional uplink IP address must not belong to the subnet of an existing uplink interface on the same Edge node.

Note If the Edge nodes are in active-passive configuration, feedback messages to add IP addresses for the additional uplinks on both the Edge nodes are displayed. The IP addresses of the additional uplinks on both the Edge nodes must be in the same subnet and are different.

For the example above:

- Enter **30.30.30.2/24** for the additional uplink on Edge-1.
- Enter **10.10.10.2/24** for the additional uplink on Edge-2.

4 Click the following feedback message:

`Enter VLAN ID for additional uplink.`

The VLAN ID for additional uplinks must satisfy the following conditions:

- It must be different from the VLAN that is used for the Edge TEP.
- It must be different from the VLAN that is used for the existing uplink on the same Edge node.

Note If the Edge nodes are in an active-passive configuration, you must add the VLAN ID for the additional uplinks on only one Edge node. The same VLAN ID for the additional uplinks on the other Edge node will automatically be added.

For the example above, enter VLAN ID **18** for 30.30.30.2/24 uplink on Edge-1, and **20** for 10.10.10.2/24 uplink on Edge-2.

5 (Optional) After you have resolved all the feedback messages for additional uplinks, if you want to change the total number of additional uplinks to **3**, do the following.

Click the feedback message mentioned in step 1, and edit the total number of uplinks to **3**.

After you submit the change, additional feedback messages are displayed to select the Edge transport node, uplink IP, and VLAN ID for this additional uplink. Respond to the feedback messages as explained in steps 2, 3, and 4.

If you decide to change the total number of uplinks back to the original value **2**, you will get warning messages about uplink-3. You can accept the default recommendation to skip uplink-3 and the associated uplink IP and VLAN ID.

What to do next

After finishing the **Migrate Configuration** step, verify that the required uplinks and the additional uplinks are created on the tier-0 gateway.

For this example, navigate to the tier-0 gateway in the NSX Manager UI, and verify that the tier-0 gateway shows the following uplinks with status as `Success`:

- uplink1: 40.40.40.2/24
- uplink2: 20.20.20.2/24
- uplink3: 30.30.30.2/24
- uplink4: 10.10.10.2/24

Migrate the NSX-V Configuration

After you have resolved all configuration issues, you can migrate the configuration. When the configuration is migrated, configuration changes are made in the NSX-T environment to replicate the NSX-V configuration.

If needed, you can roll back the configuration that is migrated. Rolling back does the following:

- Remove the migrated configuration from NSX-T.
- Roll back all the resolved issues in the previous step.

See Roll Back a Migration for more information.

Prerequisites

Verify that you have completed the **Resolve Configuration** step.

Procedure

1 From the **Migrate Configuration** page, click **Start**.

  The NSX-V configuration is migrated to NSX-T.

2 Verify that all the migrated configurations are displayed in your NSX-T environment.

You can verify the migrated configurations either in the NSX-T NSX Manager interface or by running the NSX-T APIs.

Important

- During the **Migrate Configuration** step, Security Tags from NSX-V are not migrated to NSX-T. Therefore, the Security Tag-based migrated dynamic Groups and Groups with static memberships in NSX-T are empty after this step is finished. The reason is that in NSX-V, a Security Tag is an object, whereas in NSX-T, a tag is an attribute of a VM. The tags are applied to the workload VMs only after the workloads are migrated to NSX-T during the **Migrate Hosts** step.
- When the configuration is migrated to NSX-T, the configuration changes are made in the NSX-T NSX Manager database, but it might take some time for the configuration to take effect. You must verify that all expected NSX-V configurations appear on the NSX Manager interface or API in NSX-T before you proceed to the **Migrate Edges** step. For example, firewall configuration, logical switches, transport zones.

Caution This caution note applies only when you are using an NSX for vShield Endpoint license in your virtualized environment. This license is included in all vSphere editions. It enables you to use NSX-V to offload only anti-virus processing (Guest Introspection service) on VMs to service appliances that are provided by VMware partners.

In a Guest Introspection service migration, Security Groups with only dynamic VM membership criteria are supported. If tags-based dynamic Security Groups are used in your NSX-V environment, the Security Tags are not migrated to NSX-T. As no host migration is involved, the migrated dynamic Groups in NSX-T are empty. After the **Migrate Configuration** step is finished, you must manually create the equivalent tags in NSX-T, and attach them to the VMs that require an endpoint protection.

Results

After this procedure is completed, the fp-eth0, fp-eth1, and fp-eth2 interfaces on the Edge nodes are set to "administratively down". Do not reboot the Edge nodes. Otherwise the fp interfaces will come up and inadvertently advertise the NSX-T networks.

Migrate NSX-V Edges

After you have migrated the configuration, you can migrate the NSX-V Edge Services Gateway to NSX-T.

If you have no Edge Services Gateway appliances in your topology, you must still click **Start** so that you can proceed to the **Migrate Hosts** step.

If needed, you can roll back the Edge migration to use the Edge Services Gateway in the NSX-V environment. See Roll Back a Migration for more information.

Caution If you roll back the **Migrate Edges** step, verify that the traffic is going back through the NSX-V Edge Services Gateways. You might need to take manual action to assist the rollback.

Prerequisites

- All configuration issues must be resolved.
- The NSX-V configuration must be migrated to NSX-T.
- Verify that the migrated configurations are shown in the NSX Manager UI or API of NSX-T.
- Verify that you have a backup of NSX-V and vSphere since the most recent configuration changes were made.
- If you are using new IP addresses for the NSX-T Edge node uplinks, you must configure the northbound routers with these new BGP neighbor IP addresses.
- Verify that you have created an IP pool for Edge Tunnel End Points (TEP). See Create an IP Pool for Edge Tunnel End Points.
- Logical router interfaces created in NSX-T use the global default MTU setting, which is 1500. If you want to ensure that all logical router interfaces have a larger MTU, you can change the global default MTU setting. For more information, see Change the Global MTU Setting.

  If MTU setting other than 1500 is used on peering routers, the same should be configured on NSX-T Data Center. In case of OSPF topologies, OSPF adjacencies can get stuck if MTU setting is different from peering routers' MTU setting.

Procedure

1. From the **Migrate Edges** page, click **Start**.

   All Edges are migrated. The uplinks on the NSX-V Edge Services Gateways are internally disconnected, and the uplinks on the NSX-T Edge nodes are brought online.
2. Verify that routing and services are working correctly in the new NSX-T environment.

   If so, you can migrate the hosts. Before migrating the hosts, see Configuring NSX-V Host Migration.

Results

The following changes result from the migration process:

- The routing and service configuration from NSX-V Edge Services Gateway (ESG) are transferred to the newly created NSX-T Edge nodes.
- The new TEP IP addresses for the newly created NSX-T Edge nodes are configured from a newly created IP pool for Edge Tunnel End Points.

Configuring NSX-V Host Migration

The clusters in the NSX-V environment are displayed on the **Migrate Hosts** page. The clusters are arranged into migration groups, each migration group contains one vSphere host cluster. There are several settings which control how the host migration is performed.

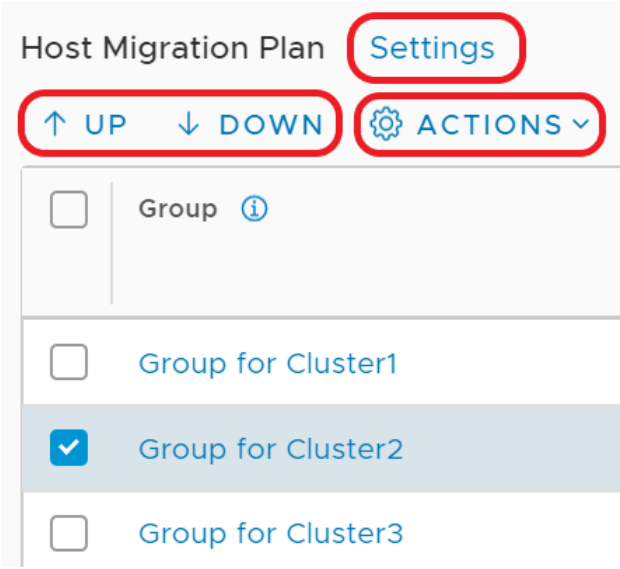

- Click **Settings** to change the global settings: **Pause Between Groups** and **Migration Order Across Groups**.
- Select a single host group (cluster) and use the arrows to move it up or down in the migration sequence.
- Select one or more host groups (clusters) and click **Actions** to change these host groups settings: **Migration Order Within Groups**, **Migration State**, and **Migration Mode**.

Pause Between Groups

When migrating multiple host groups, you can pause the migration between groups by enabling the **Pause Between Groups** setting. After a group is migrated, you must click **Continue** to migrate the next host group. This setting is disabled by default. You can enable it if you want to verify the status of the applications running on each cluster before proceeding to the next one.

Pause Between Hosts

Starting with NSX-T 3.2.2, you can pause the migration of hosts within a group by enabling the **Pause Between Hosts** setting. After a host is migrated, you must click **Continue** to migrate the next host in the group. With this feature, you can check the host that was migrated before continuing. This setting is only available if the migration order is **Serial**. You can change this setting before starting host migration or when host migration is paused. You cannot change this setting while host migration is in progress. This setting is disabled by default. To change this setting, select a group. Click **Actions > Change Migration Order Within Group**. Under **Serial**, click the **Pause Between Hosts** toggle.

Serial or Parallel Migration Order

You can specify whether migration happens in a serial or parallel order. There are two ordering settings:

- **Migration Order Across Groups** is a global setting that applies to all host groups.
  - **Serial**: One host group (cluster) at a time is migrated.
  - **Parallel**: Up to five host groups at a time are migrated. After those five host groups are migrated, the next batch of up to five host groups are migrated.

    Important If you are migrating from NSX-V 6.4.4, 6.4.5, or 6.4.6, and your environment uses vSphere Distributed Switch 7.0 or later, do not select parallel migration order across groups.

    If you are migrating from NSX-V 6.4.8 or later, and your environment uses vSphere Distributed Switch 7.0 or later, parallel migration order across groups is supported.
- **Migration Order Within Groups** is a host group (cluster) specific setting, so can be configured separately on each host group.
  - **Serial**: One host within the host group (cluster) at a time is migrated.
  - **Parallel**: Up to five hosts within the host group are migrated at a time. After those hosts are migrated, the next batch of up to five hosts are migrated.

    Important Do not select parallel migration order within groups for a cluster if you plan to use **Maintenance** migration mode for that cluster.

By default, both settings are set to **Serial**. Together, the settings determine how many hosts are migrated at a time.

Table 5-13. Effects of Migration Settings on Number of Hosts Attempting Migration Simultaneously

Migration Order Across Groups (Clusters)	Migration Order Within Groups (Clusters)	Maximum Number of Hosts Attempting Migration Simultaneously
Serial	Serial	1 One host from one host group
Serial	Parallel	5 Five hosts from one host group

Table 5-13. Effects of Migration Settings on Number of Hosts Attempting Migration Simultaneously (continued)

Migration Order Across Groups (Clusters)	Migration Order Within Groups (Clusters)	Maximum Number of Hosts Attempting Migration Simultaneously
Parallel	Serial	5 One host from five host groups
Parallel	Parallel	25 Five hosts from five host groups

Important If there is a failure to migrate a host, the migration process will pause after all in-progress host migrations have finished. If **Parallel** is selected for both migration across groups and migration within groups, there might be a long outage for the failed host before you can retry migration.

Sequence of Migration Groups

You can select a host group (cluster) and use the arrows to move it up or down in the list of groups.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

Migration State

Host groups (clusters) can have one of two migration states:

- **Enabled**

  Hosts groups with a migration state of **Enabled** are migrated to NSX-T when you click **Start** on the **Migrate Hosts** page.

- **Disabled**

  You can temporarily exclude host groups from migration by setting the migration state for the groups to **Disabled**. Hosts in disabled groups are not migrated to NSX-T when you click **Start** on the **Migrate Hosts** page. However, you must enable and migrate all **Disabled** host groups before you can click **Finish**. Finish all host migration tasks and click **Finish** within the same maintenance window.

In the **Resolve Configuration** step, hosts that are ineligible for migration are identified. In the **Migrate Hosts** step, these hosts have the migration state **Do not migrate**. For example, hosts that do not have NSX-V installed are not eligible for migration.

Migration Mode

Migration Mode is a host group (cluster) specific setting, and can be configured separately on each host group. In the **Migrate Hosts** step, you select whether to use **In-Place** or **Maintenance** mode.

There are two types of Maintenance migration modes:

- Automated
- Manual

In the **Resolve Configuration** step of the migration process, you select which type of Maintenance migration mode to use. You select a Maintenance mode even if you plan to migrate hosts using **In-Place** mode. When you select Maintenance migration mode in the **Migrate Hosts** step, the value you specified in the **Resolve Configuration step** determines whether Automated Maintenance mode or Manual Maintenance mode is used. However, in the **Migrate Hosts** step, if you select **In-Place** mode, your selected choice of Maintenance mode in the **Resolve Configuration** step does not take effect.

If you select the maintenance migration mode (manual or automatic), all new hosts must be put in maintenance mode before you add them to a cluster.

In-Place migration mode is not supported if your NSX-V installation uses vSphere Distributed Switch 7.0 or later.

If your environment uses Distributed Firewall, select **Automated Maintenance** migration mode. If you select a different migration mode, the following limitations apply to environments with Distributed Firewall:

- If you use **Manual Maintenance** migration mode, all VMs must be moved to NSX-T hosts, connected to NSX-T segments, and powered on before the last NSX-V host starts migrating. When you migrate your last NSX-V host, do not power off the VMs on the host. Move them to an NSX-T host using vMotion.
- If you use **Manual Maintenance** migration mode, VMs have a gap in firewall protection for up to 5 minutes after they move to an NSX-T host.
- If you use **In-Place** migration mode, and you have Distributed Firewall rules that are applied to a VM, those rules are not pushed to the host until the host and all its VMs are migrated. Until the rules are pushed to the host, the following applies:
  - If the NSX-T default rule is `deny`, the VM is not accessible.
  - If the NSX-T default rule is `accept`, the VM is not protected by the applied-to rules.

The migration process is different for each migration mode:

- **In-Place** migration mode

  NSX-T is installed and NSX components are migrated while VMs are running on the hosts. Hosts are not put in maintenance mode during migration. Virtual machines experience a short network outage and network storage I/O outage during the migration.

- **Automated Maintenance** migration mode

A task of entering maintenance mode is automatically queued. VMs are moved to other hosts using vMotion. Depending on availability and capacity, VMs are migrated to NSX-V or NSX-T hosts. After the host is evacuated, the host enters maintenance mode, NSX-T is installed, and NSX components are migrated. VMs are migrated back to the newly configured NSX-T host. Note that VMs that are powered off will not be reconfigured. After migration, you need to manually configure these VMs before powering them on.

When migrating NSX-V hosts that join VDS of 7.0 or later, the migrator will use/create a different VLAN transport zone for each VDS. The VLAN DVPGs in each VDS will be migrated to VLAN segments with the same VLAN in the VDS's own VLAN transport zone. VLAN segments of the same VLAN are not considered the same network because they have different policy paths. DRS does not considered two VLAN segments of the same VLAN the same or compatible network, so DRS will not map a DVPG in one VDS to a VLAN segment in anther VDS's VLAN transport zone when it vmotions VMs from one host to another host. In case two hosts in a cluster join a different VDS each, DRS will not choose one host as the target host when it tries to vmotion the VMs in the other host.

- **Manual Maintenance** migration mode

  A task of entering maintenance mode is automatically queued. To allow the host to enter maintenance mode, do one of the following tasks:

  - Power off all VMs on the hosts.
  - Move the VMs to another host using vMotion or cold migration.

  Once the host is in maintenance mode, NSX-T is installed on the host and NSX components are migrated. After the host is migrated, for the powered-off VMs and the VMs that you moved, you will need to change their network connection from the NSX-V logical switch to an NSX-T segment.

In the NSX-V environment, if the ESXi host's vmk0 management interface is connected to a VSS (vSphere Standard Switch) portgroup that does not have an uplink, and the portgroup is bridged to a VDS portgroup, and the VDS version is 6.5, 6.6 or 6.7, you must migrate using the **Maintenance** mode. If you use the **In-Place** mode, the migration will fail.

Adding or Removing a Host During Migration

Starting with NSX-T 3.2.1, during the host migration step, you can add or remove a host to be migrated when there is a pause in the migration.

Starting with NSX-T 3.2.2, you can pause the migration of hosts within a group. For more information, see Configuring NSX-V Host Migration.

Host migration will pause if you enable the setting **Pause Between Groups** or **Pause Between Hosts**, or if there is a failure to migrate a host.

You can add a host to a cluster before or after the cluster has been migrated, or while the cluster is being migrated. You can also remove a host from a cluster that has not migrated or is being migrated.

The host you want to add can be a standalone host or in a cluster. The host must not have NSX-V or NSX-T configured. If the host is in a cluster, the cluster must not have NSX-V or NSX-T configured.

You must prepare the host as a transport node first and later move it into a target cluster that has been migrated to NSX-T, is being migrated or has not started the migration. If the target cluster already has a host migrated to an NSX-T transport node, you can use the transport node's configuration as a reference to prepare the host as a transport node. If the host will support overlay traffic and a VTEP IP pool will be used to prepare the host as a transport node, the VTEP IP pool cannot be or overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated. It can be an existing IP pool (such as the one used by NSXT Edge nodes) or a newly created IP pool. For more information, see the section "Preparing ESXi and KVM Hosts as Transport Nodes" in the *NSX-T Data Center Installation Guide*.

When preparing an ESXi host as a transport node, you can choose N-VDS or VDS as the host switch. Choose VDS if the version of the VDS being migrated is 7.0 or later. Otherwise, choose N-VDS. If you prepare a host with N-VDS when you should choose VDS, the host will still be migrated but it may have network issues.

Adding a host to a cluster

1 From vCenter Server, put the host in Maintenance Mode.

2 If there is no VTEP IP pool that does not overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated, or such pools do not have enough free IPs for the VTEPs to be created for the host to be added, then go to **Networking > IP Address Pools** and create a new VTEP IP pool.

3 Follow the instructions in the installation guide to prepare the host as a transport node. When you select the transport zone, if an overlay transport zone is chosen for the host switch, choose the new IP pool that was created in step 2 or choose an existing IP pool that does not overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated. Select an uplink profile. Do not choose the one whose name contains "VXLAN" if an overlay transport zone is chosen for the host switch.

4 Wait for the status of the host node to be "Success". Do not move the host out of Maintenance Mode.

5 Choose a cluster into which the host will be added. In NSX-T Manager UI, go to **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later) and check if the cluster has a Transport Node Profile (TNP) attached. Detach the TNP if it does.

6 In vCenter Server UI, move the host into the chosen cluster.

7 Invoke the sync host groups API or click the **Refresh** button on the NSX-T Manager UI host migration screen so that the migration group for the cluster contains the new host.

8 Call the following NSX-T API to accept the new host:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/setup?
action=migrate_newly_added_host_transport_node
{
    "host_transport_node_id" : "<transport-node-uuid>"
}
```

If the API returns an error, fix the error and retry the API. If the API returns success, then make the host exit maintenance mode.

9 vMotion VMs to the host. If any VM is moved from an NSX-V host, be sure to change the network to map the source virtual-wire to NSX-T overlay segments. For example, virtual-wire vxw-dvs-64-virtualwire-4-sid-10787-1-switch-191 must map to 1-switch-191-LS.

Removing a host from a cluster

1 From vCenter Server, migrate all VMs off the host and enter the host into maintenance mode.

2 If the host is in a cluster that has not started migrating, go to the next step. Otherwise, from NSX-T Manager UI, go to **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later). If the host has NSX configured, delete it. If the host is not reachable by NSX Manager, delete it with the force option.

3 From vCenter Server, remove the host. Wait until the task is complete.

4 Click the Refresh button on the host migration screen to remove the host from the migration group.

5 Restart the host migration.

Migrate NSX-V Hosts

After you have migrated Edge Services Gateway VMs to NSX-T Edge nodes, and verified that routing and services are working correctly, you can migrate your NSX-V hosts to NSX-T host transport nodes.

You can configure several settings related to the host migration, including migration order and enabling hosts. Make sure that you understand the effects of these settings. See Configuring NSX-V Host Migration for more information. Understanding the host migration settings is especially important if you use Distributed Firewall or vSphere Distributed Switch 7.0 or later.

For more information about what happens during host migration, see Changes Made During Host Migration in an End-to-End Migration.

If the Security Policies in your NSX-V environment use a partner service for Guest Introspection or Network Introspection or both, choose the host migration mode as shown in the following table.

Partner Service	Host Migration Mode
Only Guest Introspection	In-Place and Maintenance modes are supported.
Only Network Introspection	Maintenance mode is supported. However, Automated Maintenance mode is recommended. In-Place mode is not supported.
Both Guest Introspection and Network Introspection	Maintenance mode is supported. In-Place mode is not supported.

Important Consult the VMware partner before migrating their service that is running on the NSX-V workloads. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their services to NSX-T.

Caution Host migration should be completed during the same maintenance window as Edge migration.

You must disable IPFIX and reboot the ESXi hosts before migrating them.

If the partner service in your NSX-V environment provides Guest Introspection or both Guest Introspection and Network Introspection service, follow the procedure in this topic to migrate cluster-by-cluster. After all the host clusters are migrated to NSX-T, do a host-based service deployment in each NSX-T cluster.

If the partner service in your NSX-V environment provides only Network Introspection service, use the host migration approaches that are explained in Migrate Hosts with Network Introspection Service.

Prerequisites

- Verify that Edge migration has finished and all routing and services are working correctly.
- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.

Procedure

1 On the **Migrate Hosts** page, click **Start**.

If you selected the **In-Place** or **Automated Maintenance** migration mode for all hosts groups, the host migration starts. Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

2 If you selected the **Manual Maintenance** migration mode for any host groups, you must complete one of the following tasks for each VM so that the hosts can enter maintenance mode.

Option	Action
Power off or suspend VMs.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b After the host has migrated, attach the VM interfaces to the appropriate NSX-T segments and power on the VM.
Move VMs using vMotion.	Right click the VM and select Migrate. Follow the prompts to move the VM to a different host. Note that Migration Coordinator maintains security during migration by vMotioning VMs to specific ports that are protected by temporary rules. In the case of manual vMotion, the VMs will not be moved to those ports and there could be a security breach. To vMotion manually, the VMs must be migrated using vSphere API where the networking backing must point to the OpaqueNetwork ID corresponding to the NSX Segment when using NVDS or the VDS portgroup ID when using VDS 7. In both cases, the network device's externalId must be set to the the string "VM_UUID:vNIC_ID", where VM_UUID is the VM's instance UUID and vNIC_ID is the VM's vNIC index where the first vNIC is 4000.
Move VMs using cold migration.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b Right click the VM and select Migrate. Follow the prompts to move the VM to a different host, connecting the VM interfaces to the appropriate NSX-T segments.

Here is python code to specify an external-id for each vNIC in a VM and then vMotion the VM so that the vNICs will connect to an NSX-T segment of ID “ls_id” at the correct ports:

```
devices = vmObject.config.hardware.device
nic_devices = [device for device in devices if isinstance(device,
vim.Vm.device.VirtualEthernetCard)]
vnic_changes = []
for device in nic_devices:
    vif_id = vmObject.config.instanceUuid + ":" + str(device.key)
    vnic_spec = self._get_nsxt_vnic_spec(device, ls_id, vif_id)
    vnic_changes.append(vnic_spec)
relocate_spec = vim.Vm.RelocateSpec()
relocate_spec.SetDeviceChange(vnic_changes)
# set other fields in the relocate_spec
vmotion_task = vmObject.Relocate(relocate_spec)
WaitForTask(vmotion_task)

def _get_nsxt_vnic_spec(self, device, ls_id, vif_id):
    nsxt_backing = vim.Vm.Device.VirtualEthernetCard.OpaqueNetworkBackingInfo()
    nsxt_backing.SetOpaqueNetworkId(ls_id)
    nsxt_backing.SetOpaqueNetworkType('nsx.LogicalSwitch')
    device.SetBacking(nsxt_backing)
    device.SetExternalId(vif_id)
```

```
    dev_spec = vim.Vm.Device.VirtualDeviceSpec()
    dev_spec.SetOperation(vim.Vm.Device.VirtualDeviceSpec.Operation.edit)
    dev_spec.SetDevice(device)
    return dev_spec
```

For an example of a complete script, see https://github.com/dixonly/samples/blob/main/vmotion.py

The host enters maintenance mode after all VMs are moved, powered off, or suspended. If you want to use cold migration to move the VMs to a different host before the migrating host enters maintenance mode, you must leave at least one VM running while you move VMs. When the last VM is powered off or suspended, the host enters maintenance mode, and migration of the host to NSX-T starts.

Results

After a host has migrated to NSX-T using **In-Place** migration mode, you might see a critical alarm with message `Network connectivity lost`. This alarm occurs when a vSphere Distributed Switch (VDS) 6.5 or 6.7 migrates to an N-VDS because the host no longer has a physical NIC connected to the VDS it was previously connected to. To restore the migrated hosts to the Connected state, click **Reset to Green** on each host, and suppress the warnings, if any.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

If migration fails for a host, the migration pauses after all in-progress host migrations finish. When you have resolved the problem with the host, click **Retry** to retry migration of the failed host. If the host still fails to migrate, you can configure NSX-T on the host manually or remove the host from the system. In this case, at the end of the host migration step, the **Finish** button will not be enabled because of the host that failed to migrate. You need to call the REST API `POST https://<nsx-mgr-IP>/api/v1/migration?action=finalize_infra` (<nsx-mgr-IP> is the IP address of the NSX Manager where the migration service is running) using a REST API client (for example, postman or curl) to finish the migration, and then perform the post-migration tasks.

For information about troubleshooting other host migration problems, see Chapter 13 Troubleshooting Migration Issues.

What to do next

If the migrated Security Policies use a third-party partner service, deploy an instance of the partner service in NSX-T. For detailed instructions, see:

- Deploy a Partner Service for Endpoint Protection

  Click this link to deploy a partner service that provides only Endpoint Protection or both Endpoint Protection and Network Introspection services to the NSX-T workload VMs.

- Deploy a Partner Service for Network Introspection

  Click this link to deploy a partner service that provides only Network Introspection service to the NSX-T workload VMs.

Migrate Hosts with Network Introspection Service

When Security Policies in your NSX-V environment use only a Network Introspection service that is provided by a partner, two approaches are available to migrate the NSX-V prepared hosts to NSX-T.

Both the approaches discussed in this topic assume that the partner service virtual machines (SVMs) in your NSX-V environment are not deleted before starting the migration coordinator. Depending on how much security protection downtime you are willing to accept during host migration, choose the host migration approach that best suits your needs.

Note Consult the VMware partner before migrating the hosts by using any of the two approaches. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their service to NSX-T.

When only Network Introspection service is running on your NSX-V hosts, **In-Place** host migration mode is not supported. Only **Maintenance** migration mode is supported. However, Automated Maintenance migration mode is recommended.

Approach 1: Involves more security protection downtime

This approach is the simpler of the two host migration approaches. However, it involves more security protection downtime compared to Approach 2. Let us say that you have three clusters in your NSX-V environment: Cluster 1, Cluster 2, and Cluster 3.

In this approach, enable the **Pause between groups** migration setting and migrate Cluster 1 by using the standard host migration procedure that is explained in Migrate NSX-V Hosts. After Cluster 1 is migrated to NSX-T, the migration pauses. Deploy the partner service in Cluster 1 by doing either a host-based or a clustered service deployment. Now, disable the **Pause between groups** migration setting, and continue migrating Clusters 2 and 3. After the workload VMs in Clusters 2 and 3 are migrated to NSX-T, these workloads can start redirecting packets to the partner service virtual machines (SVM) in Cluster 1.

In this approach, security protection downtime is expected during migration of Cluster 1.

When workload VMs migrate to an NSX-T host, existing data traffic during a host migration is expected to have a security protection downtime. However, new data traffic does not have a security protection downtime.

Approach 2: Involves minimal security protection downtime

This approach requires some manual intervention with an NSX-T API to create a temporary host group. Enable the **Pause between groups** migration setting and migrate any one host from Cluster 1. After this host in Cluster 1 is migrated to NSX-T, the migration coordinator pauses. Deploy a partner service on this migrated host by doing a clustered service deployment. Continue migrating the remaining hosts in Cluster 1. After all the hosts in Cluster 1 are migrated to NSX-T, you can optionally deploy additional partner SVMs in Cluster 1 by doing either a host-based or a clustered service deployment.

The detailed procedure in this topic explains the host migration workflow for a single NSX-V prepared cluster, which has three hosts, as shown in the following figure. The procedure uses Approach 2 to migrate this Cluster 1 to NSX-T.

Example:

Figure 5-7. Host Group 1 (Cluster 1) Before Migration

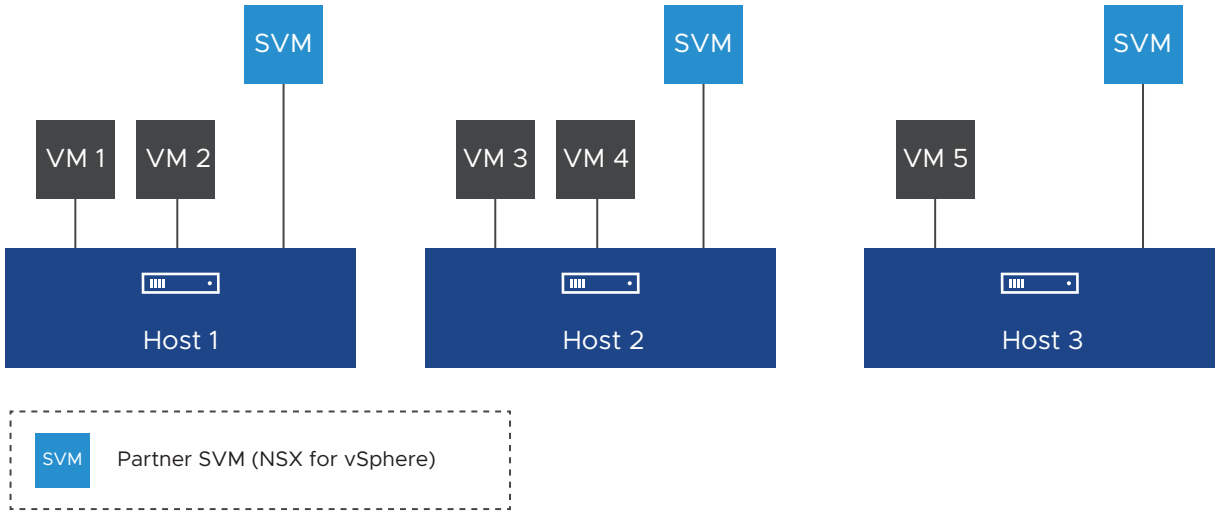

All hosts in this cluster are ESXi hosts. The Security Policies in your NSX-V environment redirect data traffic to partner service virtual appliances that provide a network introspection service to workloads. As NSX-V supports only a host-based service deployment, each host has a single partner service VM.

The following configuration settings are required for migrating hosts using Approach 2:

- Host migration mode is set to Automated Maintenance.
- Pause between groups is enabled.
- Migration order across groups is set to serial.

Prerequisites

- Verify that Edge migration has finished and all routing and services are working correctly.

- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.
- Enable vSphere DRS on the cluster that is being migrated.
- Enable vMotion on the VMkernel adapter of each host in the cluster.
- Ensure that adequate spare capacity is available in the NSX-V cluster so that the migrating hosts can enter into a maintenance mode. If enough spare capacity is unavailable to migrate NSX-V workload VMs to other hosts in the cluster, additional security protection downtime is expected.

Procedure

1 Run the following API request to create a temporary host group and move hosts 2 and 3 to this temporary group.

```
POST https://{nsxt-mgr-ip}/api/v1/migration/migration-unit-groups
```

In the request body of this POST API, specify the following details:

- Name of the temporary host group
- Migration units (IDs of hosts 2 and 3)
- Migration state of the temporary group (must be disabled)

For a detailed information about this API and an example POST API request, see the *NSX-T Data Center API Guide*.

You can obtain the host IDs from the vCenter Server Managed Object Browser (MOB) at `http://{vCenter-IP-Address}/mob`, or run the following GET API to retrieve the host IDs:

```
GET https://{nsxt-mgr-ip}/api/v1/fabric/discovered-nodes
```

A temporary host group is created and displayed on the **Migrate Hosts** page. The original host group 1 (cluster 1) now contains only host 1.

2 On the **Migrate Hosts** page, next to **Host Migration Plan**, click **Settings** and ensure that the settings are configured as follows:

- Pause between groups: Enabled
- Migration order across groups: Serial

3 Migrate host 1 to NSX-T.

a Click **Start** to start the host migration.

Workload VMs 1 and 2 are migrated to other hosts so that host 1 can enter into a maintenance mode. NSX-V partner SVM on host 1 is powered off before host 1 enters into a maintenance mode.

Assume that VMs 1 and 2 are migrated to host 3 that is prepared with NSX-V. After the migration of host 1 is successful, the migration coordinator pauses for your next input.

b (Required) Deploy a partner service on host 1 by using the clustered deployment approach.

At this stage, host-based service deployment is not supported. Deploying a partner service on host 1 is necessary to minimize security protection downtime. Remember, the security protection for NSX-V workloads that are running on hosts 2 and 3 is still intact. The partner must ensure that the migrated partner-specific Security Policies are available on the newly deployed partner SVMs on host 1.

For detailed steps on deploying a partner service in NSX-T, see Deploy a Partner Service for Network Introspection. For example, specify the following configuration settings to deploy two partner service virtual machines (SVMs) on host 1:

Configuration	Value
Deployment Type	Clustered
Host	Host 1
Clustered Deployment Count	2

The value that you enter in the **Clustered Deployment Count** text box depends on the resource capacity that is available on the host. This scenario assumes that two partner SVMs can be deployed on Host 1. This value can be different in your environment.

After this step, the cluster looks as shown in the following figure. The green colored host represents the migrated host.

Figure 5-8. Host 1 is Migrated to NSX-T

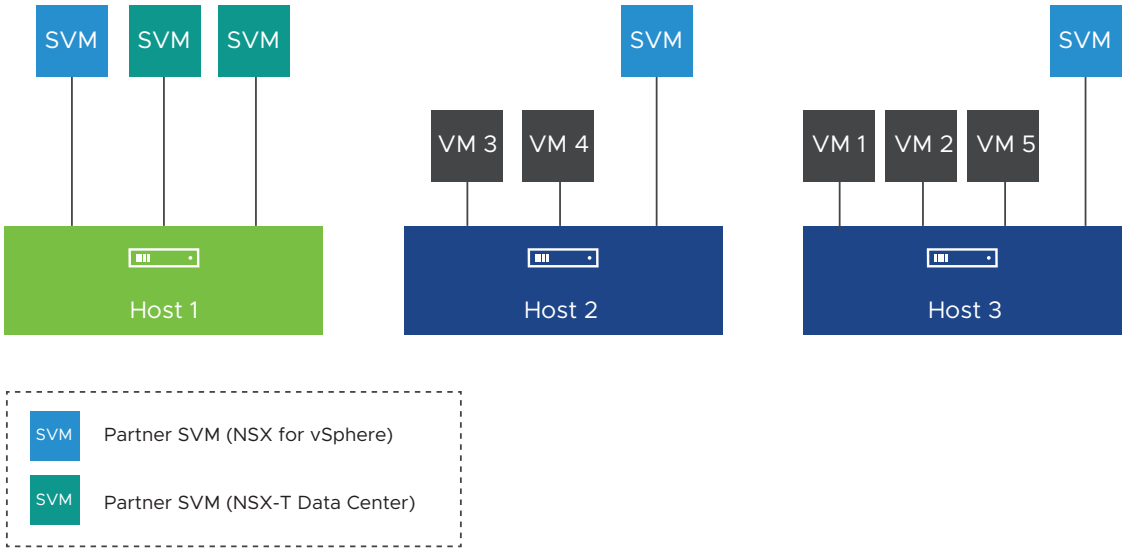

4 Migrate host 2 and host 3 to NSX-T.

a Move host 2 and host 3 from the temporary host group to the original host group 1 by running the following POST API request:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/migration-unit-groups/
{group-id}?action=add_migration_units
```

Where: *group-id* is the ID of the destination host group (host group 1). In the POST API request body, specify the ID of hosts 2 and 3 that you want to add to the original host group 1.

For a detailed information about this POST API and an example POST API request, see the *NSX-T Data Center API Guide*.

Now, the original host group 1 contains hosts 1, 2 and 3 (in the given order), and the temporary host group is deleted.

b Select the check box next to the original host group 1, and then click **Actions > Change Migration Order Within Group**. Verify that the migration order within the group is set to **Serial**.

If necessary, you can set the migration order in the original host group 1 to **Parallel**.

c Click **Continue** to resume the host migration.

Host 2 is first migrated to NSX-T, and then host 3 is migrated. To put each migrating host into a maintenance mode, workload VMs on the migrating host are moved to either NSX-V hosts or NSX-T hosts. NSX-V partner SVM on the migrating host is also powered off before the host enters into a maintenance mode.

After this step, all the hosts in host group 1 (cluster 1) are prepared with NSX-T. The cluster looks as shown in the following figure.

Figure 5-9. All Hosts are Migrated to NSX-T

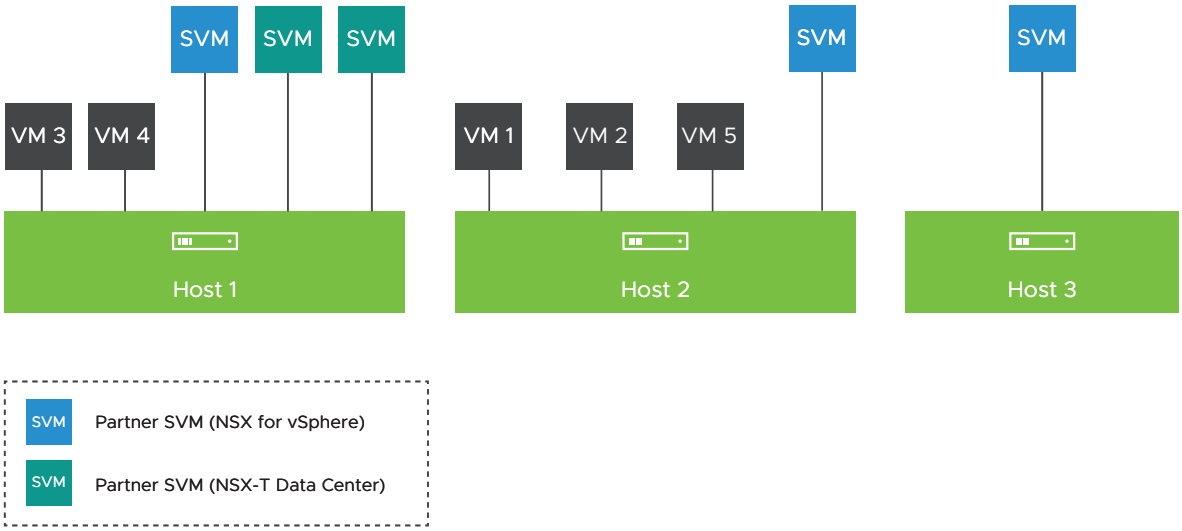

5 (Optional) Migrate some workload VMs from hosts 1 and 2 to host 3.

For example, migrate VMs 4 and 5 to host 3, as shown in the following figure.

Figure 5-10. Final Cluster 1 After Migration

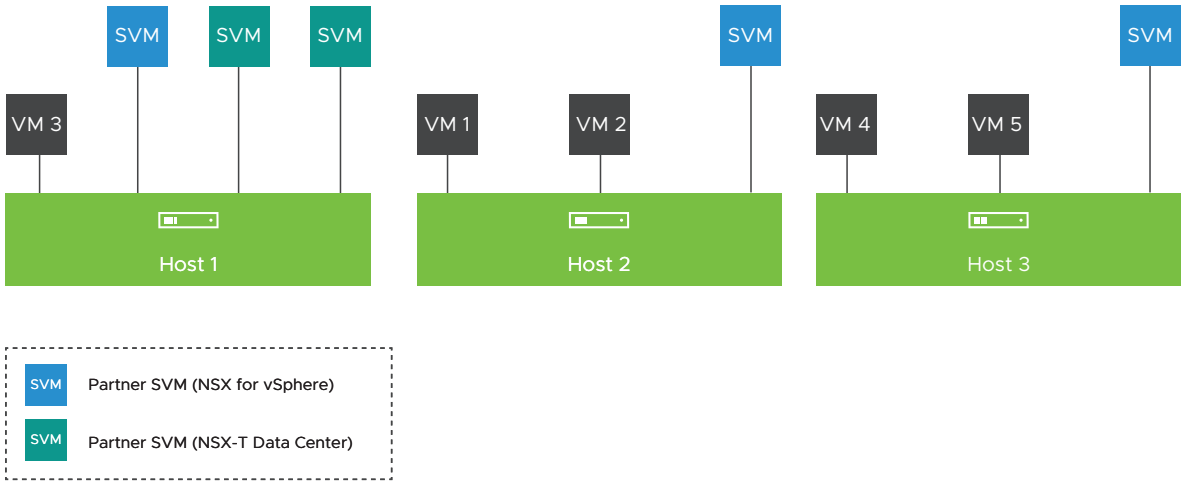

6 (Optional) After all the hosts in host group 1 are migrated to NSX-T, you can do either a host-based or a clustered service deployment.

A host-based service deployment allows new network traffic to be protected by a local partner SVM on each host.

Note If you have network introspection service running on more than one NSX-V prepared cluster, you do not have to deploy the partner SVMs in the other clusters. The network traffic though the NSX-T workload VMs in the other clusters can use the partner SVMs in cluster 1 that you just migrated. The host migration workflow covered in this procedure is required only for the first cluster. You can migrate the remaining clusters by using the standard host migration procedure.

What to do next

Delete the partner service deployment in NSX-V. Remember, you can delete the partner SVMs only at a cluster level. That is, you can delete service deployment only after all the hosts in the host group 1 are migrated to NSX-T. Complete the following steps to delete the service deployment in NSX-V:

1 Log in to the vSphere Client and navigate to **Networking and Security > Installation and Upgrade > Service Deployment**.

2 Select the deployed partner service, and click **Delete**.

Deploy a Partner Service for Endpoint Protection

When migrated Security Policies in NSX-T use a partner service that provides only Endpoint Protection or both Endpoint Protection and Network Introspection, deploy an instance of the partner service after all the clusters are migrated to NSX-T.

Only a host-based service deployment is supported.

In a host-based service deployment, one partner service virtual machine is installed on each host of the migrated cluster. In the vCenter Server, the vSphere ESX Agency Manager (EAM) service is internally used to deploy a partner service VM on each host of the cluster.

Prerequisites

- All the hosts in the cluster are migrated to NSX-T.
- All the migrated hosts are managed by a vCenter Server.
- A transport node profile is applied to the cluster.

Procedure

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

2 Navigate to **System > Service Deployments > Deployment**.

3 In the **Partner Service** drop-down menu, select the partner service to be deployed, and click **Deploy Service**.

4 Enter the service deployment name.

5 Select the vCenter Server that is registered as a compute manager in NSX-T.

6 Select the cluster where you want to deploy the partner service.

7 To specify the datastore, do one of the following actions:

- Select a datastore as the repository for the service virtual machines.
- Select **Specified on Host**.

  The **Specified on Host** option means that you do not need to select a datastore and network on the **Deploy Service** page. Before deploying the partner service, you must configure Agent VM settings on each ESXi host to point to a specific datastore and network.

  To know more about configuring Agent VM settings, see the vSphere product documentation.

8 Under Networks, click **Set** and select the NICs you want to use for deployment.

a Select the network for the Management interface.

  In a host-based deployment, if you set the datastore as **Specified on Host**, you must set the network also as **Specified on Host**.

b Set the Network type to **DHCP** or **Static IP Pool**. If you set the network type to a Static IP Pool, select from the list of available IP pools.

9 In the **Deployment Template** drop-down menu, select the registered deployment template and click **Save**.

  The deployment process might take some time depending on the vendor's implementation.

10 Check the deployment status on the **Deployment** page. Wait until the status changes to Up.

  You might have to refresh the **Deployment** page a few times to retrieve the latest status.

  If the Status column shows Down, click the icon next to Down. All deployment errors are displayed. Take the required actions to fix the errors, and click **Resolve**. The status changes to In Progress. Wait until the status changes to Up.

Results

A partner service VM is now deployed on all the hosts of the cluster.

Note When you add a new host in the cluster, EAM automatically deploys the partner service VM on the new host.

What to do next

Go to the Partner Console and verify whether the endpoint protection service is activated. Now, the migrated endpoint protection rules are enforced on the workload VMs that are running on the NSX-T prepared cluster.

For more information about activating the endpoint protection service in the Partner Console, see the partner documentation.

Deploy a Partner Service for Network Introspection

When migrated Security Policies in NSX-T use a third-party partner service only for Network Introspection, deploy an instance of the partner service either by using a clustered service deployment or a host-based service deployment approach.

Prerequisites

For a clustered service deployment approach:

- At least one host in the first cluster is migrated to NSX-T.

For a host-based service deployment approach:

- All the hosts in a cluster are migrated to NSX-T.
- A transport node profile is applied to the cluster.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Navigate to **System > Service Deployments > Deployment**.
3. In the **Partner Service** drop-down menu, select the partner service to be deployed, and click **Deploy Service**.
4. Enter the service deployment name.
5. Select the vCenter Server that is registered as a compute manager in NSX-T.
6. Select a deployment type: **Host-Based** or **Clustered**.
7. Select the cluster where you want to deploy the partner service.
8. (Clustered deployment only): In the **Host** drop-down menu, select a host, or select **Any** to allow the NSX-T NSX Manager to select a host.
9. In the **Data Store** drop-down menu, select a data store as the repository for the partner service virtual machine (SVM).
    - Clustered deployment: If you selected **Any** for the host, select a shared data store. If you specified a particular host, select a local data store.

- Host-based deployment: Select a specific datastore or select **Specified on Host**. The **Specified on Host** option means that you do not need to select a datastore and network on the **Deploy Service** page. Before deploying the partner service, you must configure Agent VM settings on each ESXi host to point to a specific datastore and network.

  To know more about configuring Agent VM settings, see the vSphere product documentation.

10 Under Networks, click **Set** and select the NICs you want to use for deployment.

   a Select the network for the Management interface.

     In a host-based deployment, if you set the datastore as **Specified on Host**, you must set the network also as **Specified on Host**.

   b Set the Network type to **DHCP** or **Static IP Pool**. If you set the network type to a Static IP Pool, select from the list of available IP pools.

11 In the **Deployment Template** drop-down menu, select the required template.

   Typically, the deployment specification and the deployment template fields are automatically selected with the information that is pushed from the Partner Console as part of the service definition.

12 In the **Service Segment** drop-down menu, select the service segment that the migration coordinator has created in the overlay transport zone.

13 (Clustered deployment only): In the **Clustered Deployment Count** text box, specify the number of service VMs to deploy in the cluster, and click **Save**.

14 Check the deployment status on the **Deployment** page. Wait until the status changes to Up.

   You might have to refresh the **Deployment** page a few times to retrieve the latest status.

   If the Status column shows Down, click the icon next to Down. All deployment errors are displayed. Take the required actions to fix the errors, and click **Resolve**. The status changes to In Progress. Wait until the status changes to Up.

Finish the Migration

After you have migrated all Edge Services Gateway VMs and hosts to the NSX-Tenvironment, confirm that the new environment is working correctly. If everything is functioning correctly, you can finish the migration.

Important Verify that everything is working before clicking **Finish**. Then perform the post-migration tasks. Do not make any vSphere life cycle operations such as upgrading ESXi hosts, VDS, or VC before the post-migration tasks are completed.

You will see errors on hosts after the migration. The error message is: `UserVars.RmqHostId' is invalid or exceeds the maximum number of characters permitted`. The error occurs because this host is still part of the NSX-V inventory.

Prerequisites

- Verify that all expected items have been migrated to the NSX-T environment.
- Verify that the NSX-T environment is working correctly.

Procedure

1 Navigate to the **Migrate Hosts** page of the migration coordinator.

2 Click **Finish**

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page, or which hosts were excluded from the migration.

Post-Migration Tasks

After migration has finished, some additional actions might be required.

- If you migrated from NSX-V 6.4.4, perform a reboot of all hosts that have migrated to NSX-T. The reboot must be done before you upgrade to a later version of NSX-T.
- During migration, all transport nodes are added to a group called `NSGroup with TransportNode for CPU Mem Threshold`. This group ensures that the transport nodes have the correct CPU memory threshold settings in NSX-T. This group is required after migration has completed. If you need to remove a transport node from NSX-T after migration and you are running NSX-T 3.2.0, you must first remove the transport node from this group. If you are running NSX-T 3.2.1 or later, you do not need to remove the transport node from this group.

  To remove the transport node from the group, make sure you are in **Manager** mode and then select **Inventory > Groups** to remove the transport node from the `NSGroup with TransportNode for CPU Mem Threshold` group. For more information about Manager mode, see the topic "NSX Manager" in the *NSX-T Data Center Administration Guide*.
- Verify that you have a valid backup and restore configuration. See "Backing Up and Restoring the NSX Manager" in the *NSX-T Data Center Administration Guide*.

Finish Deploying the NSX Manager Cluster

Deploy two additional NSX Manager appliances before using your NSX-T environment in production.

See the *NSX-T Data Center Installation Guide* for the following information:

- *NSX Manager Cluster Requirements*
- *Deploy NSX Manager Nodes to Form a Cluster from UI*
- *Configure a Virtual IP (VIP) Address for a Cluster*

Uninstalling NSX-V After Migration

When you have verified that the migration is successful, and have clicked **Finish** to finish the migration, you can uninstall your NSX-V environment.

The process for uninstalling NSX-V after migration to NSX-T is different from the standard uninstall for NSX-V.

Important If you have vCenter Enhanced Linked Mode (ELM) configured, you must migrate all the NSX-V instances associated with the vCenter ELM chain before executing steps 6, 7, and 8 in the procedure below.

Prerequisites

- Verify that the migration is successful, and all functionality is working in the NSX-T environment.
- Verify that you have clicked **Finish** on the **Migrate Hosts** page.

Procedure

1. In the vSphere client, navigate to **Networking and Security > NSX Edges** and delete all the NSX Edges.
2. In the vSphere client, navigate to **Networking and Security > Logical Switches** and delete all the logical switches.
3. In the vSphere client, navigate to **Networking and Security > Installation and Upgrade > Logical Network Settings > Transport Zones** and delete all the transport zones.
4. In the vSphere client, navigate to **Networking and Security > Installation and Upgrade > Management > NSX Controller Nodes** and delete all the NSX Controllers.
5. Clear all stale VTEPs that may remain in the NSX-V Manager database:
   a. SSH into NSX-V Manager as **root**.
   b. Run the following command to clear the database table:

      ```
      psql -U secureall -d secureall -c "delete from xvs_vmknic_info;"
      ```

   c. Run the following command to confirm that the output shows zero row:

      ```
      psql -U secureall -d secureall -c "select * from xvs_vmknic_info;"
      ```

6. Delete the ESX Agent Manager agencies that are associated with the NSX-V environment.
   a. In the vSphere Client, navigate to **Menu > Administration**. Under **Solutions**, click **vCenter Server Extensions**. Double-click **vSphere ESX Agent Manager** and click the **Configure** tab.
   b. For each agency that has a name starting with `_NSX_`, select the agency, then click the three-dot menu () and select **Delete Agency**.

7 Remove the NSX-V plug-in from vCenter Server.

a Access the Extension Manager from the Managed Object Browser at `https://<vcenter-ip>/mob/?moid=ExtensionManager`.

b Click **UnregisterExtension**.

c In the **UnregisterExtension** dialog box, enter **`com.vmware.vShieldManager`** in the **Value** text box and click **Invoke Method**.

d In the **UnregisterExtension** dialog box, enter **`com.vmware.nsx.ui.h5`** in the **Value** text box and click **Invoke Method**.

e You can verify that you unregistered the extensions by going to the Extension Manager page at `https://<vcenter-ip>/mob/?moid=ExtensionManager` and viewing the values for the **extensionList** property.

8 Delete the vSphere Web Client directories and vSphere Client (HTML5) directories for NSX for vSphere and then restart the client services.

a Connect to the vCenter Server system command line.

- If you are using a vCenter Server Appliance, log in as root using the console or SSH. You must log in as root and run the commands from the Bash shell. You can start the Bash shell using the following commands.

```
> shell.set --enabled True
> shell
```

- If you are using vCenter Server for Windows, log in as an administrator using the console or RDP.

b Delete all NSX for vSphere plug-in directories.

Note A plug-in directory might not be present if you have never launched the associated client.

On vCenter Server Appliance, delete the following directories:

- To remove the vSphere Web Client plug-in, delete the `/etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/com.vmware.vShieldManager-<version>-<build>` directory.
- To remove the vSphere Client plug-in, delete the `/etc/vmware/vsphere-ui/vc-packages/vsphere-client-serenity/com.vmware.nsx.ui.h5-<version>-<build>` directory.

On vCenter Server for Windows, delete the following directories:

- To remove the vSphere Web Client plug-in, delete the `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity\com.vmware.vShieldManager-<version>-<build>` directory.
- To remove the vSphere Client plug-in, delete the `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-ui\vc-packages\vsphere-client-serenity\com.vmware.nsx.ui.h5-<version>-<build>` directory.

c Restart the client services on the vCenter Server Appliance or vCenter Server on Windows.

Table 5-14. Client Service Commands

Client Service	vCenter Server Appliance	vCenter Server for Windows
Restart vSphere Web Client On vSphere 6.0, 6.5, and 6.7	`> shell.set --enabled True`	`> cd C:\Program Files\VMware\vCenter`

Table 5-14. Client Service Commands (continued)

Client Service	vCenter Server Appliance	vCenter Server for Windows

	`> shell`
	`# service-control --stop vsphere-client`
	`# service-control --start vsphere-client`
		`Server\bin`
		`> service-control --stop vspherewebclientsvc`
		`> service-control --start vspherewebclientsvc`

Restart vSphere Client
On vSphere 6.5, and 6.7
	`> shell.set --enabled True`
	`> shell`
	`# service-control --stop vsphere-ui`
	`# service-control --start vsphere-ui`
		`> cd C:\Program Files\VMware\vCenter Server\bin`
		`> service-control --stop vsphere-ui`
		`> service-control --start vsphere-ui`

Restart vSphere Client
On vSphere 7.0
	`> shell.set --enabled True`
	`> shell`
	`# service-control --stop vsphere-ui`
	`# service-control --start vsphere-ui`
		vSphere 7.0 does not support vCenter Server for Windows

9 Power off and delete the NSX Manager VM.

a In vSphere client, navigate to **Hosts and Clusters**.

b Locate the NSX Manager VM. Right click and select **Power Off** then right click and select **Delete from Disk**.

Migrating a User-Defined Topology

6

If the topology of your NSX-V environment is not one of the supported fixed topologies, you can migrate using the user-defined topology option. You must also use this option if you want to migrate a cross-vCenter environment to NSX Federation.

When migrating a user-defined topology, you can do an end-to-end migration or a lift-and-shift migration (see Chapter 1 Migration Modes).

In a lift-and-shift migration, the default segments in NSX-T do no support DHCP servers and will result in the servers being down after migration.

To minimize disruption during migration, ensure that:

- NSX-V and NSX-T edges are on different ESXi hosts.
- Workload VMs directly connected to an Edge Services Gateway (ESG) are on a different ESXi host than the ESG.

Read the following topics next:

- Overview - Migrating a User-Defined Topology
- Preparing for a User-Defined Topology Migration
- Configuration and Edge Migration Workflow
- Migrating a Cross-vCenter Environment to NSX Federation
- Preparing the NSX-V Environment for a User-Defined Topology End-to-End Migration
- Preparing the NSX-V Environment for a User-Defined Topology Lift-and-Shift Migration
- Preparing the NSX-T Data Center Environment for a User-Defined Topology Migration
- Preparing the NSX-T Data Center Environment for a User-Defined Topology End-to-End Migration
- Preparing the NSX-T Environment for a User-Defined Topology Lift-and-Shift Migration
- Import Configuration
- Translate Configuration Layer 2
- Resolve Configuration Layer 2
- Migrate Configuration Layer 2

- Check Realization Layer 2
- Define a Topology
- Translate Configuration Layer 3 and Above
- Resolve Configuration Layer 3 and Above
- Migrate Configuration Layer 3 and Above
- Check Realization Layer 3 and Above
- Migrate NSX-V Edges in End-to-End Migration
- Migrating Hosts in End-to-End Migration
- Post-Migration Tasks in End-to-End Migration
- Prepare Infrastructure in Lift-and-Shift Migration
- Migrate Edges
- Migrate Workloads in Lift-and-Shift Migration

Overview - Migrating a User-Defined Topology

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.
- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.

- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 6-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.

Table 6-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. The migration of DHCP IP pools on Global Manager is not supported. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. Migration Coordinator will only migrate from an NSX-V Manager with the role of Primary or Standalone. You can modify the NSX-V environment by changing the status of the secondary managers in order to migrate each NSX-V environment independently.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	

NSX-V Configuration	Supported	Details
Network I/O Control (NIOC) having vNIC with reservation	No	
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	
Host profiles	No	
ESXi lockdown mode	No	Not supported in NSX-T.
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.

NSX-V Configuration	Supported	Details
VXLAN port number other than 4789	No	
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes for in-place migration No for lift-and-shift migration	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	

NSX-V Configuration	Supported fro Migration	Details
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	

NSX-V Configuration	Supported	Details
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	

NSX-V Configuration	Supported	Details
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description

NSX-V Configuration	Supported	Details
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	

NSX-V Configuration	Supported	Details
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 6-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. DHCP Relay configured on multiple Distributed Logical Routers pointing to the same DHCP Server configured on an Edge Services Gateway is not supported. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 6-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.

DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration.
		For example, dhcp option '80' is not migrated.
		<pre><dhcpOptions>
		  <other>
		  <code>80</code>
		  <value>2f766172</value>
		  </other>
		</dhcpOptions></pre>

Table 6-3. DHCP Features (continued)

NSX-V Configuration	Supported	Details
Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 6-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	

NSX-V Configuration	Supported	Details
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ VMs not on NSX-V hosts	No	NSX-T does not support referencing objects not connected to NSX-T portgroups or networks. Those objects will be lost from the source or destination when migration is completed. To avoid this issue use IP addresses in NSX-V to reference those objects before migration.
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	

NSX-V Configuration	Supported	Details
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.

NSX-V Configuration	Supported	Details
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 6-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of 'Groups' after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 6-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.

Table 6-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.

Table 6-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 6-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 6-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.

Table 6-8. Services and Service Groups (continued)

NSX-V Configuration	Supported	Details
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 6-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 6-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000

Table 6-10. Single-Site Limits (continued)

Feature	Limit
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500

Table 6-10. Single-Site Limits (continued)

Feature	Limit
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 6-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Changes Made During Host Migration in an End-to-End Migration

During the host migration step in an end-to-end migration, changes are made to migrate NSX-V hosts to NSX-T hosts.

- NSX-V software is uninstalled.
- NSX-T software is installed.
- For vSphere Distributed Switch versions 6.5.0 and 6.6.0:

  Hosts are configured with N-VDS to replace vSphere Distributed Switches:

  - Each N-VDS is created with a name that references the distributed switch name. For example, distributed switch `ComputeSwitchA` is created as N-VDS `nvds.ComputeSwitchA`.
  - If different clusters use different distributed switches to back logical switches, an N-VDS is created with a name that combines all the distributed switch names. For example, if `ComputeCluster1` and `ComputeCluster2` use distributed switch `ComputeSwitchA` to back logical switches and `ComputeCluster3` uses `ComputeSwitchB` to back logical switches, the N-VDS is created as `nvds.ComputeSwitchA.ComputeSwitchB`.
  - PNICs and vmks in the vSphere Distributed Switch are migrated to N-VDS.
  - NSX-V VTEPs are migrated to NSX-T TEPs.
- For vSphere Distributed Switches version 7.0:

Hosts configured for vSphere Distributed Switch version 7.0 continue using the same switch after migration.

- PNICs and vmks in the vSphere Distributed Switch remain connected on the same vSphere Distributed Portgroups.
- NSX-V VTEPs are migrated to NSX-T TEPs and connected to standalone ports on the same vSphere Distributed Switch.

Note: If NSX-V has multiple VTEPs and a single LAG configuration, after migration the LAG will have a single TEP and in failover mode.

Virtual Machine Deployment During an End-to-End Migration

After you start an end-to-end migration, do not change the NSX-V environment. If you want to deploy VMs during the migration, wait until some of the NSX-V hosts have migrated to NSX-T and deploy the VMs on NSX-T hosts. Connect the VMs to NSX-T segments and install VMware Tools on the VMs.

Deploying on NSX-T with VMware Tools installed ensures that the VMs are populated into security groups and receive the intended Distributed Firewall policies.

Caution VMs deployed without VMware Tools installed, or deployed on NSX-V do not receive the intended Distributed Firewall policies.

If you use vSphere templates to deploy VMs, update the templates to use NSX-T segments for the VM network configuration. Specifying NSX-T segments ensures that any VMs deployed using the templates are deployed on NSX-T hosts.

If you use automation tools to deploy VMs on vSphere, but do not use vSphere templates, you might need to change your automation tool configuration to ensure that the VMs are deployed on NSX-T.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Preparing for a User-Defined Topology Migration

Before migrating a user-defined topology, you must prepare the NSX-V and NSX-T environments.

If you are doing an end-to-end migration, perform the following tasks:

- Preparing the NSX-V Environment for a User-Defined Topology End-to-End Migration
- Preparing the NSX-T Data Center Environment for a User-Defined Topology Migration
- Change the MAC Address of NSX-T Virtual Distributed Router

If you are doing a lift-and-shift migration, perform the following tasks:

- Preparing the NSX-V Environment for a User-Defined Topology Lift-and-Shift Migration
- Preparing the NSX-T Environment for a User-Defined Topology Lift-and-Shift Migration
- Chapter 14 Preparing Layer-2 Bridging for Lift-and-Shift Migration (This task is not needed if the migration mode is Configuration and Edge Migration (available in NSX-T 3.2.2 as a tech preview feature.)

Configuration and Edge Migration Workflow

In NSX-T Data Center 3.2.2 and later, when you migrate a user-defined topology, you can choose the **Configuration and Edge Migration** mode. This mode migrates configurations, bridges the NSX-V logical switches to their corresponding NSX-T Data Center segments, and migrates Edge nodes for north-south traffic cutover.

Note This is a tech preview feature in NSX-T Data Center 3.2.2. This feature is fully supported in NSX-T Data Center 3.2.3 and later.

Note This migration mode requires dedicated hosts ready to be added to the NSX cluster. If new hosts are not available, hosts from NSX-V can be re-used. For more information, see the sections "(Optional) Re-using an NSX-V Host as an NSX Transport Node" and "(Optional) Adding an NSX Transport Node After the **Migrate Edges** Step Has Started" below.

End-to-end Workflow of Configuration and Edge Migration

Perform the following steps when you choose this mode:

1 Prepare the NSX-V environment. See Preparing the NSX-V Environment for a User-Defined Topology Lift-and-Shift Migration.

2 Prepare the NSX-T Data Center environment. See Preparing the NSX-T Data Center Environment for a User-Defined Topology Migration.

3 Perform the migration steps **Import Configuration**, **Translate Configuration Layer 2**, **Resolve Configuration Layer 2**, **Migrate Configuration Layer 2**, and **Check Realization Layer 2**. The transport zones and segments are created and ready to be used.

4 From the NSX Manager UI, configure NSX on ESXi hosts in the target site. Choose the correct transport zone for each host switch so that VMs migrated to the hosts will be connected to the correct segments. If a VTEP IP pool is used to configure NSX, make sure that there is no overlap in IP addresses between this IP pool and IP pools configured on NSX-V.

5 On NSX, create a topology that maps to the NSX-V topology, including tier-0 and tier-1 gateways.

6 Perform the migration steps **Define a Topology**, **Translate Configuration Layer 3 and Above**, **Resolve Configuration Layer 3 and Above**, **Migrate Configuration Layer 3 and Above**, **Check Realization Layer 3 and Above**, and **Migrate Edges**.

7 In the **Migrate Workloads** step, the following methods are available to migrate workload VMs:

- HCX - For more information, see the HCX documentation.
- vMotion - Follow the instruction in the section "Migrating Workload VMs" below.

8 Verify that everything works as expected.

9 Perform post-migration tasks to clean up the source site on NSX-V if needed.

Migrating Workload VMs from HCX

Note: Verify that your version of HCX supports this capability. If HCX is not available, see the section "Migrating Workload VMs" below for information on how to migrate workload VMs.

HCX can migrate VMs by group. To migrate VMs from HCX, first perform the migration steps from NSX Manager until the last step, **Migrate Workloads**. Then migrate the VMs from HCX. When the VMs are migrated, click **Finish** on the **Migrate Workloads** screen in NSX Manager.

Migrating Workload VMs

If you do not migrate workload VMS using HCX, you can follow the steps in Migrate Workloads in Lift-and-Shift Migration.

(Optional) Re-using an NSX-V Host as an NSX-T Data Center Transport Node

If you want to re-use an NSX-V host, follow the steps below to first prepare the host as an NSX-T Data Center transport node.

If the migration target VC is the same as the source VC:

1 Enter the host into maintenance mode in VC.

2 Move the host out of its cluster so that it becomes a standalone host. Uninstall NSX-V from the host.

3 Exit the host out of maintenance mode.

4 In NSX Manager UI, find the host in **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later). Add the VDSes with the correct transport zones and host-switch uplink profiles. Wait for the state of the node to display **Success**.

If the migration target VC is different from the source VC:

1 Enter the host into maintenance mode in VC.

2 Move the host out of its cluster so that it becomes a standalone host. Uninstall NSX-V from the host.

3 Move the host out of all the VDSes in the source VC.

4 Note down the host IP and then remove the host from the inventory in the source VC.

5 Add the host to a cluster in the target VC and then to the VDSes in the target VC.

6 Exit the host out of maintenance mode.

7 In NSX Manager UI, find the host in **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later). Add the VDSes with the correct transport zones and host-switch uplink profiles. Wait for the state of the node to display **Success**.

It is recommended to re-use an NSX-V host before the **Migrate Edges** step starts. If you want to re-use an NSX-V host after the **Migrate Edges** step has started, follow the steps in this section and then follow the steps in "(Optional) Adding an NSX-T Data Center Transport Node After the **Migrate Edges** Step Has Started".

(Optional) Adding an NSX-T Data Center Transport Node After the **Migrate Edges** Step Has Started

It is recommended to add all NSX-T Data Center Transport Nodes to the target NSX-T Data Center system before the **Migrate Edges** step starts. Follow the steps below as a workaround to add an NSX-T Data Center transport node after the **Migrate Edges** step has started.

1 In NSX Manager UI, go to **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later) and copy the node's UUID.

2 Use `ssh` to log in to the NSX Manager where you launched the migration.

3 Search for the host's IP in `/var/log/migration-coordinator/v2t/nsxv-config/hosts.json` and copy the `moId` (for example, `host-12`). If no host is found in the file, find the host's `moId` through `https://<VC-IP>/mob`.

4 Run the following commands:

```
cd /opt/vmware/migration-coordinator-tomcat/bin/v2t/config-collector

python3 vc_collector.py -s <VC-IP> -o <VC-HTTPS-port> -u <VC-user-name> -t /var/log/
migration-coordinator/v2t/nsxv-config/ -c /var/log/migration-coordinator/v2t/config.json
-hId <TN-uuid> -hmoId <host-moId>
```

Note: `<VC-HTTPS-port>` is 443 by default. For example,

```
python3 vc_collector.py -s 10.78.129.191 -o 443 -u administrator@vsphere.local -t /var/log/
migration-coordinator/v2t/nsxv-config/ -c /var/log/migration-coordinator/v2t/config.json
-hId 1bad5da6-8093-4496-a910-3dc224e6ac11 -hmoId host-12
```

5 Run the following commands:

```
cd /opt/vmware/migration-coordinator-tomcat/bin/v2t/config-migrator

python3 main.py -c /var/log/migration-coordinator/v2t/config.json -s accept-tn -t tn -i
<TN-uuid> runtime
```

For example,

```
python3 main.py -c /var/log/migration-coordinator/v2t/config.json -s accept-tn -t tn -i
1bad5da6-8093-4496-a910-3dc224e6ac11 runtime
```

If you are prompted for the authentication token of NSX-V manager, you can get the token by running the following command:

```
curl -i -k -u <admin-user> -X POST https://<nsx-v-ip>/api/2.0/services/auth/token?
expiresInMinutes=720
```

Note: `<admin-user>` is the administrator account of NSX-V manager and `<nsx-v-ip>` is the NSX-V manager's IP address.

This step will take a long time when the first transport-node is accepted because it needs to power on the NSX-V controllers. This step will fail if no NSX-V controller can be powered on and is in the `connected` state. In this situation, fix the NSX-V controllers to make sure that at least one NSX-V controller is in the `connected` state. Then retry this step.

Migrating a Cross-vCenter Environment to NSX Federation

Starting with NSX-T 3.2.1, you can migrate an NSX-V cross-vCenter environment to an NSX Federation environment in NSX-T.

A cross-vCenter environment in NSX-V is a multi-site deployment with one primary NSX-V and one or more secondary NSX-Vs. The primary NSX-V can have both universal objects and local objects, whereas the secondary NSX-V can have local objects only. The primary NSX-V can also act as a site because it can have both universal and local objects.

An NSX Federation deployment in NSX-T has one Global Manager (GM) and one or more Local Managers (LMs). The GM has global objects only and the LM's have local objects. The GM cannot act as a site because it cannot have local objects. Note that the term Local Manager refers to either a single NSX Manager or a cluster of NSX Managers for a site. In a production environment, each site should have a cluster of NSX Managers configured.

To migrate cross-vCenter to NSX Federation, you must initiate the migration from the Global Manager. From the manager UI, go to the **System > Migrate** screen and select **Migrate NSX for vSphere** and **User-Defined Topology**. You can then select either **Complete Migration** or **Configuration Migration**. If you choose **Complete Migration**, you will be doing an end-to-end migration. If you choose **Configuration Migration**, you will be doing a lift-and-shift migration.

Before the migration, you must configure NSX Federation in NSX-T. You must have one LM for each NSX-V site. For more information about configuring NSX Federation, see the section "Getting Started with NSX Federation" in the *NSX-T Data Center Installation Guide*.

In NSX-T 3.2.1, the NSX-V load balancer will not be migrated. Starting with NSX-T 3.2.2, The NSX-V load balancer will be migrated to an NSX-T load balancer.

Migrating VCF Workload Domains

You can migrate VCF (VMware Cloud Foundation) workload domains from nsxv to NSX Federation. The steps are:

- From VCF SDDC manager, deploy the necessary NSX-T infrastructure (Global Manager, Local Managers, Edge nodes) and create a layer-3 topology.
- Perform the pre-migration preparation tasks documented in this guide.
- From the Global Manager, migrate the NSX-V environment.

Migrating Universal Security Groups

In NSX-V, in an active-active environment, Universal Security Groups can contain the following objects only: security groups, IP sets, and MAC sets. You cannot configure dynamic membership or excluded objects. In an active-standby environment, Universal Security Groups can contain the following objects: security groups, IP sets, MAC sets, and universal security tags. You can also configure dynamic membership using the VM name.

Note that in a cross-vCenter environment, a dynamic membership condition, such as vm.name="abc", will only be applied to objects in a specific site and not to all the sites.

In NSX Federation, dynamic membership conditions in Global Groups will always be applied to all the sites. Therefore, when you migrate a Universal Security Group in NSX-V to a Global Group in NSX-T, the Global Group might have more members than the Universal Security Group.

Before the Universal Security Groups are migrated, the migration wizard will give you the option to either migrate or not migrate the Universal Security Groups.

Supported Limits

The following limits are supported when migrating a cross-vCenter environment.

Object	Limit
Sites	4
Hosts	512
Logical Switches (universal plus non-universal)	8500
Universal Distributed Firewall Rules	24000
Universal Firewall Sections	500
Universal Security Groups	4000
Universal IP Sets	4000
Universal IP Sets per Universal Security Group	10
Universal Security Tags	750

Preparing the NSX-V Environment for a User-Defined Topology End-to-End Migration

Check the Configurations of the NSX-V Environment

Check your NSX-V environment before starting an end-to-end migration.

System State

Check the following system states:

- If your environment is vSphere 7.0 or later, upgrade the VDS to 7.0 or later.
- Verify that the NSX-V components are in a green state on the NSX Dashboard.
- Verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering maintenance mode.
- Verify that no NSX-V upgrades are in progress.
- Verify the publish status of Distributed Firewall and Service Composer to make sure that there are no unpublished changes.
- You can have vSphere High Availability (HA) enabled if the NSX-V environment has VDS 7.0 or later.

  Note: HA is not supported for previous versions of VDS This is because if the NSX-V environment has VDS 6.5 or 6.7, and the vmkernel ports (vmk's) are attached to VDSes, during an in-place migration, the hosts and VMs may lose network connectivity for a period of time long enough to trigger HA. The HA mechanism will try to power off, migrate and restart VMs. This might fail because the NSX-V environment is being migrated to NSX-T. As a result, after the migration, VMs might remain in a powered-off state or have no network connectivity if powered on. To avoid this situation, disable HA or attach the management vmk to a VSS before starting the migration.

General Configuration

- Back up the NSX-V and vSphere environments. See "NSX Backup and Restore" in the *NSX Administration Guide*.
- The VXLAN port must be set to 4789. If your NSX-V environment uses a different port, you must change it before you can migrate. See "Change VXLAN Port" in the NSX-V *NSX Administration Guide*.

Controller Configuration

- NSX-V transport zones using multicast or hybrid replication mode are not supported for migration. An NSX Controller cluster is required if VXLAN is in use. VLAN-backed micro-segmentation topologies do not use VXLAN and so do not require an NSX Controller cluster.

Host Configuration

- On all host clusters in the NSX-V environment, check these settings and update if needed:
  - Set vSphere DRS accordingly.

    Disable vSphere DRS if one of the following apply:

    - **In-Place** migration mode will be used. In this mode hosts are not put in maintenance mode during migration and VMs will experience a network outage and network storage outage during the migration. This mode is only available if the environment is vSphere 6.x (VDS will be migrated to N-VDS).
    - **Manual Maintenance** migration mode will be used. If you decide to use vMotion for migrating VMs, you can disable vSphere DRS, or set the vSphere DRS automation level to Manual, Partially Automated, or Fully Automated.
    - **Automated Maintenance** migration mode will be used and the VDS version is 6.5 or 6.7.

    Set vSphere DRS mode to Fully Automated if:

    - **Automated Maintenance** migration mode will be used and the VDS version is 7.0.

    Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.
- To migrate Network Introspection service rules, use the **Maintenance** host migration mode. **In-Place** migration mode is not supported.
- If you have hosts that have NSX-V installed, but are not added to a vSphere Distributed Switch, you must add them to distributed switches if you want to migrate them to NSX-T. See Configure Hosts Not Attached to vSphere Distributed Switches for more information.
- On each cluster that has NSX-V installed, check whether Distributed Firewall is enabled. You can view the enabled status at **Installation & Upgrade > Host Preparation**.

  If Distributed Firewall is enabled on any NSX-V clusters before migration, Distributed Firewall is enabled on all clusters when they migrate to NSX-T. Determine the impact of enabling Distributed Firewall on all clusters and change the Distributed Firewall configuration if needed.

Edge Services Gateway Configuration

- You might need to make changes to your NSX-V route redistribution configuration before migration starts.
  - Redistribution filters are not migrated. For BGP, filters can be moved to the BGP neighbor level.

    - After migration, dynamically learned routes between Distributed Logical Router and Edge Services Gateway are converted to static routes and all static routes are redistributed in BGP or OSPF. If you need to filter any of these routes, you can configure them at the BGP neighbor level or manually configure the redistribution rules on NSX-T after the configuration migration is completed and before cutover. Note that if you roll back, the manual configuration of redistribution rules will also be removed.
    - The default MTU setting is 1500 on NSX-T. If you have non-default MTU setting requirements, you can change the setting. See Change the Global MTU Setting.
- NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the **Migrate Configuration** step fails.
- If you have an Edge Services gateway performing one-armed load balancer function, you must change the following configurations if present before you import the configuration:
    - If the Edge Services Gateway has an interface configured for management, you must delete it before migration. You can have only one connected interface on an Edge Services Gateway providing one-arm load balancer function. If it has more than one interface, the **Migrate Configuration** step fails.
    - If the Edge Services Gateway firewall is disabled, and the default rule is set to deny, you must enable the firewall and change the default rule to accept. After migration the firewall is enabled on the tier-1 gateway, and the default rule accept takes effect. Changing the default rule to accept before migration prevents incoming traffic to the load balancer from being blocked.
- Verify that Edge Services Gateways are all connected correctly to the topology being migrated. If Edge Services Gateways are part of the NSX-V environment, but are not correctly attached to the rest of the environment, they are not migrated.

  For example, if an Edge Services Gateway is configured as a one-armed load balancer, but has one of the following configurations, it is not migrated:
    - The Edge Services Gateway does not have an uplink interface connected to a logical switch.
    - The Edge Services Gateway has an uplink interface connected to a logical switch, but the uplink IP address not does match the subnet associated with the distributed logical router that connects to the logical switch.

Security Configuration

- If you plan to use vMotion to move VMs during the migration, disable all SpoofGuard policies in NSX-V to prevent packet loss.
    - Automated Maintenance mode uses DRS and vMotion to move VMs during migration.

- In Manual Maintenance mode, you can optionally use vMotion to move VMs during migration.
- In-Place migration mode does not use vMotion.

Security Group Configuration

If existing Security Policies contain Guest Introspection service rules that are applied to Security Groups with static VM members or dynamic members other than VMs, do these steps:

1 Create new Security Groups with VMs only in the dynamic membership criteria. Make sure that the dynamic membership criteria produces the same effective VM members as your original Security Groups.

2 Before starting the migration, update the existing Security Policies to apply the new Security Groups to the Guest Introspection service rules.

   If you prefer not to update your existing Security Policies before the migration, you can still keep the new Security Groups ready with the correct dynamic membership criteria in your NSX-V environment. In the **Resolve Configuration** step of the migration process, you will be prompted to provide alternative Security Groups.

Service Composer Synchronization

Ensure that the Service Composer is in sync with Distributed Firewall before you start the migration. A manual synchronization ensures that if you make any last-minute changes in the policy configuration before starting the migration, these changes are applied to the Security Policies that are created using Security Composer too. For example, you edit the name of the Security Group that is used in a firewall rule before starting the migration.

To verify whether the Service Composer status is in sync, do these steps:

1 In the vSphere Client, navigate to **Networking and Security > Security > Service Composer**.

2 Click the **Security Policies** tab.

3 Verify that the Sync Status is **In Sync**. If it is not in sync, click **Synchronize**.

As a best practice, always click the **Synchronize** button before starting the migration even when the sync status is green. Do this manual synchronization regardless of whether you performed any last-minute changes in the policy configuration.

During migration, if the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of the Security Policies created using the Service Composer by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get the Service Composer in sync with Distributed Firewall, and restart the migration.

Configure Hosts Not Attached to vSphere Distributed Switches

An NSX-V environment can contain hosts that have NSX-V installed, but are not added to a vSphere Distributed Switch. You must add the hosts to a vSphere Distributed Switch before you can migrate them.

You can use a distributed switch you already have in your environment, or create a new distributed switch for this purpose. Right click the distributed switch and select **Add and Manage Hosts** to add the hosts to the distributed switch. You do not need to assign physical uplinks or VMkernel network adapters to the distributed switch.

See "Add Hosts to a vSphere Distributed Switch" in the *vSphere Networking Guide* for more information.

If you import the configuration before you make this change, you must restart the migration to import the updated configuration. See Make Changes to the NSX-V Environment.

After the migration has finished, the hosts are no longer required to be attached to the distributed switch.

- If you added the hosts to an existing distributed switch, you can remove them from the distributed switch.
- If you added the hosts to a new distributed switch that you are not using for another purpose, you can delete the distributed switch.

Tag Management VMs in a Collapsed Cluster Environment

You can migrate an environment that uses a collapsed cluster.

In a collapsed cluster design, all management VMs, workload VMs, and optionally edges run on the same vSphere cluster that is prepared for NSX-V. The management VMs of the NSX-T must be initially attached to dvPortgroups. After migration, the management VMs of NSX-T will be attached to NSX-T VLAN segments.

The management VMs in the NSX-T include appliances such as NSX Manager, vCenter Server, VMware Identity Manager, and so on. The NSX-T VLAN segment ports to which these management VMs connect are blocked in two cases: when these management VMs are rebooted after they were migrated in-place by the Migration Coordinator, or when they are moved from NSX-V hosts to NSX-T hosts by vMotion in maintenance migration mode. Therefore, the management VMs might lose connectivity in such cases.

To prevent this problem, create a "management_vms" tag category, and add tags in this category. Assign a tag from this category to all the management VMs in the NSX-T environment. These VMs will be attached to unblocked VLAN segment ports.

Procedure

1. Log in to the vSphere Client.
2. Click **Menu > Tags & Custom Attributes**.

3 Click **Categories**, and then click **New** to add a category.

   Create a category with name **`management_vms`**.

4 Click the **Tags** tab and add a tag in the management_vms category.

5 Navigate to **Menu > Hosts and Clusters**.

6 Expand the collapsed cluster from the left Navigator view, right-click the name of the NSX Manager VM, and select **Tags & Custom Attributes > Assign Tag**.

7 Assign a tag from the management_vms category to the NSX Manager VM.

8 Repeat steps 6 and 7 for all the management VMs in the cluster.

   For a detailed information about tag categories and tags, see the *vCenter Server and Host Management* documentation.

9 Log in to NSX Manager.

10 Navigate to **Inventory > Groups** and create a group.

11 Click **Set** to set members.

12 Set **Group Type** to **IP Addresses Only** and include the IP addresses of the management VMs that will be migrating from the dvPortgroups to the NSX-T segments.

13 Navigate to **Security > Distributed Firewall > Actions > Exclusion List** and add this group.

Delete Partner Service Deployments

If your NSX-V environment uses a partner service for Guest Introspection, or both Guest Introspection and Network Introspection, delete the partner service deployment before migration.

You must also delete the Guest Introspection instance (GI-SVM) so that the Guest Introspection module is uninstalled from the clusters.

If your NSX-V environment uses a partner service only for Network Introspection, you have the flexibility to decide whether to delete the partner service deployment before or after the migration. When a partner service deployment is deleted, the partner service virtual machines (SVMs) are removed from the NSX-V-prepared host cluster, and security protection is lost.

Important Consult the VMware partner before migrating their service that is running on the NSX-V workloads. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their services to NSX-T.

Procedure

1 Log in to the vSphere Client and navigate to **Networking and Security > Installation and Upgrade > Service Deployment**.

2 Select the deployed service and click **Delete**.

Preparing the NSX-V Environment for a User-Defined Topology Lift-and-Shift Migration

Configure Export Version of Distributed Firewall Filter

The export version of a Distributed Firewall (DFW) filter is a property of a vNIC. Before you start some migrations, the export version of DFW filters must be set to 1000 for the vNICs of all the VMs that will be migrated.

You must make this configuration change in the following situations:

- You are doing a lift-and-shift migration.
- You are doing an in-place migration and you need to manually migrate VMs from some NSX-V hosts to NSX-T hosts. Follow the procedure below to change the export version for only those NSX-V hosts before migrating the VMs.

Procedure

1. Based on the VMs that will be migrated, determine the hosts that the VMs are running on.
2. Perform either step 3 or step 4 below.
3. For each host, perform the following steps to update, if necessary, the export version of DFW filters for all VM vNICs.

   Note: In https://github.com/dixonly/samples, the script `updateDfwFilters.py` will print out and optionally update the DFW filter's export version for the vNICs of all the VMs in a specific cluster or all clusters. Using the script can save some time if you have a large number of VMs to migrate.

   a. Log into the command-line interface.

   b. Get the DFW filter names for all the VM vNICs. For example,

   ```
   [root@esxi:~] vsipioctl getfilters | grep "Filter Name" | grep "sfw.2"
   Filter Name: nic-2112467-eth0-vmware-sfw.2
   Filter Name: nic-2112467-eth1-vmware-sfw.2
   Filter Name: nic-2112467-eth2-vmware-sfw.2
   ```

   c. For each filter, get the export version. For example,

   ```
   [root@esxi:~] vsipioctl getexportversion -f nic-2112467-eth0-vmware-sfw.2
   Current export version: 500
   ```

d If the version is not 1000, set it to 1000. For example,

```
[root@esxi:~] vsipioctl setexportversion -f nic-2112467-eth0-vmware-sfw.2 -e 1000
```

e Verify that the export version is updated. For example,

```
[root@esxi:~] vsipioctl getexportversion  -f nic-2112467-eth0-vmware-sfw.2
Current export version: 1000
```

4 Based on the hosts that you noted in step 1, determine the clusters that contain the hosts. For each cluster, do the following:

From the vSphere Client, navigate to **Networking and Security > Installation and Upgrade > Host Preparation**. Select the cluster and click **Actions > Disable Firewall**. After the firewall is disabled, click **Actions > Enable Firewall**.

Preparing the NSX-T Data Center Environment for a User-Defined Topology Migration

Perform the tasks in the following sections for either an end-to-end migration or a lift-and-shift migration.

Deploy an NSX Manager Appliance

You must deploy a new NSX Manager appliance to run the migration coordinator. Do not deploy an NSX Global Manager.

In other words, you cannot merge your NSX-V environment into an existing NSX-T environment, which has NSX-T already installed on the vSphere host clusters.

For details on deploying a licensed version of the NSX Manager appliance, see *Install NSX Manager and Available Appliances* in the *NSX-T Data Center Installation Guide*.

Install one appliance to perform the migration. Deploy additional appliances to form a cluster after the migration is finished. See Finish Deploying the NSX Manager Cluster.

If you install the NSX Manager appliance on an ESXi host that is a part of the NSX-V environment that is migrating, do not attach the appliance interfaces to an NSX-V logical switch. To prevent the management VMs in NSX-T from losing connectivity after the VMs are rebooted post migration, tag the management VMs. For more information, see Tag Management VMs in a Collapsed Cluster Environment.

Add a Compute Manager

Before you can start the migration process, you must add the vCenter Server that is associated with NSX-V as a compute manager in NSX-T.

Prerequisites

Log into the NSX-V NSX Manager web interface to retrieve the settings used for vCenter Server registration. You must use the same settings. For example, if an IP address is specified, use the IP address and not the FQDN. Note that the FQDN of the vCenter Server is case-sensitive. If you enter the FQDN in the procedure below, be sure that it matches the vCenter Server's FQDN exactly.

Procedure

1. From a browser, log in with admin privileges to NSX Manager at `https://<nsx-manager-ip-address>`.
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 vCenter Server. You can optionally describe any special details such as, the number of clusters in the vCenter Server.
FQDN or IP Address	Type the FQDN or IP address of the vCenter Server.
Type	The default compute manager type is set to vCenter Server.
HTTPS Port of Reverse Proxy	The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances. Set the reverse proxy port to register the compute manager in NSX-T Data Center.
Username and Password	Type the vCenter Server login credentials.
SHA-256 Thumbprint	Type the vCenter Server SHA-256 thumbprint algorithm value.
Create Service Account	Enable this field for features such as vSphere Lifecycle Manager that need to authenticate with NSX-T Data Center APIs. Log in with the administrator@vsphere.local credential to register a compute manager. After registration, the compute manager creates a service account. **Note** Service account creation is not supported on a global NSX Manager. If service account creation fails, the compute manager's registration status is set to `Registered with errors`. The compute manager is successfully registered. However, vSphere Lifecycle Manager cannot be enabled on NSX-T Data Center clusters. If a vCenter Server admin deletes the service account after it was successfully created, vSphere Lifecycle Manager tries to authenticate the NSX-T Data Center APIs and the compute manager's registration status is set to `Registered with errors`.

Option	Description
Enable Trust	Enable this field to establish trust between NSX-T Data Center and compute manager, so that services running in vCenter Server can establish trusted communication with NSX-T Data Center. For example, for vSphere Lifecycle Manager to be enabled on NSX-T Data Center clusters, you must enable this field. Supported only on vCenter Server 7.0 and later versions.
Access Level	Enable one of the options based on your requirement: ■ Full Access to NSX: Is selected by default. This access level gives the compute manager complete access to NSX-T Data Center. Full access ensures vSphere for Kubernetes and vSphere Lifecycle Manager can communicate with NSX-T Data Center. The vCenter Server 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-T Data Center. The vCenter Server 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-T Data Center to discover and register the vCenter Server 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 vCenter Server credentials and click **Resolve**.

If an existing registration exists, it will be replaced.

Results

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

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

After the vCenter Server is successfully registered, do not power off and delete the NSX Manager VM without deleting the compute manager first. Otherwise, when you deploy a new NSX Manager, you will not be able to register the same vCenter Server again. You will get the error that the vCenter Server is 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, NSX Intelligence VM, 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, any NSX Intelligence VM, 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 sucessfully remove the VC, ensure that you re-deploy an NSX Manager node from the VC.

This restriction applies to a fresh installation of NSX-T Data Center as well as an upgrade.

Change the Global MTU Setting

When an Edge Services Gateway (ESG) is migrated, the MTU setting of the interfaces is not migrated. A default value of 1500 is used. You can change the default value using the API.

You can also modify the MTU setting for the interfaces after the migration.

Procedure

1 Make the following API call to retrieve the current configuration.

```
GET /api/v1/global-configs/RoutingGlobalConfig
```

2 Change the value for `logical_uplink_mtu` and make the following call.

```
PUT /api/v1/global-configs/RoutingGlobalConfig
```

Create an IP Pool for Edge Tunnel End Points

If your NSX-V environment uses Edge Services Gateways, you must create an IP pool in the NSX-T environment for the Edge Tunnel End Points (TEP) before you start the migration.

Prerequisites

- Identify existing IP pools or DHCP ranges for NSX-V VTEPs.
- Determine which IP addresses to use to create an IP pool for Edge TEPs.

  The IP range and VLAN must not already be in use in the NSX-V environment.
- Verify that the NSX-T TEP IP addresses have network connectivity to the NSX-V VTEP IP addresses.

Procedure

1. From a browser, log in with admin privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Select **Networking > IP Management > IP Address Pools**.
3. Click **Add IP Address Pool**.
4. Enter a name for the new IP pool.
5. (Optional) Enter a description.
6. In the **Subnets** column, click **Set** to add subnets.
7. Specify the IP ranges.
   a. Select **Add Subnets > IP Ranges**.
   b. Enter IPv4 or IPv6 ranges.
   c. Enter the subnet address in a CIDR format.
   d. Enter the Gateway IP address for this subnet.
   e. (Optional) Enter DNS servers.
   f. (Optional) Enter DNS suffix.
   g. Click **Add**, and then click **Apply**.
8. Click **Save**.

Plan the Mapping of the NSX-V Topology to the NSX-T Topology

Review the NSX-V topology and decide how to map it to the NSX-T topology. During the migration, the "Define a Topology" step will prompt you for the mapping.

If you are running NSX-T 3.2.0 or 3.2.1, you can migrate an NSX-V load balancer to NSX-T Advanced Load Balancer (ALB). Starting with NSX-T 3.2.2, you can only migrate an NSX-V load balancer to an NSX-T load balancer.

Specifically, determine how many NSX-T Edge clusters you need and how Edge Service Gateways (ESGs) and Distributed Logical Routers (DLRs) should map to gateways in NSX-T. The northbound ESGs without any L4-L7 services should be skipped. These are usually the ESGs peering with northbound routers and are in ECMP path. If you are using VPN on a northbound ESG, migrating to active-standby tier-0 is recommended. In other cases, migrating the ESGs/DLRs to tier-1 is recommended. An ESG and a DLR can be merged in one mapping entry in the mapping file.

You cannot map multiple tier-1 gateways without an edge cluster or a DR-only tier-1 gateway under a parent tier-0 gateway to DLRs. You also cannot map the parent tier-0 gateway to a DLR if you are mapping to a DR-only tier-1 gateway. If your topology requires mapping multiple DLRs, you must use an active-standby tier-1 gateway with an edge cluster assigned.

If you have a DHCP server configured on an ESG and DHCP relay server configured on a DLR, you must map the ESG and DLR to the same NSX-T gateway.

An example of a mapping file that maps ESGs to a tier-0 gateway:

```
[
  {
    "name":"nsxv-to-nsxt-mapping",
    "v_edges_to_policy_gateways_mappings":[
      {
        "v_edges":[
          "edge-1",
          "edge-2"
         ],
         "policy_gateway_name": "tier0-gateway"
         "policy_gateway_path": "/infra/tier-0s/tier0-gateway"
      }
    ]
  }
]
```

If you are running NSX-T 3.2.0 or 3.2.1, and you are doing a configuration migration, the above-mentioned mapping is not used for Advanced Load Balancer (ALB). Instead there is another optional mapping that you may provide for ALB. To specify that mapping, you must upload a JSON file. An example of a mapping file that maps ESGs to Service Engine groups:

```
{
  "alb": {
    "service_engine_group_per_esg": false,
    "esgs": [
      {
        "name": "edge-4",
        "interfaces": [
          {
            "name": "mgmt",
            "tier1_id": "London_Tier1Gateway1"
          },
          {
            "name": "vnic1",
```

```
            "placement_network_subnet": "172.16.1.10/16",
            "service_engine_group": "Test-SE-group"
          }
        ]
      }
    ]
  }
}
```

For more information about creating a mapping file for the load balancer, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

Starting with NSX-T 3.2.1, you can migratie a cross-vCenter environment to NSX Federation. Here is a sample mapping file for such a migration:

```
[
  {
    "name": "site-GM",
    "nsxv_id": "20.20.0.131",
    "nsxt_site_id": "",
    "v_edges_to_policy_gateways_mappings": [
      {
        "v_edges": [
          "edge-3",                                          <== Dhcp server
          "edge-1b4b70c4-27da-4ee9-91cb-bd6389dadca2"        <== Dhcp Relay
        ],
        "policy_gateway_name": "tier0_1_global",
        "policy_gateway_path": "/global-infra/tier-0s/tier0_1_global"
      },
      {
        "v_edges": [
          "edge-a529c168-56c0-4e1a-98e4-e1f0312d82a4"
        ],
        "policy_gateway_name": "tier1_0_global",
        "policy_gateway_path": "/global-infra/tier-1s/tier1_0_global".     <== mapping
corresponding to UDLR
      },
      {
        "v_edges": [
          "edge-5"
        ],
        "policy_gateway_name": "tier0_2_global",
        "policy_gateway_path": "/global-infra/tier-0s/tier0_2_global"
      }
    ]
  },
  {
    "name": "london",
    "nsxv_id": "20.20.0.131",
    "nsxt_site_id": "23790eb0-201a-48c3-8e34-8a03be03b61a",
    "v_edges_to_policy_gateways_mappings": [
      {
        "v_edges": [
          "edge-4"
        ],
```

```
        "policy_gateway_name": "site0_tier1_0",
        "policy_gateway_path": "/infra/tier-1s/site0_tier1_0"
      }
    ]
  },
  {
    "name": "paris",
    "nsxv_id": "20.20.0.132",
    "nsxt_site_id": "27db83d3-aa10-4a40-b15c-bf6e234b9e74",
    "v_edges_to_policy_gateways_mappings": [
      {
        "v_edges": [
          "edge-1",
          "edge-5"
        ],
        "policy_gateway_name": "site1_tier0_1_local",
        "policy_gateway_path": "/infra/tier-0s/site1_tier0_1_local"
      }
    ]
  }
]
```

If you are migrating a cross-vCenter environment to NSX Federation, note the following:

- Tier-0 gateways created on Local Manager must have an Edge cluster assigned.
- Tier-1 gateways created on Local Manager must either have an Edge cluster assigned or be connected to a tier-0 gateway that has an Edge cluster assigned.
- Tier-0 and tier-1 gateways created on Global Manager must span all the sites.
- In the mapping file, when specifying attributes for the Global Manager, `nsxt_site_id` must be an empty string.
- Universal DLR (UDLR) should not be merged with any other ESG.
- Northbound ESGs with L4-L7 services peering with UDLR must be mapped to a local site's active-standby tier-0 gateway.

Deploy NSX Edge Nodes

You can deploy NSX Edge nodes using an OVA or OVF file or from the NSX Manager user interface.

Do not deploy on bare metal.

Snapshots of NSX appliances (including Edge node VMs) are not supported and must be disabled. For information on how to disable snapshots, see the topic Disable Snapshots on an NSX Appliance in the *NSX-T Data Center Installation Guide*.

NSX Edge nodes must be connected to trunk portgroups. To learn more about NSX Edge networking, see "NSX Edge Networking Setup" in the *NSX-T Data Center Installation Guide*.

Caution If you deploy the NSX Edge node VM on an NSX-V-prepared host, connectivity of the Edge node might be affected by a Distributed Firewall deny rule in NSX-V. To avoid this issue, add the Edge node VM to the Distributed Firewall's exclusion list.

Prerequisites

- You must have sufficient ESXi hosts with appropriate resources available to accommodate the NSX Edge appliances.

Procedure

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

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

2 In the vSphere Client, select the host on which to install NSX Edge 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 vCenter Server and vSphere inventory.

6 Select a compute resource for the NSX Edge node appliance, and click **Next**.

7 Review and verify the OVF template details, and click **Next**.

8 Select a deployment configuration and click **Next**.

   See the **Import Configuration** step for details on the size of Edge nodes you must deploy.

9 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.

10 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.

   Post-migration, the NSX Edge node is connected to one of these three trunk networks using only a single fastpath interface. The network settings can be adjusted or verified after the NSX Edge node is deployed.

11 Configure IP Allocation settings.

   a For IP allocation, specify **`Static - Manual`**.

   b For IP protocol, select **`IPv4`**.

12 Click **Next**.

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

13 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.

14 Enter the hostname of the NSX Edge.

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

   Skip any VMC network settings.

16 Enter the DNS Server list, the Domain Search list, and the NTP Server IP or FQDN list.

17 (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.

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

   The installation might take 7-8 minutes.

19 Start the NSX Edge node VM manually.

20 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.

21 After the NSX Edge node starts, log in to the CLI 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.

22 Run the `get interface eth0` (without VLAN) or `get interface eth0.<vlan_ID>` (with a VLAN) command to verify that the IP address was applied as expected.

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

Interface: eth0.100
  Address: 192.168.110.37/24
  MAC address: 00:50:56:86:62:4d
```

```
MTU: 1500
Default gateway: 192.168.110.1
Broadcast address: 192.168.110.255
...
```

23 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 IP or FQDN list.

24 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 use DHCP 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 Type the **`set interface`** ***`interface`*** **`dhcp plane mgmt`** command.

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

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

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

Join NSX Edge Node VM with the Management Plane

You must join the NSX Edge node VM you created to the management plane.

Do not join the NSX Edge node VM to the management plane using any other method. Do not create transport nodes from the NSX Edge node VM.

Procedure

1 Open an SSH session or console session to the NSX Manager appliance.

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 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.

5 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)
```

6 In the NSX Manager UI, you can navigate to **System > Fabric > Nodes > Edge Transport Nodes** and see the NSX Edge node. The **Configuration State** column will display **Pending**. If you click the name of the Edge node, you will be prompted to configure the node. Do not configure the node. The configuration will occur during the migration.

Configure NSX-T for a User-Defined Topology Migration

Before migrating a user-defined topology, you must configure NSX-T.

Regarding the configuration of Edge clusters, VLAN transport zone, and segments used for tier-0 gateway uplinks, you have two options:

- Let the migration coordinator configure an Edge cluster, VLAN transport zone, and other layer-2 entities needed for northbound connectivity. Note that only one Edge cluster will be created.

- Manually configure Edge clusters, VLAN transport zone, and other layer-2 entities needed for northbound connectivity. You must use this option if you want more than one Edge cluster. For more information about creating an Edge cluster, see the section "Create an NSX Edge Cluster" in the *NSX-T Data Center Installation Guide*.

After the Edge cluster is ready, do the following:

- Create tier-0 and tier-1 gateways and configure dynamic or static routing on the tier-0 gateways towards the northbound routers. For dynamic routing, configure BGP or OSPF, route redistribution as well as any filtering or route-maps based on your requirements. For more information about configuring static or dynamic routing (BGP or OSPF), see the section "Tier-0 Gateways" in the *NSX-T Data Center Administration Guide*.
- After configuring the dynamic routing on tier-0 gateways, check that the dynamic routing has converged, that is, BGP sessions are established or OSPF neighborships are FULL as applicable.

Other than the configurations mentioned above, no other configurations should be performed.

Preparing the NSX-T Data Center Environment for a User-Defined Topology End-to-End Migration

Register Third-Party Guest Introspection Service with NSX-T

If Security Policies in your NSX-V environment use third-party Guest Introspection service provided by a partner, register the partner service with NSX-T before you start the migration.

You might need to upgrade the Partner Console to register the service with the version of NSX-T that is used for this migration. For more information, see the partner documentation.

Complete the following procedure to register a partner service for endpoint protection with the NSX Manager in your NSX-T environment.

Procedure

1. Log in to the Partner Console with **Admin** privileges.
2. Update the NSX endpoint in the Partner Console. Specify the following details:
   - IP address of NSX-T NSX Manager
   - Port (default is 443)
   - User name and password of the NSX-T NSX Manager

   Make sure to test the connection before proceeding to the next step. If you need help with using the Partner Console, see the partner documentation.

   The partner service and the vendor templates that are associated with this partner service are now created in NSX-T.

3 Verify that the partner service is registered with NSX-T.

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

   b Navigate to **System > Service Deployments > Deployment**.

   c Click the **Partner Service** drop-down menu, and check that the partner service is listed.

Register Third-Party Network Introspection Services with NSX-T

If Security Policies in your NSX-V environment use third-party Network Introspection services provided by partners, partner services must be registered with NSX-T before you start the migration.

You might need to upgrade the Partner Console to ensure that the partner service is registered with the version of NSX-T that is used for this migration. For more information, see the partner documentation.

The following types of east-west network introspection services are supported for migration:

- Intrusion detection services (IDS)
- Intrusion protection services (IPS)
- Network monitoring services
- Next-generation firewall services

A partner registers the service, vendor template, and the Partner Management Console/Partner Service Manager. Then, either you or the partner can create the service profile. It can vary from one partner to another. See the partner documentation.

In the following procedure, step 2 is required when your NSX-V environment uses only Network Introspection service.

If your environment uses a combination of both Guest Introspection and Network Introspection services from a single partner (partner A), partner does step 1. Step 2 is not required.

If your environment uses Guest Introspection service from one partner (partner A) and Network Introspection service from another partner (partner B), then:

- Use the Partner Console of partner A to register the Guest Introspection service. See the partner documentation for help on registering the service.
- Partner B registers the Network Introspection service (step 1 of the procedure). Either you or the partner can create the service profile, as explained in step 2.

Procedure

1 Partner registers the partner service, vendor template, and the partner Service Manager using NSX-T APIs.

2 Create a service profile to specify attributes of a vendor template for a given partner service.

For a network introspection service, multiple service profiles can be associated with a single vendor template.

You can create a service profile either by using the NSX-T API or the NSX Manager UI. For detailed steps on creating the service profile by using the NSX Manager UI, see the *NSX-T Data Center Administration Guide*.

When you use the NSX Manager UI to create a service profile, the service reference is internally created, if it is not already present.

If you decide to use the NSX-T APIs to create a service profile, do the following steps:

a Create a service reference.

```
PATCH https://{policy-mgr-ip}/policy/api/v1/infra/service-references/
{service-reference-id}
```

b Use the *service-reference-id* from the previous step to create the service profile.

```
PATCH https://{policy-mgr-ip}/policy/api/v1/infra/service-references/
{service-reference-id}/service-profiles/{service-profile-id}
```

For a detailed information about these APIs, see the *NSX-T Data Center API Guide*.

Preparing the NSX-T Environment for a User-Defined Topology Lift-and-Shift Migration

To prepare for a lift-and-shift migration, you may need to create a layer-2 bridge between NSX-V and NSX-T.

For information on creating a layer-2 bridge between NSX-V and NSX-T, see Chapter 14 Preparing Layer-2 Bridging for Lift-and-Shift Migration.

If you are migrating the NSX-V load balancer to NSX-T Advanced Load Balancer (ALB), see Migrating NSX-V Load Balancer to Advanced Load Balancer.

Migrating NSX-V Load Balancer to Advanced Load Balancer

You can migrate the NSX-V load balancer to NSX-T Advanced Load Balancer (ALB).

To migrate to ALB, from NSX Manager, go to the **System > Migrate** screen and select **Migrate NSX for vSphere**, **User-Defined Topology** and **Configuration Migration**. This is the only way to migrate the NSX-V load balancer to ALB. All other migration modes will migrate the NSX-V load balancer to NSX-T load balancer. Only a single-site deployment is supported.

Pre-migration Tasks

Before the migration, perform the following tasks to set up ALB in NSX-T:

- Deploy Avi Controllers from the NSX Manager UI (navigate to **System > Appliances**). For more information, see the topic Install NSX Advanced Load Balancer Appliance Cluster in the *NSX-T Data Center Installation Guide*.
- Avi Controller configurations:
  - Set up the license mode and upload the required Avi licenses on the Avi Controller using the cross-launch UI.
  - Set up cloud configurations for the corresponding transport zone in NSX-T.
  - Make sure that the management network for the Service Engine (SE) has connectivity to the controller IP.
  - Make sure that ports 22, 443, 8443, and 123 are accessible from SE to the controller.

Supported Topologies

You can migrate load balancers either deployed on VLAN or on overlay. The following deployment types are supported.

Overlay:

- Inline load balancer without transparent mode
- One-arm load balancer with a single logical interface (LIF)
- One-arm load balancer across two LIFs
- Inline load balancer with a one-arm load balancer

VLAN:

- One-arm load balancer
- Inline load balancer without transparent mode
- Inline nested load balancer

The following diagrams show the topologies before the migration (on the left) and after (on the right). In the diagrams, C stands for client VM, S stands for server VM, and E stands for Edge Services Gateway.

Overlay topologies:

Topology 1: On the left: NSX-V inline LB without transparent mode. On the right: NSX-T single-arm ALB.

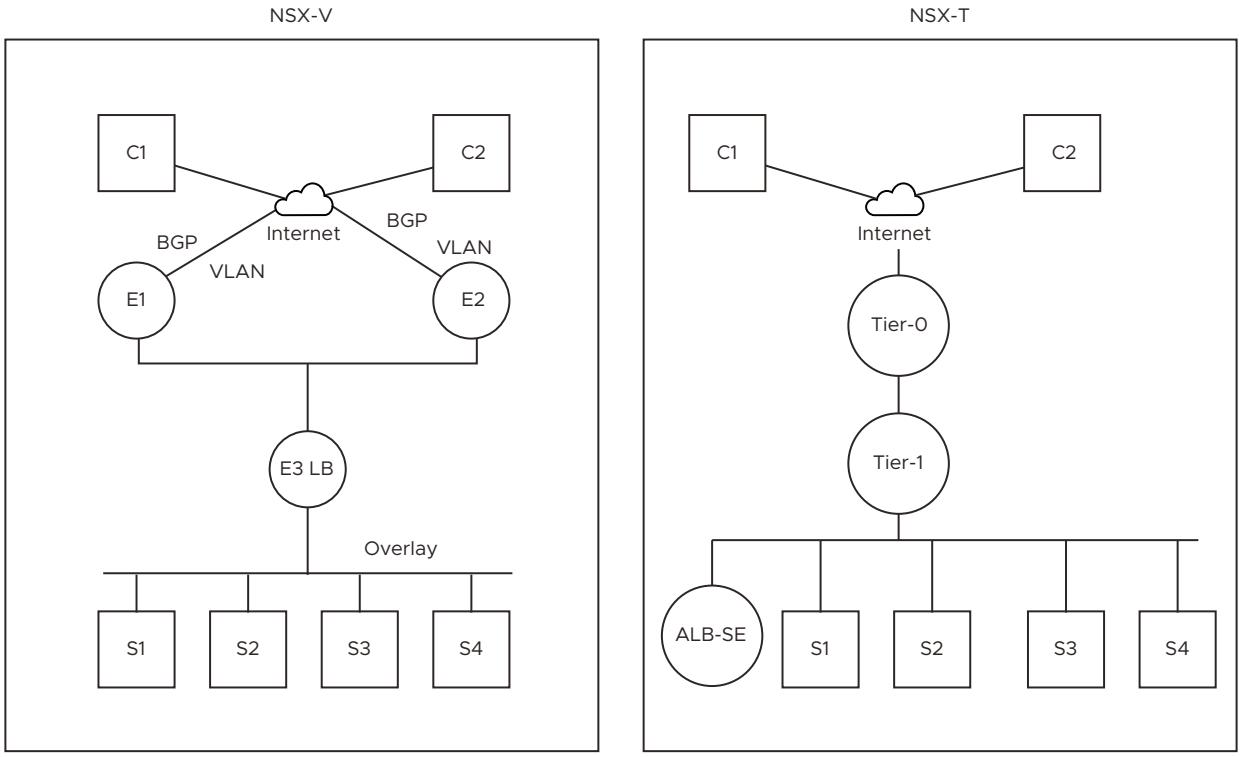

Topology 2: On the left: NSX-V single-arm LB. On the right: NSX-T single-arm ALB.

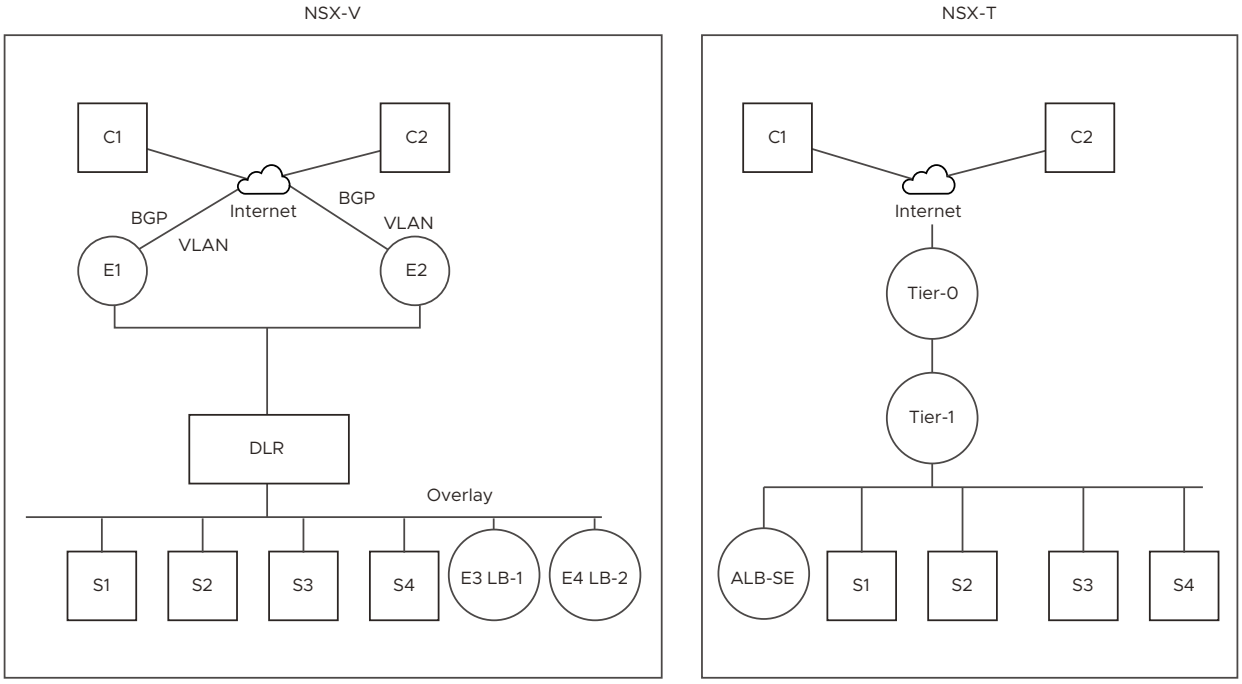

Topology 3: On the left: NSX-V. Two single-arm LBs across two LIFs (1, 2) to one LIF (1). On the right: NSX-T single-arm ALB across different LIFs (1, 2).

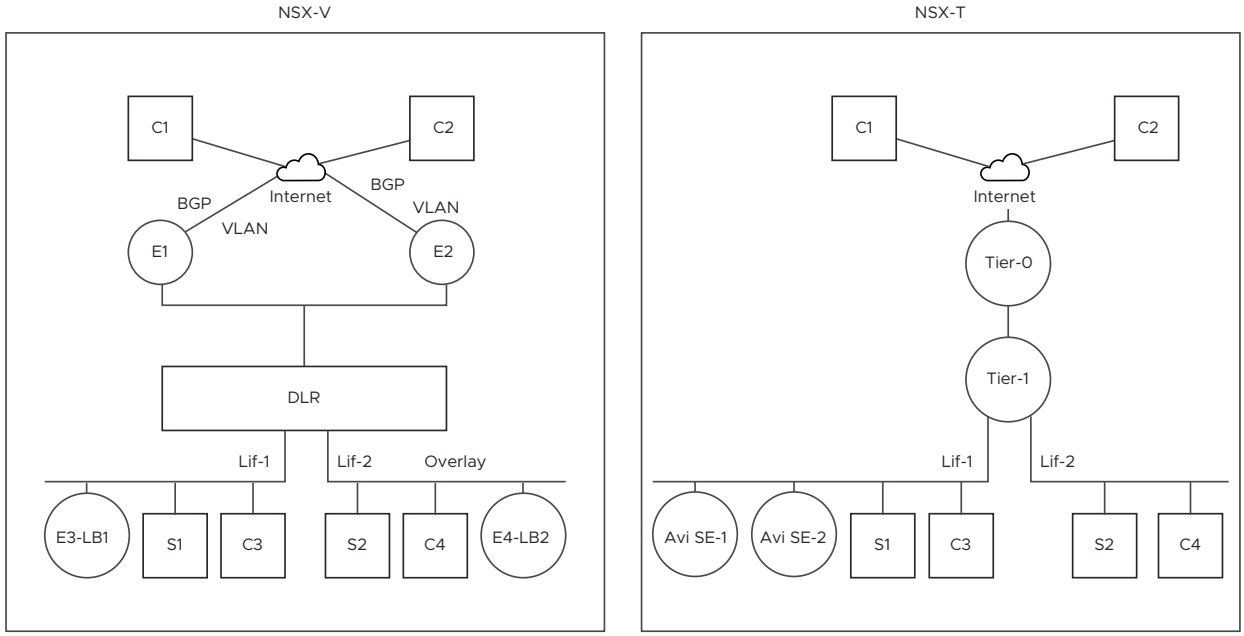

Topology 4: On the left: NSX-V inline LB with single-arm LB. On the right: NSX-T inline ALB with single-arm LB.

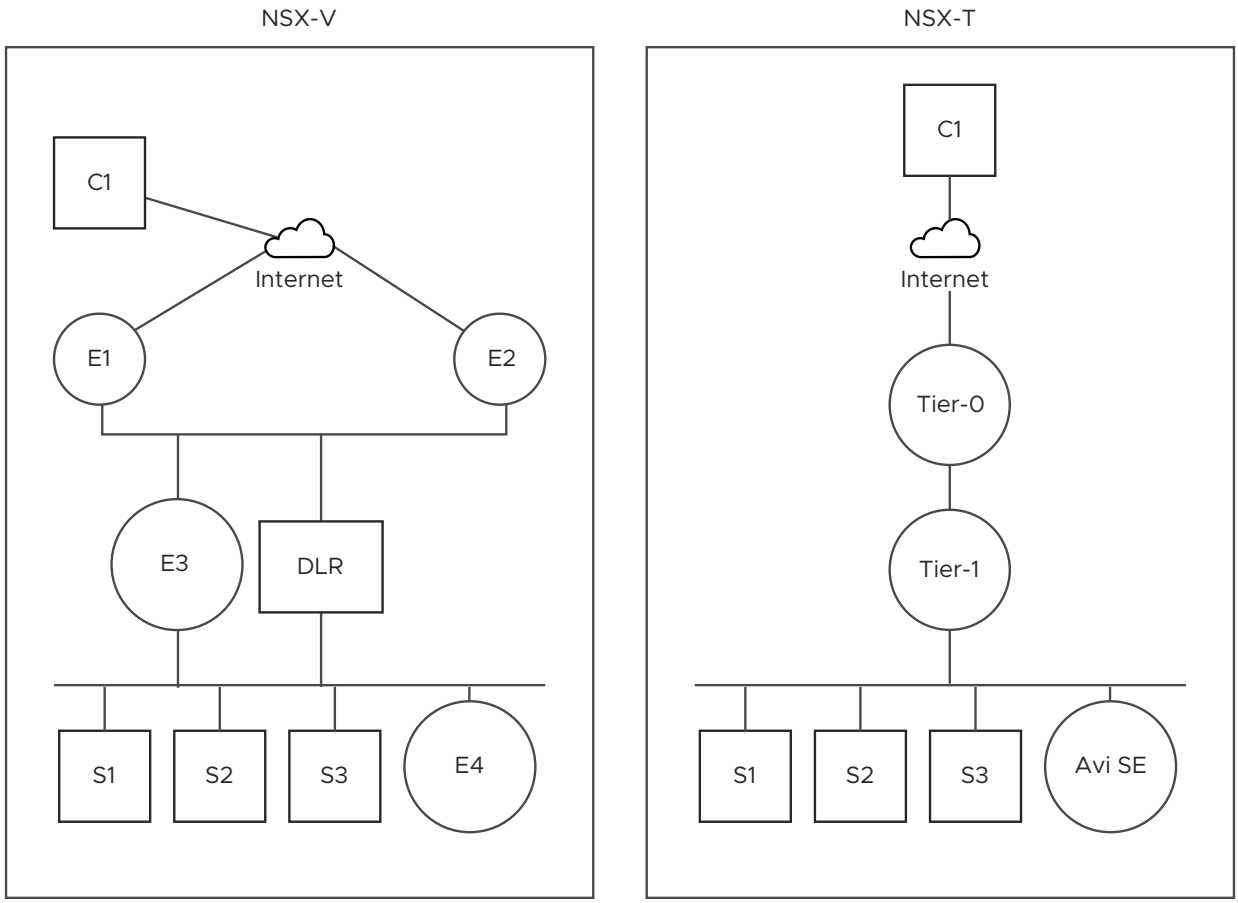

VLAN topologies:

Topology 5: On the left: NSX-V single-arm LB. On the right: NSX-T single-arm ALB.

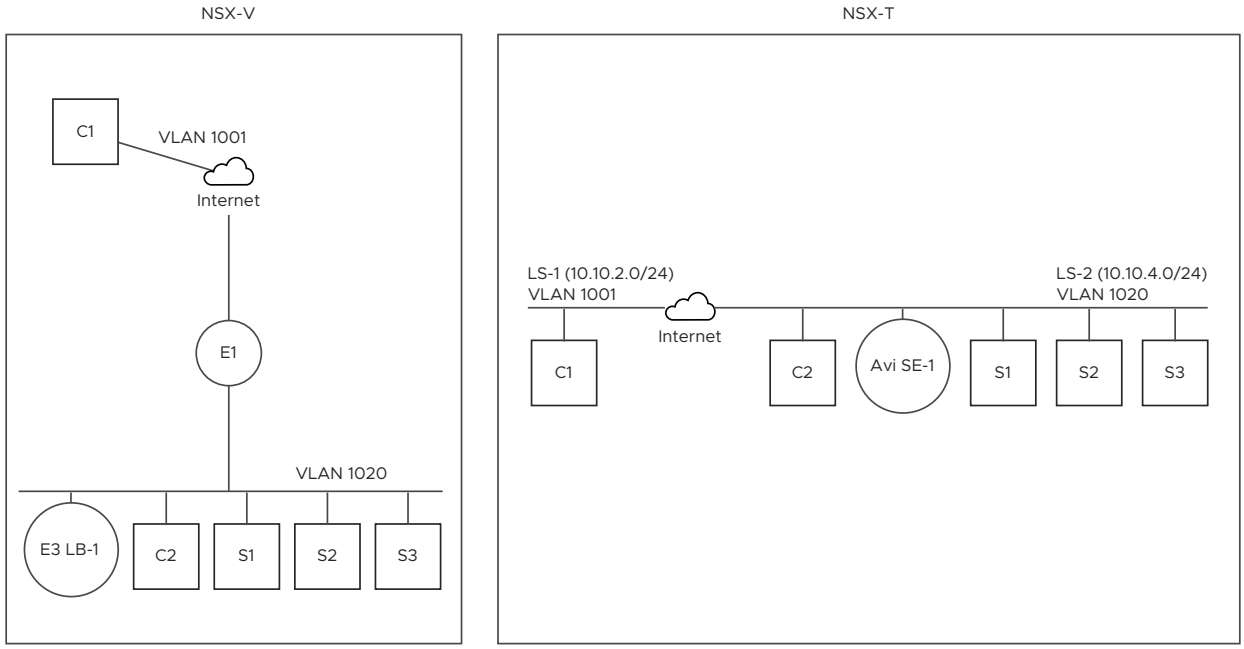

Topology 6: On the left: NSX-V inline LB with transparent mode. On the right: NSX-T inline ALB with transparent mode.

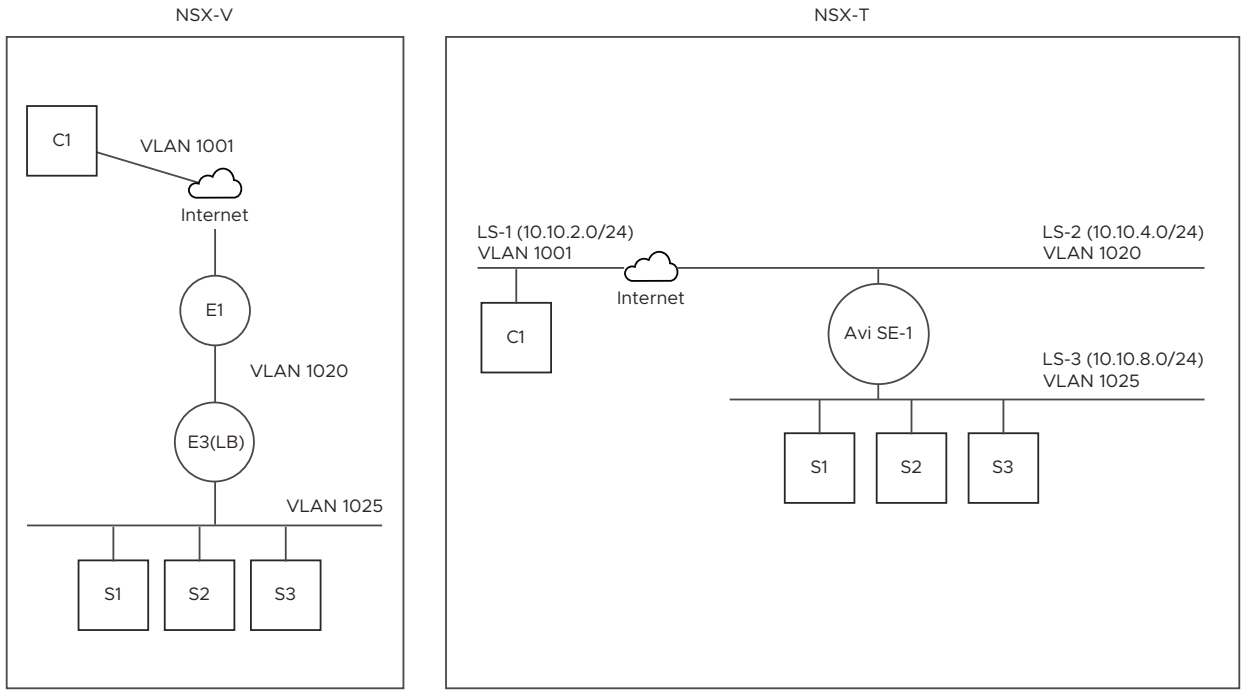

Topology 7: On the left: NSX-V inline nested LB. On the right: NSX-T inline nested ALB.

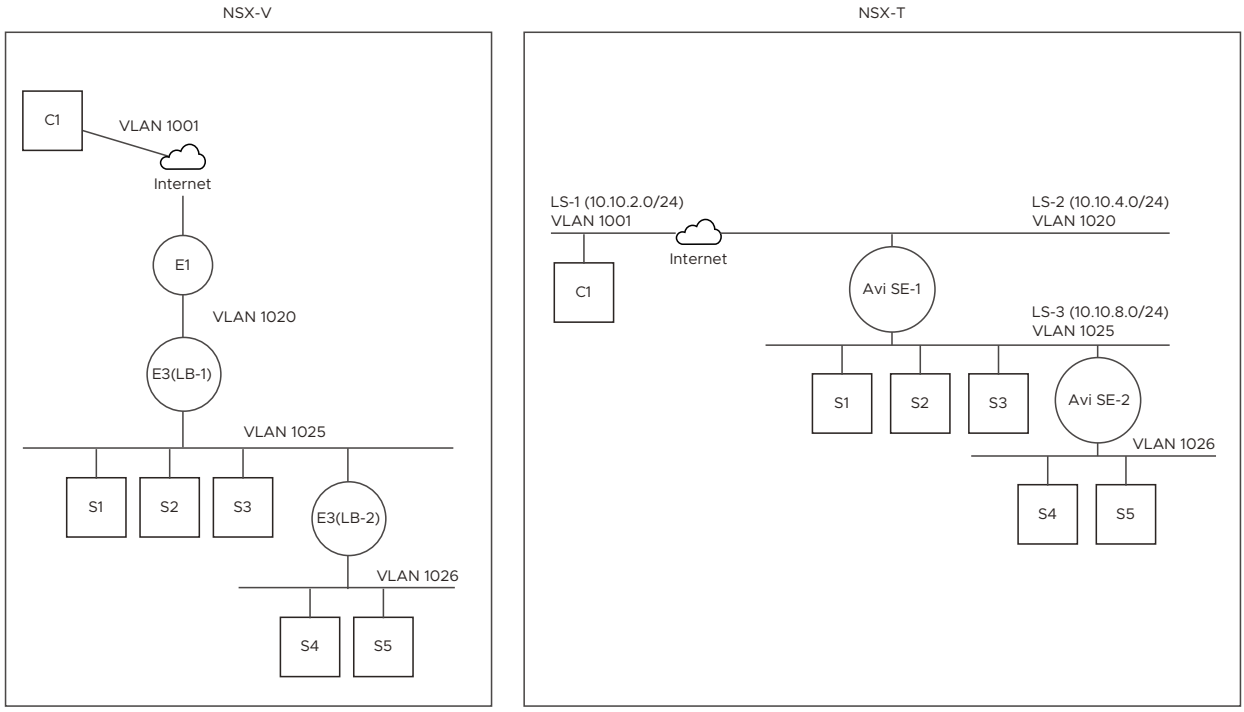

ALB Mapping File

For the migration, you can prepare a mapping file in JSON format that specifies how Edge Services Gateways (ESGs) should be mapped to SE groups. Providing a mapping file is optional. If you do not provide it, you will be prompted for any required information. The following is an example mapping file for an overlay topology:

```
{
  "alb": {
    "service_engine_group_per_esg": false,
    "esgs": [
      {
        "name": "edge-4",
        "interfaces": [
          {
            "name": "mgmt",
            "tier1_id": "London_Tier1Gateway1"
          },
          {
            "name": "vnic1",
            "placement_network_subnet": "172.16.1.10/16",
            "service_engine_group": "Test-SE-group"
          }
        ]
      }
    ]
  }
}
```

You can specify the following fields in the mapping file:

Field	Description
alb	Top level title
service_engine_group_per_esg	■ If not specified, with a basic license, the default SE group will be cloned if needed to handle the SE groups. With an enterprise license, all ESGs will be mapped to the default SE Group. ■ If set to `false`, you must specify `service_engine_group` which will be used for the virtual services belonging to that interface. ■ If set to `true`, an SE group will be created for every ESG.
esgs	The list of ESGs in your environment.
name	Name of the ESG.
default_tier1_id	If this is set all overlay interfaces of the ESG will be mapped to this tie-1 gateway.
interfaces	The list of interfaces on the ESG.
interfaces.name	The name of the interface.
interfaces.tier1_id	■ All the virtual services connected to the interface will be mapped to this tier-1 gateway. ■ This setting is for overlay virtual services. ■ This setting overrides the `default_tier1_id` value.
interfaces.placement_network_subnet	■ This specifies the subnets that are to be configured in the cloud network on Avi controller. ■ This setting is for VLAN virtual services.
interfaces.service_engine_group	■ This specifies the SE group for the virtual services connected to the interface. ■ This setting has precedence over other SE group settings.

Common Issues that Generate Feedback

The migration wizard will check the environment and provide feedback about issues to resolve before migration can proceed. The following table lists the issues and the actions you can take.

Issue	Action	Example
Missing configuration	Update the configuration	Missing cloud configuration on Avi Controller
Unsupported object	Skip	MSSQL Monitor
Unsupported configuration	Partially migrated	Transparent mode is skipped but the virtual services are migrated.
Missing information in mapping file	Provide input	Tier-1 or segment mapping missing
Orphan object	Skip	Pools not referenced in virtual services
Layer-2 dependencies	Skip	virtual services skipped if the corresponding segment is not migrated

The migration wizard will also provide informational messages about conditions that do not prevent migration. For example, if your environment has a basic license, you will see a message recommending an enterprise license.

Tier-1/Segment Mapping

For an overlay topology:

- Only single-arm load balancer is supported.
- The "esgs" section in the mapping file specifies the mapping.
- All virtual services and their pools will be migrated to the tier-1 specified in the mapping file.

For an VLAN topology:

- Both single and inline load balancers are supported.
- The VLAN segment for virtual services is derived from the migrated interface of the corresponding virtual service.
- For pools, the migration wizard will provide feedback about VLAN segment mapping.

Service Engine (SE) Group Mapping

Default mappings:

- Basic license
  - Supports only active-standby mode.
  - A new SE group is cloned for every 10th vNIC interface.
  - A new SE group is cloned if the number of virtual services per SE group exceeds the configured limit.
- Advanced license
  - Mapped to default SE group

In the mapping file, you can specify the following to override the default:

- Set the parameter "service_engine_group_per_esg" to true or false.
- Specify specific ESG mapping that will override the "service_engine_group_per_esg" parameter.

Note the following:

- All virtual services sharing the same VIP are mapped to the same SE group.
- If the shared virtual service count is greater than the number of virtual services per SE group, feedback will be generated to resolve the issue.

Import Configuration

This step imports the configuration of the NSX-V environment.

In NSX-T 3.2.0 and 3.2.1, when migrating a single-site environment, the NSX-V load balancer will be migrated to NSX-T Advanced Load Balancer (ALB). Starting with NSX-T 3.2.2, The NSX-V load balancer will be migrated to an NSX-T load balancer.

Starting with NSX-T 3.2.1, you can migrate an NSX-V cross-vCenter environment to an NSX Federation environment in NSX-T. If you are migrating a cross-vCenter environment to NSX Federation, you must log in to the Global Manager to start the migration. Note that in NSX-T 3.2.1, when migrating a cross-vCenter environment, the NSX-V load balancer will not be migrated. Starting with NSX-T 3.2.2, The NSX-V load balancer will be migrated to an NSX-T load balancer.

Procedure

1 From a browser, log in to NSX Manager or Global Manager as **admin**.

   Log in to Global Manager only if you are migrating a cross-vCenter environment to NSX Federation

2 Navigate to **System > Migrate**.

3 In the **NSX for vSphere** pane, click **Get Started** and select **User Defined Topology**.

4 Select a migration mode and click **Next**..

   The options are:

   - **Complete migration** - Configurations, edges, hosts and workloads are migrated. The migration happens in-place with no need for additional hosts. This mode will migrate the NSX-V Load Balancer configuration to NSX-T Load Balancer.
   - **Configuration migration** - Only configurations are migrated. Workload VMs are not migrated. If you logged in to NSX Manager (and not Global Manager), this mode will migrate the NSX-V Load Balancer to NSX-T Advanced Load Balancer (ALB).

5 Under **Authentication**, provide the required credentials.

   - If you are not migrating a cross-vCenter environment, provide the credentials for vCenter Server and NSX-V.
   - If you are migrating a cross-vCenter environment to NSX Federation, provide the credentials for the primary NSX-V and the NSX-T Global Manager. The list of NSX-V sites will then be listed. For each NSX-V site, provide the NSX-V credentials, select an NSX-T location, the vCenter Server for that location, and the credentials for that vCenter Server. Click **Check Status** to validate the information you provided.

6 (This step is applicable if you are running NSX-T 3.2.0 or 3.2.1, and you logged in to NSX Manager (and not Global Manager) and if you selected **Configuration migration** in step 4) Under **Advanced Load Balancing (ALB) Authentication**, if you plan to migrate to ALB, set the **LB Migration** toggle to **On** and provide the credentials for the ALB appliance. Otherwise, set the **LB Migration** toggle to **Off**.

7 Click **Start** to start the import.

8 In the confirmation dialog, click **Import** and wait for the process to complete.

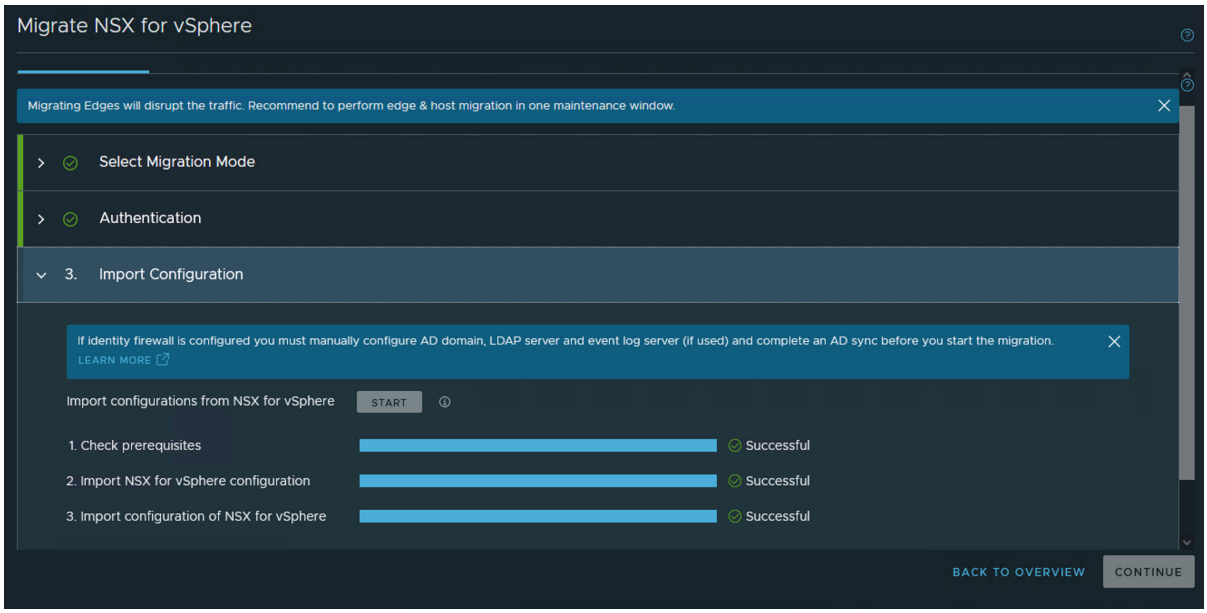

9 If the status is **Successful**, click **Continue** to go to the next step.

10 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Results

If the NSX-V topology is imported successfully, you can click the **View Imported Topology** link to view the imported topology. However, the topology viewer might not work properly for a large-scale NSX-V environment.

Translate Configuration Layer 2

This step translates the layer-2 NSX-V configuration that was imported.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Translate** and wait for the process to complete.

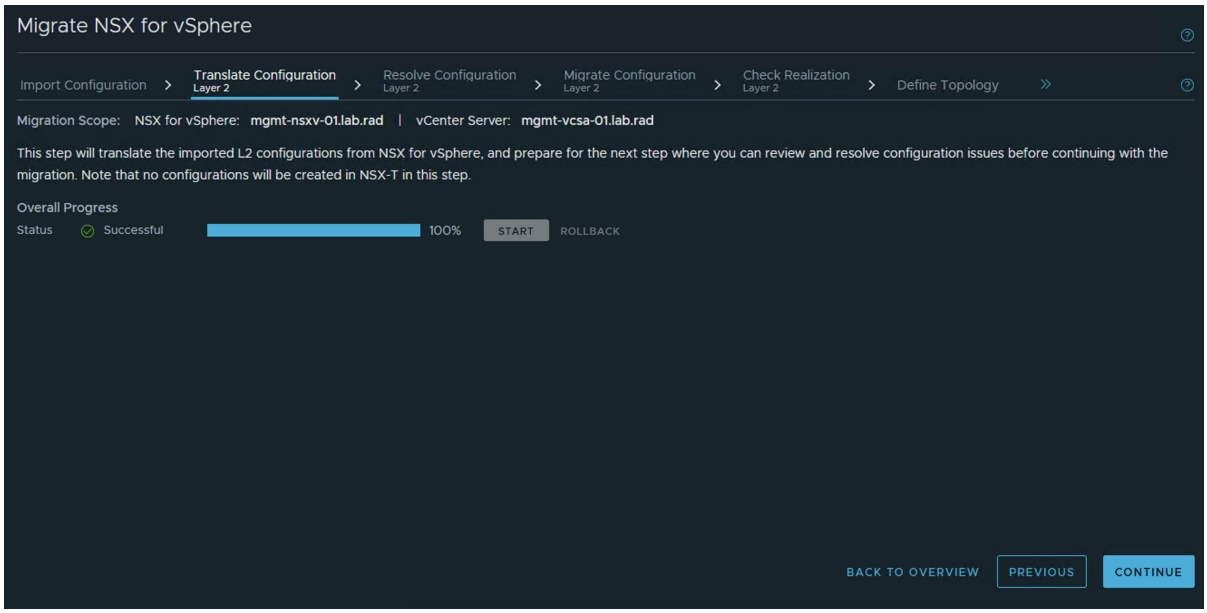

3 If the status is **Successful**, click **Continue** to go to the next step.

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Resolve Configuration Layer 2

This step resolves layer-2 issues in the configuration that was imported.

For more information about resolving issues, see Resolve Configuration Issues.

Procedure

1 Click the message for each issue to see the details. You can click **Accept** to accept the recommendation. You can also select all issues and click **Accept Recommendations** to accept the recommendations for all the issues.

2 Click **Submit** to confirm that you want to proceed with the resolution of the issues.

3 Click **Continue** to go to the next step.

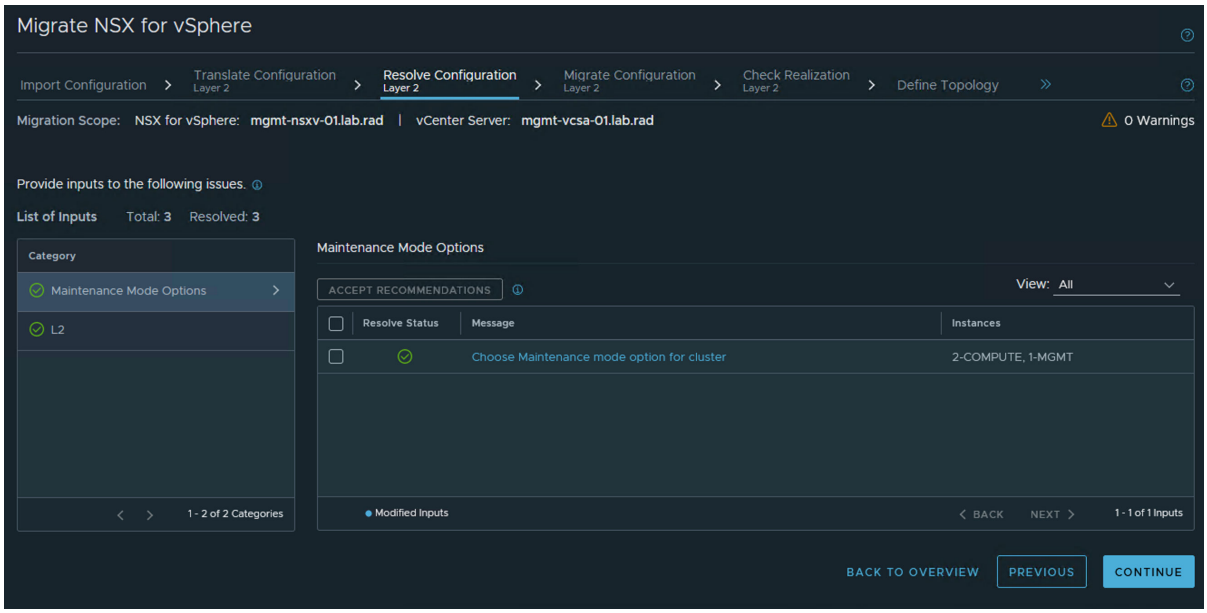

Migrate Configuration Layer 2

This step migrates the layer-2 configuration.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

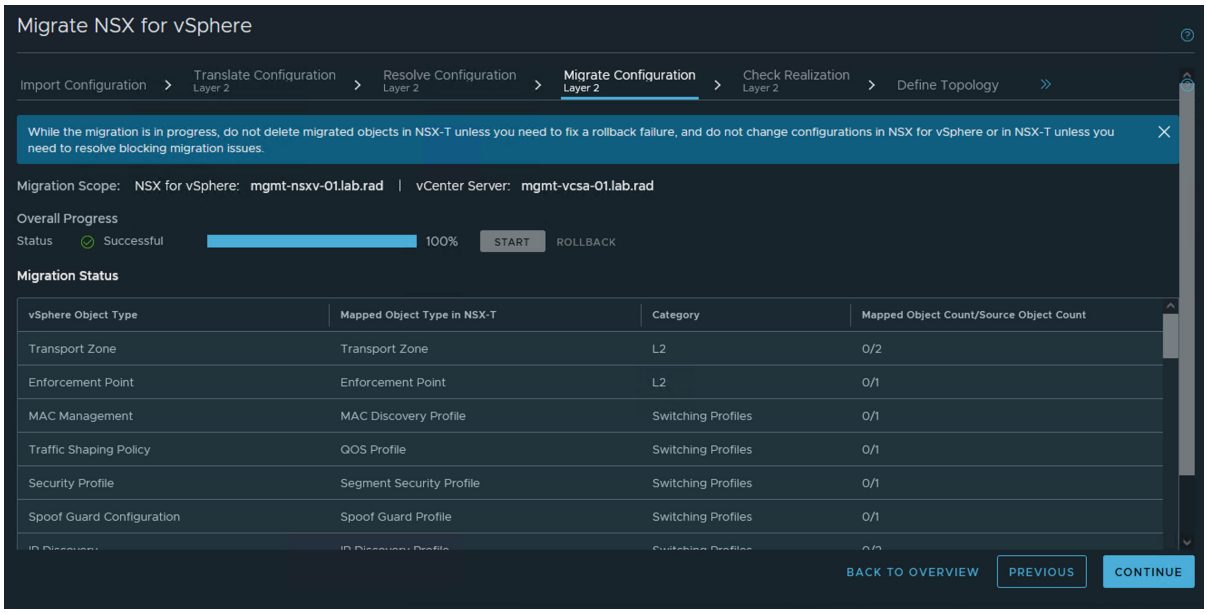

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Check Realization Layer 2

This step checks that the configuration that was migrated is realized in NSX-T.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

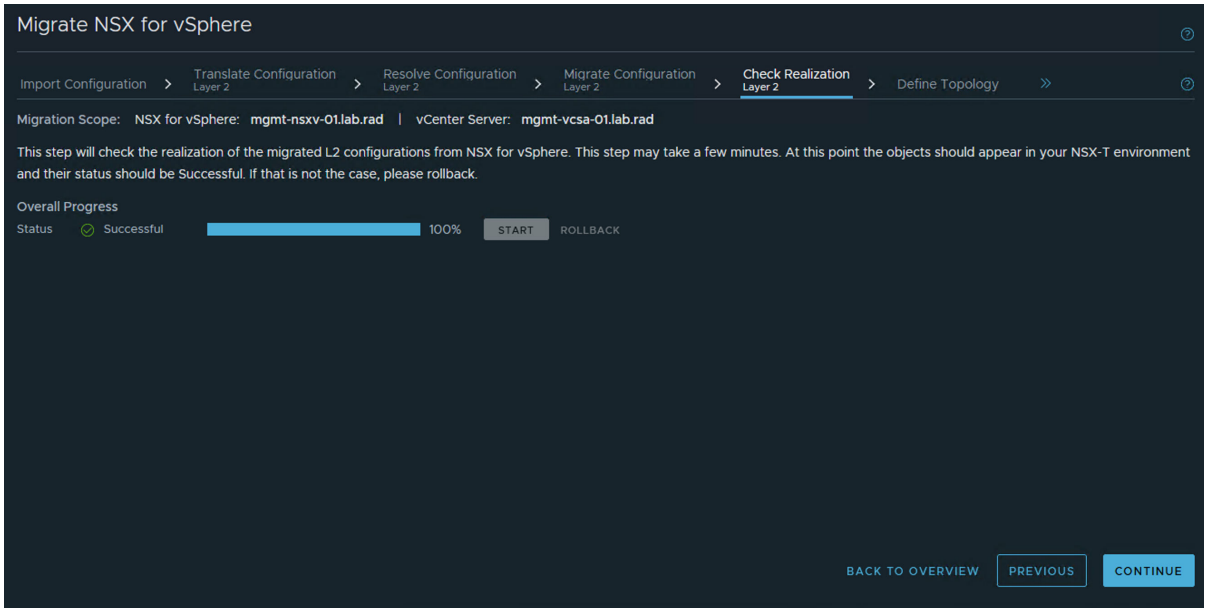

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Define a Topology

In this step, you specify the topology that will be migrated. This can be done through the NSX Manager UI or by using a mapping file in JSON format.

This mapping specifies how Edge Service Gateways (ESGs) and Distributed Logical Routers (DLRs) should map to gateways in NSX-T. Before providing the mapping, evaluate your topology requirements and plan to do the following:

1 Determine how you will map the ESGs and DLRs. The northbound ESGs without any L4-L7 services should be skipped. These are usually the ESGs peering with northbound routers and are in ECMP path. If you are using VPN on a northbound ESG, migrating to active-standby tier-0 is recommended. In other cases, migrating the ESGs/DLRs to tier-1 is recommended. An ESG and a DLR can be merged in one mapping entry.

2 Create tier-0 and tier-1 gateways and configure dynamic or static routing on the tier-0 gateways towards the northbound routers based on your requirements. For dynamic routing you can choose to configure either BGP or OSPF, based on your NSX-V configuration. You need to manually configure this northbound routing.

3 When configuring northbound routing, you must create and configure uplink interfaces on the NSX-T tier-0 gateway. The uplink interface subnet can be the same subnet as the NSX-V ESG northbound uplinks. The uplink interface IP addresses must be different from the IP addresses of the ESG uplinks.

4 After configuring the dynamic routing on tier-0 gateways, check that the dynamic routing has converged, that is, BGP sessions are established or OSPF neighborships are FULL as applicable. After this, proceed with providing the mapping.

When defining a mapping, make sure that the following conditions are met. Note that a tier-1 distributed router (DR)-only gateway is a tier-1 gateway without any Edge cluster.

- A tier-0 gateway must have an uplink interface.
- A tier-1 DR-only gateway must be connected to a tier-0 gateway that has an uplink interface.
- When UDLR is mapped to a stretched tier-1 DR only, the stretched tier-0 to which it is connected must have an uplink on all sites.
- When UDLR is mapped to a stretched tier-0, the stretched tier-0 must have an uplink on all sites.
- When UDLR is mapped to an active-standby stretched tier-1, the primary site of this gateway must match the primary site of the connected stretched tier-0.

An example of a mapping file that maps ESGs to a tier-0 gateway:

```
[
  {
    "name":"nsxv-to-nsxt-mapping",
    "v_edges_to_policy_gateways_mappings":[
      {
        "v_edges":[
          "edge-1",
          "edge-2"
         ],
         "policy_gateway_name": "tier0-gateway"
         "policy_gateway_path": "/infra/tier-0s/tier0-gateway"
      }
    ]
  }
]
```

If you are doing a configuration migration, the above-mentioned mapping is not used for Advanced Load Balancer (ALB). Instead there is another optional mapping that you may provide for ALB. To specify that mapping, you must upload a JSON file. An example of a mapping file that maps ESGs to Service Engine groups:

```
{
  "alb": {
    "service_engine_group_per_esg": false,
    "esgs": [
      {
        "name": "edge-4",
        "interfaces": [
          {
            "name": "mgmt",
            "tier1_id": "London_Tier1Gateway1"
          },
          {
            "name": "vnic1",
            "placement_network_subnet": "172.16.1.10/16",
            "service_engine_group": "Test-SE-group"
          }
        ]
      }
    ]
  }
}
```

Starting with NSX-T 3.2.1, you can migratie a cross-vCenter environment to NSX Federation. Here is a sample mapping file for such a migration:

```
[
    {
        "name": "london",
        "nsxv_id": "10.206.106.163",
        "nsxt_site_id": "1722c659-b0a9-4e70-b7ba-f264e057e1ea",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-2"
                ]
                "policy_gateway_name": "Tier1Gateway1",
                "policy_gateway_path": "/infra/tier-1s/Tier1Gateway1"
            }
        ]
    },
    {
        "name": paris",
        "nsxv_id": "10.206.96.206",
        "nsxt_site_id": "00d3802e-5673-4791-b86d-71805a2c0aa6",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-2"
                ]
```

```
                "policy_gateway_name": "Tier1Gateway1",
                "policy_gateway_path": "/infra/tier-1s/Tier1Gateway1"
            }
        ]
    },
    {
        "name": "site-GM",
        "nsxv_id": "10.206.106.163",
        "nsxt_site_id": "",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-4e5065d6-d12d-49b1-a7da-5d9fcc7888f0"
                ]
                "policy_gateway_name": "Tier1Gateway1",
                "policy_gateway_path": "/infra/tier-1s/Tier1Gateway1"
            }
        ]
    },
]
```

If you are migrating a cross-vCenter environment to NSX Federation, note the following:

- Tier-0 gateways created on Local Manager must have an Edge cluster assigned.
- Tier-1 gateways created on Local Manager must either have an Edge cluster assigned or be connected to a tier-0 gateway that has an Edge cluster assigned.
- Tier-0 and tier-1 gateways created on Global Manager must span all the sites.

If you are using the Manager UI to do the mapping, tier-0 and tier-1 gateways that do not satisfy the above conditions are not shown in the dropdown list.

For more information about creating a mapping file for the load balancer, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

Procedure

1 Choose one of the following options:

- Do not migrate L3-L7 entities and services.
- Upload a mapping file (*.json).
- Select a Tier-0 or Tier-1 gateway for each entity that needs to be migrated.

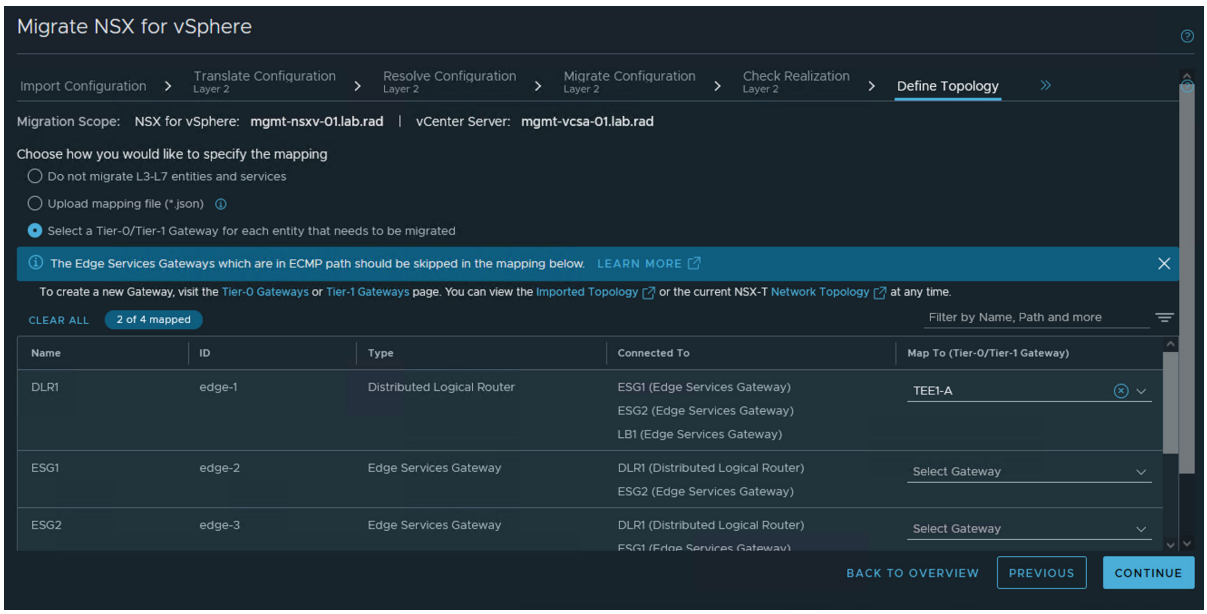

2 After you have specified a mapping, click **Continue** to go to the next step.

Translate Configuration Layer 3 and Above

This step translates the layer-3 and L4-L7 services of the NSX-V topology that you defined.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Translate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

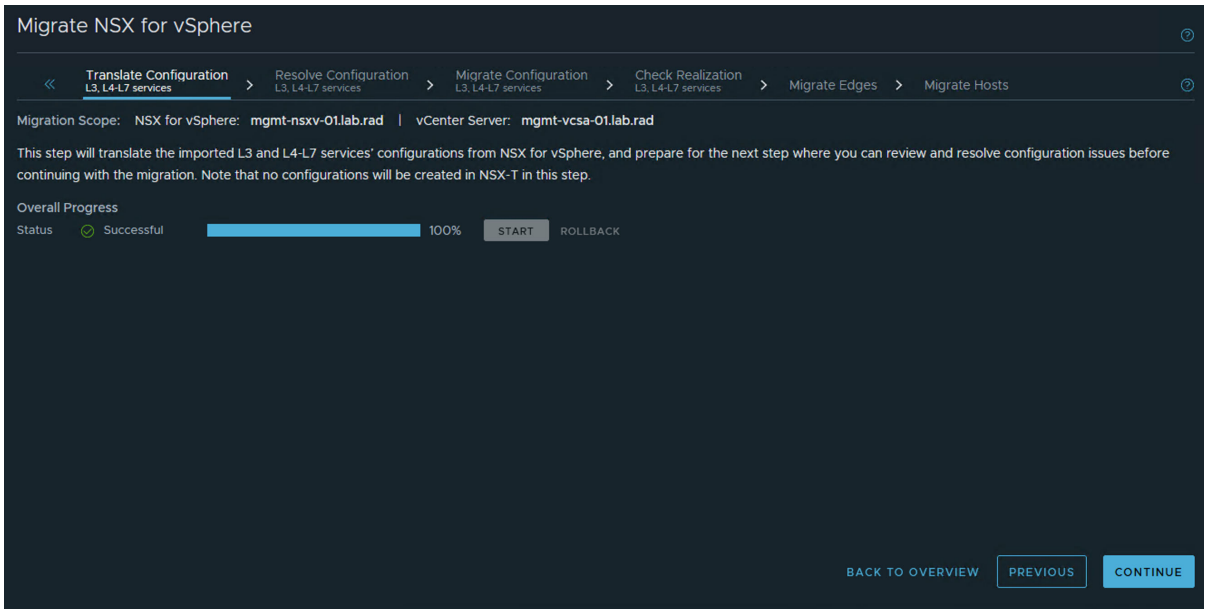

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Resolve Configuration Layer 3 and Above

This step resolves issues in the translation of the topology that you specified.

Issues for the following categories of objects will be displayed:

- NS Service
- Edge
- Other
- Profile
- RBAC

Note that for some features such as VPN, you might get the message "Feature ABC on Edge-XYZ cannot be migrated" even though the feature is not configured on NSX-V. You can simply accept the input and proceed with the migration.

Procedure

1 For each category, click the message for each issue to see the details. You can click **Accept** to accept the recommendation. You can also select all issues and click **Accept Recommendations** to accept the recommendations for all the issues.

2 Click **Submit** to confirm that you want to proceed with the resolution of the issues.

3 Click **Continue** to go to the next step.

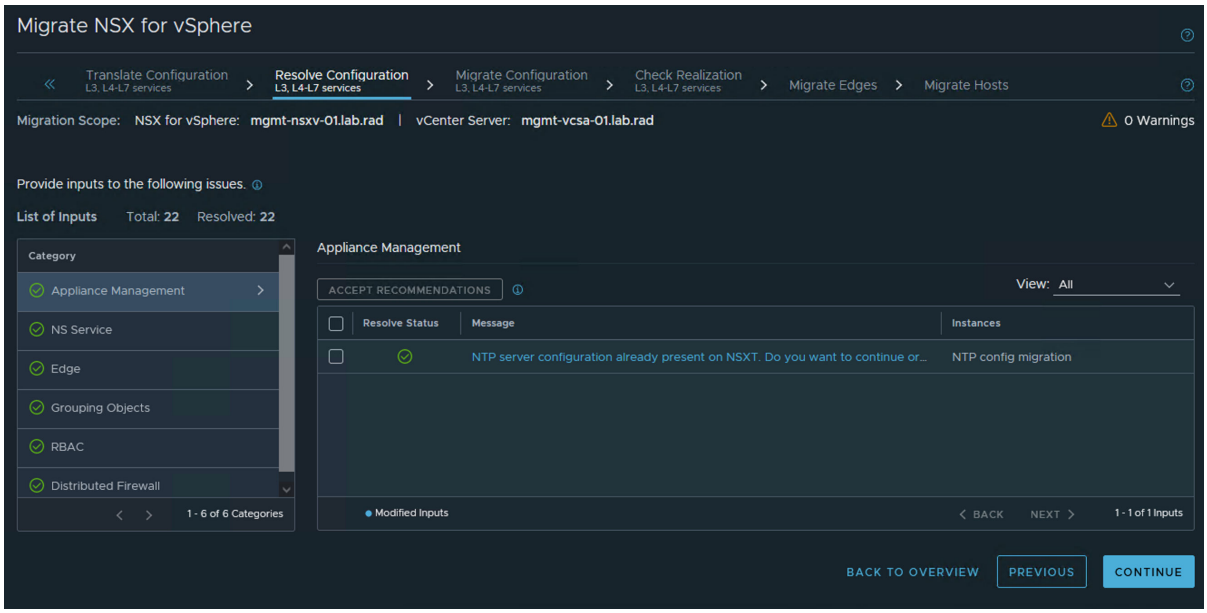

Migrate Configuration Layer 3 and Above

This step migrates the topology that you specified.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

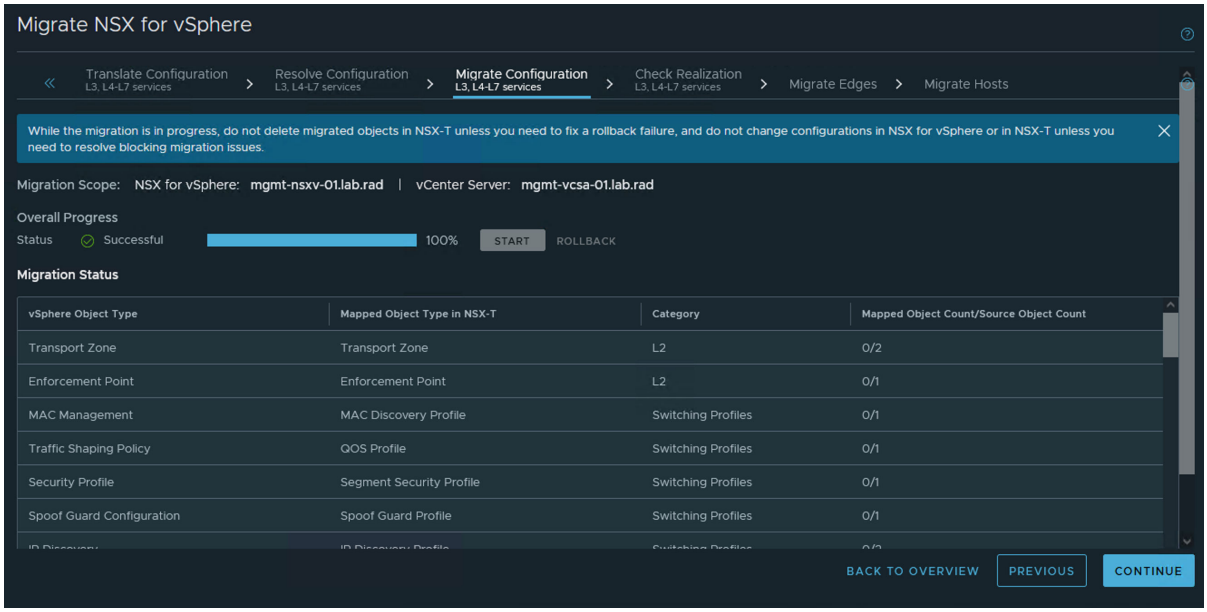

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Check Realization Layer 3 and Above

This step checks that the topology that was migrated is realized in NSX-T.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

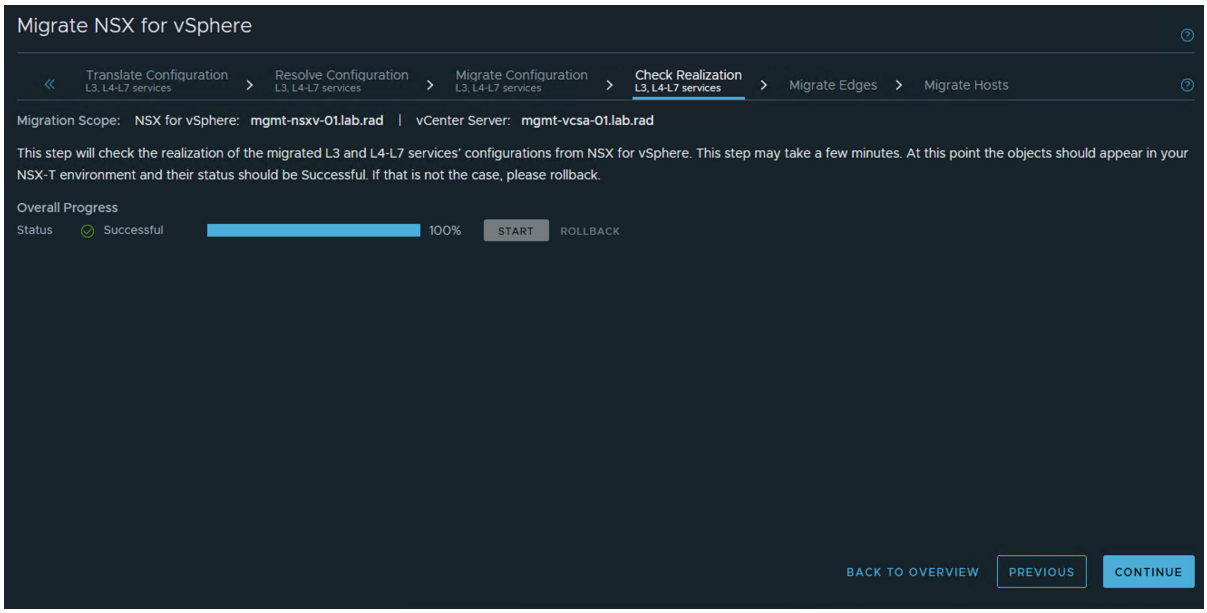

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** and resolve the issues. Then try the migration step again.

Migrate NSX-V Edges in End-to-End Migration

In this step, you migrate the Edge Services Gateways (ESGs).

If you have no Edge Services Gateway appliances in your topology, you must still click **Start** so that you can proceed to the **Migrate Hosts** step.

Caution If you roll back the **Migrate Edges** step, verify that the traffic is going back through the NSX-V Edge Services Gateways. You might need to take manual action to assist the rollback.

Prerequisites

- All configuration issues must be resolved.
- The NSX-V configuration must be migrated to NSX-T.
- Verify that the migrated configurations are shown in the NSX Manager UI or API of NSX-T.
- Verify that you have a backup of NSX-V and vSphere since the most recent configuration changes were made.
- If you are using new IP addresses for the NSX-T Edge node uplinks, you must configure the northbound routers with these new BGP neighbor IP addresses.
- Verify that you have created an IP pool for Edge Tunnel End Points (TEP). See Create an IP Pool for Edge Tunnel End Points.

- Logical router interfaces created in NSX-T use the global default MTU setting, which is 1500. If you want to ensure that all logical router interfaces have a larger MTU, you can change the global default MTU setting. For more information, see Change the Global MTU Setting.

  If MTU setting other than 1500 is used on peering routers, the same should be configured on NSX-T. In case of OSPF topologies, OSPF adjacencies can get stuck if MTU setting is different from peering routers' MTU setting.

Procedure

1 From the **Migrate Edges** page, click **Start**.

  All Edges are migrated. The uplinks on the NSX-V Edge Services Gateways are internally disconnected, and the uplinks on the NSX-T Edge nodes are brought online.

2 Verify that routing and services are working correctly in the new NSX-T environment.

  If so, you can migrate the hosts. Before migrating the hosts, see Configuring NSX-V Host Migration.

3 Click **Continue** to go to the next step.

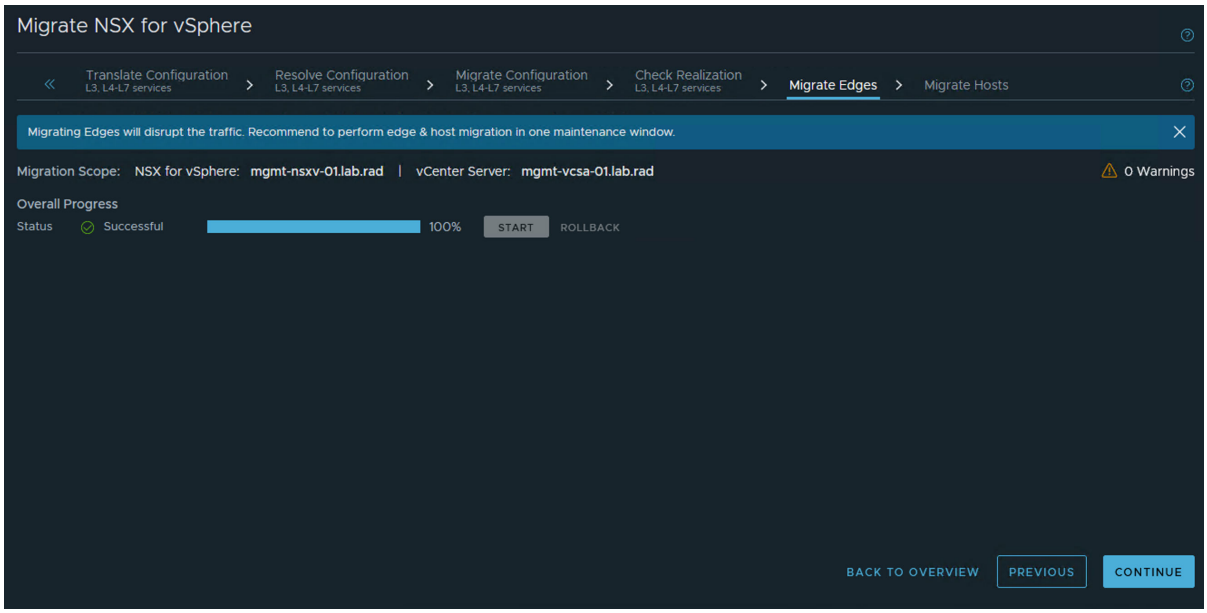

Results

The following changes result from the migration process:

- The routing and service configuration from NSX-V Edge Services Gateway (ESG) are transferred to the newly created NSX-T Edge nodes.
- The new TEP IP addresses for the newly created NSX-T Edge nodes are configured from a newly created IP pool for Edge Tunnel End Points.

Migrating Hosts in End-to-End Migration

Before migrating the hosts, select a host migration plan.

For more information about what happens during host migration, see Changes Made During Host Migration in an End-to-End Migration.

If the Security Policies in your NSX-V environment use a partner service for Guest Introspection or Network Introspection or both, choose the host migration mode as shown in the following table.

Partner Service	Host Migration Mode
Only Guest Introspection	In-Place and Maintenance modes are supported.
Only Network Introspection	Maintenance mode is supported. However, Automated Maintenance mode is recommended. In-Place mode is not supported.
Both Guest Introspection and Network Introspection	Maintenance mode is supported. In-Place mode is not supported.

Important Consult the VMware partner before migrating their service that is running on the NSX-V workloads. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their services to NSX-T.

Caution Host migration should be completed during the same maintenance window as Edge migration.

You must disable IPFIX and reboot the ESXi hosts before migrating them.

If the partner service in your NSX-V environment provides Guest Introspection or both Guest Introspection and Network Introspection service, follow the procedure in Migrate Hosts with Guest Introspection Service.

If the partner service in your NSX-V environment provides only Network Introspection service, follow the procedure in Migrate Hosts with Network Introspection Service.

Select a Host Migration Plan

The clusters in the NSX-V environment are displayed on the **Migrate Hosts** page. The clusters are arranged into migration groups, each migration group contains one vSphere host cluster. There are several settings which control how the host migration is performed.

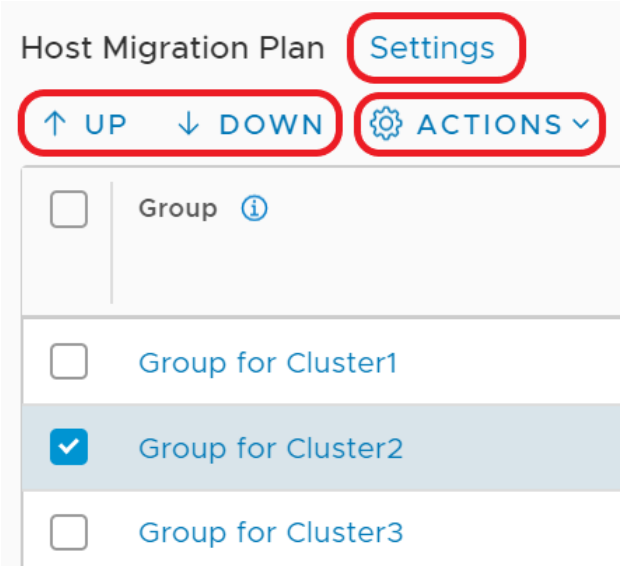

- Click **Settings** to change the global settings: **Pause Between Groups** and **Migration Order Across Groups**.
- Select a single host group (cluster) and use the arrows to move it up or down in the migration sequence.
- Select one or more host groups (clusters) and click **Actions** to change these host groups settings: **Migration Order Within Groups**, **Migration State**, and **Migration Mode**.

Pause Between Groups

When migrating multiple host groups, you can pause the migration between groups by enabling the **Pause Between Groups** setting. After a group is migrated, you must click **Continue** to migrate the next host group. This setting is disabled by default. You can enable it if you want to verify the status of the applications running on each cluster before proceeding to the next one.

Serial or Parallel Migration Order

You can specify whether migration happens in a serial or parallel order. There are two ordering settings:

- **Migration Order Across Groups** is a global setting that applies to all host groups.
  - **Serial**: One host group (cluster) at a time is migrated.
  - **Parallel**: Up to five host groups at a time are migrated. After those five host groups are migrated, the next batch of up to five host groups are migrated.

    Important If you are migrating from NSX-V 6.4.4, 6.4.5, or 6.4.6, and your environment uses vSphere Distributed Switch 7.0 or later, do not select parallel migration order across groups.

    If you are migrating from NSX-V 6.4.8 or later, and your environment uses vSphere Distributed Switch 7.0 or later, parallel migration order across groups is supported.

- **Migration Order Within Groups** is a host group (cluster) specific setting, so can be configured separately on each host group.
  - **Serial**: One host within the host group (cluster) at a time is migrated.
  - **Parallel**: Up to five hosts within the host group are migrated at a time. After those hosts are migrated, the next batch of up to five hosts are migrated.

    Important Do not select parallel migration order within groups for a cluster if you plan to use **Maintenance** migration mode for that cluster.

By default, both settings are set to **Serial**. Together, the settings determine how many hosts are migrated at a time.

Table 6-12. Effects of Migration Settings on Number of Hosts Attempting Migration Simultaneously

Migration Order Across Groups (Clusters)	Migration Order Within Groups (Clusters)	Maximum Number of Hosts Attempting Migration Simultaneously
Serial	Serial	1 One host from one host group
Serial	Parallel	5 Five hosts from one host group
Parallel	Serial	5 One host from five host groups
Parallel	Parallel	25 Five hosts from five host groups

Important If there is a failure to migrate a host, the migration process will pause after all in-progress host migrations have finished. If **Parallel** is selected for both migration across groups and migration within groups, there might be a long outage for the failed host before you can retry migration.

Sequence of Migration Groups

You can select a host group (cluster) and use the arrows to move it up or down in the list of groups.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

Migration State

Host groups (clusters) can have one of two migration states:

- **Enabled**

  Hosts groups with a migration state of **Enabled** are migrated to NSX-T when you click **Start** on the **Migrate Hosts** page.

- **Disabled**

  You can temporarily exclude host groups from migration by setting the migration state for the groups to **Disabled**. Hosts in disabled groups are not migrated to NSX-T when you click **Start** on the **Migrate Hosts** page. However, you must enable and migrate all **Disabled** host groups before you can click **Finish**. Finish all host migration tasks and click **Finish** within the same maintenance window.

In the **Resolve Configuration** step, hosts that are ineligible for migration are identified. In the **Migrate Hosts** step, these hosts have the migration state **Do not migrate**. For example, hosts that do not have NSX-V installed are not eligible for migration.

Migration Mode

Migration Mode is a host group (cluster) specific setting, and can be configured separately on each host group. In the **Migrate Hosts** step, you select whether to use **In-Place** or **Maintenance** mode.

There are two types of Maintenance migration modes:

- Automated
- Manual

In the **Resolve Configuration** step of the migration process, you select which type of Maintenance migration mode to use. You select a Maintenance mode even if you plan to migrate hosts using **In-Place** mode. When you select Maintenance migration mode in the **Migrate Hosts** step, the value you specified in the **Resolve Configuration step** determines whether Automated Maintenance mode or Manual Maintenance mode is used. However, in the **Migrate Hosts** step, if you select **In-Place** mode, your selected choice of Maintenance mode in the **Resolve Configuration** step does not take effect.

In-Place migration mode is not supported if your NSX-V installation uses vSphere Distributed Switch 7.0 or later.

If your environment uses Distributed Firewall, select **Automated Maintenance** migration mode. If you select a different migration mode, the following limitations apply to environments with Distributed Firewall:

- If you use **Manual Maintenance** migration mode, all VMs must be moved to NSX-T hosts, connected to NSX-T segments, and powered on before the last NSX-V host starts migrating. When you migrate your last NSX-V host, do not power off the VMs on the host. Move them to an NSX-T host using vMotion.
- If you use **Manual Maintenance** migration mode, VMs have a gap in firewall protection for up to 5 minutes after they move to an NSX-T host.
- If you use **In-Place** migration mode, and you have Distributed Firewall rules that are applied to a VM, those rules are not pushed to the host until the host and all its VMs are migrated. Until the rules are pushed to the host, the following applies:
  - If the NSX-T default rule is `deny`, the VM is not accessible.

- If the NSX-T default rule is `accept`, the VM is not protected by the applied-to rules.

The migration process is different for each migration mode:

- **In-Place** migration mode

  NSX-T is installed and NSX components are migrated while VMs are running on the hosts. Hosts are not put in maintenance mode during migration. Virtual machines experience a short network outage and network storage I/O outage during the migration.

- **Automated Maintenance** migration mode

  A task of entering maintenance mode is automatically queued. VMs are moved to other hosts using vMotion. Depending on availability and capacity, VMs are migrated to NSX-V or NSX-T hosts. After the host is evacuated, the host enters maintenance mode, NSX-T is installed, and NSX components are migrated. VMs are migrated back to the newly configured NSX-T host. Note that VMs that are powered off will not be reconfigured. After migration, you need to manually configure these VMs before powering them on.

- **Manual Maintenance** migration mode

  A task of entering maintenance mode is automatically queued. To allow the host to enter maintenance mode, do one of the following tasks:

  - Power off all VMs on the hosts.
  - Move the VMs to another host using vMotion or cold migration.

  Once the host is in maintenance mode, NSX-T is installed on the host and NSX components are migrated. After the host is migrated, for the powered-off VMs and the VMs that you moved, you will need to change their network connection from the NSX-V logical switch to an NSX-T segment.

In the NSX-V environment, if the ESXi host's vmk0 management interface is connected to a VSS (vSphere Standard Switch) portgroup that does not have an uplink, and the portgroup is bridged to a VDS portgroup, and the VDS version is 6.5, 6.6 or 6.7, you must migrate using the **Maintenance** mode. If you use the **In-Place** mode, the migration will fail.

Adding or Removing a Host During Migration

Starting with NSX-T 3.2.1, during the host migration step, you can add or remove a host to be migrated when there is a pause in the migration.

Starting with NSX-T 3.2.2, you can pause the migration of hosts within a group. For more information, see Configuring NSX-V Host Migration.

Host migration will pause if you enable the setting **Pause Between Groups** or **Pause Between Hosts**, or if there is a failure to migrate a host.

You can add a host to a cluster before or after the cluster has been migrated, or while the cluster is being migrated. You can also remove a host from a cluster that has not migrated or is being migrated.

The host you want to add can be a standalone host or in a cluster. The host must not have NSX-V or NSX-T configured. If the host is in a cluster, the cluster must not have NSX-V or NSX-T configured.

You must prepare the host as a transport node first and later move it into a target cluster that has been migrated to NSX-T, is being migrated or has not started the migration. If the target cluster already has a host migrated to an NSX-T transport node, you can use the transport node's configuration as a reference to prepare the host as a transport node. If the host will support overlay traffic and a VTEP IP pool will be used to prepare the host as a transport node, the VTEP IP pool cannot be or overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated. It can be an existing IP pool (such as the one used by NSX-T Edge nodes) or a newly created IP pool. For more information, see the section "Preparing ESXi and KVM Hosts as Transport Nodes" in the *NSX-T Data Center Installation Guide*.

When preparing an ESXi host as a transport node, you can choose N-VDS or VDS as the host switch. Choose VDS if the version of the VDS being migrated is 7.0 or later. Otherwise, choose N-VDS. If you prepare a host with N-VDS when you should choose VDS, the host will still be migrated but it may have network issues.

Adding a host to a cluster

1 From vCenter Server, put the host in Maintenance Mode.

2 If there is no VTEP IP pool that does not overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated, or such pools do not have enough free IPs for the VTEPs to be created for the host to be added, then go to **Networking > IP Address Pools** and create a new VTEP IP pool.

3 Follow the instructions in the installation guide to prepare the host as a transport node. When you select the transport zone, if an overlay transport zone is chosen for the host switch, choose the new IP pool that was created in step 2 or choose an existing IP pool that does not overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated. Select an uplink profile. Do not choose the one whose name contains "VXLAN" if an overlay transport zone is chosen for the host switch.

4 Wait for the status of the host node to be "Success". Do not move the host out of Maintenance Mode.

5 Choose a cluster into which the host will be added. In NSX-T Manager UI, go to **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later) and check if the cluster has a Transport Node Profile (TNP) attached. Detach the TNP if it does.

6 In vCenter Server UI, move the host into the chosen cluster.

7 Invoke the sync host groups API or click the **Refresh** button on the NSX-T Manager UI host migration screen so that the migration group for the cluster contains the new host.

8 Call the following NSX-T API to accept the new host:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/setup?
action=migrate_newly_added_host_transport_node
{
    "host_transport_node_id" : "<transport-node-uuid>"
}
```

If the API returns an error, fix the error and retry the API. If the API returns success, then make the host exit maintenance mode.

9 vMotion VMs to the host. If any VM is moved from an NSX-V host, be sure to change the network to map the source virtual-wire to NSX-T overlay segments. For example, virtual-wire vxw-dvs-64-virtualwire-4-sid-10787-1-switch-191 must map to 1-switch-191-LS.

Removing a host from a cluster

1 From vCenter Server, migrate all VMs off the host and enter the host into maintenance mode.

2 If the host is in a cluster that has not started migrating, go to the next step. Otherwise, from NSX-T Manager UI, go to **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later). If the host has NSX configured, delete it. If the host is not reachable by NSX Manager, delete it with the force option.

3 From vCenter Server, remove the host. Wait until the task is complete.

4 Click the Refresh button on the host migration screen to remove the host from the migration group.

5 Restart the host migration.

Migrate Hosts with Guest Introspection Service

After you migrate the Edge Services Gateways successfully, you can migrate the NSX-V hosts to NSX-T host transport nodes.

Prerequisites

- Verify that Edge migration has finished and all routing and services are working correctly.
- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.

Procedure

1 On the **Migrate Hosts** page, click **Start**.

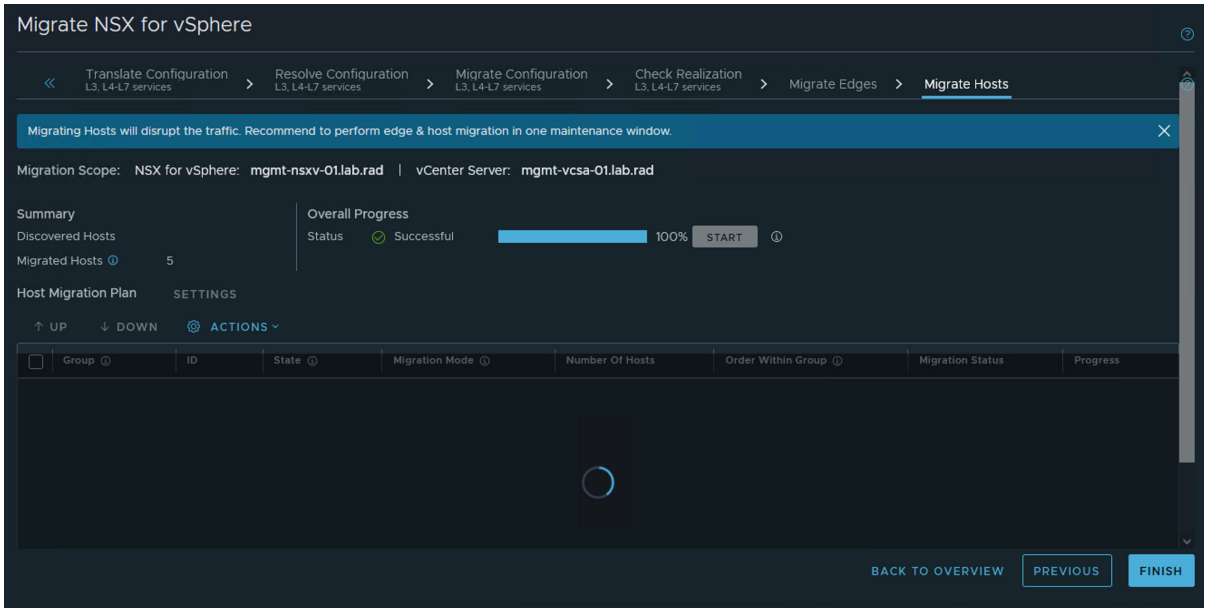

If you selected the **In-Place** or **Automated Maintenance** migration mode for all hosts groups, the host migration starts. Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

2 If you selected the **Manual Maintenance** migration mode for any host groups, you must complete one of the following tasks for each VM so that the hosts can enter maintenance mode.

Option	Action
Power off or suspend VMs.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b After the host has migrated, attach the VM interfaces to the appropriate NSX-T segments and power on the VM.
Move VMs using vMotion.	Right click the VM and select Migrate. Follow the prompts to move the VM to a different host. Note that Migration Coordinator maintains security during migration by vMotioning VMs to specific ports that are protected by temporary rules. In the case of manual vMotion, the VMs will not be moved to those ports and there could be a security breach. To vMotion manually, the VMs must be migrated using vSphere API where the networking backing must point to the OpaqueNetwork ID corresponding to the NSX Segment when using NVDS or the VDS portgroup ID when using VDS 7. In both cases, the network device's externalId must be set to the the string "VM_UUID:vNIC_ID", where VM_UUID is the VM's instance UUID and vNIC_ID is the VM's vNIC index where the first vNIC is 4000.
Move VMs using cold migration.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b Right click the VM and select Migrate. Follow the prompts to move the VM to a different host, connecting the VM interfaces to the appropriate NSX-T segments.

Here is python code to specify an external-id for each vNIC in a VM and then vMotion the VM so that the vNICs will connect to an NSX-T segment of ID “ls_id” at the correct ports:

```
devices = vmObject.config.hardware.device
nic_devices = [device for device in devices if isinstance(device,
vim.Vm.device.VirtualEthernetCard)]
vnic_changes = []
for device in nic_devices:
    vif_id = vmObject.config.instanceUuid + ":" + str(device.key)
    vnic_spec = self._get_nsxt_vnic_spec(device, ls_id, vif_id)
    vnic_changes.append(vnic_spec)
relocate_spec = vim.Vm.RelocateSpec()
relocate_spec.SetDeviceChange(vnic_changes)
# set other fields in the relocate_spec
vmotion_task = vmObject.Relocate(relocate_spec)
WaitForTask(vmotion_task)

def _get_nsxt_vnic_spec(self, device, ls_id, vif_id):
    nsxt_backing = vim.Vm.Device.VirtualEthernetCard.OpaqueNetworkBackingInfo()
    nsxt_backing.SetOpaqueNetworkId(ls_id)
    nsxt_backing.SetOpaqueNetworkType('nsx.LogicalSwitch')
    device.SetBacking(nsxt_backing)
    device.SetExternalId(vif_id)
```

```
    dev_spec = vim.Vm.Device.VirtualDeviceSpec()
    dev_spec.SetOperation(vim.Vm.Device.VirtualDeviceSpec.Operation.edit)
    dev_spec.SetDevice(device)
    return dev_spec
```

For an example of a complete script, see https://github.com/dixonly/samples/blob/main/vmotion.py

The host enters maintenance mode after all VMs are moved, powered off, or suspended. If you want to use cold migration to move the VMs to a different host before the migrating host enters maintenance mode, you must leave at least one VM running while you move VMs. When the last VM is powered off or suspended, the host enters maintenance mode, and migration of the host to NSX-T starts.

Results

After a host has migrated to NSX-T using **In-Place** migration mode, you might see a critical alarm with message `Network connectivity lost`. This alarm occurs when a vSphere Distributed Switch (VDS) 6.5 or 6.7 migrates to an N-VDS because the host no longer has a physical NIC connected to the VDS it was previously connected to. To restore the migrated hosts to the Connected state, click **Reset to Green** on each host, and suppress the warnings, if any.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

If migration fails for a host, the migration pauses after all in-progress host migrations finish. When you have resolved the problem with the host, click **Retry** to retry migration of the failed host. If the host still fails to migrate, you can configure NSX-T on the host manually or remove the host from the system. In this case, at the end of the host migration step, the **Finish** button will not be enabled because of the host that failed to migrate. You need to call the REST API `POST https://<nsx-mgr-IP>/api/v1/migration?action=finalize_infra` (<nsx-mgr-IP> is the IP address of the NSX Manager where the migration service is running) using a REST API client (for example, postman or curl) to finish the migration, and then perform the post-migration tasks.

For information about troubleshooting other host migration problems, see Chapter 13 Troubleshooting Migration Issues.

What to do next

If the migrated Security Policies use a third-party partner service, deploy an instance of the partner service in NSX-T. For detailed instructions, see:

- Deploy a Partner Service for Endpoint Protection

  Click this link to deploy a partner service that provides only Endpoint Protection or both Endpoint Protection and Network Introspection services to the NSX-T workload VMs.

- Deploy a Partner Service for Network Introspection

  Click this link to deploy a partner service that provides only Network Introspection service to the NSX-T workload VMs.

Migrate Hosts with Network Introspection Service

When Security Policies in your NSX-V environment use only a Network Introspection service that is provided by a partner, two approaches are available to migrate the NSX-V prepared hosts to NSX-T.

Both the approaches discussed in this topic assume that the partner service virtual machines (SVMs) in your NSX-V environment are not deleted before starting the migration coordinator. Depending on how much security protection downtime you are willing to accept during host migration, choose the host migration approach that best suits your needs.

Note Consult the VMware partner before migrating the hosts by using any of the two approaches. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their service to NSX-T.

When only Network Introspection service is running on your NSX-V hosts, **In-Place** host migration mode is not supported. Only **Maintenance** migration mode is supported. However, Automated Maintenance migration mode is recommended.

Approach 1: Involves more security protection downtime

This approach is the simpler of the two host migration approaches. However, it involves more security protection downtime compared to Approach 2. Let us say that you have three clusters in your NSX-V environment: Cluster 1, Cluster 2, and Cluster 3.

In this approach, enable the **Pause between groups** migration setting and migrate Cluster 1 by using the standard host migration procedure that is explained in Migrate NSX-V Hosts. After Cluster 1 is migrated to NSX-T, the migration pauses. Deploy the partner service in Cluster 1 by doing either a host-based or a clustered service deployment. Now, disable the **Pause between groups** migration setting, and continue migrating Clusters 2 and 3. After the workload VMs in Clusters 2 and 3 are migrated to NSX-T, these workloads can start redirecting packets to the partner service virtual machines (SVM) in Cluster 1.

In this approach, security protection downtime is expected during migration of Cluster 1.

When workload VMs migrate to an NSX-T host, existing data traffic during a host migration is expected to have a security protection downtime. However, new data traffic does not have a security protection downtime.

Approach 2: Involves minimal security protection downtime

This approach requires some manual intervention with an NSX-T API to create a temporary host group. Enable the **Pause between groups** migration setting and migrate any one host from Cluster 1. After this host in Cluster 1 is migrated to NSX-T, the migration coordinator pauses. Deploy a partner service on this migrated host by doing a clustered service deployment. Continue migrating the remaining hosts in Cluster 1. After all the hosts in Cluster 1 are migrated to NSX-T, you can optionally deploy additional partner SVMs in Cluster 1 by doing either a host-based or a clustered service deployment.

The detailed procedure in this topic explains the host migration workflow for a single NSX-V prepared cluster, which has three hosts, as shown in the following figure. The procedure uses Approach 2 to migrate this Cluster 1 to NSX-T.

Example:

Figure 6-1. Host Group 1 (Cluster 1) Before Migration

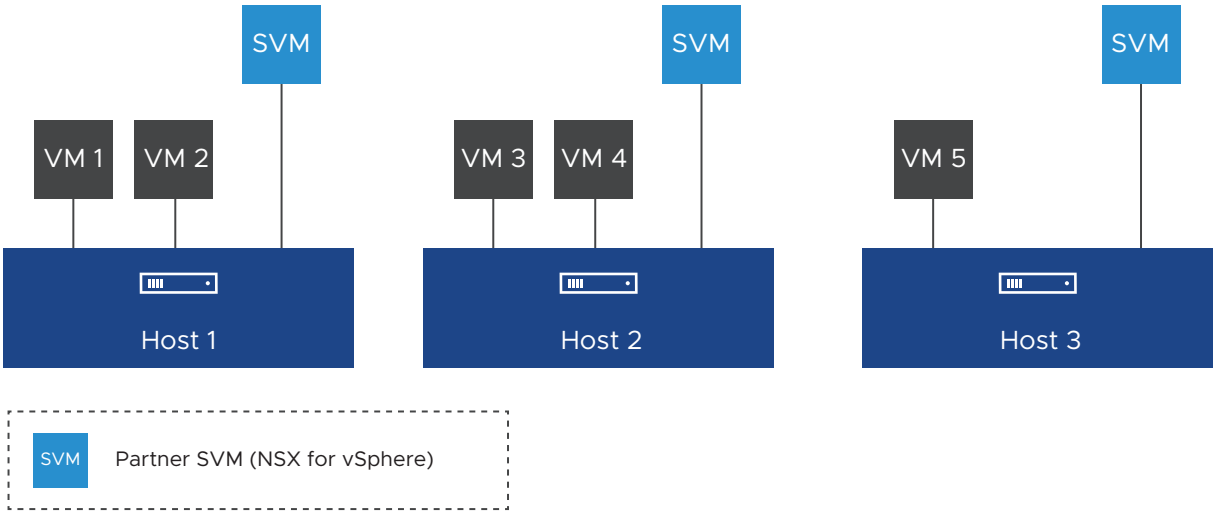

All hosts in this cluster are ESXi hosts. The Security Policies in your NSX-V environment redirect data traffic to partner service virtual appliances that provide a network introspection service to workloads. As NSX-V supports only a host-based service deployment, each host has a single partner service VM.

The following configuration settings are required for migrating hosts using Approach 2:

- Host migration mode is set to Automated Maintenance.
- Pause between groups is enabled.
- Migration order across groups is set to serial.

Prerequisites

- Verify that Edge migration has finished and all routing and services are working correctly.

- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.
- Enable vSphere DRS on the cluster that is being migrated.
- Enable vMotion on the VMkernel adapter of each host in the cluster.
- Ensure that adequate spare capacity is available in the NSX-V cluster so that the migrating hosts can enter into a maintenance mode. If enough spare capacity is unavailable to migrate NSX-V workload VMs to other hosts in the cluster, additional security protection downtime is expected.

Procedure

1 Run the following API request to create a temporary host group and move hosts 2 and 3 to this temporary group.

```
POST https://{nsxt-mgr-ip}/api/v1/migration/migration-unit-groups
```

In the request body of this POST API, specify the following details:

- Name of the temporary host group
- Migration units (IDs of hosts 2 and 3)
- Migration state of the temporary group (must be disabled)

For a detailed information about this API and an example POST API request, see the *NSX-T Data Center API Guide*.

You can obtain the host IDs from the vCenter Server Managed Object Browser (MOB) at `http://{vCenter-IP-Address}/mob`, or run the following GET API to retrieve the host IDs:

```
GET https://{nsxt-mgr-ip}/api/v1/fabric/discovered-nodes
```

A temporary host group is created and displayed on the **Migrate Hosts** page. The original host group 1 (cluster 1) now contains only host 1.

2 On the **Migrate Hosts** page, next to **Host Migration Plan**, click **Settings** and ensure that the settings are configured as follows:

- Pause between groups: Enabled
- Migration order across groups: Serial

3 Migrate host 1 to NSX-T.

a Click **Start** to start the host migration.

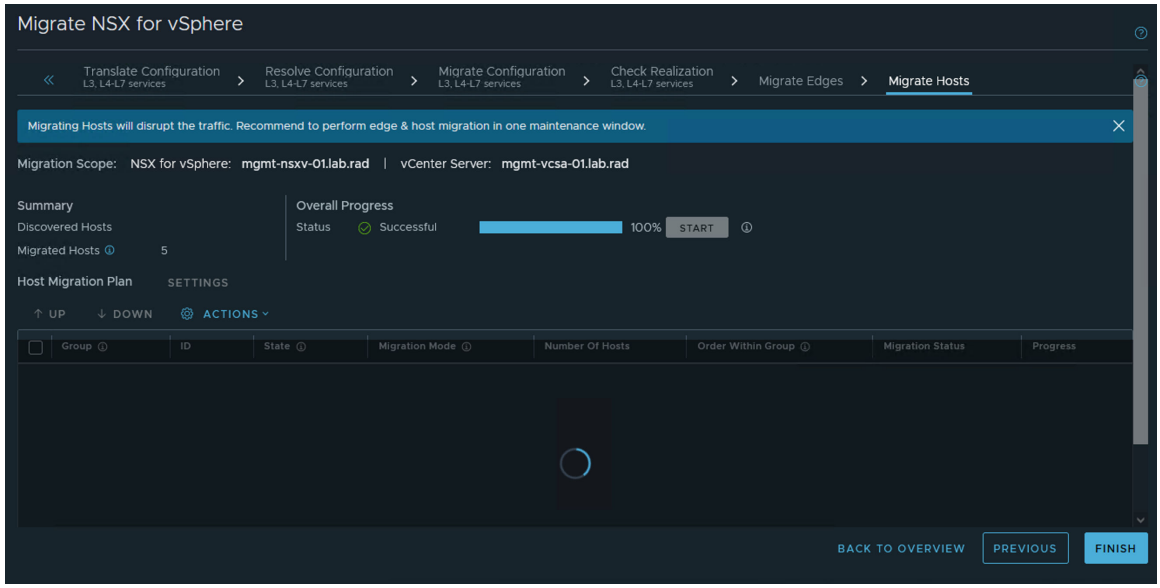

Workload VMs 1 and 2 are migrated to other hosts so that host 1 can enter into a maintenance mode. NSX-V partner SVM on host 1 is powered off before host 1 enters into a maintenance mode.

Assume that VMs 1 and 2 are migrated to host 3 that is prepared with NSX-V. After the migration of host 1 is successful, the migration coordinator pauses for your next input.

b (Required) Deploy a partner service on host 1 by using the clustered deployment approach.

At this stage, host-based service deployment is not supported. Deploying a partner service on host 1 is necessary to minimize security protection downtime. Remember, the security protection for NSX-V workloads that are running on hosts 2 and 3 is still intact. The partner must ensure that the migrated partner-specific Security Policies are available on the newly deployed partner SVMs on host 1.

For detailed steps on deploying a partner service in NSX-T, see Deploy a Partner Service for Network Introspection. For example, specify the following configuration settings to deploy two partner service virtual machines (SVMs) on host 1:

Configuration	Value
Deployment Type	Clustered
Host	Host 1
Clustered Deployment Count	2

The value that you enter in the **Clustered Deployment Count** text box depends on the resource capacity that is available on the host. This scenario assumes that two partner SVMs can be deployed on Host 1. This value can be different in your environment.

After this step, the cluster looks as shown in the following figure. The green colored host represents the migrated host.

Figure 6-2. Host 1 is Migrated to NSX-T

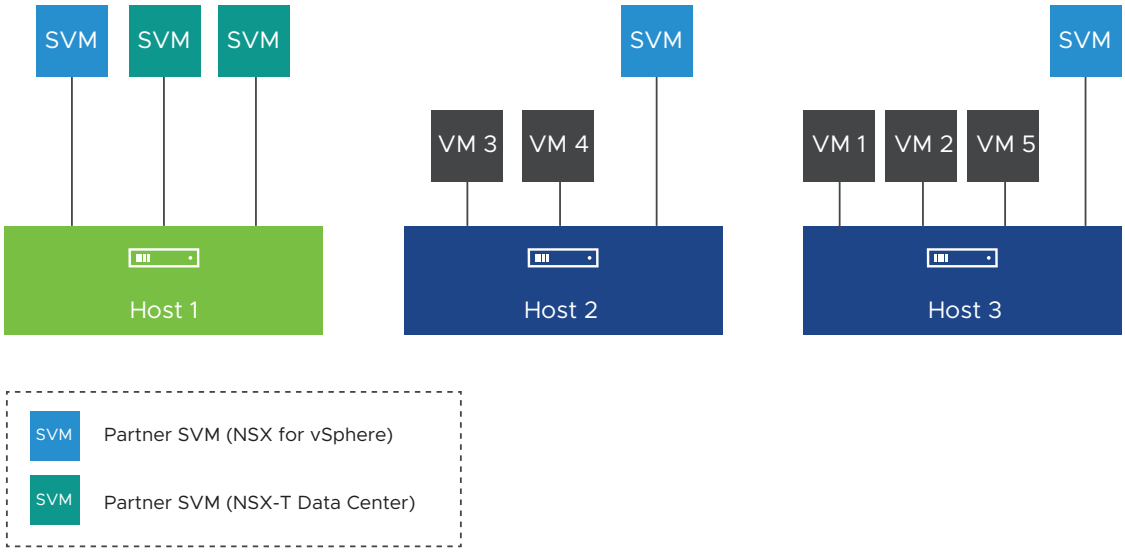

4 Migrate host 2 and host 3 to NSX-T.

a Move host 2 and host 3 from the temporary host group to the original host group 1 by running the following POST API request:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/migration-unit-groups/
{group-id}?action=add_migration_units
```

Where: *group-id* is the ID of the destination host group (host group 1). In the POST API request body, specify the ID of hosts 2 and 3 that you want to add to the original host group 1.

For a detailed information about this POST API and an example POST API request, see the *NSX-T Data Center API Guide*.

Now, the original host group 1 contains hosts 1, 2 and 3 (in the given order), and the temporary host group is deleted.

b Select the check box next to the original host group 1, and then click **Actions > Change Migration Order Within Group**. Verify that the migration order within the group is set to **Serial**.

If necessary, you can set the migration order in the original host group 1 to **Parallel**.

c Click **Continue** to resume the host migration.

Host 2 is first migrated to NSX-T, and then host 3 is migrated. To put each migrating host into a maintenance mode, workload VMs on the migrating host are moved to either NSX-V hosts or NSX-T hosts. NSX-V partner SVM on the migrating host is also powered off before the host enters into a maintenance mode.

After this step, all the hosts in host group 1 (cluster 1) are prepared with NSX-T. The cluster looks as shown in the following figure.

Figure 6-3. All Hosts are Migrated to NSX-T

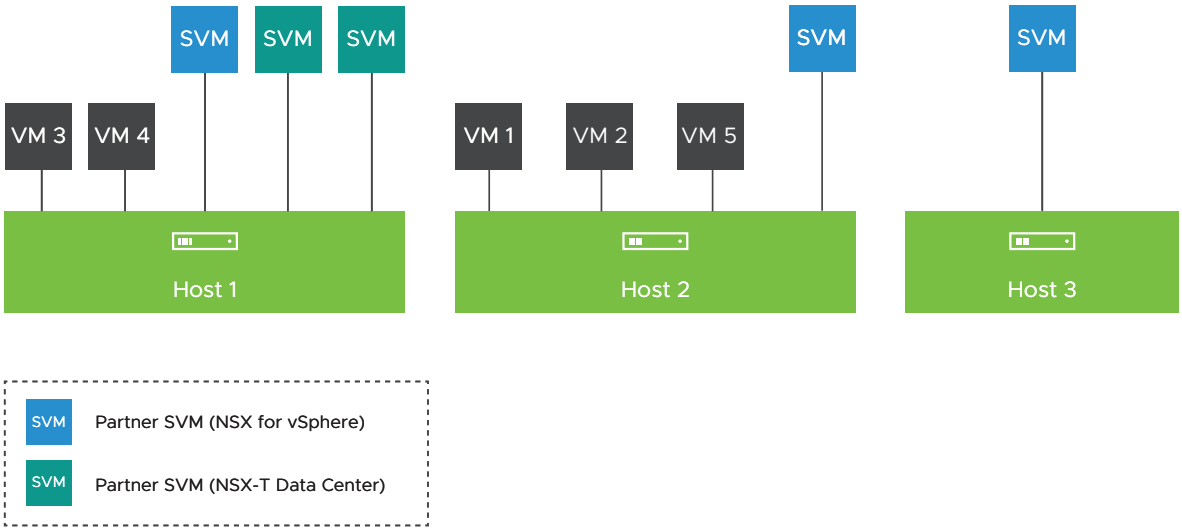

5 (Optional) Migrate some workload VMs from hosts 1 and 2 to host 3.

For example, migrate VMs 4 and 5 to host 3, as shown in the following figure.

Figure 6-4. Final Cluster 1 After Migration

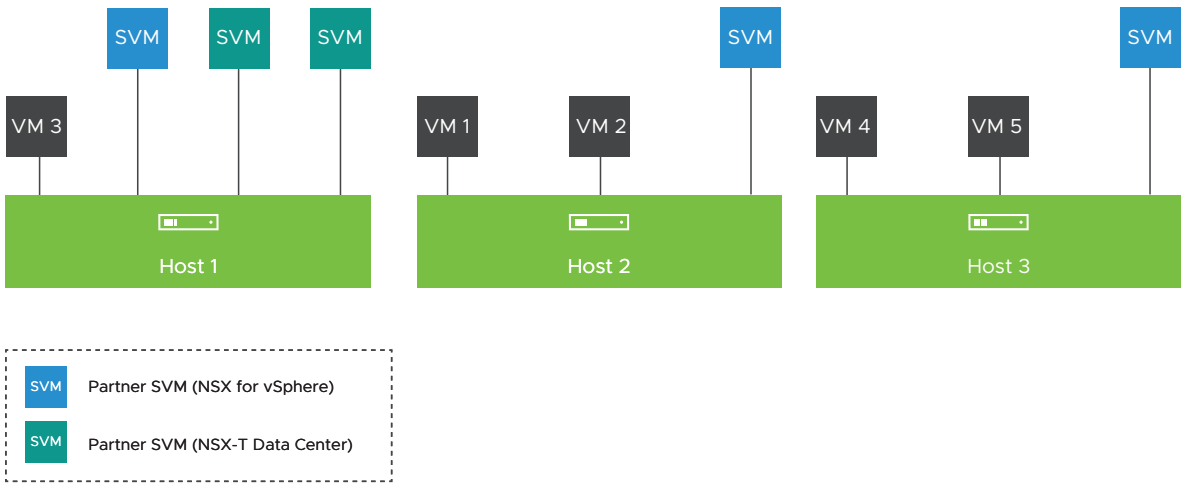

6 (Optional) After all the hosts in host group 1 are migrated to NSX-T, you can do either a host-based or a clustered service deployment.

A host-based service deployment allows new network traffic to be protected by a local partner SVM on each host.

Note If you have network introspection service running on more than one NSX-V prepared cluster, you do not have to deploy the partner SVMs in the other clusters. The network traffic though the NSX-T workload VMs in the other clusters can use the partner SVMs in cluster 1 that you just migrated. The host migration workflow covered in this procedure is required only for the first cluster. You can migrate the remaining clusters by using the standard host migration procedure.

What to do next

Delete the partner service deployment in NSX-V. Remember, you can delete the partner SVMs only at a cluster level. That is, you can delete service deployment only after all the hosts in the host group 1 are migrated to NSX-T. Complete the following steps to delete the service deployment in NSX-V:

1 Log in to the vSphere Client and navigate to **Networking and Security > Installation and Upgrade > Service Deployment**.

2 Select the deployed partner service, and click **Delete**.

Finish the End-to-End Migration

After you have migrated all Edge Services Gateway VMs and hosts to the NSX-T environment, confirm that the new environment is working correctly. If everything is functioning correctly, you can finish the migration.

Important Verify that everything is working before clicking **Finish**. Then perform the post-migration tasks. Do not make any vSphere life cycle operations such as upgrading ESXi hosts, VDS, or VC before the post-migration tasks are completed.

You will see errors on hosts after the migration. The error message is: `UserVars.RmqHostId' is invalid or exceeds the maximum number of characters permitted`. The error occurs because this host is still part of the NSX-V inventory.

Prerequisites

- Verify that all expected items have been migrated to the NSX-Tenvironment.
- Verify that the NSX-T environment is working correctly.

Procedure

1 Navigate to the **Migrate Hosts** page of the migration coordinator.

2 Click **Finish**

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page, or which hosts were excluded from the migration.

Post-Migration Tasks in End-to-End Migration

After the end-to-end migration has completed, some additional actions might be required.

- If you migrated from NSX-V 6.4.4, perform a reboot of all hosts that have migrated to NSX-T. The reboot must be done before you upgrade to a later version of NSX-T.
- During migration, all transport nodes are added to a group called `NSGroup with TransportNode for CPU Mem Threshold`. This group ensures that the transport nodes have the correct CPU memory threshold settings in NSX-T. This group is required after migration has completed. If you need to remove a transport node from NSX-T after migration and you are running NSX-T 3.2.0, you must first remove the transport node from this group. If you are running NSX-T 3.2.1 or later, you do not need to remove the transport node from this group.

  To remove the transport node from the group, make sure you are in **Manager** mode and then select **Inventory > Groups** to remove the transport node from the `NSGroup with TransportNode for CPU Mem Threshold` group. For more information about Manager mode, see the topic "NSX Manager" in the *NSX-T Data Center Administration Guide*.
- Verify that you have a valid backup and restore configuration. See "Backing Up and Restoring the NSX Manager" in the *NSX-T Data Center Administration Guide*.

Deploy a Partner Service for Endpoint Protection

When migrated Security Policies in NSX-T use a partner service that provides only Endpoint Protection or both Endpoint Protection and Network Introspection, deploy an instance of the partner service after all the clusters are migrated to NSX-T.

Only a host-based service deployment is supported.

In a host-based service deployment, one partner service virtual machine is installed on each host of the migrated cluster. In the vCenter Server, the vSphere ESX Agency Manager (EAM) service is internally used to deploy a partner service VM on each host of the cluster.

Prerequisites

- All the hosts in the cluster are migrated to NSX-T.
- All the migrated hosts are managed by a vCenter Server.
- A transport node profile is applied to the cluster.

Procedure

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

2 Navigate to **System > Service Deployments > Deployment**.

3 In the **Partner Service** drop-down menu, select the partner service to be deployed, and click **Deploy Service**.

4 Enter the service deployment name.

5 Select the vCenter Server that is registered as a compute manager in NSX-T.

6 Select the cluster where you want to deploy the partner service.

7 To specify the datastore, do one of the following actions:

- Select a datastore as the repository for the service virtual machines.
- Select **Specified on Host**.

  The **Specified on Host** option means that you do not need to select a datastore and network on the **Deploy Service** page. Before deploying the partner service, you must configure Agent VM settings on each ESXi host to point to a specific datastore and network.

  To know more about configuring Agent VM settings, see the vSphere product documentation.

8 Under Networks, click **Set** and select the NICs you want to use for deployment.

a Select the network for the Management interface.

  In a host-based deployment, if you set the datastore as **Specified on Host**, you must set the network also as **Specified on Host**.

b Set the Network type to **DHCP** or **Static IP Pool**. If you set the network type to a Static IP Pool, select from the list of available IP pools.

9 In the **Deployment Template** drop-down menu, select the registered deployment template and click **Save**.

  The deployment process might take some time depending on the vendor's implementation.

10 Check the deployment status on the **Deployment** page. Wait until the status changes to Up.

  You might have to refresh the **Deployment** page a few times to retrieve the latest status.

  If the Status column shows Down, click the icon next to Down. All deployment errors are displayed. Take the required actions to fix the errors, and click **Resolve**. The status changes to In Progress. Wait until the status changes to Up.

Results

A partner service VM is now deployed on all the hosts of the cluster.

Note When you add a new host in the cluster, EAM automatically deploys the partner service VM on the new host.

What to do next

Go to the Partner Console and verify whether the endpoint protection service is activated. Now, the migrated endpoint protection rules are enforced on the workload VMs that are running on the NSX-T prepared cluster.

For more information about activating the endpoint protection service in the Partner Console, see the partner documentation.

Deploy a Partner Service for Network Introspection

When migrated Security Policies in NSX-T use a third-party partner service only for Network Introspection, deploy an instance of the partner service either by using a clustered service deployment or a host-based service deployment approach.

Prerequisites

For a clustered service deployment approach:

- At least one host in the first cluster is migrated to NSX-T.

For a host-based service deployment approach:

- All the hosts in a cluster are migrated to NSX-T.
- A transport node profile is applied to the cluster.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Navigate to **System > Service Deployments > Deployment**.
3. In the **Partner Service** drop-down menu, select the partner service to be deployed, and click **Deploy Service**.
4. Enter the service deployment name.
5. Select the vCenter Server that is registered as a compute manager in NSX-T.
6. Select a deployment type: **Host-Based** or **Clustered**.
7. Select the cluster where you want to deploy the partner service.
8. (Clustered deployment only): In the **Host** drop-down menu, select a host, or select **Any** to allow the NSX-T NSX Manager to select a host.

9 In the **Data Store** drop-down menu, select a data store as the repository for the partner service virtual machine (SVM).

- Clustered deployment: If you selected **Any** for the host, select a shared data store. If you specified a particular host, select a local data store.
- Host-based deployment: Select a specific datastore or select **Specified on Host**. The **Specified on Host** option means that you do not need to select a datastore and network on the **Deploy Service** page. Before deploying the partner service, you must configure Agent VM settings on each ESXi host to point to a specific datastore and network.

  To know more about configuring Agent VM settings, see the vSphere product documentation.

10 Under Networks, click **Set** and select the NICs you want to use for deployment.

a Select the network for the Management interface.

  In a host-based deployment, if you set the datastore as **Specified on Host**, you must set the network also as **Specified on Host**.

b Set the Network type to **DHCP** or **Static IP Pool**. If you set the network type to a Static IP Pool, select from the list of available IP pools.

11 In the **Deployment Template** drop-down menu, select the required template.

  Typically, the deployment specification and the deployment template fields are automatically selected with the information that is pushed from the Partner Console as part of the service definition.

12 In the **Service Segment** drop-down menu, select the service segment that the migration coordinator has created in the overlay transport zone.

13 (Clustered deployment only): In the **Clustered Deployment Count** text box, specify the number of service VMs to deploy in the cluster, and click **Save**.

14 Check the deployment status on the **Deployment** page. Wait until the status changes to Up.

  You might have to refresh the **Deployment** page a few times to retrieve the latest status.

  If the Status column shows Down, click the icon next to Down. All deployment errors are displayed. Take the required actions to fix the errors, and click **Resolve**. The status changes to In Progress. Wait until the status changes to Up.

Uninstalling NSX-V After Migration

When you have verified that the migration is successful, and have clicked **Finish** to finish the migration, you can uninstall your NSX-V environment.

The process for uninstalling NSX-V after migration to NSX-T is different from the standard uninstall for NSX-V.

Important If you have vCenter Enhanced Linked Mode (ELM) configured, you must migrate all the NSX-V instances associated with the vCenter ELM chain before executing steps 6, 7, and 8 in the procedure below.

Prerequisites

- Verify that the migration is successful, and all functionality is working in the NSX-T environment.
- Verify that you have clicked **Finish** on the **Migrate Hosts** page.

Procedure

1. In the vSphere client, navigate to **Networking and Security > NSX Edges** and delete all the NSX Edges.
2. In the vSphere client, navigate to **Networking and Security > Logical Switches** and delete all the logical switches.
3. In the vSphere client, navigate to **Networking and Security > Installation and Upgrade > Logical Network Settings > Transport Zones** and delete all the transport zones.
4. In the vSphere client, navigate to **Networking and Security > Installation and Upgrade > Management > NSX Controller Nodes** and delete all the NSX Controllers.
5. Clear all stale VTEPs that may remain in the NSX-V Manager database:
   a. SSH into NSX-V Manager as **root**.
   b. Run the following command to clear the database table:

   ```
   psql -U secureall -d secureall -c "delete from xvs_vmknic_info;"
   ```

   c. Run the following command to confirm that the output shows zero row:

   ```
   psql -U secureall -d secureall -c "select * from xvs_vmknic_info;"
   ```

6. Delete the ESX Agent Manager agencies that are associated with the NSX-V environment.
   a. In the vSphere Client, navigate to **Menu > Administration**. Under **Solutions**, click **vCenter Server Extensions**. Double-click **vSphere ESX Agent Manager** and click the **Configure** tab.
   b. For each agency that has a name starting with `_NSX_`, select the agency, then click the three-dot menu icon **(**  **)** and select **Delete Agency**.

7 Remove the NSX-V plug-in from vCenter Server.

a Access the Extension Manager from the Managed Object Browser at `https://<vcenter-ip>/mob/?moid=ExtensionManager`.

b Click **UnregisterExtension**.

c In the **UnregisterExtension** dialog box, enter **`com.vmware.vShieldManager`** in the **Value** text box and click **Invoke Method**.

d In the **UnregisterExtension** dialog box, enter **`com.vmware.nsx.ui.h5`** in the **Value** text box and click **Invoke Method**.

e You can verify that you unregistered the extensions by going to the Extension Manager page at `https://<vcenter-ip>/mob/?moid=ExtensionManager` and viewing the values for the **extensionList** property.

8 Delete the vSphere Web Client directories and vSphere Client (HTML5) directories for NSX-V and then restart the client services.

a Connect to the vCenter Server system command line.

- If you are using a vCenter Server Appliance, log in as root using the console or SSH. You must log in as root and run the commands from the Bash shell. You can start the Bash shell using the following commands.

```
> shell.set --enabled True
> shell
```

- If you are using vCenter Server for Windows, log in as an administrator using the console or RDP.

b Delete all NSX-V plug-in directories.

Note A plug-in directory might not be present if you have never launched the associated client.

On vCenter Server Appliance, delete the following directories:

- To remove the vSphere Web Client plug-in, delete the `/etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/com.vmware.vShieldManager-<version>-<build>` directory.
- To remove the vSphere Client plug-in, delete the `/etc/vmware/vsphere-ui/vc-packages/vsphere-client-serenity/com.vmware.nsx.ui.h5-<version>-<build>` directory.

On vCenter Server for Windows, delete the following directories:

- To remove the vSphere Web Client plug-in, delete the `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity\com.vmware.vShieldManager-<version>-<build>` directory.
- To remove the vSphere Client plug-in, delete the `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-ui\vc-packages\vsphere-client-serenity\com.vmware.nsx.ui.h5-<version>-<build>` directory.

c Restart the client services on the vCenter Server Appliance or vCenter Server on Windows.

Table 6-13. Client Service Commands

Client Service	vCenter Server Appliance	vCenter Server for Windows
Restart vSphere Web Client On vSphere 6.0, 6.5, and 6.7	`> shell.set --enabled True`	`> cd C:\Program Files\VMware\vCenter`

Table 6-13. Client Service Commands (continued)

Client Service	vCenter Server Appliance	vCenter Server for Windows

	`> shell`
	`# service-control --stop vsphere-client`
	`# service-control --start vsphere-client`
		`Server\bin`
		`> service-control --stop vspherewebclientsvc`
		`> service-control --start vspherewebclientsvc`

Restart vSphere Client
On vSphere 6.5, and 6.7
	`> shell.set --enabled True`
	`> shell`
	`# service-control --stop vsphere-ui`
	`# service-control --start vsphere-ui`
		`> cd C:\Program Files\VMware\vCenter Server\bin`
		`> service-control --stop vsphere-ui`
		`> service-control --start vsphere-ui`

Restart vSphere Client
On vSphere 7.0
	`> shell.set --enabled True`
	`> shell`
	`# service-control --stop vsphere-ui`
	`# service-control --start vsphere-ui`
		vSphere 7.0 does not support vCenter Server for Windows

9 Power off and delete the NSX Manager VM.

a In vSphere client, navigate to **Hosts and Clusters**.

b Locate the NSX Manager VM. Right click and select **Power Off** then right click and select **Delete from Disk**.

Prepare Infrastructure in Lift-and-Shift Migration

This step prepares the infrastructure in the NSX-T environment for the **Configuration Migration** mode when you migrate a user-defined topology..

Prerequisites

If the migration includes migrating NSX-V load balancer to Advanced Load Balancer (ALB), perform the following tasks in preparation for a workload migration:

- Edit the overlay segment that is part of the layer-2 bridge that you created. Set the **Admin State** toggle to Down (gray instead of green), and set the **Connectivity** toggle to Off.
- Make sure that route advertisement on the tier-1 gateways is enabled.

Procedure

1 Click **Start** and wait for the process to complete.

2 If the status is **Successful**, click **Continue** to go to the next step.

3 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

4 Remove the `v2t-arp-proxy-services` tag from the mapped tier-0/tier-1 router.

5 Enable `arp_proxy`.

Migrate Edges

This step creates a bridge between the NSX-V logical switches and the NSX-T Data Center segments and migrates Edge nodes for north-south traffic cutover.

This step is part of the workflow when you choose the migration mode **Configuration and Edge Migration** (available in NSX-T Data Center 3.2.2 as a tech preview feature).

Procedure

1 Click **Start**.

2 When the migration of the Edge nodes is completed, click **Continue** to go to the next step.

Migrate Workloads in Lift-and-Shift Migration

This is the last step in a lift-and-shift migration. You migrate the workload VMs from NSX-V to NSX-T.

If the migration mode is **Configuration and Edge Migration** (available as a tech preview feature in NSX-T Data Center 3.2.2), the following methods are available to migrate workload VMs:

- HCX (for more information, see the HCX documentation)
- vMotion

If you are using the vMotion method, depending on your NSX-V environment, do one of the following to migrate workload VMs:

- Migrate Workload VMs (Simple Case)
- Migrate Workload VMs (Complex Case)

If the migration mode is **Configuration Migration**, perform the steps in Switch the Default Gateway to NSX-T and, depending on your NSX-V environment, do one of the following to migrate workload VMs:

- Migrate Workload VMs (Simple Case)
- Migrate Workload VMs (Complex Case)

Switch the Default Gateway to NSX-T

The workload VMs on the NSX-V bridged Logical Switch are currently using the Distributed Logical Router (DLR) as their default gateway for all the north-south traffic.

When you are ready to migrate workload VMs to the bridged overlay segment, you must make the following configuration changes:

- Switch the default gateway to NSX-T. In other words, the migrated VMs must connect to the tier-1 gateway as their default gateway for all north-south traffic. If your NSX-T environment has a single tier routing topology, you can switch to the tier-0 gateway.
- If you have pre-configured Layer 3 network services on the tier-1 or tier-0 gateway, and dynamic route peering between tier-0 gateway and north-facing physical routers, do the following configurations:
  - On the tier-1 gateway, turn on **Route Advertisement** and Layer 3 services.
  - On the tier-0 gateway, turn on **Route Re-distribution Status**.

When the bridged overlay segment is connected to the NSX-T gateway, a GARP (gratuitous ARP) message will be sent and all connected VMs (including the NSX-V VMs) can update their ARP table accordingly.

You can make these configurations changes either manually or automate them by running APIs in a script file. Automation can help you to minimize the data traffic outage.

The following procedure explains the manual method of switching the default gateway to the tier-1 or tier-0 gateway in NSX-T by using the UI when Layer 3 services are not configured. For a minimum data traffic outage, you can automate the switching process with APIs.

Prerequisites

Perform the procedure in Change the MAC Address of NSX-T Virtual Distributed Router so that the NSX-V VMs can reach the default gateway in NSX-T.

Procedure

1. In the NSX-V environment, disconnect the bridged Logical Switch from the DLR.
   a. In the vSphere Client, navigate to the NSX Edge (DLR).
   b. Click **Configure > Interfaces**.
   c. Select the internal interface on the DLR that is connected to the bridged Logical Switch and click **Disconnect**.
2. In the NSX-T environment, connect the bridged overlay segment to the tier-1 gateway.
   a. In NSX Manager, navigate to **Networking > Segments**.
   b. Next to the name of the bridged overlay segment, click the vertical ellipses, and then click **Edit**.
   c. Turn on the **Connectivity** of the overlay segment, and click **Save**.

Migrate Workload VMs (Simple Case)

Use this procedure if "Applied To" is not configured in any of the DFW rules (this means that "Applied To" is set to "DFW").

If "Applied To" is configured in any of the DFW rules, do not use this procedure. Follow the procedure in Migrate Workload VMs (Complex Case) instead.

Note For NSX-V to NSX-T migration, see the KB article https://kb.vmware.com/s/article/56991 for more information.

For NSX-T to NSX-V migration, migrating a workload VM back to NSX-V might not work because the distributed firewall filter in NSX-T is always higher than in NSX-V. The workaround is to place the workload VM in the NSX-T exclusion list prior to vMotion.

Prerequisites

- Ensure that:
  - vSphere vMotion is enabled on the VMkernel adapter of each host in the cluster that is involved in this migration. For detailed steps about enabling vMotion on the VMkernel adapter, see the *vSphere* product documentation.
  - The destination host in NSX-T has sufficient resources to receive the migrated VMs.
  - The source and destination hosts are in an operational state. Resolve any problems with hosts including disconnected states.

For more information about vMotion, see Migration with vMotion in the *vSphere* product documentation.

Procedure

1 Start migrating the workload VMs using vMotion in the vSphere Client. See Migrating Virtual Machines in the *vSphere* product documentation for detailed instructions.

   Note During vMotion from NSX-V to NSX-T, the workload VMs are always protected because the migration coordinator translates the existing NSX-V DFW rules and security groups into temporary IP-based rules and groups.

2 Finalize the infrastructure to finish the migration.

   ```
   POST https://{nsxt-mgr-ip}/api/v1/migration?action=finalize_infra
   ```

   This migration API deletes any temporary object configurations that were created during the migration, and ensures that the NSX-T infrastructure is in a clean state. For example, temporary IP Sets are removed from the Groups.

   This POST API does not have a request body.

3 Verify that the expected configuration items have been migrated to the NSX-T environment.

   For example, check whether the following configurations are migrated successfully:

   - User-defined Distributed Firewall rules.
   - All Grouping objects, such as IP Sets, Groups, Tags, and so on.
   - Effective members are displayed in the dynamic Groups.

- Tags are applied to migrated workload VMs.

4 On the **Migrate Workloads** page, click **Finish**.

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page.

What to do next

After the migration of workload VMs and the DFW-only configuration is successful and thoroughly verified, remove the Layer 2 bridge to release the NSX-T Edge that you used for bridging.

Migrate Workload VMs (Complex Case)

Use this procedure if "Applied To" is configured in any of the DFW rules (this means that "Applied To" is not set to "DFW").

Note For NSX-V to NSX-T migration, see the KB article https://kb.vmware.com/s/article/56991 for more information.

For NSX-T to NSX-V migration, migrating a workload VM back to NSX-V might not work because the distributed firewall filter in NSX-T is always higher than in NSX-V. The workaround is to place the workload VM in the NSX-T exclusion list prior to vMotion.

Prerequisites

- Ensure that:
  - vSphere vMotion is enabled on the VMkernel adapter of each host in the cluster that is involved in this migration. For detailed steps about enabling vMotion on the VMkernel adapter, see the *vSphere* product documentation.
  - The destination host in NSX-T has sufficient resources to receive the migrated VMs.
  - The source and destination hosts are in an operational state. Resolve any problems with hosts including disconnected states.

For more information about vMotion, see Migration with vMotion in the *vSphere* product documentation.

Procedure

1 Get the instance UUID of all the VMs that you plan to migrate.

The instance UUIDs are needed when you make the API call in the next step. See the example at the bottom of this section on how to obtain the instance UUID of a VM.

2 Run the following POST API request:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/vmgroup?action=pre_migrate
```

This API creates a logical segment port (VIF) corresponding to the VM instance UUID of each NSX-V workload VM in the VM group that you will be migrating through the Layer 2 bridge to the NSX-T overlay segment. For an example request body of this API, see the Lift and Shift Migration Process section of the NSX Tech Zone article.

3 Call the API `GetVmGroupExecutionDetails`. This API is available starting with NSX-T 3.2.2

Call the API `GetVmGroupExecutionDetails` to get the result of the pre-migrate API call with the same group_id (and federation_site_id for cross-VC migration). The result includes a "logical_switch_id_to_vm_instance_id_and_vnics_map" list and an optional "failedVmInstanceIds" list, which includes the UUIDs of VMs that are not found in the source VC. For example:

```
GET /api/v1/migration/vmgroup/actions/get_vm_group_execution_details?group_id=<group-
id>&federation_site_id=<site_id>
Response:
{
  "logical_switch_id_to_vm_instance_id_and_vnics_map":[
    {
      "ls_id":"36885723-7581-4696-a195-ef83851dc35f",
      "vm_and_vnics_mapping":[
        {
          "vm_instance_id":"52199e21-6aab-26e4-8c82-069a17d67667",
          "vnics":[
            "4001"
          ]
        },
        {
          "vm_instance_id":"52630e5d-ce6f-fac0-424c-4aa4bdf6bd56",
          "vnics":[
            "4001"
          ]
        }
      ]
    }
  ],
  "failedVmInstanceIds":[
    "501557f6-2197-1fe8-14e5-89898cee5fec"
  ]
}
```

4 Follow the pseudo python code below to write a script to vmotion the VMs.

For an example, see the Python Example Scripts section of the NSX Tech Zone article.

```
    define _get_nsx_networks_in_host(self, host):
        ls_id_to_nsx_pgs_map = {}
        for net in host.network:
            if isinstance(net, vim.dvs.DistributedVirtualPortgroup):
                if hasattr(net.config, 'backingType'):
                    if net.config.backingType == 'nsx' and net.config.logicalSwitchUuid:
                        ls_id_to_nsx_pgs_map[net.config.logicalSwitchUuid] =\
                           [net.key, net.config.distributedVirtualSwitch.uuid]
```

```
            elif isinstance(net, vim.OpaqueNetwork):
                if net.summary.opaqueNetworkType == 'nsx.LogicalSwitch':
                    ls_id_to_nsx_pgs_map[net.summary.opaqueNetworkId] = [None,
net.summary.opaqueNetworkId]
        return ls_id_to_nsx_pgs_map

     define _get_vms_vnic_to_ls_id_map(self,
logical_switch_id_to_vm_instance_id_and_vnics_map):
        vm_uuid_2_vnics_map = {}
        for ls_id_2_vm_vnics in logical_switch_id_to_vm_instance_id_and_vnics_map:
            ls_id = ls_id_2_vm_vnics['ls_id']
            for vm_vnics in ls_id_2_vm_vnics['vm_and_vnics_mapping']:
                vnic_2_ls_id = vm_uuid_2_vnics_map.get(vm_vnics['vm_instance_id'], {})
                for vnic in vm_vnics['vnics']:
                    vnic_2_ls_id[vnic] = ls_id
                vm_uuid_2_vnics_map[vm_vnics['vm_instance_id']] = vnic_2_ls_id
        return vm_uuid_2_vnics_map

    def _get_nsxt_vnic_spec(self, device, dvpg_key, switch_id, vif_id):
        If dvpg_key:
            vdsPgConn = vim.dvs.PortConnection()
            vdsPgConn.portgroupKey = dvpg_key
            vdsPgConn.switchUuid = switch_id
            device.backing =
vim.vm.device.VirtualEthernetCard.DistributedVirtualPortBackingInfo()
            device.backing.port = vdsPgConn
        else:
            device.backing = vim.vm.device.VirtualEthernetCard.OpaqueNetworkBackingInfo()
            device.backing.opaqueNetworkId = switch_id
            device.backing.opaqueNetworkType = 'nsx.LogicalSwitch'
        device.externalId = vif_id
        dev_spec = vim.Vm.Device.VirtualDeviceSpec()
        dev_spec.SetOperation(vim.Vm.Device.VirtualDeviceSpec.Operation.edit)
        dev_spec.SetDevice(device)
        return dev_spec

    def _migrate_vm(self, vmObject, vnic_2_ls_id_map, ls_id_to_nsx_pgs_map):
        devices = vmObject.config.hardware.device
        nic_devices = [device for device in devices if isinstance(device,
vim.Vm.device.VirtualEthernetCard)]
        vnic_changes = []
        for device in nic_devices:
            ls_id = vnic_2_ls_id_map.get(str(device.key))
            if ls_id:
                vif_id = vmObject.config.instanceUuid + ":" + str(device.key)
                nsx_pg = ls_id_to_nsx_pgs_map.get(ls_id)
                vnic_spec = self._get_nsxt_vnic_spec(device, nsx_pg[0], nsx_pg[1], vif_id)
                vnic_changes.append(vnic_spec)
        relocate_spec = vim.Vm.RelocateSpec()
        relocate_spec.SetDeviceChange(vnic_changes)
        # set other fields in the relocate_spec
        vmotion_task = vmObject.Relocate(relocate_spec)
        WaitForTask(vmotion_task)
```

```
    vm_uuid_2_vnics_map =
self._get_vms_vnic_to_ls_id_map(logical_switch_id_to_vm_instance_id_and_vnics_map)
    for vm_uuid, vnic_2_ls_id_map in vm_uuid_2_vnics_map.items():
        # get the vmObject by the vm_uuid
        # find a target host that has all the networks needed by this VM
        ls_id_to_nsx_pgs_map = self._get_nsx_networks_in_host(host)
        self._migrate_vm(vmObject, vnic_2_ls_id_map, ls_id_to_nsx_pgs_map)
```

5 Apply the Security Tags and VM static membership to the migrated VMs.

`POST https://{nsxt-mgr-ip}/api/v1/migration/vmgroup?action=post_migrate`

The `vmgroup` API endpoint with `post_migrate` action applies the NSX-V Security Tags to the migrated workload VMs on the NSX-T overlay segment.

For an example request body of this API, see the Lift and Shift Migration Process section of the NSX Tech Zone article.

6 Finalize the infrastructure to finish the migration.

`POST https://{nsxt-mgr-ip}/api/v1/migration?action=finalize_infra`

This migration API deletes any temporary object configurations that were created during the migration, and ensures that the NSX-T infrastructure is in a clean state. For example, temporary IP Sets are removed from the Groups.

This POST API does not have a request body.

7 Verify that the expected configuration items have been migrated to the NSX-T environment.

For example, check whether the following configurations are migrated successfully:

- User-defined Distributed Firewall rules.
- All Grouping objects, such as IP Sets, Groups, Tags, and so on.
- Effective members are displayed in the dynamic Groups.
- Tags are applied to migrated workload VMs.

8 On the **Migrate Workloads** page, click **Finish**.

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page.

Example: Obtaining VM Instance UUID from the vCenter MOB

This example shows how to obtain or confirm a VM's instance UUID from the vCenter Server Managed Object Browser (MOB) at `http://{vCenter-IP-Address}/mob`. You can also obtain or confirm a VM's instance UUID by making an API call to vSphere.

1 In a web browser, enter the vCenter Managed Object Browser at `http//{vCenter-IP-Address}/mob`.

2 Click **content**.

3 Find **rootFolder** in the Name column, and click the corresponding link in the Value column. For example, group-d1.

4 Find **childEntity** in the Name column, and click the corresponding link in the Value column. For example, datacenter-21.

5 Find **hostFolder** in the Name column, and click the corresponding link in the Value column. For example, group-h23.

6 Find **childEntity** in the Name column. The corresponding Value column contains links to host clusters. Click the appropriate host cluster link. For example, domain-c33.

7 Find **host** in the Name column. The corresponding Value column lists the hosts in that cluster by vCenter MOID and hostname. Click the appropriate host link, For example, host-32.

8 Find **vm** in the Name column. The corresponding Value column lists the virtual machines by vCenter MOID and hostname. For example, vm-216 (web-01a). Click the VM that you are interested in.

9 Find **config** in the Name column. Click **config** in the Value column.

10 Find **instanceUuid** in the Name column. The corresponding Value column lists the VM instance UUID. For example, 502e71fa-1a00-759b-e40f-ce778e915f16.

What to do next

After the migration of workload VMs, you can remove the layer-2 bridge.

Migrating Distributed Firewall Configuration

7

In this migration mode, you migrate the Distributed Firewall (DFW) configuration from NSX-V to NSX-T. If you have Identity Firewall (IDFW) configured, it will also be migrated.

For more information about migrating IDFW, see Chapter 2 Migrating Identity Firewall (End-to-End and Lift-and-Shift).

The following logical object configurations are migrated:

- User-defined Distributed Firewall (DFW) rules
- Grouping Objects
  - IP Sets
  - MAC Sets
  - Security Groups
  - Services and Service Groups
  - Security Tags
- Security Policies created using Service Composer (only DFW rule configurations are migrated)

  Guest Introspection service configuration and Network Introspection rule configurations in the Service Composer are not migrated.

Depending on the DFW configuration, there are two ways to migrate the workload VMs after you migrate the DFW configuration. If "Applied To" is not configured in any of the DFW rules (this means that "Applied To" is set to "DFW"), you can use the vSphere Client to migrate the workload VMs (follow the procedure Migrate Workload VMs (Simple Case)). Otherwise, you must use a script to migrate the VMs (follow the procedure Migrate Workload VMs (Complex Case)).

Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with only local objects.

For a detailed list of all the configurations that are supported for the migration of Distributed Firewall configuration, see the Detailed Feature Support for Migration.

You have the following migration goals:

- Migrate only the existing Distributed Firewall configuration from NSX-V to NSX-T.
- Use the layer-2 Edge bridge and vSphere vMotion to migrate workload VMs from NSX-V to NSX-T.

To extend the layer-2 networks, you can use the NSX-T native Edge bridge.

Important When you use the lift and shift approach to migrate the DFW configuration from NSX-V to NSX-T, you must run the DFW-only migration mode only once. After the DFW configuration is migrated to NSX-T, you must not update the DFW configuration in your NSX-V environment and run the DFW-only migration mode again. Running the DFW-only migration mode multiple times is not recommended.

Prerequisites for DFW-Only Migration

- A new NSX-T environment is prepared for this migration.
- No user-defined DFW rules exist in NSX-T before this migration.
- All states in the **System Overview** pane of the NSX-V dashboard are green.
- There are no unpublished changes for Distributed Firewall and Service Composer policies in the NSX-V environment.

- All hosts in the NSX-managed cluster (NSX-V as well as NSX-T) must be connected to the same version of VDS and each host within the NSX-managed cluster must be a member of a single version of VDS.

Note

- The lift and shift migration of DFW-only configuration does not involve migrating hosts from NSX-V to NSX-T. Therefore, it is not mandatory for the ESXi version that is used in your NSX-V environment to be supported by NSX-T.
- In DFW-only migration mode, the firewall state (DVFilter) for existing connection sessions is maintained throughout the migration including vMotion. The firewall state is maintained regardless of whether the VMs are migrating within a single vCenter Server or across vCenter Servers. Also, the dynamic membership in the firewall rules is maintained after the Security Tags are migrated to the workload VMs.
- Objects that are created during the migration must not be updated or deleted before the migration is finished. However, you can create additional objects in NSX-T, if necessary.
- In NSX-T, DFW is enabled out of the box. All flows with sources and destinations as "any" in the DFW rules are allowed by default. When Distributed Firewall is enabled in the NSX-T environment, you cannot migrate the workload VMs again from NSX-T to NSX-V. Roll back of migrated workload VMs is not supported. The workaround is to add the workload VMs to the NSX-T Firewall Exclusion List, and then migrate the workload VMs back to NSX-V using vSphere vMotion.
- The automated migration of the DFW configurations supports workload VMs that are attached to NSX-V logical switches. These VMs will be migrated to NSX-T overlay segments. Workload VMs in NSX-V that are attached to vSphere Distributed Virtual Port Groups are not automatically migrated to NSX Distributed Virtual Port Groups. As a workaround, you must create the NSX Distributed Virtual Port Groups manually and attach the workload VMs to them.
- DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.
- Logical ports and switches that are created during migration are not deleted when the workload VMs are deleted. You must delete these ports and switches via the NSX Manager UI or the API.
- The default segments in NSX-T do no support DHCP servers and will result in the servers being down after migration.

Read the following topics next:

- Overview - Migrating Distributed Firewall Configuration
- Preparing for a DFW Configuration Migration
- Import the NSX-V Configuration
- Resolve Configuration

- Migrate the Distributed Firewall Configuration
- Switch the Default Gateway to NSX-T
- Migrate Workload VMs (Simple Case)
- Migrate Workload VMs (Complex Case)

Overview - Migrating Distributed Firewall Configuration

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.
- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.
- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 7-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.

Table 7-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.

NSX-V Configuration	Supported	Details
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. Migration Coordinator will only migrate from an NSX-V Manager with the role of Primary or Standalone. You can modify the NSX-V environment by changing the status of the secondary managers in order to migrate each NSX-V environment independently.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	
Network I/O Control (NIOC) having vNIC with reservation	No	
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	

NSX-V Configuration	Supported	Details
Host profiles	No	
ESXi lockdown mode	No	Not supported in NSX-T.
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.
VXLAN port number other than 4789	No	
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes for in-place migration No for lift-and-shift migration	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.

NSX-V Configuration	Supported	Details
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	

NSX-V Configuration	Supported	Details
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled

NSX-V Configuration	Supported	Details
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network

NSX-V Configuration	Supported	Details
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.

NSX-V Configuration	Supported	Details
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	

NSX-V Configuration	Supported	Details
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.

NSX-V Configuration	Supported	Details
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 7-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 7-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.

Table 7-3. DHCP Features (continued)

NSX-V Configuration	Supported	Details

DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration.
		For example, dhcp option '80' is not migrated.
		<pre><dhcpOptions>
		<other>
		<code>80</code>
		<value>2f766172</value>
		</other>
		</dhcpOptions></pre>

Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 7-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.

NSX-V Configuration	Supported	Details
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ VMs not on NSX-V hosts	No	NSX-T does not support referencing objects not connected to NSX-T portgroups or networks. Those objects will be lost from the source or destination when migration is completed. To avoid this issue use IP addresses in NSX-V to reference those objects before migration.
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group

NSX-V Configuration	Supported	Details
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.

NSX-V Configuration	Supported	Details
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	

NSX-V Configuration	Supported	Details
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 7-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of ‘Groups’ after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 7-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.

Table 7-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.

Table 7-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 7-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 7-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 7-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 7-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000

Table 7-10. Single-Site Limits (continued)

Feature	Limit
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 7-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Preparing for a DFW Configuration Migration

To prepare for a distributed firewall (DFW) configuration migration, configure the export version of DFW filter on hosts and create a layer-2 bridge between NSX-V and NSX-T. Also, tag the management VMs if you have a collapsed vCenter Server cluster.

For information on creating a layer-2 bridge between NSX-V and NSX-T, see Chapter 14 Preparing Layer-2 Bridging for Lift-and-Shift Migration.

Configure Export Version of Distributed Firewall Filter

The export version of a Distributed Firewall (DFW) filter is a property of a vNIC. Before you start some migrations, the export version of DFW filters must be set to 1000 for the vNICs of all the VMs that will be migrated.

You must make this configuration change in the following situations:

- You are doing a lift-and-shift migration.
- You are doing an in-place migration and you need to manually migrate VMs from some NSX-V hosts to NSX-T hosts. Follow the procedure below to change the export version for only those NSX-V hosts before migrating the VMs.

Procedure

1. Based on the VMs that will be migrated, determine the hosts that the VMs are running on.
2. Perform either step 3 or step 4 below.

3 For each host, perform the following steps to update, if necessary, the export version of DFW filters for all VM vNICs.

Note: In https://github.com/dixonly/samples, the script `updateDfwFilters.py` will print out and optionally update the DFW filter's export version for the vNICs of all the VMs in a specific cluster or all clusters. Using the script can save some time if you have a large number of VMs to migrate.

a Log into the command-line interface.

b Get the DFW filter names for all the VM vNICs. For example,

```
[root@esxi:~] vsipioctl getfilters | grep "Filter Name" | grep "sfw.2"
Filter Name: nic-2112467-eth0-vmware-sfw.2
Filter Name: nic-2112467-eth1-vmware-sfw.2
Filter Name: nic-2112467-eth2-vmware-sfw.2
```

c For each filter, get the export version. For example,

```
[root@esxi:~] vsipioctl getexportversion -f nic-2112467-eth0-vmware-sfw.2
Current export version: 500
```

d If the version is not 1000, set it to 1000. For example,

```
[root@esxi:~] vsipioctl setexportversion -f nic-2112467-eth0-vmware-sfw.2 -e 1000
```

e Verify that the export version is updated. For example,

```
[root@esxi:~] vsipioctl getexportversion  -f nic-2112467-eth0-vmware-sfw.2
Current export version: 1000
```

4 Based on the hosts that you noted in step 1, determine the clusters that contain the hosts. For each cluster, do the following:

From the vSphere Client, navigate to **Networking and Security > Installation and Upgrade > Host Preparation**. Select the cluster and click **Actions > Disable Firewall**. After the firewall is disabled, click **Actions > Enable Firewall**.

Tag Management VMs in a Collapsed Cluster Environment

You can migrate an environment that uses a collapsed cluster.

In a collapsed cluster design, all management VMs, workload VMs, and optionally edges run on the same vSphere cluster that is prepared for NSX-V. The management VMs of the NSX-T must be initially attached to dvPortgroups. After migration, the management VMs of NSX-T will be attached to NSX-T VLAN segments.

The management VMs in the NSX-T include appliances such as NSX Manager, vCenter Server, VMware Identity Manager, and so on. The NSX-T VLAN segment ports to which these management VMs connect are blocked in two cases: when these management VMs are rebooted after they were migrated in-place by the Migration Coordinator, or when they are moved from NSX-V hosts to NSX-T hosts by vMotion in maintenance migration mode. Therefore, the management VMs might lose connectivity in such cases.

To prevent this problem, create a "management_vms" tag category, and add tags in this category. Assign a tag from this category to all the management VMs in the NSX-T environment. These VMs will be attached to unblocked VLAN segment ports.

Procedure

1 Log in to the vSphere Client.

2 Click **Menu > Tags & Custom Attributes**.

3 Click **Categories**, and then click **New** to add a category.

Create a category with name **`management_vms`**.

4 Click the **Tags** tab and add a tag in the management_vms category.

5 Navigate to **Menu > Hosts and Clusters**.

6 Expand the collapsed cluster from the left Navigator view, right-click the name of the NSX Manager VM, and select **Tags & Custom Attributes > Assign Tag**.

7 Assign a tag from the management_vms category to the NSX Manager VM.

8 Repeat steps 6 and 7 for all the management VMs in the cluster.

For a detailed information about tag categories and tags, see the *vCenter Server and Host Management* documentation.

9 Log in to NSX Manager.

10 Navigate to **Inventory > Groups** and create a group.

11 Click **Set** to set members.

12 Set **Group Type** to **IP Addresses Only** and include the IP addresses of the management VMs that will be migrating from the dvPortgroups to the NSX-T segments.

13 Navigate to **Security > Distributed Firewall > Actions > Exclusion List** and add this group.

Import the NSX-V Configuration

The first step of the migration process is to import the NSX-V configuration.

Prerequisites

- The vCenter Server must be added as a compute manager in NSX-T.

  You can share the vCenter Server that is used in NSX-V or deploy another one in NSX-T.

Procedure

1 From a browser, log in to NSX Manager as **admin**.

2 Navigate to **System > Migrate**.

3 Expand the **Advanced Migration Modes** section, and in the **Distributed Firewall** pane, click **Get Started**.

4 From the **Import Configuration** page, click **Select NSX** and provide the credentials for vCenter Server and NSX-V.

Note The drop-down menu for vCenter Server displays all vCenter Server systems that are registered as compute managers. Click **Add New** if you need to add a compute manager.

5 Click **Start** to import the configuration.

6 When the import has finished, click **Continue** to proceed to the **Resolve Configuration** page.

Results

When the NSX-V topology is imported successfully, the **View Imported Topology** link is enabled. Click this link to view a graph of the imported topology. However, the topology viewer does not display the graph of a large-scale NSX-V environment.

Resolve Configuration

After you have imported the configuration from your NSX-V environment, you must review and resolve the reported configuration issues before you can continue with the migration.

On the **Resolve Configuration** page, two types of configuration issues are reported.

Blocking Issues (if any)

As the name suggests, these issues block the migration, and they must be fixed for the migration to proceed. You might have to change the configurations in your NSX-V environment before you can migrate to NSX-T.

Warnings

These configuration issues are organized into several categories, and each category can have one or more configuration items. You should provide inputs to fix the warning messages that are displayed for the configuration items, or choose to skip the warnings, if needed.

Procedure

1 On the **Resolve Configuration** page, if you find any blocking issues, fix them in the NSX-V environment before you can proceed with the migration.

After making the required changes in the NSX-V environment, return to the migration coordinator. Go to the **Import Configuration** page, and click **Rollback**. Click **Start** to import the updated NSX-V configuration.

If you did not find any blocking issues that require a change in the NSX-V environment, proceed to the next step.

2 Review the warnings and issues reported in each category.

3 Click each issue and provide input.

Each issue can cover multiple configuration items. For each item there might be one or more possible resolutions to the issue, for example, skip, configure, or select a specific value.

For issues that apply to multiple configuration items, you can provide input for each item individually, or select all and provide one answer for all items.

4 After the input is provided, a **Submit** button is displayed on the **Resolve Configuration** page. Click **Submit** to save your progress.

5 When you have provided input for all configuration issues, click **Submit**.

The input is validated. You are prompted to update any invalid input. Additional input might be required for some configuration items.

6 After you have submitted all requested input, click **Continue** to proceed to the **Migrate Configuration** page.

Example: Configuration Issues

Migration coordinator shows warning messages if it detects the following issues in your NSX-T environment:

- A missing NSX-T VLAN segment corresponding to the Distributed Virtual Port Group (DVPG) VLAN ID in your NSX-V environment.
- A missing NSX-T overlay segment corresponding to the VNI of the Logical Switch in the NSX-V environment.

To resolve this configuration issue, you can either skip creating the missing NSX-T segments or roll back the migration. If you choose to roll back the migration, create the missing segments in NSX-T by running the following PATCH API request, and start the migration again.

```
PATCH https://{policy-manager}/policy/api/v1/infra/segments/{segment-id}?
force=true
```

For example, to create a segment "App" with overlay ID "5001", the API URL and the payload of the API request is as follows:

```
PATCH https://{policy-manager}/policy/api/v1/infra/segments/App?force=true
```

```
{ "type" : "ROUTED",
  "subnets" : [ {” "gateway_address" : "172.16.20.1/24",
    "network" : "172.16.20.0/24""
    "} ],
  "connectivity_path" : "/infra/tier-0s/T0-GW-01",
  "transport_zone_path" : "/infra/sites/default/enforcement-points/default/transport-zones/
   1b3a2f36-bfd1-443e-a0f6-4de01abc963e",
  "overlay_id":5001,
```

```
  "admin_state" : "UP",
  "replication_mode" : "MTEP",
  "resource_type" : "Segment",
  "id" : "App",
  "display_name" : "App",
  "path" : "/infra/segments/App",
  "relative_path" : "App",
  "parent_path" : "/infra" }
```

If you choose to skip creating the missing NSX-T VLAN or overlay segments, the VM vNICs that are connected to the DVPG VLANs or the Logical Switches in the NSX-V environment lose networking after the workload VMs are migrated using vSphere vMotion. In this case, the migration coordinator does not migrate the DFW rules and Security Groups that are associated with the skipped segments.

For example, you might not want to create the missing segments in the following situations:

- Port groups on the vSphere Distributed Switch (VDS) have only VMkernel ports. So, no requirement to create VLAN segments.
- Port groups from another VDS that you do not want to migrate.
- NSX-V Logical Switch with no workload VMs attached to it. For instance, a transit Logical Switch between an Edge Services Gateway and the Distributed Logical Router.

Migrate the Distributed Firewall Configuration

After you have resolved all configuration issues, you can migrate the Distributed Firewall configuration. When the configuration is migrated, logical object configurations are realized in NSX-T environment, which replicate the NSX-V logical object configurations.

In the Prepare Infrastructure step, temporary IP sets will be added to NSX-V if the NSX-V security groups are used in a distributed firewall rule. This is required to maintain security while the VMs are migrated from NSX-V to NSX-T. After the migration, during the finalize infrastructure phase, the temporary IP sets will be deleted.

You can skip the Prepare Infrastructure step. However, doing so may compromise security until the finalize infrastructure phase is complete.

Prerequisites

Verify that you have completed the **Resolve Configuration** step.

Procedure

1 From the **Migrate Configuration** page, click **Start**.

2 Verify that the Distributed Firewall configuration objects are displayed in your NSX-T environment.

You can verify the migrated configurations either in the NSX-T NSX Manager interface or by running the NSX-T APIs.

Note

- During the **Migrate Configuration** step, Security Tags from NSX-V are not migrated to NSX-T. Therefore, the Security Tag-based migrated dynamic Groups in NSX-T are empty. The reason is that in NSX-V, a Security Tag is an object, whereas in NSX-T, a tag is an attribute of a VM. The tags are applied to the workload VMs only after you migrate the workloads to NSX-T and run the `vmgroup` API endpoint with a `post_migrate` action. For more information, see step 2 in Migrate Workload VMs (Complex Case).

  If the migrated NSX-T Groups have static memberships, these Groups also are empty after this step is finished. The reason is that the static members are not available in NSX-T Groups until the workload VMs are migrated.

  If only IP-based DFW rules are used in the NSX-V environment, you do not have to run the `vmgroup` API endpoint with `pre_migrate` and `post_migrate` action.

- When the logical configurations are migrated to NSX-T, the configuration changes are made in the NSX-T NSX Manager database, but it might take some time for the configurations to take effect.

3 Click **Continue** to proceed.

If needed, you can roll back the migrated DFW configuration.

Rolling back does the following:

- Remove the migrated configuration from NSX-T.
- Roll back all the resolved issues in the previous step.

Any NSX-T objects that you manually created after the DFW migration are at risk of being lost during the rollback.

4 In the **Prepare Infrastructure** step, click **Start** to prepare the infrastructure.

If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Results

After the prepare infrastructure step is completed, the next steps are:

- Switch the Default Gateway to NSX-T
- Depending on your environment, perform the step Migrate Workload VMs (Simple Case) or Migrate Workload VMs (Complex Case).

Switch the Default Gateway to NSX-T

The workload VMs on the NSX-V bridged Logical Switch are currently using the Distributed Logical Router (DLR) as their default gateway for all the north-south traffic.

When you are ready to migrate workload VMs to the bridged overlay segment, you must make the following configuration changes:

- Switch the default gateway to NSX-T. In other words, the migrated VMs must connect to the tier-1 gateway as their default gateway for all north-south traffic. If your NSX-T environment has a single tier routing topology, you can switch to the tier-0 gateway.
- If you have pre-configured Layer 3 network services on the tier-1 or tier-0 gateway, and dynamic route peering between tier-0 gateway and north-facing physical routers, do the following configurations:
  - On the tier-1 gateway, turn on **Route Advertisement** and Layer 3 services.
  - On the tier-0 gateway, turn on **Route Re-distribution Status**.

When the bridged overlay segment is connected to the NSX-T gateway, a GARP (gratuitous ARP) message will be sent and all connected VMs (including the NSX-V VMs) can update their ARP table accordingly.

You can make these configurations changes either manually or automate them by running APIs in a script file. Automation can help you to minimize the data traffic outage.

The following procedure explains the manual method of switching the default gateway to the tier-1 or tier-0 gateway in NSX-T by using the UI when Layer 3 services are not configured. For a minimum data traffic outage, you can automate the switching process with APIs.

Prerequisites

Perform the procedure in Change the MAC Address of NSX-T Virtual Distributed Router so that the NSX-V VMs can reach the default gateway in NSX-T.

Procedure

1. In the NSX-V environment, disconnect the bridged Logical Switch from the DLR.
   a. In the vSphere Client, navigate to the NSX Edge (DLR).
   b. Click **Configure > Interfaces**.
   c. Select the internal interface on the DLR that is connected to the bridged Logical Switch and click **Disconnect**.
2. In the NSX-T environment, connect the bridged overlay segment to the tier-1 gateway.
   a. In NSX Manager, navigate to **Networking > Segments**.
   b. Next to the name of the bridged overlay segment, click the vertical ellipses, and then click **Edit**.
   c. Turn on the **Connectivity** of the overlay segment, and click **Save**.

Migrate Workload VMs (Simple Case)

Use this procedure if "Applied To" is not configured in any of the DFW rules (this means that "Applied To" is set to "DFW").

If "Applied To" is configured in any of the DFW rules, do not use this procedure. Follow the procedure Migrate Workload VMs (Complex Case) instead.

Note For NSX-V to NSX-T migration, see the KB article https://kb.vmware.com/s/article/56991 for more information.

For NSX-T to NSX-V migration, migrating a workload VM back to NSX-V might not work because the distributed firewall filter in NSX-T is always higher than in NSX-V. The workaround is to place the workload VM in the NSX-T exclusion list prior to vMotion.

Prerequisites

- Ensure that:
  - vSphere vMotion is enabled on the VMkernel adapter of each host in the cluster that is involved in this migration. For detailed steps about enabling vMotion on the VMkernel adapter, see the *vSphere* product documentation.
  - The destination host in NSX-T has sufficient resources to receive the migrated VMs.
  - The source and destination hosts are in an operational state. Resolve any problems with hosts including disconnected states.

For more information about vMotion, see Migration with vMotion in the *vSphere* product documentation.

Procedure

1. Start migrating the workload VMs using vMotion in the vSphere Client. See Migrating Virtual Machines in the *vSphere* product documentation for detailed instructions.

   Note During vMotion from NSX-V to NSX-T, the workload VMs are always protected because the migration coordinator translates the existing NSX-V DFW rules and security groups into temporary IP-based rules and groups.

2. Finalize the infrastructure to finish the migration.

   ```
   POST https://{nsxt-mgr-ip}/api/v1/migration?action=finalize_infra
   ```

   This migration API deletes any temporary object configurations that were created during the migration, and ensures that the NSX-T infrastructure is in a clean state. For example, temporary IP Sets are removed from the Groups.

   This POST API does not have a request body.

3 Verify that the expected configuration items have been migrated to the NSX-T environment.

   For example, check whether the following configurations are migrated successfully:

   - User-defined Distributed Firewall rules.
   - All Grouping objects, such as IP Sets, Groups, Tags, and so on.
   - Effective members are displayed in the dynamic Groups.
   - Tags are applied to migrated workload VMs.

4 On the **Migrate Workloads** page, click **Finish**.

   A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page.

What to do next

After the migration of workload VMs and the DFW-only configuration is successful and thoroughly verified, remove the Layer 2 bridge to release the NSX-T Edge that you used for bridging.

Migrate Workload VMs (Complex Case)

Use this procedure if "Applied To" is configured in any of the DFW rules (this means that "Applied To" is not set to "DFW").

Note For NSX-V to NSX-T migration, see the KB article https://kb.vmware.com/s/article/56991 for more information.

For NSX-T to NSX-V migration, migrating a workload VM back to NSX-V might not work because the distributed firewall filter in NSX-T is always higher than in NSX-V. The workaround is to place the workload VM in the NSX-T exclusion list prior to vMotion.

Prerequisites

- Ensure that:
  - vSphere vMotion is enabled on the VMkernel adapter of each host in the cluster that is involved in this migration. For detailed steps about enabling vMotion on the VMkernel adapter, see the *vSphere* product documentation.
  - The destination host in NSX-T has sufficient resources to receive the migrated VMs.
  - The source and destination hosts are in an operational state. Resolve any problems with hosts including disconnected states.

For more information about vMotion, see Migration with vMotion in the *vSphere* product documentation.

Procedure

1 Get the instance UUID of all the VMs that you plan to migrate.

The instance UUIDs are needed when you make the API call in the next step. See the example at the bottom of this section on how to obtain the instance UUID of a VM.

2 Run the following POST API request:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/vmgroup?action=pre_migrate
```

This API creates a logical segment port (VIF) corresponding to the VM instance UUID of each NSX-V workload VM in the VM group that you will be migrating through the Layer 2 bridge to the NSX-T overlay segment. For an example request body of this API, see the Lift and Shift Migration Process section of the NSX Tech Zone article.

3 Call the API `GetVmGroupExecutionDetails`. This API is available starting with NSX-T 3.2.2

Call the API `GetVmGroupExecutionDetails` to get the result of the pre-migrate API call with the same group_id (and federation_site_id for cross-VC migration). The result includes a "logical_switch_id_to_vm_instance_id_and_vnics_map" list and an optional "failedVmInstanceIds" list, which includes the UUIDs of VMs that are not found in the source VC. For example:

```
GET /api/v1/migration/vmgroup/actions/get_vm_group_execution_details?group_id=<group-
id>&federation_site_id=<site_id>
Response:
{
  "logical_switch_id_to_vm_instance_id_and_vnics_map":[
    {
      "ls_id":"36885723-7581-4696-a195-ef83851dc35f",
      "vm_and_vnics_mapping":[
        {
          "vm_instance_id":"52199e21-6aab-26e4-8c82-069a17d67667",
          "vnics":[
            "4001"
          ]
        },
        {
          "vm_instance_id":"52630e5d-ce6f-fac0-424c-4aa4bdf6bd56",
          "vnics":[
            "4001"
          ]
        }
      ]
    }
  ],
  "failedVmInstanceIds":[
    "501557f6-2197-1fe8-14e5-89898cee5fec"
  ]
}
```

4 Follow the pseudo python code below to write a script to vmotion the VMs.

For an example, see the Python Example Scripts section of the NSX Tech Zone article.

```
    define _get_nsx_networks_in_host(self, host):
        ls_id_to_nsx_pgs_map = {}
        for net in host.network:
            if isinstance(net, vim.dvs.DistributedVirtualPortgroup):
                if hasattr(net.config, 'backingType'):
                    if net.config.backingType == 'nsx' and net.config.logicalSwitchUuid:
                        ls_id_to_nsx_pgs_map[net.config.logicalSwitchUuid] =\
                           [net.key, net.config.distributedVirtualSwitch.uuid]
            elif isinstance(net, vim.OpaqueNetwork):
                if net.summary.opaqueNetworkType == 'nsx.LogicalSwitch':
                    ls_id_to_nsx_pgs_map[net.summary.opaqueNetworkId] = [None,
net.summary.opaqueNetworkId]
        return ls_id_to_nsx_pgs_map

     define _get_vms_vnic_to_ls_id_map(self,
logical_switch_id_to_vm_instance_id_and_vnics_map):
        vm_uuid_2_vnics_map = {}
        for ls_id_2_vm_vnics in logical_switch_id_to_vm_instance_id_and_vnics_map:
            ls_id = ls_id_2_vm_vnics['ls_id']
            for vm_vnics in ls_id_2_vm_vnics['vm_and_vnics_mapping']:
                vnic_2_ls_id = vm_uuid_2_vnics_map.get(vm_vnics['vm_instance_id'], {})
                for vnic in vm_vnics['vnics']:
                    vnic_2_ls_id[vnic] = ls_id
                vm_uuid_2_vnics_map[vm_vnics['vm_instance_id']] = vnic_2_ls_id
        return vm_uuid_2_vnics_map

    def _get_nsxt_vnic_spec(self, device, dvpg_key, switch_id, vif_id):
        If dvpg_key:
            vdsPgConn = vim.dvs.PortConnection()
            vdsPgConn.portgroupKey = dvpg_key
            vdsPgConn.switchUuid = switch_id
            device.backing =
vim.vm.device.VirtualEthernetCard.DistributedVirtualPortBackingInfo()
            device.backing.port = vdsPgConn
        else:
            device.backing = vim.vm.device.VirtualEthernetCard.OpaqueNetworkBackingInfo()
            device.backing.opaqueNetworkId = switch_id
            device.backing.opaqueNetworkType = 'nsx.LogicalSwitch'
        device.externalId = vif_id
        dev_spec = vim.Vm.Device.VirtualDeviceSpec()
        dev_spec.SetOperation(vim.Vm.Device.VirtualDeviceSpec.Operation.edit)
        dev_spec.SetDevice(device)
        return dev_spec

    def _migrate_vm(self, vmObject, vnic_2_ls_id_map, ls_id_to_nsx_pgs_map):
        devices = vmObject.config.hardware.device
        nic_devices = [device for device in devices if isinstance(device,
vim.Vm.device.VirtualEthernetCard)]
        vnic_changes = []
        for device in nic_devices:
            ls_id = vnic_2_ls_id_map.get(str(device.key))
```

```
            if ls_id:
                vif_id = vmObject.config.instanceUuid + ":" + str(device.key)
                nsx_pg = ls_id_to_nsx_pgs_map.get(ls_id)
                vnic_spec = self._get_nsxt_vnic_spec(device, nsx_pg[0], nsx_pg[1], vif_id)
                vnic_changes.append(vnic_spec)
        relocate_spec = vim.Vm.RelocateSpec()
        relocate_spec.SetDeviceChange(vnic_changes)
        # set other fields in the relocate_spec
        vmotion_task = vmObject.Relocate(relocate_spec)
        WaitForTask(vmotion_task)


    vm_uuid_2_vnics_map =
self._get_vms_vnic_to_ls_id_map(logical_switch_id_to_vm_instance_id_and_vnics_map)
    for vm_uuid, vnic_2_ls_id_map in vm_uuid_2_vnics_map.items():
        # get the vmObject by the vm_uuid
        # find a target host that has all the networks needed by this VM
        ls_id_to_nsx_pgs_map = self._get_nsx_networks_in_host(host)
        self._migrate_vm(vmObject, vnic_2_ls_id_map, ls_id_to_nsx_pgs_map)
```

5 Apply the Security Tags and VM static membership to the migrated VMs.

`POST https://{nsxt-mgr-ip}/api/v1/migration/vmgroup?action=post_migrate`

The `vmgroup` API endpoint with `post_migrate` action applies the NSX-V Security Tags to the migrated workload VMs on the NSX-T overlay segment.

For an example request body of this API, see the Lift and Shift Migration Process section of the NSX Tech Zone article.

6 Finalize the infrastructure to finish the migration.

`POST https://{nsxt-mgr-ip}/api/v1/migration?action=finalize_infra`

This migration API deletes any temporary object configurations that were created during the migration, and ensures that the NSX-T infrastructure is in a clean state. For example, temporary IP Sets are removed from the Groups.

This POST API does not have a request body.

7 Verify that the expected configuration items have been migrated to the NSX-T environment.

For example, check whether the following configurations are migrated successfully:

- User-defined Distributed Firewall rules.
- All Grouping objects, such as IP Sets, Groups, Tags, and so on.
- Effective members are displayed in the dynamic Groups.
- Tags are applied to migrated workload VMs.

8 On the **Migrate Workloads** page, click **Finish**.

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page.

Example: Obtaining VM Instance UUID from the vCenter MOB

This example shows how to obtain or confirm a VM's instance UUID from the vCenter Server Managed Object Browser (MOB) at `http://{vCenter-IP-Address}/mob`. You can also obtain or confirm a VM's instance UUID by making an API call to vSphere.

1 In a web browser, enter the vCenter Managed Object Browser at `http//{vCenter-IP-Address}/mob`.

2 Click **content**.

3 Find **rootFolder** in the Name column, and click the corresponding link in the Value column. For example, group-d1.

4 Find **childEntity** in the Name column, and click the corresponding link in the Value column. For example, datacenter-21.

5 Find **hostFolder** in the Name column, and click the corresponding link in the Value column. For example, group-h23.

6 Find **childEntity** in the Name column. The corresponding Value column contains links to host clusters. Click the appropriate host cluster link. For example, domain-c33.

7 Find **host** in the Name column. The corresponding Value column lists the hosts in that cluster by vCenter MOID and hostname. Click the appropriate host link, For example, host-32.

8 Find **vm** in the Name column. The corresponding Value column lists the virtual machines by vCenter MOID and hostname. For example, vm-216 (web-01a). Click the VM that you are interested in.

9 Find **config** in the Name column. Click **config** in the Value column.

10 Find **instanceUuid** in the Name column. The corresponding Value column lists the VM instance UUID. For example, 502e71fa-1a00-759b-e40f-ce778e915f16.

What to do next

After the migration of workload VMs, you can remove the layer-2 bridge.

In-Place Migration of Specific Parts of NSX-V

8

This section covers advanced migration modes that you can use to migrate specific parts of your NSX-V environment without the need to deploy an extra hardware, such as separate compute clusters, in your destination NSX-T environment.

For example: Your organization might prefer to create a new NSX-T topology and configure the NSX-T Edges and other networking services manually without using the migration coordinator. For instance, you can deploy NSX-T Edge nodes on existing NSX-V hosts.

Your objective is to use the migration coordinator to migrate only the existing Distributed Firewall (DFW) configuration and NSX-V compute hosts to NSX-T. You want the workload VMs to continue running on the existing compute hardware with a minimal disruption to the east-west traffic security protection when migrating to the new NSX-T environment.

As you are doing an in-place migration of only the DFW configuration and compute hosts, you have the flexibility to define your own NSX-T topology. In other words, the supported fixed topologies mentioned in Fixed Topologies Supported for End-to-End Migration are not relevant for this in-place migration.

You can have vSphere High Availability (HA) enabled if the NSX-V environment has VDS 7.0 or later. If the NSX-V environment has VDS 6.5 or 6.7, and the vmkernel ports (vmks) are attached to VDSes, during an in-place migration, the hosts and VMs may lose network connectivity for a period of time log enough to trigger HA. The HA mechanism will try to power off, migrate and restart VMs. This might fail because the NSX-V environment is being migrated to NSX-V. As a result, after the migration, VMs might remain in a powered-off state or have no network connectivity if powered on. To avoid this situation, disable HA or attach the management vmk to a VSS before starting the migration to NSX-T.

Note: Logical ports and switches that are created during migration are not deleted when the workload VMs are deleted. You must delete these ports and switches via the NSX Manager UI or the API.

To minimize disruption during migration, ensure that:

- NSX-V and NSX-T edges are on different ESXi hosts.
- Workload VMs directly connected to an Edge Services Gateway (ESG) are on a different ESXi host than the ESG.

For this migration, if you plan or need to manually migrate VMs from some NSX-V hosts to NSX-T hosts, you must configure the export version of Distributed Firewall (see Configure Export Version of Distributed Firewall Filter).

If the NSX-V environment has VDS 7.0 or later, you must enable and configure DRS for each NSX-V cluster to be migrated. Under **Automation**, set **Automation Level** to **Fully Automated**. To prevent DRS from automatically migrating any VM from an NSX-V host to an NSX-T host before the migration process starts migrating the NSX-V host, under **Automation**, set **Migration Threshold** to most conservative. Also, under **Additional Options**, do not enable **VM Distribution**. After a cluster's migration to NSX-T is completed, you can change **Migration Threshold** to aggressive and enable **VM Distribution** under **Additional Options** so that DRS will balance the VMs among the hosts.

Read the following topics next:

- Overview - In-Place Migration of Specific Parts
- Tag Management VMs in a Collapsed Cluster Environment
- Migrating Distributed Firewall Configuration, Hosts, and Workloads
- Migrating North-South Traffic to NSX-T Edges Using Edge Cutover

Overview - In-Place Migration of Specific Parts

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.

- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.
- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 8-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.

Table 8-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. Migration Coordinator will only migrate from an NSX-V Manager with the role of Primary or Standalone. You can modify the NSX-V environment by changing the status of the secondary managers in order to migrate each NSX-V environment independently.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	

NSX-V Configuration	Supported	Details
Network I/O Control (NIOC) having vNIC with reservation	No	
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	
Host profiles	No	
ESXi lockdown mode	No	Not supported in NSX-T.
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.

NSX-V Configuration	Supported	Details
VXLAN port number other than 4789	No	
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes for in-place migration No for lift-and-shift migration	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	

NSX-V Configuration	Supported fro Migration	Details
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	

NSX-V Configuration	Supported	Details
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	

NSX-V Configuration	Supported	Details
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description

NSX-V Configuration	Supported	Details
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	

NSX-V Configuration	Supported	Details
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 8-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 8-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.
DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration. For example, dhcp option '80' is not migrated. <dhcpOptions> <other> <code>80</code> <value>2f766172</value> </other> </dhcpOptions>
Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 8-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	

NSX-V Configuration	Supported	Details
Rule – Source / Destination: ■ VMs not on NSX-V hosts	No	NSX-T does not support referencing objects not connected to NSX-T portgroups or networks. Those objects will be lost from the source or destination when migration is completed. To avoid this issue use IP addresses in NSX-V to reference those objects before migration.
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.

NSX-V Configuration	Supported	Details
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	

NSX-V Configuration	Supported	Details
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 8-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of ‘Groups’ after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 8-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.

Table 8-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.

Table 8-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 8-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 8-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 8-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 8-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000

Table 8-10. Single-Site Limits (continued)

Feature	Limit
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 8-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Changes Made During Host Migration in an End-to-End Migration

During the host migration step in an end-to-end migration, changes are made to migrate NSX-V hosts to NSX-T hosts.

- NSX-V software is uninstalled.
- NSX-T software is installed.
- For vSphere Distributed Switch versions 6.5.0 and 6.6.0:

  Hosts are configured with N-VDS to replace vSphere Distributed Switches:

  - Each N-VDS is created with a name that references the distributed switch name. For example, distributed switch `ComputeSwitchA` is created as N-VDS `nvds.ComputeSwitchA`.
  - If different clusters use different distributed switches to back logical switches, an N-VDS is created with a name that combines all the distributed switch names. For example, if `ComputeCluster1` and `ComputeCluster2` use distributed switch `ComputeSwitchA` to back logical switches and `ComputeCluster3` uses `ComputeSwitchB` to back logical switches, the N-VDS is created as `nvds.ComputeSwitchA.ComputeSwitchB`.
  - PNICs and vmks in the vSphere Distributed Switch are migrated to N-VDS.
  - NSX-V VTEPs are migrated to NSX-T TEPs.
- For vSphere Distributed Switches version 7.0:

  Hosts configured for vSphere Distributed Switch version 7.0 continue using the same switch after migration.

  - PNICs and vmks in the vSphere Distributed Switch remain connected on the same vSphere Distributed Portgroups.
  - NSX-V VTEPs are migrated to NSX-T TEPs and connected to standalone ports on the same vSphere Distributed Switch.

Note: If NSX-V has multiple VTEPs and a single LAG configuration, after migration the LAG will have a single TEP and in failover mode.

Virtual Machine Deployment During an End-to-End Migration

After you start an end-to-end migration, do not change the NSX-V environment. If you want to deploy VMs during the migration, wait until some of the NSX-V hosts have migrated to NSX-T and deploy the VMs on NSX-T hosts. Connect the VMs to NSX-T segments and install VMware Tools on the VMs.

Deploying on NSX-T with VMware Tools installed ensures that the VMs are populated into security groups and receive the intended Distributed Firewall policies.

Caution VMs deployed without VMware Tools installed, or deployed on NSX-V do not receive the intended Distributed Firewall policies.

If you use vSphere templates to deploy VMs, update the templates to use NSX-T segments for the VM network configuration. Specifying NSX-T segments ensures that any VMs deployed using the templates are deployed on NSX-T hosts.

If you use automation tools to deploy VMs on vSphere, but do not use vSphere templates, you might need to change your automation tool configuration to ensure that the VMs are deployed on NSX-T.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1
	Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Tag Management VMs in a Collapsed Cluster Environment

You can migrate an environment that uses a collapsed cluster.

In a collapsed cluster design, all management VMs, workload VMs, and optionally edges run on the same vSphere cluster that is prepared for NSX-V. The management VMs of the NSX-T must be initially attached to dvPortgroups. After migration, the management VMs of NSX-T will be attached to NSX-T VLAN segments.

The management VMs in the NSX-T include appliances such as NSX Manager, vCenter Server, VMware Identity Manager, and so on. The NSX-T VLAN segment ports to which these management VMs connect are blocked in two cases: when these management VMs are rebooted after they were migrated in-place by the Migration Coordinator, or when they are moved from NSX-V hosts to NSX-T hosts by vMotion in maintenance migration mode. Therefore, the management VMs might lose connectivity in such cases.

To prevent this problem, create a "management_vms" tag category, and add tags in this category. Assign a tag from this category to all the management VMs in the NSX-T environment. These VMs will be attached to unblocked VLAN segment ports.

Procedure

1 Log in to the vSphere Client.

2 Click **Menu > Tags & Custom Attributes**.

3 Click **Categories**, and then click **New** to add a category.

Create a category with name **`management_vms`**.

4 Click the **Tags** tab and add a tag in the management_vms category.

5 Navigate to **Menu > Hosts and Clusters**.

6 Expand the collapsed cluster from the left Navigator view, right-click the name of the NSX Manager VM, and select **Tags & Custom Attributes > Assign Tag**.

7 Assign a tag from the management_vms category to the NSX Manager VM.

8 Repeat steps 6 and 7 for all the management VMs in the cluster.

For a detailed information about tag categories and tags, see the *vCenter Server and Host Management* documentation.

9 Log in to NSX Manager.

10 Navigate to **Inventory > Groups** and create a group.

11 Click **Set** to set members.

12 Set **Group Type** to **IP Addresses Only** and include the IP addresses of the management VMs that will be migrating from the dvPortgroups to the NSX-T segments.

13 Navigate to **Security > Distributed Firewall > Actions > Exclusion List** and add this group.

Migrating Distributed Firewall Configuration, Hosts, and Workloads

In this migration mode, you migrate the Distributed Firewall configuration, NSX-V hosts, and workload VMs.

If you have Identity Firewall (IDFW) configured, it will also be migrated. For more information about migrating IDFW, see Chapter 2 Migrating Identity Firewall (End-to-End and Lift-and-Shift).

The existing NSX-V prepared compute clusters are migrated to NSX-T. You do not require separate compute host clusters in your destination NSX-T environment.

The migration process will create the required infrastructure to extend the networks between hosts that are still on NSX-V and hosts that are migrated to NSX-T. The layer-2 extension allows the migration of the environment without disrupting the connectivity between the VMs on NSX-V hosts and the VMs on hosts that are migrated to NSX-T.

The following objects in the DFW configuration are migrated:

- User-defined Distributed Firewall (DFW) rules

- Grouping Objects
  - IP Sets
  - MAC Sets
  - Security Groups
  - Services and Service Groups
  - Security Tags
- Security Policies created using Service Composer (only DFW rule configurations are migrated)

  Guest Introspection service configuration and Network Introspection rule configurations in the Service Composer are not migrated.

Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with only local objects.

For a detailed list of all the configurations that are supported for the migration of Distributed Firewall configuration, see the Detailed Feature Support for Migration.

Prerequisites for DFW, Host, and Workload Migration

- A new NSX-T is deployed for this migration.
  - Deploy NSX Manager appliances.

    In a production environment, add an NSX Manager cluster with three appliances. However, for migration purposes, a single NSX Manager appliance is adequate.
  - Deploy a vCenter Server appliance.

    The vCenter Server must be added as a compute manager in NSX-T. You can share the vCenter Server that is used in NSX-V or deploy another one in NSX-T.
  - This migration mode does not require you to deploy NSX-T Edges before starting the migration. However, to provide routing, Layer 3 networking services, and north-south connectivity to the physical ToR switches, you must deploy Edges in your NSX-T environment.
  - Create overlay segments in NSX-T with the same virtual network identifier (VNI) and subnet address as the Logical Switches in NSX-V.

    That is, for each NSX-V Logical Switch, add a corresponding overlay segment in NSX-T. Same subnet address helps in ensuring that the IP addresses of the workload VMs are retained after the VMs move to NSX-T segments. Use the NSX-T APIs to create the overlay segments. You cannot create overlay segments with the same VNI in the NSX Manager UI.

You must create the segments with the `SOURCE` replication mode, and change the mode to `MTEP` only after the migration is done.

- Create VLAN segments in NSX-T with the same VLAN IDs and subnet address as the VLAN Distributed Virtual Port Groups (DVPG) in NSX-V.

  Note VLAN DVPG must be associated only with a VLAN ID. VLAN Trunk is not supported.

- No user-defined DFW rules exist in NSX-T before this migration.
- All states in the **System Overview** pane of the NSX-V dashboard are green.
- There are no unpublished changes for Distributed Firewall and Service Composer policies in the NSX-V environment.

Import the NSX-V Configuration

The first step of the migration process is to import the NSX-V configuration.

Prerequisites

- The vCenter Server must be added as a compute manager in NSX-T.

  You can share the vCenter Server that is used in NSX-V or deploy another one in NSX-T.

Procedure

1. From a browser, log in to NSX Manager as **admin**.
2. Navigate to **System > Migrate**.
3. Expand the **Advanced Migration Modes** section, and in the **Distributed Firewall, Host, and Workload** pane, click **Get Started**.
4. From the **Import Configuration** page, click **Select NSX** and provide the credentials for vCenter Server and NSX-V.

   Note The drop-down menu for vCenter Server displays all vCenter Server systems that are registered as compute managers. Click **Add New** if you need to add a compute manager.

5. Click **Start** to import the configuration.
6. When the import has finished, click **Continue** to proceed to the **Resolve Configuration** page.

Results

When the NSX-V topology is imported successfully, the **View Imported Topology** link is enabled. Click this link to view a graph of the imported topology. However, the topology viewer does not display the graph of a large-scale NSX-V environment.

Resolve Configuration

After you have imported the configuration from your NSX-V environment, you must review and resolve the reported configuration issues before you can continue with the migration.

On the **Resolve Configuration** page, two types of configuration issues are reported.

Blocking Issues (if any)

As the name suggests, these issues block the migration, and they must be fixed for the migration to proceed. You might have to change the configurations in your NSX-V environment before you can migrate to NSX-T.

Warnings

These configuration issues are organized into several categories, and each category can have one or more configuration items. You should provide inputs to fix the warning messages that are displayed for the configuration items, or choose to skip the warnings, if needed.

Procedure

1. On the **Resolve Configuration** page, if you find any blocking issues, fix them in the NSX-V environment before you can proceed with the migration.

   After making the required changes in the NSX-V environment, return to the migration coordinator. Go to the **Import Configuration** page, and click **Rollback**. Click **Start** to import the updated NSX-V configuration.

   If you did not find any blocking issues that require a change in the NSX-V environment, proceed to the next step.
2. Review the warnings and issues reported in each category.
3. Click each issue and provide input.

   Each issue can cover multiple configuration items. For each item there might be one or more possible resolutions to the issue, for example, skip, configure, or select a specific value.

   For issues that apply to multiple configuration items, you can provide input for each item individually, or select all and provide one answer for all items.
4. After the input is provided, a **Submit** button is displayed on the **Resolve Configuration** page. Click **Submit** to save your progress.
5. When you have provided input for all configuration issues, click **Submit**.

   The input is validated. You are prompted to update any invalid input. Additional input might be required for some configuration items.
6. After you have submitted all requested input, click **Continue** to proceed to the **Migrate Configuration** page.

Example: Configuration Issues

Migration coordinator shows warning messages if it detects the following issues in your NSX-T environment:

- A missing NSX-T VLAN segment corresponding to the Distributed Virtual Port Group (DVPG) VLAN ID in your NSX-V environment.
- A missing NSX-T overlay segment corresponding to the VNI of the Logical Switch in the NSX-V environment.

To resolve this configuration issue, you can either skip creating the missing NSX-T segments or roll back the migration. If you choose to roll back the migration, create the missing segments in NSX-T by running the following PATCH API request, and start the migration again.

```
PATCH https://{policy-manager}/policy/api/v1/infra/segments/{segment-id}?
force=true
```

For example, to create a segment "App" with overlay ID "5001", the API URL and the payload of the API request is as follows:

```
PATCH https://{policy-manager}/policy/api/v1/infra/segments/App?force=true
```

```
{ "type" : "ROUTED",
  "subnets" : [ {" "gateway_address" : "172.16.20.1/24",
    "network" : "172.16.20.0/24""
    "} ],
  "connectivity_path" : "/infra/tier-0s/T0-GW-01",
  "transport_zone_path" : "/infra/sites/default/enforcement-points/default/transport-zones/
   1b3a2f36-bfd1-443e-a0f6-4de01abc963e",
  "overlay_id":5001,
  "admin_state" : "UP",
  "replication_mode" : "MTEP",
  "resource_type" : "Segment",
  "id" : "App",
  "display_name" : "App",
  "path" : "/infra/segments/App",
  "relative_path" : "App",
  "parent_path" : "/infra" }
```

If you choose to skip creating the missing NSX-T VLAN or overlay segments, the VM vNICs that are connected to the DVPG VLANs or the Logical Switches in the NSX-V environment lose networking after the workload VMs are migrated using vSphere vMotion. In this case, the migration coordinator does not migrate the DFW rules and Security Groups that are associated with the skipped segments.

For example, you might not want to create the missing segments in the following situations:

- Port groups on the vSphere Distributed Switch (VDS) have only VMkernel ports. So, no requirement to create VLAN segments.
- Port groups from another VDS that you do not want to migrate.

- NSX-V Logical Switch with no workload VMs attached to it. For instance, a transit Logical Switch between an Edge Services Gateway and the Distributed Logical Router.

Migrate the Distributed Firewall Configuration

After you have resolved all configuration issues, you can migrate the Distributed Firewall configuration. When the configuration is migrated, logical object configurations are realized in NSX-T environment, which replicate the NSX-V logical object configurations.

DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Prerequisites

Verify that you have completed the **Resolve Configuration** step.

Procedure

1 From the **Migrate Configuration** page, click **Start**.

2 Verify that the Distributed Firewall configuration objects are displayed in your NSX-T environment.

You can verify the migrated configurations either in the NSX-T NSX Manager interface or by running the NSX-T APIs.

Note

- During the **Migrate Configuration** step, Security Tags from NSX-V are not migrated to NSX-T. Therefore, the Security Tag-based migrated dynamic Groups and Groups with static memberships in NSX-T are empty after this step is finished. The reason is that in NSX-V, a Security Tag is an object, whereas in NSX-T, a tag is an attribute of a VM. The tags are applied to the workload VMs only after the workloads are migrated to NSX-T during the **Migrate Hosts** step.
- When the logical configurations are migrated to NSX-T, the configuration changes are made in the NSX-T NSX Manager database, but it might take some time for the configurations to take effect.

3 Click **Continue** to proceed.

If needed, you can roll back the migrated DFW configuration.

Rolling back does the following:

- Remove the migrated configuration from NSX-T.
- Roll back all the resolved issues in the previous step.

Any NSX-T objects that you manually created after the DFW migration are at risk of being lost during the rollback.

Prepare Infrastructure to Extend Layer 2 Before Host Migration

After migrating the logical configurations to NSX-T, the next step is to prepare the required infrastructure to extend the networks between hosts that are still on NSX-V and hosts that are migrated to NSX-T.

The layer-2 extension allows the migration of the environment without disrupting the connectivity between the VMs on NSX-V hosts and the VMs on hosts that are migrated to NSX-T.

In the **Prepare Infrastructure** step, the following tasks happen in the background:

- The Controller Disconnected Operation (CDO) mode is enabled in NSX-V. The NSX-V NSX Manager creates a special CDO logical switch with VNI 4999 on the NSX-V controller. This VXLAN Network Identifier of the special CDO logical switch is unique from all other logical switches. When the CDO mode is enabled, any one NSX-V controller in the Controller Cluster collects the VXLAN Tunnel Endpoint (VTEP) information reported from all NSX-V prepared hosts, and replicates the updated VTEP information to all other NSX-V hosts. After detecting the CDO mode, broadcast packets such as, ARP, GARP, and RARP are sent to the global VTEP list. By enabling the CDO mode, vMotion of VMs can occur without any data plane connectivity issues when the NSX-V control plane fails.
- The NSX-V VTEP table is made available to the Central Control Plane (CCP) service, which is running on the NSX-T NSX Manager appliance.

Prerequisites

- In the NSX-V environment, NSX Controller Cluster is up.
- In the NSX-T environment, Central Control Plane (CCP) service on the NSX Manager appliance is up.

Procedure

1. On the **Prepare Infrastructure** page, click **Start**.
2. After the infrastructure preparation is finished, click **Continue** to proceed to the **Migrate Hosts** page.

   If the infrastructure preparation fails, resolve the errors, click **Rollback** to roll back the migration, and start the **Prepare Infrastructure** step again. For example, infrastructure preparation might fail due to the following reasons:

   - NSX-V Controllers are down.
   - NSX-T CCP service on the NSX Manager appliance is down.

Migrate NSX-V Hosts

After the infrastructure is prepared for extending Layer 2 networks, you can migrate your NSX-V hosts to NSX-T host transport nodes.

Prerequisites

- Configure several settings related to the host migration, including migration order and enabling hosts. Make sure that you understand the effects of these settings. See Configuring NSX-V Host Migration for more information. Understanding the host migration settings is especially important when you use Distributed Firewall or vSphere Distributed Switch 7.0.
- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.

Procedure

1 On the **Migrate Hosts** page, click **Start**.

If you selected the **In-Place** or **Automated Maintenance** migration mode for all hosts groups, the host migration starts. Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

2 If you selected the **Manual Maintenance** migration mode for any host groups, you must complete one of the following tasks for each VM so that the hosts can enter maintenance mode.

Option	Action
Power off or suspend VMs.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b After the host has migrated, attach the VM interfaces to the appropriate NSX-T segments and power on the VM.
Move VMs using vMotion.	Right click the VM and select Migrate. Follow the prompts to move the VM to a different host. Note that Migration Coordinator maintains security during migration by vMotioning VMs to specific ports that are protected by temporary rules. In the case of manual vMotion, the VMs will not be moved to those ports and there could be a security breach. To vMotion manually, the VMs must be migrated using vSphere API where the networking backing must point to the OpaqueNetwork ID corresponding to the NSX Segment when using NVDS or the VDS portgroup ID when using VDS 7. In both cases, the network device's externalId must be set to the the string "VM_UUID:vNIC_ID", where VM_UUID is the VM's instance UUID and vNIC_ID is the VM's vNIC index where the first vNIC is 4000.
Move VMs using cold migration.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b Right click the VM and select Migrate. Follow the prompts to move the VM to a different host, connecting the VM interfaces to the appropriate NSX-T segments.

Here is python code to specify an external-id for each vNIC in a VM and then vMotion the VM so that the vNICs will connect to an NSX-T segment of ID “ls_id” at the correct ports:

```
devices = vmObject.config.hardware.device
nic_devices = [device for device in devices if isinstance(device,
vim.Vm.device.VirtualEthernetCard)]
vnic_changes = []
for device in nic_devices:
    vif_id = vmObject.config.instanceUuid + ":" + str(device.key)
    vnic_spec = self._get_nsxt_vnic_spec(device, ls_id, vif_id)
    vnic_changes.append(vnic_spec)
relocate_spec = vim.Vm.RelocateSpec()
relocate_spec.SetDeviceChange(vnic_changes)
# set other fields in the relocate_spec
vmotion_task = vmObject.Relocate(relocate_spec)
WaitForTask(vmotion_task)

def _get_nsxt_vnic_spec(self, device, ls_id, vif_id):
    nsxt_backing = vim.Vm.Device.VirtualEthernetCard.OpaqueNetworkBackingInfo()
    nsxt_backing.SetOpaqueNetworkId(ls_id)
    nsxt_backing.SetOpaqueNetworkType('nsx.LogicalSwitch')
    device.SetBacking(nsxt_backing)
    device.SetExternalId(vif_id)
    dev_spec = vim.Vm.Device.VirtualDeviceSpec()
    dev_spec.SetOperation(vim.Vm.Device.VirtualDeviceSpec.Operation.edit)
    dev_spec.SetDevice(device)
    return dev_spec
```

For an example of a complete script, see https://github.com/dixonly/samples/blob/main/vmotion.py

The host enters maintenance mode after all VMs are moved, powered off, or suspended. If you want to use cold migration to move the VMs to a different host before the migrating host enters maintenance mode, you must leave at least one VM running while you move VMs. When the last VM is powered off or suspended, the host enters maintenance mode, and migration of the host to NSX-T starts.

Results

After a host has migrated to NSX-T using **In-Place** migration mode, you might see a critical alarm with message `Network connectivity lost`. This alarm occurs when a vSphere Distributed Switch (VDS) 6.5 or 6.7 migrates to an N-VDS because the host no longer has a physical NIC connected to the VDS it was previously connected to. To restore the migrated hosts to the Connected state, click **Reset to Green** on each host, and suppress the warnings, if any.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

If migration fails for a host, the migration pauses after all in-progress host migrations finish. When you have resolved the problem with the host, click **Retry** to retry migration of the failed host. If the host still fails to migrate, you can configure NSX-T on the host manually or remove the host from the system. In this case, at the end of the host migration step, the **Finish** button

will not be enabled because of the host that failed to migrate. You need to call the REST API `POST https://<nsx-mgr-IP>/api/v1/migration?action=finalize_infra` (<nsx-mgr-IP> is the IP address of the NSX Manager where the migration service is running) using a REST API client (for example, postman or curl) to finish the migration, and then perform the post-migration tasks.

What to do next

1 Verify that the expected configuration items have been migrated to the NSX-T environment. For example, check whether the following configurations are migrated successfully:

- User-defined Distributed Firewall rules.
- All Grouping objects, such as IP Sets, Groups, Tags, and so on.
- Effective members are displayed in the dynamic Groups.
- Tags are applied to migrated workload VMs.

Verify that the VMs running on the NSX-T hosts are connected to the correct NSX-T overlay segment and validate the following connectivity:

- VM-to-VM connectivity in the NSX-T network.
- Connectivity of the VMs to the external machines outside the NSX-T network, provided the DFW rules allow it.

2 On the **Migrate Hosts** page, click **Finish**.

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page, or which hosts were excluded from the migration.

Migrating North-South Traffic to NSX-T Edges Using Edge Cutover

In this migration, north-south traffic is migrated from NSX-V Edge Services Gateways to NSX-T Edge nodes.

During the cutover, the uplinks on the NSX-V Distributed Logical Router are internally disconnected from the Transit Logical Switch to which the Edge Services Gateways are also connected. The uplinks on the NSX-T Edge nodes are brought online.

If your NSX-V environment has DHCP service configured on an Edge Services Gateway, the DHCP leases of the workload VMs are migrated to the NSX-T Edge during the Edge cutover. After the Edge cutover is finished, the workload VMs, which are still running on NSX-V prepared hosts, have the same DHCP leases as existed before the migration.

Remember that this migration mode does not migrate the existing NSX-V topology and logical configurations, such as routing, Edge firewall, distributed firewall, L3 networking services, and so on, to your new NSX-T environment. You must pre-configure the topology and the logical objects manually in your new NSX-T environment. Or migrate the logical configurations to NSX-T before starting the Edge cutover migration.

This migration mode does not do an in-place migration of your existing NSX-V hosts to NSX-T. You must migrate the existing hosts to NSX-T after the Edge cutover migration is finished.

Note To use the Edge cutover migration mode, your NSX-V environment can be configured in any topology. In other words, the Edge cutover migration does not require the NSX-V topology to be mandatorily configured in any of the supported migration topologies that are explained in Fixed Topologies Supported for End-to-End Migration.

Prerequisites for Edge Cutover Migration

- Supported software version requirements:
  - See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
  - vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
  - The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- A new NSX-T environment is deployed and configured for this migration.
  - Deploy NSX Manager appliances.

    In a production environment, add an NSX Manager cluster with three appliances. However, for migration purposes, a single NSX Manager appliance is adequate.
  - Deploy a vCenter Server appliance.

    The vCenter Server must be added as a compute manager in NSX-T. You can share the vCenter Server that is used in NSX-V or deploy another one in NSX-T.
  - Deploy the correct number of appropriately sized NSX-T Edge appliances to replace the NSX-V Edge Service Gateways to match it feature wise and performance wise. You can deploy the NSX Edge nodes on the existing NSX-V prepared hosts.
  - Join the Edge nodes to the management plane from the command line.
  - Configure NSX-T on the NSX Edge nodes.
  - Either create an NSX-T IP pool to use for the NSX-T Edge TEPs, or configure static IP addresses for the Edge TEPs.
  - Add tier-0 and tier-1 gateways depending on the requirements of your NSX-T network topology.
  - Create overlay segments in NSX-T with the same virtual network identifier (VNI) and subnet address as the Logical Switches in NSX-V.

Use the NSX-T APIs to create the overlay segments. You cannot create overlay segments with the same VNI in the NSX Manager UI.

You must create the segments with the `SOURCE` replication mode, and change the mode to `MTEP` only after the migration is done.

- Create VLAN segments in NSX-T with the same VLAN IDs and subnet address as the VLAN Distributed Virtual Port Groups (DVPG) in NSX-V.

  Note VLAN DVPG must be associated only with a VLAN ID. VLAN Trunk is not supported.

- Connect the uplink interface of the tier-0 gateway to a transit VLAN segment.

  Configure dynamic route peering between tier-0 gateway and the north-facing physical routers.

- Attach the NSX-T overlay segments to the downlinks of the tier-0 or tier-1 gateway depending on the requirements of your NSX-T topology.

- For all Layer 3 services, such as Network Address Translation, Load Balancing, VPN, and so on, that are configured on the NSX-V Edge Services Gateway, pre-configure equivalent services on the tier-0 or tier-1 gateway of your NSX-T environment. Enable Route Advertisement and Layer 3 services on the tier-1 gateway. Enable Route Redistribution Status on the tier-0 gateway.

  Important If a DHCP service is configured on your NSX-V Edge Services Gateway, pre-configure a Gateway DHCP service on the NSX-T overlay segment. For migrating DHCP leases, Edge cutover migration mode supports only Gateway DHCP service. Local DHCP server or Local DHCP relay is not supported.

- The DPDK fast path uplink interfaces on the NSX-T edges (fp-eth0, fp-eth1, and fp-eth2) must be down.

  Use the NSX-T APIs to update the `admin_status` parameter of each DPDK fast path interface on the Edge transport node to `down`.

  For example, to change the `admin_status` parameter of fp-eth0 interface on the Edge transport node, run the following APIs:

  1 Retrieve the edge-id of the NSX-T Edge transport node:

    ```
    GET https://{nsxt-mgr-ip}/api/v1/transport-nodes
    ```

  2 Use the edge-id that you obtained in step 1 to retrieve the properties of the fp-eth0 network interface:

    ```
    GET https://{nsxt-mgr-ip}/api/v1/transport-nodes/{edge-id}/node/
    network/interfaces/fp-eth0
    ```

  3 Paste the full GET API response in a text editor and change the `admin_status` parameter to `down`.

4 Paste the full edited API response in the request body of the PUT API:

```
PUT https://{nsxt-mgr-ip}/api/v1/transport-nodes/{edge-id}/node/
network/interfaces/fp-eth0
```

For a detailed information about the parameters in these APIs, see the *NSX-T Data Center API Guide*.

- You cannot map multiple tier-1 gateways without an edge cluster or a DR-only tier-1 gateway under a parent tier-0 gateway to DLRs. You also cannot map the parent tier-0 gateway to a DLR if you are mapping to a DR-only tier-1 gateway. If your topology requires mapping multiple DLRs, you must use an active-standby tier-1 gateway with an edge cluster assigned.

Overview of Input Configuration File

To do the Edge cutover migration, create a configuration file in a `.json` format and provide it as an input to the migration coordinator.

The configuration file contains the following information:

- List of NSX-V Edge appliances, which includes a Distributed Logical Router (DLR) or an Edge Services Gateway (ESG), or both:
  - Edge ID of the Distributed Logical Router for Edge cutover.
  - Edge ID of the Edge Services Gateway for migrating the DHCP leases, if DHCP service is configured on the Edge Services Gateway.
- The mapping of the NSX-V Edges to the name of the NSX-T tier-0 or tier-1 gateway.

For example, the following figure shows a partial view of the **NSX Edges** page in the vSphere Client. The Edge IDs of the Distributed Logical Router and Edge Services Gateway are highlighted. Use these Edge IDs in the configuration file.

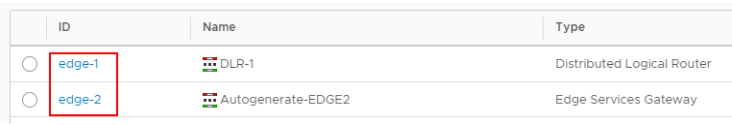

The structure of the `.json` configuration file with sample parameter values is as follows:

```
[
    {
        "name": "ns-edge-cutover",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-1",
                    "edge-2"
                ],
                "policy_gateway_name": "my_tier1"
```

```
            }
        ]
    }
]
```

Table 8-12. Configuration File Parameters

Parameter	Description	Data Type	Notes
`name`	A name for this configuration.	String	This parameter is required.
`v_edges_to_policy_gateways_mappings`	A list of mappings for the Edge cutover and DHCP lease migration. Each mapping consists of two parameters: `v_edges` and `policy_gateway_name`. See the next two rows for more details about these two parameters.	Array	This parameter is required.
`v_edges`	A list of NSX-V Edge IDs. This list includes the Edge ID of the DLR for Edge cutover, or the Edge ID of an ESG for a DHCP lease migration, or both.	Array of string values	This parameter is required. Minimum: One Edge ID Maximum: Two Edge IDs per list Each Edge ID in the list must be unique.
`policy_gateway_name`	The desired mapping of the NSX-V Edges to the name of the NSX-T tier-0 or tier-1 gateway.	String	This parameter is required. The name must match exactly with the preconfigured NSX-T tier-0 or tier-1 gateway name.

Example 1: Configuration File

Following figure shows the NSX-V environment that is configured in Fixed Topologies Supported for End-to-End Migration. The equivalent NSX-T topology is shown to the right.

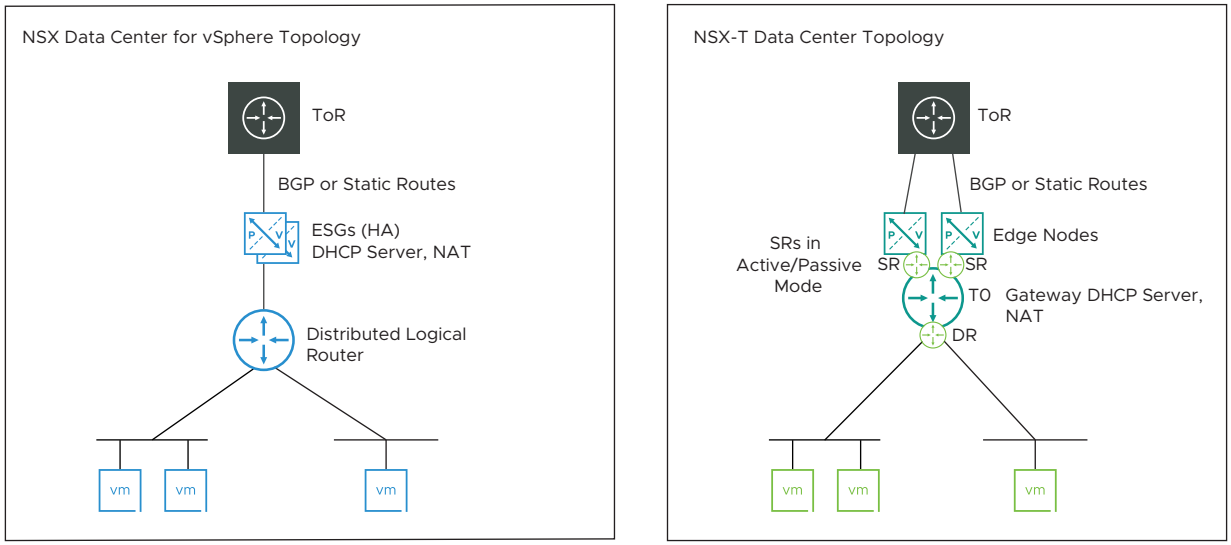

In this example, assume that NAT and DHCP server are configured on the ESG. In the NSX-T environment, you have configured these services on the tier-0 gateway. During Edge cutover, the DHCP leases on the ESG are migrated to the Gateway DHCP server on the tier-0 gateway.

The desired mapping in the input configuration file is as follows:

```
[
    {
        "name": "ns-edge-cutover",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-1",
                    "edge-2"
                ],
                "policy_gateway_name": "my_tier0"
            }
        ]
    }
]
```

In this configuration file:

- edge-1 is the Edge ID of the Distributed Logical Router for Edge cutover.
- edge-2 is the Edge ID of the Edge Services Gateway where the DHCP service is configured.
- my_tier0 is the name of the NSX-T tier-0 gateway.

Example 2: Configuration File

Following figure shows the NSX-V environment that is configured in Fixed Topologies Supported for End-to-End Migration. The equivalent NSX-T topology is shown to the right.

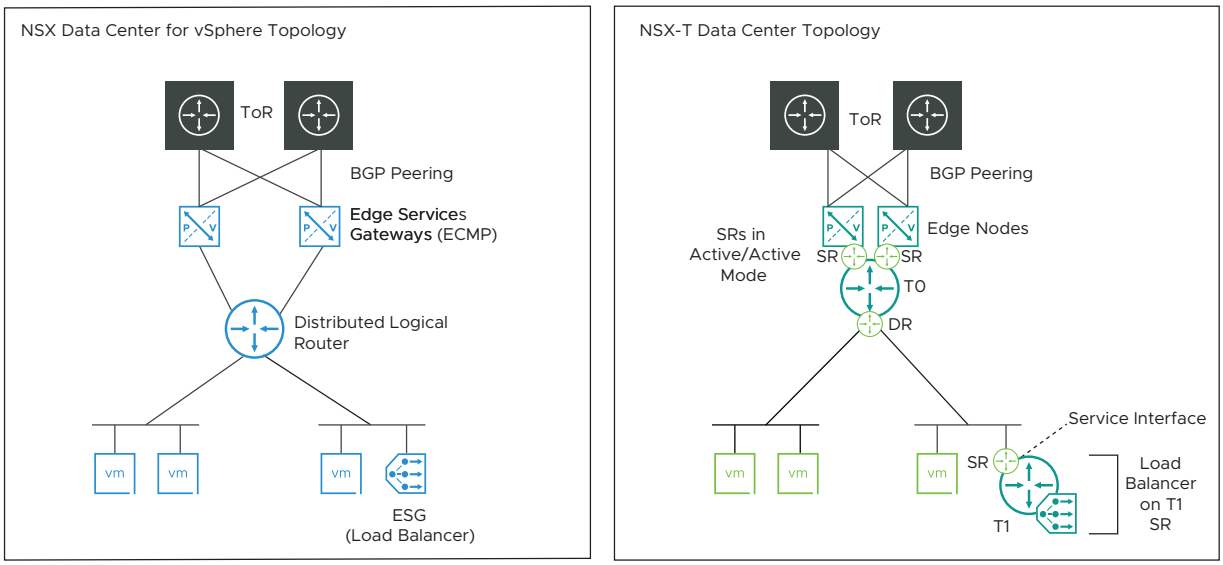

In this example, only a single-arm load balancer is configured on the ESG that is attached to the NSX-V Logical Switch. DHCP service is not running on this ESG. In the corresponding NSX-T topology, load balancer service is preconfigured on the tier-1 gateway (Service Interface) before the Edge cutover. When Edge cutover occurs, only the north-south traffic is migrated to the NSX-T Edge nodes. No DHCP lease migration is involved.

The desired mapping in the input configuration file is as follows:

```
[
    {
        "name": "ns-edge-cutover",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-1"
                ],
                "policy_gateway_name": "my_tier0"
            }
        ]
    }
]
```

In this configuration file:

- edge-1 is the Edge ID of the Distributed Logical Router for Edge cutover.
- my_tier0 is the name of the NSX-T tier-0 gateway.

Example 3: Configuration File

Following figure shows an NSX-V environment that is configured in Topology 4 (One-Armed Load Balancer). In this example, the ESG that is attached to the NSX-V Logical Switch has both load balancer and DHCP server running on it. The equivalent NSX-T topology is shown to the right. Remember that Gateway DHCP server and load balancer services are preconfigured in the NSX-T topology before the Edge cutover.

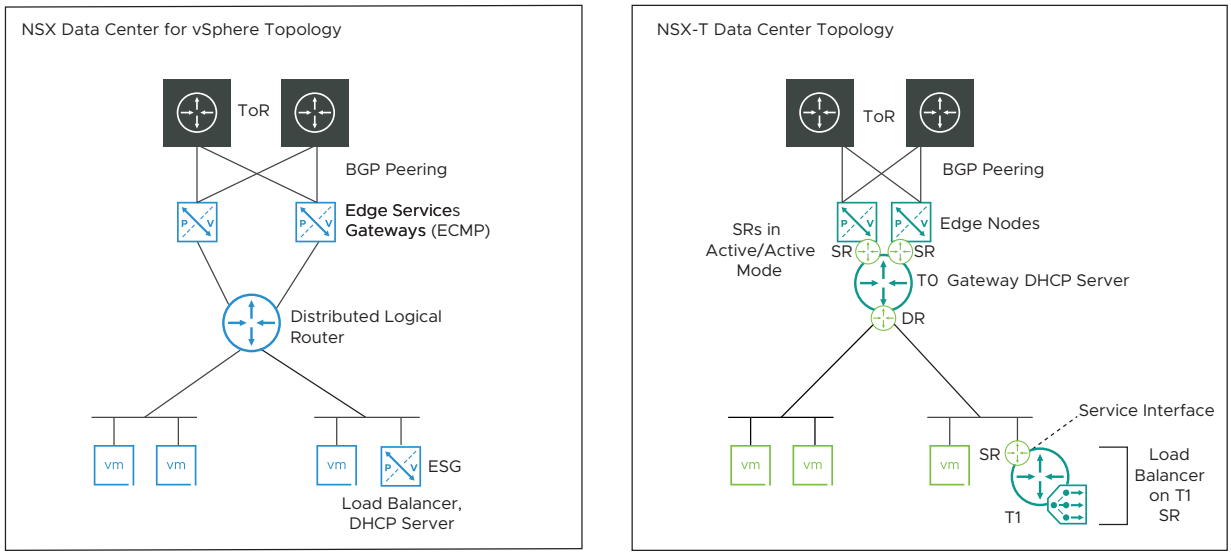

When Edge cutover occurs, the DHCP leases on the ESG that is attached to the Logical Switch are migrated to the Gateway DHCP server on the tier-0 gateway.

The desired mapping in the input configuration file is as follows:

```
[
    {
        "name": "ns-edge-cutover",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-1",
                    "edge-2"
                ],
                "policy_gateway_name": "my_tier0"
            }
        ]
    }
]
```

In this configuration file:

- edge-1 is the Edge ID of the Distributed Logical Router for Edge cutover.
- edge-2 is the Edge ID of the Edge Services Gateway where the DHCP service is configured. This ESG is attached to the NSX-V Logical Switch.
- my_tier0 is the name of the NSX-T tier-0 gateway.

Another alternative is to configure the NSX-T topology, as shown in the following figure. In this topology, the DHCP server profile is attached to a tier-1 gateway. The uplink of this tier-1 gateway is connected to the tier-0 gateway, and the NSX-T overlay segments are connected on the downlink of this tier-1 gateway

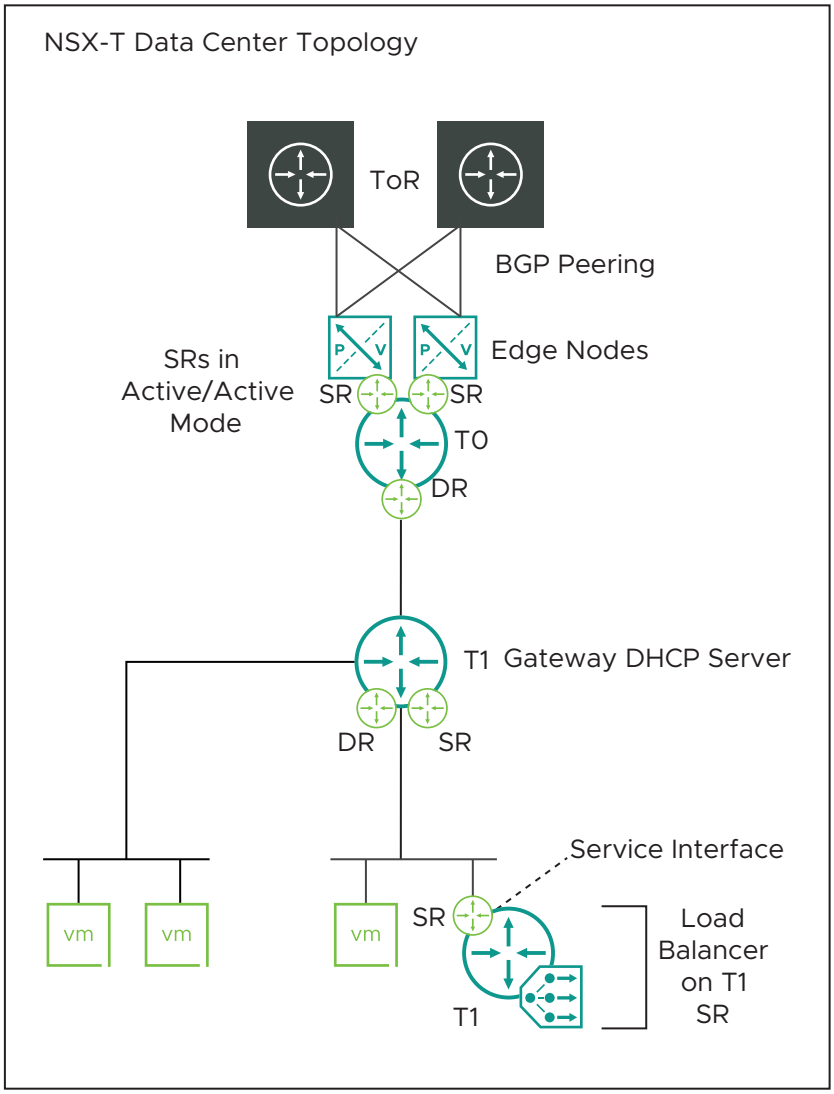

In this case, the migration coordinator migrates the DHCP leases to the Gateway DHCP server on the tier-1 gateway that is connected to the tier-0 gateway.

The desired mapping in the input configuration file is as follows:

```
[
    {
        "name": "ns-edge-cutover",
        "v_edges_to_policy_gateways_mappings": [
            {
                "v_edges": [
                    "edge-1",
                    "edge-2"
```

```
                ],
                "policy_gateway_name": "my_tier1"
            }
        ]
    }
]
```

In this configuration file:

- edge-1 is the Edge ID of the Distributed Logical Router for Edge cutover.
- edge-2 is the Edge ID of the Edge Services Gateway where the DHCP service is configured. This ESG is attached to the NSX-V Logical Switch.
- my_tier1 is the name of the NSX-T tier-1 gateway that is connected to the tier-0 gateway.

Import the NSX-V Configuration

The first step of the migration process is to import the NSX-V configuration. You must also provide a mapping of the NSX-V Edges to the NSX-T gateways.

Prerequisites

- The vCenter Server must be added as a compute manager in NSX-T.

  You can share the vCenter Server that is used in NSX-V or deploy another one in NSX-T.
- Create an input configuration file in a `.json` format.

  The migration coordinator uses the mapping information in this configuration file to do the Edge cutover and migrate the DHCP leases from NSX-V Edges to NSX-T Edges. For more information about the contents of the configuration file, see Overview of Input Configuration File.

Procedure

1. From a browser, log in to NSX Manager as **admin**.
2. Navigate to **System > Migrate**.
3. Expand the **Advanced Migration Modes** section, and in the **Edge Cutover** pane, click **Get Started**.
4. From the **Import Configuration** page, click **Select NSX** and provide the credentials for vCenter Server and NSX-V.

   Note The drop-down menu for vCenter Server displays all vCenter Server systems that are registered as compute managers. Click **Add New** if you need to add a compute manager.

5 Click **Select File** and browse to the `.json` configuration file that provides the mapping of the NSX-V Edges to the NSX-T gateway. Click **Upload**.

   After uploading a file, if necessary, you can click **Select File** again and upload a different file. The previous file is overwritten. You cannot remove the `.json` file after it is uploaded in the migration coordinator. The migration coordinator removes this configuration file from the NSX Manager appliance only when you take any of the following actions:

   - Roll back the migration on the **Import Configuration** page of the migration coordinator.
   - Finish the migration on the **Migrate Edges** page.

6 Click **Start** to import the configuration.

7 When the import has finished, click **Continue** to proceed to the **Resolve Configuration** page.

   If the following error situations occur, import configuration can fail at the translate configuration phase:

   - Syntax errors are found in the `.json` file.
   - The `.json` file does not conform to a valid JSON schema.
   - Incorrect NSX-T gateway name or NSX-V Edge IDs are mentioned in the configuration file. For example, the NSX-T gateway name that is mentioned in the `.json` file is not already realized in your NSX-T environment.

   When any of these error situations occur, roll back the migration, resolve the error, and import the configuration again. Apart from the UI displaying the errors, the error messages are also stored in the migration log file (`cm.log`) on the NSX Manager appliance where the migration coordinator service is running.

   You can find this log file at `/var/log/migration-coordinator/v2t`.

What to do next

When the NSX-V topology is imported successfully, the **View Imported Topology** link is enabled. Click this link to view a graph of the imported topology. However, the topology viewer does not display the graph of a large-scale NSX-V environment.

Resolve Configuration

After you have imported the configuration from your NSX-V environment, you must review and resolve the reported configuration issues before you can continue with the migration.

On the **Resolve Configuration** page, two types of configuration issues are reported.

Blocking Issues (if any)

As the name suggests, these issues block the migration, and they must be fixed for the migration to proceed. You might have to change the configurations in your NSX-V environment before you can migrate to NSX-T.

Warnings

These configuration issues are organized into several categories, and each category can have one or more configuration items. You should provide inputs to fix the warning messages that are displayed for the configuration items, or choose to skip the warnings, if needed.

Procedure

1 On the **Resolve Configuration** page, if you find any blocking issues, fix them in the NSX-V environment before you can proceed with the migration.

After making the required changes in the NSX-V environment, return to the migration coordinator. Go to the **Import Configuration** page, and click **Rollback**. Click **Start** to import the updated NSX-V configuration.

If you did not find any blocking issues that require a change in the NSX-V environment, proceed to the next step.

2 Review the warnings and issues reported in each category.

3 Click each issue and provide input.

Each issue can cover multiple configuration items. For each item there might be one or more possible resolutions to the issue, for example, skip, configure, or select a specific value.

For issues that apply to multiple configuration items, you can provide input for each item individually, or select all and provide one answer for all items.

4 After the input is provided, a **Submit** button is displayed on the **Resolve Configuration** page. Click **Submit** to save your progress.

5 When you have provided input for all configuration issues, click **Submit**.

The input is validated. You are prompted to update any invalid input. Additional input might be required for some configuration items.

6 After you have submitted all requested input, click **Continue** to proceed to the **Migrate Edges** page.

Example: Configuration Issues

Migration coordinator shows warning messages if it detects the following issues in your NSX-T environment:

- A missing NSX-T VLAN segment corresponding to the Distributed Virtual Port Group (DVPG) VLAN ID in your NSX-V environment.
- A missing NSX-T overlay segment corresponding to the VNI of the Logical Switch in the NSX-V environment.

To resolve this configuration issue, you can either skip creating the missing NSX-T segments or roll back the migration. If you choose to roll back the migration, create the missing segments in NSX-T by running the following PATCH API request, and start the migration again.

```
PATCH https://{policy-manager}/policy/api/v1/infra/segments/{segment-id}?
force=true
```

For example, to create a segment "App" with overlay ID "5001", the API URL and the payload of the API request is as follows:

```
{ "type" : "ROUTED",
  "subnets" : [ {" "gateway_address" : "172.16.20.1/24",
    "network" : "172.16.20.0/24""
    "} ],
  "connectivity_path" : "/infra/tier-0s/T0-GW-01",
  "transport_zone_path" : "/infra/sites/default/enforcement-points/default/transport-zones/
   1b3a2f36-bfd1-443e-a0f6-4de01abc963e",
  "overlay_id":5001,
  "admin_state" : "UP",
  "replication_mode" : "MTEP",
  "resource_type" : "Segment",
  "id" : "App",
  "display_name" : "App",
  "path" : "/infra/segments/App",
  "relative_path" : "App",
  "parent_path" : "/infra" }
```

If you choose to skip creating the missing NSX-T VLAN or overlay segments, the VM vNICs that are connected to the DVPG VLANs or the Logical Switches in the NSX-V environment lose networking when you move the workload VMs to the NSX-T environment. In this case, the DFW rules and Security Groups that are associated with the skipped segments are not applied in the NSX-T environment.

For example, you might not want to create the missing segments in the following situations:

- Port groups on the vSphere Distributed Switch (VDS) have only VMkernel ports. So, no requirement to create VLAN segments.
- Port groups from another VDS that you do not want to migrate.
- NSX-V Logical Switch with no workload VMs attached to it. For instance, a transit Logical Switch between an Edge Services Gateway and the Distributed Logical Router.

Migrate NSX-V Edges

During the **Migrate Edges** step, north-south traffic is cut off from the NSX-V Edge Services Gateways and traffic passes through the NSX-T Edges.

Prerequisites

- All NSX-V logical configurations including routing and Layer 3 networking service configurations must either be migrated or manually created in NSX-T before the Edge cutover.
- Verify the migrated or manually created configurations in the NSX-T NSX Manager UI, or use the NSX-T APIs to verify the logical configurations.
- The DPDK fast path uplink interfaces (fp-eth0, fp-eth1, and fp-eth2) on the NSX-T Edge nodes must be down.

- All configuration issues must be resolved.
- Verify that either you have either created an IP pool for Edge Tunnel End Points (TEP) or configured static IP addresses for the Edge TEPs. For information about creating an IP pool, see Create an IP Pool for Edge Tunnel End Points.

Procedure

1 From the **Migrate Edges** page, click **Start**.

Caution North-south traffic is temporarily interrupted during the **Migrate Edges** step.

The migration coordinator takes the following actions during the Edge cutover:

- The uplinks on the NSX-V Distributed Logical Router are internally disconnected from the Transit Logical Switch to which the Edge Services Gateways are also connected. The uplink interfaces on the NSX-T Edge nodes are brought online. The DPDK fast path uplink interfaces on the NSX-T edges (fp-eth0, fp-eth1, and fp-eth2) are enabled (admin_status: up). All north-south traffic that was previously passing through the Edge Services Gateways now moves through the NSX-T Edges.
- If DHCP service was configured on NSX-V Edge Services Gateway before the migration, the DHCP leases of the workload VMs are migrated to the NSX-T Edge.
- Scenario: Let us consider that in the NSX-V environment, a DHCP relay service is configured on the Distributed Logical Router (DLR). The uplink of the DLR is connected to an ESG and a DHCP server is configured on the ESG. Before the Edge cutover, this relay service forwards the DHCP requests to the DHCP server on the ESG. That is, before the Edge cutover, the relay service contains the DHCP server IP address that is configured on the ESG. When an Edge cutover occurs, the migration coordinator automatically updates the server IP in the DLR relay with the DHCP server IP in the DHCP profile that is attached to the NSX-T tier-0 or tier-1 gateway. In other words, after the Edge cutover, the DHCP requests are relayed to the Gateway DHCP server in the NSX-T environment.

If necessary, you can roll back the **Migrate Edges** step. When you roll back the migration, the north-south traffic switches back to the NSX-V ESGs. In addition, the migration coordinator brings down the fast path interfaces on the NSX-T Edge nodes. The DHCP leases are removed from the NSX-T Edges, and the workload VMs start receiving the DHCP leases from the DHCP server on the ESG.

2 Click **Finish**.

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page.

What to do next

1 Verify that the north-south traffic is flowing successfully. Run `ping` commands on the NSX-V workload VMs to check connectivity to the machines in the external physical network.

2 If DHCP service is running on NSX-T Edges, log in to the NSX-T Edge CLI, and run the `get dhcp leases` command. Observe that the DHCP leases are retained on the NSX-V workload VMs.

Migrating vSphere Networking

9

For environments with vSphere 6.X you can migrate an existing vSphere Distributed Switch (VDS) configuration to an NSX-Tenvironment backed by NSX Virtual Distributed Switch (N-VDS).

For VDS 6.5.0 and 6.6.0, this migration will move VDS, compute hosts, PNICs, vmkNICs, and vNIC backings to N-VDS.

Note For VDS 7.0, vSphere networking to N-VDS migration is not supported. You must perform a fresh install of NSX-T and configure it for use with your vSphere deployment using VDS 7.0.

You can migrate VDS configurations to NSX-T only if NSX-V is not installed on the host.

Read the following topics next:

- Overview - Migrating vSphere Networking
- Understanding the vSphere Networking Migration
- Preparing to Migrate vSphere Networking
- Migrate vSphere Networking to NSX-T

Overview - Migrating vSphere Networking

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.

- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.
- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.
- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 9-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	

Table 9-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. However, if either the primary or secondary NSX Manager is set to a standalone or transit mode, the migration is supported.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	

NSX-V Configuration	Supported	Details
Network I/O Control (NIOC) having vNIC with reservation	No	
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	
Host profiles	No	
ESXi lockdown mode	No	Not supported in NSX-T.
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.

NSX-V Configuration	Supported	Details
VXLAN port number other than 4789	No	
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	

NSX-V Configuration	Supported	Details
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction

NSX-V Configuration	Supported	Details
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network

NSX-V Configuration	Supported	Details
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.

NSX-V Configuration	Supported	Details
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	

NSX-V Configuration	Supported	Details
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 9-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 9-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.
DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration. For example, dhcp option '80' is not migrated. <dhcpOptions> <other> <code>80</code> <value>2f766172</value> </other> </dhcpOptions>
Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 9-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall

NSX-V Configuration	Supported	Details
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.

NSX-V Configuration	Supported	Details
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	

NSX-V Configuration	Supported	Details
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 9-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of ‘Groups’ after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 9-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.

Table 9-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.

Table 9-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 9-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 9-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 9-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 9-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000

Table 9-10. Single-Site Limits (continued)

Feature	Limit
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 9-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Understanding the vSphere Networking Migration

You can migrate one vSphere Distributed Switch at a time to NSX-T.

Overview of Migration Process

During the migration you will complete the following steps:

- Prepare your NSX-T environment.
  - Configure a compute manager in the NSX-T environment.
  - Add the vCenter Server system that manages the vSphere Distributed Switch (versions 6.5.0 and 6.6.0) you want to migrate.
  - Start the migration coordinator service.
- Import configuration from vSphere.
  - Enter the details of your vSphere environment.
  - The configuration is retrieved and pre-checks are run.
- Select the vSphere Distributed Switch that you want to migrate.
- Resolve issues with the configuration.

  Provide answers to configuration questions that must be resolved before you can migrate your vSphere environment to NSX-T. Resolving issues can be done in multiple passes by multiple people.
- Migrate configuration.
  - After all configuration issues are resolved, you can import the configuration to NSX-T. Configuration changes are made on NSX-T, but no changes are made to the vSphere environment yet.

- Migrate Hosts.
  - NSX-T software is installed on the hosts. VM interfaces are disconnected from vSphere Distributed Switch port groups and connected to the new NSX-T segments.

    Caution If you select **In-Place** migration mode, there is a traffic interruption during the Migrate Hosts step. However, if you select **Maintenance** migration mode, traffic interruption does not occur.
- Finish Migration.
  - After you have verified that the migrated networking is working correctly, you can click **Finish** to clear the migration state. You can now migrate another vSphere Distributed Switch to NSX-T.

Preparing to Migrate vSphere Networking

Within certain limitations, you can migrate vSphere Distributed Switches that are not part of an NSX-V environment.

Required Software and Versions

- See the *VMware Product Interoperability Matrices* for required versions of vCenter Server and ESXi: http://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php#interop&175=&1=&2=
- vSphere Distributed Switch version 6.5.0 and 6.6.0 are supported.

Add a Compute Manager

To migrate a vSphere Distributed Switch, you must configure the associated vCenter Server system as a compute manager in NSX-T before you can start the migration process.

Procedure

1. From a browser, log in with admin privileges to NSX Manager at `https://<nsx-manager-ip-address>`.
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 vCenter Server. You can optionally describe any special details such as, the number of clusters in the vCenter Server.
FQDN or IP Address	Type the FQDN or IP address of the vCenter Server.
Type	The default compute manager type is set to vCenter Server.

Option	Description
HTTPS Port of Reverse Proxy	The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances. Set the reverse proxy port to register the compute manager in NSX-T Data Center.
Username and Password	Type the vCenter Server login credentials.
SHA-256 Thumbprint	Type the vCenter Server SHA-256 thumbprint algorithm value.
Create Service Account	Enable this field for features such as vSphere Lifecycle Manager that need to authenticate with NSX-T Data Center APIs. Log in with the administrator@vsphere.local credential to register a compute manager. After registration, the compute manager creates a service account. **Note** Service account creation is not supported on a global NSX Manager. If service account creation fails, the compute manager's registration status is set to `Registered with errors`. The compute manager is successfully registered. However, vSphere Lifecycle Manager cannot be enabled on NSX-T Data Center clusters. If a vCenter Server admin deletes the service account after it was successfully created, vSphere Lifecycle Manager tries to authenticate the NSX-T Data Center APIs and the compute manager's registration status is set to `Registered with errors`.
Enable Trust	Enable this field to establish trust between NSX-T Data Center and compute manager, so that services running in vCenter Server can establish trusted communication with NSX-T Data Center. For example, for vSphere Lifecycle Manager to be enabled on NSX-T Data Center clusters, you must enable this field. Supported only on vCenter Server 7.0 and later versions.
Access Level	Enable one of the options based on your requirement: ■ Full Access to NSX: Is selected by default. This access level gives the compute manager complete access to NSX-T Data Center. Full access ensures vSphere for Kubernetes and vSphere Lifecycle Manager can communicate with NSX-T Data Center. The vCenter Server 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-T Data Center. The vCenter Server 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-T Data Center to discover and register the vCenter Server 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 vCenter Server credentials and click **Resolve**.

   If an existing registration exists, it will be replaced.

Results

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

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

After the vCenter Server is successfully registered, do not power off and delete the NSX Manager VM without deleting the compute manager first. Otherwise, when you deploy a new NSX Manager, you will not be able to register the same vCenter Server again. You will get the error that the vCenter Server is 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, NSX Intelligence VM, 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, any NSX Intelligence VM, 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 sucessfully remove the VC, ensure that you re-deploy an NSX Manager node from the VC.

This restriction applies to a fresh installation of NSX-T Data Center as well as an upgrade.

Tag Management VMs in a Collapsed Cluster

You can migrate a vSphere environment that uses a collapsed cluster.

In a collapsed vSphere cluster, all NSX-T management VMs and workload VMs must be initially attached to dvPortgroups. After migration, the management VMs will be attached to the NSX-T VLAN segments.

The management VMs in the NSX-T include appliances such as NSX Manager, vCenter Server, VMware Identity Manager, and so on. The NSX-T VLAN segment ports to which these management VMs connect are blocked in two cases: when these management VMs are rebooted after they were migrated in-place by the Migration Coordinator, or when they are moved from NSX-V hosts to NSX-T hosts by vMotion in maintenance migration mode. Therefore, the management VMs might lose connectivity in such cases.

To prevent this problem, create a "management_vms" tag category, and add tags in this category. Assign a tag from this category to all the management VMs in the NSX-T environment. These VMs will be attached to unblocked VLAN segment ports.

Procedure

1 Log in to the vSphere Client.

2 Click **Menu > Tags & Custom Attributes**.

3 Click **Categories**, and then click **New** to add a category.

Create a category with name **`management_vms`**.

4 Click the **Tags** tab and add a tag in the management_vms category.

5 Navigate to **Menu > Hosts and Clusters**.

6 Expand the collapsed cluster from the left Navigator view, right-click the name of the NSX Manager VM, and select **Tags & Custom Attributes > Assign Tag**.

7 Assign a tag from the management_vms category to the NSX Manager VM.

8 Repeat steps 6 and 7 for all the management VMs in the cluster.

For a detailed information about tag categories and tags, see the *vCenter Server and Host Management* documentation.

9 Log in to NSX Manager.

10 Navigate to **Inventory > Groups** and create a group.

11 Click **Set** to set members.

12 Set **Group Type** to **IP Addresses Only** and include the IP addresses of the management VMs that will be migrating from the dvPortgroups to the NSX-T segments.

13 Navigate to **Security > Distributed Firewall > Actions > Exclusion List** and add this group.

Migrate vSphere Networking to NSX-T

After you complete the preparations, you can migrate vSphere networking to NSX-T.

Import the vSphere Networking Configuration

The first step of the migration process is ro import the vSphere networking configuration.

Prerequisites

- Verify that the vCenter Server system associated with the vSphere Distributed Switch you want to migrate is registered as a compute manager. See Add a Compute Manager.

Procedure

1. From a browser, log in to the NSX Manager node which is running the migration coordinator service. Log in using an account with admin privileges.
2. Navigate to **System > Migrate**.
3. On the **Migrate vSphere Networking** pane, click **Get Started**.
4. From the **Import Configuration** page, click **Select vSphere** and provide the requested information about your vSphere environment.

   Note The drop-down menu for vCenter displays all vCenter Server systems that are registered as compute managers. Click **Add New** if you need to add a compute manager.
5. Click **Start** to import the configuration.
6. When the import has finished, click **Continue** to proceed to the **Resolve Issues** page.

Roll Back the vSphere Networking Migration

After you have started the migration process, you can roll back the migration to undo some or all of your progress.

You can roll back or undo the migration from some of the migration steps. After the migration has started, you can click **Rollback** on the furthest step completed. The button is disabled on all other pages.

Table 9-12. Rolling Back vSphere Networking Migration

Migration Step	Rollback Details
Import Configuration	Click **Rollback** on this page to roll back the Import Configuration step.
Resolve Configuration	Rollback is not available here. Click **Rollback** from the **Import Configuration** page.

Table 9-12. Rolling Back vSphere Networking Migration (continued)

Migration Step	Rollback Details
Migrate Configuration	Click **Rollback** on this page to roll back the migration of the configuration to NSX-T and the input provided on the **Resolve Configuration** page.
Migrate Hosts	Rollback is not available here.

Resolve Issues with the vSphere Networking Configuration

After you have imported the networking configuration from your vSphere environment, you must review and resolve the reported configuration issues before you can continue with the migration.

You must provide feedback for all configuration issues that must be resolved before the migration can continue. Multiple people can provide the feedback over multiple sessions. After you provide feedback for a given issue, you can click **Submit** to save it. You can return to a submitted input and modify it.

After you have submitted feedback for all issues, the feedback is validated. The validation might result in additional requests for feedback before the migration can proceed.

Procedure

1 From the **Resolve Configuration** page, click **Select Switch** to select which vSphere Distributed Switch to migrate.

   Once a distributed switch is selected, the configuration issues are displayed.

2 Review the reported issues.

   Issues are organized into groups. Each issue can cover multiple configuration items. For each item there might be one or more possibly resolutions to the issue, for example, skip, configure, or select a specific value.

3 Click each issue and provide feedback.

   For issues that apply to multiple configuration items, you can provide feedback for each individually, or select all and provide one answer for all items.

   Multiple people can provide the input over multiple sessions. You can return to a submitted input and modify it.

4 After some feedback has been provided, a **Submit** button appears on the **Resolve Issues** page. Click **Submit** to save your progress.

5 When you have provided feedback for all configuration issues, click **Submit**.

   The input is validated. You are prompted to update any invalid input. Additional input might be required for some configuration items.

6 After you have submitted all requested feedback, click **Continue** to proceed to the Migrate Configuration step.

Migrate vSphere Networking Configuration

After you have resolved all configuration issues, you can migrate the vSphere networking configuration. Configuration changes are made in the NSX-T environment to replicate the translated vSphere configuration.

If needed, you can roll back the configuration migration. This will do the following:

- Remove the migrated configuration from NSX-T.
- Roll back all the resolved issues in the previous step.

See Roll Back the vSphere Networking Migration for more information.

Prerequisites

Verify you have completed the **Resolve Configuration** step.

Procedure

- From the **Migrate Configuration** page, click **Start**.

  The distributed switch configuration is migrated to NSX-T.

Configuring vSphere Host Migration

The clusters in the vSphere environment are displayed on the **Migrate Hosts** page. The clusters are arranged into migration groups, each migration group contains one vSphere host cluster. There are several settings which control how the host migration is performed.

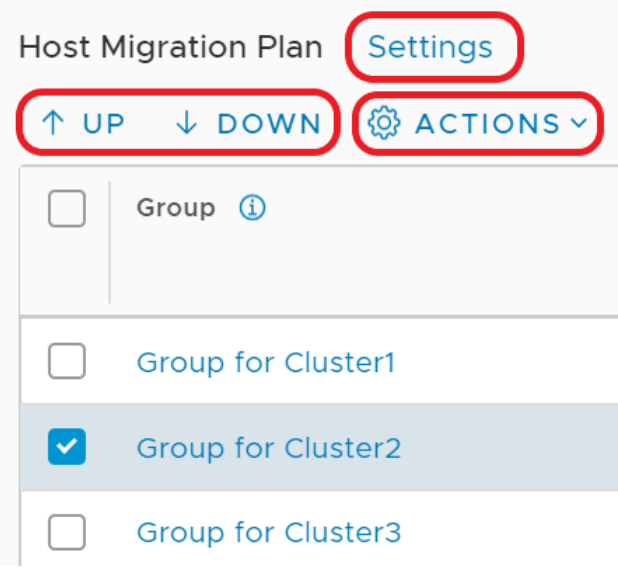

- Click **Settings** to change the global settings: **Pause Between Groups** and **Migration Order Across Groups**.
- Select a single host group (cluster) and use the arrows to move it up or down in the migration sequence.

- Select one or more host groups (clusters) and click **Actions** to change these host groups settings: **Migration Order Within Groups**, **Migration State**, and **Migration Mode**.

Pause Between Groups

When migrating multiple host groups, you can pause the migration between groups by enabling the **Pause Between Groups** setting. After a group is migrated, you must click **Continue** to migrate the next host group. This setting is disabled by default. You can enable it if you want to verify the status of the applications running on each cluster before proceeding to the next one.

Serial or Parallel Migration Order

You can define whether migration happens in a serial or parallel order. There are two ordering settings:

- **Migration Order Across Groups** is a global setting that applies to all host groups.
  - **Serial**: One host group (cluster) at a time is migrated.
  - **Parallel**: Up to five host groups at a time are migrated. After those five host groups are migrated, the next batch of up to five host groups are migrated.

    Important For migrations involving vSphere Distributed Switch 7.0, do not select parallel migration order across groups.
- **Migration Order Within Groups** is a host group (cluster) specific setting, so can be configured separately on each host group.
  - **Serial**: One host within the host group (cluster) at a time is migrated.
  - **Parallel**: Up to five hosts within the host group are migrated at a time. After those hosts are migrated, the next batch of up to five hosts are migrated.

    Important Do not select parallel migration order within groups for a cluster if you plan to use **Maintenance** migration mode for that cluster.

By default, both settings are set to **Serial**. Together, the settings determine how many hosts are migrated at a time.

Table 9-13. Effects of Migration Settings on Number of Hosts Attempting Migration Simultaneously

Migration Order Across Groups (Clusters)	Migration Order Within Groups (Clusters)	Maximum Number of Hosts Attempting Migration Simultaneously
Serial	Serial	1 One host from one host group
Serial	Parallel	5 Five hosts from one host group

Table 9-13. Effects of Migration Settings on Number of Hosts Attempting Migration Simultaneously (continued)

Migration Order Across Groups (Clusters)	Migration Order Within Groups (Clusters)	Maximum Number of Hosts Attempting Migration Simultaneously
Parallel	Serial	5 One host from five host groups
Parallel	Parallel	25 Five hosts from five host groups

Important If there is a failure to migrate a host, the migration process will pause after all in-progress host migrations have finished. If **Parallel** is selected for both migration across groups and migration within groups, there might be a long outage for the failed host before you can retry migration.

Sequence of Migration Groups

You can select a host group (cluster) and use the arrows to move it up or down in the list of groups.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

Migration State

Host groups (clusters) can have one of two migration states:

- **Enabled**

  Hosts groups with a migration state of **Enabled** are migrated to NSX-T when you click **Start** on the **Migrate Hosts** page.

- **Disabled**

  You can temporarily exclude host groups from migration by setting the migration state for the groups to **Disabled**. Hosts in disabled groups are not migrated to NSX-T when you click **Start** on the **Migrate Hosts** page. However, you must enable and migrate all **Disabled** host groups before you can click **Finish**. Finish all host migration tasks and click **Finish** within the same maintenance window.

Migration Mode

Migration Mode is a host group (cluster) specific setting, and can be configured separately on each host group. In the **Migrate Hosts** step, you select whether to use **In-Place** or **Maintenance** mode.

There are two types of Maintenance migration modes:

- Automated
- Manual

In the **Resolve Configuration** step of the migration process, you select which type of Maintenance migration mode to use. You select a Maintenance mode even if you plan to migrate hosts using **In-Place** mode. When you select Maintenance migration mode in the **Migrate Hosts** step, the value you specified in the **Resolve Configuration step** determines whether Automated Maintenance mode or Manual Maintenance mode is used. However, in the **Migrate Hosts** step, if you select **In-Place** mode, your selected choice of Maintenance mode in the **Resolve Configuration** step does not take effect.

- **In-Place** migration mode

  NSX-T is installed and hosts are migrated while VMs are running on the hosts. Hosts are not put in maintenance mode during migration. Virtual machines experience a short network outage and network storage I/O outage during the migration.

- **Automated Maintenance** migration mode

  A task of entering maintenance mode is automatically queued. VMs are moved to other hosts using vMotion. Depending on availability and capacity, VMs are migrated to vSphere or NSX-T hosts. After the host is evacuated, the host enters maintenance mode, and NSX-T is installed. VMs are moved back to the newly configured NSX-T host. Note that Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

- **Manual Maintenance** migration mode

  A task of entering maintenance mode is automatically queued. To allow the host to enter maintenance mode, do one of the following tasks:

  - Power off all VMs on the hosts.
  - Move the VMs to another host using vMotion or cold migration.

  Once the host is in maintenance mode, NSX-T is installed on the host. After the host is migrated, for the powered-off VMs and the VMs that you moved, you will need to change their network connection from the NSX-V logical switch to an NSX-T segment.

In the NSX-V environment, if the ESXi host's vmk0 management interface is connected to a VSS (vSphere Standard Switch) portgroup that does not have an uplink, and the portgroup is bridged to a VDS portgroup, and the VDS version is 6.5, 6.6 or 6.7, you must migrate using the **Maintenance** mode. If you use the **In-Place** mode, the migration will fail.

Migrate vSphere Hosts

After you have migrated the configuration, you can migrate the vSphere hosts to NSX-T.

You can configure several settings related to the host migration, including migration order and enabling hosts. Before you change any default settings, make sure that you understand the effects of these settings. See Configuring vSphere Host Migration for more information.

Caution There is a traffic interruption during the host migration. Perform this step during a maintenance window.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

If migration fails for a host, the migration pauses after all in-progress host migrations finish. When you have resolved the problem with the host, click **Retry** to retry migration of the failed host. If the host still fails to migrate, you can configure NSX-T on the host manually or remove the host from the system. In this case, at the end of the host migration step, the **Finish** button will not be enabled because of the host that failed to migrate. You need to call the REST API `POST https://<nsx-mgr-IP>/api/v1/migration?action=finalize_infra` (<nsx-mgr-IP> is the IP address of the NSX Manager where the migration service is running) using a REST API client (for example, postman or curl) to finish the migration, and then perform the post-migration tasks.

Prerequisites

- Verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering maintenance mode.

Procedure

1 On the **Migrate Hosts** page, click **Start**.

  If you selected the **In-Place** or **Automated Maintenance** migration mode for all hosts groups, the host migration starts. Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

2 If you selected the **Manual Maintenance** migration mode for any host groups, you must complete one of the following tasks for each VM so that the hosts can enter maintenance mode.

Option	Action
Power off or suspend VMs.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b After the host has migrated, attach the VM interfaces to the appropriate NSX-T segments and power on the VM.
Move VMs using vMotion.	Right click the VM and select Migrate. Follow the prompts to move the VM to a different host. Note that Migration Coordinator maintains security during migration by vMotioning VMs to specific ports that are protected by temporary rules. In the case of manual vMotion, the VMs will not be moved to those ports and there could be a security breach. To vMotion manually, the VMs must be migrated using vSphere API where the networking backing must point to the OpaqueNetwork ID corresponding to the NSX Segment when using NVDS or the VDS portgroup ID when using VDS 7. In both cases, the network device's externalId must be set to the the string "VM_UUID:vNIC_ID", where VM_UUID is the VM's instance UUID and vNIC_ID is the VM's vNIC index where the first vNIC is 4000.
Move VMs using cold migration.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b Right click the VM and select Migrate. Follow the prompts to move the VM to a different host, connecting the VM interfaces to the appropriate NSX-T segments.

Results

After a host has migrated to NSX-T using **In-Place** migration mode, you might see a critical alarm with message `Network connectivity lost`. This alarm occurs when a vSphere Distributed Switch (VDS) 6.5 or 6.7 migrates to an N-VDS because the host no longer has a physical NIC connected to the VDS it was previously connected to. To restore the migrated hosts to the Connected state, click **Reset to Green** on each host, and suppress the warnings, if any.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

If migration fails for a host, the migration pauses after all in-progress host migrations finish. When you have resolved the problem with the host, click **Retry** to retry migration of the failed host. If the host still fails to migrate, you can configure NSX-T on the host manually or remove the host from the system. In this case, at the end of the host migration step, the **Finish** button will not be enabled because of the host that failed to migrate. You need to call the REST API `POST https://<nsx-mgr-IP>/api/v1/migration?action=finalize_infra` (<nsx-mgr-IP> is the IP address of the NSX Manager where the migration service is running) using a REST API client (for example, postman or curl) to finish the migration, and then perform the post-migration tasks.

For information about troubleshooting other host migration problems, see Chapter 13 Troubleshooting Migration Issues.

Finish Migration

After you have migrated hosts to the NSX-Tenvironment, confirm that the new environment is working correctly. If everything is functioning correctly, you can finish the migration.

Important Verify everything is working and click **Finish** within the maintenance window. Clicking **Finish** performs some post-migration clean-up. Do not leave the migration coordinator in a unfinished state beyond the migration window.

Prerequisites

Verify that the NSX-T environment is working correctly.

Procedure

1. Navigate to the **Migrate Hosts** page of the migration coordinator.
2. Click **Finish**

   A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Issues** page.

Migrating NSX-V with vRealize Automation - Fixed Topology

10

Your organization uses NSX-V in its network environment. To meet the on-demand business needs of your organization, you have integrated and used vRealize Automation with your existing NSX-V environment.

vRealize Automation runs as the cloud management platform on top of your existing vCenter Server and NSX-V environment. Over multiple day two networking operations, you have created several deployments in vRealize Automation to deploy resources, such as networks, VMs, security groups, DHCP servers, and so on. The on-demand resources deployed through vRealize Automation also consume existing objects in your NSX-V environment.

Your organization has now decided to migrate to the modern NSX-T platform.

You have the following goals:

- Migrate the existing NSX-V environment to a new NSX-T environment.
- Migrate the vRealize Automation created resources that are integrated with your NSX-V topology to the new NSX-T.
- Maintain the integration with vRealize Automation after the migration to NSX-T.

To summarize, you want to migrate NSX-V to NSX-T when NSX-V is integrated and used with vRealize Automation.

The migration coordinator can help you meet these goals.

Read the following topics next:

- Migrate NSX-V Hosts
- Post-Migration Tasks

Overview - Migrating NSX-V with vRealize Automation

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.
- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.
- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 10-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.

Table 10-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.

NSX-V Configuration	Supported	Details
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. However, if either the primary or secondary NSX Manager is set to a standalone or transit mode, the migration is supported.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	
Network I/O Control (NIOC) having vNIC with reservation	No	
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	
Host profiles	No	

NSX-V Configuration	Supported	Details
ESXi lockdown mode	No	Not supported in NSX-T.
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.
VXLAN port number other than 4789	No	
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.

NSX-V Configuration	Supported	Details
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	

NSX-V Configuration	Supported	Details
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled

NSX-V Configuration	Supported	Details
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network

NSX-V Configuration	Supported	Details
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.

NSX-V Configuration	Supported	Details
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	

NSX-V Configuration	Supported	Details
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.

NSX-V Configuration	Supported	Details
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 10-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 10-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.

Table 10-3. DHCP Features (continued)

NSX-V Configuration	Supported	Details

DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration.
		For example, dhcp option '80' is not migrated.
		<pre><dhcpOptions>
		<other>
		<code>80</code>
		<value>2f766172</value>
		</other>
		</dhcpOptions></pre>

Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 10-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.

NSX-V Configuration	Supported	Details
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group

NSX-V Configuration	Supported	Details
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ VMs not on NSX-V hosts	No	NSX-T does not support referencing objects not connected to NSX-T portgroups or networks. Those objects will be lost from the source or destination when migration is completed. To avoid this issue use IP addresses in NSX-V to reference those objects before migration.
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.

NSX-V Configuration	Supported	Details
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	

NSX-V Configuration	Supported	Details
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 10-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of ‘Groups’ after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 10-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.

Table 10-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.

Table 10-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 10-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 10-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 10-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 10-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000

Table 10-10. Single-Site Limits (continued)

Feature	Limit
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 10-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Changes Made During Host Migration in an End-to-End Migration

During the host migration step in an end-to-end migration, changes are made to migrate NSX-V hosts to NSX-T hosts.

- NSX-V software is uninstalled.
- NSX-T software is installed.
- For vSphere Distributed Switch versions 6.5.0 and 6.6.0:

  Hosts are configured with N-VDS to replace vSphere Distributed Switches:

  - Each N-VDS is created with a name that references the distributed switch name. For example, distributed switch `ComputeSwitchA` is created as N-VDS `nvds.ComputeSwitchA`.
  - If different clusters use different distributed switches to back logical switches, an N-VDS is created with a name that combines all the distributed switch names. For example, if `ComputeCluster1` and `ComputeCluster2` use distributed switch `ComputeSwitchA` to back logical switches and `ComputeCluster3` uses `ComputeSwitchB` to back logical switches, the N-VDS is created as `nvds.ComputeSwitchA.ComputeSwitchB`.
  - PNICs and vmks in the vSphere Distributed Switch are migrated to N-VDS.
  - NSX-V VTEPs are migrated to NSX-T TEPs.
- For vSphere Distributed Switches version 7.0:

  Hosts configured for vSphere Distributed Switch version 7.0 continue using the same switch after migration.

  - PNICs and vmks in the vSphere Distributed Switch remain connected on the same vSphere Distributed Portgroups.
  - NSX-V VTEPs are migrated to NSX-T TEPs and connected to standalone ports on the same vSphere Distributed Switch.

Note: If NSX-V has multiple VTEPs and a single LAG configuration, after migration the LAG will have a single TEP and in failover mode.

Virtual Machine Deployment During an End-to-End Migration

After you start an end-to-end migration, do not change the NSX-V environment. If you want to deploy VMs during the migration, wait until some of the NSX-V hosts have migrated to NSX-T and deploy the VMs on NSX-T hosts. Connect the VMs to NSX-T segments and install VMware Tools on the VMs.

Deploying on NSX-T with VMware Tools installed ensures that the VMs are populated into security groups and receive the intended Distributed Firewall policies.

Caution VMs deployed without VMware Tools installed, or deployed on NSX-V do not receive the intended Distributed Firewall policies.

If you use vSphere templates to deploy VMs, update the templates to use NSX-T segments for the VM network configuration. Specifying NSX-T segments ensures that any VMs deployed using the templates are deployed on NSX-T hosts.

If you use automation tools to deploy VMs on vSphere, but do not use vSphere templates, you might need to change your automation tool configuration to ensure that the VMs are deployed on NSX-T.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1
	Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Understanding the Migration of NSX-V with vRealize Automation

You can migrate an NSX-V environment and its existing integration with vRealize Automation to a new NSX-T environment.

You must deploy a new NSX-T environment for this migration. You cannot merge vRealize Automation deployments from your NSX-V environment into an existing NSX-T environment, which is preconfigured and used.

In addition to setting up a new NSX-T environment in advance, migration preparation might also require you to modify your existing NSX-V environment.

You should be familiar with NSX-T concepts, vRealize Automation concepts, and administration tasks in both environments before you migrate.

Topologies Supported for Integration with vRealize Automation

You can migrate an NSX-V environment with vRealize Automation (vRA) if the topology is supported for migration.

The following resources that are created using vRA are supported:

- On-demand routed networks without services
- On-demand routed networks with DHCP server
- On-demand private networks (isolated networks)
- On-demand security groups
- On-demand outbound networks with NAT only
- On-demand outbound networks with DHCP, NAT, one-arm load balancer
- On-demand outbound networks with DHCP, NAT, inline load balancer
- On-demand outbound networks with DHCP, NAT, one-arm load balancer, inline load balancer
- On-demand one-arm load balancer on existing networks
- On-demand security groups or existing security groups with load balancer
- Existing security groups

In addition, the following pre-created objects in NSX-V that are consumed in vRA deployments are supported for migration:

- Existing networks (Logical Switches, VLAN Distributed Virtual Port Groups)
- Existing security groups
- Distributed Logical Router (DLR)

The version of vRA that you are using in your integrated environment must support the topologies mentioned below. For information about topologies that vRA supports for migration, see the vRealize documentation about migrating NSX-V to NSX-T.

Caution In topologies that contain a vRA-created Edge Services Gateway (ESG), the firewall rules are not migrated to NSX-T Gateway during migration. This is in line with vRA's behavior when it creates similar topologies directly in an NSX-T environment.

In NSX-V, the default action for Edge firewall rules is to deny all traffic and then specific rules are added by vRA on the Edge to allow traffic for the services that it configures. However, in NSX-T the default behavior on the gateway is to allow all traffic.

In the network topology diagrams that follow later in this topic, the following guidelines are used:

Before-migration topology

Objects that are pre-created or existing in NSX-V are shown in a dark gray color. The pre-created topology in NSX-V can match any one of the following four NSX-V topologies that are supported for migration:

- ESG with High Availability and L4-L7 Services (Topology 1)
- ESG with No L4-L7 Services (Topology 2)

- Two Levels of ESG with L4-L7 Services on Second-Level ESG (Topology 3)
- One-Armed Load Balancer (Topology 4)

For details about these supported topologies, see Fixed Topologies Supported for End-to-End Migration.

The on-demand routed networks, outbound networks, and private networks that are shown in orange represent the resources that were created in vRealize Automation.

After-migration topology

The existing or pre-created objects in the NSX-V topology map to NSX-T objects in a dark gray color. The on-demand resources created in vRealize Automation map to NSX-T objects in an orange color.

Note

- In all the supported topologies, vRealize Automation does not create Distributed Logical Router (DLR) or north-facing Edge Services Gateways (ESG). vRealize Automation can only consume an existing DLR in its Network Profile and use it for creating the on-demand routed networks.
- In all the supported topologies, both BGP and OSPF are supported between north-facing Edge Services Gateways and physical ToR switches.

On-Demand Routed Networks Without Services (Topology A)

This topology supports the following configurations:

- Logical Switches created in vRealize Automation connect to an existing Distributed Logical Router in NSX-V.
- Workload VMs created using vRealize Automation are attached to the Logical Switches created in vRealize Automation.

The configurations are migrated to NSX-T as follows:

- A tier-1 gateway is created for each vRealize Automation created Logical Switch that is connected to the Distributed Logical Router.
- The downlink interfaces of the Distributed Logical Router map to the downlink interfaces on the tier-1 gateway.
- An NSX-T segment is created for each vRealize Automation created Logical Switch.
- Workload VMs created in vRealize Automation are attached to the NSX-T segments.

Figure 10-1. Topology A: Before and After Migration

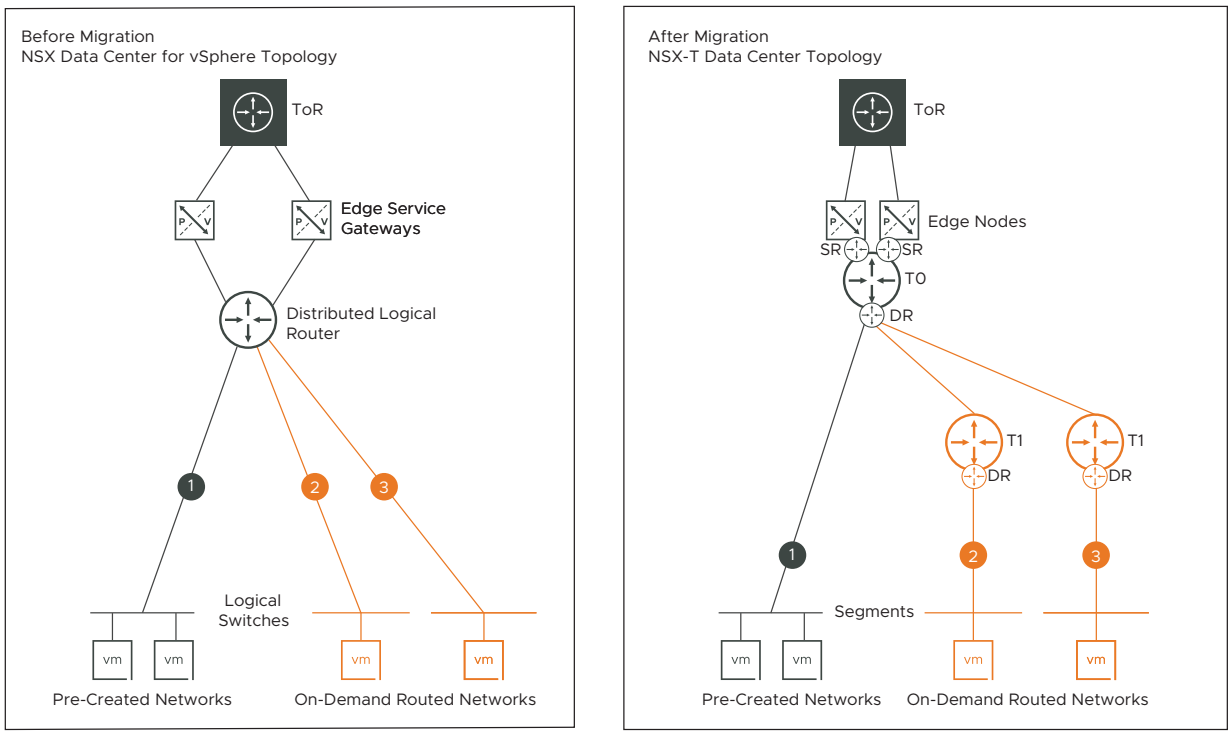

On-Demand Routed Networks with DHCP Server Only (Topology B)

This topology supports the following configurations:

- Logical Switches created in vRealize Automation connect to an existing Distributed Logical Router in NSX-V.
- An Edge Services Gateway with a DHCP server configuration is created using vRealize Automation.
- The Edge Services Gateway is attached to the vRealize Automation created Logical Switch.

The configurations are migrated to NSX-T as follows:

- A tier-1 gateway is created for each vRealize Automation created Logical Switch that is connected to the Distributed Logical Router.
- The downlink interfaces of the Distributed Logical Router map to the downlink interfaces on the tier-1 gateway.
- An NSX-T segment is created for each vRealize Automation created Logical Switch.
- Workload VMs created in vRealize Automation are attached to the NSX-T segments.
- DHCP server configuration on the Edge Services Gateway is migrated to a Local DHCP server configuration on the NSX-T segment.
- DHCP leases of the workload VMs that are attached to the vRealize Automation created Logical Switch are migrated to NSX-T.

Figure 10-2. Topology B: Before and After Migration

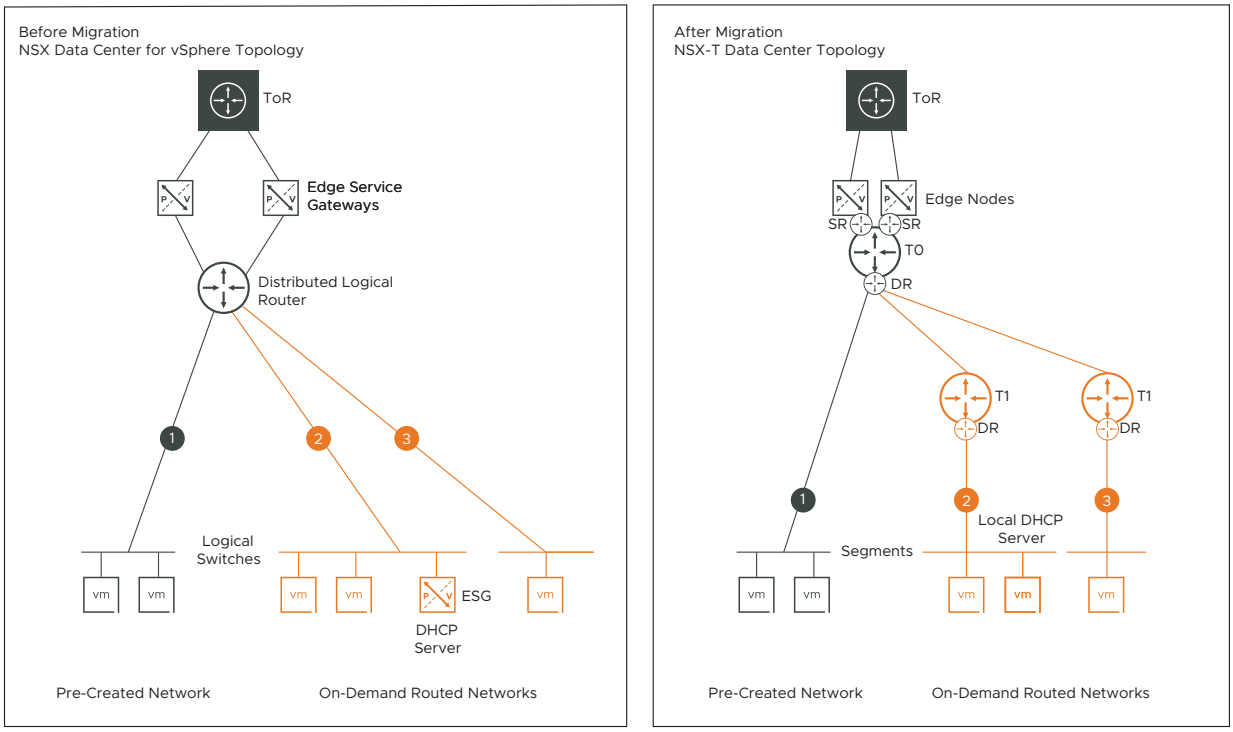

On-Demand Private Networks (Topology C)

This topology contains vRealize Automation created Logical Switches that are not connected to either the Distributed Logical Router or the Edge Services Gateway. That is, there is no routing in this topology. The isolated on-demand private networks allow Layer 2 connectivity between workload VMs on the same Logical Switch. This topology supports the following configurations:

- Workload VMs are created in vRealize Automation and attached to the vRealize Automation created Logical Switches.
- Optional: An Edge Services Gateway is created in vRealize Automation, and DHCP server is configured on this ESG. The ESG is attached to the vRealize Automation created Logical Switch.

The configurations are migrated to NSX-T as follows:

- An NSX-T segment is created for each vRealize Automation created Logical Switch.
- Workload VMs created in vRealize Automation are attached to the NSX-T segment.
- Local DHCP server is configured on the NSX-T segment.

Figure 10-3. Topology C: Before and After Migration

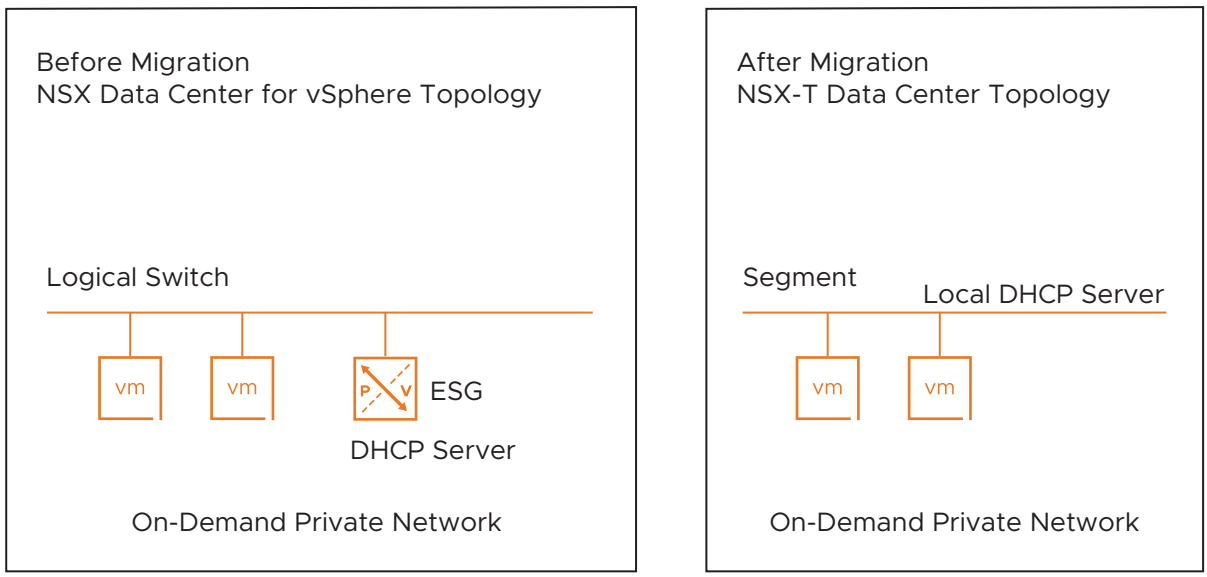

On-Demand Security Groups (Topology D)

This topology supports creating Security Groups using vRealize Automation. You can do the following configurations in the Security Groups through vRealize Automation:

- Add only VM vNICs as static members.
- Add Security Policies.
- Add firewall rules in the Security Policy. The rules can have a combination of IP addresses and applications, with or without port and protocol combinations.
- Apply Security Policies to the vRealize Automation created Security Groups.

If the firewall rules in the vRealize Automation created Security Policy contain IP addresses, vRealize Automation creates an IP set. If the rule contains port and protocol combinations, vRealize Automation creates the necessary applications in the firewall rule.

After migration, the vRealize Automation created Security Groups map to Groups in NSX-T. An L4 application in the firewall rule on NSX-V maps to an L4 service in the NSX-T firewall rule. However, if the rule contains an L7 application, the mapping can be any one of the following:

- An L7 application without port and protocol combination in an NSX-V firewall rule maps to a single Context Profile in an NSX-T firewall rule.
- An L7 application with port and protocol combination in an NSX-V firewall rule maps to a single Context Profile and a L4 service in an NSX-T firewall rule.

Existing Security Groups (Topology E)

In this topology, vRealize Automation consumes or references existing or pre-created Security Groups in the NSX-V environment. However, with existing Security Groups, limited functionality is supported. Using vRealize Automation, you can add only vNICs of workload VMs as static members in the existing Security Groups. You cannot create Distributed Firewall rules in vRealize Automation and apply them to existing Security Groups.

After migration, the Security Groups in NSX-V map to Groups in NSX-T. Security Tags created through vRealize Automation are not supported for migration to NSX-T. If you have assigned pre-created NSX-V Security Tags to workload VMs, these tags are migrated to NSX-T, but this information is not updated in vRealize Automation.

On-Demand Outbound Networks with NAT Only (Topology F)

This topology supports the following configurations:

- On-demand outbound networks (Logical Switches) are created using vRealize Automation. For example, networks 2 and 3 in the before migration topology represents the on-demand outbound networks.
- The uplink interface of the vRealize Automation created Edge Services Gateways (ESG-3 and ESG-4) are connected to an existing transit Logical Switch (192.178.14.0/24).
- The vRealize Automation created VMs on the on-demand outbound networks have a static IP address.
- Network address translation (NAT) rules are created using vRealize Automation. SNAT action is configured on the uplink interface of ESG-3 and ESG-4. The SNAT rules allow only outgoing traffic from the VMs on the outbound networks to the external clients on the public network. Port forwarding is not supported on vRealize Automation created SNAT rules.

  Example: NAT rule configuration on ESG-3 and ESG-4.

ESG-3	ESG-4
Original Action: SNAT Applied on: vNIC (uplink interface) Protocol: any Source IP range: 192.168.200.1-192.168.200.14 Source ports: any Destination IP range: any Destination ports: any	**Original** Action: SNAT Applied on: vNIC (uplink interface) Protocol: any Source IP range: 192.168.210.1-192.168.210.14 Source ports: any Destination IP range: any Destination ports: any
Translated IP address: 192.178.14.4 Port range: any	**Translated** IP address: 192.178.14.5 Port range: any

Figure 10-4. Topology F: Before Migration

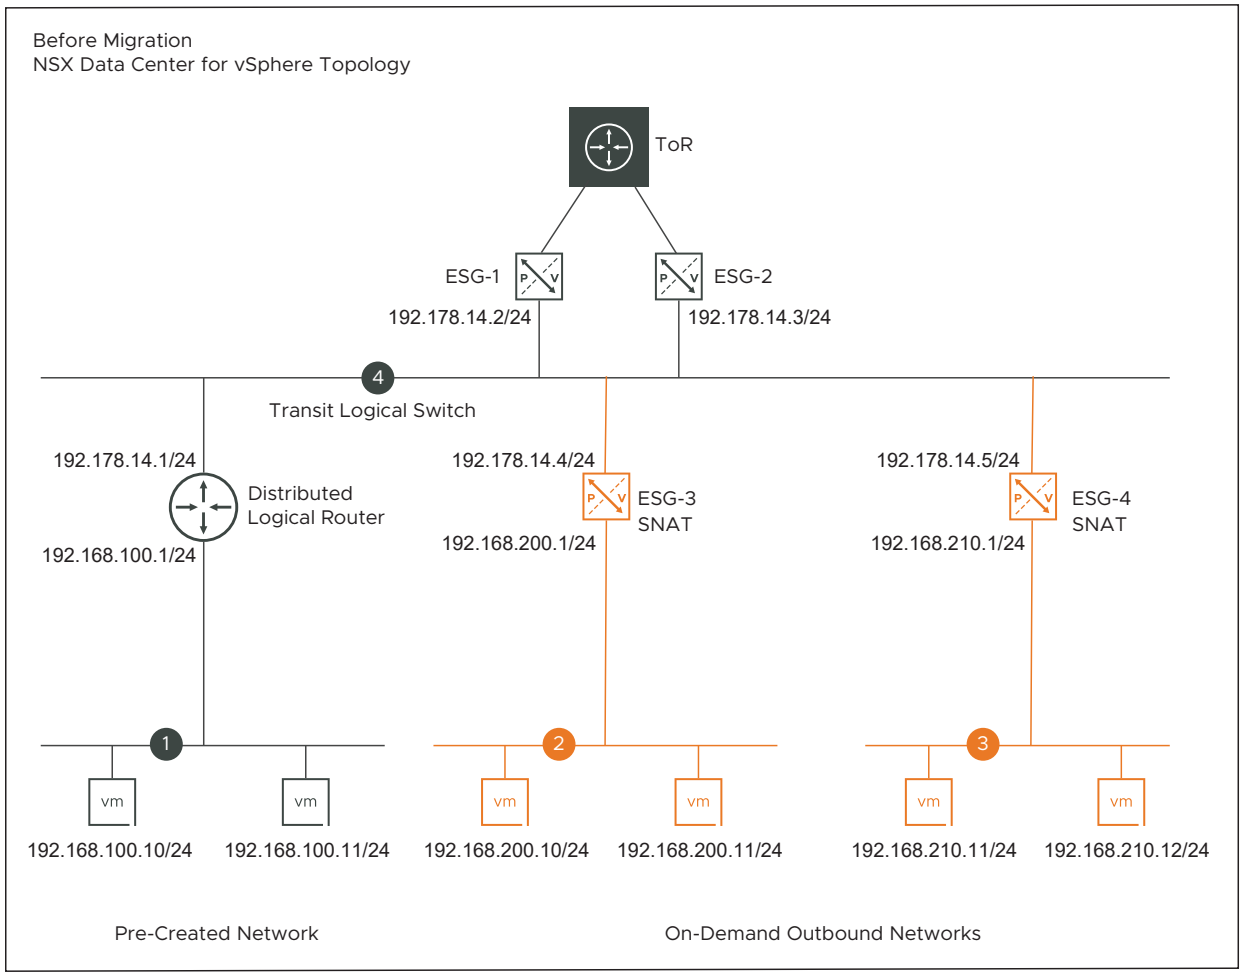

The configurations are migrated to NSX-T as follows:

- For each vRealize Automation created outbound network that is connected to an ESG, a tier-1 gateway is created and an NSX-T overlay segment is attached to the downlink of this tier-1 gateway.
- NAT rules with SNAT action are configured on the tier-1 gateways. These SNAT rules have the same configuration as the SNAT rules on the vRealize Automation created ESGs.
- Tier-1 gateways are created in the same edge cluster where the tier-0 gateway is created.

Figure 10-5. Topology F: After Migration

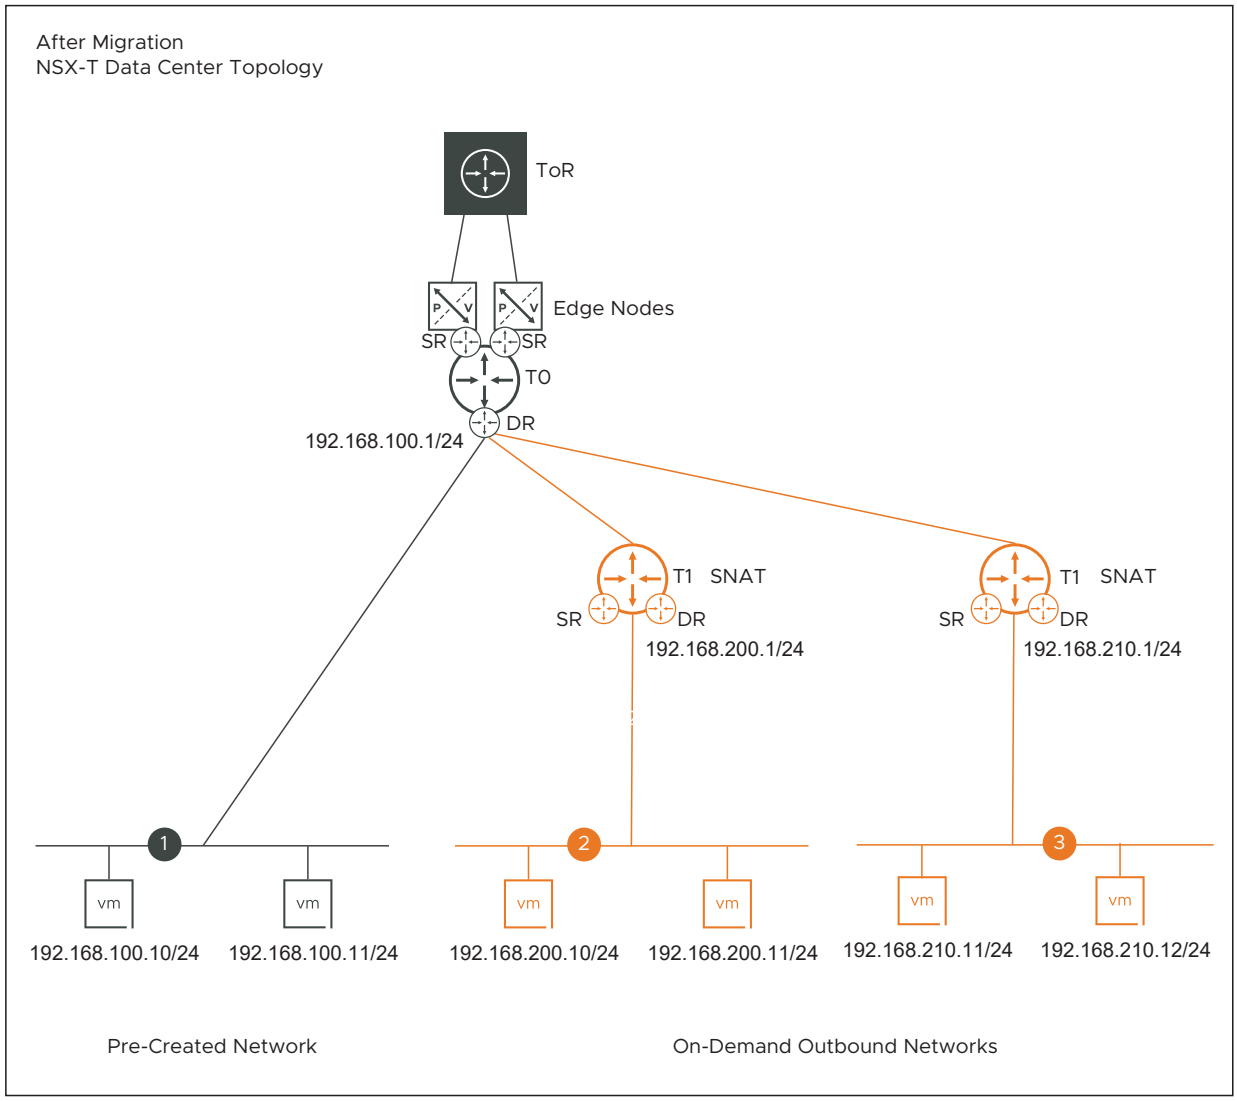

On-Demand Outbound Networks with DHCP, NAT, One-Arm Load Balancer (Topology G)

This topology supports the following configurations:

- On-demand outbound networks (Logical Switches) are created using vRealize Automation. For example, networks 2 and 3 in the before migration topology diagram represents the on-demand outbound networks.
- The uplink interface of the vRealize Automation created Edge Services Gateways (ESG-3 and ESG-4) are connected to an existing transit Logical Switch in NSX-V (for example, 192.178.14.0/24).
- ESG-3 is configured with a one-arm load balancer and NAT rules (SNAT action).
- ESG-4 is configured with a DHCP server and NAT rules (SNAT action).
- The vRealize Automation created VMs on the on-demand outbound network 2 have a static IP address.
- The vRealize Automation created workload VMs on the outbound network 3 are assigned an IP address by the DHCP server.
- Network address translation (NAT) rules are created using vRealize Automation. SNAT action is configured on the uplink interface of ESG-3 and ESG-4. SNAT rules allow only outgoing traffic from the VMs on the outbound networks to the external clients on the public network.

  Example: NAT rule configuration on ESG-3 and ESG-4.

ESG-3	ESG-4
Original	**Original**
Action: SNAT	Action: SNAT
Applied on: vNIC (uplink interface)	Applied on: vNIC (uplink interface)
Protocol: any	Protocol: any
Source IP range: 192.168.200.1-192.168.200.14	Source IP range: 192.168.210.1-192.168.210.14
Source ports: any	Source ports: any
Destination IP range: any	Destination IP range: any
Destination ports: any	Destination ports: any
Translated	**Translated**
IP address: 192.178.14.4	IP address: 192.178.14.5
Port range: any	Port range: any

- Example: One-arm load balancer configuration on ESG-3:
  - Virtual Server IP: 192.168.200.14
  - Default Pool: Pool-1
  - Pool-1 has three VMs: 192.168.200.10/24, 192.168.200.11/24, 192.168.200.12/24

  The other load balancer configuration settings are not listed here because they are not of interest in this migration example.

- Example: Internal interface of ESG-3 has two IP addresses configured: primary IP address and secondary IP address.
  - Primary IP address: 192.168.200.1/24. It connects the outbound network to the downlink (internal) interface of the ESG-3.
  - Secondary IP address: 192.168.200.14/24. It is used as the virtual server IP.

  If necessary, you can configure a single IP address on the internal interface of the ESG-3. In that case, the virtual server IP is the same as the primary IP address.
- Example: DHCP server configuration on ESG-4:
  - DHCP pool range: 192.168.210.20-192.168.210.40
  - Default gateway: 192.168.210.1

  The other DHCP server settings are not listed here because they are not of interest in this migration example.

Figure 10-6. Topology G: Before Migration

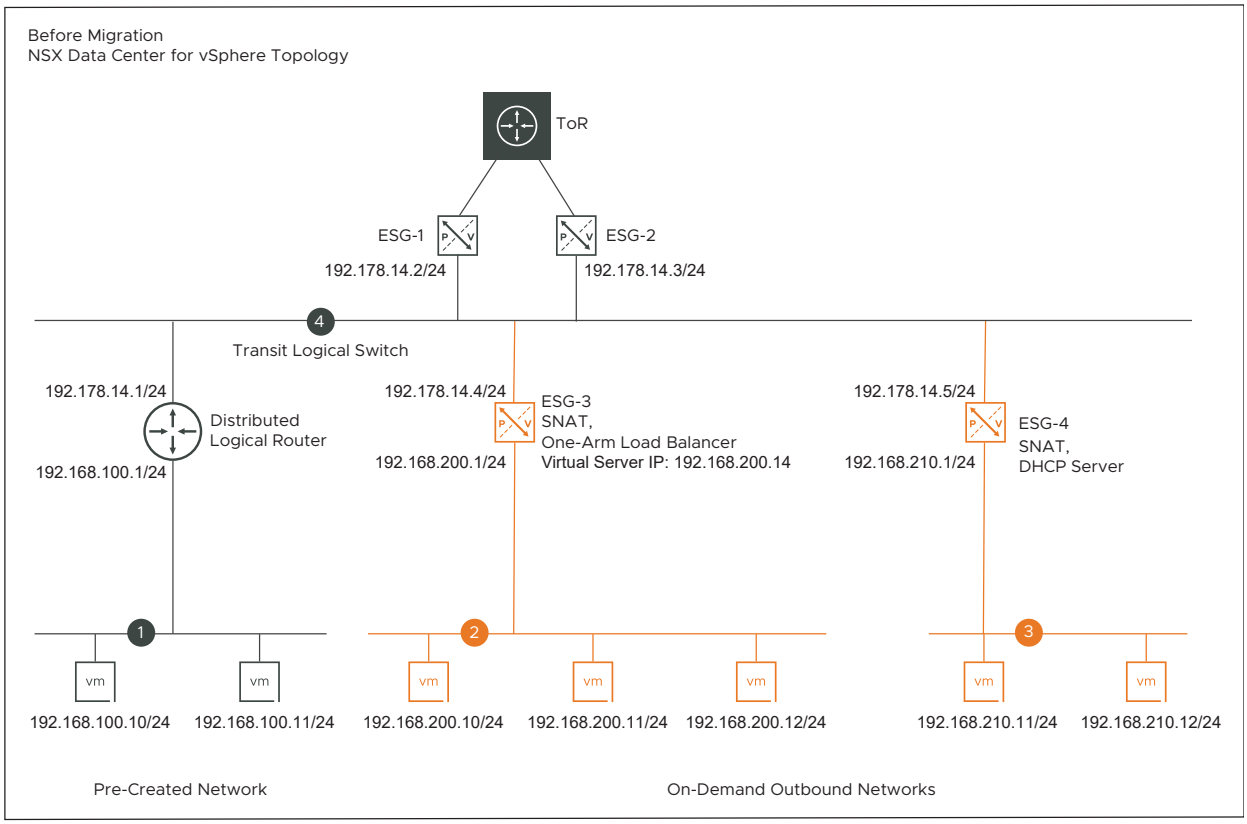

The configurations are migrated to NSX-T as follows:

- For each vRealize Automation created outbound network that is connected to an ESG, a tier-1 gateway is created and an NSX-T overlay segment is attached to the downlink of this tier-1 gateway.

- NAT rules with SNAT action are configured on the tier-1 gateways. These SNAT rules have the same configuration as the SNAT rules on the vRealize Automation created ESGs.
- Tier-1 gateways are created in the same edge cluster where the tier-0 gateway is created.
- Load balancer service configuration on ESG-3 is migrated to a load balancer service configuration on the tier-1 gateway. This tier-1 gateway is connected to the NSX-T overlay segment 2 on the downlink.
- DHCP service configuration on ESG-4 is migrated to a Gateway DHCP server configuration on the tier-1 gateway. This tier-1 gateway is connected to the NSX-T overlay segment 3 on the downlink.

  Migration coordinator uses a default DHCP server IP address to configure the Gateway DHCP server in NSX-T. If necessary, you can enter a different server IP address during the **Resolve Configuration** step of the migration. Follow the instructions that are shown on the **Submit Input** page of the migration coordinator UI to specify a different server IP address.
- In NSX-V, lease time is configured at the level of DHCP IP pool, whereas in NSX-T lease time is configurable at the level of each segment. If the DHCP service on an NSX-V network is configured with multiple DHCP IP pools, migration coordinator takes the highest lease time from among all the DHCP IP pools and configures it on the NSX-T segment.
- DHCP leases of the workload VMs that are attached to the vRealize Automation created outbound network 3 are migrated to NSX-T.

Figure 10-7. Topology G: After Migration

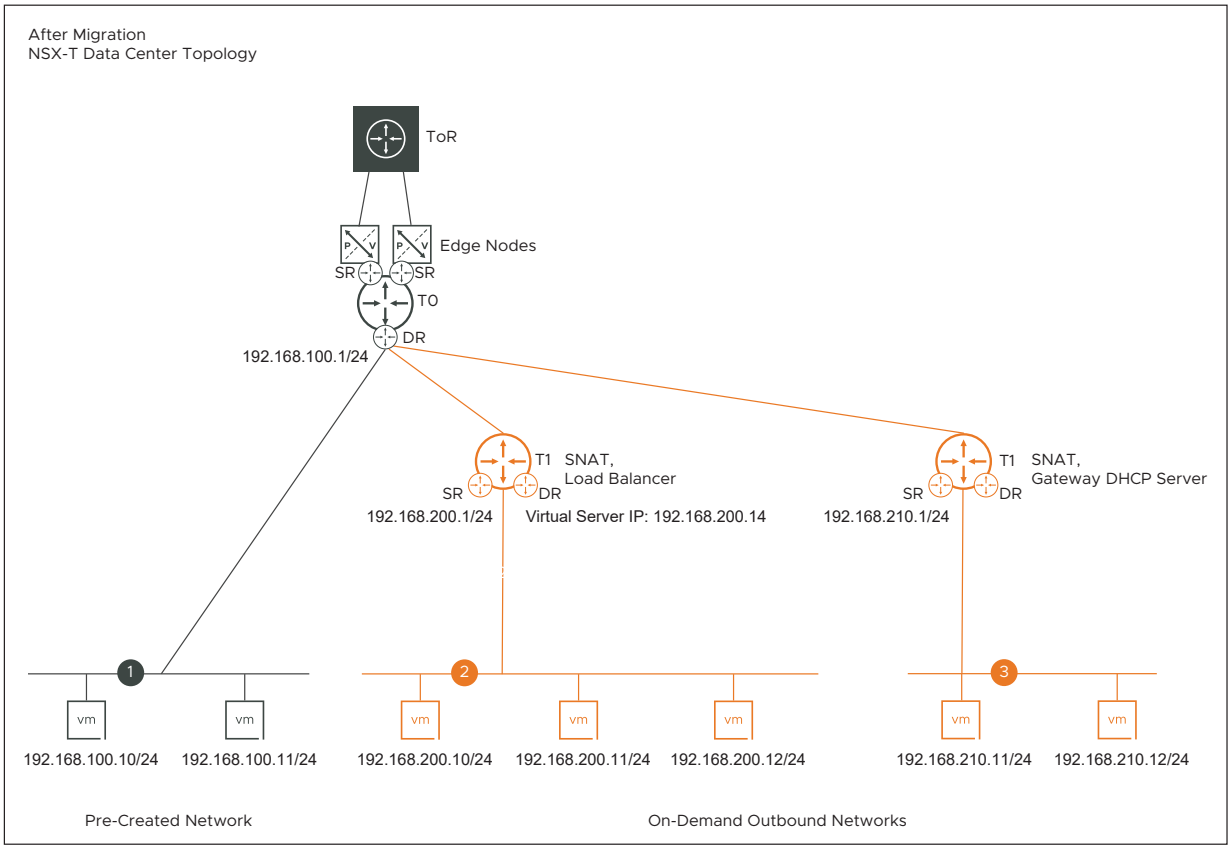

On-Demand Outbound Networks with DHCP, NAT, Inline Load Balancer (Topology H)

This topology supports the following configurations:

- On-demand outbound networks (Logical Switches) are created using vRealize Automation. For example, networks 2 and 3 in the before migration topology diagram represents the on-demand outbound networks.
- The uplink interfaces of the vRealize Automation created Edge Services Gateways (ESG-3 and ESG-4) are connected to an existing transit Logical Switch in NSX-V (for example, 192.178.14.0/24).
- ESG-3 is configured with an inline load balancer, DHCP server, and NAT rules (SNAT action).
- ESG-4 is configured with only NAT rules (SNAT action).
- The vRealize Automation created VMs on the on-demand outbound network 2 are assigned an IP address by the DHCP server.
- The vRealize Automation created VMs on the on-demand outbound network 3 have a static IP address.

- Network address translation (NAT) rules are created using vRealize Automation. SNAT action is configured on the uplink interface of ESG-3 and ESG-4. SNAT rules allow only outgoing traffic from the VMs on the outbound networks to the external clients on the public network.

  For an example of NAT rule configuration on ESG-3 and ESG-4, see the before migration topology description of Topology G that is explained earlier in this topic.
- Example: Inline load balancer configuration on ESG-3:
  - Virtual Server IP: 192.178.14.6
  - Default Pool: Pool-1
  - Pool-1 has two VMs: 192.168.200.10/24, 192.168.200.11/24

  The other load balancer configuration settings are not listed here because they are not of interest in this migration example.
- Example: Uplink interface of ESG-3 has two IP addresses configured: primary IP address and secondary IP address.
  - Primary IP address: 192.178.14.4/24. It connects ESG-3 to the transit Logical Switch.
  - Secondary IP address: 192.178.14.6/24. It is used as the virtual server IP.
- Example: DHCP server configuration on ESG-3:
  - DHCP pool range: 192.168.200.20-192.168.200.40
  - Default gateway: 192.168.200.1

  The other DHCP server settings are not listed here because they are not of interest in this migration example.

Figure 10-8. Topology H: Before Migration

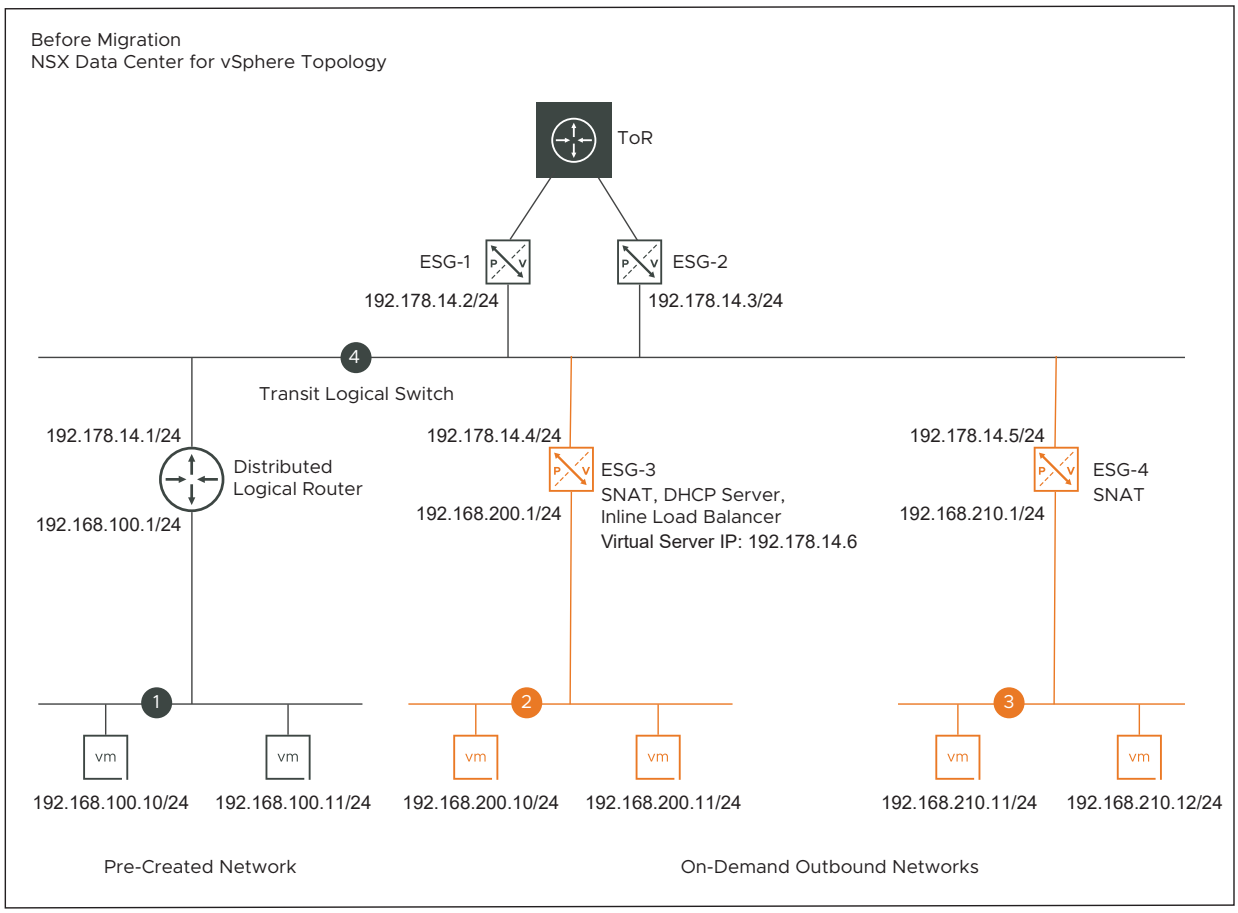

The configurations are migrated to NSX-T as follows:

- For each vRealize Automation created outbound network that is connected to an ESG, a tier-1 gateway is created and an NSX-T overlay segment is attached to the downlink of this tier-1 gateway.
- NAT rules with SNAT action are configured on the tier-1 gateways. These SNAT rules have the same configuration as the SNAT rules on the vRealize Automation created ESGs.
- Tier-1 gateways are created in the same edge cluster where the tier-0 gateway is created.
- Load balancer service on ESG-3 is migrated to a load balancer service on the tier-1 gateway. This tier-1 gateway is connected to the NSX-T overlay segment 2 on the downlink.
- DHCP service configuration on ESG-3 is migrated to a Gateway DHCP server on the tier-1 gateway.

  Migration coordinator uses a default DHCP server IP address to configure the Gateway DHCP server in NSX-T. If necessary, you can enter a different server IP address during the **Resolve Configuration** step of the migration. Follow the instructions that are shown on the **Submit Input** page of the migration coordinator UI to specify a different server IP address.

- In NSX-V, lease time is configured at the level of DHCP IP pool, whereas in NSX-T lease time is configurable at the level of each segment. If the DHCP service on an NSX-V network is configured with multiple DHCP IP pools, migration coordinator takes the highest lease time from among all the DHCP IP pools and configures it on the NSX-T segment.
- DHCP leases of the workload VMs that are attached to the vRealize Automation created outbound network 2 are migrated to NSX-T.

Figure 10-9. Topology H: After Migration

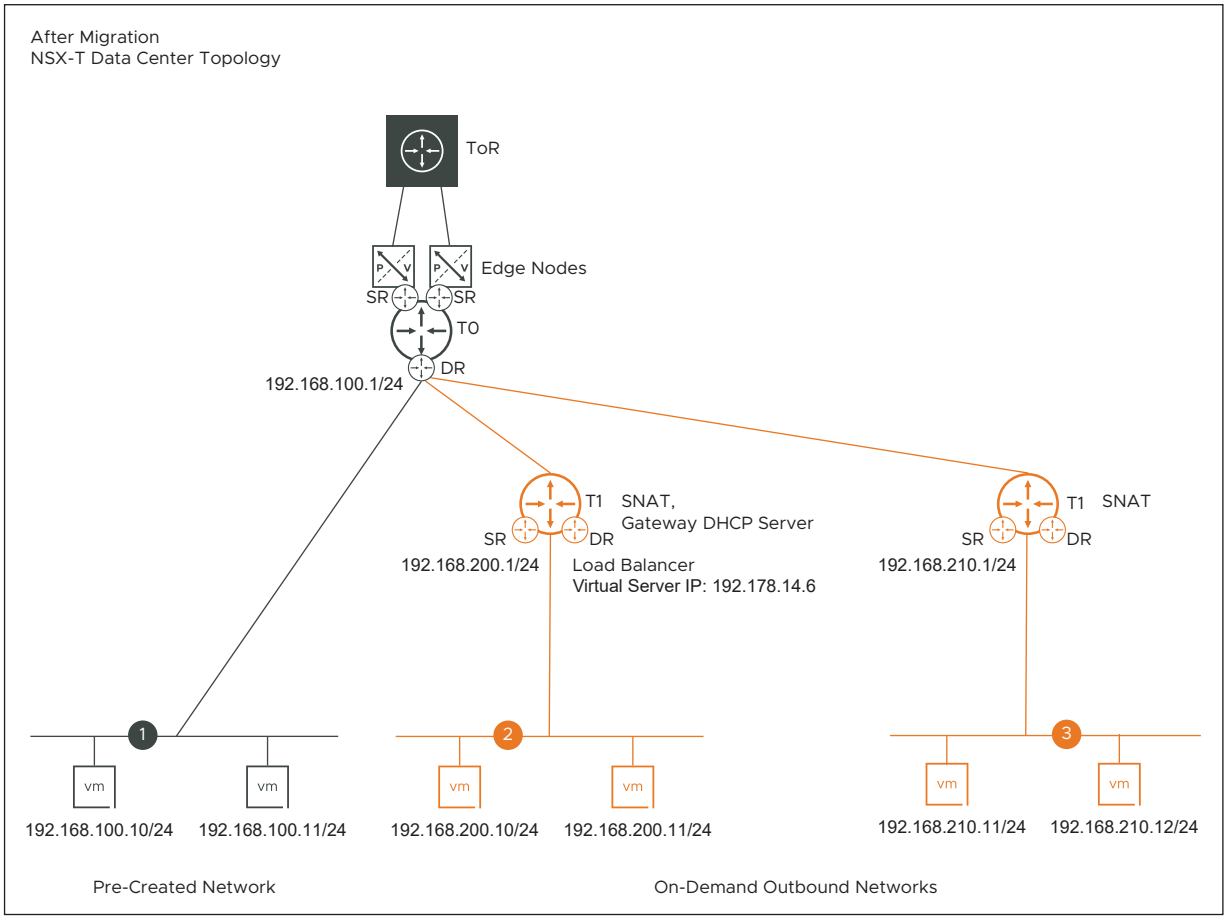

On-Demand Outbound Networks with DHCP, NAT, One-Arm Load Balancer, Inline Load Balancer (Topology I)

This topology supports the following configurations:

- On-demand outbound networks (Logical Switches) are created using vRealize Automation. For example, networks 2 and 3 in the before migration topology diagram represents the on-demand outbound networks.
- Both outbound networks are connected to the downlink interfaces of a single vRealize Automation created Edge Services Gateway (ESG-3). This topology diagram shows two outbound networks connected to ESG-3. However, you can have more than two outbound networks connected to the same ESG.

- The uplink interface of ESG-3 is connected to an existing transit Logical Switch in NSX-V (for example, 192.178.14.0/24).
- The following services are configured on ESG-3:
  - NAT rules with SNAT and DNAT action.
  - DHCP server
  - One-arm load balancer
  - Inline load balancer
- The vRealize Automation created VMs on the on-demand outbound networks 2 and 3 are assigned an IP address by the DHCP server.
- Example: Inline load balancer configuration on ESG-3. It load balances traffic on network 2:
  - Virtual Server IP: 192.178.14.5
  - Default Pool: Pool-1
  - Pool-1 has two VMs: 192.168.200.10/24, 192.168.200.11/24

  The other load balancer configuration settings are not listed here because they are not of interest in this migration example.
- Example: One-arm load balancer configuration on ESG-3. It load balances traffic on network 3:
  - Virtual Server IP: 192.168.210.2
  - Default Pool: Pool-2
  - Pool-2 has two VMs: 192.168.210.11/24, 192.168.210.12/24
- NAT rules with SNAT action and DNAT action are configured on ESG-3. Port forwarding is supported for both SNAT and DNAT rules.
- Example: SNAT rules configuration on ESG-3.

SNAT Rule 1	SNAT Rule 2
Original	**Original**
Applied on: vNIC (uplink interface)	Applied on: vNIC (uplink interface)
Protocol: any	Protocol: any
Source IP range: 192.168.200.1-192.168.200.14	Source IP range: 192.168.210.1-192.168.210.14
Source ports: any	Source ports: any
Destination IP range: any	Destination IP range: any
Destination ports: any	Destination ports: any
Translated	**Translated**
IP address: 192.178.14.5	IP address: 192.178.14.4
Port range: any	Port range: any

- Example: DNAT rule configuration on ESG-3.

  Original
  Protocol: any
  Source IP range: any
  Source ports: any
  Destination IP range: 192.178.14.5
  Destination ports: 100

  Translated
  IP address: 192.178.200.35
  Port range: 80

- Example: DHCP server configuration on ESG-3 has two IP pools:
  - DHCP IP pool 1: 192.168.200.20-192.168.200.40
  - Default gateway for pool 1: 192.168.200.1
  - DHCP IP pool 2: 192.168.210.20-192.168.210.40
  - Default gateway for pool 2: 192.168.210.1

  The other DHCP server settings are not listed here because they are not of interest in this migration example.

Figure 10-10. Topology I: Before Migration

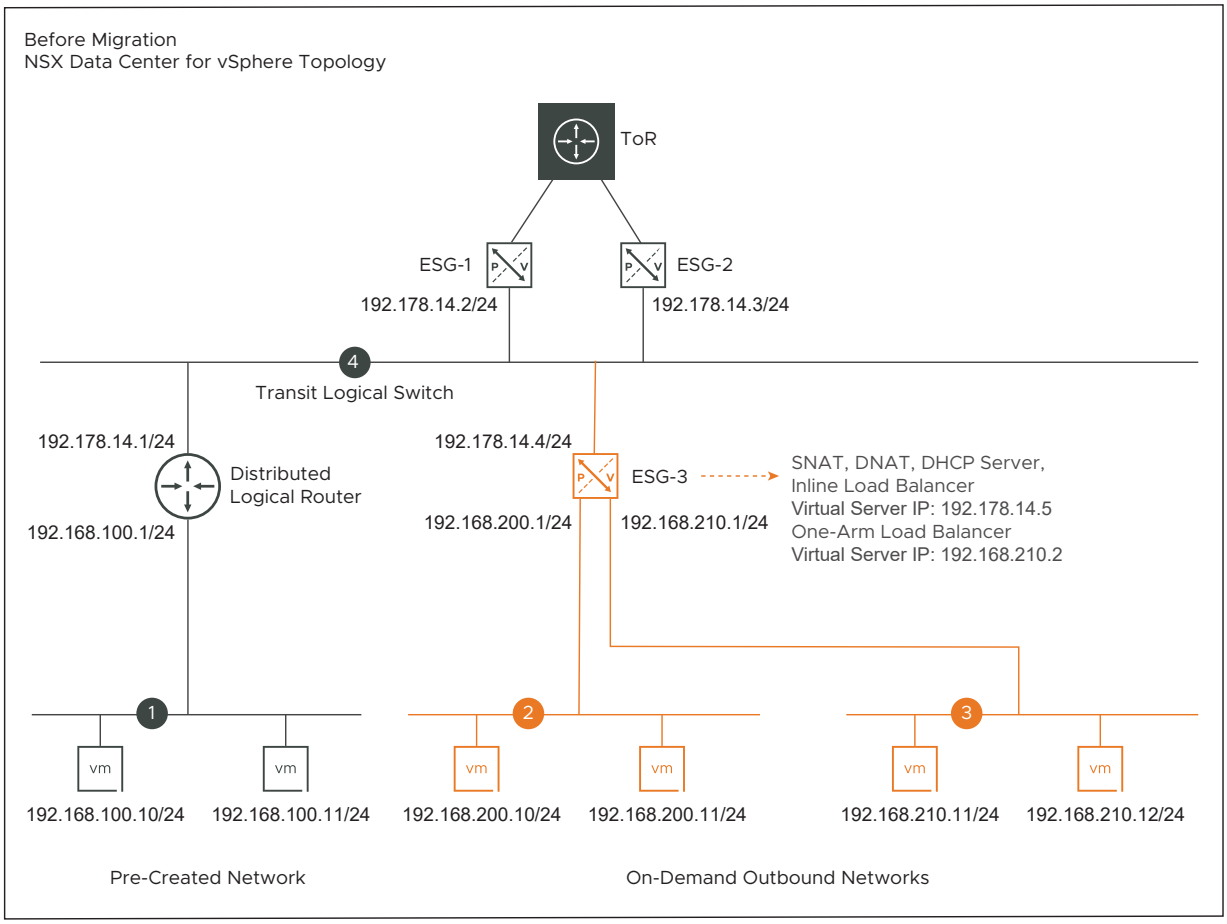

The configurations are migrated to NSX-T as follows:

- For the vRealize Automation created outbound networks that are connected to a single ESG, a tier-1 gateway is created and NSX-T overlay segments are attached to the downlink of this tier-1 gateway.
- SNAT and DNAT rules are configured on the tier-1 gateway. These rules have the same configuration as the SNAT and DNAT rules on the vRealize Automation created ESG-3.
- Tier-1 gateway is created in the same edge cluster where the tier-0 gateway is created.
- Load balancer service configurations on ESG-3 are migrated to the load balancer service configurations on the tier-1 gateway.
- DHCP service configuration on ESG-3 is migrated to a Gateway DHCP server on the tier-1 gateway.

  Migration coordinator uses a default DHCP server IP address to configure the Gateway DHCP server in NSX-T. If necessary, you can enter a different server IP address during the **Resolve Configuration** step of the migration. Follow the instructions that are shown on the **Submit Input** page of the migration coordinator UI to specify a different server IP address.

- In NSX-V, lease time is configured at the level of DHCP IP pool, whereas in NSX-T lease time is configurable at the level of each segment. When the DHCP service on an NSX-V network is configured with multiple DHCP IP pools, migration coordinator takes the highest lease time from among all the DHCP IP pools and configures it on the NSX-T segment.
- DHCP leases of the workload VMs that are attached to the vRealize Automation created outbound networks are migrated to NSX-T.

Figure 10-11. Topology I: After Migration

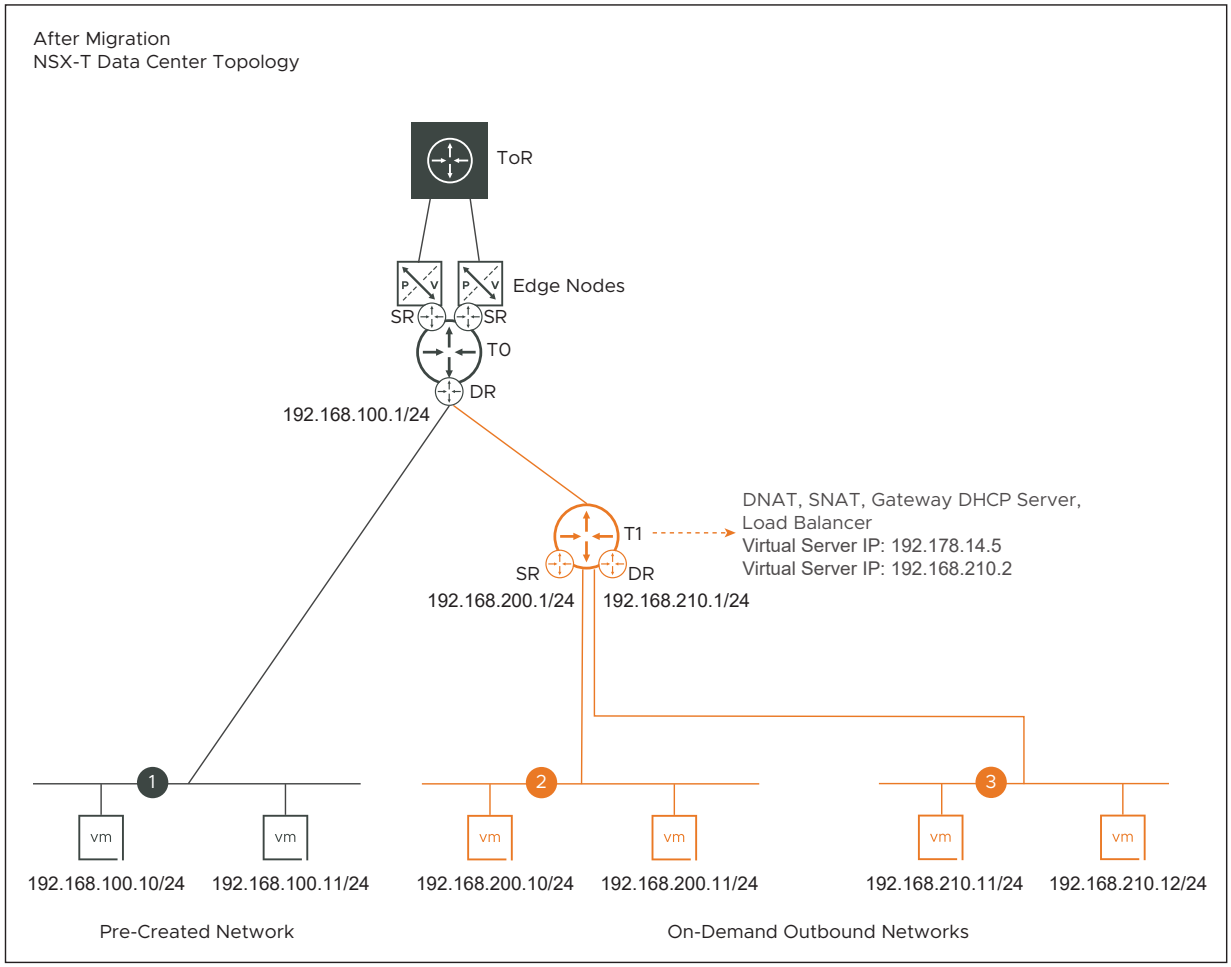

On-Demand One-Arm Load Balancer on Existing Networks (Topology J)

This topology has the following configurations:

- A one-arm load balancer is configured on a vRealize Automation created ESG. This ESG is connected to an existing network in NSX-V. The existing network can either be a VLAN or a Logical Switch.

The before topology diagram shows a single one-arm load balancer on an existing network 1. For the purposes of this topology description, a single one-arm load balancer on a vRealize Automation created ESG is considered. However, this topology also supports the following configurations:

- Multiple vRealize Automation created one-arm load balancers connected to a single existing network. One load balancer is configured on each vRealize Automation created ESG. All ESGs are deployed as part of a single vRealize Automation deployment.
- Multiple vRealize Automation created one-arm load balancers, where one load balancer is connected to each different existing network. One load balancer is configured on each vRealize Automation created ESG. All ESGs are deployed as part of a single vRealize Automation deployment.

- The vRealize Automation created workload VMs that are connected to an existing network have a static IP address.
- Example: One-arm load balancer configuration on ESG-3.
  - Virtual Server IP: 192.168.100.2
  - Default Pool: Pool-1
  - Pool-1 has two VMs: 192.168.100.10/24, 192.168.100.11/24

Figure 10-12. Topology J: Before Migration

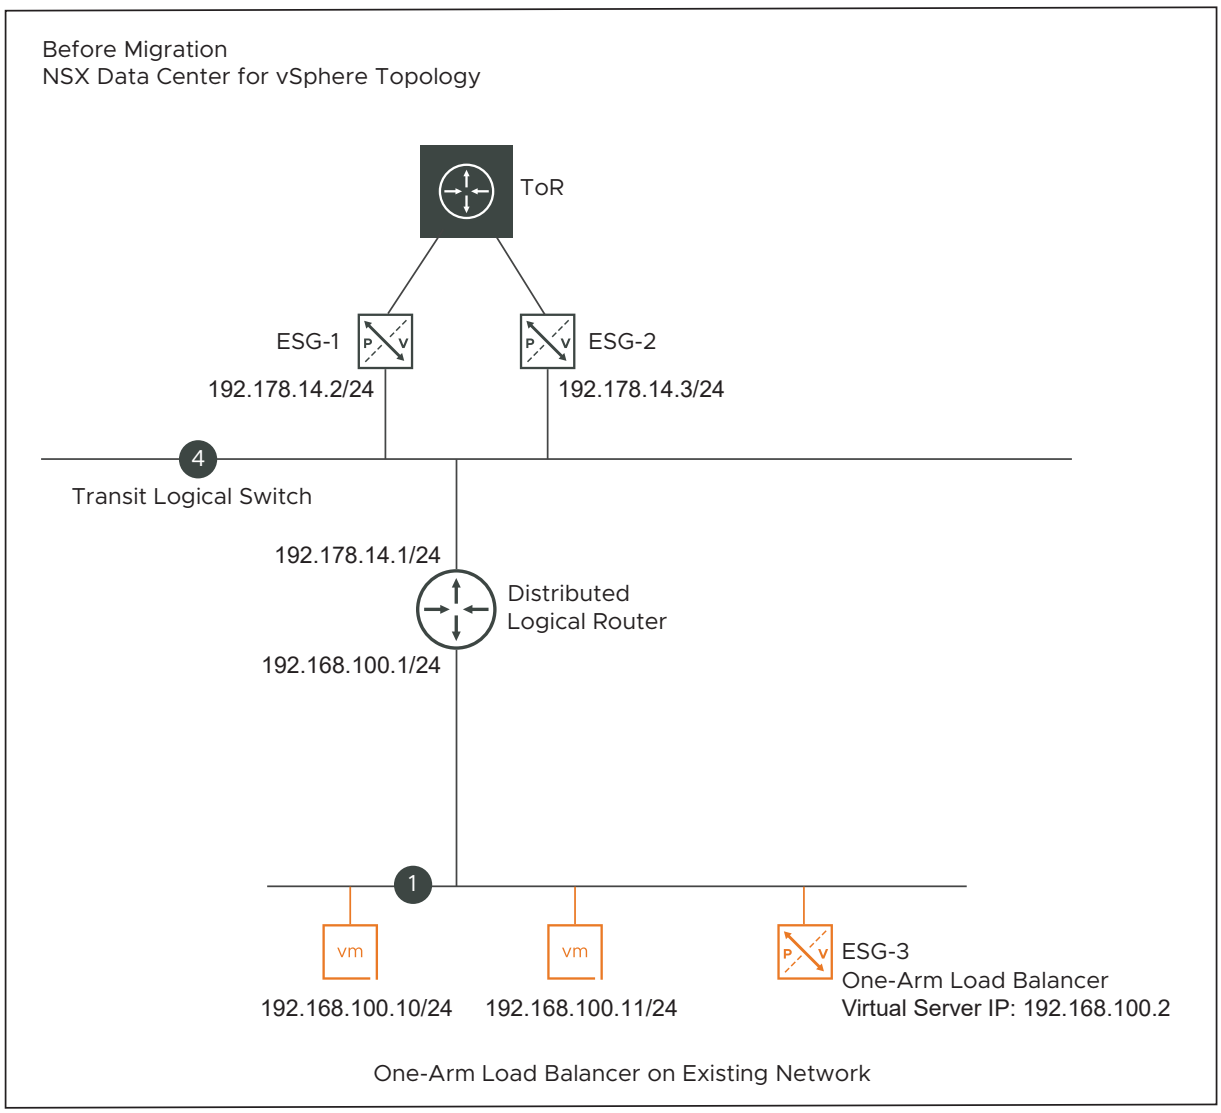

The configurations are migrated to NSX-T as follows:

- ESG-3 in the NSX-V topology is migrated to a standalone tier-1 gateway in the NSX-T topology. This tier-1 gateway is not connected to the tier-0 gateway.
- The load balancer configuration on ESG-3 is migrated to a load balancer configuration on the tier-1 gateway.

Figure 10-13. Topology J: After Migration

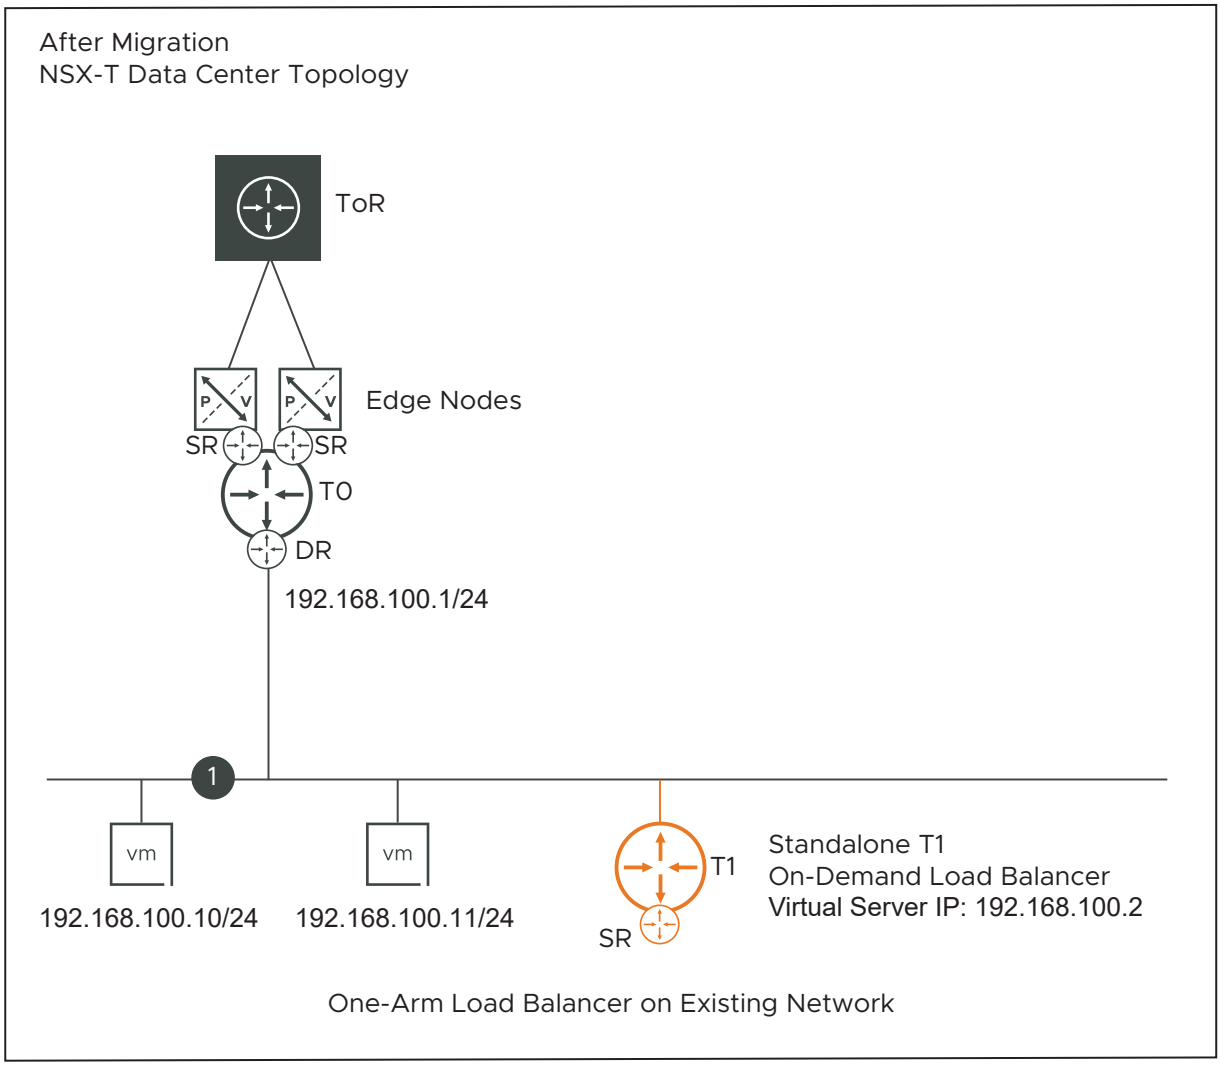

On-Demand Security Groups or Existing Security Groups with Load Balancer (Topology K)

This topology supports a vRealize Automation created ESG with a single load balancer configured on it. The ESG can be connected either to an existing network or to an on-demand network that is created through vRealize Automation.

An existing Security Group is a part of the vRealize Automation Network Profile, or an on-demand/existing Security Group is present in the vRealize Automation deployment that contains on-demand networks.

In the input deployment configuration file, vRealize Automation automatically adds the virtual server IP of the load balancer inside an IP set object. This IP set is added as a member of the Security Group in the vRealize Automation Network Profile. Or the IP set is added as a member of the on-demand or existing Security Group from the blueprint that is associated with the load balanced pool members.

When such a topology is migrated to NSX-T, the ESG maps to a tier-1 gateway in the NSX-T topology. Load balancer service configuration on ESG is migrated to a load balancer service configuration on the tier-1 gateway. The Security Group in NSX-V maps to a Group in NSX-T. This NSX-T Group contains a nested Group that has the virtual server IP of the load balancer.

Deployment Configuration File

A deployment configuration file is an input to the migration coordinator tool. The migration coordinator reads the `.json` configuration file, validates it against a pre-defined JSON schema, and migrates the vRealize Automation resource configurations from the existing NSX-V environment to a new NSX-T environment.

The deployment configuration file contains the following configuration information:

- List of resources that vRealize Automation has created.
- List of resources that vRealize Automation references from the existing NSX-V environment.
- List of network interfaces of the workload VMs that vRealize Automation has created.
- Desired mapping of the vRealize Automation resources to NSX-T objects.

Mapping example: Consider that your topology has a vRealize Automation created Logical Switch that connects to an existing Distributed Logical Router in NSX-V. The mapping information in the deployment configuration file tells the migration coordinator to do the following:

- Create a relevant tier-1 gateway in NSX-T that maps to the Network Profile deployed in vRealize Automation.
- Attach an NSX-T overlay segment to the downlink of this tier-1 gateway.

The deployment configuration file is generated from the vRealize Automation environment, and it provides configuration information about all the vRealize Automation created on-demand resources that are deployed over several days of network operations. This configuration file must be valid and conform to a pre-defined JSON schema.

Typically, the cloud administrator generates the deployment configuration file and provides it to the network administrator, who runs the migration coordinator tool in the NSX-T environment. In some organizations, a single administrator might perform both roles.

Remember, each vRealize Automation deployment can create multiple on-demand networks, and each network can have multiple network interfaces. The networks inside a deployment can consume resources that are either created by vRealize Automation or that are existing in NSX-V.

For example, let us assume that over several days of network operations, you have created multiple deployments in vRealize Automation to deploy the following resources in your NSX-V topology:

- Deployment 1: For creating two on-demand private networks and one on-demand routed network.
- Deployment 2: For creating on-demand security groups.

- Deployment 3: For creating on-demand routed networks with no services.

When you are ready to migrate NSX-V with vRealize Automation to NSX-T, a single deployment configuration file is generated from vRealize Automation, and uploaded as an input file to the migration coordinator.

Output Mapping File

After the resource configurations in the vRealize Automation deployment configuration file are migrated from NSX-V to NSX-T, the migration coordinator creates an output mapping file in a `.json` format.

This mapping file contains information about how vRealize Automation created resources in NSX-V map to NSX-T objects.

vRealize Automation can use this mapping file to update or refresh its own database after the migration is finished. The migration coordinator places this mapping file at the following path on the NSX Manager appliance where the migration coordinator service runs: `/var/log/migration-coordinator/v2t`.

The migration coordinator creates a partial or an intermediate mapping file for the first time after the Check Realization step is completed. In the subsequent steps of the migration process, the mapping file changes. For example, mapping information about vRealize Automation created Security Policies, VM network interfaces, Security Groups, and so on, are available in the output mapping file after host migration is completed. The final mapping file is generated after all the hosts in the clusters, which are enabled for migration, are migrated to NSX-T.

You can download the output mapping file from the migration coordinator UI after the following stages in the migration process are complete:

- Check Realization step
- Migrate Edges step
- Migrate Hosts step

The information in the mapping file varies depending on the migration stage that is completed.

For example:

- The details in the mapping file before host migration and after host migration can change. Before the hosts are migrated, the mapping file does not contain the network interfaces (vNIC IDs) of the workload VMs that are created in vRealize Automation. The network interfaces of the workload VMs are mapped only after the VMs connect to NSX-T segments during host migration.
- The mapping file shows the new network interfaces of the static VM members in the migrated Security Groups only after all the hosts are migrated to NSX-T.

Scenario: Deleted Resources List in Mapping File

Consider the following sequence of events:

1 Your deployment configuration file contains several vRealize Automation created resources. One of the resources is an on-demand Security Group.

2 Before running the migration coordinator tool, during a day-two operation, you delete this vRealize Automation created Security Group in the NSX-V UI. The NSX-V inventory does not have the Security Group.

3 Let us assume that the data about this deleted Security Group is not yet reconciled with the vRealize Automation database. Therefore, the vRealize Automation database continues to reference this deleted Security Group.

4 Now, you run the migration coordinator tool.

After migration to NSX-T, the migration coordinator shows the deleted Security Group in the list of "output deleted resources" of the final mapping file. Because the migration coordinator cannot find the ID of the deleted Security Group, the mapping file does not have the path to the corresponding NSX-T Group.

High-Level View of Migrating NSX-V with vRealize Automation

The migration process includes setting up a new NSX-T environment and running the migration coordinator. You also might need to change your existing NSX-V environment to ensure that it can migrate to NSX-T.

Caution Deploy a new NSX-T environment to be the destination for the NSX-V migration.

During the **Import Configuration** step, all NSX Edge node interfaces in the destination NSX-T environment are shut down. If the destination NSX-T environment is already configured and is in use, starting the configuration import will interrupt traffic.

During the migration, you will complete the following steps:

1 Create a new NSX-T environment.

- Deploy a single NSX Manager appliance to create the NSX-T environment.
- If you plan to use Maintenance Mode migration for hosts, configure a shared storage on the cluster to be migrated from NSX-V to NSX-T. This enables automated vMotion for the migration process. Any VMs that do not meet this criterion must be manually powered off prior to migration or manually vMotioned.
- Configure a compute manager in the NSX-T environment. Add the vCenter Server as a compute resource.

  Important Use the exact IP or hostname specified in NSX-V vCenter Server registration.
- Start the migration coordinator service.
- If you want to import users from NSX-V, set up VMware Identity Manager.

- If your NSX-V topology uses Edge Services Gateways, create an NSX-T IP pool to use for the NSX-T Edge TEPs. These IPs must be able to communicate with all the existing NSX-V VTEPs.
- Deploy NSX Edge nodes.
  - Deploy the correct number of appropriately sized NSX-T Edge appliances.
  - Join the Edge nodes to the management plane from the command line.
- If Security Policies in your NSX-V environment use third-party Guest Introspection or Network Introspection services provided by partners, the partner services must be registered with NSX-T.

2 Import configuration from NSX-V.
- Enter the details of your NSX-V environment.
- Upload a vRealize Automation deployment configuration file in `.json` format as an input to the migration coordinator.
- The configuration is retrieved and pre-checks are run.

3 Resolve issues with the configuration.
- Review messages and the reported configuration issues to identify any blocking issues or other issues that require a change to the NSX-V environment.
  - If you make any changes to the NSX-V environment while a migration is in progress, you must restart the migration and import the configuration again.
- Provide inputs to configuration issues that must be resolved before you can migrate your NSX-V environment to NSX-T. Resolving issues can be done in multiple passes by multiple people.

4 Migrate configuration.
- After all configuration issues are resolved, you can migrate the configuration to NSX-T. Configuration changes are made on NSX-T, but no changes are made to the NSX-V environment yet.

5 Check realization of migrated configurations in NSX-T.
- Check whether the migration coordinator shows any unrealized configuration errors.
- Download an intermediate output mapping file and share it with the vRealize Automation cloud administrator for an initial review.
- Take appropriate actions to resolve any missing configurations in NSX-T before proceeding to Edge migration.

6 Migrate Edges.

- Routing and Edge services are migrated from NSX-V to NSX-T.

  Caution North-South traffic is interrupted during the Migrate Edges step. All traffic that was previously traversing through the Edge Services Gateways (North-South traffic) moves to the NSX-T Edges.

7 Migrate Hosts.

- NSX-V software is removed from the hosts, and NSX-T software is installed. VM interfaces are connected to the new NSX-T segments.

  Caution If you select **In-Place** migration mode, there is a traffic interruption for a few seconds during the Migrate Hosts step. However, if you select **Maintenance** migration mode, traffic interruption does not occur.

- Download the final output mapping file after the hosts in all the clusters are migrated to NSX-T.

8 Finish Migration.

- After you have verified that the new NSX-T environment is working correctly, you can finish the migration, which clears the migration state.

9 Perform post-migration tasks.

- Deploy two additional NSX Manager appliances before you use your NSX-T environment in a production environment.
- Uninstall NSX-V environment.

Note: Logical ports and switches that are created during migration are not deleted when the workload VMs are deleted. You must delete these ports and switches via the NSX Manager UI or the API.

Migration-Supported Operations in NSX-V for vRealize Automation Resources

Over several days of network operations in NSX-V, you can update the configurations of the on-demand resources that you defined earlier in vRealize Automation.

You can directly use the NSX-V UI or APIs to update the previously defined vRealize Automation resource configurations. The migration coordinator fetches the updated configurations from NSX-V and migrates them to NSX-T.

In the context of migration, the following network operations in NSX-V are supported for vRealize Automation created resources.

Table 10-12. Supported Operations in NSX-V for vRealize Automation Resources

vRealize Automation Created Resource	Supported Network Operations in NSX-V
On-Demand Networks	■ Add or delete VMs from existing on-demand networks that are created using vRealize Automation. ■ Delete an existing on-demand network after deleting or moving workload VMs that are created using vRealize Automation.
On-Demand Security Groups	■ Delete Security Groups that are created using vRealize Automation. ■ Add Security Groups that are created with vRealize Automation to an existing Security Group (out-of-band configuration). ■ Add or change the Security Group in Security Policies that are created with vRealize Automation. ■ Add new firewall rules in Security Policies that are created with vRealize Automation.
DHCP Server configuration	■ Delete or disable DHCP server configuration that is created with vRealize Automation.
Workload VMs	■ Change the network interface (vNIC) assignment of a vRealize Automation created VM from dynamic to static, or reverse, or from static IP1 to static IP2. ■ Change from a single vNIC on a vRealize Automation created VM to multiple vNICs, and reverse.
Load balancer configuration	■ Edit health check related parameters in the load balancer ■ Add or delete pool members ■ Disable load balancer service on the ESG
NAT configuration	■ Modify existing rules ■ Delete existing rules ■ Reorder rules ■ Add port forwarding rules

Preparing to Migrate NSX-V with vRealize Automation

Before you migrate, review the documentation and verify that you have the required software versions. Modify your existing NSX-V environment if needed, and deploy a new NSX-T environment.

Required Software and Versions

- The version of vRealize Automation that you are using must support this migration. Support for this migration is available starting in vRealize Automation 8.3. See the Release Notes in the vRealize Automation documentation at https://docs.vmware.com/en/vRealize-Automation/index.html.

In vRealize Automation, the target NSX-T cloud account must be associated with NSX-T 3.1.1 or later.

Prepare vRealize Automation Environment for Migration

Do the following tasks in vRealize Automation to prepare for the migration:

- Add a new target NSX-T cloud account in vRealize Automation for the source NSX-V cloud account in your vRealize Automation projects.
- Create a migration plan to migrate from NSX-V to NSX-T. Complete the following four steps in the migration plan:
  - Step 1: Specify the source and target cloud accounts.
  - Step 2: Assess the NSX cloud accounts for migration readiness.
  - Step 3: Enter maintenance mode for the cloud accounts.
  - Step 4: Generate a deployment configuration file.

  For more information, see the vRealize documentation about migrating NSX-V to NSX-T.

The deployment configuration file in `.json` format becomes an input to the NSX-T migration coordinator.

Import Configuration of NSX-V with vRealize Automation

The first step of the migration process is to import the NSX-V with vRealize Automation configuration.

Prerequisites

- Generate a deployment configuration file in a `.json` format from the vRealize Automation environment. The configuration file must be valid and conform to a predefined JSON schema.

Procedure

1. From a browser, log in to NSX Manager as **admin**.
2. Navigate to **System > Migrate**.
3. In the **vRealize Automation** pane, click **Get Started** and select **Fixed Topology**.
4. In the **Upload Deployment Configuration File** section, click **Select File**. Select the `.json` configuration file that was generated from the vRealize Automation environment, and click **Upload**.

   After uploading a file, if necessary, you can click **Select File** again and upload a different file. The previous file is overwritten. You cannot remove the `.json` file after it is uploaded. The configuration file is removed only if you take any of the following actions:

   - Roll back the migration on the **Import Configuration** page or the **Migrate Configuration** page.

- Finish the migration on the **Migrate Hosts** page.

5 Click **Start** to import the configuration.

6 When the import has finished, click **Continue** to proceed to the **Resolve Configuration** page.

- If the import fails due to incorrect edge node configuration translation, click the **Failed** flag to view information about the number and size of the required NSX Edge resources. After you deploy the correct number and size of edge nodes, click **Rollback** to roll back this migration attempt and restart the configuration import.
- If syntax errors are found in the deployment configuration file, import fails during the translate configuration step. Roll back the migration, resolve the syntax errors in the `.json` file, and import the configuration again.
- If import configuration fails due to schema mismatch errors, review the error messages, roll back the imported configuration, and upload a valid configuration file to start a new migration. The error message in the migration coordinator UI might not display all the possible reasons for the schema mismatch errors. For a complete information about the schema mismatch errors, see the migration log (`cm.log` file) on the NSX Manager appliance where the migration coordinator service is running.

  You can find this log file at `/var/log/migration-coordinator/v2t`, and use it for further troubleshooting purposes.

Results

When the NSX-V with vRealize Automation topology is imported successfully, the **View Imported Topology** link is enabled. Click this link to view a graph of the imported topology. However, the topology viewer does not work for a scaled NSX-V environment.

Roll Back the NSX-V with vRealize Automation Migration

After you have started the migration process, you can roll back the migration to undo some or all your progress.

You can roll back or undo the migration from some of the migration steps. After the migration has started, you can click **Rollback** on the furthest step completed. The button is disabled on all other pages.

Table 10-13. Rolling Back NSX-V with vRealize Automation Migration

Migration Step	Roll Back Details
Import Configuration	Click **Rollback** on this page to roll back the **Import Configuration** step.
Resolve Configuration	Rollback is not available in this step. Click **Rollback** from the **Import Configuration** page.

Table 10-13. Rolling Back NSX-V with vRealize Automation Migration (continued)

Migration Step	Roll Back Details
Migrate Configuration	Click **Rollback** on this page to undo or remove the migrated configurations in NSX-T. The inputs provided on the **Resolve Configuration** page are removed. The migration log files are removed from the NSX Manager appliance. You must start a new migration. Before rolling back the migration from the **Migrate Configuration** page, it is recommended to collect the Support Bundle. For more information, see *Collect Support Bundles* in the *NSX-T Data Center Administration Guide*. Verify that the rollback was successful before you start a new migration. Log into the NSX Manager web interface and switch to **Manager** mode. Verify that all configurations have been removed. For more information about Manager mode, see *Overview of the NSX Manager* in the *NSX-T Data Center Administration Guide*.
Check Realization	Click **Rollback** on this page to undo this step and return to the **Migrate Configuration** page. The migrated configurations in NSX-T are retained.
Migrate Edges	Click **Rollback** on this page to roll back the migration of Edge routing and services to NSX-T. **Caution** If you roll back the **Migrate Edges** step, verify that the traffic is going back through the NSX-V Edge Services Gateways. You might need to take manual action to assist the rollback.
Migrate Hosts	Rollback is not available in this step. However, you can still do a manual rollback to remove NSX-T from the migrated hosts and reinstall NSX-V on the hosts. If you are migrating from NSX-V 6.4.8 or later, run the following REST API on the NSX-V NSX Manager before doing the manual rollback, and then reinstall NSX-V on the hosts. `POST api/2.0/nwfabric/blockEamEvents?action=unblock` This API enables the vSphere ESX Agent Manager (EAM) on the hosts so that the NSX-V VIBs can be installed correctly. If you are migrating from NSX-V 6.4.4, 6.4.5, or 6.4.6, this API is not needed. **Note** It is better to do a manual rollback of a failed host after the cluster, which contains the failed host, has stopped. If you are doing an In-Place host migration, and chosen parallel migration order within the cluster, then wait until the migration of all hosts in the cluster stops, regardless of failure or success. If you have chosen serial migration order of hosts within the cluster, migration stops when a host migration fails.

When you roll back a migrated configuration that contains Network Introspection redirection rules, you might see the following error message:

```
Service Insertion failed with '400: The object path=[/infra/segments/service-segments/
vdnscope-1]
cannot be deleted as either it has children or it is being referenced by other objects path=
[/infra/service-chains/Service-Chain-serviceprofile-1].
```

This error occurs because a service segment in NSX-T depends on a service chain. A service chain is not deleted until all the redirection rules referenced by it are deleted. Wait for approximately five minutes and try rolling back the migrated configuration again.

Resolve Configuration Issues

After you have imported the configuration from your NSX-V environment, you must review and resolve the reported configuration issues before you can continue with the migration.

On the **Resolve Configuration** page, two types of configuration issues are reported.

Blocking Issues

As the name suggests, these issues block the migration, and they must be fixed for the migration to proceed. You might have to change the configurations in your NSX-V environment before you can migrate to NSX-T.

Warnings

These configuration issues are organized into several categories, and each category can have one or more configuration items. You should provide inputs to fix the warning messages that are displayed for the configuration items, or choose to skip the warnings, if needed. The input you provide determines how the NSX-T environment is configured.

- Review Migration Information

  The **Resolve Configuration** page contains information about the features and configurations that are not supported for migration, and the issues that must be fixed in the NSX-V environment before you can migrate.

- Make Changes to the Environment

  If you find blocking issues or other configuration issues that must be fixed in your NSX-V environment, fix those issues before you can proceed with the migration. After you make the configuration changes, you must import the configuration again.

- Provide Inputs for Configuration Issues

  After you have reviewed the migration information and are ready to proceed with the migration, you can provide inputs for the reported configuration issues. The input you provide determines how the NSX-T environment is configured.

Review Migration Information

The **Resolve Configuration** page contains information about the features and configurations that are not supported for migration, and the issues that must be fixed in the NSX-V environment before you can migrate.

After reviewing the blocking issues and warnings, you might need to change configurations in your NSX-V environment before you can migrate to NSX-T. If you change the NSX-V environment, you must restart the migration to pick up the new configuration. Review all migration feedback before you provide input to avoid duplication of work.

Note For some NSX-V features, there might be automatic configurations such as certificates present. If these configurations are for features that are not supported for the specific topology, these automatic configurations are flagged as issues that need to be skipped from migration. For example, in topologies that don't support L4-L7 services on Edge Services Gateways, the certificates present for VPN and DNS will raise issues to skip these configurations from migration.

Procedure

1 From the **Resolve Configuration** page, review the reported issues in the **Blocking** category to identify blocking issues that require changes to your NSX-V environment.

Figure 10-14. Blocking Issues on the Resolve Configuration Page

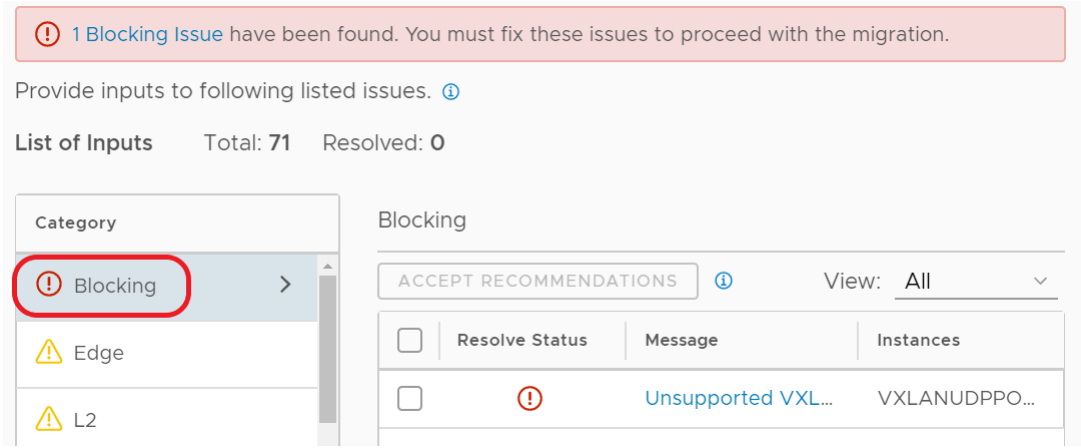

Some examples of blocking issues are:

- Incorrect DRS configuration of Maintenance mode migration.
- vMotion vmknics not configured on host for Maintenance mode migration.
- Unsupported VXLAN UDP port.

2 Review the warnings and issues reported in each category.

Figure 10-15. Warnings and Categories of Issues on the Resolve Configuration Page

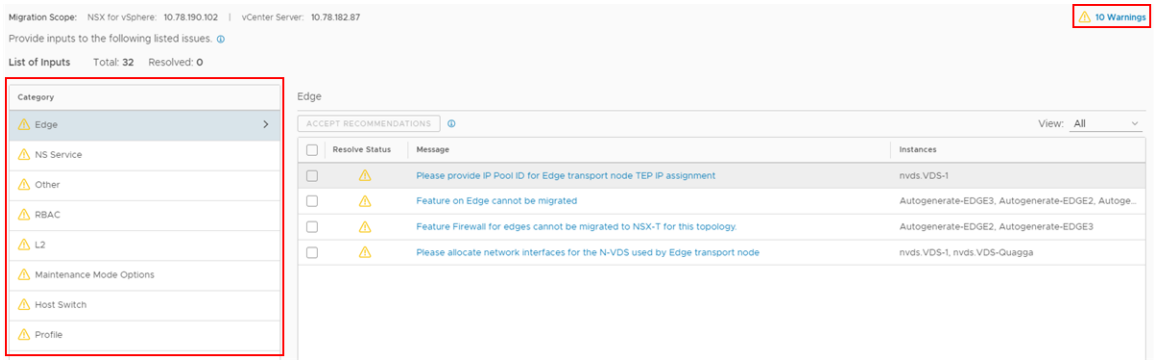

a Click **Warnings** and review the information there.

b Review the reported issues in all categories.

What to do next

- If you find blocking issues, fix them in the NSX-V environment before you can proceed with the migration. See Make Changes to the Environment.
- If you did not find any blocking issues or other configurations that require a change in the NSX-V environment, you can proceed with the migration. See Provide Inputs for Configuration Issues.

Make Changes to the Environment

If you find blocking issues or other configuration issues that must be fixed in your NSX-V environment, fix those issues before you can proceed with the migration. After you make the configuration changes, you must import the configuration again.

Prerequisites

Verify that Host or Edge migration has not started. See Roll Back the NSX-V with vRealize Automation Migration for more information about restarting the migration.

Procedure

1 Make the required changes in the NSX-V environment.

2 Navigate to the **Import Configuration** page and click **Rollback**.

The vRealize Automation deployment configuration input file and the migration log file are removed. If any inputs were submitted on the **Resolve Configuration** page, they are cleared.

3 Upload the vRealize Automation deployment configuration file again.

4 Click **Start** to import the updated NSX-V configuration.

Results

The migration starts with the new NSX-V configuration.

What to do next

Continue the migration process. See Resolve Configuration Issues.

Provide Inputs for Configuration Issues

After you have reviewed the migration information and are ready to proceed with the migration, you can provide inputs for the reported configuration issues. The input you provide determines how the NSX-T environment is configured.

Multiple people can provide the input over multiple sessions. You can return to a submitted input and modify it. Depending on your configuration, you might run through the **Resolve Issues** process multiple times, update your NSX-V environment as needed, and restart the migration.

Important If you have changed the NSX-V environment for any reason since you last imported the configuration, you must restart the migration. For example, if you have connected a new VM to a logical switch, made a firewall rule change, or installed NSX-V on new hosts. See Make Changes to the Environment for more information on restarting the migration.

Note For some NSX-V features, there might be automatic configurations such as certificates present. If these configurations are for features that are not supported for the specific topology, these automatic configurations are flagged as issues that need to be skipped from migration. For example, in topologies that don't support L4-L7 services on Edge Services Gateways, the certificates present for VPN and DNS will raise issues to skip these configurations from migration.

Prerequisites

- Verify that you have reviewed all issues and warning messages and are ready to continue with the migration.
- Verify that you have addressed all blocking issues and other issues requiring a change to the NSX-V.

Procedure

1. Navigate to **System > Migrate**.
2. Go to the **Migrate NSX for vSphere with vRealize Automation** pane, and click **Resolve Configuration**.
3. Click each issue and provide input.

   Each issue can cover multiple configuration items. For each item there might be one or more possible resolutions to the issue, for example, skip, configure, or select a specific value.

   For issues that apply to multiple configuration items, you can provide input for each item individually, or select all and provide one answer for all items.

4 After the input is provided, a **Submit** button is displayed on the **Resolve Configuration** page. Click **Submit** to save your progress.

5 When you have provided input for all configuration issues, click **Submit**.

The input is validated. You are prompted to update any invalid input. Additional input might be required for some configuration items.

6 After you have submitted all requested input, click **Continue** to proceed to the **Migrate Configuration** page.

Example: Configuration Issues

For some examples of configuration issues and their required inputs, see Example Configuration Issues.

Migrate Configuration of NSX-V with vRealize Automation

After you have resolved all configuration issues, you can migrate the configuration. When the configuration is migrated, configuration changes are made in the NSX-T environment to replicate the NSX-V configuration.

Prerequisites

Verify that you have completed the **Resolve Configuration** step.

Procedure

1 From the **Migrate Configuration** page, click **Start**.

- The configurations of the pre-created NSX-V objects are migrated to NSX-T.
- The configurations of the vRealize Automation created resources are migrated to NSX-T.

2 Verify that all the migrated configurations are displayed in your NSX-T environment.

You can verify the migrated configurations either in the NSX-T NSX Manager interface or by running the NSX-T APIs.

Note

- During the **Migrate Configuration** step, Security Tags from NSX-V are not migrated to NSX-T. Therefore, the Security Tag-based migrated dynamic Groups and Groups with static memberships in NSX-T are empty after this step is finished. The reason is that in NSX-V, a Security Tag is an object, whereas in NSX-T, a tag is an attribute of a VM. The tags are applied to the workload VMs only after the workloads are migrated to NSX-T during the **Migrate Hosts** step.
- When the configurations are migrated to NSX-T, the configuration changes are made in the NSX-T NSX Manager database, but it might take some time for the configurations to take effect.

3 Click **Continue** to proceed to the **Check Realization** page.

If needed, you can roll back the migrated configurations. Rolling back does the following:

- Remove the migrated configuration from NSX-T.
- Roll back all the resolved issues in the previous step.
- Remove the vRealize Automation deployment configuration file.

Check Realized Configurations in NSX-T

The **Check Realization** step is a preventive or a safety step during which the migration coordinator checks for any unrealized NSX-T configurations. This step provides additional time that the migration coordinator might need to ensure that all unrealized configuration errors are resolved before you can proceed to Edge migration.

Prerequisites

Verify that you have completed the **Migrate Configuration** step.

Procedure

1 On the **Check Realization** page, click **Start**.

2 Review realization errors, if any, and resolve them.

Sometimes, the configuration of some NSX-T objects might take longer to take effect. Wait for some seconds and click **Start** again.

Results

- If the Check Realization fails due to any unrealized configurations in NSX-T, a list of errors is shown in the migration coordinator UI. If the unrealized configuration errors are due to any NSX-T Policy objects, alarms are displayed on the **Alarms** page of the NSX Manager interface. Review the messages that are shown in the alarms and resolve them. After the alarms are resolved, run the **Check Realization** step again.
- When the **Check Realization** step is successful, the migration coordinator creates an intermediate or a partial output mapping file.

Migrate NSX-V Edges

After you have migrated the configuration, you can migrate the NSX-V Edge Services Gateway to NSX-T.

If you have no Edge Services Gateway appliances in your topology, you must still click **Start** so that you can proceed to the **Migrate Hosts** step.

Prerequisites

- All configuration issues must be resolved.

- The configuration of NSX-V objects and vRealize Automation created resources must be migrated to NSX-T.
- Verify that the migrated configurations are shown in the NSX Manager UI or API of NSX-T.
- Verify that you have a backup of NSX-V and vSphere since the most recent configuration changes were made.
- If you are using new IP addresses for the NSX-T Edge node uplinks, you must configure the northbound routers with these new BGP neighbor IP addresses.
- Verify that you have created an IP pool for Edge Tunnel End Points (TEP). See Create an IP Pool for Edge Tunnel End Points.
- Logical router interfaces created in NSX-T use the global default MTU setting, which is 1500. If you want to ensure that all logical router interfaces have a larger MTU, you can change the global default MTU setting. For more information, see Change the Global MTU Setting.

  If MTU setting other than 1500 is used on peering routers, the same should be configured on NSX-T. In case of OSPF topologies, OSPF adjacencies can get stuck if MTU setting is different from peering routers' MTU setting.

Procedure

1 From the **Migrate Edges** page, click **Start**.

  All Edges are migrated. The uplinks on the NSX-V Edge Services Gateways are internally disconnected, and the uplinks on the NSX-T Edge nodes are brought online.

2 Verify that routing and services are working correctly in the new NSX-T environment.

  If so, you can migrate the hosts. Before migrating the hosts, see Configuring NSX-V Host Migration.

Results

The following changes result from the migration process:

- The routing and service configuration from NSX-V Edge Services Gateway (ESG) are transferred to the newly created NSX-T Edge nodes.
- The new TEP IP addresses for the newly created NSX-T Edge nodes are configured from a newly created IP pool for Edge Tunnel End Points.

What to do next

Optional: Download and review the partial output mapping file after migrating the NSX-V Edge Services Gateways. However, no new information is added to the partial mapping file after the **Migrate Edges** step is completed.

Migrate NSX-V Hosts

After you have migrated Edge Services Gateway VMs to NSX-T Edge nodes, and verified that routing and services are working correctly, you can migrate your NSX-V hosts to NSX-T host transport nodes.

You can configure several settings related to the host migration, including migration order and enabling hosts. Make sure that you understand the effects of these settings. See Configuring NSX-V Host Migration for more information. Understanding the host migration settings is especially important if you use Distributed Firewall or vSphere Distributed Switch 7.0 or later.

For more information about what happens during host migration, see Changes Made During Host Migration in an End-to-End Migration.

If the Security Policies in your NSX-V environment use a partner service for Guest Introspection or Network Introspection or both, choose the host migration mode as shown in the following table.

Partner Service	Host Migration Mode
Only Guest Introspection	In-Place and Maintenance modes are supported.
Only Network Introspection	Maintenance mode is supported. However, Automated Maintenance mode is recommended. In-Place mode is not supported.
Both Guest Introspection and Network Introspection	Maintenance mode is supported. In-Place mode is not supported.

Important Consult the VMware partner before migrating their service that is running on the NSX-V workloads. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their services to NSX-T.

Caution Host migration should be completed during the same maintenance window as Edge migration.

You must disable IPFIX and reboot the ESXi hosts before migrating them.

If the partner service in your NSX-V environment provides Guest Introspection or both Guest Introspection and Network Introspection service, follow the procedure in this topic to migrate cluster-by-cluster. After all the host clusters are migrated to NSX-T, do a host-based service deployment in each NSX-T cluster.

If the partner service in your NSX-V environment provides only Network Introspection service, use the host migration approaches that are explained in Migrate Hosts with Network Introspection Service.

Prerequisites

- Verify that Edge migration has finished and all routing and services are working correctly.

- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.

Procedure

1 On the **Migrate Hosts** page, click **Start**.

If you selected the **In-Place** or **Automated Maintenance** migration mode for all hosts groups, the host migration starts. Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

2 If you selected the **Manual Maintenance** migration mode for any host groups, you must complete one of the following tasks for each VM so that the hosts can enter maintenance mode.

Option	Action
Power off or suspend VMs.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b After the host has migrated, attach the VM interfaces to the appropriate NSX-T segments and power on the VM.
Move VMs using vMotion.	Right click the VM and select Migrate. Follow the prompts to move the VM to a different host. Note that Migration Coordinator maintains security during migration by vMotioning VMs to specific ports that are protected by temporary rules. In the case of manual vMotion, the VMs will not be moved to those ports and there could be a security breach. To vMotion manually, the VMs must be migrated using vSphere API where the networking backing must point to the OpaqueNetwork ID corresponding to the NSX Segment when using NVDS or the VDS portgroup ID when using VDS 7. In both cases, the network device's externalId must be set to the the string "VM_UUID:vNIC_ID", where VM_UUID is the VM's instance UUID and vNIC_ID is the VM's vNIC index where the first vNIC is 4000.
Move VMs using cold migration.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b Right click the VM and select Migrate. Follow the prompts to move the VM to a different host, connecting the VM interfaces to the appropriate NSX-T segments.

Here is python code to specify an external-id for each vNIC in a VM and then vMotion the VM so that the vNICs will connect to an NSX-T segment of ID “ls_id” at the correct ports:

```
devices = vmObject.config.hardware.device
nic_devices = [device for device in devices if isinstance(device,
vim.Vm.device.VirtualEthernetCard)]
vnic_changes = []
for device in nic_devices:
    vif_id = vmObject.config.instanceUuid + ":" + str(device.key)
    vnic_spec = self._get_nsxt_vnic_spec(device, ls_id, vif_id)
    vnic_changes.append(vnic_spec)
```

```
relocate_spec = vim.Vm.RelocateSpec()
relocate_spec.SetDeviceChange(vnic_changes)
# set other fields in the relocate_spec
vmotion_task = vmObject.Relocate(relocate_spec)
WaitForTask(vmotion_task)

def _get_nsxt_vnic_spec(self, device, ls_id, vif_id):
    nsxt_backing = vim.Vm.Device.VirtualEthernetCard.OpaqueNetworkBackingInfo()
    nsxt_backing.SetOpaqueNetworkId(ls_id)
    nsxt_backing.SetOpaqueNetworkType('nsx.LogicalSwitch')
    device.SetBacking(nsxt_backing)
    device.SetExternalId(vif_id)
    dev_spec = vim.Vm.Device.VirtualDeviceSpec()
    dev_spec.SetOperation(vim.Vm.Device.VirtualDeviceSpec.Operation.edit)
    dev_spec.SetDevice(device)
    return dev_spec
```

For an example of a complete script, see https://github.com/dixonly/samples/blob/main/vmotion.py

The host enters maintenance mode after all VMs are moved, powered off, or suspended. If you want to use cold migration to move the VMs to a different host before the migrating host enters maintenance mode, you must leave at least one VM running while you move VMs. When the last VM is powered off or suspended, the host enters maintenance mode, and migration of the host to NSX-T starts.

Results

After a host has migrated to NSX-T using **In-Place** migration mode, you might see a critical alarm with message `Network connectivity lost`. This alarm occurs when a vSphere Distributed Switch (VDS) 6.5 or 6.7 migrates to an N-VDS because the host no longer has a physical NIC connected to the VDS it was previously connected to. To restore the migrated hosts to the Connected state, click **Reset to Green** on each host, and suppress the warnings, if any.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

If migration fails for a host, the migration pauses after all in-progress host migrations finish. When you have resolved the problem with the host, click **Retry** to retry migration of the failed host. If the host still fails to migrate, you can configure NSX-T on the host manually or remove the host from the system. In this case, at the end of the host migration step, the **Finish** button will not be enabled because of the host that failed to migrate. You need to call the REST API `POST https://<nsx-mgr-IP>/api/v1/migration?action=finalize_infra` (<nsx-mgr-IP> is the IP address of the NSX Manager where the migration service is running) using a REST API client (for example, postman or curl) to finish the migration, and then perform the post-migration tasks.

For information about troubleshooting other host migration problems, see Chapter 13 Troubleshooting Migration Issues.

What to do next

Download the final output mapping file after all the hosts in the cluster are migrated successfully to NSX-T. You must download the final mapping file before clicking **Finish**.

Important Verify that everything is working before clicking **Finish**. Then perform the post-migration tasks. Do not make any vSphere life cycle operations such as upgrading ESXi hosts, VDS, or VC before the post-migration tasks are completed.

For more information about the mapping file, see Output Mapping File.

If the migrated Security Policies use a third-party partner service, deploy an instance of the partner service in NSX-T. For detailed instructions, see:

- Deploy a Partner Service for Endpoint Protection

  Click this link to deploy a partner service that provides only Endpoint Protection or both Endpoint Protection and Network Introspection services to the NSX-T workload VMs.

- Deploy a Partner Service for Network Introspection

  Click this link to deploy a partner service that provides only Network Introspection service to the NSX-T workload VMs.

Post-Migration Tasks

After migration has finished, some additional actions are required.

- If you migrated from NSX-V 6.4.4, perform a reboot of all hosts that have migrated to NSX-T. The reboot must be done before you upgrade to a later version of NSX-T.
- During migration, all transport nodes are added to a group called `NSGroup with TransportNode for CPU Mem Threshold`. This group ensures that the transport nodes have the correct CPU memory threshold settings in NSX-T. This group is required after migration has completed. If you need to remove a transport node from NSX-T after migration and you are running NSX-T 3.2.0, you must first remove the transport node from this group. If you are running NSX-T 3.2.1 or later, you do not need to remove the transport node from this group.

  To remove the transport node from the group, make sure you are in **Manager** mode and then select **Inventory > Groups** to remove the transport node from the `NSGroup with TransportNode for CPU Mem Threshold` group. For more information about Manager mode, see the topic "NSX Manager" in the *NSX-T Data Center Administration Guide*.

- Verify that you have a valid backup and restore configuration. See "Backing Up and Restoring the NSX Manager" in the *NSX-T Data Center Administration Guide*.

Import Final Output Mapping File into vRealize Automation

Share the final output mapping file with the vRealize Automation administrator. The vRealize Automation administrator must import this output mapping file in step 4 of the vRealize Automation NSX-V to NSX-T migration plan.

The vRealize Automation administrator must complete the following steps in the migration plan:

- Step 5: Migrate the NSX-V cloud account and its related objects to NSX-T.
- Step 6: Test the migration results.
- Step 7: Remove cloud accounts from maintenance mode and exit the migration plan.

For more information, see the vRealize documentation about migrating NSX-V to NSX-T.

Uninstall NSX-V After Migration

After the migration is verified and successful, uninstall your NSX-V environment. The process for uninstalling NSX-V after migration to NSX-T is different from the standard uninstall for NSX-V.

For more information, see Uninstalling NSX-V After Migration.

Finish Deploying the NSX Manager Cluster

You can run the migration coordinator tool with only one NSX Manager appliance deployed. Deploy two additional NSX Manager appliances before you use your NSX-T environment in production.

See the *NSX-T Data Center Installation Guide* for the following information:

- NSX Manager Cluster Requirements
- Deploy NSX Manager Nodes to Form a Cluster from UI
- Configure a Virtual IP (VIP) Address for a Cluster

Migrating NSX-V with vRealize Automation - User-Defined Topology

11

If the topology of your NSX-V with vRealize Automation (vRA) environment is not one of the supported fixed topologies, you can migrate using the user-defined topology option.

Read the following topics next:

Overview

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.
- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.
- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 11-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	

Table 11-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.

Table 11-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. However, if either the primary or secondary NSX Manager is set to a standalone or transit mode, the migration is supported.

NSX-V Configuration	Supported	Details
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	
Network I/O Control (NIOC) having vNIC with reservation	No	
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	
Host profiles	No	
ESXi lockdown mode	No	Not supported in NSX-T.
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	

NSX-V Configuration	Supported	Details
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.
VXLAN port number other than 4789	No	
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.

NSX-V Configuration	Supported	Details
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	

Details	Supported	Notes
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy

NSX-V Configuration	Supported	Details
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag

NSX-V Configuration	Supported	Details
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.

NSX-V Configuration	Supported	Details
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.

NSX-V Configuration	Supported	Details
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.

NSX-V Configuration	Supported	Details
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 11-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 11-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.

Table 11-3. DHCP Features (continued)

NSX-V Configuration	Supported	Details

DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration.
		For example, dhcp option '80' is not migrated.
		<pre><dhcpOptions>
		<other>
		<code>80</code>
		<value>2f766172</value>
		</other>
		</dhcpOptions></pre>

Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 11-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.

NSX-V Configuration	Supported	Details
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ VMs not on NSX-V hosts	No	NSX-T does not support referencing objects not connected to NSX-T portgroups or networks. Those objects will be lost from the source or destination when migration is completed. To avoid this issue use IP addresses in NSX-V to reference those objects before migration.
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group

NSX-V Configuration	Supported	Details
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.

NSX-V Configuration	Supported	Details
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	

NSX-V Configuration	Supported	Details
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 11-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of ‘Groups’ after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 11-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.

Table 11-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.

Table 11-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 11-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 11-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 11-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 11-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000

Table 11-10. Single-Site Limits (continued)

Feature	Limit
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 11-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Changes Made During Host Migration in an End-to-End Migration

During the host migration step in an end-to-end migration, changes are made to migrate NSX-V hosts to NSX-T hosts.

- NSX-V software is uninstalled.
- NSX-T software is installed.
- For vSphere Distributed Switch versions 6.5.0 and 6.6.0:

  Hosts are configured with N-VDS to replace vSphere Distributed Switches:

  - Each N-VDS is created with a name that references the distributed switch name. For example, distributed switch `ComputeSwitchA` is created as N-VDS `nvds.ComputeSwitchA`.
  - If different clusters use different distributed switches to back logical switches, an N-VDS is created with a name that combines all the distributed switch names. For example, if `ComputeCluster1` and `ComputeCluster2` use distributed switch `ComputeSwitchA` to back logical switches and `ComputeCluster3` uses `ComputeSwitchB` to back logical switches, the N-VDS is created as `nvds.ComputeSwitchA.ComputeSwitchB`.
  - PNICs and vmks in the vSphere Distributed Switch are migrated to N-VDS.
  - NSX-V VTEPs are migrated to NSX-T TEPs.
- For vSphere Distributed Switches version 7.0:

  Hosts configured for vSphere Distributed Switch version 7.0 continue using the same switch after migration.

  - PNICs and vmks in the vSphere Distributed Switch remain connected on the same vSphere Distributed Portgroups.
  - NSX-V VTEPs are migrated to NSX-T TEPs and connected to standalone ports on the same vSphere Distributed Switch.

Note: If NSX-V has multiple VTEPs and a single LAG configuration, after migration the LAG will have a single TEP and in failover mode.

Virtual Machine Deployment During an End-to-End Migration

After you start an end-to-end migration, do not change the NSX-V environment. If you want to deploy VMs during the migration, wait until some of the NSX-V hosts have migrated to NSX-T and deploy the VMs on NSX-T hosts. Connect the VMs to NSX-T segments and install VMware Tools on the VMs.

Deploying on NSX-T with VMware Tools installed ensures that the VMs are populated into security groups and receive the intended Distributed Firewall policies.

Caution VMs deployed without VMware Tools installed, or deployed on NSX-V do not receive the intended Distributed Firewall policies.

If you use vSphere templates to deploy VMs, update the templates to use NSX-T segments for the VM network configuration. Specifying NSX-T segments ensures that any VMs deployed using the templates are deployed on NSX-T hosts.

If you use automation tools to deploy VMs on vSphere, but do not use vSphere templates, you might need to change your automation tool configuration to ensure that the VMs are deployed on NSX-T.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1
	Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Understanding the Migration of NSX-V with vRealize Automation

You can migrate an NSX-V environment and its existing integration with vRealize Automation to a new NSX-T environment.

You must deploy a new NSX-T environment for this migration. You cannot merge vRealize Automation deployments from your NSX-V environment into an existing NSX-T environment, which is preconfigured and used.

In addition to setting up a new NSX-T environment in advance, migration preparation might also require you to modify your existing NSX-V environment.

You should be familiar with NSX-T concepts, vRealize Automation concepts, and administration tasks in both environments before you migrate.

Deployment Configuration File

A deployment configuration file is an input to the migration coordinator tool. The migration coordinator reads the `.json` configuration file, validates it against a pre-defined JSON schema,

and migrates the vRealize Automation resource configurations from the existing NSX-V environment to a new NSX-T environment.

The deployment configuration file contains the following configuration information:

- List of resources that vRealize Automation has created.
- List of resources that vRealize Automation references from the existing NSX-V environment.
- List of network interfaces of the workload VMs that vRealize Automation has created.
- Desired mapping of the vRealize Automation resources to NSX-T objects.

Mapping example: Consider that your topology has a vRealize Automation created Logical Switch that connects to an existing Distributed Logical Router in NSX-V. The mapping information in the deployment configuration file tells the migration coordinator to do the following:

- Create a relevant tier-1 gateway in NSX-T that maps to the Network Profile deployed in vRealize Automation.
- Attach an NSX-T overlay segment to the downlink of this tier-1 gateway.

The deployment configuration file is generated from the vRealize Automation environment, and it provides configuration information about all the vRealize Automation created on-demand resources that are deployed over several days of network operations. This configuration file must be valid and conform to a pre-defined JSON schema.

Typically, the cloud administrator generates the deployment configuration file and provides it to the network administrator, who runs the migration coordinator tool in the NSX-T environment. In some organizations, a single administrator might perform both roles.

Remember, each vRealize Automation deployment can create multiple on-demand networks, and each network can have multiple network interfaces. The networks inside a deployment can consume resources that are either created by vRealize Automation or that are existing in NSX-V.

For example, let us assume that over several days of network operations, you have created multiple deployments in vRealize Automation to deploy the following resources in your NSX-V topology:

- Deployment 1: For creating two on-demand private networks and one on-demand routed network.
- Deployment 2: For creating on-demand security groups.
- Deployment 3: For creating on-demand routed networks with no services.

When you are ready to migrate NSX-V with vRealize Automation to NSX-T, a single deployment configuration file is generated from vRealize Automation, and uploaded as an input file to the migration coordinator.

Output Mapping File

After the resource configurations in the vRealize Automation deployment configuration file are migrated from NSX-V to NSX-T, the migration coordinator creates an output mapping file in a `.json` format.

This mapping file contains information about how vRealize Automation created resources in NSX-V map to NSX-T objects.

vRealize Automation can use this mapping file to update or refresh its own database after the migration is finished. The migration coordinator places this mapping file at the following path on the NSX Manager appliance where the migration coordinator service runs: `/var/log/migration-coordinator/v2t`.

The migration coordinator creates a partial or an intermediate mapping file for the first time after the Check Realization step is completed. In the subsequent steps of the migration process, the mapping file changes. For example, mapping information about vRealize Automation created Security Policies, VM network interfaces, Security Groups, and so on, are available in the output mapping file after host migration is completed. The final mapping file is generated after all the hosts in the clusters, which are enabled for migration, are migrated to NSX-T.

You can download the output mapping file from the migration coordinator UI after the following stages in the migration process are complete:

- Check Realization step
- Migrate Edges step
- Migrate Hosts step

The information in the mapping file varies depending on the migration stage that is completed.

For example:

- The details in the mapping file before host migration and after host migration can change. Before the hosts are migrated, the mapping file does not contain the network interfaces (vNIC IDs) of the workload VMs that are created in vRealize Automation. The network interfaces of the workload VMs are mapped only after the VMs connect to NSX-T segments during host migration.
- The mapping file shows the new network interfaces of the static VM members in the migrated Security Groups only after all the hosts are migrated to NSX-T.

Scenario: Deleted Resources List in Mapping File

Consider the following sequence of events:

1 Your deployment configuration file contains several vRealize Automation created resources. One of the resources is an on-demand Security Group.

2 Before running the migration coordinator tool, during a day-two operation, you delete this vRealize Automation created Security Group in the NSX-V UI. The NSX-V inventory does not have the Security Group.

3 Let us assume that the data about this deleted Security Group is not yet reconciled with the vRealize Automation database. Therefore, the vRealize Automation database continues to reference this deleted Security Group.

4 Now, you run the migration coordinator tool.

After migration to NSX-T, the migration coordinator shows the deleted Security Group in the list of "output deleted resources" of the final mapping file. Because the migration coordinator cannot find the ID of the deleted Security Group, the mapping file does not have the path to the corresponding NSX-T Group.

High-Level View of Migrating NSX-V with vRealize Automation

The migration process includes setting up a new NSX-T environment and running the migration coordinator. You also might need to change your existing NSX-V environment to ensure that it can migrate to NSX-T.

Caution Deploy a new NSX-T environment to be the destination for the NSX-V migration.

During the **Import Configuration** step, all NSX Edge node interfaces in the destination NSX-T environment are shut down. If the destination NSX-T environment is already configured and is in use, starting the configuration import will interrupt traffic.

During the migration, you will complete the following steps:

1 Create a new NSX-T environment.

- Deploy a single NSX Manager appliance to create the NSX-T environment.
- If you plan to use Maintenance Mode migration for hosts, configure a shared storage on the cluster to be migrated from NSX-V to NSX-T. This enables automated vMotion for the migration process. Any VMs that do not meet this criterion must be manually powered off prior to migration or manually vMotioned.
- Configure a compute manager in the NSX-T environment. Add the vCenter Server as a compute resource.

  Important Use the exact IP or hostname specified in NSX-V vCenter Server registration.
- Start the migration coordinator service.
- If you want to import users from NSX-V, set up VMware Identity Manager.
- If your NSX-V topology uses Edge Services Gateways, create an NSX-T IP pool to use for the NSX-T Edge TEPs. These IPs must be able to communicate with all the existing NSX-V VTEPs.
- Deploy NSX Edge nodes.
  - Deploy the correct number of appropriately sized NSX-T Edge appliances.
  - Join the Edge nodes to the management plane from the command line.

- If Security Policies in your NSX-V environment use third-party Guest Introspection or Network Introspection services provided by partners, the partner services must be registered with NSX-T.

2 Import configuration from NSX-V.

- Enter the details of your NSX-V environment.
- Upload a vRealize Automation deployment configuration file in `.json` format as an input to the migration coordinator.
- The configuration is retrieved and pre-checks are run.

3 Resolve issues with the configuration.

- Review messages and the reported configuration issues to identify any blocking issues or other issues that require a change to the NSX-V environment.
  - If you make any changes to the NSX-V environment while a migration is in progress, you must restart the migration and import the configuration again.
- Provide inputs to configuration issues that must be resolved before you can migrate your NSX-V environment to NSX-T. Resolving issues can be done in multiple passes by multiple people.

4 Migrate configuration.

- After all configuration issues are resolved, you can migrate the configuration to NSX-T. Configuration changes are made on NSX-T, but no changes are made to the NSX-V environment yet.

5 Check realization of migrated configurations in NSX-T.

- Check whether the migration coordinator shows any unrealized configuration errors.
- Download an intermediate output mapping file and share it with the vRealize Automation cloud administrator for an initial review.
- Take appropriate actions to resolve any missing configurations in NSX-T before proceeding to Edge migration.

6 Migrate Edges.

- Routing and Edge services are migrated from NSX-V to NSX-T.

  Caution North-South traffic is interrupted during the Migrate Edges step. All traffic that was previously traversing through the Edge Services Gateways (North-South traffic) moves to the NSX-T Edges.

7 Migrate Hosts.

- NSX-V software is removed from the hosts, and NSX-T software is installed. VM interfaces are connected to the new NSX-T segments.

  Caution If you select **In-Place** migration mode, there is a traffic interruption for a few seconds during the Migrate Hosts step. However, if you select **Maintenance** migration mode, traffic interruption does not occur.

- Download the final output mapping file after the hosts in all the clusters are migrated to NSX-T.

8 Finish Migration.

- After you have verified that the new NSX-T environment is working correctly, you can finish the migration, which clears the migration state.

9 Perform post-migration tasks.

- Deploy two additional NSX Manager appliances before you use your NSX-T environment in a production environment.
- Uninstall NSX-V environment.

Note: Logical ports and switches that are created during migration are not deleted when the workload VMs are deleted. You must delete these ports and switches via the NSX Manager UI or the API.

Migration-Supported Operations in NSX-V for vRealize Automation Resources

Over several days of network operations in NSX-V, you can update the configurations of the on-demand resources that you defined earlier in vRealize Automation.

You can directly use the NSX-V UI or APIs to update the previously defined vRealize Automation resource configurations. The migration coordinator fetches the updated configurations from NSX-V and migrates them to NSX-T.

In the context of migration, the following network operations in NSX-V are supported for vRealize Automation created resources.

Table 11-12. Supported Operations in NSX-V for vRealize Automation Resources

vRealize Automation Created Resource	Supported Network Operations in NSX-V
On-Demand Networks	■ Add or delete VMs from existing on-demand networks that are created using vRealize Automation. ■ Delete an existing on-demand network after deleting or moving workload VMs that are created using vRealize Automation.
On-Demand Security Groups	■ Delete Security Groups that are created using vRealize Automation. ■ Add Security Groups that are created with vRealize Automation to an existing Security Group (out-of-band configuration). ■ Add or change the Security Group in Security Policies that are created with vRealize Automation. ■ Add new firewall rules in Security Policies that are created with vRealize Automation.
DHCP Server configuration	■ Delete or disable DHCP server configuration that is created with vRealize Automation.
Workload VMs	■ Change the network interface (vNIC) assignment of a vRealize Automation created VM from dynamic to static, or reverse, or from static IP1 to static IP2. ■ Change from a single vNIC on a vRealize Automation created VM to multiple vNICs, and reverse.
Load balancer configuration	■ Edit health check related parameters in the load balancer ■ Add or delete pool members ■ Disable load balancer service on the ESG
NAT configuration	■ Modify existing rules ■ Delete existing rules ■ Reorder rules ■ Add port forwarding rules

Preparing for an NSX-V with vRA User-Defined Topology Migration

Preparing the NSX-V Environment

Perform the tasks in the following sections to prepare the NSX-V environment.

Check the Configurations of the NSX-V Environment

Check your NSX-V environment before starting an end-to-end migration.

System State

Check the following system states:

- If your environment is vSphere 7.0 or later, upgrade the VDS to 7.0 or later.
- Verify that the NSX-V components are in a green state on the NSX Dashboard.
- Verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering maintenance mode.
- Verify that no NSX-V upgrades are in progress.
- Verify the publish status of Distributed Firewall and Service Composer to make sure that there are no unpublished changes.
- You can have vSphere High Availability (HA) enabled if the NSX-V environment has VDS 7.0 or later.

  Note: HA is not supported for previous versions of VDS This is because if the NSX-V environment has VDS 6.5 or 6.7, and the vmkernel ports (vmk's) are attached to VDSes, during an in-place migration, the hosts and VMs may lose network connectivity for a period of time long enough to trigger HA. The HA mechanism will try to power off, migrate and restart VMs. This might fail because the NSX-V environment is being migrated to NSX-T. As a result, after the migration, VMs might remain in a powered-off state or have no network connectivity if powered on. To avoid this situation, disable HA or attach the management vmk to a VSS before starting the migration.

General Configuration

- Back up the NSX-V and vSphere environments. See "NSX Backup and Restore" in the *NSX Administration Guide*.
- The VXLAN port must be set to 4789. If your NSX-V environment uses a different port, you must change it before you can migrate. See "Change VXLAN Port" in the NSX-V *NSX Administration Guide*.

Controller Configuration

- NSX-V transport zones using multicast or hybrid replication mode are not supported for migration. An NSX Controller cluster is required if VXLAN is in use. VLAN-backed micro-segmentation topologies do not use VXLAN and so do not require an NSX Controller cluster.

Host Configuration

- On all host clusters in the NSX-V environment, check these settings and update if needed:
  - Set vSphere DRS accordingly.

Disable vSphere DRS if one of the following apply:

- **In-Place** migration mode will be used. In this mode hosts are not put in maintenance mode during migration and VMs will experience a network outage and network storage outage during the migration. This mode is only available if the environment is vSphere 6.x (VDS will be migrated to N-VDS).
- **Manual Maintenance** migration mode will be used. If you decide to use vMotion for migrating VMs, you can disable vSphere DRS, or set the vSphere DRS automation level to Manual, Partially Automated, or Fully Automated.
- **Automated Maintenance** migration mode will be used and the VDS version is 6.5 or 6.7.

Set vSphere DRS mode to Fully Automated if:

- **Automated Maintenance** migration mode will be used and the VDS version is 7.0.

Note that in **Automated Maintenance** mode, Migration Coordinator will not reconfigure VMs that are powered off. After migration, you need to manually configure these VMs before powering them on.

- To migrate Network Introspection service rules, use the **Maintenance** host migration mode. **In-Place** migration mode is not supported.
- If you have hosts that have NSX-V installed, but are not added to a vSphere Distributed Switch, you must add them to distributed switches if you want to migrate them to NSX-T. See Configure Hosts Not Attached to vSphere Distributed Switches for more information.
- On each cluster that has NSX-V installed, check whether Distributed Firewall is enabled. You can view the enabled status at **Installation & Upgrade > Host Preparation**.

  If Distributed Firewall is enabled on any NSX-V clusters before migration, Distributed Firewall is enabled on all clusters when they migrate to NSX-T. Determine the impact of enabling Distributed Firewall on all clusters and change the Distributed Firewall configuration if needed.

Edge Services Gateway Configuration

- You might need to make changes to your NSX-V route redistribution configuration before migration starts.
  - Redistribution filters are not migrated. For BGP, filters can be moved to the BGP neighbor level.
  - After migration, dynamically learned routes between Distributed Logical Router and Edge Services Gateway are converted to static routes and all static routes are redistributed in BGP or OSPF. If you need to filter any of these routes, you can configure them at the BGP neighbor level or manually configure the redistribution rules on NSX-T after the configuration migration is completed and before cutover. Note that if you roll back, the manual configuration of redistribution rules will also be removed.

    - The default MTU setting is 1500 on NSX-T. If you have non-default MTU setting requirements, you can change the setting. See Change the Global MTU Setting.
- NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the **Migrate Configuration** step fails.
- If you have an Edge Services gateway performing one-armed load balancer function, you must change the following configurations if present before you import the configuration:
    - If the Edge Services Gateway has an interface configured for management, you must delete it before migration. You can have only one connected interface on an Edge Services Gateway providing one-arm load balancer function. If it has more than one interface, the **Migrate Configuration** step fails.
    - If the Edge Services Gateway firewall is disabled, and the default rule is set to deny, you must enable the firewall and change the default rule to accept. After migration the firewall is enabled on the tier-1 gateway, and the default rule accept takes effect. Changing the default rule to accept before migration prevents incoming traffic to the load balancer from being blocked.
- Verify that Edge Services Gateways are all connected correctly to the topology being migrated. If Edge Services Gateways are part of the NSX-V environment, but are not correctly attached to the rest of the environment, they are not migrated.

  For example, if an Edge Services Gateway is configured as a one-armed load balancer, but has one of the following configurations, it is not migrated:
    - The Edge Services Gateway does not have an uplink interface connected to a logical switch.
    - The Edge Services Gateway has an uplink interface connected to a logical switch, but the uplink IP address not does match the subnet associated with the distributed logical router that connects to the logical switch.

Security Configuration

- If you plan to use vMotion to move VMs during the migration, disable all SpoofGuard policies in NSX-V to prevent packet loss.
    - Automated Maintenance mode uses DRS and vMotion to move VMs during migration.
    - In Manual Maintenance mode, you can optionally use vMotion to move VMs during migration.
    - In-Place migration mode does not use vMotion.

Security Group Configuration

If existing Security Policies contain Guest Introspection service rules that are applied to Security Groups with static VM members or dynamic members other than VMs, do these steps:

1 Create new Security Groups with VMs only in the dynamic membership criteria. Make sure that the dynamic membership criteria produces the same effective VM members as your original Security Groups.

2 Before starting the migration, update the existing Security Policies to apply the new Security Groups to the Guest Introspection service rules.

   If you prefer not to update your existing Security Policies before the migration, you can still keep the new Security Groups ready with the correct dynamic membership criteria in your NSX-V environment. In the **Resolve Configuration** step of the migration process, you will be prompted to provide alternative Security Groups.

Service Composer Synchronization

Ensure that the Service Composer is in sync with Distributed Firewall before you start the migration. A manual synchronization ensures that if you make any last-minute changes in the policy configuration before starting the migration, these changes are applied to the Security Policies that are created using Security Composer too. For example, you edit the name of the Security Group that is used in a firewall rule before starting the migration.

To verify whether the Service Composer status is in sync, do these steps:

1 In the vSphere Client, navigate to **Networking and Security > Security > Service Composer**.

2 Click the **Security Policies** tab.

3 Verify that the Sync Status is **In Sync**. If it is not in sync, click **Synchronize**.

As a best practice, always click the **Synchronize** button before starting the migration even when the sync status is green. Do this manual synchronization regardless of whether you performed any last-minute changes in the policy configuration.

During migration, if the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of the Security Policies created using the Service Composer by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get the Service Composer in sync with Distributed Firewall, and restart the migration.

Configure Hosts Not Attached to vSphere Distributed Switches

An NSX-V environment can contain hosts that have NSX-V installed, but are not added to a vSphere Distributed Switch. You must add the hosts to a vSphere Distributed Switch before you can migrate them.

You can use a distributed switch you already have in your environment, or create a new distributed switch for this purpose. Right click the distributed switch and select **Add and Manage Hosts** to add the hosts to the distributed switch. You do not need to assign physical uplinks or VMkernel network adapters to the distributed switch.

See "Add Hosts to a vSphere Distributed Switch" in the *vSphere Networking Guide* for more information.

If you import the configuration before you make this change, you must restart the migration to import the updated configuration. See Make Changes to the NSX-V Environment.

After the migration has finished, the hosts are no longer required to be attached to the distributed switch.

- If you added the hosts to an existing distributed switch, you can remove them from the distributed switch.
- If you added the hosts to a new distributed switch that you are not using for another purpose, you can delete the distributed switch.

Tag Management VMs in a Collapsed Cluster Environment

You can migrate an environment that uses a collapsed cluster.

In a collapsed cluster design, all management VMs, workload VMs, and optionally edges run on the same vSphere cluster that is prepared for NSX-V. The management VMs of the NSX-T must be initially attached to dvPortgroups. After migration, the management VMs of NSX-T will be attached to NSX-T VLAN segments.

The management VMs in the NSX-T include appliances such as NSX Manager, vCenter Server, VMware Identity Manager, and so on. The NSX-T VLAN segment ports to which these management VMs connect are blocked in two cases: when these management VMs are rebooted after they were migrated in-place by the Migration Coordinator, or when they are moved from NSX-V hosts to NSX-T hosts by vMotion in maintenance migration mode. Therefore, the management VMs might lose connectivity in such cases.

To prevent this problem, create a "management_vms" tag category, and add tags in this category. Assign a tag from this category to all the management VMs in the NSX-T environment. These VMs will be attached to unblocked VLAN segment ports.

Procedure

1. Log in to the vSphere Client.
2. Click **Menu > Tags & Custom Attributes**.
3. Click **Categories**, and then click **New** to add a category.

   Create a category with name **`management_vms`**.
4. Click the **Tags** tab and add a tag in the management_vms category.
5. Navigate to **Menu > Hosts and Clusters**.
6. Expand the collapsed cluster from the left Navigator view, right-click the name of the NSX Manager VM, and select **Tags & Custom Attributes > Assign Tag**.
7. Assign a tag from the management_vms category to the NSX Manager VM.

8 Repeat steps 6 and 7 for all the management VMs in the cluster.

   For a detailed information about tag categories and tags, see the *vCenter Server and Host Management* documentation.

9 Log in to NSX Manager.

10 Navigate to **Inventory > Groups** and create a group.

11 Click **Set** to set members.

12 Set **Group Type** to **IP Addresses Only** and include the IP addresses of the management VMs that will be migrating from the dvPortgroups to the NSX-T segments.

13 Navigate to **Security > Distributed Firewall > Actions > Exclusion List** and add this group.

Delete Partner Service Deployments

If your NSX-V environment uses a partner service for Guest Introspection, or both Guest Introspection and Network Introspection, delete the partner service deployment before migration.

You must also delete the Guest Introspection instance (GI-SVM) so that the Guest Introspection module is uninstalled from the clusters.

If your NSX-V environment uses a partner service only for Network Introspection, you have the flexibility to decide whether to delete the partner service deployment before or after the migration. When a partner service deployment is deleted, the partner service virtual machines (SVMs) are removed from the NSX-V-prepared host cluster, and security protection is lost.

Important Consult the VMware partner before migrating their service that is running on the NSX-V workloads. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their services to NSX-T.

Procedure

1 Log in to the vSphere Client and navigate to **Networking and Security > Installation and Upgrade > Service Deployment**.

2 Select the deployed service and click **Delete**.

Preparing the NSX-T Environment

Perform the tasks in the following sections to prepare the NSX-T environment.

Deploy an NSX Manager Appliance

You must deploy a new NSX Manager appliance to run the migration coordinator. Do not deploy an NSX Global Manager.

In other words, you cannot merge your NSX-V environment into an existing NSX-T environment, which has NSX-T already installed on the vSphere host clusters.

For details on deploying a licensed version of the NSX Manager appliance, see *Install NSX Manager and Available Appliances* in the *NSX-T Data Center Installation Guide*.

Install one appliance to perform the migration. Deploy additional appliances to form a cluster after the migration is finished. See Finish Deploying the NSX Manager Cluster.

If you install the NSX Manager appliance on an ESXi host that is a part of the NSX-V environment that is migrating, do not attach the appliance interfaces to an NSX-V logical switch. To prevent the management VMs in NSX-T from losing connectivity after the VMs are rebooted post migration, tag the management VMs. For more information, see Tag Management VMs in a Collapsed Cluster Environment.

Add a Compute Manager

Before you can start the migration process, you must add the vCenter Server that is associated with NSX-V as a compute manager in NSX-T.

Prerequisites

Log into the NSX-V NSX Manager web interface to retrieve the settings used for vCenter Server registration. You must use the same settings. For example, if an IP address is specified, use the IP address and not the FQDN.

Procedure

1. From a browser, log in with admin privileges to NSX Manager at `https://<nsx-manager-ip-address>`.
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 vCenter Server. You can optionally describe any special details such as, the number of clusters in the vCenter Server.
FQDN or IP Address	Type the FQDN or IP address of the vCenter Server.
Type	The default compute manager type is set to vCenter Server.
HTTPS Port of Reverse Proxy	The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances. Set the reverse proxy port to register the compute manager in NSX-T Data Center.
Username and Password	Type the vCenter Server login credentials.
SHA-256 Thumbprint	Type the vCenter Server SHA-256 thumbprint algorithm value.

Option	Description
Create Service Account	Enable this field for features such as vSphere Lifecycle Manager that need to authenticate with NSX-T Data Center APIs. Log in with the administrator@vsphere.local credential to register a compute manager. After registration, the compute manager creates a service account. **Note** Service account creation is not supported on a global NSX Manager. If service account creation fails, the compute manager's registration status is set to `Registered with errors`. The compute manager is successfully registered. However, vSphere Lifecycle Manager cannot be enabled on NSX-T Data Center clusters. If a vCenter Server admin deletes the service account after it was successfully created, vSphere Lifecycle Manager tries to authenticate the NSX-T Data Center APIs and the compute manager's registration status is set to `Registered with errors`.
Enable Trust	Enable this field to establish trust between NSX-T Data Center and compute manager, so that services running in vCenter Server can establish trusted communication with NSX-T Data Center. For example, for vSphere Lifecycle Manager to be enabled on NSX-T Data Center clusters, you must enable this field. Supported only on vCenter Server 7.0 and later versions.
Access Level	Enable one of the options based on your requirement: ■ Full Access to NSX: Is selected by default. This access level gives the compute manager complete access to NSX-T Data Center. Full access ensures vSphere for Kubernetes and vSphere Lifecycle Manager can communicate with NSX-T Data Center. The vCenter Server 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-T Data Center. The vCenter Server 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-T Data Center to discover and register the vCenter Server 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 vCenter Server credentials and click **Resolve**.

If an existing registration exists, it will be replaced.

Results

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

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

After the vCenter Server is successfully registered, do not power off and delete the NSX Manager VM without deleting the compute manager first. Otherwise, when you deploy a new NSX Manager, you will not be able to register the same vCenter Server again. You will get the error that the vCenter Server is 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, NSX Intelligence VM, 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, any NSX Intelligence VM, 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 sucessfully remove the VC, ensure that you re-deploy an NSX Manager node from the VC.

This restriction applies to a fresh installation of NSX-T Data Center as well as an upgrade.

Change the Global MTU Setting

When an Edge Services Gateway (ESG) is migrated, the MTU setting of the interfaces is not migrated. A default value of 1500 is used. You can change the default value using the API.

You can also modify the MTU setting for the interfaces after the migration.

Procedure

1 Make the following API call to retrieve the current configuration.

```
GET /api/v1/global-configs/RoutingGlobalConfig
```

2 Change the value for `logical_uplink_mtu` and make the following call.

```
PUT /api/v1/global-configs/RoutingGlobalConfig
```

Create an IP Pool for Edge Tunnel End Points

If your NSX-V environment uses Edge Services Gateways, you must create an IP pool in the NSX-T environment for the Edge Tunnel End Points (TEP) before you start the migration.

Prerequisites

- Identify existing IP pools or DHCP ranges for NSX-V VTEPs.
- Determine which IP addresses to use to create an IP pool for Edge TEPs.

  The IP range and VLAN must not already be in use in the NSX-V environment.
- Verify that the NSX-T TEP IP addresses have network connectivity to the NSX-V VTEP IP addresses.

Procedure

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

2 Select **Networking > IP Management > IP Address Pools**.

3 Click **Add IP Address Pool**.

4 Enter a name for the new IP pool.

5 (Optional) Enter a description.

6 In the **Subnets** column, click **Set** to add subnets.

7 Specify the IP ranges.

   a Select **Add Subnets > IP Ranges**.

   b Enter IPv4 or IPv6 ranges.

   c Enter the subnet address in a CIDR format.

   d Enter the Gateway IP address for this subnet.

   e (Optional) Enter DNS servers.

   f (Optional) Enter DNS suffix.

   g Click **Add**, and then click **Apply**.

8 Click **Save**.

Plan the Mapping of the NSX-V Topology to the NSX-T Topology

Review the NSX-V topology and decide how to map it to the NSX-T topology. During the migration, the "Define a Topology" step will prompt you for the mapping.

Specifically, determine how many NSX-T Edge clusters you need and how Edge Service Gateways (ESGs) and Distributed Logical Routers (DLRs) should map to gateways in NSX-T. The northbound ESGs without any L4-L7 services should be skipped. These are usually the ESGs peering with northbound routers and are in ECMP path. If you are using VPN on a northbound ESG, migrating to active-standby tier-0 is recommended. In other cases, migrating the ESGs/DLRs to tier-1 is recommended. An ESG and a DLR can be merged in one mapping entry in the mapping file.

An example of a mapping file that maps ESGs to a tier-0 gateway:

```
[
  {
    "name":"nsxv-to-nsxt-mapping",
    "v_edges_to_policy_gateways_mappings":[
      {
        "v_edges":[
          "edge-1",
          "edge-2"
         ],
         "policy_gateway_name": "tier0-gateway"
         "policy_gateway_path": "/infra/tier-0s/tier0-gateway"
      }
    ]
  }
]
```

Deploy NSX Edge Nodes

You must deploy NSX Edge nodes as a virtual machine on ESXi using an OVA or OVF file.

Do not deploy on bare metal. Do not deploy from the NSX Manager user interface.

Snapshots of NSX appliances (including Edge node VMs) are not supported and must be disabled. For information on how to disable snapshots, see the topic Disable Snapshots on an NSX Appliance in the *NSX-T Data Center Installation Guide*.

NSX Edge nodes must be connected to trunk portgroups. To learn more about NSX Edge networking, see "NSX Edge Networking Setup" in the *NSX-T Data Center Installation Guide*.

Prerequisites

- You must have sufficient ESXi hosts with appropriate resources available to accommodate the NSX Edge appliances.

Procedure

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

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

2 In the vSphere Client, select the host on which to install NSX Edge 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 vCenter Server and vSphere inventory.

6 Select a compute resource for the NSX Edge node appliance, and click **Next**.

7 Review and verify the OVF template details, and click **Next**.

8 Select a deployment configuration and click **Next**.

See the **Import Configuration** step for details on the size of Edge nodes you must deploy.

9 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.

10 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.

Post-migration, the NSX Edge node is connected to one of these three trunk networks using only a single fastpath interface. The network settings can be adjusted or verified after the NSX Edge node is deployed.

11 Configure IP Allocation settings.

a For IP allocation, specify **`Static - Manual`**.

b For IP protocol, select **`IPv4`**.

12 Click **Next**.

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

13 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.

14 Enter the hostname of the NSX Edge.

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

   Skip any VMC network settings.

16 Enter the DNS Server list, the Domain Search list, and the NTP Server IP or FQDN list.

17 (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.

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

   The installation might take 7-8 minutes.

19 Start the NSX Edge node VM manually.

20 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.

21 After the NSX Edge node starts, log in to the CLI 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.

22 Run the `get interface eth0` (without VLAN) or `get interface eth0.<vlan_ID>` (with a VLAN) command to verify that the IP address was applied as expected.

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

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

23 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 IP or FQDN list.

24 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 use DHCP 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 Type the **`set interface`** ***`interface`*** **`dhcp plane mgmt`** command.

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

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

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

Join NSX Edge Node VM with the Management Plane

You must join the NSX Edge node VM you created to the management plane.

Do not join the NSX Edge node VM to the management plane using any other method. Do not create transport nodes from the NSX Edge node VM.

Procedure

1 Open an SSH session or console session to the NSX Manager appliance.

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 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.

5 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)
```

6 In the NSX Manager UI, you can navigate to **System > Fabric > Nodes > Edge Transport Nodes** and see the NSX Edge node. The **Configuration State** column will display **Pending**. If you click the name of the Edge node, you will be prompted to configure the node. Do not configure the node. The configuration will occur during the migration.

Configure NSX-Tfor a User-Defined Topology Migration

Before migrating a user-defined topology, you must configure NSX-T.

Regarding the configuration of Edge clusters, VLAN transport zone, and segments used for tier-0 gateway uplinks, you have two options:

- Let the migration coordinator configure an Edge cluster, VLAN transport zone, and other layer-2 entities needed for northbound connectivity. Note that only one Edge cluster will be created.
- Manually configure Edge clusters, VLAN transport zone, and other layer-2 entities needed for northbound connectivity. You must use this option if you want more than one Edge cluster. For more information about creating an Edge cluster, see the section "Create an NSX Edge Cluster" in the *NSX-T Data Center Installation Guide*.

After the Edge cluster is ready, do the following:

- Create tier-0 and tier-1 gateways and configure dynamic or static routing on the tier-0 gateways towards the northbound routers. For dynamic routing, configure BGP or OSPF, route redistribution as well as any filtering or route-maps based on your requirements. For more information about configuring static or dynamic routing (BGP or OSPF), see the section "Tier-0 Gateways" in the *NSX-T Data Center Administration Guide*.
- After configuring the dynamic routing on tier-0 gateways, check that the dynamic routing has converged, that is, BGP sessions are established or OSPF neighborships are FULL as applicable.

Other than the configurations mentioned above, no other configurations should be performed.

Register Third-Party Guest Introspection Service with NSX-T

If Security Policies in your NSX-V environment use third-party Guest Introspection service provided by a partner, register the partner service with NSX-T before you start the migration.

You might need to upgrade the Partner Console to register the service with the version of NSX-T that is used for this migration. For more information, see the partner documentation.

Complete the following procedure to register a partner service for endpoint protection with the NSX Manager in your NSX-T environment.

Procedure

1 Log in to the Partner Console with **Admin** privileges.

2 Update the NSX endpoint in the Partner Console. Specify the following details:

- IP address of NSX-T NSX Manager
- Port (default is 443)
- User name and password of the NSX-T NSX Manager

Make sure to test the connection before proceeding to the next step. If you need help with using the Partner Console, see the partner documentation.

The partner service and the vendor templates that are associated with this partner service are now created in NSX-T.

3 Verify that the partner service is registered with NSX-T.

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

b Navigate to **System > Service Deployments > Deployment**.

c Click the **Partner Service** drop-down menu, and check that the partner service is listed.

Register Third-Party Network Introspection Services with NSX-T

If Security Policies in your NSX-V environment use third-party Network Introspection services provided by partners, partner services must be registered with NSX-T before you start the migration.

You might need to upgrade the Partner Console to ensure that the partner service is registered with the version of NSX-T that is used for this migration. For more information, see the partner documentation.

The following types of east-west network introspection services are supported for migration:

- Intrusion detection services (IDS)
- Intrusion protection services (IPS)
- Network monitoring services
- Next-generation firewall services

A partner registers the service, vendor template, and the Partner Management Console/Partner Service Manager. Then, either you or the partner can create the service profile. It can vary from one partner to another. See the partner documentation.

In the following procedure, step 2 is required when your NSX-V environment uses only Network Introspection service.

If your environment uses a combination of both Guest Introspection and Network Introspection services from a single partner (partner A), partner does step 1. Step 2 is not required.

If your environment uses Guest Introspection service from one partner (partner A) and Network Introspection service from another partner (partner B), then:

- Use the Partner Console of partner A to register the Guest Introspection service. See the partner documentation for help on registering the service.
- Partner B registers the Network Introspection service (step 1 of the procedure). Either you or the partner can create the service profile, as explained in step 2.

Procedure

1 Partner registers the partner service, vendor template, and the partner Service Manager using NSX-T APIs.

2 Create a service profile to specify attributes of a vendor template for a given partner service.

   For a network introspection service, multiple service profiles can be associated with a single vendor template.

   You can create a service profile either by using the NSX-T API or the NSX Manager UI. For detailed steps on creating the service profile by using the NSX Manager UI, see the *NSX-T Data Center Administration Guide*.

   When you use the NSX Manager UI to create a service profile, the service reference is internally created, if it is not already present.

   If you decide to use the NSX-T APIs to create a service profile, do the following steps:

   a Create a service reference.

   ```
   PATCH https://{policy-mgr-ip}/policy/api/v1/infra/service-references/
   {service-reference-id}
   ```

   b Use the *service-reference-id* from the previous step to create the service profile.

   ```
   PATCH https://{policy-mgr-ip}/policy/api/v1/infra/service-references/
   {service-reference-id}/service-profiles/{service-profile-id}
   ```

   For a detailed information about these APIs, see the *NSX-T Data Center API Guide*.

Prepare the vRealize Automation Environment

Perform the following tasks in vRealize Automation before starting the migration.

- Add a new target NSX-T cloud account in vRealize Automation for the source NSX-V cloud account in your vRealize Automation projects.

- Create a migration plan to migrate from NSX-V to NSX-T. Complete the following four steps in the migration plan:
  - Step 1: Specify the source and target cloud accounts.
  - Step 2: Assess the NSX cloud accounts for migration readiness.
  - Step 3: Enter maintenance mode for the cloud accounts.
  - Step 4: Generate a deployment configuration file. You will upload this file when you do the migration.

  For more information, see the vRealize documentation about migrating NSX-V to NSX-T.

Import Configuration

This step imports the configuration of the NSX-V environment.

Procedure

1 From a browser, log in to NSX Manager as **admin**.

2 Navigate to **System > Migrate**.

3 In the **vRealize Automation** pane, click **Get Started** and select **User Defined Topology**.

4 Under **Authentication**, provide the required credentials.

5 In the **Upload Deployment Configuration File** section, click **Select File**. Select the `.json` configuration file that was generated from the vRealize Automation environment, and click **Upload**.

   After uploading a file, if necessary, you can click **Select File** again and upload a different file. The previous file is overwritten. You cannot remove the `.json` file after it is uploaded. The configuration file is removed only if you take any of the following actions:

   - Roll back the migration on the **Import Configuration** page or the **Migrate Configuration** page.
   - Finish the migration on the **Migrate Hosts** page.

6 Click **Start** to start the import.

7 In the confirmation dialog, click **Import** and wait for the process to complete.

8 If the status is **Successful**, click **Continue** to go to the next step.

9 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Results

If the NSX-V topology is imported successfully, you can click the **View Imported Topology** link to view the imported topology. However, the topology viewer might not work properly for a large-scale NSX-V environment.

Translate Configuration Layer 2

This step translates the layer-2 NSX-V configuration that was imported.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Translate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Resolve Configuration Layer 2

This step resolves layer-2 issues in the configuration that was imported.

For more information about resolving issues, see Resolve Configuration Issues.

Procedure

1 Click the message for each issue to see the details. You can click **Accept** to accept the recommendation. You can also select all issues and click **Accept Recommendations** to accept the recommendations for all the issues.

2 Click **Submit** to confirm that you want to proceed with the resolution of the issues.

3 Click **Continue** to go to the next step.

Migrate Configuration Layer 2

This step migrates the layer-2 configuration.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Check Realization Layer 2

This step checks that the configuration that was migrated is realized in NSX-T.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Define a Topology

In this step, you specify the topology that will be migrated. This can be done through the NSX Manager UI or by using a mapping file in JSON format.

This mapping specifies how Edge Service Gateways (ESGs) and Distributed Logical Routers (DLRs) should map to gateways in NSX-T. Before providing the mapping, evaluate your topology requirements and plan to do the following:

1 Determine how you will map the ESGs and DLRs. The northbound ESGs without any L4-L7 services should be skipped. These are usually the ESGs peering with northbound routers and are in ECMP path. If you are using VPN on a northbound ESG, migrating to active-standby tier-0 is recommended. In other cases, migrating the ESGs/DLRs to tier-1 is recommended. An ESG and a DLR can be merged in one mapping entry.

2 Create tier-0 and tier-1 gateways and configure dynamic or static routing on the tier-0 gateways towards the northbound routers based on your requirements.

3 After configuring the dynamic routing on tier-0 gateways, check that the dynamic routing has converged, that is, BGP sessions are established or OSPF neighborships are FULL as applicable. After this, proceed with providing the mapping.

An example of a mapping file that maps ESGs to a tier-0 gateway:

```
[
  {
    "name":"nsxv-to-nsxt-mapping",
    "v_edges_to_policy_gateways_mappings":[
      {
        "v_edges":[
          "edge-1",
          "edge-2"
         ],
         "policy_gateway_name": "tier0-gateway"
         "policy_gateway_path": "/infra/tier-0s/tier0-gateway"
      }
    ]
  }
]
```

Procedure

1 Choose one of the following options:

- Do not migrate L3-L7 entities and services.

- Upload a mapping file (*.json).
- Select a Tier-0 or Tier-1 gateway for each entity that needs to be migrated.

2 After you have specified a mapping, click **Continue** to go to the next step.

Translate Configuration Layer 3 and Above

This step translates the layer-3 and L4-L7 services of the NSX-V topology that you defined.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Translate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Resolve Configuration Layer 3 and Above

This step resolves issues in the translation of the topology that you specified.

Issues for the following categories of objects will be displayed:

- NS Service
- Edge
- Other
- Profile
- RBAC

Note that for some features such as VPN, you might get the message "Feature ABC on Edge-XYZ cannot be migrated" even though the feature is not configured on NSX-V. You can simply accept the input and proceed with the migration.

Procedure

1 For each category, click the message for each issue to see the details. You can click **Accept** to accept the recommendation. You can also select all issues and click **Accept Recommendations** to accept the recommendations for all the issues.

2 Click **Submit** to confirm that you want to proceed with the resolution of the issues.

3 Click **Continue** to go to the next step.

Migrate Configuration Layer 3 and Above

This step migrates the topology that you specified.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Check Realization Layer 3 and Above

This step checks that the topology that was migrated is realized in NSX-T.

Procedure

1 Click **Start**.

2 In the confirmation dialog, click **Migrate** and wait for the process to complete.

3 If the status is **Successful**, click **Continue** to go to the next step.

4 If the status is **Failed**, the details of the failure are displayed. Click **Rollback** to resolve the issues.

Migrate NSX-V Edges

In this step, you migrate the Edge Services Gateways (ESGs).

If you have no Edge Services Gateway appliances in your topology, you must still click **Start** so that you can proceed to the **Migrate Hosts** step.

Caution If you roll back the **Migrate Edges** step, verify that the traffic is going back through the NSX-V Edge Services Gateways. You might need to take manual action to assist the rollback.

Prerequisites

- All configuration issues must be resolved.
- The NSX-V configuration must be migrated to NSX-T.
- Verify that the migrated configurations are shown in the NSX Manager UI or API of NSX-T.
- Verify that you have a backup of NSX-V and vSphere since the most recent configuration changes were made.
- If you are using new IP addresses for the NSX-T Edge node uplinks, you must configure the northbound routers with these new BGP neighbor IP addresses.
- Verify that you have created an IP pool for Edge Tunnel End Points (TEP). See Create an IP Pool for Edge Tunnel End Points.

- Logical router interfaces created in NSX-T use the global default MTU setting, which is 1500. If you want to ensure that all logical router interfaces have a larger MTU, you can change the global default MTU setting. For more information, see Change the Global MTU Setting.

  If MTU setting other than 1500 is used on peering routers, the same should be configured on NSX-T. In case of OSPF topologies, OSPF adjacencies can get stuck if MTU setting is different from peering routers' MTU setting.

Procedure

1 From the **Migrate Edges** page, click **Start**.

  All Edges are migrated. The uplinks on the NSX-V Edge Services Gateways are internally disconnected, and the uplinks on the NSX-T Edge nodes are brought online.

2 Verify that routing and services are working correctly in the new NSX-T environment.

  If so, you can migrate the hosts. Before migrating the hosts, see Configuring NSX-V Host Migration.

3 Click **Continue** to go to the next step.

Results

The following changes result from the migration process:

- The routing and service configuration from NSX-V Edge Services Gateway (ESG) are transferred to the newly created NSX-T Edge nodes.
- The new TEP IP addresses for the newly created NSX-T Edge nodes are configured from a newly created IP pool for Edge Tunnel End Points.

Migrating Hosts

Before migrating the hosts, select a host migration plan.

For more information about what happens during host migration, see Changes Made During Host Migration in an End-to-End Migration.

If the Security Policies in your NSX-V environment use a partner service for Guest Introspection or Network Introspection or both, choose the host migration mode as shown in the following table.

Partner Service	Host Migration Mode
Only Guest Introspection	In-Place and Maintenance modes are supported.
Only Network Introspection	Maintenance mode is supported. However, Automated Maintenance mode is recommended. In-Place mode is not supported.
Both Guest Introspection and Network Introspection	Maintenance mode is supported. In-Place mode is not supported.

Important Consult the VMware partner before migrating their service that is running on the NSX-V workloads. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their services to NSX-T.

Caution Host migration should be completed during the same maintenance window as Edge migration.

You must disable IPFIX and reboot the ESXi hosts before migrating them.

If the partner service in your NSX-V environment provides Guest Introspection or both Guest Introspection and Network Introspection service, follow the procedure in Migrate Hosts with Guest Introspection Service.

If the partner service in your NSX-V environment provides only Network Introspection service, follow the procedure in Migrate Hosts with Network Introspection Service.

Select a Host Migration Plan

The clusters in the NSX-V environment are displayed on the **Migrate Hosts** page. The clusters are arranged into migration groups, each migration group contains one vSphere host cluster. There are several settings which control how the host migration is performed.

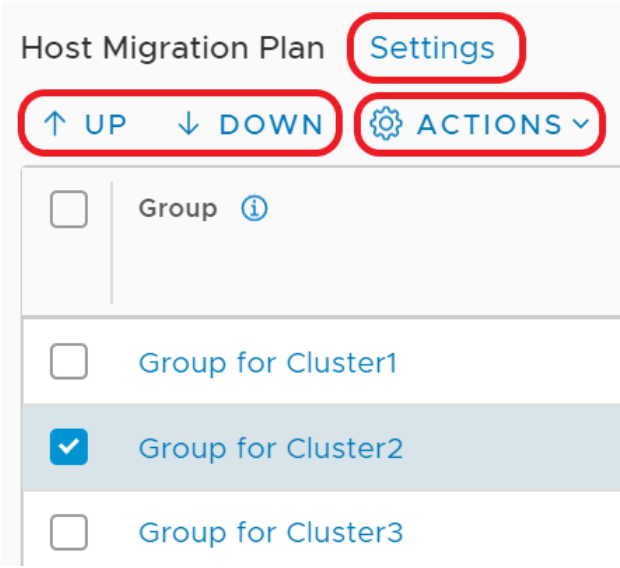

- Click **Settings** to change the global settings: **Pause Between Groups** and **Migration Order Across Groups**.
- Select a single host group (cluster) and use the arrows to move it up or down in the migration sequence.
- Select one or more host groups (clusters) and click **Actions** to change these host groups settings: **Migration Order Within Groups**, **Migration State**, and **Migration Mode**.

Pause Between Groups

When migrating multiple host groups, you can pause the migration between groups by enabling the **Pause Between Groups** setting. After a group is migrated, you must click **Continue** to migrate the next host group. This setting is disabled by default. You can enable it if you want to verify the status of the applications running on each cluster before proceeding to the next one.

Serial or Parallel Migration Order

You can specify whether migration happens in a serial or parallel order. There are two ordering settings:

- **Migration Order Across Groups** is a global setting that applies to all host groups.
  - **Serial**: One host group (cluster) at a time is migrated.
  - **Parallel**: Up to five host groups at a time are migrated. After those five host groups are migrated, the next batch of up to five host groups are migrated.

    Important If you are migrating from NSX-V 6.4.4, 6.4.5, or 6.4.6, and your environment uses vSphere Distributed Switch 7.0 or later, do not select parallel migration order across groups.

    If you are migrating from NSX-V 6.4.8 or later, and your environment uses vSphere Distributed Switch 7.0 or later, parallel migration order across groups is supported.

- **Migration Order Within Groups** is a host group (cluster) specific setting, so can be configured separately on each host group.
  - **Serial**: One host within the host group (cluster) at a time is migrated.
  - **Parallel**: Up to five hosts within the host group are migrated at a time. After those hosts are migrated, the next batch of up to five hosts are migrated.

    Important Do not select parallel migration order within groups for a cluster if you plan to use **Maintenance** migration mode for that cluster.

By default, both settings are set to **Serial**. Together, the settings determine how many hosts are migrated at a time.

Table 11-13. Effects of Migration Settings on Number of Hosts Attempting Migration Simultaneously

Migration Order Across Groups (Clusters)	Migration Order Within Groups (Clusters)	Maximum Number of Hosts Attempting Migration Simultaneously
Serial	Serial	1 One host from one host group
Serial	Parallel	5 Five hosts from one host group
Parallel	Serial	5 One host from five host groups
Parallel	Parallel	25 Five hosts from five host groups

Important If there is a failure to migrate a host, the migration process will pause after all in-progress host migrations have finished. If **Parallel** is selected for both migration across groups and migration within groups, there might be a long outage for the failed host before you can retry migration.

Sequence of Migration Groups

You can select a host group (cluster) and use the arrows to move it up or down in the list of groups.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

Migration State

Host groups (clusters) can have one of two migration states:

- **Enabled**

  Hosts groups with a migration state of **Enabled** are migrated to NSX-T when you click **Start** on the **Migrate Hosts** page.

- **Disabled**

  You can temporarily exclude host groups from migration by setting the migration state for the groups to **Disabled**. Hosts in disabled groups are not migrated to NSX-T when you click **Start** on the **Migrate Hosts** page. However, you must enable and migrate all **Disabled** host groups before you can click **Finish**. Finish all host migration tasks and click **Finish** within the same maintenance window.

In the **Resolve Configuration** step, hosts that are ineligible for migration are identified. In the **Migrate Hosts** step, these hosts have the migration state **Do not migrate**. For example, hosts that do not have NSX-V installed are not eligible for migration.

Migration Mode

Migration Mode is a host group (cluster) specific setting, and can be configured separately on each host group. In the **Migrate Hosts** step, you select whether to use **In-Place** or **Maintenance** mode.

There are two types of Maintenance migration modes:

- Automated
- Manual

In the **Resolve Configuration** step of the migration process, you select which type of Maintenance migration mode to use. You select a Maintenance mode even if you plan to migrate hosts using **In-Place** mode. When you select Maintenance migration mode in the **Migrate Hosts** step, the value you specified in the **Resolve Configuration step** determines whether Automated Maintenance mode or Manual Maintenance mode is used. However, in the **Migrate Hosts** step, if you select **In-Place** mode, your selected choice of Maintenance mode in the **Resolve Configuration** step does not take effect.

In-Place migration mode is not supported if your NSX-V installation uses vSphere Distributed Switch 7.0 or later.

If your environment uses Distributed Firewall, select **Automated Maintenance** migration mode. If you select a different migration mode, the following limitations apply to environments with Distributed Firewall:

- If you use **Manual Maintenance** migration mode, all VMs must be moved to NSX-T hosts, connected to NSX-T segments, and powered on before the last NSX-V host starts migrating. When you migrate your last NSX-V host, do not power off the VMs on the host. Move them to an NSX-T host using vMotion.
- If you use **Manual Maintenance** migration mode, VMs have a gap in firewall protection for up to 5 minutes after they move to an NSX-T host.
- If you use **In-Place** migration mode, and you have Distributed Firewall rules that are applied to a VM, those rules are not pushed to the host until the host and all its VMs are migrated. Until the rules are pushed to the host, the following applies:
  - If the NSX-T default rule is `deny`, the VM is not accessible.

- If the NSX-T default rule is `accept`, the VM is not protected by the applied-to rules.

The migration process is different for each migration mode:

- **In-Place** migration mode

  NSX-T is installed and NSX components are migrated while VMs are running on the hosts. Hosts are not put in maintenance mode during migration. Virtual machines experience a short network outage and network storage I/O outage during the migration.

- **Automated Maintenance** migration mode

  A task of entering maintenance mode is automatically queued. VMs are moved to other hosts using vMotion. Depending on availability and capacity, VMs are migrated to NSX-V or NSX-T hosts. After the host is evacuated, the host enters maintenance mode, NSX-T is installed, and NSX components are migrated. VMs are migrated back to the newly configured NSX-T host. Note that VMs that are powered off will not be reconfigured. After migration, you need to manually configure these VMs before powering them on.

- **Manual Maintenance** migration mode

  A task of entering maintenance mode is automatically queued. To allow the host to enter maintenance mode, do one of the following tasks:

  - Power off all VMs on the hosts.
  - Move the VMs to another host using vMotion or cold migration.

  Once the host is in maintenance mode, NSX-T is installed on the host and NSX components are migrated. After the host is migrated, for the powered-off VMs and the VMs that you moved, you will need to change their network connection from the NSX-V logical switch to an NSX-T segment.

In the NSX-V environment, if the ESXi host's vmk0 management interface is connected to a VSS (vSphere Standard Switch) portgroup that does not have an uplink, and the portgroup is bridged to a VDS portgroup, and the VDS version is 6.5, 6.6 or 6.7, you must migrate using the **Maintenance** mode. If you use the **In-Place** mode, the migration will fail.

Adding or Removing a Host During Migration

Starting with NSX-T 3.2.1, during the host migration step, you can add or remove a host to be migrated when there is a pause in the migration.

Starting with NSX-T 3.2.2, you can pause the migration of hosts within a group. For more information, see Configuring NSX-V Host Migration.

Host migration will pause if you enable the setting **Pause Between Groups** or **Pause Between Hosts**, or if there is a failure to migrate a host.

You can add a host to a cluster before or after the cluster has been migrated, or while the cluster is being migrated. You can also remove a host from a cluster that has not migrated or is being migrated.

The host you want to add can be a standalone host or in a cluster. The host must not have NSX-V or NSX-T configured. If the host is in a cluster, the cluster must not have NSX-V or NSX-T configured.

You must prepare the host as a transport node first and later move it into a target cluster that has been migrated to NSX-T, is being migrated or has not started the migration. If the target cluster already has a host migrated to an NSX-T transport node, you can use the transport node's configuration as a reference to prepare the host as a transport node. If the host will support overlay traffic and a VTEP IP pool will be used to prepare the host as a transport node, the VTEP IP pool cannot be or overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated. It can be an existing IP pool (such as the one used by NSXT Edge nodes) or a newly created IP pool. For more information, see the section "Preparing ESXi and KVM Hosts as Transport Nodes" in the *NSX-T Data Center Installation Guide*.

When preparing an ESXi host as a transport node, you can choose N-VDS or VDS as the host switch. Choose VDS if the version of the VDS being migrated is 7.0 or later. Otherwise, choose N-VDS. If you prepare a host with N-VDS when you should choose VDS, the host will still be migrated but it may have network issues.

Adding a host to a cluster

1 From vCenter Server, put the host in Maintenance Mode.

2 If there is no VTEP IP pool that does not overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated, or such pools do not have enough free IPs for the VTEPs to be created for the host to be added, then go to **Networking > IP Address Pools** and create a new VTEP IP pool.

3 Follow the instructions in the installation guide to prepare the host as a transport node. When you select the transport zone, if an overlay transport zone is chosen for the host switch, choose the new IP pool that was created in step 2 or choose an existing IP pool that does not overlap with the VTEP IP pool created for the NSX-V hosts being migrated or to be migrated. Select an uplink profile. Do not choose the one whose name contains "VXLAN" if an overlay transport zone is chosen for the host switch.

4 Wait for the status of the host node to be "Success". Do not move the host out of Maintenance Mode.

5 Choose a cluster into which the host will be added. In NSX-T Manager UI, go to **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later) and check if the cluster has a Transport Node Profile (TNP) attached. Detach the TNP if it does.

6 In vCenter Server UI, move the host into the chosen cluster.

7 Invoke the sync host groups API or click the **Refresh** button on the NSX-T Manager UI host migration screen so that the migration group for the cluster contains the new host.

8 Call the following NSX-T API to accept the new host:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/setup?
action=migrate_newly_added_host_transport_node
{
    "host_transport_node_id" : "<transport-node-uuid>"
}
```

If the API returns an error, fix the error and retry the API. If the API returns success, then make the host exit maintenance mode.

9 vMotion VMs to the host. If any VM is moved from an NSX-V host, be sure to change the network to map the source virtual-wire to NSX-T overlay segments. For example, virtual-wire vxw-dvs-64-virtualwire-4-sid-10787-1-switch-191 must map to 1-switch-191-LS.

Removing a host from a cluster

1 From vCenter Server, migrate all VMs off the host and enter the host into maintenance mode.

2 If the host is in a cluster that has not started migrating, go to the next step. Otherwise, from NSX-T Manager UI, go to **System > Fabric > Nodes > Host Transport Nodes** (if the NSX-T version is 3.2.0 or 3.2.1) or **System > Fabric > Hosts** (if the NSX-T version is 3.2.2 or later). If the host has NSX configured, delete it. If the host is not reachable by NSX Manager, delete it with the force option.

3 From vCenter Server, remove the host. Wait until the task is complete.

4 Click the Refresh button on the host migration screen to remove the host from the migration group.

5 Restart the host migration.

Migrate Hosts with Guest Introspection Service

After you migrate the Edge Services Gateways successfully, you can migrate the NSX-V hosts to NSX-T host transport nodes.

Prerequisites

- Verify that Edge migration has finished and all routing and services are working correctly.
- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.

Procedure

1 On the **Migrate Hosts** page, click **Start**.

2 If you selected the **Manual Maintenance** migration mode for any host groups, you must complete one of the following tasks for each VM so that the hosts can enter maintenance mode.

Option	Action
Power off or suspend VMs.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b After the host has migrated, attach the VM interfaces to the appropriate NSX-T segments and power on the VM.
Move VMs using vMotion.	Right click the VM and select Migrate. Follow the prompts to move the VM to a different host. Note that Migration Coordinator maintains security during migration by vMotioning VMs to specific ports that are protected by temporary rules. In the case of manual vMotion, the VMs will not be moved to those ports and there could be a security breach. To vMotion manually, the VMs must be migrated using vSphere API where the networking backing must point to the OpaqueNetwork ID corresponding to the NSX Segment when using NVDS or the VDS portgroup ID when using VDS 7. In both cases, the network device's externalId must be set to the the string "VM_UUID:vNIC_ID", where VM_UUID is the VM's instance UUID and vNIC_ID is the VM's vNIC index where the first vNIC is 4000.
Move VMs using cold migration.	a Right click the VM and select **Power > Power off** , **Power > Shut Down Guest OS**, or **Power > Suspend**. b Right click the VM and select Migrate. Follow the prompts to move the VM to a different host, connecting the VM interfaces to the appropriate NSX-T segments.

Here is python code to specify an external-id for each vNIC in a VM and then vMotion the VM so that the vNICs will connect to an NSX-T segment of ID “ls_id” at the correct ports:

```
devices = vmObject.config.hardware.device
nic_devices = [device for device in devices if isinstance(device,
vim.Vm.device.VirtualEthernetCard)]
vnic_changes = []
for device in nic_devices:
    vif_id = vmObject.config.instanceUuid + ":" + str(device.key)
    vnic_spec = self._get_nsxt_vnic_spec(device, ls_id, vif_id)
    vnic_changes.append(vnic_spec)
relocate_spec = vim.Vm.RelocateSpec()
relocate_spec.SetDeviceChange(vnic_changes)
# set other fields in the relocate_spec
vmotion_task = vmObject.Relocate(relocate_spec)
WaitForTask(vmotion_task)

def _get_nsxt_vnic_spec(self, device, ls_id, vif_id):
    nsxt_backing = vim.Vm.Device.VirtualEthernetCard.OpaqueNetworkBackingInfo()
    nsxt_backing.SetOpaqueNetworkId(ls_id)
    nsxt_backing.SetOpaqueNetworkType('nsx.LogicalSwitch')
    device.SetBacking(nsxt_backing)
    device.SetExternalId(vif_id)
```

```
    dev_spec = vim.Vm.Device.VirtualDeviceSpec()
    dev_spec.SetOperation(vim.Vm.Device.VirtualDeviceSpec.Operation.edit)
    dev_spec.SetDevice(device)
    return dev_spec
```

For an example of a complete script, see https://github.com/dixonly/samples/blob/main/vmotion.py

The host enters maintenance mode after all VMs are moved, powered off, or suspended. If you want to use cold migration to move the VMs to a different host before the migrating host enters maintenance mode, you must leave at least one VM running while you move VMs. When the last VM is powered off or suspended, the host enters maintenance mode, and migration of the host to NSX-T starts.

Results

After a host has migrated to NSX-T using **In-Place** migration mode, you might see a critical alarm with message `Network connectivity lost`. This alarm occurs when a vSphere Distributed Switch (VDS) 6.5 or 6.7 migrates to an N-VDS because the host no longer has a physical NIC connected to the VDS it was previously connected to. To restore the migrated hosts to the Connected state, click **Reset to Green** on each host, and suppress the warnings, if any.

If migration fails for a host, you can move its host group to the bottom of the list of groups. The migration of other host groups can proceed while you resolve the problem with the failed host.

If migration fails for a host, the migration pauses after all in-progress host migrations finish. When you have resolved the problem with the host, click **Retry** to retry migration of the failed host. If the host still fails to migrate, you can configure NSX-T on the host manually or remove the host from the system. In this case, at the end of the host migration step, the **Finish** button will not be enabled because of the host that failed to migrate. You need to call the REST API `POST https://<nsx-mgr-IP>/api/v1/migration?action=finalize_infra` (<nsx-mgr-IP> is the IP address of the NSX Manager where the migration service is running) using a REST API client (for example, postman or curl) to finish the migration, and then perform the post-migration tasks.

For information about troubleshooting other host migration problems, see Chapter 13 Troubleshooting Migration Issues.

What to do next

If the migrated Security Policies use a third-party partner service, deploy an instance of the partner service in NSX-T. For detailed instructions, see:

- Deploy a Partner Service for Endpoint Protection

  Click this link to deploy a partner service that provides only Endpoint Protection or both Endpoint Protection and Network Introspection services to the NSX-T workload VMs.

- Deploy a Partner Service for Network Introspection

  Click this link to deploy a partner service that provides only Network Introspection service to the NSX-T workload VMs.

Migrate Hosts with Network Introspection Service

When Security Policies in your NSX-V environment use only a Network Introspection service that is provided by a partner, two approaches are available to migrate the NSX-V prepared hosts to NSX-T.

Both the approaches discussed in this topic assume that the partner service virtual machines (SVMs) in your NSX-V environment are not deleted before starting the migration coordinator. Depending on how much security protection downtime you are willing to accept during host migration, choose the host migration approach that best suits your needs.

Note Consult the VMware partner before migrating the hosts by using any of the two approaches. Check with the partner whether their service is supported for migration to NSX-T and seek their inputs before the migration. Partners will have their own guidance to migrate their service to NSX-T.

When only Network Introspection service is running on your NSX-V hosts, **In-Place** host migration mode is not supported. Only **Maintenance** migration mode is supported. However, Automated Maintenance migration mode is recommended.

Approach 1: Involves more security protection downtime

This approach is the simpler of the two host migration approaches. However, it involves more security protection downtime compared to Approach 2. Let us say that you have three clusters in your NSX-V environment: Cluster 1, Cluster 2, and Cluster 3.

In this approach, enable the **Pause between groups** migration setting and migrate Cluster 1 by using the standard host migration procedure that is explained in Migrate NSX-V Hosts. After Cluster 1 is migrated to NSX-T, the migration pauses. Deploy the partner service in Cluster 1 by doing either a host-based or a clustered service deployment. Now, disable the **Pause between groups** migration setting, and continue migrating Clusters 2 and 3. After the workload VMs in Clusters 2 and 3 are migrated to NSX-T, these workloads can start redirecting packets to the partner service virtual machines (SVM) in Cluster 1.

In this approach, security protection downtime is expected during migration of Cluster 1.

When workload VMs migrate to an NSX-T host, existing data traffic during a host migration is expected to have a security protection downtime. However, new data traffic does not have a security protection downtime.

Approach 2: Involves minimal security protection downtime

This approach requires some manual intervention with an NSX-T API to create a temporary host group. Enable the **Pause between groups** migration setting and migrate any one host from Cluster 1. After this host in Cluster 1 is migrated to NSX-T, the migration coordinator pauses. Deploy a partner service on this migrated host by doing a clustered service deployment. Continue migrating the remaining hosts in Cluster 1. After all the hosts in Cluster 1 are migrated to NSX-T, you can optionally deploy additional partner SVMs in Cluster 1 by doing either a host-based or a clustered service deployment.

The detailed procedure in this topic explains the host migration workflow for a single NSX-V prepared cluster, which has three hosts, as shown in the following figure. The procedure uses Approach 2 to migrate this Cluster 1 to NSX-T.

Example:

Figure 11-1. Host Group 1 (Cluster 1) Before Migration

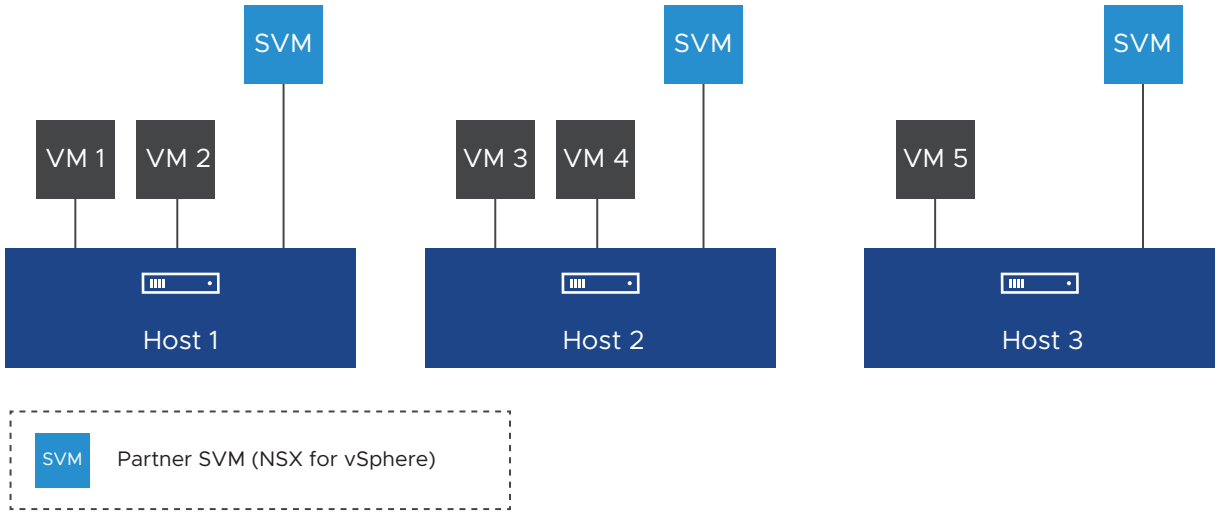

All hosts in this cluster are ESXi hosts. The Security Policies in your NSX-V environment redirect data traffic to partner service virtual appliances that provide a network introspection service to workloads. As NSX-V supports only a host-based service deployment, each host has a single partner service VM.

The following configuration settings are required for migrating hosts using Approach 2:

- Host migration mode is set to Automated Maintenance.
- Pause between groups is enabled.
- Migration order across groups is set to serial.

Prerequisites

- Verify that Edge migration has finished and all routing and services are working correctly.

- In the vCenter Server UI, go to the **Hosts and Clusters** page, and verify that all ESXi hosts are in an operational state. Address any problems with hosts including disconnected states. There must be no pending reboots or pending tasks for entering and exiting maintenance mode.
- Enable vSphere DRS on the cluster that is being migrated.
- Enable vMotion on the VMkernel adapter of each host in the cluster.
- Ensure that adequate spare capacity is available in the NSX-V cluster so that the migrating hosts can enter into a maintenance mode. If enough spare capacity is unavailable to migrate NSX-V workload VMs to other hosts in the cluster, additional security protection downtime is expected.

Procedure

1 Run the following API request to create a temporary host group and move hosts 2 and 3 to this temporary group.

```
POST https://{nsxt-mgr-ip}/api/v1/migration/migration-unit-groups
```

In the request body of this POST API, specify the following details:

- Name of the temporary host group
- Migration units (IDs of hosts 2 and 3)
- Migration state of the temporary group (must be disabled)

For a detailed information about this API and an example POST API request, see the *NSX-T Data Center API Guide*.

You can obtain the host IDs from the vCenter Server Managed Object Browser (MOB) at `http://{vCenter-IP-Address}/mob`, or run the following GET API to retrieve the host IDs:

```
GET https://{nsxt-mgr-ip}/api/v1/fabric/discovered-nodes
```

A temporary host group is created and displayed on the **Migrate Hosts** page. The original host group 1 (cluster 1) now contains only host 1.

2 On the **Migrate Hosts** page, next to **Host Migration Plan**, click **Settings** and ensure that the settings are configured as follows:

- Pause between groups: Enabled
- Migration order across groups: Serial

3 Migrate host 1 to NSX-T.

a Click **Start** to start the host migration.

Workload VMs 1 and 2 are migrated to other hosts so that host 1 can enter into a maintenance mode. NSX-V partner SVM on host 1 is powered off before host 1 enters into a maintenance mode.

Assume that VMs 1 and 2 are migrated to host 3 that is prepared with NSX-V. After the migration of host 1 is successful, the migration coordinator pauses for your next input.

b (Required) Deploy a partner service on host 1 by using the clustered deployment approach.

At this stage, host-based service deployment is not supported. Deploying a partner service on host 1 is necessary to minimize security protection downtime. Remember, the security protection for NSX-V workloads that are running on hosts 2 and 3 is still intact. The partner must ensure that the migrated partner-specific Security Policies are available on the newly deployed partner SVMs on host 1.

For detailed steps on deploying a partner service in NSX-T, see Deploy a Partner Service for Network Introspection. For example, specify the following configuration settings to deploy two partner service virtual machines (SVMs) on host 1:

Configuration	Value
Deployment Type	Clustered
Host	Host 1
Clustered Deployment Count	2

The value that you enter in the **Clustered Deployment Count** text box depends on the resource capacity that is available on the host. This scenario assumes that two partner SVMs can be deployed on Host 1. This value can be different in your environment.

After this step, the cluster looks as shown in the following figure. The green colored host represents the migrated host.

Figure 11-2. Host 1 is Migrated to NSX-T

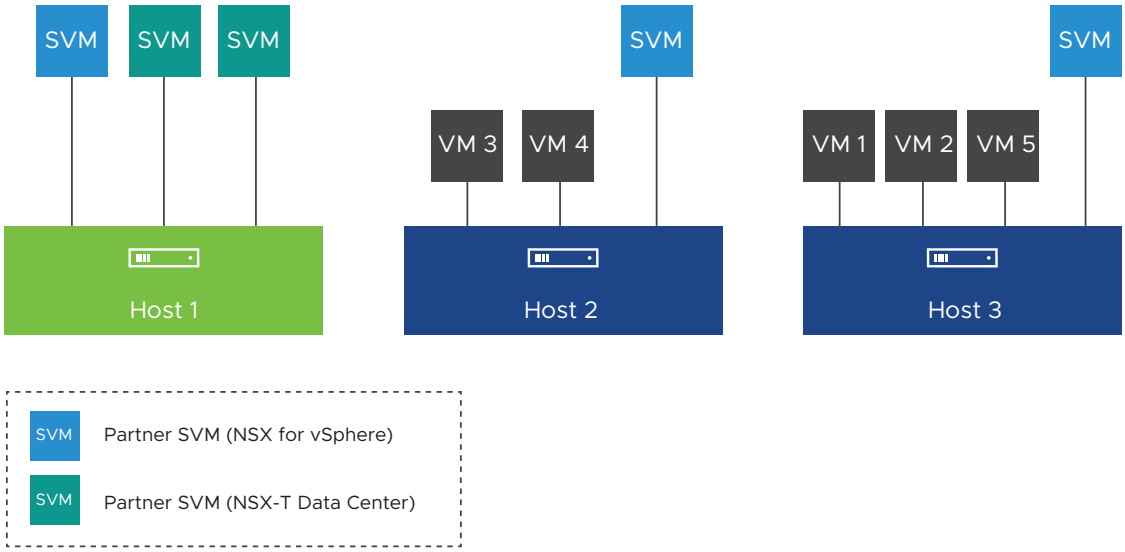

4 Migrate host 2 and host 3 to NSX-T.

a Move host 2 and host 3 from the temporary host group to the original host group 1 by running the following POST API request:

```
POST https://{nsxt-mgr-ip}/api/v1/migration/migration-unit-groups/
{group-id}?action=add_migration_units
```

Where: *group-id* is the ID of the destination host group (host group 1). In the POST API request body, specify the ID of hosts 2 and 3 that you want to add to the original host group 1.

For a detailed information about this POST API and an example POST API request, see the *NSX-T Data Center API Guide*.

Now, the original host group 1 contains hosts 1, 2 and 3 (in the given order), and the temporary host group is deleted.

b Select the check box next to the original host group 1, and then click **Actions > Change Migration Order Within Group**. Verify that the migration order within the group is set to **Serial**.

If necessary, you can set the migration order in the original host group 1 to **Parallel**.

c Click **Continue** to resume the host migration.

Host 2 is first migrated to NSX-T, and then host 3 is migrated. To put each migrating host into a maintenance mode, workload VMs on the migrating host are moved to either NSX-V hosts or NSX-T hosts. NSX-V partner SVM on the migrating host is also powered off before the host enters into a maintenance mode.

After this step, all the hosts in host group 1 (cluster 1) are prepared with NSX-T. The cluster looks as shown in the following figure.

Figure 11-3. All Hosts are Migrated to NSX-T

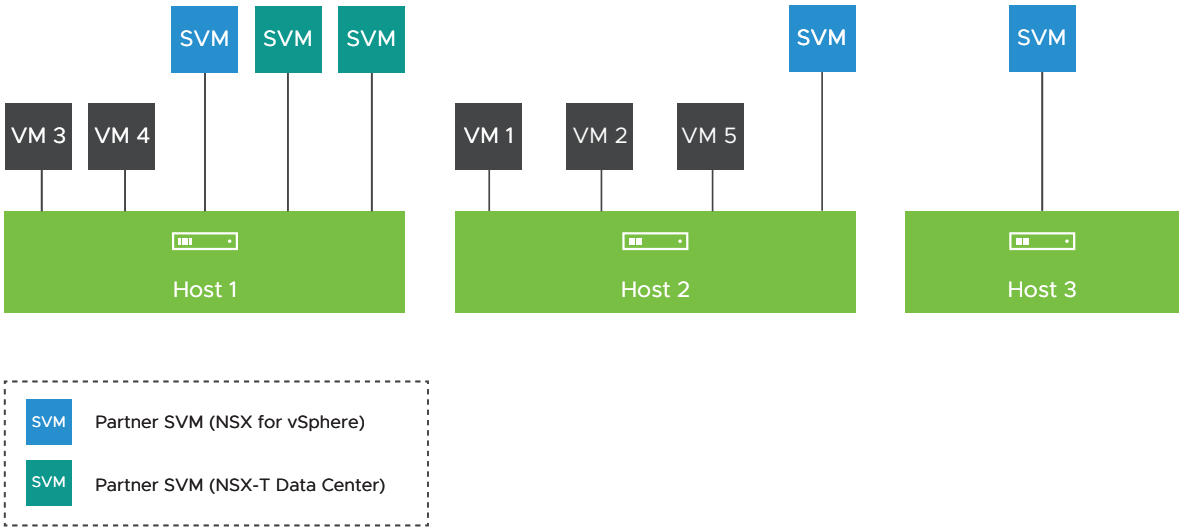

5 (Optional) Migrate some workload VMs from hosts 1 and 2 to host 3.

For example, migrate VMs 4 and 5 to host 3, as shown in the following figure.

Figure 11-4. Final Cluster 1 After Migration

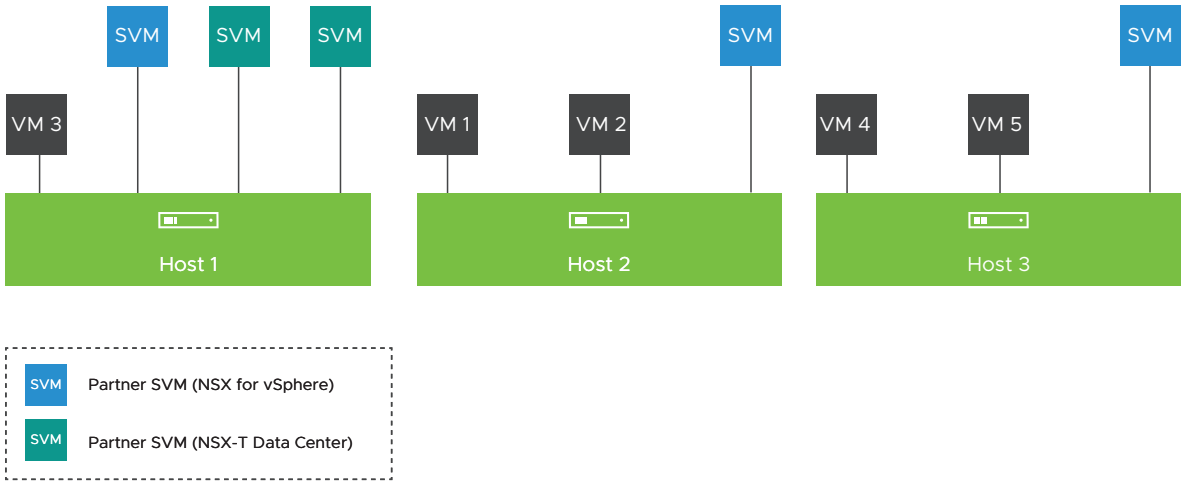

6 (Optional) After all the hosts in host group 1 are migrated to NSX-T, you can do either a host-based or a clustered service deployment.

A host-based service deployment allows new network traffic to be protected by a local partner SVM on each host.

Note If you have network introspection service running on more than one NSX-V prepared cluster, you do not have to deploy the partner SVMs in the other clusters. The network traffic though the NSX-T workload VMs in the other clusters can use the partner SVMs in cluster 1 that you just migrated. The host migration workflow covered in this procedure is required only for the first cluster. You can migrate the remaining clusters by using the standard host migration procedure.

What to do next

Delete the partner service deployment in NSX-V. Remember, you can delete the partner SVMs only at a cluster level. That is, you can delete service deployment only after all the hosts in the host group 1 are migrated to NSX-T. Complete the following steps to delete the service deployment in NSX-V:

1 Log in to the vSphere Client and navigate to **Networking and Security > Installation and Upgrade > Service Deployment**.

2 Select the deployed partner service, and click **Delete**.

Finish the End-to-End Migration

After you have migrated all Edge Services Gateway VMs and hosts to the NSX-T environment, confirm that the new environment is working correctly. If everything is functioning correctly, you can finish the migration.

Important Verify that everything is working before clicking **Finish**. Then perform the post-migration tasks. Do not make any vSphere life cycle operations such as upgrading ESXi hosts, VDS, or VC before the post-migration tasks are completed.

You will see errors on hosts after the migration. The error message is: `UserVars.RmqHostId' is invalid or exceeds the maximum number of characters permitted`. The error occurs because this host is still part of the NSX-V inventory.

Prerequisites

- Verify that all expected items have been migrated to the NSX-T environment.
- Verify that the NSX-T environment is working correctly.

Procedure

1 Navigate to the **Migrate Hosts** page of the migration coordinator.

2 Click **Finish**

A dialog box appears to confirm finishing the migration. If you finish the migration, all migration details are cleared. You can no longer review the settings of this migration. For example, which inputs were made on the **Resolve Configuration** page, or which hosts were excluded from the migration.

Post-Migration Tasks

After migration has finished, some additional actions are required.

- If you migrated from NSX-V 6.4.4, perform a reboot of all hosts that have migrated to NSX-T. The reboot must be done before you upgrade to a later version of NSX-T.
- During migration, all transport nodes are added to a group called `NSGroup with TransportNode for CPU Mem Threshold`. This group ensures that the transport nodes have the correct CPU memory threshold settings in NSX-T. This group is required after migration has completed. If you need to remove a transport node from NSX-T after migration and you are running NSX-T 3.2.0, you must first remove the transport node from this group. If you are running NSX-T 3.2.1 or later, you do not need to remove the transport node from this group.

  To remove the transport node from the group, make sure you are in **Manager** mode and then select **Inventory > Groups** to remove the transport node from the `NSGroup with TransportNode for CPU Mem Threshold` group. For more information about Manager mode, see the topic "NSX Manager" in the *NSX-T Data Center Administration Guide*.
- Verify that you have a valid backup and restore configuration. See "Backing Up and Restoring the NSX Manager" in the *NSX-T Data Center Administration Guide*.

Import Final Output Mapping File into vRealize Automation

Share the final output mapping file with the vRealize Automation administrator. The vRealize Automation administrator must import this output mapping file in step 4 of the vRealize Automation NSX-V to NSX-T migration plan.

The vRealize Automation administrator must complete the following steps in the migration plan:

- Step 5: Migrate the NSX-V cloud account and its related objects to NSX-T.
- Step 6: Test the migration results.
- Step 7: Remove cloud accounts from maintenance mode and exit the migration plan.

For more information, see the vRealize documentation about migrating NSX-V to NSX-T.

Uninstall NSX-V After Migration

After the migration is verified and successful, uninstall your NSX-V environment. The process for uninstalling NSX-V after migration to NSX-T is different from the standard uninstall for NSX-V.

For more information, see Uninstalling NSX-V After Migration.

Finish Deploying the NSX Manager Cluster

You can run the migration coordinator tool with only one NSX Manager appliance deployed. Deploy two additional NSX Manager appliances before you use your NSX-T environment in production.

See the *NSX-T Data Center Installation Guide* for the following information:

- NSX Manager Cluster Requirements
- Deploy NSX Manager Nodes to Form a Cluster from UI
- Configure a Virtual IP (VIP) Address for a Cluster

Deploy a Partner Service for Endpoint Protection

When migrated Security Policies in NSX-T use a partner service that provides only Endpoint Protection or both Endpoint Protection and Network Introspection, deploy an instance of the partner service after all the clusters are migrated to NSX-T.

Only a host-based service deployment is supported.

In a host-based service deployment, one partner service virtual machine is installed on each host of the migrated cluster. In the vCenter Server, the vSphere ESX Agency Manager (EAM) service is internally used to deploy a partner service VM on each host of the cluster.

Prerequisites

- All the hosts in the cluster are migrated to NSX-T.
- All the migrated hosts are managed by a vCenter Server.
- A transport node profile is applied to the cluster.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Navigate to **System > Service Deployments > Deployment**.
3. In the **Partner Service** drop-down menu, select the partner service to be deployed, and click **Deploy Service**.
4. Enter the service deployment name.
5. Select the vCenter Server that is registered as a compute manager in NSX-T.
6. Select the cluster where you want to deploy the partner service.
7. To specify the datastore, do one of the following actions:
   - Select a datastore as the repository for the service virtual machines.
   - Select **Specified on Host**.

The **Specified on Host** option means that you do not need to select a datastore and network on the **Deploy Service** page. Before deploying the partner service, you must configure Agent VM settings on each ESXi host to point to a specific datastore and network.

To know more about configuring Agent VM settings, see the vSphere product documentation.

8 Under Networks, click **Set** and select the NICs you want to use for deployment.

a Select the network for the Management interface.

In a host-based deployment, if you set the datastore as **Specified on Host**, you must set the network also as **Specified on Host**.

b Set the Network type to **DHCP** or **Static IP Pool**. If you set the network type to a Static IP Pool, select from the list of available IP pools.

9 In the **Deployment Template** drop-down menu, select the registered deployment template and click **Save**.

The deployment process might take some time depending on the vendor's implementation.

10 Check the deployment status on the **Deployment** page. Wait until the status changes to Up.

You might have to refresh the **Deployment** page a few times to retrieve the latest status.

If the Status column shows Down, click the icon next to Down. All deployment errors are displayed. Take the required actions to fix the errors, and click **Resolve**. The status changes to In Progress. Wait until the status changes to Up.

Results

A partner service VM is now deployed on all the hosts of the cluster.

Note When you add a new host in the cluster, EAM automatically deploys the partner service VM on the new host.

What to do next

Go to the Partner Console and verify whether the endpoint protection service is activated. Now, the migrated endpoint protection rules are enforced on the workload VMs that are running on the NSX-T prepared cluster.

For more information about activating the endpoint protection service in the Partner Console, see the partner documentation.

Deploy a Partner Service for Network Introspection

When migrated Security Policies in NSX-T use a third-party partner service only for Network Introspection, deploy an instance of the partner service either by using a clustered service deployment or a host-based service deployment approach.

Prerequisites

For a clustered service deployment approach:

- At least one host in the first cluster is migrated to NSX-T.

For a host-based service deployment approach:

- All the hosts in a cluster are migrated to NSX-T.
- A transport node profile is applied to the cluster.

Procedure

1. From your browser, log in with **admin** privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Navigate to **System > Service Deployments > Deployment**.
3. In the **Partner Service** drop-down menu, select the partner service to be deployed, and click **Deploy Service**.
4. Enter the service deployment name.
5. Select the vCenter Server that is registered as a compute manager in NSX-T.
6. Select a deployment type: **Host-Based** or **Clustered**.
7. Select the cluster where you want to deploy the partner service.
8. (Clustered deployment only): In the **Host** drop-down menu, select a host, or select **Any** to allow the NSX-T NSX Manager to select a host.
9. In the **Data Store** drop-down menu, select a data store as the repository for the partner service virtual machine (SVM).
   - Clustered deployment: If you selected **Any** for the host, select a shared data store. If you specified a particular host, select a local data store.
   - Host-based deployment: Select a specific datastore or select **Specified on Host**. The **Specified on Host** option means that you do not need to select a datastore and network on the **Deploy Service** page. Before deploying the partner service, you must configure Agent VM settings on each ESXi host to point to a specific datastore and network.

     To know more about configuring Agent VM settings, see the vSphere product documentation.
10. Under Networks, click **Set** and select the NICs you want to use for deployment.
    a. Select the network for the Management interface.

       In a host-based deployment, if you set the datastore as **Specified on Host**, you must set the network also as **Specified on Host**.

    b. Set the Network type to **DHCP** or **Static IP Pool**. If you set the network type to a Static IP Pool, select from the list of available IP pools.

11 In the **Deployment Template** drop-down menu, select the required template.

   Typically, the deployment specification and the deployment template fields are automatically selected with the information that is pushed from the Partner Console as part of the service definition.

12 In the **Service Segment** drop-down menu, select the service segment that the migration coordinator has created in the overlay transport zone.

13 (Clustered deployment only): In the **Clustered Deployment Count** text box, specify the number of service VMs to deploy in the cluster, and click **Save**.

14 Check the deployment status on the **Deployment** page. Wait until the status changes to Up.

   You might have to refresh the **Deployment** page a few times to retrieve the latest status.

   If the Status column shows Down, click the icon next to Down. All deployment errors are displayed. Take the required actions to fix the errors, and click **Resolve**. The status changes to In Progress. Wait until the status changes to Up.

Uninstalling NSX-V After Migration

When you have verified that the migration is successful, and have clicked **Finish** to finish the migration, you can uninstall your NSX-V environment.

The process for uninstalling NSX-V after migration to NSX-T is different from the standard uninstall for NSX-V.

Important If you have vCenter Enhanced Linked Mode (ELM) configured, you must migrate all the NSX-V instances associated with the vCenter ELM chain before executing steps 6, 7, and 8 in the procedure below.

Prerequisites

- Verify that the migration is successful, and all functionality is working in the NSX-T environment.
- Verify that you have clicked **Finish** on the **Migrate Hosts** page.

Procedure

1 In the vSphere client, navigate to **Networking and Security > NSX Edges** and delete all the NSX Edges.

2 In the vSphere client, navigate to **Networking and Security > Logical Switches** and delete all the logical switches.

3 In the vSphere client, navigate to **Networking and Security > Installation and Upgrade > Logical Network Settings > Transport Zones** and delete all the transport zones.

4 In the vSphere client, navigate to **Networking and Security > Installation and Upgrade > Management > NSX Controller Nodes** and delete all the NSX Controllers.

5 Clear all stale VTEPs that may remain in the NSX-V Manager database:

   a SSH into NSX-V Manager as **root**.

   b Run the following command to clear the database table:

   ```
   psql -U secureall -d secureall -c "delete from xvs_vmknic_info;"
   ```

   c Run the following command to confirm that the output shows zero row:

   ```
   psql -U secureall -d secureall -c "select * from xvs_vmknic_info;"
   ```

6 Delete the ESX Agent Manager agencies that are associated with the NSX-V environment.

   a In the vSphere Client, navigate to **Menu > Administration**. Under **Solutions**, click **vCenter Server Extensions**. Double-click **vSphere ESX Agent Manager** and click the **Configure** tab.

   b For each agency that has a name starting with `_NSX_`, select the agency, then click the three-dot menu icon **(**  **)** and select **Delete Agency**.

7 Remove the NSX-V plug-in from vCenter Server.

   a Access the Extension Manager from the Managed Object Browser at `https://<vcenter-ip>/mob/?moid=ExtensionManager`.

   b Click **UnregisterExtension**.

   c In the **UnregisterExtension** dialog box, enter **`com.vmware.vShieldManager`** in the **Value** text box and click **Invoke Method**.

   d In the **UnregisterExtension** dialog box, enter **`com.vmware.nsx.ui.h5`** in the **Value** text box and click **Invoke Method**.

   e You can verify that you unregistered the extensions by going to the Extension Manager page at `https://<vcenter-ip>/mob/?moid=ExtensionManager` and viewing the values for the **extensionList** property.

8 Delete the vSphere Web Client directories and vSphere Client (HTML5) directories for NSX for vSphere and then restart the client services.

a Connect to the vCenter Server system command line.

- If you are using a vCenter Server Appliance, log in as root using the console or SSH. You must log in as root and run the commands from the Bash shell. You can start the Bash shell using the following commands.

```
> shell.set --enabled True
> shell
```

- If you are using vCenter Server for Windows, log in as an administrator using the console or RDP.

b Delete all NSX for vSphere plug-in directories.

Note A plug-in directory might not be present if you have never launched the associated client.

On vCenter Server Appliance, delete the following directories:

- To remove the vSphere Web Client plug-in, delete the `/etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/com.vmware.vShieldManager-<version>-<build>` directory.
- To remove the vSphere Client plug-in, delete the `/etc/vmware/vsphere-ui/vc-packages/vsphere-client-serenity/com.vmware.nsx.ui.h5-<version>-<build>` directory.

On vCenter Server for Windows, delete the following directories:

- To remove the vSphere Web Client plug-in, delete the `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity\com.vmware.vShieldManager-<version>-<build>` directory.
- To remove the vSphere Client plug-in, delete the `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-ui\vc-packages\vsphere-client-serenity\com.vmware.nsx.ui.h5-<version>-<build>` directory.

c Restart the client services on the vCenter Server Appliance or vCenter Server on Windows.

Table 11-14. Client Service Commands

Client Service	vCenter Server Appliance	vCenter Server for Windows
Restart vSphere Web Client On vSphere 6.0, 6.5, and 6.7	`> shell.set --enabled True`	`> cd C:\Program Files\VMware\vCenter`

Table 11-14. Client Service Commands (continued)

Client Service	vCenter Server Appliance	vCenter Server for Windows
	> shell # service-control --stop vsphere-client # service-control --start vsphere-client	Server\bin > service-control --stop vspherewebclientsvc > service-control --start vspherewebclientsvc
Restart vSphere Client On vSphere 6.5, and 6.7	> shell.set --enabled True > shell # service-control --stop vsphere-ui # service-control --start vsphere-ui	> cd C:\Program Files\VMware\vCenter Server\bin > service-control --stop vsphere-ui > service-control --start vsphere-ui
Restart vSphere Client On vSphere 7.0	> shell.set --enabled True > shell # service-control --stop vsphere-ui # service-control --start vsphere-ui	vSphere 7.0 does not support vCenter Server for Windows

9 Power off and delete the NSX Manager VM.

a In vSphere client, navigate to **Hosts and Clusters**.

b Locate the NSX Manager VM. Right click and select **Power Off** then right click and select **Delete from Disk**.

Migrating VMware Integrated Openstack

12

Read the following topics next:

- Overview - Migrating VMware Integrated Openstack
- Migrating VMware Integrated OpenStack

Overview - Migrating VMware Integrated Openstack

System Requirements

Before the migration, verify that your environment meets the following requirements.

- NSX-V versions 6.4.4, 6.4.5, 6.4.6, 6.4.8 and later are supported for all migration modes with the following exceptions:
  - VMware Integrated OpenStack migration supports NSX-V 6.4.11 and later.
  - Cross-vCenter to NSX Federation migration supports NSX-V 6.4.11 and later.
- See the VMware Product Interoperability Matrices for required versions of vCenter Server and ESXi.
- The version of ESXi used in your NSX-V environment must be supported by NSX-T.
- vSphere Distributed Switch versions 6.5.0, 6.6.0, and 7.0 are supported.
- The NSX-V environment must match the NSX-T system requirements for ESXi, vCenter Server, and vSphere Distributed Switch.
- DNS must be functioning in the NSX-V environment.
- If you are migrating VMs that are DHCP servers, you must configure a Segment Security profile appropriately on NSX-T. For more information, see "Create a Segment Security Segment Profile" in the *NSX-T Data Center Administration Guide*.
- If your Edge Service Gateways have their MTU setting changed, you must change the global MTU setting in NSX-T. See Change the Global MTU Setting.

- NSX-V does not support running another transport node inside it running VXLAN on the same transport VLAN. If you have an NSX-T Data Center Edge on an NSX-V host, the NSX-T Data Center Edge VTEP VLAN must be different.

Summary of Features Supported for Migration

A subset of NSX-V features is supported for migration.

Most features have some limitations. When you import your NSX-V configuration for migration, you will get detailed feedback about what features and configurations in your environment are supported or not supported.

See Detailed Feature Support for Migration for detailed information about what is supported for migration.

Table 12-1. Support Matrix for Migration

NSX-V Feature	Supported	Notes
VLAN-backed logical switches	Yes	
Overlay-backed logical switches	Yes	
L2 Bridges	No	
Transport Zones	Yes	
Routing	Yes	
East-West Micro-Segmentation	Yes	
Edge Firewall	Yes	
NAT	Yes	
L2 VPN	Yes	
L3 VPN	Yes	
Load Balancer	Yes	
DHCP and DNS	Yes	
Distributed Firewall	Yes	
Service Composer	Yes	
Grouping objects	Yes	Limitations include number of items, and dynamic expressions making up security groups.

Table 12-1. Support Matrix for Migration (continued)

NSX-V Feature	Supported	Notes
Guest Introspection	Yes	Supported for end-to-end migration only. Only Guest Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service must be registered with NSX-T to create the respective service and vendor templates in NSX-T. Consult the VMware partner before proceeding with migration.
Network Introspection	Yes	Supported for end-to-end migration only. Only Network Introspection service configuration is migrated. Partner service registration, partner service VMs, and vendor templates are not migrated. Before migration, the partner service, vendor templates, and Partner Management Console/Partner Service Manager must be registered with NSX-T. Consult the VMware partner before proceeding with migration.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes (with no secondary NSX Managers)	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. The NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution.	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and VMware Integrated Openstack ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

Detailed Feature Support for Migration

NSX-V features and configurations that can be migrated are listed below.

Platform Support

See the VMware Interoperability Matrix for supported versions of ESXi and vCenter Server.

NSX-V Configuration	Supported	Details
NSX-V with vSAN or iSCSI on vSphere Distributed Switch	Yes	
Pre-existing NSX-T configuration	No	You must deploy a new NSX-T environment. If the migration mode is not for a user-defined topology, during the **Import Configuration** step, all NSX-T Edge node interfaces in the NSX-T environment are shut down. If the NSX-T environment is in use, this will interrupt traffic.
Cross-vCenter NSX	(NSX-T 3.2.1 and later) Yes (NSX-T 3.2.0) Yes, but with no secondary NSX Managers	(NSX-T 3.2.1 and later) Only supported if you choose the **Migrate NSX for vSphere** mode and **User Defined Topology**. (NSX-T 3.2.0) Migration of a single site NSX-V deployment that contains an NSX Manager in primary mode, no secondary NSX Managers, and with universal objects on the primary site, is supported. Such a single site NSX-V deployment is migrated to a single site NSX-T environment (non-federated) with local objects. Migration of a cross-vCenter NSX-V deployment with a primary NSX Manager and multiple secondary NSX Managers is not supported. However, if either the primary or secondary NSX Manager is set to a standalone or transit mode, the migration is supported.
VCF Workload Domains	(NSX-T 3.2.1 and later) Yes	
NSX-V with a Cloud Management Platform, Integrated Stack Solution, or PaaS Solution	Yes	Migration of NSX-V with vRealize Automation is supported. Contact your VMware representative before proceeding with migration. Scripts and integrations might break if you migrate the integrated environments: For example: ■ NSX-V and vCloud Director ■ NSX-V with Integrated Stack Solution ■ NSX-V with PaaS Solution such as Pivotal Cloud Foundry, RedHat OpenShift ■ NSX-V with vRealize Operations workflows

vSphere and ESXi Features

NSX-V Configuration	Supported	Details
ESXi host already in maintenance mode (no VMs)	Yes	
Network I/O Control (NIOC) version 3	Yes	
Network I/O Control (NIOC) version 2	No	
Network I/O Control (NIOC) having vNIC with reservation	No	

NSX-V Configuration	Supported	Details
vSphere Standard Switch	No	VMs and VMkernel interfaces on VSS are not migrated. NSX-V features applied to the VSS cannot be migrated.
vSphere Distributed Switch	Yes	
Stateless ESXi	No	
Host profiles	No	
ESXi lockdown mode	No	Not supported in NSX-T.
ESXi host pending maintenance mode task.	No	
Disconnected ESXi host in vCenter cluster	No	
vSphere FT	No	
vSphere DRS fully automated	Yes	Supported starting in vSphere 7.0
vSphere High Availability	Yes	
Traffic filtering ACL	No	
vSphere Health Check	No	
SRIOV	No	
vmknic pinning to physical NIC	No	
Private VLAN	No	
Ephemeral dvPortGroup	No	
DirectPath IO	No	
L2 security	No	
Learn switch on virtual wire	No	
Hardware Gateway (Tunnel endpoint integration with physical switching hardware)	No	
SNMP	No	
Disconnected vNIC in VM	No	Due to ESX 6.5 limitation, stale entries might present on DVFilter for disconnected VMs. Reboot the VM as a workaround.
VXLAN port number other than 4789	No	

NSX-V Configuration	Supported	Details
Multicast Filtering Mode	No	
Hosts with multiple VTEPs	Yes	

NSX Manager Appliance System Configuration

NSX-V Configuration	Supported	Details
NTP server/time setting	Yes	
Syslog server configuration	Yes	
Backup configuration	Yes	If needed, change NSX-V passphrase to match the NSX-T requirements. It must be at least 8 characters long and contain the following: ■ At least one lowercase letter ■ At least one uppercase letter ■ At least one numeric character ■ At least one special character
FIPS	No	FIPS on/off not supported by NSX-T.
Locale	No	NSX-T only supports English locale
Appliance certificate	No	

Role-Based Access Control

NSX-V Configuration	Supported	Details
Local users	No	
NSX roles assigned to a vCenter user added via LDAP	Yes	VMware Identity Manager must be installed and configured to migrate user roles for LDAP users.
NSX roles assigned to a vCenter group	No	

Certificates

NSX-V Configuration	Supported	Details
Certificates (Server, CA signed)	Yes	This applies to certificates added through truststore APIs only.
Certificate changes during migration	(NSX-T 3.2.0) No (NSX-T 3.2.1 and later) Yes	(NSX-T 3.2.1 and later) Certificate changes are supported when the migration is paused for all migration modes except migrating vSphere networking. Not supported when hosts and workloads are being migrated.

Operations

Details	Supported	Notes
Discovery protocol CDP	See notes.	Yes if migrating to VDS 7.0. No if migrating to N-VDS.
Discovery protocol LLDP	Yes	The listen mode is turned on by default and can't be changed in NSX-T. Only the Advertise mode can be modified.
PortMirroring: ■ Encapsulated remote Mirroring Source (L3)	Yes	Only L3 session type is supported for migration
PortMirroring: ■ Distributed PortMirroring ■ Remote Mirroring Source ■ Remote Mirroring Destination ■ Distributed Port Mirroring (legacy)	No	
L2 IPFIX	Yes	LAG with IPFIX is not supported
Distributed Firewall IPFIX	No	
MAC Learning	Yes	You must enable (accept) forged transmits.
Hardware VTEP	No	
Promiscuous Mode	No	
Resource Allocation	No	vNIC enabled with resource allocation is not supported
IPFIX – Internal flows	No	IPFIX with InternalFlows is not supported

Switch

NSX-V Configuration	Supported	Details
L2 Bridging	No	
Trunk VLAN	Yes	Trunk uplink portgroups must be configured with a VLAN range of 0-4094.
VLAN Configuration	Yes	Configuration with only VLAN (no VXLAN) is supported.
Teaming and Failover: ■ Load Balancing ■ Uplink Failover Order	Yes	Supported options for load balancing (teaming policy): ■ Use explicit failover order ■ Route based on source MAC hash Other load balancing options are not supported.

NSX-V Configuration	Supported	Details
Teaming and Failover: ■ Network Failure Detection ■ Notify Switches ■ Reverse Policy ■ Rolling Order	No	
LACP	Yes	For VDS 7.0 and later, the LACP functionality is not modified during migration. For earlier versions of VDS, a new N-VDS switch replaces the VDS. This will lead to traffic loss during host migration. IPFIX configured on DVS (not DFW IPFIX) is not supported with LACP

Switch Security and IP Discovery

NSX-V Configuration	Supported fro Migration	Details
IP Discovery (ARP, ND, DHCPv4 and DHCPv6)	Yes	The following binding limits apply on NSX-T for migration: ■ 128 for ARP discovered IPs ■ 128 for DHCPv4 discovered IPs ■ 15 for DHCPv6 discovered IPs ■ 15 for ND discovered IPs
SpoofGuard (Manual, TOFU, Disabled)	Yes	
Switch Security (BPDU Filter, DHCP client block, DHCP server block, RA guard)	Yes	
Migrating datapath bindings from Switch Security module in NSX-V to Switch security module in NSX-T	Yes	If SpoofGuard is enabled, bindings are migrated from the Switch Security module to support ARP suppression. VSIP – Switch security not supported as VSIP bindings are migrated as statically configured rules.
Discovery profiles	Yes	The ipdiscovery profiles are created after migration using the IP Discovery configuration for the logical switch and the global and cluster ARP and DHCP configuration.

Central Control Plane

NSX-V Configuration	Supported	Details
VTEP replication per logical switch (VNI) and routing domain	Yes	
MAC/IP replication	No	

NSX-V Configuration	Supported	Details
NSX-V transport zones using multicast or hybrid replication mode	No	
NSX-V transport zones using unicast replication mode	Yes	

NSX Edge Features

NSX-V Configuration	Supported	Details
Routing between Edge Service Gateway and northbound router	Yes	BGP is supported. Static routes are supported. OSPF is supported.
Routing between Edge Services Gateway and Distributed Logical Router	Yes	Routes are converted to static routes after migration.
Load balancer	Yes	See Chapter 1 Migration Modes for details.
VLAN-backed Micro-Segmentation environment	Yes	
NAT64	No	Not supported in NSX-T.
Node level settings on Edge Services Gateway or Distributed Logical Router	No	Node level settings, for example, syslog or NTP server, are not supported. You can configure syslog and NTP manually on the NSX-T edge nodes.
IPv6	No	
Unicast Reverse Path Filter (URPF) configuration for Edge Services Gateway interfaces	No	URPF on NSX-T gateway interfaces is set to Strict.
Maximum Transmission Unit (MTU) configuration Edge Services Gateway interfaces	No	See Change the Global MTU Setting for information about changing the default MTU on NSX-T.
IP Multicast routing	No	
Route Redistribution Prefix Filters	No	
Default originate	No	Not supported in NSX-T.

Edge Firewall

NSX-V Configuration	Supported	Details
Firewall Section: Display name	Yes	Firewall sections can have a maximum of 1000 rules. If a section contains more than 1000 rules, it is migrated as multiple sections.
Action for default rule	Yes	NSX-V API: GatewayPolicy/action NSX-T API: SecurityPolicy.action
Firewall Global Configuration	No	Default timeouts are used
Firewall Rule	Yes	NSX-V API: firewallRule NSX-T API: SecurityPolicy
Firewall Rule: name	Yes	
Firewall Rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: Rule_tag
Sources and destinations in firewall rules: ■ Grouping objects ■ IP addresses	Yes	NSX-V API: ■ source/groupingObjectId ■ source/ipAddress NSX-T API: ■ source_groups NSX-V API: ■ destination/groupingObjectId ■ destination/ipAddress NSX-T API: ■ destination_groups
Firewall rule sources and destinations: ■ vNIC Group	No	
Services (applications) in firewall rules: ■ Service ■ Service Group ■ Protocol/port/source port	Yes	NSX-V API: ■ application/applicationId ■ application/service/protocol ■ application/service/port ■ application/service/sourcePort NSX-T API: ■ Services
Firewall Rule: Match translated	No	Match translated must be ‘false’.
Firewall Rule: Direction	Yes	Both APIs: direction
Firewall Rule: Action	Yes	Both APIs: action
Firewall Rule: Enabled	Yes	Both APIs: enabled

NSX-V Configuration	Supported	Details
Firewall Rule: Logging	Yes	NSX-V API: logging NSX-T API: logged
Firewall Rule: Description	Yes	Both APIs: description

Edge NAT

NSX-V Configuration	Supported	Details
NAT rule	Yes	NSX-V API: natRule NSX-T API: /nat/USER/nat-rules
NAT rule: rule tag	Yes	NSX-V API: ruleTag NSX-T API: rule_tag
NAT rule: action	Yes	NSX-V API: action NSX-T API: action
NAT rule: original address (Source address for SNAT rules, and the destination address for DNAT rules.)	Yes	NSX-V API: originalAddress NSX-T API: source_network for SNAT rule or destination_network for DNAT rule
NAT rule: translatedAddress	Yes	NSX-V API: translatedAddress NSX-T API: translated_network
NAT rule: Applying NAT rule on a specific interface	No	Applied on must be “any”.
NAT rule: logging	Yes	NSX-V API: loggingEnabled NSX-T API: logging
NAT rule: enabled	Yes	NSX-V API: enabled NSX-T API: disabled
NAT rule: description	Yes	NSX-V API: description NSX-T API: description
NAT rule: protocol	Yes	NSX-V API: protocol NSX-T API: Service
NAT rule: original port (source port for SNAT rules, destination port for DNAT rules)	Yes	NSX-V API: originalPort NSX-T API: Service
NAT rule: translated port	Yes	NSX-V API: translatedPort NSX-T API: Translated_ports
NAT rule: Source address in DNAT rule	Yes	NSX-V API: dnatMatchSourceAddress NSX-T API: source_network
NAT rule: Destination address in SNAT rule	Yes	NSX-V API: snatMatchDestinationAddress NSX-T API: destination_network

NSX-V Configuration	Supported	Details
NAT rule: Source port in DNAT rule	Yes	NSX-V API: dnatMatchSourcePort NSX-T API: Service
NAT rule: Destination port in SNAT rule	Yes	NSX-V API: snatMatchDestinationPort NSX-T API: Service
NAT rule: rule ID	Yes	NSX-V API: ruleID NSX-T API: id and display_name

L2VPN

NSX-V Configuration	Supported	Details
L2VPN configuration based on IPSec using pre-shared key (PSK)	Yes	Supported if the networking being stretched over L2VPN is an overlay logical switch. Not supported for VLAN networks.
L2VPN configuration based on IPSec using certificate-based authentication	No	
L2VPN configuration based on SSL	No	
L2VPN configurations with local egress optimizations	No	
L2VPN client mode	No	

L3VPN

NSX-V Configuration	Supported	Details
Dead Peer Detection	Yes	Dead Peer Detection supports different options on NSX-V and NSX-T. You might want to consider using BGP for faster convergence or configure a peer to perform DPD if it is supported.
Changed Dead Peer Detection (dpd) default values for: ■ dpdtimeout ■ dpdaction	No	In NSX-T, dpdaction is set to “restart” and cannot be changed. If NSX-V setting for dpdtimeout is set to 0, dpd is disabled in NSX-T. Otherwise, any dpdtimeout settings are ignored and the default value is used.
Changed Dead Peer Detection (dpd) default values for: ■ dpddelay	Yes	NSX-V dpdelay maps to NSX-T dpdinternal.
Overlapping local and peer subnets of two or more sessions.	No	NSX-V supports policy-based IPSec VPN sessions where the local and peer subnets of two or more sessions overlap with each other. This behavior is not supported on NSX-T. You must reconfigure the subnets so they do not overlap before you start the migration. If this configuration issue is not resolved, the Migrate Configuration step fails.

NSX-V Configuration	Supported	Details
IPSec sessions with peer endpoint set as any.	No	Configuration is not migrated.
Changes to the extension securelocaltrafficbyip.	No	NSX-T Service Router does not have any local generated traffic that needs to be sent over tunnel.
Changes to these extensions: auto, sha2_truncbug, sareftrack, leftid, leftsendcert, leftxauthserver, leftxauthclient, leftxauthusername, leftmodecfgserver, leftmodecfgclient, modecfgpull, modecfgdns1, modecfgdns2, modecfgwins1, modecfgwins2, remote_peer_type, nm_configured, forceencaps,overlapip, aggrmode, rekey, rekeymargin, rekeyfuzz, compress, metric,disablearrivalcheck, failureshunt,leftnexthop, keyingtries	No	Those extensions are not supported on NSX-T and changes to them are not migrated.

Load Balancer

The following table applies to migrating NSX-V load balancer to NSX-T load balancer (NLB) or NSX-T Advanced Load Balancer (ALB). For information about migrating NSX-V load balancer to ALB, see Migrating NSX-V Load Balancer to Advanced Load Balancer.

NSX-V Configuration	Supported	Details
Monitor / health-checks for: ■ LDAP ■ DNS ■ MSSQL	See details.	(NLB) The monitors are not migrated. (ALB) LDAP and MSSQL monitors are not migrated. DNS monitor is migrated if ALB has an enterprise license, but not if it has a basic license.
Application rules	No	NSX-V uses application rules based on HAProxy to support L7. In NSX-T, the rules are based on NGINX. The application rules cannot be migrated. You must create new rules after migration.
L7 virtual server port range	(NLB) No (ALB) Yes	

NSX-V Configuration	Supported	Details
IPv6	No	If IPv6 is used in virtual server, the whole virtual server would be ignored. If IPv6 is used in pool, the pool would be still migrated, however, the related pool member would be removed.
URL, URI, HTTPHEADER algorithms	See details.	(NLB) Pools with these algorithms are not migrated. (ALB) Supported if ALB has an enterprise license. With a basic license, you will be provided feedback to select a different algorithm.
Isolated pool	No	
LB pool member with different monitor port	See details.	(NLB) The pool member which has a different monitor port is not migrated. (ALB) The pool member is migrated but will not be in the monitor port configuration.
Pool member minConn	No	
Monitor extension	No	
SSL sessionID persistence / table	(NLB) No (ALB) Yes	
MSRDP persistence / session table	No	
Cookie app session / session table	(NLB) No (ALB) Yes	
App persistence	(NLB) No (ALB) Yes	
Monitor for: ■ Explicit escape ■ Quit ■ Delay	No	
Monitor for: ■ Send ■ Expect ■ Timeout ■ Interval ■ maxRetries	Yes	
Haproxy Tuning/IPVS Tuning	No	
Pool IP filter ■ IPv4 addresses	Yes	IPv4 IP addresses are supported. If Any is used, only the IPv4 addresses of the IP pool are migrated.

NSX-V Configuration	Supported	Details
Pool IP Filter ■ IPv6 addresses	No	
Pool containing unsupported grouping object: ■ Cluster ■ Datacenter ■ Distributed port group ■ MAC set ■ Virtual App	No	If a pool includes an unsupported grouping object, those objects are ignored, and the pool is created with supported grouping object members. If there are no supported grouping object members, then an empty pool is created.

DHCP and DNS

Table 12-2. DHCP Configuration Topologies

NSX-V Configuration	Supported	Details
DHCP Relay configured on Distributed Logical Router pointing to a DHCP Server configured on a directly connected Edge Services Gateway	Yes	The DHCP Relay server IP must be one of the Edge Services Gateway's internal interface IPs. The DHCP Server must be configured on an Edge Services Gateway that is directly connected to the Distributed Logical Router configured with the DHCP relay. It is not supported to use DNAT to translate a DHCP Relay IP that does not match an Edge Services Gateway internal interface.
DHCP Relay configured on Distributed Logical Router only, no DHCP Server configuration on connected Edge Services Gateway	No	
DHCP Server configured on Edge Services Gateway only, no DHCP Relay configuration on connected Distributed Logical Router	No	

Table 12-3. DHCP Features

NSX-V Configuration	Supported	Details
IP Pools	Yes	
Static bindings	Yes	
DHCP leases	Yes	
General DHCP options	Yes	
Disabled DHCP service	No	In NSX-T you cannot disable the DHCP service. If there is a disabled DHCP service on NSX-V it is not migrated.

Table 12-3. DHCP Features (continued)

NSX-V Configuration	Supported	Details

DHCP option: "other"	No	The "other" field in dhcp options is not supported for migration.
		For example, dhcp option '80' is not migrated.
		<pre><dhcpOptions>
		  <other>
		  <code>80</code>
		  <value>2f766172</value>
		  </other>
		</dhcpOptions></pre>

Orphaned ip-pools/ bindings	No	If ip-pools or static-bindings are configured on a DHCP Server but are not used by any connected logical switches, these objects are skipped from migration.
DHCP configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DHCP service on a centralized service port, so the DHCP service configuration is not migrated for these interfaces.

Table 12-4. DNS Features

NSX-V Configuration	Supported	Details
DNS views	Yes	Only the first dnsView is migrated to the NSX-T default DNS forwarder zone.
DNS configuration	Yes	You must provide available DNS listener IPs for all Edge Nodes. A message is displayed during Resolve Configuration to prompt for this.
DNS – L3 VPN	Yes	You must add the newly configured NSX-T DNS listener IPs into the remote L3 VPN prefix list. A message is displayed during Resolve Configuration to prompt for this.
DNS configured on Edge Service Gateway with directly connected logical switches	No	During migration, directly connected Edge Service Gateway interfaces are migrated as centralized service ports. However, NSX-T does not support DNS Service on a centralized service port, so the DNS Service configuration is not migrated for these interfaces.

Distributed Firewall (DFW)

NSX-V Configuration	Supported	Details
Identity Firewall	Yes	
Section - ■ Name ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	If a firewall section has more than 1000 rules, then the migrator will migrate the rules in multiple sections of 1000 rules each.

NSX-V Configuration	Supported	Details
Universal Sections	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ IP Address / Range / CIDR ■ Logical Port ■ Logical Switch	Yes	
Rule – Source / Destination: ■ VM ■ Logical Port ■ Security Group / IP Set / MAC Set	Yes	maps to Security Group
Rule – Source / Destination: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Source / Destination: ■ Universal Logical Switch	Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rule – Source / Destination: ■ VMs not on NSX-V hosts	No	NSX-T does not support referencing objects not connected to NSX-T portgroups or networks. Those objects will be lost from the source or destination when migration is completed. To avoid this issue use IP addresses in NSX-V to reference those objects before migration.
Rule – Applied To: ■ ANY	Yes	maps to Distributed Firewall
Rule – Applied To: ■ Security Group ■ Logical Port ■ Logical Switch ■ VM	Yes	maps to Security Group

NSX-V Configuration	Supported	Details
Rule – Applied To: ■ Cluster ■ Datacenter ■ DVPG ■ vSS ■ Host	No	
Rule – Applied To: ■ Universal Logical Switch	No Yes if the NSX-V deployment has an NSX Manager in primary mode and no secondary NSX Managers.	
Rules Disabled in Distributed Firewall	Yes	
Disabling Distributed Firewall on a cluster level	No	When Distributed Firewall is enabled on NSX-T, it is enabled on all clusters. You cannot enable it on some clusters and disable on others.
DFW Exclusion List	No	DFW exclusion lists are not migrated. You need to re-create them on NSX-T after migration.

Partner Services: East-West Network Introspection

NSX-V Configuration	Supported	Details
Service	No	Service registration is not migrated. Partner must register the service with NSX-T before migration.
Vendor Template	No	Vendor template is not migrated. Partner must register the vendor template with NSX-T before migration.
Service Profile	No	Service profiles are not migrated. Either you or the partner must create the service profiles before migration. In the Resolve Configuration step of the migration, you will be prompted to map each NSX-V service profile to an NSX-T service profile. If you skip the mapping of service profiles, the rules that use these service profiles are not migrated. A service chain in NSX-T will be created for each service profile in NSX-V. The service chain is created with the following naming convention: Service-Chain-*service_profile_name* The same service profile is used in the forward path and reverse path of the service chain.
Service Instance	No	Partner service virtual machines (SVMs) are not migrated. The NSX-V partner SVMs cannot be used in NSX-T. For east-west Network Introspection service in NSX-T, partner service VMs must be deployed on an overlay segment.

NSX-V Configuration	Supported	Details
Section ■ Name ■ ID ■ Description ■ TCP Strict ■ Stateless Firewall	Yes	A section maps to a redirection policy. ID is user-defined, and not auto-generated in NSX-T. If a firewall section in NSX-V has more than 1000 rules, the rules will be migrated in multiple sections of 1000 rules each. For example, if a section contains 2500 rules, three policies will be created: Policy 1 with 1000 rules, Policy 2 with 1000 rules, and Policy 3 with 500 rules. Stateful or stateless firewall rules in NSX-V are migrated to stateful or stateless redirection rules in NSX-T.
Partner Services: Rules		
Name	Yes	
Rule ID	Yes	Rule ID is system generated. It can be different from the rule ID in NSX-V.
Negate Source	Yes	
Negate Destination	Yes	
Source/Destination ■ VM ■ Security Group ■ IP Set ■ vNIC	Yes	
Services/Service Groups	Yes	For details, see the Services and Service Groups table.
Advanced Settings ■ Direction ■ Packet Type ■ Rule Tag ■ Comments ■ Log	Yes	

NSX-V Configuration	Supported	Details
Service Profile and Action ■ Service Name ■ Service Profile ■ Action ■ Service Profile Bindings	Yes	A service profile binding can have Distributed Virtual Port Groups (DVPG), Logical Switches, and Security Groups as its members. A service profile binding in NSX-V maps to the Applied To field of a redirection rule in NSX-T. Applied To field accepts only Groups, and this field determines the scope of the rule. In NSX-T, rule redirection is at the level of a policy. All rules in a redirection policy have the same scope (Applied To). Applied To field in an NSX-T redirection rule can have a maximum of 128 members. If the number of members in a service profile binding exceeds 128, reduce them to <= 128 before starting the migration. For example, assume that a service profile binding has 140 members (Security Groups). Do the following steps in NSX-V before starting the migration: 1 Create a dummy Security Group. 2 Move 13 Security Groups to this dummy Security Group. In other words, the dummy Security Group has 13 members. 3 Remove the binding of these 13 Security Groups from the service profile. You now have 127 members in the service profile binding (140-13). 4 Add the dummy Security Group to the service profile binding. Now, the total number of members in the service profile binding is 128 (127 + 1).
Enable/Disable Rule	Yes	

Service Segment

A service segment will be created in the overlay transport zone that you select in the Resolve Configuration step of the migration. In the NSX-V environment, if the VXLAN transport zone is not prepared with NSX-V, you have the option to select the default overlay transport zone in NSX-T to create the service segment. If one or multiple VXLAN transport zones are prepared with NSX-V, you must select any one overlay transport zone to create the service segment in NSX-T.

Service Profile Priority

In NSX-V, a service profile has a priority. If a service has multiple service profiles, and multiple profiles are bound to the same vNIC, the service profile with higher priority is applied first on the vNIC. However, in NSX-T, service profile does not have a priority. When multiple redirection rules have the same Applied To setting, the rule order decides which rule is hit first. In other words, the rules with a higher profile priority will be placed before the rules with a lower profile priority in the NSX-T rule table. For a detailed example, see scenario 2 in Order of Migrated Network Introspection Rules in NSX-T.

Service Precedence

To redirect traffic to multiple services, NSX-V uses multiple DVFilter slots in the service insertion data path. One DVFilter slot is used to redirect traffic to one service. A service with high precedence is placed higher in the slot compared to a service with low precedence. In NSX-T, only a single DVFilter slot is used and it redirects traffic to a service chain. After migration to NSX-T, the rules that use a partner service with higher precedence are placed before the rules that use a partner service with a lower precedence. For a detailed example, see scenario 3 in Order of Migrated Network Introspection Rules in NSX-T.

Redirection of traffic on a vNIC to multiple partner services is not supported. Redirection to only a single partner service is supported. Although, all the NSX-V rules are migrated to NSX-T, the migrated rule configurations use a service chain with only one service profile. You cannot modify an existing service chain that is used in redirection rules.

Workaround: To redirect traffic on a vNIC to multiple services, create a new service chain and define the order of service profiles in the service chain. Update the migrated rules to use this new service chain.

Network Introspection Service on VMs Connected to a VM Network

In the NSX-V environment, if Network Introspection service rules are running on VMs that are connected to a VM Network, these VMs lose security protection after host migration. To ensure that the Network Introspection rules are enforced on the vNICs of these VMs post host migration, you must connect these VMs to an NSX-T segment.

Grouping Objects and Service Composer

IP Sets and MAC Sets are migrated to NSX-T as groups. See **Inventory > Groups** in the NSX-T Manager web interface.

Table 12-5. IP Sets and MAC Sets

NSX-V Configuration	Supported	Details
IP Sets	Yes	IP sets with up to 2 million members (IP addresses, IP address subnets, IP ranges) can be migrated. IP sets with more members are not migrated.
Mac Sets	Yes	MAC sets with up to 2 million members can be migrated. MAC sets with more members are not migrated.

Security Groups are supported for migration with the limitations listed. Security Groups are migrated to NSX-T as Groups. See **Inventory > Groups** in the NSX-T Manager web interface.

NSX-V has system-defined and user-defined Security Groups. These are all migrated to NSX-T as user-defined Groups.

The total number of ‘Groups’ after migration might not be equal to the number of Security Groups on NSX-V. For example, a Distributed Firewall rule containing a VM as its source would be migrated into a rule containing a new Group with the VM as its member. This increases the total number of groups on NSX-T after migration.

Table 12-6. Security Groups

NSX-V Configuration	Supported	Details
Security Group with members that don't exist	No	If any of the members of the Security Group do not exist, then the Security Group is not migrated.
Security Group that contains a Security Group with unsupported members	No	If any members of the Security Group are not supported for migration, the Security Group is not migrated. If a Security Group contains a Security Group with unsupported members, the parent Security Group is not migrated.
Exclude membership in Security Group	No	Security Groups with an exclude member directly or indirectly (via nesting) are not migrated
Security Group Static Membership	Yes	A Security Group can contain up to 500 static members. However, system-generated static members are added if the Security Group is used in Distributed Firewall rules, lowering the effective limit to 499 or 498. ■ If the Security Group is used in either layer 2 or layer 3 rules, one system-generated static member is added to the Security Group. ■ If the Security Group is used in both layer 2 and layer 3 rules, two system-generated static members are added. If any members do not exist during the Resolve Configuration step, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Cluster ■ Datacenter ■ Directory Group ■ Distributed Port Group ■ Legacy Port Group / Network ■ Resource Pool ■ vApp	No	If a security group contains any of the unsupported member types, the security group is not migrated.
Security Group Member Types (Static or Entity Belongs To): ■ Security Group ■ IP Sets ■ MAC Sets	Yes	Security groups, IP sets, and MAC sets are migrated to NSX-T as Groups. If an NSX-V security group contains an IP set, MAC set, or nested security group as a static member, the corresponding Groups are added to the parent Group. If one of these static members was not migrated to NSX-T, the parent security group does not migrate to NSX-T. For example, an IP set with more than 2 million members cannot migrate to NSX-T. Therefore, a security group that contains an IP set with more than 2 million members cannot migrate.
Security Group Member Types (Static or Entity Belongs To): ■ Logical Switch (Virtual Wire)	Yes	If a security group contains logical switches that do not migrate to NSX-T segments, the security group does not migrate to NSX-T.

Table 12-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Security Group Member Types (Static or Entity Belongs To): ■ Security Tag	Yes	If a security tag is added to the security group as a static member or as a dynamic member using Entity Belongs To, the security tag must exist for the security group to be migrated. If the security tag is added to the security group as a dynamic member (not using Entity Belongs To), the existence of the security tag is not checked before migrating the security group.
Security Group Member Types (Static or Entity Belongs To): ■ vNIC ■ Virtual Machine	Yes	■ vNICs and VMs are migrated as an ExternalIDExpression. ■ Orphaned VMs (VMs deleted from hosts) are ignored during Security Group migration. ■ Once the Groups appear on NSX-T, the VM and vNIC memberships are updated after some time. During this intermediate time, there can be temporary groups and their temporary groups might appear as members. However, once the Host Migration has finished, these extra temporary groups are no longer seen.
Using “Matches regular expression” operator for dynamic membership	No	This affects Security Tag and VM Name only. “Matches regular expression” is not available for other attributes.
Using other available operators for dynamic membership criteria for attributes: ■ Security Tag ■ VM Name ■ Computer Name ■ Computer OS Name	Yes	Available operators for VM Name, Computer Name, and Computer OS Name are Contains, Ends with, Equals to, Not equals to, Starts with. Available operators for Security Tag are Contains, Ends with, Equals to, Starts with.

Table 12-6. Security Groups (continued)

NSX-V Configuration	Supported	Details
Entity Belongs to criteria	Yes	The same limitations for migrating static members apply to Entity Belongs to criteria. For example, if you have a Security Group that uses Entity Belongs to a cluster in the definition, the Security Group is not migrated. Security Groups that contain Entity Belongs to criteria that are combined with AND are not migrated.
Dynamic membership criteria operators (AND, OR) in Security Group	Yes.	When you define dynamic membership for an NSX-V Security Group, you can configure the following: ■ One or more dynamic sets. ■ Each dynamic set can contain one or more dynamic criteria. For example, “VM Name Contains web”. ■ You can select whether to match Any or All dynamic criteria within a dynamic set. ■ You can select to match with AND or OR across dynamic sets. NSX-V does not limit the number of dynamic criteria, dynamic sets, and you can have any combinations of AND and OR. In NSX-T, you can have a group with five expressions. NSX-V security groups which contain more than five expressions are not migrated. Examples of security groups that can be migrated: ■ Up to 5 dynamic sets related with OR where each dynamic set contains up to 5 dynamic criteria related with AND (All in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with OR (Any in NSX-V). ■ 1 dynamic set containing 5 dynamic criteria related with AND (All in NSX-V). All member types must be the same. ■ 5 dynamic sets related with AND and each dynamic set containing exactly 1 dynamic criteria. All member types must be the same. Using “Entity belongs to” criteria with AND operators is not supported. All other combinations or definitions of a security group containing unsupported scenarios are not migrated.

In NSX-V, security tags are objects which can be applied to VMs. When migrated to NSX-T security tags are attributes of a VM.

Table 12-7. Security Tags

NSX-V Configuration	Supported	Details
Security Tags	Yes	If a VM has 25 or fewer security tags applied, migration of security tags is supported. If more than 25 security tags are applied, no tags are migrated. Note: If security tags are not migrated, the VM is not included in any groups defined by tag membership. Security tags that are not applied to any VM are not migrated.

Services and Service Groups are migrated to NSX-T as Services. See **Inventory > Services** in the NSX-T Manager web interface.

Table 12-8. Services and Service Groups

NSX-V Configuration	Supported	Details
Services and Service Groups (Applications and Application Groups)	Yes	Most of the default Services and Service Groups are mapped to NSX-T Services. If any Service or Service Group is not present in NSX-T, a new Service is created in NSX-T.
APP_ALL and APP_POP2 Service Groups	No	These system-defined service groups are not migrated.
Services and Service Groups with naming conflicts	Yes	If a name conflict is identified in NSX-T for a modified Service or Service Group a new Service is created in NSX-T with a name in format: <NSXv-Application-Name> migrated from NSX-V
Service Groups that combine layer 2 services with services in other layers	No	
Empty Service Groups	No	NSX-T does not support empty Services.
Layer 2 Services	Yes	NSX-V layer 2 Services are migrated as NSX-T Service Entry EtherTypeServiceEntry.
Layer 3 Services	Yes	Based on the protocol, NSX-V layer 3 Services are migrated to NSX-T Service Entry as follows: ■ TCP/UDP protocol: L4PortSetServiceEntry ■ ICMP / IPV6ICMP protocol: ICMPTypeServiceEntry ■ IGMP protocol: IGMPTypeServiceEntry ■ Other protocols: IPProtocolServiceEntry
Layer 4 Services	Yes	Migrated as NSX-T Service Entry ALGTypeServiceEntry.
Layer 7 Services	Yes	Migrated as NSX-T Service Entry PolicyContextProfile If an NSX-V Layer 7 application has a port and protocol defined, a Service is created in NSX-T with the appropriate port and protocol configuration and mapped to the PolicyContextProfile.
Layer 7 Service Groups	No	
Distributed Firewall, Edge Firewall, or NAT rules that contain port and protocol	Yes	NSX-T requires a Service to create these rules. If an appropriate Service exists, it is used. If no appropriate Service exists, a Service is created using the port and protocol specified in the rule.

Table 12-9. Service Composer

NSX-V Configuration	Supported	Details
Service Composer Security Policies	Yes	Firewall rules defined in a Security Policy are migrated to NSX-T as Distributed Firewall rules. Disabled firewall rules defined in a Service Composer Security Policy are not migrated. Guest Introspection rules or Network Introspection rules defined in a Service Composer Security Policy are migrated. If the Service Composer status is not in sync, the Resolve Configuration step shows a warning. You can skip the migration of Service Composer policies by skipping the relevant Distributed Firewall sections. Alternatively, you can roll back the migration, get Service Composer in sync with Distributed Firewall, and restart the migration.
Service Composer Security Policies not applied to any Security Groups	No	

Active Directory Server Configuration

Configuration	Supported	Details
Active Directory (AD) server	No	

Limits Supported for Migration

Before the migration, make sure that the NSX-V configurations do not exceed the limits listed below.

Table 12-10. Single-Site Limits

Feature	Limit
Hosts per NSX Manager (Single vCenter - Transport Zone)	512
vCenter Clusters	64
Distributed Logical Router interfaces per Distributed Logical Router	999
Logical Switches	10000
ECMP paths	8
Static routes per Edge Services Gateway	10000
Virtual interfaces per hypervisor host	150
NAT rules per Edge Services Gateway	4096
Edge firewall rules per Edge Services Gateway	2000

Table 12-10. Single-Site Limits (continued)

Feature	Limit
Security Groups per NSX Manager	10000
Distributed Firewall rules per host	10000
Distributed Firewall sections	5600
Distributed Firewall rules per NSX Manager	100000
Networks per L2VPN client-server pair	100
L2VPN clients (spoke) handled by a single L2VPN server (hub)	2
IPSec tunnels per Edge Services Gateway	500
DHCP leases per Edge Services Gateway	2048
IP Sets	10000
Security Tags	9000
Security Tags per VM	30
Grouping Objects (Security Groups to which a virtual machine can be a member)	100
Grouping Objects per NSX Manager	10000 Grouping Objects include Security Groups, IP Sets, and Security Tags. The sum of all their values in your NSX-V environment must not exceed 10000.
Virtual servers per load balancer	1024
AD groups	70000
Users in the AD domain	100000
AD groups per individual user	200
VMs using terminal service per host	8
Hosts (when migrating IDFW)	250
VMs with Network Introspection enabled	1000
Network Introspection rules per NSX Manager	3500
VMs per Security Group with Network Introspection enabled	1000
VMs per host with Network Introspection enabled	125
Guest Introspection VMs per host	40

Table 12-11. Load Balancer (if migrating to NSX-T Load Balancer)

Feature	Limit
Pools per Load Balancer (large and quad large ESG)	64
Pools per Load Balancer (x-large ESG)	1000
Load Balancer servers per pool	32
Load Balancer pools per ESG	64
Load Balancer VIPs per ESG	64

For load balancer limits if you are migrating to Advanced Load Balancer, see https://avinetworks.com/docs/21.1/avi-controller-sizing.

Changes Made During Host Migration in an End-to-End Migration

During the host migration step in an end-to-end migration, changes are made to migrate NSX-V hosts to NSX-T hosts.

- NSX-V software is uninstalled.
- NSX-T software is installed.
- For vSphere Distributed Switch versions 6.5.0 and 6.6.0:

  Hosts are configured with N-VDS to replace vSphere Distributed Switches:

  - Each N-VDS is created with a name that references the distributed switch name. For example, distributed switch `ComputeSwitchA` is created as N-VDS `nvds.ComputeSwitchA`.
  - If different clusters use different distributed switches to back logical switches, an N-VDS is created with a name that combines all the distributed switch names. For example, if `ComputeCluster1` and `ComputeCluster2` use distributed switch `ComputeSwitchA` to back logical switches and `ComputeCluster3` uses `ComputeSwitchB` to back logical switches, the N-VDS is created as `nvds.ComputeSwitchA.ComputeSwitchB`.
  - PNICs and vmks in the vSphere Distributed Switch are migrated to N-VDS.
  - NSX-V VTEPs are migrated to NSX-T TEPs.
- For vSphere Distributed Switches version 7.0:

  Hosts configured for vSphere Distributed Switch version 7.0 continue using the same switch after migration.

  - PNICs and vmks in the vSphere Distributed Switch remain connected on the same vSphere Distributed Portgroups.
  - NSX-V VTEPs are migrated to NSX-T TEPs and connected to standalone ports on the same vSphere Distributed Switch.

Note: If NSX-V has multiple VTEPs and a single LAG configuration, after migration the LAG will have a single TEP and in failover mode.

Virtual Machine Deployment During an End-to-End Migration

After you start an end-to-end migration, do not change the NSX-V environment. If you want to deploy VMs during the migration, wait until some of the NSX-V hosts have migrated to NSX-T and deploy the VMs on NSX-T hosts. Connect the VMs to NSX-T segments and install VMware Tools on the VMs.

Deploying on NSX-T with VMware Tools installed ensures that the VMs are populated into security groups and receive the intended Distributed Firewall policies.

Caution VMs deployed without VMware Tools installed, or deployed on NSX-V do not receive the intended Distributed Firewall policies.

If you use vSphere templates to deploy VMs, update the templates to use NSX-T segments for the VM network configuration. Specifying NSX-T segments ensures that any VMs deployed using the templates are deployed on NSX-T hosts.

If you use automation tools to deploy VMs on vSphere, but do not use vSphere templates, you might need to change your automation tool configuration to ensure that the VMs are deployed on NSX-T.

Order of Migrated Network Introspection Rules in NSX-T

In NSX-V, traffic redirection to partner services is at a rule level, and not at the section level. That is, a single section in NSX-V can have rules redirecting the network traffic to multiple service profiles of a single partner service or multiple partner services.

However, in NSX-T, redirection is at a policy level. Therefore, if a single firewall section in NSX-V has rules redirecting to multiple service profiles, multiple NSX-T policies will be created.

Read the scenarios in this topic for examples about rule ordering in NSX-T.

This topic uses the following acronyms:

- SP: Service Profile
- SG: Security Group
- SC: Service Chain

Scenario 1: Single Partner Service, Single Service Profile

A single network introspection partner service is running. This partner service contains a single service profile.

Rule configuration in NSX-V is as follows:

- SP1 is bound to SG-1 and SG-2.
- Network traffic from SG-A to SG-B is redirected to SP-1.
- Network traffic from SG-P to SG-Q is redirected to SP-1.

Migrated rule configuration in NSX-T is as follows:

- SC-1 contains SP-1 in the forward and reverse path of the traffic.
- Network traffic from SG-A to SG-B is redirected to SC-1. This rule is applied on SG-1 and SG-2.
- Network traffic from SG-P to SG-Q is redirected to SC-1. This rule is applied on SG-1 and SG-2.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-P to SG-Q, Redirect to SC-1

Scenario 2: Single Partner Service, Multiple Service Profiles

A partner service has two service profiles SP-1 and SP-2.

Case 2A: SP-1 has higher priority than SP-2

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-1 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1
	Policy 2 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2

Case 2B: SP-2 has higher priority than SP-1

In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

In NSX-T, SC-1 contains SP-1, and SC-2 contains SP-2 in the forward and reverse path of the traffic.

In this case, rules redirecting to SC-2 are placed first in the NSX-T rule table.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-P to SG-Q, Redirect to SP-2 Section 2 ■ Rule 3: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-2) ■ Rule 2: SG-P to SG-Q, Redirect to SC-2 Policy 2 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 Policy 3 (Redirect to SC-1) ■ Rule 3: SG-P to SG-Q, Redirect to SC-1

Scenario 3: Two Partner Services, One Service Profile Per Partner

Service-1 from partner 1 has higher precedence than Service-2 from partner 2. Service-1 contains SP-1 and Service-2 contains SP-2. In NSX-V, SP-1 is bound to SG-1, and SP-2 is bound to SG-2 and SG-3.

NSX-V	NSX-T
Section 1 ■ Rule 1: SG-A to SG-B, Redirect to SP-1 ■ Rule 2: SG-A to SG-C, Redirect to SP-1 ■ Rule 3: SG-P to SG-Q, Redirect to SP-2 ■ Rule 4: SG-A to SG-D, Redirect to SP-1 Section 2 ■ Rule 5: SG-P to SG-Q, Redirect to SP-1	Policy 1 (Redirect to SC-1) ■ Rule 1: SG-A to SG-B, Redirect to SC-1 ■ Rule 2: SG-A to SG-C, Redirect to SC-1 ■ Rule 4: SG-A to SG-D, Redirect to SC-1 Policy 2 (Redirect to SC-1) ■ Rule 5: SG-P to SG-Q, Redirect to SC-1 Policy 3 (Redirect to SC-2) ■ Rule 3: SG-P to SG-Q, Redirect to SC-2

Migrating VMware Integrated OpenStack

You can migrate your VMware Integrated OpenStack (VIO) deployment from NSX-V to NSX-T. During the migration, the VIO control plane must be in read-only mode.

Datapath connectivity for VMs is unaffected during the migration, except for brief interruptions during north-south cutover and host migration. This migration must be performed during a single maintenance window.

Overview of the Migration Process

1 Install NSX-T.

2 Prepare NSX-T for VIO. This requires setting up tier-0 gateways or VRF-lites for external networks, as well as configuring edge clusters, DHCP server profiles, and metadata proxies. For more information, see https://docs.vmware.com/en/VMware-Integrated-OpenStack/index.html.

3 Get the neutron migrator bundle, which is part of the VIO deliverables.

4 Configure the neutron migrator.

5 Deploy the neutron migrator pod.

6 From the NSX Manager UI:

- Start the edge cutover migration.
- Handle feedback and complete the migration.
- Start the host migration.
- Handle feedback and complete the migration.

7 Wait for the neutron migrator pod to complete.

8 Delete the neutron migrator deployment.

9 Remove the VIO installation in NSX-V.

Prerequisites

- VIO 7.2.0 or later
- NSX-V 6.4.11 or later
- vSphere 6.7 or later (It is recommended that you upgrade ESXi hosts to 7.0 or later before the migration.)

The neutron migrator pod will run the following validation checks. Checks producing a warning can be bypassed via the neutron migrator configuration.

- Number of address pairs allowed in Neutron (number of manual address bindings must not exceed 128)
- Number of multiple subnets with DHCP per logical switch (only one allowed in NSX-T)
- Number of router uplinks per network (only one in NSX-T)
- Host groups - If HA for NSX Edge nodes is enabled and host groups are specified for the edge nodes to be placed. This will generate a warning.
- Edge HA is ignored in NSX-T as it does not apply. This will generate a warning.
- Provider networks or external networks based on a DVS port group are not supported in the NSX-T plugin.
- Multi-provider VLAN networks are not supported.
- Load balancing topologies not supported by the NSX-T plugin (for example, a load balancer with members from various subnets not uplinked to the same edge router or a load balancer on a network not connected to a Neutron router).
- Usage of invalid addresses for NSX-T (for example, overlap with transit network).
- VMs deployed on external networks. They do not work on NSX-T.
- Reachability of subnets for load balancing members. NSX-T requires that all the load balancer's subnets are attached to the same gateway.

On NSX-T there must not be any openstack-owned resources (for instance resources from a previous VIO deployments on the NSX-T instance).

See Chapter 8 In-Place Migration of Specific Parts of NSX-V for any preparations that are needed for the edge cutover migration and the host migration.

Preparing for the Migration - Sizing NSX-T Edge Cluster

The NSX-T edge cluster must have enough slots for OpenStack load balancers (LBs). To determine the list of tier-1 gateways that will host an LB service, do the following:

- For each OpenStack VIP, find the corresponding subnet, and retrieve the router it is uplinked to, unless the subnet is on an external network.
- For each OpenStack LB pool, list the members. Find the subnet they belong to and retrieve the router the subnet is uplinked to.

The number of routers found, together with the size of the largest OpenStack LB, will determine the number of LB slots required on the NSX-T edge cluster. For each LB, two slots will be required, for the active and standby service routers. Refer to https://configmax.vmware.com for the maximum number of load balancers that can run on each NSX-T edge appliance.

Preparing for the Migration - Configuring TEP IP Pool

During host migration, NSX-V and NSX-T TEPs must be able to reach each other to ensure connectivity. You must configure the NSX-T TEP IP pool so that it can route traffic to NSX-V TEPs.

NSX-V Configuration Parameters not Supported in NSX-T

The following table lists the unsupported NSX-V parameters and the reasons.

Parameter	Description	Reason
cluster_moid	Lists IDs of clusters used by Openstack.	Not applicable in NSX-T.
datacenter_moid	Identifies the datacenter for deploying NSX-V edge appliances.	Not applicable in NSX-T.
deployment_container_id	Identifies deployment container for NSX-V edges.	Not applicable in NSX-T.
resource_pool_id	Identifies resource pool for NSX-V edges.	Not applicable in NSX-T.
datastore_id	Identifies datastore for NSX-V edges.	Not applicable in NSX-T.
ha_datastore_id	Additional datastore if edge HA is enabled.	Not applicable in NSX-T.
ha_placement_random	Divide active edges between primary and secondary datastore.	Not applicable in NSX-T.
edge_host_groups	Ensure active/backup edges are placed in listed host groups.	Not applicable in NSX-T.
external_network	ID of DVPG to use for physical network uplink.	Not applicable in NSX-T.

Parameter	Description	Reason
task_status_check_interval	Interval for checking for task completion.	Not applicable in NSX-T.
vdn_scope_id	ID of network scope object for VXLAN virtual wires.	VDN scopes are replaced by NSX-T overlay transport zones.
dvs_id	ID of DVS connected to management and edge cluster. Also used by default for VLAN networks.	DVS is replaced by VLAN transport zone in NSX-T.
maximum_tunnels_per_vnic	Maximum number of sub-interface supported by a VNIC on an edge appliance.	Not applicable in NSX-T.
backup_edge_pool	Defines the size for NSX-V edge pool to be used by the Openstack deployment.	Not applicable in NSX-T.
mgm_net_moid	Portgroup ID for metadata proxy management network.	Not applicable in NSX-T.
mgt_net_proxy_ips	Comma-separated list of management network IP addresses.	Not applicable in NSX-T.
mgt_net_proxy_netmask	Management network netmask for metadata proxy.	Not applicable in NSX-T.
mgt_net_default_gateway	Management network default gateway for metadata proxy.	Not applicable in NSX-T.
nova_metadata_ips	IP addresses used by Nova metadata service.	Provided in NSX-T metadata proxy configuration.
nova_metadat_port	Port used by the Nova metadata service.	Provided in NSX-T metadata proxy configuration.
spoofguard_enabled	By default spoofguard is enabled in NSX-V but if you disable spoofguard in NSX-V, spoofguard will be enabled in NSX-T after migration.	Enabled by default in NSX-T (cannot be globally turned off).
use_exclude_list	Use NSX-V exclude list component when port security is disabled and spoofguard is enabled.	Enabled by default in NSX-T (cannot be globally turned off).
tenant_router_types	Ordered list of router types to allocate as tenant routers.	Not applicable in NSX-T.
edge_appliance_user	Username to configure for Edge appliance login.	Not applicable in NSX-T.
metadata_initializer	Initialize metadata access infrastructure	Not applicable in NSX-T.
shared_router_appliance_size	Edge appliance size to be used for creating a shared router edge.	Not applicable in NSX-T.
use_dvs_features	Allow for directly configuring DVS backing NSX-V.	Not applicable in NSX-T.
service_insertion_profile_id	The profile ID of the redirect firewall rules that will be used for service insertion.	Feature does not exist in NSX-T integration.
service_insertion_redirect_all	Creates a firewall rule to redirect all traffic to a third-party firewall.	Feature does not exist in NSX-T integration.

Parameter	Description	Reason
use_nsx_policies	Use NSX policies for implementing Neutron security groups.	Feature does not exist in NSX-T integration.
default_policy_id	If `use_nsx_policies` is `True`, this policy will be used as the default policy for new tenants	Feature does not exist in NSX-T integration.
bind_floatingip_to_all_interfaces	Bind floating IPs to downlink interfaces when set to `True`.	In NSX-T, NAT for floating IP is always processed for east-west traffic as well.
vdr_transit_network	Network range for distributed router TLR/PLR connectivity.	In NSX-T the range for DR/SR connectivity cannot be configured from OpenStack.
exclusive_dhcp_edge	Have exclusive DHCP edge per network	Does not apply to NSX-T as DHCP is implemented on edge cluster.
bgp_neighbour_hold_down_timer	Interval for BGP neighbour hold down time.	Feature does not exist in NSX-T integration. BGP peering is configured on NSX tier-0 gateway routing configuration.
bgp_neighbour_keep_alive_timer	Interval for neighbour keep alive time.	Feature does not exist in NSX-T integration. BGP peering is configured on NSX tier-0 gateway routing configuration.
share_edges_between_tenants	Use same DHCP or router edge for multiple tenants.	Not applicable in NSX-T.
use_routers_as_lbaas_platform	Use subnet's exclusive router as a platform for LBaaS.	Not applicable in NSX-T, where LB services are always attached to rotuers used for forwarding.
nsx_sg_name_format	Format for the NSX name of an OpenStack security group.	Backend resource naming is implicit in NSX-T.
loadbalancer_pool_transparency	Create LBaaS pools in transparent mode.	Transparent mode is not supported in NSX-T.
default_edge_size	Defines the default edge size for router, DHCP, and LB edges.	Not applicable in NSX-T.

Configuring the Neutron Migrator

Before launching the neutron migrator, create a JSON file called `migrator.conf.json` to specify the NSX-T environment and the hosts that need to be migrated. This file will be mounted in the migrator pod and validated by the migration process. The following is a sample `migrator.conf.json` file:

```
{
 "strict_validation": true,
 "edge_migration": true,
 "host_migration": true,
 "edge_migration_interfaces_down": true,
 "post_migration_cleanup": true,
 "rollback": false,
 "nsxv_token_lifetime": 1440,
 "compute_clusters": [
   "domain-c17",
   "domain-c29",
```

```
    "domain-c71",
   ],
   "nsx_manager_ips": [
     "192.168.16.32",
     "192.168.16.64",
     "192.168.16.96",
   ],
   "nsx_manager_user": "admin",
   "nsx_manager_password": "<NSX password>",
   "metadata_proxy": "VIO_mdproxy",
   "dhcp_profile": "VIO_dhcp_profile",
   "default_overlay_tz": "0b3d2a91-2dfc-40a7-ac6b-fbd62b0e4c79",
   "default_vlan_tz": "b87c7a69-6d1a-4857-badd-0d0e4d4e924f",
   "default_tier0_router": "VIO_Tier0",
   "availability_zones": [
   {
     "name": "az1",
     "metadata_proxy": "VIOAZ1_mdproxy",
     "dhcp_profile": "VIOAZ1_dhcp_profile",
     "default_vlan_tz": "6320d1e3-45a1-4f37-87b4-6d35d19cafef",
     "default_tier0_router": "VIOAZ1_Tier0VRFLite"
    }
    ],
    "external_networks_map": {
     "61282e88-0abb-4036-9ea8-22418f85cdf3": "VIO_Tier0",
     "39db1d0f-4279-462b-a17e-1995a5c00ae8": "VIOAZ1_Tier0VRFLite"
   },
   "transit_network": "100.64.0.0/16"
 }
```

The configuration parameters are:

Parameter	Default Value	Description
post_migration_cleanup	True	After the migration is completed, remove additional NSX-T entities created by the migration process that are not used by VIO or duplicated by other VIO resources.
rollback	True	Automatically roll back upon Failure (if possible).
nsxv_token_lifetime	1440	Duration in minutes of the token for NSX-V access. Token is provided to NSX-T. Duration should be chosen according to the deployment size and time expected to complete the migration. Token should not expire before the migration is completed.
compute_clusters		List of vSphere compute clusters that will be migrated. This should include only the clusters where VIO VM instances are deployed. Edge clusters and VIO management clusters should not be included.
nsx_manager_ips		IP or FQDN for NSX-T manager. If a manager cluster is used, this parameter can either specify a VIP or the list of NSX manager instances. In the latter case client-side load balancing will be used when accessing NSX Manager.
nsx_manager_user	admin	User for NSX Manager access. Authentication with principal identities is not supported by VIO.

Parameter	Default Value	Description
nsx_manager_password		Password for NSX Manager access.
metadata_proxy		Identifier of the metadata proxy for the VIO default availability zone. The identifier is last segment of the resource's policy path.
dhcp_profile		Identifier of the DHCP profile for the VIO default availability zone.
default_tier0_router		Identifier of the tier-0 gateway for the VIO default availability zone. Will be used for north-south traffic by neutron routers whose gateway is the default external network.
default_overlay_tz		Overlay NSX-T transport zone to be used for the VIO deployment.
default_vlan_tz		VLAN NSX-T transport zone for the default availability zone.
transit_network	100.64.0.0/16	CIDR for the NSX-T transit network. Modify only if it was changed from NSX-T default.
external_networks_map	Empty list	
availability_zones	Empty list	

Deploying the Neutron Migrator

In the migrator bundle is a script called script `build_yaml.sh`. When the migrator configuration is ready, run the script to create the deployment specification and deploy it on the VIO control plane. For example:

```
./build_yaml.sh -t 7.1.1.1899999
```

The script accepts the following parameters:

-k	Optional. Do not include vCenter Server certificate in deployment. Specify this only when VIO is using an insecure vCenter connection.
-t <full VIO version>	Required. The VIO version must include the build number and match tag for existing VIO images.

The `build_yaml.sh` script creates <YAML-FILE-NAME> which contains all the information for deploying the neutron migration control plane.

Starting the Migration

To start the migration, run the following command:

```
kubectl apply -f <YAML-FILE-NAME>
```

This will create the neutron-migrator deployment in the openstack namespace. This deployment has a single replica. The migration pod is automatically started when the deployment's pod is created.

Migration Pod Startup

During startup the migrator pod will read the configuration file and the current status of the migration. Based on this information it will decide the next step of the migration, which could be one of the following:

- API replay
- Starting migration from NSX Manager
- VIO reconfiguration

The migration pod will terminate if the configuration file is not found or if some required parameter has not been specified.

The migration pod will also terminate with an error if the current state of the migration is inconsistent, for example, if API replay has not completed, but a migration is already in progress.

When the migrator job is started, configuration files for Neutron NSX plugins are mounted into the pod. Any change made to Neutron configuration once the migrator is started will not be process by the migrator job. You must not make changes to Neutron configuration while the migrator is running. If you need to make changes, the migrator job must be restarted.

API Replay

In this state the migration process will create all the necessary configurations on NSX-T and populate the VIO Neutron database for use with NSX-T.

At the end of this process, all logical networking entities required by VIO will be configured in NSX-T, even if workloads are still running on NSX-V.

Before implementing VIO configuration on NSX-T, the following checks are performed:

- Pre validation checks. These are the checks listed in the Prerequisites section above.
- NSX-T version check. The NSX-T version must be 3.2 or later.
- Ensure that compute manager is configured. The migration requires VIO's vCenter to be registered as a compute manager in NSX-T. This check verifies this has been done.
- No Neutron resource should be configured on NSX-T. If the rollback option is set to True the migrator process will cleanup any (likely stale) neutron resource found on NSX-T.

After the checks are completed, the migration process initializes the Neutron NSX-T database and prepares its structure. Then a temporary neutron server is started within the migrator pod. This temporary Neutron server has been configured to run with NSX-T. After the temporary neutron server is up, the migration process collects information about the network VNI mappings and port/VIF mappings.

The API migration process is then started and will migrate the following resources:

- Routers (to tier-1 gateways)
- Networks (to segments)

- Subnets (to segment subnets and segments' DHCP configuration)
- Port (to segment ports and DHCP static bindings)
- Security groups (to security policies, rules, groups, and services)
- Floating IPs (to NAT rules)
- QoS policies and rules
- FWaaS groups, policies, and rules
- Octavia load balancers, listeners, pools, members, and health monitors

After the API replay is completed, the temporary neutron server pod is shut down.

Monitor the migrator pod logs with the `tail` command. When the logs show that the migrator pod is waiting for the NSX-T migration process to be started, perform the next task (Edge Cutover).

Edge Cutover

Make the following API call to get the ID of the Edge node:

```
curl -v -s -X GET -k -u admin:<password> https://<nsx-mgr-ip>/api/v1/transport-nodes/ -H
content-type:application/json
```

Make the following API call to modify the parameter `v2t-migration-config` on all the Edge nodes:

```
curl -v -s -X PUT -k -u admin:<password> https://<nsx-mgr-ip>/api/v1/transport-nodes/<edge-
nodeid>/node/v2t-migration-config -H content-type:application/json -d '{"enabled": true}'
```

Follow the procedure in Migrating North-South Traffic to NSX-T Edges Using Edge Cutover. After this migration, the north-south traffic will be handled by NSX-T. The migration process will:

- Bring NSX-V edge appliance interfaces down.
- Enable ARP on NSX-T tier-1 downlink to ensuring seamless east-west and north-south traffic transition during migration.
- Connect to vCenter to retrieve an NSX-V authentication token.
- Prepare a mapping file for distributed routers (NSX-V DLRs).
- Set up Edge migration on NSX-T and wait for its completion.

During the north-south cutover, VMs might briefly lose connectivity as connectivity is switched from NSX-V ESGs or DLRs to NSX-T tier-1 gateways. After the north-south cutover is complete, the NSX-V and metadata Edges will be powered off. The next step is host migration.

IMPORTANT: Before starting the north-south cutover, after a rollback, make sure that the edge mapping file is present. The file is automatically deleted after a rollback. The migrator job will restore it within 10 seconds of a rollback completion. This does not apply if there are no distributed routers in the NSX-V VIO environment.

Note: The NSX-V access token is renewed at each pod execution. Its duration should be long enough to ensure that the migration is completed within the migrator pod lifecycle. If the migrator pod is restarted for any reason, a new token will be fetched.

Host Migration

Follow the procedure in Migrating Distributed Firewall Configuration, Hosts, and Workloads .

The VIO migration utility will:

- Power off all NSX-V edge appliances.
- Set up the host migration on NSX-T.
- Wait for host migration to successfully complete.

Powering off the edge appliances is necessary to ensure host migration completes successfully. Do not power on the NSX-V edge appliances during host migration.

After host migration is completed, make the following API call to reset the parameter `v2t-migration-config` for the Edge nodes. This parameter was set at the beginning of the Edge cutover step.

```
curl -v -s -X PUT -k -u admin:<password> https://<nsx-mgr-ip>/api/v1/transport-nodes/<edge-nodeid>/node/v2t-migration-config -H content-type:application/json -d '{"enabled":false}'
```

Post-Migration Cleanup

The migrator job reconfigures the Neutron CR to use NSX-T but does not remove the NSX-V configuration parameters so that you can view them for reference. These parameters are harmless. After the migration is completed you can remove them using the `viocli update neutron` command.

Logging

The neutron migrator process produces detailed logging for every phase in the process. Logs written to the pod's stdout is at the INFO level. Debug level logs are at `/var/log/migration/vio-v2t.log` on the VIO controller node where the migrator pod is running.

You can find out on which node the neutron-migrator pod is running with the following command:

```
osctl get pods neutron-migrator -o wide
```

You can then use the command `viossh` to open a shell on the controller node.

The `/var/log/migration` directory also contains the temporary neutron server log.

Rollback

Rollback can happen at various stages during migration.

If a failure occurs during the API replay stage, there is no need for an explicit rollback. The VIO neutron migrator utility will automatically remove resources that were created and then retry the migration.

If you choose to interrupt the migration by destroying the neutron migrator pod, the VIO control plane will still be functional in NSX-V. There may be NSX-T resources created by API replay. These resources will be removed.

Note that NSX-T does not allow rolling back host migration. After hosts have been migrated to NSX-T, it will not be possible to move them back to NSX-V.

If a failure occurs during host migration, you can review the logs and address the issue accordingly.

Alternatively, if a host consistently fails to migrate to NSX-T, you can remove it from the vSphere cluster and retry the migration. VMs running on the affected host will be migrated to other hosts in the cluster. After the migration, install NSX-T on the host and add it to the original vSphere cluster.

Error Codes

Code	Description
0001, 0002, 0003, 0004	Bad system state or configuration. There are major issues with the migration such as: ■ Host migration already completed, but API replay not performed. ■ VIO running already with NSX-T but API replay or migration not performed. ■ Hosts on NSX-T, VIO running with NSX-T, but API replay not performed.
0101	Unable to create configuration files for the temporary Neutron server, which needs to be up for API replay. Check the migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors. This error can usually be fixed by addressing the root cause with configuration file changes.
1001	NSX migration coordinator not running. To fix this error, start the migration coordinator service on the first node specified in migrator.conf.json. If using HA VIP, make sure the active manager instance is the one where the migration coordinator is running. For the migration, it is recommended to use a specific NSX manager, or use client-side load balancing. NSX-T manager FQDNs can be changed once the migration is completed.
1002	Invalid NSX-T version. NSX 3.2.0 or higher is required.
1003	Cannot retrieve NSX-T version. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.
1004	Failure in compute manager validation. There must be at least one compute manager defined in NSX-T. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.
1005	Must run cleanup on NSX-T. The NSX-T setup already has resources created by VIO. Ensure `rollback` is set to `True` in `migrator.conf.json`.
1006	Cannot start NSX-T migration. This is probably the result of a previous migration attempt. Roll back any migration in progress and retry.
1007	Cannot prepare NSX-T for north-south cutover. There was an error while setting up north-south cutover on NSX. This could either be an error in generating the "edge mappings" file or an error while preparing the migration plan. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.

Code	Description
1008	The migrator pod is unable to bring down interfaces on NSX-T edge apliances. This is a required step for north-south cutover. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors. To workaround this issue set `edge_migration_interfaces_down` to `False` in `migrator.conf.json` and manually ensure edge interfaces are down, or disconnected, before starting north-south cutover.
1009	Cannot migrate routers without downlinks. There are neutron router without downlinks. These cannot be migrated. If the operator believe this error is returned by mistake, it can be skipped by setting `advanced_router_validation` to `False` in `migrator.conf.json`.
1100	Invalid mode in migration plan. The NSX-T migration coordinator is already configured with a different plan. This is probably the result of a previous migration attempt. Roll back any migration in progress and retry.
1101	NSX-T Migration not acknowledged in configuration. Ensure `edge_migration` and/or `host_migration` are set to `True` in `migration.conf.json`.
1105	Cannot patch routers without gateway. The process for ensuring neutron routers without gateway can be seamlessy migrated to NSX-T failed. Check migrator job's pod logs or /var/log/migration/vio-v2t.log for errors. By setting advanced_router_validation to False this process will be skipped. It will be however up to the operator to ensure that each tier-1 gateway is connected to a tier-0 router before starting north-south cutover on NSX-T.
1106	Cannot restore routers without gateway. The process for restoring neutron routers without gateway after north-south cutover failed. Check migrator job's pod logs or /var/log/migration/vio-v2t.log for errors. By setting `advanced_router_validation` to `False` this process will be skipped. It will be however up to the operator ensuring the tier-1 gateways are disconnected from the tier-0 for neutron routers without gateways.
1110	Cannot start north-south cutover migration to NSX-T. There was an error while applying the migration plan. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.
1114	Missing VM for edge appliances. Some edge appliances do not have an associated VM appliance. Remove the corresponding neutron router so that the edge is removed.
1115	Cannot power off NSX-V edge VMs before starting host migration. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors. You can consider powering off VMs manually. This is necessary to avoid issues during host migration's runtime phase. You must power off at least DHCP and metadata proxy edge appliances.
1120	Cannot start host migration. There was an error while applying the migration plan. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for error details.
1130, 1131	Cannot complete migration. Error while setting migration as finished. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.
1132	Timeout during migration. Timeout for north-south cutover is 12 hours. Timeout for host migration is 48 hours. If the migrator's job pod is left waiting for a migration to start it will eventually timeout. Operator just need to restart it.
2001	Unable to retrieve neutron CR from VIO control plane. This could either be an authorization issue or a problem in reaching VIO's Kubernetes control plane. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.
2002	Unable to parse neutron CR. Make sure there is a `'manifests'` attribute in the `'spec'` section.
2003	Invalid contents in neutron CR. Make sure the NSX-V plugin is enabled and all the other plugins, including the NSX-T Policy plugin are disabled.

Code	Description
2004	Cannot update neutron CR. There was an error while updating Neutron CR. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors. This could either be an error in updating the Neutron CR, creating the VIOSecret instance for the NSX-T password, or creating resources for NSX managers. Verify these resources have not been left stale from some previous failed attempt.
2011	There was a failure while creating a database for NSX-T with policy. This is likely a SQL error. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.
2012	There was a failure while renaming the '`neutron_policy`' database to '`neutron`' This is likely a SQL error. Check migrator job's pod logs or `/var/log/migration/vio-v2t.log` for errors.
2111	The temporary neutron server used for API replay could not be started. This is likely a mistake in the configuration of the temporary neutron server. Check `/var/log/neutron-server-tmp.log` for errors.
2112	API replay failed. This indicates an error while creating resources in NSX-T. Check migrator job's pod logs or /var/log/migration/vio-v2t.log for errors. Logs will reveal which resource failed to create. Then check /var/log/neutron-server-tmp.log for failure details. Common failure reasons: ■ Incorrect transport zones in temporary neutron server configuration ■ Non-Openstack networks using the same VLAN as some Openstack network ■ Edge cluster running out of slots for load balancers

Troubleshooting Migration Issues 13

Information in this section might be useful in troubleshooting issues during the migration.

Import Configuration Problems

Problem	Solution
Import configuration fails.	Click **Retry** to try importing again. Only the failed import steps are retried.

Host Migration Problems

Problem	Solution
Host migration fails due to a missing compute manager configuration.	The compute manager configuration is a prerequisite for migration. However, if the compute manager configuration is removed from the NSX Manager after the migration is started, the migration coordinator retains the setting. The migration proceeds until the host migration step, which fails. Add a compute manager to NSX Manager and enter the same vCenter Server details that were used for the initial NSX-V configuration import.

Host migration fails due to stale dvFilters present.
Example error message: `Stale dvFilters present: ['port 33554463 (disconnected)', 'port 33554464 (disconnected)'] Stale dvfilters present. Aborting ]`
	Log in to the host which failed to migrate, identify the disconnected ports, and either reboot the appropriate VM or connect the disconnected ports. Then you can retry the Host Migration step.
	1 Log into the command-line interface of the host which failed to migrate.
	2 Run `summarize-dvfilter` and look for the ports reported in the error message.
	`world 1000057161 vmm0:2-vm_RHEL-srv5.6.0.9-32-local-258-963adcb8-ab56-41d6-bd9e-2d1c329e7745 vcUuid:'96 3a dc b8 ab 56 41 d6-bd 9e 2d 1c 32 9e 77 45'`
	`port 33554463 (disconnected)`
	`vNic slot 2`
	`name: nic-1000057161-eth1-vmware-sfw.2`
	`agentName: vmware-sfw`
	`state: IOChain Detached`
	`vmState: Detached`
	`failurePolicy: failClosed`
	`slowPathID: none`
	`filter source: Dynamic Filter Creation`
	3 Locate the affected VM and port.
	For example, the error message says port `33554463` is disconnected.
	a Find the section of the `summarize-dvfilter` output that corresponds to this port. The VM name is listed here. In this case, it is `2-vm_RHEL-srv5.6.0.9-32-local-258-963adcb8-ab56-41d6-bd9e-2d1c329e7745`.
	b Look for the `name` entry to determine which VM interface is disconnected. In this case, it is `eth1`. So the second interface of `2-vm_RHEL-srv5.6.0.9-32-local-258-963adcb8-ab56-41d6-bd9e-2d1c329e7745` is disconnected.
	4 Resolve the issue with this port. Do one of the following steps:
	■ Reboot the affected VM.
	■ Connect the disconnected vnic port to any network.
	5 On the **Migrate Hosts** page, click **Retry**.

Problem	Solution
After host migration using vMotion, VMs might experience traffic outage if SpoofGuard is enabled in NSX-V. Symptoms: The `vmkernel.log` file on the host at `/var/run/log/` shows a drop in traffic due to SpoofGuard. For example, the log file shows: `WARNING: swsec.throttle: SpoofGuardMatchWL:296:[nsx@6876 comp="nsx-esx" subcomp="swsec"]Filter 0x8000012 [P]DROP sgType 4 vlan 0 mac 00:50:56:84:ee:db` Cause: The logical switch and the logical switch port configuration are migrated through the migration coordinator, which migrates the SpoofGuard configuration. However, the discovered port bindings are not migrated through vMotion. Therefore, SpoofGuard drops the packets.	If SpoofGuard is enabled in NSX-V before migration, do any one of these workaround steps after vMotion of VMs: ■ Disable SpoofGuard policies. ■ Add the port IP and MAC address bindings as manual bindings. ■ If ARP snooping is enabled, wait for the VM IP addresses to be snooped by ARP. In the first two options, network traffic is restored immediately. In the third option: ■ Traffic downtime is observed until the VM sends an ARP request or reply. ■ If DHCP snooping is also enabled and the VM IP address was assigned by the DHCP server, then it will most likely be snooped as an ARP first and later as a DHCP-snooped IP address.

Problem	Solution

In the middle of a cluster migration, host migration has failed due to some hardware failure in the host.
For example, let us say that a cluster has 10 hosts, and four hosts have migrated successfully. The fifth host has a hardware failure and the host migration fails.
	If the host hardware failure cannot be fixed, skip this failed host for migration, and retry the host migration. Complete the following workaround steps:
	1 In the vCenter Server UI, remove the failed host from the inventory.
	Wait for a few minutes until the host is removed.
	2 Log in to the NSX Manager appliance where the migration coordinator service is running, and run the following API request:
	`GET https://{nsxt-policy-ip}/api/v1/migration/migration-unit-groups?component_type=HOST&sync=true`
	3 Return to the NSX-T NSX Manager UI, and refresh the browser. Observe that the failed host is no longer visible.
	4 Click **Retry** to restart the host migration.
	If you need to restart the migration coordinator service for any reason, the clusters that are already migrated to NSX-T become available for migration again on the **Migrate Hosts** page. This behavior is a known issue. In this case, the workaround is to skip the migrated clusters by doing these steps:
	1 Open an SSH session to the NSX-T NSX Manager appliance where the migration coordinator service is running.
	2 Edit the `/var/log/migration-coordinator/v2t/clusters-to-migrate.json` file to remove the clusters that are already migrated.
	For example, if the file has the following content and cluster-1 has been migrated, then remove the element {"modId":"domain-c9", "name":"cluster-1"}.
	`"clusters":[ { "modId":"domain-c9", "name":"cluster-1" }, { "modId":"domain-c19", "name":"cluster-2" } ]`
	3 Run the same API request on the NSX Manager appliance as mentioned in the earlier workaround.
	4 Return to the NSX-T NSX Manager UI, and refresh the browser. Go to the **Migrate Hosts** page, and observe that the clusters that you removed from the `clusters-to-migrate.json` file are shown as **Do not migrate**.
	5 Click **Retry** to restart the host migration.

Host migration is blocked after the recommendation is accepted because NSX-V controller VM is in power-off state.	In the host migration step, the feedback recommends that you abort the migration. If you accept the recommendation, the migration will fail. Because the Edge cutover is done, you can change the action to `skip` and continue the migration with the following steps: 1 Make the following API call and search the result for `NoNsxvControllerInRunningSate` to find the feedback request and get its ID: `GET https://$NSX_MANAGER_IP/api/v1/migration/feedack-requests?state=UNRESOLVED`

Problem	Solution

	2 Accept all recommendations by making the following API call:
	`POST https://$NSX_MANAGER_IP/api/v1/migration/feedback-response?action=accept-recommended`
	3 Provide a feedback response with the action `skip` with the following API call (note that `$FEEDBACK_ID` is the ID you obtained in step 1):
	`PUT https://$NSX_MANAGER_IP/api/v1/migration/feedback-response -d '{"response_list":[{"id": $FEEDBACK_ID, "action": "skip" }]}'`

Rolling Back a Migration

Problem	Solution
With some NSX-V OSPF deployments, if you perform a rollback after the Edge migration phase, you might see the error "Reason: NSCutover failed with '400: Configuration failed on NSX Edge VM vm-XXXX".	Re-deploy the relevant NSX-V Edge VM. After the VM is successful re-deployed, perform the rollback again.

Retrying a Migration

Problem	Solution
If a host reboots for any reason during a migration, retrying the migration fails with an error such as "The requested object : TransportNode/ 42178ba8-49fb-9545-2b78-5e9c64fddda7 could not be found. Object identifiers are case sensitive."	Perform the following steps: 1 From the VC UI, remove the host from its cluster and make it a standalone host. 2 From the NSX Manager UI, configure NSX on the standalone host using the same VDS. Make the transport node join the same overlay and VLAN transport zones that other migrated hosts join. 3 From the NSX Manager UI, go back to the migration screen, refresh it to make sure that the host is not in the cluster being migrated. Retry the migration of the cluster. 4 After the migration, add the host back to the cluster.

Removing Stale VTEP Data

Problem	Solution

If the migration is aborted after migrating Edge Services Gateways, there might be stale VTEP tables in NSX-T. If there are transport nodes in NSX-T, their tunnel status will remain down for these stale VTEPs.	To remove the stale VTEP data, make the following API call:
	`GET https://<nsx-manager-IP>/api/v1/global-configs/SwitchingGlobalConfig`
	If the parameter `global_replication_mode_enabled` in the result payload is true, take this payload, set `global_replication_mode_enabled` to false, and use the payload to make the following API call:
	`PUT https://<nsx-manager-IP>/api/v1/global-configs/SwitchingGlobalConfig`
	.

Partner Service Migration Problems

Problem	Solution
Migration coordinator does not display the feedback messages for the Service Insertion category on the **Resolve Configuration** page even though the Security Policies in your NSX-V environment contain Network Introspection rules. This problem occurs when you are migrating a combination of Guest Introspection and Network Introspection services from the same partner. If a service profile for the partner service is already created in NSX-T, migration coordinator does not initiate the migration of the Network Introspection rules.	Check whether a service profile is already created in your NSX-T environment. If yes, do these steps: 1 Roll back the migration. 2 Delete the partner service profile and service reference in NSX-T. 3 Restart the migration.

Post-Migration Issues

Problem	Solution
After a migration, and after ESGs are removed from the network, NSX-T raises alarms about OSPF neighbors being down for these ESGs. If you resolve the alarms, they are raised again.	Acknowledge the alarms but do not resolve them. This will keep the alarms from being raised again.

Preparing Layer-2 Bridging for Lift-and-Shift Migration

14

A lift-and-shift migration may require a layer-2 bridge for overlay logical switches.

A layer-2 bridge is required if overlay is used in the environment and you want to minimize downtime. A layer-2 bridge is not required if VLAN is used in the environment. A layer-2 bridge is also not required if you plan to migrate all the workload VMs in a single maintenance window and you can accept downtime during the migration.

If a bridge is required, perform the following tasks.

Read the following topics next:

Extending Layer 2 Networks with NSX-T Edge Bridge

You can deploy an NSX-T Edge on a host that is prepared for NSX-V, and configure it as a bridge to extend a Layer 2 network between both network environments.

This approach allows the Edge bridge to take advantage of decapsulation of VXLAN frames that happens on the NSX-V prepared host. The bridge performs only GENEVE encapsulation. This bridging approach requires an NSX-T Edge in a virtual machine form factor.

Overview of Edge Bridging in NSX-T

You can configure an Edge bridge on an overlay segment in NSX-T to extend it logically to a VXLAN Logical Switch in NSX-V (local or universal).

After the Layer 2 bridge is set up and connectivity is established on either side of the bridge, you can use vSphere vMotion to migrate the workload VMs from NSX-V to NSX-T with a minimum downtime.

Logical View of Bridging Using NSX-T Edge

The following diagram shows the logical view of a bridged network. The left side of the diagram shows an NSX-V environment, and the right side shows an NSX-T environment.

Figure 14-1. Layer 2 Bridge Created Using NSX-T Edge

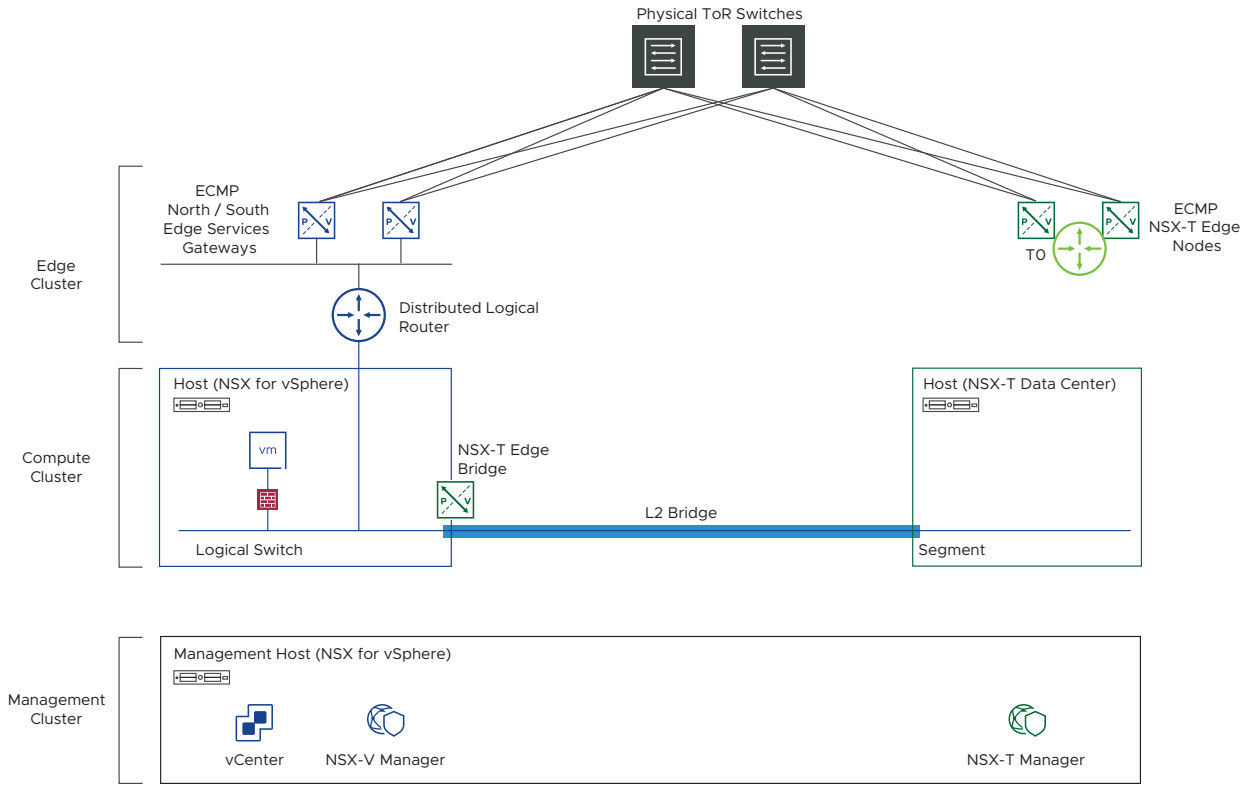

In this logical diagram, the bridged environment uses a shared Management cluster where the vCenter Server is shared between NSX-V and NSX-T. If needed, you can use a separate Management cluster for both of them. The Compute clusters must be separate, whereas the Edge cluster might be shared or different.

A recommended practice is to use one NSX-T Edge node to extend one NSX-V Logical Switch. This Edge node must be a virtual appliance and not a bare metal Edge because the NSX-T Edge must be deployed on a compute host that is prepared for NSX-V.

The Edge node extends the VXLAN Logical Switch to the GENEVE overlay segment. The Edge Services Gateways in the NSX-V environment serve as the default gateway for all north-south traffic from the workload VMs on the Logical Switch.

The preceding diagram shows a single-tier topology in the NSX-T environment. In a single-tier topology, the overlay segment in NSX-T must be initially disconnected from the tier-0 gateway. If you are using a two-tier topology in your NSX-T environment, the NSX-T overlay segment must be initially disconnected from the tier-1 gateway. When the Edge bridge is configured, and connectivity is tested across the bridge, you are ready to migrate the workload VMs by using the bridge and vSphere vMotion. Now, you can disconnect the NSX-V Logical Switch from the Distributed Logical Router (DLR), and connect the NSX-T overlay segment to the tier-0 or tier-1 gateway. This action switches the default gateway for north-south traffic to the tier-0 or tier-1 gateway in the NSX-T environment. After all the workload VMs are migrated from NSX-V to NSX-T, you can optionally remove the bridge.

Figure 14-2. Default Gateway Switched to NSX-T

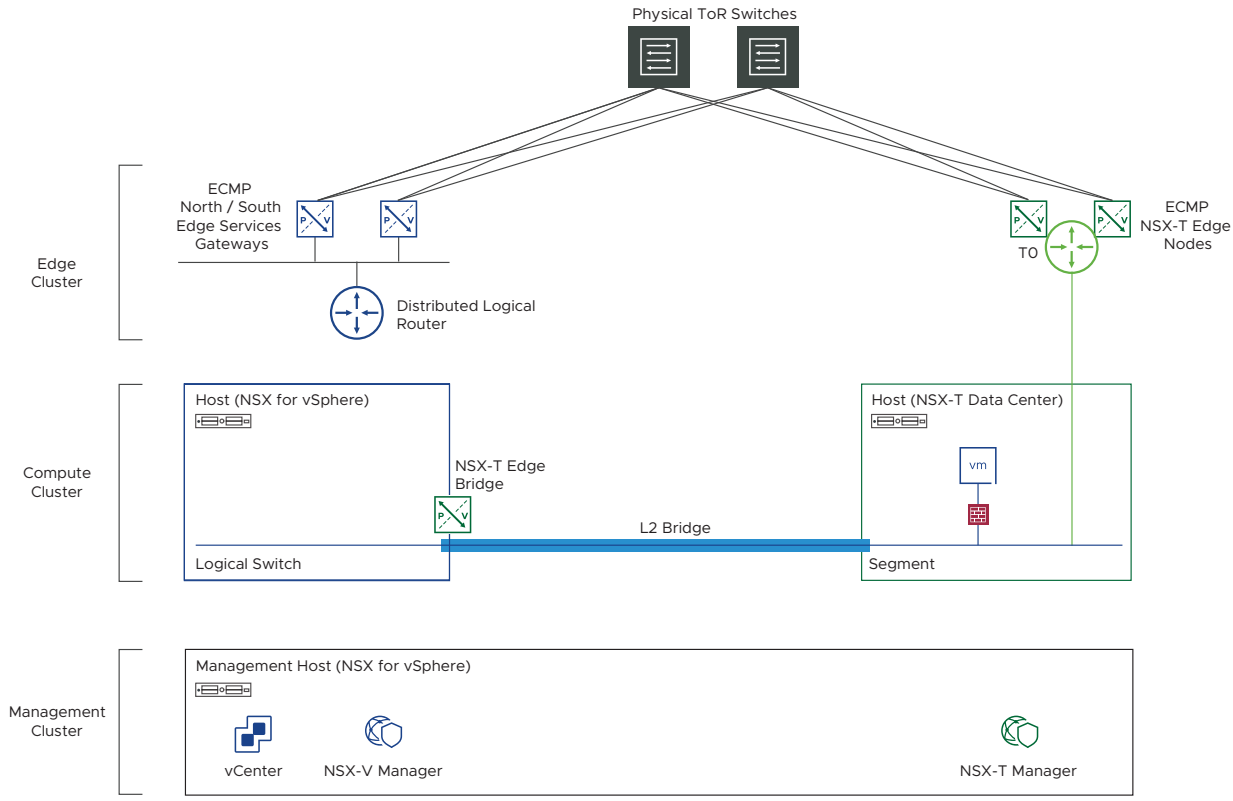

Note If you want to move workloads from a VLAN Distributed Virtual Port Group in NSX-V to a VLAN segment in NSX-T, you do not require a bridge. You must only ensure that the VLAN ID of the Distributed Virtual Port Group and the VLAN segment are the same.

Prepare the NSX-T Environment to Bridge Layer 2 Networks

Review your existing NSX-V environment and plan your bridging requirements.

Prepare a plan to determine the infrastructure that you will require for your bridging needs. For example:

- How many Logical Switches do you have in your existing NSX-V environment?
- How many Logical Switches you want to extend with NSX-T overlay segments?
- Do you plan to extend all Logical Switches in a single batch or one network at a time?
- Do you plan to configure High Availability (HA) on the NSX-T Edge nodes?
- Do you have enough capacity on the NSX-V prepared host to deploy the number and size of NSX-T Edge nodes you require for bridging?

Deploy a New NSX-T Environment

Create a new NSX-T environment on a separate hardware with a topology of your choice. This new NSX-T environment coexists while you are still using the existing NSX-V environment.

NSX-T environment must use its own dedicated Compute cluster. It can either use separate Edge and Management clusters or share them with the existing NSX-V environment.

The following procedure outlines the high-level workflow to prepare a new NSX-T environment. For detailed steps about deploying the NSX-T infrastructure, see the *NSX-T Data Center Installation Guide*.

Procedure

1. Deploy NSX Manager appliances.

   In a production environment, add an NSX Manager cluster with three appliances. However, for migrating workloads using a bridge and vSphere vMotion, a single NSX Manager appliance is adequate.

2. Deploy a vCenter Server appliance.

   The vCenter Server must be added as a compute manager in NSX-T. You can share the vCenter Server that is used in NSX-V or deploy another one in NSX-T.

3. Deploy an appropriate number and size of NSX Edge nodes for north-south routing based on the requirements of your topology.
4. Join the NSX Edge nodes to the management network.
5. Configure NSX-T on the NSX Edge nodes.
6. Set up one or more Compute clusters based on your requirements.
7. Install NSX-T on the Compute clusters. The hosts in the cluster now become host transport nodes.

Create the NSX-T Topology

Add the networking components to create the logical network topology in the NSX-T. You can create the same logical topology as your existing NSX-V or create a new topology, if necessary.

You must also pre-configure the networking services that are required for your applications to run before the VMs are moved to the new NSX-T.

The following procedure outlines the workflow for creating the NSX-T logical topology. For a detailed information about creating and configuring the networking objects, see the *NSX-T Data Center Administration Guide*. If you plan to create the topology using APIs, see the *NSX-T Data Center API Guide* for more information.

Procedure

1. Add tier-0 and tier-1 gateways depending on the requirements of your NSX-T network topology.
2. Add NSX-T overlay segments with the same subnet address as the Logical Switches in NSX-V. Similarly, add NSX-T VLAN segments with the same subnet address as the Distributed Virtual Port Group (DVPG) VLANs in NSX-V.

   The same subnet address helps in ensuring that the IP addresses of the workload VMs are retained after the VMs move to NSX-T segments.

   You must create the segments with the `SOURCE` replication mode, and change the mode to `MTEP` only after the migration is done.
3. To migrate Distributed Firewall configuration from your NSX-V environment, ensure that the following requirements are met:
   - The overlay segments in NSX-T must have the same virtual network identifier (VNI) as the Logical Switches in NSX-V. You must use the NSX-T APIs to create the overlay segments. You cannot create overlay segments with the same VNI in the NSX Manager UI.
   - The VLAN segments in NSX-T must have the same VLAN IDs as the VLAN Distributed Virtual Port Groups in NSX-V.

     Note VLAN Distributed Virtual Port Group must be associated only with a VLAN ID. VLAN Trunk is not supported.

4 If Layer 3 services such as Network Address Translation, Load Balancing, VPN, and so on, are configured on your NSX-V Edge Services Gateway, configure equivalent services on the tier-1 or tier-0 gateway of your NSX-T environment. Make sure that both steps 4 and 5 are done.

If Layer 3 services are not configured, skip steps 4 and 5 and proceed directly to step 6.

Caution Be careful not to enable route advertisement and Layer 3 services on the tier-1 gateway while the north-south traffic is being routed through the Edge Services Gateway. It can conflict with the NSX-V environment. Also, remember that your workload VMs are not yet moved to NSX-T. The best time to enable route advertisement and Layer 3 network services is when you are ready to switch the default gateway for north-south traffic to the NSX-T side.

a In NSX Manager, navigate to **Networking > Tier-1 Gateways**.

b Click the vertical ellipses next to the tier-1 gateway, and then click **Edit**.

c Expand the **Route Advertisement** section, and turn off all the toggle buttons for the L3 services.

For example:

5 Connect the uplink interface of the tier-0 gateway to a transit VLAN segment.

Optionally, configure dynamic route peering between tier-0 gateway and the north-facing physical routers. If you configure dynamic routing, ensure that **Route Redistribution Status** is turned off on the tier-0 gateway so that no subnets are advertised in the NSX-T environment. You must enable **Route Redistribution Status** when you are ready to switch the default gateway to the NSX-T side for routing the north-south traffic.

a In NSX Manager, navigate to **Networking > Tier-0 Gateways**.

b Click the vertical ellipses next to the tier-0 gateway, and then click **Edit**.

c Expand the **Route Re-Distribution** section, and turn-off the **Route Re-distribution Status** toggle button.

6 Attach the overlay segments to the downlinks of the tier-0 or tier-1 gateway.

Turn off **Connectivity** on the segment while the north-south traffic is being routed through the Edge Services Gateway in your NSX-V environment. Turn on the segment connectivity only when you are ready to switch the default gateway to the NSX-T side for routing the north-south traffic.

a In NSX Manager, navigate to **Networking > Segments**.

b Click the vertical ellipses next to the segment, and then click **Edit**.

c Turn off the **Connectivity** option to disconnect the segment from the network topology.

Change the MAC Address of NSX-T Virtual Distributed Router

You must change the default MAC address of the NSX-T virtual distributed router so that it does not use the same MAC address that is used by the Distributed Logical Router (DLR) of NSX-V.

The virtual distributed routers (VDR) in all the transport nodes of an NSX-T environment use the default global MAC address. You can change the global MAC address of the NSX-T VDR by updating the global gateway configuration with the following PUT API:

```
PUT https://{policy-manager}/policy/api/v1/infra/global-config
```

For a detailed explanation about all the parameters in this global gateway configuration API, see the *NSX-T Data Center API Guide*.

Caution While changing the MAC address of NSX-T Virtual Distributed Router, if a VM is attached to an NSX-T overlay segment, you might observe a short disruption in the data plane.

The default global MAC address is: **02:50:56:56:44:52**. For example, you want to change it to **02:50:56:56:44:62**.

Procedure

1. Retrieve the output of the global gateway configuration API by running the following GET API:

   ```
   GET https://{policy-manager}/policy/api/v1/infra/global-config
   ```

2. Copy and paste the GET API response in a text editor.
3. Edit the following two parameter values:
   - `vdr_mac`: **02:50:56:56:44:62**
   - `allow_changing_vdr_mac_in_use`: true

   Retain the existing values of the other parameters.
4. Copy and paste the complete edited API response in the request body of the following PUT API:

   ```
   PUT https://{policy-manager}/policy/api/v1/infra/global-config
   ```

5. Run the PUT API to change the default global MAC address of the NSX-T VDR.

Deploy NSX Edge Nodes for Bridging

To extend a Layer 2 network from NSX-V to NSX-T, create an Edge cluster for bridging. Add one Edge node in this Edge cluster if Edge HA is not required, or two Edge nodes if Edge HA is required.

You must deploy an NSX-T Edge VM on a host that is prepared for NSX-V.

The Edge bridge cluster is independent of the Edge cluster that you use for routing and north-south connectivity with the physical routers. A recommended practice is to deploy one NSX Edge node to bridge one L2 network. The number of Edge nodes required in the Edge bridge cluster depends on the number of L2 networks you want to extend from NSX-V to NSX-T, and whether HA is configured on the Edge bridge nodes.

Caution Remember that the NSX-T Edge VM, which is used as a bridge, is deployed on an NSX-V prepared host. This Edge VM might potentially hit a Distributed Firewall deny rule in NSX-V. To avoid this issue, add the NSX-T Edge VM to the NSX-V Distributed Firewall Exclusion List, or make sure that the NSX-V Distributed Firewall rule allows communication to the NSX-T Edge VM.

Prepare a plan to determine the infrastructure required for your bridging requirements. For example:

- How many L2 networks you want to bridge?
- Do you plan to extend all networks in a single batch or one network at a time?
- Do you plan to configure High Availability on the Edge nodes for bridging?
- Do you have enough capacity on the NSX-V prepared host to deploy the number of Edge nodes you require for bridging?

For detailed instructions on deploying an NSX Edge node in NSX-T by using an OVA file, see Deploy NSX Edge Nodes.

In step 10 of this hyperlinked topic, select the networks for the Edge interfaces, as explained in the following example.

Example

Consider that your goal is to extend a Logical Switch named **`Vwire-1`** in NSX-V with an overlay segment named **`Segment-1`** in NSX-T.

To do this bridging, deploy a single NSX Edge node using an OVA file or from the NSX Manager user interface and name this node as **`EN1`**.

NSX-V configuration is as follows:

- Vwire-1 is connected to vSphere Distributed Switch **`VDS-1`**.
- The virtual wire port group on VDS-1 for Vwire-1 is **`vxw-dvs-36-virtualwire-1-sid-10600-Vwire-1`**.

During deployment of Edge node (EN1), select the VDS port groups (networks) for the Edge interfaces as follows.

Edge Interface	Source Network	Destination Network
vNIC1	Network 0	dvPG-Management
vNIC2	Network 1	dvPG-Trunk

Edge Interface	Source Network	Destination Network
vNIC3	Network 2	vxw-dvs-36-virtualwire-1-sid-10600-Vwire-1
vNIC4	Network 3	Not used for this bridging example

Figure below shows the logical view of the vNIC configuration on the NSX-T Edge VM.

Figure 14-3. Logical View of vNIC Configuration on NSX-T Edge VM

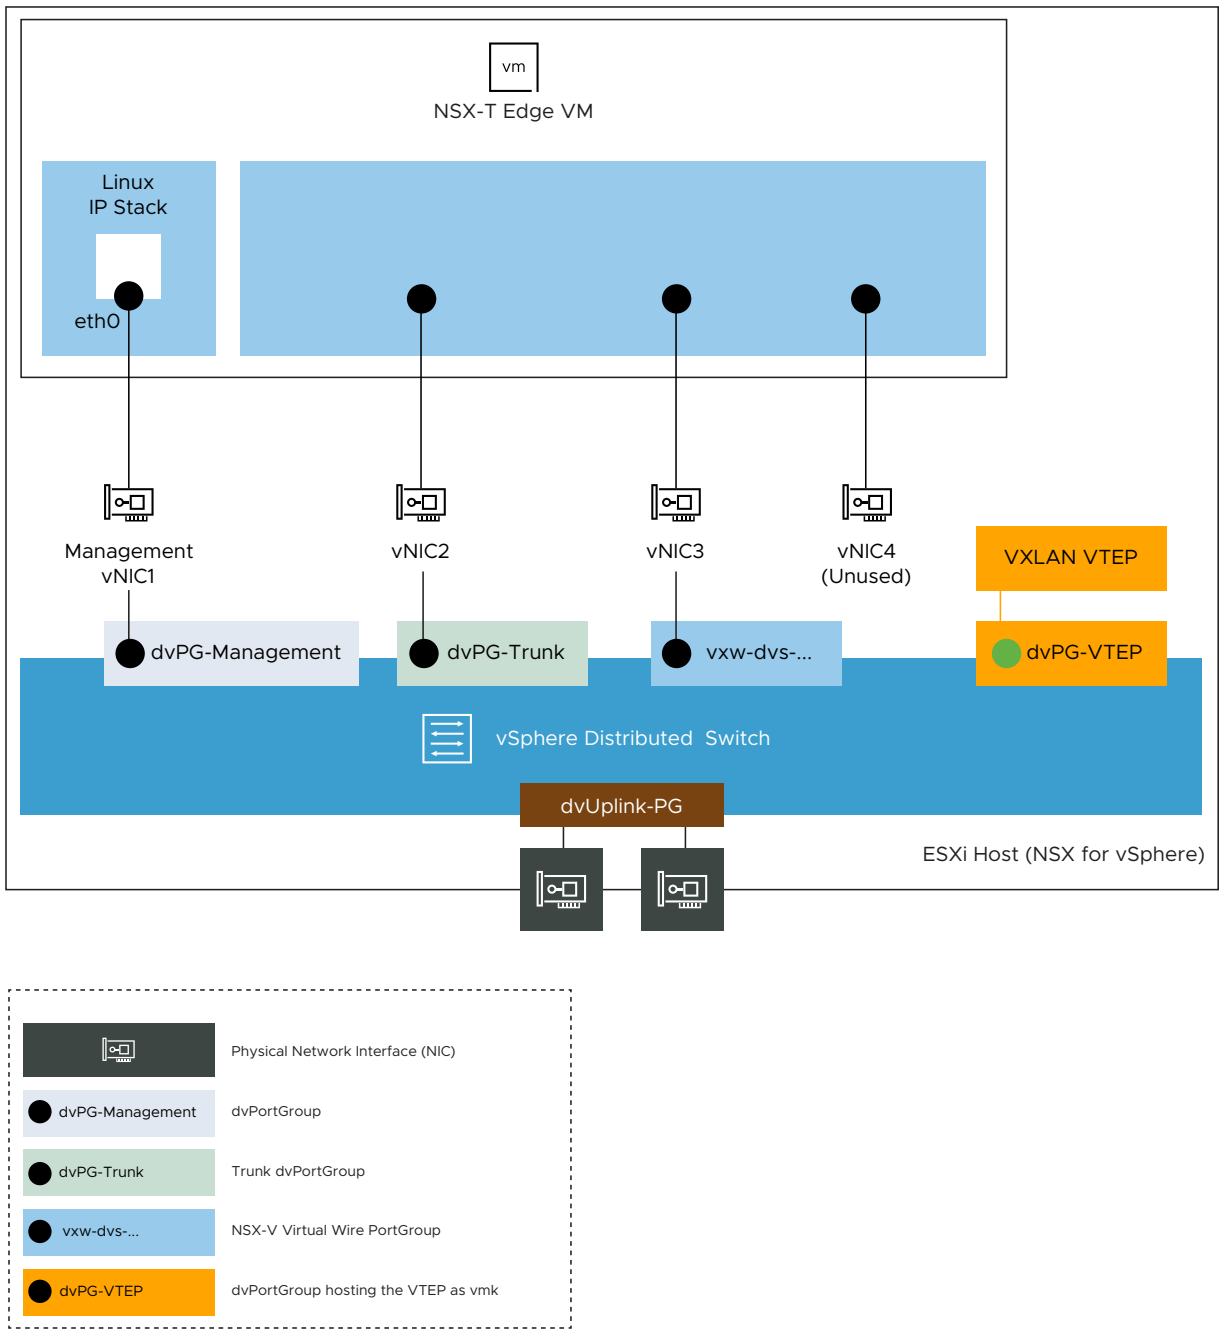

Configure an NSX Edge Bridge as a Transport Node

Add the N-VDS switches on the NSX-T Edge VM that you have deployed on the NSX-V prepared host.

Logical View of N-VDS Switches on the NSX-T Edge VM

Following figure shows a logical view of the N-VDS switches on an NSX-T Edge VM (EN1) that is deployed on an NSX-V prepared host.

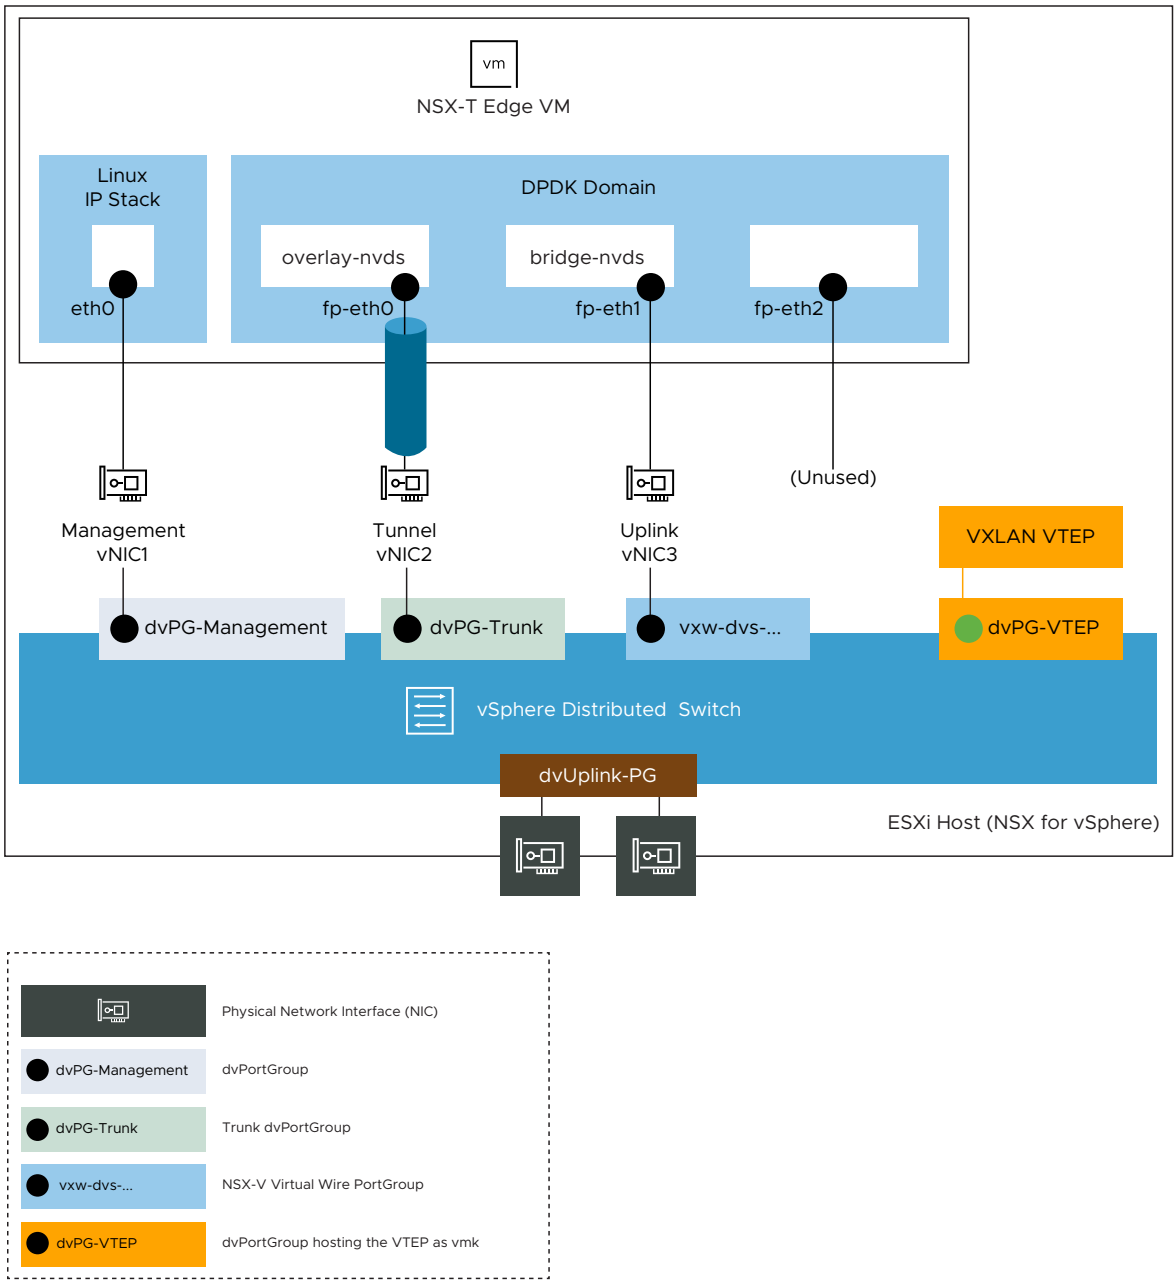

As this topic is focused on configuring the Edge VM for bridging, the port groups for vMotion, Storage, VMkernel interfaces, and so on, are not shown.

NSX-T Edge VM (EN1) is a transport node in two transport zones: Overlay-TZ and VLAN-TZ. This Edge VM is used for bridging and it has four vNICs. However, for this bridging example, three vNICs are used:

- vNIC1 is dedicated to the management traffic.
- vNIC2 is the uplink of the N-VDS switch named **overlay-nvds**. This Edge switch is attached to the overlay transport zone and used for tunneling overlay traffic.
- vNIC3 is the uplink of the N-VDS switch named **bridge-nvds**. This Edge switch is attached to the VLAN transport zone. vNIC3 is directly connected to the virtual wire port group that you want to bridge.

  In this example, the virtual wire port group on the vSphere Distributed Switch **VDS-1** is **vxw-dvs-36-virtualwire-1-sid-10600-Vwire-1**.

An NSX-T Edge VM has four internal interfaces: eth0, fp-eth0, fp-eth1, and fp-eth2. Eth0 is reserved for management, while the other interfaces are assigned to Data Plan Development Kit (DPDK) Fast Path. The fp-eth interfaces are connected to physical ToR switches for north-south connectivity and to provide NSX-T overlay tunneling. You have complete flexibility in assigning the Fast Path interfaces (fp-eth) for overlay connectivity or external connectivity.

In this example, because the NSX-T Edge VM is used for L2 bridging, fp-eth1 interface is not connected to the physical ToR switch, but it is connected to the virtual wire port group on VDS-1 for the **Vwire-1** Logical Switch. The fp-eth interfaces are configured as follows:

- fp-eth0 is assigned for overlay traffic.
- fp-eth1 is assigned for external traffic (connected to the virtual wire port group for **Vwire-1** Logical Switch)
- fp-eth2 is unused.

Because the NSX-T Edge VM (EN1) is deployed on an NSX-V prepared host, the Edge node can take advantage of the following:

- NSX-T Edge node that serves as the bridge has its own Tunnel Endpoint (TEP) and does not have to be on an NSX-T prepared host.
- The VXLAN Tunnel Endpoint (VTEP) on an NSX-V prepared host decapsulates the VXLAN frames before the frames reach the Edge VM. That is, the frames that reach the NSX-T Edge VM are not encapsulated with VXLAN.

To configure the Edge Tunnel End Point (TEP) in the overlay-nvds switch configuration, both static IP list and IP pool are supported. In this bridging example, an IP pool is used.

Prerequisites

Create an IP pool, for example **Edge_TEP_Pool**.

For detailed instructions, see Create an IP Pool for Edge Tunnel End Points.

Procedure

1 For this bridging example, create two uplink profiles that define how the two N-VDS switches on the NSX-T Edge VM (EN1) connect to the physical network.

a From a browser, log in with **admin** privileges to an NSX Manager in your NSX-T environment at `https://nsx-manager-ip-address`.

b Click **System > Fabric > Profiles > Add**.

c Specify the properties of the uplink profile to use for the bridge-nvds.

Example: uplink profile for bridge-nvds

Option	Description
Name	nsx-edge-nic-bridge-uplink-profile
Transport VLAN	0
MTU	1600
Teaming Policy	Failover Order (Default teaming)

d On similar lines, add another uplink profile to use for the overlay-nvds and specify its properties.

Example: uplink profile for overlay-nvds

Option	Description
Name	nsx-edge-nic-overlay-uplink-profile
Transport VLAN	*Edge TEP VLAN* Replace *Edge TEP VLAN* with a VLAN ID that is preferably different from the VXLAN VTEP VLAN.
MTU	1600
Teaming Policy	Failover Order (Default teaming)

2 Add the N-VDS switches on the NSX-T Edge VM for overlay transport zone and VLAN transport zone.

a Go to the NSX Manager in your NSX-T environment.

b Click **System > Fabric > Nodes > Edge Transport Nodes**.

c Click the Edge transport node, and then click **Edit**.

For example, click **EN1**.

d Click **Add Switch** and define the N-VDS switch properties to attach the Edge VM to the overlay transport zone.

Example:

Option	Description
Edge Switch Name	overlay-nvds
Transport Zone	Overlay-TZ
Uplink Profile	nsx-edge-nic-overlay-uplink-profile
IP Assignment (TEP)	Use IP Pool
IP Pool	Edge_TEP_Pool
Uplink	fp-eth0

If you want to configure static IP addresses for Edge TEP instead of using an IP pool, in the **IP Assignment (TEP)** drop-down menu, select **Use Static IP List**.

e Again click **Add Switch** and define the N-VDS switch properties to attach the Edge VM to the VLAN transport zone.

Example:

Option	Description
Edge Switch Name	bridge-nvds
Transport Zone	VLAN-TZ
Uplink Profile	nsx-edge-nic-bridge-uplink-profile
Uplink	fp-eth1

Although the bridge-nvds switch is attached to the VLAN-TZ, the Edge internally uses the fp-eth1 interface to connect directly to the VXLAN Logical Switch (Vwire-1) in NSX-V.

Results

The NSX-T Edge VM that you configured for bridging is now an Edge transport node in your NSX-T environment.

Create an NSX Edge Cluster

After the Edge bridge node is configured to become an NSX-T Edge transport node, you can add the node to an NSX Edge cluster. This Edge bridge cluster is independent of the Edge cluster that you use for routing and north-south connectivity with the physical routers.

For example, let us create an NSX Edge cluster named **`Cluster1`** and add one NSX Edge transport node **`EN1`** to this cluster.

Prerequisites

- At least one NSX Edge node is deployed.

- NSX Edge is joined to the management plane.
- NSX Edge is configured to become an NSX-T Edge transport node.

Procedure

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

2 Select **System > Fabric > Nodes > Edge Clusters > Add**.

3 Enter the NSX Edge cluster name.

For example, name the cluster as **`Cluster1`**.

4 (Optional) Select an NSX Edge cluster profile.

The default NSX Edge High Availability (HA) profile is preselected. In this example, Edge HA is not required.

5 In the **Member Type** drop-down menu, ensure that **Edge Node** is selected.

6 From the **Available** list, select NSX Edges, and click the right-arrow to move them to the **Selected** list.

Select one Edge bridge VM when Edge HA is not required or two Edge VMs when Edge HA is required.

In this example, move **`EN1`** to the **Selected** list.

Create an Edge Bridge Profile

An Edge bridge profile makes an NSX Edge cluster capable of providing Layer 2 bridging to an overlay segment.

When you create an Edge bridge profile, you can specify one of the following failover modes:

Preemptive

When a failover occurs, the backup node becomes the active node. After the failed node recovers, it becomes the active node again.

Non-Preemptive

When a failover occurs, the standby node becomes the active node. After the failed node recovers, it becomes the standby node.

You can manually set the standby edge node to be the active node by running the `set l2bridge-port <uuid> state active` CLI command on the standby edge node.

You can run this command only in the non-preemptive mode. Otherwise, an error occurs. In the non-preemptive mode, the command triggers an HA failover when applied on a standby node, and it is ignored when applied on an active node. For more information, see the *NSX-T Data Center Command-Line Interface Reference*.

Prerequisites

An NSX Edge cluster with NSX Edge transport nodes was created. For example, in the previous topic, you created an Edge cluster named **`Cluster1`** and added the NSX Edge transport node **`EN1`** to this cluster.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Select **Networking** > **Segments** > **Edge Bridge Profiles**.
3. Click **Add Edge Bridge Profile**.
4. Enter a name for the Edge bridge profile and optionally a description.

   For example, name the profile as **`Edge_Bridge_Profile1`**.
5. Select an NSX Edge cluster.

   For example, select **Cluster1**.
6. Select a primary node.

   For example, select **EN1**.
7. (Optional) Select a backup node.

   In this example, backup node is not selected. However, if you plan to enable Edge HA, you can deploy another NSX Edge node, say **`EN2`**, on a different NSX-V host. Add EN2 to Cluster1, and select EN2 as the backup node in the Edge bridge profile. When a backup node is specified, change the default failover mode, if necessary. The default failover mode is Preemptive.
8. Click **Save**.

What to do next

You can now configure an Edge bridge on an NSX-T overlay segment by using this Edge bridge profile.

Configure an Edge Bridge on an Overlay Segment

Use the Edge bridge profile that you created earlier to configure bridge settings on the overlay segment. The Edge bridge helps you to extend the overlay segment with the Logical Switch in your NSX-V environment.

The NSX-V prepared host on which you have deployed the NSX-T Edge bridge node serves as the transport node. This host has workload VMs that require connectivity with the overlay segment in your NSX-T environment.

Prerequisites

- An Edge bridge profile is added. For example, **`Edge_Bridge_Profile1`**.
- A segment is added in the overlay transport zone. For example, **`Segment-1`** is added in the overlay transport zone **`Overlay-TZ`**.
- A Logical Switch is available in the NSX-V environment to extend Layer 2 connectivity with this NSX-T overlay segment.

  For example, you are extending **`LogicalSwitch-1`** with **`Segment-1`**.

Procedure

1. From your browser, log in with admin privileges to an NSX Manager at https://*nsx-manager-ip-address*.
2. Select **Networking** > **Segments**.
3. Click the vertical ellipses next to the overlay segment that you want to configure Layer 2 bridging on and select **Edit**.

   For example, edit the settings of **`Segment-1`**.
4. Expand **Additional Settings**. Next to **Edge Bridges**, click **Set**.
5. Click **Add Edge Bridge**.
6. Select the Edge bridge profile that you created earlier.

   For example, select **Edge_Bridge_Profile1**.
7. Select the VLAN transport zone.

   For example, select **VLAN-TZ**.
8. Enter VLAN ID as **0**.

   The untagged VLAN 0 is used so that the packets coming into and out of the Edge bridge node are not tagged.
9. (Optional) Select a teaming policy.
10. Click **Add**.

Configure the Logical Switch to Connect to the Edge Bridge

Configure some additional settings on the virtual wire port group of the NSX-V Logical Switch to enable connectivity with the NSX-T Edge bridge.

This additional configuration is required only on the distributed port group of the NSX-V Logical Switch that connects to the NSX-T Edge bridge. If an NSX-V transport zone spans multiple vSphere Distributed Switches, each Logical Switch creates one virtual wire port group per Logical Switch. The other distributed port groups on the vSphere Distributed Switch (VDS) do not require this configuration.

For example, the virtual wire port group on VDS-1 for Vwire-1 Logical Switch is **vxw-dvs-36-virtualwire-1-sid-10600-Vwire-1**. This virtual wire port group connects to the Edge bridge node **EN1**.

There are two methods to enable this connectivity:

Method 1: Enable Promiscuous Mode and Forged Transmit

Enable these two configuration settings on the distributed port group of the Logical Switch where the NSX-T Edge bridge node is connected. The drawback of enabling promiscuous mode is that all the VMs on the Logical Switch can access the packets even if a single VM receives the packet. Therefore, enabling promiscuous mode might impact network performance.

Method 2: Enable MAC Learning and Forged Transmit

MAC Learning is more efficient as compared to promiscuous mode. MAC Learning is a native feature in vSphere Distributed Switch. This feature is available starting in vSphere 6.7, and it is supported in vSphere Distributed Switch 6.6.0 or later. However, you can enable MAC Learning only with the vSphere API, and you must be familiar with scripting to enable this feature on the port group.

See an example Python script in the NSX Tech Zone article to enable MAC Learning and Forged Transmit on a port group. After you have enabled MAC Learning, you can verify the `macLearningPolicy` settings in the vCenter Managed Object Browser (MOB) at `http://{vCenter-IP-Address}/mob`.

As method 2 requires technical knowledge of scripting, you can use the simpler method 1 to enable the configuration settings on the virtual wire port group of the Logical Switch.

Procedure

1 Enable Promiscuous Mode and Forged Transmit on the distributed port group.

   a In the vSphere Client, navigate to **Hosts and Clusters**, and from the left Navigator view, click **Networking**.

   b Under **VDS-1**, right-click the virtual wire port group that is connected to the NSX-T Edge bridge node, and click **Edit Settings**.

      For example, right-click the **vxw-dvs-36-virtualwire-1-sid-10600-Vwire-1** virtual wire port group on VDS-1.

   c Click **Security**.

   d In the **Promiscuous mode** drop-down menu, select **Accept**.

   e In the **Forged Transmits** drop-down menu, select **Accept**.

   f Click **OK**.

2 If you have enabled MAC Learning using the Python script, verify whether the `macLearningPolicy` property is enabled on the distributed port group.

a In the vSphere Client, check the dvportgroupId of the virtual wire port group.

To obtain the dvportgroupId, navigate to **Hosts and Clusters**, and from the left Navigator view, click **Networking**. Click the virtual wire port group in the Navigator view. Retrieve the dvportgroupId from the URI path in the browser. You can see something like this in the URI: `DistributedVirtualPortgroup:dvportgroup-id`

For example, click the **vxw-dvs-36-virtualwire-1-sid-10600-Vwire-1** virtual wire port group on VDS-1. The dvportgroupId is `dvportgroup-72`. The dvportgroupId might be different in your environment.

b Log in to the vCenter Server MOB and go to the following URL to view the configuration properties of the dvportgroup:

```
https://{vcenter-ip}/mob/?
moid={dvportgroupId}&doPath=config%2edefaultPortConfig
```

Make sure to replace *vcenter-ip* and *dvportgroupId* with the actual values, as applicable in your environment.

The properties of the `VMwareDVSPortSetting` object type are displayed.

c Click **macManagementPolicy** from the Value column.

For example:

Data Object Type: **VMwareDVSPortSetting**
Parent Managed Object ID: **dvportgroup-72**
Property Path: **config.defaultPortConfig**

Properties

NAME	TYPE	VALUE
blocked	BoolPolicy	blocked
filterPolicy	DvsFilterPolicy	filterPolicy
inShapingPolicy	DVSTrafficShapingPolicy	inShapingPolicy
ipfixEnabled	BoolPolicy	ipfixEnabled
lacpPolicy	VMwareUplinkLacpPolicy	lacpPolicy
macManagementPolicy	DVSMacManagementPolicy	macManagementPolicy
networkResourcePoolKey	StringPolicy	networkResourcePoolKey
outShapingPolicy	DVSTrafficShapingPolicy	outShapingPolicy
qosTag	IntPolicy	qosTag
securityPolicy	DVSSecurityPolicy	securityPolicy
txUplink	BoolPolicy	txUplink
uplinkTeamingPolicy	VmwareUplinkPortTeamingPolicy	uplinkTeamingPolicy
vendorSpecificConfig	DVSVendorSpecificConfig	vendorSpecificConfig
vlan	VmwareDistributedVirtualSwitchVlanSpec	vlan
vmDirectPathGen2Allowed	BoolPolicy	vmDirectPathGen2Allowed

The properties of the `DVSMacManagementPolicy` object type are displayed.

d Verify that the `forgedTransmits` property is set to `true`.

For example:

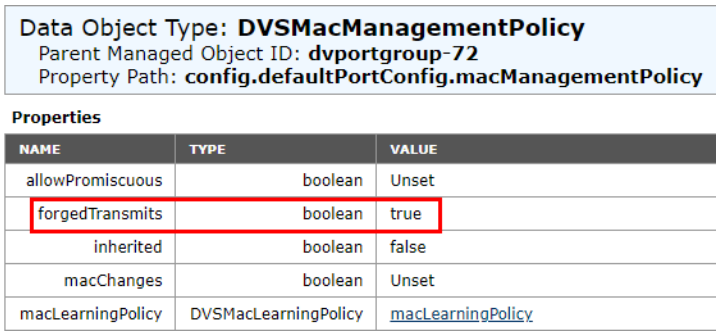

Data Object Type: **DVSMacManagementPolicy**
Parent Managed Object ID: **dvportgroup-72**
Property Path: **config.defaultPortConfig.macManagementPolicy**

Properties

NAME	TYPE	VALUE
allowPromiscuous	boolean	Unset
forgedTransmits	boolean	true
inherited	boolean	false
macChanges	boolean	Unset
macLearningPolicy	DVSMacLearningPolicy	macLearningPolicy

e Again click **macManagementPolicy** from the Value column.

The properties of the `DVSMacLearningPolicy` object type are displayed.

f Verify that `macLearningPolicy` is configured as follows:

- `enabled`: true
- `limit`: 4096
- `limitPolicy`: Drop

For example:

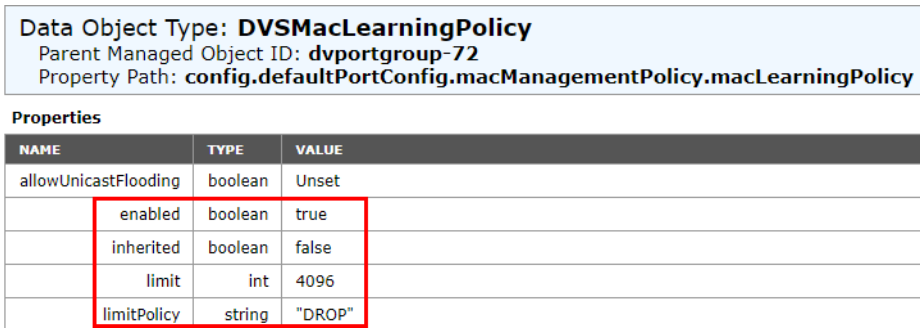

Data Object Type: **DVSMacLearningPolicy**
Parent Managed Object ID: **dvportgroup-72**
Property Path: **config.defaultPortConfig.macManagementPolicy.macLearningPolicy**

Properties

NAME	TYPE	VALUE
allowUnicastFlooding	boolean	Unset
enabled	boolean	true
inherited	boolean	false
limit	int	4096
limitPolicy	string	"DROP"

Note If the value of `limit` and `limitPolicy` is `Unset`, MAC Learning does not work even when `macLearningPolicy` is enabled.

Bridging a Federated Segment for the Migration

You can create a bridge on Federated segments for the migration.

Prerequisites

- Configure an NSX Edge bridge as a transport node. See Configure an NSX Edge Bridge as a Transport Node .
- Create an NSX Edge cluster. See Create an NSX Edge Cluster.

- Configure the NSX Manager user interface settings and set **Toggle Visibility** to **Visible to All Users**. See "Configure the User Interface Settings" in the *NSX-T Data Center Administration Guide*.

Procedure

1 Log in to NSX Manager with admin privileges.

2 Click the **Manager** button in the upper-right corner to select Manager mode.

3 Create an Edge bridge profile by navigating to **Networking > Logical Switches > Edge Bridge Profiles**.

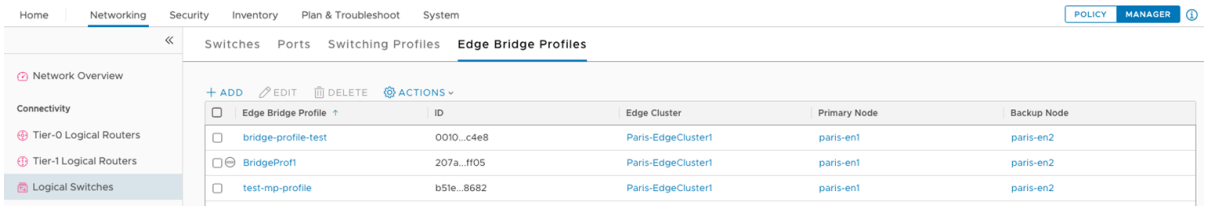

4 Make the following API call to tie the logical switch created from the GM (Global Manager) segment. For example:

```
POST https://<nsx-manager>/api/v1/bridge-endpoints
{
    "bridge_endpoint_profile_id": "0010b881-0c1c-4829-98f3-8389697dc4e8", <- ID of the
Bridge Endpoint Profile
    "vlan_transport_zone_id": "a95c914d-748d-497c-94ab-10d4647daeba", <- ID of the VLAN
Transport Zone
    "vlan": 0
}

Response:
{
    "vlan": 0,
    "vlan_trunk_spec": {
        "vlan_ranges": []
    },
    "ha_enable": true,
    "bridge_endpoint_profile_id": "0010b881-0c1c-4829-98f3-8389697dc4e8", <- ID of the
Bridge Endpoint Profile
    "vlan_transport_zone_id": "a95c914d-748d-497c-94ab-10d4647daeba", <- ID of the Bridge
Endpoint Profile
    "resource_type": "BridgeEndpoint",
    "id": "7e9c3517-f15b-490b-b14e-5ec356e92655",
    "display_name": "7e9c3517-f15b-490b-b14e-5ec356e92655", <- ID of the Bridge Endpoint
    "tags": [],
    "_create_user": "admin",
    "_create_time": 1638555736780,
    "_last_modified_user": "admin",
    "_last_modified_time": 1638555736780,
```

```
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 0
}
```

You can find the UUID of the transport zone by navigating to **System > Fabric > Transport Zones**. Look for the VLAN transport zone of your Edge node.

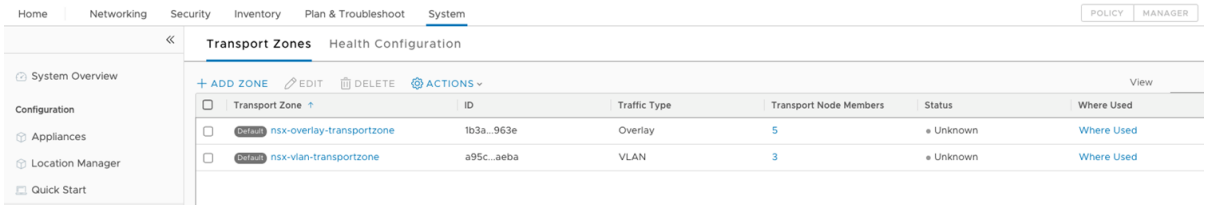

5 Make the following API call to create a port on the bridge.

```
POST https://<nsx-manager>/api/v1/logical-ports
{
    "logical_switch_id": "dd2841db-dff9-4927-834f-11b5ac8803d4", <- ID of the Logical
Switch
    "admin_state": "UP",
    "attachment": {
        "attachment_type": "BRIDGEENDPOINT",
        "id": "7e9c3517-f15b-490b-b14e-5ec356e92655" <- ID of the Bridge Endpoint
    }
}

Response:
{
    "logical_switch_id": "dd2841db-dff9-4927-834f-11b5ac8803d4", <- ID of the Logical
Switch
    "attachment": {
        "attachment_type": "BRIDGEENDPOINT",
        "id": "7e9c3517-f15b-490b-b14e-5ec356e92655"  <- ID of the Bridge Endpoint
    },
    "admin_state": "UP",
    "address_bindings": [],
    "switching_profile_ids": [
        {
            "key": "SwitchSecuritySwitchingProfile",
            "value": "47ffda0e-035f-4900-83e4-0a2086813ede"
        },
        {
            "key": "SpoofGuardSwitchingProfile",
            "value": "fad98876-d7ff-11e4-b9d6-1681e6b88ec1"
        },
        {
            "key": "IpDiscoverySwitchingProfile",
            "value": "64814784-7896-3901-9741-badeff705639"
        },
        {
            "key": "MacManagementSwitchingProfile",
            "value": "1e7101c8-cfef-415a-9c8c-ce3d8dd078fb"
```

```
        },
        {
            "key": "PortMirroringSwitchingProfile",
            "value": "93b4b7e8-f116-415d-a50c-3364611b5d09"
        },
        {
            "key": "QosSwitchingProfile",
            "value": "f313290b-eba8-4262-bd93-fab5026e9495"
        }
    ],
    "ignore_address_bindings": [],
    "internal_id": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "resource_type": "LogicalPort",
    "id": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3", <- ID of the Port created with the
Bridge
    "display_name": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "_create_user": "admin",
    "_create_time": 1638556071051,
    "_last_modified_user": "admin",
    "_last_modified_time": 1638556071051,
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 0
}
```

You can see the UUID by navigating to **Networking > Logical Switches**.

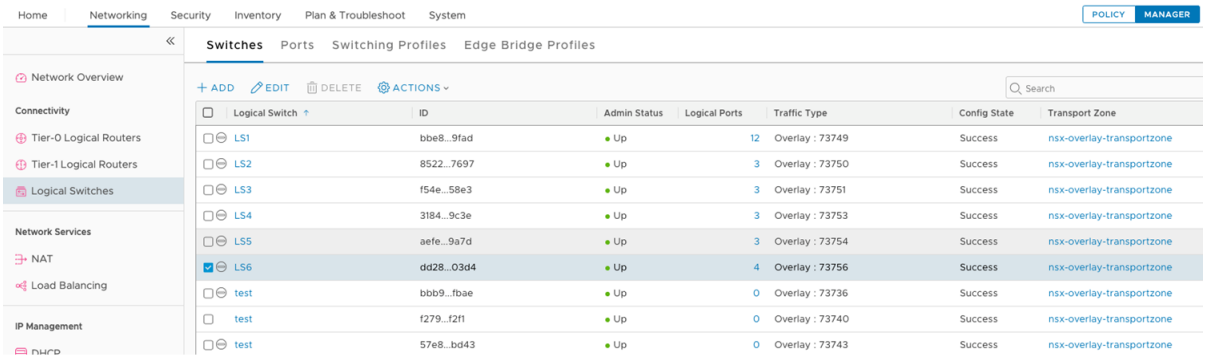

This logical switch is the stretched segment created from Global Manager. You can see it in NSX Manager UI after selecting Policy mode.

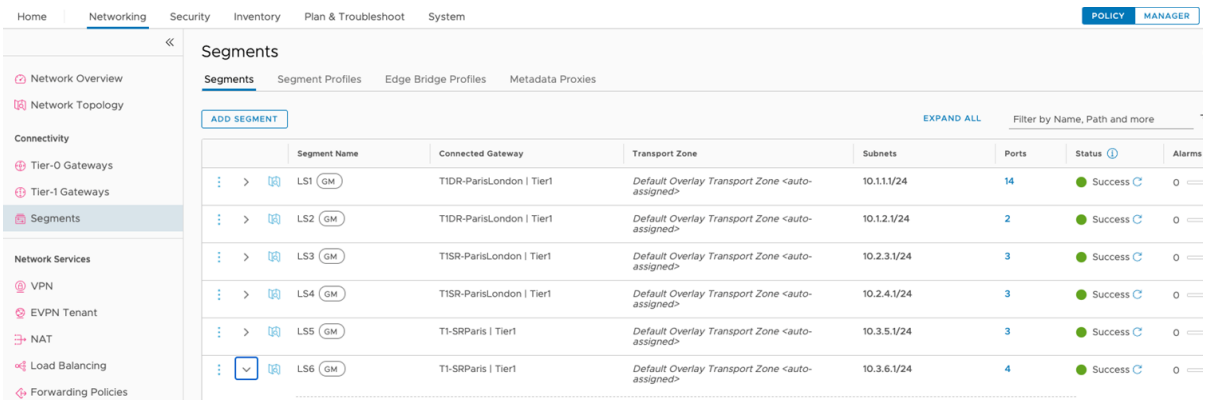

In Manager mode, you can see that the bridge is successful.

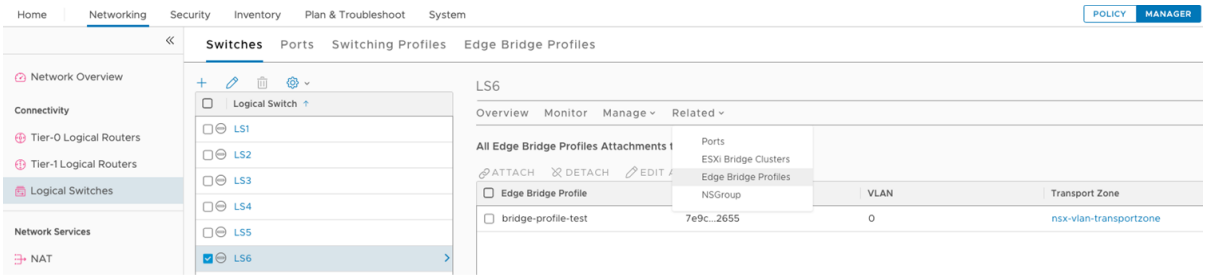

In Policy mode, you can see the VIF that was created. It shows that it is connected to the bridge endpoint but does not have more details.

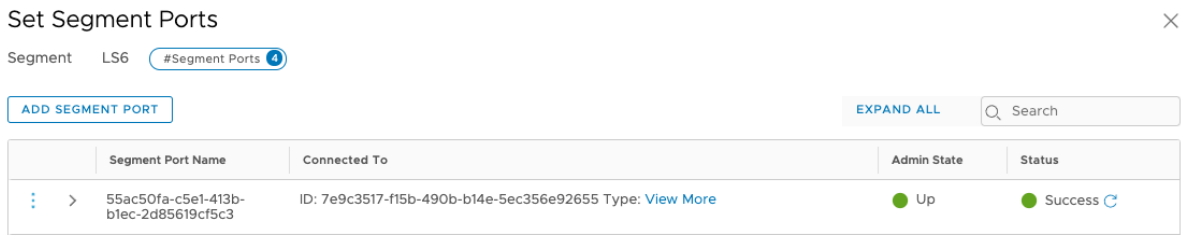

You can get more information on this port with the following API call:

```
GET https://<nsx-manager>/api/v1/logical-ports/55ac50fa-c5e1-413b-b1ec-2d85619cf5c3

Response:
{
    "logical_switch_id": "dd2841db-dff9-4927-834f-11b5ac8803d4",
    "attachment": {
        "attachment_type": "BRIDGEENDPOINT",
        "id": "7e9c3517-f15b-490b-b14e-5ec356e92655"
    },
    "admin_state": "UP",
    "address_bindings": [],
    "switching_profile_ids": [
        {
            "key": "SwitchSecuritySwitchingProfile",
            "value": "47ffda0e-035f-4900-83e4-0a2086813ede"
        },
        {
            "key": "SpoofGuardSwitchingProfile",
            "value": "fad98876-d7ff-11e4-b9d6-1681e6b88ec1"
        },
        {
            "key": "IpDiscoverySwitchingProfile",
            "value": "64814784-7896-3901-9741-badeff705639"
        },
        {
            "key": "MacManagementSwitchingProfile",
            "value": "1e7101c8-cfef-415a-9c8c-ce3d8dd078fb"
        },
        {
            "key": "PortMirroringSwitchingProfile",
```

```
            "value": "93b4b7e8-f116-415d-a50c-3364611b5d09"
        },
        {
            "key": "QosSwitchingProfile",
            "value": "f313290b-eba8-4262-bd93-fab5026e9495"
        }
    ],
    "ignore_address_bindings": [],
    "internal_id": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "resource_type": "LogicalPort",
    "id": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "display_name": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "_create_user": "admin",
    "_create_time": 1638556071051,
    "_last_modified_user": "admin",
    "_last_modified_time": 1638556071051,
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 0
}
```

Results

Now that you have configured the bridge for the Federated segment, perform the following tasks:

- Configure the logical switch to connect to the Edge bridge. See Configure the Logical Switch to Connect to the Edge Bridge.
- Test the connectivity across the bridge. See Test the Connectivity Across the Layer 2 Bridge.

Post Migration Bridge Removal

After the migration, remove the objects that are no longer needed.

1 Make the following PUT API call. The logical port ID (55ac50fa-c5e1-413b-b1ec-2d85619cf5c3) is in the response from the POST API call in step 5 above. Use the response from the POST API call as the body of this call but without the "`attachment`" parameter.

```
PUT https://<nsx-manager>/api/v1/logical-ports/55ac50fa-c5e1-413b-b1ec-2d85619cf5c3
{
    "logical_switch_id": "dd2841db-dff9-4927-834f-11b5ac8803d4",
    "admin_state": "UP",
    "address_bindings": [],
    "switching_profile_ids": [
        {
            "key": "SwitchSecuritySwitchingProfile",
            "value": "47ffda0e-035f-4900-83e4-0a2086813ede"
        },
        {
            "key": "SpoofGuardSwitchingProfile",
            "value": "fad98876-d7ff-11e4-b9d6-1681e6b88ec1"
        },
        {
            "key": "IpDiscoverySwitchingProfile",
```

```
            "value": "64814784-7896-3901-9741-badeff705639"
        },
        {
            "key": "MacManagementSwitchingProfile",
            "value": "1e7101c8-cfef-415a-9c8c-ce3d8dd078fb"
        },
        {
            "key": "PortMirroringSwitchingProfile",
            "value": "93b4b7e8-f116-415d-a50c-3364611b5d09"
        },
        {
            "key": "QosSwitchingProfile",
            "value": "f313290b-eba8-4262-bd93-fab5026e9495"
        }
    ],
    "ignore_address_bindings": [],
    "internal_id": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "resource_type": "LogicalPort",
    "id": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "display_name": "55ac50fa-c5e1-413b-b1ec-2d85619cf5c3",
    "_create_user": "admin",
    "_create_time": 1638556071051,
    "_last_modified_user": "admin",
    "_last_modified_time": 1638556071051,
    "_system_owned": false,
    "_protection": "NOT_PROTECTED",
    "_revision": 0
}
```


2 Delete the port and bridge endpoints with the following API calls. The port ID and the bridge endpoint ID are in the response from the POST API call in step 5 above.

```
DEL https://<nsx-manager>/api/v1/logical-ports/55ac50fa-c5e1-413b-b1ec-2d85619cf5c3
DEL https://<nsx-manager>/api/v1/bridge-endpoints/7e9c3517-f15b-490b-b14e-5ec356e92655
```

Test the Connectivity Across the Layer 2 Bridge

Check the status of the Edge bridge and perform basic connectivity tests to ensure that the bridge is ready to use.

Prerequisites

- A Layer 2 Edge bridge is created and configured, as explained earlier.
- Make sure that you have a VM attached to the bridged overlay segment to run the connectivity tests.

Procedure

1 Verify that the Edge bridge state is Up.

   a Log in to the Edge CLI as an **admin** user where the bridge is configured. You can connect to the CLI either using an SSH session or vSphere Web console.

   b Display configuration and state of all L2 bridges on the Edge.

   ```
   nsx-edge-1> get bridge
   ```

   Verify that the Device State is Up.

   For more information about this Edge CLI command, see the *NSX-T Data Center Command-Line Interface Reference*.

2 Do the following tests to check the connectivity across the bridge.

   a Run ping commands to verify VM-to-VM connectivity across the bridge.

   Remember that Distributed Firewall (DFW) rules are still running on the NSX-V workload VMs. Depending on the Security Policy, connectivity to VM on the overlay segment might be blocked. You might need to edit the DFW rules temporarily only to allow VM-to-VM connectivity across the bridge.

   b Run ping commands to verify connectivity of the VM on the bridged overlay segment to the default gateway, which currently is the Distributed Logical Router in your NSX-V environment.