Dell Networking: Multitenancy Across Physical and Logical Environments with VRF-lite and VMware NSX

Dell Networking Technical Marketing
December 2014

A Dell Technical White Paper

Revisions

Date	Description	Authors
12/21/14	Version 1.2	Humair Ahmed, Sr. Technical Marketing Engineer, Dell Networking
12/20/14	Version 1.1	Humair Ahmed, Sr. Technical Marketing Engineer, Dell Networking
12/18/14	Version 1 - Initial release	Humair Ahmed, Sr. Technical Marketing Engineer, Dell Networking

Table of contents

1 Overview

The objective of this paper is to demonstrate multitenancy across physical and logical environments leveraging Network Virtualization Overlay (NVO) and VRF-lite.

Virtual Routing and Forwarding (VRF) allows a Dell Networking L3 switch/router to be partitioned into multiple Virtual Routers (VRs). The control and data plane are isolated in each VR so traffic does not flow across VRs; this allows different routing tables to simultaneously exist within the same physical L3 switch/router. VRF-lite also supports route leaking, which enables routes to be distributed across VRs in a controlled manner. VRF-lite is supported on the following Dell Networking switches: S4810, S4820T, S5000, S6000, Z9500, and C-Series chassis. In this white paper VRF-lite is used on Dell S6000 switches.

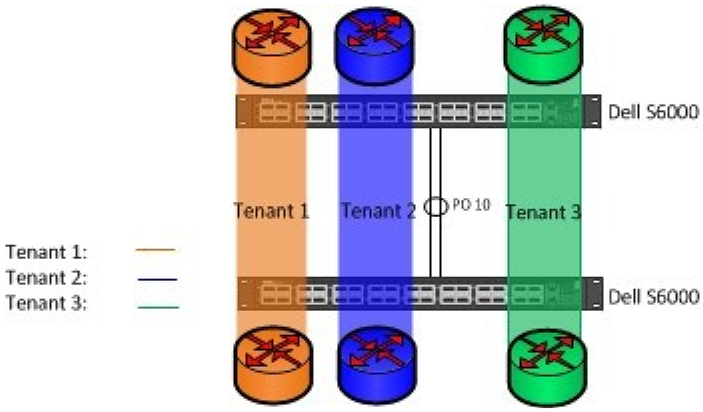

Figure 1 Multitenancy with VRF-lite on Dell S6000s

NVO encapsulates L2 frames within L3 packets and allows the creation of logical networks over physical networks. In this paper, VMware NSX-vSphere is used to create logical networks via Virtual Extensible LAN (VXLAN). VXLAN is a standard network overlay technology where MAC frames are encapsulated into a VXLAN and UDP header; communication occurs between two endpoints called Virtual Tunnel Endpoints (VTEPs). VMware NSX uses VXLAN to build logical L2 networks over any L2/L3 physical IP infrastructure.

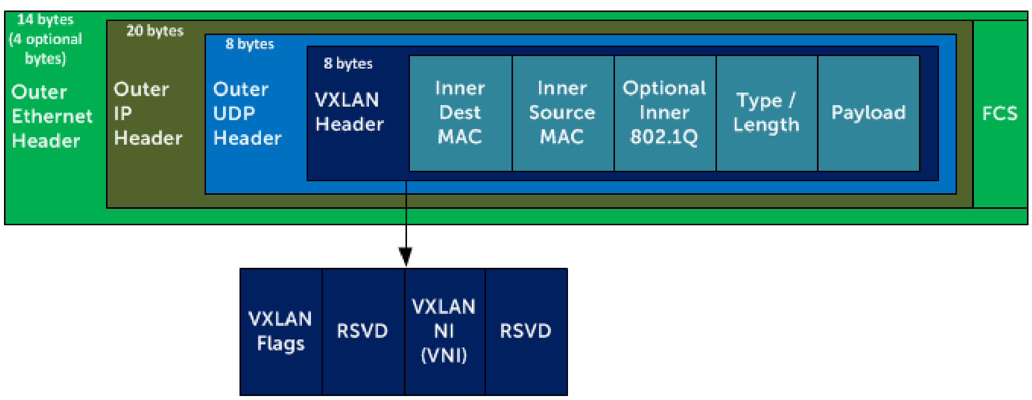

Figure 2 VXLAN Encapsulated Frame

Multitenancy can easily be achieved with VMware NSX by creating multiple virtual routers; this allows for segmentation of routing information for different tenants.

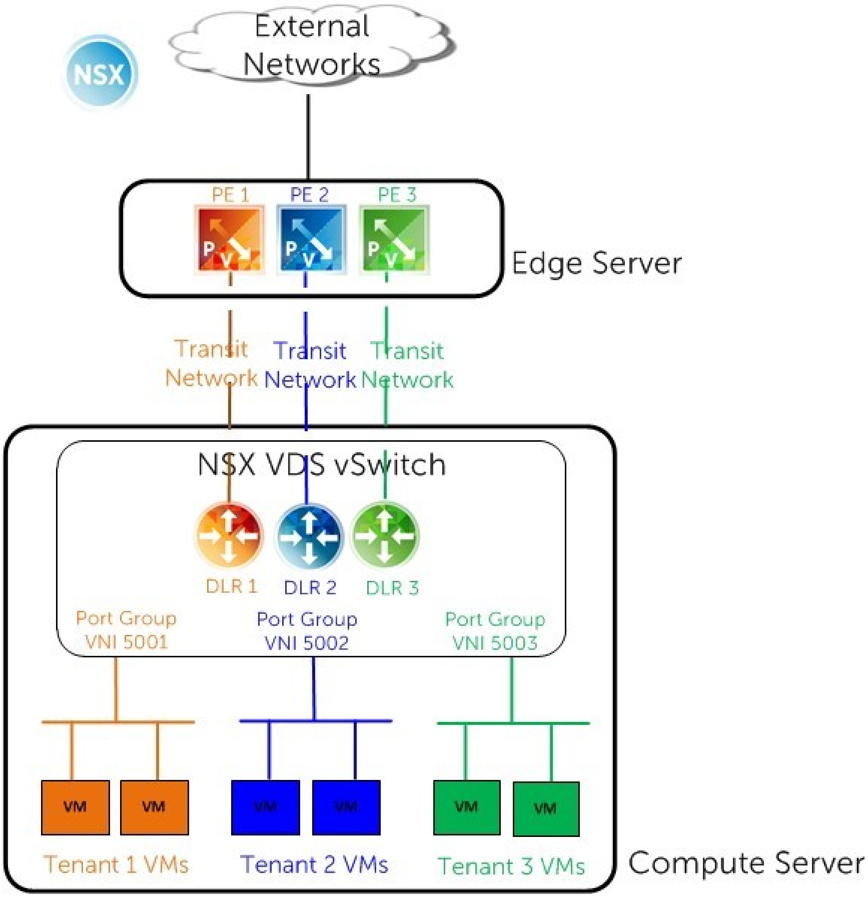

Figure 3 Multitenancy with VMware NSX

Together, NVO and VRF-lite allow for creating multitenancy across physical and logical environments while allowing tenants to have overlapping IP addresses.

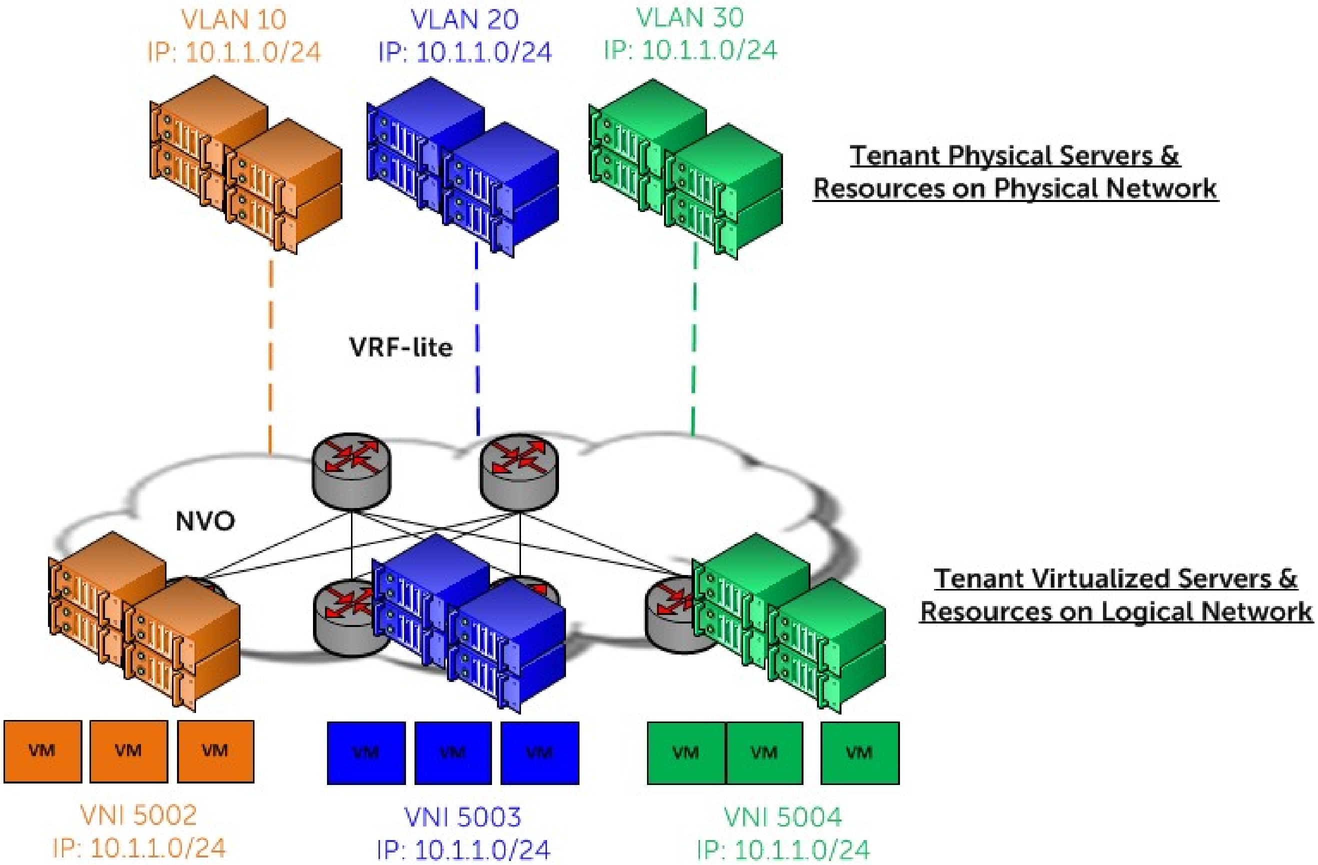

Figure 4 VRF-lite used in conjunction with VXLAN Network Overlay in a virtualized environment

For more specific information on VRF-lite, please see the Dell Networking: Multitenancy with VRF-lite white paper. For more information on Dell Networking with VMware-NSX-vSphere, see the Network Virtualization with Dell Infrastructure and VMware NSX Reference Architecture white paper.

Goals

1. Configure logical components/networks via NSX for multitenancy
2. Configure VRF instances on Dell S6000 switches via VRF-lite
3. Validate multitenancy across physical and logical networks

2 Use Case for Multitenancy Across Logical and Physical Networks via NSX and VRF-lite

VRF-lite can be useful in virtualized environments leveraging Network Virtualization Overlays (NVOs) where multitenancy is also needed on the physical network when bridging between logical and physical environments.

For example, VMware NSX can be utilized to provide network virtualization and multitenancy in the logical space via NVO. At the same time, if tenant workloads will be traversing both the logical and physical network, VRF-lite can be used on the physical network. This allows for multitenancy which supports overlapping IP addresses across both the logical and physical network/resources.

Examples and applications include:

- Mapping of non-virtualized resources (Ex: databases, file servers, etc.) in the physical environment that need to be part of a tenant's environment where the physical resources for each tenant will also utilize the same IP address space across all tenants.
- Multitenancy across two remote locations with one location employing NVO and the other location using traditional VLANs with VRF-lite on the physical environment.
- Using VFR-lite as a migration strategy for network virtualization where multitenancy between logical and physical networks/resources needs to be maintained.

3 NSX and VRF-lite - Example Network Designs

Figure 5 below presents an example where each VXLAN Network Identifier (VNI) on the logical network maps to a VLAN on the physical network. Tenant logical switches (represented by a VNI) connect to a tenant Distributed Logical Router (DLR) in the logical space and maps to a VLAN in the physical space. In this example, an IP address is placed on the physical VLAN on the Dell S6000 switch and tied to a specific VR for the respective tenant. This is only needed in this use case to allow for overlapping IP addresses on the physical network. The physical resources connecting to the VRs can be on the same subnet or on different subnets. In either case, this setup allows for overlapping IP addresses for both the logical and physical resources. In this setup, if routing to a physical resource is needed, it is handled by the VR for the respective tenant.

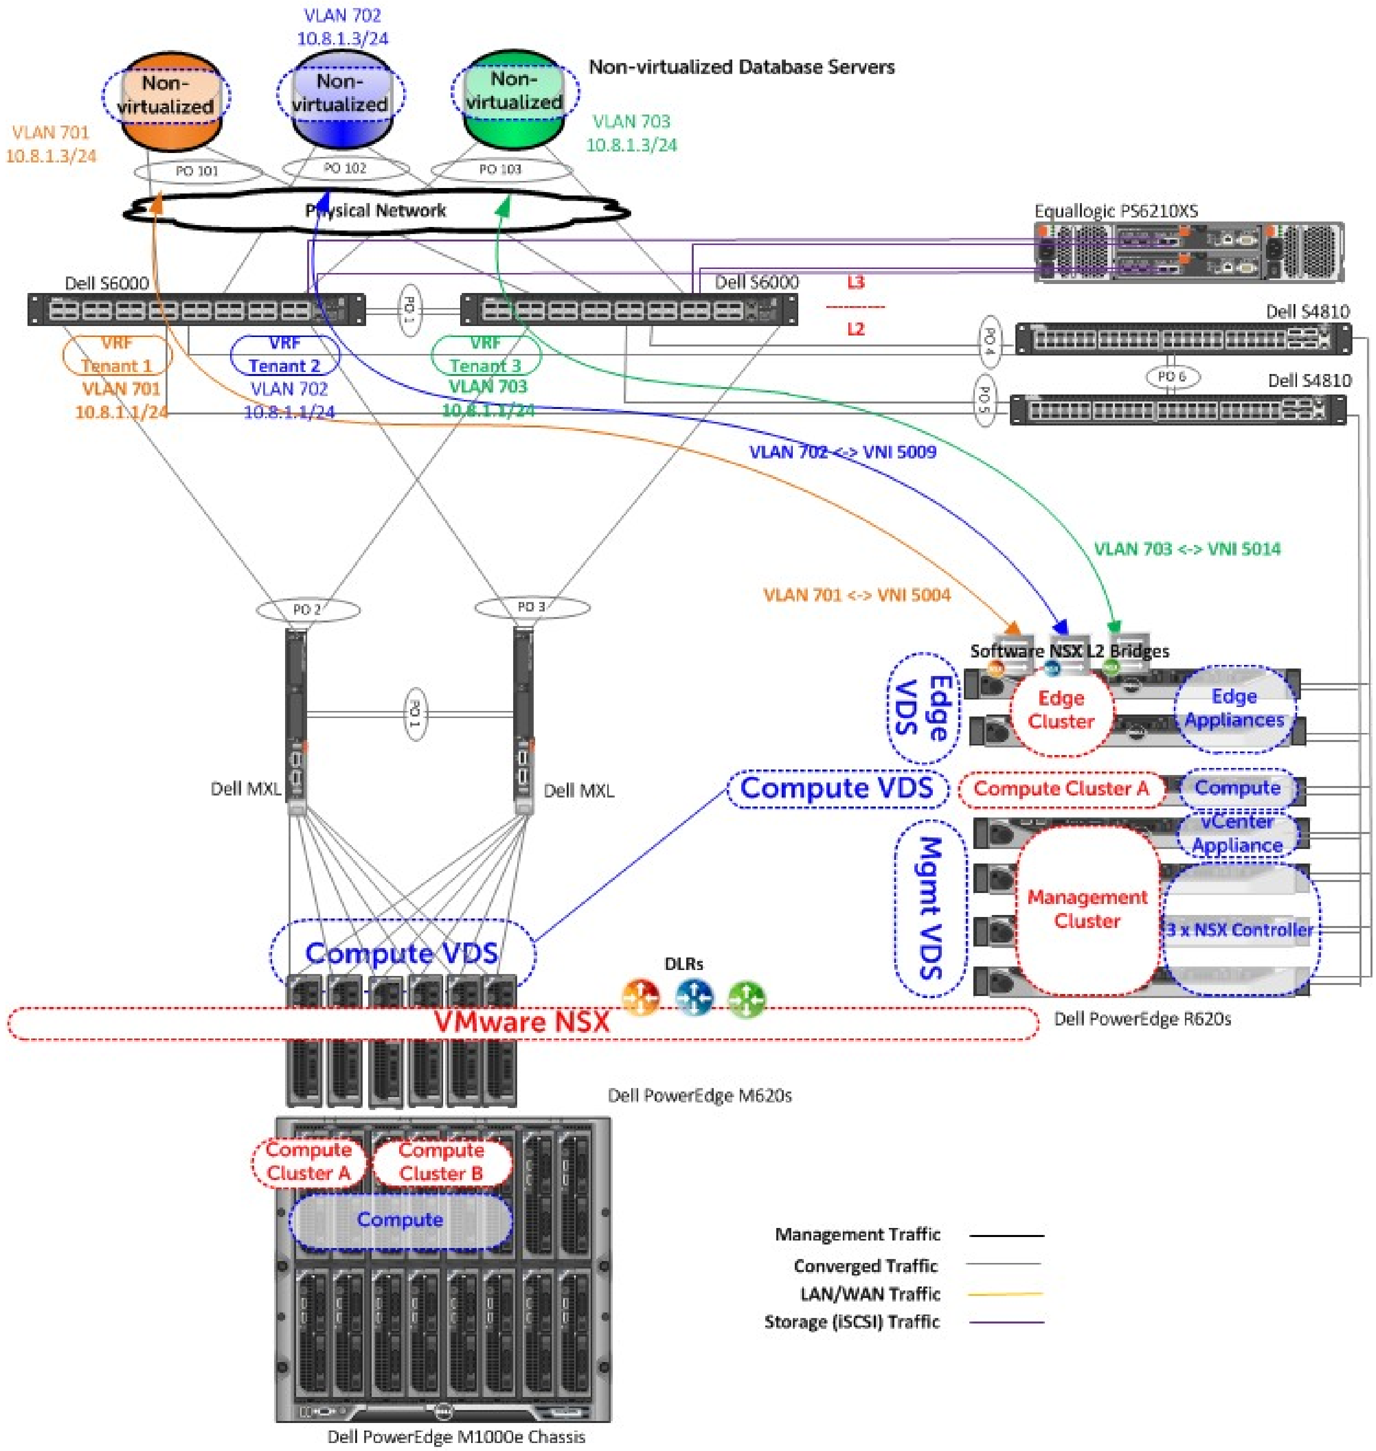

Figure 5 Multitenenacy across logical and physical networks with NSX and VRF-lite using NSX L2 Gateway

It's also possible to use the VMware Edge Services Router (ESR) / Perimeter Edge (PE) to route to external physical resources and still use VRF on the physical routers to allow for overlapping IP addresses. Figure 6 shows this specific scenario, which is discussed in more detail in this white paper.

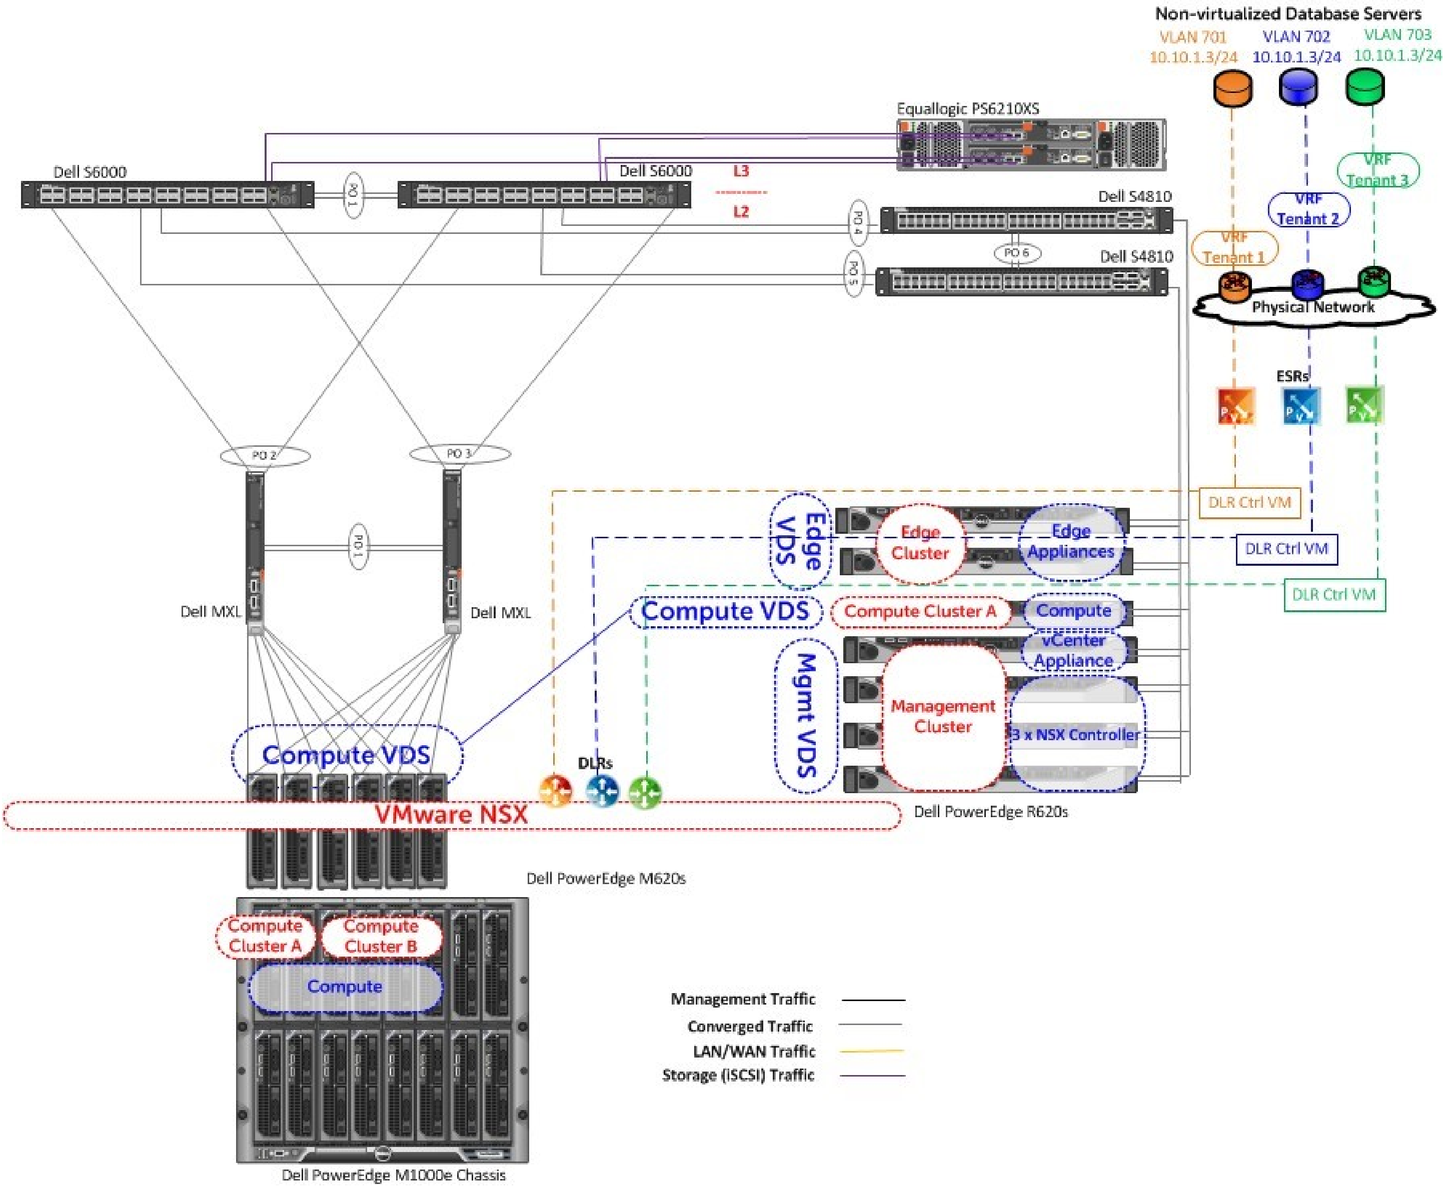

Figure 6 Multitenenacy across logical and physical networks with NSX and VRF-lite using NSX ESR

The difference between Figure 6 and the previous example is, instead of using the VMware NSX L2 gateway to bridge between VXLAN and VLAN and have the VR on the physical switch do the routing, the ESR for the respective tenant does the routing, while the VRs on the physical switches still allow for overlapping IP addresses for the physical network/resources.

In this design, there are three tenants represented by different colors (orange, blue, and green). A DLR is created for the logical network via NSX for each tenant. The DLR is a kernel level module within each hypervisor participating in the logical network; it isolates routing in the logical environment for each tenant. In this setup, a separate ESR or Perimeter Gateway is also used for each tenant; the ESR peers with the physical external networks.

VRF-lite is used on the spine Dell S6000 switches to create a VR for each of the three tenants. The VRs isolate the routing for each tenant and allow for overlapping IPs similar to how DLRs provide isolation and overlapping IPs on the logical network.

Multitenancy is achieved because each tenant in the logical environment has a different Edge VLAN that also exists on the physical switches and is mapped to the respective tenant VRF. Figure 7 shows an example of a separate Edge VLAN for Tenant 1 (VLAN 501) and a separate Edge VLAN for Tenant 2 (VLAN

502) created on the VMware Virtual Distributed Switch (VDS). In this setup, each tenant has two ESR virtual appliances in their respective VLAN deployed as Active/Standby; starting with NSX 6.1, ESRs can also be deployed as Active/Active allowing for equal-cost multi-path (ECMP) routing. The ESR will peer with the external network and share routing information. VRFs on the physical switches/routers will map to the respective tenant edge VLAN to keep routing and IP addressing isolated for each tenant.

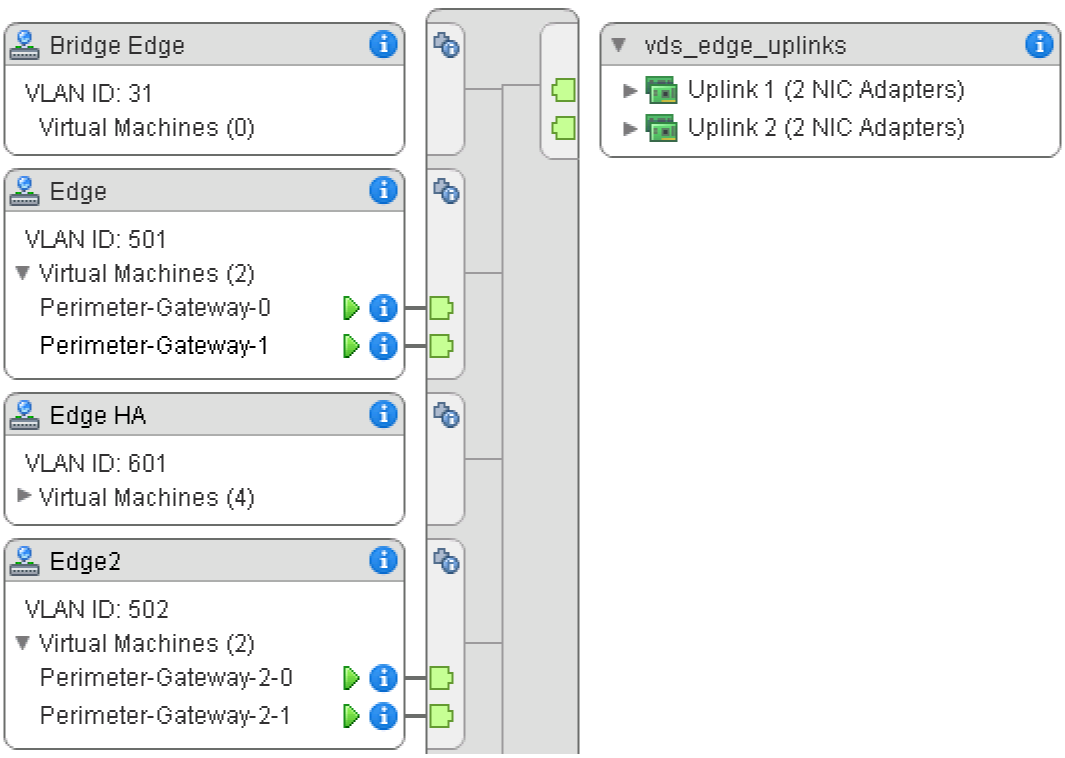

Figure 7 Image of part of VDS switch showing Tenant 1 VLAN 501 and Tenant 2 VLAN 502

The Dell S6000s use Dell's Virtual Link Trunking (VLT) technology to provide an Active/Active connection to the ToR/access layer switches where the respective tenant databases are connected. Note, the IP addressing scheme is overlapping and consistent across all tenants.

4 VMware NSX Configuration

4.1 Creating Multitenancy in the Logical Environment

Initially, for one tenant, one ESR and one DLR with high-availability can be deployed as shown below.

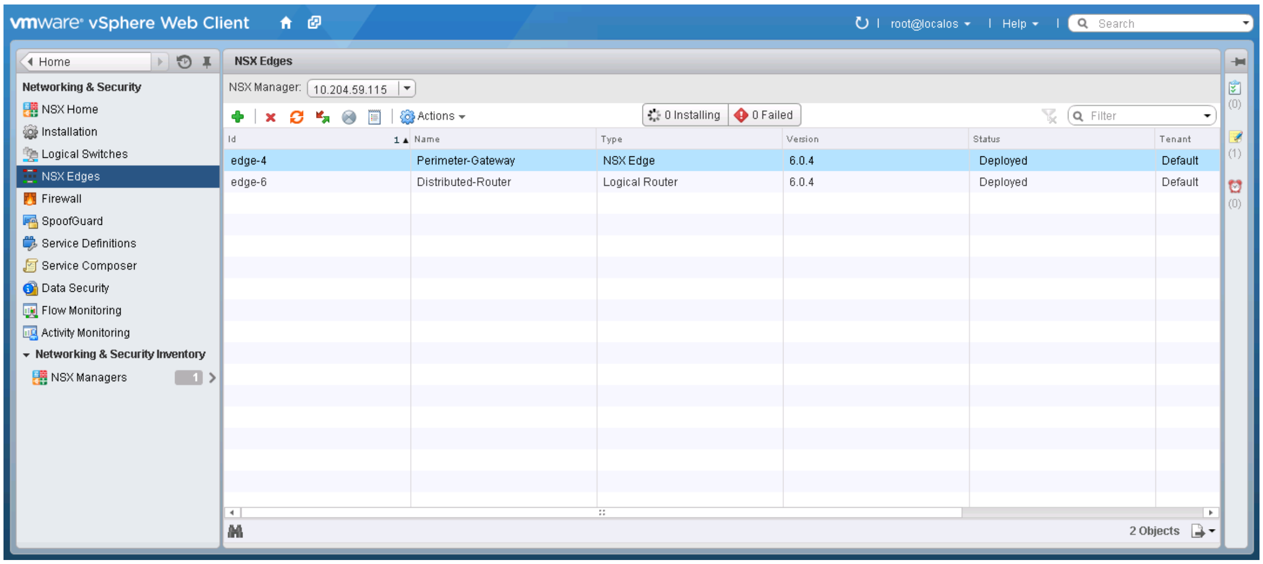

Figure 8 NSX Edges deployed for one tenant

For multitenancy in the logical environment, a separate router must be created for each tenant to isolate tenant routing and allow for overlapping IP addresses. Clicking the green + sign allows for creating additional ESRs and DLRs as shown below in Figure 9.

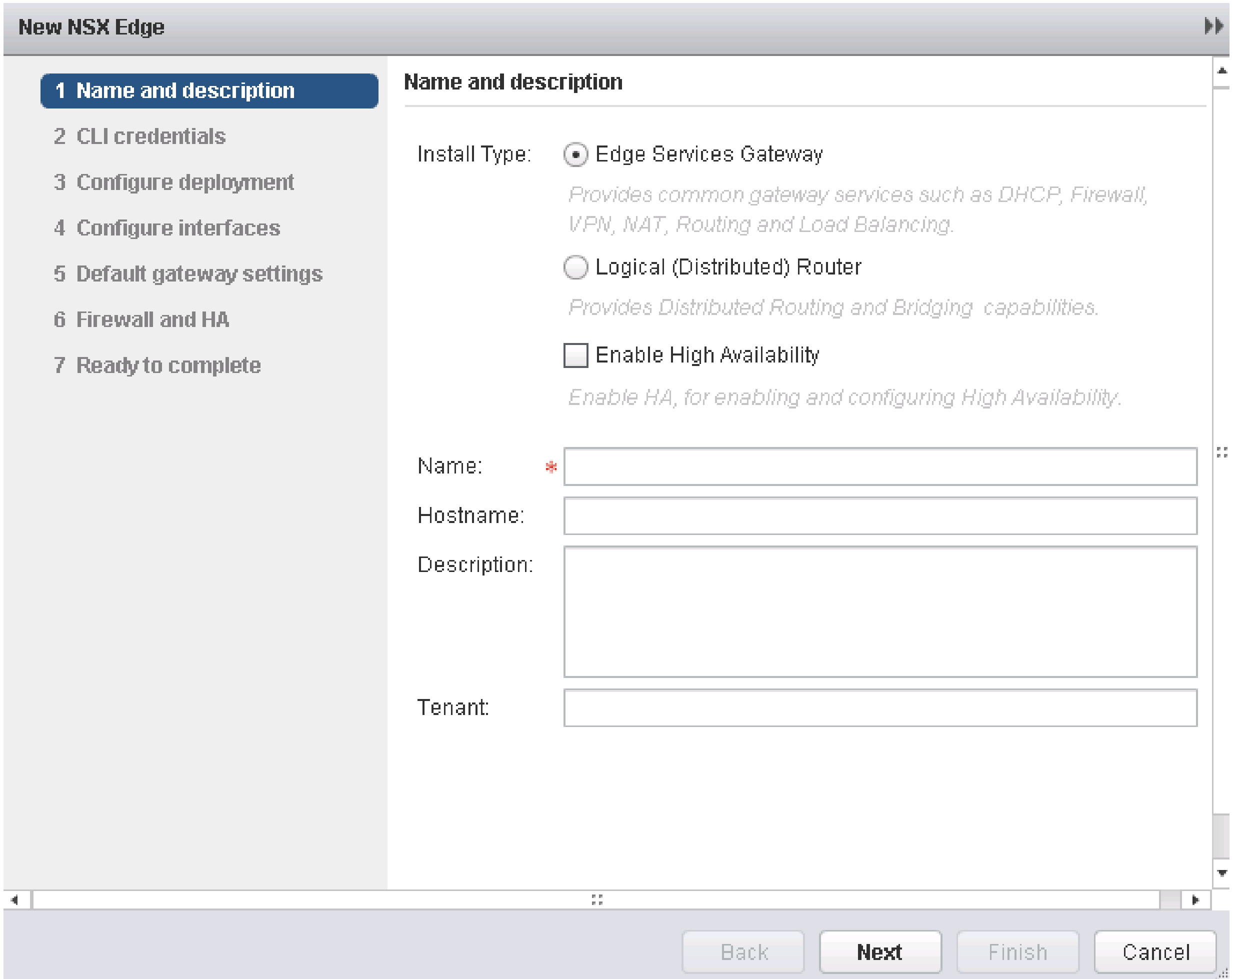

Figure 9 Adding a NSX Edge

Figure 10 shows the final step of creating a new ESR for Tenant 2. The ESR is deployed in HA mode and each active and standby ESR respectively is installed on a separate ESX host and utilizes a different datastore on the EqualLogic PS6210SX iSCSI array. The appliance can be installed in four different sizes based on the system resources available and additional L4-L7 services which will be deployed. The sizes that can be deployed are Compact, Large, X-Large, and Quad Large. The larger NSX Edge sizes provide for more CPU, memory, and disk space. As this appliance will be utilized for several L4-L7 services, it is deployed with a X-Large size.

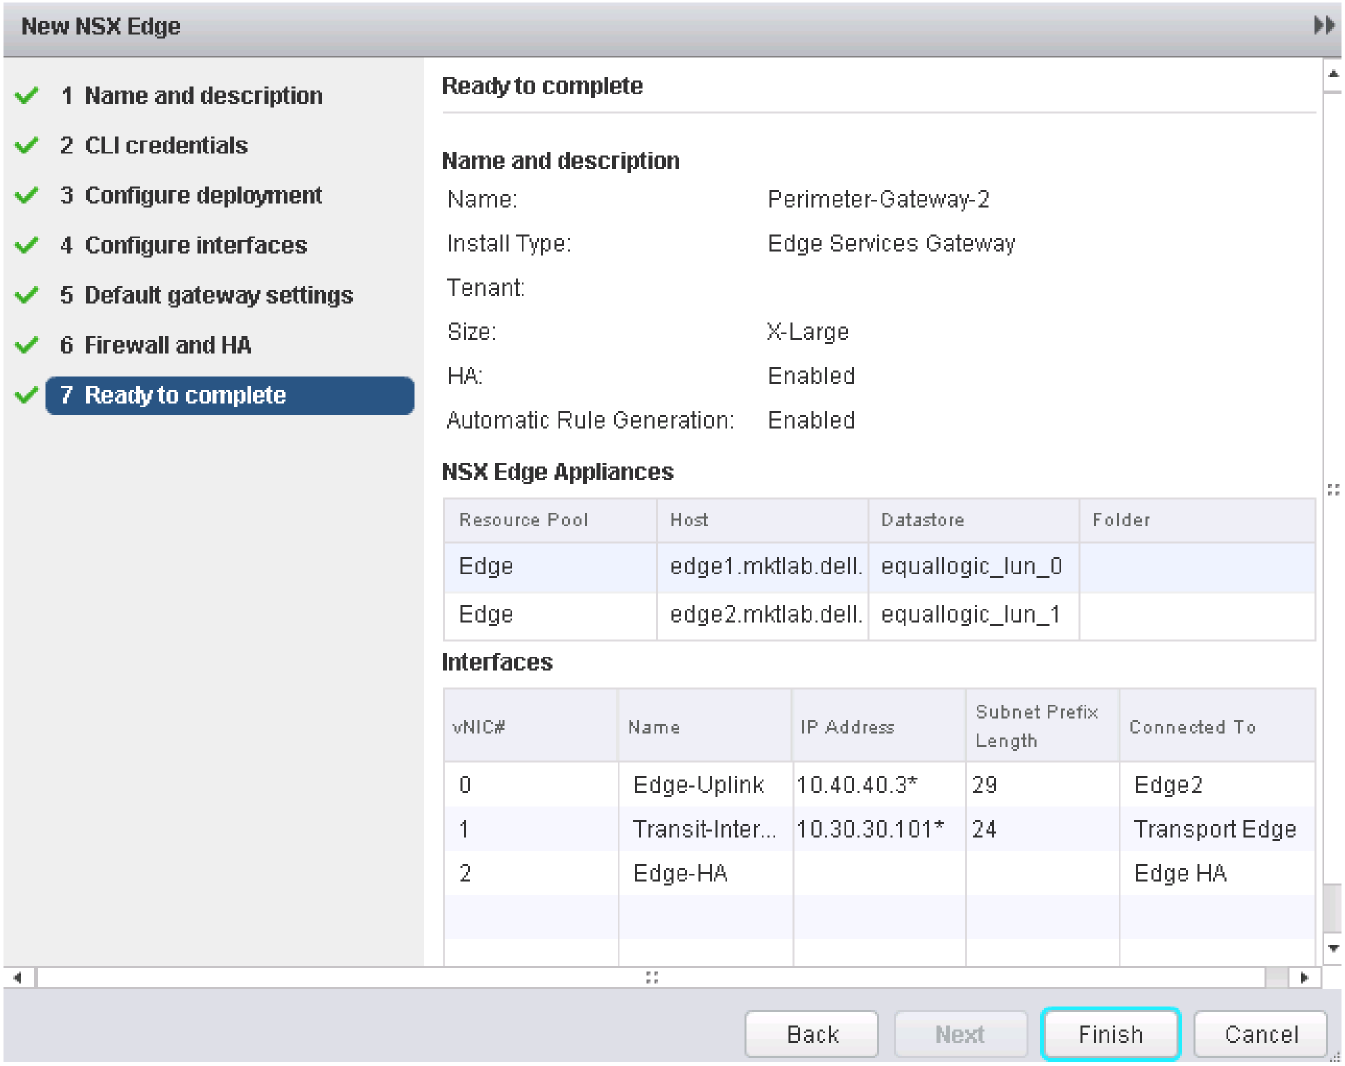

Figure 10 Final step of deploying an ESR for Tenant 2

The following interfaces on the ESR have been configured.

- vNIC 0 is connected to the Edge2 distributed virtual port group which is mapped to VLAN 502. This is important, because VLAN 502 is the edge VLAN and associated to the Tenant 2 VRF on the physical switches/routers, and this is how the tenant-specific routing and IP addresses are isolated on the physical network. This interface is used to peer with the external/physical network.
- vNIC 1 is connected to the Transport Edge distributed virtual port group which is used for the Transport Network (VLAN 401) where VTEP/VXLAN communication takes place and where the DLR must be connected to.
- vNIC 2 is connected to the Edge HA port group (VLAN 601). This is used by the edge appliance for heartbeat and high availability.

A DLR for each tenant must be deployed. The DLR kernel level module has already been installed in each host for distributed routing during the initial NSX setup. A DLR Control VM virtual appliance is installed when a new DLR is created; it provides the control plane and central management of the DLR and facilitates peering with the external physical network. Figure 11 below shows the final step of deploying a DLR for Tenant 2. In this setup, the DLR Control VM is deployed in High Availability (HA) as Active/Standby with each appliance installed on a separate ESX host and utilizing a different datastore on the EqualLogic PS6210SX iSCSI array.

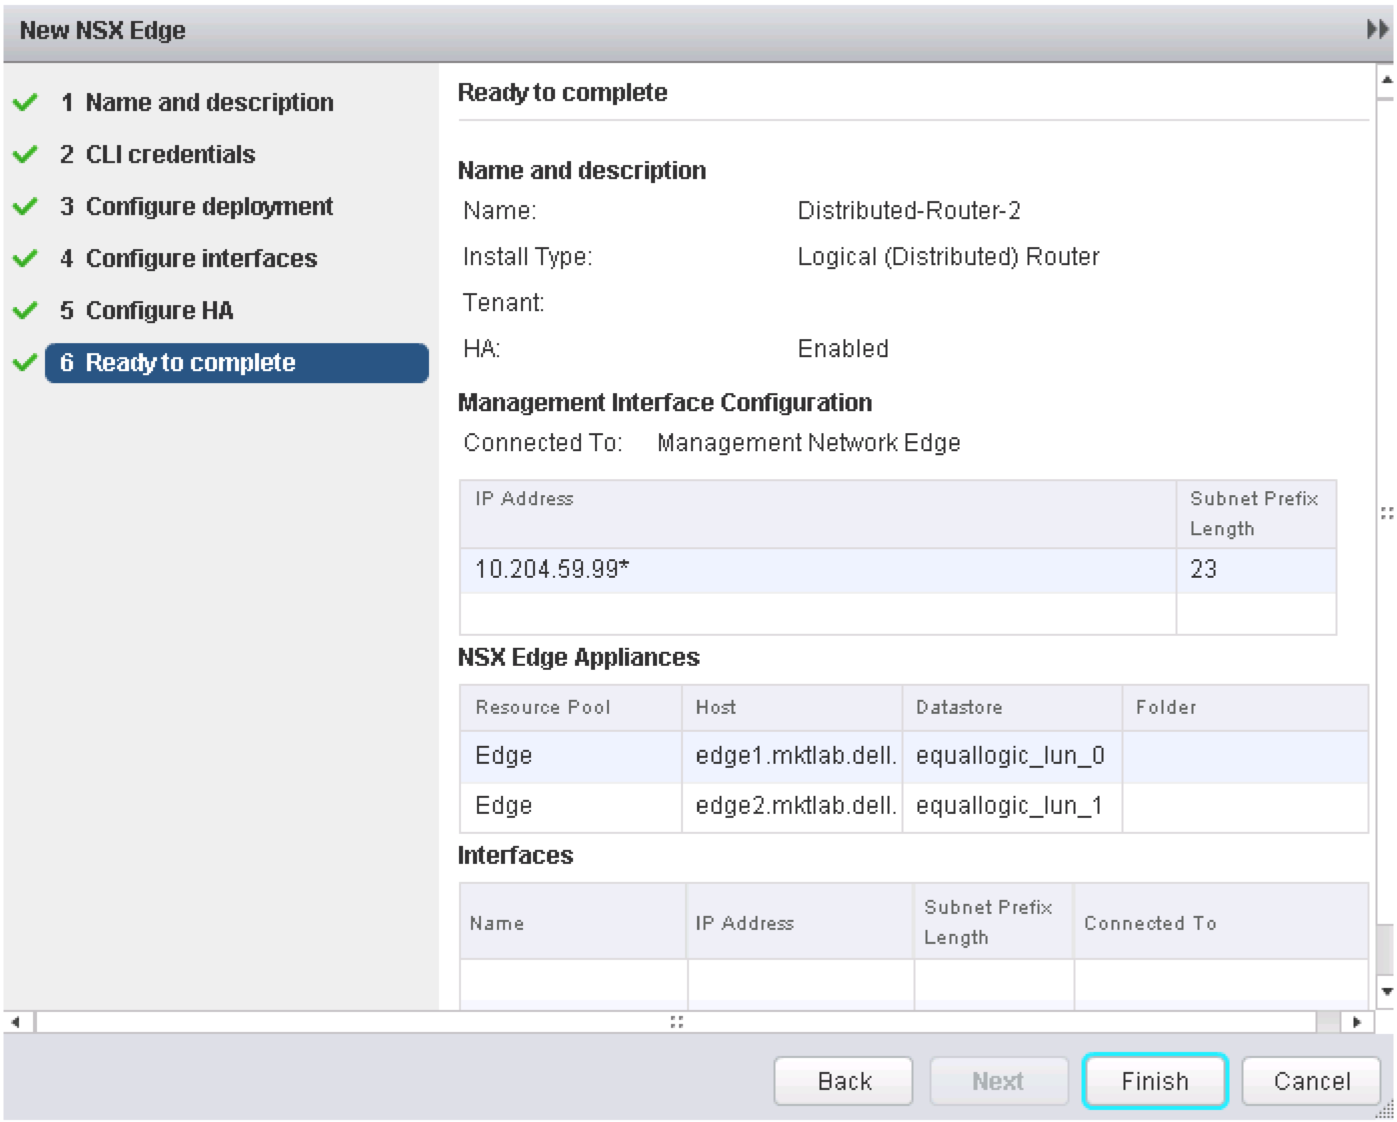

Figure 11 Final step of deploying a DLR for Tenant 2

In this example, when creating a DLR, no interfaces are initially added. However, additional interfaces are added to both the DLR and ESR later in accordance to the logical network diagram shown in Figure 12, The respective logical switches need connectivity to the respective routers for routing between subnets.

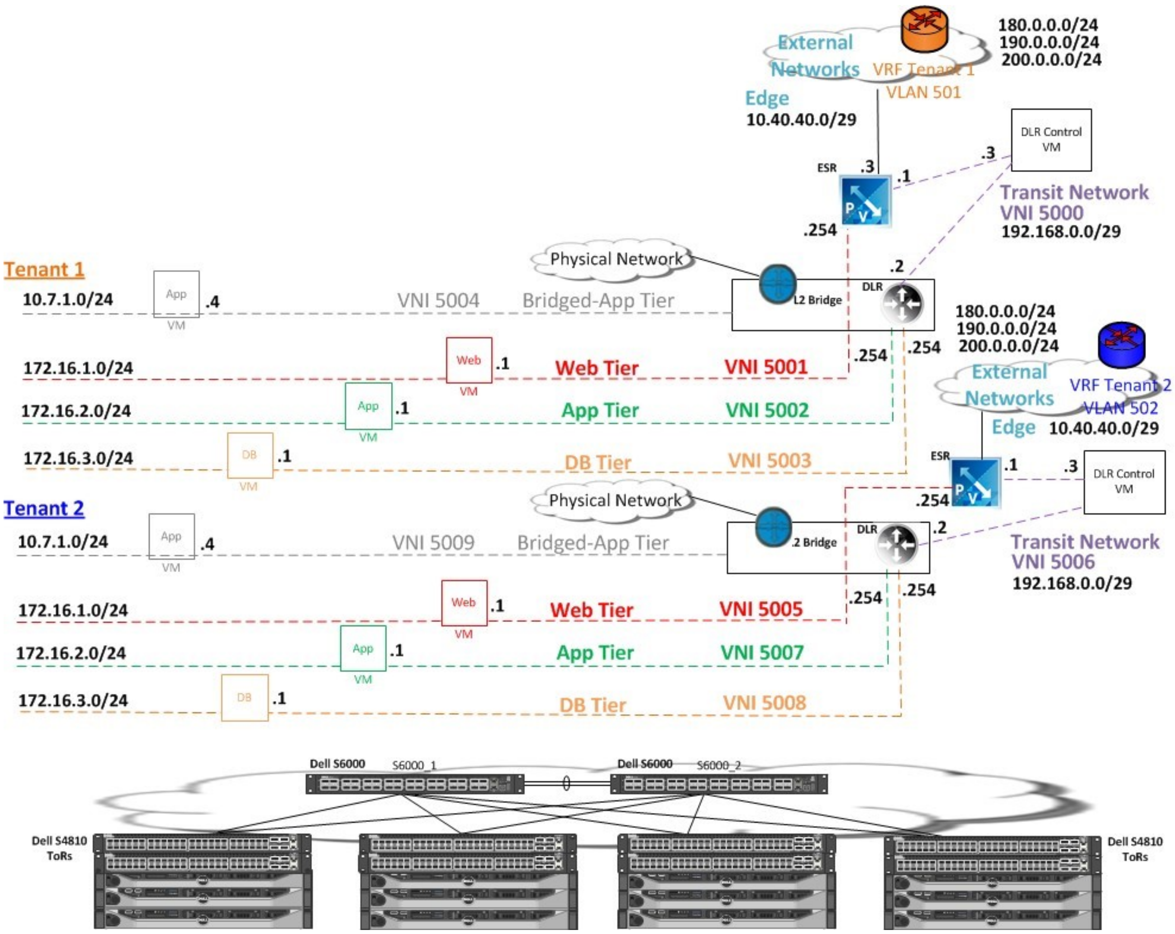

Figure 12 Logical network diagram of multitenancy across logical and physical networks

Figure 13 shows the respective logical switches which are created (for clarity, not all fields are shown).

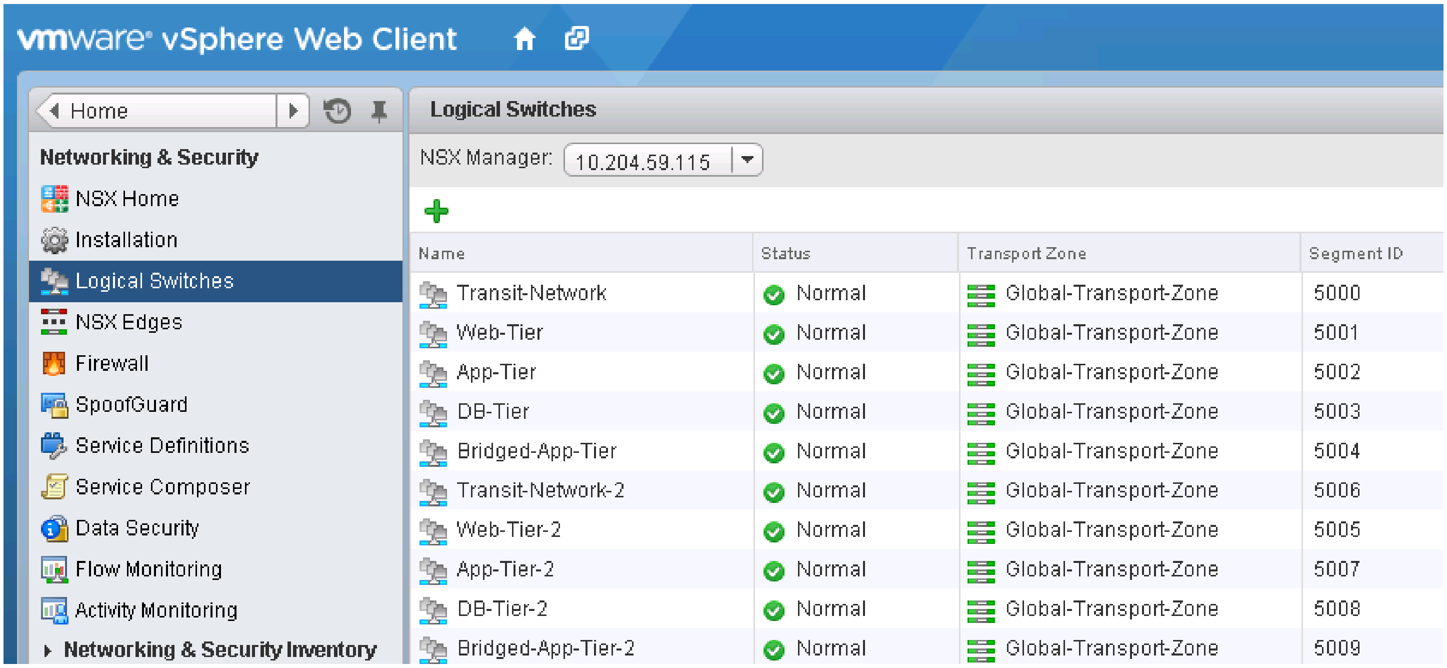

Figure 13 Logical switches created in VMware NSX for Tenant 1 and Tenant 2

Each tenant has its own set of logical switches. The Tenant 2 switches are identified by the 2 at the end of their names (Ex:. Web-Tier-2, App-Tier-2 etc.). The Tenant 1 switch names do not end with a number.

The logical switches are then connected to the respective tenant NSX DLR/ESR routers. Figure 14 shows the DLR and ESR routers created for Tenant 1 and Tenant 2. Note, the screenshot in Figure 14 displays the DLR Control VM, as the DLR is a kernel-level module on the respective hosts.

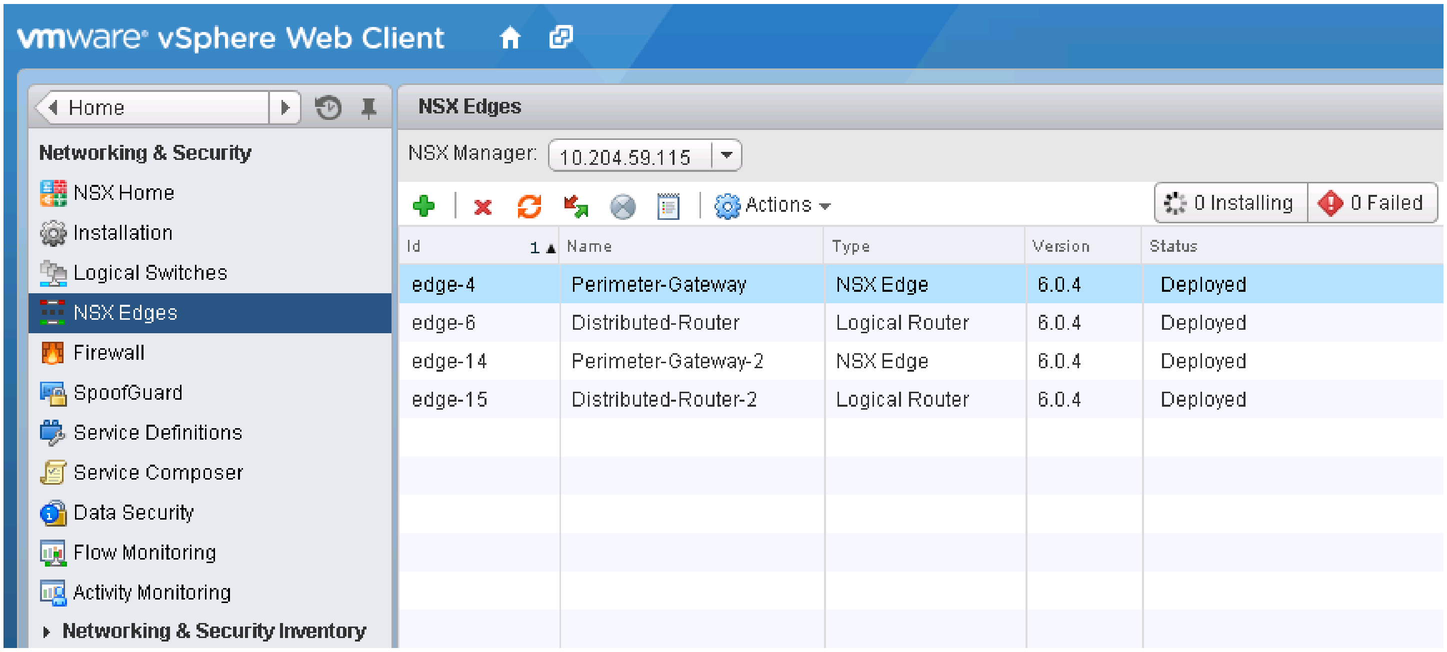

Figure 14 Logical routers (ESR and DLR Control VM) created in VMware NSX for Tenant 1 and Tenant 2

The ESR created for each tenant peers with the external network and learns external routes which are then learned by the DLR. In this setup OSPF is being utilized; OSPF can be configured for both the DLR and ESR under the routing tab for each respective appliance. Figure 15 shows OSPF being configured for the Tenant 2 ESR.

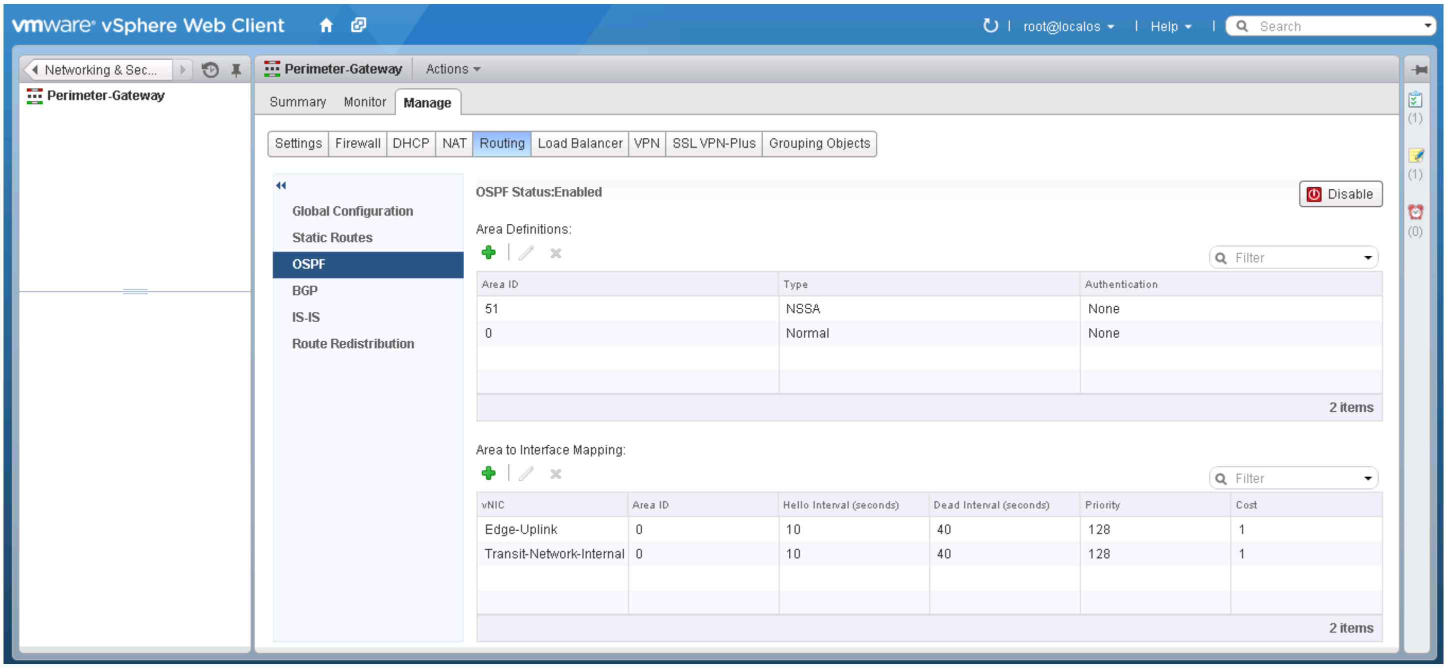

Figure 15 Configuring OSPF on the ESR for Tenant 1

This completes the details on setting up a multitenant environment in the logical environment with VMware NSX. Next, we look at the details of setting-up multitenancy on the physical network.

4.2 Creating Multitenancy in the Physical Environment

VRFs are used on the Dell Networking physical L3 switches/routers to allow for Multitenancy. Assuming servers are connected to ToR S4810s at L2 and all L3 and routing is done at the spine/core-level Dell S6000 switches, VLANs can be trunked up to the S6000s and a VR for each tenant only needs to be created on the S6000s.

Loopback interfaces are used in this example to simulate physical devices that have overlapping IP addresses on the network for each tenant. The network addresses for these loopback interfaces that OSPF will advertise are: 180.0.0.0/24, 190.0.0.0/24, and 200.0.0.0/24

Assuming the respective VLANs are already created and the interfaces connecting to the tenant-specific physical resources are already tagged with the corresponding tenant VLANs, the commands for configuring the VRs and routing on the S6000s for Tenant 1 and Tenant 2 as shown in Figure 12 are displayed below.

Tenant 2 setup is similar to Tenant 1 configuration except a different VLAN and VR is used; the IP addresses are overlapping and remain consistent.

4.2.1 S6000_1 Configuration

```
enable
configure
feature vrf

ip vrf tenant 1
exit

ip vrf tenant2 2
exit
```

Create loopback interfaces to imitate tenant-specific physical devices on the network and tie them to the respective VRF for each tenant. Note, the interfaces for the respective tenants have overlapping IP addresses.

```
interface Loopback 0
ip vrf forwarding tenant
ip address 180.0.0.1/24
no shutdown
exit

interface Loopback 1
ip vrf forwarding tenant
ip address 190.0.0.1/24
no shutdown
exit

interface Loopback 2
ip vrf forwarding tenant2
ip address 180.0.0.1/24
no shutdown
exit

interface Loopback 3
ip vrf forwarding tenant2
ip address 190.0.0.1/24
no shutdown
exit
```

Tie the tenant-specific VLANs to the respective VRFs and assign an IP address to the VLANs. Note, the VLANs for the respective tenants have overlapping IP addresses.

```
interface vlan 501
ip vrf forwarding tenant
ip address 10.40.40.1/29
no shutdown
exit


interface vlan 502
ip vrf forwarding tenant2
ip address 10.40.40.1/29
no shutdown
exit
```

Create a separate VRF instance for each tenant, tie it to the respective VRF, and add the networks for OSPF.

```
router ospf 1 vrf tenant
router-id 9.9.9.9
network 180.0.0.0/24 area 0
network 190.0.0.0/24 area 0
network 10.40.40.0/29 area 0
exit

router ospf 2 vrf tenant2
router-id 10.10.10.10
network 180.0.0.0/24 area 0
network 190.0.0.0/24 area 0
network 10.40.40.0/29 area 0
exit
write
```

Next, the peer S6000 switch is configured with similar steps. For demonstration purposes, this S6000 advertises a different network via OSPF. When it's validated that the NSX virtual routers have learned the external physical networks, the 180.0.0.0/24 and 190.0.0.0/24 networks on S6000_1 and the 200.0.0.0/24 network on S6000_2 should be learned.

4.2.2 S6000_2 Configuration

```
enable
configure
feature vrf

ip vrf tenant 1
exit

ip vrf tenant2 2
exit
```

Create loopback interfaces to simulate tenant-specific physical devices on the network and tie them to the respective VRF for each tenant. Note, the interfaces for the respective tenants have overlapping IP addresses.

```
interface Loopback 0
ip vrf forwarding tenant
ip address 200.0.0.1/24
no shutdown
exit

interface Loopback 1
ip vrf forwarding tenant2
ip address 200.0.0.1/24
no shutdown
exit
```

Tie the tenant-specific VLANs to the respective VRFs and assign an IP address to the VLANs. Note, the VLANs for the respective tenants have overlapping IP addresses.

```
interface vlan 501
ip vrf forwarding tenant
ip address 10.40.40.2/29
no shutdown
exit

interface vlan 502
ip vrf forwarding tenant2
ip address 10.40.40.2/29
no shutdown
exit
```

Create a separate VRF instance for each tenant, tie it to the respective VRF, and add the networks for OSPF.

```
router ospf 1 vrf tenant
router-id 9.9.9.8
network 200.0.0.0/24 area 0
network 10.40.40.0/29 area 0
exit

router ospf 2 vrf tenant2
router-id  10.10.10.9
network 200.0.0.0/24 area 0
network 10.40.40.0/29 area 0
exit
write
```

This completes the details on setting up a multitenant environment in the physical environment with VRF. Next, it's validated that the Tenant 1 and Tenant 2 ESRs and DLRs are learning the respective physical networks and tenant communication across the logical and physical network works.

5 Test and Validate Multitenancy Across Logical and Physical Networks

By using the vSphere Web Client to console on to the ESR and DLR routers for each tenant, it can be confirmed that the virtual routers in NSX are learning the physical networks for each tenant and multitenancy across the logical and physical network has been configured correctly.

5.1 Validating Tenant 1 Configuration

The **show ip route ospf** command is used on the Tenant 1 NSX ESR to confirm the tenant's physical networks were learned (Figure 16). The **ping** command is also used to confirm connectivity. Note, the App and DB logical networks connected to the DLR are also learned as OSPF E2 routes; this is because redistribution of connected routes was configured on the DLR.

Perimeter-Gateway-0

```
vShield-edge-4-0> show ip route ospf

Codes: O - OSPF derived, i - IS-IS derived, B - BGP derived,
C - connected, S - static, L1 - IS-IS level-1, L2 - IS-IS level-2,
IA - OSPF inter area, E1 - OSPF external type 1, E2 - OSPF external type 2


O    E2   10.204.58.0/23        [110/1]        via 192.168.0.2
O    E2   172.16.2.0/24         [110/1]        via 192.168.0.2
O    E2   172.16.3.0/24         [110/1]        via 192.168.0.2
O         180.0.0.1/32          [30/1]         via 10.40.40.1
O         190.0.0.1/32          [30/1]         via 10.40.40.1
O         200.0.0.1/32          [30/1]         via 10.40.40.2
vShield-edge-4-0>
vShield-edge-4-0> ping 180.0.0.1
PING 180.0.0.1 (180.0.0.1) 56(84) bytes of data.
64 bytes from 180.0.0.1: icmp_seq=1 ttl=255 time=0.583 ms
64 bytes from 180.0.0.1: icmp_seq=2 ttl=255 time=0.495 ms
64 bytes from 180.0.0.1: icmp_seq=3 ttl=255 time=0.579 ms
^C
--- 180.0.0.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 0.495/0.552/0.583/0.044 ms
vShield-edge-4-0>
vShield-edge-4-0> _
```

Figure 16 Validation that the Tenant 1 ESR learned the physical networks via OSPF

The same test is used to confirm configuration and connectivity on the Tenant 1 DLR (Figure 17).

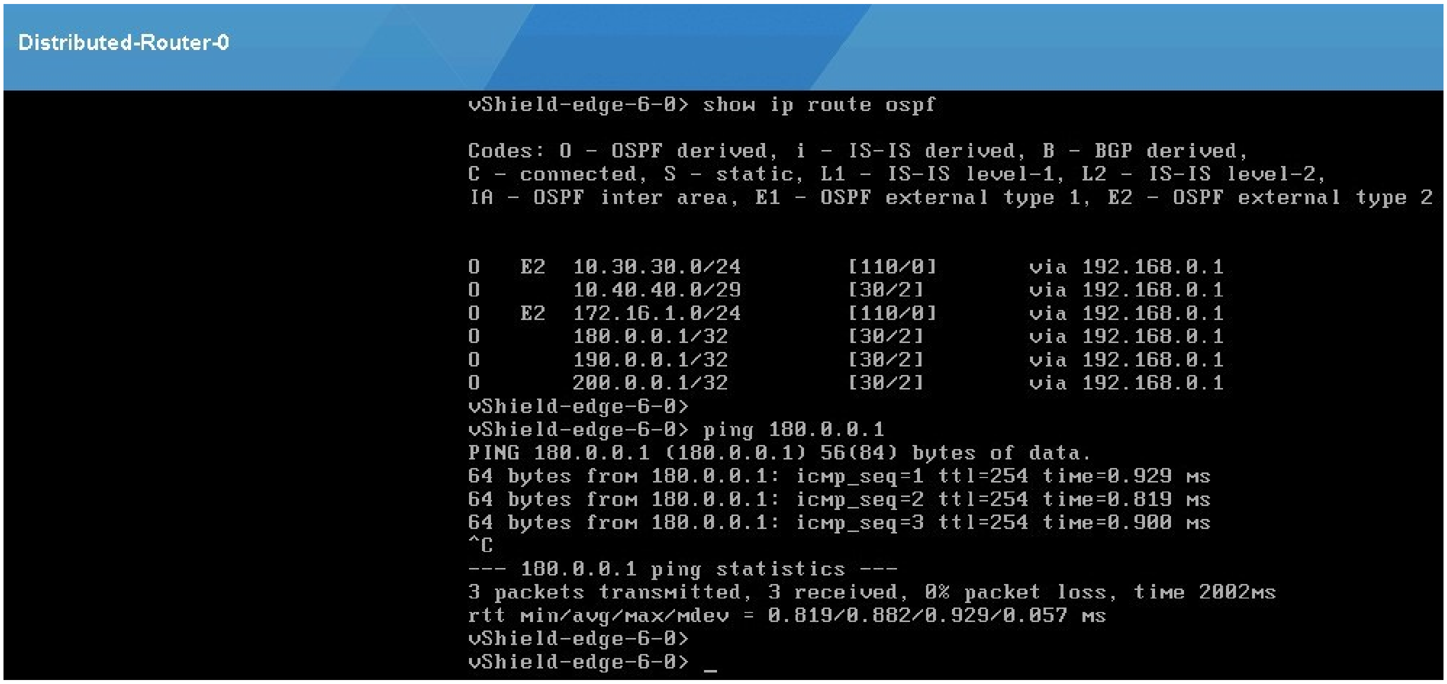

Figure 17 Validation that the Tenant 1 DLR learned the physical networks via OSPF

5.2 Validating Tenant 2 Configuration

The **show ip route ospf** command is used on the Tenant 2 NSX ESR to confirm the tenant's physical networks were learned (Figure 18). The **ping** command is also used to confirm connectivity. Note, the App and DB logical networks connected to the DLR are also learned as OSPF E2 routes; this is because redistribution of connected routes was configured on the DLR.

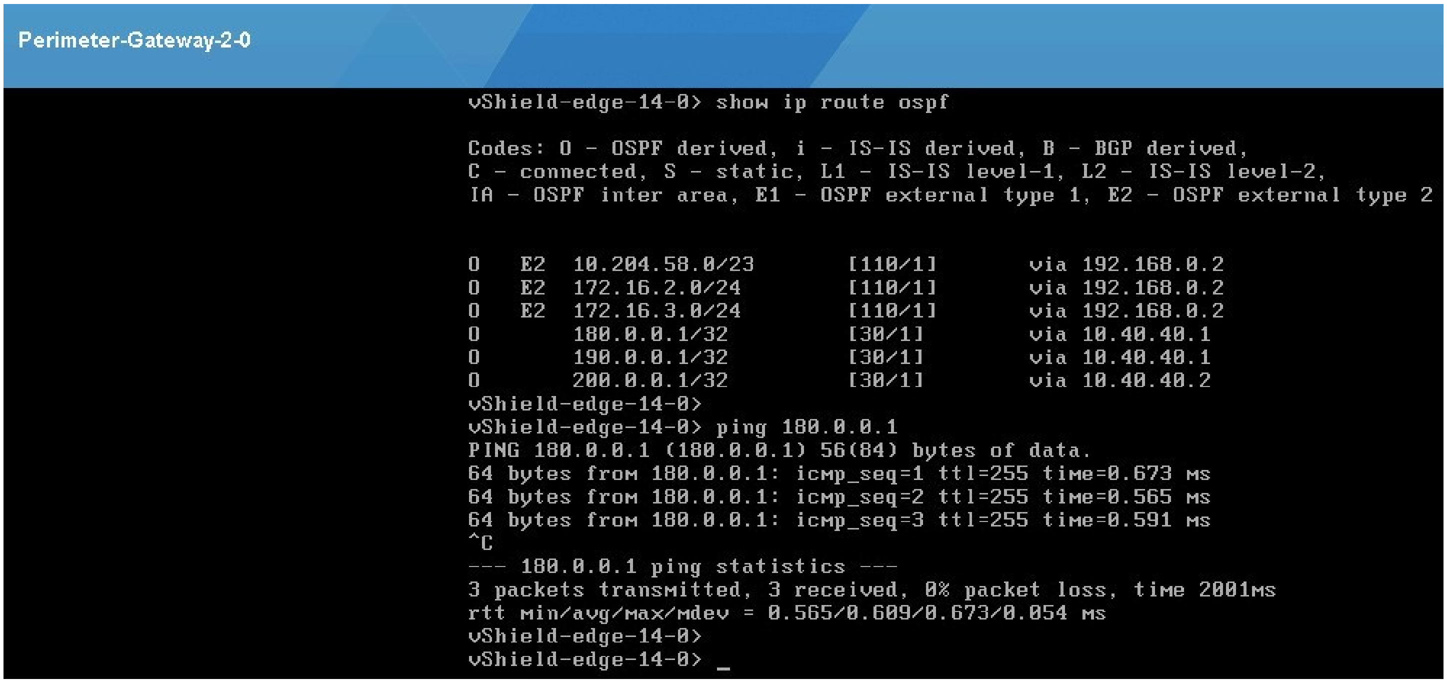

Figure 18 Validation that the Tenant 2 ESR learned the physical networks via OSPF

The same test is used to confirm configuration and connectivity on the Tenant 2 DLR (Figure 19).

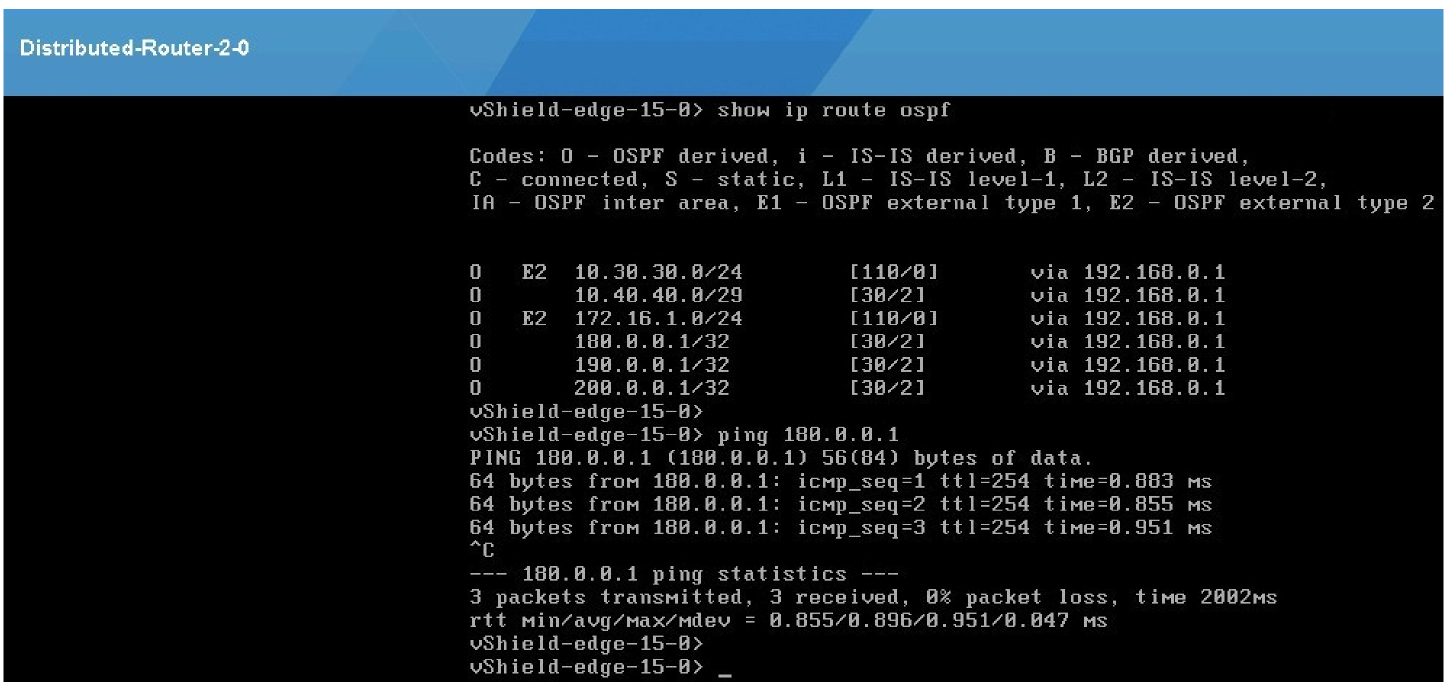

Figure 19 Validation that the Tenant 2 DLR learned the physical networks via OSPF

6 Conclusion

In this white paper, we looked at use cases of creating a multitenant environment across logical and physical networks and how this can be accomplished with VMware NSX and Dell Networking.

VMware NSX provides a robust and flexible infrastructure for network virtualization and multitenancy in the logical environment, and Dell Networking provides VRF-lite to allow for multitenancy on the physical network. Together, with VMware NSX utilizing NVO and Dell Networking utilizing VRF-lite, it’s possible to deploy a consistent multitenant framework across logical and physical networks.