

Dell EMC VxRail™ with SmartFabric Services Planning and Preparation Guide

Physical and Logical Requirements and Considerations for the Planning and Preparation of VxRail with Dell EMC SFS

Abstract

The implementation and management of the VxRail cluster with the Dell EMC SmartFabric Services networking feature is unique in contrast with a deployment on a standard network. For better understanding, this guide provides the unique data center and networking requirements for a VxRail cluster integrated with a networking solution based on Dell EMC SmartFabric Services. This document is not a replacement for the necessary planning and design effort for the supporting network, or the services effort required for implementation and management of the VxRail cluster. It should not be used to for these purposes.

April 2020

Revision history

Date	Description
April 2020	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.

Copyright © 2020 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. [12/3/2020] [Planning Guide] [H18290]

Table of contents

1	Introduction to VxRail with Dell EMC SFS	7
1.1	VxRail logical networks	7
1.2	Dell EMC SFS and VxRail	8
1.3	SmartFabric Network Management	10
2	Multi-Rack VxRail with SFS	12
3	VxRail with SFS Network Design	14
3.1	SmartFabric Network Topology Standards and Concepts	14
3.2	Network Design Layer 2/Layer 3 Boundary Decision	14
3.3	VxRail-SmartFabric Topology: Single Rack	15
3.4	VxRail-SmartFabric Topology: Single Rack to Multi-Rack	16
3.5	VxRail-SmartFabric Topology: Multi-Rack with Leaf Border	17
3.6	VxRail-SmartFabric Topology: Multi-Rack with Spine Border	18
3.7	VxRail-SmartFabric Topology: Multi-Rack with Dedicated Leaf Border	19
4	VxRail physical network connectivity planning	20
4.1	Consult the VxRail Support Matrix	20
4.2	VxRail hardware selection	20
4.3	Dell EMC networking switch selection	20
4.4	Data center core network physical connectivity	22
5	Task outline for VxRail cluster on SmartFabric-based network	23
6	VxRail and SmartFabric Network Planning	26
6.1	VxRail and SmartFabric Networks Overview	26
6.2	Data Center Components and Services	27
6.2.1	IT management switched network	27
6.2.2	Domain Name Services (DNS)	27
6.2.3	Network Time Protocol (NTP)	27
6.2.4	External vCenter	27
6.2.5	FTP/SFTP Server	27
6.2.6	Console Access Server	27
6.3	IP Address Reservations for IT Management Network	27
6.3.1	Network Settings for SmartFabric switches	28

6.3.2	Network Settings for VxRail Node Management	28
6.3.3	Network Settings for FTP/SFTP Server	28
6.3.4	Network Settings for OMNI vCenter plug-in	28
6.4	Update Default SmartFabric Network Settings.....	28
7	VxRail Cluster Initial Build Planning	30
7.1	VxRail Cluster Network VLAN Reservation	30
7.2	VxRail Management Network Reservation.....	30
7.3	VxRail Cluster Settings	31
7.3.1	VxRail Management Network IP Address Assignments.....	31
7.3.2	VxRail Cluster Global Settings	31
7.3.3	VxRail node hostnames.....	31
7.3.4	VxRail vCenter Identification	32
7.3.5	vMotion Settings	32
7.3.6	vSAN Settings.....	32
7.3.7	VM Networks Settings	32
7.3.8	Solutions Settings.....	32
7.4	Prepare DNS for VxRail cluster deployment	32
8	Data Center Network Preparation	33
8.1	Enable Jumbo Frames	33
8.2	Layer 2 Network Integration	33
8.3	Layer 3 Network Integration	34
8.3.1	Point-to-Point Links	34
8.3.2	Layer 3 Integration with Static Routing.....	35
8.3.3	Layer 3 Integration with Border Gateway Protocol (BGP).....	35
8.4	Enable IP Addressing for a Virtual Network	36
8.5	Layer 3 Integration using a VLAN.....	37
8.6	Spanning Tree Integration	38
8.6.1	Spanning Tree Options for OS10 and SFS	38
8.6.2	Spanning Tree Guidelines for VxRail and SmartFabric networks	39
9	Device Onboarding to SmartFabric Network.....	Error! Bookmark not defined.
9.1	Create server interface profile	Error! Bookmark not defined.
9.2	Identify switch ports for non-VxRail devices	Error! Bookmark not defined.
9.3	Identify networks to attach to non-VxRail devices.....	Error! Bookmark not defined.

A VxRail with SmartFabric Network Configuration Tables	41
B VxRail Cluster Network Configuration Tables	44
C Server Interface Profile Table	Error! Bookmark not defined.

Intended use and audience

There are many differentiators between a VxRail™ cluster that is implemented on a standard network infrastructure and one implemented on a network infrastructure that is placed under the control of Dell EMC SmartFabric Services (SFS). These differentiators impact the decision points to be made on the physical and logical settings and administrative considerations for the VxRail cluster. This guide discusses those essential data center and networking details for a VxRail cluster implementation with an integrated network infrastructure based on Dell EMC SFS. This guide has been prepared for anyone involved in planning and implementation efforts of a VxRail cluster on a network infrastructure based on Dell EMC SFS, including Dell EMC field engineers, customer system and network administrators, and IT personnel. It should be used for planning and preparation purposes only. It introduces requirements, best practices, and recommendations for both physical and virtual network environments. It should not be used to perform the actual installation and set-up of VxRail cluster or the underlying network infrastructure. Work with your Dell EMC service representative to perform the actual planning and design for the overall solution, and the subsequent implementation effort.

1 Introduction to VxRail with Dell EMC SFS

Dell EMC VxRail is a hyperconverged infrastructure (HCI) solution that consolidates compute, storage, and network into a single, highly available, unified system. All physical compute, network, and storage resources in VxRail are managed as a single shared pool. They are allocated to applications and services based on customer-defined business and operational requirements. The foundation for VxRail is a collection of uniquely engineered and manufactured server nodes that are bonded together in a cluster, and placed under a single point of management.

1.1 VxRail logical networks

The pool of nodes that comprise a VxRail cluster are dependent on an external physical network infrastructure to serve as a backplane to enable essential HCI services. They provide connectivity to customer resources and end users. The logical networking in a VxRail cluster is enabled by the configuration of a single virtual distributed switch to support required VxRail management, virtual machine interconnectivity, and access to external application and end-user resources. Each node connects to a redundant set of adjacent leaf Ethernet switches in the data center to enable interconnectivity between the physical and logical networks through the virtual distributed switch.

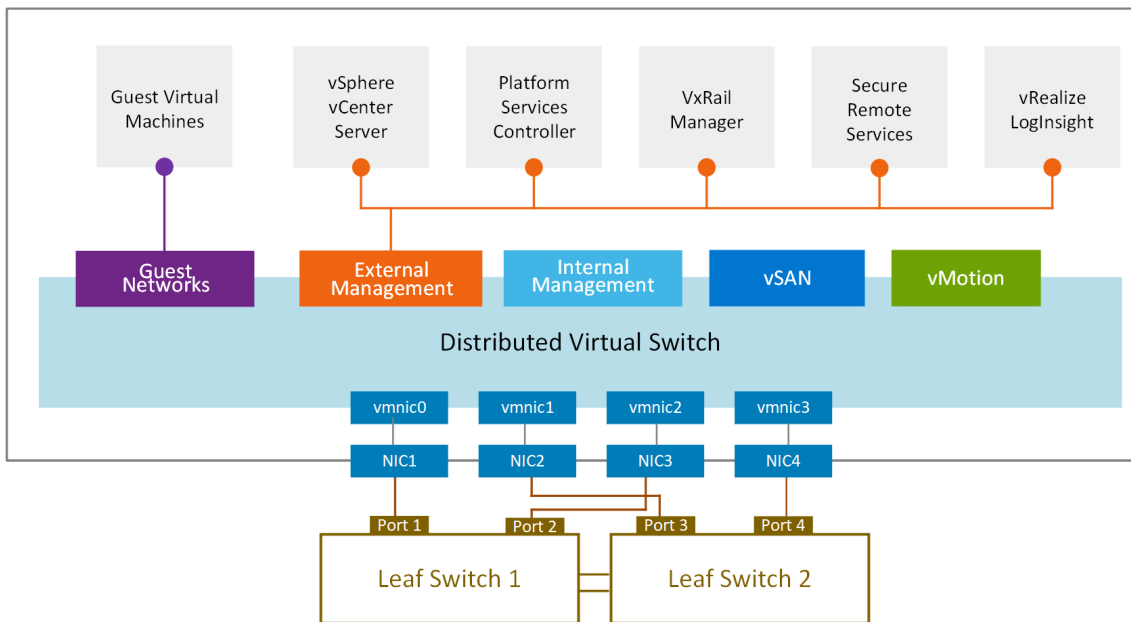


Figure 1 VxRail Management Components and Logical Networks with four physical connections

When the VxRail cluster is formed from the pool of nodes during initial implementation, the nodes establish physical network connectivity with the adjacent leaf switches. Portgroups are configured on the virtual distributed switch that is created during this process. Each of these portgroups enables connectivity for the essential networking services required by the VxRail cluster.

The VxRail management components in the form of virtual machines are deployed as a part of the implementation process. They are assigned to a portgroup on the virtual distributed switch that supports the external management network. These VxRail management components are dependent on the underlying physical networking infrastructure to provide connectivity to IT administration. In addition, virtual machines created on the VxRail cluster after implementation must be able to communicate with other virtual machines and applications both on the local virtual network, and with applications and end users external to the

VxRail cluster. This is accomplished through the configuration of guest networks on the VxRail cluster using defined portgroups on the virtual distributed switch.

The remaining logical networks shown in the graphic that are required by the VxRail cluster at the time of implementation are internal to the VxRail cluster. They do not require connectivity outside of the adjacent leaf switches. The internal management network is used for node and switch discovery, the vSAN network support storage services, and the vMotion network which enables virtual machine mobility between the nodes of the cluster.

1.2 Dell EMC SFS and VxRail

The network infrastructure supporting the VxRail cluster must be properly configured before initial implementation to support VxRail connectivity. Continuous management and reconfiguration of the physical network is expected to support ongoing operation of the applications and virtual machines on the cluster. These tasks enable continuous management of the leaf switches to support the VxRail required networks. They also provide the upstream infrastructure to extend the networks east and west across data center racks, and northbound and southbound to and from the core network.

VxRail is an integrated infrastructure solution with a validated architecture and control points derived from a single point of management. The procedures and processes to establish connectivity between the physical and logical networks are prescriptive and standardized. This simplifies the tasks that must be performed on the physical network to support VxRail applications and operations. Many of those repetitive tasks can be automated.




Figure 2 Overview of Dell EMC SFS with VxRail

Dell EMC SFS is an optional feature integrated into OS10, the operating system engineered for Dell EMC enterprise-class Ethernet switches. When SFS is enabled on the switches, control and management no longer originate from the default method of the console, outside of a few core management functions. Instead, it is reconfigured to establish a new personality profile at enablement. It transitions to a mode to support the applications or products using its networking resources. When a set of interconnected switches

are set to a common personality profile, the switches form a united fabric, with one switch being selected as the primary switch. The fabric can be easily expanded by connecting additional switches into the unified infrastructure, and enabling the personality profile through SFS. The fabric can be expanded linearly at the leaf layer, or spine switches can be added to form a two-tier network that can span across data center racks.

Support for a VxRail personality profile is engineered into both SFS and into VxRail Manager, and can be enabled with a minimum of two switches to form a fabric. When the VxRail personality profile mode is set on the pair of switches, the primary switch is reconfigured so that a browser can be used in a simplified interface to configure the properties for implementing a VxRail cluster on the fabric.

When a VxRail cluster is implemented, VxRail Manager discovers the powered-on nodes on the internal management network, places them under VxRail control, and uses them to form the managed cluster. VxRail Manager uses this same internal network to discover a switch fabric with the VxRail personality profile enabled, thereby establishing a link and point of control between the VxRail logical network and the physical network.

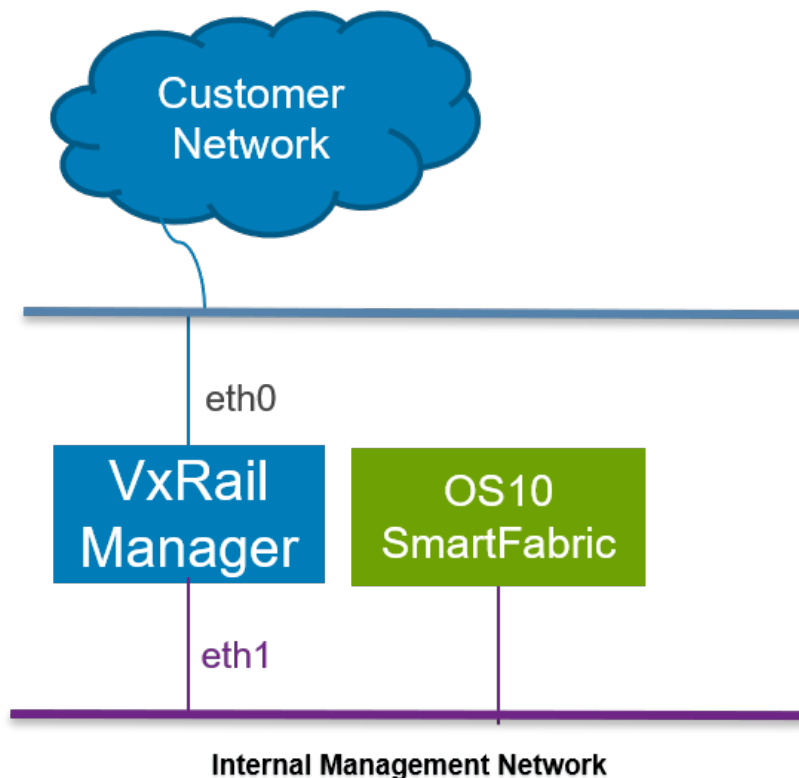


Figure 3 VxRail internal management network can discover SmartFabric-enabled switches.

The primary benefit of this established link is the synchronization of the physical switch fabric and VxRail logical networks during the cluster build process to ensure uninterrupted network connectivity. During initial cluster build in VxRail Manager, the network settings and properties are entered to form the VxRail cluster. After the data entry process is completed, VxRail Manager sends the network settings and properties to the switch fabric, where SFS automatically configures the switch fabric to support the VxRail cluster.

1.3 SmartFabric Network Management

The VxRail personality profile within SFS supports expansion to a multi-rack, multi-tier network topology into a single managed switch fabric. The default method of using the console for switch and network configuration is disabled by SFS, except for a handful of basic management functions. The OMNI vCenter plug-in is used to support ongoing management of the single switch fabric.

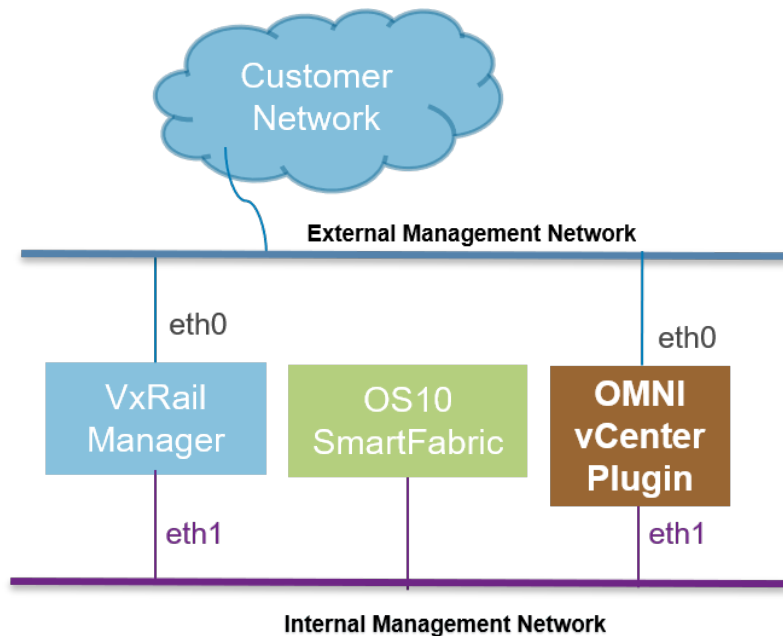


Figure 4 OMNI vCenter plug-in integration into SFS

Both vCenter and VxRail Manager do not have control or visibility to the supporting physical network enabled by SFS without the OMNI plug-in. The OMNI plug-in is a free download from Dell. It is deployed as a virtual appliance on a vCenter instance in the data center to support SmartFabric-enabled network administration tasks. The tool can manage and operate multiple instances of SmartFabric-enabled networks, and provides support for tasks such as uplink connectivity, on-boarding of non-integrated devices, and life cycle management. The OMNI plug-in can be accessed directly through a vClient with a user interface on the vCenter instance. When registered with the vCenter supporting the VxRail cluster, the OMNI plug-in enables a single, unified management view for the physical network fabric and the virtualized network environment.

The link between the physical and virtual networks by the OMNI plug-in also enables the automatic synchronization of the physical and virtual networks to ensure uninterrupted network connectivity after the VxRail cluster is built. For instance, if a new network portgroup with a new VLAN is created on the virtual distributed switch in the cluster, that network state change is detected, and the new VLAN is sent to the switch fabric, where it is applied to the fabric using SFS. There is no need to manually configure the switch fabric using the console in this instance.

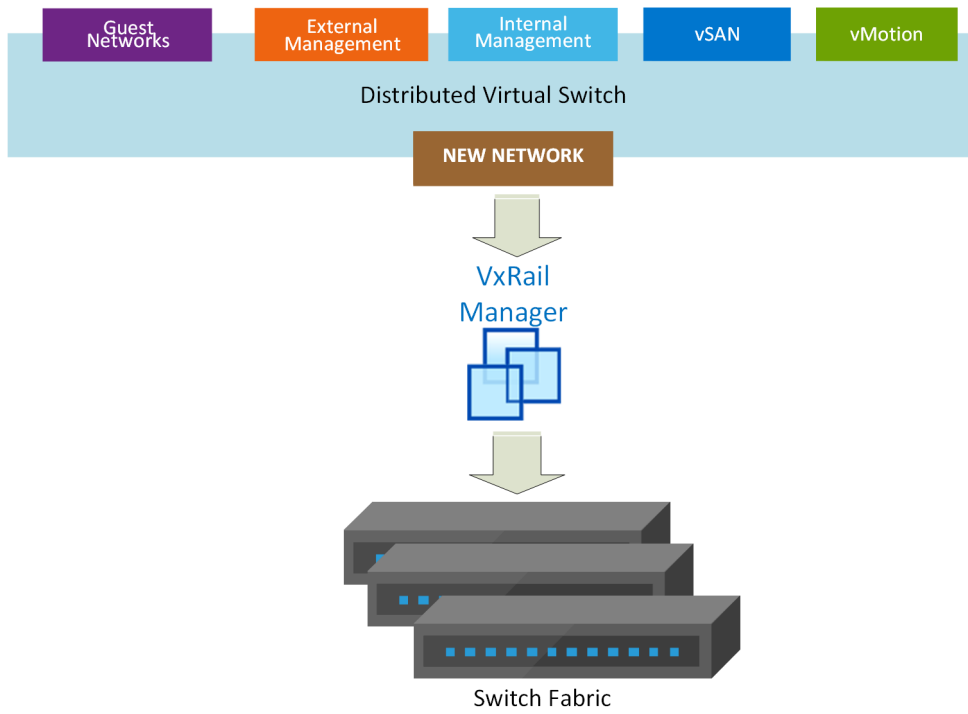


Figure 5 Network state reconfiguration between VxRail and SmartFabric-enabled switches

2 Multi-Rack VxRail with SFS

Both VxRail and a switch fabric based on SFS are architected to expand past the capacity of a single rack. VxRail supports up to 64 nodes in a single cluster, and can support the expansion of a single cluster across as many as six physical racks. A switch fabric based on SFS can support a single managed fabric consisting of up to twenty switches in eight racks.

By default, all VxRail networks are configured as Layer 2. To expand a VxRail cluster past a single rack, Layer 3 services are required in the expanded network to enable node discovery and cluster connectivity. In a standard network infrastructure, this means enabling routing for the required VxRail networks between the racks. When using a switch infrastructure with the VxRail personality profile enabled for SFS, the set of tasks required to connect the VxRail networks across racks is automated.

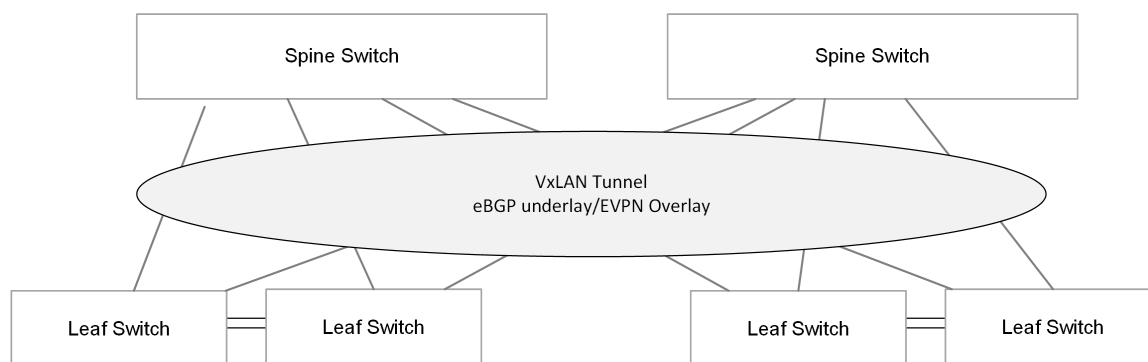


Figure 6 VxLAN network with SFS

SFS supports the automatic creation and expansion of a VxLAN tunnel network across the switch fabric in multiple racks. VxLAN (Virtual Extensible Local Area Network) is an overlay technology that allows you to extend a Layer 2 “overlay” network over a Layer 3 “underlay” network by adding a VxLAN header to the original Ethernet frame and encapsulating it. Encapsulation occurs by adding a VxLAN header on the original Layer 2 Ethernet frame, and placing it into an IP/UDP packet to be transported across the Layer 3 underlay network.




Figure 7 VxLAN packet format

The function of VxLAN is similar to VLAN regarding segmentation, but overcomes the 4096 VLAN limitation, and allows for up to 16 million VxLANs to be created. This key enabler for a multi-rack environment, as well as the transport ability, is why VxRail networks created in a SmartFabric-enabled network are configured as virtual networks, rather than VLANs. While a VLAN has limited mobility within the constraints of a Layer 2 network, a virtual network is configured automatically for transport on the VxLAN-based network, and can therefore enable connectivity across racks.

When a switched fabric is formed using SmartFabric personality profile supported for VxRail, two virtual networks are created: the client management network and the client control network. The client control network is used as the foundation for the VxRail internal management network, whose purpose is node discovery. This allows node discovery for cluster creation and expansion purposes to occur across racks without any additional intervention.

Powered-on VxRail nodes discover each other on the client control network after SFS are enabled and the switch fabric is formed. One VxRail node is elected as the primary, and VxRail Manager is started on that primary node. The client management network is then used to perform initial connectivity through a configured jump port to VxRail Manager for the purpose of discovering the switch fabric and building the cluster from the discovered nodes.

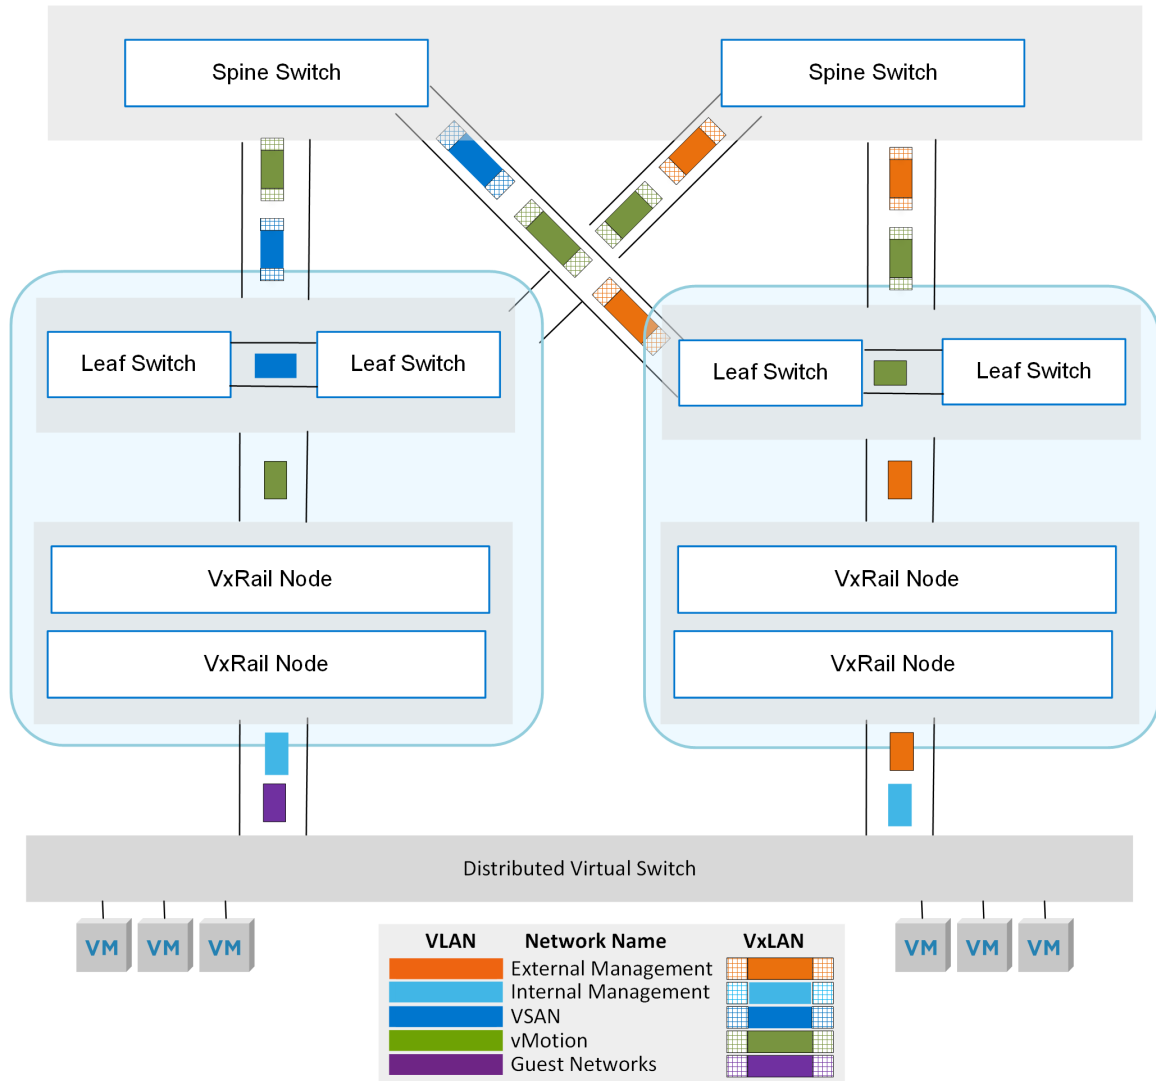


Figure 8 VxRail networks transport over the VxLan tunnel

The cluster build process configures the required VxRail networks on the SmartFabric-enabled network. SmartFabric captures those VxRail networks and automatically configure those as virtual networks for transport on the VxLAN tunnel network. The leaf switches in the switch fabric depend on the spine switches on the upper tier for transport on the VxLAN tunnel network across multiple racks, as there is no direct connectivity between leaf switches in separate racks.

3 VxRail with SFS Network Design

A VxRail cluster deployed on a SmartFabric-based network requires adherence to a set of supported network topologies. These topologies drive the overall design process and influence decision points on data center network integration. Acceptance of these requirements and adoption of the supported network architectures and interconnectivity options as part of the overall planning process is imperative before starting implementation.

Note: A switch on a network infrastructure enabled for SFS does not retain the network and fabric configuration settings if reverted to full switch mode and removed from the switch fabric. In this instance, if all switches are removed from the fabric, any operational VxRail clusters on that network infrastructure will lose connectivity.

A switch fabric supporting VxRail can start as small as two switches in a single rack. It can expand out to support multiple VxRail clusters or other infrastructure with up to twenty switches in a single fabric. As the switch fabric expands from a single rack to a multi-rack topology, the expansion to a multi-rack, multi-tier network must follow a prescriptive process and align with the defined VxRail and SmartFabric topologies.

3.1 SmartFabric Network Topology Standards and Concepts

SmartFabric-based network follows the same rules and requirements as a standard multi-tier network that spans across racks. The following standards apply to a SmartFabric-based network topology:

- Connections within racks from the VxRail nodes to the leaf switches is Layer 2
- The two leaf switches in each rack are configured with Virtual Link Trunking (VLT) to support active connectivity with multipathing and load balancing, and provides fault tolerance.
- VxRail nodes connect only to the leaf switches, and never to the spine switches
- Each leaf switch connects to every spine switch in the network topology. This means the total number of connections between the leaf and spine tiers is equal to the total number of leaf switches multiplied by the number of spine switches in the network.
- The overlay and underlay networks created by SFS to support VxLAN tunneling use isolated Layer 3 routed networks between the leaf layer and the spine layer.
- Spine switches are never connected to each other in the Layer 3 based network topology required for SmartFabric.
- Equal cost multipath routing (ECMP) is used to load balance traffic across the leaf switches and spine switches in the isolated Layer 3 routed network.

3.2 Network Design Layer 2/Layer 3 Boundary Decision

A key decision is how the SmartFabric-based network will interconnect with the upstream network. All switches supported for SFS and supported qualified for VxRail include support for Layer 2 and Layer 3, including static routing and Border Gateway Protocol (BGP) for Layer 3. A single rack deployment supports either Layer 2 or Layer 3 from the leaf switches. A multi-rack deployment supports either Layer 2 or Layer 3 from the leaf switch layer, or Layer 3 from the spine switch layer.

If the uplinks from the SmartFabric switches are Layer 2:

- Connectivity options from the leaf switches upstream are either a static Link Aggregation Group (LAG) or Link Aggregation Control Protocol (LACP).
- Upstream switches can be VLT/vPC peers or individual Layer 3 switches. SmartFabric automatically creates single MLAG or individual LAGs to each uplink switch.
- The boundary for the Layer 2 networks derived from the VxRail cluster is in the upstream data center network. Layer 3 routing services must be configured at this layer to enable connectivity for VxRail networks that require external access.

If the uplinks from the SmartFabric switches are Layer 3:

- Point-to-point links are required between the SmartFabric switches with the Layer 3 uplinks and the upstream switches.
- Connectivity options from the leaf switches or spine switches upstream are static routing or BGP.
- For static routing
 - Next-hop static routing statements must be configured on the upstream network for the Layer 3 networks passed upstream from the SmartFabric network.
- For BGP
 - SmartFabric switches must peer with upstream switches running eBGP.
 - The point-to-point links must include the ASN value dependent on whether the uplinks are derived from the leaf switches or spine switches.
 - A default ASN value is set at the spine layer and the leaf layer at the time of SFS enablement. These values can be modified before deploying the VxRail cluster on the switch fabric.
- SmartFabric networks support both Layer 3 routing per interface or Layer 3 routing per VLAN.
- The SmartFabric switches will run RPVST+ (Rapid Per-VLAN Spanning Tree) if the Layer 3 routing per VLAN option is enabled.

For any SmartFabric-enabled network for VxRail, the user interface connected with the primary switch is used to configure the uplinks to the upstream network.

3.3 VxRail-SmartFabric Topology: Single Rack

A switch fabric enabled with SFS can start with just a pair of leaf switches. Multiple VxRail clusters can be deployed on a pair of leaf switches, depending on the size of the clusters and the number of open switch data ports. The leaf switch pair will interconnect with a pair of links to form a VLT. The best practice is for each leaf switch to have a minimum of two connections to the redundant upstream network.

For the Layer 2 connectivity option, static LAG or LACP is required on the receiving ports on the upstream switches from the SmartFabric-enabled switches. If Layer 3 using BGP is selected, a peer network will be established with receiving ports on the upstream switches, and eBGP must be configured on the upstream network.

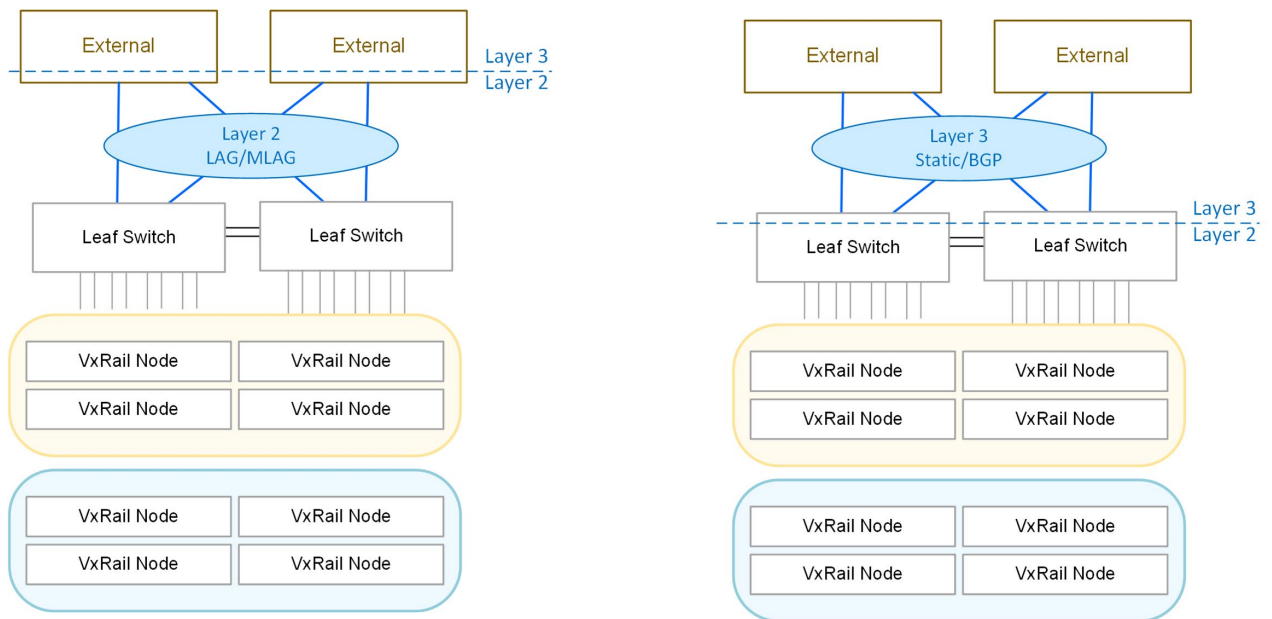


Figure 9 SmartFabric topologies for a single rack supporting Layer 2 and Layer 3 uplinks

3.4 VxRail-SmartFabric Topology: Single Rack to Multi-Rack

In a SmartFabric-based network, the transition from a single rack to multi-rack requires the deployment of a minimum of two qualified Dell EMC switches at the spine layer to interconnect the leaf switches in the individual racks. Each leaf switch requires a cable connection to each spine switch.

When the spine switches are enabled for SFS, they are discovered by the primary switch in the fabric. As part of their integration into the switch fabric, VxLAN tunnels are automatically set up for Layer 2 extensions between racks using Layer 3 VLAN over an eBGP underlay network and EVPN overlay network. Specific networks, such as those used by the VxRail cluster, are automatically added to the VxLAN tunnel network by SFS integration with VxRail Manager. Virtual networks can also be manually configured as a day 2 operation with the OMNI plug-in. This VxLAN-based tunnel network enables nodes in a VxRail cluster to reside in different physical racks, and be discovered by VxRail Manager across racks.

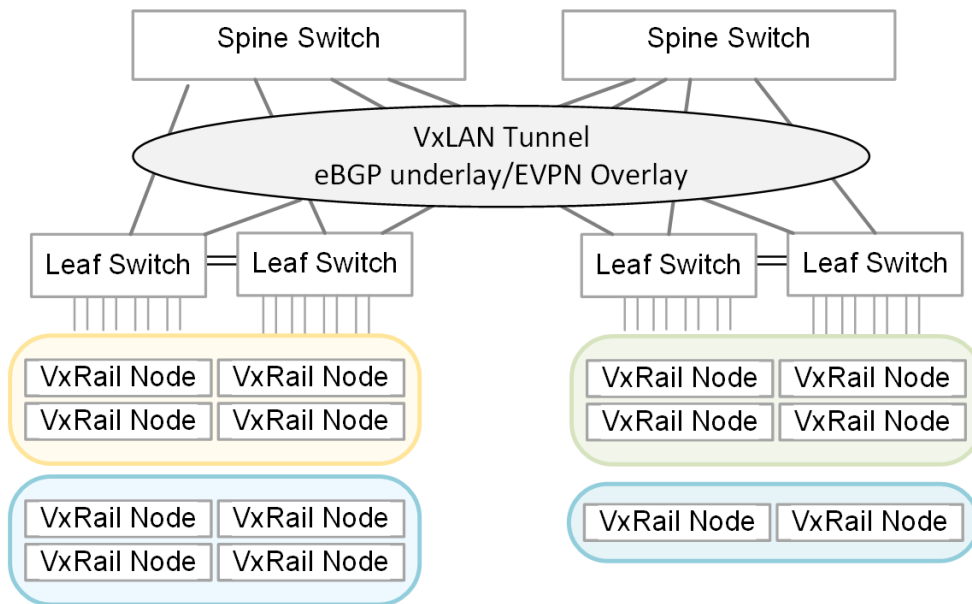


Figure 10 VxLAN tunnel network for SmartFabric multi-rack deployments

3.5 VxRail-SmartFabric Topology: Multi-Rack with Leaf Border

The expansion from a single rack fabric to a multi-rack fabric provides the option to keep upstream connectivity at the leaf layer or move upstream connectivity to the spine layer. For upstream connectivity at the leaf layer, a set of leaf switches can be selected to serve as the boundary between the switch fabric and upstream network.

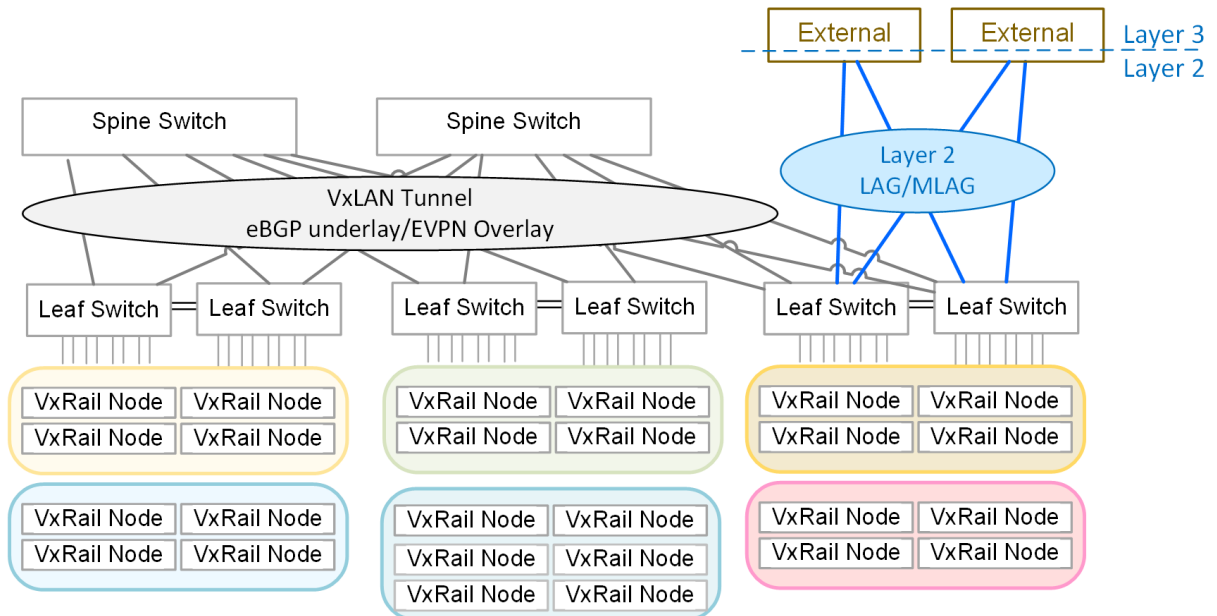


Figure 11 SmartFabric topology for multi-rack with Layer 2 uplinks at leaf layer

For connectivity to the upstream network, either Layer 2 or Layer 3 protocols can be enabled.

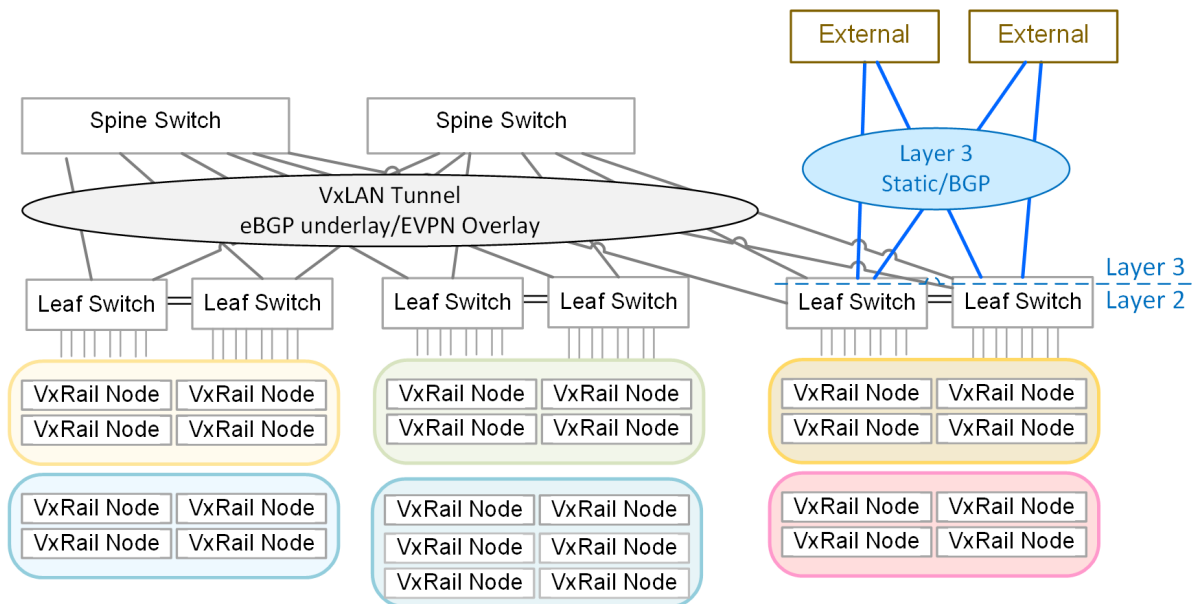


Figure 12 VxRail-SmartFabric Topology: Multi-Rack with Layer 3 Leaf Border

3.6 VxRail-SmartFabric Topology: Multi-Rack with Spine Border

The transition from a single rack to a multi-rack topology expands the upstream connectivity options to include the switches at the spine layer. Layer 3 protocols are the only option at the spine layer.

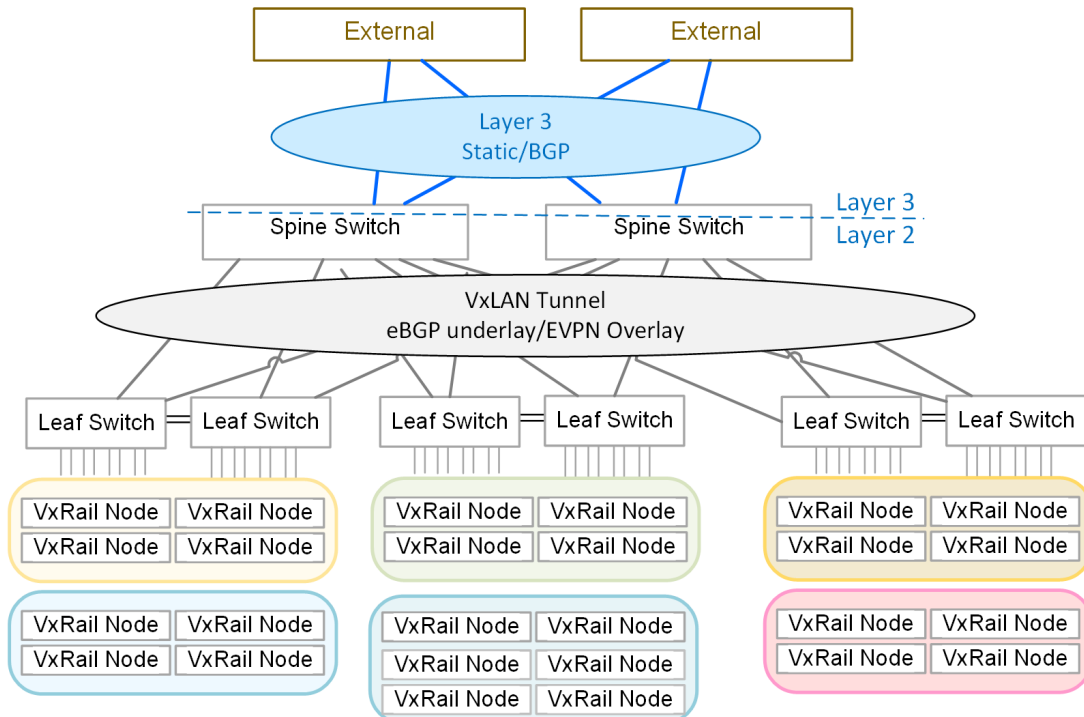


Figure 13 VxRail-SmartFabric Topology: Multi-Rack with Layer 3 Spine Border

3.7 VxRail-SmartFabric Topology: Multi-Rack with Dedicated Leaf Border

In certain circumstances, such as data center network policies or for expansion capacity reasons, the option to deploy a pair of switches dedicated to enable routing to the upstream network is also supported.

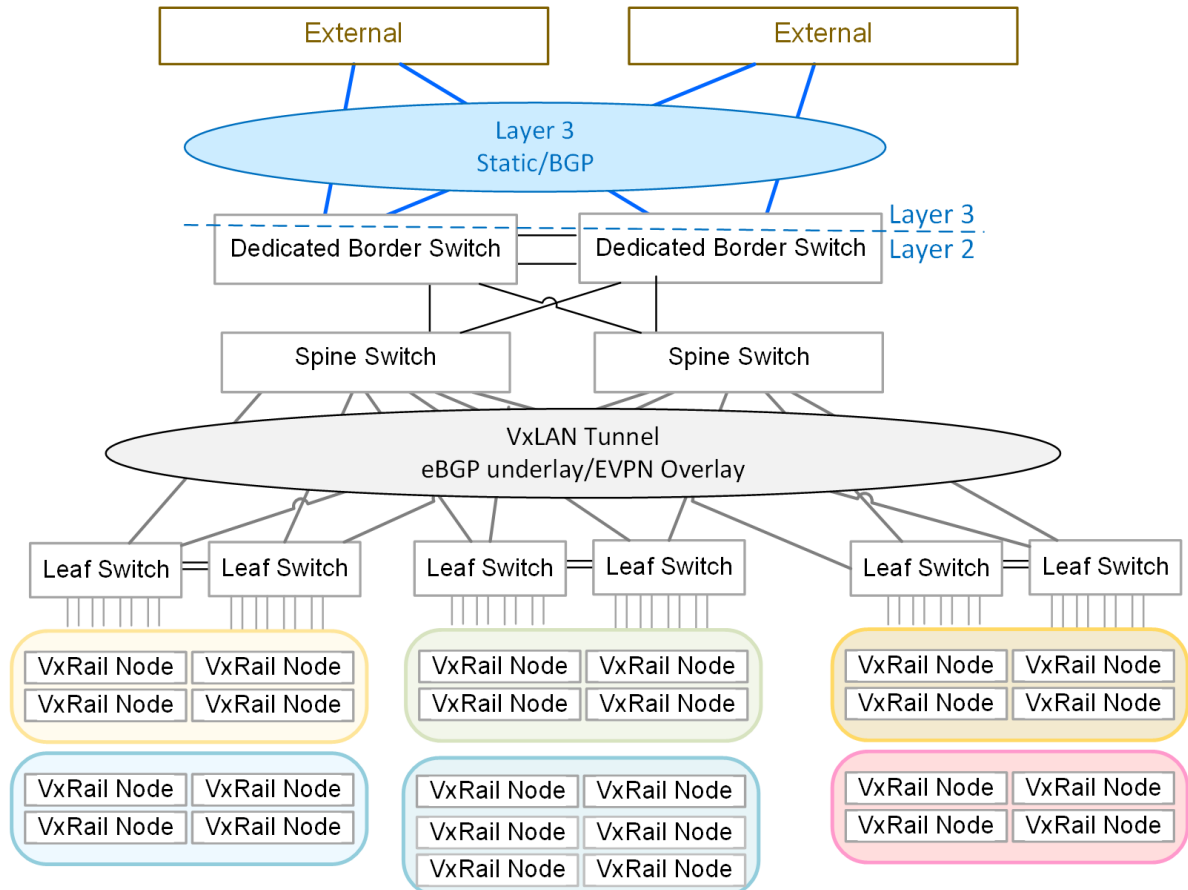


Figure 14 VxRail-SmartFabric Topology: Multi-Rack with Dedicated Border Leafs.

4 VxRail physical network connectivity planning

Deploying VxRail cluster with SFS involves careful planning for hardware and software compatibility between the nodes and switches, and also involves considerations for future expansion.

4.1 Consult the VxRail Support Matrix

The section under the heading 'Dell Networking SmartFabric Support' in the [VxRail Support Matrix](#) on the Dell EMC support site lists all the hardware, software, and feature compatibility for VxRail partnered with a switch fabric based on SFS. It is important to note that there are more restrictions when implementing a VxRail cluster on a network with SFS than to a standard network infrastructure.

4.2 VxRail hardware selection

An existing VxRail cluster cannot be migrated to a SmartFabric-enabled network. Plan on a new VxRail hardware to build the cluster. Existing VxRail hardware can be repurposed if they are unused in an inventory location.

For any planning efforts for a new VxRail cluster, a sizing exercise using the VxRail sizer tool is the best practice. This sizing exercise will produce a report with the bill of materials listing the recommended nodes to be deployed to meet your business and operational requirements. There is some flexibility with regard to the Ethernet port selection for the individual VxRail nodes.

- SmartFabric-enabled switches support either 10 GbE or 25 GbE for connectivity for the VxRail nodes.
- SFP28, SFP+, or RJ45 can be selected for the Ethernet ports on the VxRail nodes. If you select 10 GbE connectivity, select either SFP+ or RJ45, as these combinations cannot be mixed in a VxRail cluster.
- Two or four Ethernet ports can be connected from the NDC on each VxRail node to a leaf switch to support VxRail network traffic. If your requirements include network services beyond what is provided by VxRail (NSX-T, for instance), it is best practice to reserve additional Ethernet ports on each VxRail node for this purpose.

You will need to plan rack space to accommodate both the VxRail nodes and the switches, depending on the selected network topology. The VxRail E-Series nodes require 1U of rack space, while the P-Series, V-Series, G-Series, and S-Series nodes each require 2U of rack space.

4.3 Dell EMC networking switch selection

It is best practice to deploy a new network environment to support VxRail with SmartFabric networking services, unless you have compatible and unused Dell EMC Ethernet switches in inventory. Enabling SFS on compatible Dell EMC Ethernet switches reverts the switch back to its factory default state, and any existing network configuration on the switch is removed.

- See the [VxRail Support Matrix](#) for the specific switch models supported for SmartFabric with VxRail.
- For the leaf layer, either the S4100 (10 GbE) series switches or the S5200 (25 GbE) series switches can be deployed.
- The switch model selected for the leaf layer is dependent on the Ethernet ports to be configured on the VxRail nodes. Select the S4100 or S4200 series switches for 10 GbE data port connectivity, and the S5200 series for 25 GbE data port connectivity.

- The port type (SFP28, SFP+, RJ45) on the switch must match the port type on the selected VxRail nodes.
- For the spine layer, the Z9200 series switch or the S52532 switch models are supported.

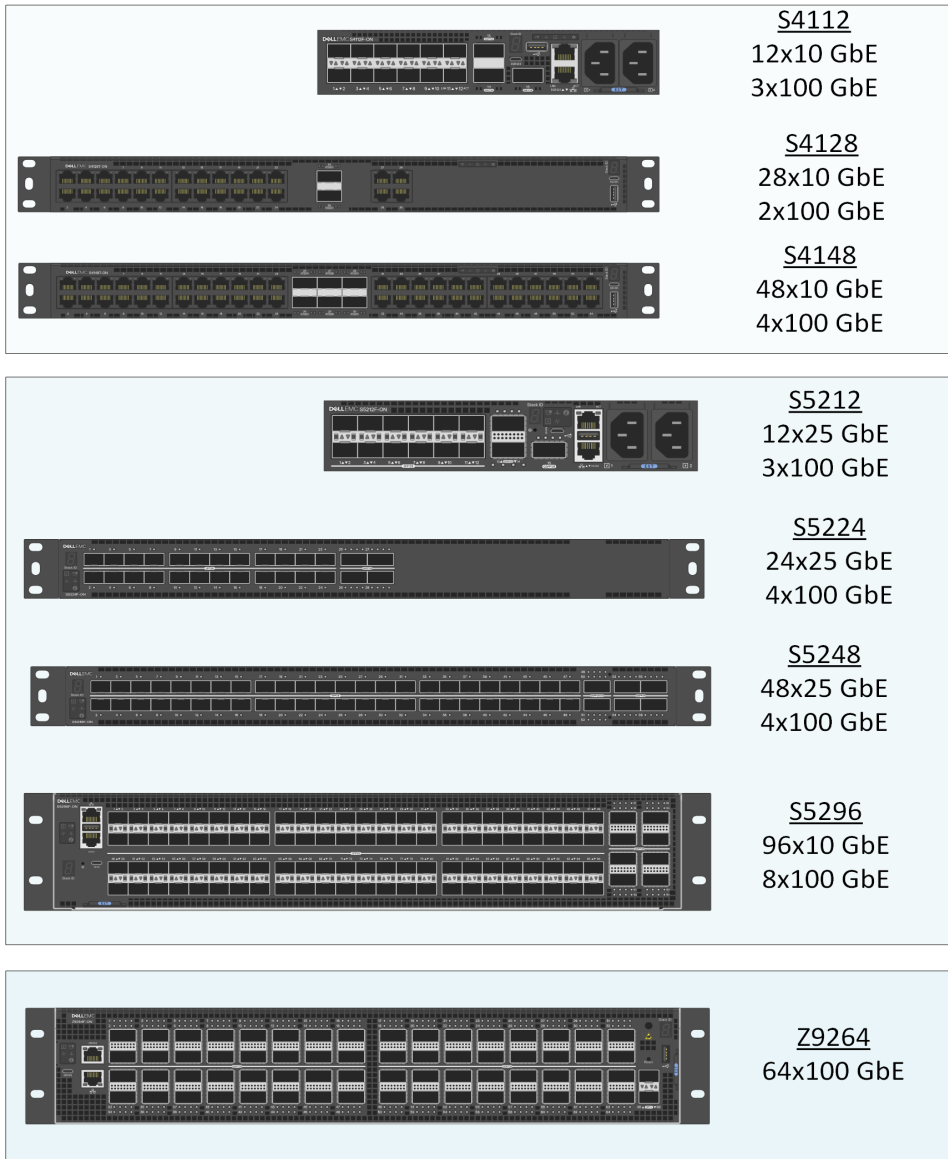


Figure 15 Dell EMC 10GbE/25GbE/100GbE switches supported for VxRail with SFS

The following points must be considered for switch model selection:

- Each VxRail node requires either two or four data ports on the leaf switch.
- Two ports on each leaf switch must be reserved for VLT (Virtual Link Trunking), a requirement for SFS.
- For multi-rack, reserve two ports on each leaf switch for redundant connectivity to the spine switches

For rack space planning purposes, each physical rack will require space for a pair of leaf switches. All leaf switches except the S5296 require 1U of rack space, while the S5296 requires 2U of rack space. If the

switch fabric expands to encompass multiple racks, additional space will be needed in one of the racks for the spine switches.

See the [Dell Networking Optics Specification Sheet](#) for more information.

4.4 Data center core network physical connectivity

Each of the Dell EMC switch models supported for SFS with VxRail includes ports targeted for uplink connections to the data center network.

- You can choose the egress point from the SmartFabric infrastructure to your upstream core network at either the leaf layer, spine layer, or with dedicated border switches per the selected network topology.
- Changing the upstream connectivity option between the leaf layer and spine layer if the SmartFabric-enabled network expands from a leaf-only architecture to a leaf-spine architecture is supported.
- For the ports on the switches targeted for upstream connectivity, the 100GbE Ethernet ports on the Dell switches are based on QSFP28 and the 40GbE ports are based on QSFP+.
- If the 100 GbE or 40 GbE ports are not compatible with the planned data center network, port breakout is supported on the Dell EMC Ethernet switches for VxRail with SFS. The port breakout feature reconfigures selected 40 GbE ports to run at 10 GbE speed, or reconfigure selected 100 GbE ports to run at either 25 GbE, 40 GbE, or 50 GbE. This feature can be set from the user interface derived from SmartFabric primary switch.
- See the [Dell Networking Optics Specification Sheet](#) for complete details on qualified cables and adapters for Dell EMC Ethernet switches.

5 Task outline for VxRail cluster on SmartFabric-based network

An understanding of the implementation experience is key to completing the tasks needed to support this effort, and the prerequisites that must be met before starting the implementation effort. The implementation experience differs depending on decisions made for connectivity upstream to the data center network.

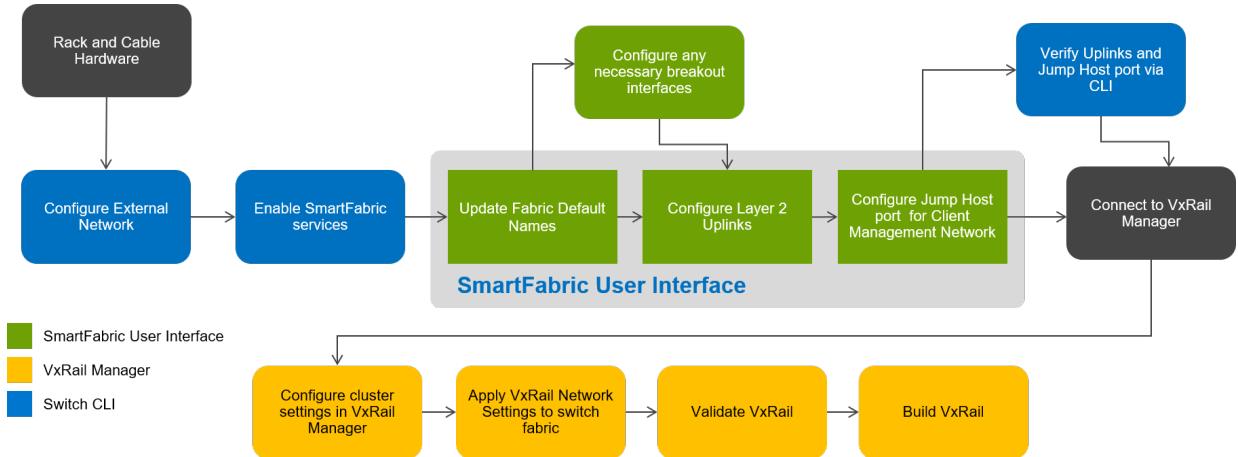


Figure 16 VxRail-SmartFabric network deployments milestones with Layer 2 uplinks

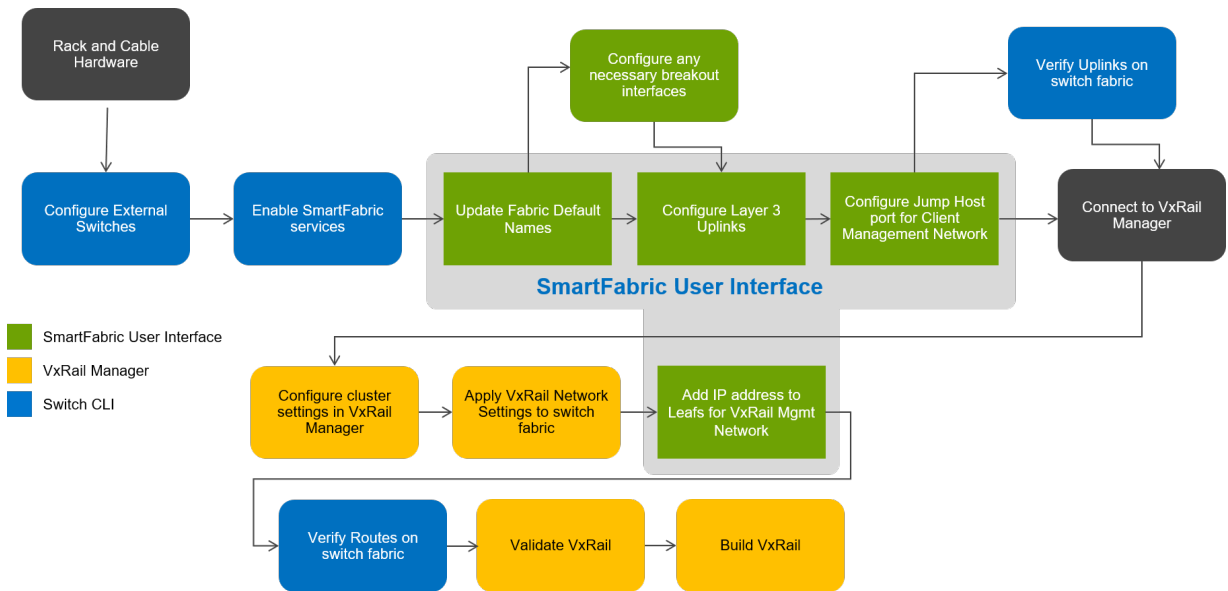


Figure 17 VxRail-SmartFabric network deployments milestones with Layer 3 uplinks

Following is a more detailed outline of the major tasks that are performed to deploy a VxRail cluster onto a SmartFabric-based network:

1. First rack installation
 - a. Install two leaf switches into first rack.
 - b. Install VxRail nodes in first rack to perform initial cluster build.
2. Physical Network Connectivity
 - a. Connect VxRail nodes to data ports on each leaf switch.
 - b. Connect VxRail iDRAC ports to IT management switched network.

- c. Connect two leaf switches to form VLT.
 - d. Connect leaf switches to upstream core network.
 - e. Connect leaf switch out-of-band management ports to IT management switched network.
3. Initial Power-on
 - a. Power on the VxRail nodes
 - b. Configure initial iDRAC settings on each VxRail node.
 - c. Power on the leaf switches
 - d. Configure initial management settings on two leaf switches.
 - e. Verify OS10 software version and update if applicable.
 4. Create SmartFabric network.
 - a. Apply OS10 Enterprise license to each leaf switch.
 - b. Verify VxRail nodes and peer leaf switch visible on each leaf switch.
 - c. Enable SFS on each leaf switch and reboot each switch.
 5. Initialize SmartFabric network for VxRail cluster implementation.
 - a. Connect to the primary switch user interface.
 - b. Configure unique names for SmartFabric components.
 - c. Configure breakout ports, if applicable.
 - d. Configure uplinks for connectivity to upstream network.
 - e. Configure port on leaf switch for jump host access.
 6. Perform Initial Build of VxRail cluster.
 - a. Connect laptop or jump host to leaf switch jump port.
 - b. Connect to VxRail Manager on the primary VxRail node.
 - c. Discover VxRail nodes and switch fabric.
 - d. Enter settings and properties to build VxRail cluster.
 - e. Enable Layer 3 upstream connectivity for VxRail external management network (if applicable)
 - f. Build VxRail cluster on switch fabric.
 7. Deploy OMNI (Open Management Network Interface) plug-in.
 - a. Connect to vCenter instance supporting VxRail cluster.
 - b. Deploy OMNI plug-in on vCenter instance.
 - c. Register OMNI plug-in with vCenter instance

The remaining optional tasks outlined below are performed depending on the scale and scope of the implementation.

These are the milestones for a SmartFabric network that scales beyond a single rack:

1. Install two spine switches in the first rack or one of the new racks.
2. Install two leaf switches in the expansion racks.
3. Connect out-of-band management port of each switch to IT management switched network.
4. In the expansion racks, make two connections between the leaf switches to form a VLT.
5. Make a connection from each leaf switch in the first rack and each new rack to each spine switch.
6. Power on switches
7. Enable SFS on each spine switch and reboot each switch.
8. Enable SFS on each new leaf switch and reboot each switch.

9. Configure initial management settings on each switch.
10. Use the primary switch user interface or OMNI plug-in to enter friendly names and perform uplink management tasks, if necessary.

These are the milestones for a VxRail cluster that scales beyond a single rack:

1. Install VxRail nodes into expansion rack.
2. Connect VxRail nodes to data ports on each leaf switch.
3. Connect VxRail iDRAC ports to IT management switched network.
4. Power on VxRail nodes
5. Connect to VxRail Manager.
6. Select powered-on nodes visible in VxRail Manager and add to VxRail cluster.

6 VxRail and SmartFabric Network Planning

A set of pre-requisites and preparation steps must be fulfilled before moving forward with an implementation of the SmartFabric-based network and VxRail cluster. These steps include best practices to ensure a seamless integration into the target data center.

6.1 VxRail and SmartFabric Networks Overview

VxRail clusters are dependent on interconnectivity with the physical network for overall management and virtual machine traffic. VxRail depends on a properly configured SmartFabric-based network to gain access to required applications and services, and to enable end-user access.

The following image displays the networks that are created when implementing a VxRail cluster on a SmartFabric-based network. From the VxRail cluster, only the external management network and guest networks created for guest virtual machines require access to the upstream production network. In the following example, it is at the leaf layer.

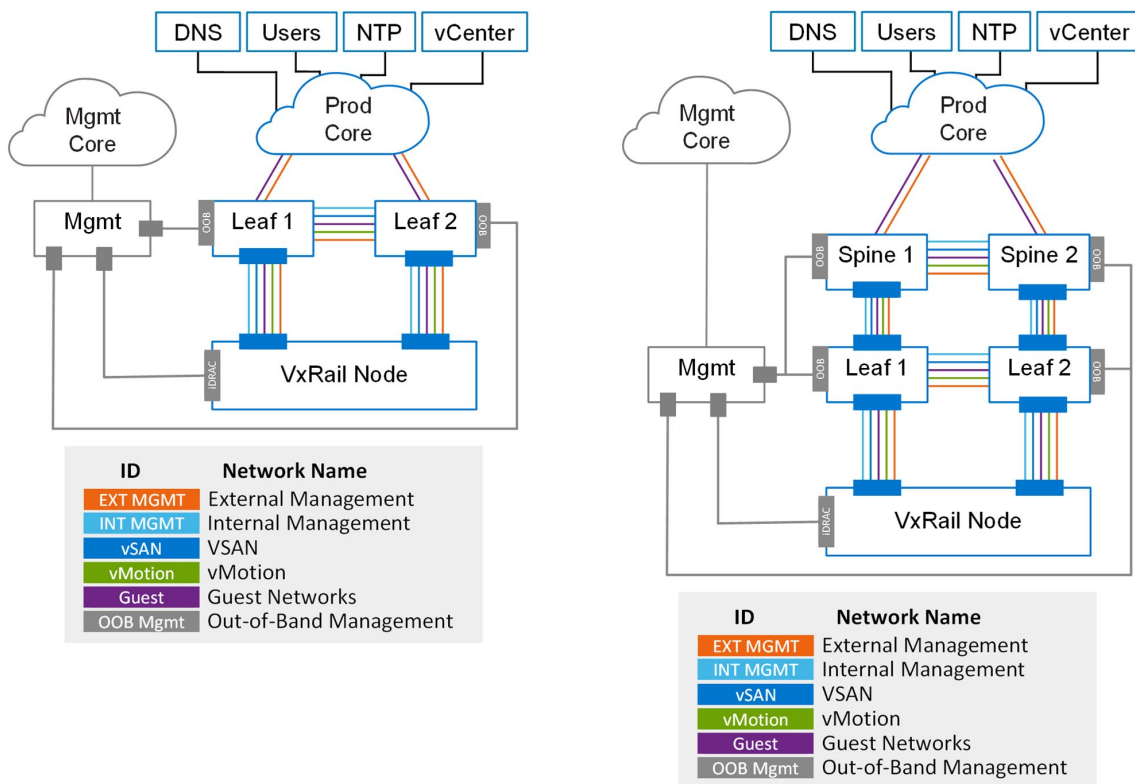


Figure 18 Leaf and Spine uplinks for VxRail networks with out-of-band management network

The VxRail external management network must be able to connect to the required data center services at the time of initial build and for continued operations, including DNS and NTP. The traffic for the other VxRail required networks is configured by default to stay within the confines of the VxRail cluster.

A VxRail cluster can be deployed to either an embedded vCenter instance that is included with the product, or to a customer-supplied external vCenter. VxRail Manager is a vCenter plug-in, so it requires access to the vCenter instance.

6.2 Data Center Components and Services

6.2.1 IT management switched network

Network connectivity is required for access to the iDRAC interfaces for the VxRail nodes, which are based on Dell EMC PowerEdge servers, and SSH and HTTPS access to the leaf and spine switches. HTTPS access to the primary switch is required to perform the initial settings in the SmartFabric network to support the VxRail cluster. For both the nodes and the switches, a 1 GbE switch or pair of 1 GbE switches not enabled for SmartFabric is required.

6.2.2 Domain Name Services (DNS)

Every VxRail node and VxRail management component requires a forward and reverse DNS entry. VxRail Manager must be able to connect to the external DNS service over the external VxRail management network at the time of initial build of the cluster, and for ongoing operations.

6.2.3 Network Time Protocol (NTP)

It is best practice for VxRail to leverage an NTP service to assure a consistent date and time across all cluster components and services. Certain services within the cluster depend on a consistent and standardized date and time for proper operations. It is best practice and strongly recommended for VxRail Manager to connect to the external NTP service over the external VxRail management network at the time of initial build of the cluster, and for ongoing cluster operations.

6.2.4 External vCenter

At the time of initial VxRail build, a decision must be made whether to use the embedded vCenter instance included with the VxRail software bundle as the foundation for the cluster, or to use a compatible vCenter instance. If an external option is chosen, VxRail Manager must be able to link with the vCenter instance over the external VxRail management network.

6.2.5 FTP/SFTP Server

FTP/SFTP services are the preferred method for uploading OS10 software bundles and OS10 licenses locally to the Dell-branded Ethernet switches. FTP/SFTP services are also the preferred method for uploading the software upgrade bundle to the OMNI plug-in. The best practice is to deploy an FTP/SFTP server with an IP address accessible to the IT management network.

6.2.6 Console Access Server

While assigning an IP address to the out-of-band management port of the switches is required, it is a best practice to deploy a serial concentrator to provide console access to the pool of Dell EMC Ethernet switches to ease management responsibilities.

6.3 IP Address Reservations for IT Management Network

Fixed IP addresses must be assigned to the components and applications that need to connect to the IT management network. Each row in the 'IT Management Network IP Address Reservation' table in the [VxRail with SmartFabric Network Configuration Tables](#) section.

6.3.1 Network Settings for SmartFabric switches

Each leaf and spine switch that will be a member of the SmartFabric network must be assigned a fixed IP address for out-of-band management. When a SmartFabric network is formed, one switch is selected as the primary switch. In the event of an outage of that primary switch, another switch in the fabric is assigned primary status. Connectivity to that primary switch using a browser is required to enter the initial settings on the fabric before building the VxRail cluster.

6.3.2 Network Settings for VxRail Node Management

The underlying server hardware for VxRail nodes are Dell EMC PowerEdge servers. All PowerEdge servers include a management interface called the 'integrated Dell Remote Access Controller' (iDRAC) for access to the server hardware and firmware. A fixed IP address must be assigned for iDRAC network port for connectivity to the IT management network.

6.3.3 Network Settings for FTP/SFTP Server

It is best practice to deploy an FTP/SFTP server in the data center and enable connectivity to the IT management network. The server is used for uploading software images and licenses to the Dell EMC switches, and software updates to the OMNI plug-in.

6.3.4 Network Settings for OMNI vCenter plug-in

The OMNI plug-in is a free download from the Dell EMC networking support site. It must be deployed on a vCenter instance in the data center with access to the SmartFabric management network either through the IPV4 out-of-band management network or IPV6 internal discovery network. If the plug-in is deployed on the embedded vCenter instance where the VxRail cluster is deployed, the IPV6 internal discovery network can be used for this purpose.

To enable access to the OMNI plug-in from a vClient, the plug-in must be able to register with the vCenter instance using its assigned hostname or IPV4 address. A single OMNI plug-in can manage multiple SmartFabric network instances, provided it can connect to the primary switch in each SmartFabric instance through the out-of-band management network.

6.4 Update Default SmartFabric Network Settings

The 'SmartFabric Network Default Settings' table in the [VxRail with SmartFabric Network Configuration Tables](#) section lists the default values that are applied when the VxRail personality profile is enabled and a switch fabric is formed. The default values can be updated from the SmartFabric primary switch user interface. Any changes made to the default values resets the switch fabric and will force a reload of the fabric switches. This task should be performed before initial build of the VxRail cluster.

The value set for the client control network is isolated to the switch fabric and should not cause a conflict. The default value of 3939 is the same value configured in the factory into every VxRail node for discovery purposes. If the default value for the client control network is changed, this new value must be set on each standard virtual switch on each of the VxRail nodes to enable discovery on the internal VxRail management network.

The value set for the client management network is isolated to the switch fabric by default. This virtual network is configured as 'untagged' on the switch fabric, and is configured as an 'untagged' virtual network on each switch port that is connected to a powered-on VxRail node. This virtual network is used primarily to

enable a laptop or jump host to connect to VxRail Manager through a jump port on a leaf switch to perform initial build.

The Autonomous System Number (ASN) value set for the leaf switches and spine switches is relevant only if BGP will be used for routing purposes. These values only need to be changed if BGP will be used, and if there is a conflict with the data center network.

The network settings for the underlay and overlay network are used to enable connectivity across racks in the switch fabric. These networks are isolated on the switch fabric, and should not need to be changed.

7 VxRail Cluster Initial Build Planning

The VxRail Cluster Network Configuration Tables section has tables describing the properties and settings that are collected and used for the deployment of a VxRail cluster. These properties and settings are the same whether the cluster is built on a standard network configured manually, or built on a SmartFabric-enabled network configured with the VxRail personality profile.

NOTE: Dell EMC professional services will capture and record the values in these tables, and use them to deploy the VxRail cluster in the data center. Capturing these values before VxRail cluster deployment is critical for preparing the data center network for that phase.

7.1 VxRail Cluster Network VLAN Reservation

Whether a decision is made to enable Layer 2 or Layer 3 connectivity upstream, VLANs for the required VxRail cluster networks must be reserved. The 'VxRail Required Network VLANs' table in the Settings' in the [VxRail Cluster Network Configuration Tables](#) section lists the VxRail networks required when performing initial build.

- It is best practice to use a VLAN entered for the VxRail External Management Network that is different than the VLAN representing the virtual network for the Client Management Network. This will allow for additional VxRail clusters to be deployed on the switch fabric and preserve the virtual network assigned to the client management network for its intended purpose of access to discovered nodes and client onboarding. This VLAN captured in VxRail Manager during initial build will be configured as a tagged virtual network on the switch fabric.
- The default VLAN for the VxRail Internal Management Network (3939) is configured into each VxRail node in the factory to standardize the network ID used for node and switch discovery. This VLAN must match the value representing the virtual network for the Client Control Network configured into the SmartFabric-based network. If a new VLAN ID is used for the VxRail Internal Management Network, then the value for the Client Control Network must be changed to match this VLAN. In addition, this VLAN must be changed on each VxRail node to run in the switch fabric to enable discovery. This can only be done through the ESXi console on a powered-on VxRail node.

7.2 VxRail Management Network Reservation

The VxRail management components and the VxRail nodes are specific guidelines for network properties:

- The VxRail management components and the VxRail nodes are connected to the VxRail External Management Network at the time of VxRail cluster build.
- The VxRail External Management Network must be able to connect to required data center services (DNS, NTP, and so forth).
- Each VxRail management component and the VxRail node in the VxRail cluster must have an assigned IP address and hostname.
- The IP address scheme must be a public IP address range that can pass through the SmartFabric network uplinks to the upstream network.
- The IP address VxRail management components and the VxRail nodes must be fixed (no DHCP).
- The IP addresses assigned to the VxRail management components and the VxRail nodes cannot be in use.
- The IP addresses assigned to the VxRail nodes must be contiguous.
- The IP address range assigned must all be in the same subnet.

The 'VxRail Management Network Properties' table in the [VxRail Cluster Network Configuration Tables](#) section shows what needs to be captured.

7.3 VxRail Cluster Settings

The table 'VxRail Cluster Initial Build Settings' in the [VxRail Cluster Network Configuration Tables](#) section shows what needs to be captured.

7.3.1 VxRail Management Network IP Address Assignments

- For the category 'VxRail Management Network IP Address Assignments', assign an individual IP address to each the VxRail management components and an IP address range for the VxRail nodes.
- Do not assign an IP address for the embedded vCenter and Platform Service Controller components if the cluster is to be built on an external, customer-supplied vCenter instance.
- Do not assign an IP address to Log Insight if it will not be deployed.

7.3.2 VxRail Cluster Global Settings

- For the category 'System, capture the IP address of the DNS and NTP services to be used to support the VxRail cluster.
- Capture the time zone the VxRail cluster is to be deployed.
- Verify the VxRail External Management network will be able to reach these system services.

7.3.3 VxRail node hostnames

All VxRail node hostnames in a VxRail cluster are defined by a naming scheme that comprises: an ESXi hostname prefix (an alphanumeric string), a separator ("None" or a dash "-"), an iterator (Alpha, Num X, or Num 0X), an offset (empty or numeric), a suffix (empty or alphanumeric string with no.) and a domain.

These settings are captured under 'VxRail node hostnames' under the category 'Management' in the table. Follow the samples in the table below to define a supported format.

	Example 1	Example 2	Example 3
Prefix	host	myname	esxi-host
Separator	None	-	-
Iterator	Num 0X	Num X	Alpha
Offset		4	
Suffix		lab	
Domain	local	college.edu	company.com
Resulting hostname	host01.local	myname-4lab.college.edu	esxi-host-a.company.com

- Count the number of VxRail nodes planned for the cluster.
- Make a list of the hostnames based on the selected format.
- Assign an IP address to each hostname. Assign the IP address with the lowest octet to the first hostname in order, and perform the assignment in sequential order.

7.3.4 VxRail vCenter Identification

- See the [VxRail vCenter Server Planning Guide](#) for complete details on the planning steps for vCenter to support VxRail.
- For the category 'Management', a decision needs to be made whether the vCenter instance bundled with VxRail will be used to support the cluster, or if an external, customer-supplied vCenter will support the cluster.
- For an embedded vCenter:
 - Select a vCenter Server hostname.
 - Select a Platform Services Controller hostname.
- For an external vCenter:
 - Capture the information listed in the table.
 - Note that the VxRail External Management network must be able to reach the external vCenter instance.

7.3.5 vMotion Settings

- Each VxRail node must be assigned an IP address for vMotion.
- VxRail using ESXi supports a TCP-IP stack for vMotion that is separate than the default TCP-IP stack.
- The IP address range can be a public IP address range or a private IP address range.

7.3.6 vSAN Settings

- Each VxRail node must be assigned an IP address for vSAN.
- vSAN traffic is isolated on the VxRail cluster, so a private IP address range is the default.

7.3.7 VM Networks Settings

- A minimum of one guest network must be created at the time of VxRail cluster initial build.

7.3.8 Solutions Settings

- Logging supports either the Log Insight appliance included with VxRail, or an external syslog instance can be used instead.

7.4 Prepare DNS for VxRail cluster deployment

Lookup records must be created in the DNS that will support the VxRail cluster for every VxRail management component and every VxRail node. These components include VxRail Manager, embedded (VxRail) vCenter Server, embedded (VxRail) Platform Service Controller, Log Insight, and each VxRail node. The DNS entries must be both forward and reverse records.

8 Data Center Network Preparation

After decisions have been made about the selected network architecture and scale, and the planning efforts for the VxRail cluster are completed, certain tasks must be undertaken in the data center network based on those past efforts to prepare for the integration of the VxRail cluster and the SmartFabric-enabled network into the destination data center. The preparation steps will vary depending upon the decisions made for the network settings assigned to the VxRail cluster and the uplink integration into the data center network.

VxRail traffic traverses the SmartFabric network using VxLAN tunneling. Decision needs to be made about which networks are passed upstream, and by what method these networks are passed upstream.

8.1 Enable Jumbo Frames

Jumbo frames are enabled in the switch fabric to support the added MTU size to support VxLAN. If the upstream network switches support jumbo frames, then that setting should be configured on the switches.

8.2 Layer 2 Network Integration

A SmartFabric network supporting VxRail will support Layer 2 upstream connectivity from the leaf switch layer. The primary switch user interface or OMNI plug-in can configure either a static or dynamic (LACP) port channel on the selected leaf switch ports.




Figure 19 Layer 2 uplinks from SmartFabric network using link aggregation

Networks created on the SmartFabric switch fabric for VxRail are represented as virtual networks, and backed by VLANs. The numerical value used to represent a virtual network in the SmartFabric network is the same as the numerical value used to represent a VLAN within the VxRail cluster, and is the same as the numerical value passed upstream in a Layer 2 uplink.

To integrate the upstream network with the SmartFabric network to support Layer 2 connectivity, the following tasks need to be performed:

- A virtual link tunnel configured between the receiving upstream switches to enable redundancy
- The VLANs for the VxRail external management network and any guest networks planned for upstream access must be configured on the upstream network switches.
- The VLANs must be configured on the port channels used to support the uplinks from the downstream SmartFabric-enabled switches.
- The settings (static or dynamic) must be compatible between port channels on the upstream and downstream switches.
- The port channels must be assigned to the switch ports connected to the uplinks from the downstream SmartFabric-enabled switches.
- Routing services must be configured on the upstream layer where the Layer 2 networks representing the VxRail External Management Network and guest networks terminate.

After cluster initial build is complete and operational, VxRail integration with the SmartFabric network includes support for the automated creation of new virtual networks on the switch fabric driven by updates made to the virtual distributed switches within the VxRail cluster. For instance, if a new port group for a new guest network is created on the VxRail virtual distributed switch, the new VLAN setting on that portgroup is captured and applied to the switch fabric as a virtual network. The switch fabric can be directed to automatically add these virtual networks to the port channel. In this event, these same VLAN values must be configured on the upstream network switches and added to the port channels.

8.3 Layer 3 Network Integration

A SmartFabric-enabled network for VxRail will support Layer 3 upstream connectivity from either the leaf layer, or from the spine layer in a multi-rack network deployment. The initial uplinks for Layer 3 routing can be applied to interfaces on the first set of switches in the fabric using the user interface from the primary switch. This method will be necessary for initial build of the VxRail cluster in order to apply an IP address to the VxRail external management virtual network ID to enable routing.

Layer 3 routing can be enabled by selecting the interfaces on either the leaf or spine switches, and applying the policy for static routing or using Border Gateway Protocol. Assigning Layer 3 routing properties directly to a VLAN used for upstream network traffic is also supported.

8.3.1 Point-to-Point Links

Layer 3 integration between the upstream network and SmartFabric network will require point-to-point links to carry the routed networks. Because the point-to-point links are between two ports, a private IP addressing scheme of '192.168.x.x/31' is recommended.

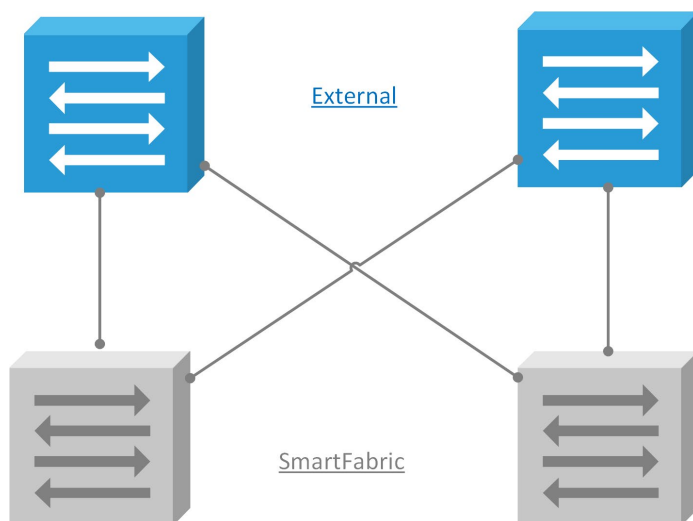


Figure 20 Point-to-point links between SmartFabric switches and external network

Dedicated physical cabling between the ports on the upstream and downstream switches must be in place to support the point-to-point links.

8.3.2 Layer 3 Integration with Static Routing

Static routing can be configured for a SmartFabric network. At a minimum, the VxRail external management network must have passage upstream to enable connectivity to data center services and for connectivity to end users.

The 'Static Point-to-Point Links' in the [VxRail with SmartFabric Network Configuration Tables](#) section can be used to capture the values to enable static routing.

- Identify the VxRail network addresses that will need passage to the upstream network. The IP address range should be defined in CIDR format (<IP Address>/<Prefix Length>). The VxRail external management network and VxRail guest networks will require passage upstream.
- Identify the interfaces on the switch fabric and on the upstream network switches are connected to support Layer 3 the point-to-point links.
- Capture the IP addresses for the point-to-point links. The next-hop address is the IP address assigned to the port upstream in the point-to-point link.
- Configure static routes on the upstream network to support passage for the VxRail external management network and VxRail guest networks.

8.3.3 Layer 3 Integration with Border Gateway Protocol (BGP)

A SmartFabric network will support Layer 3 upstream connectivity from either the leaf layer, or the spine layer, in a multi-rack network deployment using Border Gateway Protocol. Internal BGP (iBGP) is enabled within the SmartFabric network to share routing information within the ASN. External BGP is used for router sharing with the upstream network with a unique ASN representation.

The 'BGP Point-to-Point Links' and 'BGP Autonomous System Numbers' tables in the [VxRail with SmartFabric Network Configuration Tables](#) section can be used to capture the settings for the BGP peering.

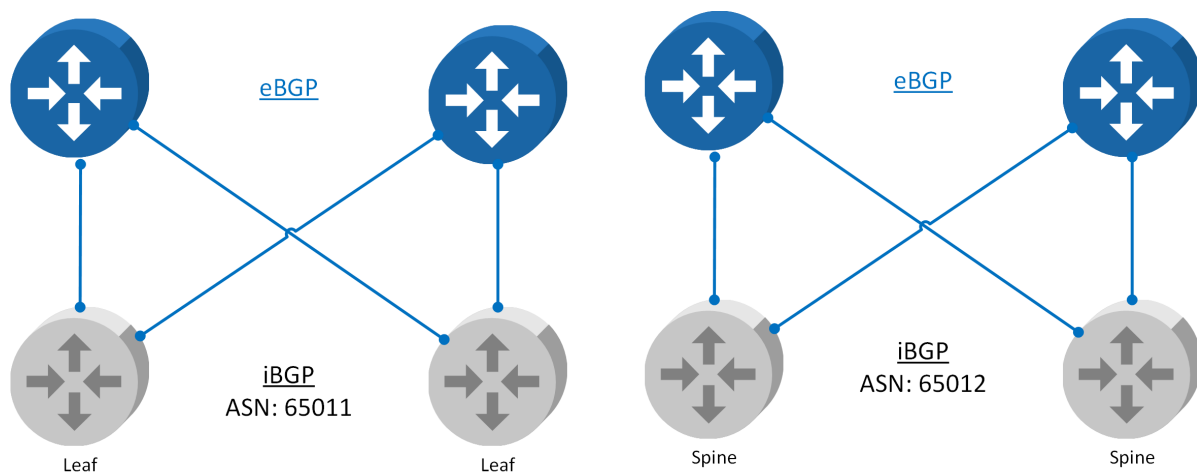


Figure 21 BGP Point-to-point from leaf layer or spine layer

To prepare the data center upstream network to support BGP integration, the following tasks must be performed.

- Decide whether to enable BGP peering at either the leaf layer or the spine layer. For a new VxRail implementation with a new SmartFabric network, BGP peering must be done initially at the leaf layer. Routing can later be migrated to the spine layer if desired as a day 2 operation.
- Capture point-to-point link settings between the upstream network switches and switches at either the leaf layer or the spine layer. A private IP address range of '192.168.xxx.xxx/31' is recommended for peer-to-peer links.
- Enable interfaces on upstream network switches for point-to-point connectivity with switches at either the leaf layer or the spine layer.
- Configure BGP on the upstream network that the SmartFabric switches will peer with. The settings for linking with downstream SmartFabric switches with a neighbor IP address and a neighbor ASN must be configured on the upstream network. Be aware that the SmartFabric ASN value differs between the leaf layer and spine layer, so be sure to capture and apply the correct ASN value.

8.4 Enable IP Addressing for a Virtual Network

The SmartFabric user interface from the primary switch or the OMNI plug-in can be used to assign IP addressing to a virtual network. This process will allow the traffic from this virtual network to be routable to the upstream network. This assignment process should be done after Layer 3 routing planning is complete and Layer 3 routing is established between the switch fabric and upstream network.

The 'Virtual Network – Layer 3 Settings' table in the [VxRail with SmartFabric Network Configuration Tables](#) section can be used to capture the settings for a virtual network that requires passage upstream with Layer 3 network. The VxRail external management network, which is configured as a virtual network by SFS, needs to be assigned Layer 3 properties to enable connectivity upstream to data center services and to end users. For the VxRail external management network, the VLAN to be used for the VxRail external management network will be assigned an equivalent virtual network ID by SFS.

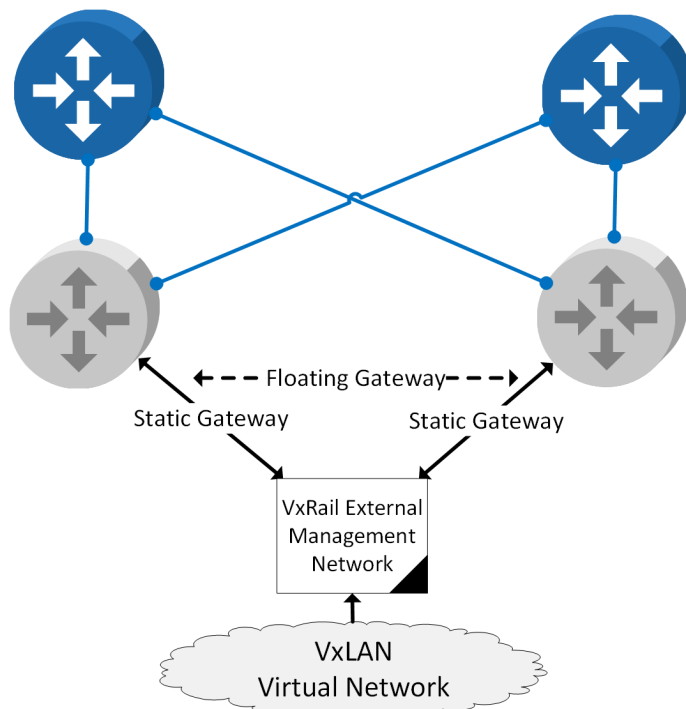


Figure 22 Gateway address assignments for Layer 3 enablement for the VxRail external management virtual network

The Layer 3 settings for a virtual network include configuring a static IP address to each leaf or spine switch (depending on whether upstream connectivity is from the leaf or spine layer) for passageway for the virtual network, and a floating IP address that represents the default gateway for the network. A typical addressing policy for gateways is to reserve IP addresses at either the beginning (x.x.x.1-x.x.x.5 for example) or ending (x.x.x.254-x.x.x.250) of the address range.

The 'Virtual Network – Layer 3 Settings' table in the [VxRail with SmartFabric Network Configuration Tables](#) section can be used to capture the settings needed to enable upstream passage for a virtual network on the switch fabric.

8.5 Layer 3 Integration using a VLAN

In limited use cases, there may be a requirement for VLAN to be configured to enable routing services with the upstream network. In this event, the same rules apply for preparing for Layer 3 integration using either static routing or BGP, but the exception is that a VLAN is assigned for the transport between the two network layers.

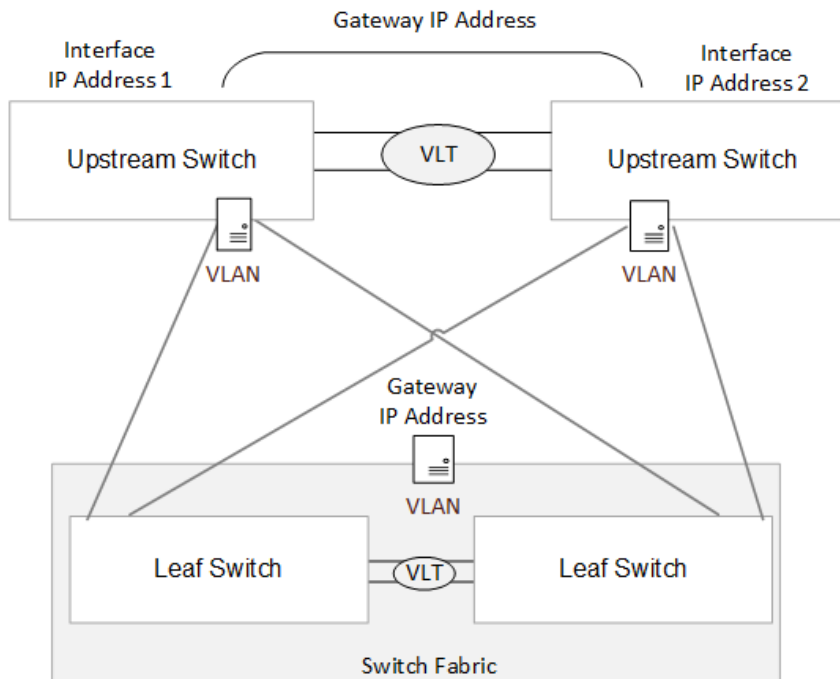


Figure 23 Layer 3 integration using a VLAN

To prepare for Layer 3 integration using a VLAN, the following tasks need to be performed:

- Identify the VLAN to be used for transport services between the two network layers.
- Configure the upstream network for routing services between the two network layers. Configure static routes if static routing will be used, or configure BGP if BGP peering will be used.
- Configure the VLAN to be used for Layer 3 transport services on the upstream network.
- Identify the network to be routed using VLAN-enabled routing services.

The 'VLAN – Layer 3 Settings' table in the [VxRail with SmartFabric Network Configuration Tables](#) section can be used to capture the setting to enable routing services with the upstream network using a VLAN.

8.6 Spanning Tree Integration

The function of spanning tree is to prevent loops in Layer 2 networks, which can cause congestion and reach such a high saturation level that network performance is impacted. Spanning tree enablement is part of the SmartFabric automation process when a switch fabric is created.

8.6.1 Spanning Tree Options for OS10 and SFS

Dell Ethernet switches running OS10 in full switch mode support RSTP (Rapid Spanning Tree Protocol), MST (Multiple Spanning Tree) and RPVST+ (Rapid Per-VLAN Spanning Tree Plus). RPVST+ is the spanning tree protocol configured by default in SmartFabric mode, with RSTP as an optional configuration setting.

The primary difference with RPVST+ to RSTP is that RPVST+ can create separate spanning tree instances on a per-VLAN instance. This feature makes it more efficient than other static spanning tree protocols that support a single network topology, and why it is the default setting for a SmartFabric-enabled network.

Another similar spanning tree protocol, PVST (Per-VLAN Spanning Tree), also creates a network topology on a per-VLAN basis. However, RPVST+ differs from PVST in that it sends industry standard BPDU (Bridge Protocol Data Unit) packets between the switches. The BPDU packets exchanged across the switches contain information about key metrics like priority and cost, and are used to manage an adaptive network topology. PVST is not currently a supported protocol when running SFS.




Figure 24 Static RSTP topology compared to per-VLAN support for RPVST+

8.6.2 Spanning Tree Guidelines for VxRail and SmartFabric networks

At the time SFS is enabled on a set of switches, the initial spanning tree setting is applied to the switch fabric. SFS will disable spanning tree on the switch ports connected to the VxRail nodes. These switch ports are configured by SFS as trunked ports. When the VxRail cluster is then built on the switch fabric, all the networks created for VxRail are VxLAN-backed virtual networks, and not VLAN-backed networks. The VxLAN protocol does not support the broadcast of BPDU packets, so there is no requirement of switch ports connected to VxRail nodes needing to process BPDU packets as edge ports.

However, once a switch fabric expands to a multi-rack configuration, the recommendation is that RSTP, if used, should be disabled and RPVST+ enacted for a SmartFabric-enabled switch network. RPVST+ is the default setting in SFS, and is the spanning tree setting configured to the switch fabric. In a multi-rack topology with leaf and spine switches, RPVST+ will enforce that the links between the leaf and spine switches will not be active/active. This is due to the common VLAN requirements supporting the VxLAN network across racks in a leaf and spine topology.

In a multi-rack switch fabric, RPVST+ applies updated topology rules if changes are made to the switch fabric as a part of day 2 operations to prevent loops in the network. In addition, SmartFabric also enforces a single path topology for the Layer 2 networks between the ingress and egress points on the switch fabric to prevent looping with the upstream network.

For Layer 3 connectivity, enabling static routing or peering through BGP is the recommended method. Spanning tree settings on the upstream network are not impacted with the Layer 3 routed option. The default RPVST+ setting can stay in effect on the leaf and spine switch network.

Spanning tree rules can come into effect in either of these conditions:

- Layer 2 VLAN-based networks from the leaf switch layer to the upstream network
- Layer 3 properties are assigned to a VLAN for passage to the upstream network.

For Layer 2 VLAN-based networks:

- The uplink created by SFS is a port channel, either LAG or VLTi MLAG (multi-chassis LAG). Traffic distribution follows standard load balancing principles across all the physical links within the port channel.
- SFS enforces a single logical uplink is active to prevent looping with the upstream network.
- SFS enforces that a Layer 2 network is only allowed to be added on one of the Layer 2 uplinks in the SmartFabric network, negating the requirement for spanning tree to enforce a single path rule.

For Layer 3 VLAN-backed networks, RPVST+ is the only spanning tree option. In these instances:

- The receiving switches upstream must be running RPVST+ to enable spanning tree integration with the SmartFabric-based network.
- Configure root bridges in the upstream network to optimize distribution of the VLAN traffic.




Figure 25 Per-VLAN root bridges at distribution layer for RPVST+

A VxRail with SmartFabric Network Configuration Tables

The Dell EMC service representative will use a data collection workbook to capture the settings needed to build the VxRail cluster. The workbook will include the following information:

IT Management Network IP Address Reservation		
Setting	Item	Description
Out-of-band network	Leaf Switch 1	IP Address
	Leaf Switch 2	IP Address
	Spine Switch 1	IP Address
	Spine Switch 2	IP Address
	Leaf Switch 1	IP Address
	Leaf Switch 2	IP Address
	Forwarding Route	IP Address
iDRAC network	VxRail Node 1	IP Address
	VxRail Node 2	IP Address
	VxRail Node 3	IP Address
	VxRail Node 4	IP Address
	VxRail Node 5	IP Address
	VxRail Node 6	IP Address
	Subnet Mask	Subnet mask for iDRAC network
	Gateway	Gateway for iDRAC network
Connection to IT management network	FTP/SFTP Server	IP Address
		Subnet Mask
		Gateway
Connection to External VxRail Management Network	OMNI plug-in	Hostname
		IP Address
		Subnet Mask
		Gateway

SmartFabric Network Default Settings		
Item	Default Value	Description
Client Control Network	3939	The client control network is a private, isolated network used for node and switch discovery.
Client Management Network	4091	The client management network is a private, isolated network used for client onboarding for VxRail cluster initial build. This virtual network is configured as untagged on all switch ports with powered-on VxRail nodes at fabric initialization.

Leaf ASN	65011	The ASN for Layer 3 BGP peering from the leaf layer of a SmartFabric network
Spine ASN	65012	The ASN for Layer 3 BGP peering from the spine layer of a SmartFabric network
Private Subnet Prefix	172.16.0.0	IP range used to establish connectivity between leaf switches and spine switches for VTEP.
Private Prefix Length	16	Prefix length for IP range used to establish connectivity between leaf switches and spine switches.
Global Subnet Prefix	172.30.0.0	IP range for overlay network to support VTEP
Global Prefix Length	16	Prefix length for IP range for overlay network to support VTEP

Static Point-to-Point Links					
External Switch		SmartFabric Switch		Point-to-Point Network	Routed Address
Interface	IP Address	Interface	IP Address	IP CIDR (192.168.x.x/31)	xxx.xxx.xxx.xxx/xx

BGP Point-to-Point Links				
Upstream Switch		SmartFabric Switch		Point-to-Point Network
Interface	IP Address	Interface	IP Address	IP CIDR (192.168.x.x/31)

BGP Autonomous System Numbers	
External Network	SmartFabric Network

Virtual Network – Layer 3 Settings		
Item	Setting	Description
Virtual Network	Virtual network ID	The virtual network ID assigned to SFS
Default Gateway	IP Address	Default gateway for IP address range assigned to the virtual network. This gateway is configured as floating gateway between the SmartFabric switches.

Switch 1	IP Address	IP address assigned to first switch. Must be in same subnet as default gateway.
Switch 2	IP Address	IP address assigned to second switch. Must be in same subnet as default gateway.
DHCP Helper	IP Address	Optional. Only needed if DHCP helper to be configured.
VLAN – Layer 3 Settings		
Item	Setting	Description
Uplink VLAN	VLAN	VLAN to be used for Layer 3 network transport
Default Gateway	IP Address	Default gateway for IP address range assigned to the virtual network. This gateway is configured as floating gateway between the SmartFabric switches.
Switch 1	IP Address	IP address assigned to first switch. Must be in same subnet as default gateway.
Switch 2	IP Address	IP address assigned to second switch. Must be in same subnet as default gateway.
DHCP Helper	IP Address	Optional. Only needed if DHCP helper to be configured.

B VxRail Cluster Network Configuration Tables

The Dell EMC service representative will use a data collection workbook to capture the settings required to build the VxRail cluster. The workbook will include the following information:

VxRail Required Network VLANs	
VLAN	Description
External Management Network	Use value unique from Client Management Network VLAN. Network must be able to reach required external services (DNS, NTP, and so forth).
Internal Management Network	This network traffic is isolated in the switch fabric. The default VLAN of 3939 matches the Client Control Network VLAN.
vSAN	Support VxRail storage services. Network is isolated on switch fabric.
vMotion	Supports VM mobility across nodes in cluster. Network by default is isolated on switch fabric.
Guest Network	For non-management general usage. Network must be configured to pass upstream to reach end users and external applications.

VxRail Management Network Properties	
Item	Description
Subnet (in x.x.x.x/x format)	All of the VxRail management components and VxRail nodes must be assigned IP addresses from the same subnet
Subnet Mask	Subnet mask for VxRail management network
Gateway	Gateway for VxRail management network

VxRail Cluster Initial Build Settings		
Category	Item	Description
VxRail Management Network IP Address Assignments	VxRail node starting IP address	IP Address
	VxRail node ending IP address	IP Address
	VxRail Manager	IP Address
	Embedded vCenter	IP Address
	Embedded PSC	IP Address
	Log Insight	IP Address
System	Global settings	Time zone
		NTP servers
		DNS servers

Management	VxRail node hostnames	Prefix
		Separator
		Iterator
		Offset
		Suffix
		Domain
	Embedded vCenter Server (Leave blank if using external vCenter)	vCenter Server hostname
		Platform Services Controller hostname
	External vCenter Server (Leave blank if using embedded VxRail vCenter)	Customer-supplied Platform Services Controller (PSC) Hostname (FQDN) (Leave blank if PSC is embedded in external vCenter Server).
		Customer-supplied vCenter Server hostname (FQDN)
		Customer-supplied vCenter Server SSO domain
		Admin username/password or the newly created VxRail non-admin username and password
		New VxRail management username and password
		Customer-supplied data center name
	VxRail Manager	New cluster name
VxRail hostname		
vMotion		Starting address for IP pool
		Ending address for IP pool
		Subnet mask
		Default Gateway or vMotion Gateway
vSAN		Starting address for IP pool
		Ending address for IP pool
		Subnet mask
VM Networks	(minimum one)	VM Network name for every guest network VLAN
Solutions	Logging (select one)	vRealize Log Insight™ hostname
		External Syslog Server (FQDN)
		None

