Juniper Networks | Engineering Simplicity

# 3-Stage Data Center Design with Juniper Apstra and VMware NSX-T (Inline Mode)—Juniper Validated Design Extension (JVDE)

Published
2024-04-10

# Table of Contents

About this Document | 1

Solution Benefits | 1

Use Case and Reference Architecture | 3

Solution Architecture | 9

Configuration Walkthrough | 13

Validation Framework | 92

Test Objectives | 96

Results Summary and Analysis | 98

Recommendations | 99

# 3-Stage Data Center Design with Juniper Apstra and VMware NSX-T (Inline Mode)—Juniper Validated Design (JVD)

Juniper Networks Validated Designs provide customers with a comprehensive, end-to-end blueprint for deploying Juniper solutions in their network. These designs are created by Juniper's expert engineers and tested to ensure they meet the customer's requirements. Using a validated design, customers can reduce the risk of costly mistakes, save time and money, and ensure that their network is optimized for maximum performance.

## About this Document

This document is intended for an audience familiar with VMware technologies such as VMware NSX-T and VMware vSphere. The audience is also expected to be proficient with the 3-Stage data center design and Juniper products such as QFX Series Switches and Juniper Apstra.

**NOTE**: Nomenclature Note: Edge-Routed Bridging (ERB) is the Juniper terminology for a network architecture that is referred to as Distributed VXLAN Routing with EVPN, or the distributed gateways model elsewhere in the industry.

## Solution Benefits

**IN THIS SECTION**


- Juniper Validated Design Extensions Benefits | 2
- Juniper Apstra and VMware Integration Benefits | 2

This Juniper Validated Design Extension (JVDE) document is an extension of the 3-Stage Data Center Design with Juniper Apstra JVD; for more information on deploying 3-stage data center fabric with Juniper Apstra, refer to the JVD. The document provides detailed instructions for deploying VMware-NSX-T in inline mode for integration with Juniper Apstra. The solution is designed to meet the needs of Juniper's customers who run data centers with the VMware platform with vSphere, ESXi servers, and virtual machines.

It is based on best practices as determined by Juniper's subject matter experts, and Juniper support teams have extensive training and resources necessary to support networks based on JVDEs.

## Juniper Validated Design Extensions Benefits

- Qualified Deployments—JVDEs are a prescriptive blueprint for building upon a JVD data center fabric to meet the requirements of a specific use case. This approach makes building blocks JVDEs "known quantities" that can be deployed quickly, simply, and reliably.
- JVDEs are designed to meet the needs of most of Juniper's data center customers and are based on customer feedback. It is designed to scale beyond the initial design and support the adoption of different hardware platforms based on customer requirements.
- Risk Mitigation—Each JVDE goes through the New Product Initiative (NPI) testing framework to achieve validation. JVDEs contain the configuration necessary to extend a JVD data center network fabric with new functionality based on best practices and common use cases.
- JVDEs are verified by a suite of automated testing tools that can be used to validate the performance and reliability of Juniper solutions.
- Predictability—NPI testing verifies that all products in the JVDE work together as expected, using the explicitly defined versions of hardware and software documented therein. Common use cases are tested to determine the capabilities and limitations of the JVDE's constituent products when working together.

The underlying JVD data center network fabric, as well as any products and services listed in the JVDE, is tested for end-to-end functionality. This ensures that the specific combination of hardware, software, and features function as expected with the prescribed Junos OS releases.

## Juniper Apstra and VMware Integration Benefits

Integrating Juniper Apstra with VMware NSX-T and VMware vCenter simplifies data center network operations. This integration accelerates the deployment of fabric VLANs needed for deploying NSX-T in the data center and connects the NSX-T overlay network with the fabric underlay. This, in turn,

expedites troubleshooting and remediation of VLAN misconfiguration by automatically suggesting the correct network fabric changes.

Combining Juniper Apstra, VMware NSX-T, and VMware vCenter provides network administrators with visibility into the networking details of virtual machines (VMs) and containers hosted by ESXi servers, which are connected to leaf switches managed by Juniper Apstra as part of the JVD data center network fabric.

# Use Case and Reference Architecture

**IN THIS SECTION**


- VRF Characteristics: | 4
- Juniper Hardware and Software Components | 5
- VMware Software Components | 7
- Apstra Resources: ASN, Fabric, and Loopback IP Address | 8
- VMware NSX-T Manager Resources | 8


This JVDE utilizes an edge-routed bridging (ERB) network architecture. ERB uses lean spines that only perform IP forwarding and do not terminate VXLAN tunnel endpoints (VTEPs). This approach allows for spine switches with simpler configurations and reduced demands leading to higher network stability.

In an ERB architecture, leaf switches focus on learning and advertising the local MAC addresses to other remote switches through the BGP EVPN control plane. This means leaf switches can discover all the “remote” hosts without flooding the overlay with ARP or ND requests. Border Leaf switches serve as the gateway to external networks. With this design philosophy, in this document, the VMware NSX-T edge node terminates on the border leaf switches.

This JVDE is built upon the 3-Stage Data Center Design with Juniper Apstra JVD, which is the underlying network fabric for the purposes of this document. The underlying JVD network uses Juniper QFX, PTX, and ACX Series switches, which are managed by Juniper Apstra. Figure 1 on page 4 depicts the topology of the 3-Stage Data Center Design with Juniper Apstra JVD and is referenced throughout this JVDE.

**Figure 1: 3-Stage Reference Design with VMware NSX-T**

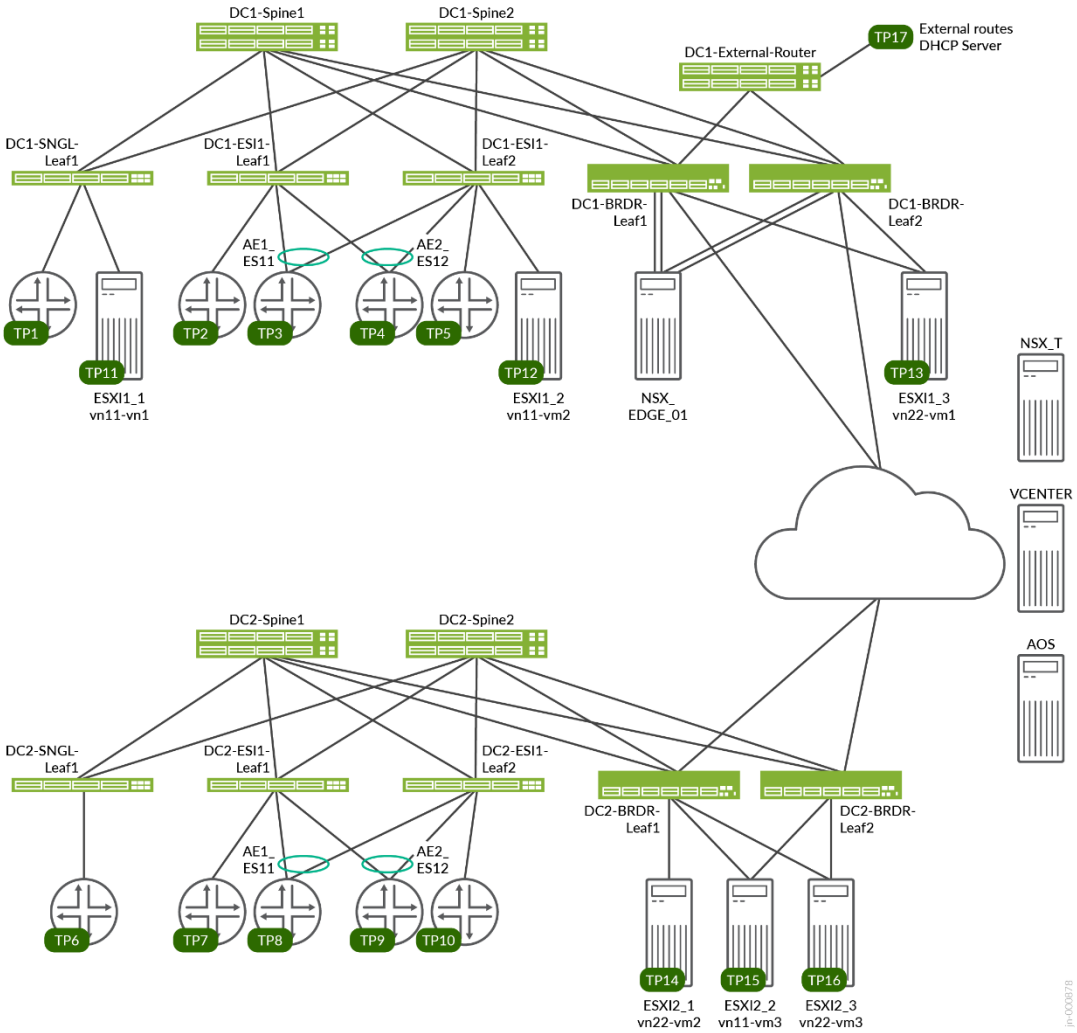


## VRF Characteristics:

### RED VRF

- VLANs 400–649 with IRB v4/v6
- on DC1-SNGL-LEAF1 single access port
- on DC1-ESI-LEAF1 single access port, AE1 and AE2
- on DC1-ESI1-LEAF2 single access port, AE1 and AE2
- on DC1-BRDR-LEAF1 to distribute routes to external-router

- on DC1-BRDR-LEAF2 to distribute routes to external-router
- VLANs 400–649 on each test port with 10 unique MAC/IP per VLAN
- DHCP client on TP3
- External DHCP server on TP17

**Blue VRF**

- VLANs 3500–3749 with IRB v4/v6
- on DC1-SNGL-LEAF1 single access port
- on DC1-ESI-LEAF1 single access port, AE1 and AE2
- on DC1-ESI1-LEAF2 single access port, AE1 and AE2
- on DC1-BRDR-LEAF1 to distribute routes to external-router
- on DC1-BRDR-LEAF2 to distribute routes to external-router
- VLANs 3500–3749 on each test port with 10 unique MAC/IP per VLAN
- DHCP client on TP3, TP4, TP5
- External DHCP server on TP2

## Juniper Hardware and Software Components

For this solution, the Juniper products and software versions are listed below. The listed architecture is the recommended base representation for the validated solution. As part of a complete solutions suite, we routinely swap hardware devices with other models during iterative use case testing. Each platform also goes through the same tests for each specified version of Junos OS.

**Juniper Hardware Components**

The following switches are tested and validated to work with the 3-Stage Fabric with Juniper Apstra JVD in the following roles:

**Table 1: Validated Devices and Positioning**

| Validated Devices and Positioning | | | |
|---|---|---|---|
| **Solution** | **Server Leaf Switches** | **Border Leaf Switches** | **Spine** |
| 3-stage EVPN/VXLAN (ERB) | QFX5120-48Y-8C* | QFX5130-32CD* | QFX5220-32CD* |
| | QFX5110-48S | QFX5700 | QFX5120-32C |
| | | ACX7100-48L | |
| | | ACX7100-32C | |
| | | PTX10001-36MR | |
| | | QFX10002-36Q | |

* marked are baseline devices

**Table 2: Baseline Devices and Positioning**

| Baseline Devices and Positioning | | | |
|---|---|---|---|
| **Juniper Devices** | **Role** | **Hostname** | **Software or Image Version** |
| QFX5220-32CD | Spine | dc1-spine1 and dc1-spine2 | Junos OS Evolved 22.2R3-S3.13 |
| QFX5120-48Y | Server Leaf | dc1-single-001-leaf1, dc1-esi-001-leaf1, and dc1-esi-001-leaf2 | Junos OS 22.2R3-S3.18 |
| QFX5130-32CD | Border Leaf | dc1-border-001-leaf1 and dc1-border-001-leaf2 | Junos OS Evolved 22.2R3-S3.13 |

**NOTE**: The 3-stage qualified devices are validated against Junos OS 22.2R3-S3 release, see Feature list for more information.

**Table 3: Juniper Software and Version**

| Juniper Software | |
|---|---|
| **Juniper Products** | **Software or Image version** |
| Juniper Apstra | AOS 4.2.1-207 |

## VMware Software Components

For the purposes of this document, the VMware products and their software versions are below. The listed architecture is the recommended base representation for the validated solution. As part of a complete solutions suite, we routinely swap hardware devices with other models during iterative use case testing. Each platform also goes through the same tests for each specified version of Junos OS.

**Table 4: VMware Products and Software Version**

| VMware Products | |
|---|---|
| **VMware Products** | **Software or Image Version** |
| NSX-T Edge | nsx-edge-3.2.1.0.0.19232403 |
| NSX-Manager | Version: 3.2.0.1.0.19232396 |
| vSphere Client | Version: 7.0.2 |
| ESXi | VMware ESXi, 7.0.2, 17630552 or later |

**NOTE**: Installing and upgrading of VMware components are not within the scope of this document.

## Apstra Resources: ASN, Fabric, and Loopback IP Address

Apstra resources for this JVDE are listed below. Resource assignments are based on the 3-Stage Data Center Design with Juniper Apstra JVD. To learn more about creating Resources in Apstra, see the Juniper Apstra User Guide.

**Table 5: Apstra Resources Used**

| Resources | Range |
|---|---|
| Fabric IP | 10.0.1.0/24 |
| Fabric Loopback IP | 192.168.255.0/24 |
| ASN | 64512 – 64999 |
| Routed Interface IP to NSX-T Edge Node (Border Leaf1to Left Link) | 192.168.100.0/24 |
| Routed Interface IP to NSX-T Edge Node (Border Leaf2 – Right Link) | 192.168.200.0/24 |
| VLAN from Border Leaf1 to NSX-T Edge Node (Left) | 100 |
| VLAN from Border Leaf1 to NSX-T Edge Node (Right) | 200 |

## VMware NSX-T Manager Resources

VMware resources for the validated solution are listed below.

**Table 6: VMware Resources Configured**

| Resources | Range | Notes |
|---|---|---|
| TEP Pool | 10.10.10.0/24 | Assigned by NSX-T manager to ESXi Host |
| vn11 | 10.9.11.0/24 | Assigned to VMs created in this document |
| vn22 | 10.9.22.0/24 | Assigned to VMs created in this document |

**Table 6: VMware Resources Configured *(Continued)***

| Resources | Range | Notes |
|---|---|---|
| ASN | 65000 | ASN for T0 Gateway |
| Loopback IP of T0 Gateway | 10.0.0.1/32 | Assigned while configuring T0 Gateway |
| Interface IP for T0 Interfaces to Border Leaf1 | 192.168.100.0/24 | Assigned while configuring T0 Gateway |
| Interface IP for T0 Interfaces to Border Leaf2 | 192.168.200.0/24 | Assigned while configuring T0 Gateway |
| Uplink1 segment VLAN | 100 | Uplink VLAN for Left |
| Uplink2 Segment VLAN | 200 | Uplink VLAN for Right |

# Solution Architecture

**IN THIS SECTION**


- 3-Stage Blueprint Topology | **10**
- VMware NSX-T Edge Cluster | **12**
- Transport Node Networking | **13**


The 3-stage fabric topology created in JVD 3-Stage Data Center Design with Juniper Apstra will be used for the purposes of this JVDE. There are two spines, three server leaf switches, and two border leaf switches. For brevity of this document, the whole process of setting up the data center blueprint is not covered in this document. Figure 2 on page 10 shows the data center blueprint from Juniper Apstra.

## 3-Stage Blueprint Topology

In the below topology, Figure 2 on page 10, racks are shown for border leaf switches, ESI leaf switches, and single leaf (non-ESI) switches. The leaf switches are connected to generic servers, including ESXi servers. When NSX-T is configured on the ESXi servers (as discussed in walkthrough sections), the virtual machines running on these servers can be micro-segmented based on security needs. This helps to improve security by reducing the attack surface and limiting the lateral movement of threats within your environment.

**Figure 2: Three-Stage Fabric Blueprint Topology**

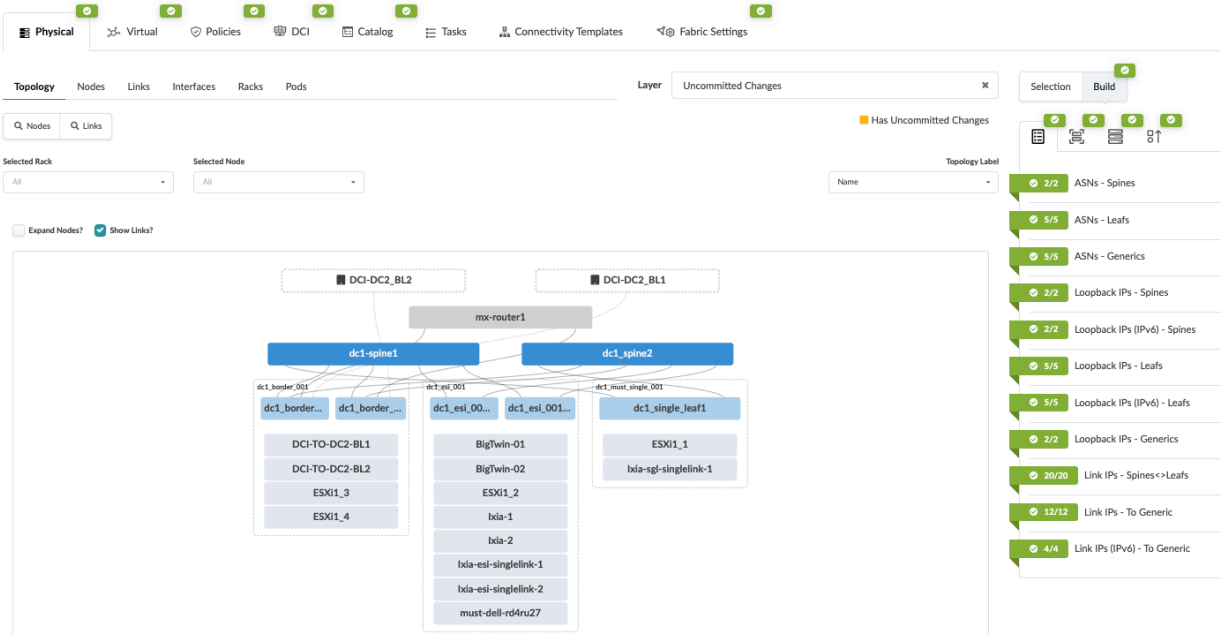


From the 3-stage topology above the border leaf switches are the gateways to the NSX-T edge routers. The ESXi servers connected to the border leaf rack host the NSX-T Edge routers as virtual machines. The following graphic shows the detailed border leaf connectivity towards the ESXi servers.

## Figure 3: Border Leaf1 Connectivity

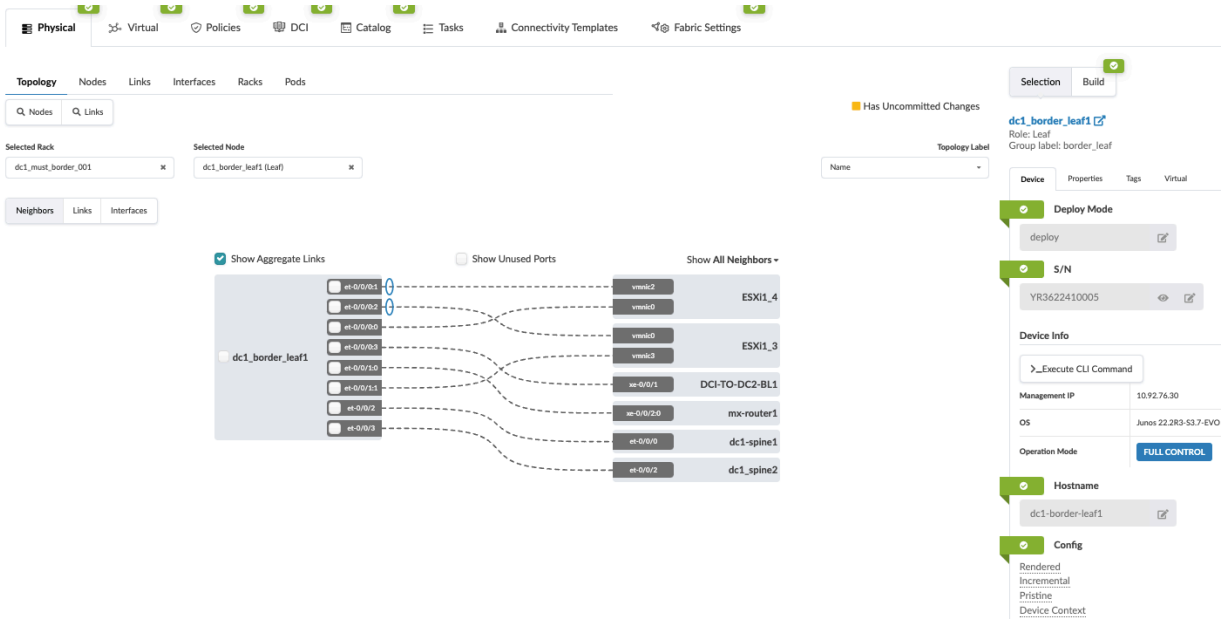


## Figure 4: Border Leaf2 Connectivity

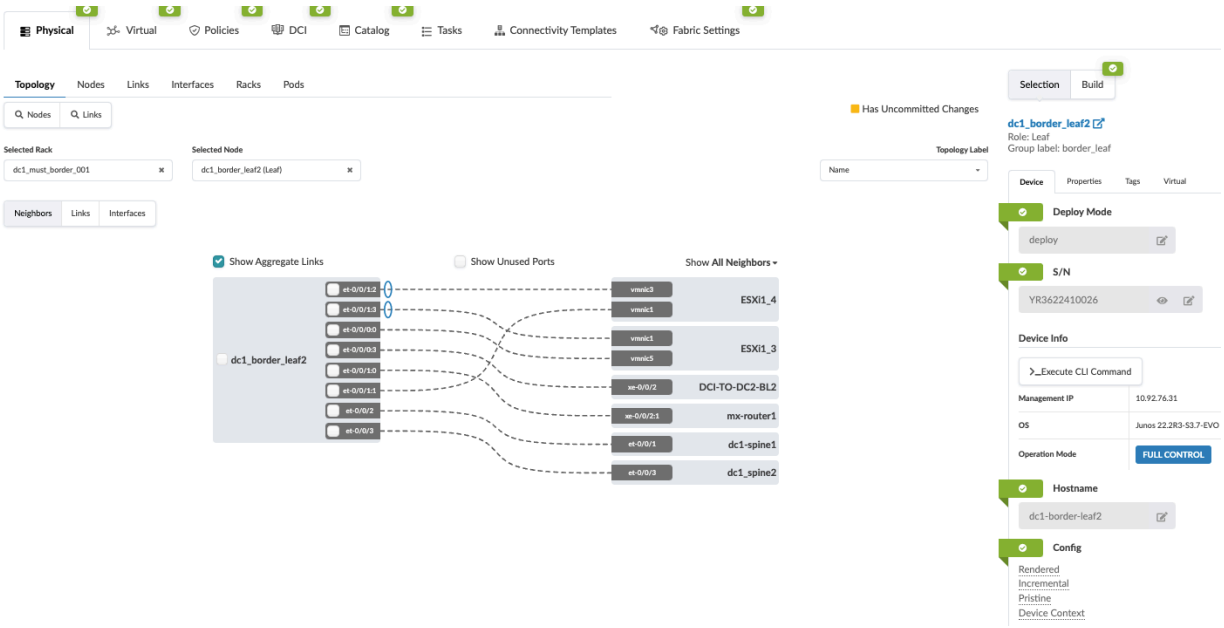

**NOTE**: ESXi1_4 (as shown in Figure 3 on page 11 and Figure 4 on page 11) is the ESXi server on which the NSX-T edge node will be hosted as will covered later in this document. Its connected to Border-Leaf1 and Border-Leaf2 for left and right uplinks.

## VMware NSX-T Edge Cluster

From VMware's documentation: *An NSX Edge Node is a transport node that runs the local control plane daemons and forwarding engines implementing the NSX-T data plane. It runs an instance of the NSX-T virtual switch called the NSX Virtual Distributed Switch, or N-VDS.*

For the purpose of this document, a single NSX-T Edge cluster is deployed with one edge node. An edge node is a virtual machine created on the ESXi server in vSphere. We highly recommend deploying more than one edge node on separate ESXi hosts.

**NOTE**: In production data center implementation to guarantee service high availability, we recommend that multiple NSX-T Edge Nodes or Edge Node clusters are deployed. This can be a single cluster with a maximum of 10 Edge Nodes, depending on the number of ECMP paths available. For more information, refer to the VMware NSXT Reference Design guide. **Edge Nodes** are 'service appliances' that provide pools of capacity. Edge Nodes provide network services that are not distributed down to the hypervisors and cannot be used for any other purposes. They provide the physical network uplinks (pNICs) that connect to data center fabric (underlay). In NSX-T 3.0 and later, the types of Edge Nodes supported are bare metal edge and VM edge. **Edge Clusters** are a group of Edge Transport nodes that provide a scalable, high-throughput, and highly available (redundant) gateway for logical networks created in NSX-T. This helps to guarantee the service availability of tier-0 and tier-1 gateways. For more information on high availability failover, refer to the VMware NSXT Reference Design guide.

In an ERB architecture, such as the one used in the underlying JVD used by this JVDE, border leaf switches serve as the gateway to external networks. Pursuant to this design philosophy, in this JVDE, the VMware NSX-T edge node terminates on the border leaf switches.

Figure 5 on page 13 shows how the Edge Node terminates on the border leaf switches. The setup of the NSX-T Edge VDS port connectivity is discussed later in this document.

**Figure 5: Edge Node Connectivity**

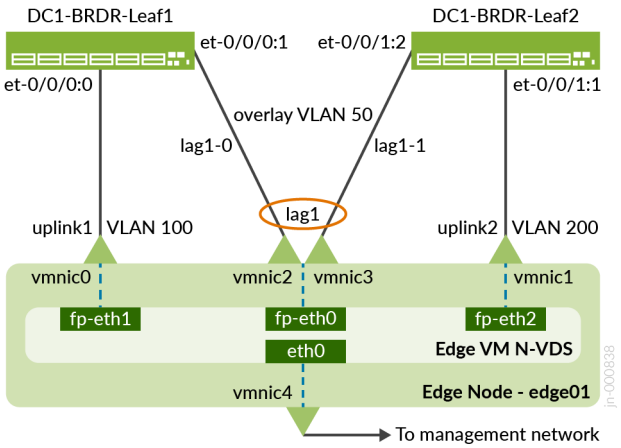


## Transport Node Networking

From VMware's documentation: Transport nodes are hypervisor hosts and NSX Edges that will participate in an NSX-T data center overlay. For a hypervisor host, this means that it hosts VMs that will communicate over NSX-T data center logical switches. For NSX-T Edge nodes, this means that it will have logical router uplinks and downlinks. All ESXi hosts in this document will be added to the NSX-T transport zone.

# Configuration Walkthrough

**IN THIS SECTION**


- VMware NSX-T and Juniper Apstra Integration Setup - Prerequisite Steps | **15**
- VMware NSX-T Manager: Create Tunnel Endpoint Pools | **16**
- VMware NSX-T Manager: Add vSphere to NSX-T Manager | **16**
- VMware vSphere: Configure VDS On ESXi Host | **17**
- VMware NSX-T: Create Overlay and VLAN Uplink Transport Zones | **29**
- VMware NSX-T: Configure the Uplink Profiles for the Host and Edge Nodes | **30**
- VMware vSphere: Add Transport Node Hosts to VDS | **34**

- VMware NSX-T: Prepare the Compute Cluster | **36**
- VMware NSX-T: Transport Nodes-Tunnel IPs | **38**
- VMware NSX-T: Deploy NSX Edge Node and Create Edge Cluster | **39**
- VMware NSX-T: Create a T-1 Gateway | **46**
- VMware NSX-T: Create Logical Segments | **46**
- VMware NSX-T: Create VLAN Backed Logical Segments | **48**
- VMware vSphere: Confirm the Creation of the Logical Segments | **49**
- VMware vSphere: Create VMs in the Segments | **49**
- VMware NSX-T: Create Tier0 Gateways T0-1 | **51**
- VMware NSX-T: Configure the Interfaces on the T0 Gateway | **52**
- VMware NSX-T: Configure Loopback Interface on T0 Gateway | **54**
- VMware NSX-T: Configure BGP on the T0 Gateway | **55**
- VMware NSX-T: Configure In-Line Mode and Route-Redistribution on the T0 Gateway | **58**
- VMware NSX-T: Create a Static Route to Loopback on Border Leaf Switches | **61**
- VMware NSX-T: Create IP Prefix lists on T0 Gateway | **64**
- VMware NSX-T: Connect the T1 and T0 Gateways | **65**
- Juniper Apstra: Add the NSX-Manager | **66**
- Juniper Apstra: Create a Routing Policy for NSX-T in the Blueprint | **67**
- Juniper Apstra: Create a Routing Zone in the Blueprint | **68**
- Juniper Apstra: Assign the Loopback IPs to the Routing Zone | **69**
- Juniper Apstra: Add the NSX-Manager into the Blueprint | **70**
- Juniper Apstra: Add the NSX-T-Overlay as a VN | **71**
- Juniper Apstra: Verify the Connectivity Templates | **72**
- Juniper Apstra: Assign Interface to the Connectivity Templates | **73**
- Juniper Apstra: Commit the Configuration | **74**
- VMware NSX-T: GENEVE Tunnels | **74**
- Juniper Apstra: Add Connectivity Templates for Connectivity from Edge Node to the Fabric | **74**
- Juniper Apstra: Add IP Link, BGP Peering and Static Route | **75**
- Juniper Apstra: Renaming the Generic System and Adding Links from Border Leaf switches | **78**
- Juniper Apstra: Assign the Interfaces to the Connectivity Template | **80**
- Juniper Apstra: Assign the IPs and VLAN IDs to the Interfaces | **83**
- Juniper Apstra: Commit the Configuration | **83**

- Juniper Junos OS: Verify Configs | **83**
- VMware NSX-T: Verify BGP Session on Edge | **84**
- VMware NSX-T: Verify BGP Session on ToR | **86**
- Verify Physical Fabric Configuration | **87**
- VMware NSX-T: Verify Overlay Connectivity (East-West) | **89**
- (Optional) Juniper Apstra: Adding vSphere Server to Juniper Apstra | **91**


## VMware NSX-T and Juniper Apstra Integration Setup - Prerequisite Steps

- Ensure NSX-T Manager is installed, and the management address is assigned to access it through SSH or User Interface (UI). For this solution, only one node cluster for NSX-T Manager is deployed. However, in production deployment, adding more than one node is recommended based on the NSX-T Data Center Installation Guide.

**Figure 6: Transport Node Networking**

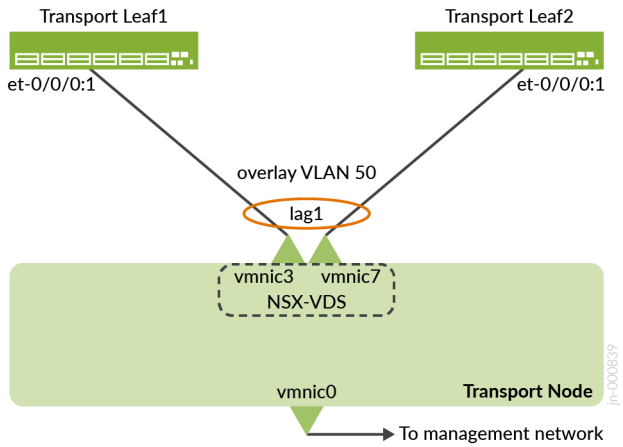


- Add all requisite ESXi hosts in VMware vSphere. Note that only four ESXi hosts are required for this use case, as shown in Figure 2 on page 10.

**NOTE**: Configuring VMware vSphere and adding hosts to the vSphere Client is not in the scope of this guide.

For detailed information on VMware NSX-T, refer to the VMware NSX-T 3.2 Guide.

## VMware NSX-T Manager: Create Tunnel Endpoint Pools

Tunnel Endpoints (TEPs) are used in the header of the outer (external) IP encapsulation to uniquely identify the hypervisor hosts originating and terminating the NSX-T encapsulation of overlay frames.

To create a TEP pool, log into VMware NSX-T Manager and navigate to **Networking** > **IP Management** > **IP Address Pools** > **IP Address Pools.**

**Figure 7: Create TEP Pool in NSX-T**

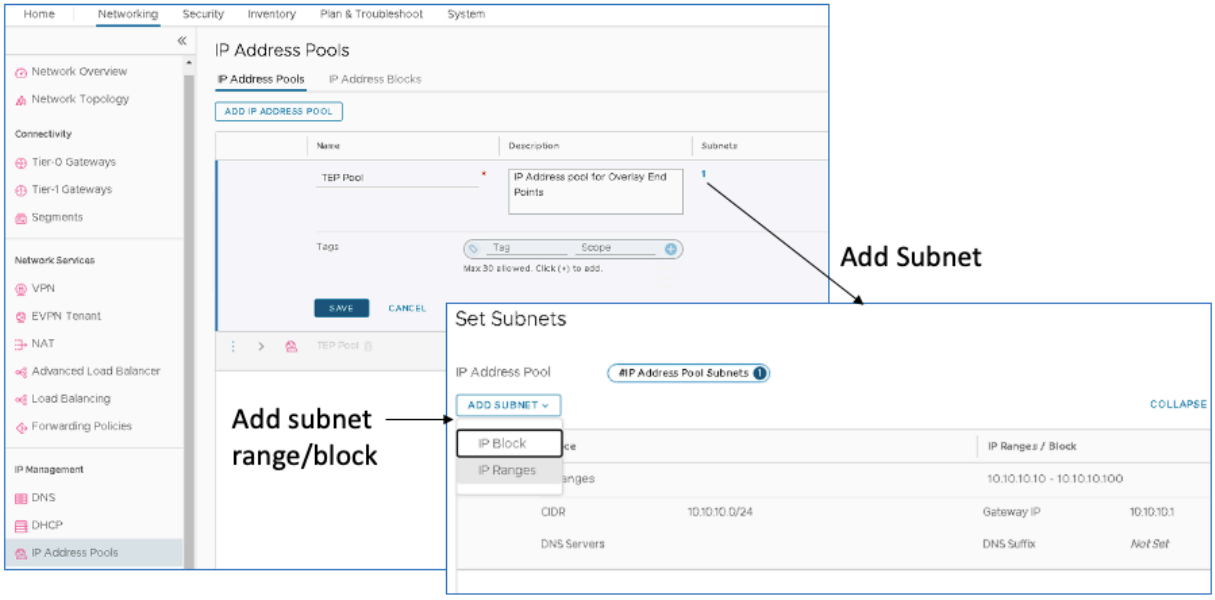


## VMware NSX-T Manager: Add vSphere to NSX-T Manager

Next, add the vSphere server as a Compute manager to the NSX-T manager. The NSX-T manager gets the inventory of ESXi hosts that will be used for Compute and Edge transport nodes.

To add the vSphere Server, log into NSX-T Manager and navigate to **System** > **Fabric** > **Compute Managers**.

**Figure 8: Add vSphere Details as Compute Manager**

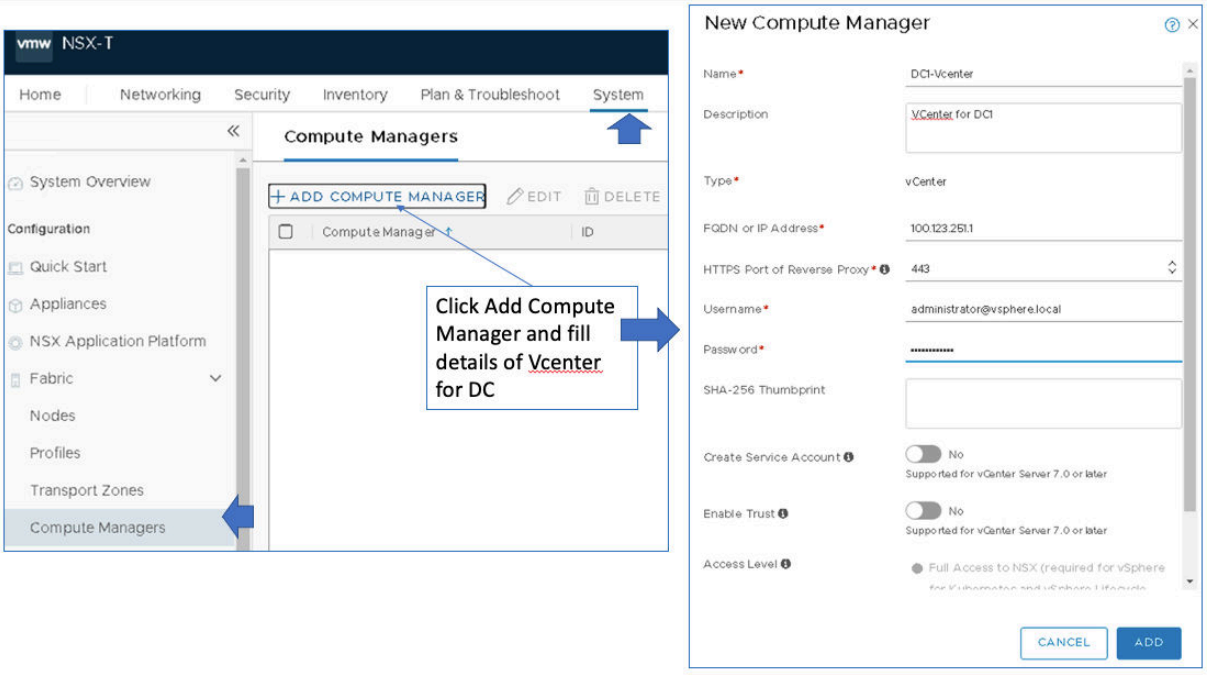


**NOTE**: If adding the vSphere in the NSX-T Manager fails because of the *Certificate of Compute Manager not valid* error, then follow the VMware KB article to validate the vSphere certificate.

## VMware vSphere: Configure VDS On ESXi Host

A vSphere Distributed Switch (VDS) provides centralized management and monitoring of the networking configuration of all hosts associated with the switch. For more information, please refer to the VMware documentation.

In VMware vSphere, under **Networking**, right-click the data center and select **Distributed Switch** > **New Distributed Switch**.

For this virtual distributed switch (VDS), create three Distributed Port Groups (DPGs) on the VDS. Enable **VLAN Trunking** on all the nodes. The configuration should be as follows:

On the ESXi1_4 device (that hosts the NSX-T-Edge Node) that is connected to the border leaf switches, the following VMNICs are configured in the lab:

1. Overlay VLAN: LAG1 (vmnic3+vmnic7). This is the aggregate Ethernet interface to border leaf switches.

2. Left-uplink: vmnic2. This is the routed interface to border leaf1 et-0/0/0:0.

3. Right-uplink: vmnic6. This is the routed interface to border leaf2 et-0/0/0:0.

> **NOTE**: The above VMNICs can be different depending on the setup. Ensure to select the appropriate switch interface to VMNIC mapping.

**Figure 9: Physical Adapter Configured on ESXi Host**

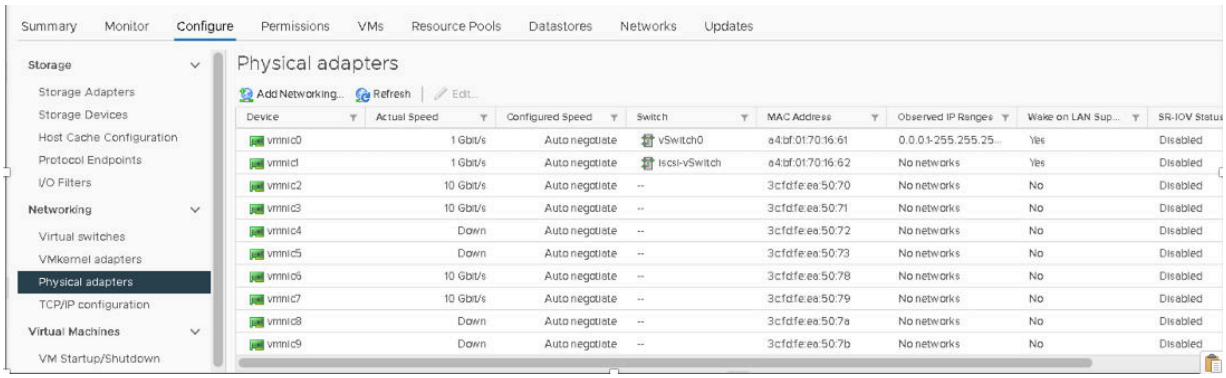


### 1. Create a VDS

In VMware vSphere, navigate to **Networking** and perform the following: Create a VDS (named NSX-T-Edge-VDS in this case) on the Edge Node Host and assign default uplinks. There are four default uplinks.

**Figure 10: Adding NSX-T-Edge VDS**

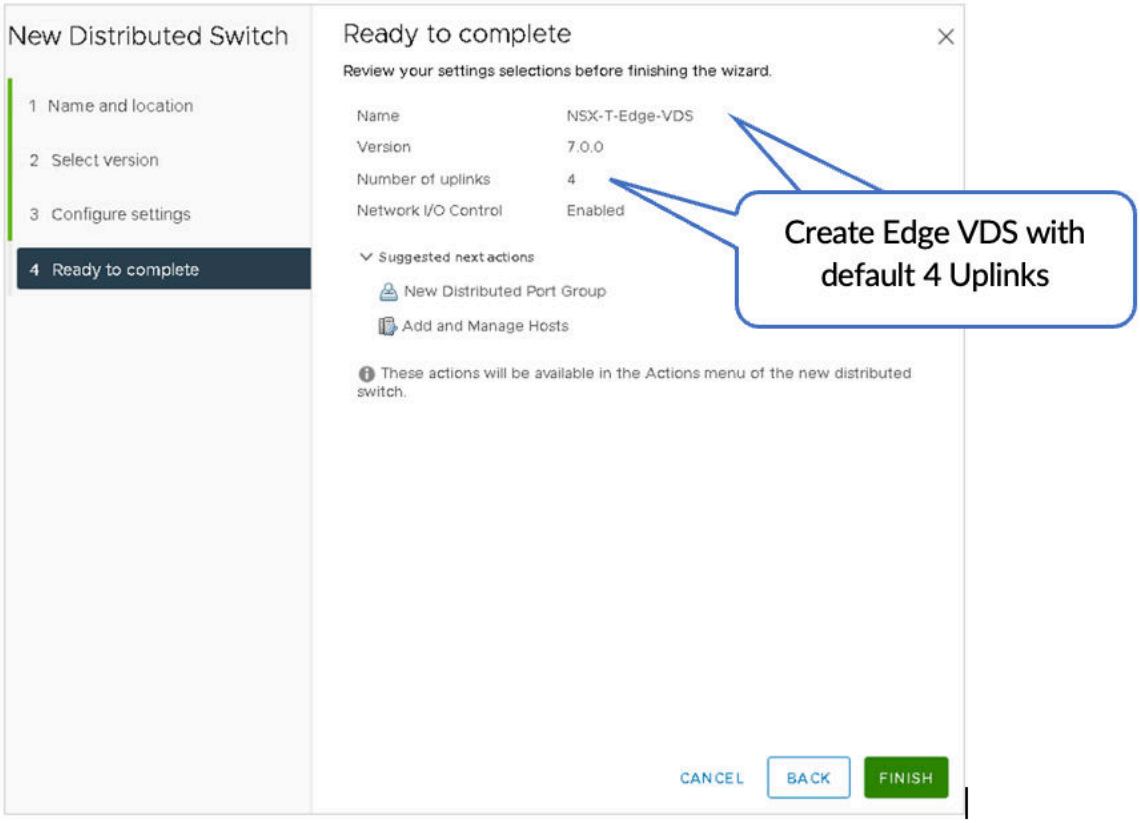


**2. Configure LACP lag1**

In VMware vSphere, navigate to **Networking** and perform the following:

Navigate to **Networks** > **NSX-T-Edge-VDS** > **Configure** > **LACP** to configure LACP lag1.

**Figure 11: Configuring LACP on NSX-T-EDGE VDS**

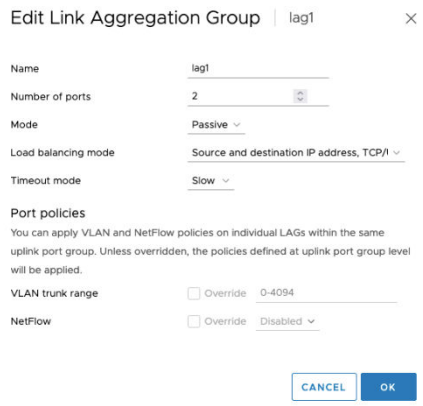


### 3. Assign VMNICs to Uplinks and LAG Ports

In VMware vSphere, perform the following:

On NSX-T-Edge-VDS actions, use **Add and Manage Hosts** to assign the VMNICs to the uplinks and LAG ports.

Each host must be connected to the fabric as described in the "Solution Architecture" on page 9 section.

**Figure 12: Add Hosts to NSX-T-Edge-VDS Showing VMNICs Assigned**

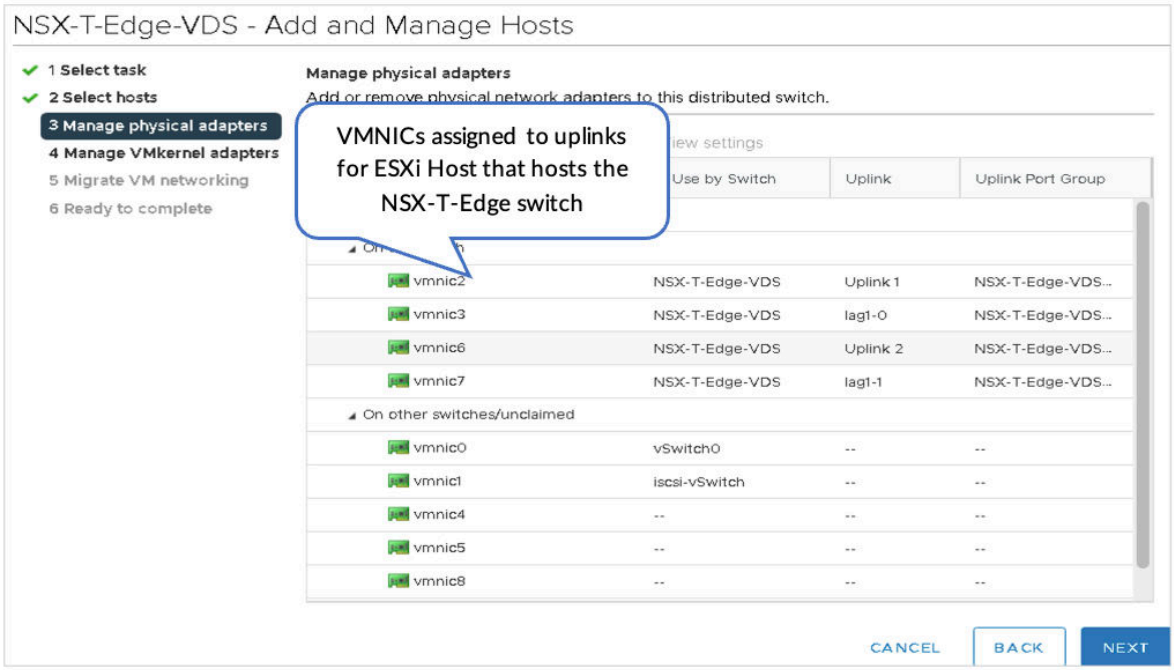

**Figure 13: Add Hosts to NSX-T-Edge-VDS -Assign VMNIC to Uplinks**

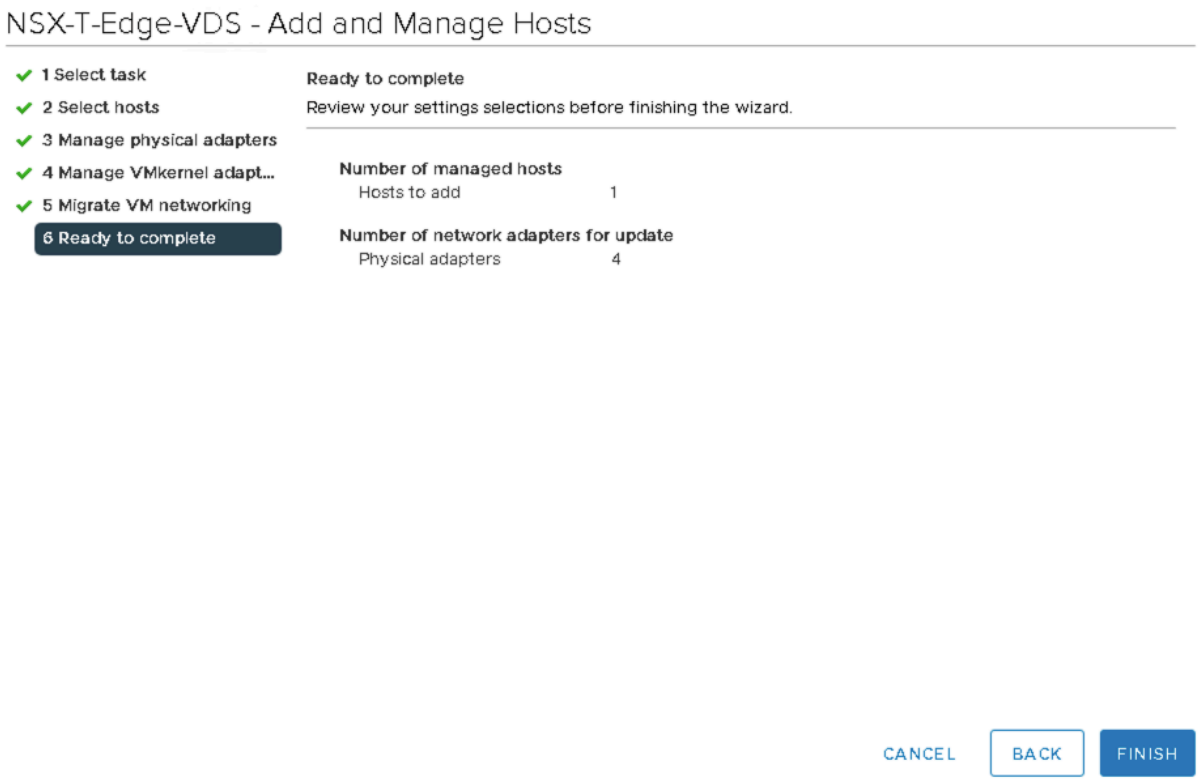


### 4. Create Three Distributed Port Groups

In VMware vSphere, perform the following:

1. On the VDS, create the following:
   - Port
   - Distributed Port Group:
     - Left-Uplink
     - Overlay
     - Right-Uplink
2. Enable **VLAN Trunking** on all the port groups.

**Figure 14: Adding Distributed Port Groups**

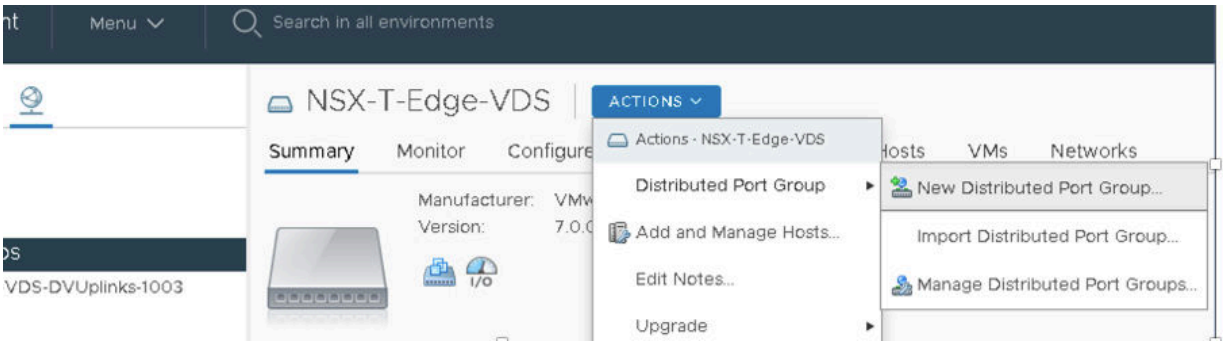


Below are the three Distributed Port Groups added to the NSX-T-Edge-VDS switch.

- **Left-UplinkConnects to the Border Leaf1**

**Figure 15: Adding Left-Uplink Distributed Port**

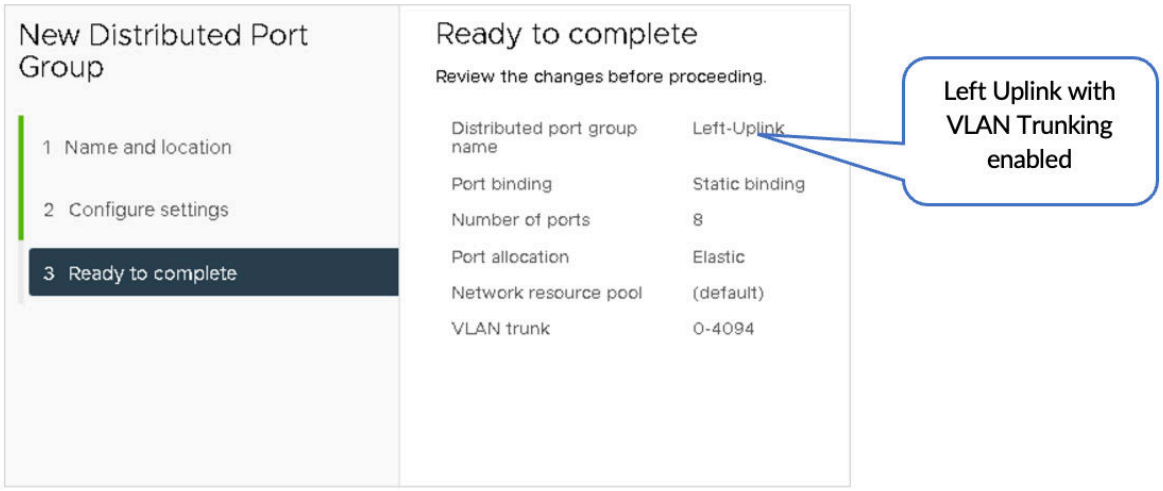

**Figure 16: Editing Left Uplink to Assign the Active Uplink**

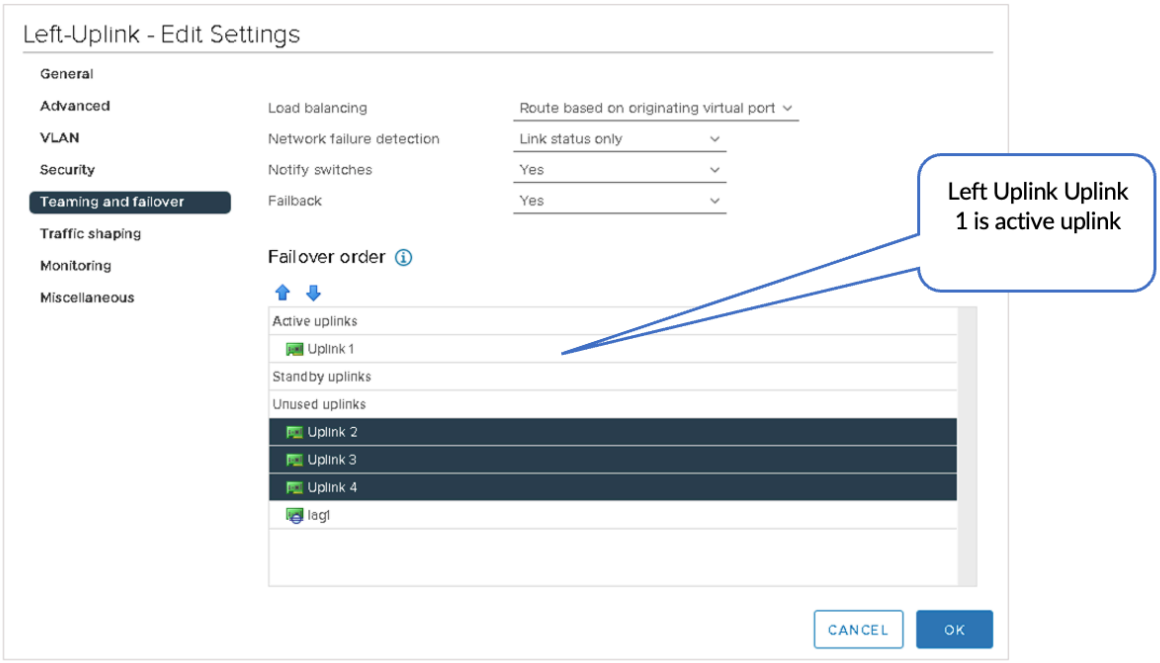


- **Right-Uplink Connects to the Border Leaf2**

**Figure 17: Adding Right-Uplink Distributed Port Group**

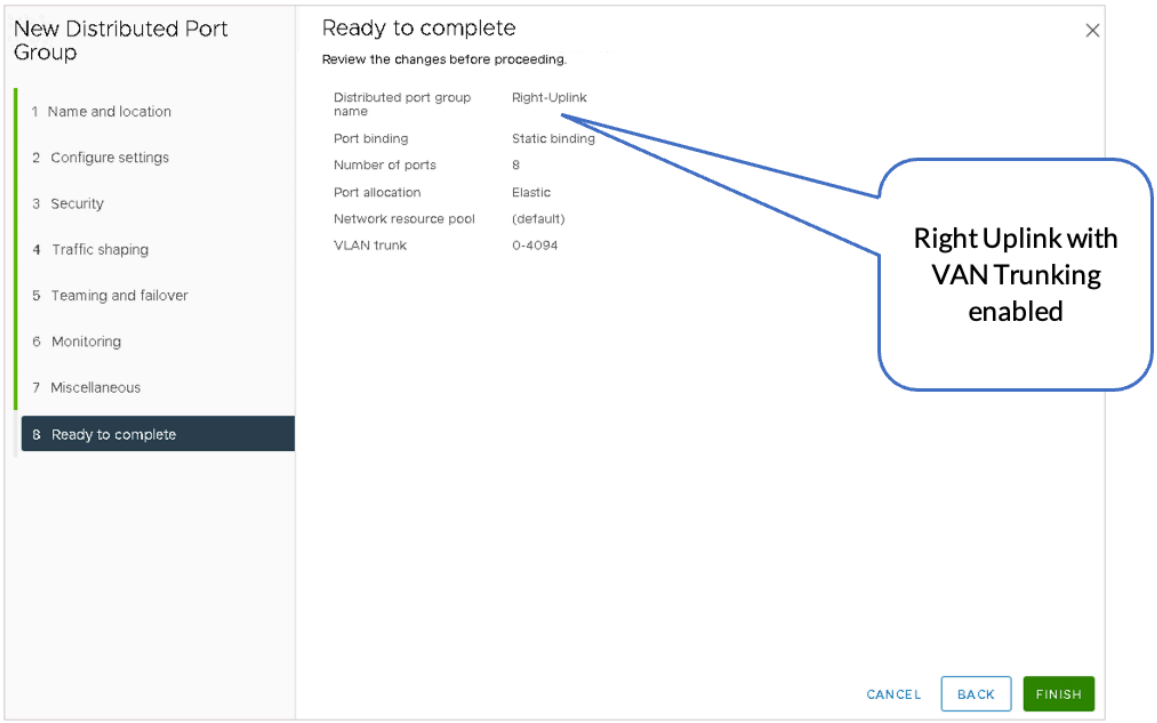

**Figure 18: Assigning Uplink in Active to Unused Order**

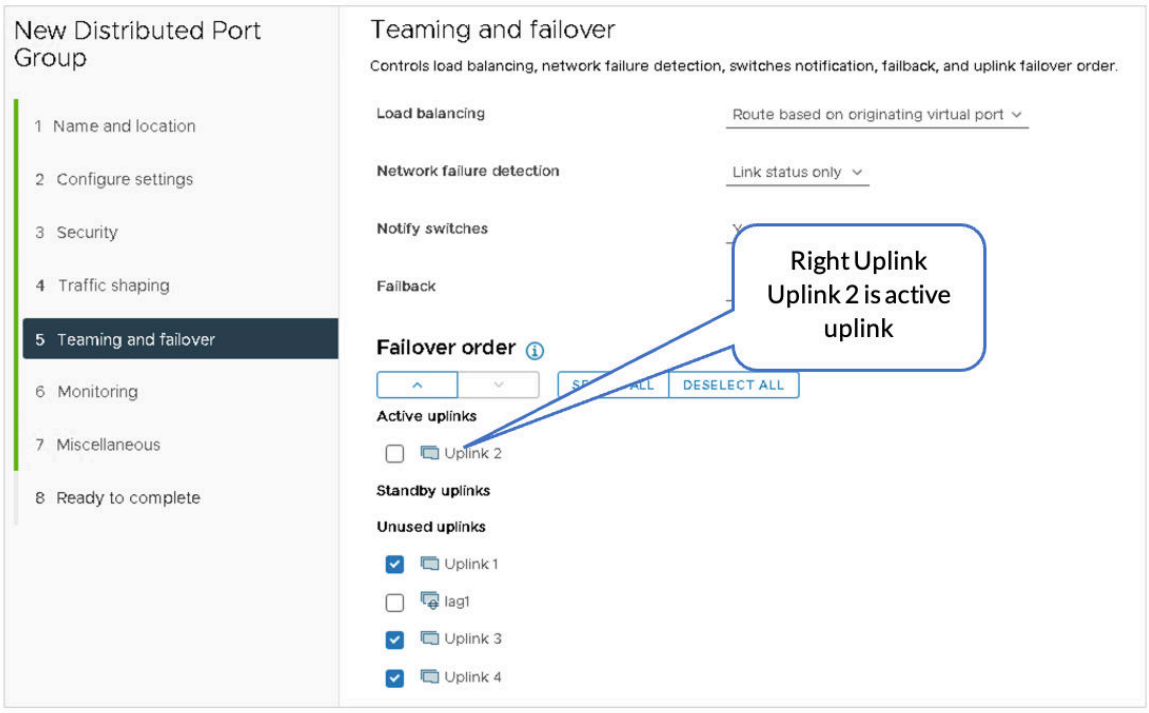


- **Overlay Link is Used for VLAN Transport Traffic.**

**Figure 19: Adding Overlay Distributed Port Group**

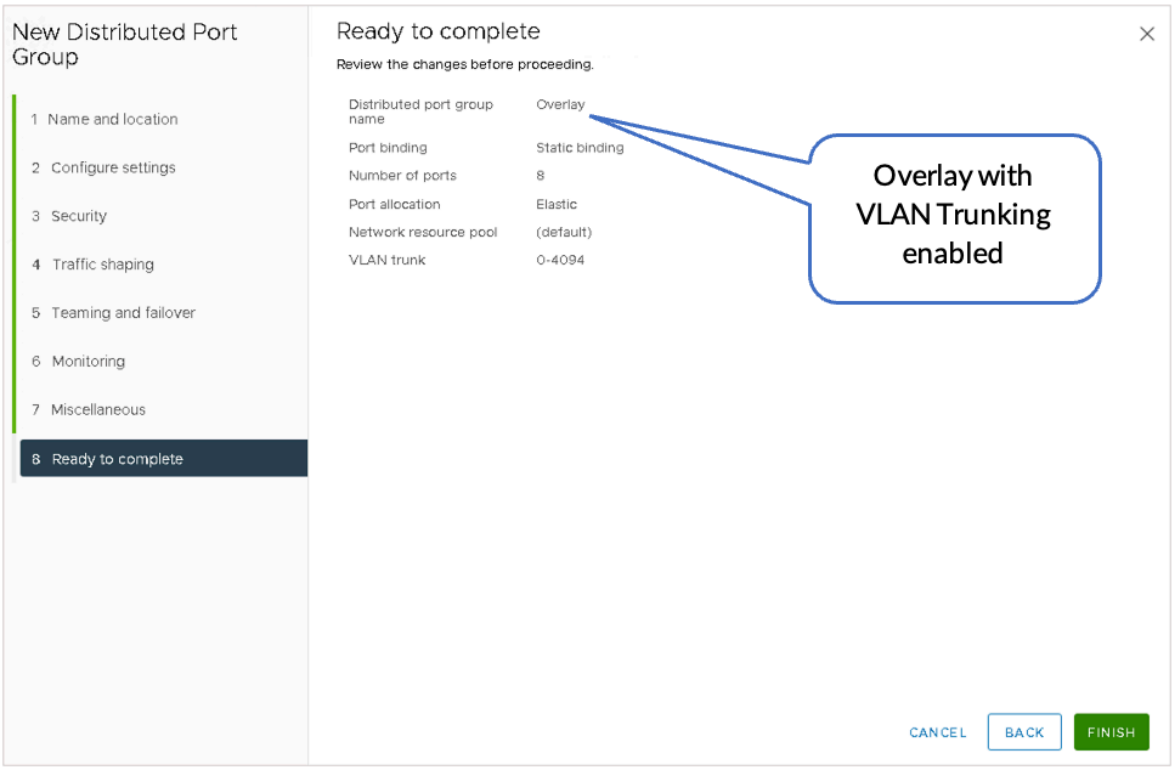

**Figure 20: Assigning Uplink Failover Order for Overlay**

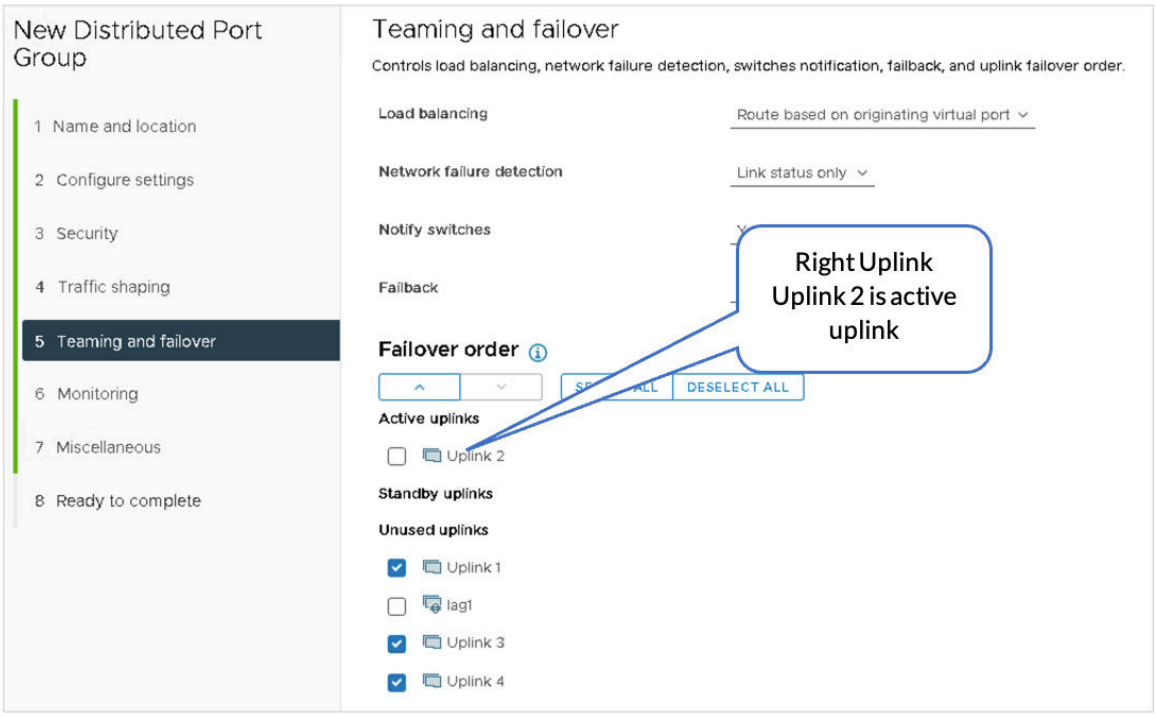


## 5. Review VDS

After configuring the NSXT-Edge-VDS switch, the switch should show the following port groups as configured. The Edge switch VM is assigned to the relevant port groups. Also, the physical adapters on the ESXi host are now allocated to the NSX-T-Edge-VDS.

**Figure 21: NSX-T Edge VM Created in vSphere**

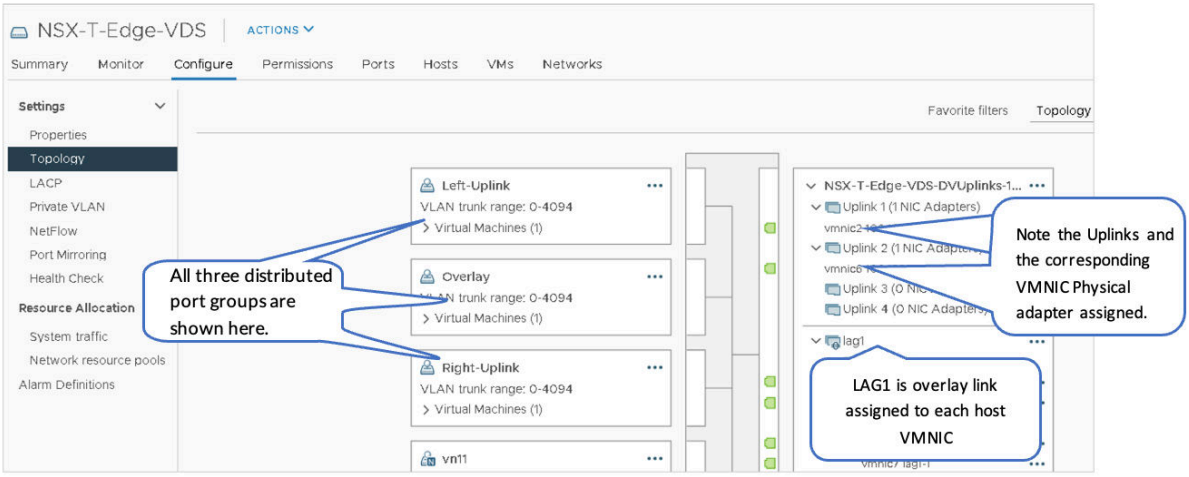


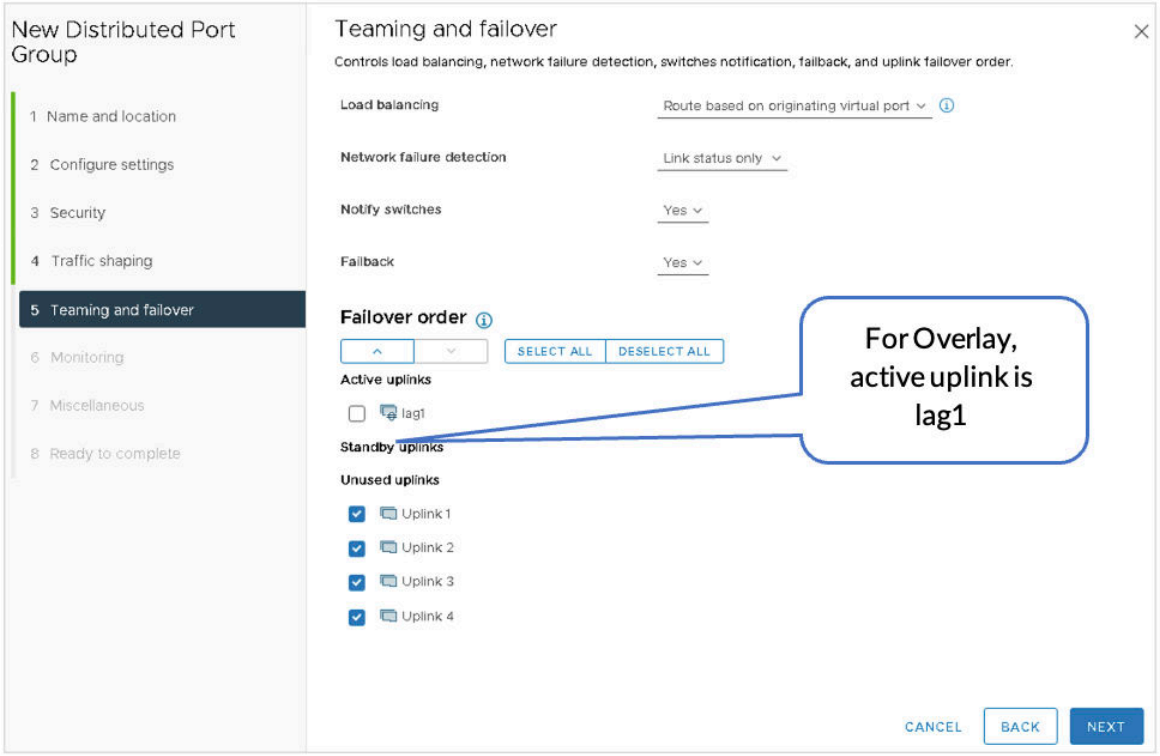

**Figure 22: Physical Adapters on ESXi Host Assigned to NSX-T-Edge-VDS**

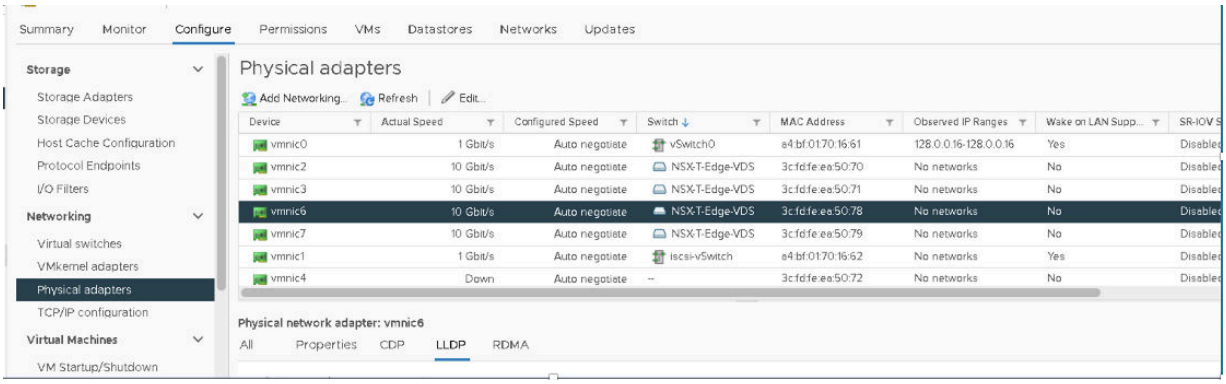


## VMware NSX-T: Create Overlay and VLAN Uplink Transport Zones

Three Transport Zones need to be created:

- Overlay-TZ—Used for GENEVE encapsulation of data packets.
- Uplink-TZ—Used for the NSX Edge uplink to connect to the switch fabric.
- Uplink-TZ-2.

### Create Three Transport Zones

To create three Transport Zones:

1. In NSX-T Manager, navigate to **System > Fabric > Transport Zones** and then click **+ADD ZONE**.
2. Select **Overlay** traffic type for Overlay-TZ.
3. Select **VLAN** traffic type for Uplink-TZ.
4. Select **VLAN** traffic type for Uplink-TZ-2.

**Figure 23: Transport Zones in NSX-T**

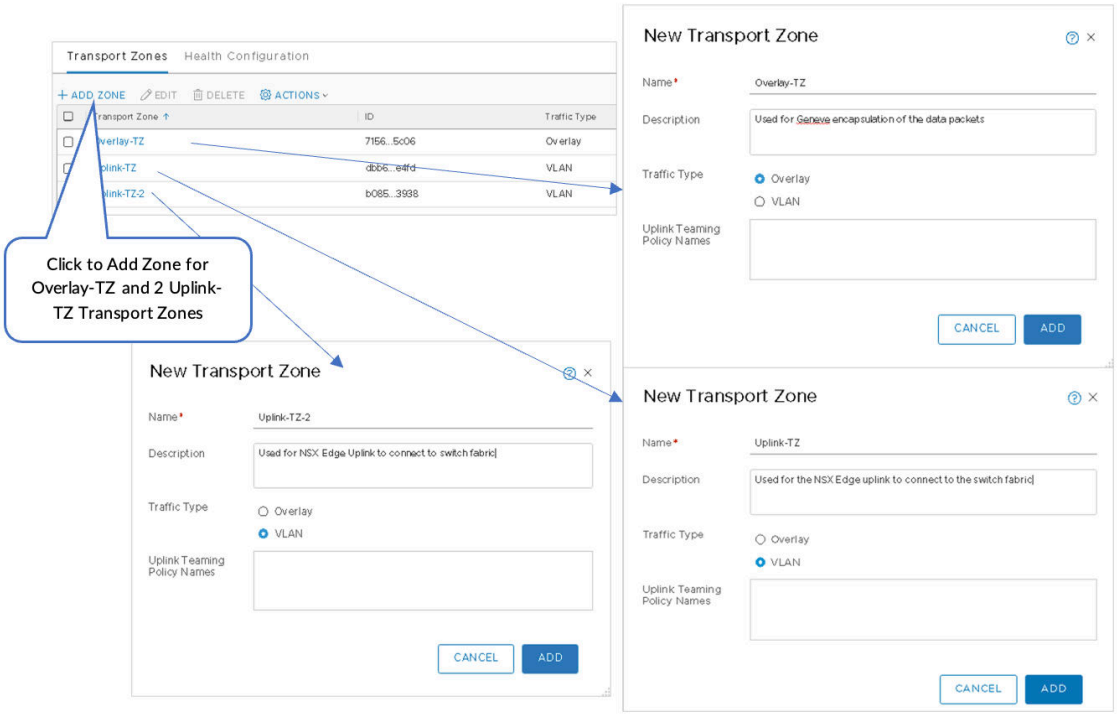


## VMware NSX-T: Configure the Uplink Profiles for the Host and Edge Nodes

Now that the VDS is configured within vSphere, the uplink profiles must be created within NSX-T. The profiles created correspond to the uplinks Overlay, Edge-Right, and Edge-Left.

**NOTE**: NSX-T Edge-left connects to Border Leaf-1 and NSX-T Edge-right connects to Border Leaf-2 for uplink/BGP redundancy.

- **Create Three Uplink Profiles**

1. In NSX-T Manager, navigate to **System** > **Fabric** > **Profiles** > **Uplink Profiles**, and then click **ADD**.
2. Overlay-profile: VLAN 50 maps to lag1.

**Figure 24: Overlay Profile with VLAN 50**

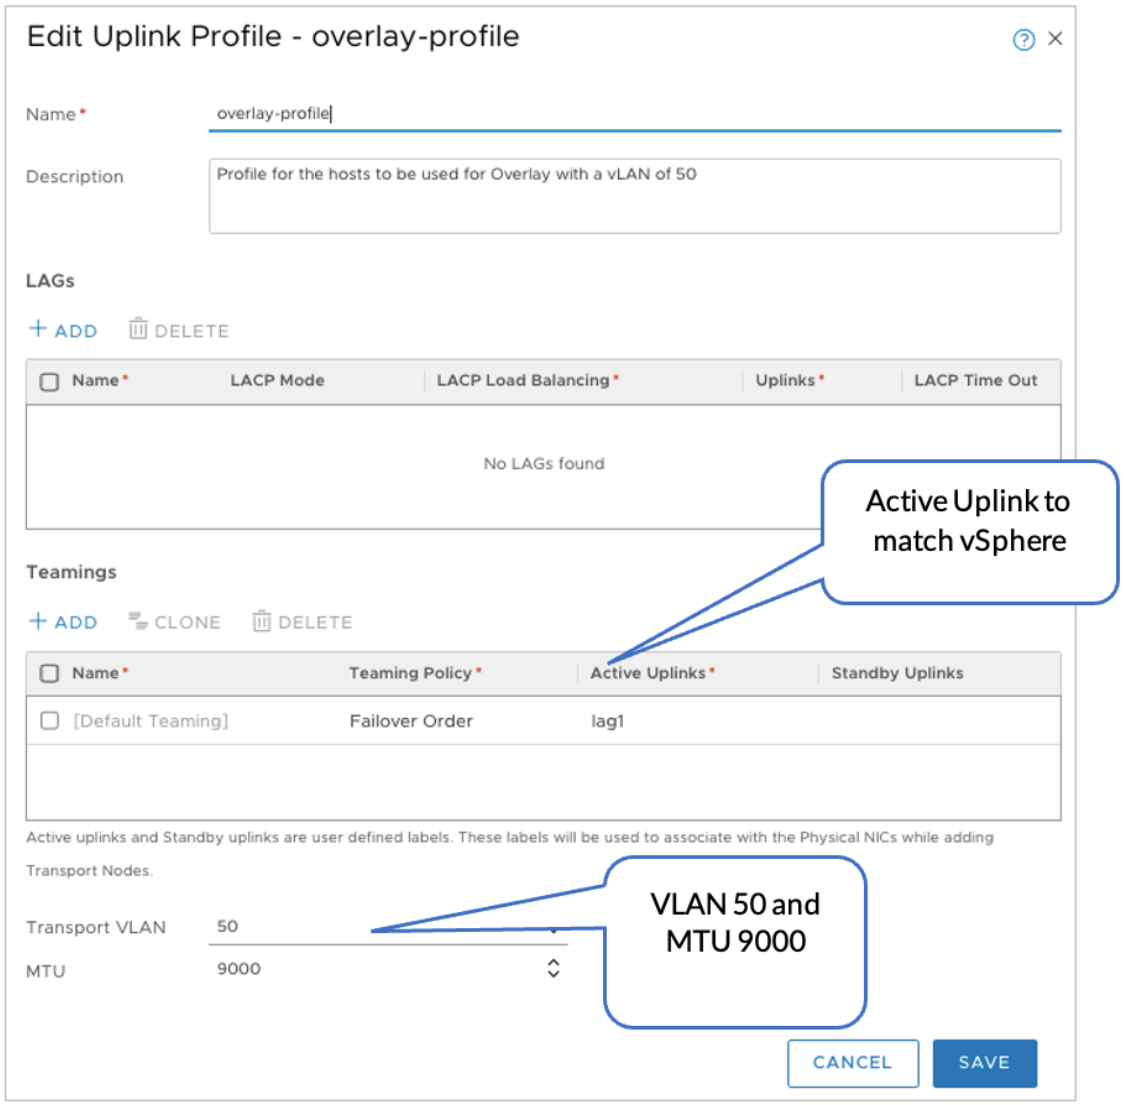


3. Edge-left-uplink profile: VLAN 100 maps to uplink1.

**Figure 25: Edge-Left-uplink Profile with VLAN 100**

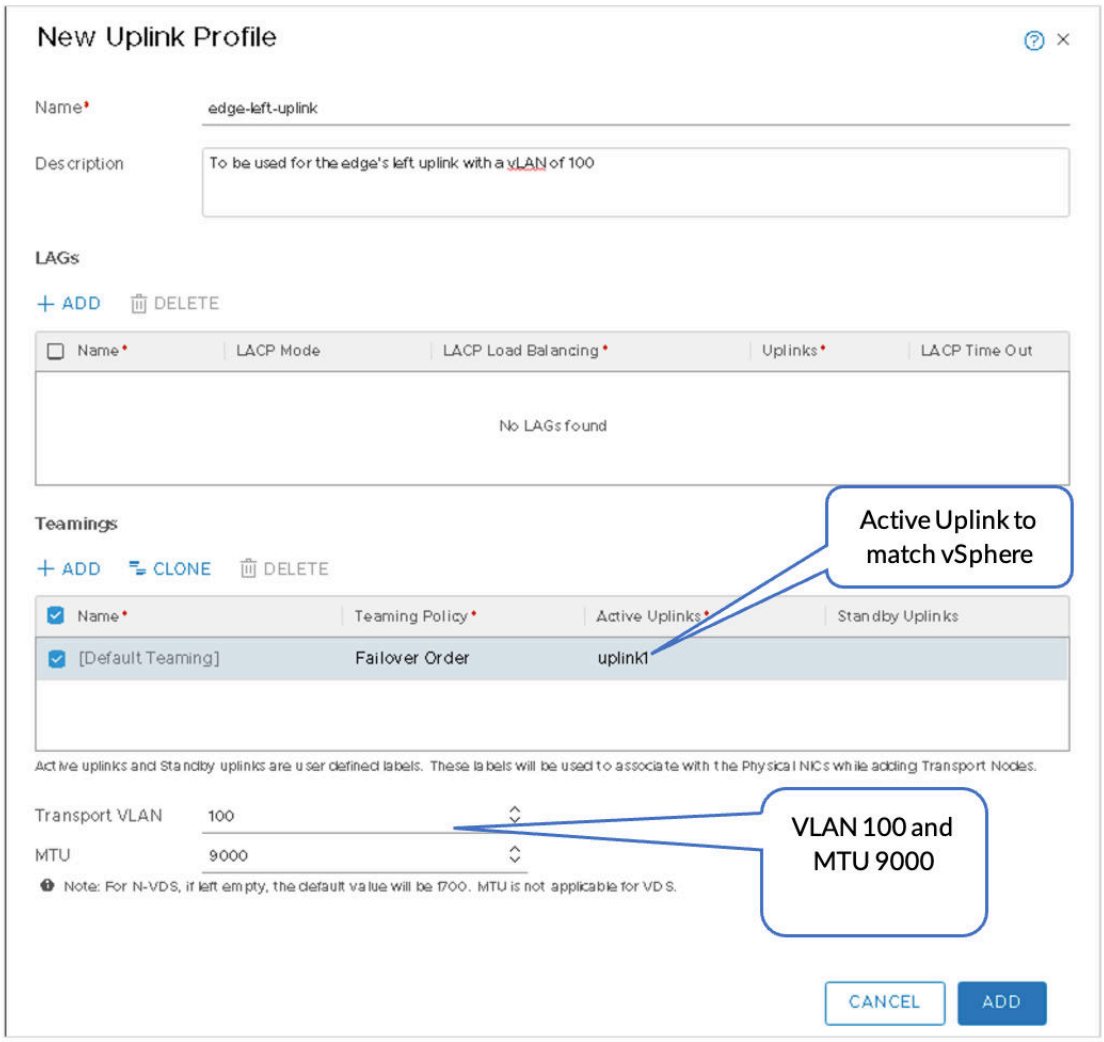


4. Edge-right-uplink profile: VLAN 200 maps to uplink2.

**Figure 26: Edge-Right-uplink Profile with VLAN 200**

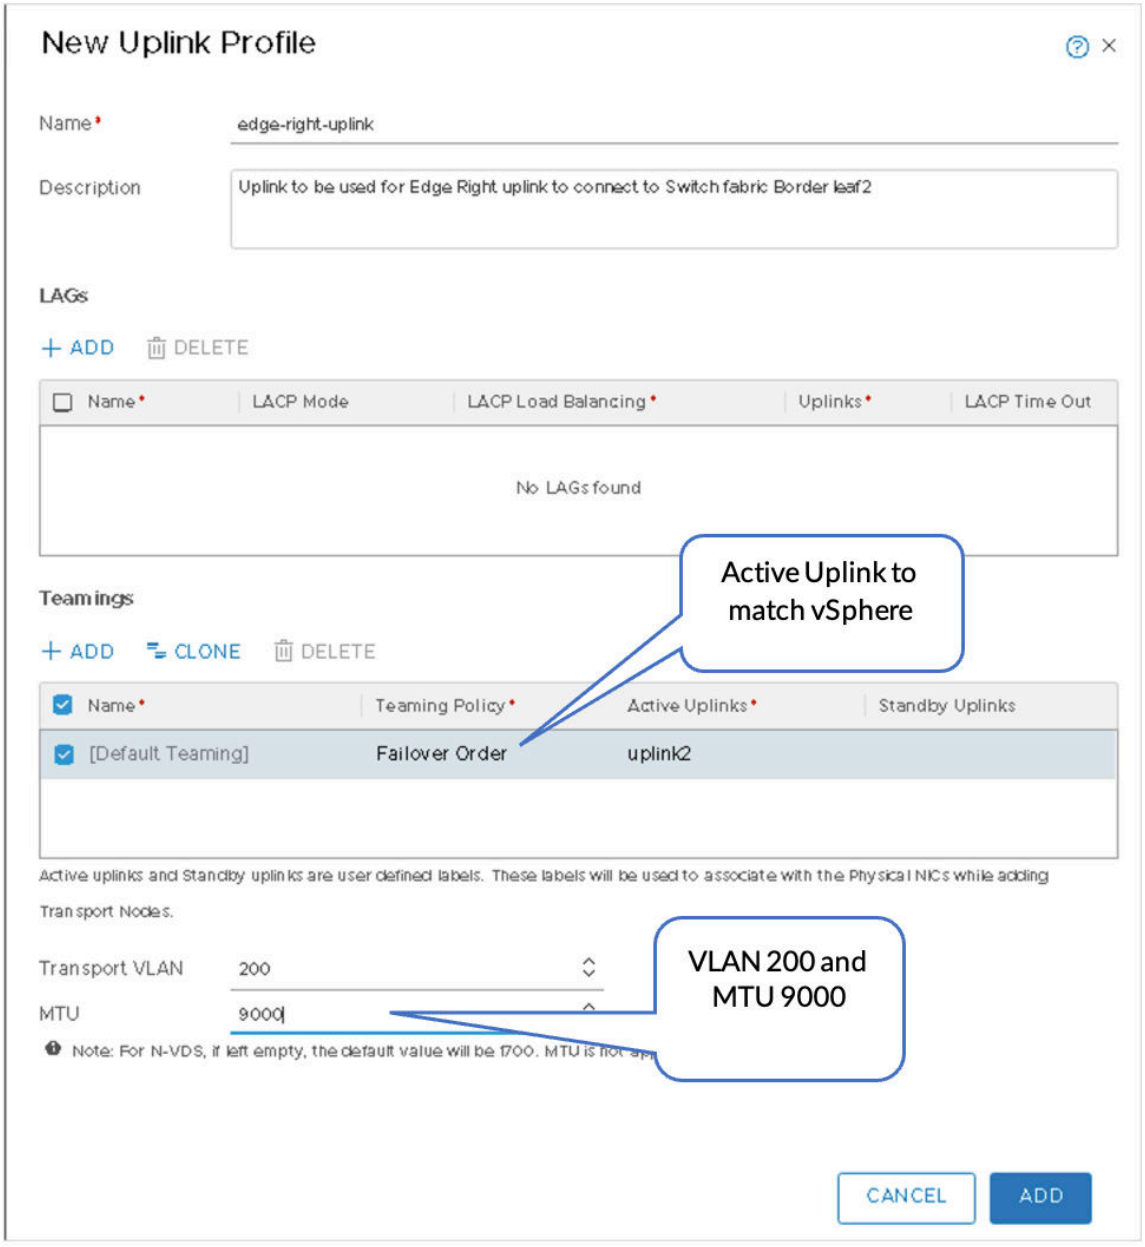

**Figure 27: Uplink Profiles for Overlay, Left and Right Uplink**

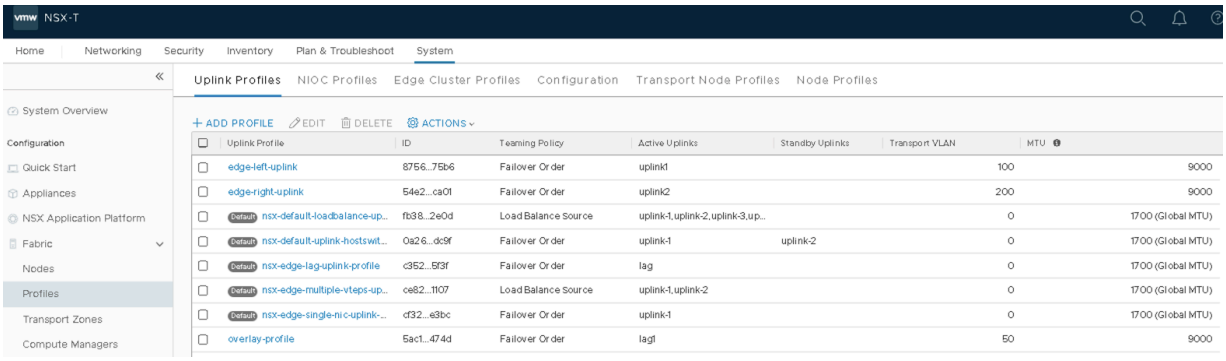


## VMware vSphere: Add Transport Node Hosts to VDS

For transport node hosts, VDS should be configured so that the hosts can connect to the overlay transport network.

In VMware vSphere, under **Networking**, right-click the VDS created in the Juniper Apstra: Add the NSX-Manager section and add all the hosts that form part of the transport node to the VDS. Assign the respective VMNICs that will be used for overlay uplink (LAG link).

**Figure 28: Assign the Physical Adapters on All Hosts to the Uplink**

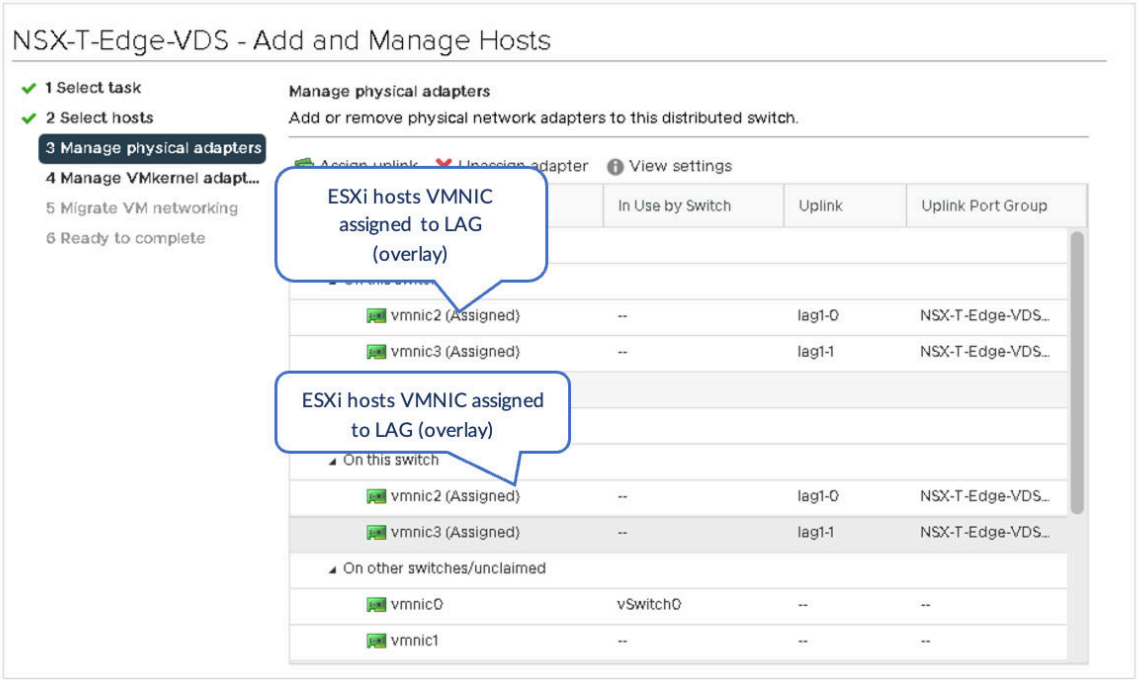


**Figure 29: Configure MTU on NSX-T-Edge-VDS**

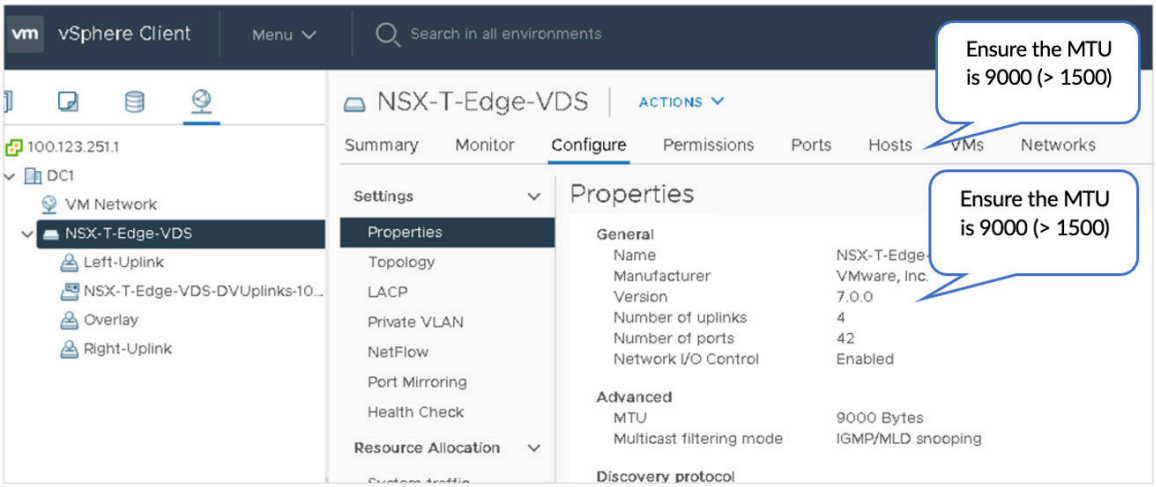

## VMware NSX-T: Prepare the Compute Cluster

For an ESXi host to be part of the NSX-T overlay, it must first be added to the NSX-T fabric.

A fabric node is a node that is registered with the NSX-T management plane and has NSX-T modules installed.

1. In NSX-T Manager, navigate to **System** > **Fabric** > **Nodes** > **Host Transport** Nodes. Select the appropriate vSphere instance from the Managed by list under Host Transport Nodes.

**Figure 30: NSX-Manager Transport Nodes**

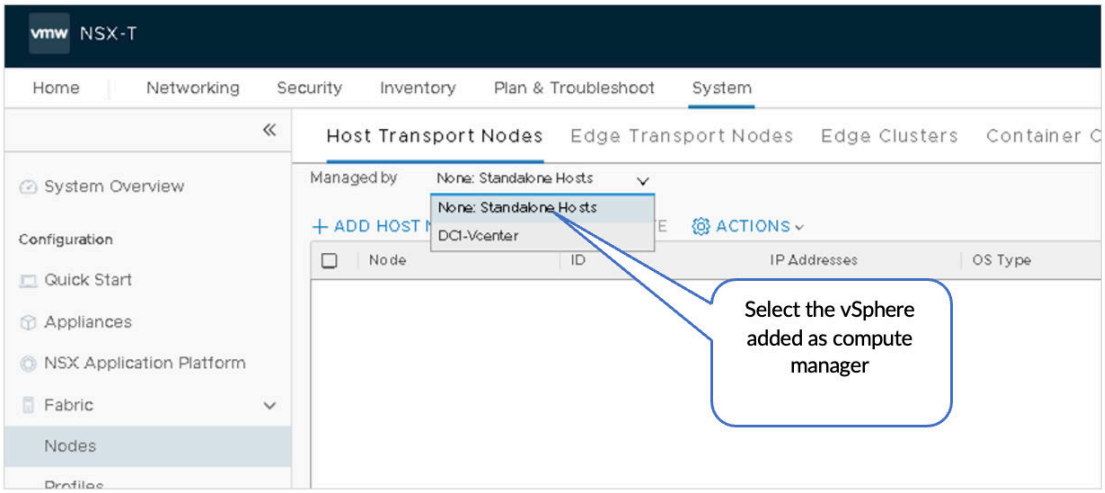


2. Select the ESXi host and then click **Configure NSX**.

**Figure 31: Configure NSX VDS on Transport Nodes**

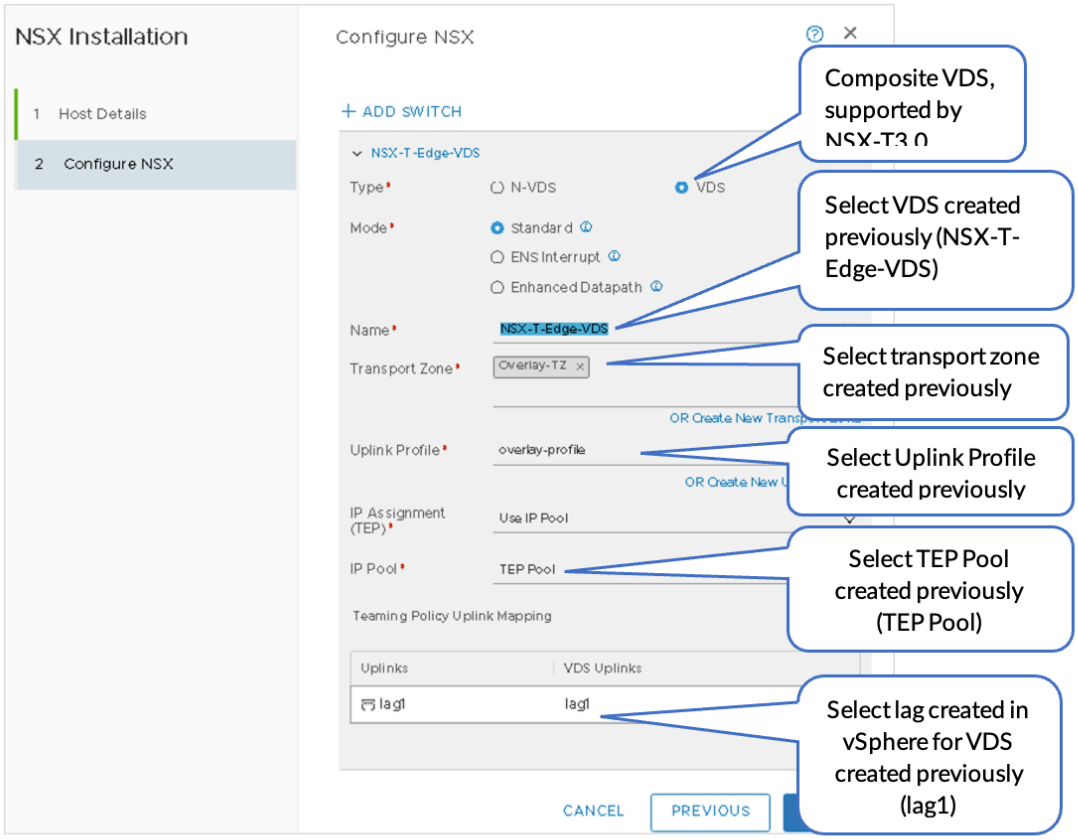

**Figure 32: ESXi Configure NSX**

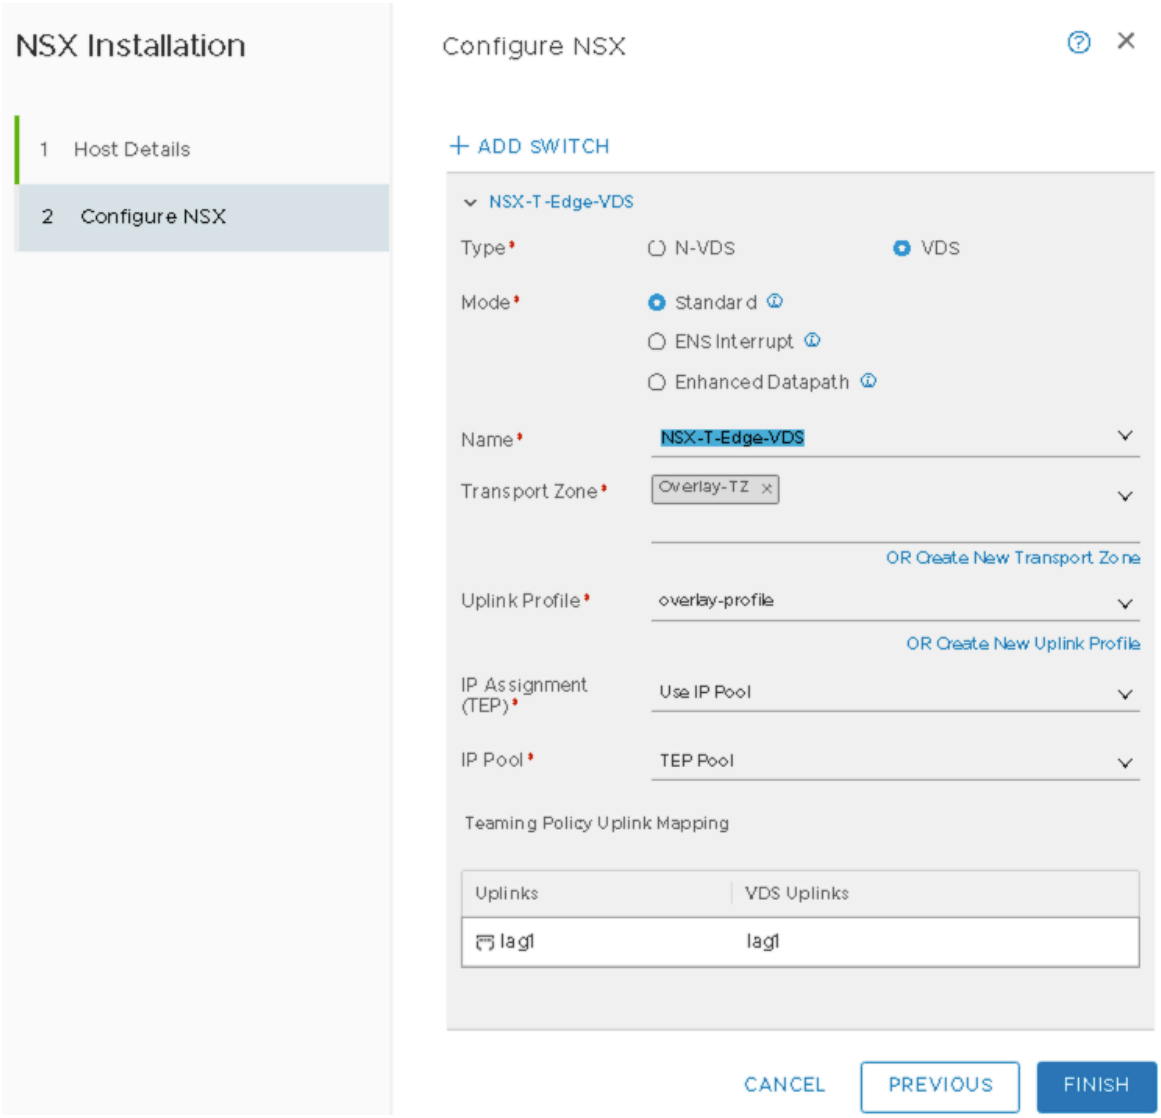


3. Click **Finish**.
4. Repeat steps 1 through 3 for all the ESXi hosts that need to be configured as Transport Nodes of the NSX-T cluster.

## VMware NSX-T: Transport Nodes-Tunnel IPs

After NSX-T is configured on the nodes, the hosts should report the NSX configuration as “Success” and the node status as “Up.” The NSX version will also be displayed.

Remember the TEP IP addresses as they will be required in the later steps. These are the IP addresses assigned to each node. These IP addresses should be set from the TEP Address Pool that was configured earlier.

**Figure 33: TEP Pools assigned to ESXi**

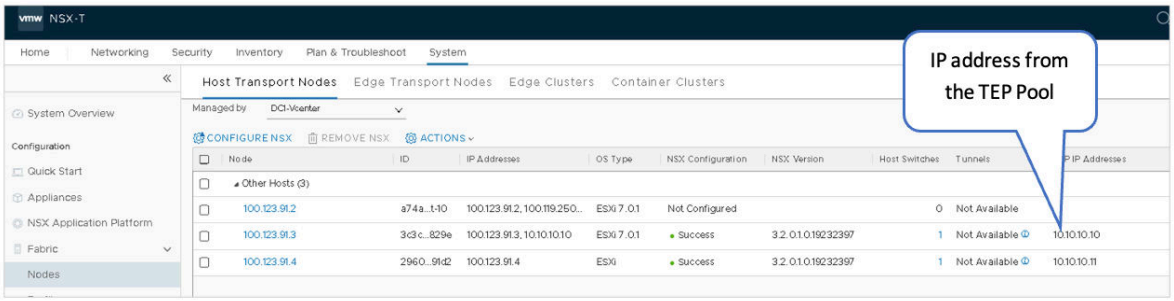


## VMware NSX-T: Deploy NSX Edge Node and Create Edge Cluster

Next, the NSX Edge VM must be created. This will be used for north-south communication and BGP peering with the fabric.

- **Create the Edge VM**

1. Log on to NSX-T and navigate to **System** > **Fabric** > **Edge transport Nodes**.
2. Click **+ADD EDGE NODE**.

**Figure 34: NSX-T Edge Transport Node**

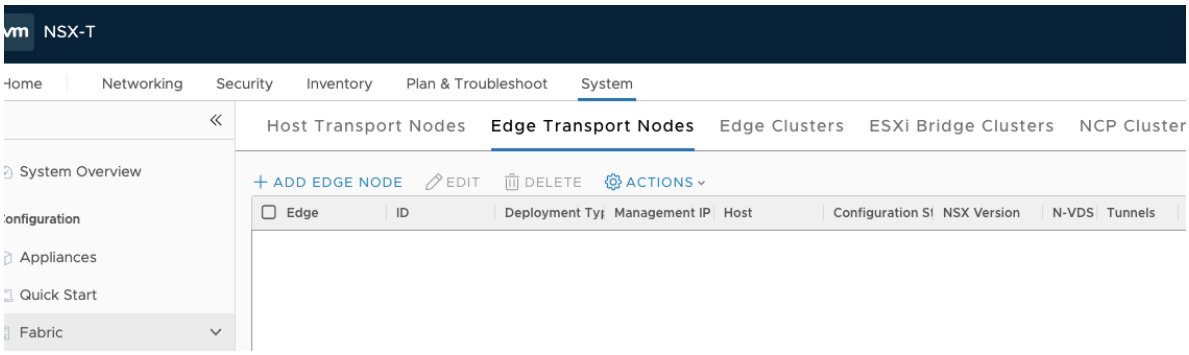


3. Name the Edge VM edge01.

**Figure 35: Adding Edge Node**

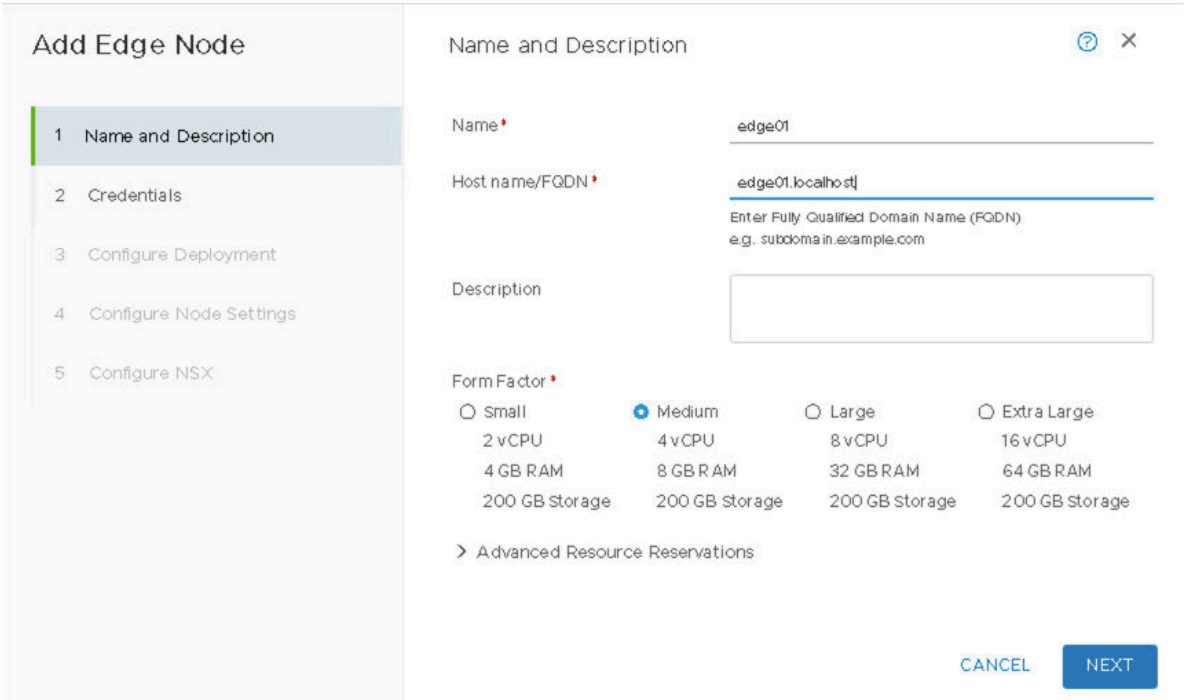


4. Enter the Credentials for the NSX Edge VM.
   Note down the credentials to use them in the later steps.

**Figure 36: Adding Credentials for NSX Edge VM**

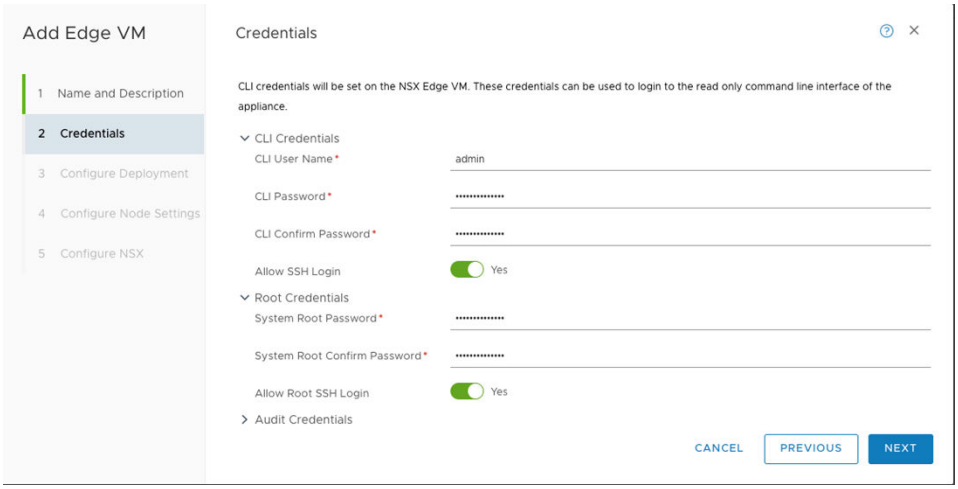


5. In the next step, Select the Compute Manager, Cluster, and Datastore to deploy the Edge VM.

6. Next, configure the Edge VM Network Settings, such as the management IP, default gateway IP, DNS, and NTP server for the edge node.
   - Within vSphere, a VDS was created With three Uplinks. Create one NSX-T VDS for each of the uplinks in the vSphere VDS in NSX-T, Overlay, Edge-Right, and Edge-Left VDS.
     - Name the first NSX-T VDS:
       a. Name the first NSX-T VDS as **nvds-overlay**.
       b. Set transport zone to **Overlay-TZ**.
       c. Set the Uplink Profile to **overlay-profile**.
       d. Select **TEP-Pool** for the IP Pool.

**Figure 37: NVDS Overlay for Edge Node**

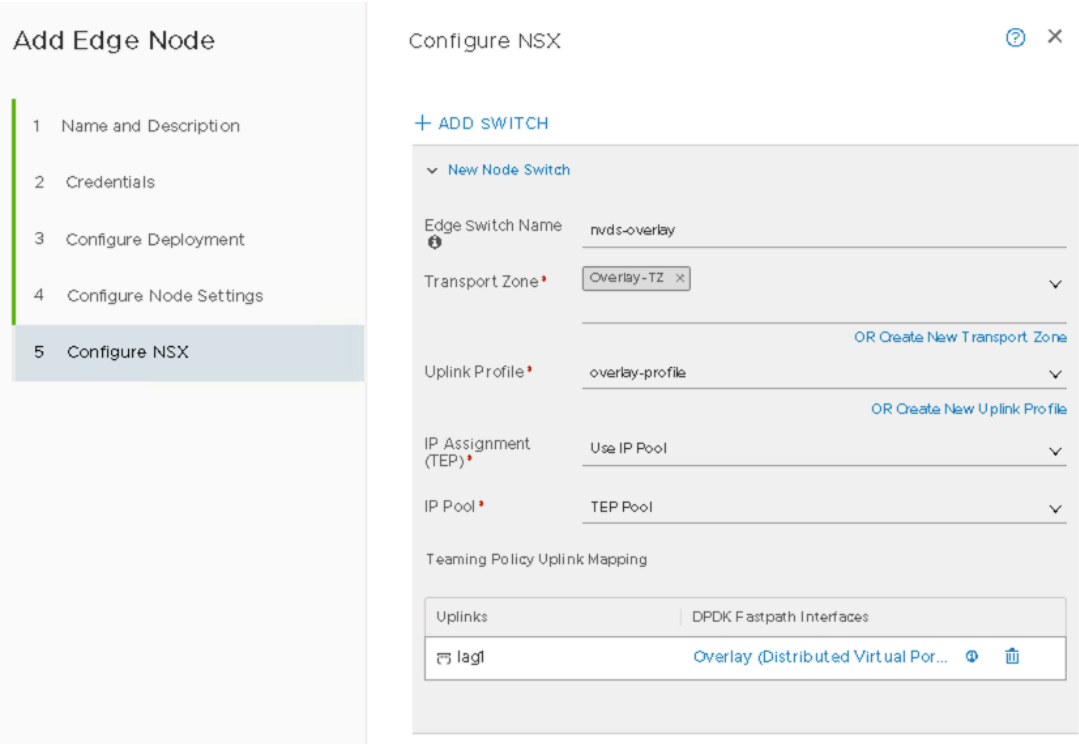


   - Name the Second NSX-T VDS:
     a. Name the second NSX-T VDS as nvds-left.
     b. Set the Transport Zone to Uplink-TZ.
     c. Set the Uplink Profile to edge-left-uplink.

**Figure 38: NVDS Left for Edge Node**

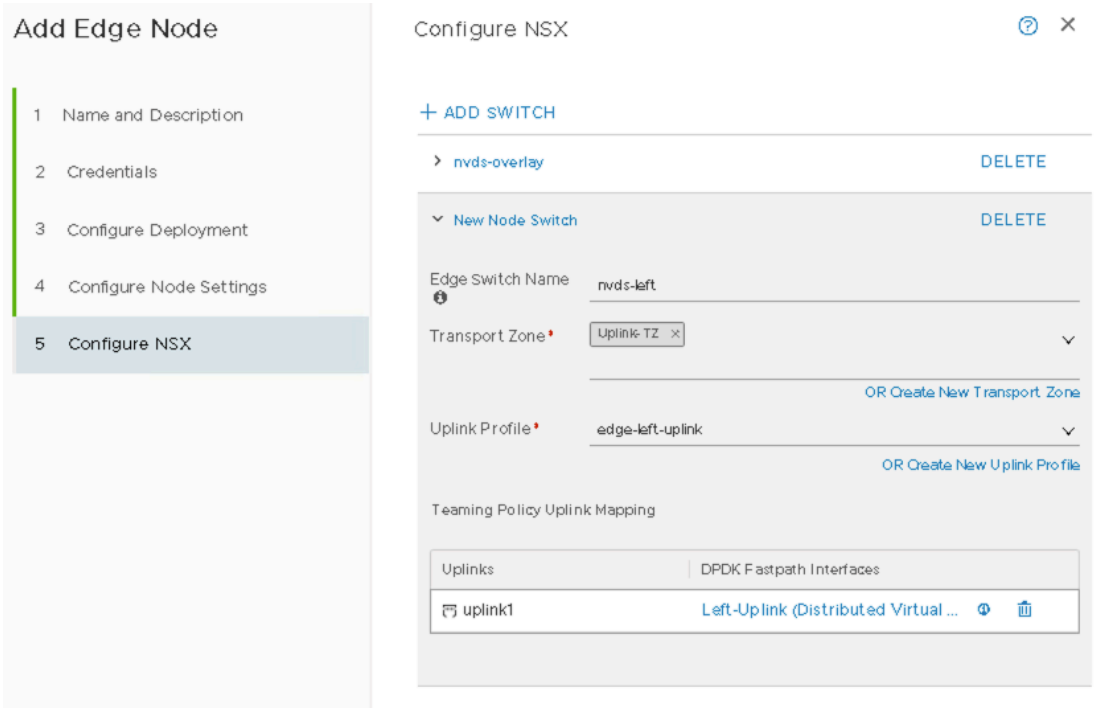


- Name the Third NSX-T VDS:
  a. Name the third NSX-T VDS as nvds-right.
  b. Set the Transport Zone to Uplink-TZ.
  c. Set the Uplink Profile to edge-right-uplink.

**Figure 39: NVDS Right for Edge Node**

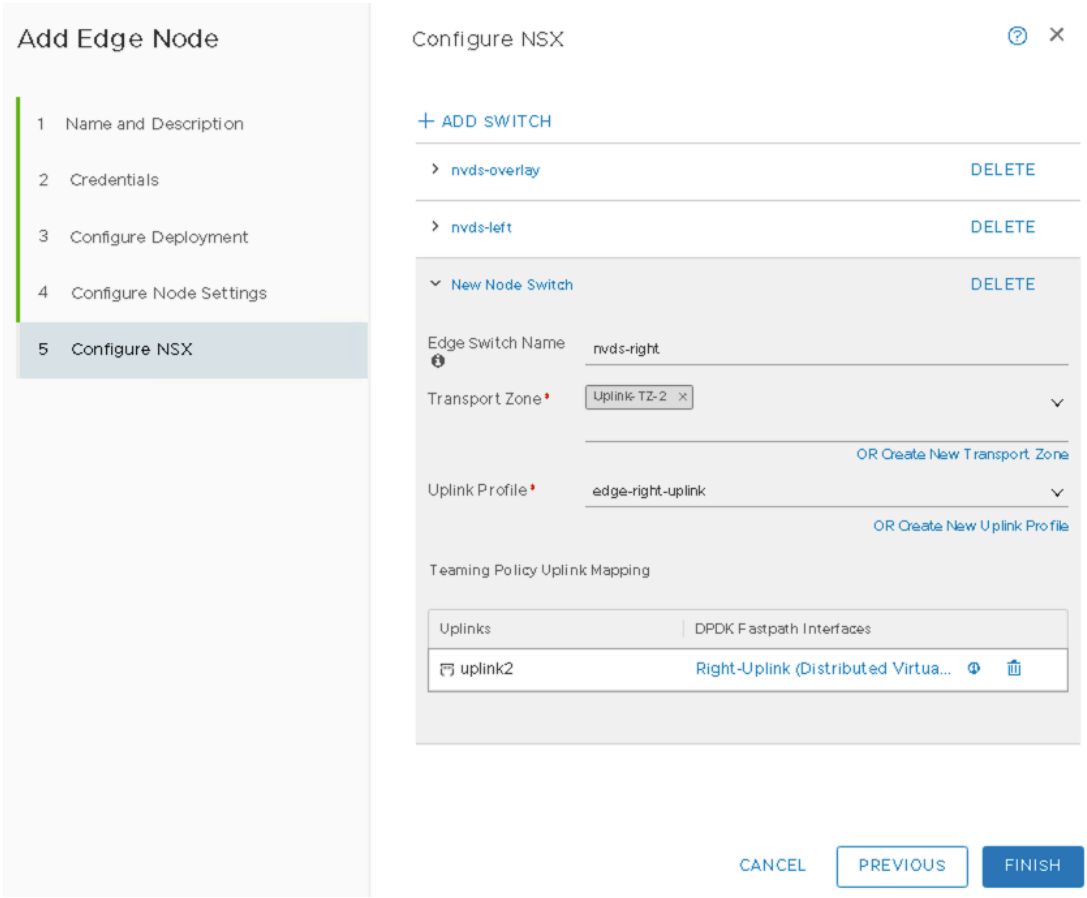


- **Verify NSX-T Edge Creation**

A successful message is displayed once the NSX-T Edge is created. The TEP IP address must be assigned from the previously configured TEP pool.

- **Add Edge Cluster**

1. In NSX-T Manager, navigate to **System** > **Fabric** > **Nodes** > **Edge Clusters** > **ADD.**
2. Name the cluster as **edge-cluster1**.
3. Select nsx-default-edge-high-availability-profile for the Edge Cluster Profile.
4. Select **Edge Node** for the Member Type **edge01** in edge-cluster1.

**Figure 40: Edge01 Node Added to Edge-cluster1**

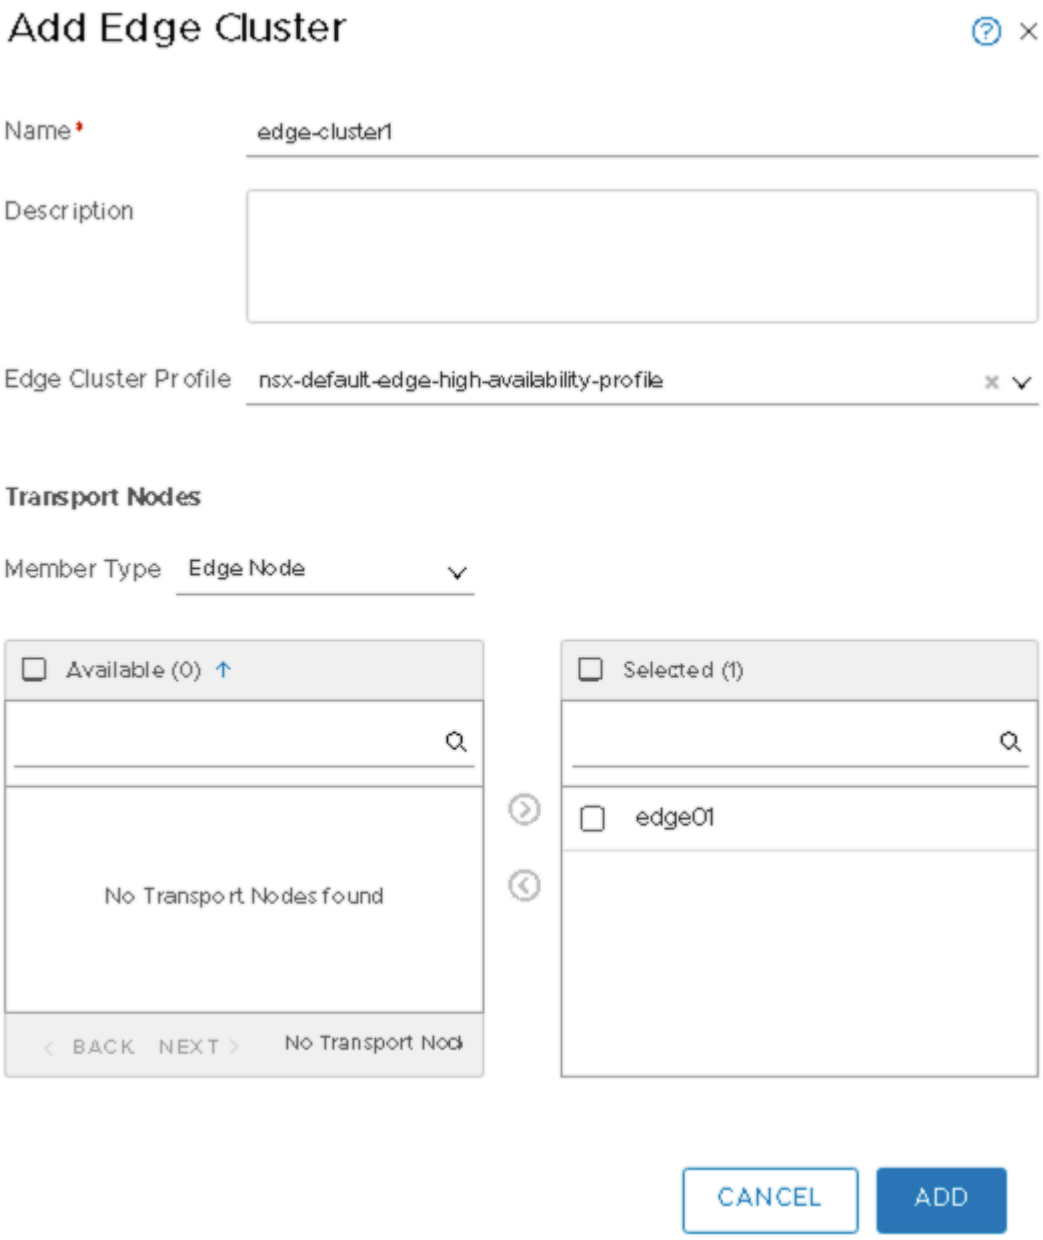

**Figure 41: Edge01 Node Added to Edge-Cluster1**

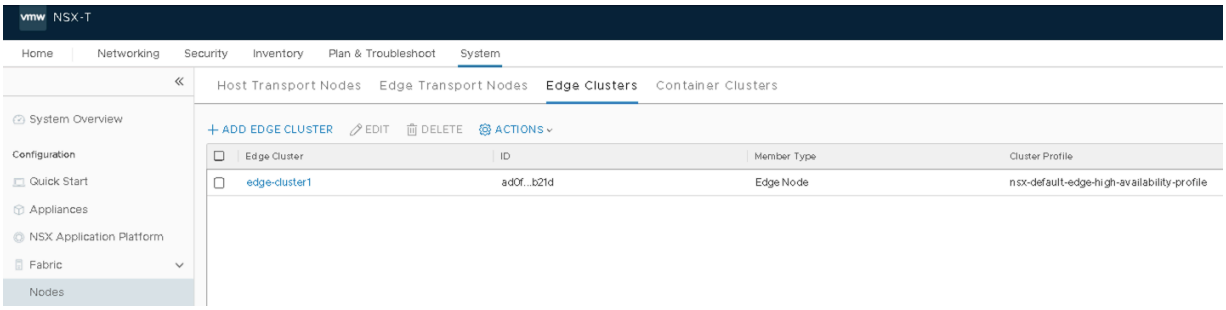


- **Verify Edge Cluster**

To verify that SSH connectivity exists, and the credentials are set up correctly, SSH to the edge node and login.

**Figure 42: SSH connectivity to Edge Node**

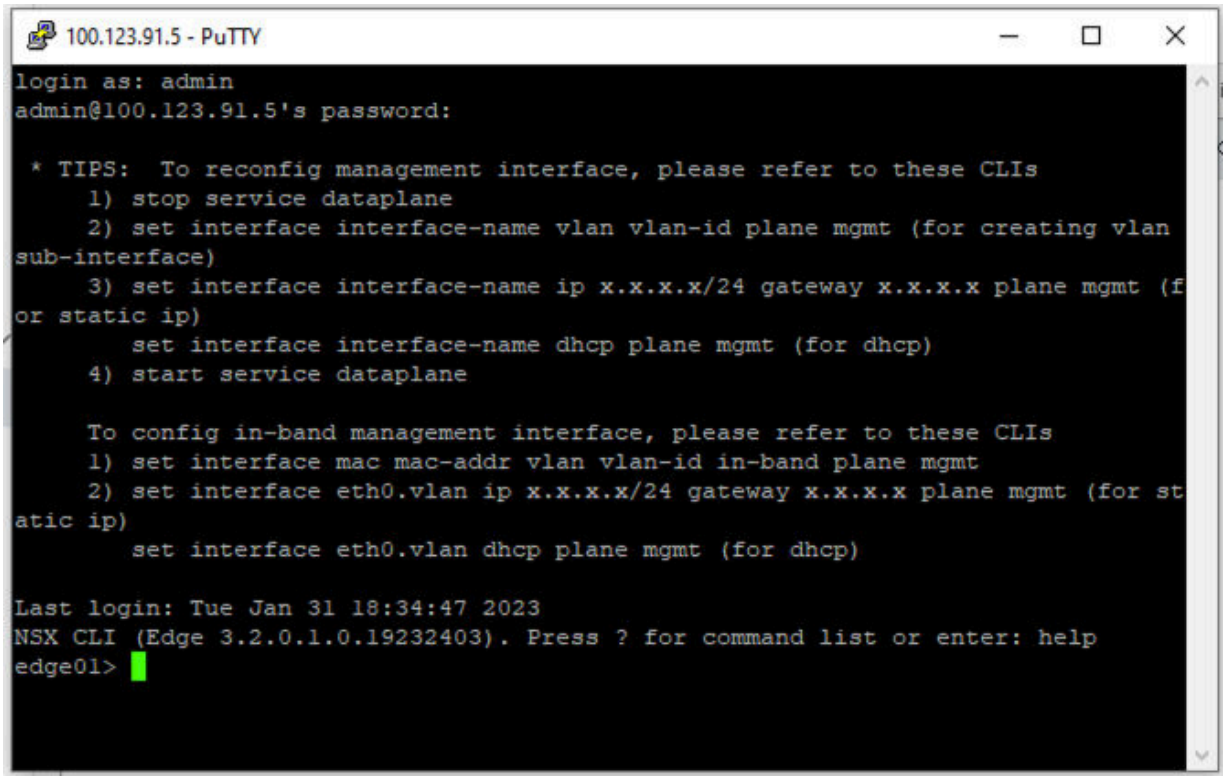

## VMware NSX-T: Create a T-1 Gateway

A Tier-1 Gateway is a logical router that provides East-West communication between VMs in the NSX-T domain.

### Create a Tier-1 Gateway

In NSX-T Manager, navigate to **Networking > Connectivity > Tier-1 Gateways > ADD TIER-1 GATEWAY** and then enter the gateway name as **T1-1**.

**Figure 43: NSX-T Tier-1 Gateway**

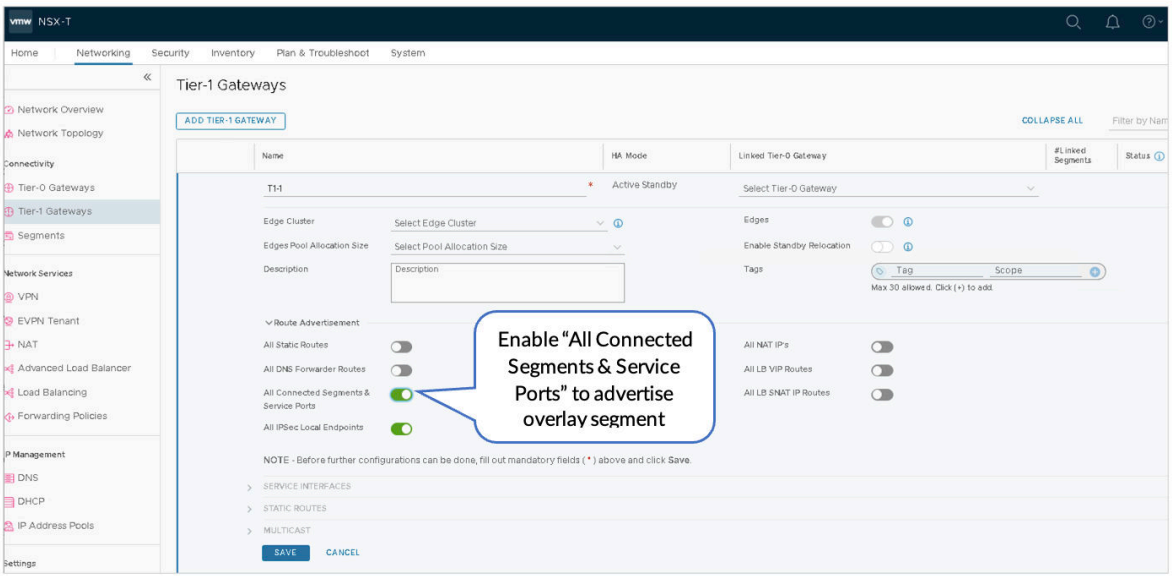


## VMware NSX-T: Create Logical Segments

Segments are virtual L2 networks, and VMs are launched in segments. Segments are connected to the T1 Gateway to enable connectivity between the VMs.

1. In NSX-T Manager, navigate to **Networking** > **Connectivity** > **Segments** > **Segments** > **ADDSEGMENT**.
2. Name the first segment as **vn11**.
3. Select **T1-1 | Tier1** under Connected Gateway list to designate Tier 1 Gateway.
4. Select **Overlay-TZ** under Transport Zone to specify the transport zone for the overlay.

5. Enter the subnet to be used: **10.9.11.1/24**.

**Figure 44: Create Logical Segments**

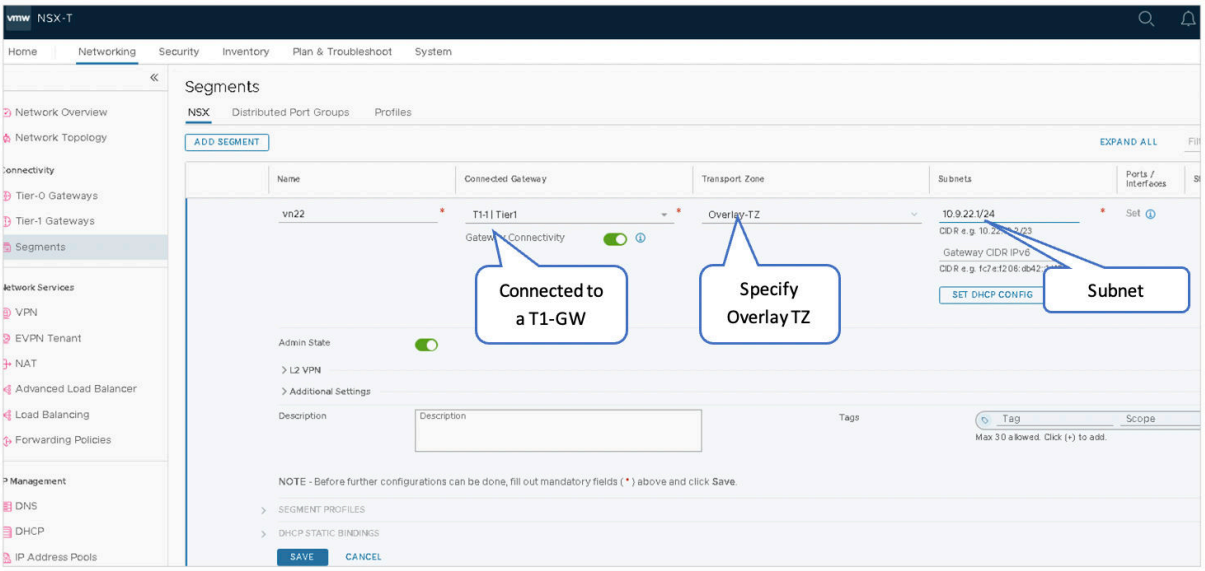


## Add Another Segment

1. Name the second segment as **vn22**.
2. Select **T1-1 | Tier1** under Connected Gateway to designate Tier 1 Gateway.
3. Select **Overlay-TZ** under Transport Zone to specify the transport zone for the overlay.
4. Enter the subnet to be used: **10.9.22.1/24**.

**Figure 45: Create Logical Segments**

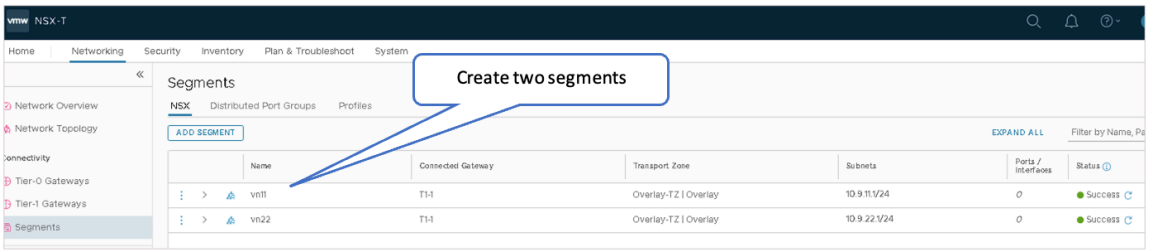

## VMware NSX-T: Create VLAN Backed Logical Segments

A VLAN-backed segment enables the Tier-0 gateway to establish BGP sessions with the fabric. The VLAN-backed segment serves as the North-South data path of the VMs in NSX to/from the rest of the Data Center Fabrics.

- **Create One VLAN-Backed Segment for Each Uplink:**

1. In NSX-T Manager, navigate to **Networking** > **Connectivity** > **Segments** > **Segments** > **ADD SEGMENT**.
2. Name the first segment as **uplink-seg-100**.
3. Do not select a gateway under Connected Gateway.
4. Select **Uplink-TZ** under Transport Zone to specify the transport zone for the overlay.
5. Do not enter a subnet to be used.
6. Under VLAN, associate **VLAN 100**.

- **Add Another VLAN Segment**

1. Name the second segment as **uplink-seg-200**.
2. Do not select a gateway under Connected Gateway.
3. Select **Uplink-TZ** under Transport Zone to specify the transport zone for the overlay.
4. Do not enter a subnet to be used.
5. Under VLAN, associate **VLAN 200**.

**Figure 46: Create Uplink Segments for Left and Right Links to Border Leaf Switches**

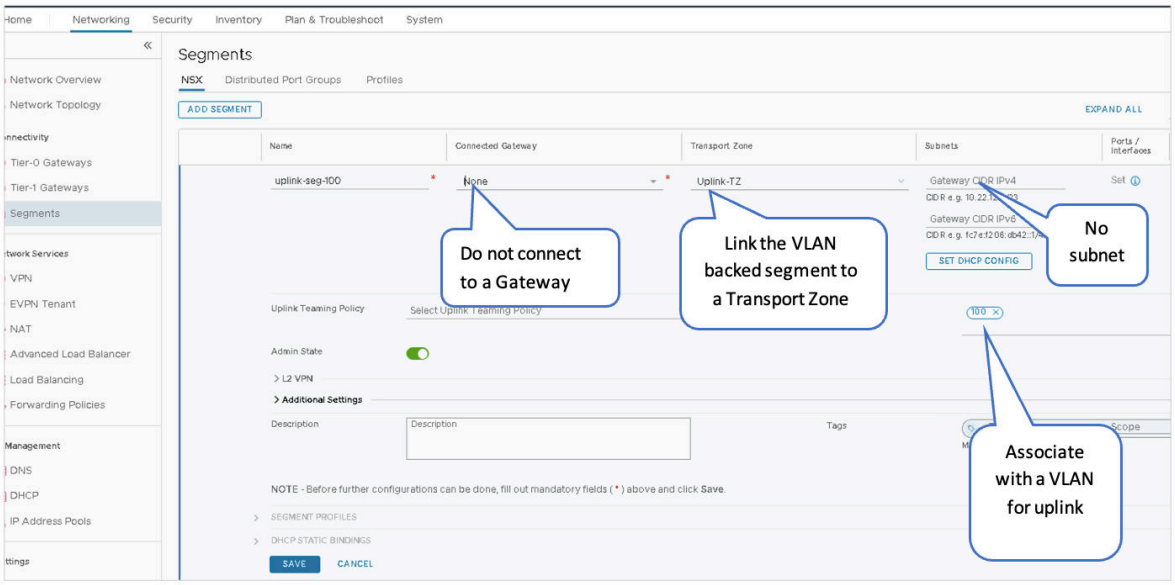


**Figure 47: Uplink Segments for left and right links to Border Leaf Switches**

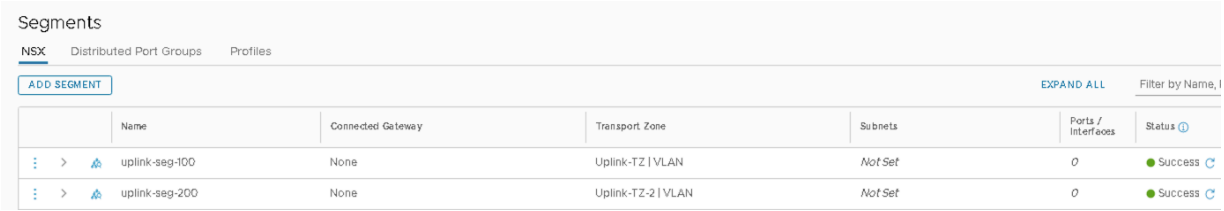


## VMware vSphere: Confirm the Creation of the Logical Segments

The Logical Segments created in the previous steps should be reflected in the vSphere client. Verify that logical segments created in NSX-T are present in the vSphere client. In the vSphere Client, navigate to **Distributed vSwitch > Configure > Topology.**

## VMware vSphere: Create VMs in the Segments

Create two test VMs on each of the Transport Nodes in the cluster:

1. Connect the first VM on each Transport Node to the vn11 logical segment on NSX-T-Edge-VDS, which will allow testing of the vn11 overlay segment for that Transport Node.
2. Connect the second VM on each Transport Node to the vn22 logical segment on NSX-T-Edge-VDS, which will allow testing of the vn22 overlay segment for that Transport Node.

For more information, refer to the VMware vSphere guide for creating a VM and setting up a network adapter.

**Figure 48: Overlay Segments Created in NSX-T are Visible in vSphere**

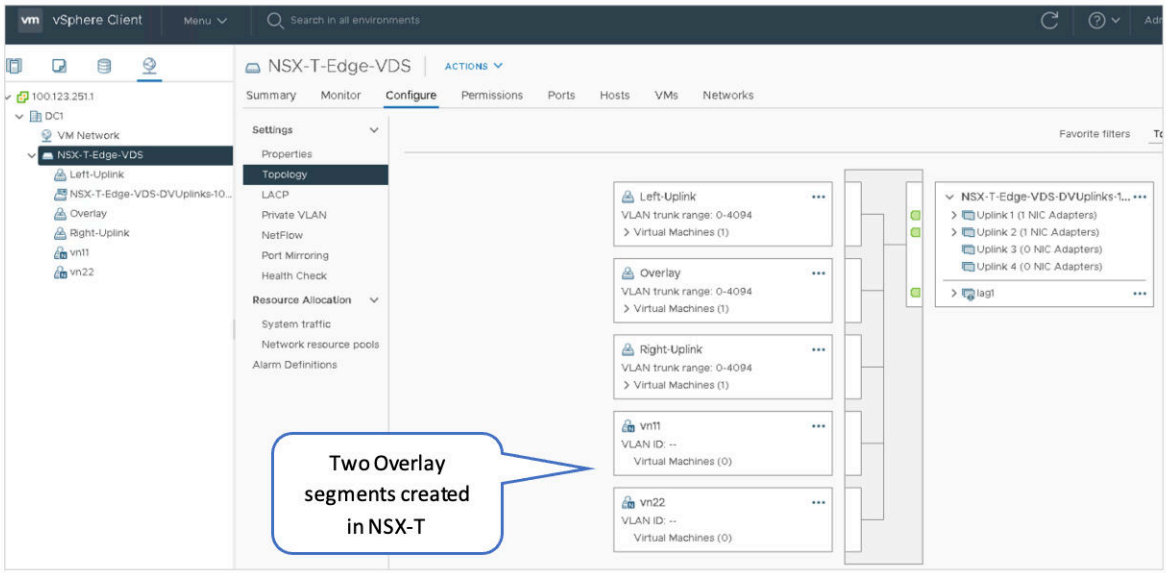

**Figure 49: VMs Created with vn11 and vn22 Ports Connected**

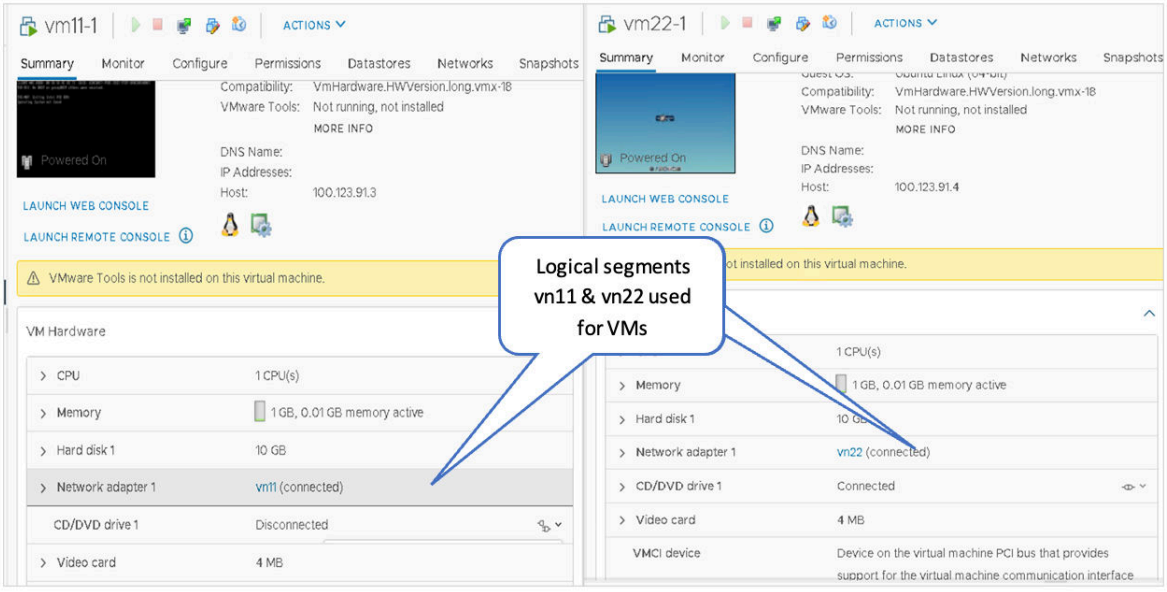


## VMware NSX-T: Create Tier0 Gateways T0-1

A Tier-0 Gateway connects the NSX-T virtual fabric with the physical switch fabric. This is accomplished by using BGP to communicate with the Top of Rack (ToR) switches. In this document, each Transport Node is connected to a pair of QFX-5120 leaf switches, while the Edge VM host (Edge01) is connected to a pair of QFX5130 border leaf switches.

To add a Tier-0 Gateway:

1. In NSX-T Manager, navigate to **Networking** > **Connectivity** > **Tier-0 Gateways** > **ADD GATEWAY**.
2. Name the Tier-0 Gateway as **T0-1**.
3. Set HA-Mode to **Active-Active**.
4. Set the Edge-cluster on T0-1 to **edge-cluster1**.
5. Save and proceed through the next steps to add interfaces, BGP, and route-redistribution.

**Figure 50: Create T0-1 Gateway and Connect to Edge-Cluster**

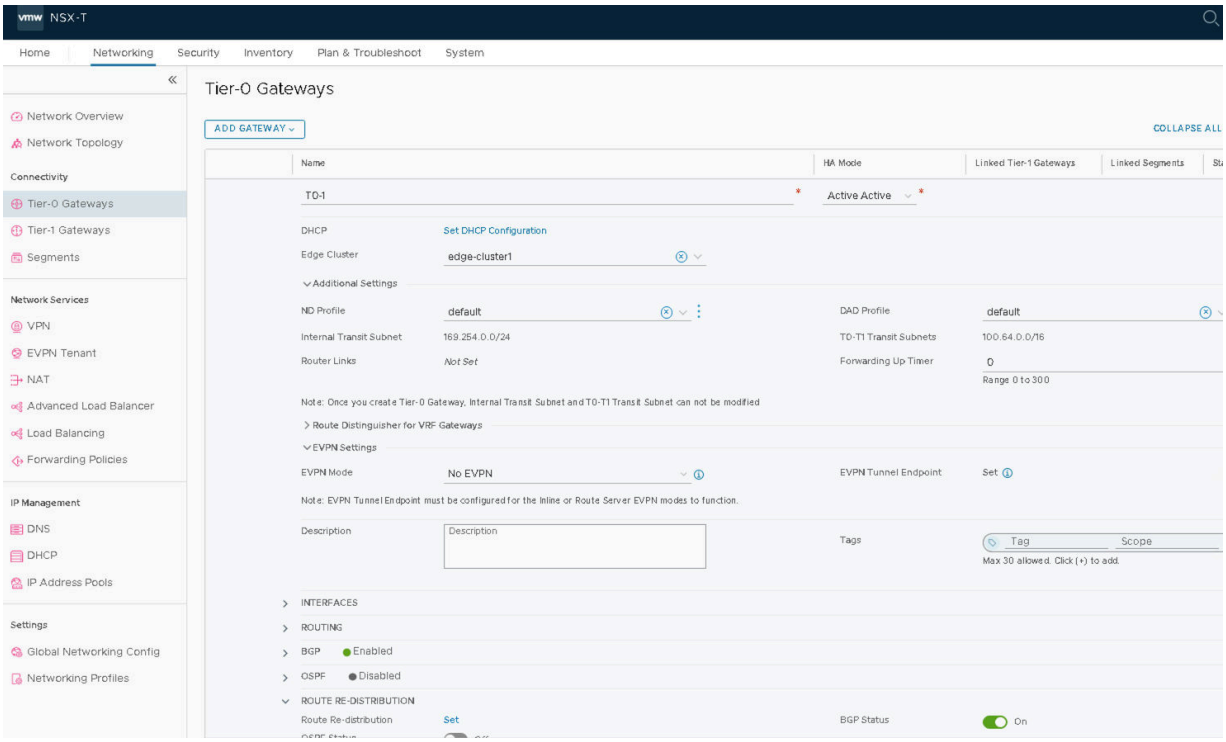


## VMware NSX-T: Configure the Interfaces on the T0 Gateway

Tier-0 Gateway (T0-1 gateway) requires one interface for each uplink segment.

1. In NSX-T Manager, navigate to **Networking** > **Connectivity** > **Tier-0 Gateways** > **Edit T0-1**.
2. To add two external tnterfaces for the fabric within the Tier-0 Gateway screen for the T0-1 Tier-0 Gateway created above:
   a. Click **Set**.
   b. Click **ADD INTERFACE** to add the first interface and configure the following:
      i. Name the interface as left-uplink.
      ii. Set type **External**.
      iii. Set the IP Address/Mask as **192.168.100.2/24**.
      iv. Connect to segment as **uplink-seg-100**.
      v. Set the Edge Node as **edge01**.

c. Click **ADD INTERFACE** to add the second interface and configure the following:

i. Name the interface as right-uplink.

ii. Set type **External**.

iii. Set the IP Address/Mask as **192.168.200.2/24**.

iv. Connect to segment as **uplink-seg-200**.

v. Set the Edge Node as **edge01**.

**Figure 51: Add Two Interfaces for T0-1 for Left-Uplink and Right-Uplink**

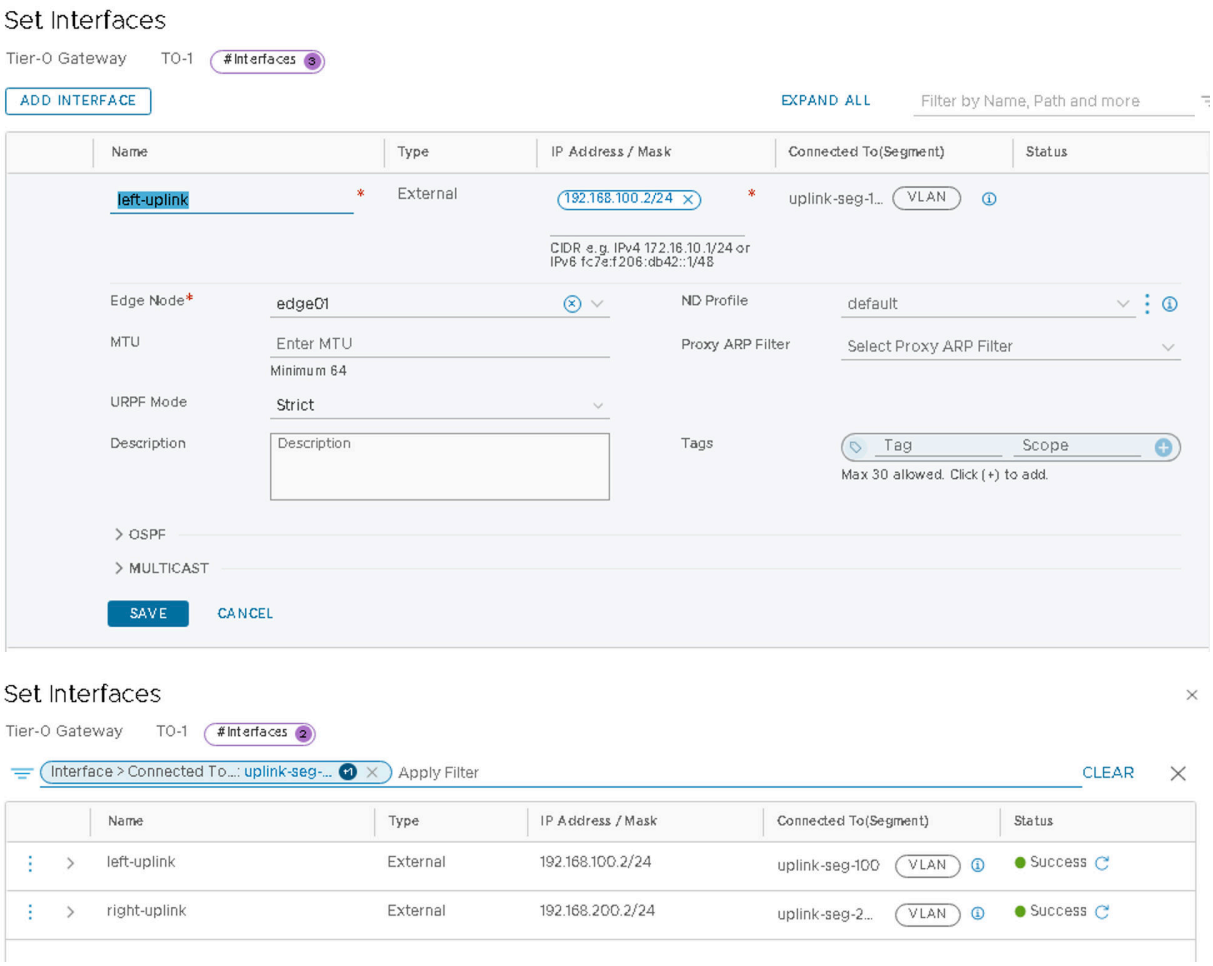

## VMware NSX-T: Configure Loopback Interface on T0 Gateway

The loopback interface is used to create a BGP session with switch Fabric border leaf switches.

1. In NSX-T Manager, navigate to **Networking** > **Connectivity** > **Tier-0 Gateways** > **Edit T0-1**.
2. To add loopback interface towards the Fabric within the Tier-0 Gateway page for the T0-1 Tier-0 Gateway created above:
   a. Click **Set**.
   b. Click **ADD INTERFACE** to add loopback:
      i. Name the loopback interface as **Loopback**.
      ii. Set type **Loopback**.
      iii. Set the IP Address/Mask as **10.0.0.1/32**.

**Figure 52: Configure Loopback Interface to Connect to the Border leaf switches**

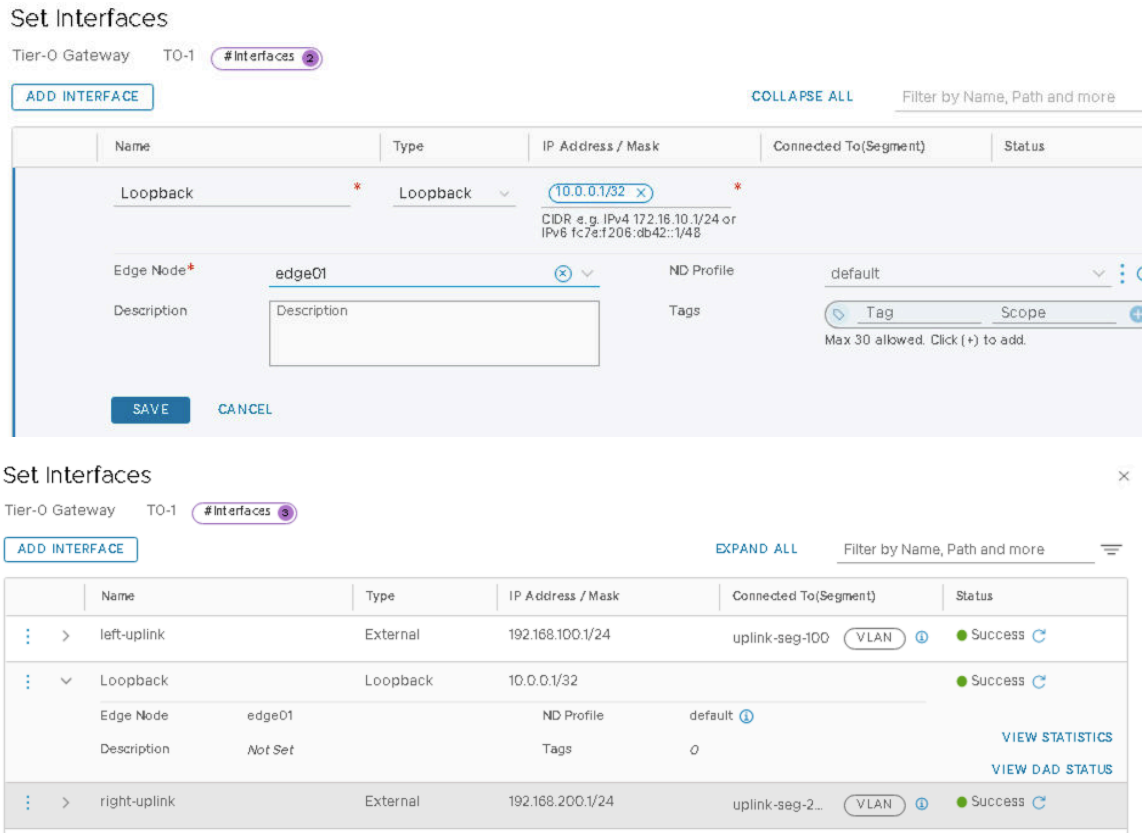

## VMware NSX-T: Configure BGP on the T0 Gateway

In NSX-T Manager, navigate to **Networking** > **Connectivity** > **Tier-0 Gateways** > **Edit T0-1**

**Configure BGP**

1. Within the Tier-0 Gateway page for the T0-1 Tier-0 Gateway created above:
   a. Click **BGP**.
   b. Set Local AS to **65000**.
   c. Enable **BGP**.
   d. Click **Set** next to BGP Neighbors.
   e. Click **ADD BGP NEIGHBOUR** to add the first BGP neighbor.

**Configure the Loopback IP Address of Border Leaf1**

1. In the Juniper Apstra UI navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Physical** > **Nodes**.
   a. Refer to column name Loopback IPv4.
   b. In the following figures, the loopback IP Address from Juniper Apstra is 192.168.255.2, but this can vary.
   c. Set BFD to be **Disabled**.

**Configure the Remote AS Number as Border Leaf1 ASN number**

1. In the Juniper Apstra UI navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Physical** > **Nodes**.
   a. Refer to column name ASN.
   b. In the below figures the ASN from Juniper Apstra is 64514, but this can vary.
   c. Set Route Filter to be **1** with IPV4 Route Filter enabled and Out filter as **prefixlist-out-default**.
   Figure 55 Setting Route Prefix List

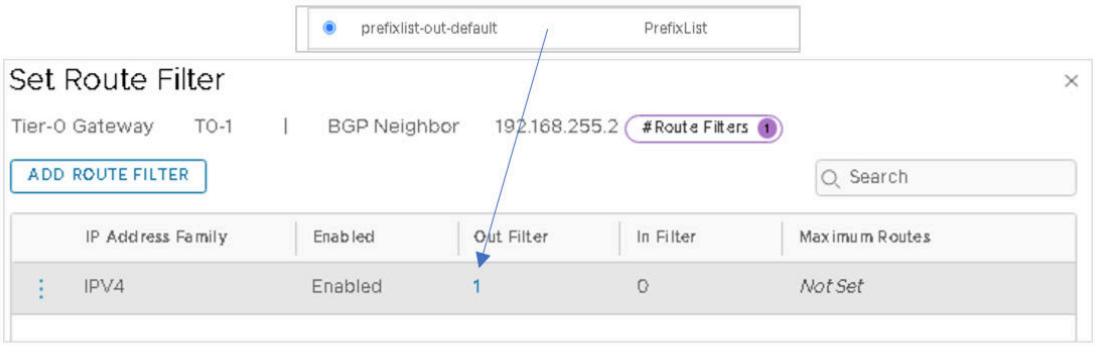


**d.** Set Allowas-in as **Disabled**.

**e.** Under Timers & Password, set Hold Down Time as **90** and Keep Alive Time as **30**.

**Figure 53: Add Border Leaf1 Loopback as the Neighbor**

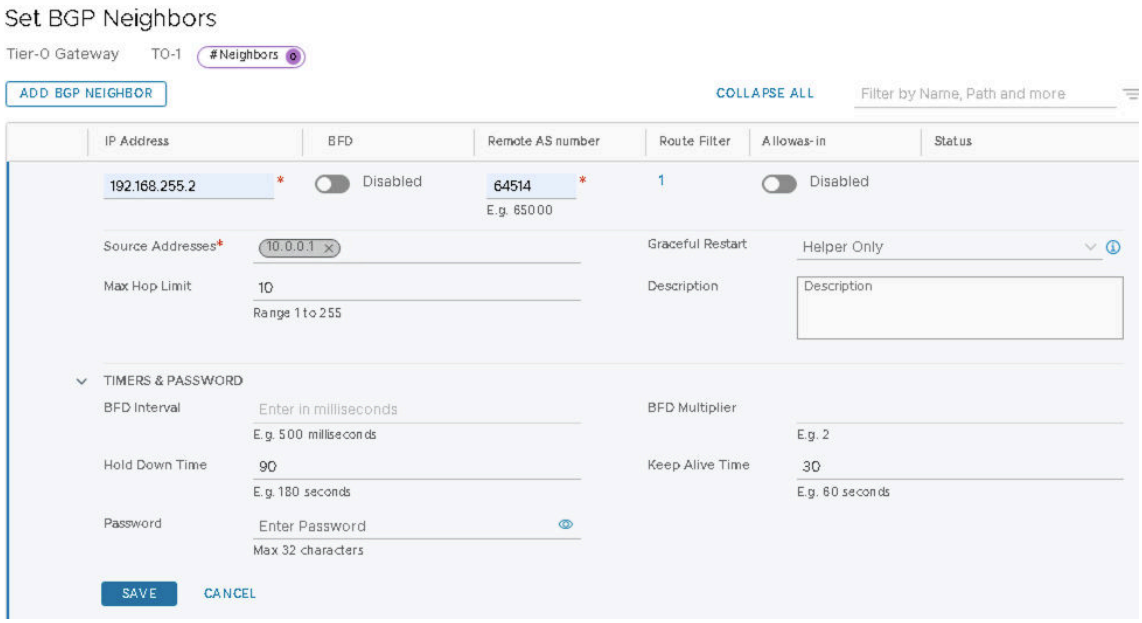


2. Click **ADD BGP NEIGHBOUR** to add the second BGP neighbor.

**Configure the Loopback IP Address of Border Leaf2**:

1. In the Juniper Apstra UI navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Physical** > **Nodes**.
2. In the following figures, the loopback IP Address from Juniper Apstra is 192.168.255.3, but this can vary.

3. Set BFD to be **Disabled**.

**Configure the Remote AS Number as Border Leaf2 ASN Number**:

1. In the Juniper Apstra UI navigate to **Blueprints** > ***<blueprint-name>*** > **Staged** > **Physical** > **Nodes**.
2. Refer to column name ASN.
3. In the following figures, the ASN from Juniper Apstra is 64554, but this can vary.
4. Select **1** for Route Filter.
5. Enable IPV4 Route Filter.
6. Select **prefixlist-out-default** for Out Filter.
7. Set Allowas-in as **Disabled**.
8. Select Max Hop Limit as **10**.
9. Set Hold Down Time as **90** and Keep Alive Time as **30** under Timers & Password:

**NOTE**: BGP status for the two neighbors will be down until Apstra is configured.

**Figure 54: BGP neighbors on T0-1 are the Border Leaf loopbacks**

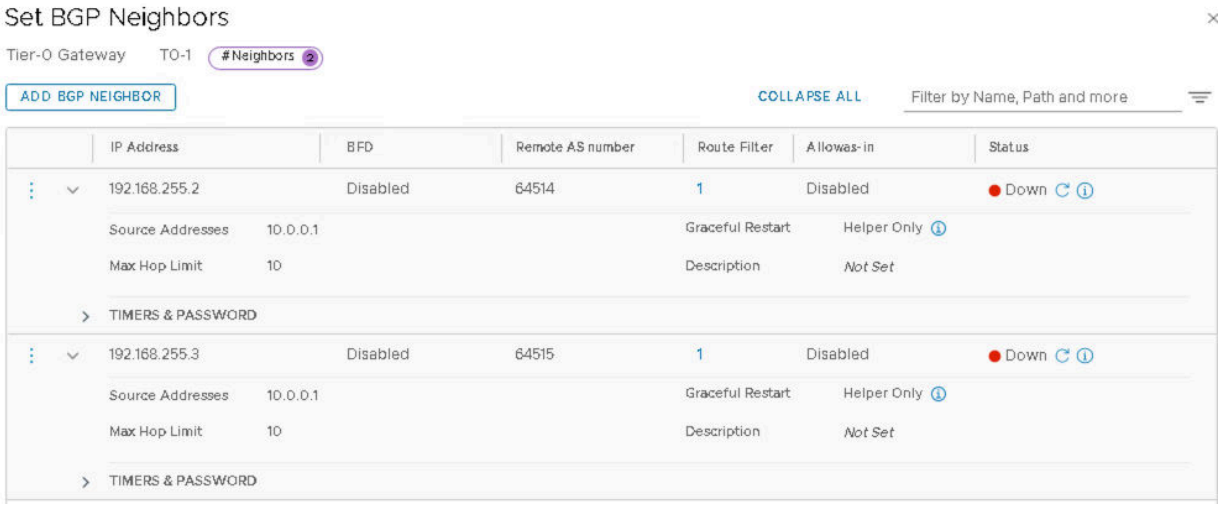

## VMware NSX-T: Configure In-Line Mode and Route-Redistribution on the T0 Gateway

1. In NSX-T Manager, navigate to **Networking** > **Connectivity** > **Tier-0 Gateways** > **Edit T0-1**.
2. Within the Tier-0 Gateway, screen for the T0-1 Tier-0 Gateway created above:
   a. Select EVPN Mode as **inline**.
   b. Click the three vertical dots and create a new VNI-Pool for the EVPN/VXLAN VNI-Pool.

**Figure 55: T0 Gateway EVPN VXLAN VNI Pool**


3. Click **Set** near EVPN Tunnel Endpoint and configure the following:
   a. Name EVPN local tunnel endpoint as **edge-vtep**.
   b. Edge-Node name**: edge01** (created as per VMware NSX-T: Deploy NSX Edge Node and Create an Edge Cluster on page 39).
   c. Local Address**: 10.0.0.1** (this is the loopback address of T0 Gateway as configured in VMware NSX-T: Configure Loopback interface on T0 Gateway).
   d. MTU: **9000**.
   e. Save changes to the Tier-0 Gateway.

**Figure 56: Configuring Local EVPN Tunnel Endpoint**

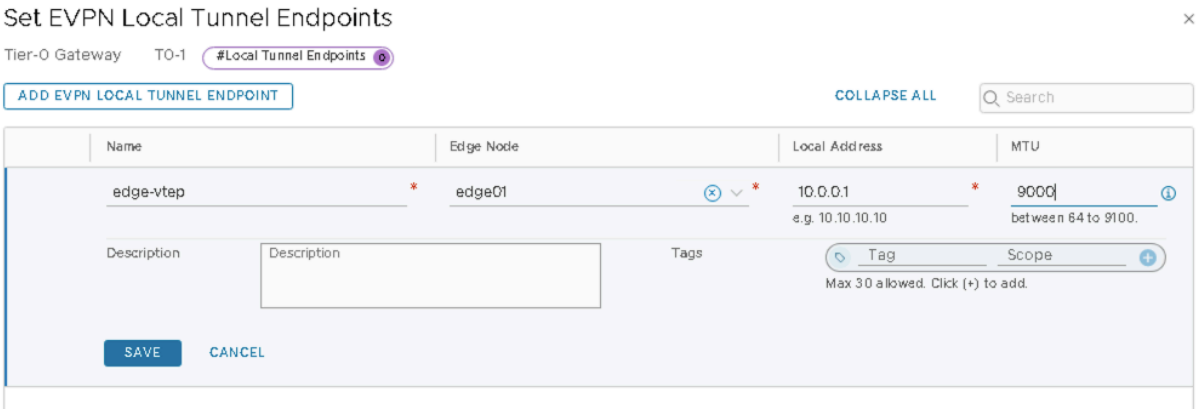


**Figure 57: T0-1 Tier-0 Gateway EVPN Inline Mode**

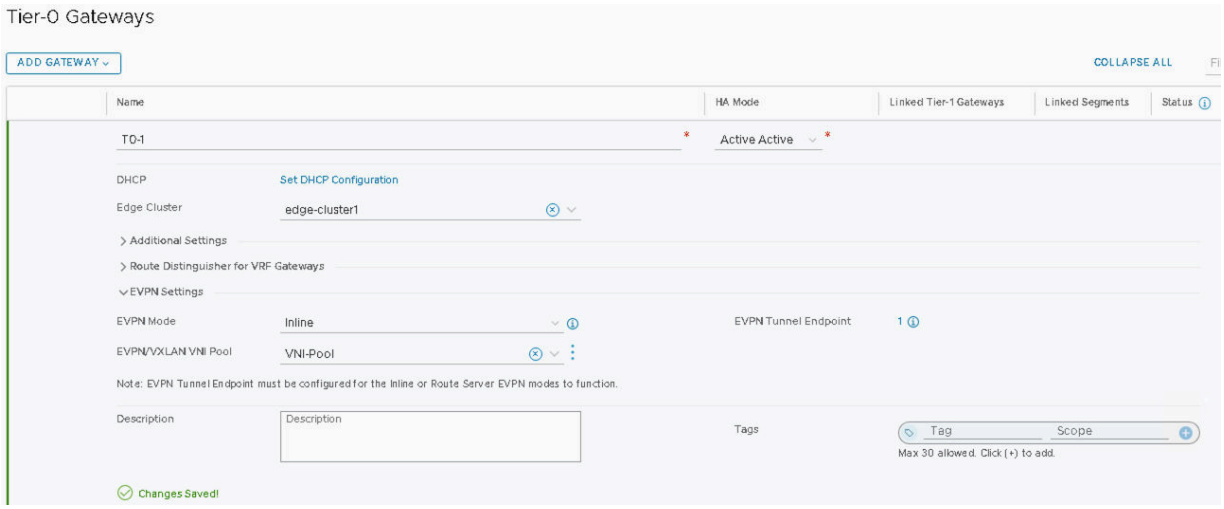


Within the Tier-0 Gateway screen for the T0-1 Tier-0 Gateway created above:

1. Expand and click **Route Re-Distribution**.

**Figure 58: Configure Route Re-distribution**

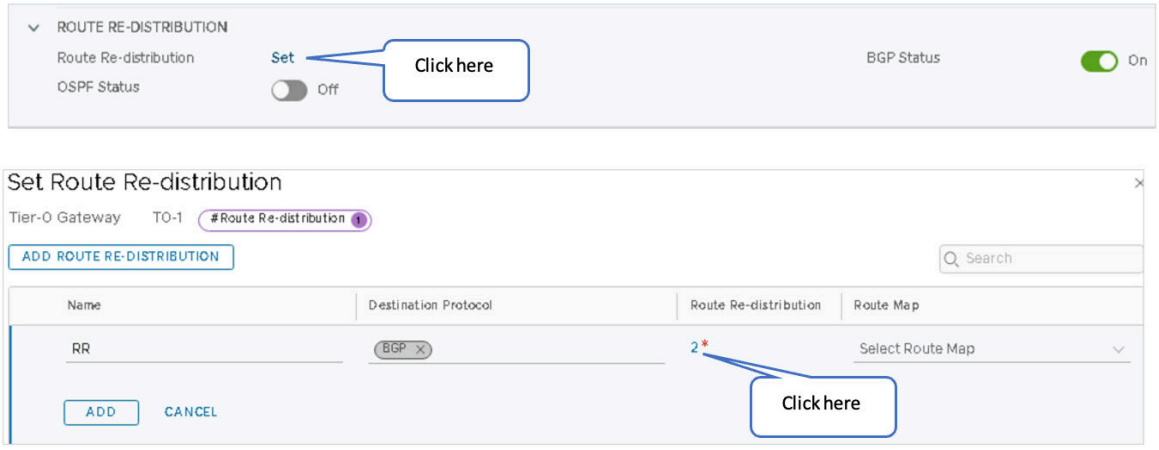


2. Under Advertised Tier-1 Subnets, check only the following sources:
   a. Connected Interfaces & Segments
   b. Service Interface Subnet
   c. Static Routes
   d. Connected Segment

   Ensure no other boxes are checked.

**Figure 59: Set Route-Redistribution**

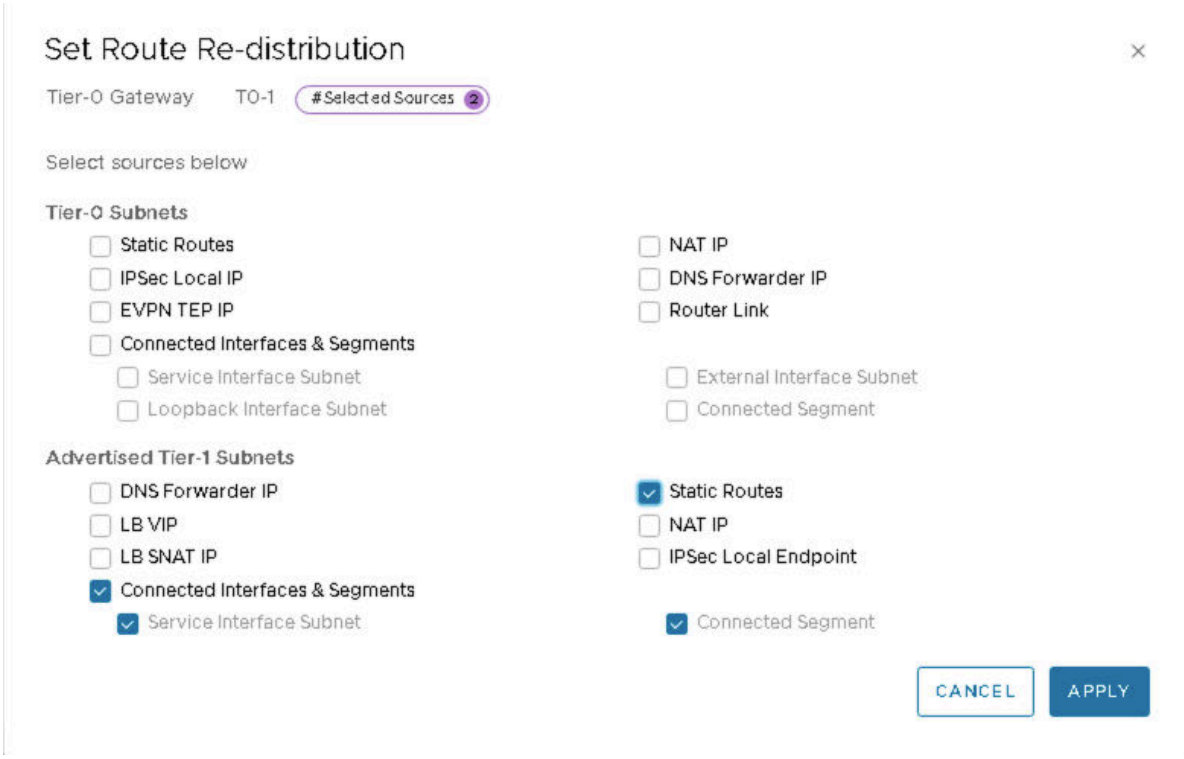


## VMware NSX-T: Create a Static Route to Loopback on Border Leaf Switches

A static route needs to be created for the reachability of the loopbacks on the border leaf switches. These interfaces are used to establish BGP neighbors between NSX-T and the border leaf switches.

1. In NSX-T Manager, navigate to **Networking** > **Tier-0 Gateways**.
2. Select **T0-1 Gateway** and edit, then select **Set**.
3. Add Loopback **192.168.255.2/32** of Border Leaf1 and **192.168.255.3/32** of Border Leaf2 as static routes.

**NOTE**: Ensure to check Apstra for the correct loopbacks for the border leaf switches.

**Static Route to Border Leaf1**

**Figure 60: Static Route to Border Leaf1**

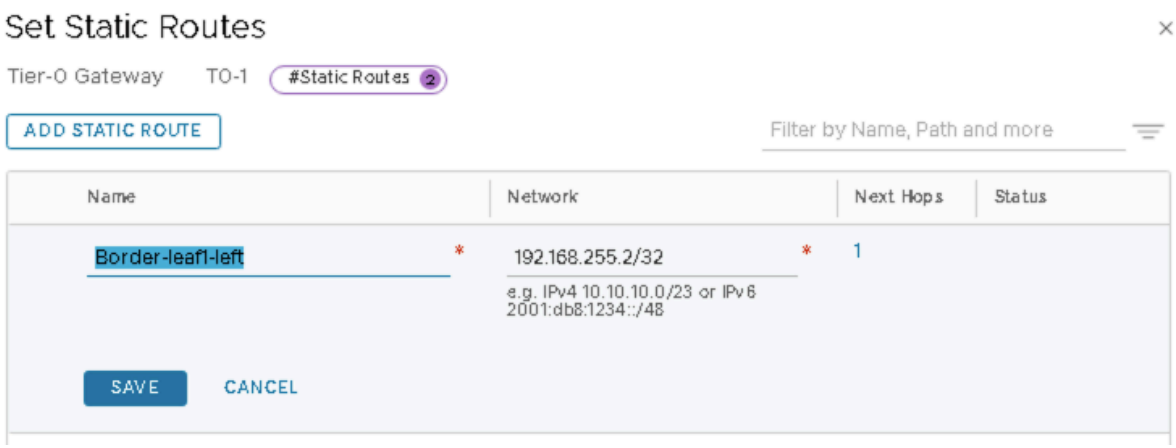


Click **SET NEXT HOPS** and add **192.168.100.1** (IP of the border leaf1 switch interface).

**Figure 61: Next Hop as Border Leaf1 Interface IP (left Uplink)**

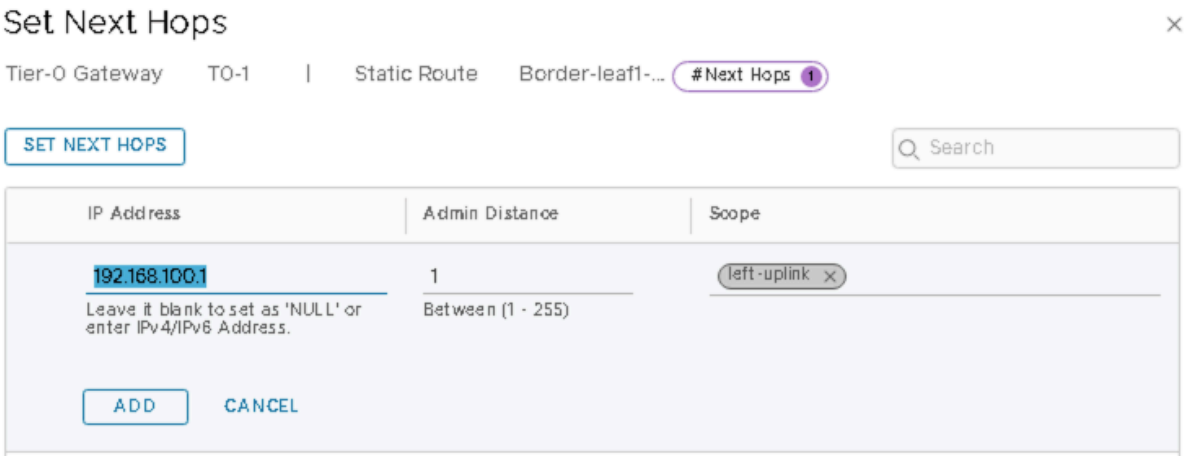


**Static Route to Border Leaf2**

For the static route to border leaf2, click **SET NEXT HOPS** and add **192.168.200.1** (IP of the border leaf2 switch interface).

**Figure 62: Border Leaf2 Static Route**

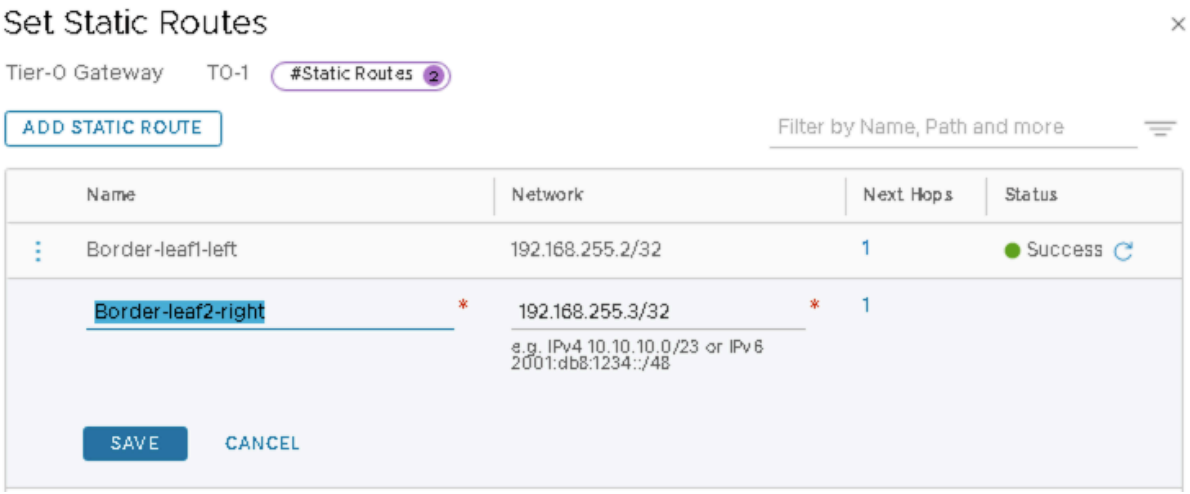


**Figure 63: Next Hop as Border Leaf2 Interface IP (Right Uplink)**

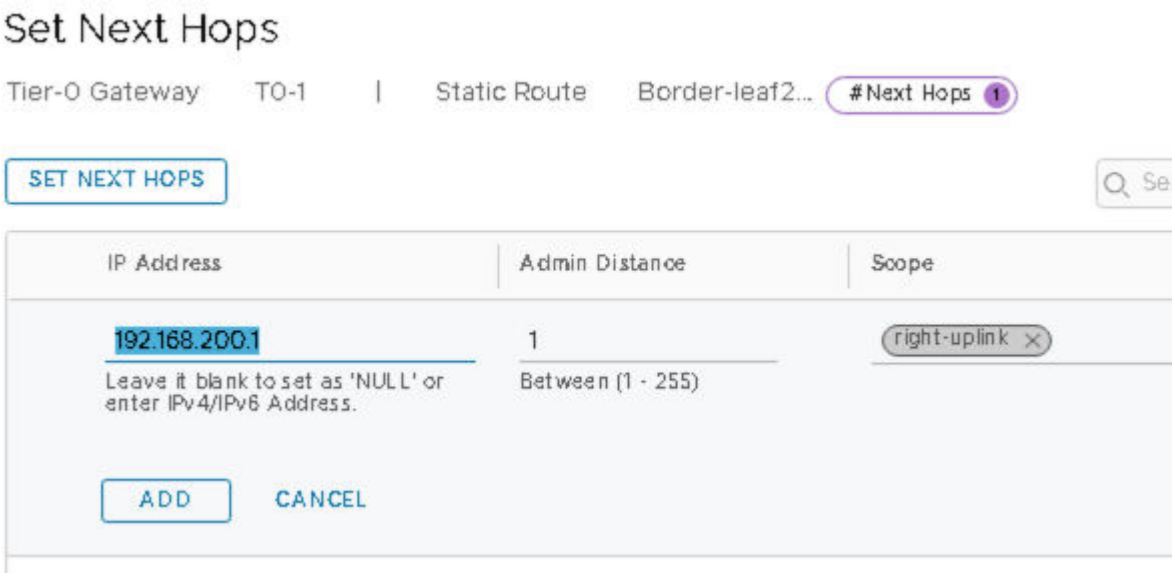

**Figure 64: Static Routes Towards Border Leaf Switches**

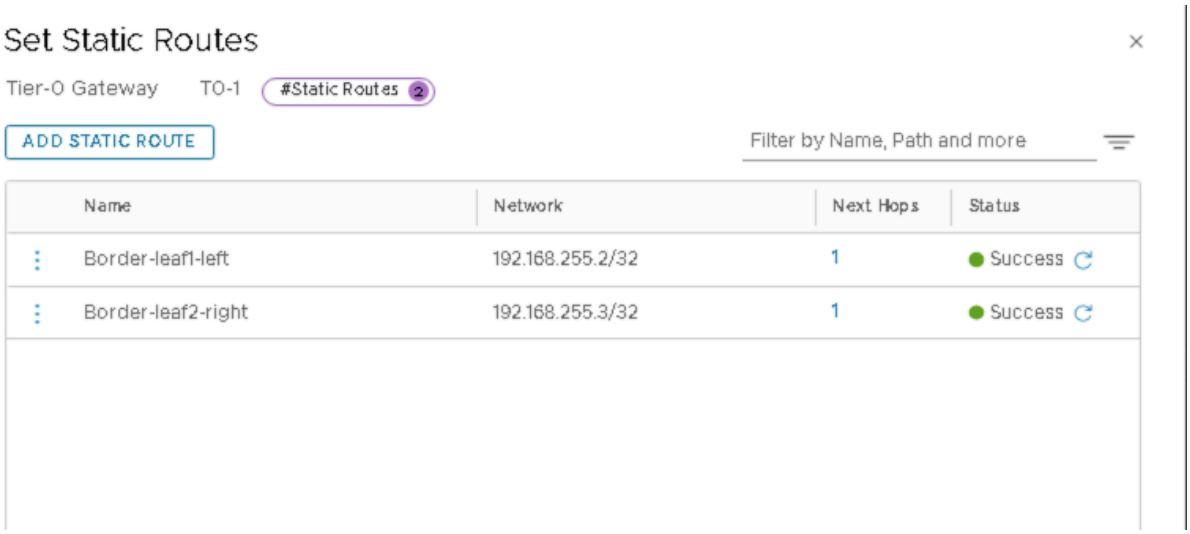


## VMware NSX-T: Create IP Prefix lists on T0 Gateway

The IP Prefix list must be set up not to allow advertising of the fabric IPs.

Add the IP Prefixes as below:

1. In NSX-T Manager, navigate to **Networking** > **Tier-0 Gateways**.
2. Select **T0-1 Gateway**, expand **Routing**, and click on the number beside IP Prefix Lists to add or edit the prefix list.

   prefixlist-out-default is the prefix-list set on the T0 Gateway BGP, as mentioned in "VMware NSX-T: Configure BGP on the T0 GW" on page 55.

**Figure 65: Adding Prefixes to prefixlist-out-default**


3. Click on **1** (or any number) under Prefixes to add prefixes.
4. Click **Edit** and add the following prefixes.

**Figure 66: VM Prefixes Permitted to be Advertised, Rest Denied**

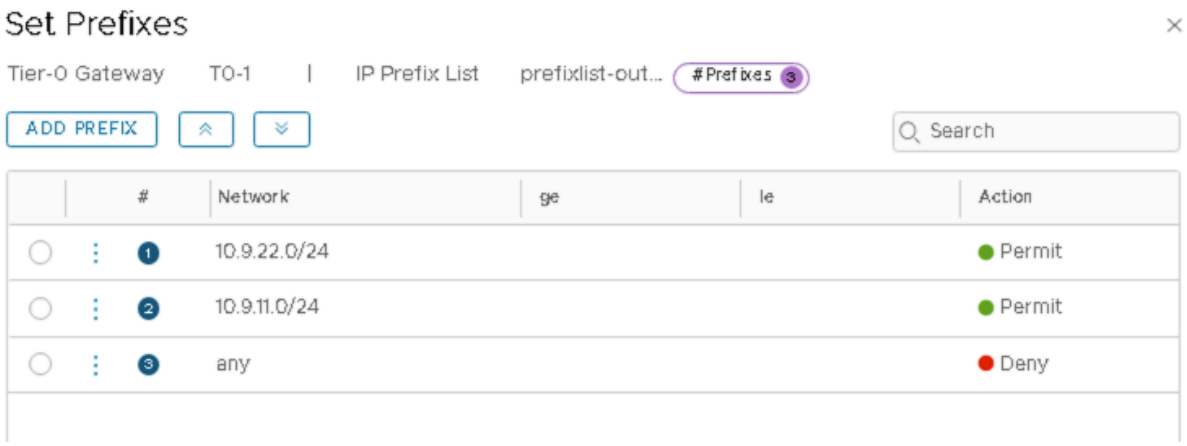


## VMware NSX-T: Connect the T1 and T0 Gateways

Connecting the T1 and T0 Gateways enables east-west connectivity and north-south communication to the VMs in the NSX-T Domain.

### Connect the Gateways

1. In NSX-T Manager, navigate to **Networking** > **Tier-1 Gateways**.
2. Select **T1-1 Gateway**.

3. Under Linked Tier-0 Gateway, select **T0-1**.
4. Add the edge-cluster1 setup in VMware NSX-T: Deploy NSX Edge Node and Create an Edge Cluster on page 39.
5. Select the following under the Route Advertisement:
   - All Static Routes
   - All LB VIP Routes
   - All LB SNAT IP Routes

**Figure 67: Link T0-1 Gateway to T1-1**

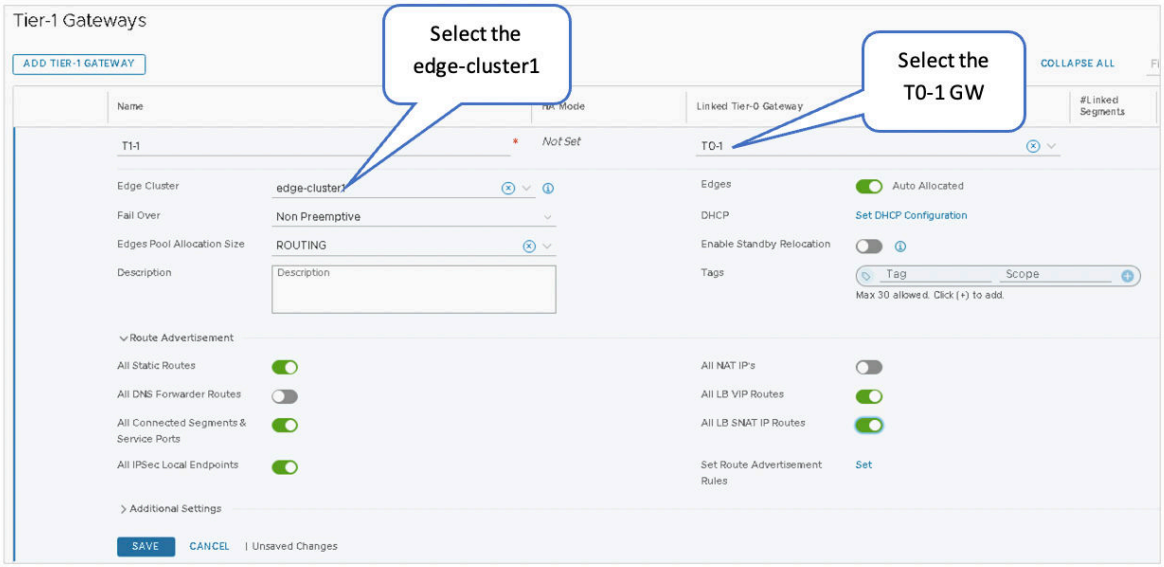


## Juniper Apstra: Add the NSX-Manager

For Juniper Apstra to remediate inconsistencies between the virtual infrastructure and the physical IP fabric, the NSX-T manager should be added in Juniper Apstra.

In the Juniper Apstra UI, navigate to **External Systems** > **Virtual Infra Managers** > **Create Virtual Infra Manager** and add the NSX-T manager details and credentials.

**Figure 68: Adding NSX-T and Vsphere in Apstra**

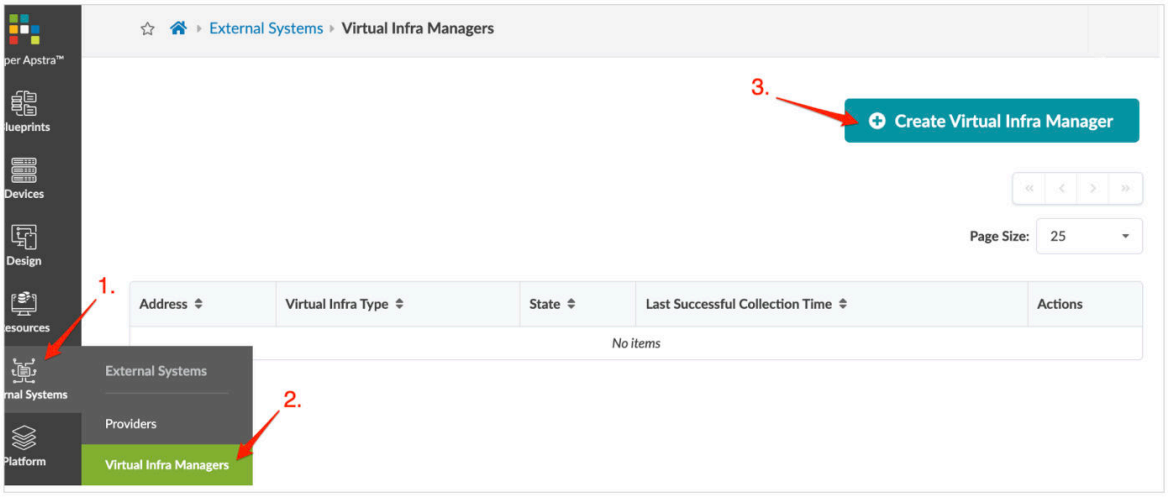


## Juniper Apstra: Create a Routing Policy for NSX-T in the Blueprint

Add a Routing policy for the NSX-T routing zone created in the next step.

In the Juniper Apstra UI, navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Policies** > **Routing Policies** > **Create Routing Policy**.

**Figure 69: NSX-T Routing Policy**

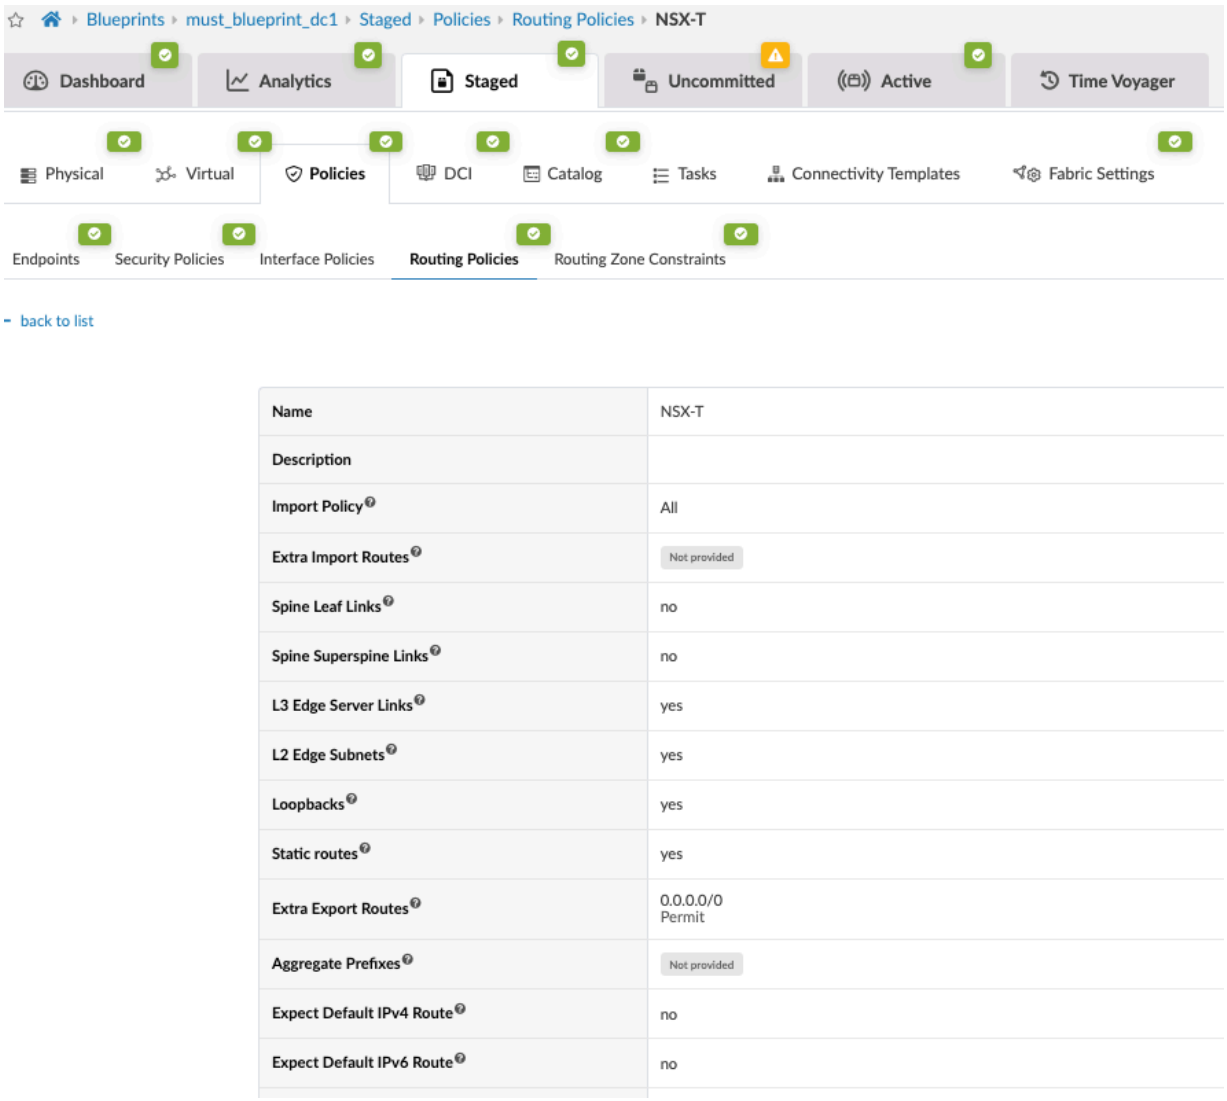


## Juniper Apstra: Create a Routing Zone in the Blueprint

**Add a Routing Zone That Maps to a VRF in the Blueprint:**

1. In the Juniper Apstra UI, navigate to **Blueprints** > ***<blueprint-name>*** > **Staged** > **Virtual** > **Routing Zones** > **Create Routing Zone** and following details:
   - VRF Name: NSX-T
   - VLAN ID: 9

- VNI: 20000
- Route Target: 20000:1
- Routing Policies: NSX-T

**Figure 70: Routing Zone to Communicate with NSX-T**

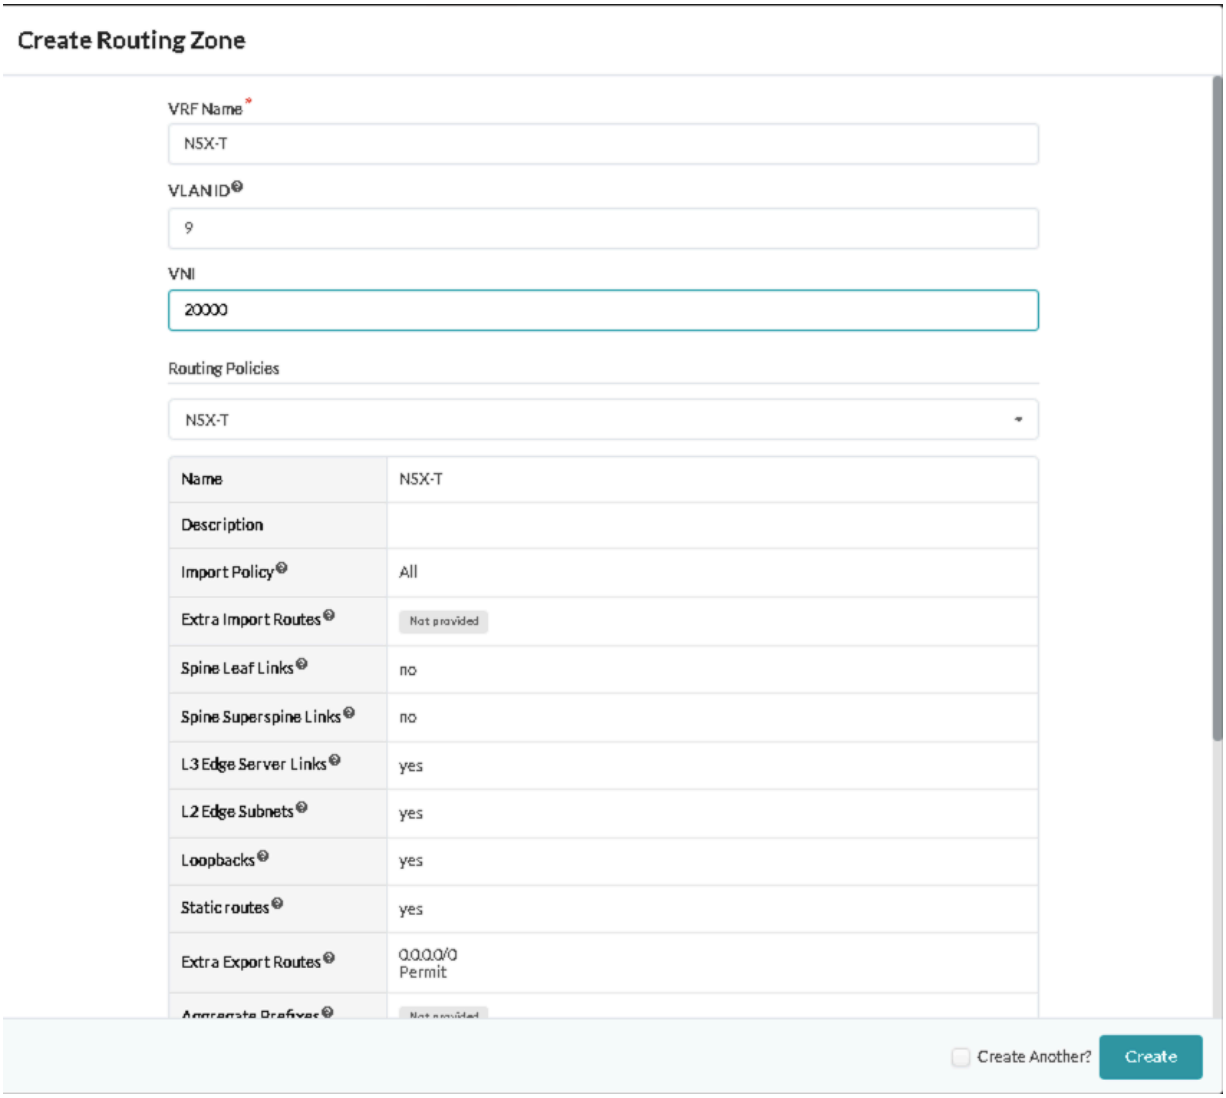


## Juniper Apstra: Assign the Loopback IPs to the Routing Zone

After creating the NSX-T routing zone, assign the loopback IPs for the routing zone. The loopback IP is allocated from an IP Pool in Resources. In the following figures, the pool MUST-EVPN-Loopbacks DC1 is

already created under Resources. This is as per the section "Apstra Resources: ASN, Fabric, and Loopback IP Address" on page 8.

The loopback IP is assigned to the routing instance and used to extend EVPN and the NSX-T overlay VLAN between the leaf switches.

**Figure 71: IP Pool from Resources**

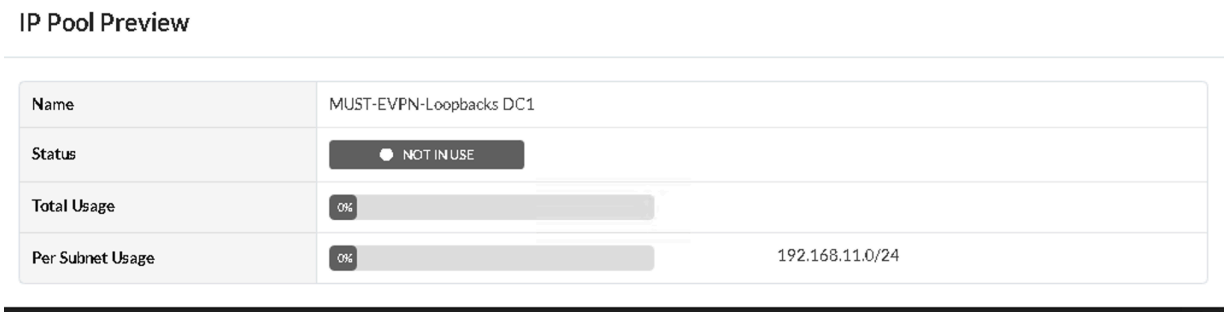


**Figure 72: Assign IP Pool for the Leaf switches in NSX-T Routing Zone**

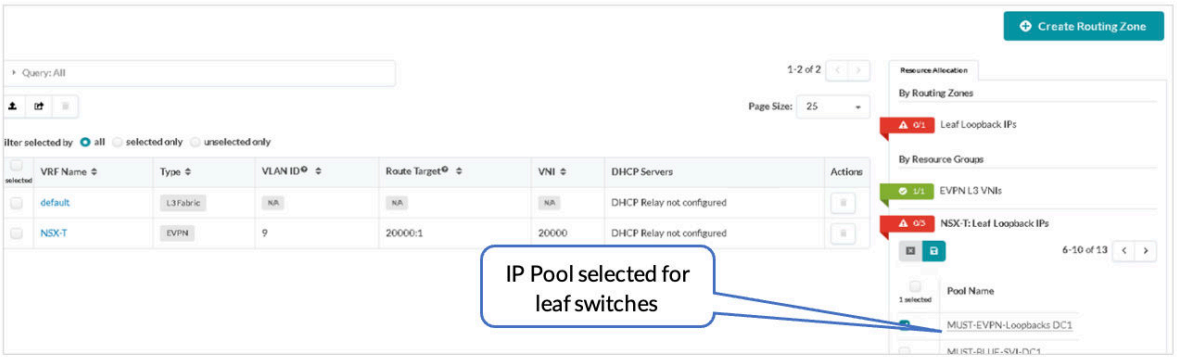


## Juniper Apstra: Add the NSX-Manager into the Blueprint

Add the NSX-Manager into the Blueprint that is managing the fabric:

1. In the Juniper Apstra UI navigate to **Blueprints** > **<blueprint-name >** > **Staged** > **Virtual** > **Virtual Infra** > **Add Virtual Infra**.
2. From Virtual Infra Manager, select the vSphere added in Juniper Apstra: Add the NSX-Manager.
3. Set the VLAN Remediation Policy VN Type as **VXLAN**.

4. Set the Routing Zone as **NSX-T**.

**Figure 73: Add NSX-T Manager into Blueprint**

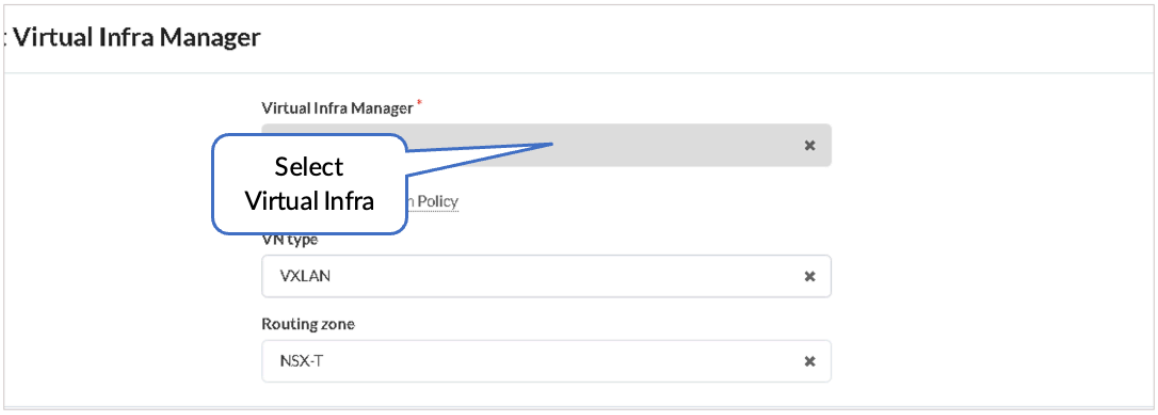


## Juniper Apstra: Add the NSX-T-Overlay as a VN

For the GENEVE Tunnels to come up between the Transport Nodes in NSX-T, connectivity needs to be established through Juniper Apstra Fabric. This is ensured by creating VXLAN Virtual Network in Apstra and assigning correct port mapping in ToR leaf switches towards Transport Node. Ensure that VLAN ID for Overlay VXLAN VN defined in Apstra match the one mapped in Overlay Profile in NSX-T for Transport Nodes.

VLAN **50** is configured in NSX-T Managed for Overlay, which maps to the VNI **10050**. Connectivity Template **Tagged** should be selected while creating a virtual network. The virtual network is assigned to all leaf switches. The IPv4 subnet (IRB) is disabled as NSX-T is already assigning the TEPs to the hosts in NSX-T.

**Figure 74: NSX-T Overlay Profile Transport VLAN Configured 50**

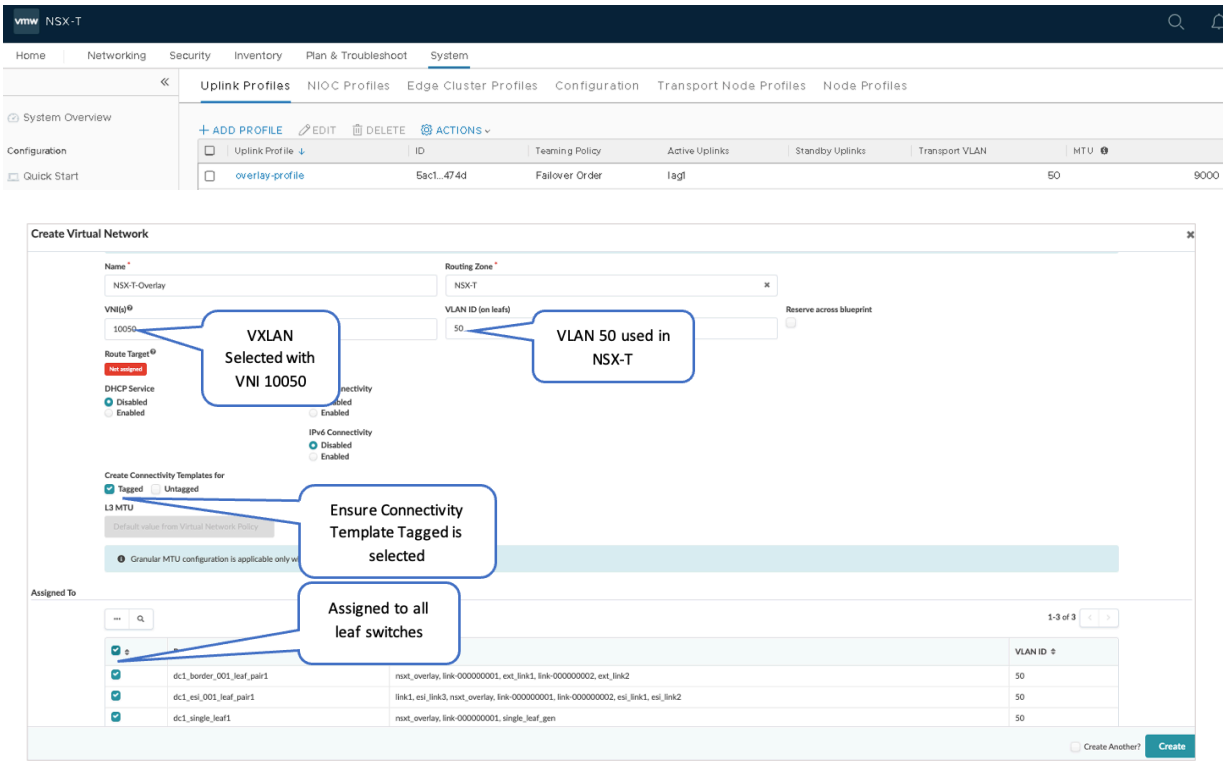


## Juniper Apstra: Verify the Connectivity Templates

With the virtual network created, if you select **Create Connectivity Templates**, then a Connectivity Template type **Virtual Network[Single]** is created.

**Verify the Creation of a Virtual Network:**

1. In the Juniper Apstra UI, navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Connectivity Templates**.
2. Scroll to look for the connectivity template for the Virtual Network.
3. Click **Edit** to view the connectivity template.

**Figure 75: Apstra Connectivity Template Generated for Overlay Virtual Network**

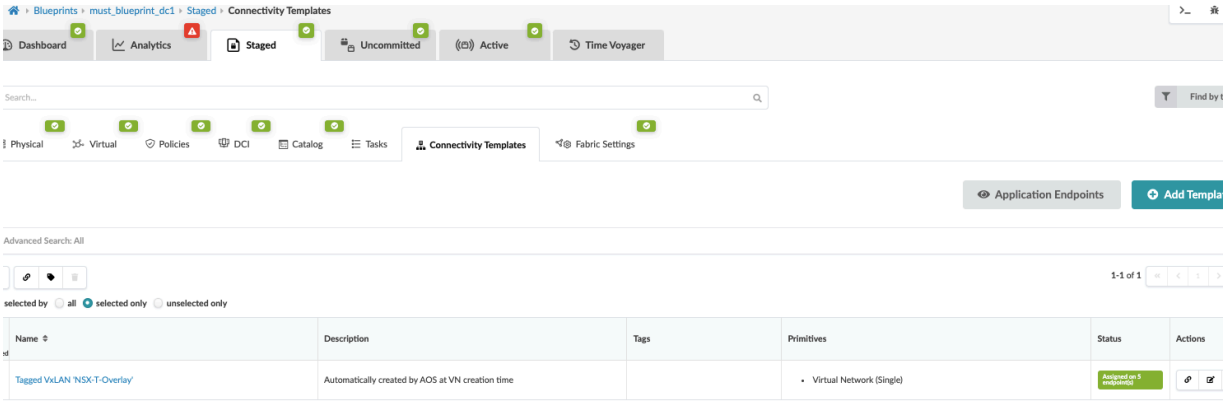


## Juniper Apstra: Assign Interface to the Connectivity Templates

The connectivity template is assigned to the aggregate Ethernet (AE) interfaces facing the ESXi hosts.

**Figure 76: AE Interfaces are assigned to Connectivity Template**

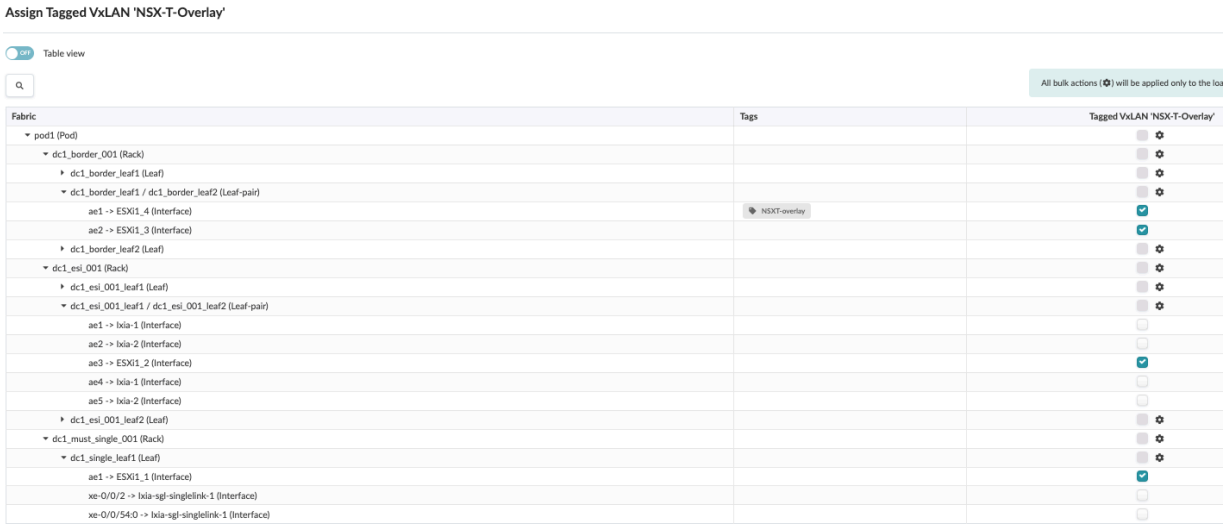

## Juniper Apstra: Commit the Configuration

From Blueprint, navigate to **Blueprints** > **<blueprint-name>** > **Uncommitted** and commit the configuration. This pushes the VN and routing zone to NSX-T to all the fabric devices.

## VMware NSX-T: GENEVE Tunnels

Once Juniper Apstra has pushed the configurations to the Junos OS devices, observe that the GENEVE tunnels between the Transport Nodes and the Edge Nodes are up:

1. On the edge01 Edge-VM , view the Tunnel Endpoints and verify status is **UP**.
2. In NSX-T Manager, navigate to **NSX-T Manager** > **System** > **Fabric** > **Nodes** > **Edge** Transport Nodes.
3. Click **edge01**, then click **Tunnels**.

**Figure 77: NSX-T Edge Node Tunnels are Up**

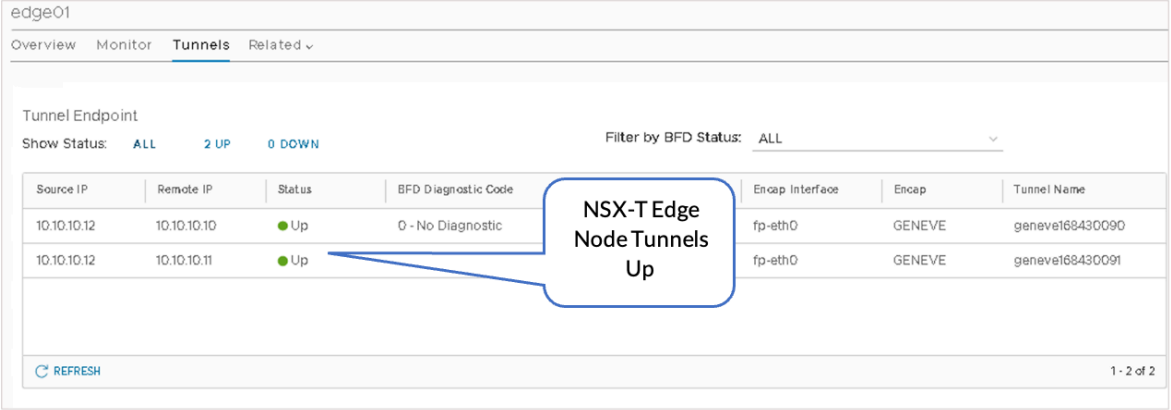


## Juniper Apstra: Add Connectivity Templates for Connectivity from Edge Node to the Fabric

The connectivity templates specify the IP link for the connectivity from the Edge Node to the fabric and the BGP peering session with a user-specified BGP neighbor-addressed peer.

1. In the Juniper Apstra UI, navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Connectivity Templates**.

2. Click **Add Template**.

## Juniper Apstra: Add IP Link, BGP Peering and Static Route

A Connectivity template is used to create the NSX-T uplinks towards NSX-T Edge Node edge01 for both left and right uplinks.

Here, the default routing zone is selected to connect to the Edge Nodes as the NSX-T traffic within the fabric is an overlay traffic. Once the traffic reaches NSX-T, it is an underlay traffic.

The peer ASN is the NSX-T T0-1 ASN 65000.

- **Create Two IP Link Connectivity Templates for Left Uplink and Right Uplink**

**Figure 78: IP Link Connectivity Template for Left and Right Uplinks**

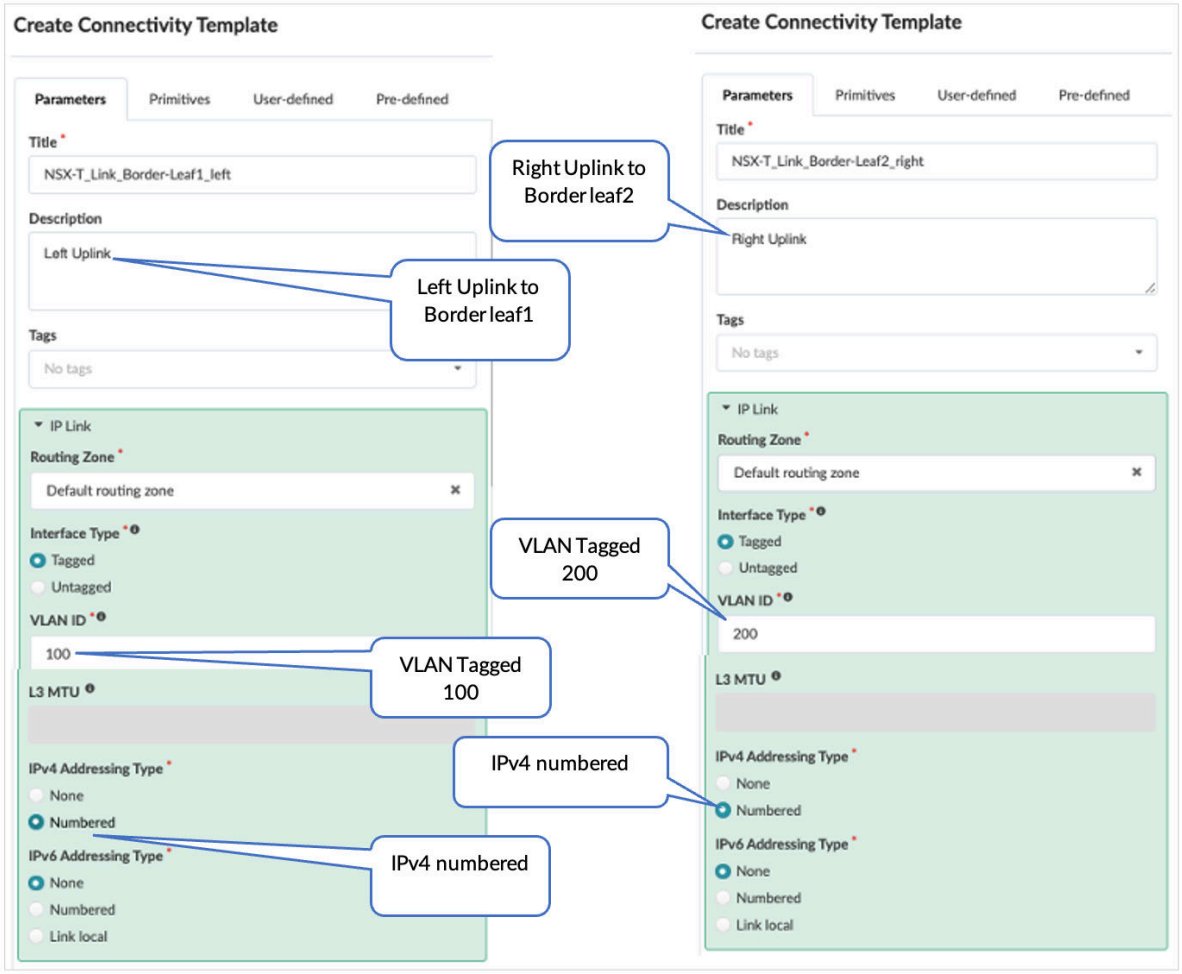

- **Create BGP Peering and Assign NSX-T Routing Policy**

1. In the Juniper Apstra UI, navigate to **Blueprints** > ***<blueprint-name>*** > **Staged** > **Connectivity Templates**.
2. Click **Add Template**.
3. Create Connectivity Templates for BGP peering and assign the NSX-T routing policy.

**Figure 79: Connectivity Template for Uplinks Towards Edge Node**

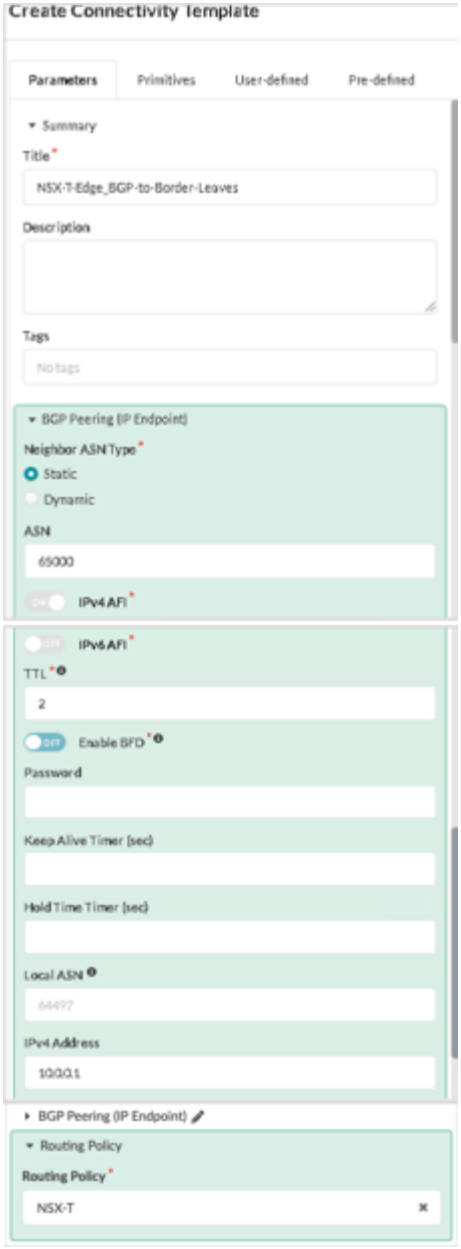

- **Create Two Custom Static Route Connectivity Templates**

The static route is created in a separate Connectivity Template as the primitive ‘Custom Static Route’ used here is generated at the system level (border-leaf level). The static route for the left uplink starts from Border Leaf1 to the Edge Node and for the right uplink from Border Leaf2 to the Edge Node.

1. In the Juniper Apstra UI, navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **ConnectivityTemplates**.
2. Click **Add Template**. Repeat this process to create Connectivity Templates for the left and right uplinks.

- **Create Two Connectivity Templates for Static Route**
  - Left static route from the Border Leaf1 to the Edge Node.
  - Right static route from the Border Leaf2 to the Edge Node.

**Figure 80: Connectivity Template for Left Uplink Static Route**

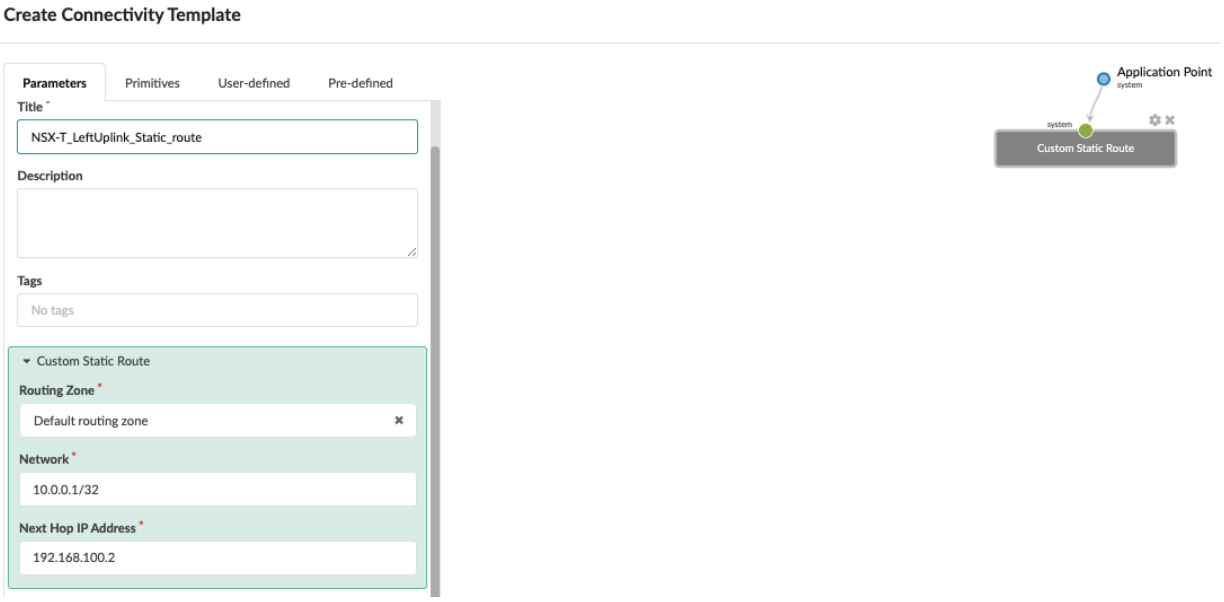

**Figure 81: Connectivity Template for Right Uplink Static Route**

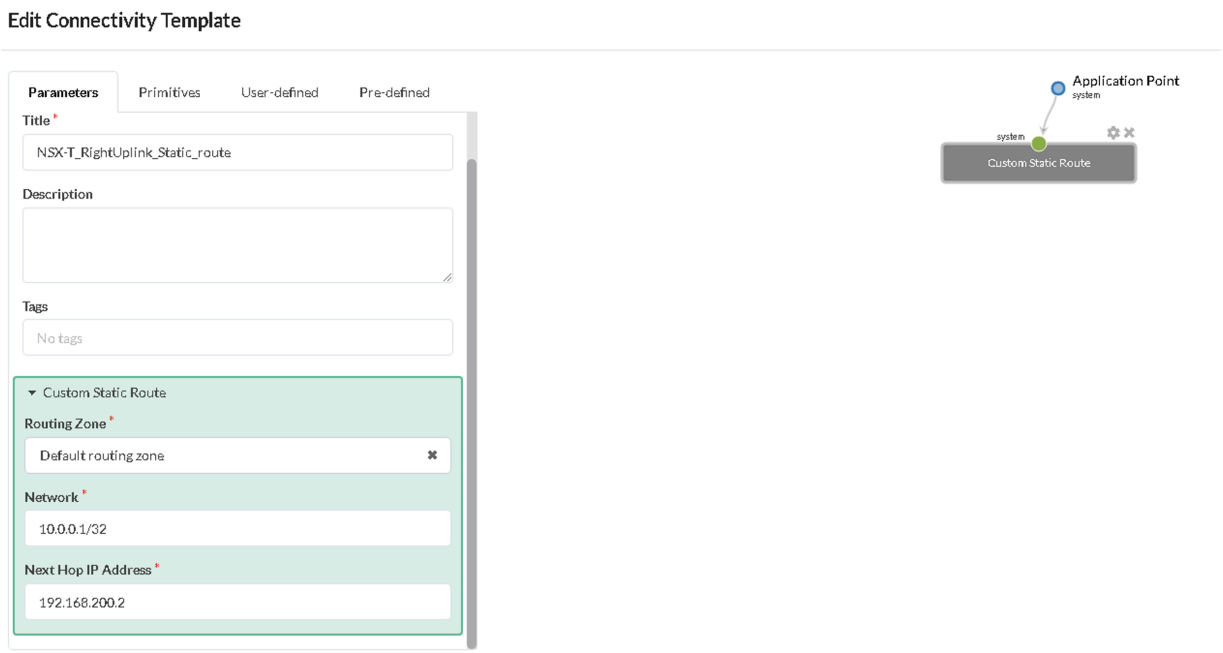


## Juniper Apstra: Renaming the Generic System and Adding Links from Border Leaf switches

Once the connectivity template for IP_Link is created for the uplink from the Border leaf switches towards the NSX-T-Edge device, Apstra adds a generic system node with the ASN (65000 is allocated to T0 in the earlier steps).

1. In the Juniper Apstra UI, navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Physical** > **Nodes**.
2. Locate the generic system added as shown below. Click the pencil icon to change this name to **NSXT-Edge-01**.

**Figure 82: Renaming Generic System Added after Adding the IP link Connectivity Template**

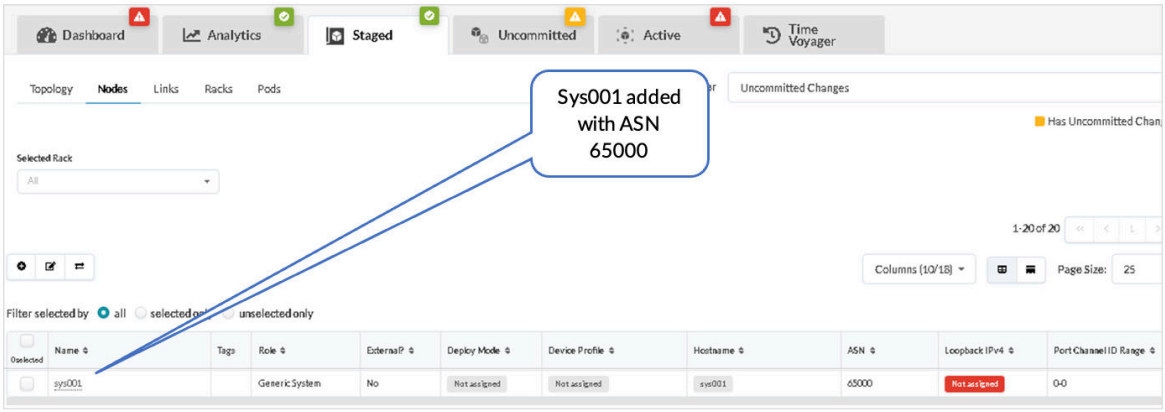


After changing the generic system name, create links from both border leaf switches towards ESXi for the left and right uplink.

1. In the Juniper Apstra UI, navigate to **Blueprints** > ***<blueprint-name>*** > **Staged** > **Physical** > **Topology**.
2. Select **Border leaf1**, then select 'Add links to a generic system' and create a link from dc1_border_leaf1.
3. For dc1_border_leaf2, the steps to create the right uplink link are similar.

**Figure 83: Add Interface Link for Left Uplink Towards NSXT-Edge**

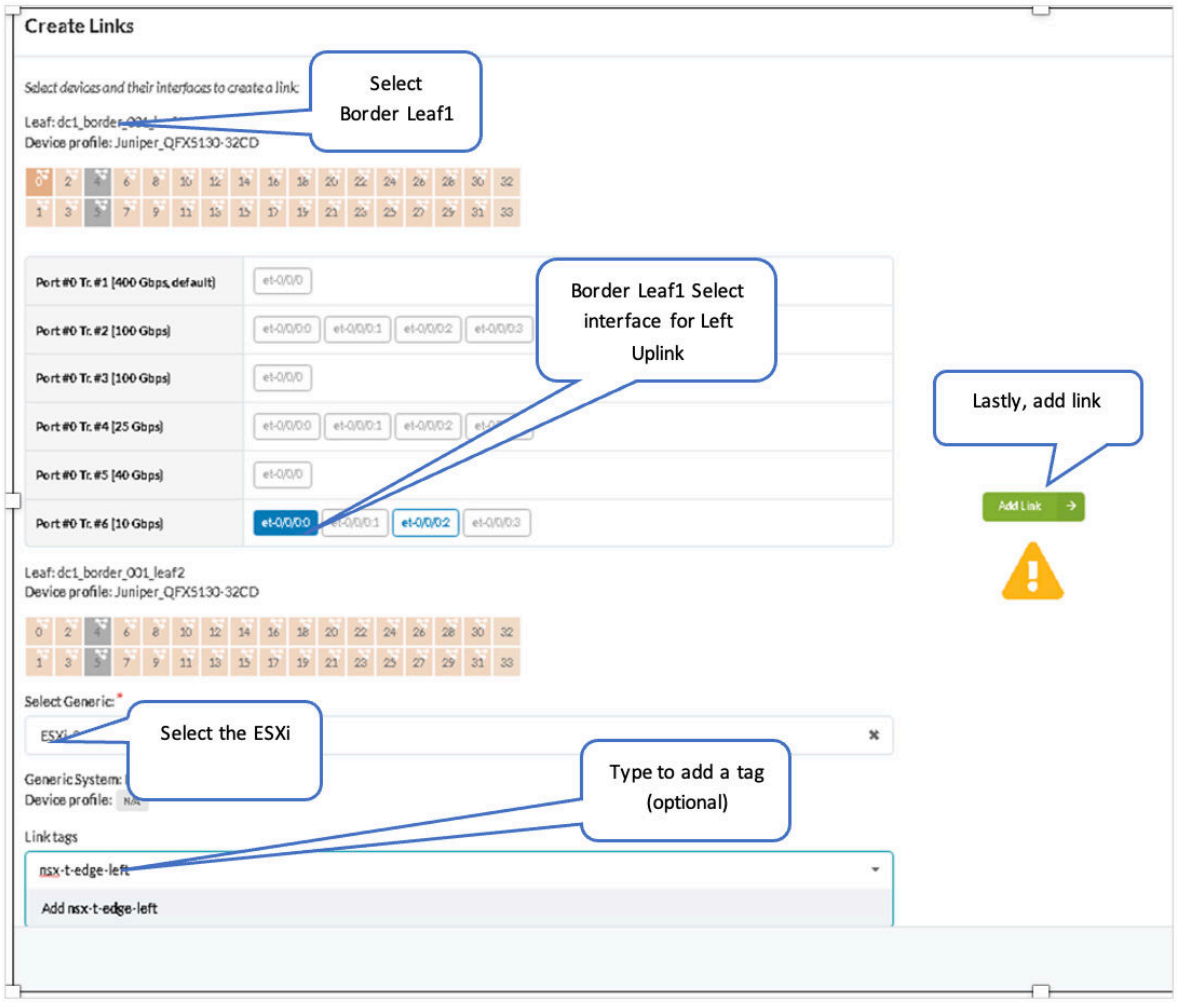


Once both border leaf switches uplinks are created, the links are shown on the topology. After this, interfaces can be assigned from the connectivity templates, as covered in the next steps.

## Juniper Apstra: Assign the Interfaces to the Connectivity Template

Assign connectivity templates created in "Juniper Apstra: Add IP Link, BGP Peering and Static Route" on page 75 section.

1. Assign Uplinks.
2. For Uplinks from the border leaf switches, assign the appropriate ethernet interface for dc1_border_leaf1 and dc1_border_leaf2:
   - Left uplink from dc1_border_leaf1

**Figure 84: Left Uplink Interface on Border Leaf1**

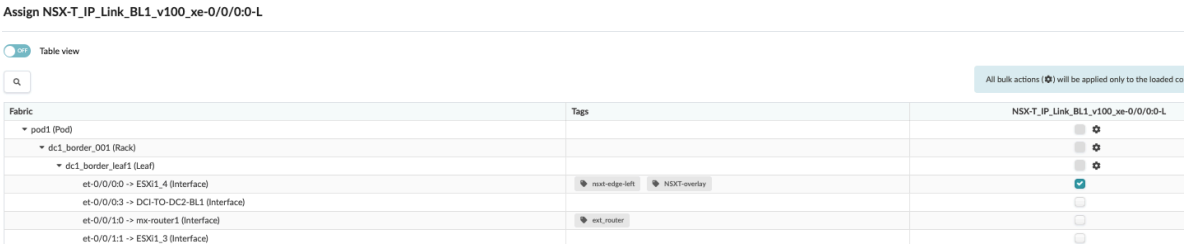


- Right uplink from dc1_border_leaf2

**Figure 85: Right Uplink interface on Border Leaf2**

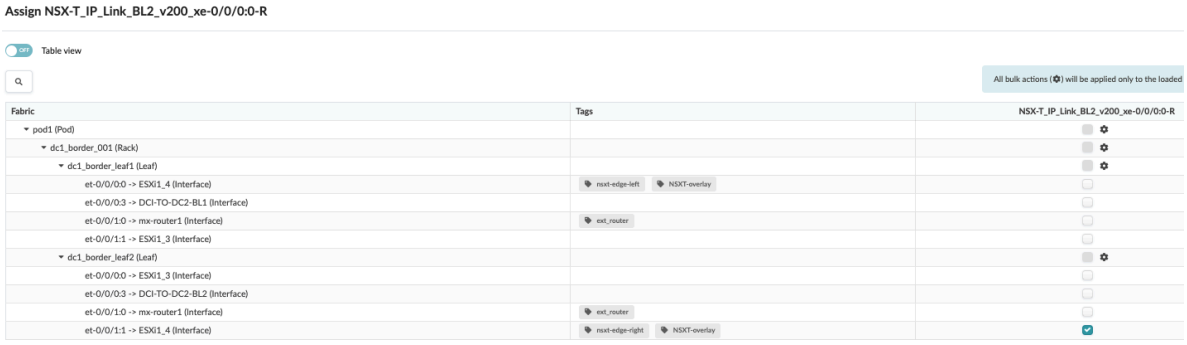


3. Assign BGP.

4. For BGP peering, select dc1_border_leaf1 and dc1_border_leaf2 loopback Interfaces lo0.0.

**Figure 86: Assign BGP Peering to Both Border Leaf switches Default Loopback Interface for Left and Right Uplink**

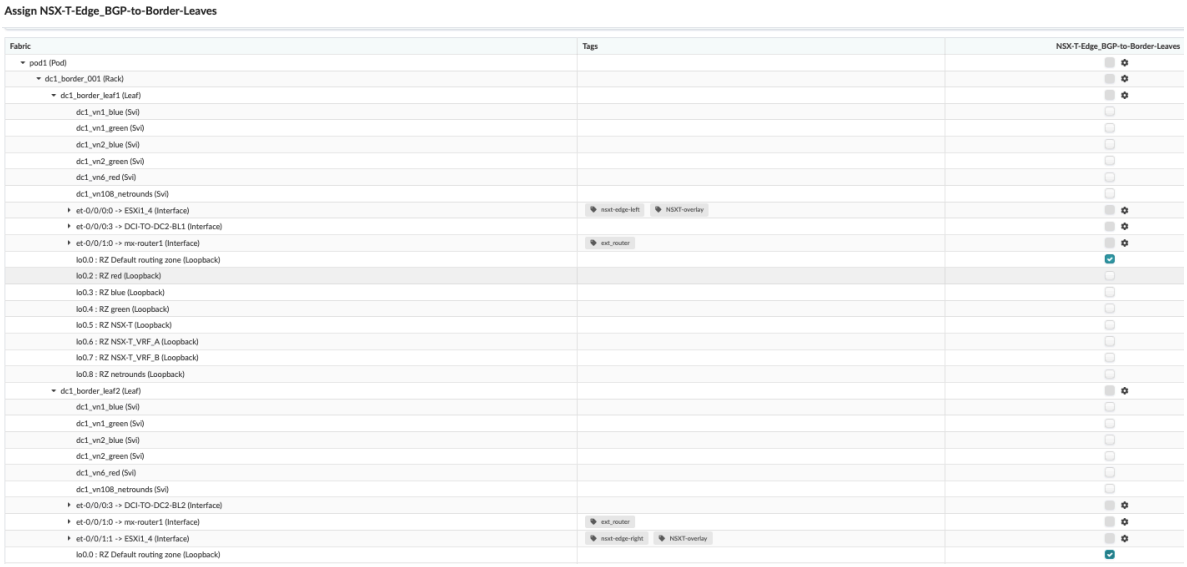


5. Assign Static Routes

- Static route for left uplink to Border Leaf1

**Figure 87: Assign Static Route for Left Uplink to Border Leaf1**

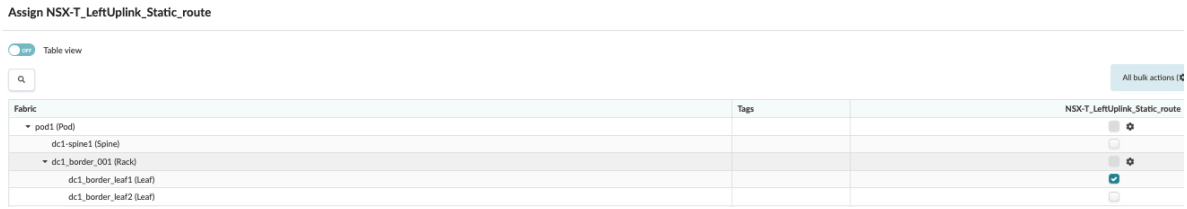


- Static route for right uplink to Border Leaf2

**Figure 88: Assign Static Route for Right Uplink to Border Leaf2**

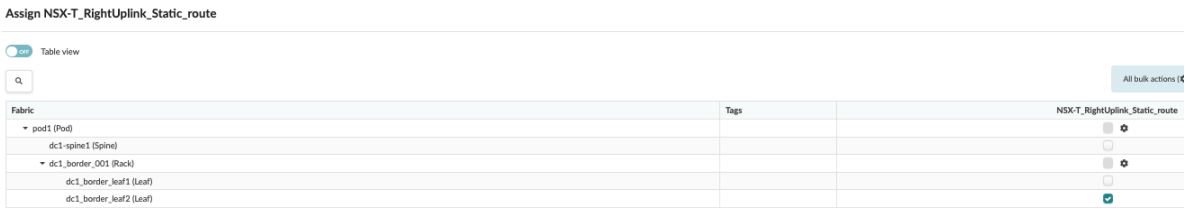

## Juniper Apstra: Assign the IPs and VLAN IDs to the Interfaces

Now that a Connectivity Template has been created and physical interfaces are assigned, you must assign IP addresses and VLAN IDs to the interfaces.

Edit the Interface connecting the Border Leaf and the Edge Node:

1. In the Juniper Apstra UI, navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Virtual** > **RoutingZones** > **DefaultRoutingZone** > **Interfaces** > **EditIPAddresses**.
2. Enter the IP address for Border Leaf switches and the IP address of the Edge node interfaces on the host.

**Figure 89: Assign IPs to the Interface connected to Edge Host**

| | | Endpoint 1 | | | | Interface 1 | | | | Endpoint 2 | | | | Interface 2 | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Routing Zone | VLAN ID | Name | Role | Interface | L3 MTU | IPv4 Address | IPv4 Address Type | IPv6 Address | IPv6 Address Type | Name | Role | Interface | L3 MTU | IPv4 Address | IPv4 Address Type | IPv6 Address | IPv6 Address Type |
| Default routing zone | 200 | dc1_border_leaf2 | Leaf | et-0/0/1:1.200 | Not provided | 192.168.200.1/24 | Numbered | N/A | None | ESXi1_4 | Generic System | n/a | Not provided | 192.168.200.2/24 | Numbered | N/A | None |
| Default routing zone | 100 | dc1_border_leaf1 | Leaf | et-0/0/0:0.100 | Not provided | 192.168.100.1/24 | Numbered | N/A | None | ESXi1_4 | Generic System | n/a | Not provided | 192.168.100.2/24 | Numbered | N/A | None |

## Juniper Apstra: Commit the Configuration

Navigate to **Blueprints** > **Uncommitted** and commit the configuration. This pushes all the uplinks created using connectivity templates to the devices.

## Juniper Junos OS: Verify Configs

Log onto one of the Border Leaf switches and verify that the configuration is being pushed.

### Verify Physical Fabric Configuration

SSH into one of the Border Leaf switches and run the following commands:

- **show configuration interfaces et-0/0/0:0 | display set**

```
root@dc1-border-leaf1> show configuration interfaces et-0/0/0:0 | display set
set interfaces et-0/0/0:0 description to.esxi-16
set interfaces et-0/0/0:0 flexible-vlan-tagging
set interfaces et-0/0/0:0 native-vlan-id 1
set interfaces et-0/0/0:0 mtu 9100
set interfaces et-0/0/0:0 unit 100 description "VRF default to
set interfaces et-0/0/0:0 unit 100 vlan-id 100
set interfaces et-0/0/0:0 unit 100 family inet address 192.168.100.1/24
```

Configuration applied based on IP link connectivity template and interfaces and IP assigned.

- **show configuration protocols bgp group l3rtr | match 10.0.0.1 | display set**

```
root@dc1-border-leaf1> show configuration protocols bgp group l3rtr | match 10.0.0.1 |
display set
set protocols bgp group l3rtr neighbor 10.0.0.1 description facing_n
set protocols bgp group l3rtr neighbor 10.0.0.1 multihop ttl 2
set protocols bgp group l3rtr neighbor 10.0.0.1 local-address 192.168
set protocols bgp group l3rtr neighbor 10.0.0.1 import ( RoutesFromExt-default-NSX-T )
set protocols bgp group l3rtr neighbor 10.0.0.1 family inet unicast
set protocols bgp group l3rtr neighbor 10.0.0.1 export ( RoutesToExt-default-NSX-T )
set protocols bgp group l3rtr neighbor 10.0.0.1 peer-as 65000
```

BGP peering interface applied.

- **show configuration routing-instances NSX-T | display set**

```
root@dc1-border-leaf1> show configuration routing-instances NSX-T      isplay set
set routing-instances NSX-T instance-type vrf
set routing-instances NSX-T routing-options rib NSX-T.inet6.0 multi
set routing-instances NSX-T routing-options multipath
set routing-instances NSX-T routing-options auto-export
set routing-instances NSX-T protocols evpn ip-prefix-routes adv
set routing-instances NSX-T protocols evpn ip-prefix-routes enc
set routing-instances NSX-T protocols evpn ip-prefix-
set routing-instances NSX-T protocols evpn ip-prefix-routes expo
set routing-instances NSX-T interface lo0.5
set routing-instances NSX-T route-distinguisher 192.168.255.2:9
set routing-instances NSX-T vrf-target target:20000:1
set routing-instances NSX-T vrf-table-label
```

Loopback interface applied and using NSX-T policy.

- **show configuration routing-options | display set**

```
root@dc1-border-leaf1> show configuration routing-options | display set
set routing-options rib inet.0 static route 10.0.0.1/32 next-hop 192.168.100.2
set routing-options router-id 192.168.255.2
set routing-options autonomous-system 64514
set routing-options forwarding-table export PFE-LB
set routing-options forwarding-table ecmp-fast-reroute
set routing-options forwarding-table chained-composite-next-hop ingress evpn
```

Static route from BL1 to Edge Node1.

## VMware NSX-T: Verify BGP Session on Edge

On the Edge Node, verify that the BGP sessions are established, and the overlay routes exchanged.

### Verify NSX-T Configuration

SSH into the NSX-T edge01 Edge-VM and run the following commands:

1. Firstly, to determine the VRF (SR-T0-1), run the command, get logical-router and pick the one that has the name “SR-T0-1”. (service router Tier0) The corresponding VRF number is in the VRF column.

```
10> get logical-routers
Fri Jan 26 2024 UTC 18:11:32.003
Logical Router
UUID                                   VRF    LR-ID  Name
Type                        Ports   Neighbors
736a80e3-23f6-5a2d-81d6-bbefb2786666   0      0
TUNNEL                      3       6/5000
c9b6e428-86e8-457a-b961-574e4222d62b   1      17     DR-T0-1
DISTRIBUTED_ROUTER_TIER0    5       2/50000
5e7295b5-9370-4dcc-bfda-5c5d8c43a997   2      8193   SR-T0-1
SERVICE_ROUTER_TIER0        9       3/50000
3468dc14-32f1-423c-ac4c-ebf9835deffd   4      8199   SR-VRF-VRF-B
VRF_SERVICE_ROUTER_TIER0    5       0/50000
bb12844c-e93e-434c-a8a8-bc41f0a97559   5      2060   DR-VRF-VRF-B
VRF_DISTRIBUTED_ROUTER_TIER0 4       0/50000
dbbb8e77-02c2-496b-b9f7-d14b95c360cd   6      8198   SR-VRF-VRF-A
VRF_SERVICE_ROUTER_TIER0    6       0/50000
8f4fad8c-bb61-4b17-8772-7b1797f599eb   7      2058   DR-VRF-VRF-A
VRF_DISTRIBUTED_ROUTER_TIER0 3       0/50000
d686dc86-3101-493f-9e21-65f94e45f73b   8      16     DR-T1-1
DISTRIBUTED_ROUTER_TIER1    5       0/50000
9e796408-7732-4fdc-9de6-2d286572f89b   9      1025   SR-T1-1
SERVICE_ROUTER_TIER1        5       2/50000
```

2. **get bgp neighbor summary**

```
 10> vrf 2
 10(tier0_sr[2])> get bgp neighbor summary
 BFD States: NC - Not configured, DC - Disconnected
             AD - Admin down, DW - Down, IN - Init, UP - Up
 BGP summary information for VRF default for address-family: ipv4Unicast
 Router ID: 10.0.0.1  Local AS: 65000

 Neighbor                        AS          State Up/DownTime  BFD InMsgs  OutMsgs
InPfx  OutPfx

 192.168.255.2                   64514       Estab 02w1d22h           50058
12     8
 192.168.255.3                   64515       Estab 02w1d22h           50045
12     8

 BFD States: NC - Not configured, DC - Disconnected
             AD - Admin down, DW - Down, IN - Init, UP - Up
 BGP summary information for VRF default for address-family: l2VpnEvpn
 Router ID: 10.0.0.1  Local AS: 65000

 Neighbor                        AS          State Up/DownTime  BFD InMsgs  OutMsgs
InPfx  OutPfx

 192.168.255.2                   64514       Estab 02w1d22h     NC  54129   50058
0      0
 192.168.255.3                   64515       Estab 02w1d22h     NC  54122   50045
0      0

 Fri Jan 26 2024 UTC 18:12:41.843
```

Uplink BGP sessions towards border leaf switches Established

3. **get route**

```
 10(tier0_sr[2])> get route

 Flags: t0c - Tier0-Connected, t0s - Tier0-Static, b - BGP, o - OSPF
 t0n - Tier0-NAT, t1s - Tier1-Static, t1c - Tier1-Connected,
 t1n: Tier1-NAT, t1l: Tier1-LB VIP, t1ls: Tier1-LB SNAT,
 t1d: Tier1-DNS FORWARDER, t1ipsec: Tier1-IPSec, isr: Inter-SR,
 > - selected route, * - FIB route

 Total number of routes: 45

 b  > * 0.0.0.0/0 [20/0] via 192.168.100.1, uplink-288, 02w1d22h
 b  > * 0.0.0.0/0 [20/0] via 192.168.200.1, uplink-286, 02w1d22h
 t0c> * 10.0.0.1/32 is directly connected, loopback-290, 05w2d23h
 isr> * 10.0.0.3/32 [200/0] via 169.254.0.130, inter-sr-278, 05w1d13h
 b  > * 10.1.1.1/32 [20/0] via 192.168.100.1, uplink-288, 02w1d22h
 b  > * 10.1.1.1/32 [20/0] via 192.168.200.1, uplink-286, 02w1d22h
 t1c> * 10.9.11.0/24 [3/0] via 100.64.64.1, downlink-280, 05w2d23h
 t1c> * 10.9.22.0/24 [3/0] via 100.64.64.1, downlink-280, 05w2d23h
 t0c> * 100.64.64.0/31 is directly connected, downlink-280, 05w2d23h
```

Default routes advertised from ToR for Internet access

Overlay routes advertised from T1-GW

## VMware NSX-T: Verify BGP Session on ToR

Juniper Apstra should detect the no BGP anomaly on the blueprint.

In the Juniper Apstra UI, navigate to **Blueprints** > **<Blueprint-name>** > **Active**.

**Figure 90: Juniper Apstra Detects No Anomaly**

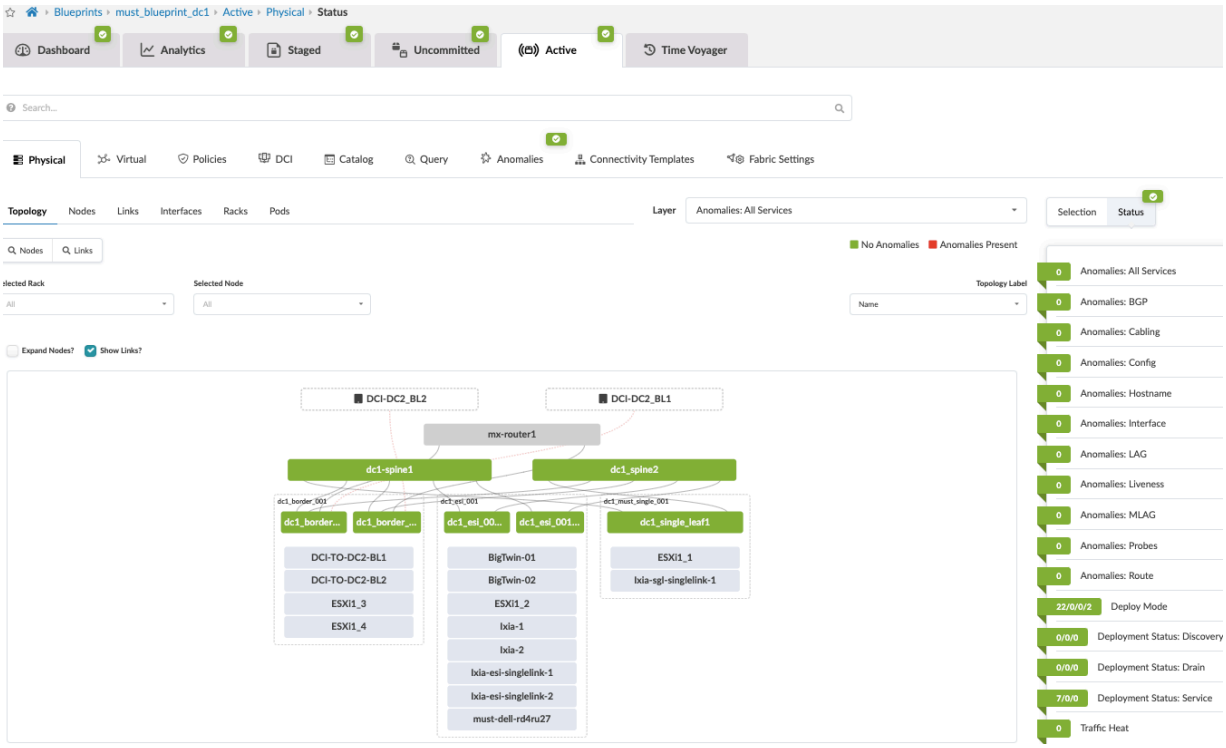


Log onto one of the border leaf switches and verify that the BGP sessions are up, and the overlay routes are exchanged.

## Verify Physical Fabric Configuration

SSH into one of the border leaf switches and run the following commands:

- **show bgp summary group l3rtr**

```
root@dc1-border-leaf1> show bgp summary group l3rtr

Warning: License key missing; requires 'bgp' license

Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 8 Peers: 14 Down peers: 2
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      56         41          0          0          0          0
bgp.evpn.0
                   14228       7114          0          0          0          0
inet6.0
                       2          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1              65000      54319        587       0       0 2w4d 20:37:13 Establ
  inet.0: 8/8/8/0
```

Uplink BGP session established towards Edge Node

- **show route receive-protocol bgp 10.0.0.1 table inet.0**

```
root@dc1-border-leaf1> show route receive-protocol bgp 10.0.0.1 table inet.0

Warning: License key missing; requires 'bgp' license

inet.0: 46 destinations, 69 routes (45 active, 0 holddo          idden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.9.11.0/24            10.0.0.1               0                 65000 ?
* 10.9.22.0/24            10.0.0.1               0                 65000 ?
* 192.168.30.0/24         10.0.0.1               0                 65000 ?
* 192.168.31.0/24         10.0.0.1               0                 65000 ?
* 192.168.32.0/24         10.0.0.1               0                 65000 ?
* 192.168.33.0/24         10.0.0.1               0                 65000 ?
* 192.168.34.0/24         10.0.0.1               0                 65000 ?
* 192.168.35.0/24         10.0.0.1               0                 65000 ?
```

Overlay route. Advertisement from Edge Node

- **show route advertising-protocol bgp 10.0.0.1 table inet.0**

```
root@dc1-border-leaf1> show route advertising-protocol bgp 10.0.0.1 table inet.0

Warning: License key missing; requires 'bgp' license

inet.0: 46 destinations, 69 routes (45 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    65003 I
* 10.0.0.1/32                                                     I
* 192.168.255.0/32        Self                                    64512 I
* 192.168.255.1/32        Self                                    64513 I
* 192.168.255.2/32        Self                                    I
* 192.168.255.3/32        Self                                    64512 64515 I
* 192.168.255.4/32        Self                                    64512 64516 I
* 192.168.255.5/32        Self                                    64512 64517 I
* 192.168.255.6/32        Self                                    64512 64518 I
* 192.168.255.7/32        Self                                    64512 64519 I
* 192.168.255.8/32        Self                                    64512 64520 I
```

Default route advertised to Edge Node

## VMware NSX-T: Verify Overlay Connectivity (East-West)

To test east-west traffic, run ping tests between the VMs across segments and between the Linux VMs created in VMware vSphere: Create VMs in the Segments.

Following is the flow shown in Figure 91 on page 89:

1. The ping from VM11-1 (on ESXi1_2 host) traverses the Fabric from ESI leaf to reach Border Leaf1.
2. From border-leaf, it's sent towards the Edge nodes hosted on ESXi1_4.
3. From Edge node the traffic is sent towards T1-GW which in turn sends the ping traffic using the TEP port on ESXi1_4 to reach TEP port of ESXi1_3, which then reaches the VM22-1.

**Figure 91: VM to VM Traffic Flow (East-West)**

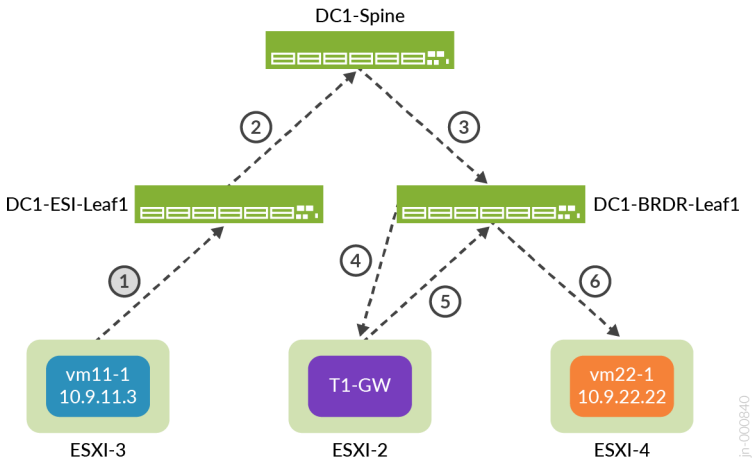

**Figure 92: VM11 Connected to vn11 Able to Ping VM22 Connected to vn22 Logical Segment**

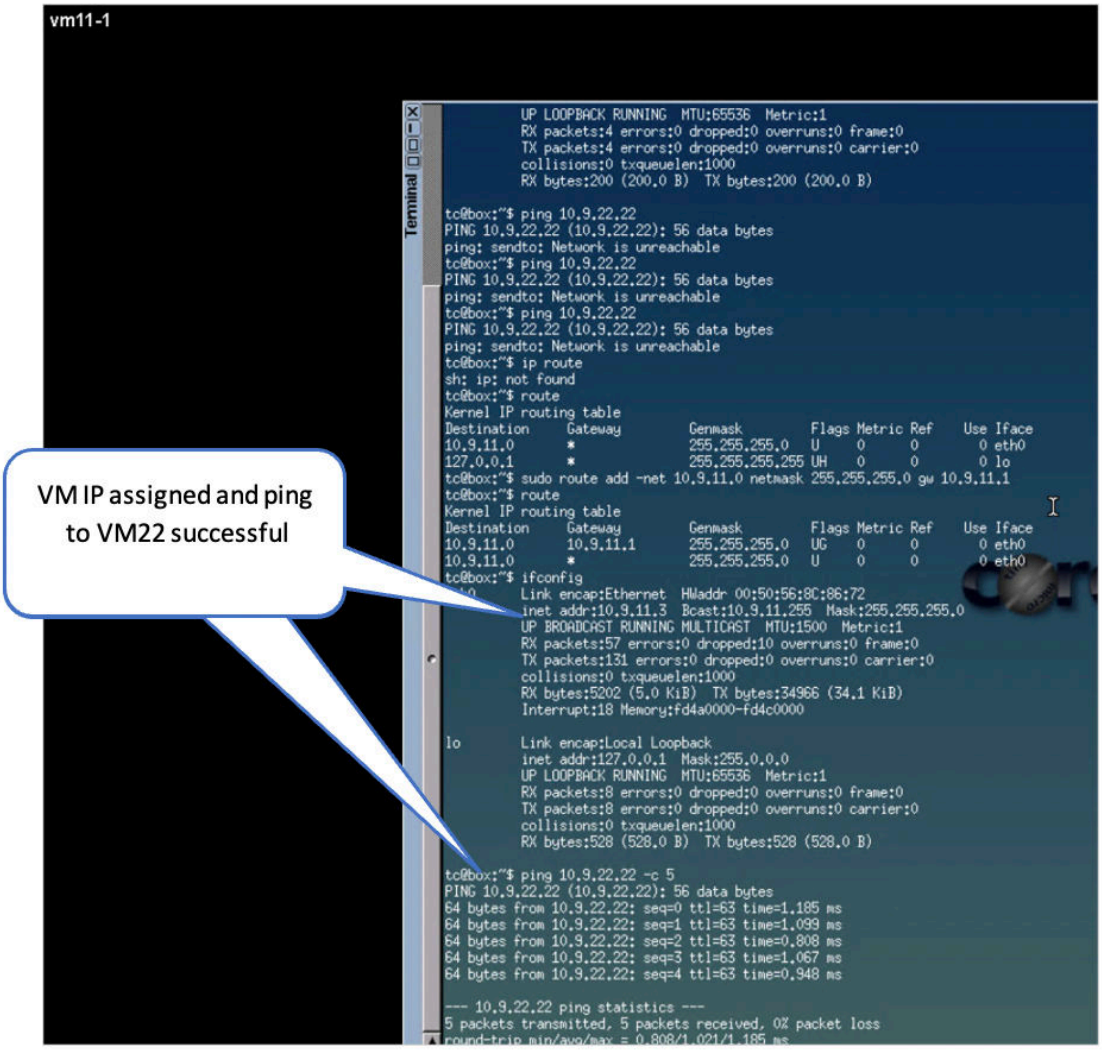

**Figure 93: VM22 Connected to vn22 Logical Segment Pinging VM11 Connected to vn11**

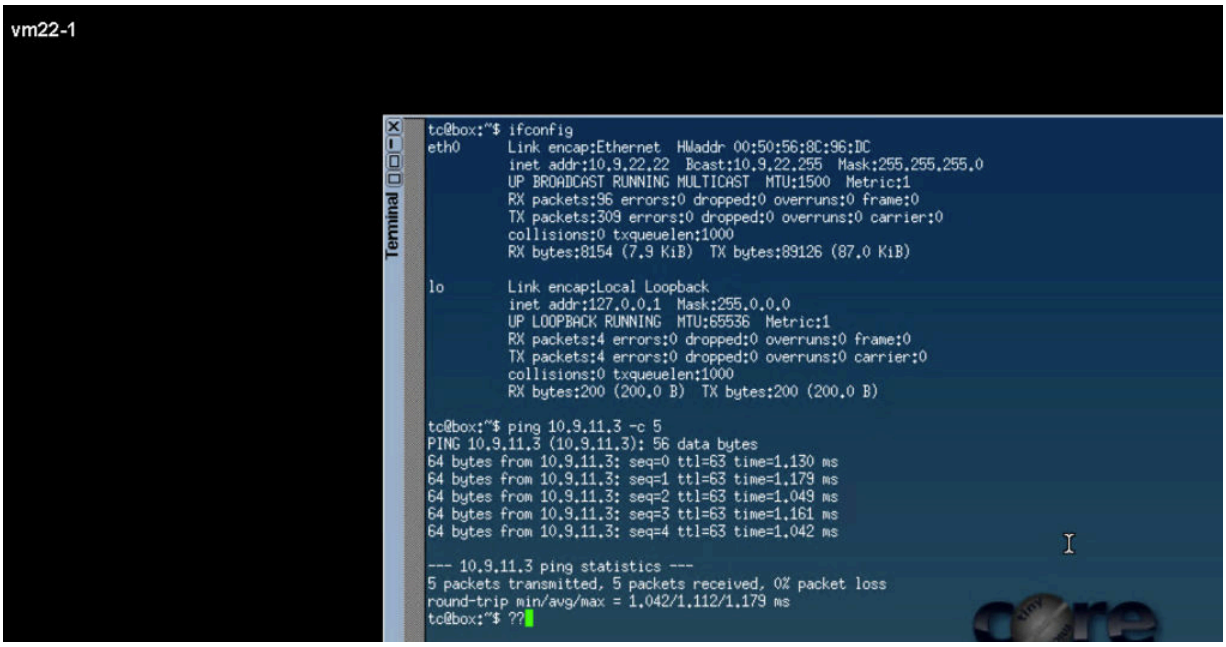


Traceroute to the overlay VMs from the Border Leaf Switches shows the path taken.

**Figure 94: Path towards VM11 from Border Leaf Switches**

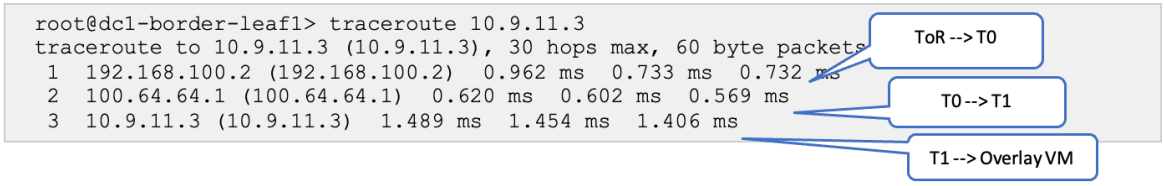


## (Optional) Juniper Apstra: Adding vSphere Server to Juniper Apstra

This step is optional, but the integration of vSphere provides an added layer of visibility into MACs, VMs, and ARPs. It also enables you to view all underlying VMs and docker containers connected with each fabric leaf device connected through the ESXi server.

Add the vSphere server into the blueprint that is managing the fabric:

1. In the Juniper Apstra UI, navigate to **Blueprints** > **<blueprint-name>** > **Staged** > **Virtual** > **VirtualInfra** > **AddVirtualInfra**.

2. Set the VLAN Remediation Policy VN Type to **VXLAN**.
3. Set the Routing Zone to **NSX-T**.
4. Navigate to **Blueprints** > **<blueprint-name>** > **Active** > **Query**. All the VMs associated with the fabric can be viewed here. For more information, refer to the Juniper Apstra 4.2 User Guide.

# Validation Framework

**IN THIS SECTION**


- Test Bed | **93**
- VRF Characteristics | **94**
- Platforms / Devices Under Test (DUT) | **95**
- Test Bed Configuration | **96**


Extensive testing of best practice architectures is the key to the Juniper Validated Design program. JVDs qualify and quantify these best practice architectures, allowing you to know exactly what you're buying and to spend your time deploying and managing your network instead of designing it.

JVDs employ a layered testing approach to deliver reliability and repeatability. Individual features receive functional testing. Multifunction testing builds on this functional testing to see if multiple features work together. Product delivery testing builds upon multifunctional testing to validate that these features combined perform as expected for tested use cases, and JVD testing builds upon product delivery testing by testing multiple products together (including third-party integrations where appropriate) to ensure that all these products combined make an industry-leading solution.

**Figure 95: Validation Framework**

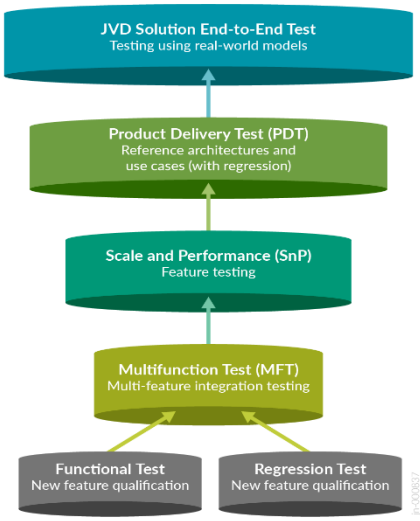


Testing with real-world applications and traffic provides more accurate data regarding performance and response to different configurations. The standardized nature of JVDs ensures the same network architecture is deployed in multiple testing environments, and the use of JVDs by multiple customers allows for any lessons learned in production deployments to rapidly benefit all JVD customers. The more JVDs that are deployed worldwide, the greater the value they provide to all.

## Test Bed

The test bed environment consists of two 3-stage EVPN/VXLAN fabric (DC1 and DC2) with Juniper Apstra with two ESI Server leaf switches, one single Server leaf (non-ESI), and two border leaf switches connected to two spines. The JVDE steps covered in this document to set up NSX-T refer to one of the DC (DC1 blueprint) as an example. However, the DC2 setup also follows a similar setup for the NSX-T configuration, but this is not shown for document brevity.

VMware NSX-T manager is hosted on VM ESXi managed by VMware Vsphere. VMware NSX-T edge switch (NSX_EDGE_01) is hosted as VM on the ESXi, connecting the two border leaf switches for left uplink (to border leaf switch one), right uplink (to border leaf switch two), and overlay (multi-homed).

The second three-stage DC (DC2 Blueprint) setup is DCI connected (using Layer 2 stretch) to the DC1 to test for remote VMs on separate Pods can be reached. The DCI connectivity is out of the scope of this JVDE document and will be covered separately in another JVDE.

The VDS in DC2 was created similarly to the steps explained in this JVDE and the hosts (ESXI2_1, ESXI2_2 and ESXI2_3) as shown in Figure 96 on page 94 was also added to NSX-T so that the VMs in DC2 could be reached from DC1 VMs which use the same logical segments (vn11, vn22).

**Figure 96: The 3-Stage Fabric with Juniper Apstra and NSX-T Test Environment**

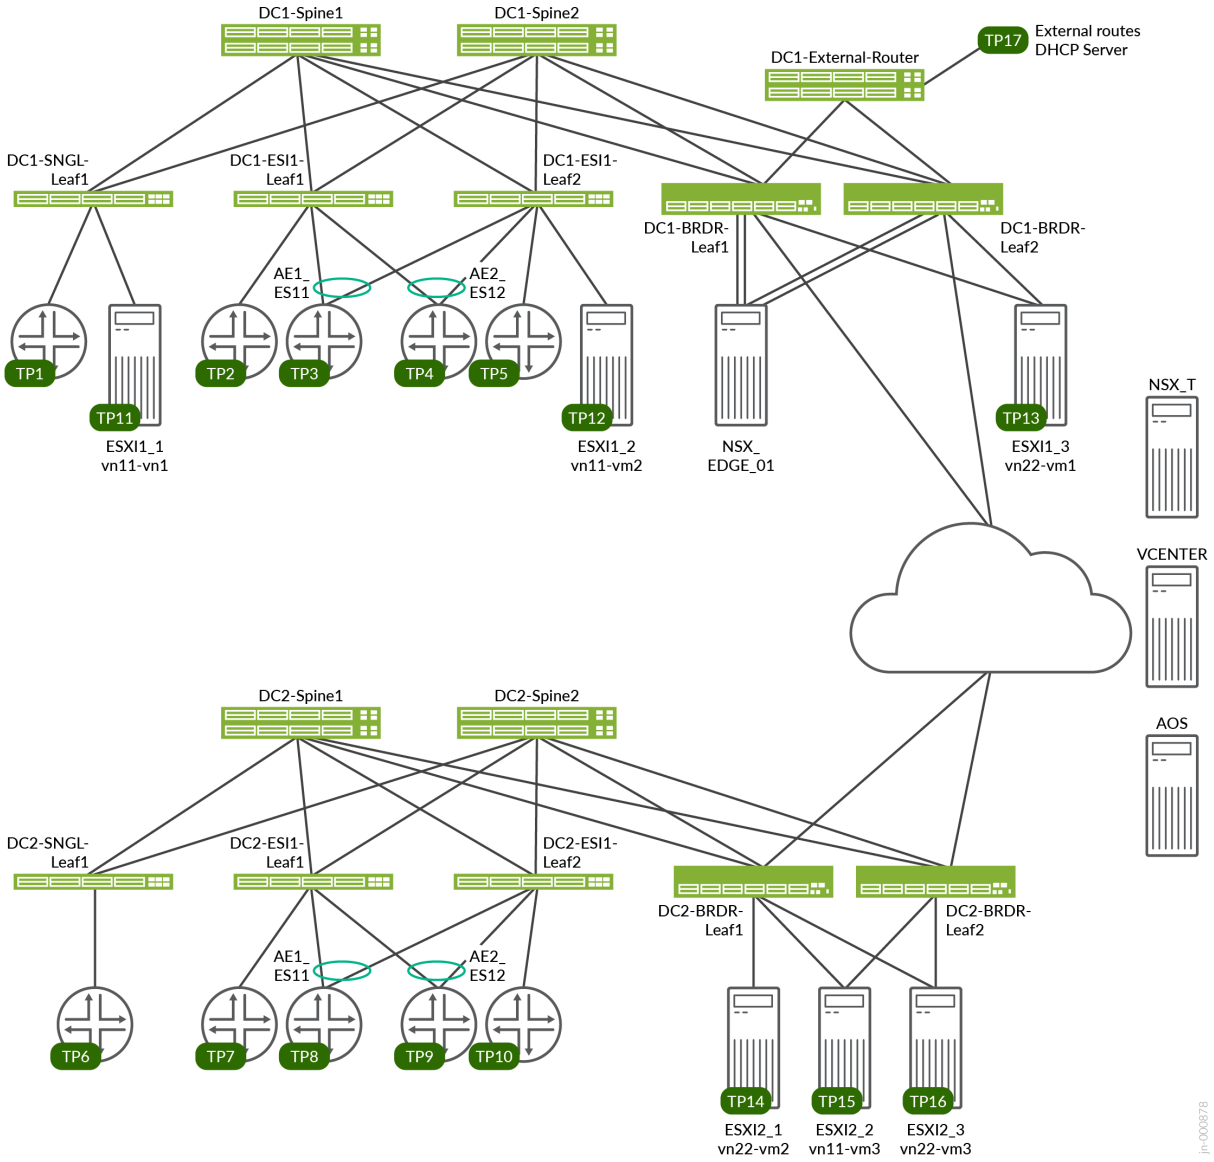


## VRF Characteristics

### RED VRF

- VLANs 400–649 with IRB v4/v6
- on DC1-SNGL-LEAF1 single access port
- on DC1-ESI-LEAF1 single access port, AE1 and AE2

- on DC1-ESI1-LEAF2 single access port, AE1 and AE2
- on DC1-BRDR-LEAF1 to distribute routes to external-router
- on DC1-BRDR-LEAF2 to distribute routes to external-router
- VLANs 400–649 on each test port with 10 unique MAC/IP per VLAN
- DHCP client on TP3
- External DHCP server on TP17

**Blue VRF**

- VLANs 3500–3749 with IRB v4/v6
- on DC1-SNGL-LEAF1 single access port
- on DC1-ESI-LEAF1 single access port, AE1 and AE2
- on DC1-ESI1-LEAF2 single access port, AE1 and AE2
- on DC1-BRDR-LEAF1 to distribute routes to external-router
- on DC1-BRDR-LEAF2 to distribute routes to external-router
- VLANs 3500–3749 on each test port with 10 unique MAC/IP per VLAN
- DHCP client on TP3, TP4, TP5
- External DHCP server on TP2

## Platforms / Devices Under Test (DUT)

**Table 7: Devices Under Test (Validated Devices)**

| Devices Under Test (Validated Devices) | | | |
|---|---|---|---|
| Solution | Server Leaf Switches | Border Leaf Switches | Spine |
| 3-stage EVPN/VXLAN (ERB) | QFX5120-48Y-8C | QFX5130-32CD | QFX5220-32CD |
| | QFX5110-48S | QFX5700 | QFX5120-32C |
| | | ACX7100-48L | |

**Table 7: Devices Under Test (Validated Devices)** ***(Continued)***

| Devices Under Test (Validated Devices) | |
|---|---|
| | ACX7100-32C |
| | PTX10001-36MR |
| | QFX10002-36Q |

## Test Bed Configuration

Contact Juniper or your Juniper account representative to obtain the full archive of the test bed configuration used for this JVD.

# Test Objectives

**IN THIS SECTION**


- Test Goals | **97**
- Test Non-Goals | **97**


The primary objective of this JVDE testing is the qualification testing of the 3-Stage Data Center Reference Design with Juniper Apstra and VMware NSX-T. The design is based on an ERB (Type 2 and Type 5) EVPN/VXLAN fabric with spine, server leaf, border leaf, and NSX-T Edge Gateway devices.

In the same way that server virtualization programmatically creates and manages virtual machines, NSX-T Data Center network virtualization programmatically creates and manages virtual networks. NSX-T Data Center works by implementing three separate but integrated planes: management, control, and data. These planes are implemented as a set of processes, modules, and agents residing on two types of nodes: NSX Manager and transport nodes.

The Juniper data center fabric is viewed as an external network, which is defined as a physical network or VLAN not managed by NSX-T Data Center. The NSX logical network or overlay network can be linked to an external network through a Tier-0 router or NSX Edge.

The goal is to ensure the design is well-documented and will produce a reliable, predictable deployment for the customer.

The qualification objectives include validation of blueprint deployment, device upgrade, incremental config pushes/provisioning, telemetry and analytics checking, failure mode analysis, and verification of host traffic.

## Test Goals

This JVDE for integration of NSX-T with 3-Stage Fabric using Juniper Apstra focuses on the following:

- Design and deployment of NSX-T and vSphere components such as NSX-T Edge Node, ESXi servers, VDS switch, and port groups.
- Configuring NSX-Manager with Edge Nodes, Transport Nodes, Tier-0 and Tier-1 gateways.
- Connectivity with fabric border leaf switches using loopback and links.
- Apstra configuration for connectivity NSX-Edge Nodes, using Virtual Networks, Routing-Zone (VRF), connectivity templates for IP link, BGP, and static routes
- Validation of end-to-end traffic flow
- Test for any anomalies, view the VM, ARP, and MAC collection from Juniper Apstra.

In order to pass validation, the 3-stage Fabric with Juniper Apstra and NSX-T integration must also pass the following scenarios:

- East-West connectivity - Host connectivity between fabric connected hosts created by Apstra towards NSX managed hosts, covered in "VMware NSX-T: Verify Overlay Connectivity (East-West)" on page 89.
- North -South connectivity – Connectivity between NSX Edge Node and border leaf switches proving external connectivity as covered in "VMware NSX-T: Verify BGP Session on Edge" on page 84.

## Test Non-Goals

Test non-goals for this JVD were to test the following:

- Establish fabric connectivity
- Deployment and configuration of the fabric.
- DCI Interconnectivity between data centers

# Results Summary and Analysis

This JVDE focuses on the validation of the VMware NSX-T integration with data center fabric using Apstra provisioning. The validation carried out shows how external systems like VMware NSX-T can be integrated using Apstra providing visibility into the networking details of virtual machines (VMs) and containers hosted by ESXi servers, which are connected to leaf switches managed by Juniper Apstra as part of the JVD data center network fabric.

**Configuration and Integration Tests**

The JVDE test validation involved creating NSX-T components needed for connectivity with the data center fabric, here are some of the Vmware NSX-T components on NSX-T manager:

- Adding the vSphere as compute manager to configure NSX-T on selected ESXi servers connected to the TOR switches.
- Configure left and right uplink VLANs and overlay uplink.
- Configure VDS switch and associate it to the uplinks and overlay uplink VLANs.
- Configure the edge Node as VM on the ESXi host connected to the border leaf.
- Create and configure logical segments for microsegmenting VM networks, the logical segments can then be seen in vSphere to use as network adapters in vSphere for VMs to communicate East-West.
- Create and configure T0 (including loopback interface to bgp peer with border leaf switches, Left and right uplink interface to connect to border Leaf switches) and T1 gateways for north-south communication with border leaf switches (serving as gateways to the data center fabric). The geneve tunnels terminate on border leaf switches and the network packets are converted to EVPN VXLAN packets.

Furthermore, the JVDE also shows the configuration required on Juniper Apstra so as to allow the network traffic flow for inter-host (ESXi) VMs and intra-host (ESXi) VMs scenarios which are connected to leaf switches using the overlay VLAN. Following are the configurations that were successfully configured on Apstra:

- Add vSphere and NSX-T manager as External Network Providers and then add them as Virtual Infra in the blueprint.

- Create Routing Zone for NSX-T traffic and then associate with the overlay VLAN layer 2 Virtual Network assigned to all the fabric leaf switches.
- Create Connectivity Templates for:
  - Creating IP Links using routed interfaces on border leaf switches to Edge Node VM hosted on ESXi server.
  - Creating BGP peering between NSX-T T0 and border leaf switches (border leaf switch one as the left uplink and border leaf switch two as the right uplink) providing resiliency. The BGP peering is also verified from the NSX-T Edge Node VM and border leaf switches.
  - Creating static routes between NSX-T T0 and border leaf switches (same is also configured on NSX-T manager)

Operational and Trigger Tests with NSX-T setup are as follows:

- Verification of NSX-managed host connectivity:
  - Intra-VLAN within NSX
  - Inter-VLAN across NSX
- Change MTU of the overlay transport node to test for configuration anomaly.
- BGP flapping on the border leaf switch is used to detect if other border leaf routes traffic correctly and traffic loss is minimal.
- Rebooting fabric switches one a time to verify if NSX-T detects the tunnels going down for the affected DUT (device under test).

Apart from the above tests described, other non-test goal tests such as reboot switches, reset DHCP bindings, deactivating BGP on leaf switches to cause BFD sessions to converge traffic to minimize traffic loss were also conducted.

The JVDE validation aim was to ensure that the Juniper data center switches can integrate with NSX-T using Apstra to configure configurations required for this setup. This has been successfully validated using the Junos OS release 22.2R3-S3 or Apstra 4.2.1.

# Recommendations

Juniper Apstra and NSX-T integration provides below benefits:

1. Apstra software can connect to the NSX-T API to gather information about the inventory in terms of hosts, clusters, VMs, portgroups, vDS/N-vDS, and NICs within the NSX-T environment.

2. NSX-T integration with Apstra provides operator visibility into the VMs, VM ports, and their connectivity to the ToR switch.
3. NSX-T integration helps identify issues on the fabric and on the virtual infrastructure.
4. Apstra Virtual Infrastructure visibility helps provide underlay/overlay correlation visibility and uses IBA analytics for overlay/underlay.
5. By using Apstra, NSX-T deployments can be accelerated as the fabric is ready in terms of LAG, MTU, and VLAN configuration as per NSX-T transport node requirements.
6. Juniper Hardware tested for the Junos release Junos 22.2R3-S3 are listed in Table 7: Devices under Test (Validated devices) on page 95.
7. In terms of software versions, all software versions listed in Table 3 on page 7 and Table 4 on page 7 are recommended. Any deviations from these should be tested thoroughly.