Dell EMC VxRail Multirack Deployment Guide

Multirack deployment of VxRail cluster using Network Virtualization Overlays (NVO)

Abstract

This document provides step-by-step deployment instructions for Dell EMC Networking OS10EE static VXLAN tunnels in an OSPF routed environment. This provides the foundation that is needed for multirack VxRail host discovery and deployment in a modern data center.

March 2019

Revisions

Date	Description
January 2019	Initial publication
February 2019	FDC addition, vSAN stretched cluster example
March 2019	Changed switch interface setting from “flowcontrol transmit on” to “flowcontrol transmit off” as a best practice.

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions 2
1 Introduction 6
1.1 Objective 7
1.2 Fabric Design Center (FDC) 7
1.3 Supported switches and operating systems 8
1.4 Typographical conventions 8
1.5 Attachments 8
2 Hardware overview 9
2.1 Dell EMC VxRail P570 9
2.2 Dell EMC VxRail E560 9
2.3 Dell EMC Networking Z9264F-ON 10
2.4 Dell EMC Networking S3048-ON 10
3 Topology 11
3.1 VxRail node connectivity 11
3.2 Underlay network design 13
3.3 Virtual extensible LAN (VXLAN) overlay 15
3.4 External vCenter connectivity 16
4 Switch configuration values and prerequisites 17
4.1 Pre-planning VLANs, IP address, and switch values 17
4.1.1 VLANs and IP addresses 17
4.1.2 Switch settings 18
4.2 Switch prerequisites 18
4.2.1 Verify OS10EE version 18
4.2.2 Verify license installation 19
4.2.3 Factory default configuration 19
5 Configure switches 20
5.1 Configure general settings 20
5.2 Configure OSPF routing and upstream network-facing ports 21
5.3 Configure Virtual Link Trunk (VLT) 22
5.4 Configure Uplink Failure Detection (UFD) 22
5.5 Configure Virtual Extensible LANs (VXLAN) 23
5.6 Configure Virtual LANs (VLAN) 25
5.7 Configure downstream VxRail node interfaces 26
6 Switch validation 28
6.1 General validation commands 28

6.1.1 show interface status 28
6.1.2 show ip interface brief 28
6.1.3 show lldp neighbors 29
6.1.4 show interface 29
6.2 OSPF validation commands 30
6.2.1 show ip ospf neighbor 30
6.2.2 show ip route ospf 30
6.2.3 show ip ospf topology 31
6.3 VXLAN validation commands 32
6.3.1 show vlan 32
6.3.2 show nve remote-vtep 32
6.3.3 show nve vxlan-vni 33
6.3.4 show virtual-network 33
6.4 VLT validation commands 34
6.4.1 show vlt all 34
6.4.2 show vlt all backup-link 34
6.4.3 show vlt all mismatch 35
6.4.4 show vlt mac-inconsistency 36
7 Perform initialization to create VxRail cluster 37
7.1 VxRail initialization 37
7.2 VxRail validation 38
8 VMware vSAN stretched clusters 39
8.1 Deploying VMware vSAN stretched clusters 40
8.1.1 Deploy the vSAN witness appliance 40
8.1.2 Configure vSAN witness appliance routing 41
8.1.3 Configure network interfaces for witness traffic 42
8.1.4 Configure fault domains 43
8.2 Witness component validation 44
A Routing in a Layer 2 VXLAN overlay 45
B Maximum Transmission Unit (MTU) considerations 47
C VxRail deployment values 48
D Connecting a workstation or laptop for VxRail initialization 49
E Validated components 51
E.1 Dell EMC Networking switches 51
E.2 VxRail E560 nodes 51
E.3 VxRail P570 nodes 51

E.4 VxRail appliance software 52
F Technical resources 53
G Support and feedback 54

1 Introduction

VxRail appliances have no backplane, therefore, communication between its nodes is facilitated using the network switches. Communication between the nodes uses auto-discovery capabilities. New VxRail nodes advertise themselves on the network and are discovered by the VxRail Manager.

Modern data centers commonly use a routed IP environment that is based on either the Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP) routing protocols. In these environments, each rack is a unique IP subnet. For a successful VxRail multirack deployment, all nodes must be able to reach each other through a single Layer 2 (L2) domain. Network Virtualization solves this problem by carving a single physical network (underlay) into multiple virtual networks (overlays), or Network Virtualization Overlays (NVOs). The standards-based protocol used to create NVOs is Virtual Extensible LAN (VXLAN). VXLAN based solutions offer one of the most cost-effective and straightforward paths to enable the routed underlay to forward L2 traffic between separate subnets.

Figure 1 shows Dell EMC Networking switches in Virtual Link Trunking (VLT) pairs connected over an OSPF enabled IP underlay using VXLAN tunnels. This topology enables the creation of multiple virtual networks over one common IP underlay network. In this example, the five required VLANs for a successful VxRail deployment are each encapsulated in separate VXLANs. The Internal Management VLAN is used to discover adjacent VxRail nodes and perform initialization to create or expand a VxRail cluster. NVOs allows multirack VxRail discovery and deployment to take place.

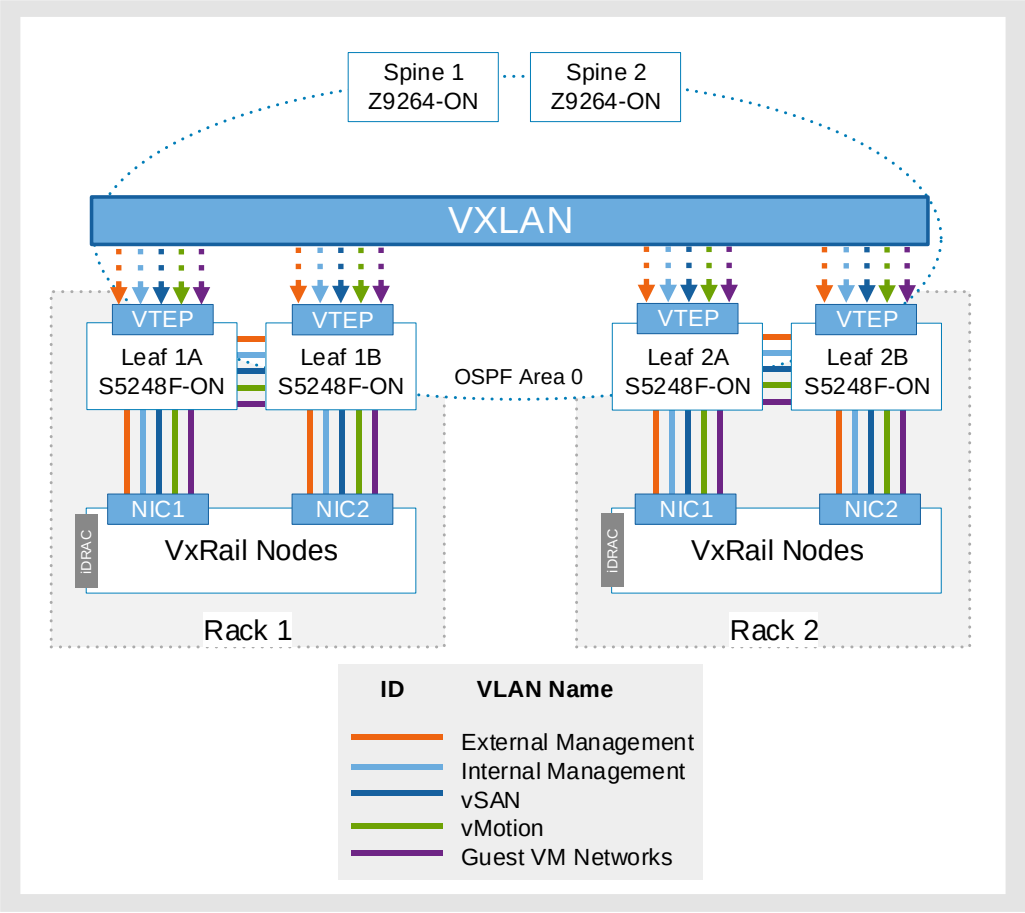

Figure 1 Logical diagram showing VXLAN encapsulation over an IP underlay network

1.1 Objective

This example uses a typical leaf-spine topology with static VXLAN tunnel endpoints (VTEPs) in VLT dual-homing domains. The individual switch configuration shows how to set up an end-to-end virtual network using a static L2 VXLAN configuration and OSPF version 2 to route IP packets.

The deployment consists of eight VxRail nodes, four in one rack and four in a second rack. The VxRail Manager and a combination of switch show commands are used to validate the deployment.

A working example of vSAN stretched cluster is also provided and highlights how to segment vSAN witness traffic from vSAN storage traffic.

Note: This guide does not provide step-by-step guidance on creating an Open Shortest Path First (OSPF) routed underlay or the specific steps for deploying VxRail post-node discovery. For detail instructions on creating a leaf-spine underlay, including alternative configurations, see *Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE.*

1.2 Fabric Design Center (FDC)

The Dell EMC Fabric Design Center (FDC) is a cloud-based application that automates the planning, design and deployment of network fabrics that power Dell EMC compute, storage and hyper-converged infrastructure solutions, including VxRail. The FDC is ideal for turnkey solutions and automation based on validated deployment guides like this one.

FDC allows design customization and flexibility to go beyond validated deployment guides. For additional information, visit the Dell EMC Fabric Design Center.

1.3 Supported switches and operating systems

The examples provided in this Deployment Guide use VxRail 4.7.0 nodes connected to Dell EMC Networking S5248F-ON switches running Dell EMC Networking Dell EMC OS10 Enterprise Edition (OS10EE) 10.4.2.

Dell EMC Networking supports the following switch and OS combinations for VxRail 4.7 and later:

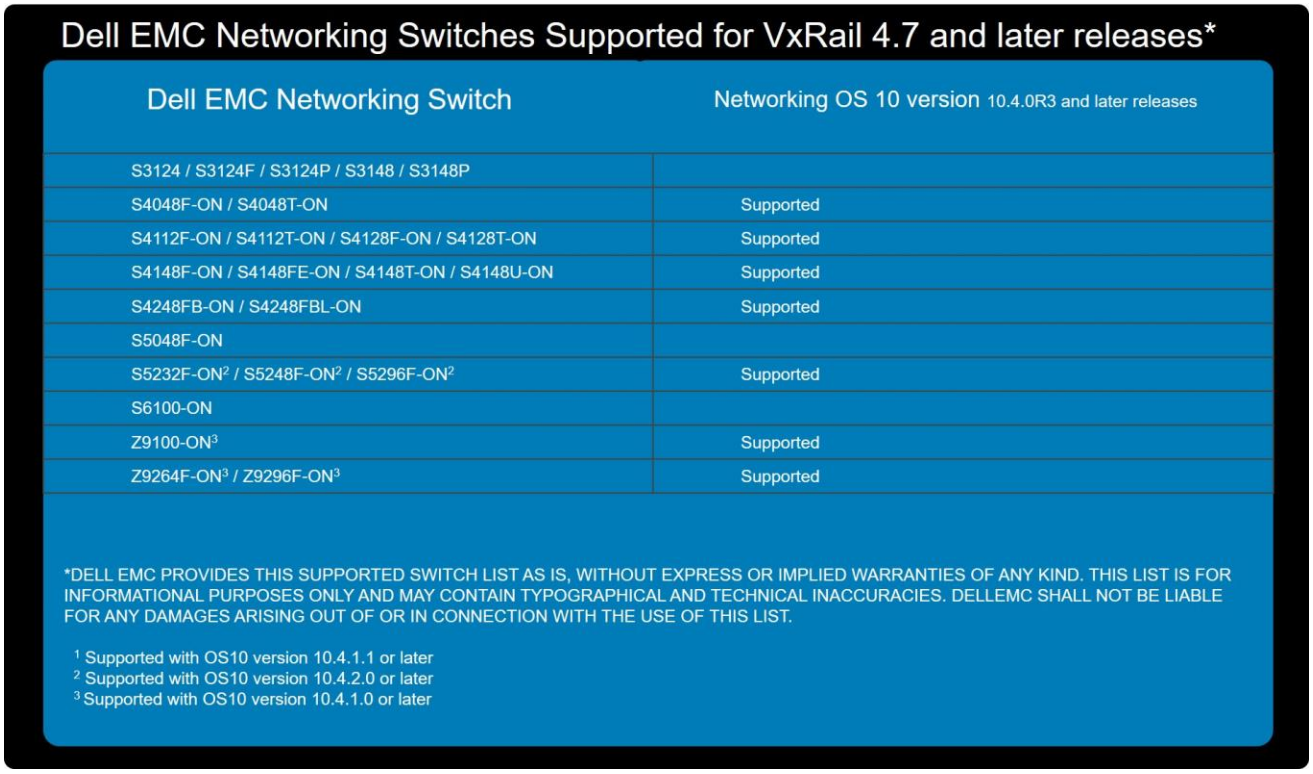

Dell EMC Networking Switches Supported for VxRail 4.7 and later releases*

Dell EMC Networking Switch	Networking OS 10 version 10.4.0R3 and later releases
S3124 / S3124F / S3124P / S3148 / S3148P	
S4048F-ON / S4048T-ON	Supported
S4112F-ON / S4112T-ON / S4128F-ON / S4128T-ON	Supported
S4148F-ON / S4148FE-ON / S4148T-ON / S4148U-ON	Supported
S4248FB-ON / S4248FBL-ON	Supported
S5048F-ON	
S5232F-ON[2] / S5248F-ON[2] / S5296F-ON[2]	Supported
S6100-ON	
Z9100-ON[3]	Supported
Z9264F-ON[3] / Z9296F-ON[3]	Supported

*DELL EMC PROVIDES THIS SUPPORTED SWITCH LIST AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND. THIS LIST IS FOR INFORMATIONAL PURPOSES ONLY AND MAY CONTAIN TYPOGRAPHICAL AND TECHNICAL INACCURACIES. DELLEMC SHALL NOT BE LIABLE FOR ANY DAMAGES ARISING OUT OF OR IN CONNECTION WITH THE USE OF THIS LIST.

[1] Supported with OS10 version 10.4.1.1 or later
[2] Supported with OS10 version 10.4.2.0 or later
[3] Supported with OS10 version 10.4.1.0 or later

Figure 2 Supported Dell EMC Networking switches and operating systems

1.4 Typographical conventions

The CLI and GUI examples in this document use the following conventions:

`Monospace Text`	CLI examples
`Underlined Monospace Text`	CLI examples that wrap the page
`Italic Monospace Text`	Variables in CLI examples
`Bold Monospace Text`	Commands entered at the CLI prompt, or to highlight information in CLI output
Bold text	GUI fields and information entered in the GUI

1.5 Attachments

This document includes switch configuration file attachments. To access attachments in Adobe Acrobat Reader, click the icon in the left pane halfway down the page, then click the icon.

2 Hardware overview

This section provides an overview of the hardware used to validate this deployment. Appendix D contains a complete listing of hardware and software that is validated for this guide.

Note: While the steps in this document were validated using the specified Dell EMC Networking switches and operating systems, they may be used for other Dell EMC Networking switch models using the same networking operating system version or later assuming the switch has the available port numbers, speeds, and types.

2.1 Dell EMC VxRail P570

The Dell EMC VxRail P series consists of high-performance nodes that are optimized for heavy workloads, such as databases. Each appliance in the series has one node per 2-Rack Unit (RU) chassis. The models within this series are the Dell EMC VxRail P570 (hybrid), and the Dell EMC VxRail P570F (all-flash). These models are single or dual processor models based on the Dell EMC PowerEdge R740 rack server. The example within this document uses four VxRail P570 nodes.

Figure 3 Dell EMC VxRail 2-RU node

2.2 Dell EMC VxRail E560

The Dell EMC VxRail E series consists of nodes that are best suited for remote office or entry workloads. The E series nodes support up to 40 CPU cores, 1536GB memory, and 16TB hybrid or 30TB all-flash storage in a 1-RU form factor. The example within this document uses four VxRail E560 nodes.

Figure 4 Dell EMC VxRail 1-RU node

The Dell EMC Networking S5248F-ON is a 1-RU fixed switch with 48x 25 GbE, 4x multirate 100 GbE, and 2x 200 GbE ports, and supports L2 static VXLAN with VLT. The example within this document uses four S5248F-ON switches in VLT pairs, as leaf switches.

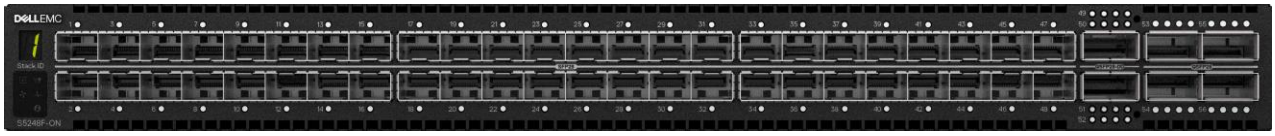

Figure 5 Dell EMC Networking S5248F-ON

2.3 Dell EMC Networking Z9264F-ON

The Dell EMC Networking Z9264F-ON is a 2-RU 100 GbE aggregation/spine switch with up to 64 ports of multirate 100 GbE, or up to 128 ports of 10/25/40/50 GbE ports, using supported breakout cables. The switch is a scalable L2 and L3 Ethernet switch with QoS and a full complement of standards-based IPv4 and IPv6 features, including OSPF and BGP routing support. The example within this document uses two Z9264F-ON switches as spine switches.

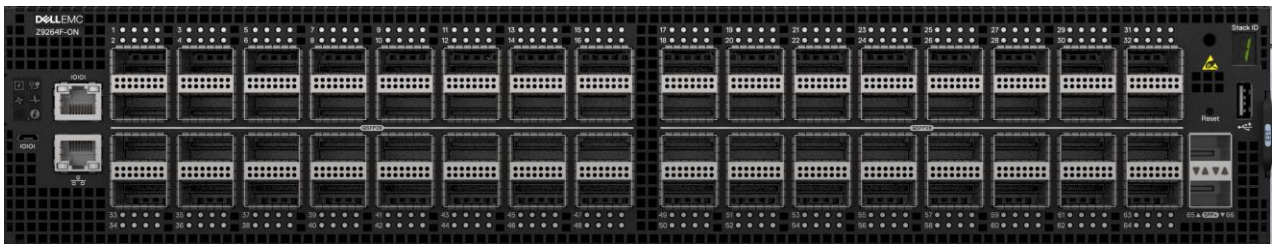

Figure 6 Dell EMC Networking Z9264F-ON

2.4 Dell EMC Networking S3048-ON

The Dell EMC Networking S3048-ON is a 1-RU switch with 48x1GbE BASE-T ports and 4x 10GbE SFP+ ports. This guide uses one S3048-ON switch for out-of-band (OOB) management traffic.

Figure 7 Dell EMC Networking S3048-ON

3 Topology

3.1 VxRail node connectivity

Each VxRail node connection uses two 25 GbE SFP28 interfaces to the Dell EMC Networking S5248F-ON switches. The switches are located in the rack for ease of manageability. All switch interfaces are configured as a VLAN trunk interface. The External Management VLAN is untagged while the remaining four are tagged on each downstream interface.

Figure 8 shows the connections and the associated VLANs, with the S5248F-ON switches configured as a VLT pair. As shown in Figure 8, the VxRail node iDRAC connects to the Dell EMC Networking S3048-ON, which serves as the external management switch.

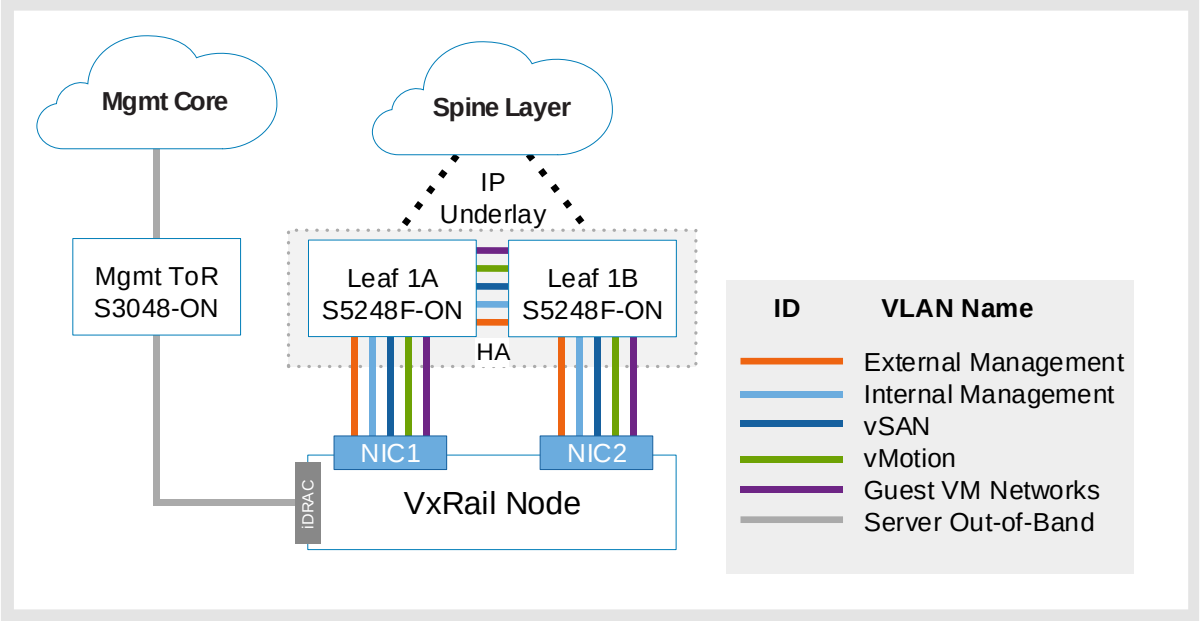

Figure 8 VxRail node connectivity to Dell EMC Networking S5248F-ON leaf switches

Figure 9 shows a physical view of Rack 1. The VxRail P series nodes, sfo01w01vxrail01 through 04 each have 2x 25 GbE links with each being connected to one of the two S5248F-ON leaf switches in the rack. Each VxRail node has an iDRAC connected to a S3048-ON OOB management switch. This connection is used for the initial node configuration. The S5248F-ON leaf switches are connected using two QSFP28-DD 200 GbE direct access cables (DAC) forming a VLT interconnect (VLTi) for a total throughput of 400 GbE . Upstream connections to the spine switches are not shown but are configured using two QSFP28 100 GbE uplinks.

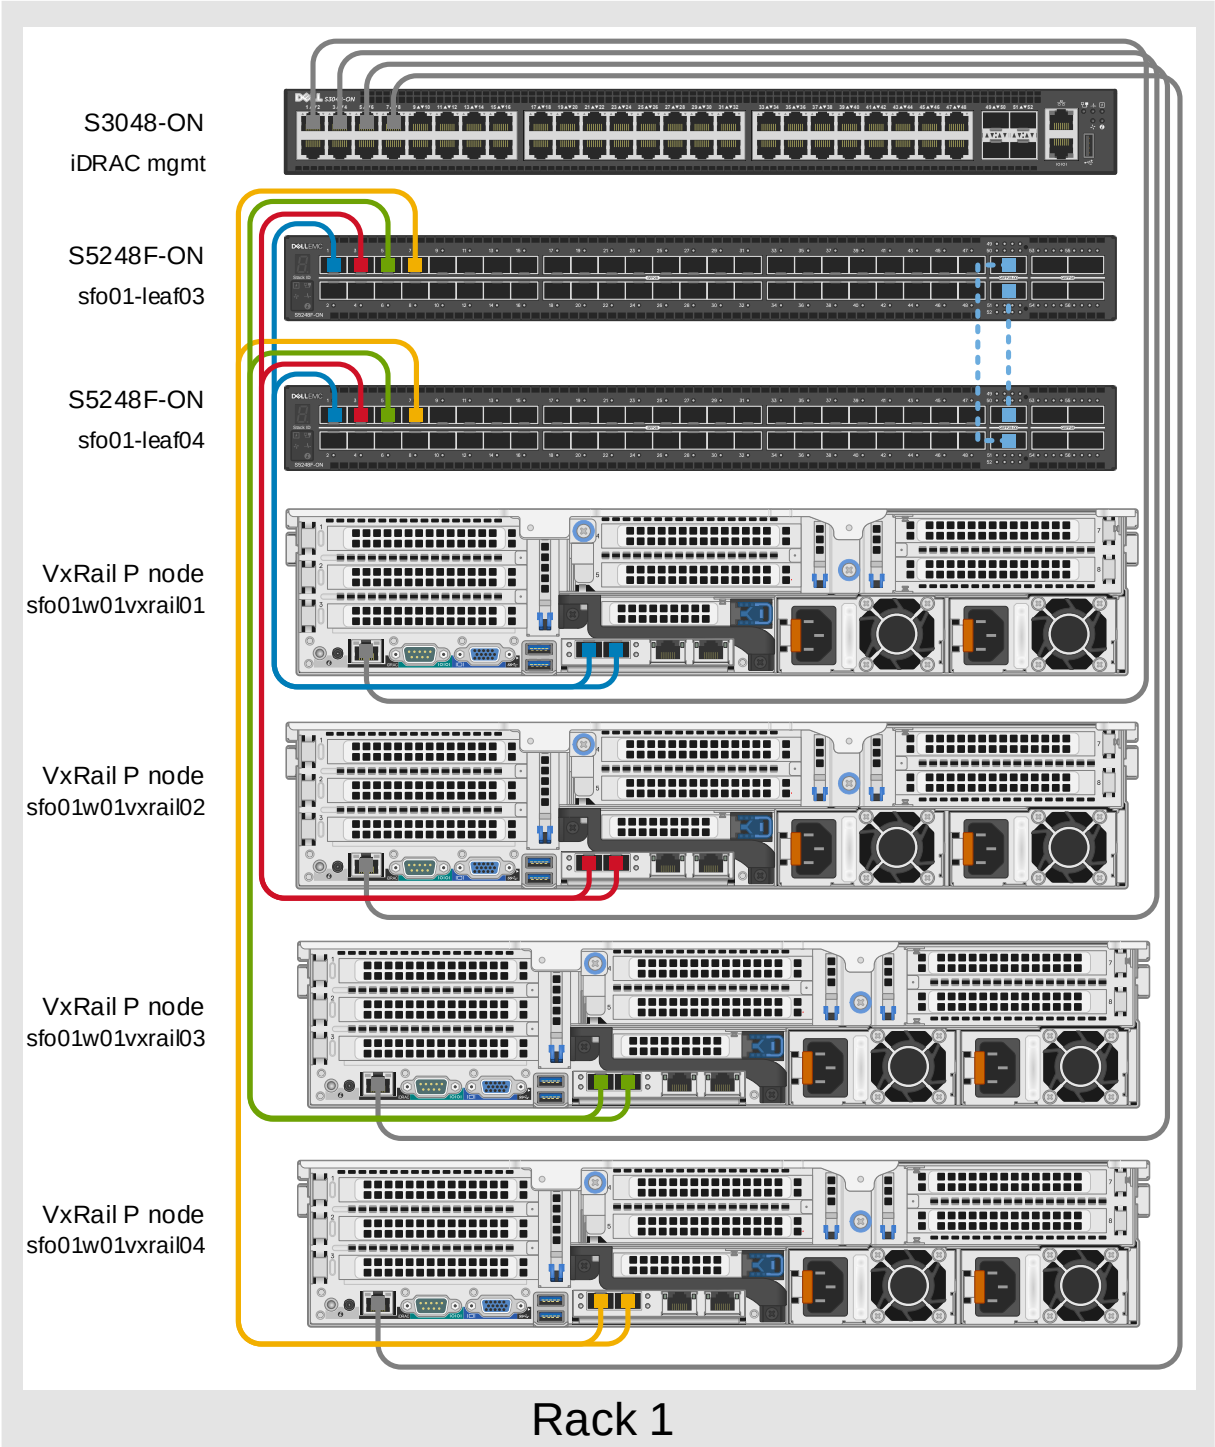

Figure 9 Dell EMC VxRail multirack Rack 1 physical connectivity

3.2 Underlay network design

In this Deployment Guide, a leaf-spine network has already been established as the underlay for the data center. Figure 10 shows that a single OSPFv2 area has been configured using point-to-point interfaces to establish router adjacency.

Changing the interface type to point-to-point avoids DR/BDR election process, reducing the time required to bring up the OSPF adjacency between the leaf switches and spines. In addition, with point-to-point interface mode, there are no Type 2 link-state advertisements (LSAs). Only type 1 LSAs are needed and this configuration model keeps the OSPF LSA database at a minimum size.

Note: For detail instructions on creating a leaf-spine underlay, including considerations and alternative configurations, see *Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE*.

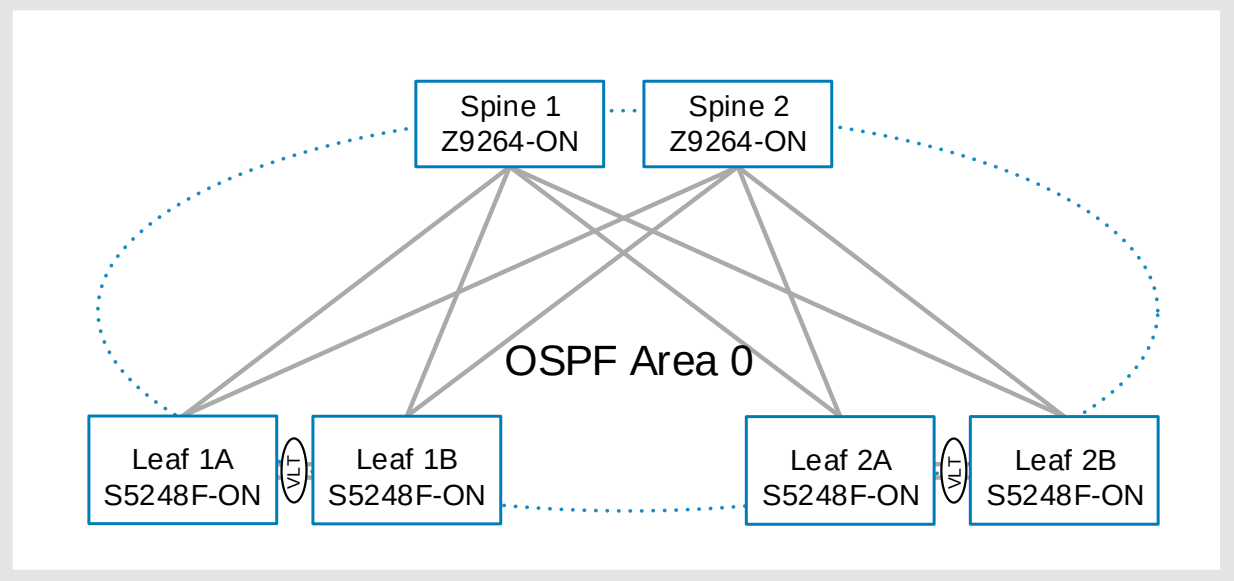

Figure 10 Leaf-spine OSPF IP underlay network diagram

Figure 11 shows the wiring configuration for the six switches that comprise the leaf-spine network. The colored solid lines are 100 GbE links and the light blue dashed lines are two QSFP28-DD 200 GbE cable pairs are used for VLTi.

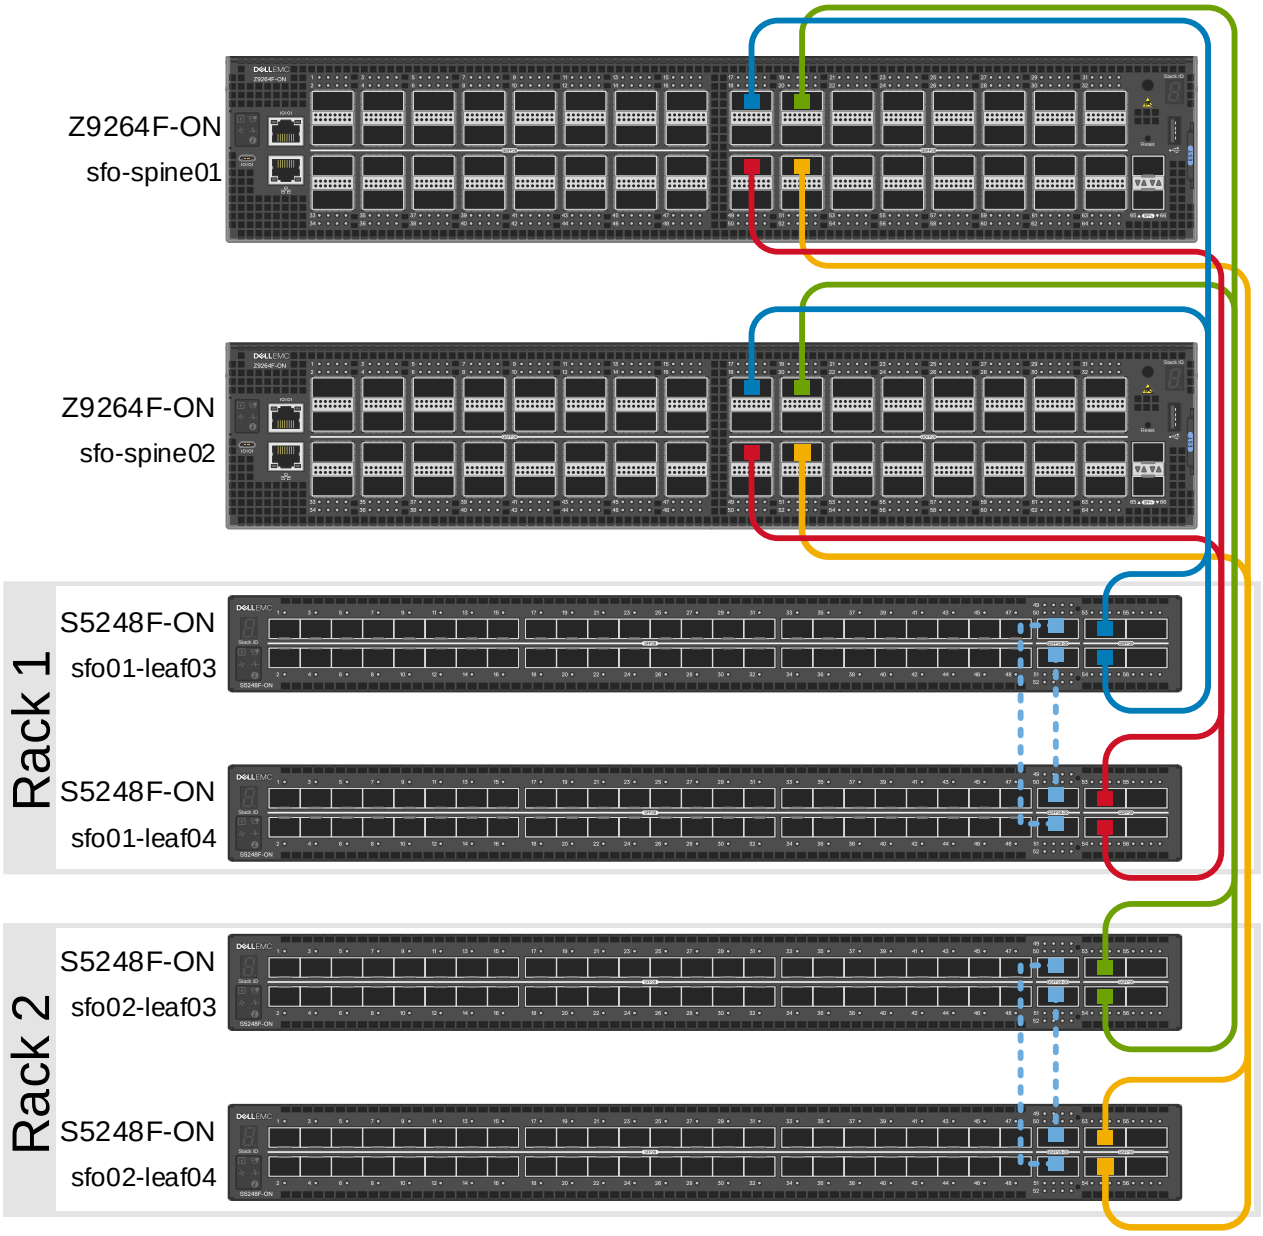

Figure 11 Physical switch topology

Note: All switch configuration commands are provided in the file attachments. See Section 1.4 for instructions on accessing the attachments.

3.3 Virtual extensible LAN (VXLAN) overlay

In this guide, two pairs of leaf switches are configured for static VXLANs. A VXLAN is a type of overlay that encapsulates a payload into UDP packets for transport across the IP underlay. Each leaf switch is configured as a Network Virtualization Edge (NVE) and has a tunnel address, which is the IP addresses used in the VXLAN tunnel header. This IP tunnel address is called a VXLAN tunnel endpoint, or VTEP, and is assigned to the loopback device of the node. 802.1Q is also enabled, but only on the edge ports facing the VxRail nodes.

Figure 12 shows the five minimum required VxRail VLANs, which are shown as solid colored lines that are attached to virtual networks, and that are associated with VTEP/VXLANs at the leaf layer and shown as corresponding dashed lines. Each leaf pair is configured with an identical VTEP and uses the same IP address on a loopback address. Each VTEP is associated statically on the leaf switches. The dashed lines represent the VXLAN tunnel for a given VxRail VLAN.

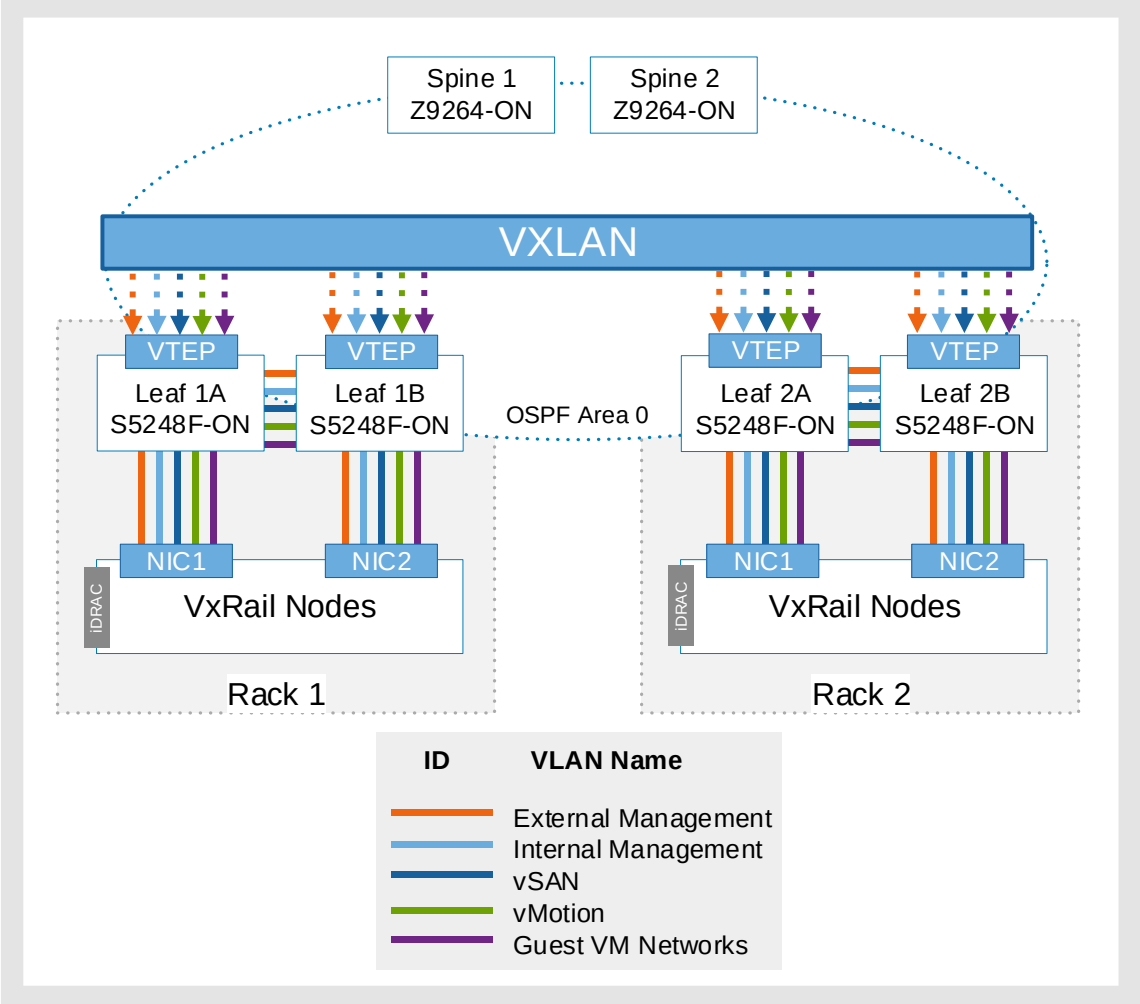

Figure 12 Static VXLAN logical diagram

Note: For more information about static VXLAN concepts, see the *OS10 Enterprise Edition User Guide.*

3.4 External vCenter connectivity

In this document, the VxRail multirack cluster is attached to an external vCenter server. Figure 13 shows the packet flow for a the VxRail cluster, which is shown in orange, connected to the vCenter server, `sfo01m01`, which is shown in blue. The dashed lines represent different VXLAN tunnels connecting the different segments to the switch `sfo-edge01`.

Note: For more information about how routing between tunnels is configured, see Appendix A.

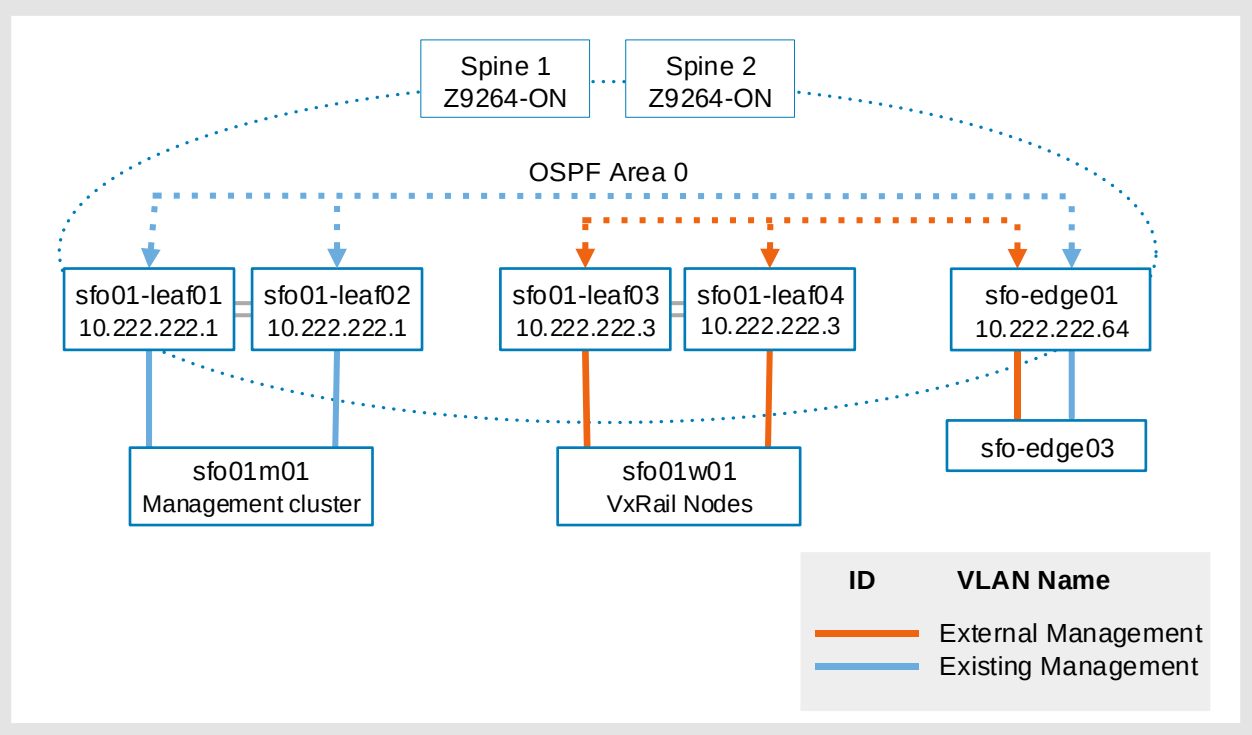

Figure 13 Accessing existing data center services

4 Switch configuration values and prerequisites

This section covers prerequisites to ensure a successful multirack VxRail cluster deployment.

4.1 Pre-planning VLANs, IP address, and switch values

Before configuring the switches or deploying VxRail, VLANs, IP address, and switch specific settings should be planned.

4.1.1 VLANs and IP addresses

VLANs and IP addresses used for VxRail node traffic must be planned before switch configuration, and VxRail deployment can begin. Table 1 shows the five VxRail VLANs and their purpose.

Table 1 VLANs used for VxRail nodes

VLAN	Purpose
External management	VxRail Manager and ESXi management traffic
Internal management	Node discovery
vMotion	Virtual machine migration
vSAN	Distributed storage traffic
Guest VM networks	One or more VLANs for VM data traffic

Table 2 shows six VLANs, VLAN IDs, and IP network addresses planned for this deployment. Two guest VM networks are defined as VLAN 10 and 20.

Table 2 VLANs and IP addresses

VLAN ID	Description	Network (CIDR)	Gateway	VLAN Traffic
1631	External management	172.16.31.0/24	172.16.31.253	Untagged
3939	Internal management	n/a	n/a	Tagged
1632	vMotion	172.16.32.0/24	172.16.32.253	Tagged
1633	vSAN	172.16.33.0/24	172.16.33.253	Tagged
10	Guest VM Network A	192.168.10.0/24	192.168.10.253	Tagged
20	Guest VM Network B	192.168.20.0/24	192.168.20.253	Tagged

Note: Gateway addresses are provided for each network for future expansion of the environment, including stretched vSAN and vMotion of VMs to other independent clusters. For information about routing L2 VXLANs, see Appendix A.

4.1.2 Switch settings

Table 3 shows all unique values for the four S5248F-ON switches. The table is listed to provide a summary of the configuration differences between each switch and between each VLT switch pair. Switches `sfo01-leaf01` and `sfo01-leaf02` are part of the existing infrastructure and shown in Appendix A.

Table 3 Unique switch settings

Setting	S5248F-Leaf1A	S5248F-Leaf1B	S5248F-Leaf2A	S5248F-Leaf2B
Hostname	sfo01-leaf03	sfo01-leaf04	sfo02-leaf03	sfo02-leaf04
Management IP	100.67.167.36/24	100.67.167.35/24	100.67.166.36/24	100.67.166.35/24
P2P IP #1	192.168.1.9/31	192.168.1.11/31	192.168.1.13/31	192.168.1.15/31
P2P IP #2	192.168.2.9/31	192.168.2.11/31	192.168.2.13/31	192.168.2.15/31
Router ID	10.0.2.5	10.0.2.6	10.0.2.7	10.0.2.8
Loopback0 IP	10.222.222.3/32	10.222.222.3/32	10.222.222.4/32	10.222.222.4/32
RSTP priority	0	4096	0	4096
VLAN 4000 IP	10.255.2.4/31	10.255.2.5/31	10.255.2.6/31	10.255.2.7/31
Remote VTEPS	10.222.222.4 10.222.222.64	10.222.222.4 10.222.222.64	10.222.222.3 10.222.222.64	10.222.222.3 10.222.222.64

Note: All switch configuration commands are provided in the file attachments. See Section 1.4 for instructions on accessing the attachments.

4.2 Switch prerequisites

Before starting, verify that all of the switches meet the following requirements:

- The latest version of Dell EMC OS10 Enterprise Edition (OS10EE) is installed
- All OS10EE licenses are entered
- Each switch is in a factory default state

4.2.1 Verify OS10EE version

The Dell EMC Networking S5248F-ON switches must have OS10EE version 10.4.2 or later to support Network Virtualization Overlays (NVOs). Run the **`show version`** command to check the operating system version.

Note: Dell EMC recommends upgrading to the latest release available on Dell Digital Locker (account required).

```
OS10# show version
Dell EMC Networking OS10-Enterprise
Copyright (c) 1999-2018 by Dell Inc. All Rights Reserved.
OS Version: 10.4.2.0
Build Version: 10.4.2.0.241
Build Time: 2018-12-03T17:28:37-0800
```

Note: Figure 2 provides a list of the supported switches and operating systems for VxRail deployments.

4.2.2 Verify license installation

Run the command `show license status` command to verify license installation. Locate the `License Type:` field and verify that `PERPETUAL` displays in the field.

```
OS10# show license status

System Information
---------------------------------------------------------
Vendor Name     :        Dell EMC
Product Name    :        S5248F-ON
Hardware Version:        A00
Platform Name   :        x86_64-dellemc_s5248f_c3538-r0
PPID            :        CN046MRJCES0085N0006
Service Tag     :        AAAAAAA
License Details
----------------
Software        :        OS10-Enterprise
Version         :        10.4.2.0
License Type    :        PERPETUAL
License Duration:        Unlimited
License Status  :        Active
License location:        /mnt/license/GPZQG02.lic
---------------------------------------------------------
```

Note: If an evaluation license is installed, licenses purchased from Dell EMC are available for download on Dell Digital Locker. See the *OS10 Enterprise Edition User Guide* for installation instructions.

4.2.3 Factory default configuration

The configuration commands begin with the Dell EMC Networking switches at their factory default settings. Dell EMC Networking switches running the Dell EMC OS10 Enterprise Edition (OS10EE) can be reset to their default configuration using the following commands:

```
OS10# delete startup-configuration
Proceed to delete startup-configuration [confirm yes/no(default)]:y

OS10# reload
System configuration has been modified. Save? [yes/no]:n
Proceed to reboot the system? [confirm yes/no]:y
```

When complete, the switch reboots to the factory default configuration.

Note: By default, OS10EE has Telnet disabled, SSH enabled, and the OOB management interface that is configured to get an IP address using DHCP. The default username and password are both `admin`. Dell EMC recommends changing the admin password to a complex password when logging in for the first time.

5 Configure switches

The following sections cover the configuration for S5248F-ON switch with the hostname `sfo01-leaf03`. All switch configuration commands are provided in the file attachments. See Section 1.4 for instructions on accessing the attachments.

5.1 Configure general settings

1. Configure the management interface with a static IP address. This IP address is for remote access, Network Time Protocol (NTP) synchronization, and for virtual link trunking (VLT) domain backup. A management route is then configured and in this example specific to the management network.

```
OS10# configure terminal
OS10(config)# interface mgmt 1/1/1
OS10(conf-if-ma-1/1/1)# no ip address dhcp
OS10(conf-if-ma-1/1/1)# ip address 100.67.167.36/24
OS10(conf-if-ma-1/1/1)# exit
OS10(config)# management route 100.64.0.0/13 managementethernet
OS10(config)# end
OS10# write memory
```

Note: At this point, the configuration can continue from the console connection or an SSH session can be established using the management IP address. The default username and password for OS10EE is admin/admin. It is suggested to change this to a more secure password.

2. Set the hostname, Unified Forwarding Table (UFT) mode, RSTP priority value, and the NTP server. The Unified Forwarding Table (UFT) setting is a performance modification and is optional. Setting the value to scaled-L2 increases the internal L2/L3 forwarding table sizes.

```
OS10# configure terminal
OS10(config)# hostname sfo01-leaf03
sfo01-leaf03(config)# hardware forwarding-table mode scaled-l2
sfo01-leaf03(config)# spanning-tree rstp priority 0
sfo01-leaf03(config)# ntp server 100.67.10.20
```

Note: For more information about UFT, see the *OS10 Enterprise Edition User Guide*.

3. It is a best practice to restrict Telnet or SSH connections to the switch by applying an access list (ACL) on the Virtual Terminal Line (VTY). The ACL `permit-mgmt-access` is created and applied to the VTY. In this example, access is limited to hosts on the 100.67.0.0/13 network.

```
sfo01-leaf03(config)# ip access-list permit-mgmt-access
sfo01-leaf03(config-ipv4-acl)# seq 10 permit tcp 100.64.0.0/13 any eq 22
sfo01-leaf03(config-ipv4-acl)# exit
sfo01-leaf03(config)# line vty
sfo01-leaf03(config-line-vty)# ip access-class permit-mgmt-access
sfo01-leaf03(config-line-vty)# exit
```

5.2 Configure OSPF routing and upstream network-facing ports

To configure Open Shortest Path First (OSPF) routing and upstream network-facing ports, perform the following steps:

1. Enable OSPF routing and assign a router ID.

```
sfo01-leaf03(config)# router ospf 1
sfo01-leaf03(config-router-ospf-1)# router-id 10.0.2.5
sfo01-leaf03(config-router-ospf-1)# exit
```

Note: For guidance on router ID IP addresses, see *Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE*.

2. Configure the upstream connection to the first spine switch, `sfo-spine01`. Enter the following commands to set a description, disable bridging, and set the MTU to the maximum size, assign an IP address, and add the interface to OSPF area 0:

```
sfo01-leaf03(config)# interface ethernet 1/1/53
sfo01-leaf03(conf-if-eth1/1/53)# description sfo-spine01
sfo01-leaf03(conf-if-eth1/1/53)# no shutdown
sfo01-leaf03(conf-if-eth1/1/53)# no switchport
sfo01-leaf03(conf-if-eth1/1/53)# mtu 9216
sfo01-leaf03(conf-if-eth1/1/53)# ip address 192.168.1.9/31
sfo01-leaf03(conf-if-eth1/1/53)# ip ospf network point-to-point
sfo01-leaf03(conf-if-eth1/1/53)# ip ospf 1 area 0.0.0.0
sfo01-leaf03(conf-if-eth1/1/53)# exit
```

3. Configure the upstream connection to the second spine switch, `sfo-spine02`. Enter the following commands to set a description, disable bridging, and set the MTU to the maximum size, assign an IP address, and add the interface to OSPF area 0:

```
sfo01-leaf03(config)# interface ethernet 1/1/54
sfo01-leaf03(conf-if-eth1/1/54)# description sfo-spine02
sfo01-leaf03(conf-if-eth1/1/54)# no shutdown
sfo01-leaf03(conf-if-eth1/1/54)# no switchport
sfo01-leaf03(conf-if-eth1/1/54)# mtu 9216
sfo01-leaf03(conf-if-eth1/1/54)# ip address 192.168.2.9/31
sfo01-leaf03(conf-if-eth1/1/54)# ip ospf network point-to-point
sfo01-leaf03(conf-if-eth1/1/54)# ip ospf 1 area 0.0.0.0
sfo01-leaf03(conf-if-eth1/1/54)# exit
```

Note: See Appendix B for information about MTU considerations.

5.3 Configure Virtual Link Trunk (VLT)

To configure a Virtual Link Trunking (VLT) domain perform the following steps:

1. To configure the VLT interconnect (VLTi) member links enter the following commands:

```
sfo01-leaf03(config)# interface range ethernet 1/1/49-1/1/52
sfo01-leaf03(conf-range-eth1/1/49-1/1/52)# description VLTi
sfo01-leaf03(conf-range-eth1/1/49-1/1/52)# no shutdown
sfo01-leaf03(conf-range-eth1/1/49-1/1/52)# no switchport
sfo01-leaf03(conf-range-eth1/1/49-1/1/52)# exit
```

2. Use the following commands to configure a VLT domain:

```
sfo01-leaf03(config)# vlt-domain 1
sfo01-leaf03(conf-vlt-1)# backup destination 100.67.166.34
sfo01-leaf03(conf-vlt-1)# discovery-interface ethernet 1/1/49-1/1/52
sfo01-leaf03(conf-vlt-1)# peer-routing
sfo01-leaf03(conf-vlt-1)# exit
```

3. Enter the following commands to configure a dedicated L3 underlay path to reach the VLT peer if there is a VLTi link failure:

```
sfo01-leaf03(config)# interface vlan 4000
sfo01-leaf03(conf-if-vl-4000)# description L3-VLT-peer
sfo01-leaf03(conf-if-vl-4000)# no shutdown
sfo01-leaf03(conf-if-vl-4000)# mtu 9216
sfo01-leaf03(conf-if-vl-4000)# ip address 10.255.2.4/31
sfo01-leaf03(conf-if-vl-4000)# ip ospf 1 area 0.0.0.0
sfo01-leaf03(conf-if-vl-4000)# exit
```

Note: This guide uses the 2x 200 GbE (QSFP-DD) ports to provide a 400 GbE VLTi link. As an alternative, QSFP28 100 GbE cables can be used in the same interfaces provided a 200 GbE VLTi link.

5.4 Configure Uplink Failure Detection (UFD)

Configure Uplink Failure Detection (UFD) with uplinks upstream network interfaces and downstream for all VxRail node interfaces. If all upstream connectivity to the spine layer is lost, the VxRail node interfaces are shut down.

```
sfo01-leaf03(config)# uplink-state-group 1
sfo01-leaf03(conf-uplink-state-group-1)# enable
sfo01-leaf03(conf-uplink-state-group-1)# downstream ethernet
1/1/1,1/1/3,1/1/5,1/1/7
sfo01-leaf03(conf-uplink-state-group-1)# upstream ethernet1/1/53-1/1/54
sfo01-leaf03(conf-uplink-state-group-1)# exit
```

5.5 Configure Virtual Extensible LANs (VXLAN)

To configure Virtual Extensible LANs (VXLAN) perform the following steps:

1. Enter the following commands to configure a loopback interface and add it to OSPF area 0:

```
sfo01-leaf03(config)# interface loopback 0
sfo01-leaf03(conf-if-lo-0)# ip address 10.222.222.3/32
sfo01-leaf03(conf-if-lo-0)# ip ospf 1 area 0.0.0.0
sfo01-leaf03(conf-if-lo-0)# exit
```

2. To configure the loopback interface as the VXLAN source tunnel interface, enter the following commands:

```
sfo01-leaf03(config)# nve
sfo01-leaf03(conf-nve)# source-interface loopback0
sfo01-leaf03(conf-nve)# exit
```

3. Enter the following commands to configure the VXLAN virtual network 1631, VxRail external management traffic, with a virtual network interface (VNI) and the appropriate remote VTEPs:

```
sfo01-leaf03(config)# virtual-network 1631
sfo01-leaf03(conf-vn-1631)# description sfo01-w01-
vxrail_external_management
sfo01-leaf03(conf-vn-1631)# vxlan-vni 1631
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.4
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.64
sfo01-leaf03(conf-vn-vxlan-vni)# exit
sfo01-leaf03(conf-vn-1631)# exit
```

4. Enter the following commands to configure the VXLAN virtual network 3939, VxRail internal management traffic, with a virtual network interface (VNI) and the appropriate remote VTEPs:

```
sfo01-leaf03(config)# virtual-network 3939
sfo01-leaf03(conf-vn-3939)# description sfo01-w01-
vxrail_internal_management
sfo01-leaf03(conf-vn-3939)# vxlan-vni 3939
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.4
sfo01-leaf03(conf-vn-vxlan-vni)# exit
sfo01-leaf03(conf-vn-3939)# exit
```

5. Enter the following commands to configure the VXLAN virtual network 1632, vMotion traffic, with a virtual network interface (VNI) and the appropriate remote VTEPs:

```
sfo01-leaf03(config)# virtual-network 1632
sfo01-leaf03(conf-vn-1632)# description sfo01-w01-vmotion
sfo01-leaf03(conf-vn-1632)# vxlan-vni 1632
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.4
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.64
sfo01-leaf03(conf-vn-vxlan-vni)# exit
sfo01-leaf03(conf-vn-1632)# exit
```

6. Enter the following commands to configure the VXLAN virtual network 1633, vSAN traffic, with a virtual network interface (VNI) and the appropriate remote VTEPs:

```
sfo01-leaf03(config)# virtual-network 1633
sfo01-leaf03(conf-vn-1633)# description sfo01-w01-vsan
sfo01-leaf03(conf-vn-1633)# vxlan-vni 1633
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.4
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.64
sfo01-leaf03(conf-vn-vxlan-vni)# exit
sfo01-leaf03(conf-vn-1633)# exit
```

7. Enter the following commands to configure the VXLAN virtual network 10, for the 1st guest network, with a virtual network interface (VNI) and the appropriate remote VTEPs:

```
sfo01-leaf03(config)# virtual-network 10
sfo01-leaf03(conf-vn-10)# description sfo01-w01-guest_vm_network_a
sfo01-leaf03(conf-vn-10)# vxlan-vni 10
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.4
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.64
sfo01-leaf03(conf-vn-vxlan-vni)# exit
sfo01-leaf03(conf-vn-10)# exit
```

8. Enter the following commands to configure the VXLAN virtual network 20, for the 2nd guest network, with a virtual network interface (VNI) and the appropriate remote VTEPs:

```
sfo01-leaf03(config)# virtual-network 20
sfo01-leaf03(conf-vn-20)# description sfo01-w01-guest_vm_network_b
sfo01-leaf03(conf-vn-20)# vxlan-vni 20
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.4
sfo01-leaf03(conf-vn-vxlan-vni)# remote-vtep 10.222.222.64
sfo01-leaf03(conf-vn-vxlan-vni)# exit
sfo01-leaf03(conf-vn-20)# exit
```

Note: The remote VTEP 10.222.222.64 is used to access external sources outside of the VXLAN tunnel. For more information about routing between VXLAN tunnels, see Appendix A.

5.6 Configure Virtual LANs (VLAN)

To configure the Virtual LANs (VLAN) that are associated with the virtual network, perform the following steps:

1. Enter the following commands to configure VLAN 1631, VxRail external management traffic, and assign the VLAN to the VXLAN virtual network:

```
sfo01-leaf03(config)# interface vlan 1631
sfo01-leaf03(conf-if-vl-1631)# description sfo01-w01-
vxrail_external_management
sfo01-leaf03(conf-if-vl-1631)# virtual-network 1631
sfo01-leaf03(conf-if-vl-1631)# no shutdown
sfo01-leaf03(conf-if-vl-1631)# exit
```

2. Enter the following commands to configure VLAN 3939, VxRail internal management traffic, and assign the VLAN to the VXLAN virtual network:

```
sfo01-leaf03(config)# interface vlan 3939
sfo01-leaf03(conf-if-vl-3939)# description sfo01-w01-
vxrail_internal_management
sfo01-leaf03(conf-if-vl-3939)# virtual-network 3939
sfo01-leaf03(conf-if-vl-3939)# no shutdown
sfo01-leaf03(conf-if-vl-3939)# exit
```

3. Enter the following commands to configure VLAN 1632, vMotion traffic, and assign the VLAN to the VXLAN virtual network:

```
sfo01-leaf03(config)# interface vlan 1632
sfo01-leaf03(conf-if-vl-1632)# description sfo01-w01-vmotion
sfo01-leaf03(conf-if-vl-1632)# virtual-network 1632
sfo01-leaf03(conf-if-vl-1632)# no shutdown
sfo01-leaf03(conf-if-vl-1632)# exit
```

4. Enter the following commands to configure VLAN 1633, vSAN traffic, and assign the VLAN to the VXLAN virtual network:

```
sfo01-leaf03(config)# interface vlan 1633
sfo01-leaf03(conf-if-vl-1633)# description sfo01-w01-vsan
sfo01-leaf03(conf-if-vl-1633)# virtual-network 1633
sfo01-leaf03(conf-if-vl-1633)# no shutdown
sfo01-leaf03(conf-if-vl-1633)# exit
```

5. Enter the following commands to configure VLAN 10 for the first guest network traffic, and assign the VLAN to the VXLAN virtual network:

```
sfo01-leaf03(config)# interface vlan 10
sfo01-leaf03(conf-if-vl-10)# description sfo01-w01-guest_vm_network_a
sfo01-leaf03(conf-if-vl-10)# virtual-network 10
sfo01-leaf03(conf-if-vl-10)# no shutdown
sfo01-leaf03(conf-if-vl-10)# exit
```

6. Enter the following commands to configure VLAN 20 for the second guest network traffic, and assign the VLAN to the VXLAN virtual network:

```
sfo01-leaf03(config)# interface vlan 20
sfo01-leaf03(conf-if-vl-20)# description sfo01-w01-guest_vm_network_b
sfo01-leaf03(conf-if-vl-20)# virtual-network 20
sfo01-leaf03(conf-if-vl-20)# no shutdown
sfo01-leaf03(conf-if-vl-20)# exit
```

5.7 Configure downstream VxRail node interfaces

To configure the four downstream interfaces connecting to the Dell EMC VxRail P-Series nodes, perform the following steps:

1. Enter the following commands to configure the downstream interface for the first VxRail node Configure the interface as a trunk link. Set the untagged VLAN to the VxRail external management VLAN 1631. Allow all remaining VLANs on the trunk port. Set MTU to the maximum size. Set flow control to “receive on” and “transmit off” on node-connected ports as a best practice. Spanning tree edge fast is enabled.

```
sfo01-leaf03(config)# interface ethernet 1/1/1
sfo01-leaf03(conf-if-eth1/1/1)# description sfo01w01vxrail01
sfo01-leaf03(conf-if-eth1/1/1)# no shutdown
sfo01-leaf03(conf-if-eth1/1/1)# switchport mode trunk
sfo01-leaf03(conf-if-eth1/1/1)# switchport access vlan 1631
sfo01-leaf03(conf-if-eth1/1/1)# switchport trunk allowed vlan 10,20,1632-
1633,3939
sfo01-leaf03(conf-if-eth1/1/1)# mtu 9216
sfo01-leaf03(conf-if-eth1/1/1)# flowcontrol receive on
sfo01-leaf03(conf-if-eth1/1/1)# flowcontrol transmit off
sfo01-leaf03(conf-if-eth1/1/1)# spanning-tree port type edge
sfo01-leaf03(conf-if-eth1/1/1)# exit
```

2. Enter the following commands to configure the downstream interface for the second VxRail node. Configure the interface as a trunk link. Set the untagged VLAN to the VxRail external management VLAN 1631. Allow all remaining VLANs on the trunk port. Set MTU to the maximum size. Set flow control to “receive on” and “transmit off” on node-connected ports as a best practice. Spanning tree edge fast is enabled.

```
sfo01-leaf03(config)# interface ethernet 1/1/3
sfo01-leaf03(conf-if-eth1/1/3)# description sfo01w01vxrail02
sfo01-leaf03(conf-if-eth1/1/3)# no shutdown
sfo01-leaf03(conf-if-eth1/1/3)# switchport mode trunk
sfo01-leaf03(conf-if-eth1/1/3)# switchport access vlan 1631
sfo01-leaf03(conf-if-eth1/1/3)# switchport trunk allowed vlan 10,20,1632-
1633,3939
sfo01-leaf03(conf-if-eth1/1/3)# mtu 9216
sfo01-leaf03(conf-if-eth1/1/3)# flowcontrol receive on
sfo01-leaf03(conf-if-eth1/1/3)# flowcontrol transmit off
sfo01-leaf03(conf-if-eth1/1/3)# spanning-tree port type edge
sfo01-leaf03(conf-if-eth1/1/3)# exit
```

3. Enter the following commands to configure the downstream interface for the third VxRail node. Configure the interface as a trunk link. Set the untagged VLAN to the VxRail external management VLAN 1631. Allow all remaining VLANs on the trunk port. Set MTU to the maximum size. Set flow

control to “receive on” and “transmit off” on node-connected ports as a best practice. Spanning tree edge fast is enabled.

```
sfo01-leaf03(config)# interface ethernet 1/1/5
sfo01-leaf03(conf-if-eth1/1/5)# description sfo01w01vxrail03
sfo01-leaf03(conf-if-eth1/1/5)# no shutdown
sfo01-leaf03(conf-if-eth1/1/5)# switchport mode trunk
sfo01-leaf03(conf-if-eth1/1/5)# switchport access vlan 1631
sfo01-leaf03(conf-if-eth1/1/5)# switchport trunk allowed vlan 10,20,1632-
1633,3939
sfo01-leaf03(conf-if-eth1/1/5)# mtu 9216
sfo01-leaf03(conf-if-eth1/1/5)# flowcontrol receive on
sfo01-leaf03(conf-if-eth1/1/5)# flowcontrol transmit off
sfo01-leaf03(conf-if-eth1/1/5)# spanning-tree port type edge
sfo01-leaf03(conf-if-eth1/1/5)# exit
```

4. Enter the following commands to configure the downstream interface for the fourth VxRail node. Configure the interface as a trunk link. Set the untagged VLAN to the VxRail external management VLAN 1631. Allow all remaining VLANs on the trunk port. Set MTU to the maximum size. Set flow control to “receive on” and “transmit off” on node-connected ports as a best practice. Spanning tree edge fast is enabled.

```
sfo01-leaf03(config)# interface ethernet 1/1/7
sfo01-leaf03(conf-if-eth1/1/7)# description sfo01w01vxrail04
sfo01-leaf03(conf-if-eth1/1/7)# no shutdown
sfo01-leaf03(conf-if-eth1/1/7)# switchport mode trunk
sfo01-leaf03(conf-if-eth1/1/7)# switchport access vlan 1631
sfo01-leaf03(conf-if-eth1/1/7)# switchport trunk allowed vlan 10,20,1632-
1633,3939
sfo01-leaf03(conf-if-eth1/1/7)# mtu 9216
sfo01-leaf03(conf-if-eth1/1/7)# flowcontrol receive on
sfo01-leaf03(conf-if-eth1/1/7)# flowcontrol transmit off
sfo01-leaf03(conf-if-eth1/1/7)# spanning-tree port type edge
sfo01-leaf03(conf-if-eth1/1/7)# exit
```

Note: See Appendix B for information about MTU considerations.

6 Switch validation

After switches are configured and devices are connected, the Dell EMC OS10 Enterprise Edition (OS10EE) CLI is used to validate the network configuration. The following commands and outputs are for `sfo01-leaf03`. The output of its peer, `sfo01-leaf04`, and the remaining S5248F-ON switches are similar. This section provides a list of the most common commands and their output for the examples that are used in this guide.

Note: In this section, different filters are used to view specific information and to keep the output brief. The following commands can be run without the use of filters. For more information, see *OS10 Enterprise Edition User Guide*.

6.1 General validation commands

6.1.1 show interface status

The `show interface status` command is shows the interfaces that are up, operational link speed, and port modes such as access, trunk, or L3 mode (–). For L2 interfaces, the untagged and tagged VLANs are listed. In this example, the four VxRail node interfaces show the six VLANs, VLTi, and upstream network-facing ports as operational.

```
sfo01-leaf03# show interface status | grep up
Port       Description  Status  Speed  Duplex  Mode Vlan Tagged-Vlans
Eth 1/1/1   sfo01w01..    up     25G    full    T   1631 10,20,1632-1633,3939
Eth 1/1/3   sfo01w01..    up     25G    full    T   1631 10,20,1632-1633,3939
Eth 1/1/5   sfo01w01..    up     25G    full    T   1631 10,20,1632-1633,3939
Eth 1/1/7   sfo01w01..    up     25G    full    T   1631 10,20,1632-1633,3939
Eth 1/1/49  VLTi          up     100G   full    -
Eth 1/1/50  VLTi          up     100G   full    -
Eth 1/1/51  VLTi          up     100G   full    -
Eth 1/1/52  VLTi          up     100G   full    -
Eth 1/1/53  sfo-spine01   up     100G   full    -
Eth 1/1/54  sfo-spine02   up     100G   full    -
```

6.1.2 show ip interface brief

The `show ip interface brief` command verifies the IP addresses assigned to interfaces.

```
sfo01-leaf03# show ip interface brief | except unassigned
Interface Name    IP-Address         OK   Method  Status  Protocol
================================================================
Ethernet 1/1/97   192.168.1.11/31    YES  manual  up      up
Ethernet 1/1/98   192.168.2.11/31    YES  manual  up      up
Management 1/1/1  100.67.167.36/24   YES  manual  up      up
Vlan 4000         10.255.2.4/31      YES  manual  up      up
Loopback 0        10.222.222.3/32    YES  manual  up      up
```

6.1.3 show lldp neighbors

The `show lldp neighbors` command identifies devices that are directly connected to the switch. `Vmnic0` is shown along with the VLTi links, `ethernet 1/1/49-1/1/52`, and both upstream network-facing ports, `ethernet1/1/53` and `ethernet1/1/54`.

```
sfo01-leaf03# show lldp neighbors | except Broadcom
Loc PortID       Rem Host Name        Rem Port Id       Rem Chassis Id
-------------------------------------------------------------------------
ethernet1/1/1   sfo01w01vxrail01  00:0a:f7:f6:2d:e0    vmnic0
ethernet1/1/3   sfo01w01vxrail02  00:0a:f7:f6:1b:a0    vmnic0
ethernet1/1/49  sfo02-leaf04      ethernet1/1/49       54:bf:64:be:f7:40
ethernet1/1/5   sfo01w01vxrail03  00:0a:f7:f5:ee:40    vmnic0
ethernet1/1/50  sfo02-leaf04      ethernet1/1/50       54:bf:64:be:f7:40
ethernet1/1/51  sfo02-leaf04      ethernet1/1/51       54:bf:64:be:f7:40
ethernet1/1/52  sfo02-leaf04      ethernet1/1/52       54:bf:64:be:f7:40
ethernet1/1/53  sfo-spine01       ethernet1/1/17       20:04:0f:11:42:62
ethernet1/1/54  sfo-spine02       ethernet1/1/17       20:04:0f:11:47:62
ethernet1/1/7   sfo01w01vxrail04  00:0a:f7:f5:fa:c0    vmnic0
```

6.1.4 show interface

The `show interface` command verifies that the flow control is enabled for receive and transmit. The following example shows that the interfaces are connected to the four VxRail nodes (`ethernet 1/1/1`, `1/1/3`, `1/1/5`, and `1/1/7`) and set flow control to “receive on” and “transmit off” on node-connected ports as a best practice.

```
sfo01-leaf03# show interface | grep "Flowcontrol rx on tx off"
Flowcontrol rx on tx off
Flowcontrol rx on tx off
Flowcontrol rx on tx off
Flowcontrol rx on tx off
```

6.2 OSPF validation commands

In this section, there are three primary tables:

- Routing
- Topology
- Neighbor

6.2.1 show ip ospf neighbor

The `show ip ospf neighbor` command verifies full adjacency with neighboring switches. In the following example, both `sfo-spine01` and `sfo-spine02` are in the `FULL/BDR` state, with the peer adjacency across `vlan4000` in the `FULL/DR` state.

```
sfo01-leaf03# show ip ospf neighbor
Neighbor ID  Pri  State     Dead Time  Address       Interface       Area
------------------------------------------------------------------------
10.0.2.6     1    FULL/DR   00:00:39   10.255.2.5    vlan4000        0.0.0.0
10.0.1.1     1    FULL/BDR  00:00:37   192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.1.2     1    FULL/BDR  00:00:39   192.168.2.10  ethernet1/1/54  0.0.0.0
```

6.2.2 show ip route ospf

The `show ip route ospf` command shows the active routes for a network, where the `10.222.222.0/24` subnet is used for VTEP endpoint IP addresses. Each of the following IP addresses has two paths - one through `sfo-spine01,` and another through `sfo-spine02`. The remaining addresses are to the various L3 interfaces that are configured in the leaf-spine underlay.

```
sfo01-leaf03# show ip route ospf
Codes: C - connected
       S - static
       B - BGP, IN - internal BGP, EX - external BGP
       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, * - candidate default,
       + - summary route, > - non-active route
Gateway of last resort is not set
  Destination            Gateway                  Dist/Metric  Last Change
-------------------------------------------------------------------------------
-
  O  10.222.222.1/32   via 192.168.1.10  eth1/1/53  110/4        21:52:55
                       via 192.168.2.10  eth1/1/54
  O  10.222.222.4/32   via 192.168.1.10  eth1/1/53  110/3        21:52:55
                       via 192.168.2.10  eth1/1/54
  O  10.222.222.64/32  via 192.168.1.10  eth1/1/53  110/4        21:52:55
                       via 192.168.2.10  eth1/1/54
  O  192.168.1.0/31    via 192.168.1.10  eth1/1/53  110/3        21:53:01
  O  192.168.1.2/31    via 192.168.1.10  eth1/1/53  110/3        21:53:01
  O  192.168.1.12/31   via 192.168.1.10  eth1/1/53  110/2        21:53:01
  O  192.168.1.14/31   via 192.168.1.10  eth1/1/53  110/2        21:53:01
  O  192.168.1.252/31  via 192.168.1.10  eth1/1/53  110/3        21:53:01
  O  192.168.2.0/31    via 192.168.2.10  eth1/1/54  110/3        21:52:55
  O  192.168.2.2/31    via 192.168.2.10  eth1/1/54  110/3        21:52:55
  O  192.168.2.12/31   via 192.168.2.10  eth1/1/54  110/2        21:52:55
  O  192.168.2.14/31   via 192.168.2.10  eth1/1/54  110/2        21:52:55
  O  192.168.2.252/31  via 192.168.2.10  eth1/1/54  110/3        21:52:55
```

6.2.3 show ip ospf topology

The `show ip ospf topology` command verifies the structure of the network.

```
sfo01-leaf03# show ip ospf topology
Router ID  Flags   Cost  Nexthop        Interface       Area
-------------------------------------------------------------
10.0.1.1   -/-/-/  1     192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.1.2   -/-/-/  1     192.168.2.10  ethernet1/1/54  0.0.0.0
10.0.2.1   -/-/-/  3     192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.2.1   -/-/-/  3     192.168.2.10  ethernet1/1/54  0.0.0.0
10.0.2.2   -/-/-/  3     192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.2.2   -/-/-/  3     192.168.2.10  ethernet1/1/54  0.0.0.0
10.0.2.6   -/-/-/  2     192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.2.6   -/-/-/  2     192.168.2.10  ethernet1/1/54  0.0.0.0
10.0.2.7   -/-/-/  2     192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.2.7   -/-/-/  2     192.168.2.10  ethernet1/1/54  0.0.0.0
10.0.2.8   -/-/-/  2     192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.2.8   -/-/-/  2     192.168.2.10  ethernet1/1/54  0.0.0.0
10.0.2.64  -/-/-/  3     192.168.1.10  ethernet1/1/53  0.0.0.0
10.0.2.64  -/-/-/  3     192.168.2.10  ethernet1/1/54  0.0.0.0
```

6.3 VXLAN validation commands

This section provides a listing of the commands that are used to validate VXLAN.

6.3.1 show vlan

The `show vlan` command verifies that the VLANs are attached to virtual networks. In the following example, the “at” symbol (@) represents active VLANs. This command also confirms that all VLANs are tagged on port channel 1000 (`po1000`), which is the default interface for VLTi.

```
sfo01-leaf03# show vlan | except default
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs,
       @ - Attached to Virtual Network
Q: A - Access (Untagged), T - Tagged
   NUM   Status    Description                Q Ports
                                              A Po1000
@  10    Active    sfo01-w01-guest_vm_netwo   T Eth1/1/1,1/1/3,1/1/5,1/1/7
                                              T Po1000
@  20    Active    sfo01-w01-guest_vm_netwo   T Eth1/1/1,1/1/3,1/1/5,1/1/7
                                              T Po1000
@  1631  Active    sfo01-w01-vxrail_externa   T Po1000
                                              A Eth1/1/1,1/1/3,1/1/5,1/1/7
@  1632  Active    sfo01-w01-vmotion          T Eth1/1/1,1/1/3,1/1/5,1/1/7
                                              T Po1000
@  1633  Active    sfo01-w01-vsan             T Eth1/1/1,1/1/3,1/1/5,1/1/7
                                              T Po1000
@  3939  Active    sfo01-w01-vxrail_interna   T Eth1/1/1,1/1/3,1/1/5,1/1/7
                                              T Po1000
   4000  Active    L3-VLT-peer                T Po1000
   4094  Active                               T Po1000
```

6.3.2 show nve remote-vtep

The `show nve remote-vtep` command displays information about remote VXLAN tunnel endpoints. The command shows the state, encapsulation type, and the associated VNIs. The `State` should list `up`.

```
sfo01-leaf03# show nve remote-vtep
IP Address: 10.222.222.4,  State: up, Encap: VxLAN
   VNI list: 10,20,1631-1633,3939
IP Address: 10.222.222.64,  State: up, Encap: VxLAN
   VNI list: 10,20,1631-1633,3939
```

Note: The remote VTEP 10.222.222.64 is used to access external sources outside of the VXLAN tunnel. For more information routing between L2 VXLAN tunnels, see Appendix A.

6.3.3 show nve vxlan-vni

The `show nve vxlan-vni` command displays virtual-network configurations including the virtual network interface (VNI), the associated loopback adapter address (`Source-IP`), and available remote-VTEP address.

```
sfo01-leaf03# show nve vxlan-vni
VNI    Virtual-Network   Source-IP      Remote-VTEPs
-----------------------------------------------------
10     10                10.222.222.3   10.222.222.4,10.222.222.64
20     20                10.222.222.3   10.222.222.4,10.222.222.64
1631   1631              10.222.222.3   10.222.222.4,10.222.222.64
1632   1632              10.222.222.3   10.222.222.4,10.222.222.64
1633   1633              10.222.222.3   10.222.222.4,10.222.222.64
3939   3939              10.222.222.3   10.222.222.4
```

6.3.4 show virtual-network

The `show virtual-network` command displays members of the specified virtual-network, what source interface is being used to establish the VTEP endpoint, and what remote VTEPs the virtual-network terminates at.

```
sfo01-leaf03# show virtual-network 1631
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-
Unicast-Drop
Virtual Network: 1631
   Description: sfo01-w01-vxrail_external_management
   Members:
      Untagged: ethernet1/1/1, ethernet1/1/3, ethernet1/1/5, ethernet1/1/7
      VLAN 1631: port-channel1000
   VxLAN Virtual Network Identifier: 1631
      Source Interface: loopback0(10.222.222.3)
      Remote-VTEPs (flood-list): 10.222.222.4(DP),10.222.222.64(DP)

sfo01-leaf03# show virtual-network 3939
Codes: DP - MAC-learn Dataplane, CP - MAC-learn Controlplane, UUD - Unknown-
Unicast-Drop
Virtual Network: 3939
   Description: sfo01-w01-vxrail_internal_management
   Members:
      VLAN 3939: port-channel1000, ethernet1/1/1, ethernet1/1/3, ethernet1/1/5,
ethernet1/1/7
   VxLAN Virtual Network Identifier: 3939
      Source Interface: loopback0(10.222.222.3)
      Remote-VTEPs (flood-list): 10.222.222.4(DP)
```

6.4 VLT validation commands

The following commands verify that the VLT connection is operational. These commands can be used to verify any switch participating in a VLT domain.

6.4.1 show vlt all

The `show vlt all` command displays information for the configured VLT domain. The `Role` for one switch in the VLT pair is `primary`, and its peer switch is assigned the `secondary` role. The `VLTi Link Status` and `VLT Peer status` should both be `up`.

```
sfo01-leaf03# show vlt all
Domain ID                  : 1
Unit ID                    : 1
Role                       : primary
Version                    : 2.0
Local System MAC address   : 54:bf:64:c7:d4:c0
Role priority              : 32768
VLT MAC address            : 54:bf:64:c7:d4:c0
IP address                 : fda5:74c8:b79e:1::1
Delay-Restore timer        : 90 seconds
Peer-Routing               : Enabled
Peer-Routing-Timeout timer : 0 seconds
VLTi Link Status
    port-channel1000       : up

VLT System MAC Address  Status  IP Address           Version
------------------------------------------------------------
  2 54:bf:64:c7:e5:40    up     fda5:74c8:b79e:1::2   2.0
```

6.4.2 show vlt all backup-link

The `show vlt all backup-link` command displays that VLT peers are communicating on the backup link over the OOB management network. The `Destination` is the management IP address of the peer. The `Peer HeartBeat status` must be `Up`.

Note: The `Peer HeartBeat status` must be `Up`.

```
sfo01-leaf03# show vlt all backup-link
VLT Backup Link
------------------------
Destination                    : 100.67.167.34
Peer Heartbeat status          : Up
Heartbeat interval             : 30
Heartbeat timeout              : 90
Destination VRF                : default
```

6.4.3 show vlt all mismatch

The `show vlt all mismatch` command displays any configuration issues between the VLT peers.

Note: All items should indicate `No mismatch`.

```
sfo01-leaf03# show vlt all mismatch
VLT-MAC mismatch:
No mismatch

Peer-routing mismatch:
No mismatch

VLAN mismatch:
No mismatch

VLT VLAN mismatch:
No mismatch

VLT Virtual Network Mismatch:
Virtual Network Name Mismatch:
No mismatch

Virtual Network VLTi-VLAN Mismatch:
No mismatch

Virtual Network Mode Mismatch:
No mismatch

Virtual Network Tagged Interfaces Mismatch:
No mismatch

Virtual Network Untagged Interfaces Mismatch:
No mismatch

Virtual Network VNI Mismatch:
No mismatch

Virtual Network Remote-VTEP Mismatch:
No mismatch
```

6.4.4 show vlt mac-inconsistency

The `show vlt mac-inconsistency` command displays the inconsistencies in dynamic MAC addresses that are learned between VLT peers across spanned-VLANs or virtual networks.

Note: Verify that the `No inconsistencies found` status displays.

```
sfo01-leaf03# show vlt mac-inconsistency virtual-network

Inconsistency check for Virtual-Network based MAC
-------------------------------------------------

Fetching VN MACs from unit 2

Fetching VN MACs from unit 1

Comparing MACs of VLT Peers ..

No inconsistencies found
```

7 Perform initialization to create VxRail cluster

This guide does not provide detailed steps to initialize the VxRail cluster. The information provided is a list of the general steps that are used to initialize the multirack cluster using two different VxRail nodes, four E-series nodes, and four P-series nodes. Virtual Extensible LAN (VXLAN) allows these eight nodes, while in separate subnets in the IP underlay, to communicate as if connected through a single broadcast domain. This results in a deployment that can use existing documentation without any modification to deployment steps.

7.1 VxRail initialization

1. Install the VxRail nodes, by model, into the two racks in the data center.

Note: In this document, Rack 1 houses four E-series nodes, where Rack 2 houses four P-series nodes. For ease of manageability, each rack contains a pair of Dell EMC Networking S5248F-ON switches.

2. Attach the appropriate cabling between the ports of the VxRail nodes and the switch ports.
3. Power on the four primary E-series nodes in Rack 1 to form the initial VxRail cluster.

Note: Do not turn on the nodes in Rack 2.

4. To access the VxRail external management on VLAN 1631, connect a workstation or laptop that is configured for VxRail.
5. Using a web browser, go to the VxRail default IP address to begin the VxRail initialization process.
6. Complete the steps provided within the initialization wizard.

Note: The values used to validate this guide are found in Appendix C.

Using the values provided, VxRail performs the verification process. Once the validation is complete, the initialization process builds a new VxRail cluster. The building progress of the cluster displays in the status window provided. When the **Hooray!** message displays, the VxRail initialization is complete and the new VxRail cluster is built.

7. Click the **Manage VxRail** button to continue to VxRail management.
8. Power on the P-series VxRail nodes in Rack 2. The VxRail management page displays the four nodes once discovered across the VXLAN/VLAN associated to the VxRail internal management VLAN (3939).

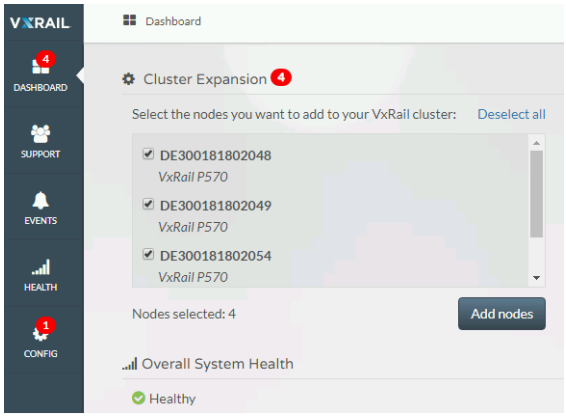

Figure 14 VxRail Manager showing four available P series nodes

9. Using the on-screen prompts, complete the **Cluster Expansion** wizard.

Note: In this document, three of the four nodes were used, however, all four can be selected in a typical cluster expansion scenario.

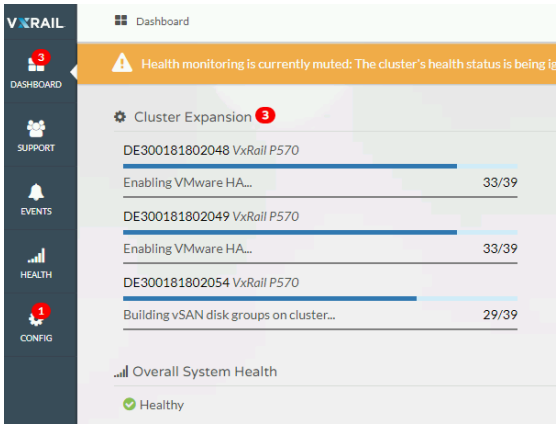

Figure 15 VxRail manager cluster expansion

7.2 VxRail validation

Once the deployment and expansion of the initial cluster are completed, use the VxRail Manager to verify that all eight hosts are in the single cluster.

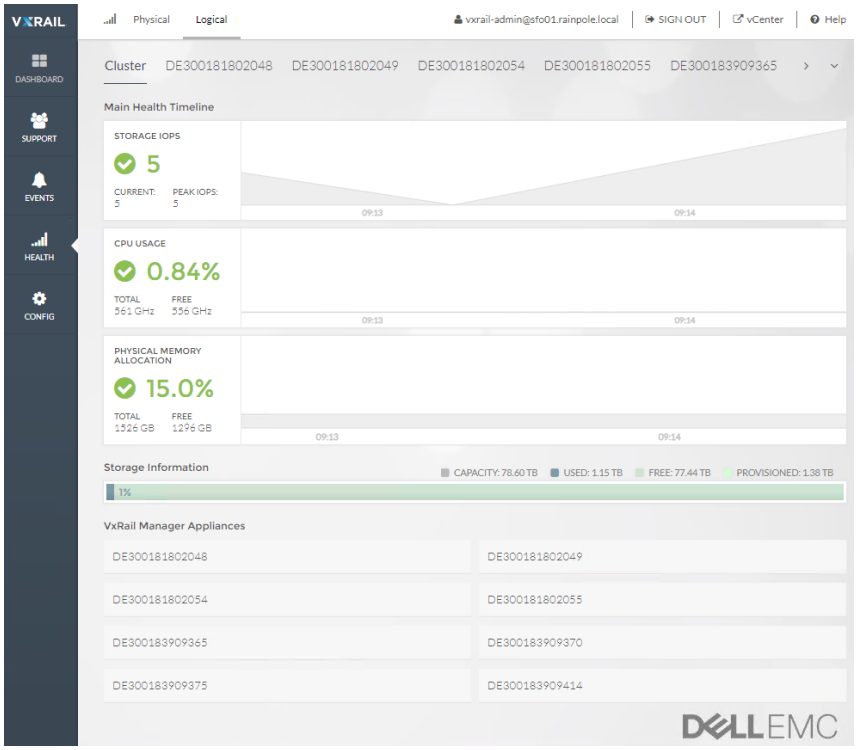

Figure 16 VxRail Manager logical view

8 VMware vSAN stretched clusters

Stretched clusters extend a vSAN cluster from a single site (geographical location, floors, or racks), to two sites for a higher level of availability. Stretched clusters are typically deployed in environments where the distance between the sites is limited, such as metropolitan or campus environments.

In this Deployment Guide, VMware vSAN stretched clusters is deployed using the two racks containing the eight total VxRail nodes. Each rack is configured as a separate fault domain achieving rack failure tolerance. Figure 17 shows a diagram illustrating the flow between the three locations: Rack 1, Rack 2, and Rack 3. Rack 3 contains the vSAN witness appliance, whose purpose is to host the witness components of virtual machine objects. Rack 1 and Rack 2 have already been configured using VXLANs stretched the Layer 2 (L2) domains between the two racks. Traffic to the vSAN witness appliance is done through L3 routing between the two data centers.

Note: For more information about how routing between tunnels is configured, see Appendix A.

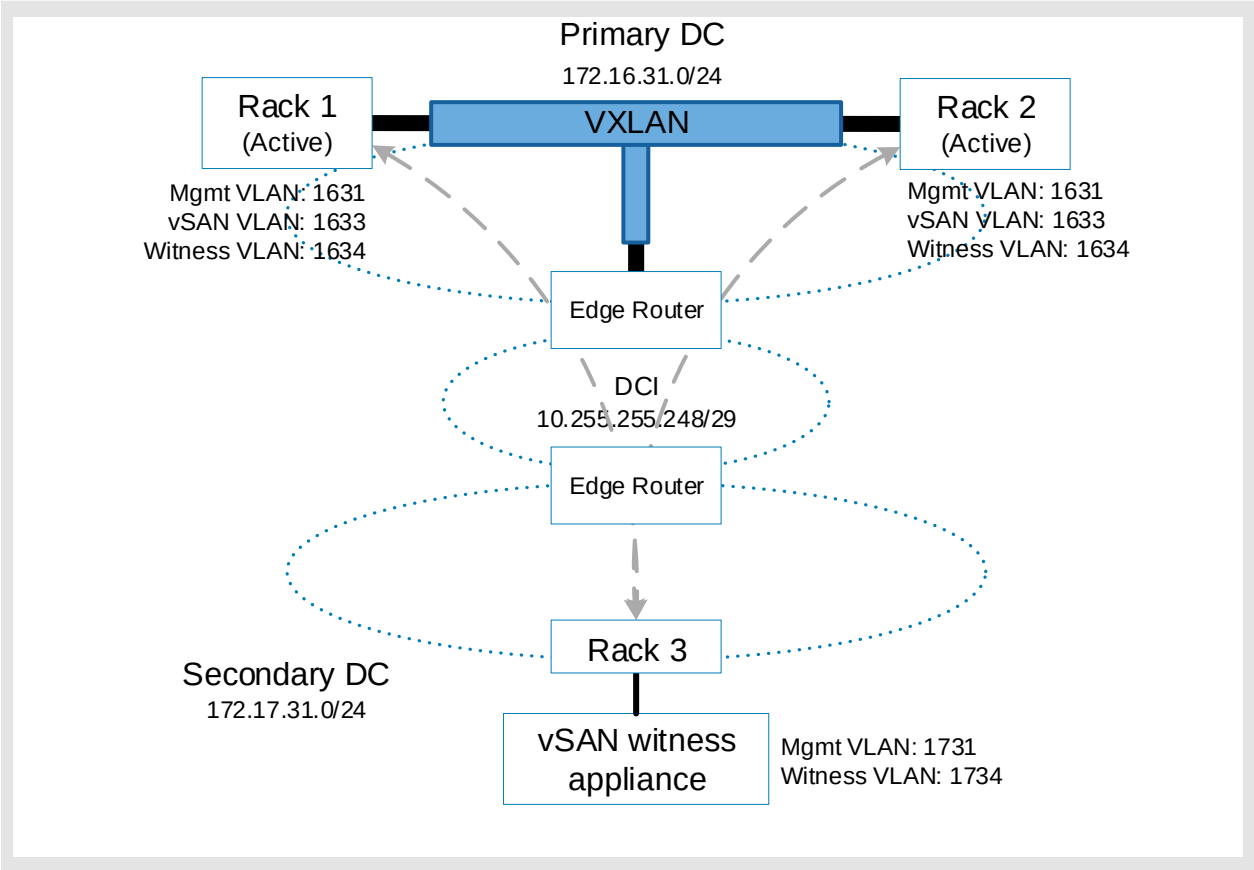

Figure 17 vSAN stretched cluster network diagram

Note: The vSAN witness appliance requires available through routing. The witness appliance can be located in the Primary data center, depending on requirements.

8.1 Deploying VMware vSAN stretched clusters

This section covers the values that are used to deploy vSAN stretched clusters in the environment that is configured in this Deployment Guide.

Note: For step-by-step instructions on deploying the vSAN witness appliance, see Deploying a *vSAN Witness Appliance*.

Deploy the vSAN witness appliance

The vSAN witness appliance is used. Figure 18 shows the vSAN witness appliance `sfo03w01vsanw01` is deployed in the secondary data center to an existing VMware cluster, `lax01-m01-mgmt01`.

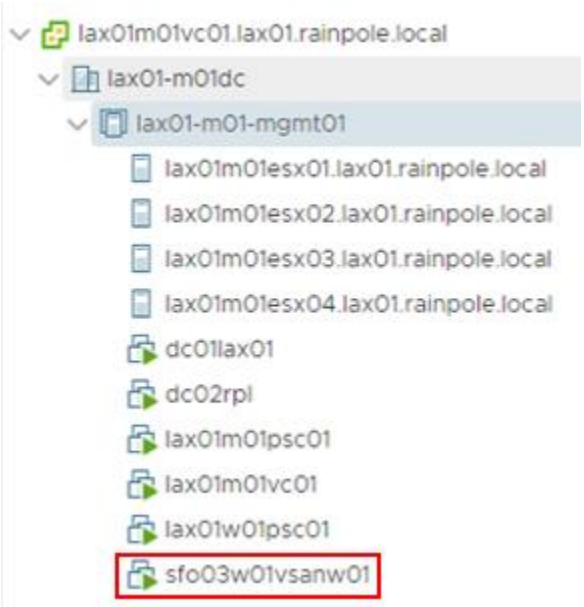

Figure 18 vSAN witness appliance

Table 4 shows the values that are used during the vSAN witness appliance deployment.

Table 4 vSAN witness appliance configuration details

VLAN	Network label	IP address	Purpose
1731	Management Network	172.17.31.201/24	ESXi management VMkernel
1733	witnessPg	172.17.33.201/24	vSAN VMkernel

Once deployed `sfo03w01vsan01` is then added to the vCenter data center hosting the exiting VxRail cluster. Figure 19 shows the appliance after it is added to the sfo01-w01dc data center object using the Management Network IP address. The icon for a vSAN witness appliance is blue, which represents that it is an imbedded ESXi appliance and not a physical server.

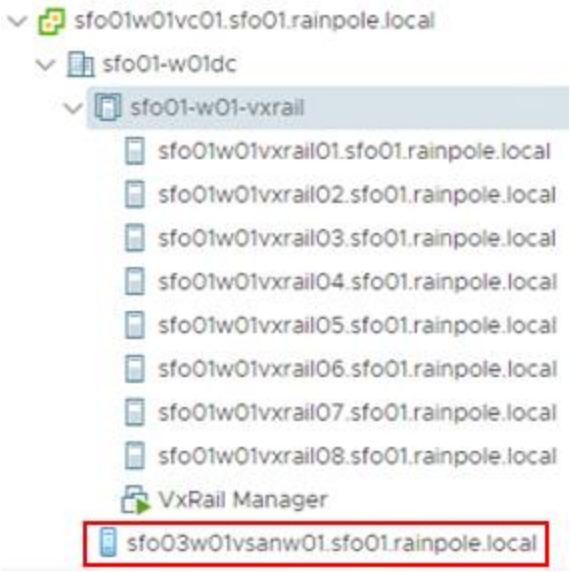

Figure 19 vSAN witness appliance added as host to data center sfo01-w01dc

Configure vSAN witness appliance routing

For the vSAN witness appliance, configure a static route to use the vSAN VMkernel interface to reach the VxRail management VMkernel in Rack 1 and Rack 2 in the primary data center.

1. Open an SSH session to the vSAN Witness Appliance, `sfo03w01vsanw01`.
2. Use the `esxcli network ip route ipv4 add command` to specific a static route to the VxRail management network (`172.16.31.0/24`) using the next hop address for second data center vSAN network (`172.17.33.0/24`).

   ```
   [root@sfo03w01vsanw01:~] esxcli network ip route ipv4 add -n
   172.16.31.0/24  -g 172.17.33.253
   ```

3. Use the `vmkping` command to verify reachability from the vSAN Witness Appliance vmk1 interface to the management IP address for the first VxRail node.

   ```
   [root@sfo03w01vsanw01:~] vmkping -I vmk1 172.16.31.101
   PING 172.16.31.101 (172.16.31.101): 56 data bytes
   64 bytes from 172.16.31.101: icmp_seq=0 ttl=60 time=0.641 ms
   64 bytes from 172.16.31.101: icmp_seq=1 ttl=60 time=0.807 ms
   64 bytes from 172.16.31.101: icmp_seq=2 ttl=60 time=0.747 ms

   --- 172.16.31.101 ping statistics ---
   3 packets transmitted, 3 packets received, 0% packet loss
   round-trip min/avg/max = 0.641/0.732/0.807 ms
   ```

Configure network interfaces for witness traffic

In VMware vSAN 6.7U1 and later, witness traffic can be separated from data traffic. vSAN data traffic requires a low-latency, high-bandwidth link. Witness traffic can use a high-latency, low-bandwidth, and routable link. To separate data traffic from witness traffic, configure a dedicated VMkernel network adapter for vSAN witness traffic.

1. Open an SSH connection to the first VxRail node, for example `sfo01w01vxrail01`.
2. Use the `esxcli network ip interface ipv4 get` command to determine which VMkernel network adapter is used for management traffic. In this Deployment Guide management traffic for the VxRail nodes uses the 172.16.31.0/24 subnet.

```
[root@sfo01w01vxrail01:~] esxcli network ip interface ipv4 get
Name  IPv4 Address   IPv4  Netmask   IPv4 Broadcast   Address Type  Gateway
----  -------------  ------------  --------------  ------------  ------
vmk2  172.16.31.101  255.255.255.0  172.16.31.255  STATIC        0.0.0.0
vmk0  0.0.0.0        0.0.0.0        0.0.0.0        NONE          0.0.0.0
vmk3  172.16.33.101  255.255.255.0  172.16.33.255  STATIC        0.0.0.0
vmk4  172.16.32.101  255.255.255.0  172.16.32.255  STATIC        0.0.0.0
vmk1  169.254.0.2    255.255.255.0  169.254.0.255  STATIC        0.0.0.0
```

3. Use the esxcli vsan network ip add command to configure the management VMkernel network adapter to support witness traffic.

```
[root@sfo01w01vxrail01:~] esxcli vsan network ip add -i vmk2 -T=witness
```

4. Use the `esxcli vsan network list` command to verify the new network configuration. Traffic Type should list `witness`.

```
[root@sfo01w01vxrail01:~] esxcli vsan network list | sed '1,/vmk2/d'
IP Protocol: IP
Interface UUID: 0c2a425c-ae19-5c96-dc9d-000af7f61ba0
Agent Group Multicast Address: 224.2.3.4
Agent Group IPv6 Multicast Address: ff19::2:3:4
Agent Group Multicast Port: 23451
Master Group Multicast Address: 224.1.2.3
Master Group IPv6 Multicast Address: ff19::1:2:3
Master Group Multicast Port: 12345
Host Unicast Channel Bound Port: 12321
Multicast TTL: 5
Traffic Type: witness
```

5. Repeat steps 1 through 4 for all remaining VxRail nodes.

Figure 20 shows the vSphere Client and the management VMkernel network interface is now enabled for vSAN Witness traffic but not vSAN traffic. Do no enable the interface for vSAN traffic in the vSphere client.

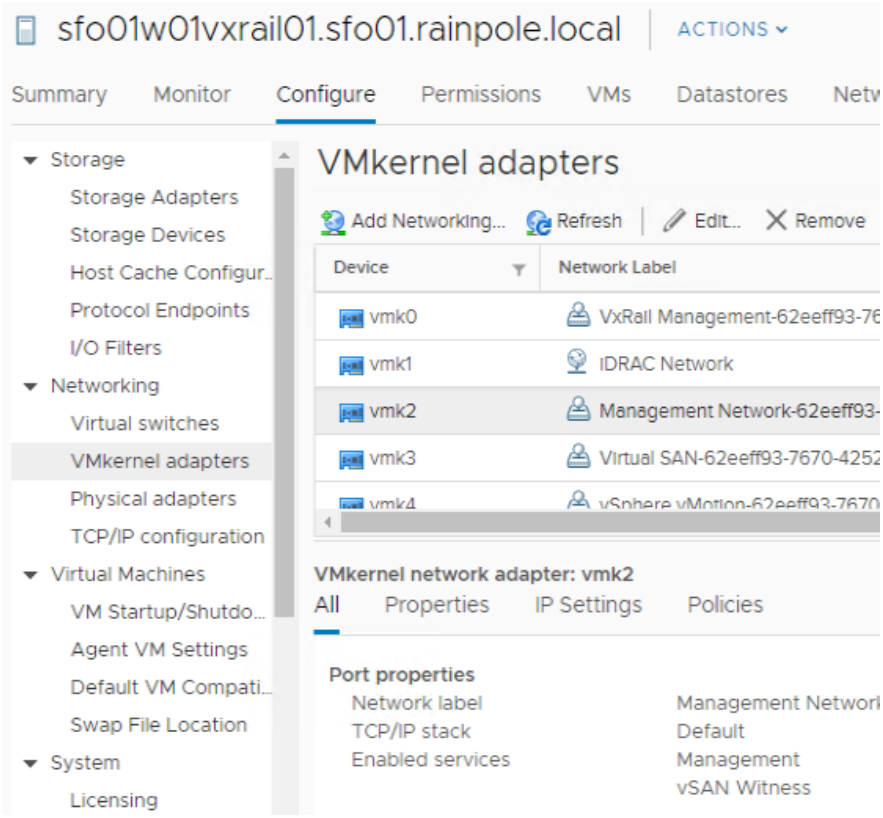

Figure 20 vSAN witness traffic as an enabled VMkernel service

Configure fault domains

In this section, the two fault domains that are required for vSAN stretched clusters are created. Table 5 shows all VxRail nodes in Rack 1 are placed in the Preferred fault domain. The table also shows all nodes in Rack 2 are placed in the Secondary fault domain.

Table 5 Fault domains

Fault domain	Rack ID	Node membership
Preferred	1	sfo01w01vxrail01, sfo01w01vxrail02, sfo01w01vxrail03, sfo01w01vxrail04
Secondary	2	sfo01w01vxrail05, sfo01w01vxrail06, sfo01w01vxrail07, sfo01w01vxrail08

Table 6 shows the vSAN witness appliance disk settings that are selected during the fault domain creation wizard. These settings are specific to the vSAN witness appliance.

Table 6 Witness appliance settings

Setting	Value
Local VMware Disk (mpx.vmhba1:C0:T2:L0)	cache tier
Local VMware Disk (mpx.vmhba1:C0:T1:L0)	capacity tier

Note: For step-by-step instructions on deploying fault domains, see *Manually Configure vSAN Stretched Cluster*.

8.2 Witness component validation

To verify vSAN component placement for a virtual machine, complete the following steps:

1. Open and login to the vSphere Web Client
2. From the Home menu, select Hosts and Clusters and expand the sfo01m01vc01.sfo01.rainpole.local tree.
3. Select the **sfo01-w01-vxrail** cluster and click the **Monitor** tab.
4. Under vSAN, click Virtual Objects.
5. Select the **VxRail Manager** virtual machine object.
6. Click View Placement Details.

Figure 21 shows the VxRail Manager virtual machines virtual objects. This view shows virtual object Hard disk 1 using the default RAID 1 (mirror) vSAN storage policy. Components are mirrored between VxRail node `sfo01w01vxrail01` in the Preferred fault domain and `sfo01w01vxrail06` in the Secondary fault domain, and the witness component is located on `sfo03w01vsanw01`.The remaining three objects are configured similar with some components being located on alternative hosts in the Secondary fault domain.

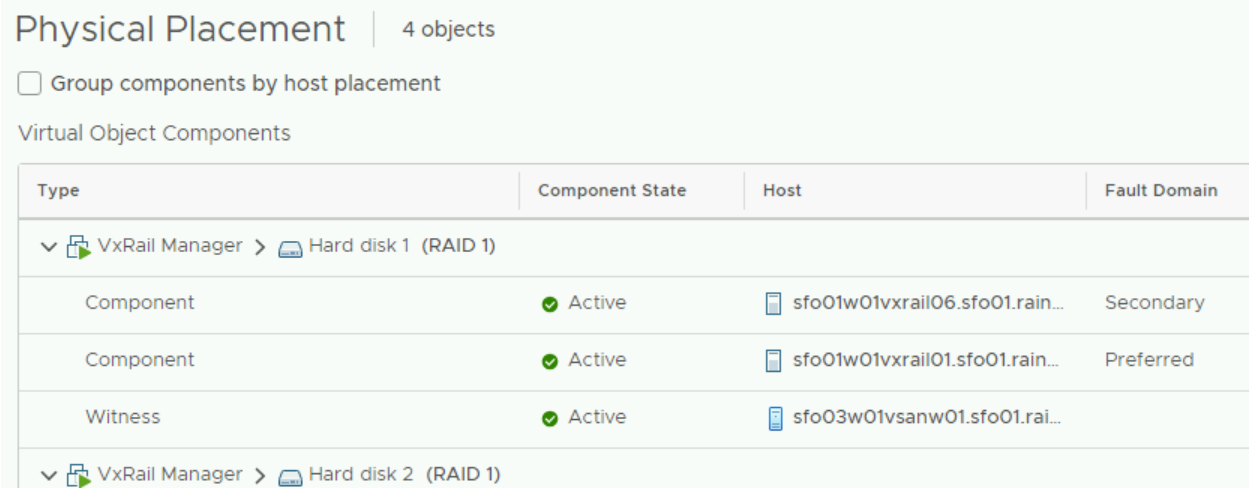

Figure 21 VxRail Manager virtual object placement

A Routing in a Layer 2 VXLAN overlay

In this document, the VxRail nodes are deployed using `sfo01w01vc01.sfo01.rainpole.local` as an external vCenter server. The vCenter server is part of an existing management cluster that is encapsulated in a separate VXLAN tunnel, virtual-network 1611. The VxRail uses virtual-network 1631 for external management which connects the hosts to the vCenter server. Figure 22 shows the two virtual networks. Both virtual-networks (shown as dashed lines) are statically configured on switch `sfo-edge01`. In this diagram `sfo01w01` represents all eight nodes in VLAN 1631.

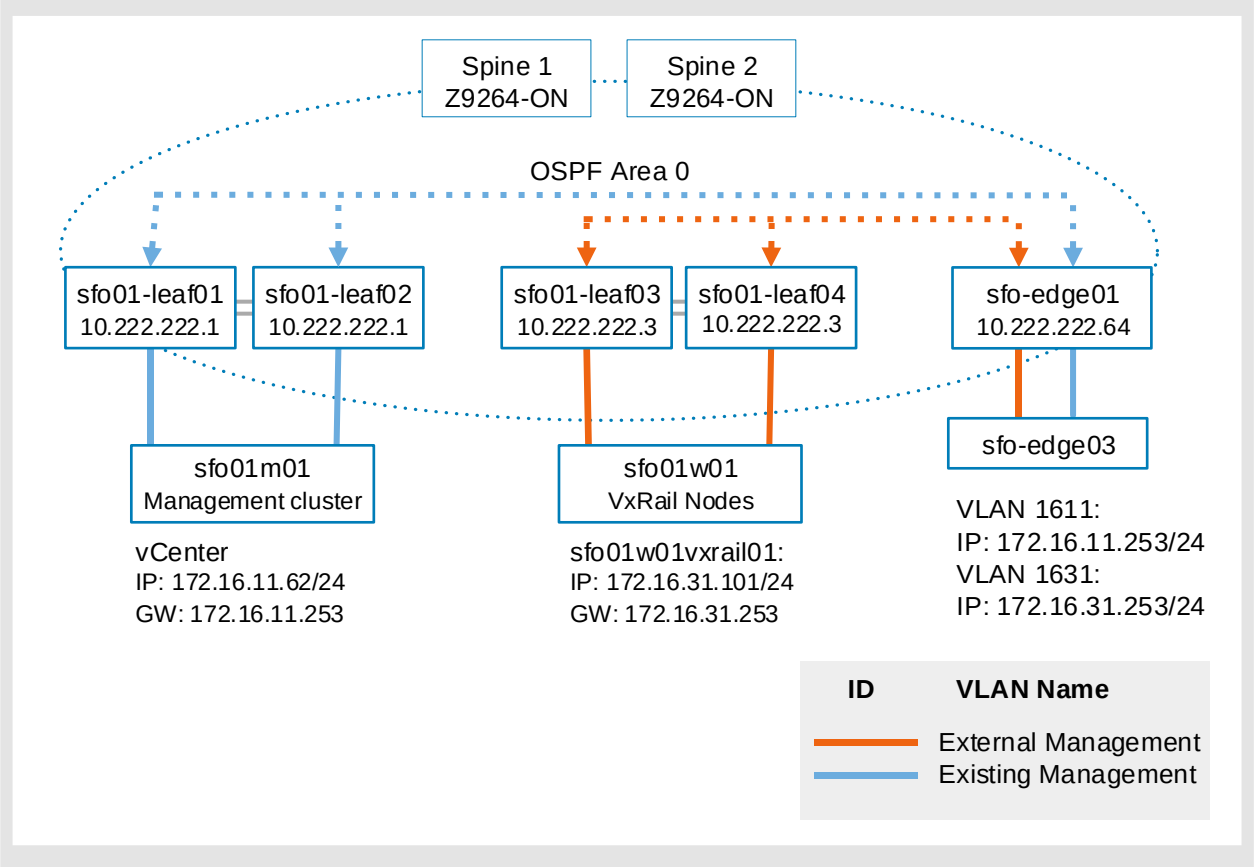

Figure 22 Routing between L2 VXLAN virtual networks

The vCenter server `sfo01w01vc01.sfo01.rainpole.local` has the IP address `172.16.11.62/24` and the first VxRail node, `sfo01w01vxrail01` has the IP address `172.16.31.101/24`. In the current implementation of a static L2 VXLAN, routing in and out of tunnels (RIOT) is not supported. Additionally, a VLAN participating in a virtual network cannot be configured as an Layer 3 (L3) interface.

To allow two hosts in two separate VXLANs to communicate, the VXLAN header must be removed so that inter-VLAN routing can take place. This process allows the payload to be encapsulated using the destination VXLAN. To achieve this, the network for both IP addresses is defined on `sfo-edge03`, which in this example, is a Dell Networking S4048-ON. This switch has an interface on each VNI and performs the inter-VLAN routing. Traffic is then re-encapsulated by `sfo-edge01` to reach the target destination

Note: Router redundancy is not shown in this example for simplicity. Best practices for router redundancy should be considered during deployment.

The command `show nve remote-vtep` is displayed to show the list of VNIs.

```
sfo-edge01# show nve remote-vtep
IP Address: 10.222.222.1,  State: up, Encap: VxLAN
   VNI list: 1611
IP Address: 10.222.222.3,  State: up, Encap: VxLAN
   VNI list: 10,20,1631-1633
IP Address: 10.222.222.4,  State: up, Encap: VxLAN
   VNI list: 10,20,1631-1633
```

The `sfo-edge01` and `sfo-edge03` switches connect using a port channel that carries the VLAN that requires routing. The `show vlan` command shows the VLAN membership as well as the virtual network association.

```
sfo-edge01# show vlan | except default
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs,
       @ - Attached to Virtual Network
Q: A - Access (Untagged), T - Tagged
    NUM    Status    Description                    Q Ports
@   10     Active    sfo01-w01-guest_vm_network_a   T Po10
@   20     Active    sfo01-w01-guest_vm_network_b   T Po10
@   1611   Active    sfo01-m01-esxi-management      T Po10
@   1631   Active    sfo01-w01-vxrail_external_ma   T Po10
@   1632   Active    sfo01-w01-vmotion              T Po10
@   1633   Active    sfo01-w01-vsan                 T Po10
```

B Maximum Transmission Unit (MTU) considerations

Virtual extensible LAN (VXLAN), as a network virtualization overlay (NVO) technology, places some additional overhead on the original data frames. This overhead manifest in the form of additional identifier information required for the functioning of VXLAN itself. As a rule, fragmentation, splitting of a frame or data packet because it is too large for a transport network, should be avoided. Fragmentation and reassembly put an additional burden on the network switches and VxRail node resources, which result in transport inefficiency.

Ethernet has a standard maximum transmission unit (MTU) of 1500 bytes. The total size of an Ethernet frame is 1518, or 1522, with the additional optional 802.1q tag, with six bytes each for the source and destination MAC address, two bytes for the Ethertype, and four bytes for the frame check sequence (FCS). This means that a server can send a payload of 1500 byte or less, which includes all the header information, from Layer 3 and above. VXLAN reduces the MTU to 1450 bytes through its default configuration. This is because VXLAN adds either 50 or 54 bytes of overhead as part of identifier information. Figure 23 shows a standard 54-byte VXLAN packet with an 802.1Q tag.

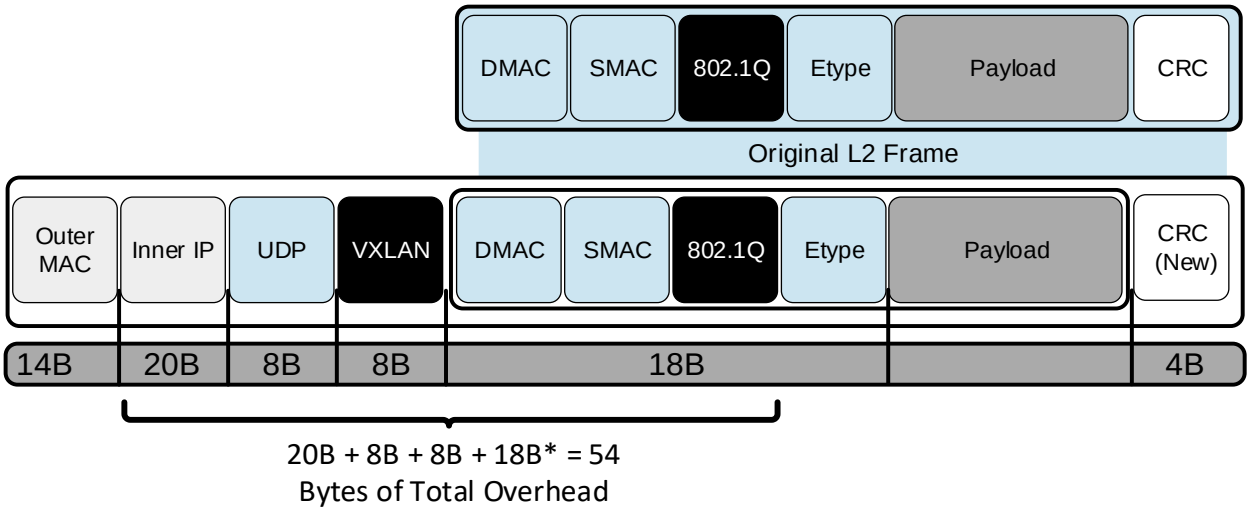

Figure 23 VXLAN Frame Format, Including Byte Count (50/54 Byte Overhead for Comparison)

VxRail configures virtual switches to accommodate jumbo frames, data units that are larger than the traditional 1500 bytes typical of Ethernet frames, and in the case of vSAN 9000 bytes. With a need to support MTUs of 9000 bytes from the VxRail nodes, and with the additional bytes from the VXLAN encapsulation, a value around 9050 or 9054 bytes needs to be allowed to avoid fragmentation.

Dell EMC Networking provides a maximum MTU value of 9216 bytes. Jumbo frame MTU of 9000 bytes from the VxRail node can be accommodated by the network switches without introducing fragmentation. The jumbo frame MTU needs to be configured consistently across all points in the network, including the ESXi virtual switches.

C VxRail deployment values

Table 7 lists the values used during the VxRail Manager initialization and expansion operation.

Note: The values are listed in order as they are entered in the GUI.

Table 7 VxRail network configuration values

	Parameter	Value
Appliance Settings	NTP server	172.16.11.251
	Domain	sfo01.rainpole.local
ESXi hostname and IP addresses	ESXi host name prefix	sfo01w01vxrail
	Separator	none
	Iterator	Num 0x
	Offset	1
	Suffix	none
	ESXi starting address	172.16.31.101
	ESXi ending address	172.16.31.165
Customer Supplied vCenter Server	vCenter Server hostname (FQDN)	sfo01w01vc01.sfo01.rainpole.local
	admin username/password	vxrail-admin@sfo01.rainpole.local
	vCenter Server SSO domain	rainpole.local
	Data center name	sfo01-w01dc
	New cluster name	sfo01-w01-vxrail01
Platform Services Controller	PSC Hostname	sfo01psc0101.sfo01.rainpole.local
VxRail Manager	VxRail Manager hostname	sfo01w01vxrail-mgr
	VxRail IP address	172.16.31.71
	Subnet mask	255.255.255.0
	Gateway	172.16.31.253
vMotion	Starting address for IP pool	172.16.32.101
	Ending address for IP pool	172.16.32.165
	Subnet mask	255.255.255.0
	VLAN ID	1632
vSAN	Starting address for IP pool	172.16.33.101
	Ending address for IP pool	172.16.33.165
	Subnet mask	255.255.255.0
	VLAN ID	1633
VM Networks	Guest_VM_Network_A	
	Guest_VM_Network_B	
Solutions Logging	vRealize Log Insight hostname	

D Connecting a workstation or laptop for VxRail initialization

A workstation/laptop with a web browser for the VxRail user interface is required to perform the initialization process. It must be either plugged into the top-of-rack switch, or able to logically reach the VxRail external management VLAN from elsewhere in the network. The steps below provide guidance on configuring a S5248F-ON to work with a 1 GbE connection from a workstation/laptop.

1. Use the `show inventory media` command to verify that the appropriate SFP-1GbE Base-T adapter is installed.

```
sfo02-leaf03# show inventory media | grep SFP
1/1/1  SFP28      SFP28 25GBASE-SR-NOF            MY013602883003F  true
1/1/3  SFP28      SFP28 25GBASE-SR-NOF            MY01360288400W6  true
1/1/5  SFP28      SFP28 25GBASE-SR-NOF            MY013602883003H  true
1/1/7  SFP28      SFP28 25GBASE-SR-NOF            MY013602883003T  true
1/1/32 SFP        SFP 1000BASE-T                  PY20PML          true
1/1/49 QSFP28-DD  QSFP28-DD 2x100GBASE-2CR4-0.5M  CNFCK0077K0016   true
1/1/50 QSFP28-DD  QSFP28-DD 2x100GBASE-2CR4-0.5M  CNFCK0077K0016   true
1/1/51 QSFP28-DD  QSFP28-DD 2x100GBASE-2CR4-0.5M  CNFCK0077K0017   true
1/1/52 QSFP28-DD  QSFP28-DD 2x100GBASE-2CR4-0.5M  CNFCK0077K0017   true
1/1/53 QSFP28     QSFP28 100GBASE-SR4             CN01360385M001F  true
1/1/54 QSFP28     QSFP28 100GBASE-SR4             CN01360364N000M  true
```

2. The show port-group command is used to identify which port-group the SFP-1GbE Base-T adapter belongs to. Here port-group 1/1/8 contains physical interfaces ethernet 1/1/29 – 1/1/32.

```
sfo02-leaf03# show port-group
Port-group            Mode           Ports      FEM
port-group1/1/1       Eth 25g-4x     1  2  3  4     -
port-group1/1/2       Eth 25g-4x     5  6  7  8     -
port-group1/1/3       Eth 10g-4x     9  10  11  12     -
port-group1/1/4       Eth 25g-4x     13  14  15  16     -
port-group1/1/5       Eth 25g-4x     17  18  19  20     -
port-group1/1/6       Eth 25g-4x     21  22  23  24     -
port-group1/1/7       Eth 25g-4x     25  26  27  28     -
port-group1/1/8       Eth 25g-4x     29  30  31  32     -
port-group1/1/9       Eth 25g-4x     33  34  35  36     -
port-group1/1/10      Eth 25g-4x     37  38  39  40     -
port-group1/1/11      Eth 25g-4x     41  42  43  44     -
port-group1/1/12      Eth 25g-4x     45  46  47  48     -
port-group1/1/13      Eth 100g-2x    49  50     -
port-group1/1/14      Eth 100g-2x    51  52     -
port-group1/1/15      Eth 100g-1x    53        -
port-group1/1/16      Eth 100g-1x    54        -
port-group1/1/17      Eth 100g-1x    55        -
port-group1/1/18      Eth 100g-1x    56        -
```

3. Use the `port-group` command `mode` to set all physical interfaces in the port group to 10 GbE.

```
sfo02-leaf03# configure
sfo02-leaf03(config)# port-group 1/1/8
sfo02-leaf03(conf-pg-1/1/8)# mode Eth 10g-4x
sfo02-leaf03(conf-pg-1/1/8)# exit
```

4.Configure the breakout interface with the appropriate management VLAN.

```
sfo02-leaf03(config)# interface ethernet 1/1/32:1
sfo02-leaf03(conf-if-eth1/1/32:1)# switchport access vlan 1631
```

5. Use `show interface` to validate that the interface is operating at the correct speed.

```
sfo02-leaf03# show interface status | grep up
Port       Description  Status  Speed  Duplex  Mode Vlan Tagged-Vlans
Eth 1/1/1  sfo01w01..   up      25G    full    T    1631 10,20,1632-1633,3939
Eth 1/1/3  sfo01w01..   up      25G    full    T    1631 10,20,1632-1633,3939
Eth 1/1/5  sfo01w01..   up      25G    full    T    1631 10,20,1632-1633,3939
Eth 1/1/7  sfo01w01..   up      25G    full    T    1631 10,20,1632-1633,3939
Eth 1/1/32:1            up      1000M  full    A    1631 –
```

E Validated components

This section contains tables detailing the specific firmware and software revisions used in this guide.

E.1 Dell EMC Networking switches

Table 8 Switches and OS versions

Qty	Item	Version
4	Dell EMC Networking S5248F-ON leaf switches	10.4.2.236
2	Dell EMC Networking Z9264F-ON spine switches	10.4.2.236
2	Dell EMC Networking S3048-ON OOB mgmt switches	10.4.2.236

E.2 VxRail E560 nodes

A cluster of four VxRail E560 nodes was used to validate the examples in this guide. The nodes were each configured using the information that is provided below.

Table 9 VxRail E560 node components

Qty per node	Item	Firmware version
2	Intel Xeon Gold 6136 CPU @ 3.00GHz, 12 cores	-
12	16GB DDR4 DIMMs (192GB total)	-
2	800GB SAS SSD	-
8	1.2TB SAS HDD	EF03
1	Dell HBA330 Storage Controller	15.17.09.06
1	Boot Optimized Storage Solution (BOSS) Controller w/ 2x240GB SATA SSDs	2.5.13.3016
1	Broadcom 57414 NDC – 2x25GbE SFP28 ports	20.08.04.04
-	BIOS	1.4.9
-	iDRAC with Lifecycle Controller	3.21.23.22

E.3 VxRail P570 nodes

A cluster of four VxRail E570 nodes was used to validate the examples in this guide. The nodes were each configured using the information that is provided above.

Table 10 VxRail E570 node components

Qty per node	Item	Firmware version
2	Intel Xeon Gold 6136 CPU @ 3.00GHz, 12 cores	-
12	16GB DDR4 DIMMs (192GB total)	-

Qty per node	Item	Firmware version
2	800GB SAS SSD	-
8	1.2TB SAS HDD	EA02
1	Dell HBA330 Storage Controller	15.17.09.06
1	Boot Optimized Storage Solution (BOSS) Controller w/ 2x240GB SATA SSDs	2.5.13.3016
1	Broadcom 57414 NDC – 2x25GbE SFP28 ports	20.08.04.04
-	BIOS	1.4.9
-	iDRAC with Lifecycle Controller	3.21.23.22

E.4 VxRail appliance software

Chapter 10 of this deployment guide was developed using VxRail appliance software 4.7.0. The software consists of the component versions that are provided below:

Table 11 VxRail appliance software component versions

Item	Version
Dell PTAgent	1.8.1-21
VMware ESXi	6.7.0 build 10764712
VxRail Manager	4.7.001 11073036
VMware vCenter Server Appliance	6.7.0 build 10244857
VxRail Manager VIB	4.7.001-11073036
VMware vSAN	vSAN 6.7 Update 1, provided with ESXi

F Technical resources

Dell EMC Networking Guides

OS10 Enterprise Edition Users Guide 10.4.2.0 Edition

Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE

Dell EMC VxRail Network Guide

VxRail Appliance 4.7 Administration Guide

VxRail Planning Guide for Virtual SAN Stretched Cluster

VMware Validated Design for Software-Defined Data Center 4.3 Release Notes

G Support and feedback

Contacting Technical Support

Support Contact Information	Web: http://www.dell.com/support
	Telephone: USA: 1-800-945-3355

Feedback for this document

To provide feedback on the quality and usefulness of this publication, send an e-mail to Dell_Networking_Solutions@Dell.com.