



# Dell Networking: Multitenancy with VRF-lite

## A Dell Deployment and Configuration Guide

Dell Networking Technical Marketing  
October 2014



## Revisions

Date	Description	Authors
10/23/14	Version 2.0	Humair Ahmed, Sr. Technical Marketing Engineer, Dell Networking
4/14/14	Version 1 - Initial release	Humair Ahmed, Sr. Technical Marketing Engineer, Dell Networking

©2014 Dell Inc., All rights reserved.

Except as stated below, no part of this document may be reproduced, distributed or transmitted in any form or by any means, without express permission of Dell.

You may distribute this document within your company or organization only, without alteration of its contents.

THIS DOCUMENT IS PROVIDED "AS-IS", AND WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED. IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE SPECIFICALLY DISCLAIMED. PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT:

<http://www.dell.com/learn/us/en/19/terms-of-sale-commercial-and-public-sector> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of EMC Corporation. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Qlogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



# Table of contents

Revisions.....	2
1 Overview.....	4
2 MultiTenancy and Use Cases with VRF-lite .....	4
3 VRF-lite Example Lab Setup.....	8
4 S6000 VRF-lite Configuration .....	8
4.1 Top S6000 Configuration.....	8
4.2 Bottom S6000 Configuration .....	11
5 Test and Validate configuration and connectivity per VRF instance .....	13
6 Configure and Validate Route Leaking .....	17
7 Conclusion.....	22



# 1 Overview

The objective of this paper is to demonstrate multitenancy via VRF-lite on supported Dell Networking switches (S4810, S4820T, S5000, S6000, Z9500). Virtual Routing and Forwarding (VRF) allows a physical router to be partitioned into multiple Virtual Routers (VRs). The control and data plane are isolated in each VR so that traffic does not flow across VRs, thus allowing for multiple different routing tables to co-exist within the same physical L3 switch/router at the same time. Route leaking is also supported and allows for routes to 'leak' or be distributed across VRs in a controlled manner.

In this whitepaper, VRF-lite is configured and validated on Dell S6000 switches. This paper is not meant to provide a comprehensive overview of all options and features available for VRF-lite. For more information on additional features and options for VRF-lite on S6000 and other supported switches, see the respective [Configuration Guide](#). For more detailed information on using VRF-lite in conjunction with network overlays in virtualized environments, see the **Dell Networking: Multitenancy Across Physical and Logical Environments with VRF-lite and VMware NSX** white paper.

## Goals

- Configure VRF instances on Dell S6000 switches via VRF-lite CLI commands.
- Demonstrate multitenancy and route leaking.
- Validate configuration and verify connectivity.

## Resources Utilized

- 2 x Dell S6000 switches (Dell Networking OS 9.6)

# 2 MultiTenancy and Use Cases with VRF-lite

VRF-lite supports multitenancy by allowing multiple different routing tables within one physical L3 switch/router. Figure 1 below demonstrates how an ISP can leverage VRF-lite to provide separate isolated networks for customers without having to acquire and install additional hardware. Each customer below utilizes the same IP address space (overlapping IP addresses). Such a setup can also be useful for development and testing where traffic isolation is needed; the overlapping IP addresses can be used for ensuring consistency between each development/test environment.



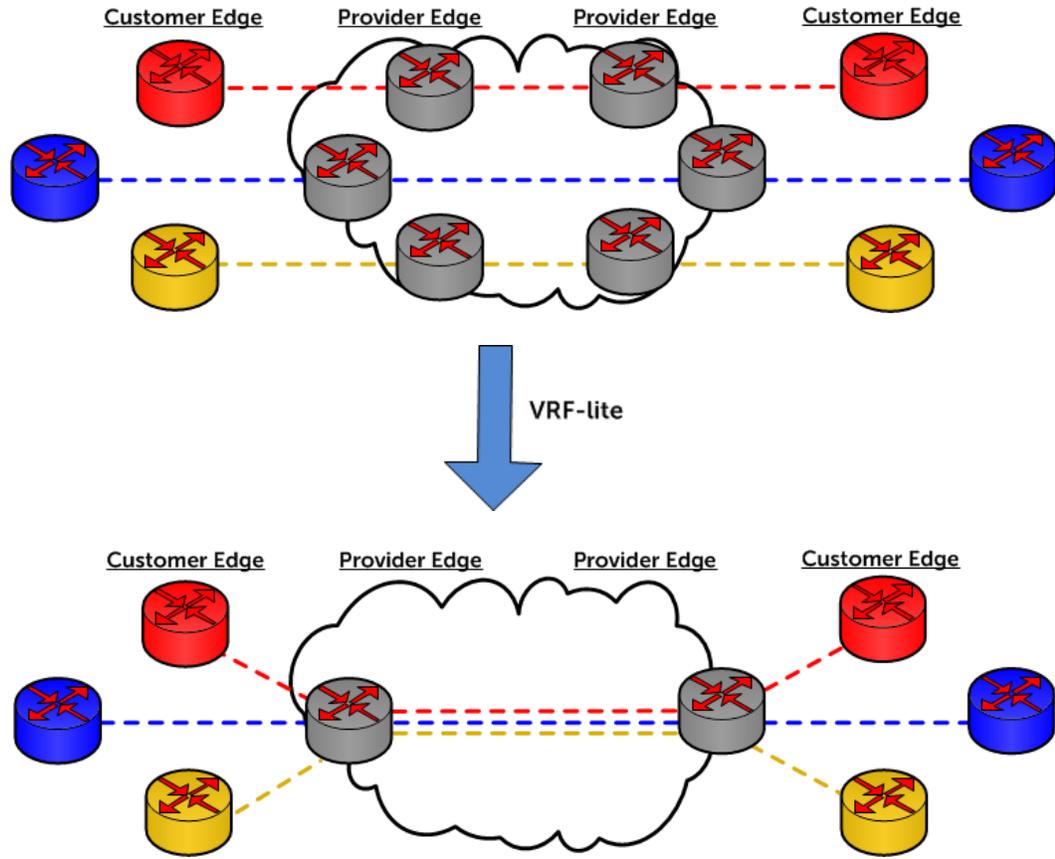


Figure 1 Example VRF-lite implementation

Another use case where VRF-lite can be incredibly helpful is 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. 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 will allow for multitenancy with overlapping IP addresses across both the logical and physical network/resources. Figure 2 presents an example where each VXLAN Network Identifier (VNI) on the logical network maps to a VLAN on the physical network. Each tenant logical switch (represented by a VNI) connects to a different distributed logical router (not shown) in the logical space and maps to a VLAN in the physical space.

Some example applications could be:

- 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.

For more detailed information on using VRF-lite in conjunction with network overlays in virtualized environments, see the **Dell Networking: Multitenancy Across Physical and Logical Environments with VRF-lite and VMware NSX** white paper.

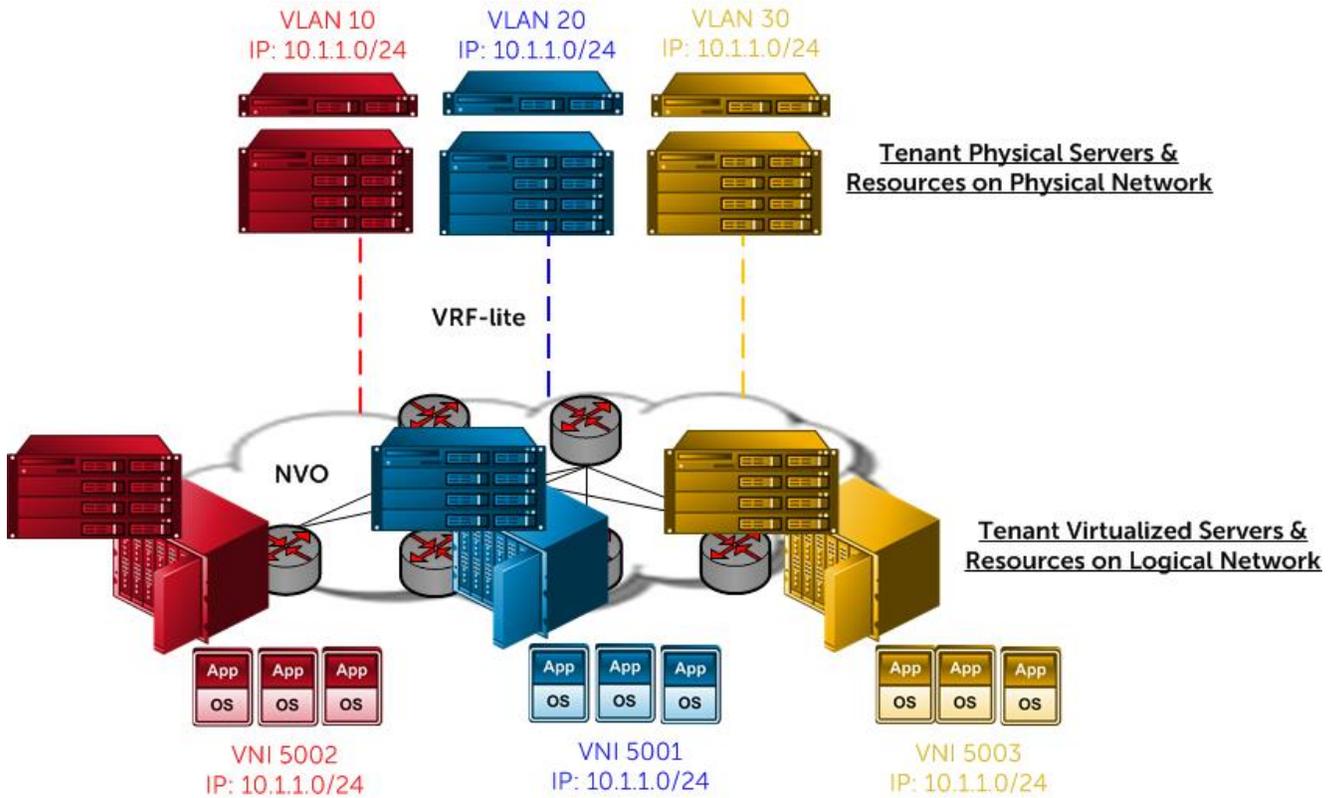


Figure 2 VRF-lite used in conjunction with VXLAN Network Overlay in a virtualized environment

Some important facts to be aware of in regards to VRF-lite are highlighted below.

- VRF-lite is a mechanism that keeps data segregated in a network by creating multiple logical routers (a routing table per VRF) within a physical router.
- VRF-lite allows network paths to be segmented (per domain or customer) without using multiple devices. This is accomplished by the ability to have multiple different routing tables co-exist within the same router at the same time.
- By automatically segregating traffic, VRF increases network security and can eliminate specific needs for encryption and authentication.

- IP addresses can be overlapped among different VRF instances with every interface being part of a VRF instance.
- Within a VRF instance, routing information can be statically provisioned and/or can be exchanged using dynamic routing protocols. All forwarding components are VRF aware and provisioned within the context of a VRF instance.
- Route leaking is supported across VRF Instances using static routes but not with dynamic protocol routes, meaning any static or dynamic route can only be leaked or distributed to another VRF via static route commands.
- VRF-lite also works seamlessly in virtualized environments. See the **Dell Networking: Multitenancy Across Physical and Logical Environments with VRF-lite and VMware NSX** white paper.
- The VRF-lite feature is supported with Dell's Virtual Link Trunking (VLT) multipathing technology.
- The VRF IDs of virtual routers running on a given physical switch are not exchanged with virtual routers running on other physical switches. Since VRF IDs are local to a router and no VRF information is exchanged between routers, VRF-lite only has local significance. Trunking of different tenant VLANs over the same link is possible since the VLANs will be unique; the L3 VLANs will belong to different routing instances specific to a tenant.
- VRRP or Proxy ARP used with VLT can be configured on interfaces that belong to a VRF instance. In a virtualized network that consists of multiple VRFs, various overlay networks can exist on a shared physical infrastructure. For more information see the **Dell Networking: Multitenancy Across Physical and Logical Environments with VRF-lite and VMware NSX** white paper.
- OSPFv2/OSPFv3, RIP, ISIS, BGP is supported both on physical and logical interfaces. Only OSPFv2 and BGP are supported on both default and non-default VRF ports; others are supported only on the default VRF ports.
- VRF supports route redistribution between different routing protocols (including static routes) only when the routes are within the same VRF.
- VRF supports up to 64 VRF instances: 1 to 63 and the default VRF (0).
- There is no restriction on the number of VLANs that can be assigned to a VRF instance.
- The total number of routes supported in VRF is limited by the size of the IPv4 CAM.
- Supported Platforms: S4810, S4820T, S5000, S6000, Z9500



### 3 VRF-lite Example Lab Setup

Figure 3 displays the lab setup that will be used in this white paper. Two Dell S6000 switches connected via a port-channel consisting of two 40 GbE ports are configured with VRF-lite. Figure 3 shows the three tenants in this setup with each color representing a different tenant. Each tenant also has unique VLANs but use the same IP addresses on their respective VLANs. A single port-channel connection between the S6000 switches is used to carry all the tenants' traffic.

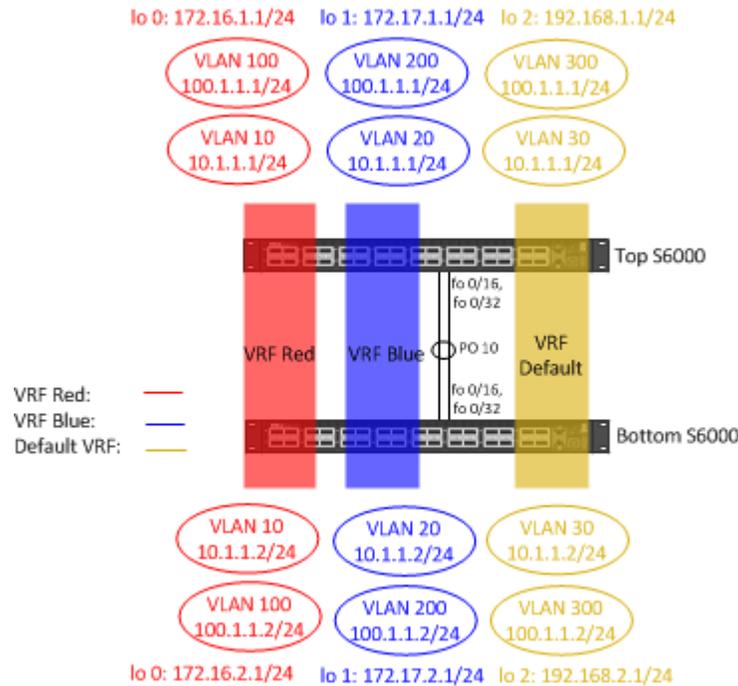


Figure 3 Lab setup for S6000 with VRF-lite and overlapping IP address scheme

## 4 S6000 VRF-lite Configuration

### 4.1 Top S6000 Configuration

For this setup, 2 x 40 GbE ports are used to create a port channel that connects to the bottom S6000 switch. The configuration entered is shown below in black; comments are in brackets. Note, when entering the configuration, the **Tab** key can be used for autocomplete.

```
enable
conf
hostname S6000_1
interface range fortyGigE 0/16 , fortyGigE 0/32
port-channel-protocol lacp
port-channel 10 mode active
```



```
no shutdown
exit
exit
```

Load VRF CAM profile and create two VRF instances.

```
feature vrf

ip vrf red 1
exit

ip vrf blue 2
exit
```

Add resources to the VRF Instances. Start with loopback Interfaces.

```
interface loopback 0 [Part of the Red VRF]
ip vrf forwarding red
ip address 172.16.1.1/24
exit

interface loopback 1 [Part of the Blue VRF]
ip vrf forwarding blue
ip address 172.17.1.1/24
exit

interface loopback 2 [Part of the Default VRF]
ip address 192.168.1.1/24
exit

interface port-channel 10
switchport
exit
```

Create VLANs for the VRFs. Note the duplicate IP addresses in the different VRF instances.

```
interface vlan 10 [Part of the Red VRF]
ip vrf forwarding red
ip address 10.1.1.1/24
tagged port-channel 10 [Trunk VLAN 10 over port-channel]
no shutdown
exit

interface vlan 100 [Part of the Red VRF]
ip vrf forwarding red
ip address 100.1.1.1/24
tagged port-channel 10 [Trunk VLAN 100 over port-channel]
```



```

no shutdown
exit

interface vlan 20 [Part of the Blue VRF]
ip vrf forwarding blue
ip address 10.1.1.1/24
tagged port-channel 10 [Trunk VLAN 20 over port-channel]
no shutdown
exit

interface vlan 200 [Part of the Blue VRF]
ip vrf forwarding blue
ip address 100.1.1.1/24
tagged port-channel 10 [Trunk VLAN 200 over port-channel]
no shutdown
exit

interface vlan 30 [Part of the Default VRF]
ip address 10.1.1.1/24
tagged port-channel 10 [Trunk VLAN 30 over port-channel]
no shutdown
exit

interface vlan 300 [Part of the Default VRF]
ip address 100.1.1.1/24
tagged port-channel 10 [Trunk VLAN 300 over port-channel]
no shutdown
exit

```

Add routing to the VRF instances.

```

router ospf 1 vrf red [Create OSPF instance on Red VRF]
router-id 100.1.1.10
log-adjacency-changes
network 10.1.1.0/24 area 0
network 100.1.1.0/24 area 0
network 172.16.1.1/24 area 0
exit

router ospf 2 vrf blue [Create OSPF instance on Blue VRF]
router-id 100.1.1.20
log-adjacency-changes
network 10.1.1.0/24 area 0
network 100.1.1.0/24 area 0
network 172.17.1.0/24 area 0
exit

router ospf 3 [Create OSPF instance on Default VRF]

```



```
router-id 100.1.1.30
log-adjacency-changes
network 10.1.1.0/24 area 0
network 100.1.1.0/24 area 0
network 192.168.1.0/24 area 0
end
write
```

This completes the configuration on the top S6000 switch. Next, configure the bottom S6000 switch.

## 4.2 Bottom S6000 Configuration

For this setup, 2 x 40 GbE ports are used to create a port channel that connects to the top S6000 switch. The configuration entered is shown below. Note, the **Tab** key can be used for autocomplete.

```
enable
conf
hostname S6000_2
interface range fortyGigE 0/16 , fortyGigE 0/32
port-channel-protocol lacp
port-channel 10 mode active
no shutdown
exit
exit
```

Load VRF CAM profile and create two VRF instances.

```
feature vrf

ip vrf red 1
exit

ip vrf blue 2
exit
```

Add resources to the VRF Instances. Start with loopback Interfaces.

```
interface loopback 0 [Part of the Red VRF]
ip vrf forwarding red
ip address 172.16.2.1/24
exit

interface loopback 1 [Part of the Blue VRF]
ip vrf forwarding blue
ip address 172.17.2.1/24
exit

interface loopback 2 [Part of the Default VRF]
```



```
ip address 192.168.2.1/24
exit

interface port-channel 10
switchport
exit
```

Create VLANs for the VRFs. Note the duplicate IP addresses on the different VRF instances.

```
interface vlan 10 [Part of the Red VRF]
ip vrf forwarding red
ip address 10.1.1.2/24
tagged port-channel 10 [Trunk VLAN 10 over port-channel]
no shutdown
exit

interface vlan 100 [Part of the Red VRF]
ip vrf forwarding red
ip address 100.1.1.2/24
tagged port-channel 10 [Trunk VLAN 100 over port-channel]
no shutdown
exit

interface vlan 20 [Part of the Blue VRF]
ip vrf forwarding blue
ip address 10.1.1.2/24
tagged port-channel 10 [Trunk VLAN 20 over port-channel]
no shutdown
exit

interface vlan 200 [Part of the Blue VRF]
ip vrf forwarding blue
ip address 100.1.1.2/24
tagged port-channel 10 [Trunk VLAN 200 over port-channel]
no shutdown
exit

interface vlan 30 [Part of the Default VRF]
ip address 10.1.1.2/24
tagged port-channel 10 [Trunk VLAN 30 over port-channel]
no shutdown
exit

interface vlan 300 [Part of the Default VRF]
ip address 100.1.1.2/24
tagged port-channel 10 [Trunk VLAN 300 over port-channel]
no shutdown
exit
```



Add routing to the VRF instances.

```
router ospf 1 vrf red [Create OSPF instance on Red VRF]
router-id 100.1.1.11
log-adjacency-changes
network 10.1.1.0/24 area 0
network 100.1.1.0/24 area 0
network 172.16.2.1/24 area 0
exit

router ospf 2 vrf blue [Create OSPF instance on Blue VRF]
router-id 100.1.1.21
log-adjacency-changes
network 10.1.1.0/24 area 0
network 100.1.1.0/24 area 0
network 172.17.2.0/24 area 0
exit

router ospf 3 [Create OSPF instance on Default VRF]
router-id 100.1.1.31
log-adjacency-changes
network 10.1.1.0/24 area 0
network 100.1.1.0/24 area 0
network 192.168.2.0/24 area 0
end
write
```

This completes the configuration for the bottom S6000 switch. Next, verify OSPF neighbor adjacencies on the switches, validate the respective VRF routing tables, and verify connectivity per VRF instance.

## 5 Test and Validate configuration and connectivity per VRF instance

While still on the bottom S6000 console, confirm OSPF neighbors for the Default, Red, and Blue VRF instances using the below commands.

```
show ip ospf neighbor
show ip ospf vrf red neighbor
show ip ospf vrf blue neighbor
```

Figure 4 below shows the output of these commands.



```

S6000_2#show ip ospf neighbor

Neighbor ID      Pri      State      Dead Time Address      Interface Area
100.1.1.30      1        FULL/BDR   00:00:37  10.1.1.1     Vl 30        0
100.1.1.30      1        FULL/BDR   00:00:37  100.1.1.1    Vl 300       0
S6000_2#
S6000_2#
S6000_2#show ip ospf vrf red neighbor

Neighbor ID      Pri      State      Dead Time Address      Interface Area
100.1.1.10      1        FULL/BDR   00:00:31  10.1.1.1     Vl 10        0
100.1.1.10      1        FULL/BDR   00:00:31  100.1.1.1    Vl 100       0
S6000_2#
S6000_2#
S6000_2#show ip ospf vrf blue neighbor

Neighbor ID      Pri      State      Dead Time Address      Interface Area
100.1.1.20      1        FULL/BDR   00:00:39  10.1.1.1     Vl 20        0
100.1.1.20      1        FULL/BDR   00:00:39  100.1.1.1    Vl 200       0
S6000_2#

```

Figure 4 Displaying OSPF neighbors for both VRF instances

View the routing tables for the Default, Red, and Blue VRF instances. Note that the VRF routing tables have some overlapping subnet entries.

```

show ip route
show ip route vrf red
show ip route vrf blue

```

Figure 5 below shows the output of these commands.



```

S6000_2#show ip route

Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route

Gateway of last resort is not set

      Destination         Gateway                Dist/Metric  Last Change
      -----
C     10.1.1.0/24         Direct, V1 30          0/0         00:05:04
C     100.1.1.0/24       Direct, V1 300         0/0         00:05:04
O     192.168.1.1/32     via 10.1.1.1, V1 30   110/1       00:03:59
                        via 100.1.1.1, V1 300
C     192.168.2.0/24     Direct, Lo 2           0/0         00:25:59
S6000_2#
S6000_2#
S6000_2#show ip route vrf red

Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route

Gateway of last resort is not set

      Destination         Gateway                Dist/Metric  Last Change
      -----
C     10.1.1.0/24         Direct, V1 10          0/0         00:05:14
C     100.1.1.0/24       Direct, V1 100         0/0         00:05:14
O     172.16.1.1/32     via 10.1.1.1, V1 10   110/1       00:04:20
                        via 100.1.1.1, V1 100
C     172.16.2.0/24     Direct, Lo 0           0/0         00:26:09
S6000_2#
S6000_2#
S6000_2#show ip route vrf blue

Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route

Gateway of last resort is not set

      Destination         Gateway                Dist/Metric  Last Change
      -----
C     10.1.1.0/24         Direct, V1 20          0/0         00:05:28
C     100.1.1.0/24       Direct, V1 200         0/0         00:05:28
O     172.17.1.1/32     via 10.1.1.1, V1 20   110/1       00:04:23
                        via 100.1.1.1, V1 200
C     172.17.2.0/24     Direct, Lo 1           0/0         00:26:23
S6000_2#

```

Figure 5 Displaying routing tables for Default, Red , and Blue VRF instances



Use the **show ip interface brief vlan <vlan #>** command to observe the overlapping IPs on different VLANs on the same switch.

```
show ip interface brief vlan 10
show ip interface brief vlan 100
show ip interface brief vlan 20
show ip interface brief vlan 200
show ip interface brief vlan 30
show ip interface brief vlan 300
```

Figure 6 below shows the output of these commands.

```
S6000_2#show ip interface brief vlan 10
Interface          IP-Address      OK Method Status      Protocol
Vlan 10            10.1.1.2       YES Manual up        up
S6000_2#
S6000_2#show ip interface brief vlan 100
Interface          IP-Address      OK Method Status      Protocol
Vlan 100           100.1.1.2      YES Manual up        up
S6000_2#
S6000_2#show ip interface brief vlan 20
Interface          IP-Address      OK Method Status      Protocol
Vlan 20            10.1.1.2       YES Manual up        up
S6000_2#
S6000_2#show ip interface brief vlan 200
Interface          IP-Address      OK Method Status      Protocol
Vlan 200           100.1.1.2      YES Manual up        up
S6000_2#
S6000_2#show ip interface brief vlan 30
Interface          IP-Address      OK Method Status      Protocol
Vlan 30            10.1.1.2       YES Manual up        up
S6000_2#
S6000_2#show ip interface brief vlan 300
Interface          IP-Address      OK Method Status      Protocol
Vlan 300           100.1.1.2      YES Manual up        up
S6000_2#
```

Figure 6 Verifying overlapping IP addresses on different VLANs on the same switch

Use the **ping** command to verify connectivity for each VRF instance.

```
ping 192.168.1.1
ping vrf red 172.16.1.1 source ip 172.16.2.1
ping vrf blue 172.17.1.1 source ip 172.17.2.1
```

Figure 7 below shows the output of these commands.



```

S6000_2#ping 192.168.1.1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/0/0 (ms)
S6000_2#
S6000_2#
S6000_2#ping vrf red 172.16.1.1 source ip 172.16.2.1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 172.16.1.1 from 172.16.2.1, timeout is 2 seconds:
!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/0/0 (ms)
S6000_2#
S6000_2#
S6000_2#ping vrf blue 172.17.1.1 source ip 172.17.2.1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 172.17.1.1 from 172.17.2.1, timeout is 2 seconds:
!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/4/20 (ms)
S6000_2#

```

Figure 7 Verifying connectivity for each VRF instance via the 'ping' command

## 6 Configure and Validate Route Leaking

It is possible to redistribute routes between a non-default to default/non-default VRF and vice-versa. This is termed **route leaking** and can be accomplished via static routes. Route leaking will be demonstrated here.

On the top S6000 switch the 172.16.2.0/24 network in the red tenant will be redistributed into the routing table of the blue tenant, and the 172.17.2.0/24 network in the blue tenant will be redistributed into the routing table of the red tenant.

On the bottom S6000 switch the 172.16.1.0/24 network in the red tenant will be redistributed into the routing table of the blue tenant, and the 172.17.1.0/24 network in the blue tenant will be redistributed into the routing table of the red tenant.

Once complete, the redistribution can be validated by taking a look at the routing table of each tenant and using the **ping** command.

First, enter the following commands on the top S6000 switch.



```
enable
config
ip route vrf red 172.17.2.0/24 172.17.2.1 vrf blue
ip route vrf blue 172.16.2.0/24 172.16.2.1 vrf red
end
write
```

Next, enter the following commands on the bottom S6000 switch.

```
enable
config
ip route vrf red 172.17.1.0/24 172.17.1.1 vrf blue
ip route vrf blue 172.16.1.0/24 172.16.1.1 vrf red
end
write
```

Finally, validate the redistribution by looking at the routing tables (Figure 8). The new routing tables can be compared to the original routing tables to confirm the changes. The **172.17.1.0/24** blue tenant network is now reachable by the red tenant and the **172.16.1.0/24** red tenant network is now reachable by the blue tenant.



```
S6000_2#show ip route vrf red
```

```
Codes: C - connected, S - static, R - RIP,  
B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,  
O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,  
N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,  
E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,  
L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,  
> - non-active route, + - summary route
```

```
Gateway of last resort is not set
```

	Destination	Gateway	Dist/Metric	Last Change
C	10.1.1.0/24	Direct, V1 10	0/0	00:13:22
C	100.1.1.0/24	Direct, V1 100	0/0	00:13:22
O	172.16.1.1/32	via 10.1.1.1, V1 10 via 100.1.1.1, V1 100	110/1	00:12:27
C	172.16.2.0/24	Direct, Lo 0	0/0	00:34:17
S	172.17.1.0/24	via blue:172.17.1.1	1/0	00:00:28

```
S6000_2#
```

```
S6000_2#
```

```
S6000_2#show ip route vrf blue
```

```
Codes: C - connected, S - static, R - RIP,  
B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,  
O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,  
N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,  
E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,  
L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,  
> - non-active route, + - summary route
```

```
Gateway of last resort is not set
```

	Destination	Gateway	Dist/Metric	Last Change
C	10.1.1.0/24	Direct, V1 20	0/0	00:17:47
C	100.1.1.0/24	Direct, V1 200	0/0	00:17:47
S	172.16.1.0/24	via red:172.16.1.1	1/0	00:04:43
O	172.17.1.1/32	via 10.1.1.1, V1 20 via 100.1.1.1, V1 200	110/1	00:16:42
C	172.17.2.0/24	Direct, Lo 1	0/0	00:38:42

```
S6000_2#
```

Figure 8 Updated routing tables on the bottom S6000 after route leaking configuration



Figure 9 below shows similar respective results on the top S6000 switch.

```
S6000_1#show ip route vrf red

Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route

Gateway of last resort is not set

      Destination            Gateway                      Dist/Metric  Last Change
      -----            -
C     10.1.1.0/24          Direct, V1 10                0/0         00:21:49
C     100.1.1.0/24        Direct, V1 100               0/0         00:21:49
C     172.16.1.0/24       Direct, Lo 0                 0/0         00:21:52
O     172.16.2.1/32       via 10.1.1.2, V1 10          110/1       00:20:50
                               via 100.1.1.2, V1 100
S     172.17.2.0/24       via blue:172.17.2.1         1/0         00:09:36
S6000_1#
S6000_1#
S6000_1#show ip route vrf blue

Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route

Gateway of last resort is not set

      Destination            Gateway                      Dist/Metric  Last Change
      -----            -
C     10.1.1.0/24          Direct, V1 20                0/0         00:22:05
C     100.1.1.0/24        Direct, V1 200               0/0         00:22:05
S     172.16.2.0/24       via red:172.16.2.1          1/0         00:09:41
C     172.17.1.0/24       Direct, Lo 1                 0/0         00:22:08
O     172.17.2.1/32       via 10.1.1.2, V1 20          110/1       00:21:05
                               via 100.1.1.2, V1 200
S6000_1#
```

Figure 9 Updated routing tables on the top S6000 after route leaking configuration



Below in Figure 10 and 11, route leaking is validated by using the **ping** command to confirm reachability across tenants.

```
S6000_2#ping vrf red 172.17.1.1 source ip 172.16.2.1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 172.17.1.1 from 172.16.2.1, timeout is 2 seconds:
!!!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/4/20 (ms)
S6000_2#
S6000_2#
S6000_2#ping vrf blue 172.16.1.1 source ip 172.17.2.1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 172.16.1.1 from 172.17.2.1, timeout is 2 seconds:
!!!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/0/0 (ms)
S6000_2#
```

Figure 10 Validating route leaking on bottom S6000 using the 'ping' command

```
S6000_1#ping vrf red 172.17.2.1 source ip 172.16.1.1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 172.17.2.1 from 172.16.1.1, timeout is 2 seconds:
!!!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/0/0 (ms)
S6000_1#
S6000_1#
S6000_1#ping vrf blue 172.16.2.1 source ip 172.17.1.1

Type Ctrl-C to abort.

Sending 5, 100-byte ICMP Echos to 172.16.2.1 from 172.17.1.1, timeout is 2 seconds:
!!!!!!
Success rate is 100.0 percent (5/5), round-trip min/avg/max = 0/0/0 (ms)
S6000_1#
```

Figure 11 Validating route leaking on top S6000 using the 'ping' command



## 7 Conclusion

In this white paper, we looked at the benefits of VRF-lite and how it can be implemented on Dell S6000 switches to achieve multitenancy. For more general information on additional features and options for VRF-lite on S6000 and other supported switches see the respective [Configuration Guide](#). For more detailed information on using VRF-lite in conjunction with network overlays in virtualized environments, see the **Dell Networking: Multitenancy Across Physical and Logical Environments with VRF-lite and VMware NSX** white paper.

