

Dell EMC VxRail Network Planning Guide

Physical and Logical Network Considerations and Planning

April 2021

Abstract

This is a planning and preparation guide for VxRail™ Appliances. It can be used to better understand the networking requirements for VxRail implementation. This document does not replace the implementation services with VxRail Appliances requirements and should not be used to implement networking for VxRail Appliances.

Copyright

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 © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA 04/21 Planning Guide H15300.15.

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

Contents

Chapter 1	Introduction	6
	Revision history	7
	Intended use and audience	7
	Introduction to VxRail	8
Chapter 2	Planning Your Data Center Network for VxRail	10
	Introduction	11
	VxRail hardware and the physical network infrastructure	11
	VxRail clusters, appliances, and nodes	12
	Network switch	13
	Data center network	16
	VxRail Ethernet port options	17
	VxRail Ethernet adapter options	18
	VxRail node connectivity options	19
	VxRail networking rules and restrictions	20
	Topology and connections	21
Chapter 3	VxRail Feature-Driven Decision Points	23
	Introduction	24
	Software-defined data center	24
	Dell EMC SmartFabric network mode	25
	vSphere with Kubernetes on VxRail	27
	vSAN stretched-cluster	28
	2-node cluster	29
Chapter 4	VxRail Hardware and Switch Selection Decision Points	31
	VxRail nodes and switches	32
Chapter 5	Planning the VxRail Implementation	34
	Introduction	35
	Plan data center routing services	35
	Plan for multi-rack VxRail cluster	37
	Plan for a vSAN HCI mesh topology	38
	Plan for VxRail custom uplink assignments	39
	Plan for link aggregation of VxRail networks	41
	Plan link aggregation on switch port pairs	44

Chapter 6	Planning the VxRail Cluster Implementation	45
Introduction		46
Decide on VxRail single point of management		46
Decide on VxRail network traffic segmentation		47
Plan the VxRail logical network		49
Plan network exclusions reserved for VxRail Manager		54
Plan network settings for VxRail management components		55
Identify IP addresses for VxRail management components		57
Select hostnames for VxRail management components		58
Identify external applications and settings for VxRail		60
Prepare customer-supplied vCenter server		61
Prepare customer-supplied virtual distributed switch		63
Reserve IP addresses for VxRail vMotion network		67
Reserve IP addresses for VxRail vSAN network		68
Decide on VxRail logging solution		68
Assign passwords for VxRail management		69
Prepare for Dell EMC SmartFabric Services enablement		70
Chapter 7	Configure the Network for VxRail	73
Introduction		74
Setting up the network switch for VxRail connectivity		74
Setting up the upstream network for VxRail connectivity		79
Confirm your data center network		80
Confirm your data center environment		82
Chapter 8	Preparing to Build the VxRail Cluster	83
Introduction		84
Configuring a workstation/laptop for VxRail initialization		84
Perform initialization to create a VxRail cluster		85
Chapter 9	VxRail Network Considerations After Implementation	87
Introduction		88
Configure teaming and failover policies for VxRail networks		88
Using unassigned physical ports for VxRail networks		90
Configure link aggregation on VxRail networks		92
Deploy second VDS for VxRail networking		93
Appendices		95
Appendix A: VxRail Network Configuration Table		96
Appendix B: VxRail Passwords		99
Appendix C: VxRail Setup Checklist		100
Appendix D: VxRail Open Ports Requirements		102

Appendix E: Virtual Distributed Switch Portgroup Default Settings	104
Appendix F: Physical Network Switch Examples	107

Chapter 1 Introduction

This chapter presents the following topics:

Revision history7

Intended use and audience7

Introduction to VxRail8

Revision history

Date	Description
April 2019	First inclusion of this version history table Support of VMware Cloud Foundation on VxRail
June 2019	Support for VxRail 4.7.210 and updates to 25 GbE networking
August 2019	Support for VxRail 4.7.300 with Layer 3 VxRail networks
February 2020	Support for new features in VxRail 4.7.410
March 2020	Support for new functionality in vSphere 7.0
April 2020	Support for <ul style="list-style-type: none"> VxRail SmartFabric multirack switch network Optional 100 GbE Ethernet and FC network ports on VxRail nodes
May 2020	Update switch requirements for VxRail IPv6 multicast
June 2020	Update networking requirements for multirack VxRail clusters
July 2020	Support for new features in VxRail 7.0.010
August 2020	Update requirement for NIC redundancy enablement
September 2020	Outline best practices for link aggregation on non-VxRail ports
October 2020	Support for new features in VxRail 7.0.100
October 2020	Remove references to Log Insight.
November 2020	Remove requirement for VxRail guest network during initial configuration.
February 2021	Support for new features in 7.0.131
April 2021	<ul style="list-style-type: none"> Additional content on mixing of node ports in VxRail clusters Option for manual node ingestion instead of IPV6 multicast Added additional content for LACP policies Update to stretched cluster node minimums

Intended use and audience

This guide discusses the essential network details for VxRail deployment planning purposes only. It introduces best practices, recommendations, and requirements for both physical and virtual network environments. This document has been prepared for anyone that is involved in planning, installing, and maintaining VxRail, including Dell Technologies

field engineers, and customer system and network administrators. Do not use this guide to perform the installation and set-up of VxRail. Work with your Dell Technologies service representative to perform the actual installation.

Introduction to VxRail

Dell EMC VxRail™ Appliances are a hyperconverged infrastructure (HCI) solution that consolidates compute, storage, and network into a single, highly available, unified system. With careful planning, VxRail Appliances can be rapidly deployed into an existing data center environment, and the end-product is immediately available to deploy applications and services.

VxRail is not a server. It is an appliance that is based on a collection of nodes and switches that are integrated as a cluster under a single point of management. All physical compute, network, and storage resources in the appliance are managed as a single shared pool. They are allocated to applications and services based on customer-defined business and operational requirements.

The compute nodes are based on Dell EMC PowerEdge servers. The G Series consists of up to four nodes in a single chassis, whereas all other models are based on a single node. An Ethernet switch is required, at speeds of either 1/10/25 GbE, depending on the VxRail infrastructure deployed. A workstation or laptop for the VxRail user interface is also required.

VxRail has a simple, scale-out architecture, leveraging VMware vSphere® and VMware vSAN™ to provide server virtualization and software-defined storage, with simplified deployment, upgrades, and maintenance through VxRail Manager. Fundamental to the VxRail clustered architecture is network connectivity. It is through the logical and physical networks that individual nodes act as a single system providing scalability, resiliency, and workload balance.

The VxRail software bundle is preloaded onto the compute nodes, and consists of the following components (specific software versions not shown):

- VxRail Manager
- VMware vCenter Server™
- VMware vSAN
- VMware vSphere
- Dell Secure Remote Support (SRS)/VE
- VMware vRealize Log Insight™

Licenses are required for VMware vSphere and VMware vSAN. The vSphere licenses can be purchased from Dell Technologies, VMware, or your preferred VMware reseller partner.

The VxRail Appliances also include the following licenses for software that can be downloaded, installed, and configured:

- Dell EMC RecoverPoint for Virtual Machines (RP4VM)

- Five full VM licenses per single node (E, V, P, D, and S series)
- Fifteen full VM licenses for the G Series chassis

Chapter 2 Planning Your Data Center Network for VxRail

This chapter presents the following topics:

- Introduction11
- VxRail hardware and the physical network infrastructure11
- VxRail clusters, appliances, and nodes12
- Network switch.....13
- Data center network.....16
- VxRail Ethernet port options17
- VxRail Ethernet adapter options18
- VxRail node connectivity options19
- VxRail networking rules and restrictions20
- Topology and connections.....21

Introduction

The network considerations for VxRail are no different than those of any enterprise IT infrastructure: availability, performance, and extensibility. VxRail Appliances are delivered to your data center ready for deployment. The nodes in the appliance can attach to any compatible network infrastructure at 1/10/25 GbE speeds with either RJ45 or SFP+ ports. Models with single processors can attach to compatible 1 GbE network infrastructure. Most production VxRail network topologies use dual top-of-the-rack (ToR) switches to eliminate the switch as a single point of failure. This document guides you through the key phases and decision points for a successful VxRail implementation. The key phases are:

Step 1: Select the VxRail hardware and physical network infrastructure that best aligns with your business and operational objectives.

Step 2: Plan and prepare for VxRail implementation in your data center before product delivery.

Step 3: Set up the network switch infrastructure in your data center for VxRail before product delivery.

Step 4: Prepare for physical installation and VxRail initialization into the final product.

Note: Follow all the guidance and decision points described in this document; otherwise, VxRail will not implement properly, and it will not function correctly in the future. If you have separate teams for network and servers in your data center, you must work together to design the network and configure the switches.

VxRail hardware and the physical network infrastructure

VxRail nodes connect to one or more network switches, with the final product forming a VxRail cluster. VxRail communicates with the physical data center network through one or more virtual distributed switches that are deployed in the VxRail cluster. The virtual distributed switches and physical network infrastructure integration provide connectivity for the virtual infrastructure, and enable virtual network traffic to pass through the physical switch infrastructure. In this relationship, the physical switch infrastructure serves as a backplane, supporting network traffic between virtual machines in the cluster, and enabling virtual machine mobility and resiliency. In addition, the physical network infrastructure enables I/O operations between the storage objects in the VxRail vSAN datastore, and provides connectivity to applications and end-users outside of the VxRail cluster.

This section describes the physical components and selection criteria for VxRail clusters:

- VxRail clusters, appliances, and nodes
- Network switch
- Data Center Network
- Topology and connections
- Workstation/laptop
- Out-of-band management (optional)

VxRail clusters, appliances, and nodes

A VxRail appliance consists of a set of server nodes that are designed and engineered for VxRail. A VxRail physical node starts as a standard Dell PowerEdge server. The Dell PowerEdge server next goes through a manufacturing process following VxRail product engineering specifications to produce a VxRail node ready for shipment. A set of prepared VxRail nodes is delivered to the customer site based on a purchase order. The set of VxRail nodes is delivered ready for data center installation and connectivity into the data center network infrastructure.

Once the data center installation and network connectivity are complete, and the equipment is powered on, the VxRail management interface is used to perform the initialization process, which forms the final product: a VxRail cluster.

A standard VxRail cluster starts with a minimum of three nodes and can scale to a maximum of 64 nodes. The selection of the VxRail nodes to form a cluster is primarily driven by planned business use cases, and factors such as performance and capacity. Five series of VxRail models are offered, each targeting specific objectives:

VxRail Series	Target Objective
E-Series	Balanced Compute and Storage, Space Optimized (1U1N chassis)
V-Series	Virtual Desktop Enablement
P-Series	High Performance
S-Series	Storage Dense
G-Series	Compute Dense, Space Optimized (2U4N chassis)
D-Series	Durable, ruggedized, short-depth platforms designed to withstand extreme conditions

Each VxRail model series offers choices for network connectivity. The following figures show some of the physical network port options for the VxRail models.




Figure 1. Back view of VxRail E-Series on Dell 14th Generation PowerEdge server




Figure 2. Back view of VxRail V-, P-, and S-Series on Dell 14th Generation PowerEdge server




Figure 3. Back view of VxRail G-Series on Dell 14th Generation PowerEdge server

In addition to network connectivity, review the physical power, space, and cooling requirements for your planned infrastructure to ensure data center compatibility.

Network switch

A VxRail cluster depends on adjacent Ethernet switches, commonly referred to as ‘top-of-rack’ switches, to support cluster operations. VxRail is broadly compatible with most Ethernet switches on the market. For best results, select a switch platform that meets the operational and performance criteria for your planned use cases.

VxRail relationship with the Ethernet switch

The VxRail product does not have a backplane, so the adjacent ‘top-of-rack’ switch enables all connectivity between the nodes that comprise a VxRail cluster. All the networks (management, storage, virtual machine movement, guest networks) configured within the VxRail cluster depend on the ‘top-of-rack’ switches for physical network transport between the nodes, and upstream to data center services and end-users.

The network traffic configured in a VxRail cluster is Layer 2. VxRail is architected to enable efficiency with the physical ‘top-of-rack’ switches through the assignment of virtual LANs (VLANs) to individual VxRail Layer 2 networks in the cluster. This functionality will ease network administration and integration with the upstream network.

VxRail node discovery and the Ethernet switch

The VxRail product has two separate and distinct management networks. One management network extends externally to connect to end users, office sites and other applications. The second management network is isolated, visible only to the VxRail nodes.

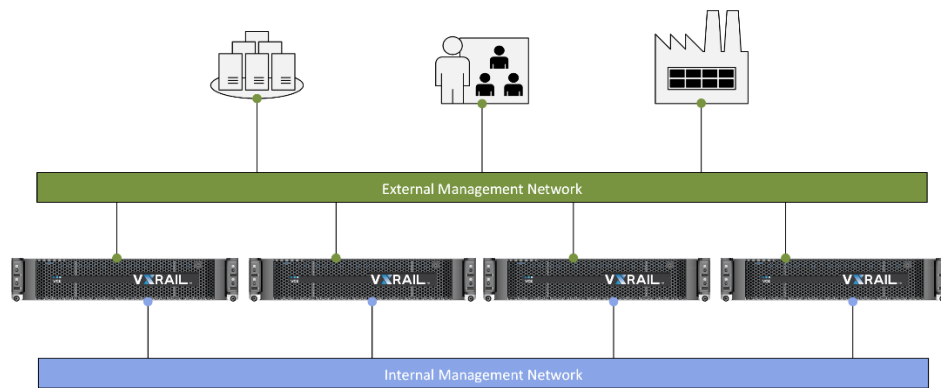


Figure 1 VxRail Management Networks

The network that is visible only to the VxRail nodes depends on IPv6 multicasting services configured on the adjacent 'top-of-rack' switches for node discovery purposes. One node is automatically designated as the primary node. It acts as the source, and listens for packets from the other nodes using multicast. A VLAN assignment on this network limits the multicast traffic to the interfaces connected to this internal management network only.

A common Ethernet switch feature, Multicast Listener Discovery (MLD) snooping and querier, is designed to further constrain the flooding of multicast traffic by examining MLD messages, and then forwarding multicast traffic only to interested interfaces. Since the traffic on this node discovery network is already constrained through the configuration of this VLAN on the ports supporting the VxRail cluster, this setting may provide some incremental efficiency benefits, but does not negatively impact network efficiency.

If your data center networking policy has restrictions for the IPV6 multicast protocol, IP addresses can be manually assigned to the VxRail nodes as an alternative to automatic discovery.

Basic switch requirements

The switch does not need to support Layer 3 services or be licensed for Layer 3 services.

A VxRail cluster can be deployed in a 'flat' network using the default VLAN on the switch, or be configured so that all the management, storage, and guest networks are segmented by virtual LANs for efficient operations. For best results, especially in a production environment, only managed switches should be deployed. A VxRail cluster that is built on a 'flat' network should be considered only for test cases or for temporary usage.

Switch performance considerations

In certain instances, additional switch features and functionality are necessary to support specific use cases or requirements.

- If your plans include deploying all-flash storage on your VxRail cluster, 10 GbE network switches are the minimum requirement for this feature. Dell Technologies recommends a 25 GbE network if that is supported in your data center infrastructure.
- Enabling advanced features on the switches planned for the VxRail cluster, such as Layer 3 routing services, can cause resource contention and consume switch buffer space. To avoid resource contention, select switches with sufficient resources and buffer capacity.
- Switches that support higher port speeds are designed with higher Network Processor Unit (NPU) buffers. An NPU shared switch buffer of at least 16 MB is

Network redundancy and performance considerations

recommended for 10 GbE network connectivity, and an NPU buffer of at least 32 MB is recommended for more demanding 25 GbE network connectivity.

- For large VxRail clusters with demanding performance requirements and advanced switch services that are enabled, consider switches with additional resource capacity and deeper buffer capacity.

Decide if you plan to use one or two switches for the VxRail cluster. One switch is acceptable, and is often used in test and development environments. To support sustained performance, high availability, and failover in production environments, two or more switches are required.

VxRail is a software-defined data center which depends on the physical top-of-rack switching for network communications, and is engineered to enable full redundancy and failure protection across the cluster. For customer environments that require protection from a single point of failure, the adjacent network supporting the VxRail cluster must also be designed and configured to eliminate any single point of failure. A minimum of two switches should be deployed to support high availability and balance the workload on the VxRail cluster, linked with a pair of cables to support the flow of Layer 2 traffic between the switches.

Consideration should also be given for link aggregation to enable load balancing and failure protection at the port level. NIC teaming, which is the pairing of a set of physical ports into a logical port for this purpose, is supported in VxRail versions 7.0.130 and later. These logical port pairings can peer with a pair of ports on the adjacent switches to enable the load balancing of demanding VxRail networks.




Figure 4. Multi-chassis link aggregation across two switches

For network-intense workloads that require high availability, consider switches that support multi-chassis link aggregation, such as Cisco's Virtual Port Channel or Dell's VLT Port Channel. This feature can be used to enable load balancing from the VxRail cluster across a logical switch port configured between the two linked switches.

Support for Link Aggregation Control Protocol (LACP) at the cluster level is also introduced in VxRail version 7.0.130. The switches supporting the VxRail cluster should support LACP for better manageability and broad load-balancing options.

Data center network

VxRail is dependent of specific data center services to implement the cluster and for day-to-day operations. The top-of-rack switches must be configured to the upstream network to enable connectivity to these data center services, and to enable connectivity to the end-user community

Data center services

- Domain Naming Services (DNS) is required to deploy the VxRail cluster and for ongoing operations.
- VxRail cluster depends on Network Time Protocol (NTP) to keep the clock settings on the various VxRail components synchronized. Dell Technologies recommends that a reliable global timing service be used for VxRail.
- Syslog service is supported with VxRail, but is not required.
- VxRail depends on VMware vCenter for cluster management and operations. You can use either the embedded vCenter instance that is included with VxRail, or an external vCenter instance in your data center.

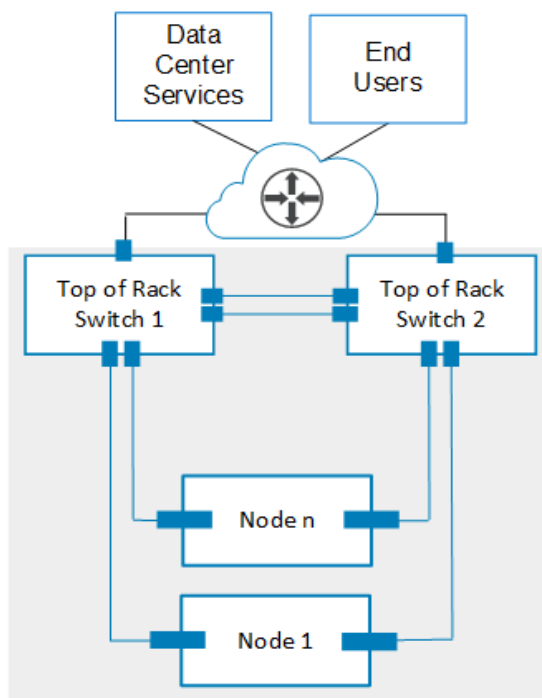


Figure 5. Connecting data center services with VxRail cluster

Routing services

VxRail cluster operations depend on a set of networks that run on both the virtual network inside the cluster and on the adjoining physical network switches. Some of these networks, specifically for VxRail management and for end-user access must be passed to the upstream network, while other VxRail networks can stay isolated on the adjoining network switches.

You must specify a set of Virtual LAN (VLAN) IDs in your data center network that will be assigned to support the VxRail networks. All the VLANs must be configured on the adjoining physical switches connected to the VxRail nodes.

One VLAN is assigned for external VxRail management access. Data center services (such as DNS and NTP) that are required by VxRail cluster must be able to connect to the external VxRail management network. Routing services must be configured to enable connectivity to these services from this network. Additional VLANs, such as those required for end-user access must also be configured to support routing end-users and external applications to the virtual machines running on the VxRail cluster.

If Layer 3 routing services are not configured on the adjacent physical switches, the VLANs that need to pass upstream must be configured on adjoining network switch uplinks, and also on the ports on the upstream network devices, so they can pass through upstream to routing layer. If Layer 3 services are enabled on the adjacent physical switches, configure the VLANs that need to pass upstream to terminate at this layer, and configure routing services for these networks to pass upstream.

VxRail Ethernet port options

The following figures show the appliance connectivity options that are supported on the Network Daughter Cards (NDCs) for each supported VxRail node model. Each diagram also shows the connectivity requirements for the physical node management port. Each diagram also shows the available PCIe-based options that are supported for each VxRail node model, including Fibre channel.

VxRail Intel Node Connectivity Comparison						
	E-Series	P-Series	S-Series	V-Series	D-Series	G-Series
VxRail Base Connectivity	4x10Gb RJ45	2x10 or 4x10Gb RJ45		4x10Gb RJ45	2x10Gb RJ45	N/A
	2x10Gb or 4x10Gb SFP+		4x10Gb SFP+		N/A	2 x10Gb SFP+
	2x25Gb SFP28					
	4x1Gb RJ45			N/A		
Management Port	1x1Gb RJ45 iDRAC 9					
VxRail Optional Connectivity	8x10Gb RJ45	16x10Gb RJ45	12x10Gb RJ45	16x10Gb RJ45	6x10Gb RJ45	4x10Gb RJ45
	4x10Gb SFP+	16x10Gb SFP+	12x10Gb SFP+	16x10Gb SFP+	2x10Gb SFP+	2x10Gb SFP+
	4x25Gb SFP28	8x25Gb SFP28	6x25Gb SFP28	6x25Gb SFP28	2x25Gb SFP28	2x25Gb SFP28
	2x100Gb SFP28	2x100Gb SFP28	2x100Gb SFP28	2x100Gb SFP28	N/A	2x100Gb SFP28
	4x16/32Gb FC	8x16/32Gb FC	6x16/32Gb FC	8x16/32Gb FC	N/A	N/A

Figure 6. VxRail Node Connectivity Options for Intel models

VxRail AMD Node Connectivity Comparison		
	E-Series	P-Series
VxRail Base Connectivity	2x10Gb RJ45	
	2x10/25Gb SFP28	
Management Port	1x1Gb RJ45 iDRAC 9	
VxRail Optional Connectivity	4x10Gb RJ45	
	2x10Gb SFP+	
	2x10/25Gb SFP28	
	N/A	2x100Gb SFP28
	2/4x16/32Gb FC	

Figure 7. VxRail Node Connectivity Options for AMD models

The total number of PCIe adapter cards that can be installed in a VxRail node is dependent on the VxRail model selected and the number of installed CPUs.

The following connectivity rules apply to VxRail nodes based on 14th Generation Dell EMC PowerEdge servers:

- The NDC built into the back of each VxRail node is required and cannot be ordered without one.
- Only one option can be selected for the NDC ports.
- More than one PCIe slot can be populated to support VxRail optional connectivity.
- The maximum number of optional ports supported per node is dependent of the number of PCIe slots supported on the VxRail node and the number of ports on each supported PCIe adapter card. 1 GbE connectivity is supported on single processor models only with a maximum of eight nodes per cluster.
- The 10GbE ports will auto-negotiate to 1 GbE when used with 1 GbE networking.

VxRail Ethernet adapter options

There are restrictions on the models of Ethernet adapters cards and ports that can be configured for VxRail nodes. Each vendor adapter card and firmware select for support with VxRail must pass a set of tests to be qualified. The following table highlights the vendors' networking products that pass qualification and are supported for VxRail:

Port Speed	Vendor
10 GbE	Broadcom
	Intel
	QLogic
25 GbE	Broadcom
	Intel
	Mellanox
	QLogic
100 GbE	Mellanox

The following guidelines should be understood to drive port adapter selection:

- When a VxRail cluster is initially built, it is recommended that all the network adapter cards that are used to form the cluster be of the same vendor. This rule does not apply to nodes added to an existing VxRail cluster, so long as the port speed and port type match the existing nodes.
- VxRail recommends using the same adapter card vendor and model for all the nodes in a cluster that support VxRail cluster operations. There is no guarantee that using optics or cables from one vendor with an adapter card from another vendor will work as expected. VxRail recommends consulting the Dell cable and optics support matrix before attempting to mix vendor equipment in a VxRail cluster.
- The feature sets supported from network adapter card suppliers do not always match. There is a dependency on the firmware and/or driver in the adapter card to support certain features. If there is a specific feature that is needed to meet a business requirement, VxRail recommends consulting with a sales specialist to verify that the needed feature is supported for a specific vendor.

VxRail node connectivity options

For VxRail clusters that can tolerate a single point of failure, and do not have demanding workload requirements, the VxRail cluster can be configured using only the Ethernet ports on the NDC. For workloads that require a higher level of failure protection, VxRail supports spreading the networks across NDC Ethernet ports and Ethernet ports on PCIe adapter cards.

Version 7.0.130 provides flexibility with the selection of the Ethernet ports to support the VxRail cluster. The cluster can be deployed using either the default network profiles supported in VxRail, or you can select the Ethernet ports to be used by the cluster and assigned those ports to the VxRail networks. The default network profile option only supports NDC-based ports, and assigns the uplinks using the predefined settings as described in Appendix F: Physical Network Switch Examples. With the customized port

selection option, you can use just the ports on the NDC, mix the ports from the NDC and a PCIe adapter card, or select only ports from PCIe adapter cards.




Figure 8. Mixing NDC and PCIe ports to support a VxRail cluster

If you later decide to migrate the VxRail networks onto other uplinks on the VxRail nodes, this activity can be performed after the cluster is built, as long as the VxRail cluster is at version 7.0.010 or later.

VxRail networking rules and restrictions

- The Ethernet ports selected during the VxRail initial build process to support the VxRail cluster are reserved exclusively for VxRail usage and cannot be reconfigured for purposes outside of VxRail networking.
- Any unused Ethernet ports on the nodes that are not reserved by the VxRail cluster can be used for other purposes, such as guest networks, NFS, and so forth.
- For VxRail nodes supplied with 10 GbE ports or 25 GbE ports, the VxRail cluster can be configured with either two ports or four ports to support VxRail network traffic.
- For VxRail clusters running all Ethernet ports at 1 GbE speed:
 - Four ports on each node must be reserved for VxRail network traffic.
 - Only hybrid VxRail models can be configured with 1 GbE speed. All-flash VxRail models cannot support 1 GbE.
- Custom Ethernet port configurations are supported with restrictions:
 - Prior to VxRail version 7.0.130, all the Ethernet ports on the VxRail nodes selected for a VxRail cluster must running at the same port type and running at the same speed.
 - Starting in VxRail version 7.0.130, the ports on the NDC and PCIe adapter cards configured in the VxRail nodes can be running at different speeds. For instance, the NDC ports can support 10GbE and the ports on the PCIe adapter cards support 25GbE.
 - The mixing RJ45 and SFP+ Ethernet ports in the VxRail nodes to support VxRail cluster network operations is not restricted, but is not recommended. Mixing different Ethernet port types invites complexity regarding firmware, drivers, and cabling with the data center network.




Figure 9. Mixing network speeds and types to support VxRail networking

Topology and connections

Various network topologies are possible with VxRail clusters. Complex production environments have multi-tier network topologies with clusters in multiple racks, and spanning across data centers. Simpler workloads can be satisfied with the nodes and adjacent switches confined to a single rack, with routing services configured further upstream. A site diagram showing the proposed network components and connectivity is highly recommended *before* cabling and powering on VxRail nodes, and before performing an initial build of the VxRail cluster.

Decide what network architecture you want to support the VxRail cluster, and what protocols will be used to connect to data center services and end users. For VxRail clusters managing production workloads, VLANs will be configured to support the VxRail networks. Determine which network tier the VxRail networking VLANs will terminate, and which tier to configure routing services.

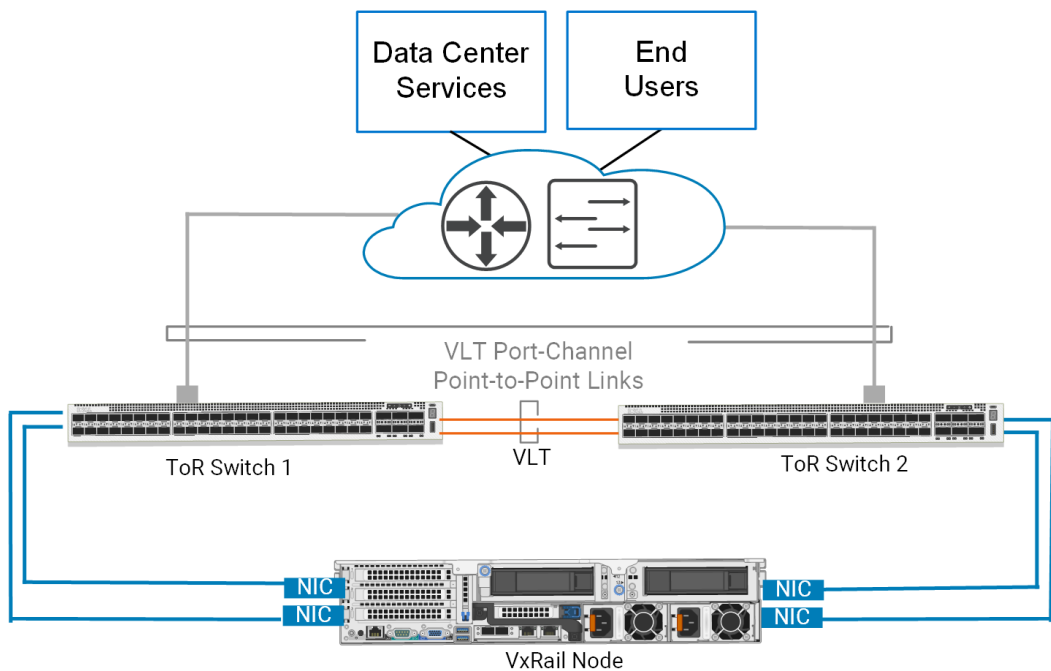


Figure 10. High-level network topology with Layer 2 and Layer 3 options

To determine the base number of ports required on each switch to support VxRail cluster operations, multiply the number of Ethernet ports on each VxRail node you select to support VxRail networking by the number of nodes to be configured into the cluster. For a dual switch configuration, ports must be reserved on each switch to form an inter-switch link for network traffic passage. You will also need to reserve additional ports to pass VxRail network traffic upstream, and one port on a switch to enable a laptop to connect to VxRail to perform initial build.

If the VxRail clusters are located at a data center that you cannot access easily, we recommend setting up an out-of-band management switch to facilitate direct communication with each node.

To use out-of-band management, connect the integrated Dell Remote Access Controller (iDRAC) port to a separate switch to provide physical network separation. Default values, capabilities, and recommendations for out-of-band management are provided with server hardware information. You must reserve an IP address for each iDRAC in your VxRail cluster (one per node).

Chapter 3 VxRail Feature-Driven Decision Points

This chapter This chapter presents the following topics:

Introduction	24
Software-defined data center	24
Dell EMC SmartFabric network mode	25
vSphere with Kubernetes on VxRail	27
vSAN stretched-cluster.....	28
2-node cluster	29

Introduction

Certain applications, software stacks, and product features that are supported on VxRail can impact the architecture, deployment, and operations of the cluster. If your plans for VxRail include any of the feature sets or software stacks that are listed in this section, make note of the requirements that each of these might have on your plans for VxRail.

Software-defined data center

If your plans include the transformation of your current data center with disparate technologies and processes towards a software-defined data center, consider that VxRail can be positioned as a building block towards that eventual outcome. The physical compute, network, and storage resources from built VxRail clusters can be allocated to VMware's cloud management and virtual desktop software solutions, and managed as a logical pool for end-user consumption. By using VxRail clusters as the underlying foundation, the software-defined data center can be designed and deployed to meet specific business and operational requirements.

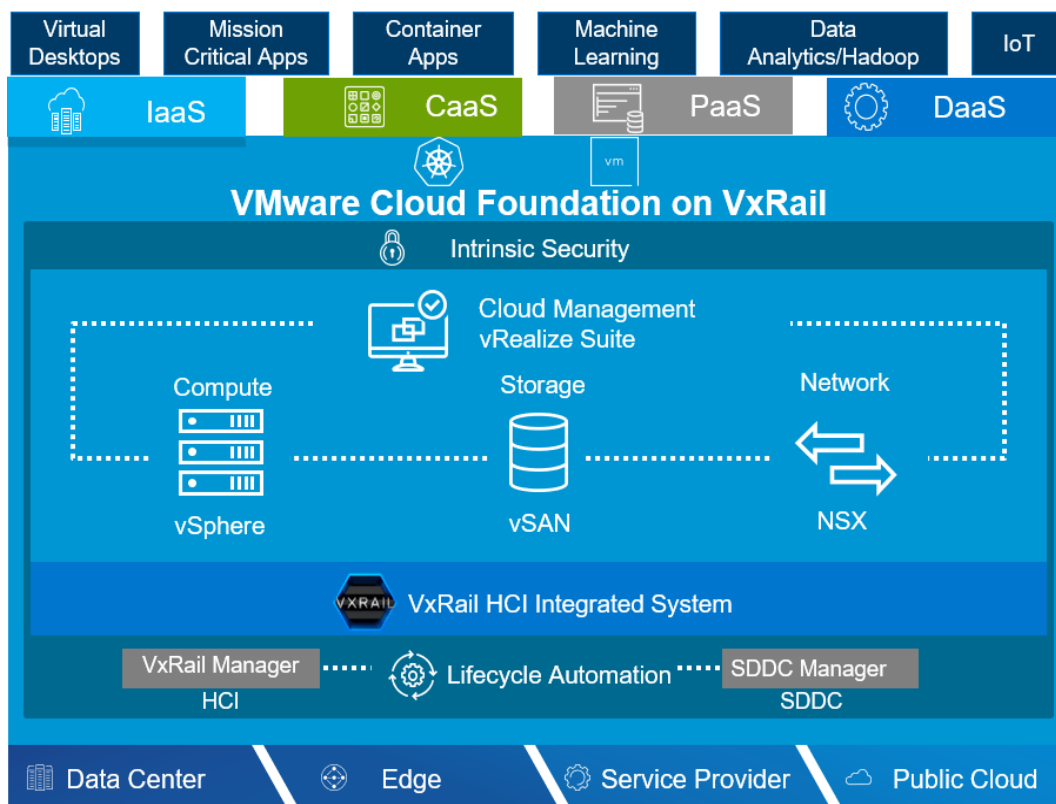


Figure 11. VxRail as the foundation for the software-defined data center

The path starts with a structured discovery and planning process that focuses on business use cases and strategic goals, and that will drive the selection of software layers that will comprise the software-defined data center. Dell Technologies implements the desired software layers in a methodical, structured manner, where each phase involves incremental planning and preparation of the supporting network.

The next phase after the deployment of the VxRail cluster is to layer the VMware Cloud Foundation software on the cluster. This will enable assigning cluster resources as the underpinning for logical domains, whose policies align with use cases and requirements.

The information that is outlined in this guide covers networking considerations for VxRail. For more information about the architecture of VMware Cloud Foundation on VxRail, and to plan and prepare for a deployment of VMware Cloud Foundation on VxRail, go to [Dell Technologies VxRail Technical Guides](#).

Dell EMC SmartFabric network mode

Dell network switches support SmartFabric Services, which enable the configuration and operation of the switches to be controlled outside of the standard management console through a REST API interface. Certain Dell EMC switch models support initializing the switches with a SmartFabric personality profile, which then forms a unified network fabric. The SmartFabric personality profile enables VxRail to become the source for the automated configuration and administration of the Dell switches.

In this profile setting, VxRail uses the SmartFabric feature to discover VxRail nodes and Dell EMC switches on the network, perform zero-touch configuration of the switch fabric to support VxRail deployment, and then create a unified hyperconverged infrastructure of the VxRail cluster and Dell EMC switch network fabric.




Figure 12. Dell EMC SmartFabric for VxRail

For ongoing VxRail cluster network management after initial deployment, the Dell EMC Open Manage Network Interface (OMNI) vCenter plug-in is provided free of charge. The Dell EMC OMNI plug-in enables the integration and orchestration of the physical and virtual networking components in the VxRail-SmartFabric HCI stack, providing deep visibility from the vClient for ease of overall management and troubleshooting. The Dell EMC OMNI plug-in serves as the centralized point of administration for SmartFabric-enabled networks in the data center, with a user interface eliminating the need to manage the switches individually at the console level.

The orchestration of SmartFabric Services with the VxRail cluster means that state changes to the virtual network settings on the vCenter instance will be synchronized to the

switch fabric using REST API. In this scenario, there is no need to manually reconfigure the switches that are connected to the VxRail nodes when an update such as a new VLAN, port group, or virtual switch, is made using the vClient.

The SmartFabric-enabled networking infrastructure can start as small as a pair of Dell EMC Ethernet switches, and can expand to support a leaf-spine topology across multiple racks. A VxLAN-based tunnel is automatically configured across the leaf and spine switches, which enable the VxRail nodes to be discovered and absorbed into a VxRail cluster from any rack within the switch fabric.

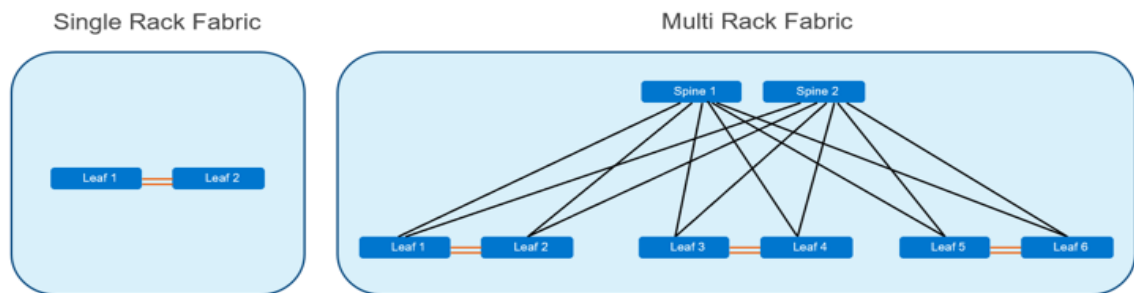


Figure 13. SmartFabric-enabled multirack network expansion

Planning for VxRail with the Dell EMC SmartFabric networking feature must be done in coordination with Dell Technologies representatives to ensure a successful deployment. The planned infrastructure must be a supported configuration as outlined in the [Dell Networking Solutions Support Matrix](#).

Using the Dell EMC SmartFabric feature with VxRail requires an understanding of several key points:

- At the time of VxRail deployment, you must choose the method of network switch configuration. Enabling the VxRail personality profile on the switches resets the switches from the factory default state and enables SmartFabric Services. If you enable SmartFabric Services, all switch configuration functionality except for basic management functions are disabled at the console, and the management of switch configuration going forward are performed with SmartFabric tools or through the automation and orchestration built into VxRail and SmartFabric Services.
- A separate Ethernet switch to support out-of-band management for the iDRAC feature on the VxRail nodes and for out-of-band management of the Dell Ethernet switches is required.
- Disabling the VxRail personality profile on the Dell network switches deletes the network configuration set up by SmartFabric services. If a VxRail cluster is operational on the Dell switch fabric, the cluster must be deployed.
- Non-VxRail devices can be attached to switches running in SmartFabric services mode using the OMNI vCenter plug-in.

For more information about how to plan and prepare for a deployment of VxRail clusters on a SmartFabric-enabled network, see the [Dell EMC VxRail with SmartFabric Planning and Preparation Guide](#). For more information about the deployment process of a VxRail cluster on a SmartFabric-enabled network, go to [VxRail Networking Solutions at Dell Technologies InfoHub](#).

vSphere with Kubernetes on VxRail

If your requirements include workload management using Kubernetes, then a VxRail cluster can be configured as a supervisor cluster for Kubernetes. Kubernetes is a portable, extensible, API-driven platform for the management of containerized workload and services. VMware's Tanzu feature enables the conversion of a VxRail cluster, whose foundation is vSphere, into a platform for running Kubernetes workloads inside dedicated resource pools. A VxRail cluster that is enabled for vSphere with Tanzu is called a Supervisor cluster.

When a VxRail cluster is enabled for vSphere with Kubernetes, the following six services are configured to support vSphere with Tanzu:

- vSphere Pod Service
- Registry Service
- Storage Service
- Network Service
- Virtual Machine Service
- Tanzu Kubernetes Grid Service for vSphere

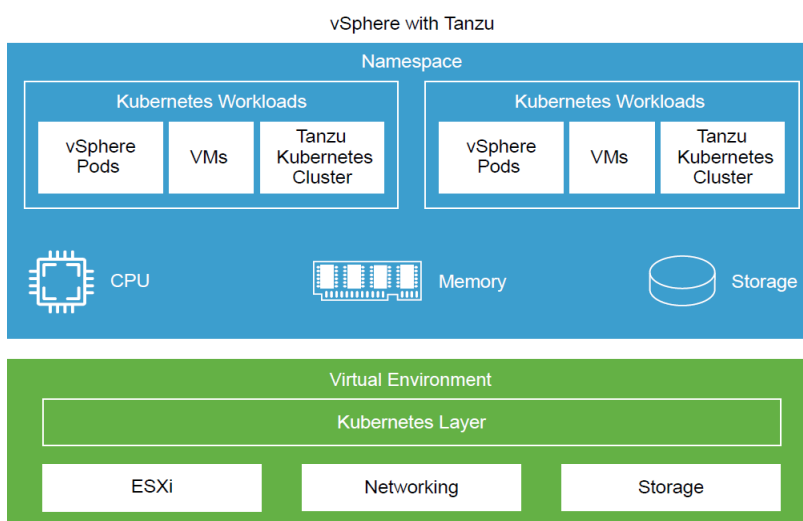


Figure 14. vSphere with Tanzu on a VxRail cluster

As a VxRail administrator using vSphere management capabilities, you can create namespaces on the Supervisor Cluster, and configure them with specified amount of memory, CPU, and storage. Within the namespaces, you can run containerized workloads on the same platform with shared resource pools.

- This feature requires each VxRail node that is part of the Supervisor cluster to be configured with a vSphere Enterprise Plus license with an add-on license for Kubernetes.
- This feature requires portgroups to be configured on the VxRail cluster virtual distributed switch to support workload networks. These networks provide

connectivity to the cluster nodes and the three Kubernetes control plane VMs. Each Supervisor Cluster must have one primary workload network.

- A virtual load balancer that is supported for vSphere must be also configured on the VxRail cluster to enable connectivity from client network to workloads running in the namespaces.
- The workload networks will require reserved IP addresses to enable connectivity for the control plane VMs and the load balancer.

For complete details on enabling a VxRail cluster to support vSphere with Tanzu, see [vSphere with Tanzu Configuration and Management Guide](#).

vSAN stretched-cluster

vSAN stretched-cluster is a VMware solution that supports synchronous I/O on a vSAN datastore over distance and is supported on VxRail. A vSAN stretched-cluster enables site-level failure protection with no loss of service or loss of data.

If you plan to deploy a vSAN stretched-cluster on VxRail, note the following requirements:

- Three data center sites: two data center sites (Primary and Secondary) host the VxRail infrastructure, and the third site supports a witness to monitor the stretched-cluster.
- A minimum of three VxRail nodes in the Primary site, and a minimum of three VxRail nodes in the Secondary site
- A minimum of one top-of-rack switch for the VxRail nodes in the Primary and Secondary sites
- An ESXi instance at the Witness site

The vSAN stretched-cluster feature has strict networking guidelines, specifically for the WAN, that must be adhered to for the solution to work.




Figure 15. vSAN Stretched Cluster Topology

More detailed information about vSAN stretched-cluster and the networking requirements can be found in the [Dell-EMC VxRail vSAN Stretched Cluster Planning Guide](#).

2-node cluster

VxRail supports a solution specifically for small-scale deployments with reduced workload and availability requirements, such as those in a remote office setting. The solution is fixed to two VxRail nodes only, and like the stretched-cluster solution, requires a third site to act as a witness for monitoring purposes.

If you plan to deploy 2-node VxRail clusters, note the following:

- The minimum VxRail software version for the 2-Node cluster is 4.7.100.
- The deployment is limited to a pair of VxRail nodes.
- You cannot expand to three or more nodes unless the cluster is running version 7.0.130 or later.
- Verify that your workload requirements do not exceed the resource capacity of this small-scale solution.
- Only one top-of-rack switch is required.
- Four Ethernet ports per node are required. Supported profiles:
 - 2 x 1 Gb and 2 x 10 Gb
 - 4 x 10 Gb
 - 2 x 25 Gb
- The switch can support either 1 GbE or 10 GbE connectivity.
- Two network topologies are supported for inter-cluster VxRail traffic:
 - All four network ports connect to the top-of-rack switch
 - A pair of network cables connect to create two links between the physical nodes, and the other two network ports connect to the top-of-rack switch
- A customer-supplied external vCenter is required. The customer-supplied external vCenter cannot reside on the 2-Node cluster.
- The Witness is a small virtual appliance that monitors the health of the 2-Node cluster. A Witness is required for the 2-Node cluster.
 - An ESXi instance is required at the Witness site.
 - Up to 64 2-node clusters can share a witness, provided the VxRail version is 7.0.100 or later. There is a 1:1 ratio of Witness per 2-Node cluster for previous versions.
 - Witness can be deployed at the same site as the data nodes but not on the 2-Node cluster.
 - For instances where there are more than one 2-Node clusters deployed at the site, the Witness can reside on a 2-Node cluster it is not monitoring. This configuration requires a VMware RPQ.

- The top-of-rack switch must be able to connect over the network with the Witness site.




Figure 16. 2-Node Cluster Topology with direct connections between nodes

Like the vSAN stretched-cluster feature, the small-scale solution has strict networking guidelines, specifically for the WAN, that must be adhered to for the solution to work. For more information about the planning and preparation for a deployment of a 2-node VxRail cluster, see the [Dell EMC vSAN 2-Node Cluster Planning and Preparation Guide](#).

Chapter 4 VxRail Hardware and Switch Selection Decision Points

This chapter presents the following topic:

VxRail nodes and switches32

VxRail nodes and switches

- Step 1.** Assess your requirements and perform a sizing exercise to determine the quantity and characteristics of the VxRail nodes you need to meet planned workload and targeted use cases.
- Step 2.** Determine the number of physical racks needed to support the quantity and footprint of VxRail nodes to meet workload requirements, including the top-of-rack switches. Verify that the data center has sufficient floor space, power, and cooling.
- Step 3.** Determine the network switch topology that aligns with your business and operational requirements. See the sample wiring diagrams in [Appendix F: Physical Network Switch Examples](#) for guidance on the options supported for VxRail cluster operations.
- Step 4.** Based on the sizing exercise, determine the number of Ethernet ports on each VxRail node you want to reserve for VxRail networking.
- Two ports might be sufficient in cases where the resource consumption on the cluster is low and will not exceed available bandwidth.
 - Workloads with a high resource requirement or with a high potential for growth will benefit from a 4-port deployment. Resource-intensive networks, such as the vSAN and vMotion networks, benefit from the 4-port option because two ports can be reserved just for those demanding networks.
 - The 4-port option is required to enable link aggregation of demanding networks for the purposes of load balancing. In this case, the two ports that are reserved exclusively for the resource-intensive networks (vSAN and possibly vMotion) are configured into a logical channel to enable load balancing.

The VxRail cluster must be at version 7.0.130 or later to support link aggregation.

- Step 5.** Determine the optimal VxRail adapter and Ethernet port types to meet planned workload and availability requirements.
- VxRail supports 1 GbE, 10 GbE, and 25 GbE connectivity options to build the initial cluster.
 - Starting with VxRail version 7.0.130, you have flexibility with the selection of Ethernet adapter types:
 - Reserve and use only ports on the NDC for VxRail cluster networking.
 - Reserve and use both NDC-based and PCIe-based ports for VxRail cluster networking.
 - Reserve and use only PCIe-based ports for VxRail cluster networking.
 - If your performance and availability requirements might change later, you can initially reserve and use just NDC ports to build the initial cluster, and then migrate certain VxRail networks to PCIe-based ports.

The VxRail cluster must be at version 7.0.010 or later to migrate VxRail networks to PCIe-based ports.

- Step 6.** Decide whether you want to attach the VxRail nodes to the switches with RJ45, SFP+ or SFP28 connections.

- VxRail nodes with RJ-45 ports require CAT5 or CAT6 cables. CAT6 cables are included with every VxRail.
- VxRail nodes with SFP+ ports require optics modules (transceivers) and optical cables, or Twinax Direct-Attach Copper (DAC) cables. These cables and optics are not included; you must supply your own. The NIC and switch connectors and cables must be on the same wavelength.
- VxRail nodes with SFP28 ports require high thermal optics for ports on the NDC. Optics that are rated for standard thermal specifications can be used on the expansion PCIe network ports supporting SFP28 connectivity.

Step 7. Determine the additional ports and port speed on the switches for the uplinks to your core network infrastructure and inter-switch links for dual switch topologies. Select a switch or switches that provide sufficient port capacity and characteristics.

Step 8. Reserve one additional port on the switch for a workstation or laptop to access the VxRail management interface for the cluster.

- The additional port for access to the management interface is removed if connectivity is available elsewhere on the logical path from a jump host on the VxRail external management VLAN. Decide whether to deploy a separate switch to support connectivity to the VxRail management port on each node.
- Dell iDRAC supports 1 GbE connectivity. Dell Technologies recommends deploying a dedicated 1 GbE switch for this purpose. In certain cases, you can also use open ports on the top-of-rack switches.

Chapter 5 Planning the VxRail Implementation

This chapter presents the following topics:

- Introduction.....35**
- Plan data center routing services35**
- Plan for multi-rack VxRail cluster37**
- Plan for a vSAN HCI mesh topology38**
- Plan for VxRail custom uplink assignments39**
- Plan for link aggregation of VxRail networks.....41**
- Plan link aggregation on switch port pairs44**

Introduction

VxRail is an entire software-defined data center in an appliance form factor. All administrative activities, including initial implementation and initialization, configuration, capacity expansion, online upgrades, as well as maintenance and support are handled within the VxRail management system. When the VxRail appliance is installed in your data center, which is connected to your network, and the physical components that are powered on, the VxRail management system automates the full implementation of the final software-defined data center based on your settings and input.

Before getting to this phase, several planning and preparation steps must be undertaken with the data center network to ensure a seamless integration of the final product into your data center environment. The decisions made in addressing these topics and performing these tasks will drive the capability and functionality of the VxRail cluster.

These planning and preparation steps include:

1. Plan data center routing services.
2. Plan for multirack VxRail cluster.
3. Plan for a vSAN HCI mesh topology.
4. Plan for VxRail custom uplink assignments.
5. Plan for link aggregation of VxRail networks.

Note: Dell Technologies advises that proper attention is given to the data center network planning and preparation topics to ensure that the VxRail cluster when deployed meets your business and operational requirements.

Plan data center routing services

Specific VxRail networks, including the VxRail external management network and any external-facing end-user networks that are configured for VxRail, must have routing services that are enabled to support connectivity to external services and applications, as well as end-users.

A leaf-spine network topology is the most common use case for VxRail clusters. A single VxRail cluster can start on a single pair of switches in a single rack. When workload requirements expand beyond a single rack, expansion racks can be deployed to support the additional VxRail nodes and switches. The switches at the top of those racks, which are positioned as a 'leaf' layer, can be connected together using switches at the adjacent upper layer, or 'spine' layer.

If you choose to use a spine-leaf network topology to support the VxRail cluster or clusters in your data center, enabling Layer 3 routing services at either the spine layer or the leaf layer can both be considered.




Figure 17. Layer 2/3 boundary at the leaf layer or spine layer

Establishing routing services at the spine layer means that the uplinks on the leaf layer are trunked ports, and pass through all the required VLANs to the switches at the spine layer. This topology has the advantage of enabling the Layer 2 networks to span across all the switches at the leaf layer. This topology can simplify VxRail clusters that extend beyond one rack, because the Layer 2 networks at the leaf layer do not need Layer 3 services to span across multiple racks. A major drawback to this topology is scalability. Ethernet standards enforce a limitation of addressable VLANs to 4094, which can be a constraint if the application workload requires a high number of reserved VLANs, or if multiple VxRail clusters are planned.

Enabling routing services at the leaf layer overcomes this VLAN limitation. This option also helps optimize network routing traffic, as it reduces the number of hops to reach routing services. However, this option does require Layer 3 services to be licensed and configured at the leaf layer. In addition, since Layer 2 VxRail networks now terminate at the leaf layer, they cannot span across leaf switches in multiple racks.

Note: If your network supports VTEP, which enables extending Layer 2 networks between switches in physical racks over a Layer 3 overlay network, that can be considered to support a multi-rack VxRail cluster.




Figure 18. VTEP tunneling between leaf switches across racks

Plan for multi-rack VxRail cluster

A VxRail cluster can be extended beyond a single physical rack, and can extend to as many as six racks. All the network addresses applied to the VxRail nodes within a single rack must be within the same subnet.

You have two options if the VxRail cluster extends beyond a single rack:

- Use the same assigned subnet ranges for all VxRail nodes in the expansion rack. This option is required if SmartFabric Services are enabled on supporting switch infrastructure.
- Assign a new subnet range with a new gateway to the VxRail nodes in the expansion racks. (Your VxRail cluster must be running version 4.7.300 or later to use this option).

If the same subnets are extended to the expansion racks, the VLANs representing those VxRail networks must be configured on the top-of-rack switches in each expansion rack and physical connectivity must be established. If new subnets are used for the VxRail nodes and management components in the expansion racks, the VLANs will terminate at the router layer and routing services must be configured to enable connectivity between the racks.

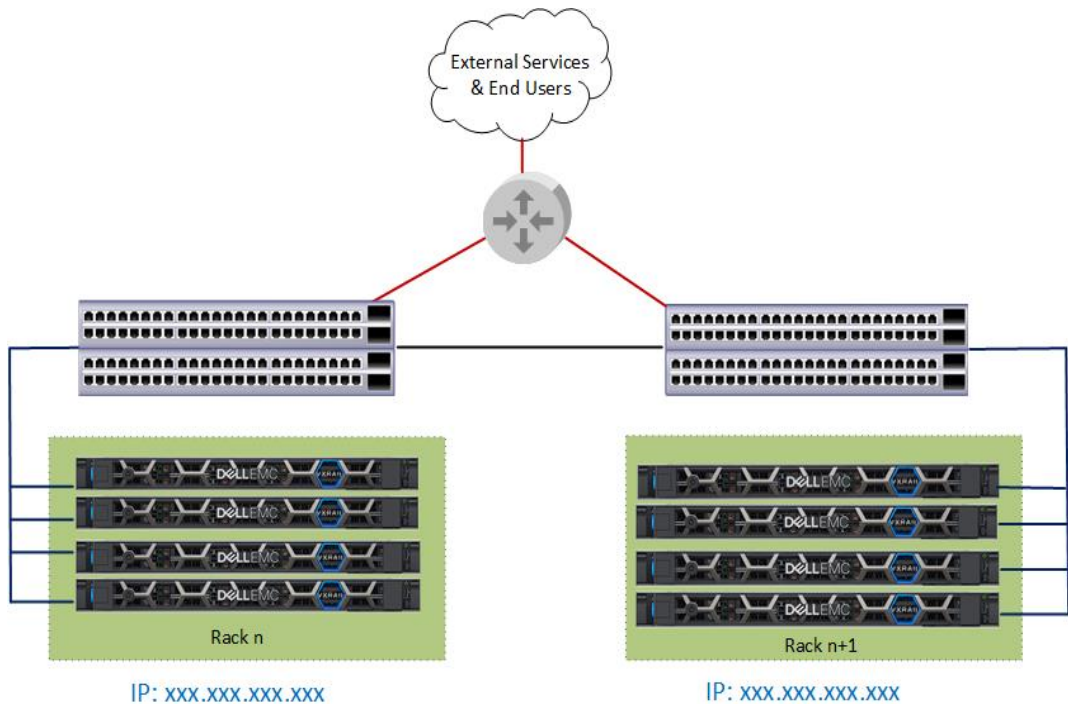


Figure 19. MultiRack VxRail sharing the same subnet

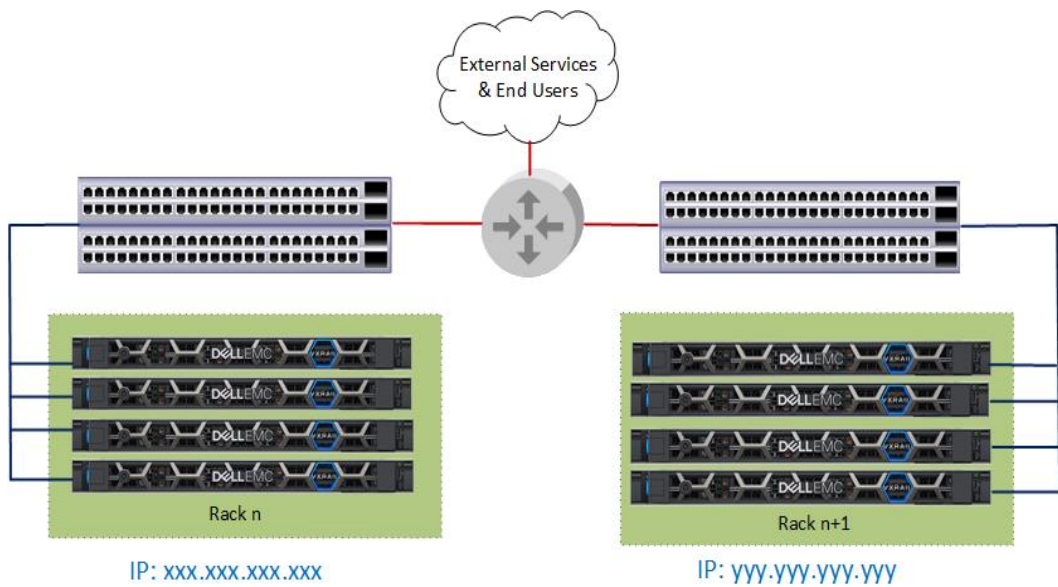


Figure 20. MultiRack VxRail with different subnets

Plan for a vSAN HCI mesh topology

Starting with VxRail version 7.0.100, the storage resources on the vSAN datastores on each VxRail cluster can be shared with other VxRail clusters. This storage sharing model is applicable only in a multi-cluster environment where the VxRail clusters are configured under a common data center on a customer-supplied vCenter instance.




Figure 21. Storage resource sharing with vSAN HCI mesh

If your plans include storage resource sharing using vSAN HCI mesh, be sure to prepare your data center to meet the following prerequisites:

- A vCenter instance at a version that supports VxRail version 7.0.100 or later.
- A vSAN Enterprise license for each VxRail cluster that will participate in a vSAN HCI mesh topology.

- A network topology that can connect the vSAN networks of the two VxRail clusters.
 - A common VLAN can be assigned to the vSAN network for each cluster to connect over a Layer 2 network. If the VxRail clusters are deployed against different top-of-rack switches, then the VLAN needs to be configured to stretch between the switch instances.
 - If the VxRail clusters are deployed against different top-of-rack switches, and the common VLAN cannot be stretched between the switch instances, then connectivity can be enabled using Layer 3 routing services. If this option is selected, be sure to assign routable IP addresses to the vSAN network on each participating VxRail cluster.

Plan for VxRail custom uplink assignments

Starting with VxRail version 7.0.130, you have flexibility in how you want to assign uplinks to the VxRail networks. You can deploy VxRail using the predefined uplink assignment templates, or you can select which uplinks on each node you want to assign to a given VxRail network. A VxRail cluster where the networking profile is pre-defined follows prescriptive rules for node port selection, and the assignment of the node ports to VxRail networks. With a custom profile, you can direct VxRail to follow a rule set you define for port selection, and determine which node ports are selected to support VxRail networking, and which uplinks are assigned to a specific VxRail network.

- If you opt to deploy VxRail with a pre-defined network profile, each VxRail node port used to support VxRail networking must be running at the same speed.
- If you choose to create a custom profile option, the following general rules are applicable:
 - The VxRail nodes can be configured with Ethernet ports running at different speeds. For instance, you can have 10 GbE ports on the NDC, and 25 GbE ports on a PCIe adapter card.
 - The Ethernet ports you select to support a VxRail network must be configured at the same speed. For instance, you can assign 10 GbE ports to the VxRail management networks, and 25 GbE ports to VxRail non-management networks such as vSAN and vMotion.
 - The Ethernet ports you select to support a VxRail network must be of the same type. For instance, you cannot assign an RJ45 port and an SFP+ port to support the same VxRail network.

If the VxRail cluster is deployed using one of the fixed network profiles, then the uplink assignments to each VxRail network are predefined based on whether two ports or four ports are selected to support the VxRail cluster. The fixed network profiles only select NDC-based ports for VxRail networking purposes.

- In a 2-port configuration, the VxRail networks share the two uplinks.
- In a 4-port configuration, the management networks are assigned two uplinks and the vMotion and vSAN networks are assigned the other two uplinks.

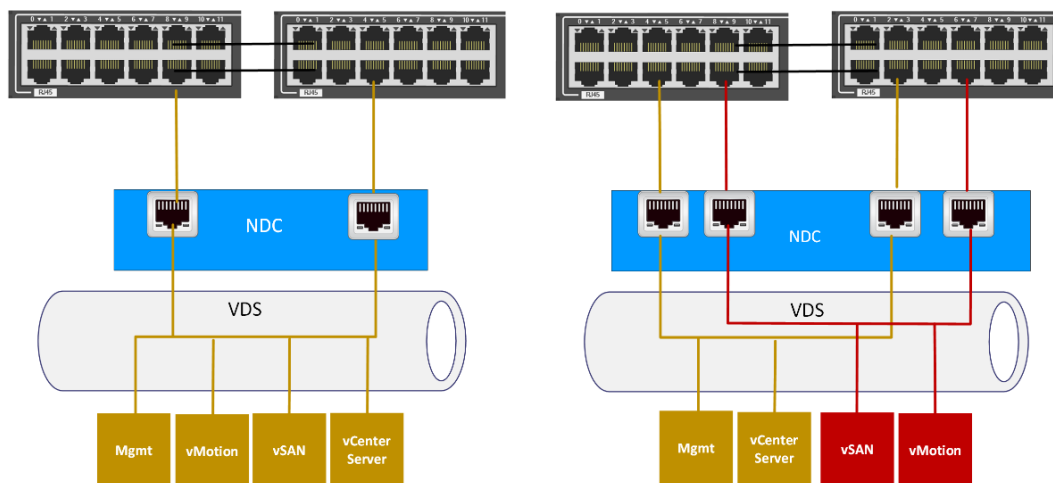


Figure 22. Default network profiles for a 2-port and a 4-port VxRail network

If you plan to use both NDC-based and PCIe-based ports to enable NIC redundancy and eliminate the NDC as a single point of failure, you can customize which ports on the VxRail nodes you want to use for each VxRail network. For example, you can select one port from the NDC and one port from a PCIe adapter card running at the same speed, and assign both of those to support the VxRail management networks. You can then select another port on the NDC, and another compatible port on the PCIe adapter card, and assign those to the non-management VxRail networks.

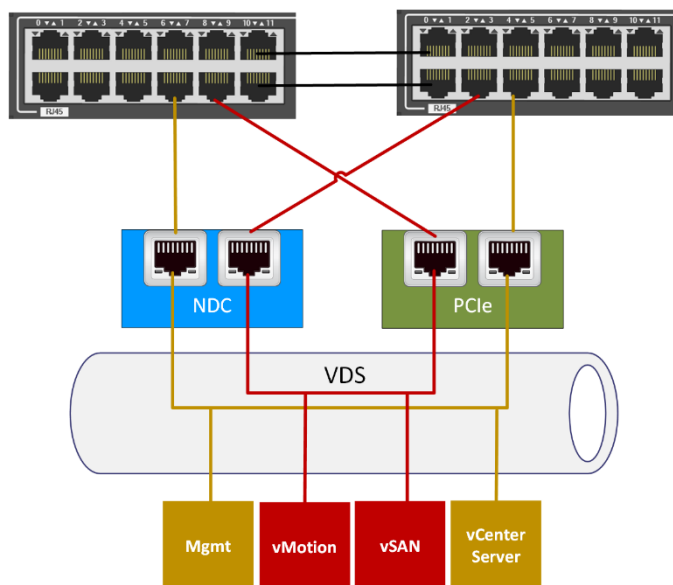


Figure 23. Custom uplink assignment across NDC-based and PCIe-based ports

If you expect the applications to be running on the VxRail cluster to be I-O intensive and require high bandwidth, you can choose to place the vMotion network on the same pair of ports as reserved for the VxRail management networks, and isolate the vSAN network on a pair of Ethernet ports.




Figure 24. Custom uplink assignment with vSAN network isolated on two Ethernet ports

The decision that you make on customizing the uplink assignments can have an impact on the top-of-rack switch configuration. For instance, with the predefined network profile using two ports, you need to configure the VLAN for each VxRail network on each of the switch ports connected to a VxRail node. If instead you choose to isolate the vSAN network, then you only need to configure the VLAN for the vSAN network on those specific switch ports.

Plan for link aggregation of VxRail networks

Link aggregation for specific VxRail networks is supported starting with VxRail version 7.0.130. NIC teaming in VxRail is the foundation for supporting link aggregation, which is the bundling of two physical network links to form a single logical channel. Link aggregation allows ports on a VxRail node to peer with a matching pair of ports on the top-of-rack switches to support load balancing and optimize network traffic distribution across the physical ports. VxRail networks with heavy resource requirements, such as vSAN and potentially vMotion, benefit most from network traffic optimization.

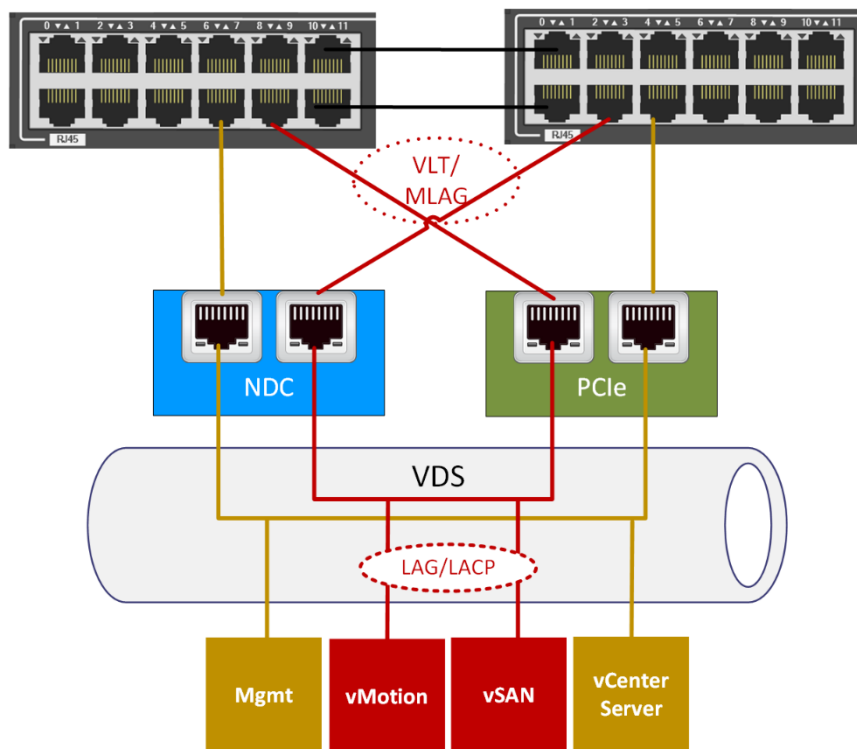


Figure 25. Overview of LAG peering relationship between VxRail and the adjacent switches

Each VxRail network is assigned two uplinks by default during the initial implementation operation. One uplink is designated as ‘active’, and the other uplink as ‘standby’. Under this model, resource-intensive networks are limited to the bandwidth of a single uplink. Enabling link aggregation allows the VxRail network to use the bandwidth of both uplinks, with the traffic flow coordinated based on the load-balancing hash algorithm.

The following guidelines must be followed when using NIC teaming and link aggregation with VxRail:

- The VxRail cluster must be running version 7.0.130 to enable NIC teaming and support link aggregation.
- Link aggregation is not supported on the VxRail management networks. These networks have low-bandwidth requirements, and load balancing would offer minimal benefit.
- The virtual distributed switch that is integrated into the VxRail cluster supports LACP.
 - If the switches that support the VxRail cluster support LACP, dynamic link aggregation can be configured on the VxRail networks.
 - If LACP is not supported on the switches, static link aggregation can be configured on the VxRail networks.
 - LACP is more favorable because it offers support for more load balancing hashing algorithms and better management capabilities.
- To enable NIC teaming and link aggregation during the VxRail initial build process:

- You must supply a compatible vCenter instance to serve as the target for the VxRail cluster.
- You must supply a compatible and preconfigured virtual distributed switch to support the VxRail networking requirements.
- Link aggregation and load balancing are preconfigured on the virtual distributed switch.
- NIC teaming and link aggregation can be configured on the VxRail-supplied virtual distributed switch after the VxRail cluster build is complete.
 - If the target version for your VxRail cluster being delivered is earlier than 7.0.130, the cluster can be upgraded after the build operation is complete.
 - Refer to [Enabling link aggregation for load balancing](#) for details on configuring link aggregation after VxRail initial implementation
- Four Ethernet ports per node must be reserved for VxRail networking.
 - Can be a mixture of NDC and PCIe Ethernet ports
 - All NDC-based or all PCIe-based Ethernet ports are also supported.
 - Two ports are reserved for the VxRail management networks. NIC teaming is not supported on this port pair.
 - Two ports are reserved for the VxRail non-management networks (vSAN/vMotion). NIC teaming is supported on this port pair.
 - All ports configured for link aggregation must be running at the same speed.

Verify the switches support link aggregation

Support for LACP, the selection of load-balancing hashing algorithms and the formation of link aggregation on the physical switches depends on the switch vendor and operating system. These features are usually branded by the vendor, using names such as 'Ether-Channel', 'Ethernet trunk', or 'Multi-Link Trunking'. Consult your switch vendor to verify that the switch models planned for the VxRail cluster supports this feature.

Verify support for multi-chassis link aggregation

If you plan to deploy a pair of switches to support the VxRail cluster, and you want to enable load balancing across both switches, you must configure multi-chassis link aggregation. This feature is usually branded by the switch vendor, such as Cisco's Virtual Port Channel or Dell's VLT Port Channel. See the guides provided by your switch vendor for the steps to complete this task.

Identify switch ports to be configured for link aggregation

Enabling load balancing for the non-management VxRail networks requires peering the pair of ports on each VxRail node with a pair of ports on the top-of-rack switches.




Figure 26. Plugging into equivalent switch ports for link aggregation

If you are enabling link aggregation across a pair of switches, and you have matching open ports on both switches, the best practice is to plug the cables into equivalent port numbers on both switches. We recommend creating a table to map each VxRail node port to a corresponding switch port. Then, identify which ports on each VxRail will be enabled for link aggregation.

For example, if you want to deploy a VxRail cluster with four nodes, and reserve and use two ports on the NDC and two ports on a PCIe adapter card for the VxRail cluster, and use the first eight ports on a pair of top-of-rack switches for connecting the cables, you could use the resulting table to identify the switch ports to be configured for link aggregation.

Switch A	Node 1 NDC 1	Node 1 PCIe 2	Node 2 NDC 1	Node 2 PCIe 2	Node 3 NDC 1	Node 3 PCIe 2	Node 4 NDC 1	Node 4 PCIe 2
Port	1	2	3	4	5	6	7	8
Switch B	Node 1 PCIe 1	Node 1 NDC 2	Node 2 PCIe 1	Node 2 NDC 2	Node 3 PCIe 1	Node 3 NDC 2	Node 4 PCIe 1	Node 4 NDC 2

Figure 27. Sample port-mapping table

Assuming that you are using the second port on the NDC and PCIe adapter card for the non-management VxRail networks, you can identify the switch port pairs, as shown in the columns shaded green, to be configured for link aggregation.

Dell Technologies recommends creating a table mapping the VxRail ports to the switch ports and part of the planning and design phase.

Plan link aggregation on switch port pairs

For each pair of ports on each node that is supporting a VxRail network that will be enabled for link aggregation, the corresponding pair of switch ports they are connected to must also be configured for link aggregation. The commands to perform this action depend on the switch model and operating system. See the guides provided by your switch vendor for the steps to complete this task.

If you are deploying the VxRail cluster against a customer-supplied virtual distributed switch, this task must be completed before the VxRail initial build operation. If you are deploying the VxRail cluster using the virtual distributed switch included with the product, this activity is performed after the initial cluster build operation is completed.

Dell-EMC will publish a procedure for the scenarios where link aggregation on the VxRail networks is performed as a Day 2 activity.

Chapter 6 Planning the VxRail Cluster Implementation

This chapter presents the following topics:

Introduction	46
Decide on VxRail single point of management	46
Decide on VxRail network traffic segmentation	47
Plan the VxRail logical network	49
Plan network exclusions reserved for VxRail Manager	54
Plan network settings for VxRail management components	55
Identify IP addresses for VxRail management components	57
Select hostnames for VxRail management components	58
Identify external applications and settings for VxRail	60
Prepare customer-supplied vCenter server	61
Prepare customer-supplied virtual distributed switch	63
Reserve IP addresses for VxRail vMotion network	67
Reserve IP addresses for VxRail vSAN network	68
Decide on VxRail logging solution	68
Assign passwords for VxRail management	69
Prepare for Dell EMC SmartFabric Services enablement	70

Introduction

VxRail is an entire software-defined data center in an appliance form factor. All administrative activities, including initial implementation and initialization, configuration, capacity expansion, online upgrades, as well as maintenance and support are handled within the VxRail management system. When the VxRail appliance is installed in your data center and connected to your network, and the physical components are powered on, the VxRail management system automates the full implementation of the final software-defined data center based on your settings and input.

Before getting to this phase, several planning and preparation steps must be undertaken to ensure a seamless integration of the final product into your data center environment. These planning and preparation steps include:

1. Decide on VxRail Single Point of Management.
2. Decide on VxRail network traffic segmentation.
3. Plan the VxRail logical network.
4. Identify IP address range for VxRail logical networks.
5. Identify unique hostnames for VxRail management components.
6. Identify external applications and settings for VxRail.
7. Create DNS records for VxRail management components.
8. Prepare customer-supplied vCenter Server.
9. Reserve IP addresses for VxRail vMotion and vSAN networks.
10. Decide on VxRail Logging Solution.
11. Decide on passwords for VxRail management.

Use the [Appendix C: VxRail Setup Checklist](#) and the [Appendix A: VxRail Network Configuration Table](#) to help create your network plan. References to rows in this document are to rows in the VxRail Network Configuration Table.

Note: Once you set up the VxRail cluster and complete the initial initialization phase to produce the final product, the configuration cannot easily be changed. We strongly recommend that you take care during this planning and preparation phase to decide on the configurations that will work most effectively for your organization.

Decide on VxRail single point of management

The unified resources of a VxRail appliance create a virtual infrastructure that is defined and managed as a vSphere cluster under a single instance of vCenter. A decision must be made to use the VxRail-supplied vCenter Server, which is deployed in the cluster as part of the initial initialization process, or a customer-supplied vCenter server, which is external to the cluster. During the VxRail initialization process which creates the final product, you must select whether to deploy the embedded VxRail-supplied vCenter Server on the cluster or deploy the cluster on an external customer-supplied vCenter server. Once the initialization process is complete, migrating to a new vCenter single point of management requires professional services assistance, and is difficult to change.

Dell Technologies recommends that you consider all the ramifications during this planning and preparation phase, and decide on the single point of management option that will work most effectively for your organization.

The following should be considered for selecting the VxRail vCenter server:

- A vCenter Standard license is included with VxRail, and does not require a separate license. This license cannot be transferred to another vCenter instance.
- The VxRail vCenter Server can manage only a single VxRail instance. This means an environment of multiple VxRail clusters with the embedded vCenter instance will require an equivalent number of points of management for each cluster.
- VxRail Lifecycle Management supports the upgrade of the VxRail vCenter server. Upgrading a customer-supplied vCenter using VxRail Lifecycle Management is not supported.
- DNS services are required for VxRail. With the VxRail vCenter option, you have the choice of using the internal DNS supported within the VxRail cluster, or leveraging external DNS in your data center.

For a customer-supplied vCenter, the following items should be considered:

- The vCenter Standard license included with VxRail cannot be transferred to a vCenter instance outside of the VxRail cluster.
- Multiple VxRail clusters can be configured on a single customer-supplied vCenter server, limiting the points of management.
- With the customer-supplied vCenter, external DNS must be configured to support the VxRail cluster.
- Ensuring version compatibility of the customer-supplied vCenter with VxRail is the responsibility of the customer.
- With the customer-supplied vCenter, you have the option of configuring the virtual distributed switch or switches yourself to support the VxRail cluster, or have VxRail deploy a virtual distributed switch and perform the configuration instead. This option is advantageous if you want better control and manageability of the virtual networking in your data center, and consolidate the number of virtual distributed switches in your vCenter instance.

Note: The option to use the internal DNS or to deploy the VxRail cluster against a preconfigured virtual distributed switch requires VxRail version of 7.0.010 or later.

For more details on the planning steps for a customer-supplied vCenter, see the [Dell EMC VxRail vCenter Server Planning Guide](#).

Decide on VxRail network traffic segmentation

You have options regarding segmenting the VxRail network traffic at the virtual distributed switch level. Prior to 7.0.130, all the required VxRail networks were confined to a single virtual distributed switch. Starting with version 7.0.130, you can decide whether you want to deploy a second virtual distributed switch to isolate the VxRail management network traffic and the VxRail non-management network traffic.

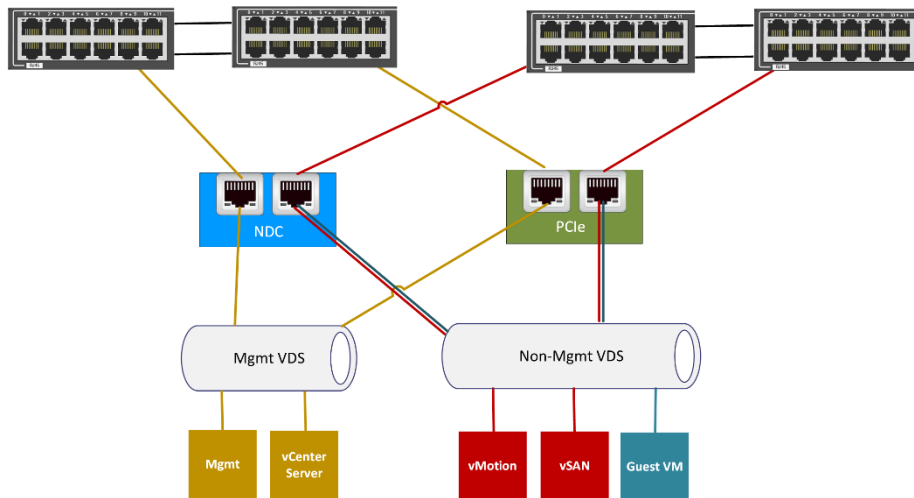


Figure 28. VxRail network segmentation with two virtual distributed switches

If your company or organization has stringent security policies regarding network separation, splitting the VxRail networks between two virtual distributed switches will enable better compliance with those policies, and simplify redirecting the VxRail management network traffic and non-management network traffic down separate physical network paths.

You can choose from the following options to align with your company or organization networking policies:

- Place all the required VxRail network traffic and guest network traffic on a single virtual distributed switch.
- Use two virtual distributed switches to segment the VxRail management network traffic from the VxRail non-management traffic and guest virtual machine network traffic.
- Deploy a separate virtual distributed switch to support guest virtual machine network traffic.




Figure 29. VxRail network segmentation with two virtual distributed switches

VxRail supports either a single virtual distributed switch or two virtual distributed switches as part of the initial implementation process. If your security posture changes after the VxRail cluster initial implementation has completed, a second virtual distributed switch can still be deployed and the VxRail network traffic can be redirected to that second virtual distributed switch. Any additional virtual distributed switches beyond two switches, such as those for user requirements outside of VxRail networking, can be deployed after initial implementation.

Plan the VxRail logical network

The physical connections between the ports on your network switches and the NICs on the VxRail nodes enable communications for the virtual infrastructure within the VxRail cluster. The virtual infrastructure within the VxRail cluster uses the virtual distributed switch to enable communication within the cluster, and out to IT management and the application user community.

VxRail has predefined logical networks to manage and control traffic within the cluster and outside of the cluster. Certain VxRail logical networks must be made accessible to the outside community. For instance, connectivity to the VxRail management system is required by IT management. VxRail networks must be configured for end-users and application owners who need to access their applications and virtual machines running in the VxRail cluster. In addition, a network supporting I/O to the vSAN datastore is required, and a network to support vMotion, which is used to dynamically migrate virtual machines between VxRail nodes to balance workload, must also be configured. Finally, an internal management network is required by VxRail for device discovery.

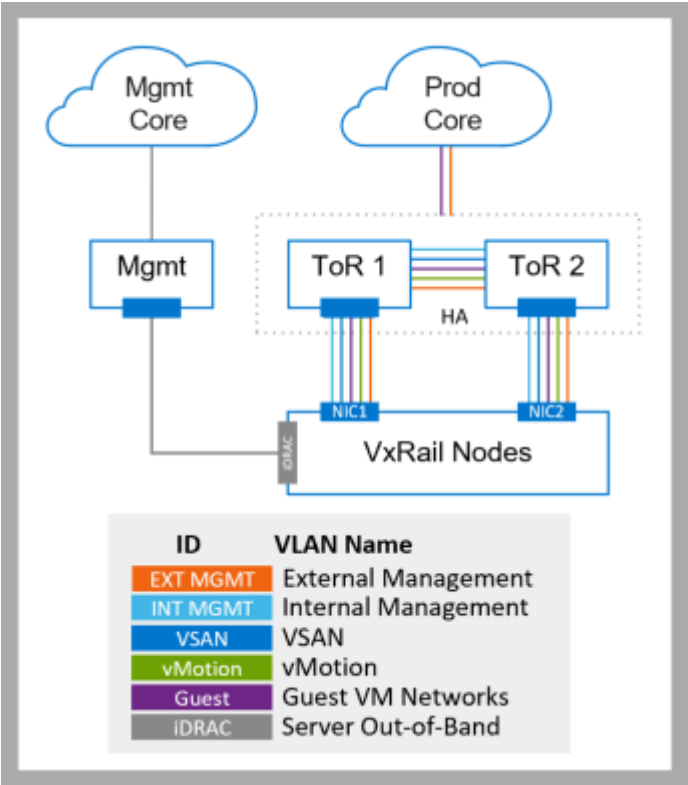


Figure 30. VxRail Logical Network Topology

All the Dell PowerEdge servers that serve as the foundation for VxRail nodes include a separate Ethernet port that enables connectivity to the platform to perform hardware-based maintenance and troubleshooting tasks. A separate network to support management access to the Dell PowerEdge servers is recommended, but not required.

IP address considerations for VxRail networks

IP addresses must be assigned to the VxRail external management network, vSAN network, vMotion network, and any guest networks you want to configure on the VxRail cluster. Decisions need to be made on the IP address ranges reserved for each VxRail network:

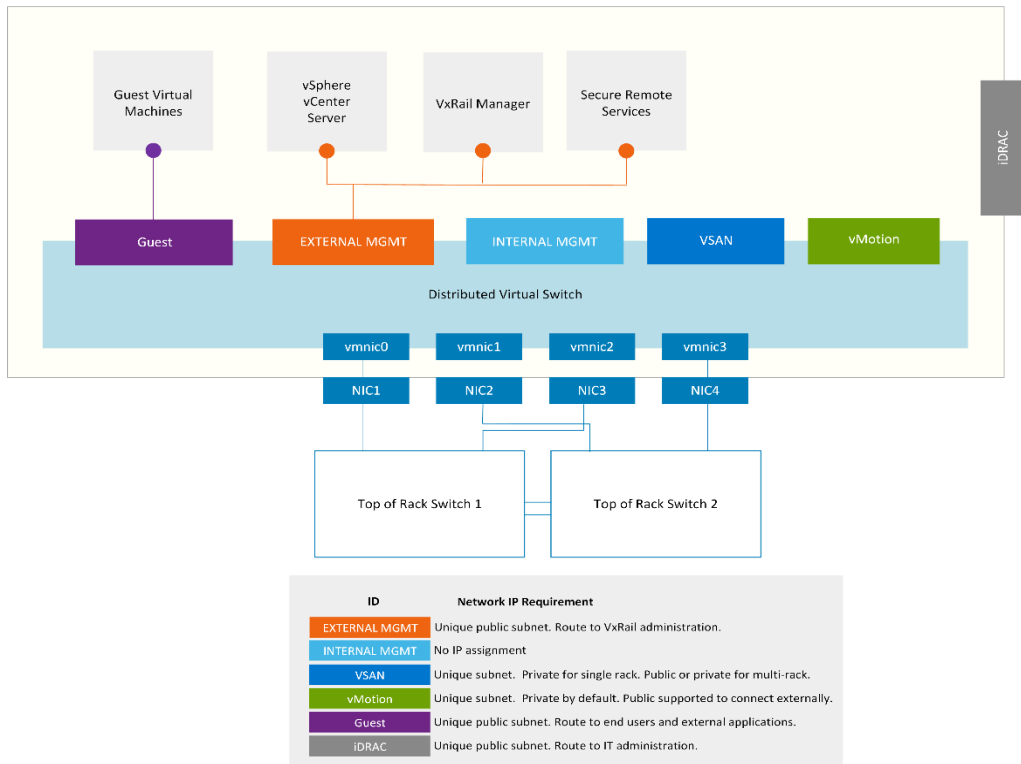


Figure 31. VxRail Network IP Requirements

- The internal management network that is used for device discovery does not require assigned IP addresses.
- Since the external management network must be able to route upstream to network services and end users, a non-private, routable IP address range must be assigned to this network.
- Traffic on the vSAN network is passed only between the VxRail nodes that form the cluster. Either a routable or non-routable IP address range can be assigned. If your plans include a multirack cluster, and you want to use a new IP subnet range in the expansion racks, then assign a routable IP address range to this network.
- If your requirements for virtual machine mobility are within the VxRail cluster, a non-routable IP address range can be assigned to the vMotion network. However, if you need to enable virtual machine mobility outside of the VxRail cluster, or have plans for a multirack expansion that will use a different subnet range on any expansion racks, reserve a routable IP address range.

Virtual LAN considerations for VxRail networks

Virtual LANs (VLANs) define the VxRail logical networks within the cluster, and the method that is used to control the paths that a logical network can pass through. A VLAN, represented as a numeric ID, is assigned to a VxRail logical network. The same VLAN ID is also configured on the individual ports on your top-of-rack switches, and on the virtual ports in the virtual-distributed switch during the automated implementation process. When an application or service in the VxRail cluster sends a network packet on the virtual-distributed switch, the VLAN ID for the logical network is attached to the packet. The packet will only be able to pass through the ports on the top-of-rack switch and the virtual distributed switch where there is a match in VLAN IDs. Isolating the VxRail logical network

traffic using separate VLANs is highly recommended, but not required. A 'flat' network is recommended only for test, non-production purposes.

As a first step, the network team and virtualization team should meet in advance to plan VxRail's network architecture.

- The virtualization team must meet with the application owners to determine which specific applications and services that are planned for VxRail are to be made accessible to specific end-users. This will determine the number of logical networks that are required to support traffic from non-management virtual machines.
- The network team must define the pool of VLAN IDs needed to support the VxRail logical networks, and determine which VLANs will restrict traffic to the cluster, and which VLANs will be allowed to pass through the switch up to the core network.
- The network team must also plan to configure the VLANs on the upstream network, and on the switches attached to the VxRail nodes.
- The network team must also configure routing services to ensure connectivity for external users and applications on VxRail network VLANs passed upstream.
- The virtualization team must assign the VLAN IDs to the individual VxRail logical networks.

VxRail groups the logical networks in the following categories: **External Management**, **Internal Management**, **vSAN**, **vSphere vMotion**, and **Virtual Machine**. VxRail assigns the settings that you specify for each of these logical networks during the initialization process.

Before VxRail version 4.7, both external and internal management traffic shared the external management network. Starting with VxRail version 4.7, the external and internal management networks are broken out into separate networks.

External Management traffic includes all VxRail Manager, vCenter Server, ESXi communications, and in certain cases, vRealize Log Insight. All VxRail external management traffic is untagged by default and should be able to go over the Native VLAN on your top-of-rack switches.

A tagged VLAN can be configured instead to support the VxRail external management network. This option is considered a best practice, and is especially applicable in environments where multiple VxRail clusters will be deployed on a single set of top-of-rack switches. To support using a tagged VLAN for the VxRail external management network, configure the VLAN on the top-of-rack switches, and then configure trunking for every switch port that is connected to a VxRail node to tag the external management traffic.

The **Internal Management** network is used solely for device discovery by VxRail Manager during initial implementation and node expansion. This network traffic is non-routable and is isolated to the top-of-rack switches connected to the VxRail nodes. Powered-on VxRail nodes advertise themselves on the Internal Management network using multicast, and discovered by VxRail Manager. The default VLAN of 3939 is configured on each VxRail node that is shipped from the factory. This VLAN must be configured on the switches, and configured on the trunked switch ports that are connected to VxRail nodes.

If a different VLAN value is used for the Internal Management network, it not only must be configured on the switches, but must also be applied to each VxRail node on-site. Device discovery on this network by VxRail Manager will fail if these steps are not followed.

Device discovery requires multicast to be configured on this network. If there are restrictions within your data center regarding the support of multicast on your switches, then you can bypass configuring this network, and instead use a manual process to select and assign the nodes that form a VxRail cluster.

Using the manual node assignment method instead of node discovery for VxRail initial implementation requires version 7.0.130 or later.

It is a best practice to configure a VLAN for the **vSphere vMotion** and **vSAN** networks. For these networks, configure a VLAN for each network on the top-of-rack switches, and then include the VLANs on the trunked switch ports that are connected to VxRail nodes.

The **Virtual Machine** networks are for the virtual machines running your applications and services. These networks can be created by VxRail during the initial build process, or created afterward using the vClient after initial configuration is complete. Dedicated VLANs are preferred to divide **Virtual Machine** traffic, based on business and operational objectives. VxRail creates one or more VM Networks for you, based on the name and VLAN ID pairs that you specify. Then, when you create VMs in vSphere Web Client to run your applications and services, you can easily assign the virtual machine to the VM Networks of your choice. For example, you could have one VLAN for Development, one for Production, and one for Staging.

Network Configuration Table ✓ Row 1	Enter the external management VLAN ID for VxRail management network (VxRail Manager, ESXi, vCenter Server/PSC, Log Insight). If you do not plan to have a dedicated management VLAN and will accept this traffic as untagged, enter "0" or "Native VLAN."
Network Configuration Table ✓ Row 2	Enter the internal management VLAN ID for VxRail device discovery. The default is 3939. If you do not accept the default, the new VLAN must be applied to each VxRail node before cluster implementation to enable discovery.
Network Configuration Table ✓ Row 3	Enter a VLAN ID for vSphere vMotion. (Enter 0 in the VLAN ID field for untagged traffic)
Network Configuration Table ✓ Row 4	Enter a VLAN ID for vSAN. (Enter 0 in the VLAN ID field for untagged traffic)
Network Configuration Table ✓ Rows 5-6	Enter a Name and VLAN ID pair for each VM guest network you want to create. You must create at least one VM Network. (Enter 0 in the VLAN ID field for untagged traffic)

Note: If you plan to have multiple independent VxRail clusters, we recommend using different VLAN IDs across multiple VxRail clusters to reduce network traffic congestion.

For a 2-Node cluster, the VxRail nodes must connect to the Witness over a separate Witness traffic separation network. The Witness traffic separation network is not required for stretched-cluster but is considered a best practice. For this network, a VLAN is

required to enable Witness network on this VLAN must be able to pass through upstream to the Witness site.

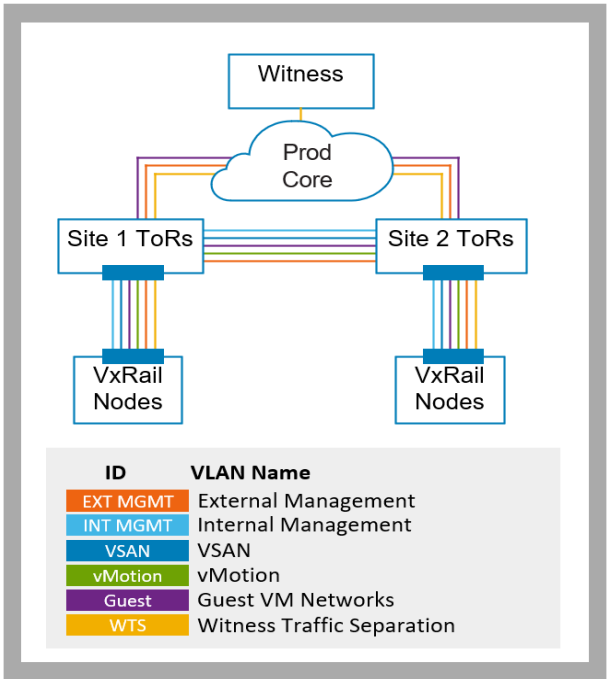


Figure 32. Logical network with Witness and Witness Traffic Separation

Network Configuration Table ✓ Row 78	Enter the Witness traffic separation VLAN ID.
---	---

Plan network exclusions reserved for VxRail Manager

VxRail Manager relies internally on a microservice model using a Docker container architecture. A set of IP addresses is reserved for use by VxRail Manager to support networking for the microservices. The IP addresses within these reserved pools are automatically assigned to the microservices initiated by VxRail Manager at the time of power-on, and assigned as needed as part of normal VxRail Manager operations. Using these reserved IP addresses for any VxRail network can potentially cause a conflict with VxRail Manager operations, and should be blocked from assignment to VxRail networks.

The reserved IP address ranges are:

- 172.28.0.0/16
- 172.29.0.0/16
- 10.0.0.0/24
- 10.0.1.0/24

Plan network settings for VxRail management components

During the initial build of the VxRail cluster, IP addresses that are entered are assigned to the VxRail components that are members of the External Management network and must follow certain rules:

- The IP address scheme must be a public IP address range.
- The IP address must be fixed (no DHCP).
- The IP addresses cannot be in use.
- The IP address range must all be in the same subnet.

You have flexibility in how the IP addresses are assigned to the VxRail management components.

- If the VxRail cluster to be deployed in at version 7.0.010 or later, you can either manually assign the IP addresses to the management components, or have the IP addresses auto-assigned during VxRail initial build.
- Before VxRail version 7.0.010, the only supported option was to auto-assign the IP addresses to the management components. The assignment process allocates IP addresses in sequential order, so a range must be provided for this method.

The decisions that you make on the final VxRail configuration that is planned for your data center impacts the number of IP addresses you will need to reserve.

- Decide if you want to reserve additional IP addresses in the VxRail management system to assign to VxRail nodes in the future for expansion purposes in a single rack. When a new node is added to an existing VxRail cluster, it will assign an IP address from the unused reserve pool, or prompt you to enter an IP address manually if none are in reserve and unused.
- Decide whether you will use the vCenter instance that is deployed in the VxRail cluster, or use an external vCenter already operational in your data center.
 - For VxRail versions 7.0 or later, if you use the vCenter instance that is deployed on the VxRail cluster, you must reserve an IP address for vCenter. The Platform Service Controller is bundled into the vCenter instance.
 - For VxRail versions earlier than version 7.0, if you have VxRail deploy vCenter, you must reserve an IP address for the vCenter instance and an IP address for the Platform Service Controller.
- Decide if you will use vSphere Log Insight that can be deployed in the VxRail cluster.
 - For VxRail version 7.0 and earlier, and you choose to use the vCenter instance that is deployed in the VxRail cluster, then you have the option to deploy vSphere Log Insight on the cluster. You can also choose to connect to an existing syslog server in your data center, or no logging at all. If you choose to deploy vSphere Log Insight in the VxRail cluster, you must reserve one IP address.
 - vRealize Log Insight is not an option for deployment during the initial VxRail configuration process starting in version 7.0.010.

- If you use an external vCenter already operational in your data center for VxRail, vSphere Log Insight cannot be deployed.
- VxRail supports the Dell EMC 'call home' feature, where alerts from the appliance are routed to customer service. The Secure Remote Services gateway is required to enable alerts from VxRail to be sent to Dell Technologies customer service.
 - Decide whether to use an existing Secure Remote Services gateway in your data center for 'call-home', deploy a virtual instance of the Secure Remote Services gateway in the VxRail cluster for this purpose, or none.
 - Reserve one IP address to deploy SRS-VE (Secure Remote Services Virtual Edition) in the VxRail cluster.
- If you are planning to deploy a VxRail cluster that requires a Witness at a remote third site, such as VxRail stretched-cluster or 2-Node cluster, two IP addresses are required to deploy the witness virtual appliance.
 - One IP address is assigned to the witness management network.
 - One IP address is assigned to the witness vSAN network.
 - Both networks must be able to route to the VxRail cluster requiring the remote site witness.

An existing vSAN witness can be shared in your remote site if the VxRail clusters are stretched clusters, and the vSAN witness can support vSAN datastores at version 7 Update 1 or later.

- For a 2-Node Cluster, the VxRail nodes must connect to the Witness over a separate Witness traffic separation network. For this network, an additional IP address is required for each of the two VxRail nodes.
 - The VxRail nodes must be able to route to the remote site Witness.
 - The traffic must be able to pass through the Witness traffic separation VLAN.

Use the following table to determine the number of public IP addresses required for the External Management logical network:

Component	Condition
VxRail Node	One per VxRail Node
VxRail Manager	One
vCenter	If you are supplying vCenter Server for VxRail: 0 If you are using vCenter on VxRail: 2
Log Insight	If you are supplying vCenter Server for VxRail: 0 If you are supplying a syslog server for VxRail: 0 If you will not enable logging for VxRail: 0 If you are using Log Insight on VxRail: 1
SRS-VE	If you are planning to deploy SRS gateway on VxRail: 1 If you will not deploy SRS gateway on VxRail: 0

Request your networking team to reserve a subnet range that has sufficient open IP addresses to cover VxRail initial build and any planned future expansion.

Network Configuration Table ✓ Row 7	Enter the subnet mask for the VxRail External Management network.
Network Configuration Table ✓ Row 8	Enter the gateway for the VxRail External Management network.

Identify IP addresses for VxRail management components

If you are choosing to auto-assign the IP addresses for the ESXi hosts that serve as the foundation for VxRail nodes, request your networking team to reserve a large enough pool of unused IP addresses.

Record the IP address range for the ESXi hosts.

Network Configuration Table ✓ Rows 24 and 25	Enter the starting and ending IP addresses for the ESXi hosts - a continuous IP range is required.
---	--

If you choose instead to assign the IP addresses to each individual ESXi host, record the IP address for each ESXi host to be included for VxRail initial build.

Network Configuration Table ✓ Rows 26 and 29	Enter the IP addresses for the ESXi hosts.
---	--

Record the permanent IP address for VxRail Manager. This is required.

Network Configuration Table ✓ Row 14	Enter the permanent IP address for VxRail Manager.
---	--

If you are going to deploy the embedded vCenter on the VxRail cluster provided with VxRail, record the permanent IP address for vCenter and Platform Service Controller (if applicable). Leave these entries blank if you will provide an external vCenter for VxRail.

Network Configuration Table ✓ Row 31	Enter the IP address for VxRail vCenter.
Network Configuration Table ✓ Row 33	Enter the IP address for VxRail Platform Service Controller (if applicable)

Record the IP address for Log Insight. Leave this entry blank if you will be deploying a version of VxRail at version 7.0.010 or later, or if you choose not to deploy Log Insight on VxRail.

Network Configuration Table ✓ Row 69	Enter the IP address for vSphere Log Insight.
---	---

Record the two IP addresses for the witness virtual appliance. Leave blank if a witness is not required for your VxRail deployment.

Network Configuration Table ✓ Row 76	Enter IP address for Witness Management Network.
Network Configuration Table ✓ Row 77	Enter IP address for Witness vSAN Network.

Record the IP addresses for each node required for Witness traffic for a 2-Node cluster deployment. Leave blank if you are not deploying a 2-Node cluster.

Network Configuration Table ✓ Row 79	Enter the IP address for the first of the two nodes in the 2-Node cluster.
Network Configuration Table ✓ Row 80	Enter the IP address for the second of the two nodes in the 2-Node Cluster.

Select hostnames for VxRail management components

Each of the VxRail management components you deploy in the VxRail cluster requires you to assign an IP address, and assign a fully qualified hostname. During initialization, each of these VxRail management components are assigned a hostname and IP address.

Determine the naming format for the hostnames to be applied to the required VxRail management components: each ESXi host, and VxRail Manager. If you deploy the vCenter Server in the VxRail cluster, that also requires a hostname. In addition, if you decide to deploy Log Insight in the VxRail cluster, that needs a hostname as well.

Note: You cannot easily change the hostnames and IP addresses of the VxRail management components after initial implementation.

Select top-level domain

Begin the process by selecting the domain to use for VxRail and assign to the fully qualified hostnames. Be aware that DNS is a requirement for VxRail, so select a domain where the naming services can support that domain.

Network Configuration Table ✓ Row 12	Enter the top-level domain.
---	-----------------------------

Select VxRail Manager hostname

A hostname must be assigned to VxRail Manager. The domain is also automatically applied to the chosen hostname. Dell Technologies recommends following the naming format that is selected for the ESXi hosts to simplify cluster management.

Network Configuration Table ✓ Row 13	Enter the hostname for VxRail Manager.
---	--

Select ESXi hostnames

All VxRail nodes in a cluster require hostnames. Starting with VxRail version 7.0.010, you have the choice of using any host naming convention you want, provided it is a legitimate format, or having VxRail auto-assign the hostnames to the ESXi nodes following VxRail rules automatically during the VxRail initial build process.

If you plan to have VxRail auto-assign the hostnames during the cluster initial build process, make sure to follow the rules stated in this section. All ESXi 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. The Preview field that is shown during VxRail initialization is an example of the hostname of the first ESXi host. For example, if the prefix is “host,” the separator is “None,” the iterator is “Num 0X,” the offset is empty, and the suffix is “lab,” and the domain is “local,” the first ESXi hostname would be “host01lab.local.” The domain is also automatically applied to the VxRail management components. (Example: my-vcenter.local).

	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

Enter the values for building and auto-assigning the ESXi hostnames if this is the chosen method.

Network Configuration Table ✓ Rows 15–19	Enter an example of your desired ESXi host-naming scheme. Be sure to show your desired prefix, separator, iterator, offset, suffix, and domain.
---	---

If the ESXi hostnames will be applied manually, capture the name for each ESXi host planned for the VxRail initial build operation.

Network Configuration Table ✓ Rows 20–23	Enter the reserved hostname for each ESXi host.
---	---

Select VxRail vCenter Server hostname

Note: You can skip this section if you plan to use an external vCenter Server in your data center for VxRail. These action items are only applicable if you plan to use the VxRail vCenter Server.

If you want to deploy a new vCenter Server on the VxRail cluster, you must specify a hostname for the VxRail vCenter Server and, if required, for the Platform Services Controller (PSC). The domain is also automatically applied to the chosen hostname. Dell Technologies recommends following the naming format that is selected for the ESXi hosts to simplify cluster management.

Network Configuration Table ✓ Row 30	Enter an alphanumeric string for the new vCenter Server hostname. The domain that is specified will be appended.
---	--

Network Configuration Table ✓ Row 32	Enter an alphanumeric string for the new Platform Services Controller hostname. The domain that is specified will be appended.
---	--

Select Log
Insight
hostname

Note: You can skip this section if you plan to deploy a VxRail cluster at version 7.0.010 or later, will use an external syslog server instead of Log Insight, or will not enable logging.

To deploy Log Insight to the VxRail cluster, the management component must be assigned a hostname. You can use your own third-party syslog server, use the vRealize Log Insight solution included with VxRail, or no logging. You can only select the vRealize Log Insight option if you also use the VxRail vCenter Server.

Network Configuration Table ✓ Row 68	Enter the hostname for Log Insight.
---	-------------------------------------

Identify external applications and settings for VxRail

VxRail depends specific applications in your data center to be available over your data center network. These data center applications must be accessible to the VxRail management network.

Set time zone
and NTP server

A **time zone** is required. It is configured on vCenter server and each ESXi host during VxRail initial configuration.

An **NTP server** is not required, but is recommended. If you provide an NTP server, vCenter server will be configured to use it. If you do not provide at least one NTP server, VxRail uses the time that is set on ESXi host #1 (regardless of whether the time is correct or not).

Note: Ensure that the NTP IP address is accessible from the VxRail External Management Network which the VxRail nodes will be connected to and is functioning properly.

Set DNS for
VxRail
management
components

Network Configuration Table ✓ Row 9	Enter your time zone.
Network Configuration Table ✓ Row 10	Enter the hostnames or IP addresses of your NTP servers.

Starting with VxRail version 7.0.010, you can either use an internal DNS included with VxRail vCenter Server, or use an external DNS in your data center. If you choose to use the internal DNS method, the steps to set up DNS as outlined in this section can be skipped.

If the internal DNS option is not selected, one or more external, customer-supplied DNS servers are required for VxRail. The DNS server that you select for VxRail must be able to support naming services for all the VxRail management components (VxRail Manager, vCenter, and so on).

Note: Ensure that the DNS IP address is accessible from the network to which VxRail is connected and functioning properly.

Network Configuration Table ✓ Row 11	Enter the IP addresses for your DNS servers.
---	--

Lookup records must be created in your selected DNS for every VxRail management component you are deploying in the cluster and are assigning a hostname and IP address. These components can include VxRail Manager, VxRail vCenter Server, VxRail Platform Service Controller, Log Insight, and each ESXi host in the VxRail cluster. The DNS entries must support both forward and reverse lookups.






 mrm-md-n1	Host (A)	192.1.0.10
 mrm-md-n2	Host (A)	192.1.0.11
 mrm-md-n3	Host (A)	192.1.0.12
 mrm-md-n4	Host (A)	192.1.0.13
 mrm-md-n5	Host (A)	192.1.0.14
 mrm-md-ivc	Host (A)	192.1.0.20
 mrm-md-vxrm	Host (A)	192.1.0.22

Figure 33. Sample DNS Forward Lookup Entries








 192.1.0.10	Pointer (PTR)	mrm-md-n1.mrmvxrail.local.
 192.1.0.11	Pointer (PTR)	mrm-md-n2.mrmvxrail.local.
 192.1.0.12	Pointer (PTR)	mrm-md-n3.mrmvxrail.local.
 192.1.0.13	Pointer (PTR)	mrm-md-n4.mrmvxrail.local.
 192.1.0.14	Pointer (PTR)	mrm-md-n5.mrmvxrail.local.
 192.1.0.20	Pointer (PTR)	mrm-md-ivc.mrmvxrail.local.
 192.1.0.22	Pointer (PTR)	mrm-md-vxrm.mrmvxrail.local.

Figure 34. Sample DNS Reverse Lookup Entries

Use the [Appendix A: VxRail Network Configuration Table](#) to determine which VxRail management components to include in your planned VxRail cluster, and have assigned a hostname and IP address. vMotion and vSAN IP addresses are not configured for routing by VxRail, so there are no entries required in the DNS server.

Prepare customer-supplied vCenter server

Note: You can skip this section if you plan to use the VxRail vCenter server. These action items are only applicable if you plan to use a customer-supplied vCenter server in your data center for VxRail.

Certain prerequisites must be completed before VxRail initial implementation if you use a customer-supplied vCenter as the VxRail cluster management platform. During the VxRail initialization process, it will connect to your customer-supplied vCenter to perform necessary validation steps, and perform configuration steps, to deploy the VxRail cluster on your vCenter instance.

- Determine if your customer-supplied vCenter server is compatible with your VxRail version.

- See the Knowledge Base article *VxRail: VxRail and External vCenter Interoperability Matrix* on the Dell product support site for the latest support matrix.
- Enter the FQDN of your selected, compatible customer-supplied vCenter server in the Appendix A: VxRail Network Configuration Table.

Network Configuration Table ✓ Row 35	Enter the FQDN of the customer-supplied vCenter Server.
---	---

- Determine whether your customer-supplied vCenter server has an embedded or external platform services controller. If the platform services controller is external to your customer-supplied vCenter, enter the platform services controller FQDN in the [Appendix A: VxRail Network Configuration Table](#).

Network Configuration Table ✓ Row 34	Enter the FQDN of the customer-supplied platform services controller (PSC). Leave this row blank if the PSC is embedded in the customer-supplied vCenter server.
---	---

- Decide on the single sign-on (SSO) domain that is configured on the customer-supplied vCenter you want to use to enable connectivity for VxRail, and enter the domain in the [Appendix A: VxRail Network Configuration Table](#).

Network Configuration Table ✓ Row 36	Enter the single sign-on (SSO) domain for the customer-supplied vCenter server. (For example, vsphere.local)
---	--

- The VxRail initialization process requires login credentials to your customer-supplied vCenter. The credentials must have the privileges to perform the necessary configuration work for VxRail. You have two choices:
 - Provide vCenter login credentials with administrator privileges.
 - Create a new set of credentials in your vCenter for this purpose. Two new roles will be created and assigned to this user by your Dell Technologies delivery services.

Network Configuration Table ✓ Row 37	Enter the administrative username/password for the customer-supplied vCenter server, or the VxRail non-admin username/password you will create on the customer-supplied vCenter server.
---	---

- A set of credentials must be created in the customer-supplied vCenter for VxRail management with no permissions and no assigned roles. These credentials are assigned a role with limited privileges during the VxRail initialization process, and then assigned to VxRail to enable connectivity to the customer-supplied vCenter after initialization completes.
 - If this is the first VxRail cluster on the customer-supplied vCenter, enter the credentials that you will create in the customer-supplied vCenter.

- If you already have an account for a previous VxRail cluster in the customer-supplied vCenter, enter those credentials.

Network Configuration Table ✓ Row 38	Enter the full VxRail management username/password. (For example, cluster1-manager@vsphere.local)
---	--

- The VxRail initialization process will deploy the VxRail cluster under an existing data center in the customer-supplied vCenter. Create a new data center, or select an existing Data center on the customer-supplied vCenter.

Network Configuration Table ✓ Row 39	Enter the name of a data center on the customer-supplied vCenter server.
---	--

- Specify the name of the cluster that will be created by the VxRail initialization process in the selected data center. This name must be unique, and not used anywhere in the data center on the customer-supplied vCenter.

Network Configuration Table ✓ Row 40	Enter the name of the cluster that will be used for VxRail.
---	---

Prepare customer-supplied virtual distributed switch

You can skip this section if your VxRail version is not 7.0.010 or later, or if you do not plan to deploy VxRail against one or more customer-supplied virtual distributed switch.

Before VxRail version 7.0.010, if you chose to deploy the VxRail cluster on an external, customer-supplied vCenter, a virtual distributed switch would be configured on the vCenter instance as part of the initial cluster build process. The automated initial build process would deploy the virtual distributed switch adhering to VxRail requirements in the vCenter instance, and then attach the VxRail networks to the portgroups on the virtual distributed switch. Depending on the target version planned for your VxRail cluster, you can decide to preconfigure one or two virtual distributed switches on your external vCenter instance to support VxRail networking.

- Starting with VxRail version 7.0.010, you have the choice of configuring a single virtual distributed switch to the external vCenter before the initial cluster build process.
- Starting with VxRail version 7.0.130, you have the choice of configuring one or two virtual distributed switches to the external vCenter instance before the initial cluster build process.

If you choose to manually configure the virtual switches and configure the network before initial cluster build, you must perform the following prerequisites:

- Unless your data center already has a vCenter instance compatible with VxRail, deploy a vCenter instance that will serve as the target for the VxRail cluster.
- You can deploy the VxRail cluster to an existing virtual distributed switch or a pair of virtual distributed switches on the target vCenter instance.

- Configure a portgroup for each of the required VxRail networks. Dell Technologies recommends using naming standards that clearly identify the VxRail network traffic type.
- Configure the VLAN assigned to each required VxRail network on the respective portgroup. The VLANs for each VxRail network traffic type can be referenced in the 'VxRail Networks' section in [Appendix A: VxRail Network Configuration Table](#).
- Configure two or four uplinks on the virtual distributed switch or pair of virtual distributed switches to support the VxRail cluster.
- Configure the teaming and failover policies for the distributed port groups. Each port group is assigned a teaming and failover policy. You can choose a simple strategy and configure a single policy that is applied to all port groups, or configure a set of policies to address requirements at the port group level.
- If you plan to enable load balancing with LACP against any non-management VxRail networks, configure the LACP policy on the virtual distributed switch, and apply the policy to the appropriate portgroup or portgroups.

Name ↑	VLAN ID	VMs
VxRail External Mgmt Network-1	VLAN access: 110	0
VxRail Internal Mgmt Network-1	VLAN access: 3939	0
VxRail vCenter Server Network-1	VLAN access: 110	0
VxRail Virtual SAN Network-1	VLAN access: 112	0
VxRail vMotion Network-1	VLAN access: 111	0

Figure 35. Sample portgroups on customer-supplied virtual distributed switch

Dell Technologies recommends referencing the configuration settings applied to the virtual distributed switch by the automated VxRail initial build process as a baseline. This ensures a successful deployment of a VxRail cluster against the customer-supplied virtual distributed switch. The settings used by the automated initial build process can be found in [Appendix E: Virtual Distributed Switch Portgroup Default Settings](#).

Network Configuration Table ✓ Row 41	Enter the name of the virtual distributed switch that will support the VxRail cluster networking.
Network Configuration Table ✓ Row 42	If a decision is made to configure two virtual distributed switches, enter the name of the second virtual distributed switch.
Network Configuration Table ✓ Row 43	Enter the name of the portgroup that will enable connectivity for the VxRail external management network.
Network Configuration Table ✓ Row 44	Enter the name of the portgroup that will enable connectivity for the VxRail vCenter Server network.

Network Configuration Table ✓ Row 45	Enter the name of the portgroup that will enable connectivity for the VxRail internal management network.
Network Configuration Table ✓ Row 46	Enter the name of the portgroup that will enable connectivity for the vMotion network.
Network Configuration Table ✓ Row 47	Enter the name of the portgroup that will enable connectivity for the vSAN network.

If your plan is to have more than one VxRail cluster deployed against a single customer-supplied virtual distributed switch, Dell Technologies recommends establishing a distinctive naming standard for the distributed port groups. This will ease network management and help distinguish the individual VxRail networks among multiple VxRail clusters.

Configuring portgroups on the virtual distributed switch for any guest networks you want to have is not required for the VxRail initial build process. These portgroups can be configured after the VxRail initial build process is complete. Dell Technologies also recommends establishing a distinctive naming standard for these distributed port groups.

Configure teaming and failover policies for customer-supplied virtual distributed switch

For a customer-supplied virtual distributed switch, you can use the default teaming and failover policy for VxRail, or customize teaming and failover policies for each portgroup. The default teaming and failover policy for VxRail is described in [Configure teaming and failover policies for VxRail networks](#).

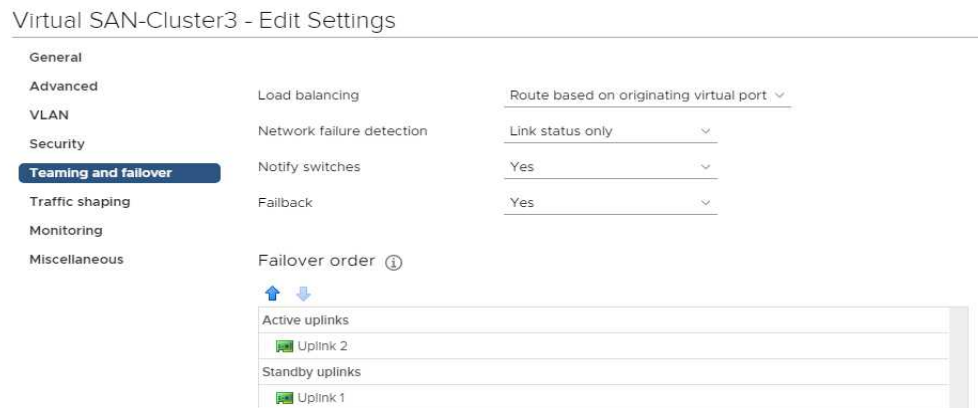


Figure 36. Sample VxRail default teaming and failover policy

Customizing the teaming and failover policies can also be performed as a post-deployment operation instead of as a prerequisite.

Prepare link aggregation on customer-supplied virtual distributed switch

You can skip this section if your VxRail version is not 7.0.130 or later, and you do not plan to enable link aggregation against one or more VxRail networks on the customer-supplied virtual distributed switch.

Starting with VxRail version 7.0.130, you can configure link aggregation against the VxRail non-management networks, which include the vSAN network and vMotion network. The following pre-requisites must be met to enable link aggregation on the VxRail non-management networks:

- Four ports from each VxRail node must be configured to support VxRail networking.

- Two ports will be configured to support VxRail management networks. Link aggregation is not supported on these networks.
- Two ports will be configured to support VxRail non-management networks. Link aggregation is supported on these networks.
- Link aggregation can be configured on the vSAN network, vMotion network, or both.
 - If link aggregation is not to be configured on the vMotion network, this network must be assigned to the same uplinks supporting the VxRail management networks.
- The adjacent top-of-rack switches must be configured to support link aggregation. See the guides provided by your switch vendor to perform this task.

The following tasks must be completed on the virtual distributed switch on your customer-supplied vCenter instance to support link aggregation:

- Configure an LACP policy on the virtual distributed switch

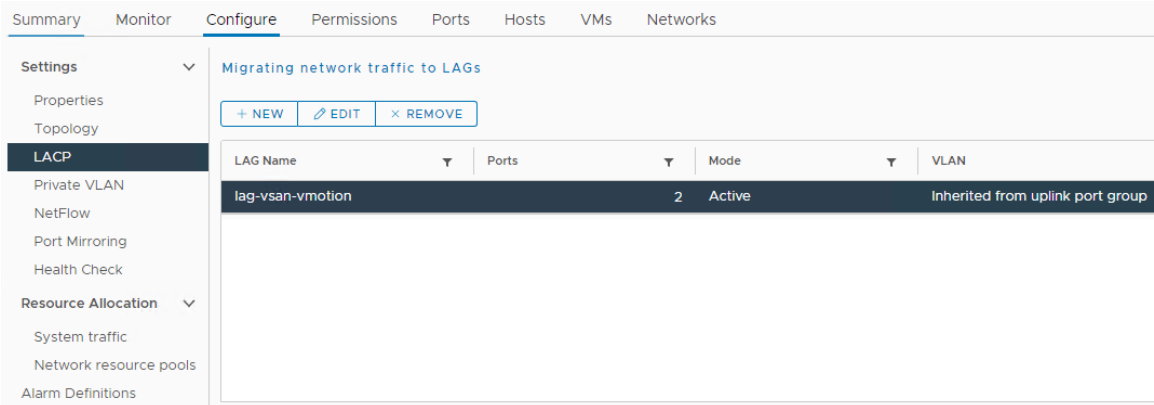


Figure 37. Sample LACP policy configured on virtual distributed switch

- Configure the teaming and failover policy on the portgroup(s) targeted for link aggregation.

General

Advanced

VLAN

Security

Teaming and failover

Traffic shaping

Monitoring

Miscellaneous

Load balancing: Route based on IP hash

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active uplinks

lag-vsan-vmotion

Standby uplinks

Unused uplinks

uplink1

uplink2

uplink3

uplink4

Figure 38. Sample LACP Policy configured as active uplink in teaming and failover policy

Reserve IP addresses for VxRail vMotion network

An IP address is required for the vMotion network for each ESXi host in the VxRail cluster. A private address range is acceptable if you decide the vMotion network will not be routable. If your plans include the ability to migrate virtual machines outside of the VxRail cluster, that needs to be considered when selecting the IP address scheme.

Starting with VxRail version 7.0.010, you can choose to have the IP addresses assigned automatically during VxRail initial build, or manually select the IP addresses for each ESXi host. If the VxRail version is earlier than 7.0.010, auto-assignment method by VxRail is the only option.

For the auto-assignment method, the IP addresses for VxRail initial build must be contiguous, with the specified range in a sequential order. The IP address range must be large enough to cover the number of ESXi hosts planned for the VxRail cluster. A larger IP address range can be specified to cover for planned expansion.

If your plans include expanding the VxRail cluster to deploy nodes in more than one physical rack, you have the option of whether to stretch the IP subnet for vMotion between the racks, or to use routing services in your data center instead.

For the IP address auto-assignment method, record the IP address range.

Network Configuration Table ✓ Rows 53-54	Enter the starting and ending IP addresses for vSphere vMotion.
---	---

For the manual assignment method, record the IP addresses.

Network Configuration Table ✓ Rows 55-58	Enter the IP addresses for vSphere vMotion.
---	---

Enter the subnet mask and default gateway. You can use the default gateway assigned to the VxRail External Management network, or enter a gateway dedicated for the vMotion network.

Reserve IP addresses for VxRail vSAN network

Network Configuration Table ✓ Row 59	Enter the subnet mask for vMotion.
Network Configuration Table ✓ Row 60	Enter the default gateway for vMotion.

An IP address is required for the vSAN network for each ESXi host in the VxRail cluster. A private address range is acceptable unless you decide you may expand beyond one rack and want to use a different subnet for expansion.

Starting with VxRail version 7.0.010, you can choose to have the IP addresses assigned automatically during VxRail initial build, or manually select the IP addresses for each ESXi host. If the VxRail version is earlier than 7.0.010, auto-assignment method by VxRail is the only option.

For the auto-assign method, the IP addresses for the initial build of the VxRail cluster must be contiguous, with the specified range in a sequential order. The IP address range must be large enough to cover the number of ESXi hosts planned for the VxRail cluster. A larger IP address range can be specified to cover for planned expansion.

For the IP address auto-assignment method, record the IP address range.

Network Configuration Table ✓ Rows 61-62	Enter the starting and ending IP addresses for vSAN.
---	--

For the manual assignment method, record the IP addresses

Network Configuration Table ✓ Rows 63-66	Enter the IP addresses for vSAN.
---	----------------------------------

Enter the subnet mask for the vSAN network.

Decide on VxRail logging solution

Network Configuration Table ✓ Row 67	Enter the subnet mask for vSAN.
---	---------------------------------

Decide whether to use your own third-party syslog server, use the vRealize Log Insight solution included with VxRail, or no logging. You can only select the vRealize Log Insight option if:

- You will deploy the vCenter instance included with the VxRail onto the VxRail cluster.
- The VxRail cluster to be deployed is version 7.0.010 or earlier.

If you use a customer-supplied vCenter server, you can either use your own third-party syslog server, or no logging. If you choose the vRealize Log Insight option, the IP address that is assigned to Log Insight must be on the same subnet as the VxRail management network.

Network Configuration Table ✓ Row 69 or Row 70	Enter the IP address for vRealize Log Insight or the hostnames of your existing third-party syslog servers. Leave blank for no logging.
---	---

Assign passwords for VxRail management

You will need to assign a password to the accounts that are members of the VxRail management ecosystem. See the

Appendix B: VxRail Passwords table to use as worksheets for your passwords.

Note: The Dell Technologies service representative will need passwords for the VxRail accounts in this table. For security purposes, you can enter the passwords during the VxRail initialization process, as opposed to providing them visibly in a document.

- For ESXi hosts, passwords must be assigned to the 'root' account. You can use one password for each ESXi host or apply the same password to each host.
- For VxRail Manager, a password must be assigned to the 'root' account [Row 1]. This credential is for access to the console.
- Access to the VxRail Manager web interface will use the 'administrator@<SSO Domain>' credentials.
 - If you deploy the VxRail vCenter Server, VxRail Manager and vCenter share the same default administrator login, 'administrator@vsphere.local'. Enter the password that you want to use [Row 2].
 - If you use a customer-supplied vCenter server, VxRail Manager will use the same 'administrator@<SSO Domain>' login credentials you use for access to the customer-supplied vCenter server.
- If you deploy the VxRail vCenter Server:
 - Enter the 'root' password for the VxRail vCenter Server [Row 3].
 - Enter a password for 'management' for the VxRail vCenter Server [Row 4].
 - A Platform Services controller will be deployed. Enter the 'root' password for the Platform Services controller [Row 5].
- If you deploy vRealize Log Insight:
 - Enter a password for 'root' [Row 6].
 - Enter a password for 'admin' [Row 7].

Passwords must adhere to VMware vSphere complexity rules. Passwords must contain between eight and 20 characters with at least one lowercase letter, one uppercase letter, one numeric character, and one special character. For more information about password requirements, see the [vSphere password](#) and [vCenter Server password](#) documentation.

Prepare for Dell EMC SmartFabric Services enablement

Note: Skip this section if you do not plan to enable Dell EMC SmartFabric Services to pass control of switch configuration to VxRail.

The planning and preparation tasks for the deployment and operations of a VxRail cluster on a network infrastructure enabled with SmartFabric Services differ from connecting a VxRail cluster to a standard data center network. The basic settings that are required for the initial buildout of the network infrastructure with SmartFabric Services are outlined in this section.

Enabling the SmartFabric personality on a Dell Ethernet switch that is qualified for SmartFabric Services initiates a discovery process for other connected switches with the

same SmartFabric personality for the purposes of forming a unified switch fabric. A switch fabric can start as small as two leaf switches in a single rack, then expand automatically by enabling the SmartFabric personality on connected spine switches, and connected leaf switches in expansion racks.

Both the Dell Ethernet switches and VxRail nodes advertise themselves at the time of power-on on this same internal discovery network. The SmartFabric-enabled network also configures an 'untagged' virtual network on the switch fabric to enable client onboarding through a jump port for access to VxRail Manager to perform cluster implementation. During VxRail initial configuration through VxRail Manager, the required VxRail networks are automatically configured on the switch fabric.

- Network connectivity to out-of-band management for each switch that is enabled with the SmartFabric personality is a requirement for VxRail. A reserved IP address is required for each switch.
- A separate Ethernet switch outside of SmartFabric is required to support connectivity to switch management through the out-of-band network.
- A reserved IP address for iDRAC connectivity to each VxRail node on this same separate management switch is recommended.

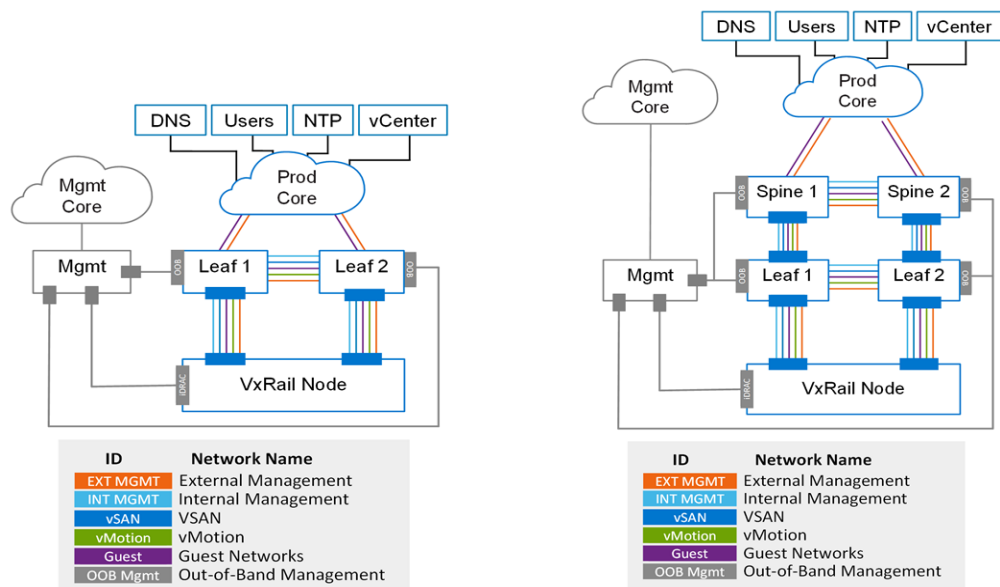


Figure 39. Logical networks for single-tier and two-tier SmartFabric deployments

The Dell EMC Open Management Network Interface (OMNI) plug-in must be deployed on the vCenter instance to support automated switch management after the VxRail cluster is built. The Dell EMC OMNI vCenter plug-in is required for each Dell EMC switch fabric pair, and requires network properties to be set during the deployment process.

Network Configuration Table ✓ Rows 71 and 72	Reserve an IP address for out-of-band management of each switch in the SmartFabric-enabled network.
Network Configuration Table ✓ Row 73	Enter the IP address for Dell EMC OMNI vCenter plug-in.

Network Configuration Table ✓ Row 74	Enter the subnet mask for Dell EMC OMNI vCenter plug-in.
Network Configuration Table ✓ Row 75	Enter the gateway for Dell EMC OMNI vCenter plug-in.

For complete details on the settings that are needed during the planning and preparation phase for a SmartFabric-enabled network, see the 'Dell EMC VxRail™ with SmartFabric Network Services Planning and Preparation Guide' on the [Dell Technologies VxRail Technical Guides](#) site.

Chapter 7 Configure the Network for VxRail

This chapter presents the following topic:

Introduction74

Setting up the network switch for VxRail connectivity.....74

Setting up the upstream network for VxRail connectivity.....79

Confirm your data center network80

Confirm your data center environment82

Introduction

For the VxRail initialization process to pass validation and build the cluster, you must configure the adjacent top-of-rack switches and upstream network **before you plug in VxRail nodes and powering them on**.

This section provides guidance on the tasks that must be undertaken on the data center network to prepare for the VxRail initial implementation. You can use the information in [Appendix C: VxRail Setup Checklist](#) for guidance. Be sure to follow your vendor's documentation for specific switch configuration activities and for best practices for performance and availability.

You can skip this section if you plan to enable Dell EMC SmartFabric Services and extend VxRail automation to the TOR switch layer.

Setting up the network switch for VxRail connectivity

Follow the steps in this section for the configuration settings required for VxRail networking.

Configure multicast for VxRail Internal Management network

Note: If you do not plan to use the auto-discover method due to multicast restrictions, and will use the manual method instead for selecting nodes for the cluster build operation, this task can be skipped.

VxRail clusters have no backplane, so communication between its nodes is facilitated through the network switch. This communication between the nodes for device discovery purposes uses VMware's Loudmouth capabilities, which are based on the RFC-recognized "Zero Network Configuration" protocol. New VxRail nodes advertise themselves on the network using the VMware Loudmouth service, and are discovered by VxRail Manager with the Loudmouth service.

VMware's Loudmouth service depends on multicasting, which is required for the VxRail internal management network. The network switch ports that connect to VxRail nodes must allow for pass-through of multicast traffic on the VxRail Internal Management VLAN. Multicast is *not* required on your entire network, just on the ports connected to VxRail nodes.

VxRail creates very little traffic through multicasting for auto-discovery and device management. Furthermore, the network traffic for the Internal Management network is restricted through a VLAN. You can choose to enable *MLD Snooping* and *MLD Querier* on the VLAN if supported on your switches.

If MLD Snooping is enabled, MLD Querier **must be** enabled. If MLD Snooping is disabled, MLD Querier **must be** disabled.

Configure unicast for VxRail vSAN network

For early versions of VxRail, multicast was required for the vSAN VLAN. One or more network switches that connected to VxRail had to allow for the pass-through of multicast traffic on the vSAN VLAN. Starting with VxRail v4.5, all vSAN traffic replaces multicast with unicast. This change helps reduce network configuration complexity and simplifies

switch configuration. Unicast is a common protocol enabled by default on most enterprise Ethernet switches.

If you are required to configure multicast, note that VxRail multicast traffic for vSAN will be limited to broadcast domain per vSAN VLAN. There is minimal impact on network overhead as management traffic is nominal. You can limit multicast traffic by enabling IGMP Snooping and IGMP Querier. We recommend enabling both IGMP Snooping and IGMP Querier if your switch supports them and you configure this setting.

IGMP Snooping software examines IGMP protocol messages within a VLAN to discover which interfaces are connected to hosts or other devices that are interested in receiving this traffic. Using the interface information, IGMP Snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding an entire VLAN. IGMP Snooping tracks ports that are attached to multicast-capable routers to help manage IGMP membership report forwarding. It also responds to topology change notifications.

IGMP Querier sends out IGMP group membership queries on a timed interval, retrieves IGMP membership reports from active members, and allows updates to group membership tables. By default, most switches enable IGMP Snooping but disable IGMP Querier. You will need to change the settings if this is the case.

If IGMP Snooping is enabled, IGMP Querier must be enabled. If IGMP Snooping is disabled, IGMP Querier must be disabled.

Configure VLANs for the VxRail networks

Configure the VLANs on the switches depending on the VxRail version being deployed and the type of cluster being deployed. The VLANs are assigned to the switch ports as a later task.

For VxRail clusters using version 4.7 or later:

- VxRail External Management VLAN (default is untagged/native).
- VxRail Internal Management VLAN – ensure that multicast is enabled on this VLAN if enabling node discovery.

For VxRail clusters using versions earlier than 4.7:

- VxRail Management VLAN (default is untagged/native) – ensure that multicast is enabled on this VLAN.

For all VxRail clusters:

- vSAN VLAN – ensure that unicast is enabled.
- vSphere vMotion VLAN
- VM Networks VLAN (need at least one for guest VM traffic)

The additional VxRail Witness traffic separation VLAN to manage traffic between the VxRail cluster and the witness. This is only needed if deploying VxRail stretched-cluster or 2-Node cluster.

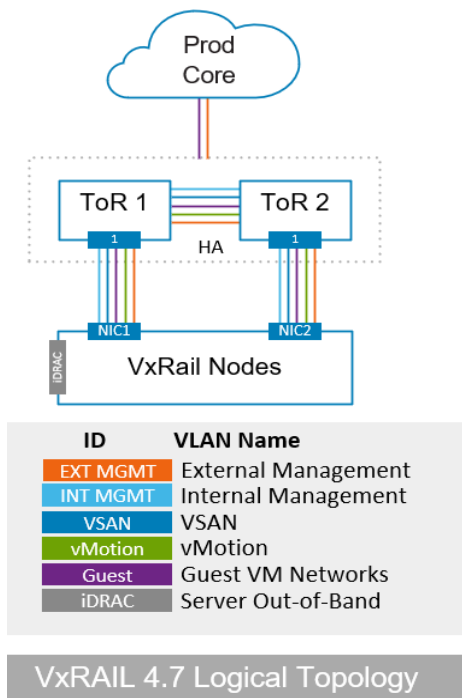


Figure 40. VxRail Logical Networks: Version 4.7 and later

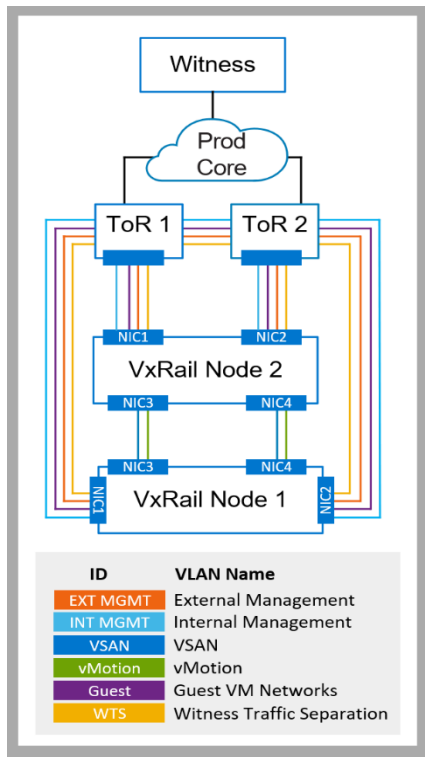


Figure 41. VxRail Logical Networks: 2-Node Cluster with Witness

Using the [VxRail Network Configuration Table](#), perform the following steps:

1. Configure the **External Management VLAN (Row 1)** on the switches. If you entered “Native VLAN,” set the ports on the switch to accept untagged traffic and

tag it to the native management VLAN ID. Untagged management traffic is the default management VLAN setting on VxRail.

2. For VxRail version 4.7 and later, configure the **Internal Management VLAN (Row 2)** on the switches.
3. Allow multicast on the Internal Management network.
4. Configure a **vSphere vMotion VLAN (Row 3)** on the switches.
5. Configure a **vSAN VLAN (Row 4)** on the switches. Unicast is required for VxRail clusters built with version 4.5 and later.
6. Configure the VLANs for your **VM Networks (Rows 6)** on the switches. These networks can be added after the cluster initial build is complete.
7. Configure the optional **VxRail Witness Traffic Separation VLAN (Row 70)** on the switches ports if required.
8. Configure the switch uplinks to allow the **External Management VLAN (Row 1)** and **VM Network VLANs (Row 6)** to pass through, and optionally the **vSphere vMotion VLAN (Row 3)** and **vSAN VLAN (Row 4)**. If a vSAN witness is required for the VxRail cluster, include the **VxRail Witness Traffic Separation VLAN (Row 70)** on the uplinks.

Configure the inter-switch links

If more than one top-of-rack switch is being deployed to support the VxRail cluster, configure inter-switch links between the switches. Configure the inter-switch links to allow the **all VLANs** to pass through.

Configure switch ports

Determine switch port mode

Configure the port mode on your switch based on the plan for the VxRail logical networks, and whether VLANs will be used to segment VxRail network traffic. Ports on a switch operate in one of the following modes:

- *Access mode* – The port accepts untagged packets only and distributes the untagged packets to all VLANs on that port. This is typically the default mode for all ports. This mode should only be used for supporting VxRail clusters for test environments or temporary usage.
- *Trunk mode* – When this port receives a tagged packet, it passes the packet to the VLAN specified in the tag. To configure the acceptance of untagged packets on a trunk port, you must first configure a single VLAN as a “Native VLAN.” A “Native VLAN” is when you configure one VLAN to use as the VLAN for all untagged traffic.
- *Tagged-access mode* – The port accepts tagged packets only.

Disable link aggregation on switch ports supporting VxRail networks

Link aggregation is supported for the VxRail initial implementation process only if the VxRail version on the nodes is 7.0.130, and you correctly follow the guidance to deploy the virtual distributed switches on your external vCenter with the proper link aggregation settings. If either of these conditions are not applicable, do not enable link aggregation, including protocols such as LACP and EtherChannel, on any switch ports that are connected to VxRail node ports before initial implementation.

During the VxRail initial build process, either 2 or 4 ports will be selected on each node to support the VxRail management networks and any guest networks configured at that time. The VxRail initial build process will configure a virtual distributed switch on the cluster, and then configure a portgroup on that virtual distributed switch for each VxRail management network.




Figure 42. Unused VxRail node ports configured for non-VxRail network traffic

When the initial implementation process completes, you can configure link aggregation on the operational VxRail cluster, as described in [Configure link aggregation on VxRail networks](#). If your requirements include using any spare network ports on the VxRail nodes that were not configured for VxRail network traffic for other use cases, then link aggregation can be configured to support that network traffic. These can include any unused ports on the NDC or on the optional PCIe adapter cards. Updates can be configured on the virtual distributed switch deployed during VxRail initial build to support the new networks, or a new virtual distributed switch can be configured. Since the initial virtual distributed switch is under the management and control of VxRail, the best practice is to configure a separate virtual distributed switch on the vCenter instance to support these networking use cases.

Limit spanning tree protocol on VxRail switch ports

Network traffic must be allowed uninterrupted passage between the physical switch ports and the VxRail nodes. Certain Spanning Tree states can place restrictions on network traffic and can force the port into an unexpected timeout mode. These conditions that are caused by Spanning Tree can disrupt VxRail normal operations and impact performance.

If Spanning Tree is enabled in your network, ensure that the physical switch ports that are connected to VxRail nodes are configured with a setting such as 'Portfast', or set as an edge port. These settings set the port to forwarding state, so no disruption occurs. Because vSphere virtual switches do not support STP, physical switch ports that are connected to an ESXi host must have a setting such as 'Portfast' configured if spanning tree is enabled to avoid loops within the physical switch network.

Enable flow control

Network instability or congestion contributes to low performance in VxRail, and has a negative effect on the vSAN I-O datastore operations. VxRail recommends enabling flow

control on the switch to assure reliability on a congested network. Flow control is a switch feature that helps manage the rate of data transfer to avoid buffer overrun. During periods of high congestion and bandwidth consumption, the receiving network will inject pause frames for a period of time to the sender network to slow transmission in order to avoid buffer overrun. The absence of flow control on a congested network can result in increased error rates and force network bandwidth to be consumed for error recovery. The flow control settings can be adjusted depending on network conditions, but VxRail recommends that flow control should be 'receive on' and 'transmit off'.

Configure ports on your switches

Now that the switch base settings are complete, the next step is the switch ports. Perform the following steps for each switch port that will be connected to a VxRail node:

1. Configure the MTU size if using jumbo frames.
2. Set the port to the appropriate speed or to auto-negotiate speed.
3. Set spanning tree mode to disable transition to a blocking state, which can cause a timeout condition
4. Enable flow control receive mode and disable flow control transmit mode.
5. Configure the **External Management VLAN (Row 1)** on the switch ports. If you entered "Native VLAN," set the ports on the switch to accept untagged traffic and tag it to the native management VLAN ID. Untagged management traffic is the default management VLAN setting on VxRail.
6. For VxRail version 4.7 and later, configure the **Internal Management VLAN (Row 2)** on the switch ports.
7. If required, allow multicast on the VxRail switch ports to support the Internal Management network.
8. Configure a **vSphere vMotion VLAN (Row 3)** on the switch ports.
9. Configure a **vSAN VLAN (Row 4)** on the switch ports. Allow unicast traffic on this VLAN.
10. Configure the VLANs for your **VM Networks (Rows 6)** on the switch ports.
11. Configure the optional **VxRail Witness Traffic Separation VLAN (Row 70)** on the switch ports, if required.

Setting up the upstream network for VxRail connectivity

The upstream network from the VxRail cluster must be configured to allow passage for VxRail networks that require external access. Using [Appendix A: VxRail Network Configuration Table](#) for reference, upstream passage is required for the **External Management VLAN (Row 1)** and any **VM Network VLANs (Row 6)**. If a vSAN witness is required for the VxRail cluster, include the **VxRail Witness Traffic Separation VLAN (Row 70)** for upstream passage. The **VxRail Internal Management VLAN (Row 2)** must be blocked from outbound upstream passage.

Optionally, the **vSphere vMotion VLAN (Row 3)** and **vSAN VLAN (Row 4)** can be configured for upstream passage. If you plan to expand the VxRail cluster beyond a single

rack, configure the VxRail network VLANs for either stretched Layer 2 networks across racks, or to pass upstream to routing services if new subnets will be assigned in expansion racks.

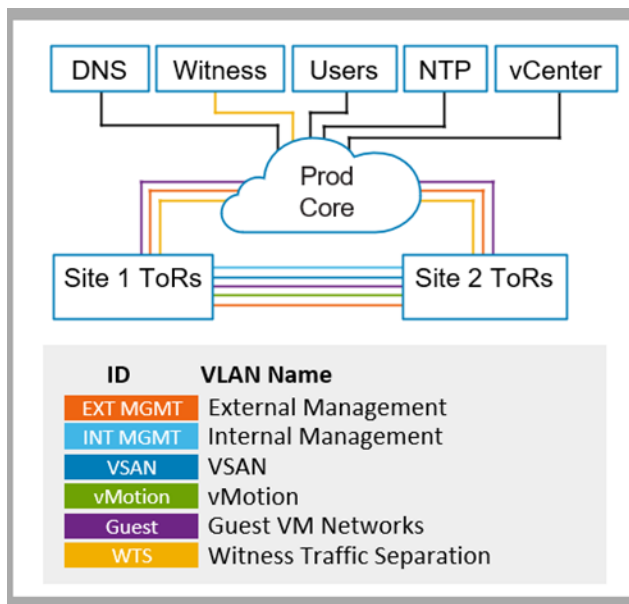


Figure 43. Logical Networks including Upstream Elements

If your Layer 2/Layer 3 boundary is at the lowest network tier (top-of-rack switch), then perform the following tasks:

- Configure point-to-point links with the adjacent upstream switches.
- Terminate the VLANs requiring upstream access on the top-of-rack switches.
- Enable and configure routing services for the VxRail networks requiring upstream passage.

If your Layer 2/Layer 3 boundary is upstream from at the lowest network tier (top-of-rack switch), then perform the following tasks:

- Connect ports on the adjacent upstream switch to the uplinks on the top-of-rack switches.
- Configure logical pairings of the ports on the adjacent upstream switch and the top-of-rack switch.
- Configure the logical port pairings, commonly known as 'port channels' or 'EtherChannels', to allow upstream passage of external VxRail networks.

Confirm your data center network

Upon completion of the switch configuration, there should be unobstructed network paths between the switch ports and the ports on the VxRail nodes. The VxRail management network and VM network should have unobstructed passage to your data center network. Before forming the VxRail cluster, the VxRail initialization process will perform several verification steps, including:

- Verifying switch and data center environment supportability
- Verifying passage of VxRail logical networks
- Verifying accessibility of required data center applications
- Verifying compatibility with the planned VxRail implementation

Certain data center environment and network configuration errors will cause the validation to fail, and the VxRail cluster will not be formed. When validation fails, the data center settings and switch configurations must undergo troubleshooting to resolve the problems reported.

Confirm the settings on the switch, using the switch vendor instructions for guidance:

1. External management traffic will be untagged on the native VLAN by default. If a tagged VLAN is used instead, the switches must be customized with the new VLAN.
2. Internal device discovery network traffic will use the default VLAN of 3939. If this has changed, all ESXi hosts must be customized with the new VLAN, or device discovery will not work.
3. Confirm that the switch ports that will attach to VxRail nodes allow passage of all VxRail network VLANs.
4. Confirm that the switch uplinks allow passage of external VxRail networks.
5. If you have two or more switches, confirm an inter-switch link is configured between them to support passage of the VxRail network VLANs.

Confirm your firewall settings

If you have positioned a firewall between the switches that are planned for VxRail and the rest of your data center network, be sure that the required firewall ports are open for VxRail network traffic.

1. Verify that VxRail can communicate with your DNS server.
2. Verify that VxRail can communicate with your NTP server.
3. Verify that your IT administrators can communicate with the VxRail management system.
4. If you plan to use a customer-supplied vCenter, verify open communication between the vCenter instance and the VxRail managed hosts.
5. If you plan to use a third-party syslog server instead of Log Insight, verify that open communication between the syslog server and the VxRail management components.
6. If you plan to deploy a separate network for ESXi host management (iDRAC), verify that your IT administrators can communicate with the iDRAC network.
7. If you plan to use an external Secure Remote Services (SRS) gateway in your data center instead of SRS-VE deployed in the VxRail cluster, verify the open communications between VxRail management and the SRS gateway.

See [Appendix D: VxRail Open Ports Requirements](#) for information of VxRail port requirements.

Confirm your data center environment

1. Confirm that you cannot ping any IP address that is reserved for VxRail management components.
2. Confirm that your DNS servers are reachable from the VxRail external management network.
3. Confirm the forward and reverse DNS entries for the VxRail management components.
4. Confirm that your management gateway IP address is accessible.
5. If you decide to use the TCP-IP stack for vMotion instead of the default TCP-IP stack, confirm that your vMotion gateway IP address is accessible.
6. If you have configured NTP servers, or a third-party syslog server, confirm that you can reach them from your configured VxRail external management network.
7. If you plan to use a customer-supplied vCenter, confirm that it is accessible from the VxRail external management network.
8. If you plan to deploy a witness at a remote site to monitor vSAN, and plan to enable Witness Traffic Separation, confirm that there is a routable path between the witness and this network.
9. If you plan to install the VxRail nodes in more than one rack, and you plan to terminate the VxRail networks at the ToR switches, verify that routing services have been configured upstream for the VxRail networks.

Chapter 8 Preparing to Build the VxRail Cluster

This chapter provides the following topics:

Introduction	84
Configuring a workstation/laptop for VxRail initialization.....	84
Perform initialization to create a VxRail cluster	85

Introduction

The steps that are outlined in this section will be performed by Dell Technologies professional services. They are described here to provide insight into the activities to be performed during the delivery engagement.

Configuring a workstation/laptop for VxRail initialization

A workstation/laptop with a web browser for the VxRail user interface is required to perform the initialization process. It must be plugged into the top-of-rack switch, or be able to logically reach the VxRail external management VLAN from elsewhere on your network; for example, a jump server ([Jump Server Description](#)). Once the VxRail initialization process is complete, the switch port or jump host is no longer required to manage VxRail.

Note: Do not try to plug your workstation/laptop directly into a VxRail server node to connect to the VxRail management interface for initialization. It must be plugged into your network or switch, and the workstation/laptop must be logically configured to reach the necessary networks.

A supported web browser is required to access VxRail management interface. The latest versions of Firefox, Chrome, and Internet Explorer 10+ are all supported. If you are using Internet Explorer 10+ and an administrator has set your browser to “compatibility mode” for all internal websites (local web addresses), you will get a warning message from VxRail.

To access the VxRail management interface to perform initialization, you must use the temporary, preconfigured VxRail initial IP address: 192.168.10.200/24. This IP address will automatically change during VxRail initialization to your desired permanent address, and assigned to VxRail Manager during cluster formation.

Example Configuration	VxRail	Workstation/laptop		
	IP address/netmask	IP address	Subnet mask	Gateway
Initial (temporary)	192.168.10.200/24	192.168.10.150	255.255.255.0	192.168.10.254
Post-configuration (permanent)	10.10.10.100/24	10.10.10.150	255.255.255.0	10.10.10.254

Your workstation/laptop must be able to reach both the temporary VxRail initial IP address and the permanent VxRail Manager IP address (Row 26 from [Appendix A: VxRail Network Configuration Table](#)). VxRail initialization will remind you that you might need to reconfigure your workstation/laptop network settings to access the new IP address.

It is best practice to give your workstation/laptop or your jump server two IP addresses on the same network port, which allows for a smoother experience. Depending on your workstation/laptop, this can be implemented in several ways (such as dual-homing or multi-homing). Otherwise, change the IP address on your workstation/laptop when instructed to and then return to VxRail Manager to continue with the initialization process.

If you cannot reach the VxRail initial IP address, Dell Technologies support team can configure a custom IP address, subnet mask, and gateway on VxRail Manager before initialization.

Note: If a custom VLAN ID will be used for the VxRail management network other than the default “Native VLAN”, ensure the workstation/laptop can also access this VLAN.

Perform initialization to create a VxRail cluster

If you have successfully followed all the steps that are listed in this document, you are ready to move to the final phase: Connect the laptop or workstation to a switch port, and perform VxRail initialization. These steps are done by Dell Technologies service representatives and are included here to help you understand the complete process.

Before coming on-site, the Dell Technologies service representative will have contacted you to capture and record the information that is described in Appendix A: VxRail Network Configuration Table and walk through Appendix C: VxRail Setup Checklist.

- Step 1.** Before coming on-site, the Dell Technologies service presentative will have contacted you to capture and record the information that is described in [Appendix A: VxRail Network Configuration Table](#) and walk through [Appendix C: VxRail Setup Checklist](#).
- Step 2.** If your planned VxRail deployment requires a Witness at a remote data center location, the Witness virtual appliance is deployed.
- Step 3.** If your planned deployment includes the purchase of Dell Ethernet switches and professional services to install and configure the switches to support the VxRail cluster, that activity is performed before VxRail deployment activities commence.
- Step 4.** Install the VxRail nodes in a rack or multiple racks in the data center. If Dell professional services are not installing the switches, install the network switches supporting the VxRail cluster into the same racks for ease of management.
- Step 5.** Attach Ethernet cables between the ports on the VxRail nodes and switch ports that are configured to support VxRail network traffic.
- Step 6.** Power on the initial nodes to form the initial VxRail cluster. Do not turn on any other VxRail nodes until you have completed the formation of the VxRail cluster with the first three or four nodes.
- Step 7.** Connect a workstation/laptop configured for VxRail initialization to access the VxRail external management network on your selected VLAN. It must be either plugged into the switch or able to logically reach the VxRail external management VLAN from elsewhere on your network.
- Step 8.** Open a browser to the VxRail initial IP address to begin the VxRail initialization process.
- Step 9.** The Dell Technologies service representative will populate the input screens on the menu with the data collected from the customer during the planning and design process.
- Step 10.** If you have enabled Dell EMC SmartFabric Services, VxRail will automatically configure the switches that are connected to VxRail nodes using the information populated on the input screens.
- Step 11.** VxRail performs the verification process, using the information input into the menus.

- Step 12.** After validation is successful, the initialization process will begin to build a new VxRail cluster.
- Step 13.** The new permanent IP address for VxRail Manager will be displayed.
- If you configured the workstation/laptop to enable connectivity to both the temporary VxRail IP address and the new permanent IP address, the browser session will make the switch automatically. If not, you must manually change the IP settings on your workstation/laptop to be on the same subnet as the new VxRail IP address.
 - If your workstation/laptop cannot connect to the new IP address that you configured, you will get a message to fix your network and try again. If you are unable to connect to the new IP address after 20 minutes, VxRail will revert to its un-configured state and you will need to re-enter your configuration at the temporary VxRail IP address.
 - After the build process starts, if you close your browser, you will need to browse to the new, permanent VxRail IP address.
- Step 14.** Progress is shown as the VxRail cluster is built. The process takes about 25-40 minutes.
- Step 15.** When you see the **Hooray!** page, VxRail initialization is complete and a new VxRail cluster is built. Click the **Manage VxRail** button to continue to VxRail management. You should also bookmark this IP address in your browser for future use.
- Step 16.** Connect to VxRail Manager using either the VxRail Manager IP address (**Row 14**) or the fully qualified domain name (FQDN) (**Row 13**) that you configured on your DNS server. This will lead you to the vCenter instance.

Chapter 9 VxRail Network Considerations After Implementation

This chapter provides the following topics:

Introduction	88
Configure teaming and failover policies for VxRail networks.....	88
Using unassigned physical ports for VxRail networks	90
Configure link aggregation on VxRail networks	92
Deploy second VDS for VxRail networking	93

Introduction

This section provides guidance on additional actions that can be performed on the VxRail cluster networks upon the completion of the initial implementation process. Each section addresses specific requirements and use cases to optimize VxRail cluster operations.

The choices you have to modify the VxRail networking after the completion of the initial implementation phase depend on the VxRail version of your cluster, and potentially the configuration settings and supported features on your data center network.

Configure teaming and failover policies for VxRail networks

Starting with VxRail version 4.7.410, the customer can customize the teaming and failover policies for the VxRail networks. If your VxRail version is 7.0.010 or later, you can also customize the failover settings for the uplinks to an active-active configuration.

The decisions that you make on the teaming and failover policies are independent of the configuration on the adjacent top-of-rack switches. Be aware that load balancing has a dependency on physical switch settings, and requires link aggregation to be configured on the uplinks supporting the VxRail networks. Updating the teaming and failover policies with the top-of-rack switches will result in incremental improvements without also enabling link aggregation for load balancing purposes. Consideration should be given to coupling this effort with link aggregation to optimize bandwidth utilization.

During the initial implementation process, VxRail will apply a default teaming and failover policy for each VxRail network during the initial build operation.

- The default load-balancing policy is 'Route based on originating virtual port' for all VxRail network traffic. This policy directs network traffic up one uplink and does not attempt to share the workload between the uplinks.
- The default network failure detection setting is 'link status only'. This setting should not be changed. VMware recommends having 3 or more physical NICs in the team for 'beacon probing' to work correctly, which is not supported with VxRail.
- The setting for 'Notify switches' is set to 'Yes'. This instructs the virtual distributed switch to notify the adjacent physical switch of a failover.
- The setting for 'Failback' is set to 'Yes'. This instructs a failed adapter to take over for the standby adapter once it is recovered and comes online again, if the uplinks are in an active-standby configuration.
- The failover order for the uplinks depends the VxRail network that is configured on the portgroup. VxRail assigns the uplinks to each portgroup depending on whether two ports or four ports were selected to support the VxRail cluster during initial implementation. For any given portgroup, one uplink is assigned 'active' mode, one uplink is assigned 'standby' mode, and any uplinks that are marked as 'unused' are in either 'active' or 'standby' mode in another portgroup.

Virtual SAN-Cluster3 - Edit Settings

The screenshot shows the 'Virtual SAN-Cluster3 - Edit Settings' window. On the left, a sidebar lists settings categories: General, Advanced, VLAN, Security, **Teaming and failover** (highlighted), Traffic shaping, Monitoring, and Miscellaneous. The main area displays settings for 'Teaming and failover':

- Load balancing:** Route based on originating virtual port (dropdown)
- Network failure detection:** Link status only (dropdown)
- Notify switches:** Yes (dropdown)
- Failback:** Yes (dropdown)
- Failover order:** Includes an information icon and a list of uplinks. Under 'Active uplinks', 'Uplink 2' is listed with a green status icon. Under 'Standby uplinks', 'Uplink 1' is listed with a green status icon.

Figure 44. Default VDS teaming and failover policy for vSAN network configured with 2 VxRail ports.

The following portgroup load-balancing policies are supported for VxRail clusters running version 4.7.410 or later:

- **Route based on the originating virtual port**—After the virtual switch selects an uplink for a virtual machine or VMkernel adapter, it always forwards traffic through the same uplink. This option makes a simple selection based on the available physical uplinks. However, this policy does not attempt to load balance based on network traffic.
- **Route based on source MAC hash**—The virtual switch selects an uplink for a virtual machine based on the virtual machine MAC address. While it requires more resources than using the originating virtual port, it has more flexibility in uplink selection. This policy does not attempt to load balance based on network traffic analysis.
- **Use explicit failover order**— Always use the highest order uplink that passes failover detection criteria from the active adapters. No actual load balancing is performed with this option.
- **Route based on physical NIC load**—The virtual switch monitors network traffic, and attempts to make adjustments on overloaded uplinks by moving traffic to another uplink. This option does use additional resources to track network traffic.

VxRail does not support the 'Route based on IP Hash' policy, as there is a dependency on the logical link setting of the physical port adapters on the switch. Link aggregation is required for this setting, which is introduced in version 7.0.130 of VxRail.

Starting with VxRail version 7.0.010, the 'Failover Order' setting on the teaming and failover policy on the VDS portgroups supporting VxRail networks can be changed. The default failover order for the uplinks each portgroup configured during VxRail initial build is described in [Default failover order policy](#). For any portgroup configured during VxRail initial build to support VxRail network traffic, an uplink in 'Standby' mode can be moved

into 'Active' mode to enable an 'active/active' configuration. This action can be performed after the VxRail cluster has completed the initial build operation.

Moving an uplink that is configured as 'Unused' for a portgroup supporting VxRail network traffic into either 'Active' mode or 'Standby' mode does not automatically activate the uplink and increase bandwidth for that portgroup. Bandwidth optimization depends on the load-balancing settings on the upstream switches, and link aggregation is required to optimize network performance.

Virtual SAN-Cluster3 - Edit Settings

General		
Advanced	Load balancing	Route based on physical NIC load
VLAN	Network failure detection	Link status only
Security	Notify switches	Yes
Teaming and failover	Failback	Yes
Traffic shaping		
Monitoring		
Miscellaneous		

Failover order ⓘ

↑ ↓

Active uplinks
Uplink 2
Uplink 1
Standby uplinks

Figure 45. Sample failover order setting set to active/active

Using unassigned physical ports for VxRail networks

For VxRail versions prior to 7.0.010, VxRail nodes ordered with extra physical network ports can be configured for non-VxRail system traffic. For instance, if you are planning to enable certain vSphere features such as fault tolerance, then plan to need additional node ports for this purpose.

Starting with version 7.0.010 of VxRail, you can configure the ports on the optional PCIe adapter cards to support VxRail network traffic. Unless you are deploying the VxRail cluster to a customer-supplied virtual distributed switch, this is only supported as a post-deployment activity.




Figure 46. VxRail node with NDC ports and ports from optional PCIe adapter card

Network redundancy across NDC and PCIe Ethernet ports can be enabled by reconfiguring the VxRail networks and migrating selected VxRail network traffic from the original NDC-based ports over to PCIe-based ports. The following table describes the supported starting and ending network reconfigurations.

Starting configuration	Ending configuration
2 NDC ports	1 NDC port & 1 PCIe port
2 NDC Ports	2 NDC ports & 2 PCIe ports
4 NDC Ports	2 NDC ports & 2 PCIe ports
4 NDC Ports	1 NDC port & 1 PCIe port

The following rules apply for migrating VxRail networks from NDC-only ports to mixed NDC-PCIe ports:

- The VxRail version on your cluster is 7.0.010 or later.
- The first port configured for VxRail networking, commonly known as 'vmnic0' or 'vmnic1', must be reserved for VxRail management and node discovery. Do not migrate VxRail management or node discovery off of this first reserved port.
- The switch ports enabling connectivity to the PCIe-based ports are properly configured to support VxRail network traffic.
- All of the network ports supporting VxRail network traffic must be running the same speed.
- The network reconfiguration requires a one-to-one swap. For example, a VxRail network that is currently running on two NDC ports can be reconfigured to run on one NDC port and one PCIe port.

Note: You must follow the official instructions/procedures from VMware and Dell Technologies for these operations.

The **supported** operations include:

- Create a new vSphere Standard Switch (VSS), and connect unused ports to the VSS.
- Connect unused ports to new port groups on the default vSphere Distributed Switch.
- Create a new vSphere Distributed Switch (VDS), add VxRail nodes to the new VDS, and connect their unused network ports to the VDS.
- Create new VMKernel adapters, and enable services of IP Storage and vSphere Replication.
- Create new VM Networks, and assign them to new port groups.

The following operations are unsupported in versions earlier than VxRail 7.0.010:

- Migrating or moving VxRail system traffic to the optional ports. VxRail system traffic includes the management, vSAN, vCenter Server, and vMotion Networks.
- Migrating VxRail system traffic to other port groups.

- Migrating VxRail system traffic to another VDS.

Note: Performing any unsupported operations will impact the stability and operations of the VxRail cluster, and may cause a failure in the VxRail cluster.

Configure link aggregation on VxRail networks

Starting with VxRail version 7.0.130, NIC teaming can be configured on the VxRail non-management networks with a both a customer-supplied and VxRail-supplied virtual distributed switch after the cluster initial implementation process is complete. NIC teaming enables the formation of a link aggregation group, which is a logical port that represents a pair of physical ports on a VxRail node. If the ports on the top-of-rack switches connected to these logical ports are also configured into a link aggregation group, peering between the two link aggregation groups will enable an active-active port configuration and support load balancing for network optimization.




Figure 47. Enabling link aggregation for load balancing

The following rules are applicable for enabling link aggregation:

- A minimum of 4 Ethernet ports per VxRail node is required for the VxRail networks.
 - 2 ports configured to support VxRail management networks (Management/vCenter/VxRail discovery)
 - 2 ports configured to support vSAN and vMotion.
- Can be all NDC ports, all PCIe ports, or a mixture of NDC and PCIe ports
 - If all of your VxRail networks are on NDC-based ports, you can migrate the non-management VxRail networks to PCIe-based ports to enable NIC redundancy.
- All VxRail node ports configured for link aggregation must be the same speed.
- Link aggregation can only be configured on the non-management VxRail networks.
 - If you want to enable load balancing only on the vSAN network, you have the option to migrate the vMotion network over to the ports supporting the VxRail management networks

Depending on your data center networking policies and features supported on the top-of-rack switches, either a static LAG or dynamic LAG can be configured for load balancing.

A static LAG requires:

- Static port channels to be configured on the ports on the adjacent top-of-rack switches
- The teaming policy on the portgroup is set to 'Route based on IP Hash'. This will calculate a hash using the source and destination IP address of each packet.

A dynamic LAG requires:

- Dynamic port channels are configured on the adjacent top-of-rack switches
- LACP support on the adjacent top-of-rack switches. If LACP is not supported on the switches, a static LAG is the only option.
- An LACP policy is configured on the virtual distributed switch.
 - The load-balancing setting on the LACP policy is compatible with the supported load-balancing hash algorithms on the adjacent top-of-rack switches.

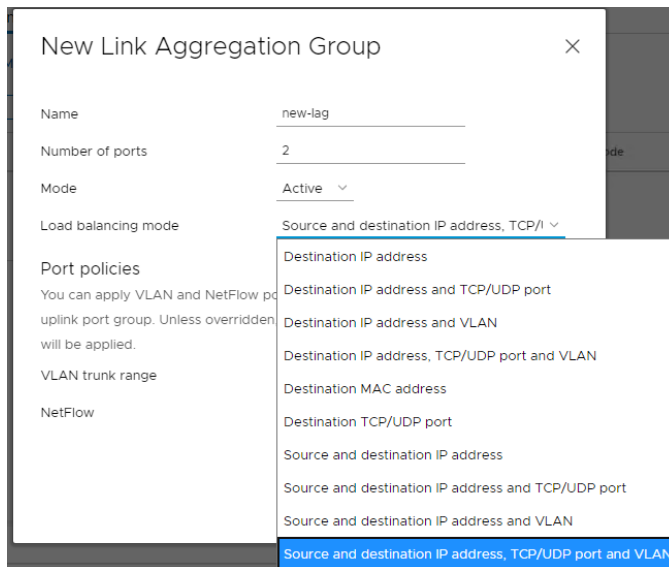


Figure 48. LACP Policy configuration on virtual distributed switch

Deploy second VDS for VxRail networking

The default VxRail configuration is a single virtual distributed switch for all VxRail networks. Support for an additional virtual distributed switch on the VxRail cluster is to address the use case of enabling network segmentation to comply with a company or organization security policy.

If your VxRail cluster is running version 7.0.130 or later, and the cluster was implemented with a single virtual distributed switch for all VxRail networks, you can deploy a second virtual distributed switch and migrate selected VxRail networks to the new switch.




Figure 49. Deploying second VDS to support network segmentation

You can also migrate uplinks from the original virtual distributed switch to the new virtual distributed switch. With support introduced for both NDC-based and PCIe-based ports, you can reconfigure the network at the virtual switch level and Ethernet port level to isolate management and non-management VxRail network traffic.

The following rules apply to deploying a second virtual distributed switch post-implementation:

- If your cluster was built with a version prior to 7.0.130, upgrade the cluster to version 7.0.130 or later.
- Both customer-supplied virtual distributed switches and VxRail-supplied virtual distributed switches are supported.
- Either 2 uplinks or 4 uplinks can be assigned to each virtual distributed switch.
- The number of uplinks assigned to each virtual distributed switch do not need to match. For example, you can assign 4 uplinks to one virtual distributed switch and 2 uplinks to the second virtual distributed switch.
- You can expand the number of uplinks reserved for VxRail networking. For example, if your cluster was originally built with 2 uplinks reserved for VxRail networking on one virtual distributed switch, and your nodes have unused Ethernet ports, you can reconfigure the virtual network with 2 uplinks on one virtual distributed switch and 2 uplinks on the second virtual distributed switch.

Appendices

This section presents the following topics:

Appendix A: VxRail Network Configuration Table.....	96
Appendix B: VxRail Passwords	99
Appendix C: VxRail Setup Checklist.....	100
Appendix D: VxRail Open Ports Requirements.....	102
Appendix E: Virtual Distributed Switch Portgroup Default Settings	104
Appendix F: Physical Network Switch Examples	107

Appendix A: VxRail Network Configuration Table

The Dell Technologies service representative uses a data collection workbook to capture the settings that are needed to build the VxRail cluster. The workbook includes the following information:

Row	Topic	Category	Description
1	VxRail Networks	External Management	Untagged traffic is recommended 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.
2		Internal Management	This network traffic should stay isolated on the top-of-rack switches. The default VLAN ID is 3939.
3		vMotion	
4		vSAN	
5		Guest Network(s)	Network Name
6			VLAN
7	VxRail Management	Subnet Mask	Subnet mask for VxRail External Management Network
8		Default Gateway	Default gateway for VxRail External Management Network
9	System	Global Settings	Time zone
10			NTP servers
11			DNS servers
12			Top Level Domain
13	VxRail Manager	Hostname	
14		IP Address	
15	ESXi Hostnames	VxRail auto-assign method	Prefix
16			Separator
17			Iterator
18			Offset
19			Suffix
20		Customer-supplied method	ESXi hostname 1
21			ESXi hostname 2
22			ESXi hostname 3
23			ESXi hostname 4
24			Starting IP Address

Row	Topic	Category	Description
25	ESXi IP Addresses	VxRail auto-assign method	Ending IP Address
26		Customer-supplied method	ESXi IP Address 1
27			ESXi IP Address 2
28			ESXi IP Address 3
29			ESXi IP Address 4
30	vCenter	VxRail vCenter Server	vCenter Server Hostname
31			vCenter Server IP Address
32			Platform Services Controller Hostname (if applicable)
33			Platform Services Controller IP address (if applicable)
34		Customer-supplied vCenter Server	Platform Services Controller Hostname (FQDN) (Leave blank if PSC is embedded in customer-supplied vCenter Server)
35			vCenter Server Hostname (FQDN)
36			vCenter Server SSO Domain
37			Admin username/password or the newly created VxRail non-admin username and password
38			New VxRail management username and password
39			vCenter Data Center Name
40			vCenter Cluster Name
41	Virtual Distributed Switch	Customer-supplied VDS Portgroups	Name of VDS portgroup supporting VxRail external management network
42			Name of VDS portgroup supporting VxRail vCenter Server network
43			Name of VDS portgroup supporting VxRail internal management network
44			Name of VDS portgroup supporting VxRail vMotion network
45			Name of VDS portgroup supporting VxRail vSAN network
46	vMotion	VxRail auto-assign method	Starting address for IP pool
47			Ending address for IP pool
48		Customer-supplied method	vMotion IP Address 1
49			vMotion IP Address 2
50			vMotion IP Address 3
51			vMotion IP Address 4
52		Subnet Mask	
53		Gateway	Gateway (Default or vMotion)
54	vSAN		Starting address for IP pool

Row	Topic	Category	Description
55		VxRail auto-assign method	Ending address for IP pool
56		Customer-supplied method	vSAN IP Address 1
57			vSAN IP Address 2
58			vSAN IP Address 3
59			vSAN IP Address 4
60		Subnet Mask	
61	Logging	VxRail Internal	vRealize Log Insight™ hostname
62			vRealize Log Insight IP address
63		VxRail External	Syslog Server (instead of Log Insight)
64	SmartFabric	Switch out-of-band management	Out-of-band management IP address for switch 1
65			Out-of-band management IP address for switch 2
66		Dell EMC OMNI plug-in	IP address
67			Subnet Mask
68			Gateway
69	Witness Site	Management IP Address	Witness management network IP address
70		vSAN IP Address	Witness vSAN network IP address
71	Witness Traffic Separation	WTS VLAN	Optional to enable Witness traffic separation on stretched-cluster or 2-Node Cluster
72	2-Node Cluster	Node 1 WTS IP address	Must be routable to Witness
73		Node 2 WTS IP address	Must be routable to Witness

Appendix B: VxRail Passwords

Item	Account	Password
VxRail Manager	Root	
VxRail vCenter Server	Administrator@<SSO Domain>	
	Root	
	Management	
VxRail Platform Service Controller	Root	
vRealize Log Insight	Root	
	Admin	

Item	Account	Password
ESXi Host #1	Root	
ESXi Host #2	Root	
ESXi Host #3	Root	
ESXi Host #4	Root	

Appendix C: VxRail Setup Checklist

Physical Network	
<p>VxRail cluster: Decide if you want to plan for additional nodes beyond the initial three (or four)-node cluster. You can have up to 64 nodes in a VxRail cluster.</p> <p>VxRail ports: Decide how many ports to configure per VxRail node, what port type, and what network speed.</p> <p>Network switch: Ensure that your switch supports VxRail requirements and provides the connectivity option that you chose for your VxRail nodes. Verify cable requirements. Decide if you will have a single or multiple switch setup for redundancy.</p> <p>Data center: Verify that the required external applications for VxRail are accessible over the network and correctly configured.</p> <p>Topology: If you are deploying VxRail over more than one rack, be sure that network connectivity is set up between the racks. Determine the Layer 2/Layer 3 boundary in the planned network topology.</p> <p>Workstation/laptop: Any operating system with a browser to access the VxRail user interface. The latest versions of Firefox, Chrome, and Internet Explorer 10+ are all supported.</p> <p>Out-of-band Management (optional): One available port that supports 1 Gb for each VxRail node.</p>	
Logical Network	
Reserve VLANs	<ul style="list-style-type: none"> ✓ One external management VLAN for traffic from VxRail, vCenter Server, ESXi ✓ One internal management VLAN with multicast for auto-discovery and device management. The default is 3939. ✓ One VLAN with unicast (starting with VxRail v4.5.0) or multicast (prior to v4.5.0) for vSAN traffic ✓ One VLAN for vSphere vMotion ✓ One or more VLANs for your VM Guest Networks ✓ If you are enabling witness traffic separation, reserve one VLAN for the VxRail witness traffic separation network.
System	<ul style="list-style-type: none"> ✓ Select the Time zone. ✓ Select the Top-Level Domain. ✓ Hostname or IP address of the NTP servers on your network (recommended) ✓ IP address of the DNS servers on your network (if external DNS) ✓ Forward and reverse DNS records for VxRail management components (if external DNS).
Management	<ul style="list-style-type: none"> ✓ Decide on your VxRail host naming scheme. The naming scheme will be applied to all VxRail management components. ✓ Reserve three or more IP addresses for ESXi hosts. ✓ Reserve one IP address for VxRail Manager. ✓ Determine default gateway and subnet mask. ✓ Select passwords for VxRail management components.
vCenter	<ul style="list-style-type: none"> ✓ Determine whether you will use a vCenter Server that is customer-supplied or new to your VxRail cluster. ✓ VxRail vCenter Server: Reserve IP addresses for vCenter Server and PSC (if applicable). ✓ Customer-supplied vCenter Server: Determine hostname and IP address for vCenter and PSC, administration user, and name of vSphere data center. Create a VxRail management user in vCenter. Select a unique VxRail cluster name. (Optional) Create a VxRail non-admin user.

Physical Network	
Virtual Distributed Switch	<ul style="list-style-type: none"> ✓ Determine whether you will preconfigure a customer-supplied virtual distributed switch or have VxRail deploy a virtual distributed switch in your vCenter instance. ✓ Customer-supplied Virtual Distributed Switch: Configure target portgroups for required VxRail networks.
vMotion	<ul style="list-style-type: none"> ✓ Decide whether you want to use the default TCP-IP stack for vMotion, or a separate IP addressing scheme for the dedicated vMotion TCP-IP stack. ✓ Reserve three or more contiguous IP addresses and a subnet mask for vSphere vMotion. ✓ Select the gateway for either the default TCP-IP stack, or the dedicated vMotion TCP-IP stack.
vSAN	<ul style="list-style-type: none"> ✓ Reserve three or more contiguous IP addresses and a subnet mask for vSAN
Solutions	<ul style="list-style-type: none"> ✓ To use vRealize Log Insight: Reserve one IP address. ✓ To use an existing syslog server: Get the hostname or IP address of your third-party syslog server.
Witness Site	<ul style="list-style-type: none"> ✓ If Witness is required, reserve one IP address for the management network and one IP address for the vSAN network.
Workstation	<ul style="list-style-type: none"> ✓ Configure your workstation/laptop to reach the VxRail initial IP address. ✓ Ensure you know how to configure the laptop to reach the VxRail Manager IP address after configuration.
Set up Switch	<ul style="list-style-type: none"> ✓ Configure your selected external management VLAN (default is untagged/native). ✓ Configure your internal management VLAN. ✓ Confirm multicast is enabled for device discovery. ✓ Configure your selected VLANs for vSAN, vSphere vMotion, and VM Guest Networks. ✓ If applicable, configure your Witness traffic separation VLAN. ✓ In dual-switch environments, configure the inter-switch links to carry traffic between switches. ✓ Configure uplinks to carry upstream network VLANs. ✓ Configure one port as an access port for laptop/workstation to connect to VxRail Manager for initial configuration. ✓ Confirm configuration and network access.
Workstation/Laptop	<ul style="list-style-type: none"> ✓ Configure your workstation/laptop to reach the VxRail Manager initial IP address. ✓ Configure the laptop to reach the VxRail Manager IP address after permanent IP address assignment.

Appendix D: VxRail Open Ports Requirements

Use the tables in this Appendix for guidance on firewall settings specific for the deployment of a VxRail cluster. Then use the links that are provided after the tables for firewall rules that are driven by product feature and use case.

The VxRail cluster needs to be able to connect to specific applications in your data center. DNS is required, and NTP is optional. Open the necessary ports to enable connectivity to the external syslog server, and for LDAP and SMTP.

Datacenter Application Access				
Description	Source Devices	Destination Devices	Protocol	Ports
DNS	VxRail Manager, Dell iDRAC	DNS Servers	UDP	53
NTP Client	Host ESXi Management Interface, Dell iDRAC, VMware vCenter Servers, VxRail Manager	NTP Servers	UDP	123
SYSLOG	Host ESXi Management Interface, vRealize Log Insight	Syslog Server	TCP	514
LDAP	VMware vCenter Servers, PSC	LDAP Server	TCP	389, 636
SMTP	SRS Gateway VMs, vRealize Log Insight	SMTP Servers	TCP	25

Open the necessary firewall ports to enable IT administrators to deploy the VxRail cluster.

Administration Access				
Description	Source Devices	Destination Devices	Protocol	Ports
ESXi Management	Administrators	Host ESXi Management Interface	TCP, UDP	902
VxRail Management GUI/Web Interfaces	Administrators	VMware vCenter Server, VxRail Manager, Host ESXi Management, Dell iDRAC port, vRealize Log Insight, PSC	TCP	80, 443
Dell server management	Administrators	Dell iDRAC	TCP	623, 5900, 5901
SSH and SCP	Administrators	Host ESXi Management, vCenter Server Appliance, Dell iDRAC port, VxRail Manager Console	TCP	22

If you plan to use a customer-supplied vCenter server instead of deploying a vCenter server in the VxRail cluster, open the necessary ports so that the vCenter instance can connect to the ESXi hosts.

vCenter and vSphere				
Description	Source Devices	Destination Devices	Protocol	Ports
vSphere Clients to vCenter Server	vSphere Clients	vCenter Server	TCP	5480, 8443, 9443, 10080, 10443
Managed Hosts to vCenter	Host ESXi Management	vCenter Server	TCP	443, 902, 5988, 5989, 6500, 8000, 8001
Managed Hosts to vCenter Heartbeat	Host ESXi Management	vCenter Server	UDP	902

Other firewall port settings may be necessary depending on your data center environment. The list of documents in this table is provided for reference purposes.

Description	Reference
VMware Ports and Protocols	VMware Ports and Protocols
Network port diagram for vSphere 6	Network Port Diagram for vSphere 6
vSAN Ports Requirements	vSAN Network Ports Requirements
Dell iDRAC Port Requirements	How to configure the iDRAC 9 for Dell PowerEdge
Secure Remote Services Port Requirements	Dell EMC Secure Remote Services Documentation

Appendix E: Virtual Distributed Switch Portgroup Default Settings

Unless you configure an external distributed virtual switch in your external vCenter for the VxRail cluster, the VxRail initial build process will configure a virtual distributed switch on the selected vCenter instance using best practices for VxRail.

Default standard settings

For each VxRail network portgroup, the initial build process will apply the following standard settings.

Setting	Value
Port Binding	Static
Port Allocation	Elastic
Number of ports	8
Network Resource Pool	(default)
Override port policies	Only 'Block ports' allowed
VLAN Type	VLAN
Promiscuous mode	Reject
MAC address changes	Reject
Forged transmits	Reject
Ingress traffic shaping	Disabled
Egress traffic shaping	Disabled
NetFlow	Disabled
Block All Ports	No

Default teaming and failover policy

VxRail will configure a teaming and failover policy for the port groups on the virtual distributed switch with the following settings:

Setting	Value
Load Balancing	Route based on originating virtual port
Network failure detection	Link status only
Notify switches	Yes
Failback	Yes

Default network I-O control (NIOC)

VxRail will enable network I-O control on the distributed switch, and configure custom Network I-O Control (NIOC) settings for the following network traffic types. The settings are dependent on whether the VxRail cluster was deployed with either 2 Ethernet ports per node reserved for the VxRail cluster, or if 4 Ethernet ports were reserved for the VxRail cluster:

Traffic Type	NIOC Shares	
	4 Ports	2 Ports
Management Traffic	40	20
vMotion Traffic	50	50
vSAN Traffic	100	100
Virtual Machine Traffic	60	30

The reservation value is set to zero for all network traffic types, with no limits set on bandwidth.

Default failover order policy

VxRail will configure an active/standby policy to the uplinks VxRail uses for the four pre-defined network traffic types that are required for operation: Management, vMotion, vSAN, and Virtual Machine.

4x10GbE Traffic Configuration

Traffic Type	Uplink1 VMNIC0	Uplink2 VMNIC1	Uplink3 VMNIC2	Uplink4 VMNIC3
Management	Standby	Active	Unused	Unused
vSphere vMotion	Unused	Unused	Standby	Active
vSAN	Unused	Unused	Active	Standby
vCenter Server Network	Active	Standby	Unused	Unused
VxRail Management	Standby	Active	Unused	Unused

2x10GbE or 2x25GbE Traffic Configuration

Traffic Type	Uplink1 VMNIC0	Uplink2 VMNIC1	Uplink3 No VMNIC	Uplink4 No VMNIC
Management	Active	Standby	Unused	Unused
vSphere vMotion	Active	Standby	Unused	Unused
vSAN	Standby	Active	Unused	Unused
vCenter Server Network	Active	Standby	Unused	Unused
VxRail Management	Active	Standby	Unused	Unused

4x25GbE Traffic Configuration

Traffic Type	Uplink1 VMNIC0	Uplink2 VMNIC1	Uplink3 VMNIC2	Uplink4 VMNIC3
Management	Active	Unused	Standby	Unused
vSphere vMotion	Unused	Standby	Unused	Active
vSAN	Unused	Active	Unused	Standby
vCenter Server Network	Standby	Active	Active	Unused
VxRail Management	Active	Unused	Standby	Unused

Appendix F: Physical Network Switch Examples

VxRail enforces a pre-defined network profile during initial implementation depending on the number of ports selected for VxRail networking, and the type of network ports. Starting with version 7.0.130, you can choose between a pre-defined network profile or choose to customize the network topology.

These diagrams show different options for physical wiring between VxRail nodes and the adjacent, top-of-rack switches, depending on your port selections. They are provided as illustrative examples to help with the planning and design process. All VxRail nodes are manufactured with Ethernet ports built into the NDC. Optional PCIe adapter cards can be installed in the VxRail nodes to provide additional Ethernet ports for redundancy and increased bandwidth.

If additional Ethernet connectivity is required to support other use cases outside of VxRail networking, then additional slots on the VxRail nodes must be reserved for PCIe adapter cards. If this is a current requirement or potential future requirement, then be sure to select a VxRail node model with sufficient PCIe slots to accommodate the additional adapter cards.

Pre-defined network profile: 2x10Gb or 2x25Gb

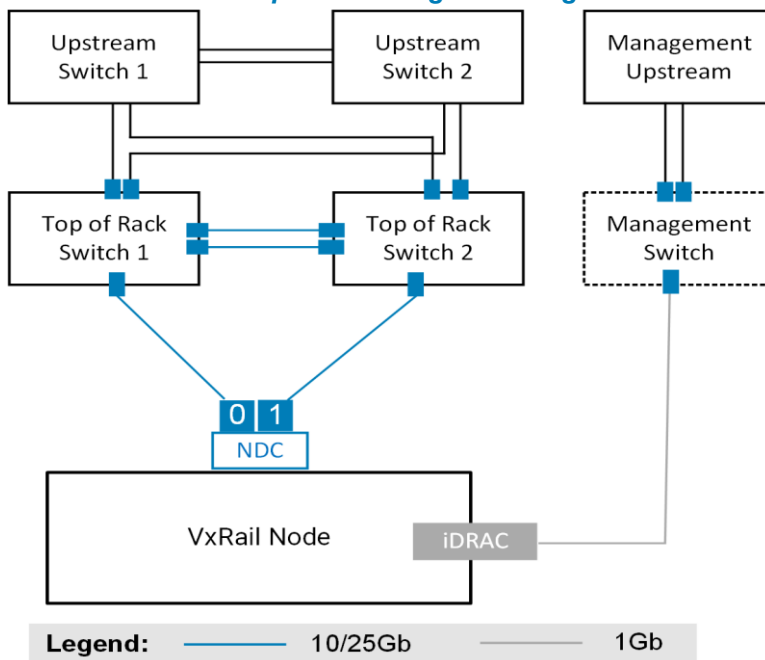


Figure 50. VxRail nodes with two 10Gb NDC ports connected to two TOR switches, and one optional connection to management switch for iDRAC

VxRail selects the two ports on the NDC to support VxRail networking. If the NDC on the VxRail nodes is shipped with four Ethernet ports, then the two leftmost ports are selected. If you choose to use only two Ethernet ports, the remaining ports can be used for other use cases. This connectivity option is the simplest to deploy. It is suitable for smaller, less demanding workloads that can tolerate the NDC as a potential single point of failure.

Pre-defined network profile: 4x10gb or 4x1gb NDC

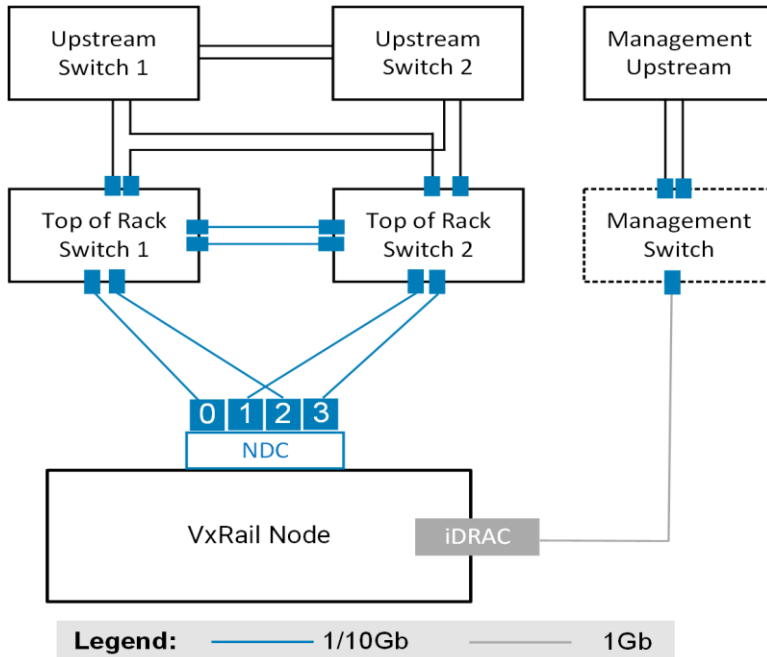


Figure 51. VxRail nodes with four 10gb NDC ports connected to 2 TOR switches, and one optional connection to management switch for iDRAC

VxRail selects all four ports on the NDC to support VxRail networking. If you are deploying VxRail with 1gb Ethernet ports, then you must connect four Ethernet ports to support VxRail networking.

Pre-defined network profile: 2x10gb NDC & 2x10gb PCIe

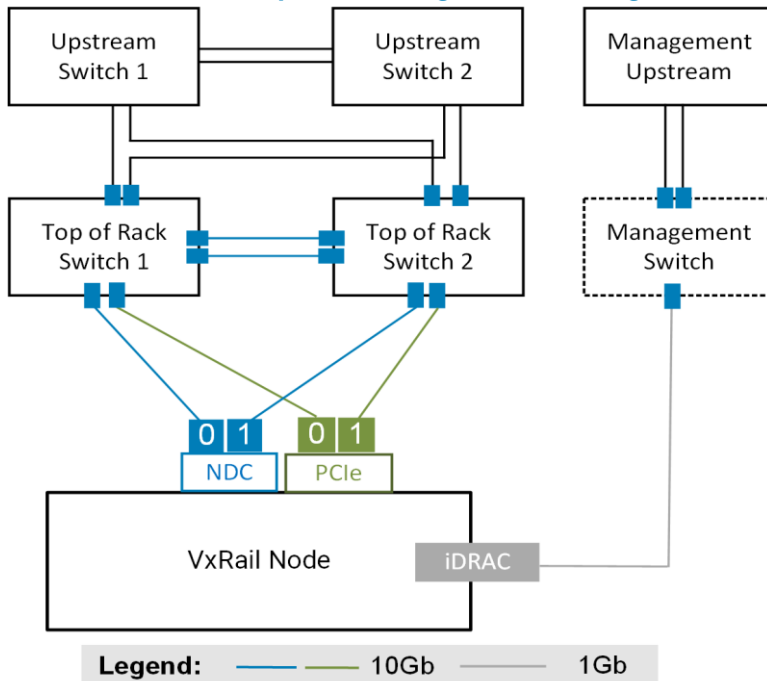


Figure 52. VxRail nodes with two 10gb NDC ports and two 10gb PCIe ports connected to 2 TOR switches, and one optional connection to management switch for iDRAC

In this option, two NDC ports and two ports on the PCIe card in the first slot are selected for VxRail networking. The network profile splits the VxRail networking workload between the NDC ports and the two switches, and splits the workload on the PCIe-based ports between the two switches. This option ensures against the loss of service with a failure at the switch level, but also with a failure in either the NDC or PCIe adapter card.

Pre-defined network profile: 2x25Gb NDC and 2x25Gb PCIe

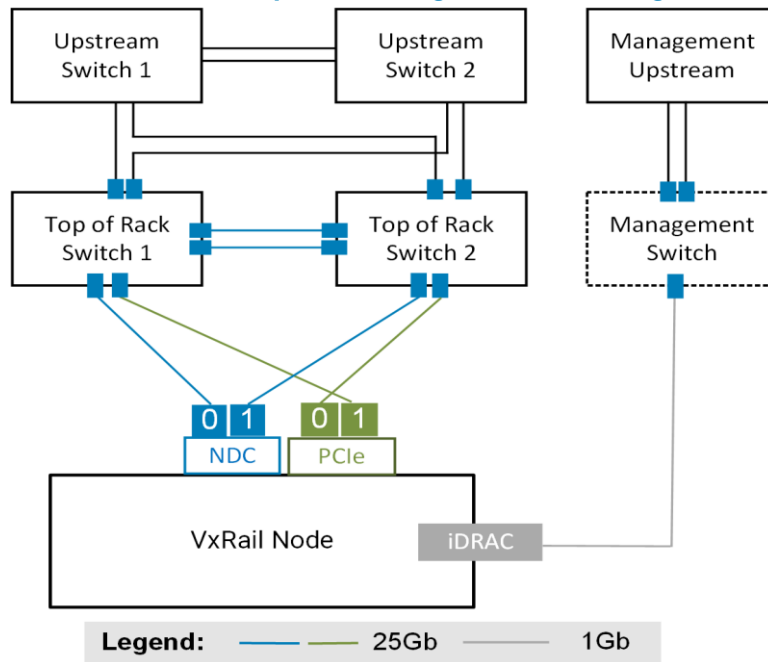


Figure 53. VxRail nodes with two 25Gb NDC ports and two 25Gb PCIe ports connected to 2 TOR switches, and one optional connection to management switch for iDRAC

In this option, two NDC ports and two ports on the PCIe card in the first slot are selected for VxRail networking. This option offers the same benefits as the 2x10Gb NDC and 2x10Gb PCIe deployment option, except for additional bandwidth available to support the workload on the VxRail cluster.

Be aware that the cabling for the 25Gb option with NDC ports and PCIe ports differs from the 10Gb option. Note that the second port on the PCIe adapter cards is paired with the first port on the NDC on the first switch, and the first port on the PCIe adapter is paired with the second port on the NDC on the second switch. This is to ensure balancing of the VxRail networks between the switches in the event of a failure at the network port layer.

Custom option: Any NDC ports paired with PCIe ports

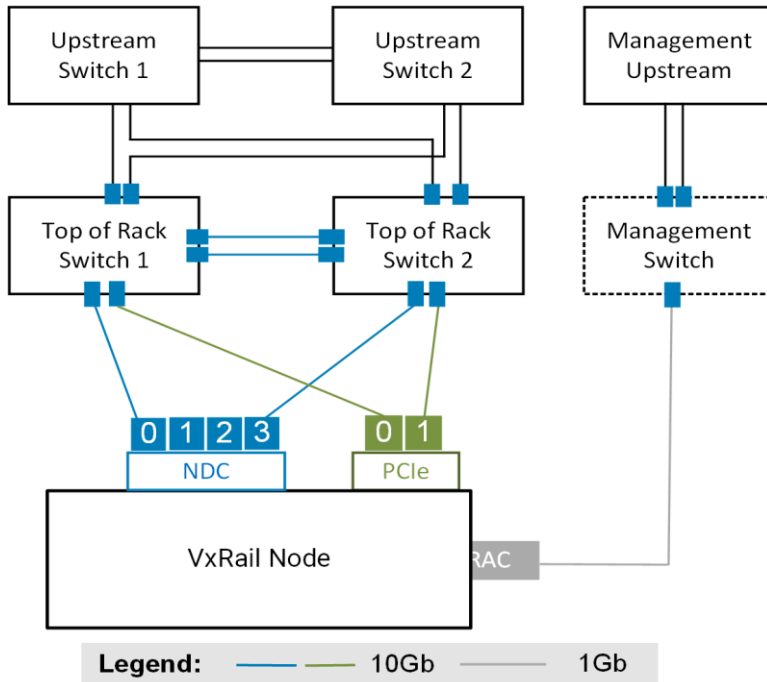


Figure 54. VxRail nodes with any two 10gb NDC ports and two 10gb PCIe ports connected to two TOR switch, and one optional connection to management switch for iDRAC

This is an example of an optional cabling setup for 2 NDC ports and 2 PCIe ports. Any NDC port and any PCIe port can be selected so long as the ports are of the same type and are running at the same speed.

Custom option: Two NDC ports paired with PCIe ports other than the first slot

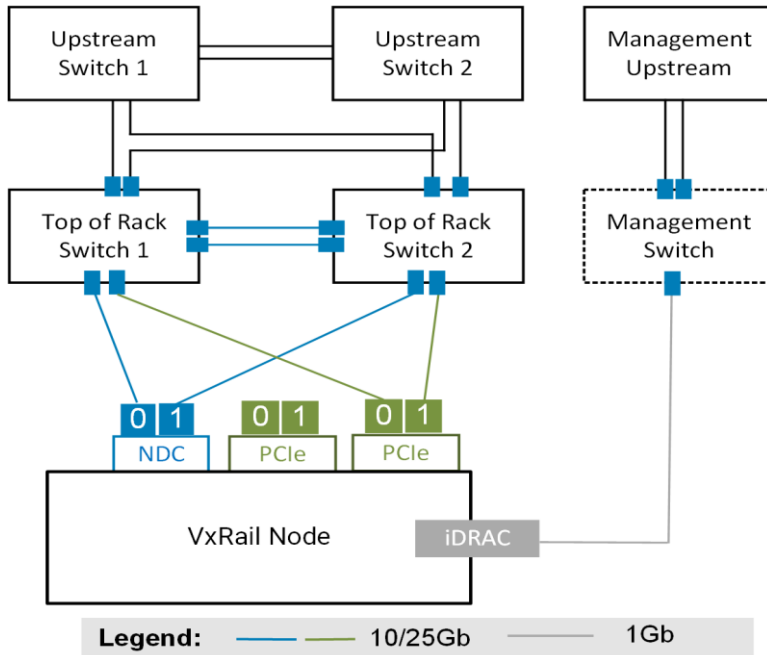


Figure 55. VxRail nodes with any two 10/25gb NDC ports and any two 10/25gb PCIe ports connected to 2 TOR switches, and one optional connection to management switch for iDRAC

With the custom option, there is no restriction that the ports selected for VxRail networking reside on the PCIe adapter card in the first slot

Custom option: PCIe ports only

In this outlier use case where there is a specific business or operational requirement, VxRail can be deployed using only the ports on PCIe adapter cards, so long as the ports are of the same type and are running at the same speed.

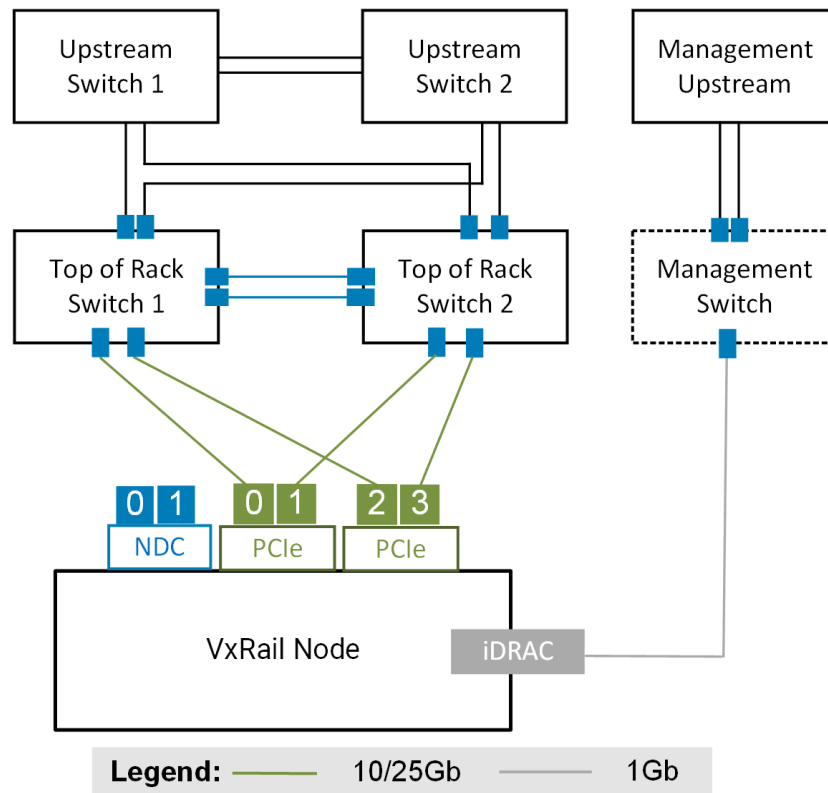


Figure 56. VxRail nodes with two or four PCIe ports connected to a pair of TOR switch, and one optional connection to management switch for iDRAC

This option supports spreading the VxRail networking across ports on more than one PCIe adapter card.

Four TOR switches to support VxRail cluster networking

For workload use cases with extreme availability, scalability and performance requirements, four TOR switches can be positioned to support VxRail networking. In this example, each Ethernet port is connected to a single TOR switch. Each pair of top-of-rack switches is logically connected using inter-switch links.

This topology also addresses the use case of physical network separation to meet specific security policies or governance requirements. For instance, the networks required for VxRail management and operations can be isolated on one pair of switches, while network traffic for guest user and application access can be targeted on the other pair of switches.

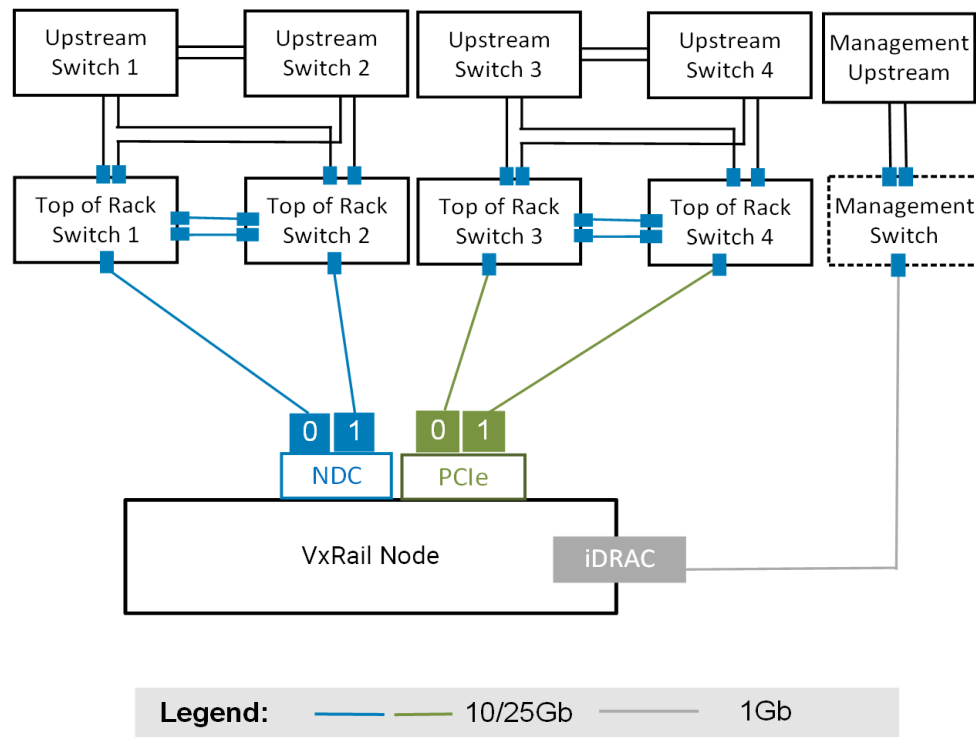


Figure 57. VxRail nodes with four ports connected to four TOR switches, and one optional connection to management switch for iDRAC