VALIDATED REFERENCE DESIGN GUIDE

AOS-CX SWITCH AND VMWARE NSX-T INTEROP SOLUTION GUIDE

aruba
a Hewlett Packard Enterprise company

www.arubanetworks.com

© Copyright 2019 Hewlett Packard Enterprise Development LP

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

CONTENTS

Aruba AOS-CX Switch and VMware NSX-T Interop Solution Guide

Introduction

VMware NSX is a network virtualization solution that virtualizes multiple network functions like Switch, Router, firewall, NAT, and VPN. This document provides guidance on setting up the VMware NSX-T and AOS-CX Switch interop. This interop solution provides L2 network connectivity between Virtual Machines in different Racks and attached to NSX-T logical switches. The NSX-T Logical Router provides L3 routing capability between L2 segments within the NSX Virtual world.
This document covers three scenarios.

1. Layer 2 Data Centers with AOS-CX switches used for the underlay
2. Layer 3 Data Centers with AOS-CX switches used for the underlay
3. BGP routing with AOS-CX and NSX-T edge VM

AOS-CX 10.3 on 8320, vSphere 6.7 and NSX-T Version 2.4.2 were used in the creation of this guide.
This document applies to all other AOS-CX switching products also.

Figure 1 describes the network topology that will be used in this document

- An Out of Band (OOB) management network (10.10.8.0/24) is used for communication between the vCenter, NSX-T Nodes. The Aruba AOS-CX switches (8320) were used to interconnect the ESXi Nodes.
- The Inband underlay network (99.99.11.0/24, 99.99.22.0/24) is used for connectivity between the ESXi Nodes, which will be used as Tunnel End Points (TEP) for Generic Network Virtualization Encapsulation (GENEVE) tunnels. The underlay network can be a L2 or L3 network since all tunnels are created and torn down by NSX-T.
- Two logical switches / virtual networks (192.10.10.0/24 & 192.20.20.0/24) are used for overlay connectivity between the VM's and the Tier-1 Logical router which is used to interconnect both logical switches.

Brief comparison between NSX-T and NSX-V

The current shipping solution is VMWare NSX-V.

- NSX-V is tightly integrated with VMware and requires VMware vCenter
- Uses VXLAN for overlay encapsulation
- Leverages Uses DLRs (distributed logical router) for centralized routing within vSphere.

NSX-Transformers (NSX-T) adds support for multi-hypervisor environments which enables NSX-T to also support KVM, Docker, Kubernetes, and OpenStack as well as AWS.

- NSX-T can be deployed without vCenter
- Uses GENEVE for overlay encapsulation, VXLAN is not required for overlay with NSX-T
- supports multiple vCenters
- uses a two-tier distributed routing model
- Supports multi-hypervisors
    - VMware vSphere(ESXi)

- Kernel-based Virtual Machine (KVM)

NSX-T supports Hybrid Cloud Networking and native AWS deployments, also it can be integrated with Docker, and Open Stack

Requirements

- Ensure DNS and NTP server infrastructure are in place
  - all devices (ESXi host, vCenter, NSX-T) nodes should point to these
  - NTP is in sync on these devices
  - DNS resolution between devices should work (all devices should have DNS host entries)
- VMware NSX-T should be deployed according to instructions stated here
  - https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.4/installation/GUID-67731519-E70F-4BC5-87CD-9F426E250349.html
- Make sure all the components (ESXi Version, ESXi Drivers, Bare metal its firmware, vCenter) are compatible with NSX-T 2.4.2 version
- Installation of ESXi's, vCenter & NSX-T Manager & Controller
- Utilize the flash based web client (FLEX) instead of HTML 5, some NSX features only exist in the FLEX client

Note : NSX-v and Aruba AOS-CX integration is documented

https://arubapedia.arubanetworks.com/arubapedia/images/5/53/VMware_NSXv_and_8325_Integration.pdf

Physical Topology

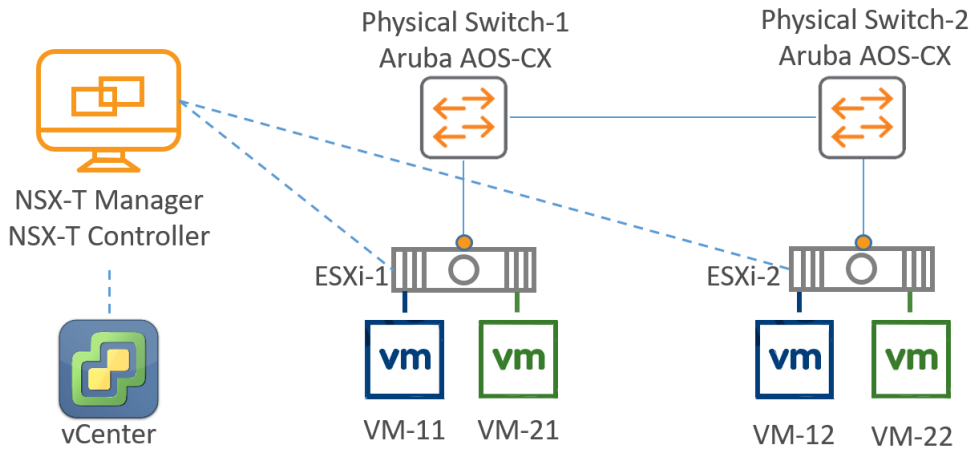

Figure 1: Physical Topology

ESXi node physical connections

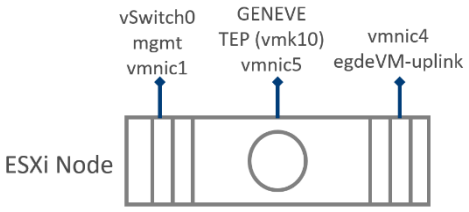

Figure 2: ESXi Node connections

vSphere information

- In this example, there are two clusters, each Hypervisor which is part of NSX fabric acts as one TEP. Usually the cluster deployments are workload clusters and Management cluster. Work load clusters hosts the tenant workload VMs (Example: Production, Development test VM's).
- In this example, Cluster (SVT-DC & SVT-DR) hosts the workload VM's and each ESXi node in the cluster terminates a GENEVE Tunnel End Point (TEP) to interconnect the VM's.
- Management cluster (Mgmt-Infra) hosts the management VMs such as NSX Manager, Controllers, vCenter VMs and vRA, vROPS appliances.

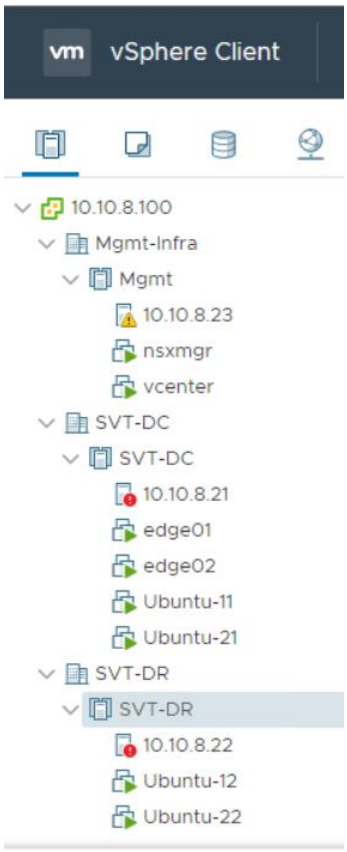

NSX-T Preparation

- Configure the underlay hardware (AOS-CX switches) and interfaces involved in the GENEVE tunnels with MTU min as 1600. In this example configured `mtu is 9000`. If the ESXi nodes are L3 apart, then configure all the transit interfaces as route interfaces with `ip mtu 9000`
- Add VMware vCenter to NSX-T as compute Manager

  Launch NSX-T

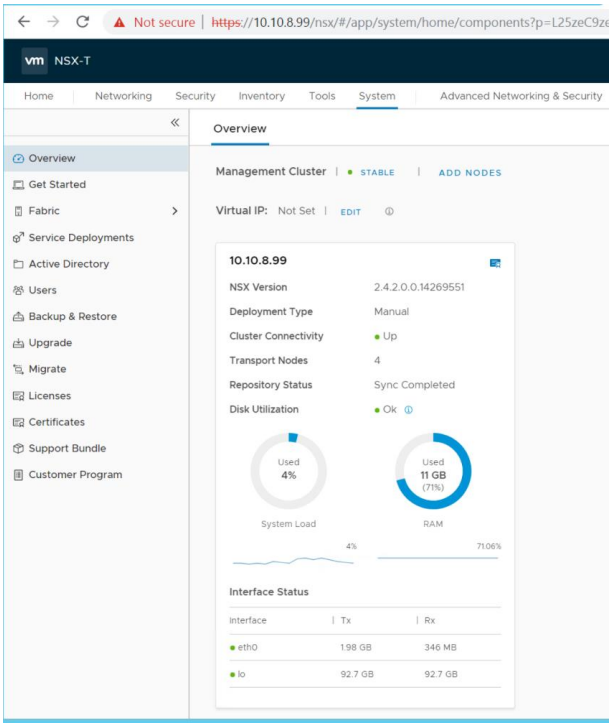

- Add vCenter as Compute Manager

  System > Fabric > Computer Managers > Add

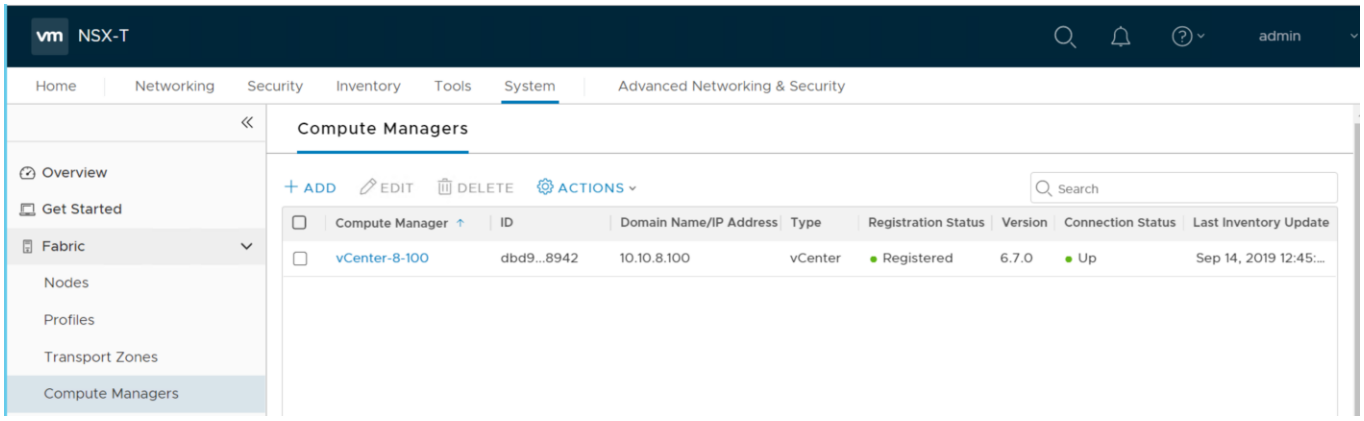

Usecase-1: Data Centers are L2-Adjacent

In this use case (ref to *Figure 1: Physical Topology)*, there is one data center, both ESXi nodes are in same data-center which are L2-Adjacent. The uplink ports from the servers are configured within a VLAN on the AOS-CX Switch. The AOS-CX Switch stretches that same VLAN to the other AOS-CX switches so that the ESXI hosts are in the same L2 VLAN Segment.

NSX-T Configuration

- Prepare **Transport Zone** (Overlay) & N-VDS

  System > Fabric > Transport Zones

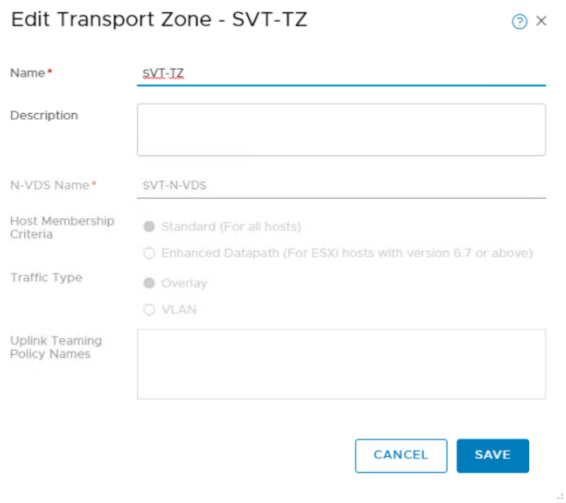

- Prepare Uplink **Profile** (select a physical interface, in this case its vmnic5)

  System > Fabric > Profiles

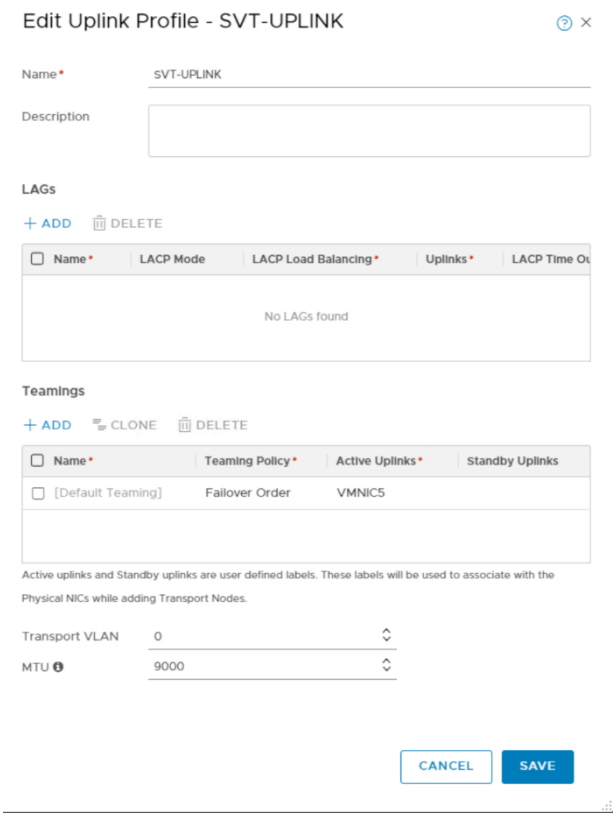

- Navigate to **System>Nodes**, select vCenter under "**Managed By**", then select specific ESXi Node/ESXi Cluster click on **Configure NSX**, associate TZ, Uplink, IP Pool.
    - Create The **IP Pool, this ip pool** is used to assign IPs to various TEP as GENEVE-TEP-IP, in this example we used 99.99.11.0/24 as the pool

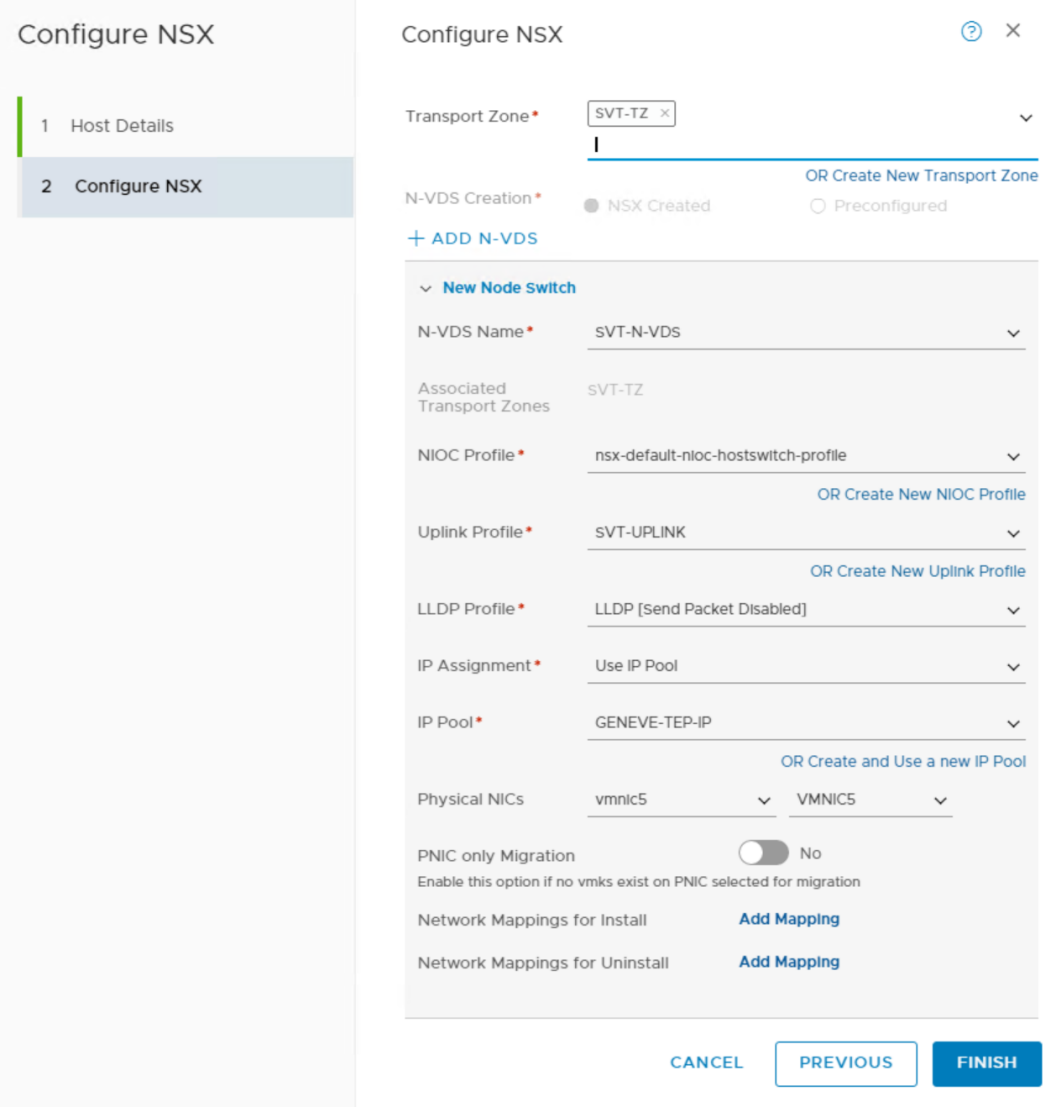

With the above steps, NSX-T will install NSX vibs on the ESXi hosts and configure the GENEVE tunnels. Here is the status after NSX configured properly.

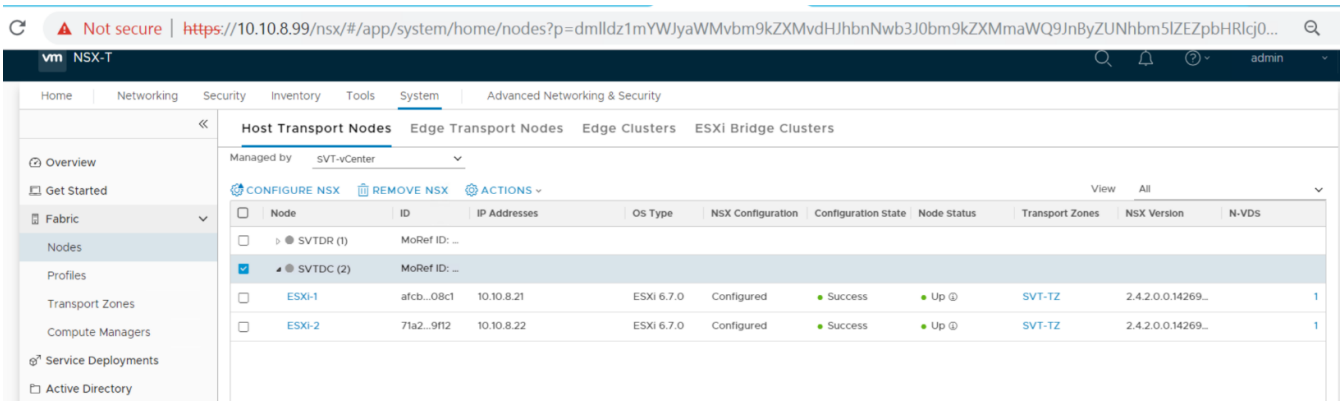

The below screenshot shows the GENEVE tunnel between each ESXi node to the other node. Notice the IP address of both ESXi servers are in same subnet (L2-Adjacent)

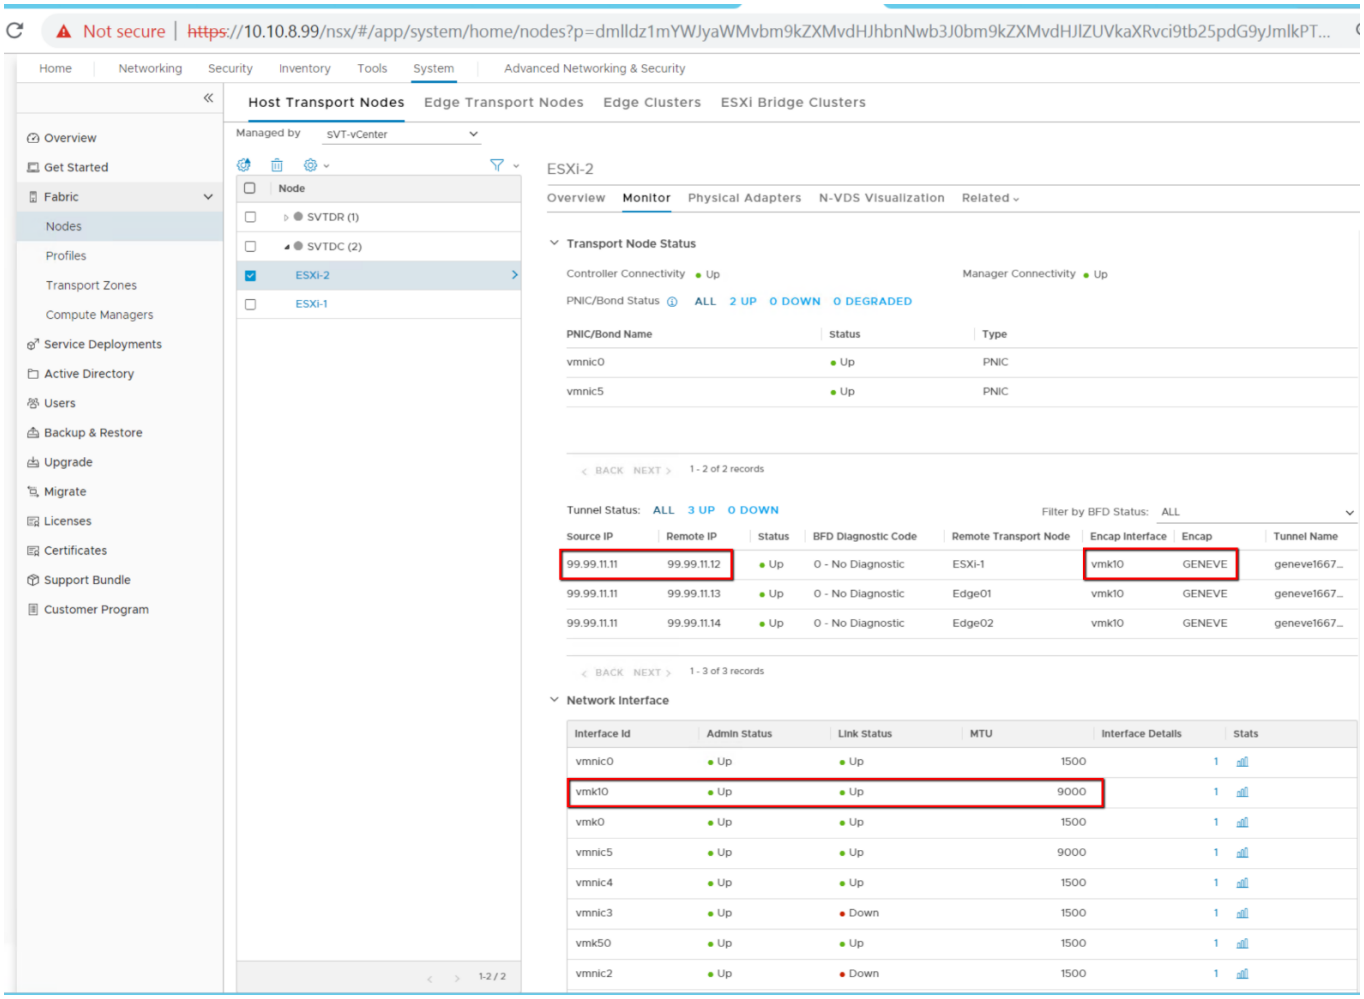

Aruba AOS-CX Switch configuration

The physical underlay infrastructure needs to support jumbo frames. The below configuration uses and MTU of 9000 between the ESXi Node and Switches.

8320-SW01 Configuration

```
interface 1/1/1
   no shutdown
   mtu 9000
   description To ESXi-1 VMNIC5
   no routing
   vlan trunk native 1
   vlan trunk allowed all
   exit

interface 1/1/46
   no shutdown
   mtu 9000
   description SW1-to-SW2
   no routing
   vlan trunk native 1
   vlan trunk allowed all
   exit
```

8320-SW02 Configuration

```
 interface 1/1/1
   no shutdown
   mtu 9000
   description To ESXi-2 VMNIC5
   no routing
   vlan trunk native 1
   vlan trunk allowed all
   exit

interface 1/1/46
   no shutdown
   mtu 9000
   no routing
   description SW2-to-SW1
   vlan trunk native 1
   vlan trunk allowed all
   exit
```

Verifications from the CX Switch

SW01- Verification

```
8320-SW1# show interface brief
1/1/1     1       trunk  SFP+DA3     yes     up                              10000
1/1/46    1       trunk  SFP+DA3     yes     up                              10000

8320-SW01# show vlan
--------------------------------------------------------------------------------
VLAN  Name                                Status  Reason     Type       Interfaces
--------------------------------------------------------------------------------
8     MGMT                                up      ok         static     1/1/1,1/1/46
```

SW02- Verification

```
8320-SW02# show int brief
1/1/1     1       trunk   SFP+DA3      yes      up                                10000
1/1/46    1       trunk   SFP+DA3      yes      up                                10000

8320-SW02# show vlan
-------------------------------------------------------------------------------------------
VLAN   Name                                Status   Reason      Type       Interfaces
-------------------------------------------------------------------------------------------
8      MGMT                                  up       ok        static     1/1/1,1/1/46
```

Connect VM's using NSX-T Logical Switch - Data Centers are L2-Adjacent

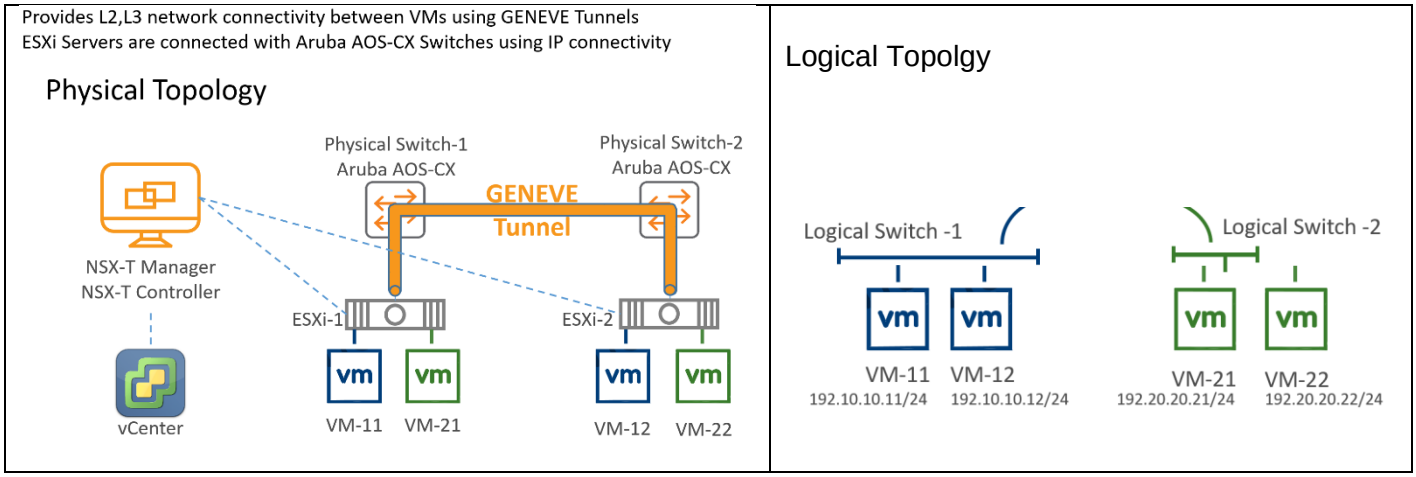

Figure 3: Interconnecting VM's using GENEVE overlay tunnel

As shown in the above topology, VM-11 and VM-12 are in same L2 segment but hosted in two diff hosts (ESXi). Similarly, VM-21 and VM-22 are in same L2 segment.

To interconnect VM-11 and VM-12, create a Logical Switch and associate the VM network interface with the appropriate Logical Switch. Interconnecting VM's across the logical Switch requires routing, which is covered in next section.

Creating Logical Switch

Navigate to **Advanced Networking & Security > Networking > Switching >** click **Add**

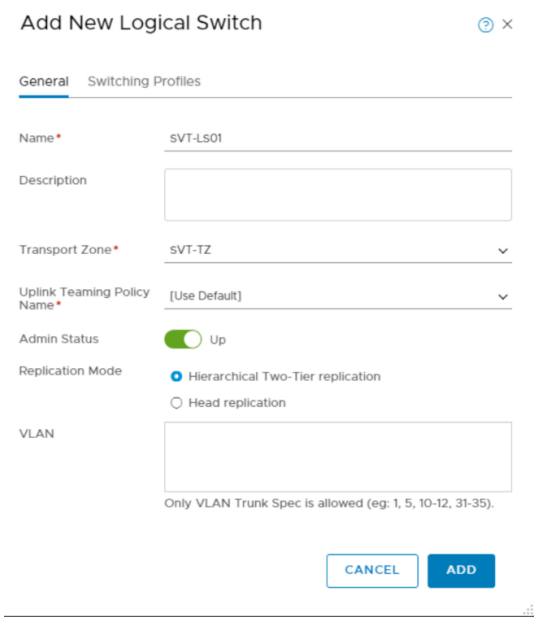

Here are the two logical Switches that were created for the two logical segments

Associate NSX Logical Switch with the respective VM

Associate each NSX Logical Switch with the respective VM's from vCenter as shown below.

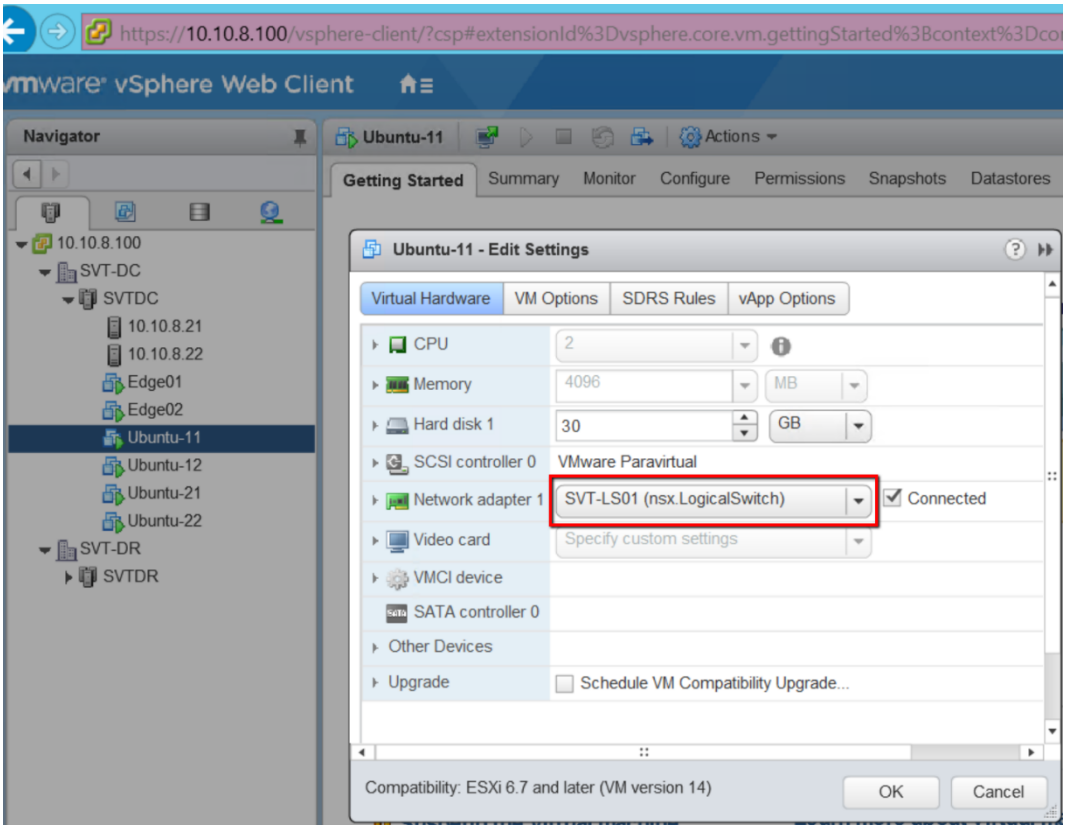

Connectivity test

Connectivity test between the VM's within same Logical Switch

```
ubuntu@ubuntu01:~$ ping 192.10.10.11
PING 192.10.10.11 (192.10.10.11) 56(84) bytes of data.
64 bytes from 192.10.10.11: icmp_seq=1 ttl=64 time=0.025 ms
64 bytes from 192.10.10.11: icmp_seq=2 ttl=64 time=0.028 ms
^C
--- 192.10.10.11 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1008ms
rtt min/avg/max/mdev = 0.025/0.026/0.028/0.005 ms
ubuntu@ubuntu01:~$ ping 192.10.10.12
PING 192.10.10.12 (192.10.10.12) 56(84) bytes of data.
64 bytes from 192.10.10.12: icmp_seq=1 ttl=64 time=0.376 ms
64 bytes from 192.10.10.12: icmp_seq=2 ttl=64 time=0.211 ms
^C
--- 192.10.10.12 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1032ms
rtt min/avg/max/mdev = 0.211/0.293/0.376/0.084 ms
```

GENEVE Wireshark Capture

As shown below, the Wireshark capture shows the inner header as well as the overlay headers created by NSX-T

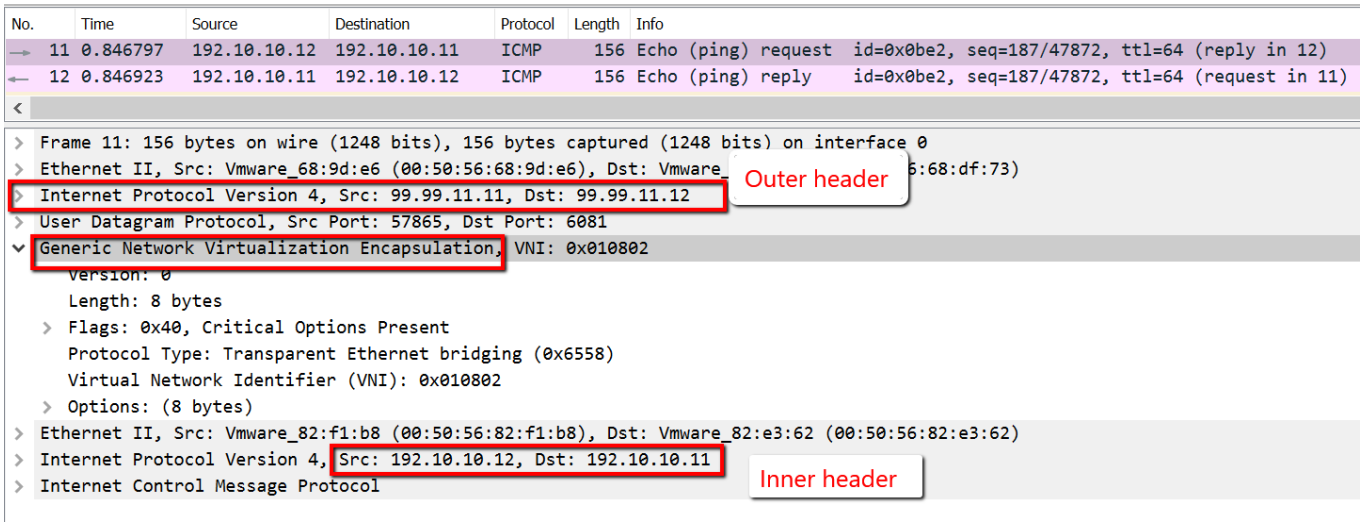

Inter Logical Switch Routing - Data Centers are L2-Adjacent

To enable inter Logical Switch routing, create a Tier 1 Logical Router which can handle the routing between the VMs on different segments.

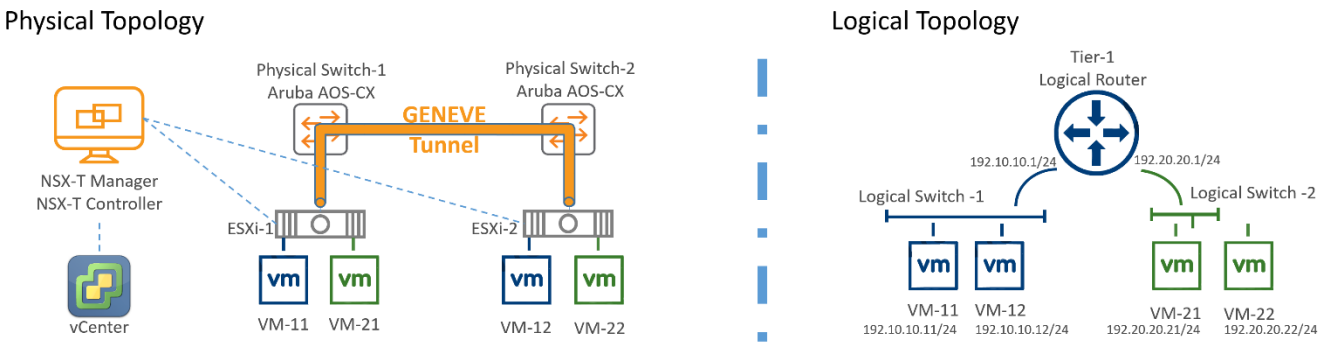

Figure 4: Interconnecting VM's between both logical Switches with NSX-T logical Router

Enable inter logical switch routing

As shown in the topology, enable routing between Logical Switches

Navigate to **Advanced Networking & Security > Networking > Routers > Add > Tier-1 Router**

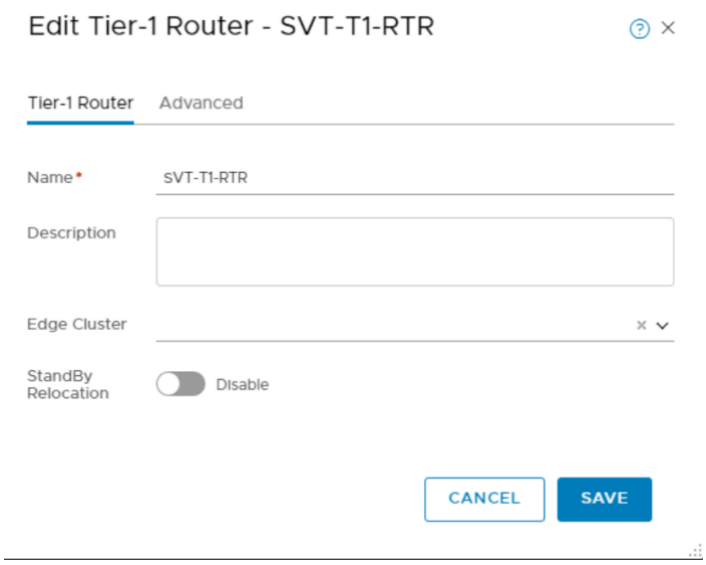

Attach logical router port

Attach a Logical Router Port to the Tier 1 Logical router that was just created for each logical switch (LS) to enable inter LS routing.

Navigate to **Advanced Networking & Security > Networking > Routers > Select the Router (double click) >** on the right pane, go to **Configuration > add** Logical router ports as below

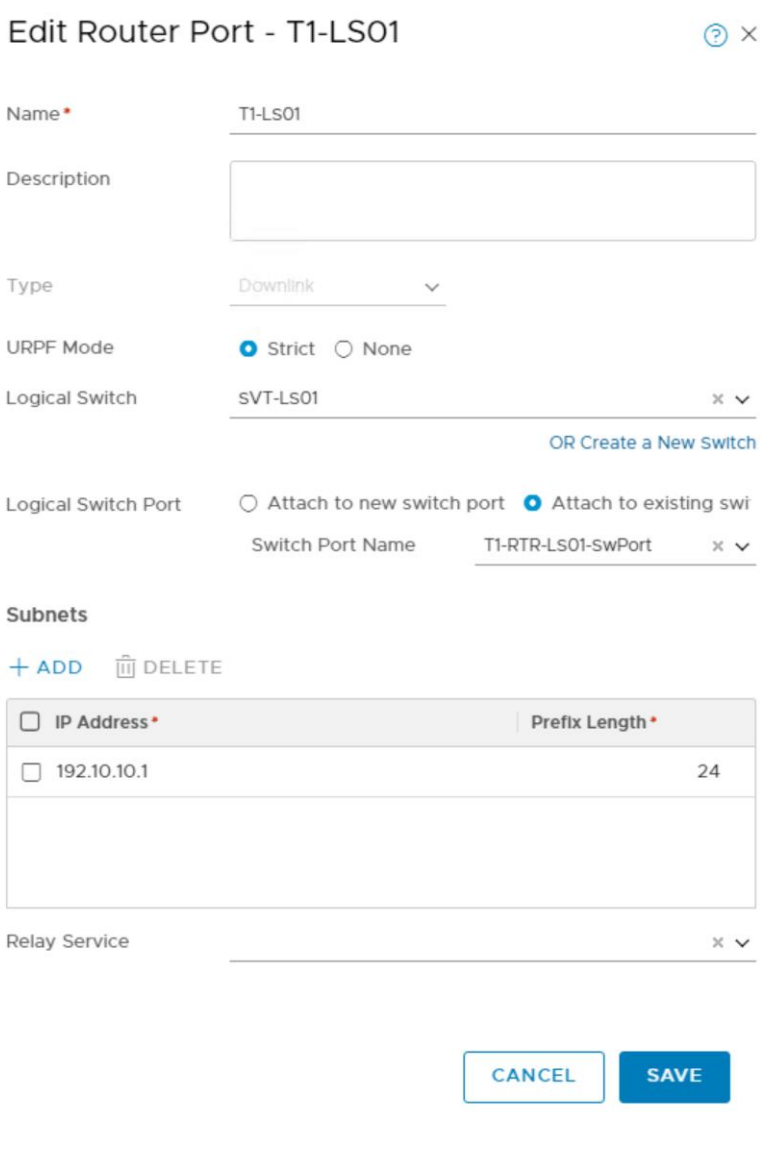

Add one more logical Router port for Logical Switch-2

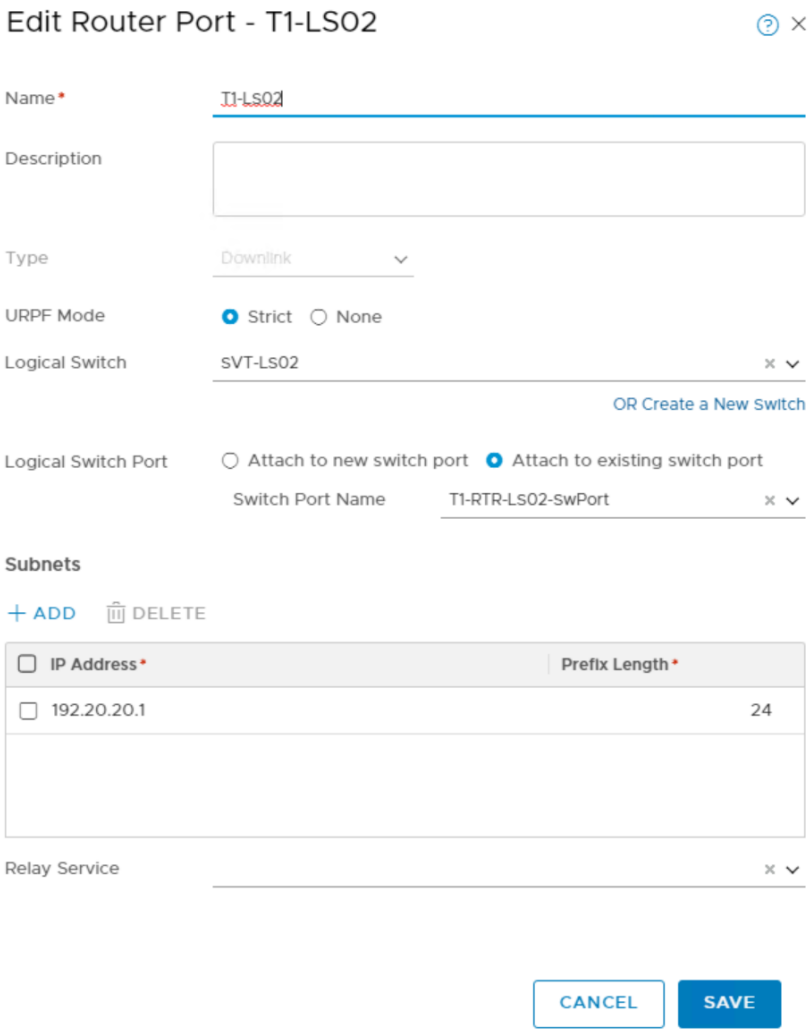

Here is the list of the router interfaces connecting to both Logical Switches.

This will enable routing between the VM's from one logical Switch to other.

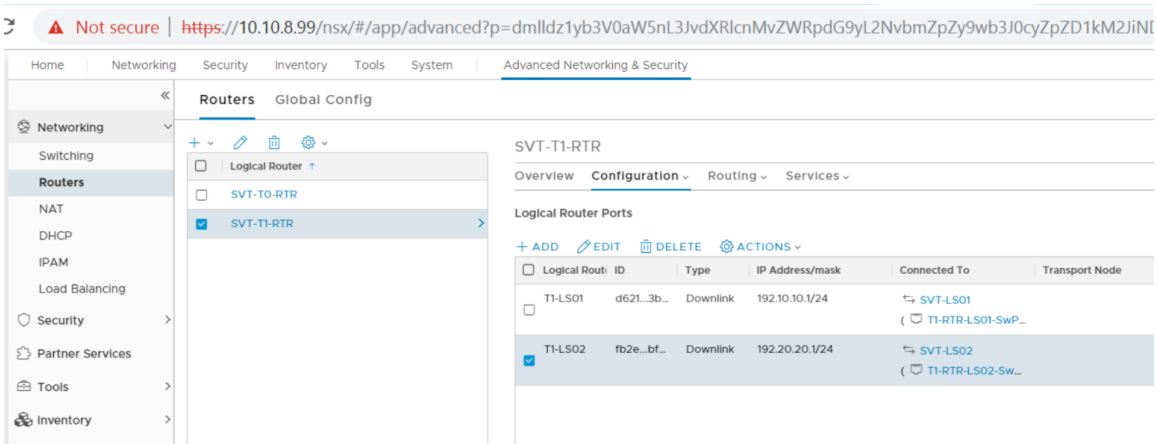

Connectivity test

Connectivity test between the VM's which are in two different Logical Switch.

```
ubuntu@ubuntu01:~$ ping 192.10.10.12
PING 192.10.10.12 (192.10.10.12) 56(84) bytes of data.
64 bytes from 192.10.10.12: icmp_seq=1 ttl=64 time=0.584 ms
64 bytes from 192.10.10.12: icmp_seq=2 ttl=64 time=0.242 ms
64 bytes from 192.10.10.12: icmp_seq=3 ttl=64 time=0.210 ms
64 bytes from 192.10.10.12: icmp_seq=4 ttl=64 time=0.210 ms
^C
--- 192.10.10.12 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3071ms
rtt min/avg/max/mdev = 0.210/0.311/0.584/0.158 ms
ubuntu@ubuntu01:~$ ping 192.20.20.22
PING 192.20.20.22 (192.20.20.22) 56(84) bytes of data.
64 bytes from 192.20.20.22: icmp_seq=1 ttl=63 time=1.36 ms
64 bytes from 192.20.20.22: icmp_seq=2 ttl=63 time=0.345 ms
64 bytes from 192.20.20.22: icmp_seq=3 ttl=63 time=0.304 ms
^C
--- 192.20.20.22 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2013ms
rtt min/avg/max/mdev = 0.304/0.672/1.368/0.492 ms
```

Usecase-2: Data Centers are L3-Apart

In this use case (as shown in below, *Figure 5: Physical Topology)*, both ESXi nodes are in same data-center which are L3-Adjacent.

Since the DC fabric is a routed L3 fabric the uplink ports from each server are configured with IP addresses in separate network segments. Traffic in the underlay between servers needs to be routed by the AOS-CX switches.

In this environment, NSX-T can be used to provide L2 connectivity between hosts and VMs.

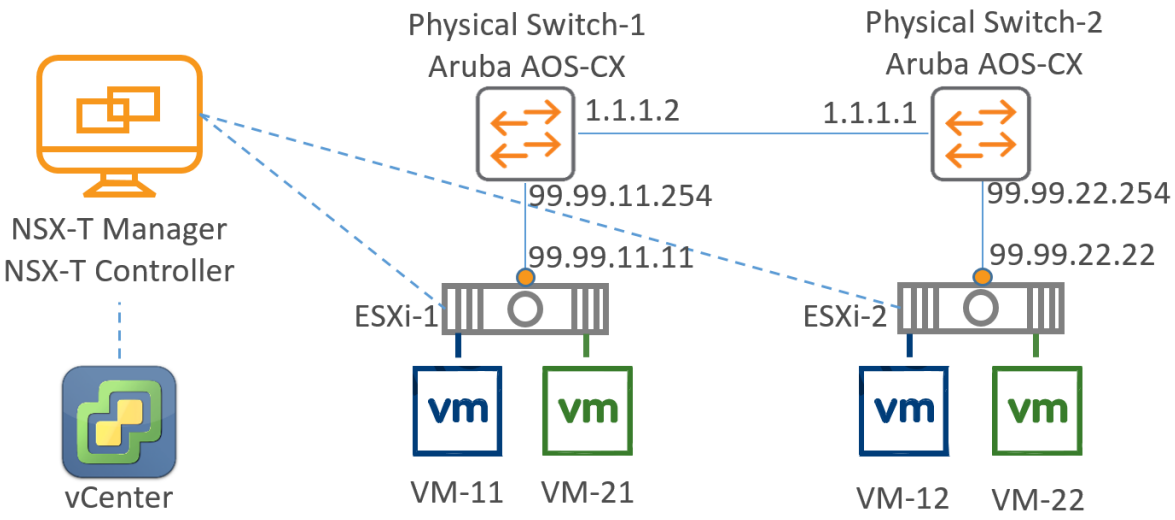

Figure 5: Physical Topology

NSX-T Configuration

- Prepare **Transport Zone** (Overlay) & N-VDS

  System > Fabric > Transport Zones

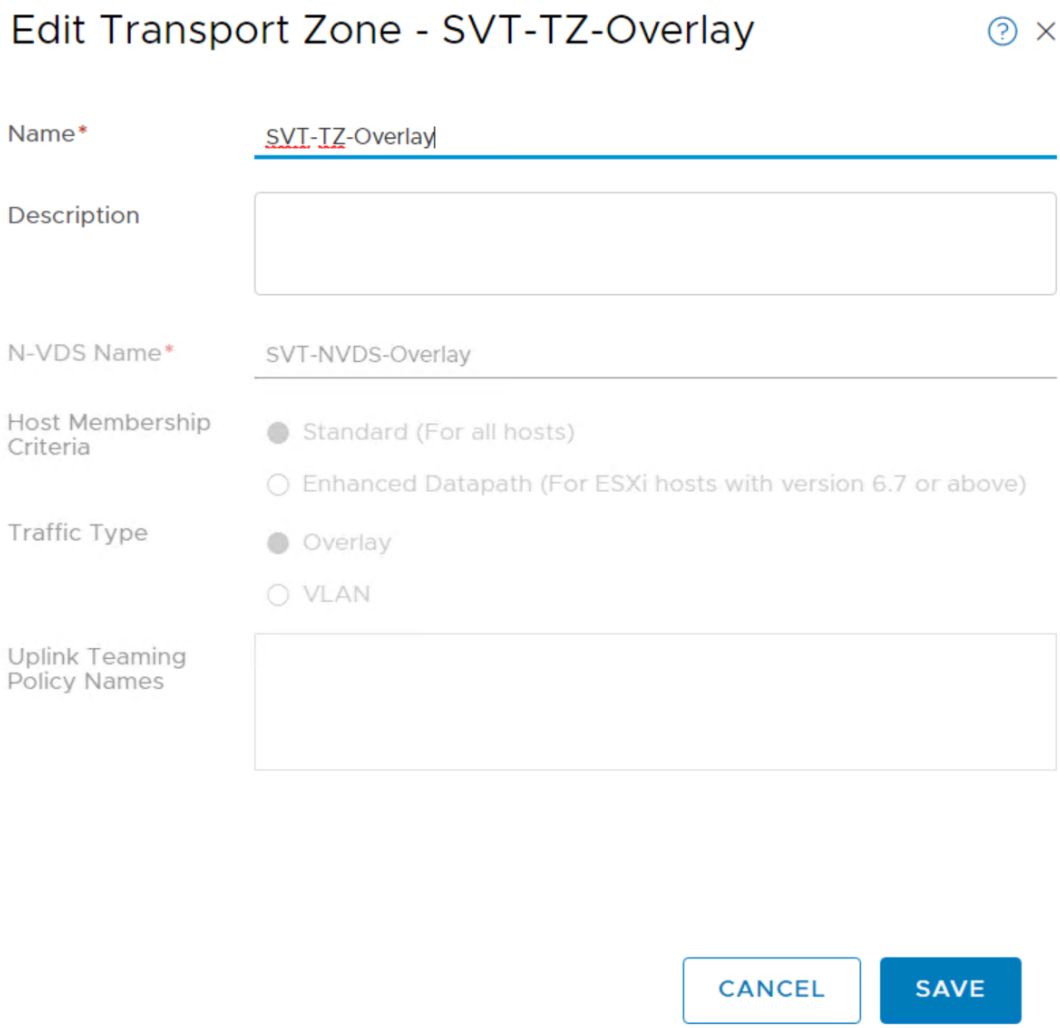

- Prepare Uplink **Profile** (select a physical interface, in this case its vmnic5), in here Uplink profile created with a name **Overlay-Uplink-v21** (vlan 21 for SVT-DC) and **Overlay-Uplink-v22** (vlan 22 for SVT-DR)

  System > Fabric > Profiles

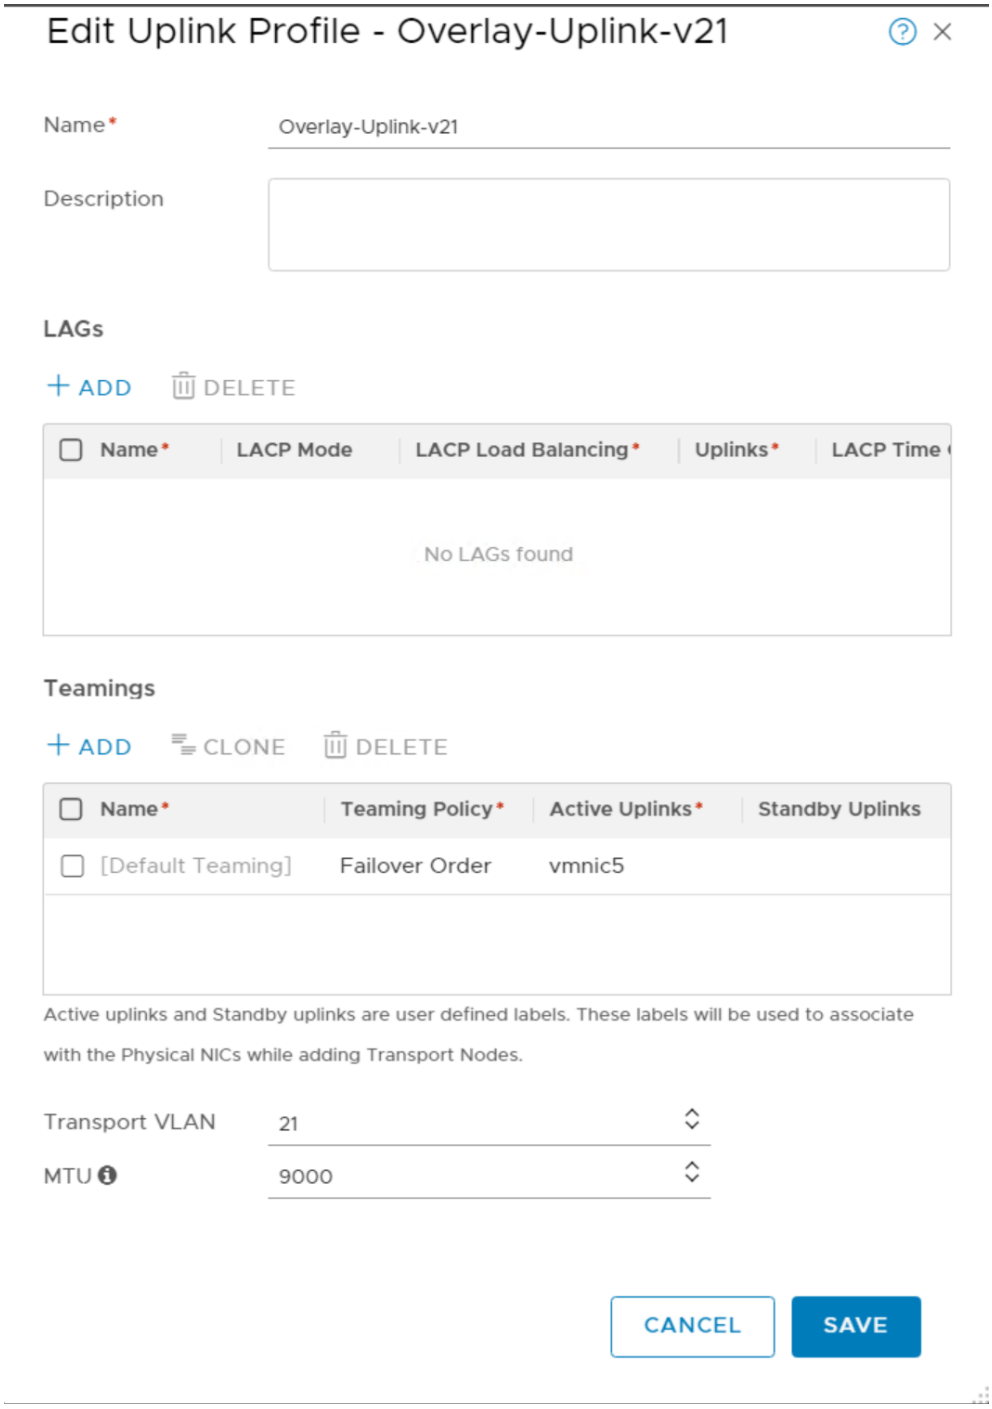

- Navigate **System> Nodes**, select vCenter under "**Managed By**" then Select specific ESXi Node/ESXi Cluster click on **Configure NSX**, associate TZ, Uplink-Profile and IP Pool
  - Create **IP Pool** for TEP as DC-TEP-POOL, in here 99.99.11.0/24 created as a pool for SVT-DC.
  - Create **IP Pool** for TEP as DR-TEP-POOL, in here 99.99.22.0/24 created as a pool for SVT-DR.

- Associate appropriate “Physical NICs” to each node (in here, vmnic5 used as overlay nic)

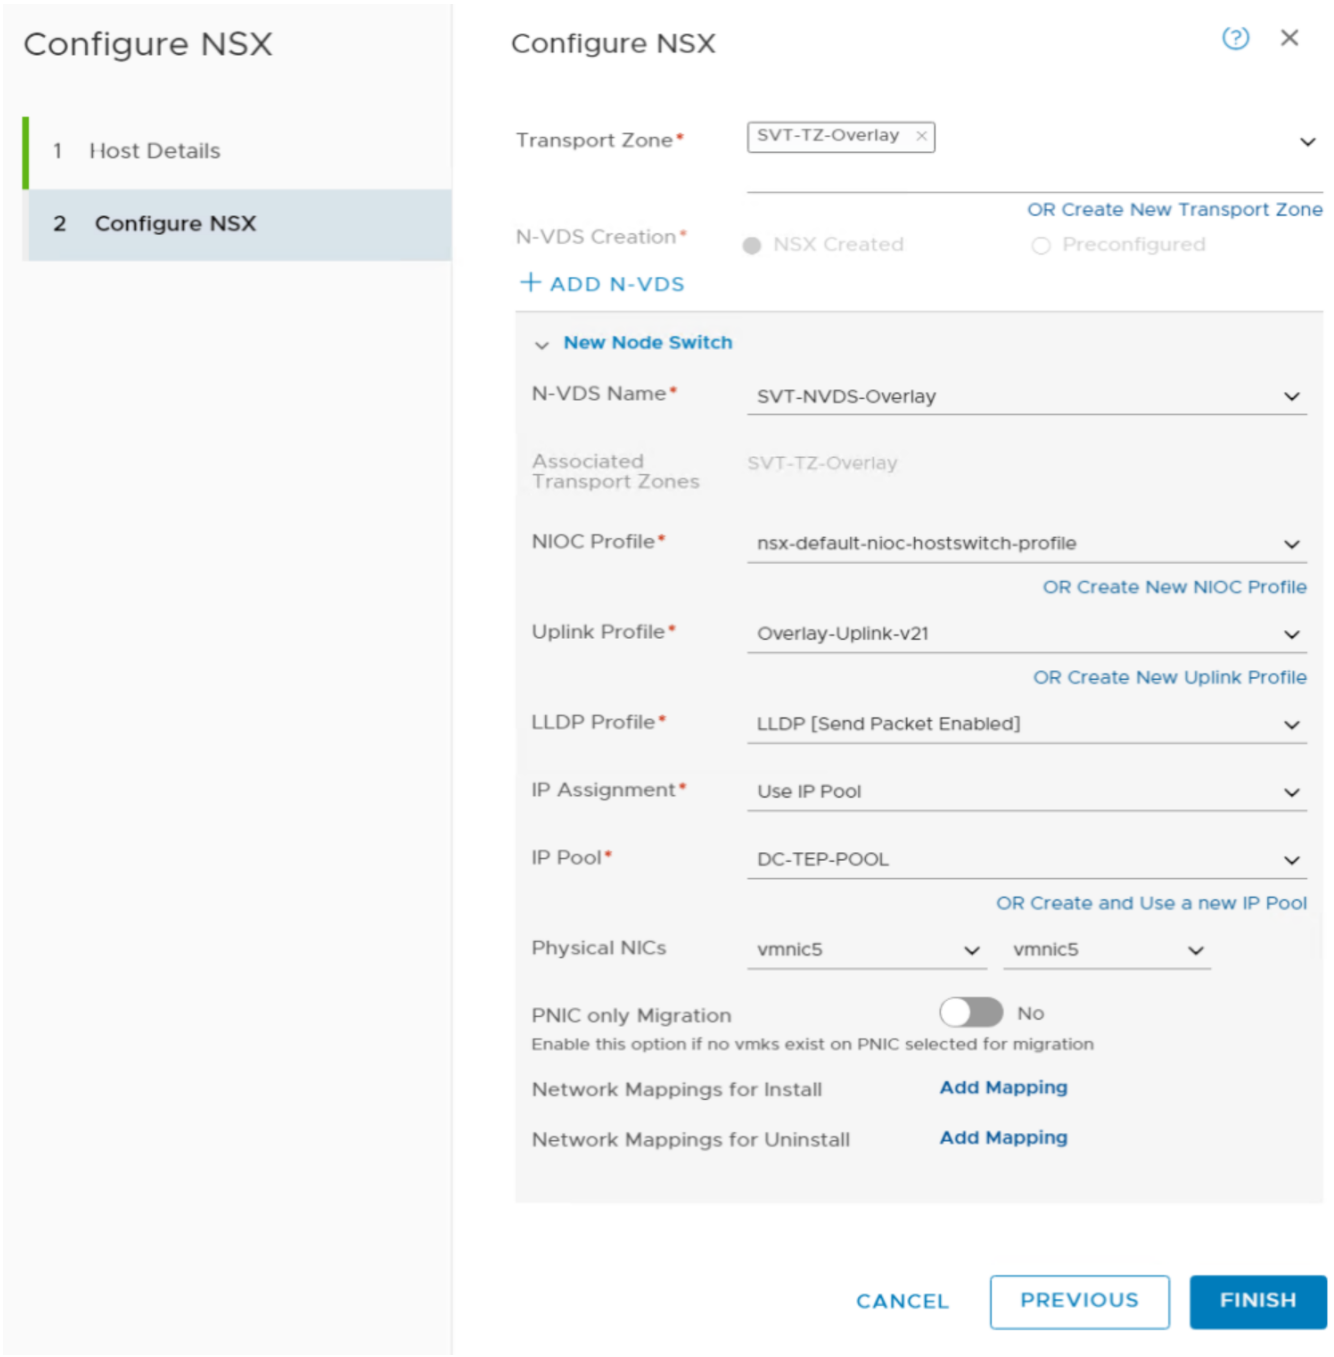

With the above steps, NSX-T will install NSX vibs on the ESXi hosts and then configure the GENEVE tunnels. Here is the

status after NSX was configured properly.

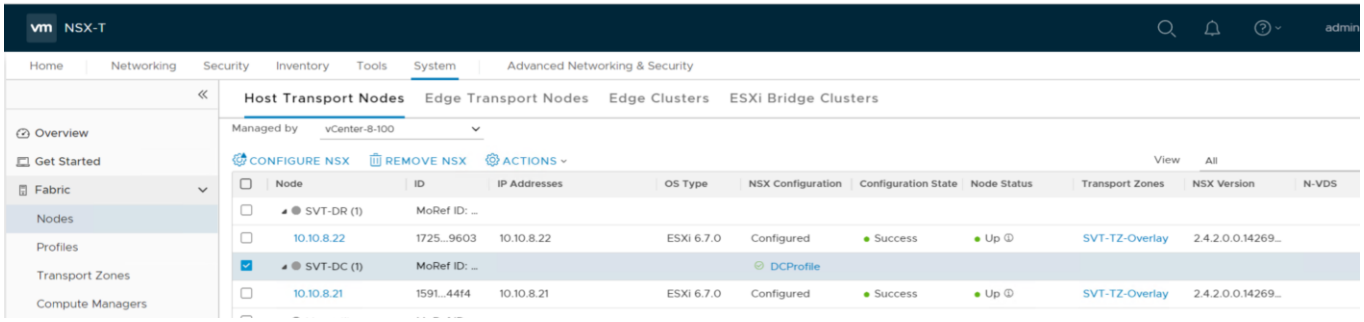

The below screenshot shows the GENEVE tunnel between each ESXi node to other node. Notice the IP address of both ESXi servers are in two different IP subnets (L3-apart)

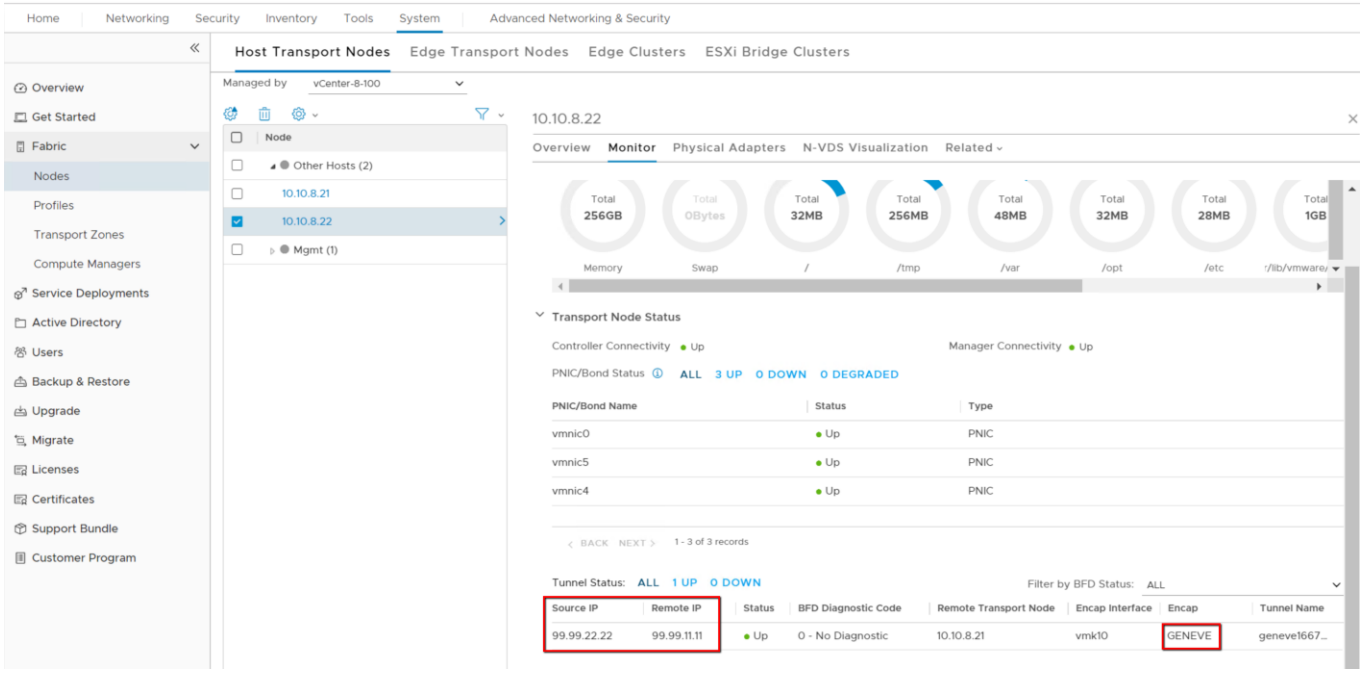

Aruba AOS-CX Switch configuration

```
8320-SW01 Configuration

interface 1/1/1
   no shutdown
   mtu 9000
   description To ESXi-1 VMNIC5
   no routing
   vlan trunk native 1
   vlan trunk allowed all
   exit
```

```
interface 1/1/46
   no shutdown
   mtu 9000
   description SW1-to-SW2
   no routing
   vlan trunk native 1
   vlan trunk allowed all
   exit

interface vlan11
    ip address 1.1.1.2/30
    ip mtu 9000
    exit

interface vlan21
    ip address 99.99.11.254/24
    ip mtu 9000
    exit

ip route 99.99.22.0/24 1.1.1.1
```

8320-SW02 Configuration

```
 interface 1/1/1
   no shutdown
   mtu 9000
   description To ESXi-2 VMNIC5
   no routing
   vlan trunk native 1
   vlan trunk allowed all
   exit

interface 1/1/46
   no shutdown
   mtu 9000
   no routing
   description SW2-to-SW1
   vlan trunk native 1
   vlan trunk allowed all

interface vlan11
    ip address 1.1.1.1/30
    ip mtu 9000
    exit

interface vlan22
    ip address 99.99.22.254/24
    ip mtu 9000
    exit

ip route 99.99.11.0/24 1.1.1.2
```

To validate the VM communication, follow the same steps from Usecase-1 for Connect VM's using NSX-T Logical Switch and Inter Logical Switch Routing

Usecase-3: BGP between NSX-T Edge VM & Aruba AOS-CX Switch

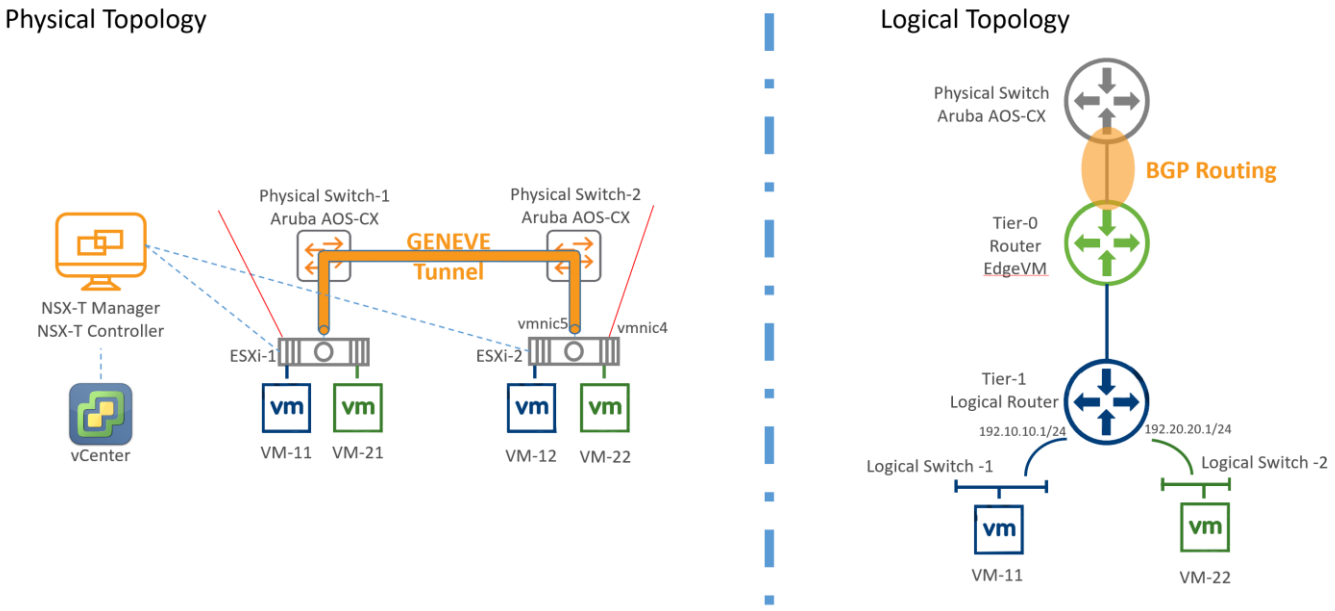

Figure 6: BGP Topology

Use case 3 details those environments that need to provide connectivity from VMs and Hosts within the overlay to targets that do not exists in the VMware NSX-T environment. In these cases (such as a VM needing to access the internet), NSX-T needs to create a routed environment between the VMware Hypervisor environment and the AOS-CX switch.

NSX-T Tier-1 routers facilitate multi-tenancy in the NSX platform. Each Tenant has their own T1 router which connects to a Tier-0 router for northbound access outside of the NSX-T environment. The link between T0 and T1 uses a reserved address space (100.64.0.0/16) and it assigns a /31 subnet on the T0-T1 link.

ESXi Host Preparation for EDGE VM

Create a DVS switch as shown below on the host where we are going to host EDGE VM's and port-groups for Transport overlay. Then configure uplinks to communicate with the AOS-CX Switch.

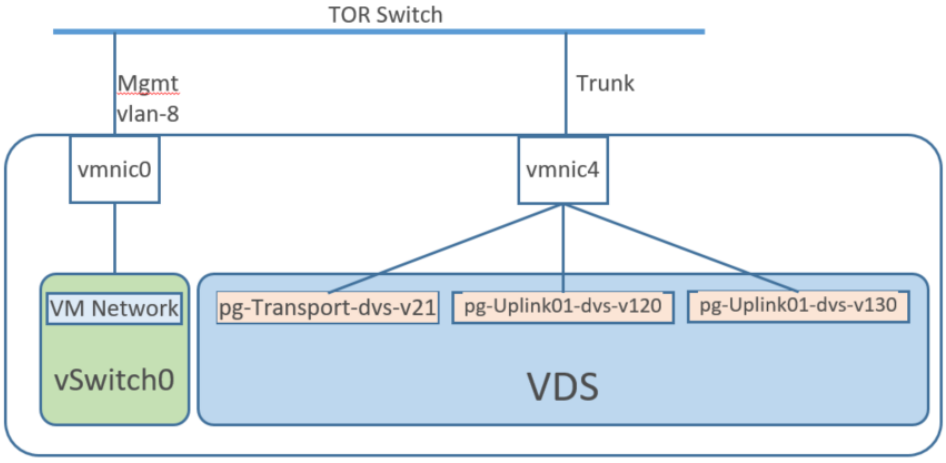

Figure 7: DVS Port-groups for EDGE VM

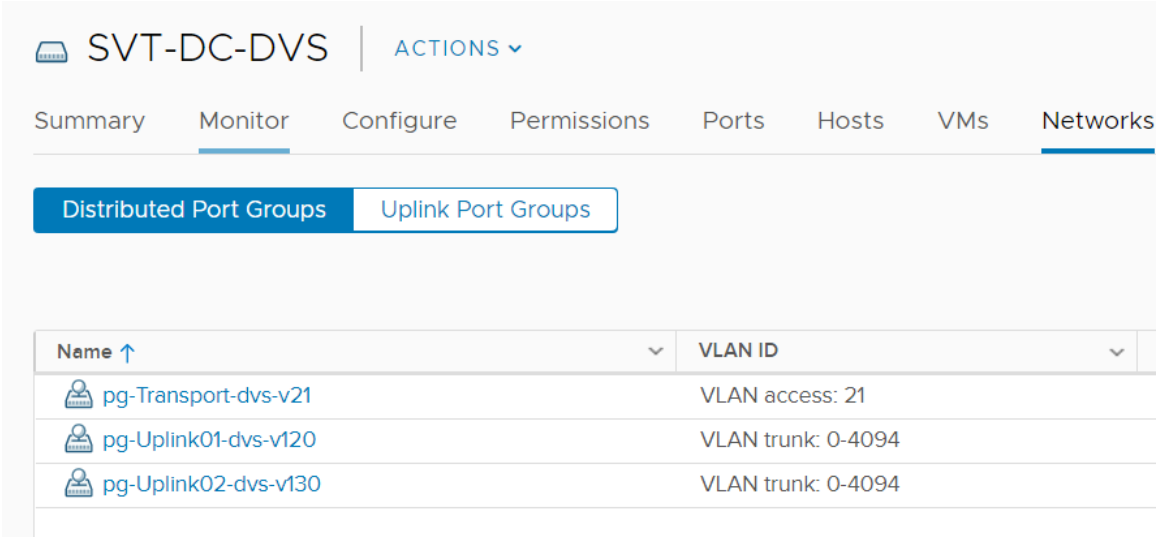

NSX-T Preparation

Create two Uplink logical switches for Edge VM in NSX-T, which enables connectivity with AOS-CX Switches

Navigate to **Advanced Networking & Security > Add Logical Switch**

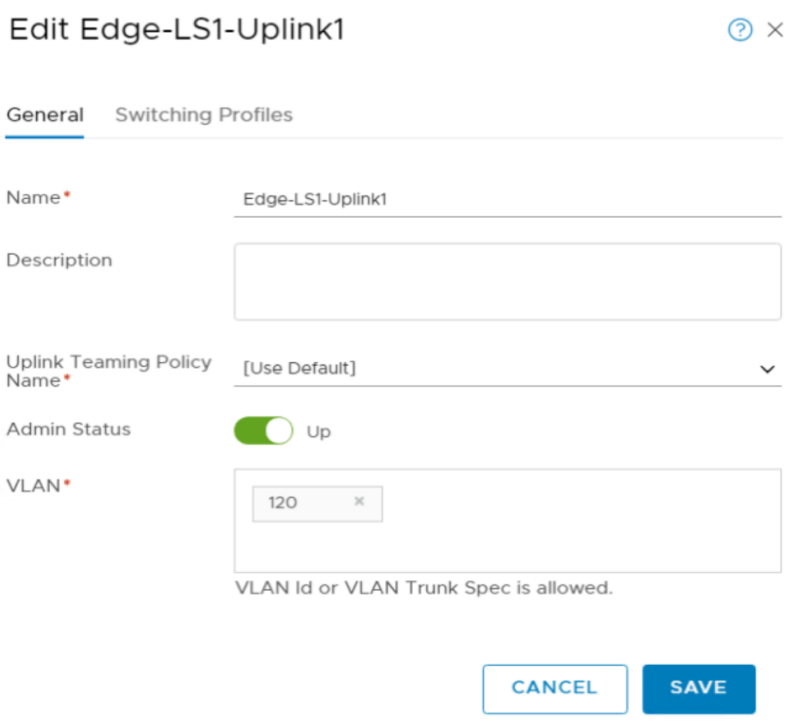

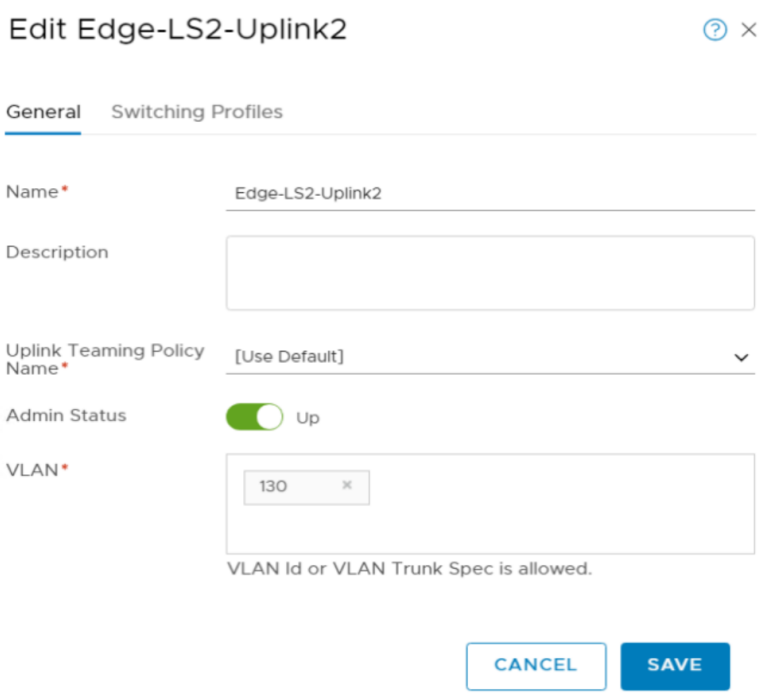

Edge Transport Nodes

Add an NSX Edge VM which helps to enable connectivity between the overlay networks and the physical network.

Here is the IP connectivity between the EDGE VM's and the AOS-CX Switches

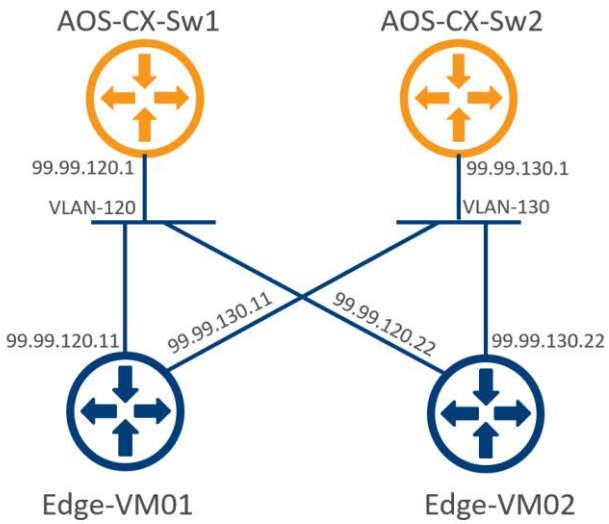

Figure 8: IP connectivity between EDGE VM's – AOS-CX Swithces

Navigate to **System > Fabric > Nodes > Edge Transport Nodes > Add Edge VM**

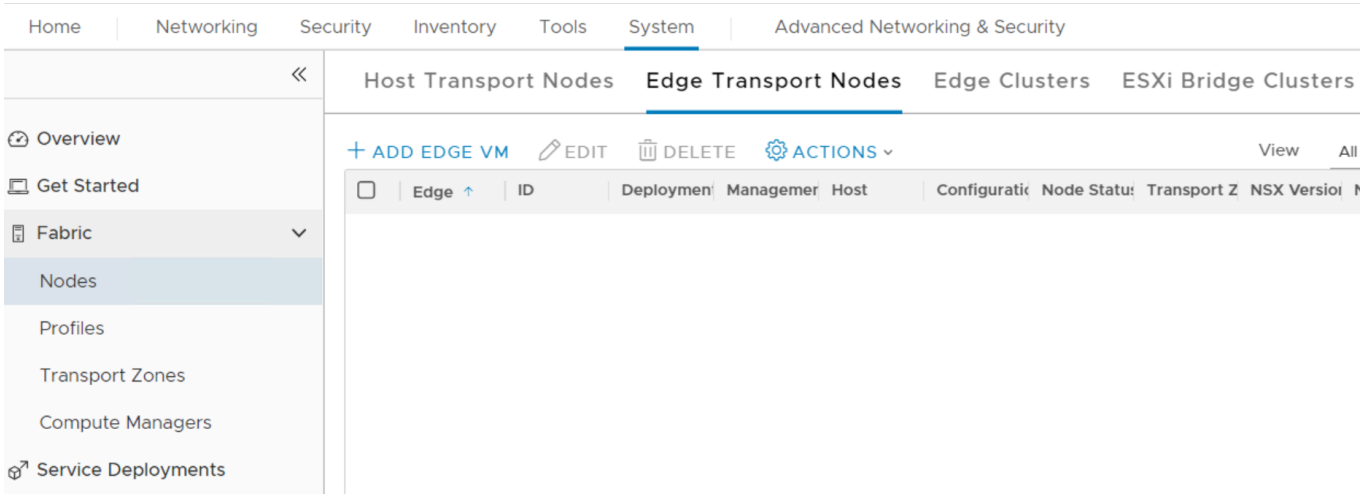

Click on "+ ADD EDGE VM"

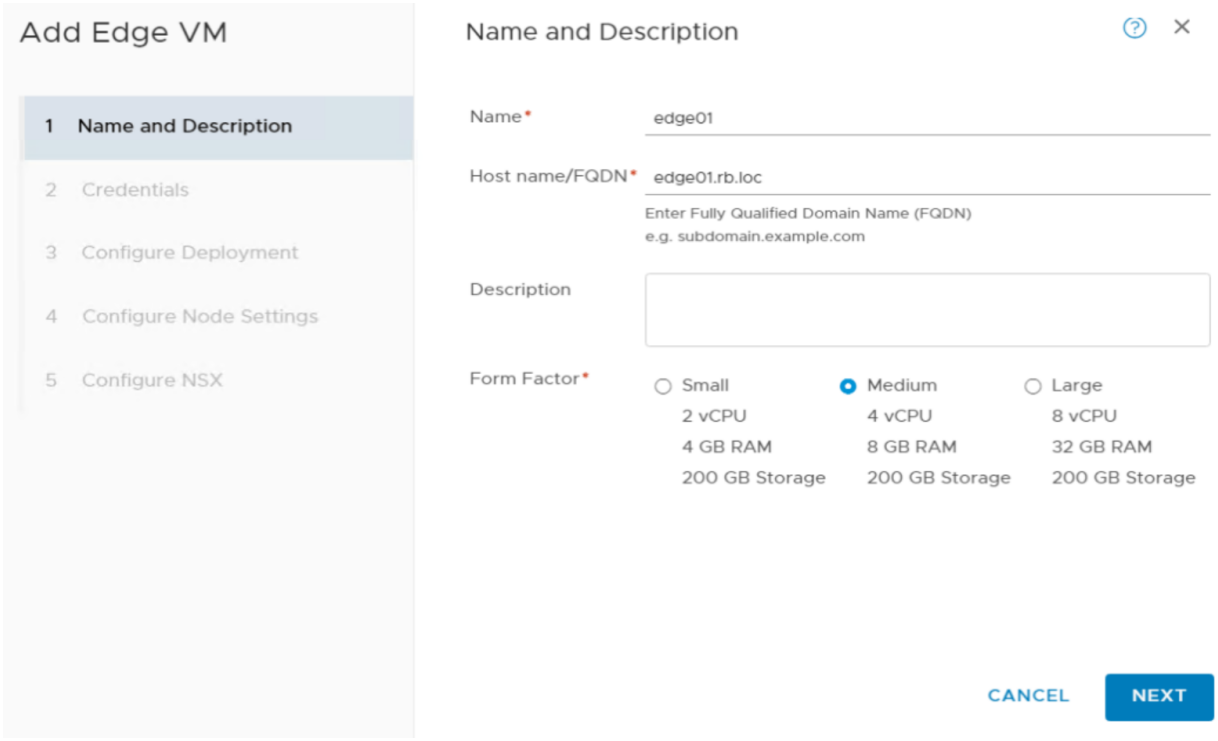
Add Edge VM
1 Name and Description
2 Credentials
3 Configure Deployment
4 Configure Node Settings
5 Configure NSX
Name and Description
Name*
edge01
Host name/FQDN*
edge01.rb.loc
Enter Fully Qualified Domain Name (FQDN)
e.g. subdomain.example.com
Description
Form Factor*
Small
2 vCPU
4 GB RAM
200 GB Storage
Medium
4 vCPU
8 GB RAM
200 GB Storage
Large
8 vCPU
32 GB RAM
200 GB Storage
CANCEL
NEXT

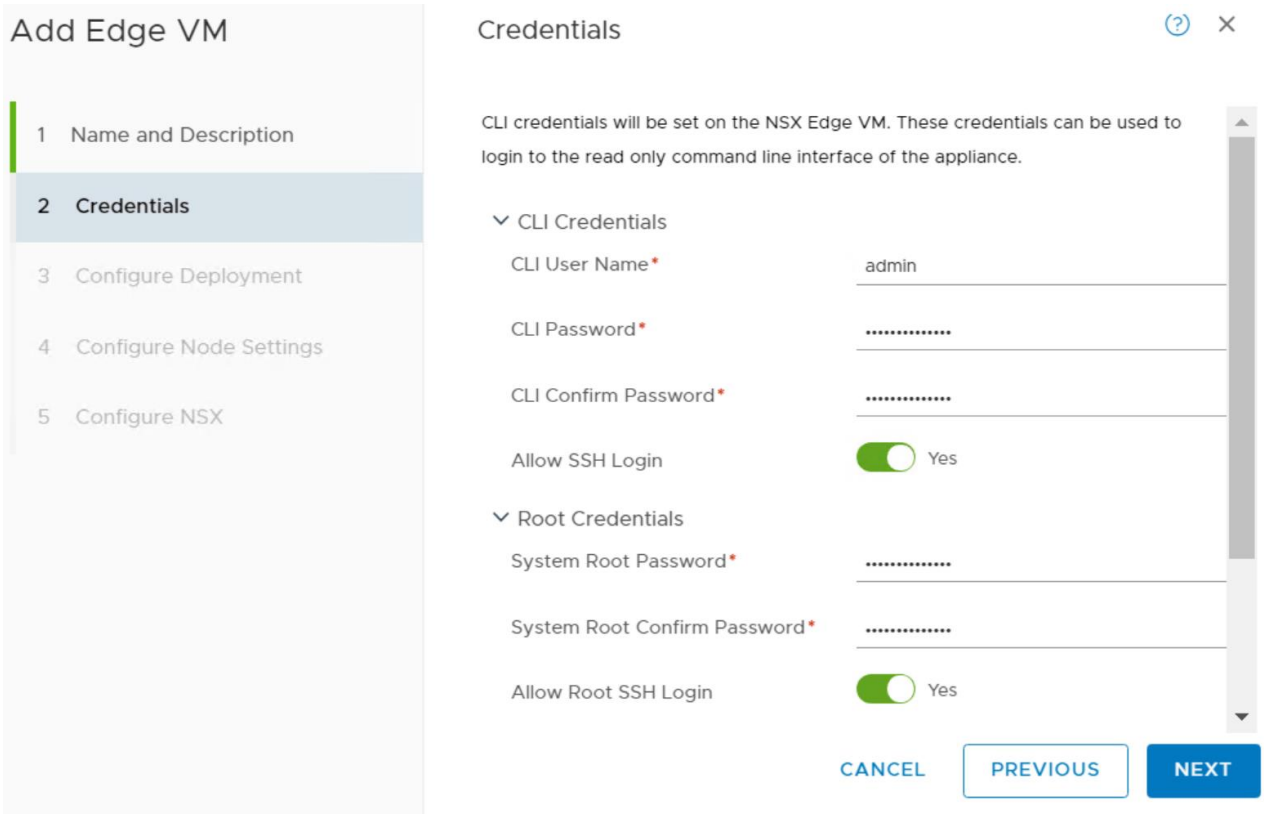
Add Edge VM
1 Name and Description
2 Credentials
3 Configure Deployment
4 Configure Node Settings
5 Configure NSX
Credentials
CLI credentials will be set on the NSX Edge VM. These credentials can be used to login to the read only command line interface of the appliance.
CLI Credentials
CLI User Name*
admin
CLI Password*
CLI Confirm Password*
Allow SSH Login
Yes
Root Credentials
System Root Password*
System Root Confirm Password*
Allow Root SSH Login
Yes
CANCEL
PREVIOUS
NEXT

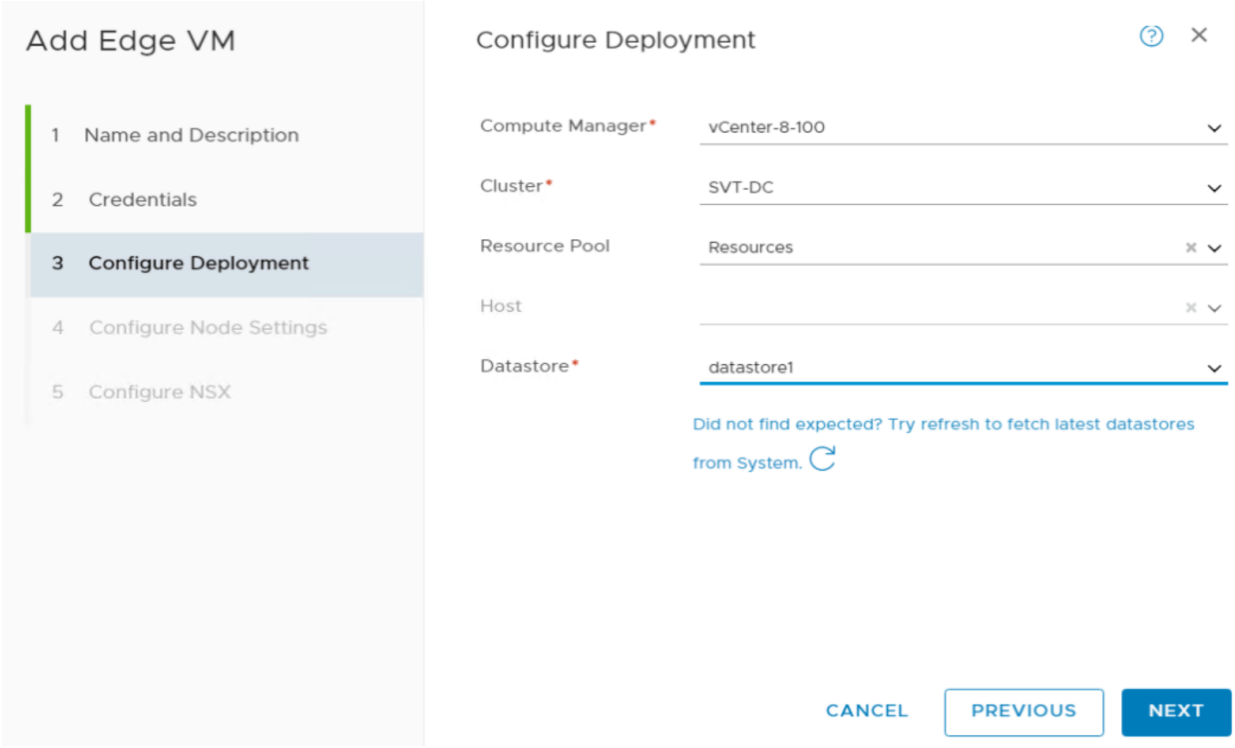
Add Edge VM
1 Name and Description
2 Credentials
3 Configure Deployment
4 Configure Node Settings
5 Configure NSX
Configure Deployment
Compute Manager*
vCenter-8-100
Cluster*
SVT-DC
Resource Pool
Resources
Host
Datastore*
datastore1
Did not find expected? Try refresh to fetch latest datastores from System.
CANCEL
PREVIOUS
NEXT

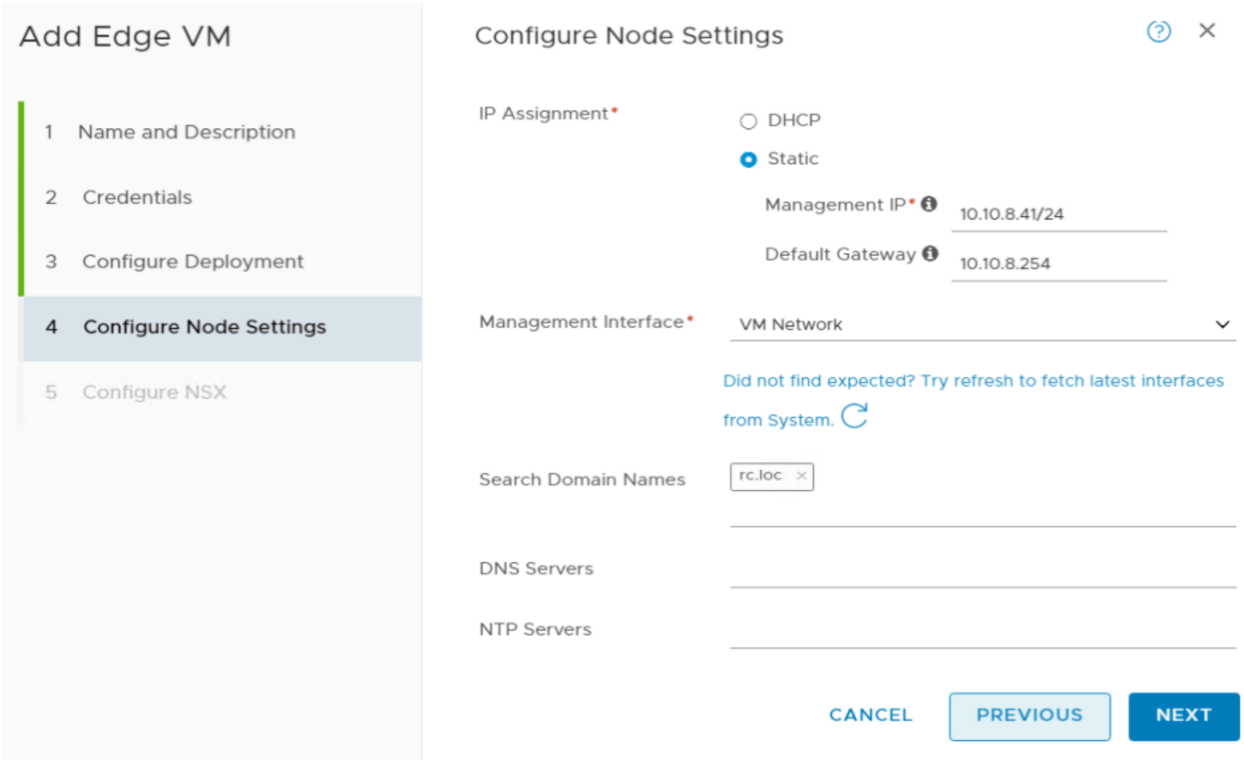
Add Edge VM
1 Name and Description
2 Credentials
3 Configure Deployment
4 Configure Node Settings
5 Configure NSX
Configure Node Settings
IP Assignment*
DHCP
Static
Management IP*
10.10.8.41/24
Default Gateway
10.10.8.254
Management Interface*
VM Network
Did not find expected? Try refresh to fetch latest interfaces from System.
Search Domain Names
rc.loc
DNS Servers
NTP Servers
CANCEL
PREVIOUS
NEXT

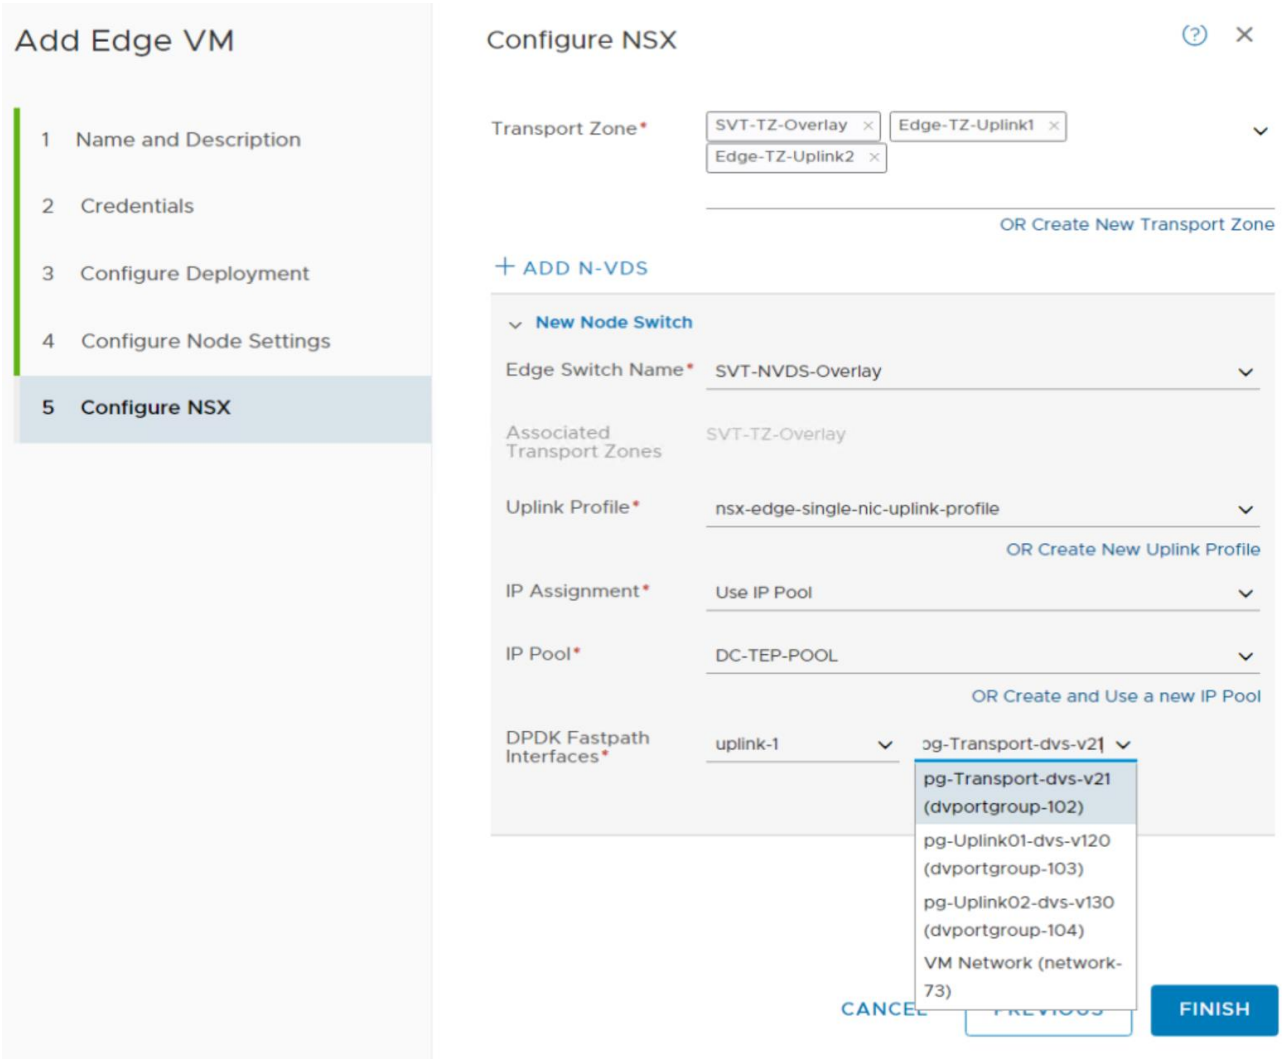
Add Edge VM
1 Name and Description
2 Credentials
3 Configure Deployment
4 Configure Node Settings
5 Configure NSX
Configure NSX
Transport Zone*
SVT-TZ-Overlay
Edge-TZ-Uplink1
Edge-TZ-Uplink2
OR Create New Transport Zone
+ ADD N-VDS
New Node Switch
Edge Switch Name*
SVT-NVDS-Overlay
Associated Transport Zones
SVT-TZ-Overlay
Uplink Profile*
nsx-edge-single-nic-uplink-profile
OR Create New Uplink Profile
IP Assignment*
Use IP Pool
IP Pool*
DC-TEP-POOL
OR Create and Use a new IP Pool
DPDK Fastpath Interfaces*
uplink-1
pg-Transport-dvs-v21
pg-Transport-dvs-v21 (dvportgroup-102)
pg-Uplink01-dvs-v120 (dvportgroup-103)
pg-Uplink02-dvs-v130 (dvportgroup-104)
VM Network (network-73)
CANCEL
FINISH

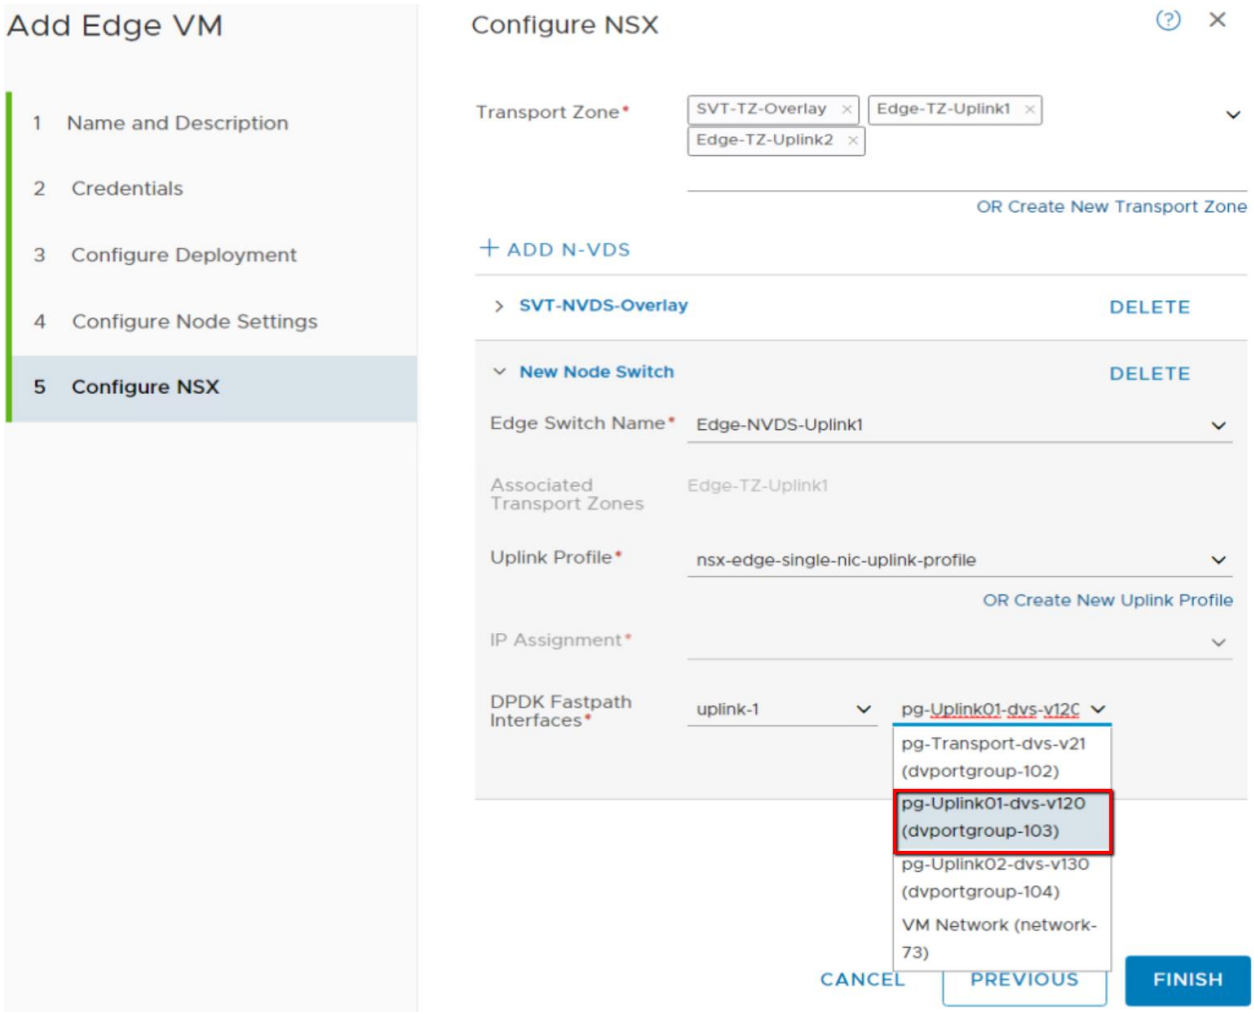
Add Edge VM
1 Name and Description
2 Credentials
3 Configure Deployment
4 Configure Node Settings
5 Configure NSX
Configure NSX
Transport Zone*
SVT-TZ-Overlay
Edge-TZ-Uplink1
Edge-TZ-Uplink2
OR Create New Transport Zone
+ ADD N-VDS
SVT-NVDS-Overlay
DELETE
New Node Switch
DELETE
Edge Switch Name*
Edge-NVDS-Uplink1
Associated Transport Zones
Edge-TZ-Uplink1
Uplink Profile*
nsx-edge-single-nic-uplink-profile
OR Create New Uplink Profile
IP Assignment*
DPDK Fastpath Interfaces*
uplink-1
pg-Uplink01-dvs-v120
pg-Transport-dvs-v21 (dvportgroup-102)
pg-Uplink01-dvs-v120 (dvportgroup-103)
pg-Uplink02-dvs-v130 (dvportgroup-104)
VM Network (network-73)
CANCEL
PREVIOUS
FINISH

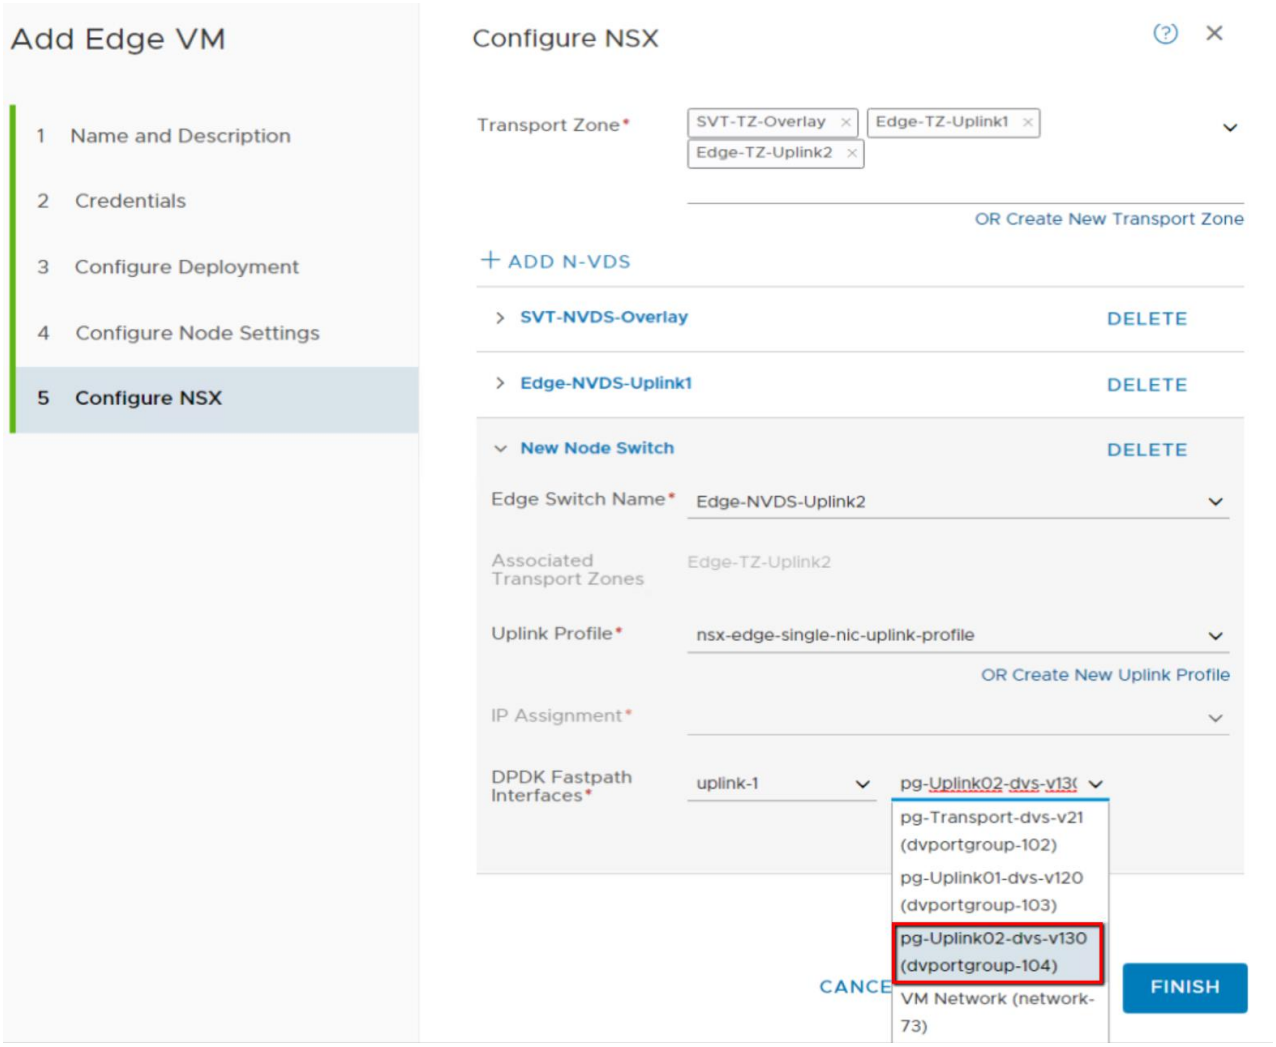

Create a second EDGE VM (edge02) in same way as above.
As these Edge VM's act like a WAN Edge for the fabric, it can be installed on a single ESXi or on a VMware Cluster for redundancy.

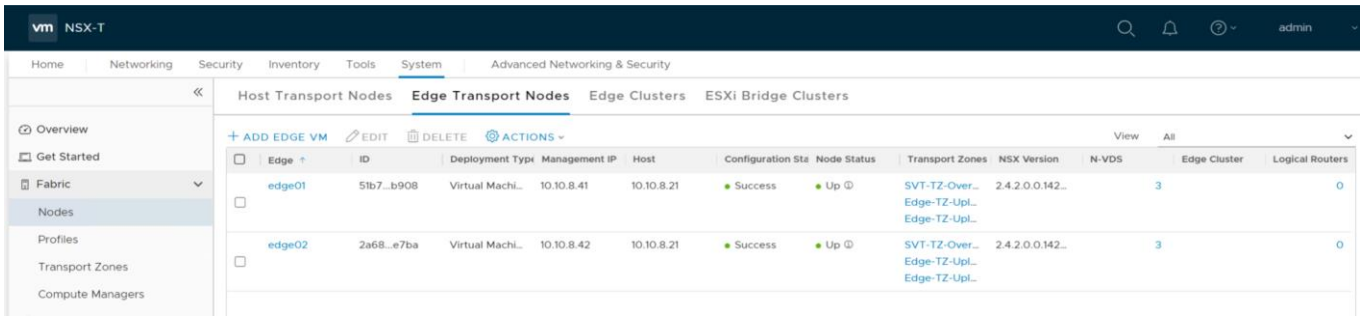

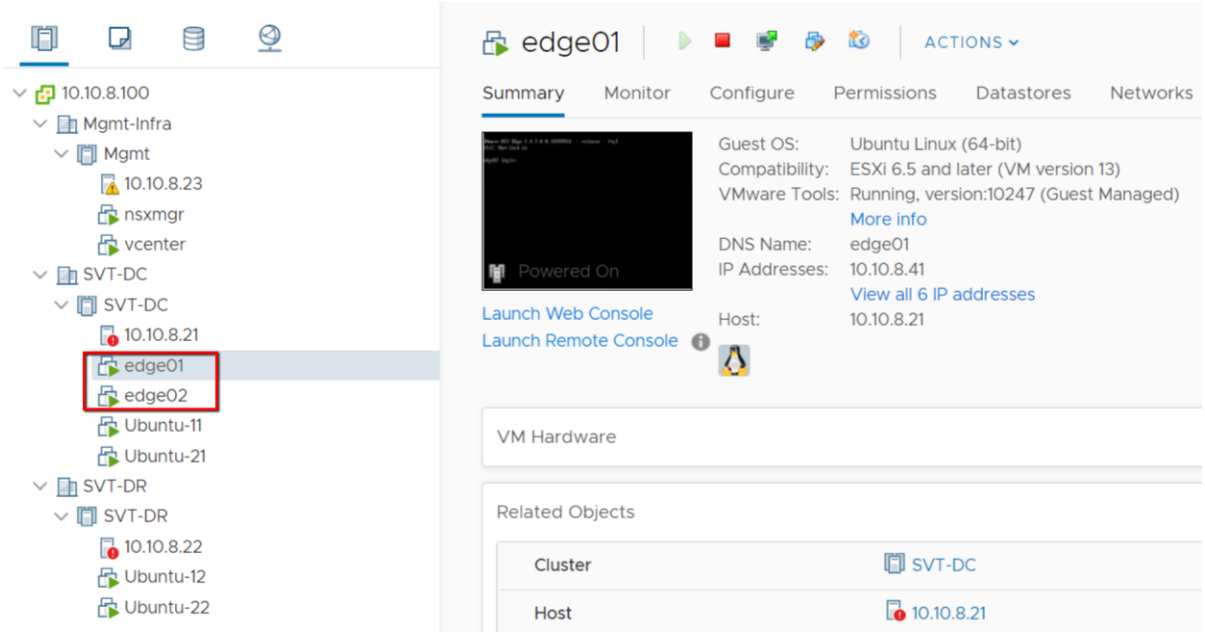

Create Edge Cluster

Group these two EDGE VM's in to an Edge Cluster.

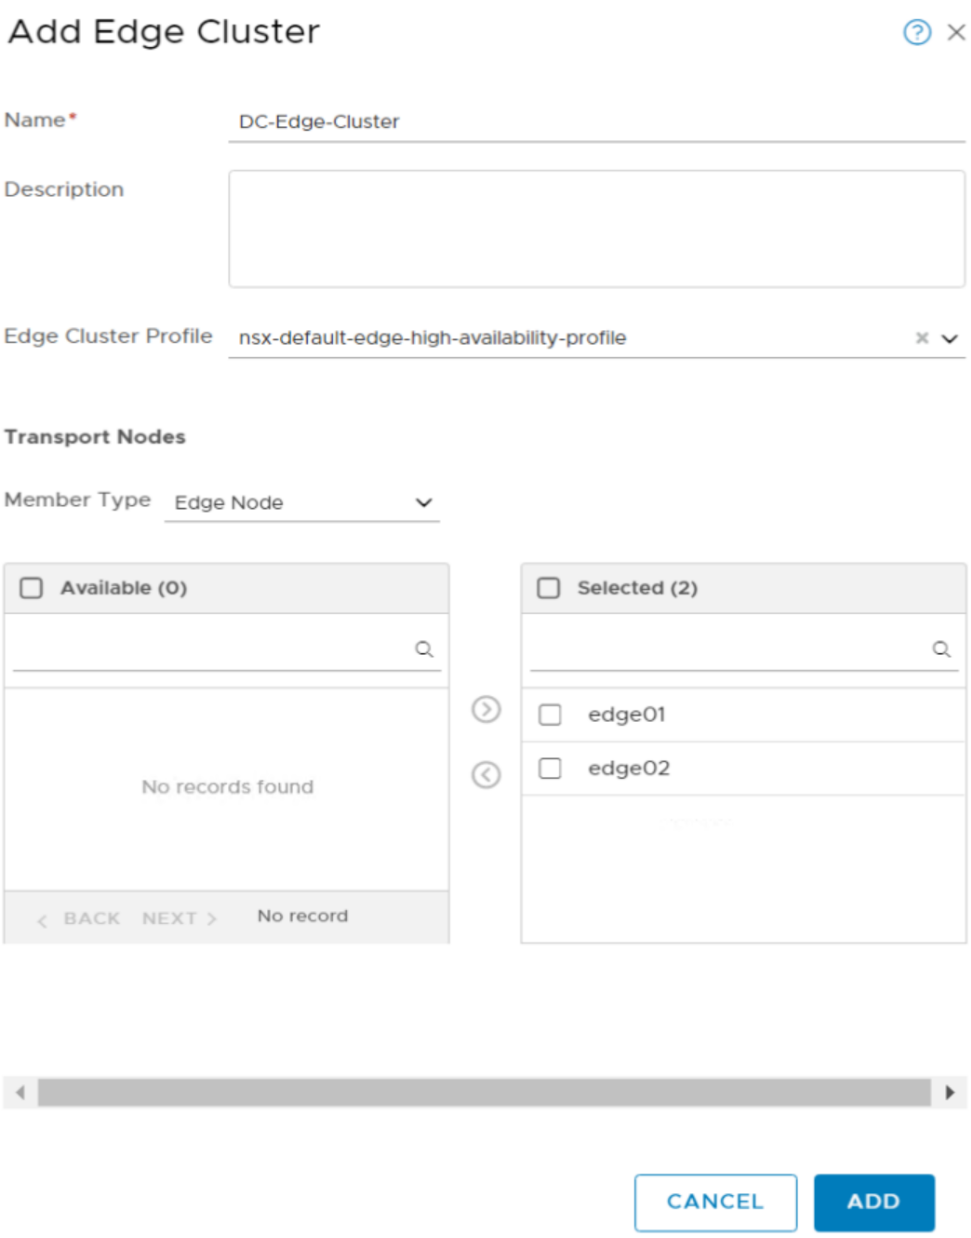

Now each edge VM's should be deployed, and the Management IP should be reachable.

Create Tier-0 Router

Create a Tier-0 Router and associate the new Edge Cluster with the Tier-0 Router.

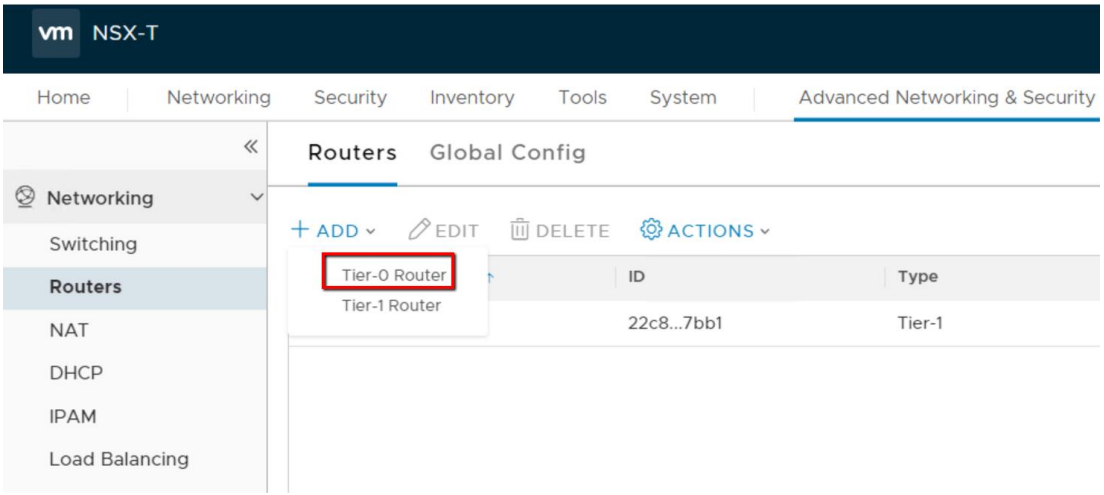

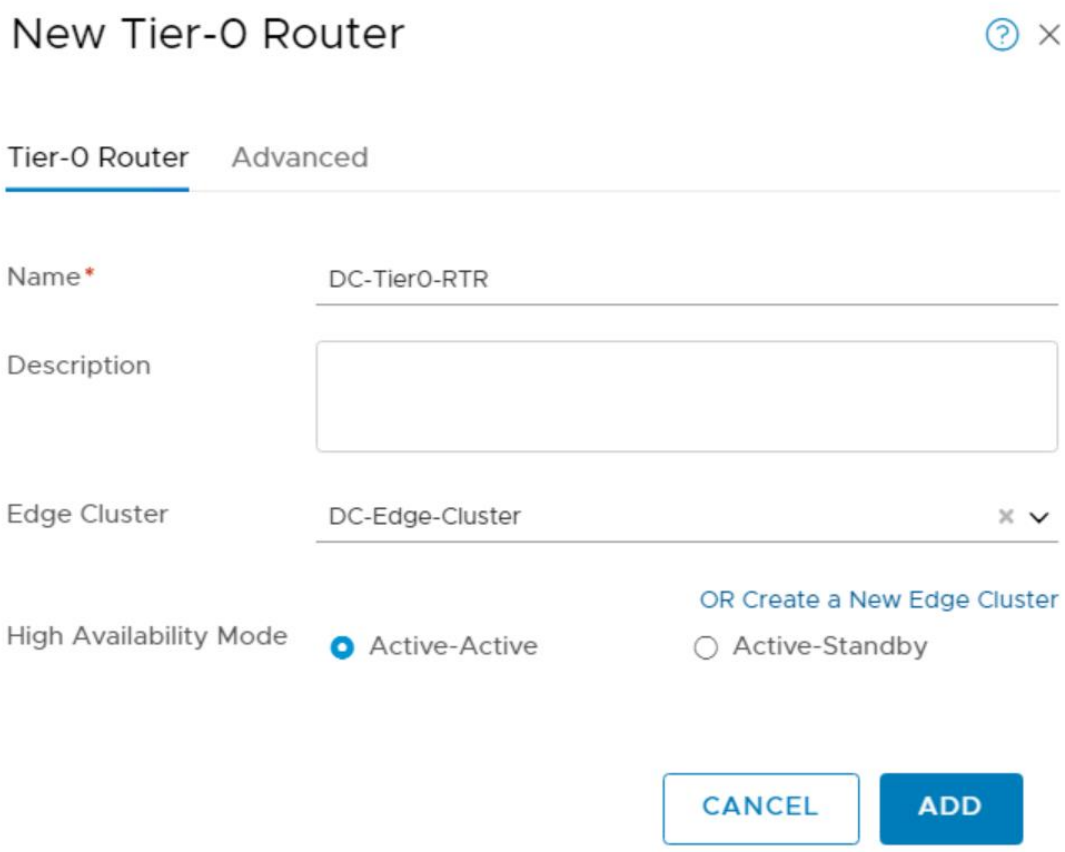

Adding router ports to the Tier-0 Router

Add router ports to the Tier-0 Router so it maps to the Mgmt., vlan120 and vlan130 port-groups using the logical switches

that are created at NSX-T Preparation section

Navigate to **Advanced Networking & Security > Networking > Routers >** Select **Tier0-RTR** (double click) **> Configuration >** Click on **Router ports > Add**

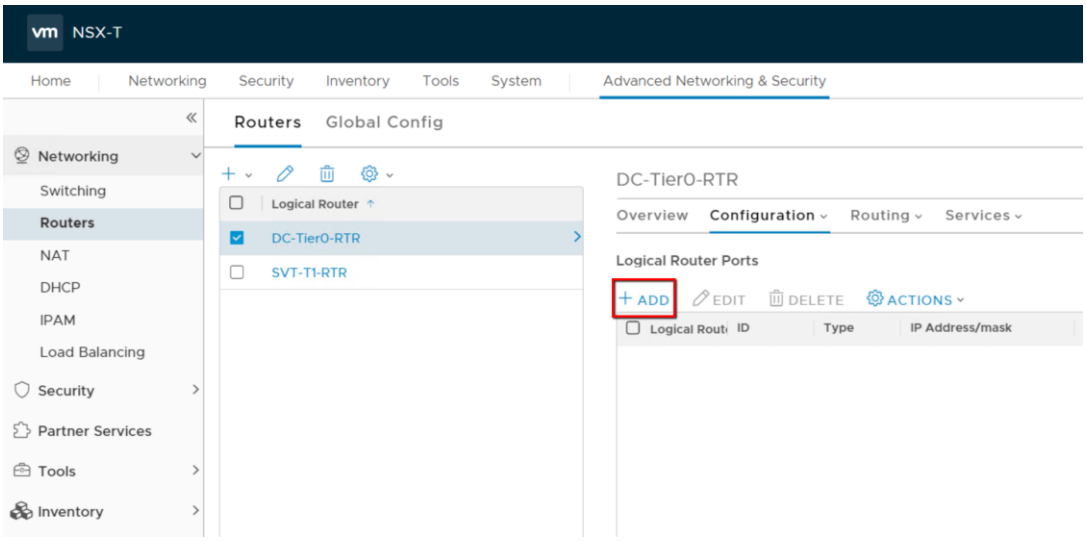

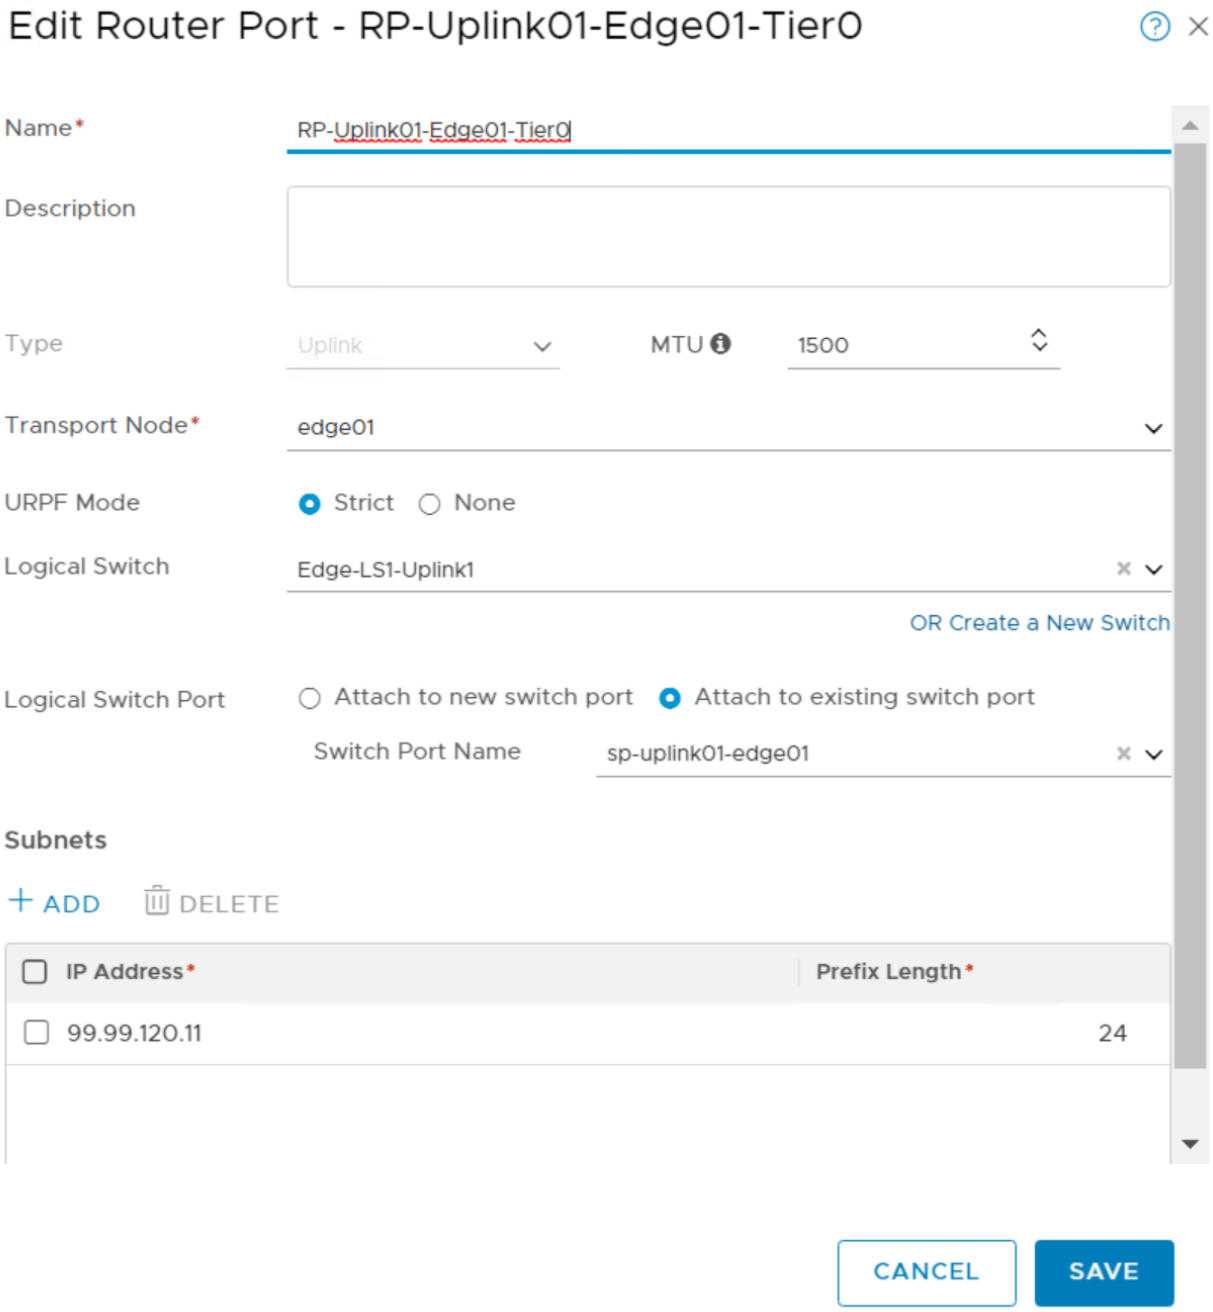
Edit Router Port - RP-Uplink01-Edge01-Tier0
Name*
RP-Uplink01-Edge01-Tier0
Description
Type
Uplink
MTU
1500
Transport Node*
edge01
URPF Mode
Strict
None
Logical Switch
Edge-LS1-Uplink1
OR Create a New Switch
Logical Switch Port
Attach to new switch port
Attach to existing switch port
Switch Port Name
sp-uplink01-edge01
Subnets
ADD
DELETE
IP Address*
Prefix Length*
99.99.120.11
24
CANCEL
SAVE

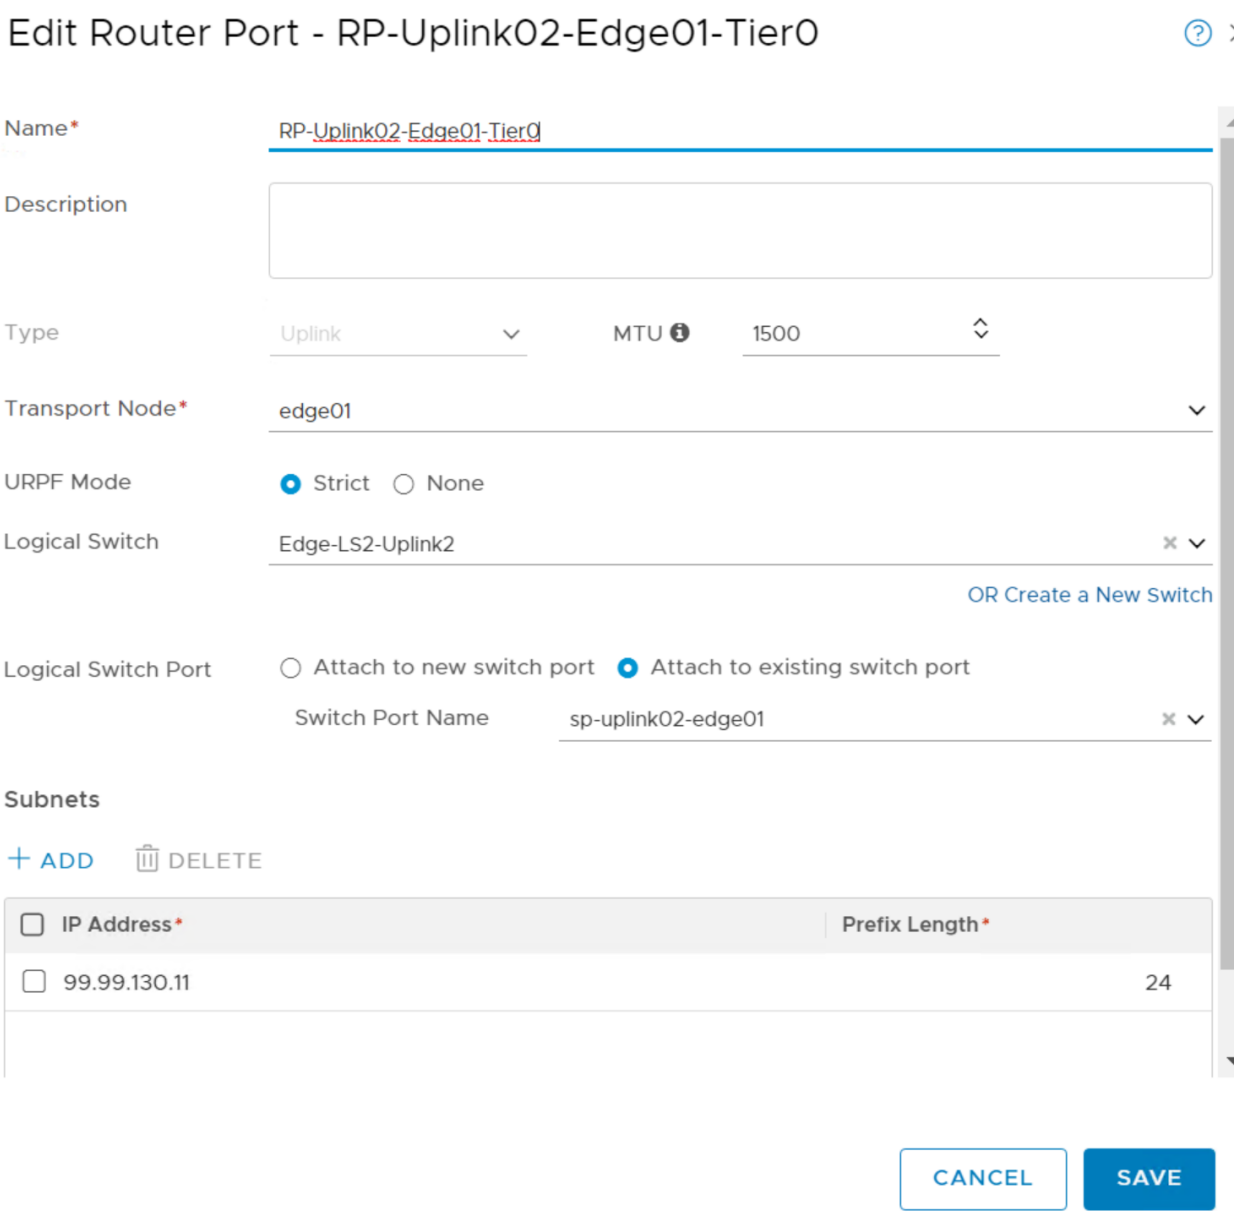
Edit Router Port - RP-Uplink02-Edge01-Tier0
Name*
RP-Uplink02-Edge01-Tier0
Description
Type
Uplink
MTU
1500
Transport Node*
edge01
URPF Mode
Strict
None
Logical Switch
Edge-LS2-Uplink2
OR Create a New Switch
Logical Switch Port
Attach to new switch port
Attach to existing switch port
Switch Port Name
sp-uplink02-edge01
Subnets
ADD
DELETE
IP Address*
Prefix Length*
99.99.130.11
24
CANCEL
SAVE

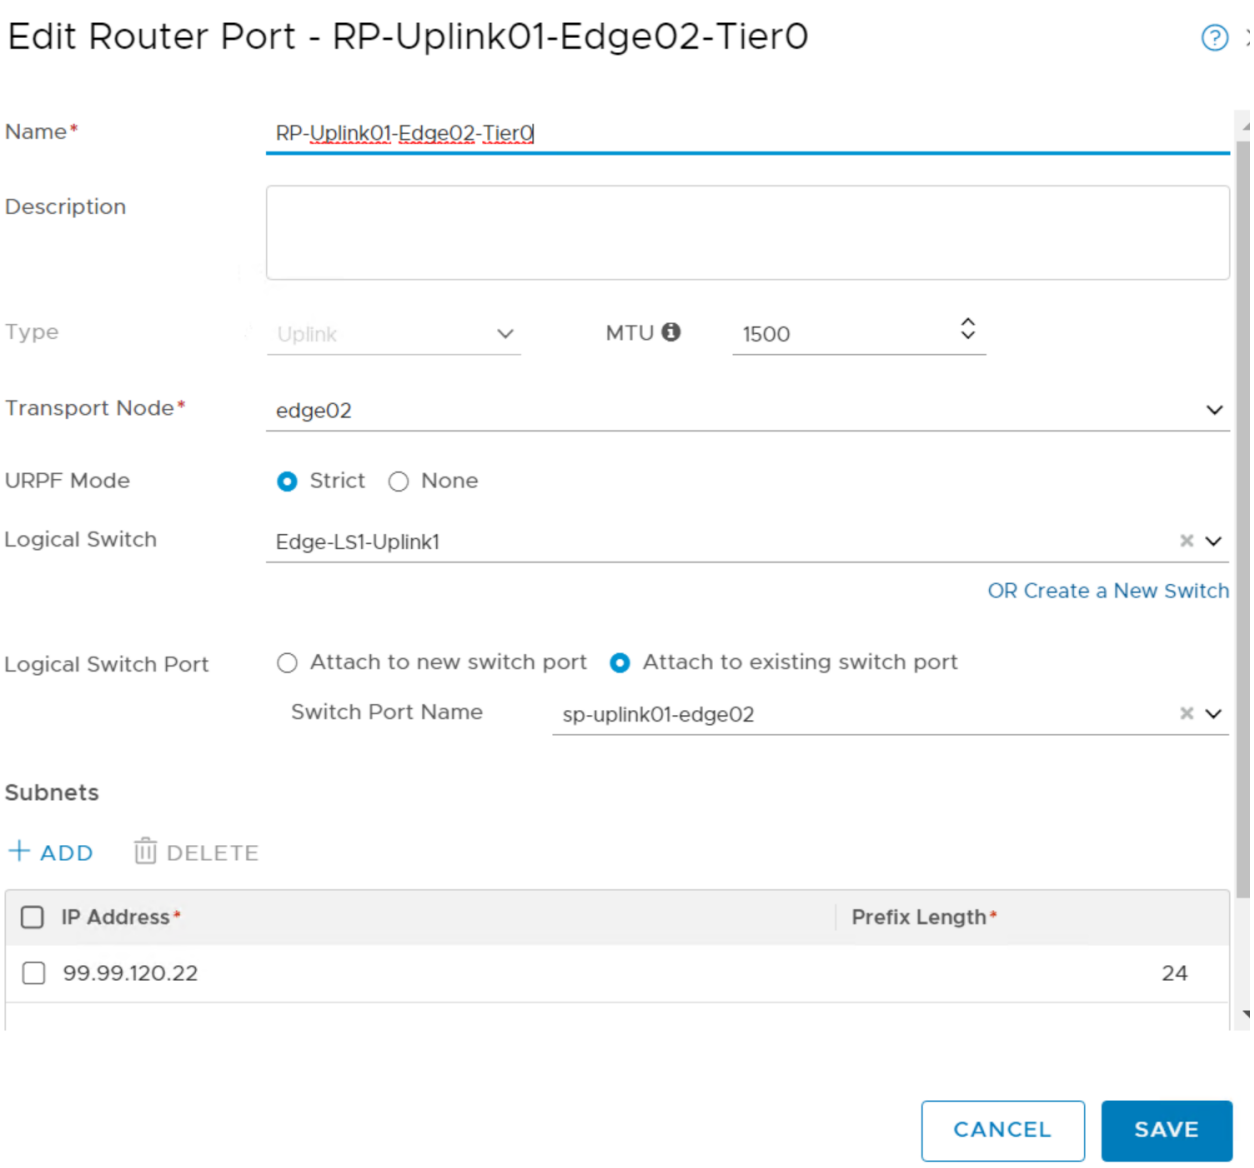
Edit Router Port - RP-Uplink01-Edge02-Tier0
Name*
RP-Uplink01-Edge02-Tier0
Description
Type
Uplink
MTU
1500
Transport Node*
edge02
URPF Mode
Strict
None
Logical Switch
Edge-LS1-Uplink1
OR Create a New Switch
Logical Switch Port
Attach to new switch port
Attach to existing switch port
Switch Port Name
sp-uplink01-edge02
Subnets
ADD
DELETE
IP Address*
Prefix Length*
99.99.120.22
24
CANCEL
SAVE

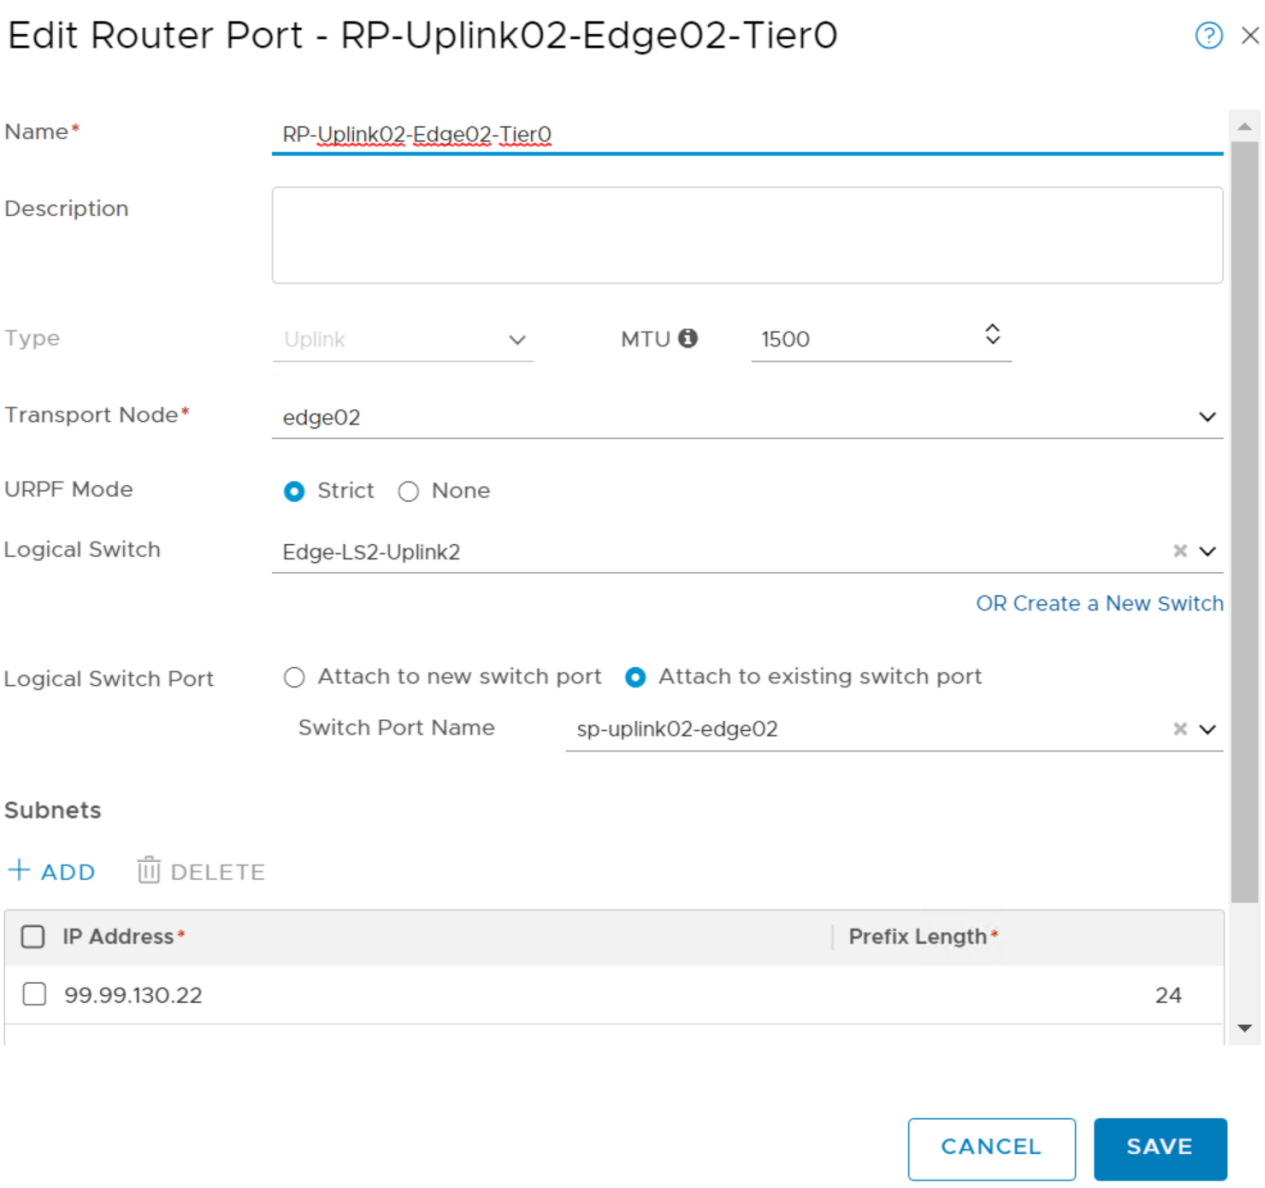

Here is the summary where we configured two interfaces with IP Address on each Edge VM

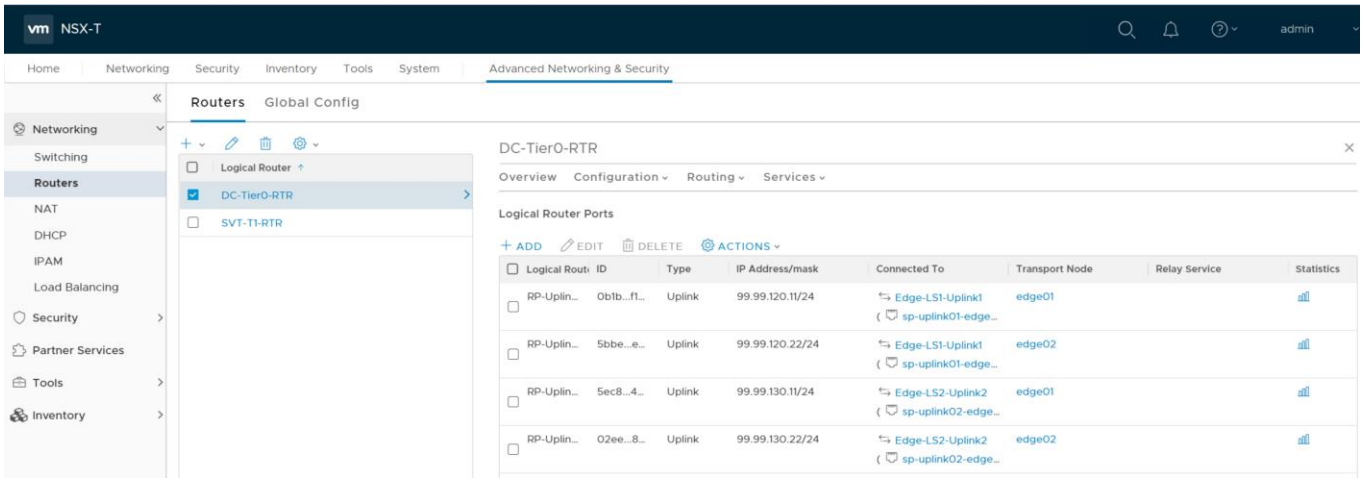

Configuring BGP on Tier-0 Router

Configure BGP to enable peering with attached AOS-CX switches.

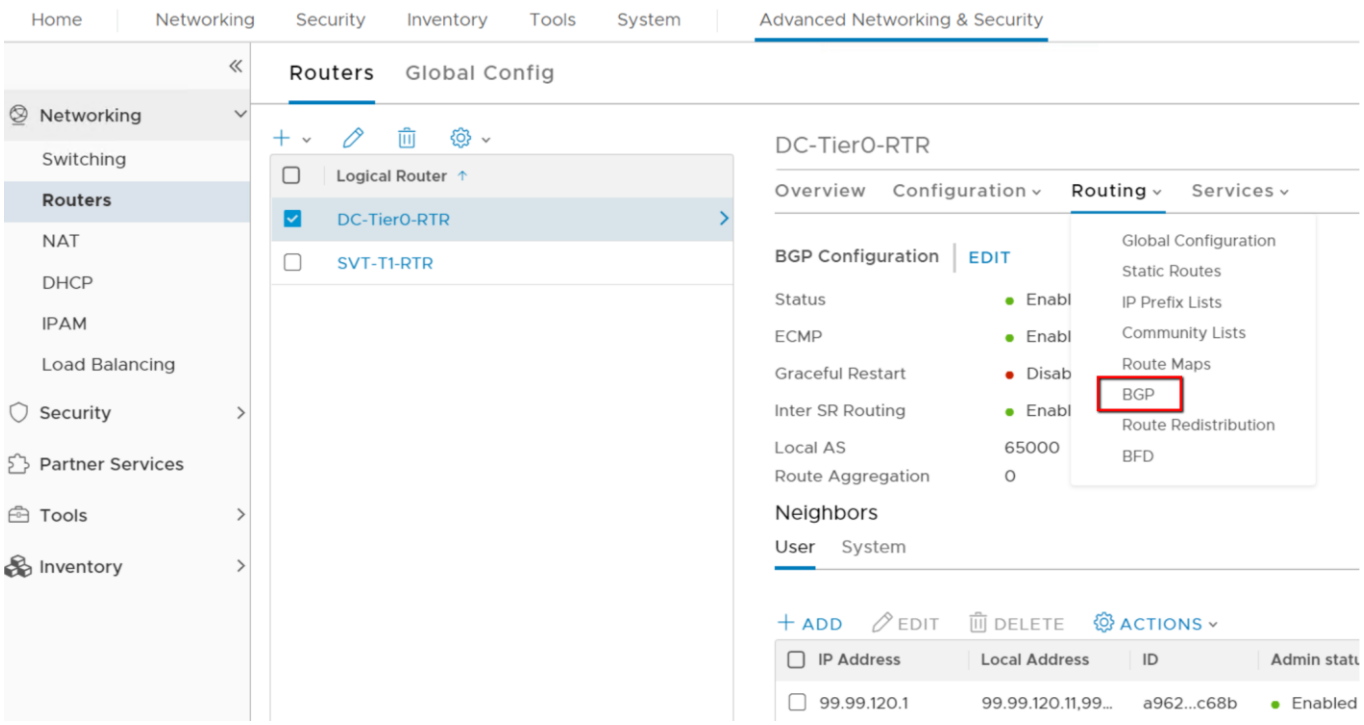

Add BGP peer

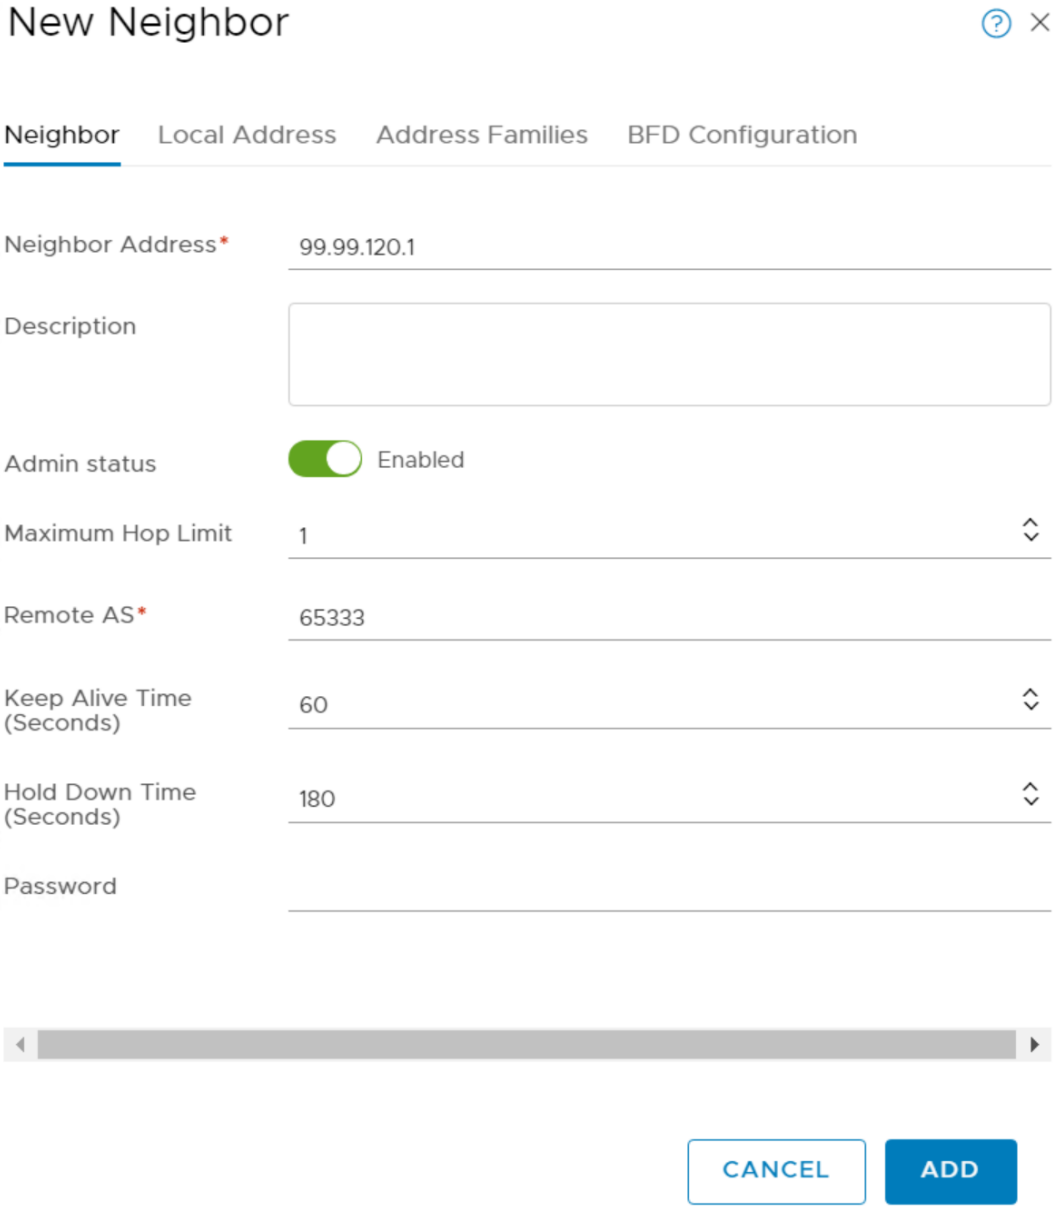
New Neighbor
Neighbor
Local Address
Address Families
BFD Configuration
Neighbor Address*
99.99.120.1
Description
Admin status
Enabled
Maximum Hop Limit
1
Remote AS*
65333
Keep Alive Time (Seconds)
60
Hold Down Time (Seconds)
180
Password
CANCEL
ADD

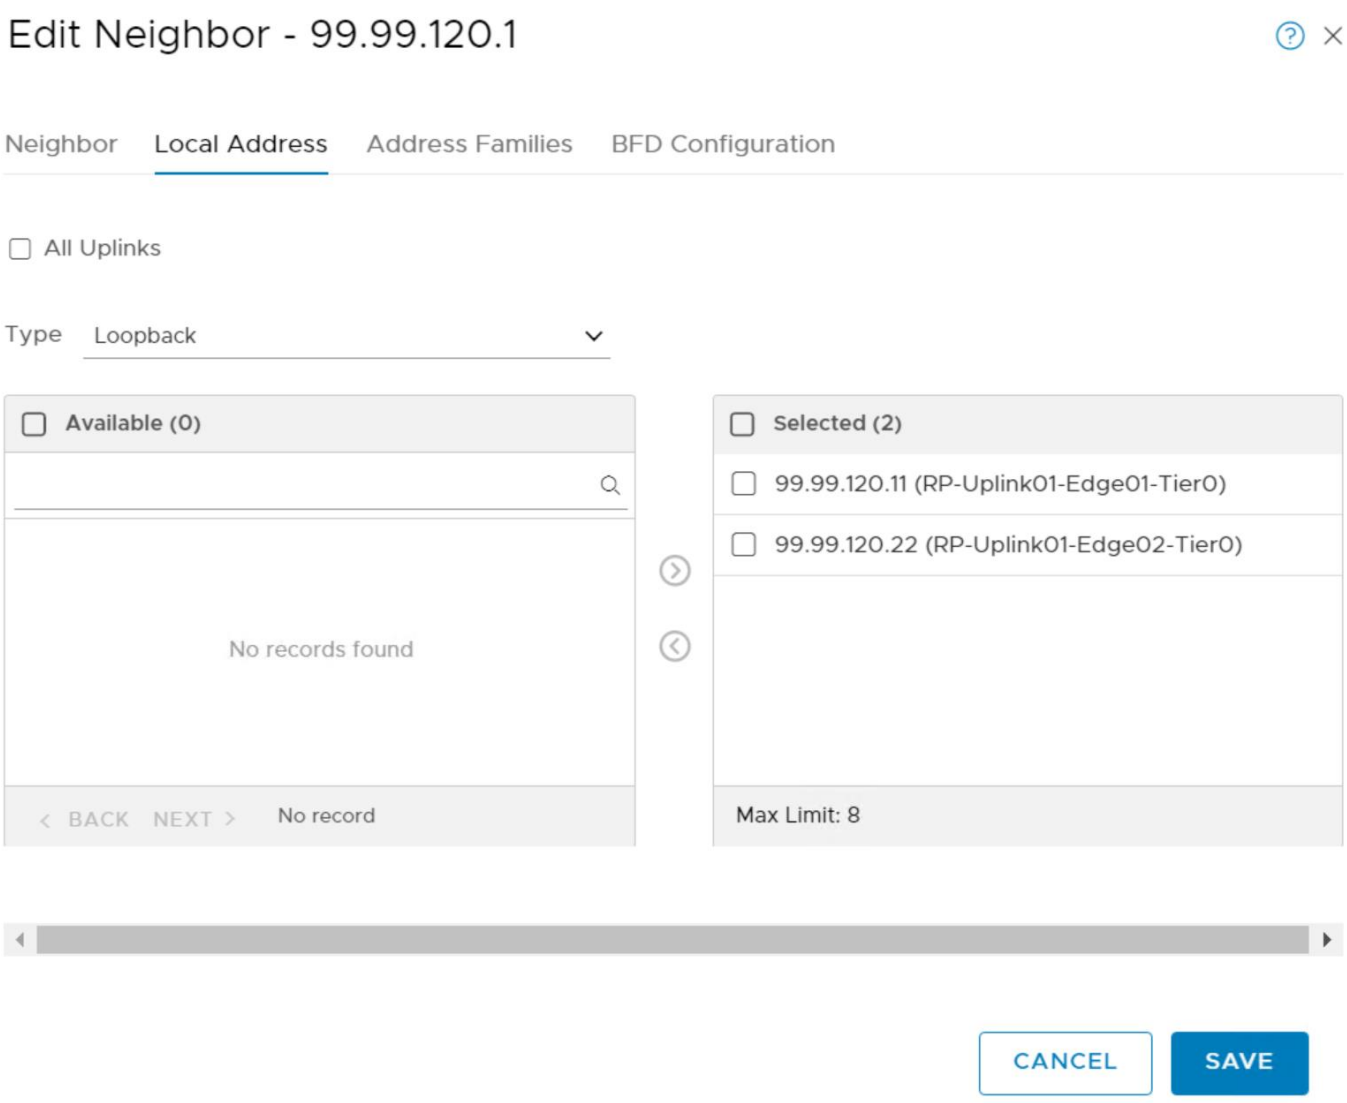
Edit Neighbor - 99.99.120.1
Neighbor
Local Address
Address Families
BFD Configuration
All Uplinks
Type
Loopback
Available (0)
No records found
BACK
NEXT
No record
Selected (2)
99.99.120.11 (RP-Uplink01-Edge01-Tier0)
99.99.120.22 (RP-Uplink01-Edge02-Tier0)
Max Limit: 8
CANCEL
SAVE

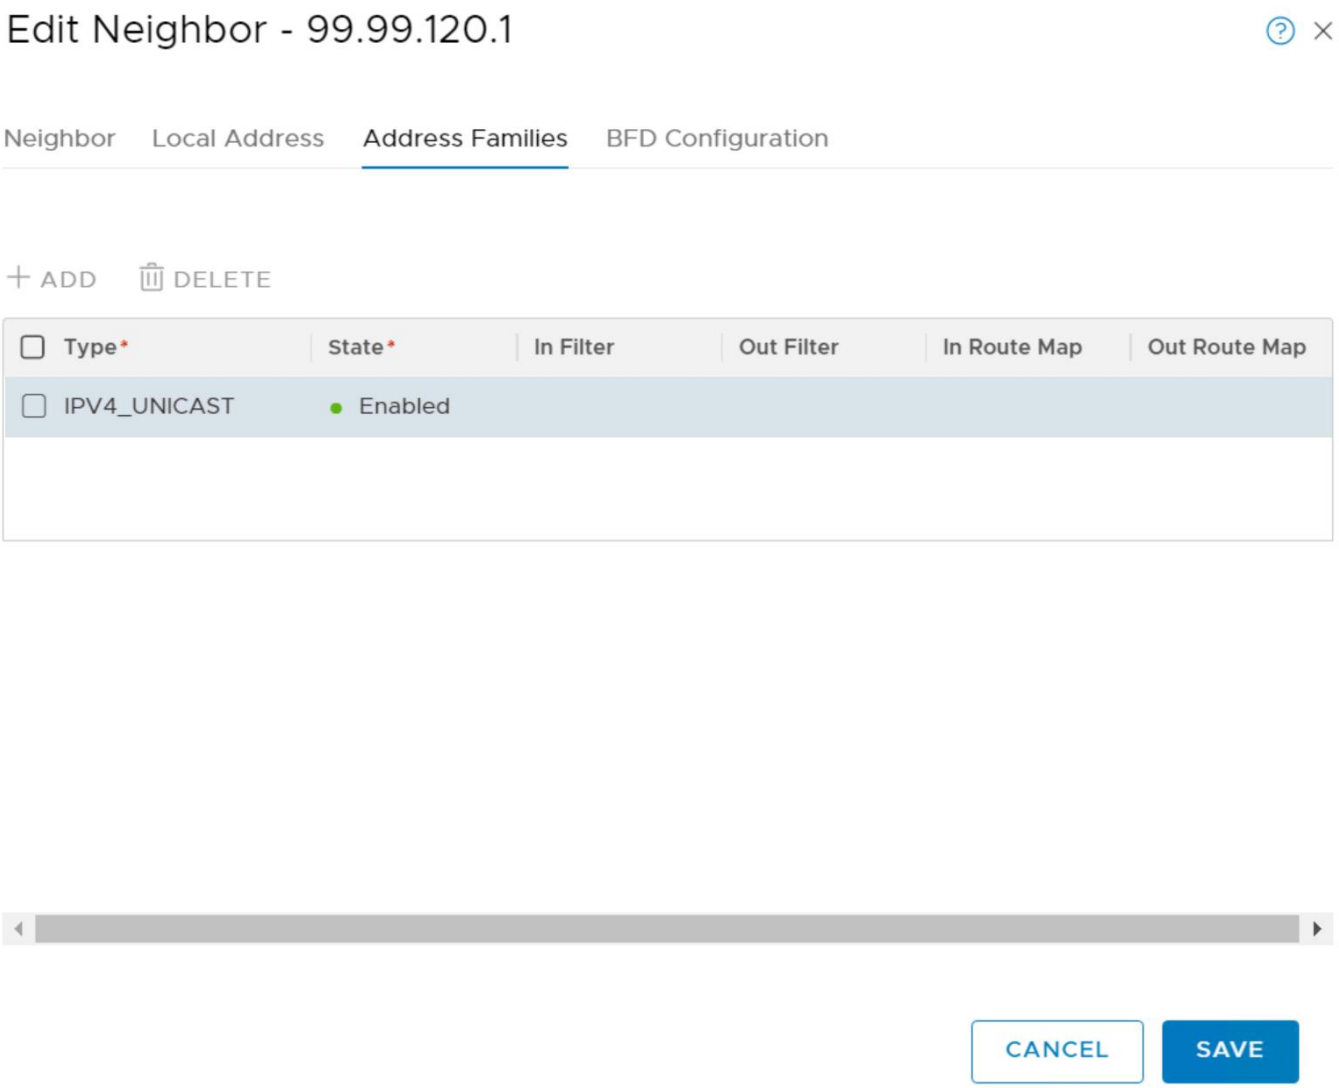
Edit Neighbor - 99.99.120.1
Neighbor
Local Address
Address Families
BFD Configuration
+ ADD
DELETE
Type*
State*
In Filter
Out Filter
In Route Map
Out Route Map
IPV4_UNICAST
Enabled
CANCEL
SAVE

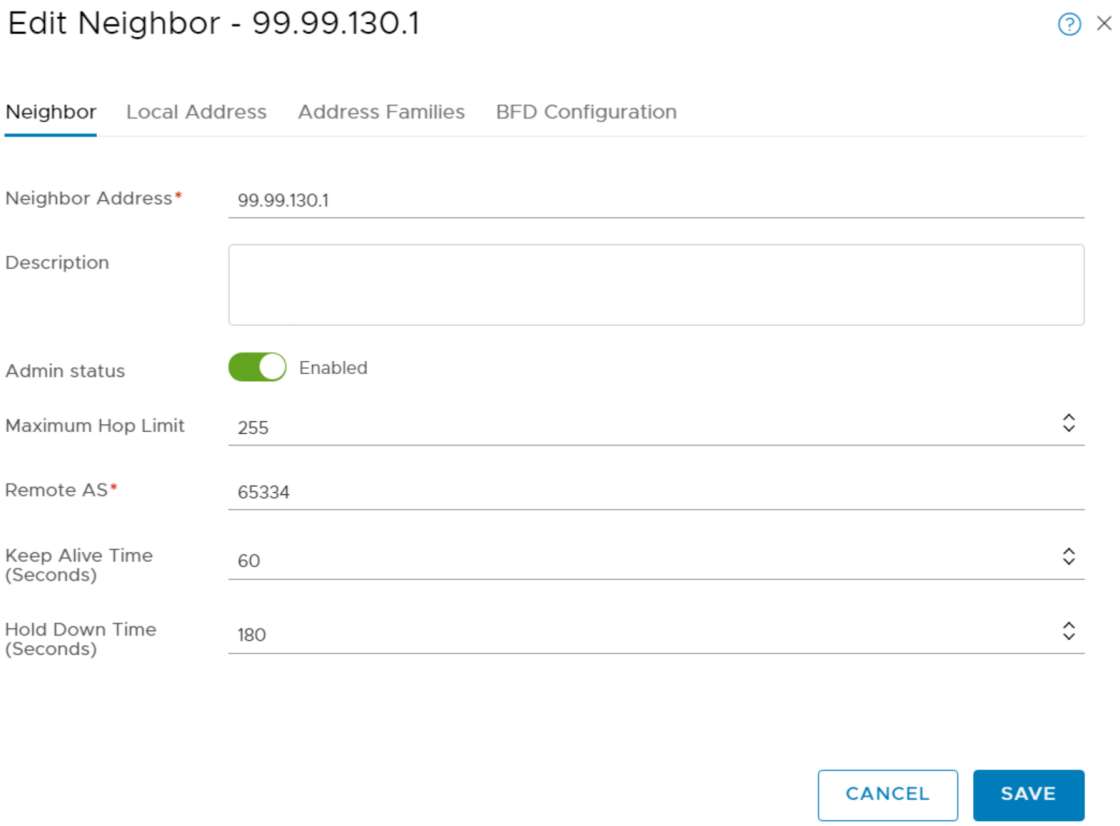
Edit Neighbor - 99.99.130.1
Neighbor
Local Address
Address Families
BFD Configuration
Neighbor Address*
99.99.130.1
Description
Admin status
Enabled
Maximum Hop Limit
255
Remote AS*
65334
Keep Alive Time (Seconds)
60
Hold Down Time (Seconds)
180
CANCEL
SAVE

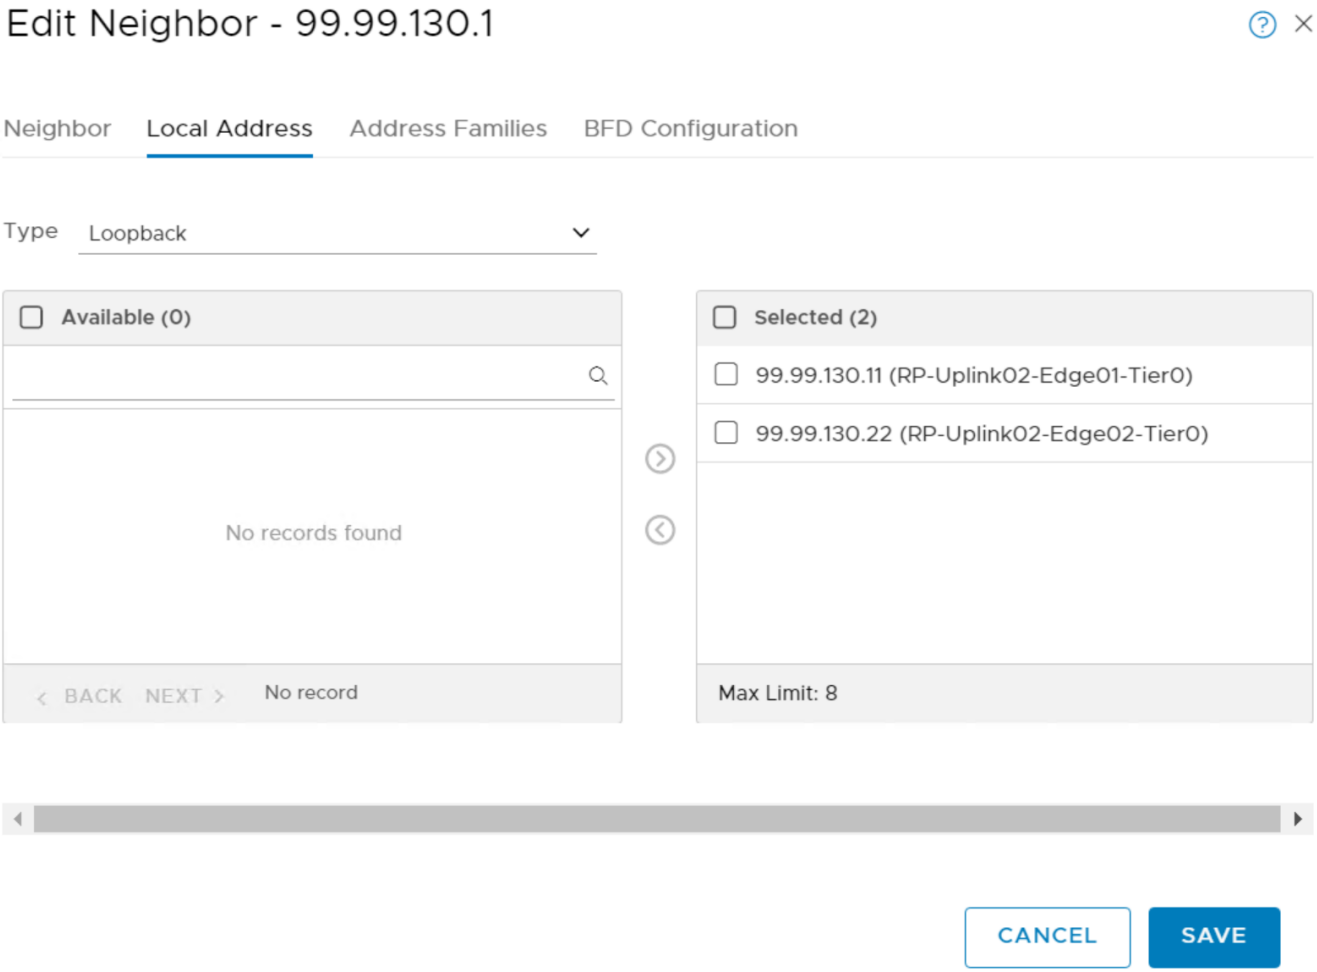
Edit Neighbor - 99.99.130.1
Neighbor
Local Address
Address Families
BFD Configuration
Type
Loopback
Available (0)
No records found
BACK
NEXT
No record
Selected (2)
99.99.130.11 (RP-Uplink02-Edge01-Tier0)
99.99.130.22 (RP-Uplink02-Edge02-Tier0)
Max Limit: 8
CANCEL
SAVE

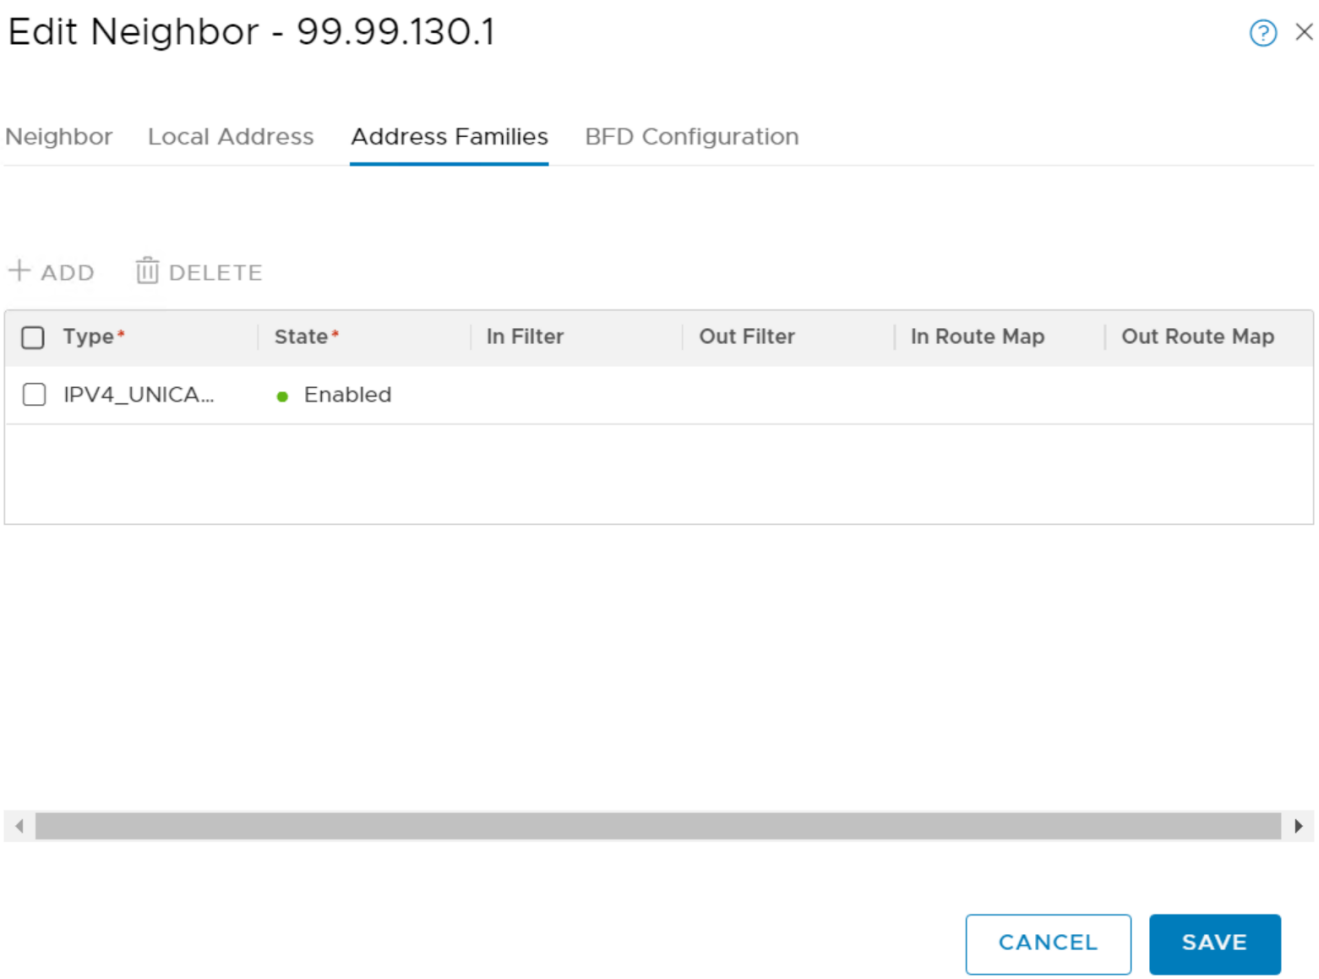

After adding the BGP peers , here is the summary

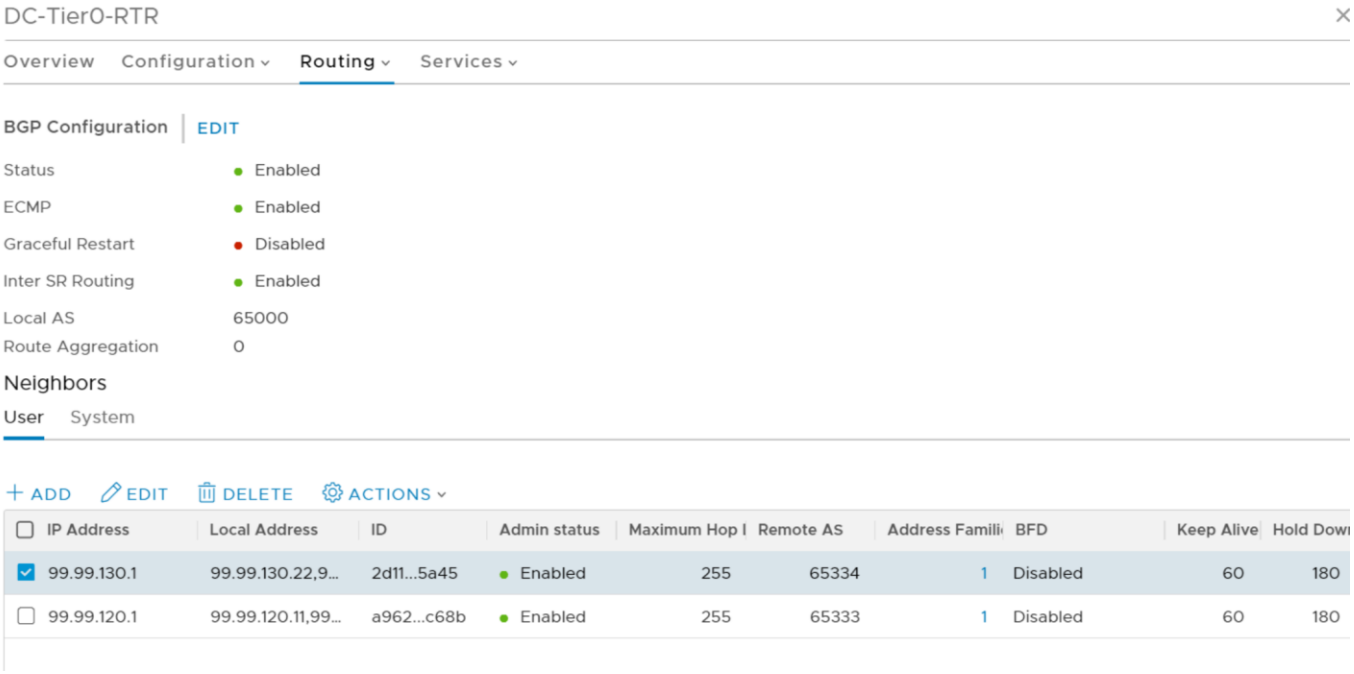

Then configure redistribution to exchange routes, e.g directly connected networks, static routes etc.

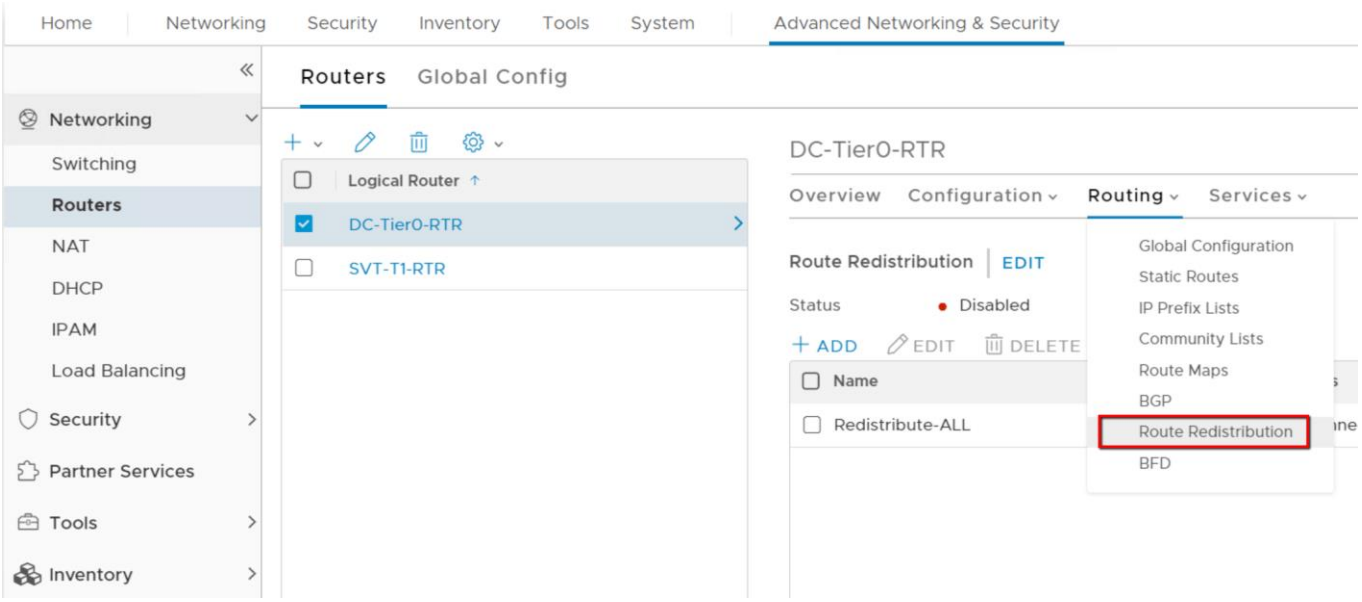

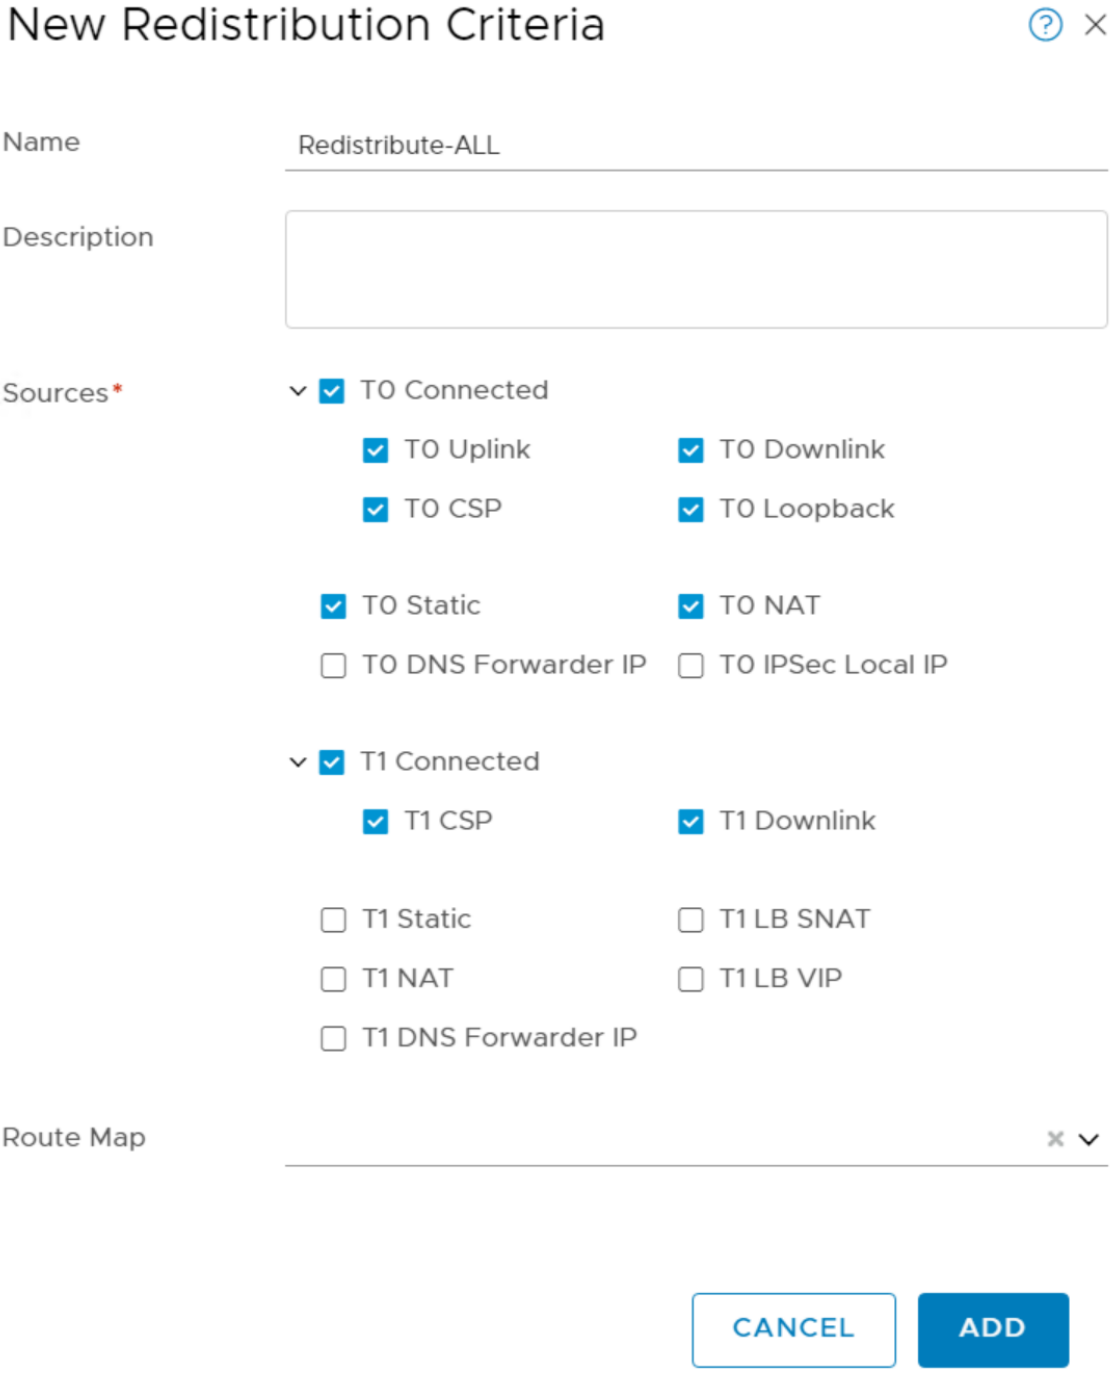
New Redistribution Criteria
Name
Redistribute-ALL
Description
Sources*
T0 Connected
T0 Uplink
T0 Downlink
T0 CSP
T0 Loopback
T0 Static
T0 NAT
T0 DNS Forwarder IP
T0 IPSec Local IP
T1 Connected
T1 CSP
T1 Downlink
T1 Static
T1 LB SNAT
T1 NAT
T1 LB VIP
T1 DNS Forwarder IP
Route Map
CANCEL
ADD

Then Enable Route Redistribution as shown below

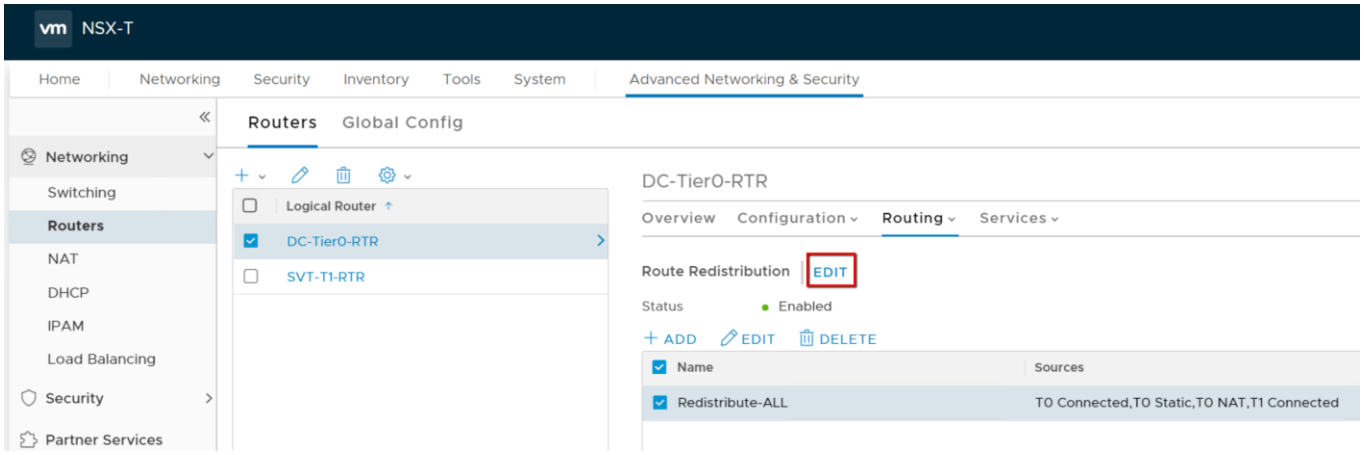

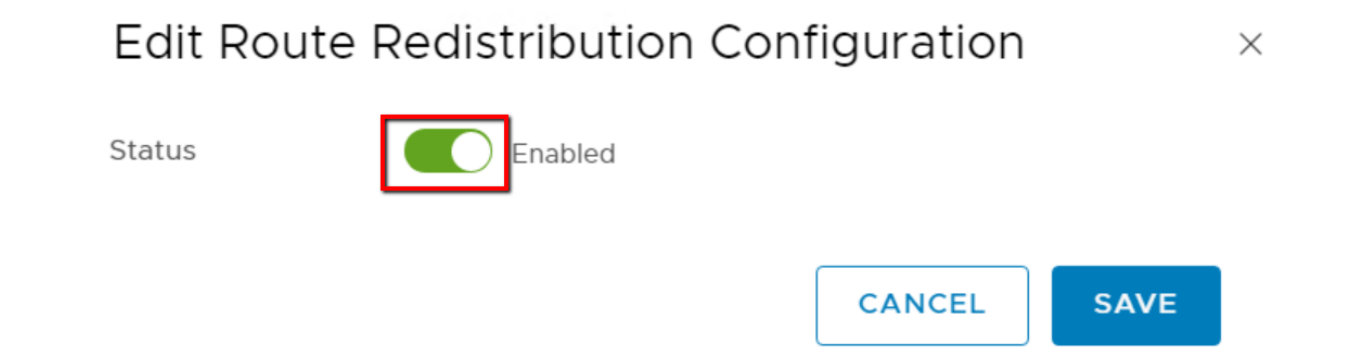

Connect Tier-0 Router with Tier-1 Router

Now connect the Tier1 router with the Tier0 router.

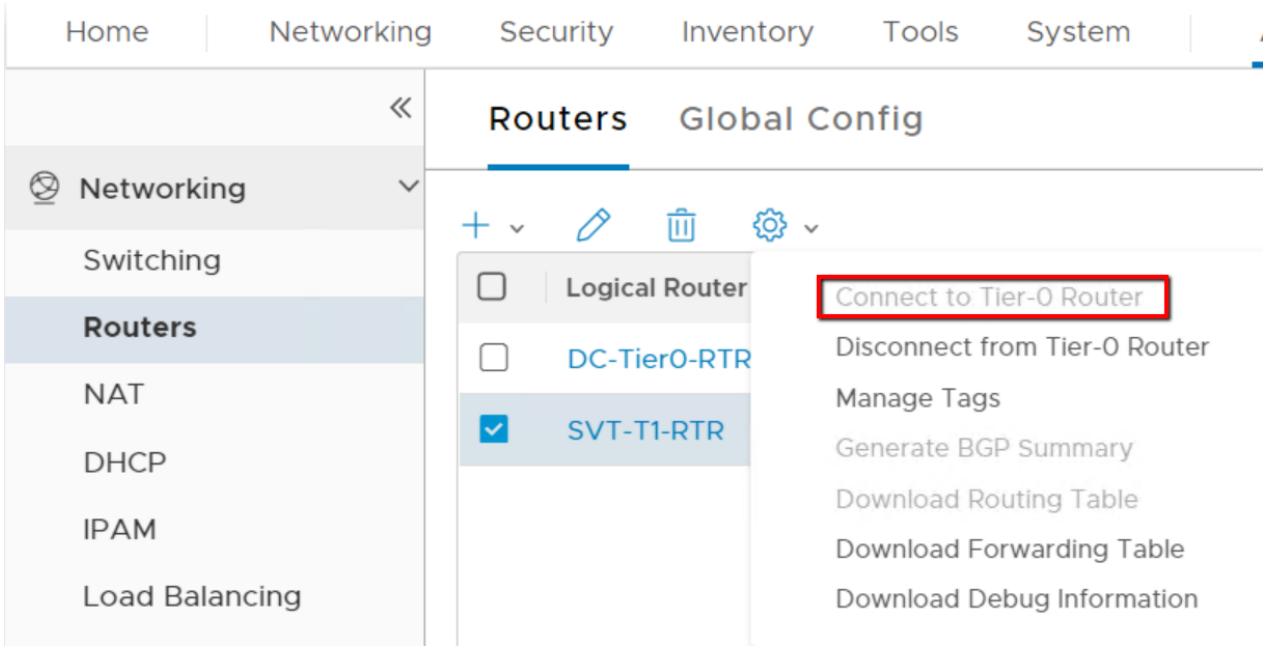

Select the Tier-0 Router from the drop down as shown below

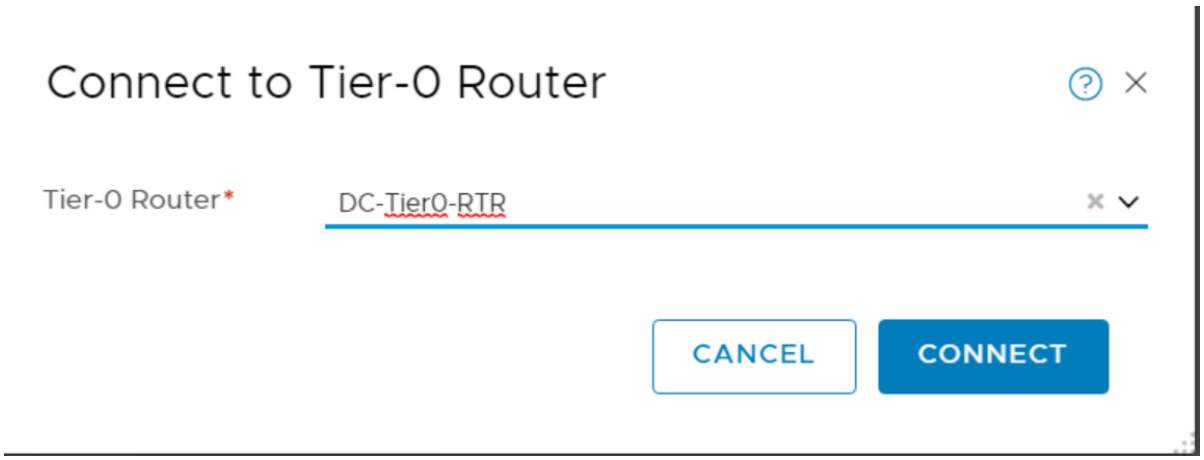

Manage the routes from this step in-case if required to filter any specific routes from redistributing it to Tier-0

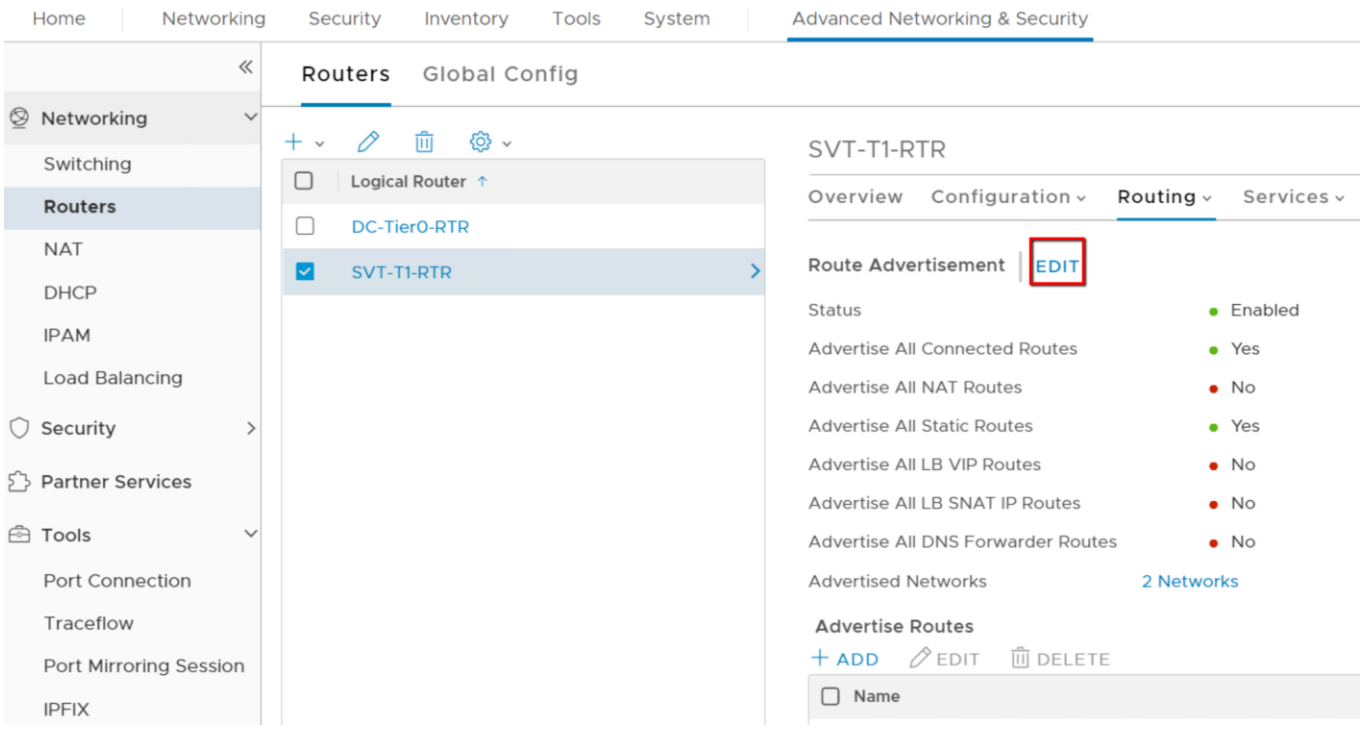
Home
Networking
Security
Inventory
Tools
System
Advanced Networking & Security
Networking
Switching
Routers
NAT
DHCP
IPAM
Load Balancing
Security
Partner Services
Tools
Port Connection
Traceflow
Port Mirroring Session
IPFIX
Routers
Global Config
Logical Router
DC-Tier0-RTR
SVT-T1-RTR
SVT-T1-RTR
Overview
Configuration
Routing
Services
Route Advertisement
EDIT
Status
Enabled
Advertise All Connected Routes
Yes
Advertise All NAT Routes
No
Advertise All Static Routes
Yes
Advertise All LB VIP Routes
No
Advertise All LB SNAT IP Routes
No
Advertise All DNS Forwarder Routes
No
Advertised Networks
2 Networks
Advertise Routes
ADD
EDIT
DELETE
Name

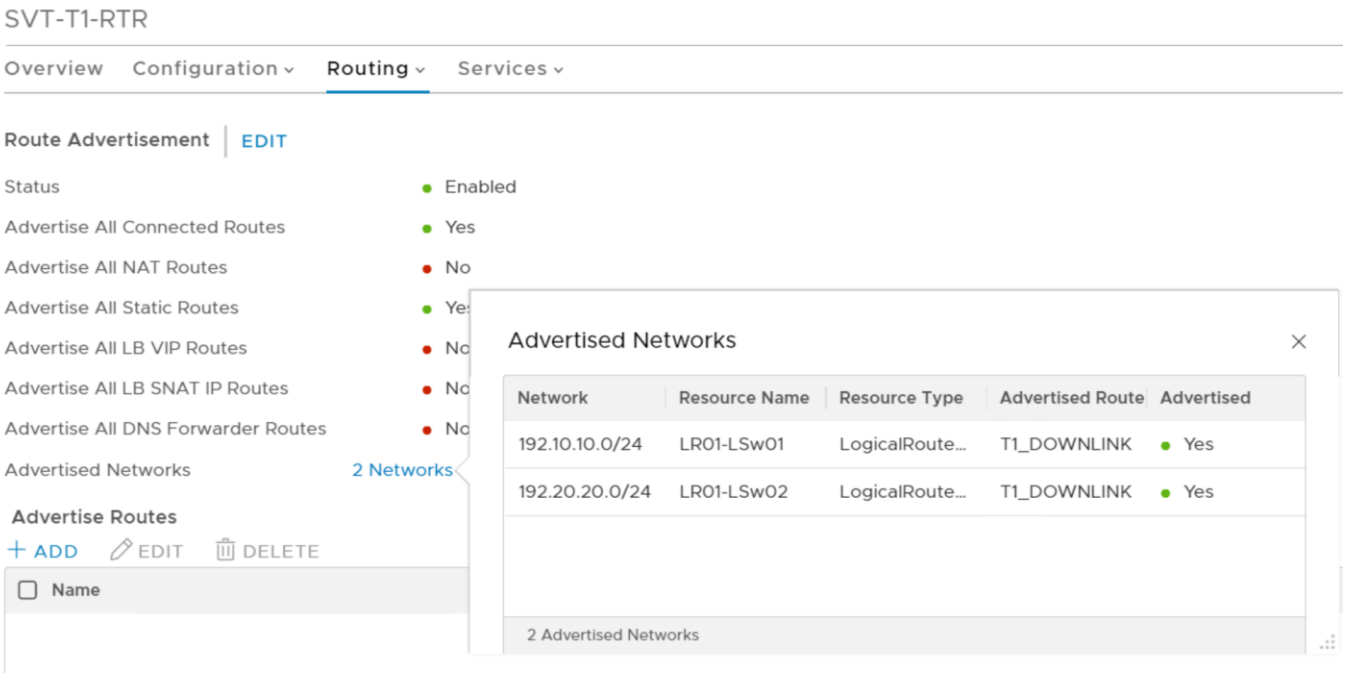
SVT-T1-RTR
Overview
Configuration
Routing
Services
Route Advertisement
EDIT
Status
Enabled
Advertise All Connected Routes
Yes
Advertise All NAT Routes
No
Advertise All Static Routes
Advertise All LB VIP Routes
Advertise All LB SNAT IP Routes
Advertise All DNS Forwarder Routes
Advertised Networks
2 Networks
Advertise Routes
ADD
EDIT
DELETE
Name
Advertised Networks
Network
Resource Name
Resource Type
Advertised Route
Advertised
192.10.10.0/24
LR01-LSw01
LogicalRoute...
T1_DOWNLINK
Yes
192.20.20.0/24
LR01-LSw02
LogicalRoute...
T1_DOWNLINK
Yes
2 Advertised Networks

AOS-CX BGP Configuration

Below is the AOS-CX switch configuration used in this example.

In BGP configuration, optionally, to allow BGP community values, use `neighbor x.x.x.x send-community`.

```
8320-SW1#

interface vlan120
    ip address 99.99.120.1/24
    no shut
    exit

router bgp 65333
    bgp router-id 99.99.120.1
    neighbor 99.99.120.11 remote-as 65000
    neighbor 99.99.120.11 update-source vlan 120
    neighbor 99.99.120.22 remote-as 65000
    neighbor 99.99.120.22 update-source vlan 120
    address-family ipv4 unicast
        neighbor 99.99.120.11 activate
        neighbor 99.99.120.11 next-hop-self
        neighbor 99.99.120.11 send-community
        neighbor 99.99.120.11 default-originate
        neighbor 99.99.120.22 activate
        neighbor 99.99.120.22 next-hop-self
        neighbor 99.99.120.22 send-community
        neighbor 99.99.120.22 default-originate
        network 99.99.120.0/24
    exit-address-family

8320-SW2#
interface vlan130
    ip address 99.99.130.1/24
    no shut
    exit

router bgp 65334
    bgp router-id 99.99.130.1
    neighbor 99.99.130.11 remote-as 65000
    neighbor 99.99.130.11 update-source vlan 130
    neighbor 99.99.130.22 remote-as 65000
    neighbor 99.99.130.22 update-source vlan 130
    address-family ipv4 unicast
        neighbor 99.99.130.11 activate
        neighbor 99.99.130.11 next-hop-self
        neighbor 99.99.130.11 send-community
        neighbor 99.99.130.11 default-originate
        neighbor 99.99.130.22 activate
        neighbor 99.99.130.22 next-hop-self
        neighbor 99.99.130.22 send-community
        neighbor 99.99.130.22 default-originate
        network 99.99.130.0/24
```

```
    exit-address-family
```

BGP verification on AOS-CX

```
8320-SW1# sh bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 65333        BGP Router Identifier  : 99.99.120.1
 Peers                  : 2            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State        AdminStatus
 99.99.120.11    65000       14      17      00h:09m:44s  Established   Up
 99.99.120.22    65000       14      15      00h:09m:45s  Established   Up

Address-family : IPv6 Unicast
------------------------------
8320-SW1# show ip route bgp

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

99.99.130.0/24, vrf default
        via  99.99.120.22,  [20/0],  bgp
        via  99.99.120.11,  [20/0],  bgp
169.254.0.128/25, vrf default
        via  99.99.120.22,  [20/0],  bgp
192.10.10.0/24, vrf default
        via  99.99.120.22,  [20/0],  bgp
        via  99.99.120.11,  [20/0],  bgp
192.20.20.0/24, vrf default
        via  99.99.120.22,  [20/0],  bgp
        via  99.99.120.11,  [20/0],  bgp


8320-SW2# sh bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 65334        BGP Router Identifier  : 99.99.130.1
 Peers                  : 2            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State        AdminStatus
 99.99.130.11    65000       5       5       00h:00m:14s  Established   Up
 99.99.130.22    65000       5       6       00h:00m:14s  Established   Up
```

```
Address-family : IPv6 Unicast
------------------------------

8320-SW2# sh ip route bgp

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

99.99.120.0/24, vrf default
        via  99.99.130.22,  [20/0],  bgp
        via  99.99.130.11,  [20/0],  bgp
169.254.0.128/25, vrf default
        via  99.99.130.22,  [20/0],  bgp
192.10.10.0/24, vrf default
        via  99.99.130.22,  [20/0],  bgp
        via  99.99.130.11,  [20/0],  bgp
192.20.20.0/24, vrf default
        via  99.99.130.22,  [20/0],  bgp
        via  99.99.130.11,  [20/0],  bgp
```

BGP verification from Edge VM (NSX-T)

```
edge02> vrf 1
edge02(tier0_sr)> get bgp neighbor summary
BFD States: NC - Not configured, AC - Activating,  DC - Disconnected
            AD - Admin down, DW - Down, IN - Init, UP - Up
BGP summary information for VRF default for address-family: ipv4Unicast

Router ID: 99.99.130.22  Local AS: 65000

Neighbor              AS          State Up/DownTime  BFD InMsgs  OutMsgs InPfx
OutPfx

99.99.130.1           65334       Estab 00:18:37     NC  27      23      1      5
169.254.0.130         65000       Estab 1d02h34m     NC  95912   95914   5      6
99.99.120.1           65333       Estab 00:29:04     NC  1524    850     1      5

BFD States: NC - Not configured, AC - Activating,  DC - Disconnected
            AD - Admin down, DW - Down, IN - Init, UP - Up
BGP summary information for VRF default for address-family: ipv6Unicast

Router ID: 99.99.130.22  Local AS: 65000

Neighbor       AS         State Up/DownTime  BFD InMsgs  OutMsgs InPfx  OutPfx

169.254.0.130  65000       Estab 1d02h34m     NC  95912   95914   1      1


edge02(tier0_sr)> get route bgp ipv4

Flags: t0c - Tier0-Connected, t0s - Tier0-Static, B - BGP,
t0n - Tier0-NAT, t1s - Tier1-Static, t1c - Tier1-Connected,
t1n: Tier1-NAT, t1l: Tier1-LB VIP, t1ls: Tier1-LB SNAT,
```

```
t1d: Tier1-DNS FORWARDER, > - selected route, * - FIB route

Total number of routes: 1

b  > * 0.0.0.0/0 [20/0] via 99.99.130.1, uplink-277, 00:00:02
b  > * 0.0.0.0/0 [20/0] via 99.99.120.1, uplink-268, 00:00:02
```

Table of Figures

References

1. NSX Data Center Installation Guide
2. NSX Data Center Admin Guide