

Dell EMC Networking OS10 Enterprise Edition Switch Configuration Guide for VxRail 4.5

Deploying a pair of S5148F-ON switches for a VxRail 4.5 cluster

Abstract

This document provides Dell EMC Networking OS10 Enterprise Edition switch configuration examples and topology options for a VxRail 4.5 cluster deployment.

March 2019

Revisions

Date	Description
March 2019	Changed switch node-facing interface settings from “flowcontrol transmit on” to “flowcontrol transmit off” as a best practice.
September 2018	Initial release

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.

© 2018 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.....	5
1.1 Supported switches and operating systems.....	6
1.2 Typographical conventions.....	7
1.3 Attachments.....	7
2 Hardware overview.....	8
2.1 Dell EMC Networking S5148F-ON.....	8
2.2 Dell EMC Networking S3048-ON.....	8
2.3 Dell EMC VxRail 14G nodes.....	8
2.3.1 Dell EMC VxRail P Series node.....	8
2.3.2 Dell EMC VxRail V Series node.....	9
2.3.3 Dell EMC VxRail S Series node.....	9
2.3.4 Dell EMC VxRail E Series node.....	9
2.3.5 VxRail node network adapters.....	9
3 Topology options.....	11
3.1 Dual switch.....	11
3.1.1 Dual switch with VLT.....	11
3.1.2 Dual switch without VLT.....	12
3.2 Single switch.....	13
4 VxRail in the data center.....	14
4.1 Leaf-spine network.....	14
4.2 OOB management network.....	15
5 Configuration planning.....	16
5.1 VLANs and IP addresses.....	16
5.2 VxRail network configuration table.....	17
5.3 Using a jump box or laptop computer.....	18
5.4 DNS and NTP servers.....	18
6 Switch configuration prerequisites.....	20
6.1 Check switch OS version.....	20
6.2 Verify license installation.....	20
6.3 Factory default configuration.....	21
7 Configure switches.....	22
7.1 Dual switch with VLT.....	22
7.1.1 General settings.....	22
7.1.2 Configure VLANs.....	23

7.1.3	Configure interfaces	24
7.1.4	Configure VLT.....	25
7.2	Dual switch without VLT	26
7.2.1	General settings.....	26
7.2.2	Configure VLANs.....	27
7.2.3	Configure interfaces	28
7.2.4	Configure the switch interconnect	29
7.3	Single switch.....	30
7.3.1	General settings.....	31
7.3.2	Configure VLANs.....	31
7.3.3	Configure interfaces	32
8	Switch validation.....	35
8.1	General validation commands	35
8.1.1	show interface status	35
8.1.2	show port-channel summary	35
8.1.3	show vlan.....	36
8.1.4	show lldp neighbors	37
8.1.5	show vrrp brief	38
8.2	VLT validation commands	38
8.2.1	show vlt <i>domain_id</i>	38
8.2.2	show vlt <i>domain_id</i> backup-link.....	38
8.2.3	show vlt <i>domain_id</i> mismatch.....	39
9	Deploy VxRail.....	40
A	Validated components.....	41
A.1	Dell EMC Networking Switches	41
A.2	VxRail P570 nodes	41
A.3	VxRail Appliance software.....	42
B	VxRail network adapter traffic optimization	43
C	Technical resources	44
D	Support and feedback	45

1 Introduction

VxRail sits at the forefront of a fundamental shift in IT infrastructure consumption – away from application-specific, “build-your-own” infrastructure and toward virtualized, general-purpose, engineered systems. Dell EMC and VMware have embraced this shift with the VxRail hyper-converged appliance. VxRail has a simple, scale-out architecture that leverages VMware vSphere and VMware vSAN to provide server virtualization and software-defined storage.

To take full advantage of the VxRail solution, one must carefully consider the network that not only connects multiple nodes into a single, cohesive cluster but also enables connectivity to the customer’s IT environment. Numerous industry studies have shown that networking is the primary source of both deployment issues and poor performance of hyper-converged solutions. In most cases, VxRail clusters (minimum of three and maximum of 64 nodes) connect to a pre-existing IP network at the customer site. The inclusion of dedicated switches for the VxRail cluster simplifies this process and avoids many of the pitfalls associated with the deployment of a hyper-converged solution that originates in network connectivity.

The audience for this document includes professional services or onsite IT personnel responsible for the deployment of a VxRail cluster when a pair of dedicated Dell EMC Networking switches is purchased with the cluster. This document covers the process of connecting a cluster of VxRail nodes to:

- A pair of Dell Networking switches configured for Virtual Link Trunking (VLT). Using VLT is the preferred topology.
- A pair of Dell Networking switches not configured for VLT.
- A single Dell Networking switch.

This document provides switch topology options and configuration examples for a VxRail 4.5 cluster using nodes built on 14th generation (14G) PowerEdge servers. Nodes in these examples use 25GbE network adapters. Switches in this guide use Dell EMC Networking OS10 Enterprise Edition (OS10EE).

1.1 Supported switches and operating systems

The examples provided in this deployment guide use VxRail 4.5 nodes connected to Dell EMC Networking S5148F-ON switches running Dell Networking OS10EE.

Dell EMC Networking supports the following switch and OS combinations for VxRail 4.0 and later:

Dell EMC Networking Switches Supported for VxRail 4.0 and later releases*		
Dell EMC Networking Switch	Networking OS 9 ver 9.12.1.0 and later releases	Networking OS 10 ver 10.4.0R3 and later releases
S3048-ON	Supported	Supported
S3124 / S3124F / S3124P / S3148 / S3148P	Supported	
S4048F-ON / S4048T-ON	Supported	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	Supported	
S5148F-ON ¹		Supported
S6010-ON	Supported	Supported
S6100-ON	Supported	
Z9100-ON ²	Supported	Supported
Z9264F-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.

¹S5148F-ON is supported with OS10 ver 10.4.1.1 or later

²Z9100-ON is supported with 10GbE connections to nodes. Z9100-ON support for 25GbE connections to nodes is planned for future releases.

³Z9264-ON is supported with OS10 ver 10.4.1.0 or later

Figure 1 Supported Dell EMC Networking switches and operating systems

1.2 Typographical conventions

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

Monospace Text	CLI examples
<u>Underlined Monospace Text</u>	CLI examples that wrap the page
<i>Italic Monospace Text</i>	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

1.3 Attachments

This document in .pdf format 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 briefly describes the hardware used to validate this deployment. [Appendix A](#) contains a complete listing of hardware and software validated for this guide.

2.1 Dell EMC Networking S5148F-ON

The S5148F-ON is a 1-Rack Unit (RU), multilayer switch with 48x25GbE, and 6x100GbE or 72x25GbE ports. This guide uses two S5148F-ONs as leaf switches.

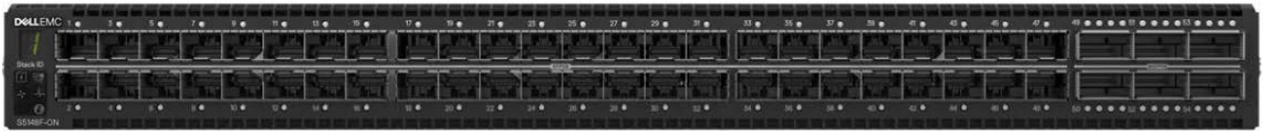


Figure 2 Dell EMC Networking S5148F-ON

2.2 Dell EMC Networking S3048-ON

The S3048-ON is a 1-RU switch with 48x1GbE BASE-T ports and 4x10GbE SFP+ ports. This guide uses one S3048-ON switch for out-of-band (OOB) management traffic.

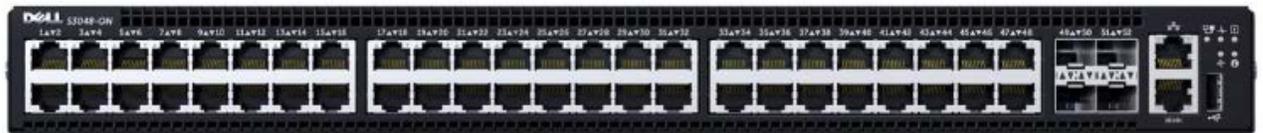


Figure 3 Dell EMC Networking S3048-ON

2.3 Dell EMC VxRail 14G nodes

Current Dell EMC VxRail P, V, S, and E Series nodes are built on 14th generation (14G) PowerEdge servers.

2.3.1 Dell EMC VxRail P Series node

VxRail P Series nodes are ideal for CPU-intensive workloads such as databases. P Series nodes support up to 44 CPU cores, 1536GB memory, and 24TB hybrid or 46TB all-flash storage in a 2-RU form factor.

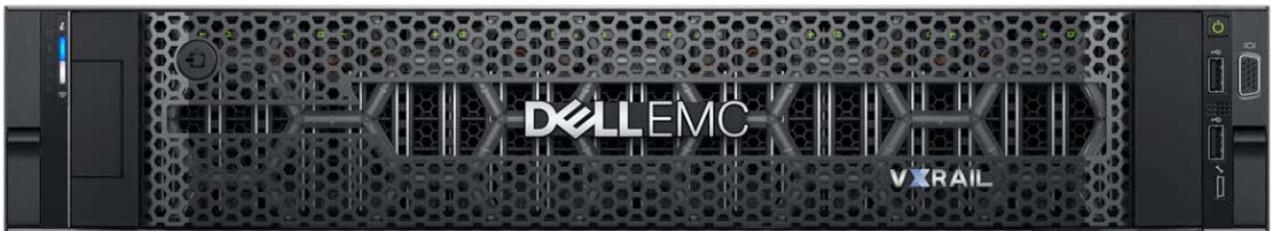


Figure 4 Dell EMC VxRail 2-RU node (P, V, or S Series)

Note: The deployment examples in this guide use a VxRail cluster consisting of four P570 nodes.

2.3.2 Dell EMC VxRail V Series node

VxRail V Series nodes are graphics-ready for uses such as high-end 2D/3D visualization. V Series nodes support up to 40 CPU cores, 1024GB memory, 3 GPUs, and 24TB hybrid or 46TB all-flash storage in a 2-RU form factor.

2.3.3 Dell EMC VxRail S Series node

VxRail S Series nodes provide expanded storage capacity for collaboration, data, and analytics. S Series nodes support up to 36 CPU cores, 1536GB memory, and 48TB hybrid storage in a 2-RU form factor.

2.3.4 Dell EMC VxRail E Series node

VxRail E Series nodes are best suited for remote office or entry workloads. 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.



Figure 5 Dell EMC VxRail E Series node

2.3.5 VxRail node network adapters

Each 14G VxRail node includes a rack server Network Daughter Card (rNDC) with one of the following port combinations:

- 2x25GbE SFP28 ports
- 2x10GbE SFP+ or BASE-T ports
- 4x10GbE SFP+ or BASE-T ports
- 4x1GbE BASE-T ports (for single CPU nodes only, max 8 nodes/cluster)

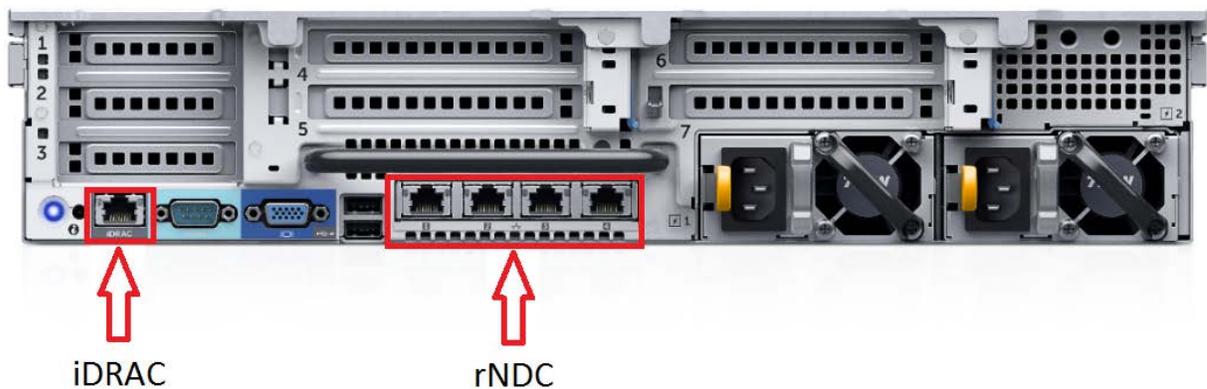


Figure 6 Rear view of VxRail 2-RU node (1-RU node is similar)

Note: Each of the VxRail P570 nodes in the deployment examples in this guide contains a Broadcom 57414 rNDC with 2x25GbE SFP28 ports.

VxRail optimizes network traffic by splitting it across rNDC uplinks and by using Network I/O control (NIOC) shares for different traffic types. See [Appendix B](#) for more information.

P, V, S, and E Series VxRail nodes also include a 1GbE BASE-T integrated Dell Remote Access Card (iDRAC) for OOB management.

Note: P, V, and S Series nodes support additional network adapters in PCIe slots. See the [Dell EMC VxRail Network Guide](#) for more information.

3 Topology options

VxRail may be deployed using a single or dual switch topology. Using a single switch provides the lowest initial cost, but creates a single point of failure. A dual switch configuration helps ensure high availability by eliminating this single point of failure.

A dual switch configuration may be used with or without VLT. Dell EMC recommends a dual switch configuration with VLT. The sections that follow explain the different options.

3.1 Dual switch

In a dual switch topology, each VxRail node has one or more connections to each of the two leaf switches providing redundancy at the NIC and switch levels. If VLT is used, the switches are connected with a VLT interconnect (VLTi). If VLT is not used, the switches are connected with a standard LACP port channel.

3.1.1 Dual switch with VLT

Dell EMC recommends using a dual switch configuration with VLT as shown in Figure 7.

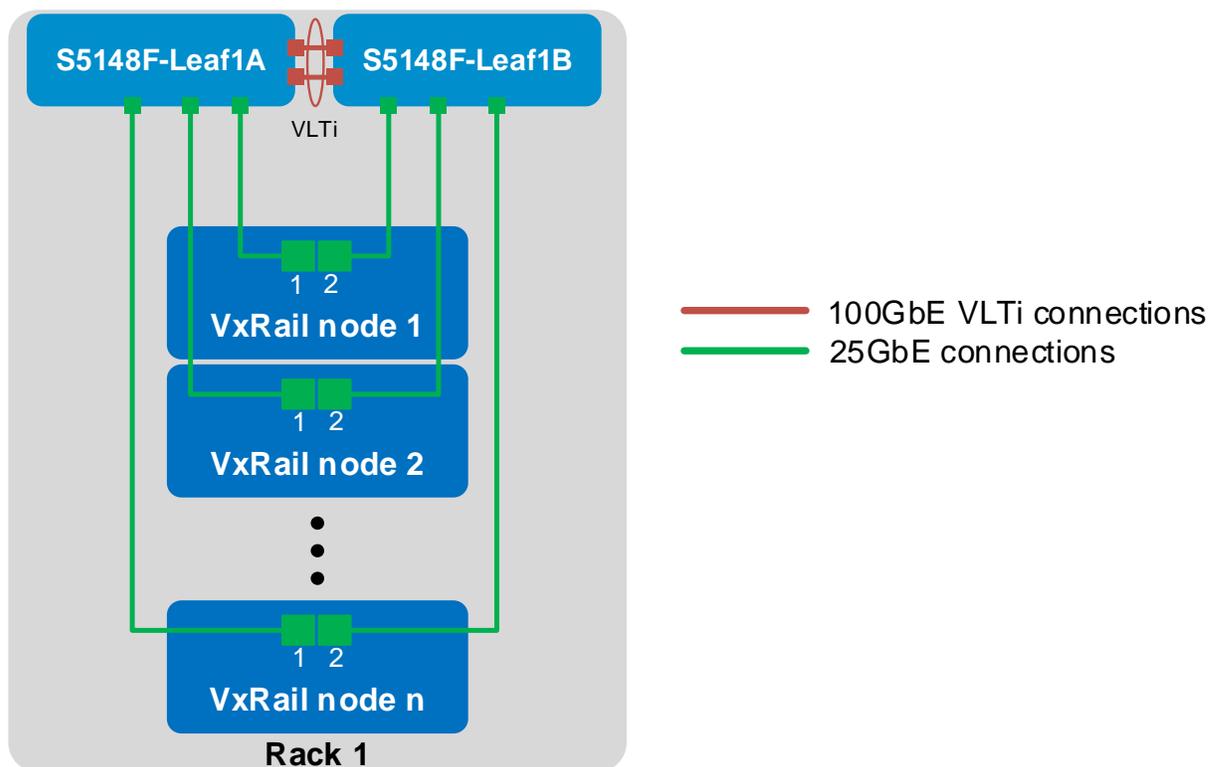


Figure 7 Dual switch topology with VLT

VLT synchronizes ARP, MAC tables, and IPv6 neighbors between the VLT peers and makes the switch pair appear as one logical unit to connected devices.

VLT provides the following benefits:

- Provides a loop-free topology and eliminates STP-blocked ports
- Optimizes the use of all available uplink bandwidth
- Guarantees fast convergence if either a link or a device fails

- Provides link-level resiliency
- Assures high availability
- Allows a single device to use a LAG across two upstream switches
- Provides Layer 2 multipathing

Note: While VxRail nodes use active and standby network adapters instead of LAGs, other servers in the rack can connect to the VLT switch pair with an LACP LAG for active/active Layer 2 multipathing. For more information on VLT, see the [OS10 Enterprise Edition User Guide Release 10.4.1.0](#).

3.1.2 Dual switch without VLT

The configuration of a switch pair without VLT is supported, but it does not provide the advantages of VLT covered in the previous section.

The switch pair is cabled as shown in Figure 8. The links connecting the two switches are configured in an LACP port channel.

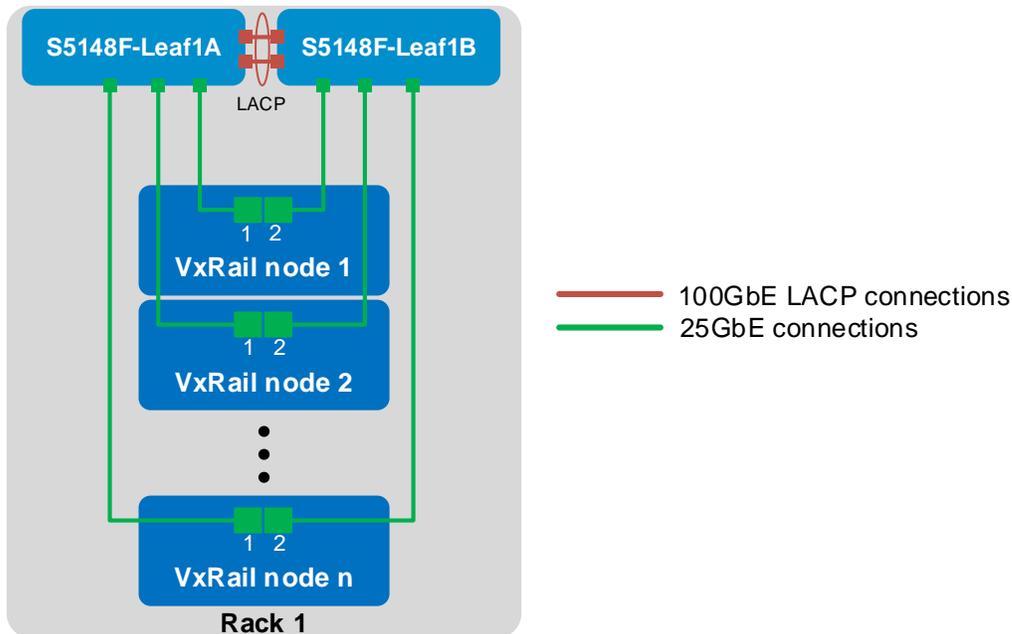


Figure 8 Dual switch topology without VLT

3.2 Single switch

In a single switch topology, all VxRail nodes connect to a single switch. This topology is lower initial cost, but it is not recommended as the switch becomes a single point of failure.

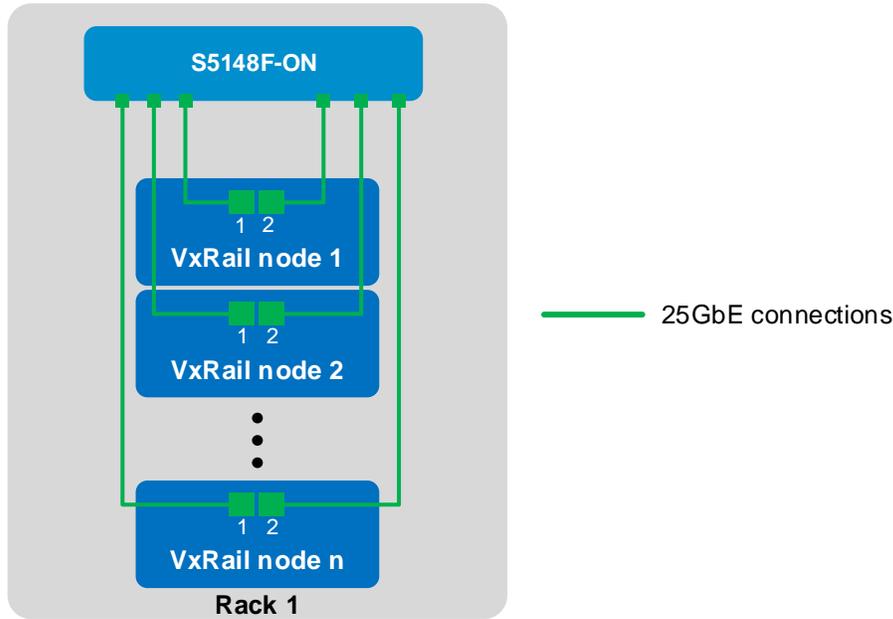


Figure 9 Single switch topology

4 VxRail in the data center

4.1 Leaf-spine network

Dell EMC recommends using a leaf-spine network in the data center with leaf switches configured as VLT peers. The switches and VxRail nodes covered in this guide are shown in Rack 1 in Figure 10 and are incorporated into a data center's leaf-spine network.

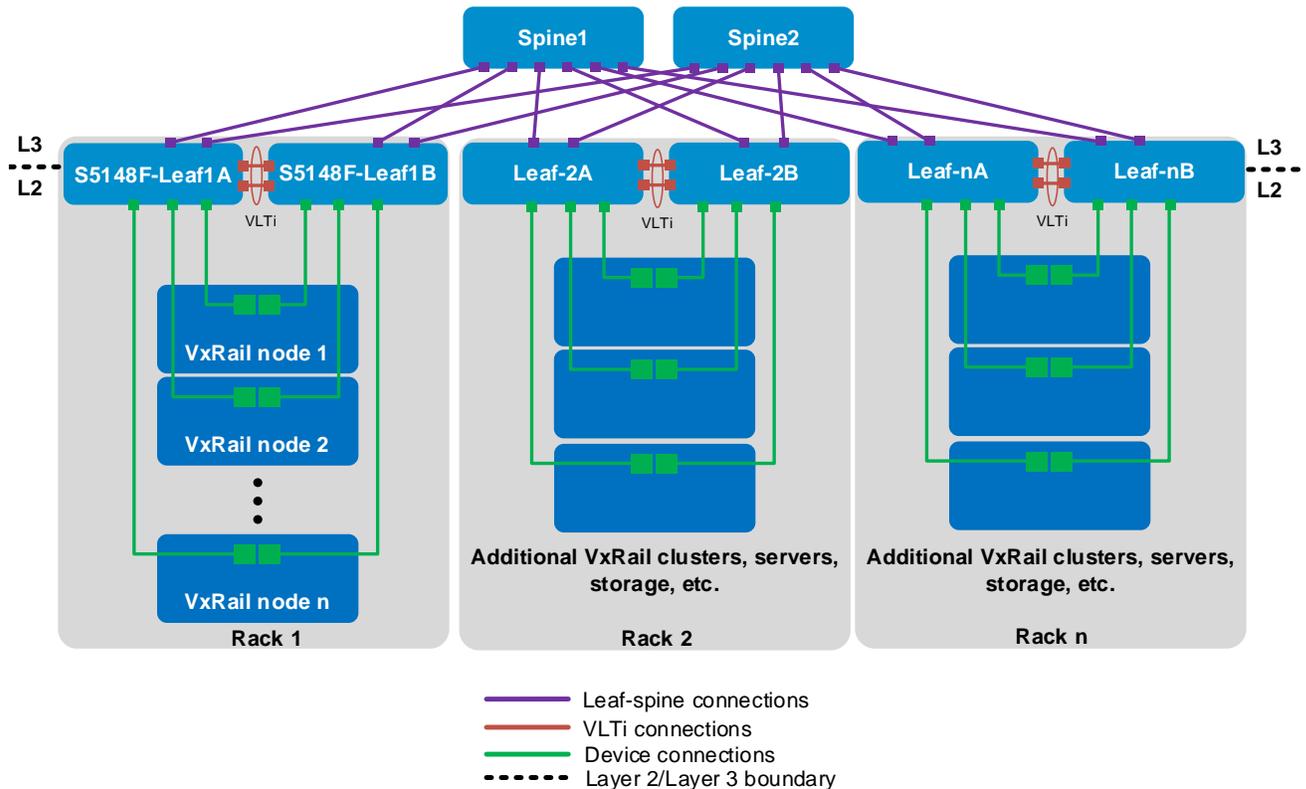


Figure 10 Dual switch topology with VLT connected to a leaf-spine network

In Figure 10, the Layer 2/Layer 3 boundary is at the leaves, meaning traffic within each rack is switched (Layer 2) and traffic between racks is routed (Layer 3). [VMware Validated Design 4.3, Architecture and Design](#) recommends isolating vSAN traffic to its own Layer 2 network segment. Therefore, for this leaf-spine topology, each vSAN should be contained within a single rack. Since a VxRail cluster contains a vSAN, a VxRail cluster is also contained within a single rack.

Note: By default, VxRail does not enable routing of vSAN or vMotion traffic.

The leaf-spine topology in Figure 10 scales to sixteen racks or more, depending on the number of ports available in each spine. Racks may contain additional VxRail clusters, switches, servers, storage arrays, and other devices as needed.

To configure the remainder of the leaf-spine network, including spine switches, connections between leafs and spines, and routing protocols, see the [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#).

4.2 OOB management network

The out-of-band (OOB) management network is an isolated network for remote management of servers, switches, and storage devices via their dedicated hardware management ports. It is also commonly used to carry heartbeat messages sent between switches configured as VLT peers.

For OOB management connections, one S3048-ON switch is installed in each rack as shown:

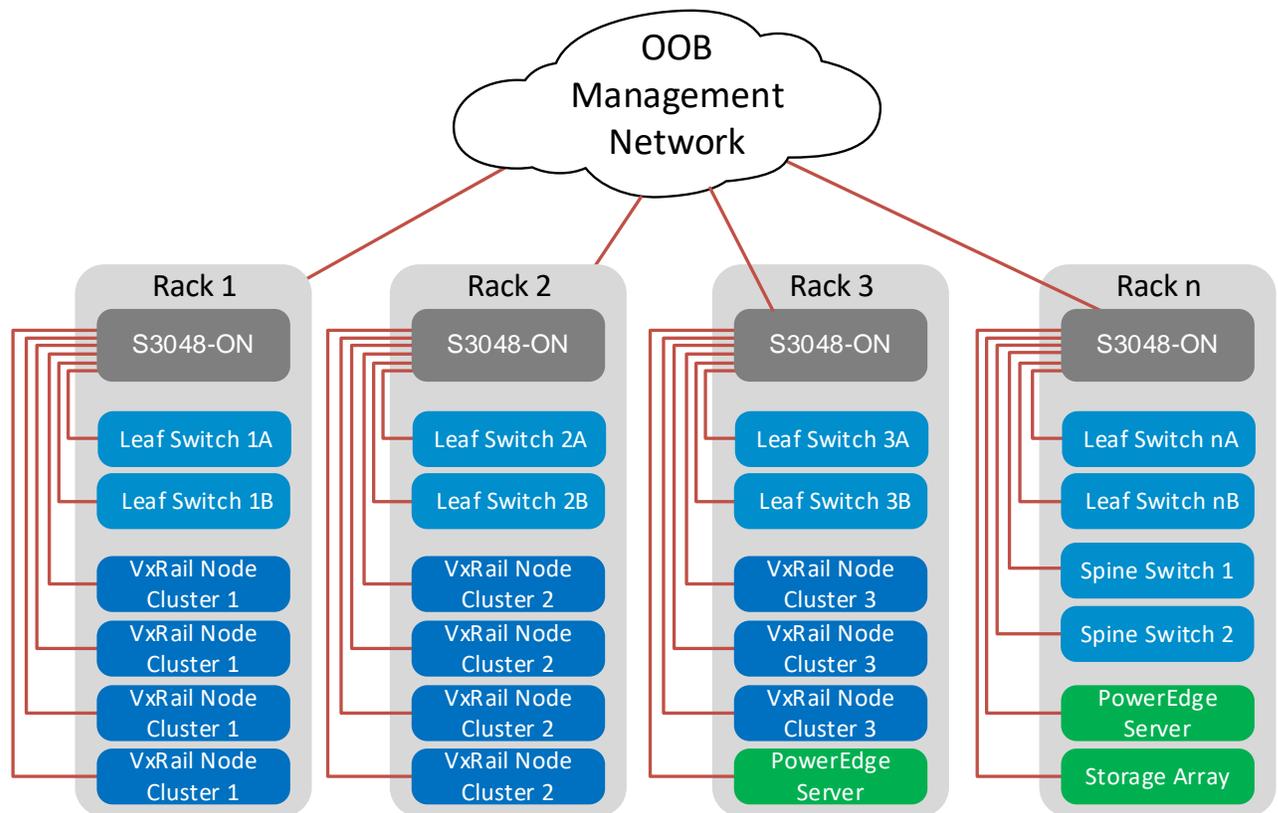


Figure 11 OOB management network connections

Note: This guide focuses on the devices in Rack 1. Devices shown in other racks are for illustration only.

1GbE BASE-T ports on each S3048-ON are connected downstream to hardware management ports on each device. This includes VxRail node iDRACs, PowerEdge Server iDRACs, and dedicated management ports on switches and storage arrays.

Four 10GbE SFP+ ports are available on each S3048-ON for use as uplinks to the OOB management network core.

Note: This is not to be confused with the in-band management network which runs on the same physical network as other VxRail and end-user traffic. In-band connections are covered in Section 3.

Devices on the OOB management network in this guide use the 100.67.0.0/16 IP address block.

5 Configuration planning

5.1 VLANs and IP addresses

VLANs and IP addresses used for VxRail node traffic must be planned before switch configuration and VxRail deployment can begin.

VxRail node traffic is divided into four or more VLANs as shown in Table 1.

Table 1 VLANs used for VxRail nodes

VLAN	Purpose
Management	Single VLAN for in-band management traffic used for VxRail node-to-node, ESXi host management, vCenter server, NTP, DNS, and VxRail management traffic
vMotion	Single VLAN for virtual machine (VM) migration
vSAN	Single VLAN for vSAN storage traffic
VM Network(s)	One or more VLANs for VM data traffic

VLAN configuration and IP network addresses planned for this deployment are shown in Table 2.

Table 2 VLANs and IP addresses

VLAN ID	Description	Network	Gateway	VLAN Traffic
1611	In-Band Management	172.16.11.0/24	172.16.11.254	Untagged
1612	vMotion	172.16.12.0/24	NA	Tagged
1613	vSAN	172.16.13.0/24	NA	Tagged
1614	VM Network A	172.16.14.0/24	Optional, not used	Tagged
1615	VM Network B	172.16.15.0/24	Optional, not used	Tagged

Note: By default, VxRail does not enable routing of vSAN or vMotion traffic.

This example uses two networks for VM data traffic, VM Network A and VM Network B, with each on a separate VLAN. The actual number of VM data networks used is based on customer requirements.

5.2 VxRail network configuration table

Information provided in the VxRail network configuration table is used during VxRail deployment. The values used for this deployment example are shown in the right column. The VLANs and IP addresses used are based on the information from Table 2.

Note: For additional information on the VxRail network configuration table, see the [Dell EMC VxRail Network Guide](#).

Table 3 VxRail network configuration table

Row		Category	Description	Values used	
1	VxRail	Management VLAN ID	The recommended is untagged traffic on the Native VLAN. If you want the host to send only tagged frames, manually configure the VLAN on each ESXi™ host using DCUI, and set tagging for your management VLAN on your switch before you deploy VxRail.	1611 (untagged)	
2		VxRail initial IP	If you cannot reach the default (192.168.10.200/24), set an alternate IP address	192.168.10.200	
3	System	Global settings	Time zone		
4			NTP server(s)	172.16.11.50	
5			DNS server(s)	172.16.11.50	
6		Proxy settings	IP address and port		
7			Username and password		
8	Management	ESXi hostnames and IP addresses	ESXi hostname prefix	vxhost	
9			Separator	none	
10			Iterator	Num 0x	
11			Offset	1	
12			Suffix	none	
13			Domain	dell.local	
14			ESXi starting address for IP pool	172.16.11.1	
15			ESXi ending address for IP pool	172.16.11.40	
16		vCenter Server	vCenter Server hostname	vxvcenter	
17		Leave blank if	vCenter Server IP address	172.16.11.100	
18		Customer	Platform Services Controller hostname	vxpsc	
19		Supplied VC	Platform Services Controller IP address	172.16.11.101	
20		Customer Supplied vCenter Server	Customer Supplied Platform Services Controller (PSC) Hostname (FQDN) Leave	Customer Supplied vCenter Server hostname (FQDN)	
21				Customer Supplied vCenter Server SSO domain	
22				admin username/password or the newly created VxRail non-admin username and password	
23				New VxRail management username and password	
24				Customer Supplied data center name	
25				New cluster name	
26				VxRail Manager	VxRail hostname
27		VxRail IP address	172.16.11.102		
28		Networking	Subnet mask	255.255.255.0	
29			Gateway	172.16.11.254	
30		Passwords	ESXi "root" passwords. Can be different for each host starting with Release 4.0.100.	Dell!234	
31			VxRail Manager and VxRail vCenter Server "administrator@vsphere.local"	P3!!!234	
32					
33		vMotion	Starting address for IP pool	172.16.12.1	
34			Ending address for IP pool	172.16.12.40	
35	Subnet mask		255.255.255.0		

Row		Category	Description	Values used
36			VLAN ID	1612
37	vSAN		Starting address for IP pool	172.16.13.1
38			Ending address for IP pool	172.16.13.40
39			Subnet mask	255.255.255.0
40			VLAN ID	1613
41	VM Networks	... (unlimited number)	VM Network name and VLAN ID	VM_Network_A, 1614
42			VM Network name and VLAN ID	VM_Network_B, 1615
43	Solutions	Logging	vRealize Log Insight™ hostname	vxinsight
44			vRealize Log Insight IP address	172.16.11.103
45			Syslog server (instead of Log Insight)	

This deployment uses four VxRail nodes. However, host IP address pools are configured for the Management, vMotion, and vSAN VLANs using addresses in the .1 - .40 range in Table 3. This ensures IP addresses are preconfigured for hosts that may be added later as the VxRail cluster grows.

5.3 Using a jump box or laptop computer

For VxRail deployment, VxRail Manager is accessed via the leaf switches connected to the nodes. This may be done by connecting a laptop computer directly to a leaf or by using a jump box (also known as a jump server or jump host) for remote access.

Reserve at least one interface on a leaf switch for this connection. This interface is configured on the in-band management VLAN during switch configuration.

Note: A PowerEdge server is used as a jump box in this deployment. More information on jump box connections is provided in Section 9.

5.4 DNS and NTP servers

VxRail nodes must be able to reach a correctly configured DNS server from the in-band management network during and after VxRail deployment. The DNS server must include forward and reverse lookup entries for ESXi hosts, vCenter, the Platform Services Controller, VxRail Manager, and vRealize Log Insight.

Add forward and reverse lookup records on the DNS server using the hostnames and IP addresses in the VxRail network configuration table (Table 3). Table 4 summarizes the DNS entries for this deployment.

Table 4 Hostnames and IP addresses used

Hostname	IP Address
vxhost01.dell.local	172.16.11.1
vxhost02.dell.local	172.16.11.2
vxhost03.dell.local	172.16.11.3
vxhost04.dell.local	172.16.11.4
vxvcenter.dell.local	172.16.11.100
vxpsc.dell.local	172.16.11.101
vxman.dell.local	172.16.11.102
vxinsight.dell.local	172.16.11.103

An NTP server is not required but is recommended. If an NTP server is not provided, VxRail uses the time that is set on VxRail node 1.

Note: For this deployment guide, the PowerEdge server used as the jump box also provides DNS and NTP services for the VxRail cluster. It is connected to both leaf switches for redundancy using Microsoft Windows NIC teaming (switch-independent load balancing).

6 Switch configuration prerequisites

6.1 Check switch OS version

S5148F-ON switches must be running OS10EE version 10.4.1.1 or later. Run the `show version` command to check the OS version. Dell EMC recommends upgrading to the latest release available on [Dell Digital Locker](#) (account required).

```
OS10# show version
S5148F-ON# show version
Dell EMC Networking OS10-Enterprise
Copyright (c) 1999-2018 by Dell Inc. All Rights Reserved.
OS Version: 10.4.1.1
Build Version: 10.4.1.1.507
```

Note: Figure 1 at the beginning of this guide lists supported switches and operating systems for VxRail deployments.

6.2 Verify license installation

Run the command `show license status` to verify license installation. The `License Type:` field should indicate `PERPETUAL`. If an evaluation license is installed, licenses purchased from Dell EMC are available for download on [Dell Digital Locker](#). Installation instructions are provided in the [OS10 Enterprise Edition User Guide Release 10.4.1.0](#).

```
OS10# show license status

System Information
-----
Vendor Name       :      Dell EMC
Product Name      :      S5148F-ON
Hardware Version  :      A00
Platform Name     :      x86_64-dellemc_s5100_c2538-r0
PPID              :      CN00Y2VTCES0082O0038
Service Tag      :      D8MSG02
License Details
-----
Software         :      OS10-Enterprise
Version          :      10.4.1.1
License Type    :      PERPETUAL
License Duration :      Unlimited
License Status   :      Active
License location :      /mnt/license/D8MSG02.lic
-----
```

Note: If OS10EE was factory installed, a perpetual license is already on the switch.

6.3 Factory default configuration

The configuration commands in the sections that follow begin with S5148F-ON switches at their factory default settings. Dell EMC Networking switches running OS10EE can be reset to their default configuration as follows:

```
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
```

The switch reboots to its factory default configuration.

Note: OS10EE at its default settings has Telnet disabled, SSH enabled, and the OOB management interface configured to get its IP address via DHCP. The default username and password are both `admin`. Dell EMC recommends changing the admin password to a complex password during the first login.

7 Configure switches

This section covers switch configuration for S5148F-ON switches in the different topologies. Commands for all examples are provided as attachments to this guide.

7.1 Dual switch with VLT

This example uses a four-node VxRail cluster connected to a pair of switches configured with VLT as shown below. Dell EMC recommends using this topology.

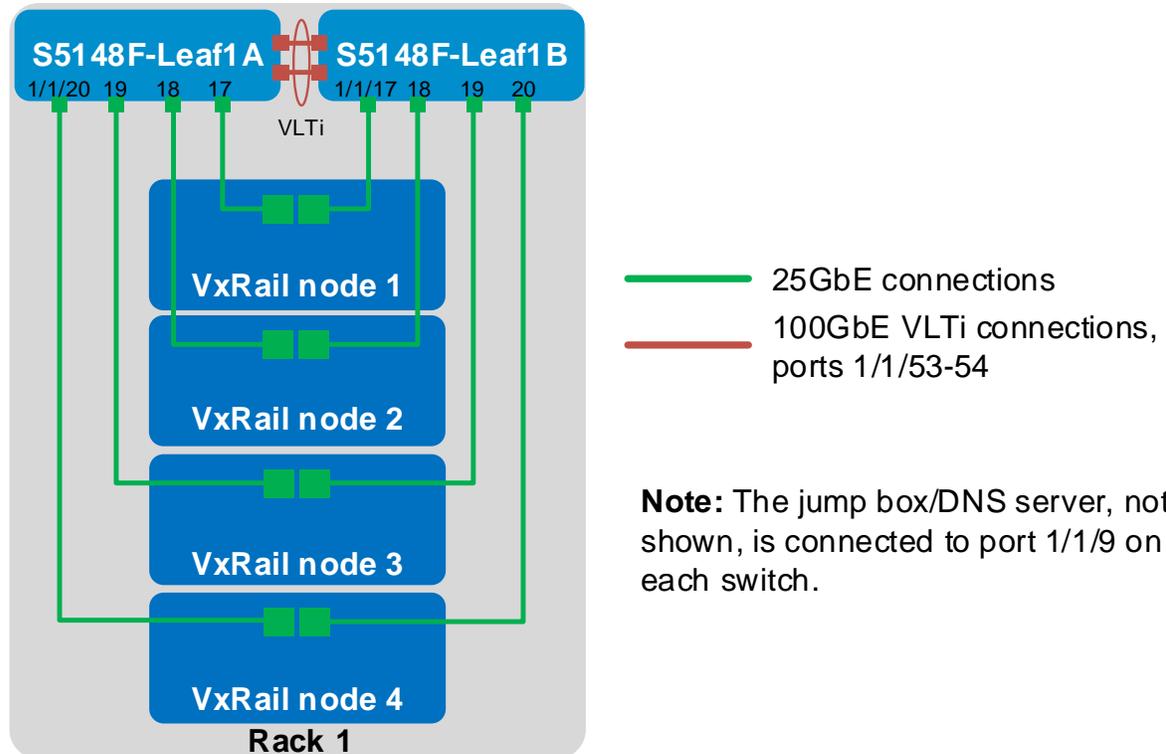


Figure 12 Four-node cluster connected to a pair of switches configured with VLT

The commands in the following sections are run to complete the configuration of both leaf switches. The port numbers used in the configuration commands correspond to those shown in Figure 12.

Note: The commands in the tables below should be entered in the order shown. All commands are provided in the file attachments named **S5148F-1A-vlt.txt** and **S5148F-1B-vlt.txt**.

7.1.1 General settings

Configure the hostname, OOB management IP address and default gateway. Specify an NTP server accessible by the switch. In this example, this is an NTP server on the OOB management network, with a different IP address than that used by the VxRail nodes on the in-band management network. Enable RSTP as a precaution against loops. Configure S5148F-Leaf1A as the primary RSTP root bridge using the `spanning-tree rstp priority 0` command and S5148F-Leaf1B as the secondary RSTP root bridge using the `spanning-tree rstp priority 4096` command.

Note: In OS10EE, LLDP is enabled on each interface and globally by default. LLDP is useful for troubleshooting and validation. Dell EMC recommends leaving it enabled.

Table 5 General settings – dual switch with VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>configure terminal hostname S5148F-Leaf1A interface mgmt1/1/1 no ip address ip address 100.67.172.38/24 no shutdown management route 0.0.0.0/0 100.67.172.254 ntp server 100.67.10.20 spanning-tree mode rstp spanning-tree rstp priority 0</pre>	<pre>configure terminal hostname S5148F-Leaf1B interface mgmt1/1/1 no ip address ip address 100.67.172.37/24 no shutdown management route 0.0.0.0/0 100.67.172.254 ntp server 100.67.10.20 spanning-tree mode rstp spanning-tree rstp priority 4096</pre>

7.1.2 Configure VLANs

In this section, VLANs are configured per Table 2.

Create the in-band management VLAN and assign a unique IP address on each switch. Configure VRRP to provide gateway redundancy. Set the VRRP priority. The switch with the largest priority value becomes the master VRRP router. Assign the same virtual address to both switches.

Create the vMotion, vSAN, and VM Network VLANs.

Table 6 Configure VLANs – dual switch with VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface vlan1611 description InBand_Mgmt ip address 172.16.11.253/24 vrrp-group 11 priority 150 virtual-address 172.16.11.254 no shutdown interface vlan1612 description vMotion no shutdown interface vlan1613 description vSAN no shutdown interface vlan1614 description VM_Network_A no shutdown interface vlan1615</pre>	<pre>interface vlan1611 description InBand_Mgmt ip address 172.16.11.252/24 vrrp-group 11 priority 100 virtual-address 172.16.11.254 no shutdown interface vlan1612 description vMotion no shutdown interface vlan1613 description vSAN no shutdown interface vlan1614 description VM_Network_A no shutdown interface vlan1615</pre>

S5148F-Leaf1A	S5148F-Leaf1B
description VM_Network_B no shutdown	description VM_Network_B no shutdown

7.1.3 Configure interfaces

Configure the interfaces for the jump box or laptop computer connections to be used during VxRail deployment. This example uses interface 1/1/9 on each switch. While only one connection is required, two are used for redundancy. Add these interfaces as access ports on VLAN 1611 (the in-band management VLAN).

Note: In this example, the jump box also provides DNS and NTP services as covered in Section 5.4. If these services are located on other devices, ensure the switch interfaces connected to these services are configured in the same manner.

Configure the interfaces for connections to the VxRail nodes. Interfaces 1/1/17 through 1/1/20 are used in this example.

Use the `switchport mode trunk` command to enable ports to carry traffic for multiple VLANs. Configure the ports as access (untagged) ports on VLAN 1611 (the in-band management VLAN). Configure the ports as trunk (tagged) ports on VLANs 1612-1615 (the vMotion, vSAN, and VM Network VLANs).

Configure ports directly connected to nodes, servers, or other endpoints as STP edge ports. Set flow control to “receive on” and “transmit off” on node-connected ports as a best practice.

Note: Flow control is enabled on all network interfaces in ESXi by default. For more information, see [VMware vSAN Network Design](#).

Table 7 Configure interfaces – dual switch with VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface ethernet1/1/9 description JumpBox_DNS switchport access vlan 1611 spanning-tree port type edge no shutdown interface ethernet1/1/17 description Node1_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/18 description Node2_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615</pre>	<pre>interface ethernet1/1/9 description JumpBox_DNS switchport access vlan 1611 spanning-tree port type edge no shutdown interface ethernet1/1/17 description Node1_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/18 description Node2_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615</pre>

S5148F-Leaf1A	S5148F-Leaf1B
<pre>spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/19 description Node3_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/20 description Node4_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown</pre>	<pre>spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/19 description Node3_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/20 description Node4_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown</pre>

7.1.4 Configure VLT

Use 100GbE interfaces for the VLTi between the two leaf switches. This example uses interfaces 1/1/53 and 1/1/54. Remove each interface from Layer 2 mode with the `no switchport` command.

Create the VLT domain. The backup destination is the OOB management IP address of the VLT peer switch. Configure the interfaces used as the VLTi with the `discovery-interface` command. Enable peer routing.

Note: For more information on VLT, see the [OS10 Enterprise Edition User Guide Release 10.4.1.0](#).

When the configuration is complete, exit configuration mode and save the configuration with the `end` and `write memory` commands.

Table 8 Configure VLT – dual switch with VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface ethernet1/1/53 description VLTi no switchport no shutdown interface ethernet1/1/54 description VLTi no switchport no shutdown vlt-domain 127</pre>	<pre>interface ethernet1/1/53 description VLTi no switchport no shutdown interface ethernet1/1/54 description VLTi no switchport no shutdown vlt-domain 127</pre>

S5148F-Leaf1A	S5148F-Leaf1B
<pre> backup destination 100.67.172.37 discovery-interface ethernet1/1/53-1/1/54 peer-routing end write memory </pre>	<pre> backup destination 100.67.172.38 discovery-interface ethernet1/1/53-1/1/54 peer-routing end write memory </pre>

7.2 Dual switch without VLT

This example uses a four-node VxRail cluster connected to a switch pair without VLT as shown below.

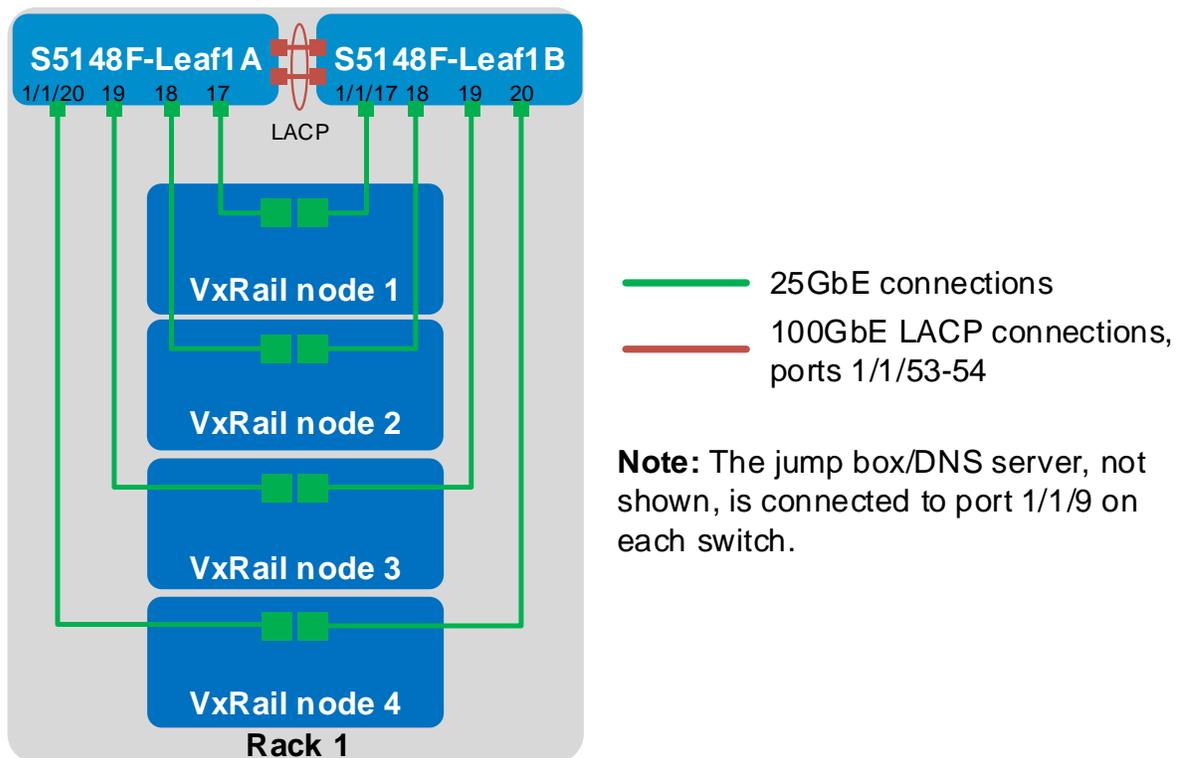


Figure 13 Four-node cluster connected to a switch pair without VLT

In this topology, an LACP port channel is used to connect the two switches.

The commands in the following sections are run to complete the configuration of both switches. The port numbers used in the configuration commands correspond to those shown in Figure 13.

Note: The commands in the tables below should be entered in the order shown. All commands are provided in the file attachments named **S5148F-1A-no-vlt.txt** and **S5148F-1B-no-vlt.txt**.

7.2.1 General settings

Configure the hostname, OOB management IP address and default gateway. Specify an NTP server accessible by the switch. In this example, this is an NTP server on the OOB management network, with a different IP address than that used by the VxRail nodes on the in-band management network. Enable RSTP

as a precaution against loops. Configure S5148F-Leaf1A as the primary RSTP root bridge using the `spanning-tree rstp priority 0` command and S5148F-Leaf1B as the secondary RSTP root bridge using the `spanning-tree rstp priority 4096` command.

Note: In OS10EE, LLDP is enabled on each interface and globally by default. LLDP is useful for troubleshooting and validation. Dell EMC recommends leaving it enabled.

Table 9 General settings - dual switch without VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>configure terminal hostname S5148F-Leaf1A interface mgmt1/1/1 no ip address ip address 100.67.172.38/24 no shutdown management route 0.0.0.0/0 100.67.172.254 ntp server 100.67.10.20 spanning-tree mode rstp spanning-tree rstp priority 0</pre>	<pre>configure terminal hostname S5148F-Leaf1B interface mgmt1/1/1 no ip address ip address 100.67.172.37/24 no shutdown management route 0.0.0.0/0 100.67.172.254 ntp server 100.67.10.20 spanning-tree mode rstp spanning-tree rstp priority 4096</pre>

7.2.2 Configure VLANs

In this section, VLANs are configured per Table 2.

Create the in-band management VLAN and assign a unique IP address on each switch. Configure VRRP to provide gateway redundancy. Set the VRRP priority. The switch with the largest priority value becomes the master VRRP router. Assign the same virtual address to both switches.

Create the vMotion, vSAN, and VM Network VLANs.

Table 10 Configure VLANs - dual switch without VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface vlan1611 description InBand_Mgmt ip address 172.16.11.253/24 vrrp-group 11 priority 150 virtual-address 172.16.11.254 no shutdown interface vlan1612 description vMotion no shutdown interface vlan1613 description vSAN no shutdown</pre>	<pre>interface vlan1611 description InBand_Mgmt ip address 172.16.11.252/24 vrrp-group 11 priority 100 virtual-address 172.16.11.254 no shutdown interface vlan1612 description vMotion no shutdown interface vlan1613 description vSAN no shutdown</pre>

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface vlan1614 description VM_Network_A no shutdown interface vlan1615 description VM_Network_B no shutdown</pre>	<pre>interface vlan1614 description VM_Network_A no shutdown interface vlan1615 description VM_Network_B no shutdown</pre>

7.2.3 Configure interfaces

Configure the interfaces for the jump box or laptop computer connections to be used during VxRail deployment. This example uses interface 1/1/9 on each switch. While only one connection is required, two are used for redundancy. Add these interfaces as access ports on VLAN 1611 (the in-band management VLAN).

Note: In this example, the jump box also provides DNS and NTP services as covered in Section 5.4. If these services are located on other devices, ensure the switch interfaces connected to these services are configured in the same manner.

Configure the interfaces for connections to the VxRail nodes. Interfaces 1/1/17 through 1/1/20 are used in this example.

Use the `switchport mode trunk` command to enable ports to carry traffic for multiple VLANs. Configure the ports as access (untagged) ports on VLAN 1611 (the in-band management VLAN). Configure the ports as trunk (tagged) ports on VLANs 1612-1615 (the vMotion, vSAN, and VM Network VLANs).

Configure ports directly connected to nodes, servers, or other endpoints as STP edge ports. Set flow control to “receive on” and “transmit off” on node-connected ports as a best practice.

Note: Flow control is enabled on all network interfaces in ESXi by default. For more information, see [VMware vSAN Network Design](#).

Table 11 Configure interfaces – dual switch without VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface ethernet1/1/9 description JumpBox_DNS switchport access vlan 1611 spanning-tree port type edge no shutdown interface ethernet1/1/17 description Node1_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown</pre>	<pre>interface ethernet1/1/9 description JumpBox_DNS switchport access vlan 1611 spanning-tree port type edge no shutdown interface ethernet1/1/17 description Node1_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612-1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown</pre>

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface ethernet1/1/18 description Node2_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612- 1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/19 description Node3_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612- 1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/20 description Node4_Port1 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612- 1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown</pre>	<pre>interface ethernet1/1/18 description Node2_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612- 1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/19 description Node3_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612- 1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown interface ethernet1/1/20 description Node4_Port2 switchport mode trunk switchport access vlan 1611 switchport trunk allowed vlan 1612- 1615 spanning-tree port type edge flowcontrol receive on flowcontrol transmit off no shutdown</pre>

7.2.4 Configure the switch interconnect

Create a port channel to use between the two leaf switches, 127 in this example. Use the `switchport mode trunk` command to enable the port channel to carry traffic for multiple VLANs, and allow all VLANs on the port channel.

Use 100GbE interfaces for the port channel, 1/1/53 and 1/1/54 in this example. Add each interface to the port channel as LACP active members with the `channel-group 127 mode active` command.

When the configuration is complete, exit configuration mode and save the configuration with the `end` and `write memory` commands.

Table 12 Configure switch interconnect - dual switch without VLT

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface port-channel 127 description To_Leaf_1B switchport mode trunk switchport trunk allowed vlan 1611- 1615 no shutdown</pre>	<pre>interface port-channel 127 description To_Leaf_1A switchport mode trunk switchport trunk allowed vlan 1611- 1615 no shutdown</pre>

S5148F-Leaf1A	S5148F-Leaf1B
<pre>interface ethernet1/1/53 description To_Leaf_1B channel-group 127 mode active no shutdown interface ethernet1/1/54 description To_Leaf_1B channel-group 127 mode active no shutdown end write memory</pre>	<pre>interface ethernet1/1/53 description To_Leaf_1A channel-group 127 mode active no shutdown interface ethernet1/1/54 description To_Leaf_1A channel-group 127 mode active no shutdown end write memory</pre>

7.3 Single switch

This example uses a four-node VxRail cluster connected to a single switch as shown below.

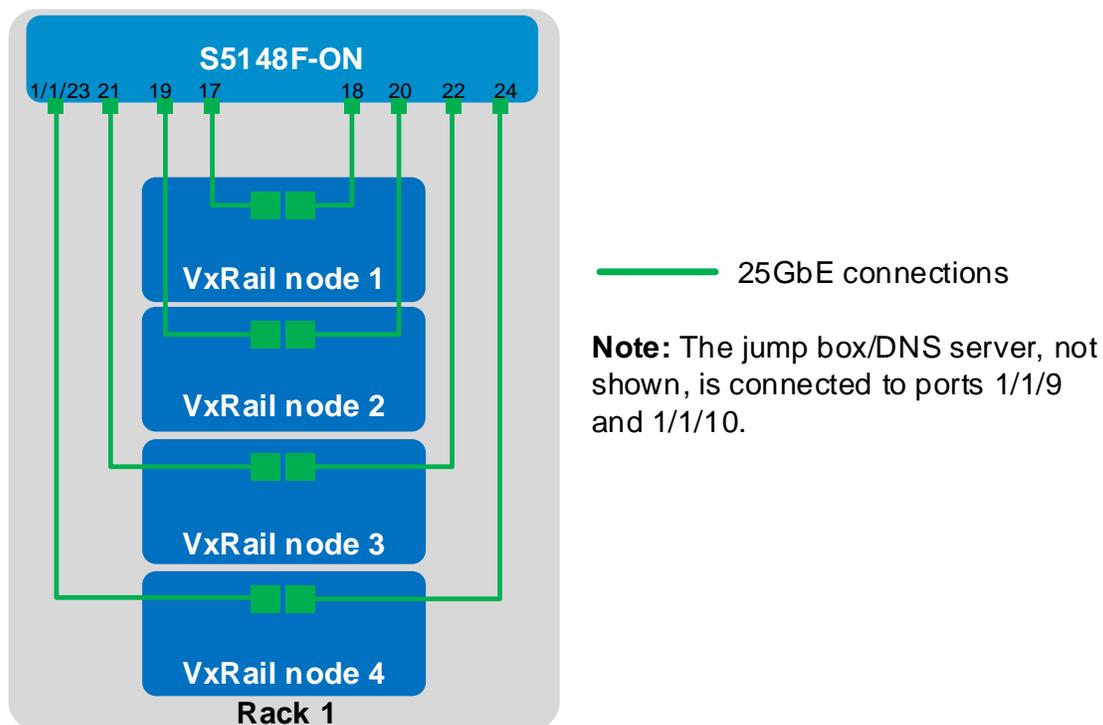


Figure 14 Four-node cluster connected to a single switch

The commands in the following sections are run to complete the switch configuration. The port numbers used in the configuration commands correspond to those shown in Figure 14.

Note: The commands in the tables below should be entered in the order shown. All commands are provided in the file attachment named **S5148F-single-switch.txt**.

7.3.1 General settings

Configure the hostname, OOB management IP address and default gateway. Specify an NTP server accessible by the switch. In this example, this is an NTP server on the OOB management network, with a different IP address than that used by the VxRail nodes on the in-band management network. Enable RSTP as a precaution against loops.

Note: In OS10EE, LLDP is enabled on each interface and globally by default. LLDP is useful for troubleshooting and validation. Dell EMC recommends leaving it enabled.

Table 13 General settings – single switch

S5148F-ON
<pre>configure terminal hostname S5148F-ON interface mgmt1/1/1 no ip address ip address 100.67.172.38/24 no shutdown management route 0.0.0.0/0 100.67.172.254 ntp server 100.67.10.20 spanning-tree mode rstp</pre>

7.3.2 Configure VLANs

In this section, VLANs are configured per Table 2.

Create the in-band management VLAN and assign an IP address. Create the vMotion, vSAN, and VM Network VLANs.

Table 14 Configure VLANs – single switch

S5148F-ON
<pre>interface vlan1611 description InBand_Mgmt ip address 172.16.11.254/24 no shutdown interface vlan1612 description vMotion no shutdown interface vlan1613 description vSAN no shutdown interface vlan1614 description VM_Network_A no shutdown interface vlan1615 description VM_Network_B</pre>

S5148F-ON

no shutdown

7.3.3 Configure interfaces

Configure the interfaces for the jump box or laptop computer connections to be used during VxRail deployment. This example uses interfaces 1/1/9 and 1/1/10. While only one connection is required, two are used for redundancy. Add these interfaces as access ports on VLAN 1611 (the in-band management VLAN).

Note: In this example, the jump box also provides DNS and NTP services as covered in Section 5.4. If these services are located on other devices, ensure the switch interfaces connected to these services are configured in the same manner.

Configure the interfaces for connections to the VxRail nodes. Interfaces 1/1/17 through 1/1/24 are used in this example.

Use the `switchport mode trunk` command to enable ports to carry traffic for multiple VLANs. Configure the ports as access (untagged) ports on VLAN 1611 (the in-band management VLAN). Configure the ports as trunk (tagged) ports on VLANs 1612-1615 (the vMotion, vSAN, and VM Network VLANs).

Configure ports directly connected to nodes, servers, or other endpoints as STP edge ports. Set flow control to “receive on” and “transmit off” on node-connected ports as a best practice.

Note: Flow control is enabled on all network interfaces in ESXi by default. For more information, see [VMware vSAN Network Design](#).

When the configuration is complete, exit configuration mode and save the configuration with the `end` and `write memory` commands.

Table 15 Configure interfaces – single switch

S5148F-ON

```
interface ethernet1/1/9
description JumpBox_DNS
switchport access vlan 1611
spanning-tree port type edge
no shutdown

interface ethernet1/1/10
description JumpBox_DNS
switchport access vlan 1611
spanning-tree port type edge
no shutdown

interface ethernet1/1/17
description Node1_Port1
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown
```

S5148F-ON

```
interface ethernet1/1/18
description Node1_Port2
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown
```

```
interface ethernet1/1/19
description Node2_Port1
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown
```

```
interface ethernet1/1/20
description Node2_Port2
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown
```

```
interface ethernet1/1/21
description Node3_Port1
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown
```

```
interface ethernet1/1/22
description Node3_Port2
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown
```

```
interface ethernet1/1/23
description Node4_Port1
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
```

S5148F-ON

```
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown

interface ethernet1/1/24
description Node4_Port2
switchport mode trunk
switchport access vlan 1611
switchport trunk allowed vlan 1612-1615
spanning-tree port type edge
flowcontrol receive on
flowcontrol transmit off
no shutdown

end
write memory
```

8 Switch validation

After switches are configured and devices are connected, the switch CLI is used to validate the network configuration. This section provides a list of the most common commands and their output for the examples used in this guide.

8.1 General validation commands

The commands and output shown below are for S5148F-Leaf1A in the dual switch with VLT topology. The output of its peer, S5148F-Leaf1B, is similar. If the dual switch without VLT or single switch topology is used, notes are added where differences exist.

Note: For additional commands and output related to the leaf-spine portion of the topology, such as BGP and Uplink Failure Detection (UFD), see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#).

8.1.1 show interface status

The `show interface status | grep up` command is used to verify required interfaces are up, and links are established at their appropriate speeds.

```
S5148F-Leaf1A#show interface status | grep up
Port          Description      Status  Speed  Duplex  Mode  Vlan  Tagged-Vlans
Eth 1/1/9     JumpBox_DNS     up      10G    full    A     1611  -
Eth 1/1/17    Node1_Port1     up      25G    full    T     1611  1612-1615
Eth 1/1/18    Node2_Port1     up      25G    full    T     1611  1612-1615
Eth 1/1/19    Node3_Port1     up      25G    full    T     1611  1612-1615
Eth 1/1/20    Node4_Port1     up      25G    full    T     1611  1612-1615
Eth 1/1/53    VLTi            up      100G   full    -     -     -
Eth 1/1/54    VLTi            up      100G   full    -     -     -
```

Note: For the dual switch without VLT topology, the output is identical to the above except for the description of ports 1/1/53-54. For the single switch topology, the output includes all eight node ports, and ports 1/1/53-54 are not used.

8.1.2 show port-channel summary

The `show port-channel summary` command is used to view port channel numbers, interfaces used, and status. In OS10EE, the VLTi is automatically configured as a static LAG using port channel 1000. Ports 1/1/53 and 1/1/54 are port channel members and (P) indicates each is up and active.

```
S5148F-Leaf1A#show port-channel summary
Flags:  D - Down      I - member up but inactive    P - member up and active
        U - Up (port-channel)  F - Fallback Activated
-----
Group  Port-Channel          Type      Protocol  Member Ports
-----
1000  port-channel1000      Eth       STATIC    1/1/53(P) 1/1/54(P)
```

Note: The command output for the dual switch without VLT topology is shown below. In this example, Port channel 127 is up (U) and DYNAMIC indicates LACP is used. Ports 1/1/53 and 1/1/54 are port channel members and (P) indicates each is up and active.

```
S5148F-Leaf1A#show port-channel summary
```

```
Flags:  D - Down      I - member up but inactive    P - member up and active
        U - Up (port-channel)    F - Fallback Activated
```

```
-----
Group Port-Channel          Type      Protocol  Member Ports
-----
127  port-channel127  (U)      Eth       DYNAMIC   1/1/53(P) 1/1/54(P)
-----
```

8.1.3 show vlan

The `show vlan` command is used to view interfaces assigned to each VLAN and whether the interfaces are access/untagged (A) or tagged (T). Port channel 1000 is the VLTi. VLAN ID 4094 is reserved as an internal control VLAN for the VLT domain, and it is not user configurable.

```
S5148F-Leaf1A#show vlan
```

```
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs
Q: A - Access (Untagged), T - Tagged
```

```
  NUM      Status      Description                                Q Ports
  ---      -
  1        Active
  1/1/16,1/1/21-1/1/52                        A Eth1/1/1-1/1/8,1/1/10-
                                                A Po1000
  1611     Active      InBand_Mgmt                               T Po1000
                                                A Eth1/1/9,1/1/17-1/1/20
  1612     Active      vMotion                                    T Eth1/1/17-1/1/20
                                                T Po1000
  1613     Active      vSAN                                        T Eth1/1/17-1/1/20
                                                T Po1000
  1614     Active      VM_Network_A                              T Eth1/1/17-1/1/20
                                                T Po1000
  1615     Active      VM_Network_B                              T Eth1/1/17-1/1/20
                                                T Po1000
  4094     Active
```

Note: For the dual switch without VLT topology, the VLAN portion of the command output is as follows:

NUM	Status	Description	Q Ports
1	Active		A Eth1/1/1-1/1/8,1/1/10-1/1/16,1/1/21-1/1/52
			A Po127
1611	Active	InBand_Mgmt	T Po127
			A Eth1/1/9,1/1/17-1/1/20
1612	Active	vMotion	T Eth1/1/17-1/1/20
			T Po127
1613	Active	vSAN	T Eth1/1/17-1/1/20
			T Po127
1614	Active	VM_Network_A	T Eth1/1/17-1/1/20
			T Po127
1615	Active	VM_Network_B	T Eth1/1/17-1/1/20
			T Po127

For the single switch topology, the VLAN portion is as follows:

NUM	Status	Description	Q Ports
1	Inactive		A Eth1/1/1-1/1/8,1/1/11-1/1/16,1/1/25-1/1/54
1611	Active	InBand_Mgmt	A Eth1/1/9-1/1/10,1/1/17-1/1/24
1612	Active	vMotion	T Eth1/1/17-1/1/24
1613	Active	vSAN	T Eth1/1/17-1/1/24
1614	Active	VM_Network_A	T Eth1/1/17-1/1/24
1615	Active	VM_Network_B	T Eth1/1/17-1/1/24

8.1.4 show lldp neighbors

The `show lldp neighbors` command is useful for identifying connected equipment. Interface 1/1/9 is connected to the jump box. Interfaces 1/1/17-20 are connected to the Broadcom adapter in each of the VxRail nodes, and 1/1/53-54 are connected to the other leaf.

S5148F-Leaf1A#`show lldp neighbors`

Loc PortID	Rem Host Name	Rem Port Id	Rem Chassis Id
ethernet1/1/9	Not Advertised	00:0a:f7:38:ba:a0	00:0a:f7:38:ba:a0
ethernet1/1/17	Broadcom Adv. Dua...	00:0a:f7:b6:15:d0	00:0a:f7:b6:15:d0
ethernet1/1/18	Broadcom Adv. Dua...	00:0a:f7:b6:1a:80	00:0a:f7:b6:1a:80
ethernet1/1/19	Broadcom Adv. Dua...	00:0a:f7:b6:11:c0	00:0a:f7:b6:11:c0
ethernet1/1/20	Broadcom Adv. Dua...	00:0a:f7:b6:1c:10	00:0a:f7:b6:1c:10
ethernet1/1/53	S5148F-Leaf1B	ethernet1/1/53	d8:9e:f3:bf:ae:00
ethernet1/1/54	S5148F-Leaf1B	ethernet1/1/54	d8:9e:f3:bf:ae:00

Note: For the dual switch without VLT topology, the output is identical to the above. For the single switch topology, the output includes all eight node-connected ports, and ports 1/1/53-54 are not used.

8.1.5 show vrrp brief

In this deployment, VRRP is configured on the in-band management VLAN, 1611, to provide gateway redundancy for management traffic. The output from the `show vrrp brief` command shows the master and virtual IP addresses and whether the switch is in the master or backup role. The switch configured with the largest priority value, shown in the `Priority` column, becomes the master.

```
S5148F-Leaf1A#show vrrp brief
```

Interface	Group	Priority	Preempt	State	Version	Master addr(s)	Virtual addr
vlan1611	IPv4 11	150	true	master	2	172.16.11.253	172.16.11.254

Note: For the dual switch without VLT topology, the output is identical to that shown above. For the single switch topology, this command is not applicable.

8.2 VLT validation commands

The following commands are only applicable to the dual switch with VLT topology.

8.2.1 show vlt domain_id

This command is used to validate the VLT configuration status. The `Role` for one switch in the VLT pair is `primary`, and its peer switch (not shown) is assigned the `secondary` role. Ensure `Peer-Routing` is `Enabled`. The `VLTi Link Status` and `VLT Peer Status` must both be `up`.

```
S5148F-Leaf1A#show vlt 127
```

```
Domain ID           : 127
Unit ID            : 1
Role               : primary
Version           : 2.0
Local System MAC address : d8:9e:f3:c0:61:00
Role priority      : 32768
VLT MAC address    : d8:9e:f3:bf:ae:00
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 Peer	Unit ID	System MAC Address	Status	IP Address	Version
2		d8:9e:f3:bf:ae:00	up	fda5:74c8:b79e:1::2	2.0

8.2.2 show vlt domain_id backup-link

This command is used to verify 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`.

```
S5148F-Leaf1A#show vlt 127 backup-link
```

VLT Backup Link

```
Destination                : 100.67.172.37
Peer Heartbeat status      : Up
Heartbeat interval         : 30
Heartbeat timeout          : 90
```

8.2.3 `show vlt domain_id mismatch`

This command highlights any potential configuration issues between VLT peers. All items must indicate No mismatch.

```
S5148F-Leaf1A#show vlt 127 mismatch
```

```
VLT-MAC mismatch:
```

```
No mismatch
```

```
Peer-routing mismatch:
```

```
No mismatch
```

```
VLAN mismatch:
```

```
No mismatch
```

```
VLT VLAN mismatch:
```

```
No mismatch
```

9 Deploy VxRail

A laptop computer or jump box with a web browser for the VxRail user interface is required. It is either plugged into a leaf switch or able to logically reach the VxRail in-band management VLAN from elsewhere on the network.

By default, the initial VxRail Manager IP address is 192.168.10.200/24. After initial configuration, the address changes to its new address on the in-band management network. The new VxRail Manager address used in this guide is 172.16.11.102/24 per the planning data in Table 3.

During installation, the laptop or jump box must be able to reach both the initial and new VxRail Manager addresses, so two addresses are configured on the network adapter, one for each network.

Note: Both addresses may be configured at the same time if the network adapter supports it, or in sequence if required.

The IP addresses configured on the laptop computer/jump box in this example are as follows:

- 192.168.10.1/24, to communicate with the initial VxRail Manager address, 192.168.10.200/24
- 172.16.11.50/24, to communicate with the new VxRail Manager address, 172.16.11.102/24

If a laptop computer is used, the connections appear as shown in Figure 15.

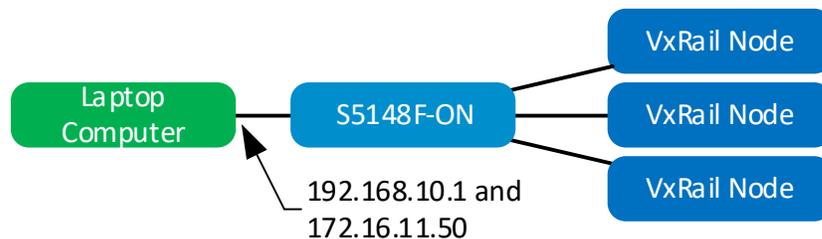


Figure 15 Use of a laptop computer for deployment

If a PowerEdge server is used as a jump box, make the connections as shown in Figure 16. Access the jump box's operating system from a workstation via the iDRAC virtual console.

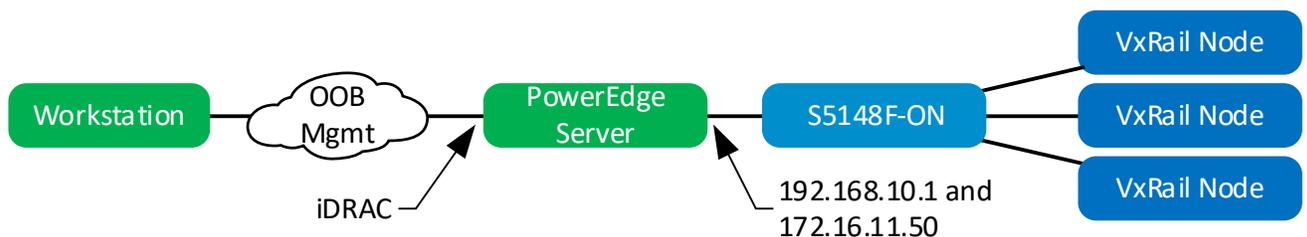


Figure 16 PowerEdge server as a jump box

Note: This deployment uses a PowerEdge server as a jump box that is also in the roles of DNS and NTP server. It may be connected to both leaf switches for redundancy using NIC teaming. Figure 16 shows a single connection from a jump box to a leaf switch for clarity.

In a web browser on the laptop computer or jump box, connect to <https://192.168.10.200> and deploy VxRail using the planning data in Table 3.

Note: VxRail installation steps are beyond the scope of this guide.

A Validated components

The following tables include the hardware, software, and firmware used to configure and validate the examples in this guide.

A.1 Dell EMC Networking Switches

Table 16 Switches and OS versions

Qty	Item	OS Version
2	Dell EMC Networking S5148F-ON leaf switches	10.4.1.1
1	Dell EMC Networking S3048-ON management switch	10.4.1.1

A.2 VxRail P570 nodes

A cluster of four VxRail P570 nodes was used to validate the examples in this guide. The nodes were each configured as shown in the table below.

Table 17 VxRail P570 node components

Qty per node	Item	Firmware version
2	Intel Xeon Gold 6136 CPU @ 3.00GHz, 12 cores	
12	16GB DDR4 DIMMs (192GB total)	
3	800GB SAS SSD	
10	1.2TB SAS HDD	
1	Dell HBA330 Storage Controller	13.17.03.05
1	Boot Optimized Storage Solution (BOSS) Controller w/ 2x240GB SATA SSDs	2.5.13.2008
1	Broadcom 57414 rNDC – 2x25GbE SFP28 ports	20.06.05.11
-	BIOS	1.2.11
-	iDRAC with Lifecycle Controller	3.15.17.15

A.3 VxRail Appliance software

This deployment guide was developed using VxRail Appliance software 4.5.200. It consists of the component versions shown in Table 18.

Table 18 VxRail Appliance software component versions

Item	Version
VxRail Manager	4.5.200.7769314
VMware ESXi	6.5.0 Update 1 Patch 36, build 7388607
VMware vCenter Server Appliance	6.5 Update 1e, build 7515524
VMware vSAN	6.6.1 Patch 02
VMware vRealize Log Insight	4.3, build 5084751

B VxRail network adapter traffic optimization

VxRail optimizes network traffic by splitting traffic across rNDC uplinks and by using NIOC shares. Based on traffic type, uplinks are configured as active or standby, and NIOC shares are used to allocate network bandwidth. Uplink load balancing and NIOC shares are automatically configured by VxRail.

Table 19 Network traffic on E, P, S and V Series 14G nodes with 2x10GbE or 2x25GbE ports

Traffic Type	Requirement*	Uplink 1 (10 or 25GbE) VMNIC0	Uplink 2 (10 or 25GbE) VMNIC1	NIOC Shares
Management	IPv6 multicast	Active	Standby	20
vSphere VMotion		Active	Standby	50
vSAN	IPv4 unicast	Standby	Active	100
Virtual Machines		Active	Standby	30

Table 20 Network traffic on E, P, S and V Series 14G nodes with 4x10GbE ports

Traffic Type	Requirement*	Uplink 1 (10GbE) VMNIC0	Uplink 2 (10GbE) VMNIC1	Uplink 3 (10GbE) VMNIC2	Uplink 4 (10GbE) VMNIC3	NIOC Shares
Management	IPv6 multicast	Active	Standby	Unused	Unused	40
vSphere VMotion		Unused	Unused	Standby	Active	50
vSAN	IPv4 unicast	Unused	Unused	Active	Standby	100
Virtual Machines		Standby	Active	Unused	Unused	60

Table 21 Network traffic on E and S Series 14G nodes with 4x1GbE ports

Traffic Type	Requirement*	Uplink 1 (1GbE) VMNIC2	Uplink 2 (1GbE) VMNIC3	Uplink 3 (1GbE) VMNIC0	Uplink 4 (1GbE) VMNIC1	NIOC Shares
Management	IPv6 multicast	Standby	Active	Unused	Unused	40
vSphere VMotion		Unused	Unused	Standby	Active	50
vSAN	IPv4 unicast	Unused	Unused	Active	Standby	100
Virtual Machines		Active	Standby	Unused	Unused	60

*Switches used must allow this traffic. Not explicitly configured on Dell EMC Networking switches.

Note: For additional VxRail nodes and rNDC configurations, see the [Dell EMC VxRail Network Guide](#).

C Technical resources

[Dell EMC Networking Guides](#)

[OS10 Enterprise Edition User Guide Release 10.4.1.0](#)

[Manuals and documentation for Dell EMC Networking S3048-ON](#)

[Manuals and documentation for Dell EMC Networking S5148F-ON](#)

[Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#)

[Dell EMC VxRail Network Guide](#)

[Dell EMC VxRail support and documentation](#) (account required)

[Dell EMC VxRail Support Matrix](#) (account required)

[Dell EMC Networking OS9 Switch Configuration Guide for VxRail 4.5](#)

[VMware vSAN Network Design](#)

[VMware Validated Design Documentation](#)

D 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

We encourage readers to provide feedback on the quality and usefulness of this publication by sending an email to Dell_Networking_Solutions@Dell.com.