

Dell VxRail Network Planning Guide

Physical and Logical Network Considerations and Planning

March 2022

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 © 2022 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 03/22 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	7
	Revision history	8
	Intended use and audience	9
	Introduction to VxRail	9
Chapter 2	Planning Your Data Center Network for VxRail	11
	Introduction.....	12
	VxRail hardware and the physical network infrastructure.....	12
	VxRail clusters, appliances, and nodes.....	13
	Network switch	14
	VxRail relationship with the Ethernet switch.....	14
	VxRail node discovery and the Ethernet switch	14
	Basic switch requirements.....	15
	Switch performance considerations.....	15
	Storage network considerations	16
	Network redundancy and performance considerations	16
	Data center network	17
	Data center services.....	17
	Routing services	18
	VxRail networking port options	19
	VxRail Ethernet adapter options.....	20
	VxRail Fibre Channel adapter options.....	21
	VxRail RoCE adapter options.....	22
	VxRail node connectivity options.....	22
	VxRail networking rules and restrictions.....	23
	Topology and connections.....	24
Chapter 3	VxRail Cluster Types	26
	Introduction.....	27
	Cluster with vSAN storage.....	27
	Dynamic cluster	28
	Fibre Channel Storage Option.....	29
	Remote vSAN Datastore Option.....	29
	PowerFlex storage option.....	30
	iSCSI and NFS Options.....	32
	vSAN stretched cluster	32

2-node cluster	33
Satellite nodes	35
Chapter 4 VxRail Feature-Driven Decision Points	37
Introduction	38
Software-defined data center	38
Dell SmartFabric network mode	39
vSphere with Kubernetes on VxRail	41
Chapter 5 VxRail Hardware and Switch Selection Decision Points	43
Step-by-step networking decision points	44
Chapter 6 Prepare data center for VxRail Implementation	46
Introduction	47
Prepare data center routing services	47
Prepare for multi-rack VxRail cluster	48
Prepare for vSAN HCI mesh topology	50
Prepare external FC storage for dynamic clusters	52
Prepare for VxRail custom uplink assignments	53
Prepare data center network MTU	55
Prepare for link aggregation of VxRail networks	55
Verify the switches support link aggregation	57
Verify support for multi-chassis link aggregation	57
Identify switch ports to be configured for link aggregation	57
Plan link aggregation on switch port pairs	58
Prepare certificate authority server for VxRail	59
Prepare VxRail remote support connectivity	59
Chapter 7 Planning the VxRail Cluster Implementation	61
Introduction	62
Decide on VxRail single point of management	62
Decide on VxRail network traffic segmentation	63
Decide on teaming and failover policies for VxRail networks	65
Plan the VxRail logical network	66
IP address considerations for VxRail networks	67
Virtual LAN considerations for VxRail networks	69
Plan network settings for VxRail management components	72
Plan network settings for vCenter Server management network	74
Identify IP addresses for VxRail management components	74
Select hostnames for VxRail management components	75
Select top-level domain	76
Select VxRail Manager hostname	76

Select ESXi hostnames	76
Select VxRail vCenter Server hostname	77
Identify external applications and settings for VxRail	77
Set time zone and NTP server	77
Set DNS for VxRail management components	78
Prepare customer-supplied vCenter server.....	78
Prepare customer-supplied virtual-distributed switch	80
Prepare link aggregation on customer-supplied virtual-distributed switch.....	82
Reserve IP addresses for VxRail vMotion network	83
Reserve IP addresses for VxRail vSAN network.....	84
Decide on VxRail logging solution	85
Assign passwords for VxRail management.....	86
Prepare for Dell SmartFabric Services enablement	87
Chapter 8 Planning for VxRail Satellite Nodes	89
Introduction.....	89
Plan networking to support VxRail satellite node management	90
Assign network settings to VxRail satellite nodes	91
Assign passwords for VxRail satellite nodes	91
Chapter 9 Configure the Network for VxRail	93
Introduction.....	94
Setting up the network switch for VxRail connectivity	94
Configure multicast for VxRail Internal Management network.....	94
Configure unicast for VxRail vSAN network	94
Configure VLANs for the VxRail networks.....	95
Configure the inter-switch links.....	97
Configure switch ports	97
Determine switch port mode.....	97
Disable link aggregation on switch ports supporting VxRail networks.....	97
Limit spanning tree protocol on VxRail switch ports	98
Enable flow control	99
Configure ports on your switches	99
Setting up the upstream network for VxRail connectivity	99
Configure network to support RoCE.....	101
Confirm your data center network	101
Confirm your firewall settings	102
Confirm your data center environment.....	102
Chapter 10 Preparing to Build the VxRail Cluster	104
Introduction.....	105

Complete pre-requisites for dynamic clusters	105
Configure nodes for tagged VxRail management VLAN	105
Configure a jump host or laptop for VxRail initialization	105
Perform initialization to create a VxRail cluster	107
Chapter 11 VxRail Network Considerations After Implementation	109
Introduction.....	110
Using unassigned physical ports for VxRail networks	110
Configure link aggregation on VxRail networks.....	111
Appendixes	114
Appendix A: VxRail Network Configuration Table	115
Appendix B: VxRail Passwords	119
Appendix C: VxRail Cluster Setup Checklist	120
Appendix D: VxRail Open Ports Requirements	122
Appendix E: Virtual Distributed Switch Portgroup Default Settings	124
Default standard settings.....	124
Default teaming and failover policy.....	124
Default network I-O control (NIOC)	125
Default failover order policy	125
Appendix F: Physical Network Switch Examples	127

Chapter 1 Introduction

This chapter presents the following topics:

Revision history	8
Intended use and audience	9
Introduction to VxRail	9

Revision history

Date	Description
April 2019	First inclusion of this version history table <ul style="list-style-type: none"> 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 multi-rack switch network Optional 100 GbE Ethernet and FC network ports on VxRail nodes
May 2020	Updated switch requirements for VxRail IPv6 multicast.
June 2020	Updated networking requirements for multirack VxRail clusters.
July 2020	Support for new features in VxRail 7.0.010
August 2020	Updated requirement for NIC redundancy enablement.
September 2020	Outlined best practices for link aggregation on non-VxRail ports.
October 2020	Support for new features in VxRail 7.0.100
October 2020	Removed references to Log Insight.
November 2020	Removed 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"> Added content on mixing of node ports in VxRail clusters Option for manual node ingestion instead of IPV6 multicast Added content for LACP policies Updated stretched cluster node minimums.
June 2021	<ul style="list-style-type: none"> Updated Intel and AMD node connectivity options for 100 GbE. Expanded network topology option to include custom networks with six Ethernet ports per node. Clarified that VxRail-supplied internal DNS cannot support naming services outside of its resident cluster. Private VLANs (PVLANS) unsupported for VxRail networking
August 2021	Support for new features in 7.0.240
October 2021	Support for satellite nodes
November 2021	Support for PowerFlex as external storage for dynamic cluster
December 2021	<ul style="list-style-type: none"> Update dynamic cluster content with link to Dell published guide

	<ul style="list-style-type: none"> Update content for VxRail Manager network exclusions
January 2022	Support for new 15G VxRail models
March 2022	Support for new features in 7.0.350

Intended use and audience

This guide describes 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 who 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 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 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.

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

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 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:

Introduction	12
VxRail hardware and the physical network infrastructure	12
VxRail clusters, appliances, and nodes	13
Network switch	14
Data center network	17
VxRail networking port options	19
VxRail Ethernet adapter options	20
VxRail Fibre Channel adapter options	21
VxRail RoCE adapter options	22
VxRail node connectivity options	22
VxRail networking rules and restrictions	23
Topology and connections	24

Introduction

The network considerations for VxRail are no different than those of any enterprise IT infrastructure: availability, performance, and extensibility. VxRail appliances are manufactured in the factory per your purchase order, and delivered to your data center ready for deployment. The nodes in the appliance can attach to any compatible network infrastructure to enable operations. 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 or 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. The set of components manufactured into VxRail nodes is based on the customer 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 VxRail cluster starts with a minimum of two 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 Node




Figure 2. Back view of VxRail V-, P-, and S-Series Node




Figure 3. Back view of VxRail G-Series Node

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 connectivity 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 eases 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 IT administration and external data center services. The second management network is isolated, visible only to the VxRail nodes.

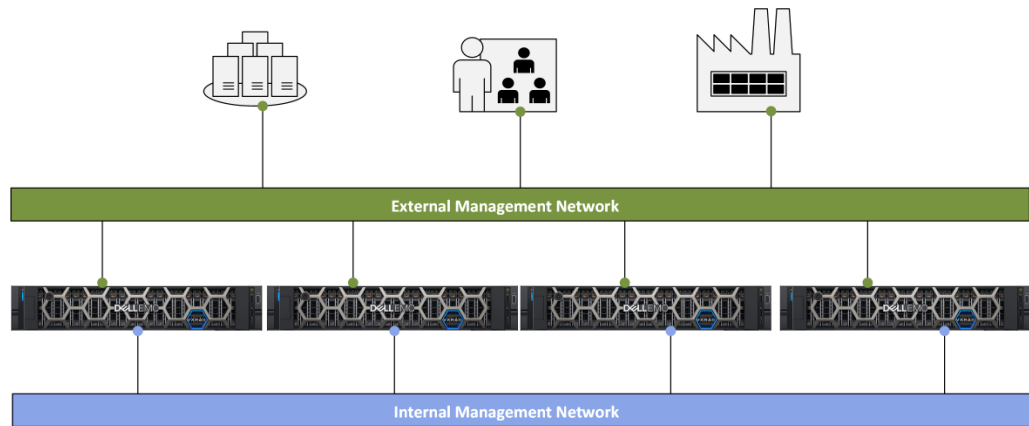


Figure 4. VxRail Management Networks

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 only to the interfaces connected to this internal management network.

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 Ethernet switch does not need to support Layer 3 services or be licensed for Layer 3 services. You can enable routing services further upstream on the network infrastructure, or enable routing services at this ‘top-of-rack’ switch.

A VxRail cluster can be deployed in a ‘flat’ network using the default VLAN on the switch. It can also 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, and VLANs should be used. 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 minimum 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 recommended for 10 GbE network connectivity. 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.

Storage network considerations

There may be additional feature considerations to account for when selecting Ethernet switches for your VxRail cluster, depending on interoperability requirements for different types of storage resources.

- If your VxRail cluster will include adapters that support RoCE (RDMA over Converged Ethernet) for vSAN storage connectivity, the supporting network must support a 'lossless' transport. A 'lossless' network is defined as one where no frames are dropped because of network congestion.
 - Select switches that support Data Center Bridging (DCB). The Data Center Bridging feature supports the elimination of packet loss due to buffer or queue overflow.
 - The Data Center Bridging must support bandwidth allocation based on priority settings, known as Class of Service (CoS).
 - Priority Flow Control (PFC) is required on the switches to provide RoCE traffic a higher priority than other network traffic.

Network redundancy and performance considerations

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. They should be 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.

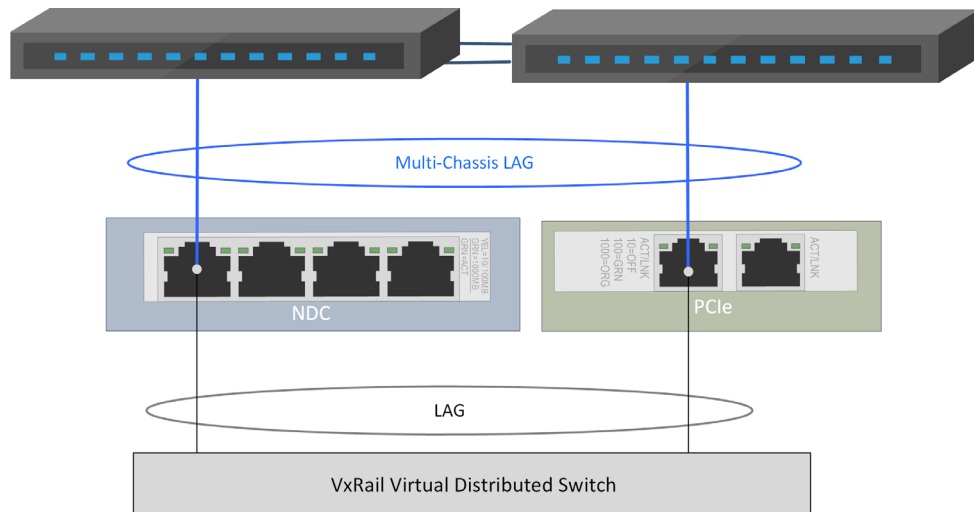


Figure 5. Multi-chassis link aggregation across two switches

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 that is 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 broader 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. You can choose a DNS service internal to VxRail, or use a DNS service in your data center.
- VxRail cluster depends on Network Time Protocol (NTP) to keep the clock settings on the various VxRail components synchronized. Dell Technologies recommends using a reliable global timing service 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.

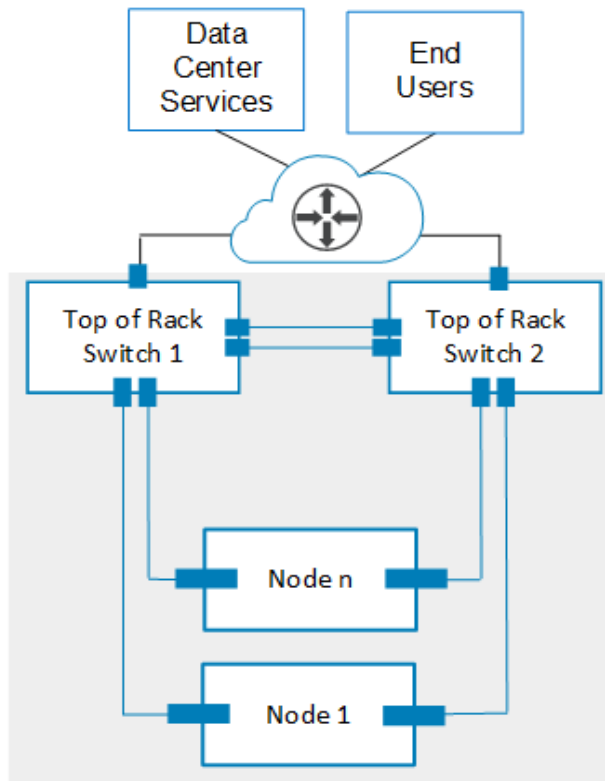


Figure 6. Connecting data center services with VxRail cluster

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.

It is best practice to reserve a set of Virtual LAN (VLAN) IDs in your data center network that will be assigned to support the VxRail networks, especially for production workloads. All these reserved VLANs must be configured on the adjoining physical switches connected to the VxRail nodes. The VLANs cannot be configured as private VLANs.

Certain VxRail management components must be able to connect to data center services, such as DNS and NTP. Routing services must be configured to enable connectivity to these services for these management components. Additional networks, 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. They must also be configured on the ports on the upstream network devices, so they can pass through upstream to Layer 2/Layer 3 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 networking port options

The following tables show the network connectivity options that are supported for base connectivity on the built-in adapter cards for each VxRail model series, and the available PCIe-based adapter options for the expansion slots on the nodes. The supported Ethernet and Fibre Channel adapter card models, the number of slots available for expansion, and the maximum number of networking ports supported per node is driven by factors such as the VxRail model series selected for the cluster, and the number of CPUs installed per node.

The following networking connectivity rules apply to VxRail nodes:

- The built-in NDC/OCP ports installed into the back of each VxRail node is required. A VxRail node cannot be ordered without an NDC/OCP adapter selection.
- There is only one NCP/OCP slot per VxRail node, and only one option can be selected per VxRail node.
- PCIe expansion slots can be populated to support VxRail cluster networking, or to support networking requirements outside of VxRail

VxRail 14G Intel-CPU Node Connectivity Options																		
VxRail Series	Storage Type	1x1 iDRAC	2x10 RJ45		2x10 SFP+		4x10 RJ45		4x10 SFP+		2x25 SFP28		4x25 SFP28		2x100 QSFP		2x16/32 FC	
			B	E	B	E	B	E	B	E	B	E	B	E	B	E	B	E
D-Series	H/F																	
E-Series	H/F																	
E-Series	N																	
G-Series	H/F																	
P-Series	H/F																	
P-Series	N																	
S-Series	H																	
V-Series	F																	

Connectivity Types: B=Built-in (NDC/OCP) E=Expansion (PCIe)
Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 7. Node Connectivity Options for 14th Generation VxRail models with Intel CPUs

VxRail 15G Intel-CPU Node Connectivity Options																		
VxRail Series	Storage Type	1x1 iDRAC	2x10 RJ45		2x10 SFP+		4x10 RJ45		4x10 SFP+		2x25 SFP28		4x25 SFP28		2x100 QSFP		2x16/32 FC	
			B	E	B	E	B	E	B	E	B	E	B	E	B	E	B	E
			E-Series	H/F														
E-Series	N																	
P-Series	F																	
S-Series	H																	
V-Series	F																	

Connectivity Types: B=Built-in (NDC/OCP) E=Expansion (PCIe)
Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 8. Node Connectivity Options for 15th Generation VxRail Models with Intel CPUs

VxRail AMD-CPU Node Connectivity Options																		
VxRail Series	Storage Type	1x1 iDRAC	2x10 RJ45		2x10 SFP+		4x10 RJ45		4x10 SFP+		2x25 SFP28		4x25 SFP28		2x100 QSFP		2x16/32 FC	
			B	E	B	E	B	E	B	E	B	E	B	E	B	E	B	E
			E-Series	H/F														
E-Series	N																	
P-Series	F																	
P-Series	N																	

Connectivity Types: B=Built-in (NDC/OCP) E=Expansion (PCIe)
Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 9. Node Connectivity Options for VxRail Models with AMD CPUs

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	Intel
	Broadcom
	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, but not required, 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 a specific feature 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 Fibre Channel adapter options

If your plans include using Fibre Channel storage as a storage resource for VxRail, the following Fibre Channel adapter cards are supported. These adapter cards can be included with your purchase order and installed on the nodes during the manufacturing process. You can also obtain the adapter cards later if your storage requirements change.

FC Speed	Vendor
16 GB	Emulex
	QLogic
32 GB	Emulex
	QLogic

VxRail RoCE adapter options

RoCE adapter support Remote Direct Memory Access (RDMA) connectivity over Converged Networks. RDMA is a technology to enable sending data over a network without involving the CPU in the transfer, thereby providing appreciable IOPS performance in comparison to other network-based storage connectivity options. Enabling RDMA over a Converged Ethernet network infrastructure provides faster data transfer for network-intensive applications through lower I-O latencies on this network.

Port Speed	Vendor
25 GbE	Mellanox

The following guidelines should be understood for RoCE-supported adapters:

- The VxRail cluster must be configured with a minimum version of 7.0.200
- All the nodes in the cluster connected to the common vSAN datastore must be configured with RoCE-supported adapter cards of the same vendor and the same model. This is strongly encouraged as a best practice to remove any possibility of slight variances in adapters from different vendors or different models disrupting I/O on the vSAN datastore.
- The Ethernet ports on the RoCE-supported adapters are reserved for vSAN traffic only.
- The physical network supporting vSAN is configured as a 'lossless' network,

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/OCP. Starting with version 7.0.130, for workloads that require a higher level of failure protection, VxRail supports spreading the networks across NDC/OCP Ethernet ports and Ethernet ports on PCIe adapter cards.

The custom network option provides flexibility with the selection of the Ethernet ports to support the VxRail cluster networking. 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 assign those ports to the VxRail networks. There is more flexibility with the customized port selection option, since you can use just the ports on the NDC/OCF, mix the ports from the NDC/OCF and a PCIe adapter card, or select only ports from PCIe adapter cards. For more details, refer to [Appendix F: Physical Network Switch Examples](#) to understand the most common node connectivity options.




Figure 10. Mixing NDC/OCF and PCIe ports to support a VxRail cluster

If you later decide to change the topology by migrating 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 for VxRail cluster networking can be used for other customer use cases, such as guest networks, external storage, and other requirements.
- Guest networks can share resources with the Ethernet ports reserved for VxRail networking, or unused Ethernet ports on the nodes can be configured to support guest networks.
- For VxRail clusters running all Ethernet ports at 1 GbE speed:
 - Four ports on each node must be reserved for VxRail network traffic.
 - Single processor VxRail models only
 - Maximum of eight nodes per cluster
 - Only hybrid VxRail models can be configured with 1 GbE speed. All-flash VxRail models cannot support 1 GbE.
- For VxRail nodes supplied with Ethernet ports greater than 1GbE:
 - The most common topology is to configure the cluster with either two ports or four ports per node to support VxRail networking traffic.
 - Adding Ethernet ports beyond the ports initially reserved for VxRail networking is not supported after the cluster is configured and operational.

- VxRail networks that become resource constrained due to increased workload demands can be migrated to higher-speed Ethernet ports, provided the VxRail cluster is running version 7.0.010 or later.
- Optionally, reserving six ports or eight ports per node for VxRail network traffic is supported. This option is best used for deployments supporting very demanding and network-intense workloads.
- Custom Ethernet port configurations are supported with restrictions:
 - Before VxRail version 7.0.130, all the Ethernet ports on the VxRail nodes selected for a VxRail cluster must be the same port type and running at the same speed.
 - Starting in VxRail version 7.0.130, the ports on the NDC/OCP and PCIe adapter cards configured in the VxRail nodes can be running at different speeds. For instance, the NDC/OCP ports can be running at 10 GbE and the ports on the PCIe adapter cards can be running at 25 GbE.
 - Any ports assigned to the same VxRail network, whether based on NDC/OCP or PCIe, must be running at the same speed. For instance, a VxRail network cannot be paired with one port running at 10 GbE and another port running at 25 GbE.
 - Individual VxRail networks can be assigned to Ethernet ports running at different speeds. For instance, one VxRail network can be assigned to ports running at 100 GbE, while another VxRail network can be assigned to ports running at 10 GbE.
 - Mixing Ethernet ports types, such as RJ45 and SFP+, 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 11. 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 12. 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 will 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 must also 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 Cluster Types

This chapter presents the following topics:

Introduction	27
Cluster with vSAN storage	27
Dynamic cluster	28
vSAN stretched cluster	32
2-node cluster	33
Satellite nodes	35

Introduction

The primary building block for VxRail is the individual node, and a collection of nodes are then used to form a VxRail cluster and placed under a single point of management. The nodes can be customized with specific components to support different VxRail cluster types based on business and operational requirements.

The VxRail nodes can be customized to provide all the physical compute, network, and storage resources for the cluster. This is accomplished by using the local disk drives on each node to form a vSAN datastore as the primary storage resource for application workload. Alternatively, the nodes can be customized without local disk drives to instead use external data center resources for primary storage.

This chapter is an introduction to the types of clusters supported with VxRail, and the types of use cases that can be supported with each cluster type.

Cluster with vSAN storage

One type of VxRail cluster is one where the VxRail nodes provide all the physical compute and storage resources to support application workload, and the primary storage resource is vSphere vSAN. For this cluster type, the slots in the nodes are filled with disk drives that meet the performance and capacity requirements for the application workload, and are then formed into a local vSAN datastore during the cluster initialization process.




Figure 13. Local vSAN datastore deployed on a VxRail cluster to be the primary storage resource

The cluster initialization process will perform an inventory of the disk drives on the nodes, and use that discovery process to form a vSAN datastore. The high-endurance SSD drives discovered on each node serve as a cache for virtual machine I-O operations, while the high-capacity drives discovered are the primary permanent storage resource for the virtual machines. The vSAN formation process will partner the cache drives with one or more capacity drives to disk groups on the nodes, with the resulting vSAN datastore consisting of the collection of disk groups.

This cluster type is designed to handle most customer application workloads and covers the most common use cases.

- With a local vSAN datastore, operational flexibility can be realized to address scalability and high availability requirements.
- This cluster type is simple to deploy and operate. The initialization process performs all the work to pool all the node resources for ease of consumption
- This cluster type is not dependent on external resources for storage, so all resources can be administered under a single point of management.

- Additional compute and storage resources can be expanded easily to the cluster through automated node and disk drive addition.
- External storage resources can be configured on the cluster as secondary storage capacity.

However, this cluster type may not be a good fit for certain use cases:

- This cluster type may be cost-prohibitive for smaller business requirements. Light workloads, such as those in a remote office, may be a better fit for a two-node cluster or a satellite node.
- This cluster type is located in a single site. It offers support for continuous operations from a failure of components within the cluster, but does not offer zero-downtime protection from a failure of a single data center or site. A stretched VxRail cluster is a better solution for very high availability requirements.
- Expansion of a cluster through node addition can potentially lead to stranded assets, where excess compute and storage resources cannot be shared outside of the cluster. For workloads which require a more precise balance of compute and storage resources, a dynamic cluster may be a better fit.

Dynamic cluster

Dynamic clusters differentiate themselves from other VxRail cluster types with the resource selected for primary storage. With other cluster types, there is a dependency on the local vSAN datastore as the primary storage resource. With a dynamic cluster, the nodes used to build the cluster do not have local disk drives. Therefore, an external storage resource is required to support workload and applications.

A dynamic cluster may be preferable to other cluster types in these situations:

- You already have an investment in compatible external storage resources in your data centers that can serve as primary storage for a dynamic cluster.
- The business and operational requirements for the applications targeted for the VxRail cluster can be better served with existing storage resources.
- The likelihood of stranded assets through node expansion is less likely with a dynamic cluster.

If a dynamic cluster is the best fit for your business and operational requirements, be aware of the following:

- The target data center for the VxRail dynamic cluster must already have deployed one of the supported options for primary storage.
- VxRail publishes a guide to assist you in preparing your data center storage for a VxRail dynamic cluster at [Configure External Storage of VxRail Dynamic Node Cluster](#) on Dell's product support site. VxRail recommends using the technical documentation provided for the selected external storage to complement this guide.
- Any performance issues that are diagnosed at the storage level may be related to infrastructure outside of VxRail, and must be managed separately.

- A dynamic cluster does not have the same level of visibility and control of an external storage resource in comparison to a local vSAN datastore.

Fibre Channel Storage Option

One of the external storage resources supported with dynamic clusters is Fibre Channel storage. With this option, a compatible Fibre Channel storage array can be configured to supply a single VMFS datastore or multiple datastores to a VxRail dynamic cluster.




Figure 14. Dynamic cluster using FC-connected VMFS for primary storage

This option has the following pre-requisites:

- Verify that the storage array in your data center you have planned for support of dynamic cluster is supported with VxRail. Consult the [VxRail E-Lab Navigator](#) to verify compatibility.
- Verify you have sufficient free capacity on the storage array. A VxRail dynamic cluster requires a VMFS device with a minimum of 800 GB to support workload and applications.
- At the time of ordering, include enough compatible Fibre Channel adapter cards. VxRail recommends a minimum of two Fibre Channel adapter cards per node for redundancy, although a single dual-port adapter card can be used.
- Verify that you have sufficient open ports on your Fibre Channel switches to accommodate the connections required from each VxRail node.

Remote vSAN Datastore Option

Another option to use for primary storage with dynamic clusters is an existing vSphere or VxRail cluster with a local vSAN datastore in your data center. The virtual machines running on the dynamic cluster will use the free storage resource on the vSAN datastore, and the compute resources on the local nodes.




Figure 15. Dynamic cluster using a remote vSAN datastore for primary storage

If you choose to pursue this option, ensure that the following guidelines are understood:

- Verify the cluster being targeted to supply storage resources to the dynamic cluster is at a support VxRail version. See the [VxRail Support Matrix](#) to verify if an upgrade is needed on this cluster.
- The cluster that is sharing its vSAN datastore and any dynamic cluster connected to this remote vSAN datastore must be configured on the same vCenter instance under a common data center object.
- If you already have a cluster that is sharing its vSAN datastore to other clusters, ensure that you have not reached the maximum of five clusters already mounted to this vSAN datastore.
- The physical Ethernet network in your data center must support connectivity between the cluster nodes of both clusters.
 - If Layer 3 connectivity is required, routing settings such as static routes or BGP must be configured.
 - The RTT latency between the cluster sharing the vSAN datastore and the dynamic cluster nodes must be less than 5 milliseconds.
 - Routable IP addresses must be used on the VMkernel adapters supporting vSAN on both the cluster nodes sharing the vSAN datastore and dynamic cluster nodes.

PowerFlex storage option

PowerFlex is a Dell storage product that is supported with VxRail dynamic cluster. PowerFlex systems provide IP-based storage to VxRail dynamic clusters that can be configured as primary storage for virtual machine workload.



Figure 16. VxRail dynamic cluster storage provided by PowerFlex virtual volume

The PowerFlex system configures pools of storage through a virtualization process, and manages the allocation of virtual volumes to connected clients. Virtual volumes can be

configured to meet certain capacity, performance, and scalability characteristics to align with the workload requirements planned for the VxRail dynamic cluster.

The PowerFlex architecture combines both the compute and storage in a fabric connected network architecture, with Dell PowerEdge servers serving as the hardware foundation for block storage capacity.




Figure 17. PowerFlex system interconnected with VxRail dynamic cluster

If you plan to leverage virtual volumes provided by a PowerFlex system to serve as the primary storage for your VxRail dynamic cluster, ensure that best practices are followed to ensure a successful deployment:

- Follow the guidance in the [Dell PowerFlex Networking Best Practices and Considerations](#) to ensure the supporting network infrastructure is properly planned and configured
- Follow the steps in [How to configure Dell PowerFlex Storage with VxRail Dynamic Nodes](#) if you are unfamiliar with provisioning PowerFlex storage for this purpose
- Reserve two Ethernet ports on each VxRail node planned for the dynamic cluster to support connectivity to the PowerFlex volumes serving as primary storage
- Enable jumbo frames on the network configured to support VxRail dynamic cluster storage
- After the VxRail dynamic cluster is built, plan to configure a separate virtual-distributed switch with new portgroups in vCenter to support connectivity to the PowerFlex front-end system

iSCSI and NFS Options

If the data center does not support Fibre Channel storage, shared vSAN resources or have a PowerFlex storage array deployed, then storage based on either iSCSI or NFS over an IP network can be used as primary storage.

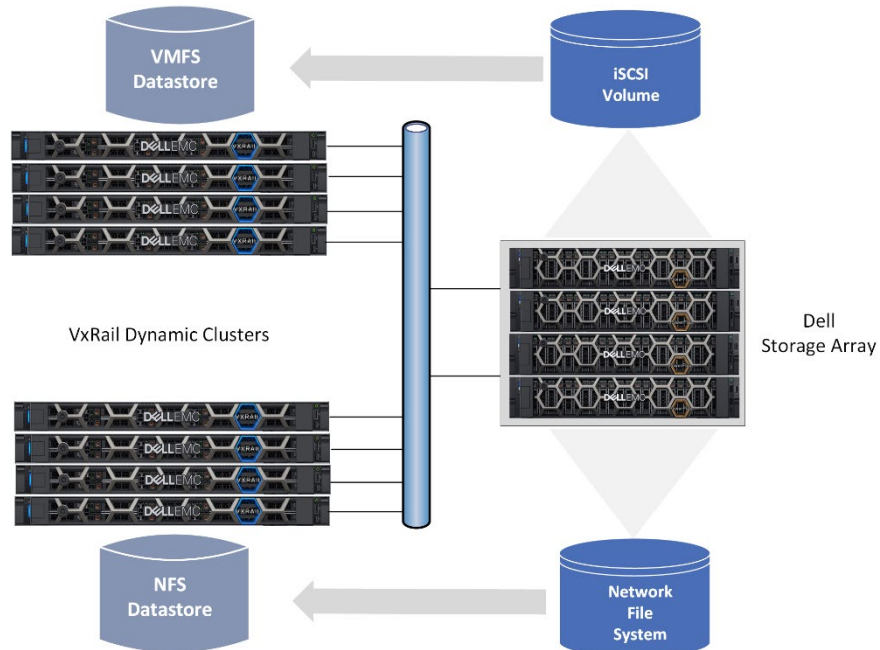


Figure 18. External storage based on iSCSI or NFS supporting VxRail dynamic clusters

With the iSCSI option, block-level storage is presented to the VxRail cluster over an IP network. iSCSI is standard feature in VMware vSphere, and is enabled by configuring a software adapter on a NIC on the VxRail nodes. The adapter serves as an iSCSI initiator by targeting external storage arrays to present LUNs back to the initiators. These LUNs are then configured as VMFS datastores to support virtual machine workload.

The NFS option also works over an IP network, except the storage presented back to the VxRail cluster is from a compatible file server, and the storage format is file-based instead of block-based. With this option, the external file system is mounted by the VxRail nodes to enable access over the IP network, and configured to serve as a datastore.

As with the other storage options for dynamic clusters, verify that the storage resource in your data center you planned to support VxRail dynamic clusters, consult the [VxRail E-Lab Navigator](#) to verify compatibility.

vSAN stretched cluster

vSAN stretched cluster is a VMware solution that supports synchronous I/O on a local vSAN datastore on two sites that are separated geographically. This type of vSAN solution is supported on VxRail. The primary benefit of a vSAN stretched cluster is that it 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:

- The VxRail nodes must be populated with the compatible disk drives so that a vSAN datastore can be configured.

- The compute and storage resources that are planned for the cluster should be doubled. It is best practice to reserve 50% of resource capacity in the cluster for failure protection.
- Three data center sites are required for this solution: 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 are required in the Primary site, and an equal number of VxRail nodes are required in the Secondary site.
- A minimum of one top-of-rack switch for the VxRail nodes in the Primary site and in the Secondary site
- 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 19. vSAN Stretched Cluster Topology

More detailed information about vSAN stretched cluster and the networking requirements can be found in the [Dell 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 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.

- Verify that your workload requirements do not exceed the resource capacity of this small-scale solution.
- You cannot expand to three or more nodes unless the cluster is running version 7.0.130 or later.
- 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
- 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 20. 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 vSAN 2-Node Cluster Planning and Preparation Guide](#).

Satellite nodes

For use cases where cost efficiency and economies of scale are essential, a VxRail cluster at a central location can be positioned to monitor and manage pools of VxRail satellite nodes, deployed locally and at remote locations.

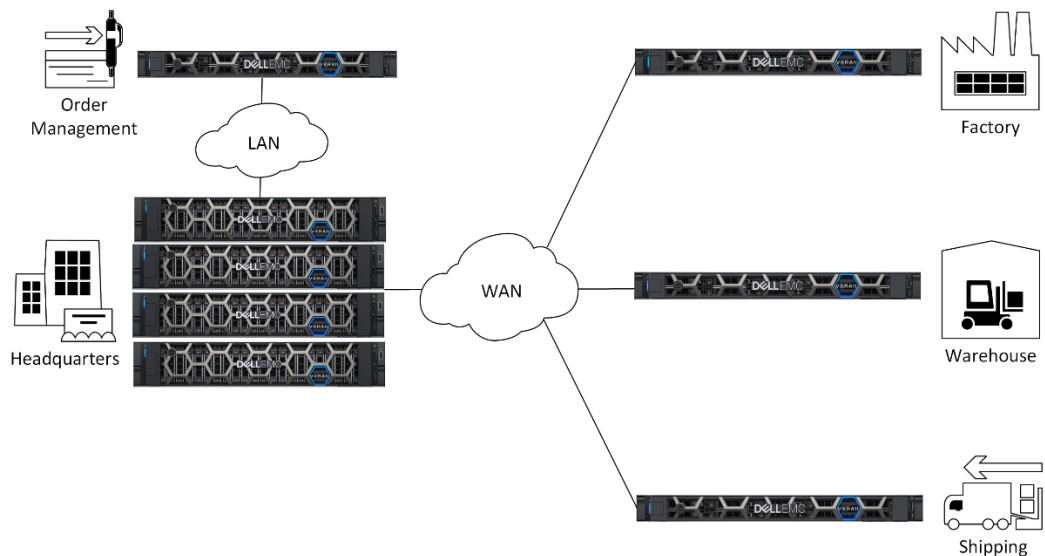


Figure 21. VxRail satellite nodes managed by VxRail cluster

The PowerEdge server models used as the hardware foundation for the other VxRail cluster types are the same for satellite nodes. Satellite nodes also go through the same engineering, qualification, and manufacturing processes as the VxRail nodes used in clusters, and software lifecycle management of satellite nodes is supported through VxRail Manager.

The primary difference from a networking perspective between satellite nodes and the nodes supporting other cluster types is that satellite nodes require only a single IP address. This IP address is used to enable connectivity to a VxRail cluster in a central location and establish communication for management purposes.

If you plan to deploy VxRail satellite nodes, note the following:

- A compatible VxRail cluster with local vSAN storage must already be deployed to support the management and monitoring of satellite nodes.
- The minimum VxRail software version to support satellite nodes is 7.0.320.
- VxRail satellite nodes are limited to a single instance and cannot be reconfigured to join a cluster.
- Verify that your workload requirements at the remote locations do not exceed the resource capacity of a satellite node.

Chapter 4 VxRail Feature-Driven Decision Points

This chapter presents the following topics:

Introduction	38
Software-defined data center	38
Dell SmartFabric network mode	39
vSphere with Kubernetes on VxRail	41

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.

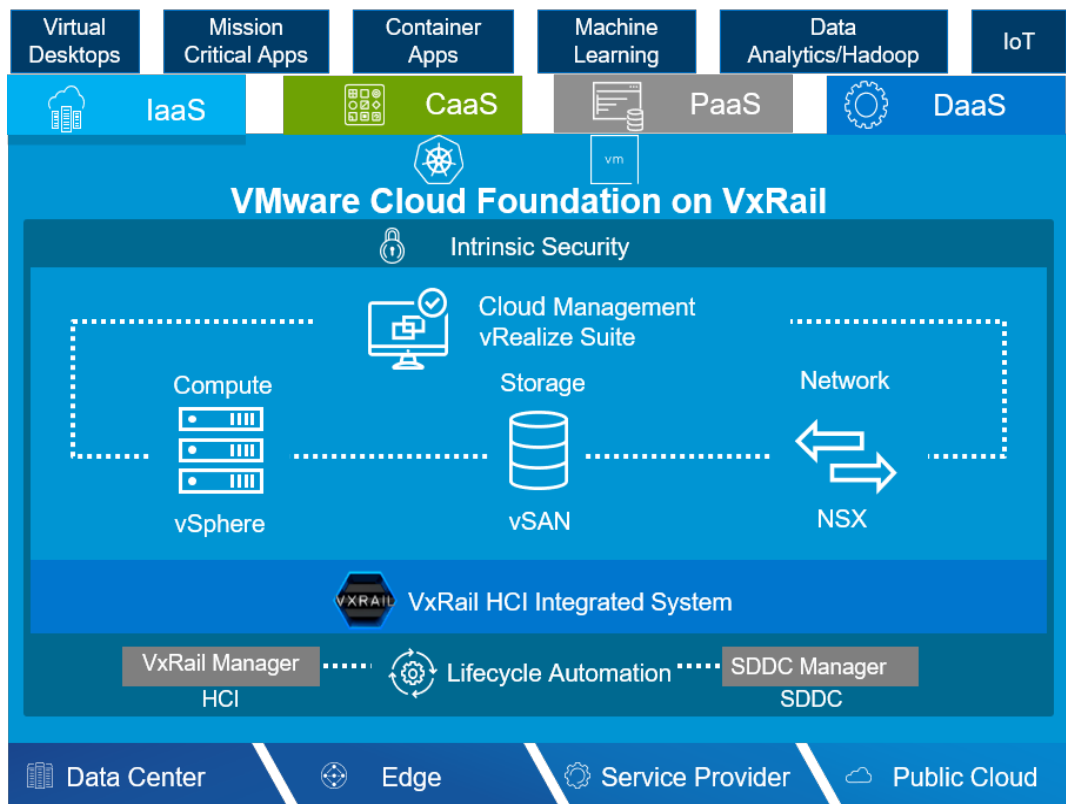


Figure 22. VxRail as the foundation for the software-defined data center

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 enables 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 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 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 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 switch network fabric.

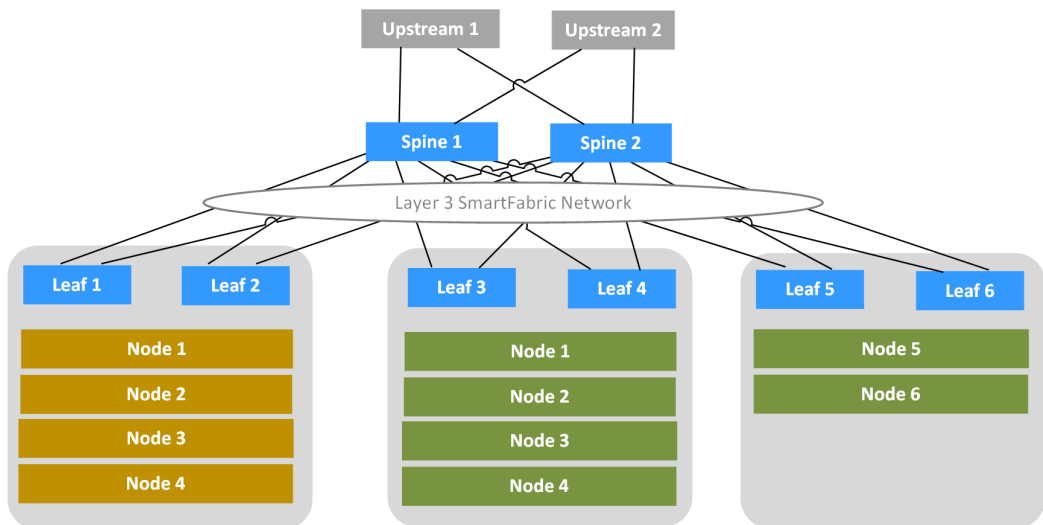


Figure 23. Dell SmartFabric for VxRail

For ongoing VxRail cluster network management after initial deployment, the Dell Open Manage Network Interface (OMNI) vCenter plug-in is provided free of charge. The Dell OMNI plug-in enables the integration and orchestration of the physical and virtual networking components in the VxRail-SmartFabric HCI stack. It provides deep visibility from the vClient for ease of overall management and troubleshooting. The Dell 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 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.

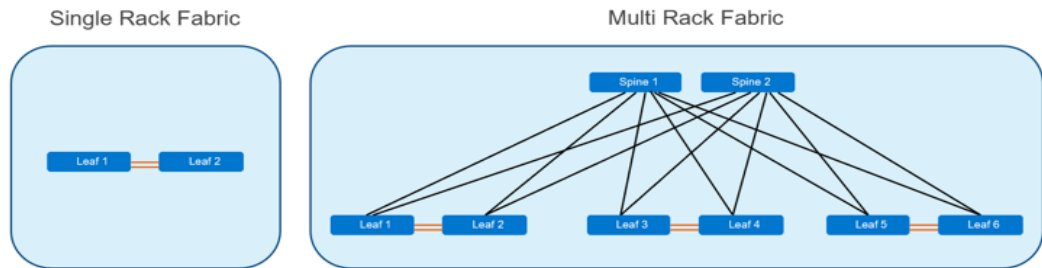


Figure 24. SmartFabric-enabled multirack network expansion

Planning for VxRail with the Dell 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 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. The management of switch configuration going forward is 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 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, 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

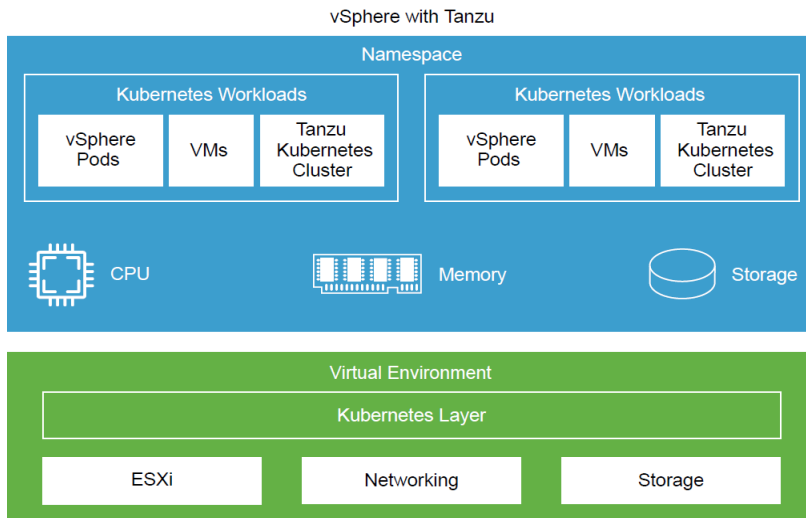


Figure 25. vSphere with Tanzu on a VxRail cluster

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 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](#).

Chapter 5 VxRail Hardware and Switch Selection Decision Points

This chapter presents the following topic:

Step-by-step networking decision points 44

Step-by-step networking decision points

- 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.
 - More than four ports per node can be reserved for VxRail networking for cases where it is desirable for certain individual VxRail networks to not share any bandwidth with other VxRail networks.

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, 25 GbE, and 100 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/OCP for VxRail cluster networking.
 - Reserve and use both NDC/OCP-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 reserve and use just NDC/OCP ports to build the initial cluster, and then migrate certain VxRail networks to PCIe-based ports.

- If your requirements include using FC storage to support VxRail workload, you can select either 16 GB or 32 GB connectivity to your Fibre Channel network.

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

- Step 6.** Select the network adapter type and cable type to connect the VxRail nodes to your switches
- VxRail nodes can connect to switches with either RJ45, SFP+, SFP28, or QSFP adapter types, depending on type of adapter card(s) selected for the nodes.
 - VxRail nodes with RJ45 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/OCP. 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.** Determine whether to enable out-of-band management. Dell iDRAC functionality is built into each VxRail node, and requires a 1GbE connection. Dell Technologies recommends deploying a dedicated 1 GbE switch for this purpose. If this is not practical, you can also use open ports on the top-of-rack switches.
- Step 9.** Decide whether to use a local laptop or a jump host to enable initial connectivity to the VxRail management interface.
- To use a local laptop for VxRail management connectivity, reserve one additional port on one of the top-of-rack switches for this purpose.
 - The need for the additional port to access the management interface is removed if connectivity is available elsewhere on the logical path from a jump host on the VxRail external management VLAN.

Chapter 6 Prepare data center for VxRail Implementation

This chapter presents the following topics:

Introduction.....	47
Prepare data center routing services	47
Prepare for multi-rack VxRail cluster.....	48
Prepare for vSAN HCI mesh topology.....	50
Prepare external FC storage for dynamic clusters.....	52
Prepare for VxRail custom uplink assignments.....	53
Prepare data center network MTU.....	55
Prepare for link aggregation of VxRail networks.....	55
Plan link aggregation on switch port pairs.....	58
Plan link aggregation on switch port pairs	59
Prepare certificate authority server for VxRail.....	59
Prepare VxRail remote support connectivity.....	59

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.

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.

Prepare data center routing services

Specific VxRail networks, including the VxRail external management network and any external-facing 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 in 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 26. 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 multirack VxRail cluster.




Figure 27. VTEP tunneling between leaf switches across racks

Prepare 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.

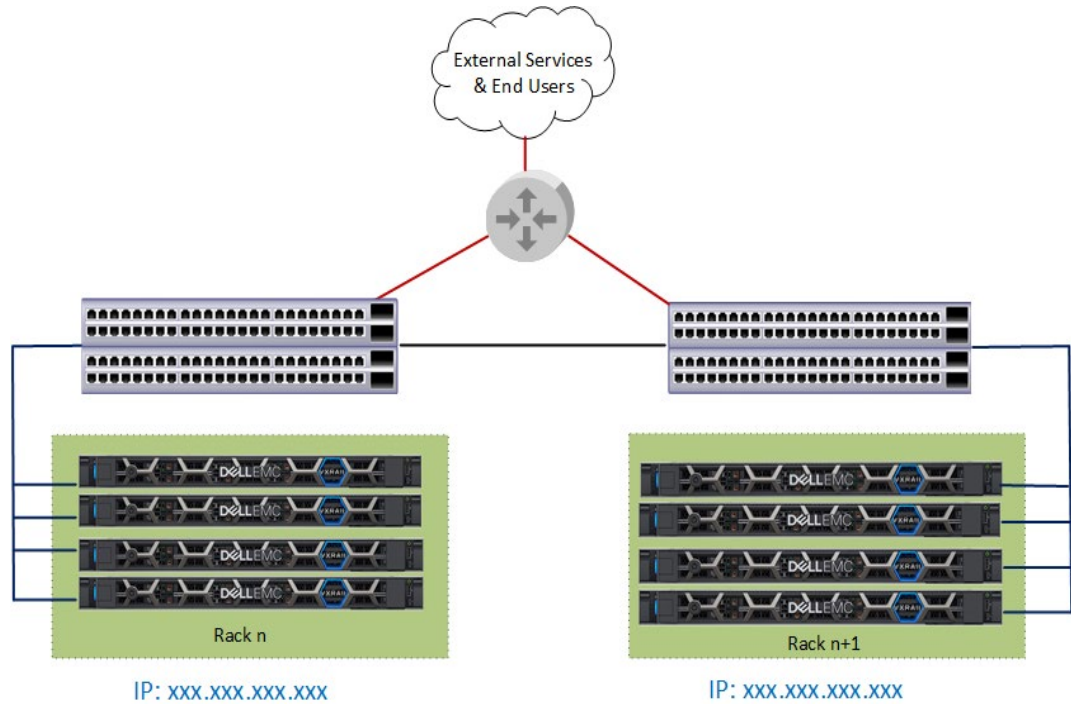


Figure 28. Multi-Rack VxRail sharing the same subnet




Figure 29. Multi-Rack VxRail with different subnets

Prepare for vSAN HCI mesh topology

This section is relevant only in situations for sharing vSAN datastore resources over the network. This section can be skipped if this is not in your plans.

With an HCI Mesh topology, the local vSAN datastore on a 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 object on a common vCenter instance.

With a vSAN HCI Mesh network in place, VxRail cluster that leverage a local vSAN datastore for primary storage can also leverage the capacity on a remote vSAN datastore for application workload. If dynamic clusters will be part of the vSAN HCI mesh network, they can mount a remote vSAN datastore on a cluster with free vSAN storage capacity and use that as their primary storage resource.

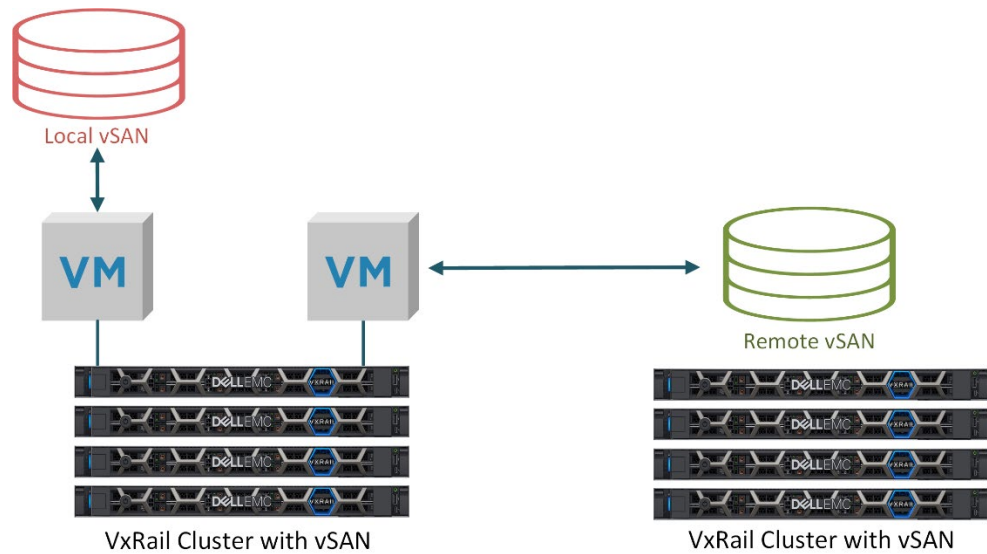


Figure 30. Storage resource sharing between VxRail clusters with vSAN HCI mesh

If your plans include sharing storage resources between clusters in a 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.

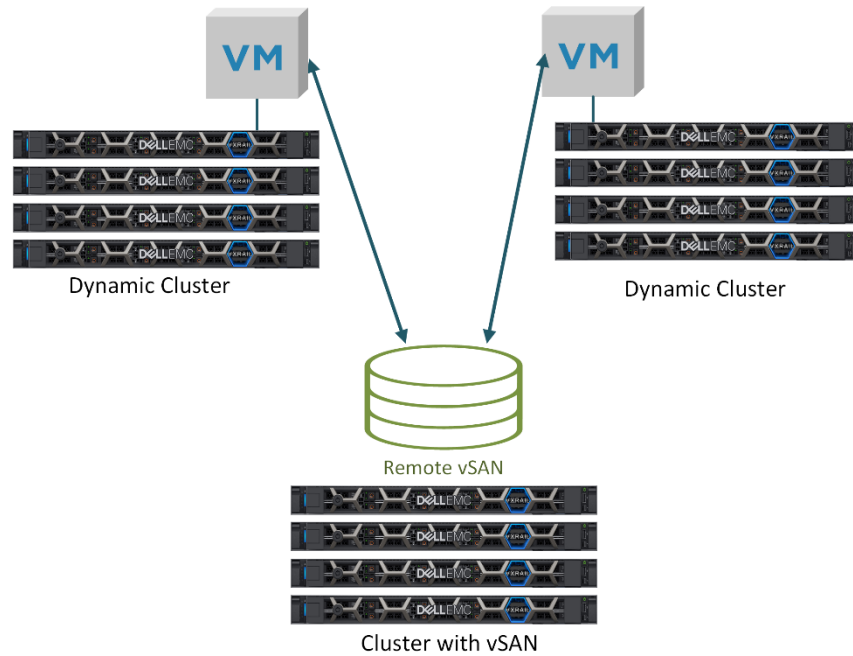


Figure 31. Storage resource sharing between a cluster with a local vSAN datastore and dynamic clusters

If your plans include sharing the vSAN resources of a VxRail cluster with one or more VxRail dynamic clusters, you must adhere to the following pre-requisites:

- A vCenter instance at a version that supports VxRail version 7.0.240 or later.
- A vSAN Enterprise license is needed only for the VxRail cluster that is sharing its vSAN storage resources. This license is not needed on a dynamic cluster.




Figure 32. Enable vSAN network connectivity between clusters

To enable vSAN HCI mesh, the data center must have a network topology that can enable connectivity of the vSAN networks on the two participating VxRail clusters.

- Two Ethernet ports on each node in the cluster will be configured to support vSAN network traffic.
- A common VLAN can be assigned to this vSAN network on each cluster so they can connect over a Layer 2 network. If the VxRail clusters are deployed against different top-of-rack switches, the VLAN must be configured to stretch between the switch instances.
- If the VxRail clusters are deployed against different top-of-rack switches, or if the common VLAN cannot be stretched between the switch instances, 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.

Prepare external FC storage for dynamic clusters

This section is relevant only in situations where Fibre Channel storage will be positioned as the primary storage resource for a VxRail dynamic cluster.

If your plans include using storage resources on a Fibre Channel array to serve as the primary storage resource for a dynamic cluster, a VMFS datastore must be preconfigured on each VxRail node before initial cluster build commences.




Figure 33. Enabling Fibre Channel connectivity for dynamic clusters

Follow these guidelines to prepare your environment before initial cluster build:

- See the [Dell E-Lab Navigator](#) on the Dell support site to verify if the Fibre Channel array you were planning to support dynamic clusters is compatible.
- Include Fibre Channel adapter cards with the VxRail order. Each node that is a member of a dynamic cluster requires Fibre Channel connectivity. At least one dual-port Fibre Channel adapter card should be installed for each VxRail node.
- Dell recommends a pair of Fibre Channel switches should be deployed for redundancy purposes to support the network topology.
- The minimum LUN size supported for VxRail dynamic clusters is 800 GB. Verify that you have sufficient capacity on your storage array for this purpose.

- If multiple VMFS datastores are detected, the largest free capacity on the primary host will be selected.
- LUN masking to all dynamic nodes must be performed before initial cluster build.
- LUNs must be formatted as a VMFS datastore in order to be supported for VxRail dynamic clusters.

Prepare for VxRail custom uplink assignments

At the time of initial cluster build, 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 predefined 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 predefined 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/OCP, 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 that 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, 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/OCP-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.

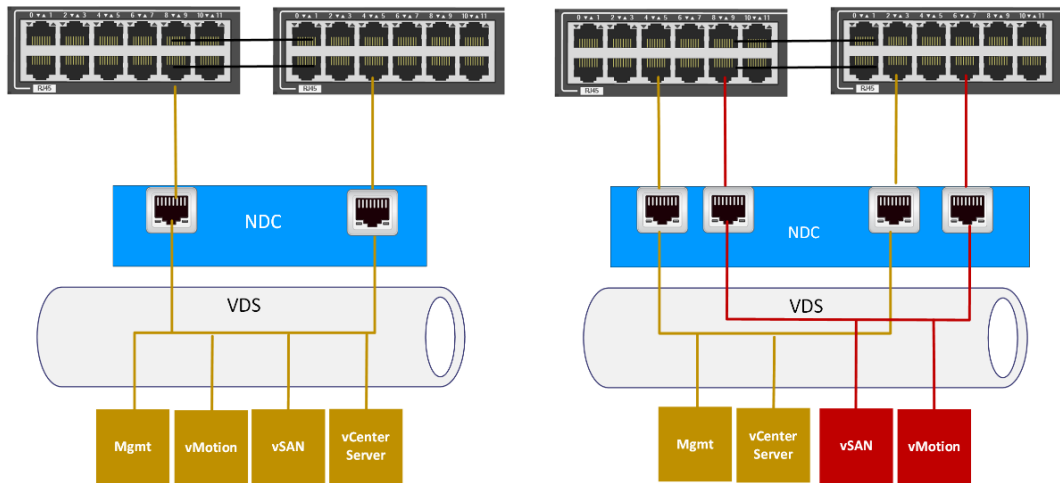


Figure 34. Default network profiles for a 2-port and a 4-port VxRail network

If you plan to use both NDC/OCP-based and PCIe-based ports to enable NIC redundancy and eliminate the NDC/OCP 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/OCP 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/OCP, and another compatible port on the PCIe adapter card, and assign those to the non-management VxRail networks.

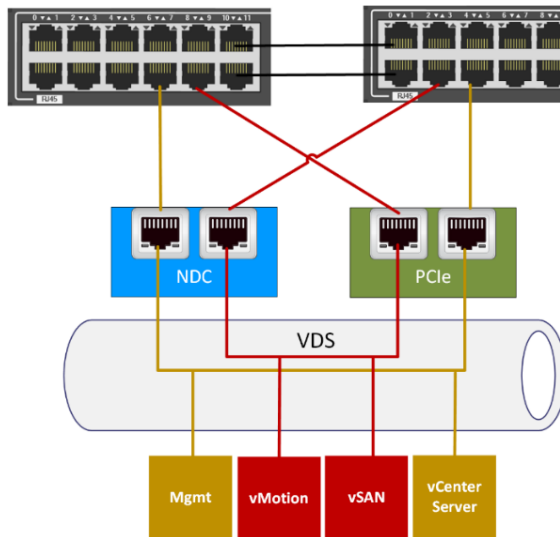


Figure 35. Custom uplink assignment across NDC/OCP-based and PCIe-based ports

If you expect the applications to be running on the VxRail cluster to be I/O intensive and require high bandwidth, you can 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 36. 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.

- With a custom uplink assignment, there is more flexibility in a data center with a mixed network. You can assign the resource-intensive networks like vSAN to higher-speed uplinks, and low-impact networks to slower uplinks, and then connect those uplinks to switches with compatible port speeds.
- On VxRail nodes with both NDC/OCP ports and PCIe Ethernet adapter cards, you can migrate certain VxRail networks off the NDC/OCP ports and onto PCIe ports after the cluster is built. This is advantageous if workload demand increases on the cluster, and additional bandwidth is required. You can later install switches that better align with adapting requirements, and migrate specific workloads to those switches.

Prepare data center network MTU

At the time of initial cluster build, you can choose the MTU size to assign to the virtual-distributed switches. This allows you to configure the switch ports supporting the VxRail cluster to match the MTU setting on the virtual-distributed switches, and reduce the potential for network fragmentation.

You can choose to run the VxRail cluster at the default MTU size of 1500, or instead choose a larger MTU if your network supports it. If your data center network supports jumbo frames, Dell recommends configuring the MTU size at 9000. A resource-intensive network such as vSAN is better suited for a network configured with jumbo frames.

Prepare 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.

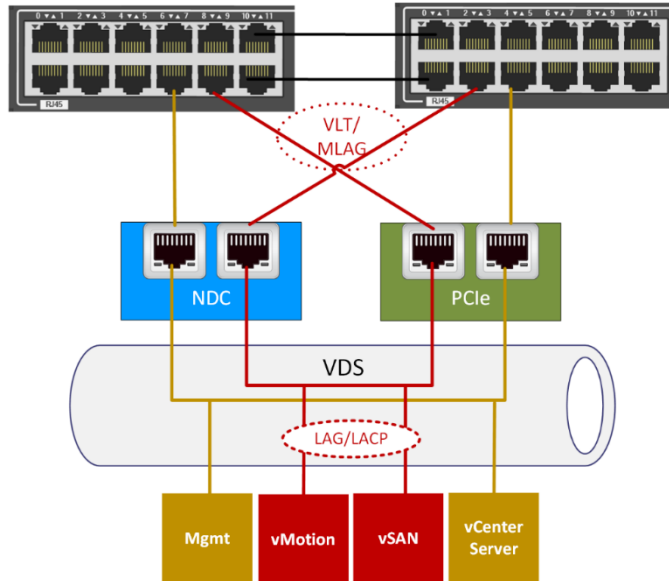


Figure 37. Overview of LAG peering relationship between VxRail and the adjacent switches

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.
- Link aggregation is not supported if the Ethernet adapters supporting the vSAN network are enabled for RDMA over Converged Ethernet (RoCE).
- 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.
 - See [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/OCP and PCIe Ethernet ports
 - All NDC/OCP-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 38. 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/OCP 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 39. Sample port-mapping table

Assuming that you are using the second port on the NDC/OCP 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 Technologies published a procedure for the scenarios where link aggregation on the VxRail networks is performed as a Day 2 activity.

Prepare certificate authority server for VxRail

VxRail Manager is deployed with a “self-signed” certificate by default. The default certificate provided with VxRail Manager is likely to be triggered as “untrusted” by most web browsers, depending on security settings. It is considered best practice to replace the default ‘self-signed” certificate with one that is signed by your organization's root Certification Authority (CA) for production VxRail environments.




Figure 40. Trusted certificate authority supporting VxRail deployment

The certificates in VxRail are configured with an expiration. Certificate replacement can be performed manually, or auto-renewed starting in VxRail version 7.0.350. When auto-renewal is enabled, VxRail Manager will automatically contact the certificate authority for new certificates before the expiration date.

Prepare VxRail remote support connectivity

This section is only relevant if you plan to enable VxRail's remote support feature to connect back to Dell Technologies Customer Support centers.

VxRail supports network connectivity to Dell's Customer Support systems, which enables health monitoring, real-time analytics, event notification, and remote support when necessary. If you plan to enable VxRail to connect to Dell Technologies back-end Customer Support centers, then there are two options to select from. One is to have VxRail Manager connect to the back-end services sites directly, or the other option is to use a centralized instance or pool of Secure Connect Gateways for this purpose.




Figure 41. Enabling ‘call-home’ connectivity with Secure Connect Gateway

If you already have Secure Connect Gateway deployed in your data center to support this service, known as ‘call-home’, for other Dell Technologies products, then this same pool of gateway servers can be used for VxRail. If this is your first Dell product you plan to connect to this service, be sure to first complete the pre-requisites:

1. Create and validate a Dell Support account
2. Configure at least one Site ID, which is used for identity purposes in the Customer Support databases
3. Register the VxRail product with a Site ID

If you do not have any Secure Connect Gateways deployed, the product is customer-installable and can be downloaded from the Dell support site. Each instance of Secure Connect Gateway will require an IP address, and VxRail Manager must be able to reach these gateways in your data center to complete the connection to the back-end customer service centers.

Versions of VxRail prior to 7.0.350 connected to a Secure Remote Services virtual appliance in the data center to enable alerts from VxRail to be sent to Dell Technologies customer service. Starting in version 7.0.350, the Secure Remote Services virtual appliance will no longer be needed for this purpose, and is replaced by a service internal to VxRail Manager. This internal VxRail service will support connectivity directly to the back-end customer service sites, or to Secure Connect Gateways.

If you are deploying VxRail with a version earlier than 7.0.350, you can decide whether to use an existing Secure Remote Services instance in your data center to support ‘call-home’, or deploy a virtual instance of Secure Remote Services on the VxRail cluster for this purpose. If you choose to deploy the Secure Remote Services virtual appliance on the VxRail cluster, you will need to connect the virtual appliance to the same network as VxRail Manager.

Chapter 7 Planning the VxRail Cluster Implementation

This chapter presents the following topics:

Introduction	62
Decide on VxRail single point of management	62
Decide on VxRail network traffic segmentation	63
Decide on teaming and failover policies for VxRail networks	65
Plan the VxRail logical network	66
Plan network settings for VxRail management components	72
Plan network settings for vCenter Server management network	74
Identify IP addresses for VxRail management components	74
Select hostnames for VxRail management components	75
Identify external applications and settings for VxRail	77
Prepare customer-supplied vCenter server	78
Prepare customer-supplied virtual-distributed switch	80
Reserve IP addresses for VxRail vMotion network	83
Reserve IP addresses for VxRail vSAN network	84
Decide on VxRail logging solution	85
Assign passwords for VxRail management	86
Prepare for Dell SmartFabric Services enablement	87

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 as described in this section.

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 requires 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. The internal DNS option will only support naming services on which the VxRail cluster it is deployed. This option cannot support naming services outside of that cluster.

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 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.
- If you are planning to deploy one or more dynamic VxRail clusters where the primary storage resource is going to be a remote vSAN datastore, a customer-supplied vCenter is preferable. All clusters participating in sharing and receiving vSAN datastore resources must reside in a common 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 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. You can configure all the required VxRail networks to a single virtual-distributed switch, or you can deploy a second virtual-distributed switch to isolate the VxRail management network traffic and the VxRail non-management network traffic.

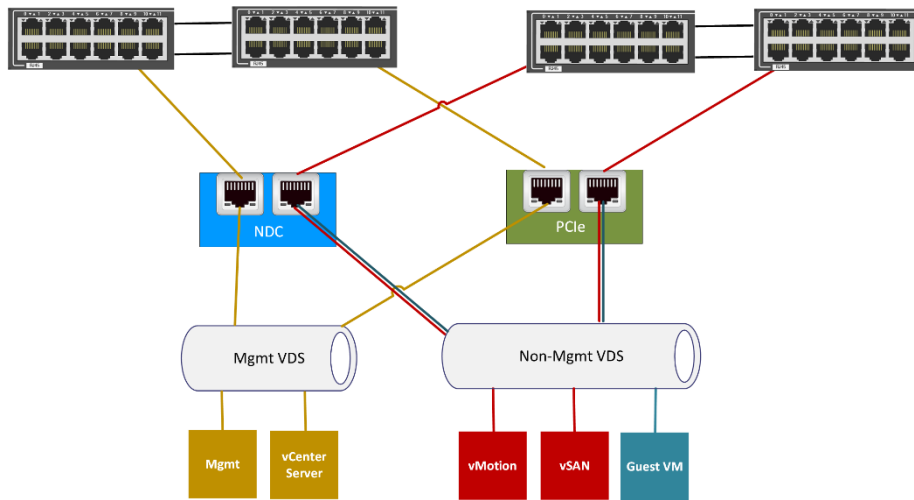


Figure 42. VxRail network segmentation with two virtual-distributed switches

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.

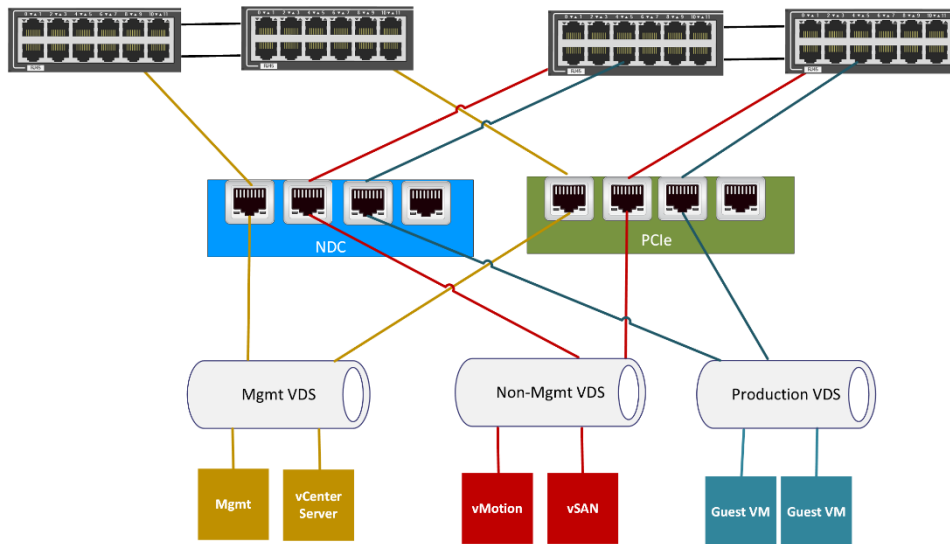


Figure 43. VxRail network segmentation with two virtual-distributed switches

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.

Decide on teaming and failover policies for VxRail networks

At the time of initial implementation, you will have the option to select and assign the teaming and failover policy on each portgroup for each required VxRail network. The following load-balancing policies are supported for VxRail clusters:

- *Route based on the originating virtual port*—After the virtual-distributed 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 adjust 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 this policy has a dependency on the logical link setting of the physical port adapters on the switch.

VxRail applies a default teaming and failover policy configuration based on the following rules:

- If the teaming policy is set to ‘Active-Active’, the default load balance policy set by VxRail is: *Route based on physical NIC load.*
- If the teaming policy is set to ‘active/standby’, the default load balance policy set by VxRail is: *Route based on originating virtual port.*

If the Ethernet adapters targeted to support vSAN support RDMA, then the following teaming policies are supported:

- Route based on originating virtual port
- Route based on source MAC hash
- Use explicit failover order

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.

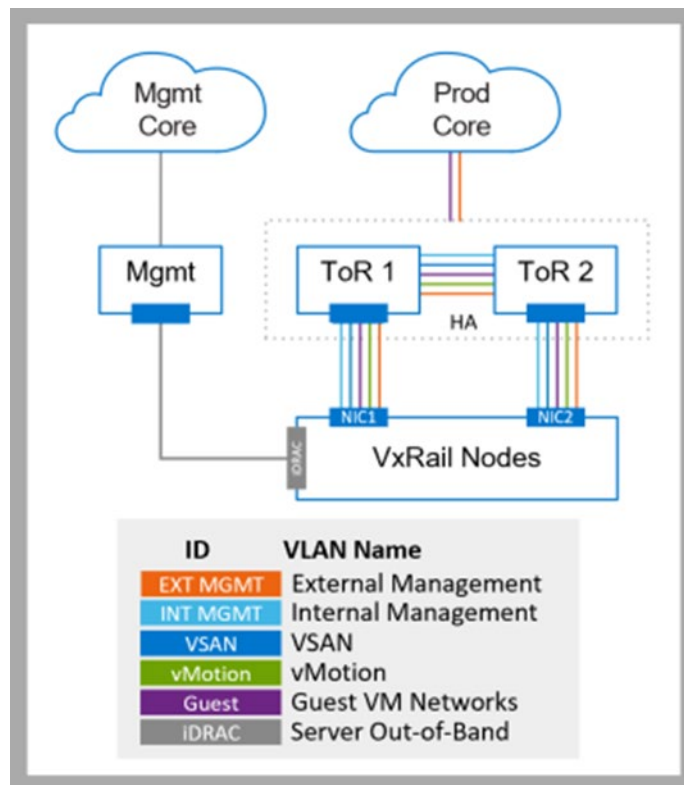


Figure 44. VxRail Logical Network Topology

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. A network to support I/O to the vSAN datastore is required unless you plan to use Fibre Channel storage as the primary storage resource with a dynamic cluster. In addition, 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, and this network can be skipped if you plan to use manual device discovery.

All the Dell PowerEdge servers that serve as the foundation for VxRail nodes include a separate Ethernet port that enables out-of-band 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, the vMotion network, and any guest networks. If your cluster will use vSAN as primary storage, then IP addresses are required for the vSAN network. You can also choose to segment the external management network to separate subnets for the physical and logical components. Decisions must be made on the IP address ranges reserved for each VxRail network.

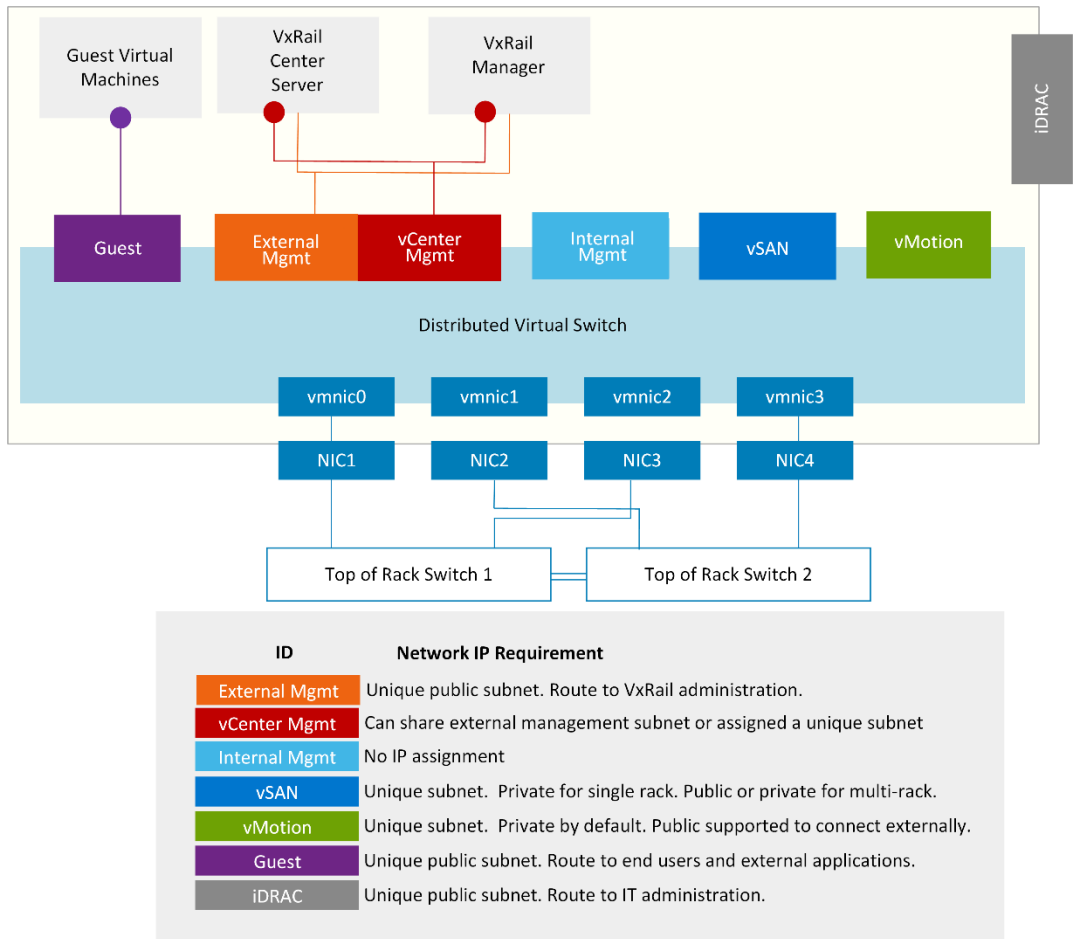


Figure 45. Overview of VxRail Core Network IP Requirements

- 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.
- By default, the external management network and vCenter management network share the same non-private, routable IP address range. The vCenter management network provides connectivity to VxRail Manager and the embedded vCenter instance. Starting with 7.0.350, if you do not want VxRail Manager and the vCenter instance to share the same subnet with the VxRail (ESXi) nodes, you can choose to assign a separate non-private, routable subnet to this network.
- The options for VxRail cluster primary storage require either Ethernet for vSAN, or another supported storage option for dynamic clusters. If your cluster will not be deployed with a vSAN datastore, then VxRail will not require a vSAN network, and this task can be skipped.
 - With a vSAN network, you can reserve a routable or non-routable (private) IP address range
 - If your plans include a multirack cluster, and you want to use a new IP subnet range in the expansion racks, reserve a routable IP address range.

- If you plan to configure the cluster into a vSAN HCI mesh topology to share vSAN storage resources, reserve a routable IP address range.
- If you are planning a VxRail cluster only with a local vSAN datastore, a non-routable IP address range is acceptable. In this case, traffic is passed only between the VxRail nodes that form the cluster.
- If you plan to configure a dynamic cluster to use external FC-based storage as primary storage, there is no need to reserve an IP address range.
- 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 determines 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.
- If you plan to have multiple, independent VxRail clusters, it is considered best practice to use different VLAN IDs across multiple VxRail clusters to reduce network traffic congestion.
- 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, vCenter Management Network, 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 network supports communications to the ESXi hosts, and also has common network settings with the **vCenter Management Network**. 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 **vCenter Management Network** hosts the VxRail Manager and the VxRail vCenter Server. By default, it also shares the same network settings as the **External Management** network. In this context, the physical ESXi hosts and the logical VxRail management components share the same subnet and share the same VLAN. Starting with version 7.0.350, this logical network can be assigned to a unique subnet and assigned a VLAN separate from the external management network.

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, 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.

If you plan to leverage vSAN for VxRail cluster storage resources, it is a best practice to configure a VLAN for the **vSAN** network. It is also a best practice to configure a VLAN for the **vSphere vMotion** network. 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 the 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, if applicable. (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. VM Network can be configured during the cluster build process, or after the cluster is built. (Enter 0 in the VLAN ID field for untagged traffic)
Network Configuration Table ✓ Row 7	Enter the vCenter Server Network VLAN ID (if different from the external management VLAN ID)

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.

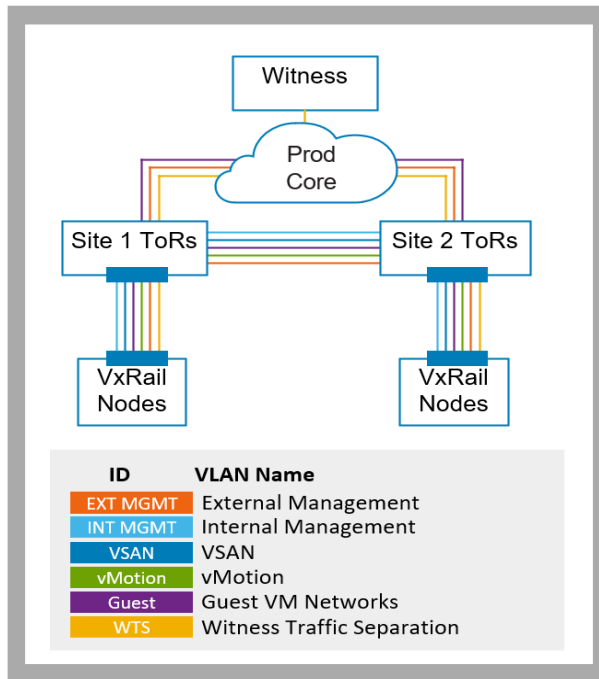


Figure 46. Logical network with Witness and Witness Traffic Separation

Network Configuration Table ✓ Row 77	Enter the Witness traffic separation VLAN ID.
---	---

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 optionally, the vCenter Server 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 some flexibility in how the IP addresses are assigned to the VxRail management components. For the ESXi hosts:

- You can manually assign the IP addresses to the ESXi hosts.
- You can have the IP addresses auto-assigned during VxRail initial build.

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 for VxRail management 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 assigns an IP address from the

unused reserve pool, or prompts 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, if you use the vCenter instance that is deployed in the VxRail cluster, 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.
- 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. 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 external connectivity:

Component	Condition
VxRail Node	One per VxRail Node

Component	Condition
VxRail Manager	One
vCenter	If you are supplying vCenter Server for VxRail: 0 If VxRail is supplying vCenter: 1 If VxRail is supplying PSC: 1
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

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 8	Enter the subnet mask for the VxRail External Management network.
Network Configuration Table ✓ Row 9	Enter the default gateway for the VxRail External Management network.

Plan network settings for vCenter Server management network

Prior to VxRail version 7.0.350, the VxRail-supplied vCenter and VxRail Manager were placed on the same VxRail external management network with the ESXi hosts, and shared the same VLAN and subnet. Starting with VxRail version 7.0.350, these virtual appliances can instead be assigned an IP address on a separate subnet and the subnet assigned a separate VLAN.

If you plan to deploy the virtual management components into a subnet separate from the ESXi hosts, capture the network properties.

Network Configuration Table ✓ Row 10	Enter the vCenter Server Network subnet mask
Network Configuration Table ✓ Row 11	Enter the vCenter Server Network gateway

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 27 and 28	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 29 and 32	Enter the IP addresses for the ESXi hosts.
---	--

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

Network Configuration Table ✓ Row 17	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 34	Enter the IP address for VxRail vCenter.
Network Configuration Table ✓ Row 36	Enter the IP address for VxRail Platform Service Controller (if applicable).

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 75	Enter IP address for Witness Management Network.
Network Configuration Table ✓ Row 76	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 78	Enter the IP address for the first of the two nodes in the 2-Node cluster.
Network Configuration Table ✓ Row 79	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. DNS is a requirement for VxRail, so select a domain where the naming services can support that domain.

Network Configuration Table ✓ Row 15	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 16	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 can use any host naming convention you want, provided that it is a legitimate format, or you can have 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 18–22	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 23–26	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-supplied 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 33	Enter an alphanumeric string for the new vCenter Server hostname. The domain that is specified will be appended.
Network Configuration Table ✓ Row 35	Enter an alphanumeric string for the new Platform Services Controller hostname. The domain that is specified will be appended.

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.

Network Configuration Table ✓ Row 12	Enter your time zone.
---	-----------------------

Set DNS for VxRail management components

Network Configuration Table ✓ Row 13	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 Manager is connected and is functioning properly.

Network Configuration Table ✓ Row 14	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 47. 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 48. 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 do not require hostnames, 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 VxRail 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 38	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 37	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 39	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 40	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 41	Enter the full VxRail management username/password. (For example, administrator@vsphere.local)
---	--

- The VxRail initialization process deploys 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 42	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 43	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 switches.

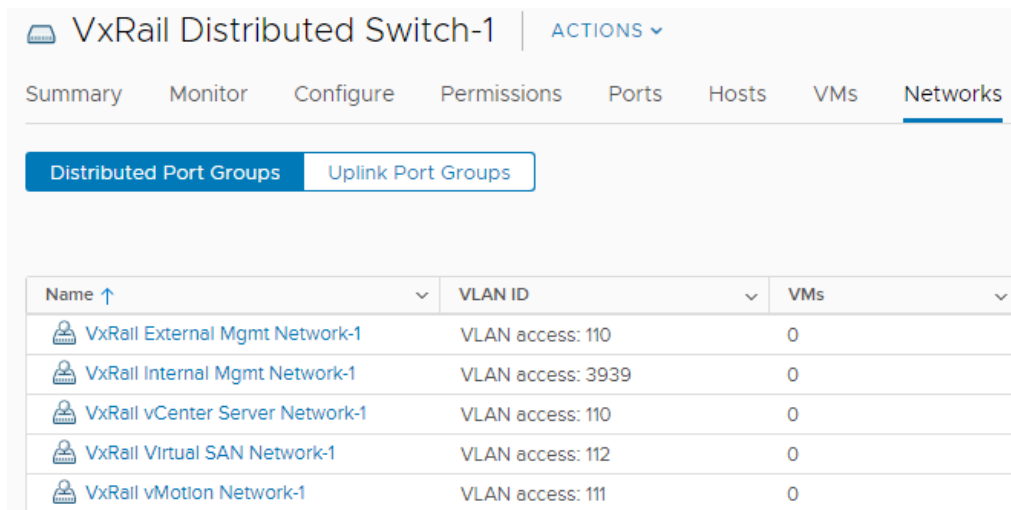
Before VxRail version 7.0.010, if you deployed the VxRail cluster on an external, customer-supplied vCenter, a virtual-distributed switch was configured on the vCenter instance as part of the initial cluster build process. The automated initial build process deployed the virtual-distributed switch adhering to VxRail requirements in the vCenter instance, and then attached the VxRail networks to the portgroups on the virtual-distributed switch. Depending on the target version planned for your VxRail cluster, you can 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 49. 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 44	Enter the name of the virtual-distributed switch that will support the VxRail cluster networking.
Network Configuration Table ✓ Row 45	If a decision is made to configure two virtual-distributed switches, enter the name of the second virtual-distributed switch.
Network Configuration Table ✓ Row 46	Enter the name of the portgroup that will enable connectivity for the VxRail external management network.
Network Configuration Table ✓ Row 47	Enter the name of the portgroup that will enable connectivity for the VxRail vCenter Server network.
Network Configuration Table ✓ Row 48	Enter the name of the portgroup that will enable connectivity for the VxRail internal management network.
Network Configuration Table ✓ Row 49	Enter the name of the portgroup that will enable connectivity for the vMotion network.
Network Configuration Table ✓ Row 50	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.

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 50. Sample LACP policy configured on virtual-distributed- switch.

- Configure the teaming and failover policy on the portgroups targeted for link aggregation.

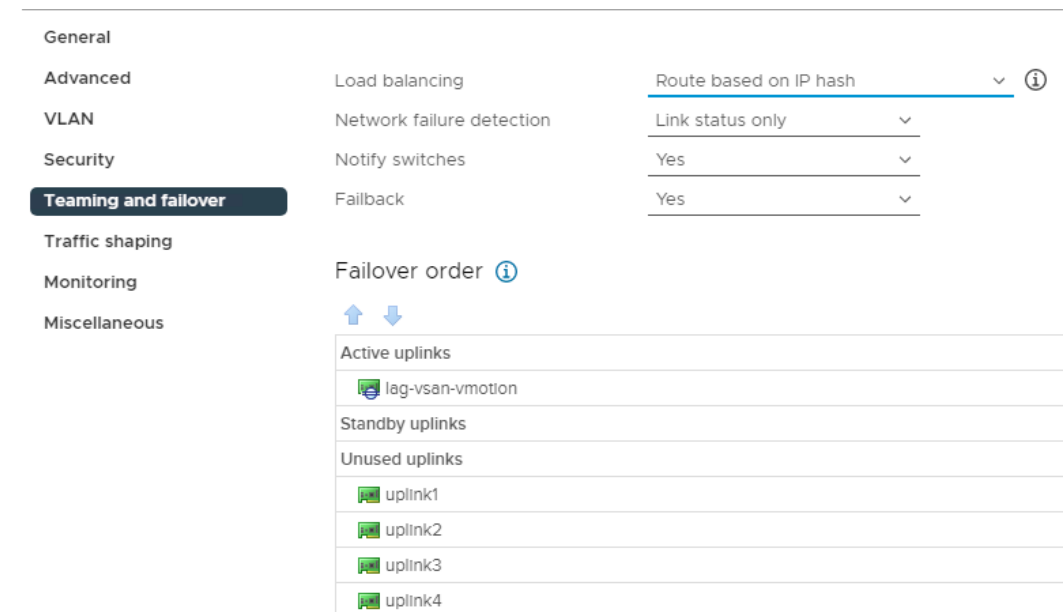


Figure 51. Sample LACP Policy configured as active uplink in teaming and failover policy

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 can either stretch the IP subnet for vMotion between the racks, or use routing services in your data center instead for a multi-rack cluster.

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

Network Configuration Table ✓ Rows 51-52	Enter the starting and ending IP addresses for vSphere vMotion.
---	---

For the manual assignment method, record the IP addresses.

Network Configuration Table ✓ Rows 53-56	Enter the IP addresses for vSphere vMotion.
---	---

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

Network Configuration Table ✓ Row 57	Enter the subnet mask for vMotion.
Network Configuration Table ✓ Row 58	Enter the gateway for vMotion.

Reserve IP addresses for VxRail vSAN network

You can skip this section if your plans do not include configuring vSAN as the primary storage for your VxRail cluster.

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

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 59-60	Enter the starting and ending IP addresses for vSAN.
---	--

For the manual assignment method, record the IP addresses.

Network Configuration Table ✓ Rows 61-64	Enter the IP addresses for vSAN.
---	----------------------------------

Enter the subnet mask and gateway for the vSAN network. You can use the default gateway assigned to the VxRail External Management network if you do not need to enable routing for this network, or enter a gateway to enable routing for the vSAN network.

Network Configuration Table ✓ Row 65	Enter the subnet mask for vSAN.
Network Configuration Table ✓ Row 66	Enter the gateway for vSAN. The default gateway can be updated if routing is required.

Decide on VxRail logging solution

Decide whether to use your own third-party syslog server, use the vRealize Log Insight solution, 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 the cluster is being deployed with a version greater than 7.0.010, VxRail will not deploy a vRealize Log Insight virtual appliance onto the cluster during initial build.
- It is possible to deploy vRealize Log Insight after the VxRail initial build operation is complete.

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 67	Enter the hostname for vRealize Log Insight
---	---

Network Configuration Table ✓ Row 68	Enter the IP address for vRealize Log Insight
---	---

If a syslog server already deployed in your data center will be used for logging, capture the IP address

Network Configuration Table ✓ Row 69	Enter the IP address of the syslog server
---	---

Assign passwords for VxRail management

You must assign a password to the accounts that are members of the VxRail management ecosystem. Use 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 uses 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 SmartFabric Services enablement

Note: Skip this section if you do not plan to enable Dell 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.

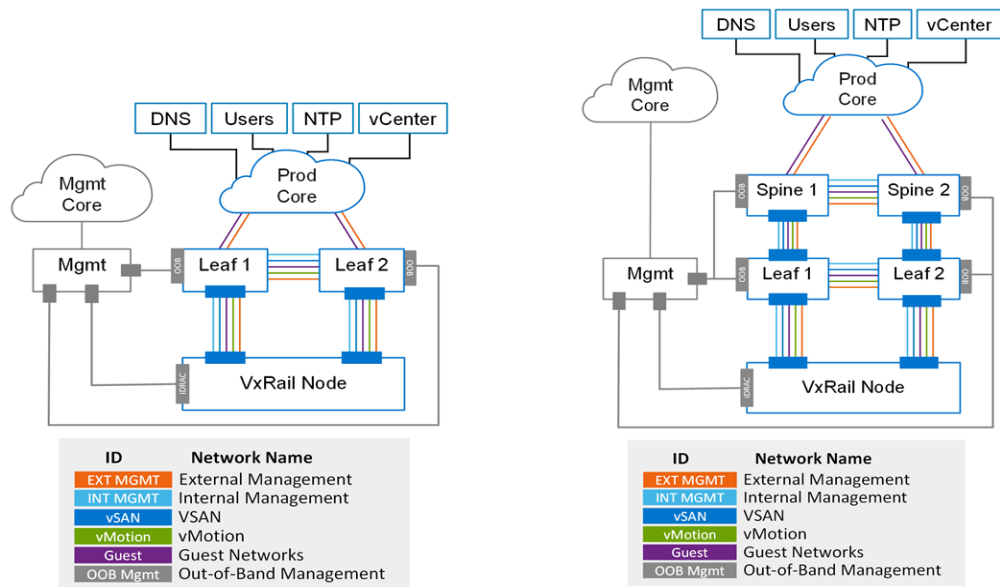


Figure 52. Logical networks for single-tier and two-tier SmartFabric deployments

The Dell 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 OMNI vCenter plug-in is required for each Dell switch fabric pair, and requires network properties to be set during the deployment process.

Network Configuration Table ✓ Rows 70 and 71	Reserve an IP address for out-of-band management of each switch in the SmartFabric-enabled network.
Network Configuration Table ✓ Row 72	Enter the IP address for Dell OMNI vCenter plug-in.
Network Configuration Table ✓ Row 73	Enter the subnet mask for Dell OMNI vCenter plug-in.
Network Configuration Table ✓ Row 74	Enter the gateway for Dell 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 [Planning Guide—Dell VxRail with SmartFabric Services](#) on infohub.delltechnologies.com.

Chapter 8 Planning for VxRail Satellite Nodes

This chapter presents the following topic:

Introduction.....	94
Plan networking to support VxRail satellite node management.....	90
Assign network settings to VxRail satellite nodes.....	91
Assign passwords for VxRail satellite nodes.....	91

Introduction

You can skip this section if your plans do not include the deployment of VxRail satellite nodes

VxRail satellite nodes are an option to support workload requirements at remote, or 'edge' locations. The server hardware and components for VxRail satellite nodes is the same product family as used for VxRail nodes to build clusters. The major difference is that the software installed at the factory into satellite nodes is not the same image as the software installed on nodes used for cluster formation. Therefore, satellite nodes can only be deployed as single instances, and cannot be used to form VxRail clusters.

The planning and preparation steps for VxRail satellite nodes differentiate from VxRail clusters in several areas:

- VxRail satellite nodes are dependent on an operational VxRail cluster for management and monitoring purposes. A VxRail satellite node that is not under the management of a VxRail Manager instance can be viewed as a stranded ESXi instance.
- vSAN networks are not supported on VxRail satellite nodes. VxRail satellite nodes depend on local disk drives and a local PERC controller for primary storage resources for virtual machines.
- vMotion networks for virtual machine mobility purposes are not supported on VxRail satellite nodes.
- VxRail Manager cannot discover VxRail satellite nodes using the internal management network. A management IP address must be assigned to a VxRail satellite node in order for it to be discovered by VxRail Manager.

Plan networking to support VxRail satellite node management

VxRail satellite nodes deployed at remote locations must be able to connect to a VxRail Manager instance on a previously deployed VxRail cluster to enable centralized monitoring and management. VxRail Manager, which is deployed on the external management network, must be able to discover every planned satellite node by IP address.




Figure 53. Routing between VxRail Manager and VxRail satellite nodes

Each VxRail satellite node is assigned a management IP address during the installation process. At initial power-on, the VxRail satellite node is considered stranded until such time as a connection with a VxRail Manager instance is established. Plan your data center network to allow the VxRail external management network to route externally, and work with your network administrators or service providers to support this network to route to each planned VxRail satellite node.

Assign network settings to VxRail satellite nodes

Every VxRail satellite node must be assigned an IP address and hostname for management purposes. Use the table in [Appendix A: VxRail Network Configuration Table](#) to capture these settings for up to three satellite nodes

Network Configuration Table ✓ Rows 80 and 81	Enter the management IP address and hostname for the first satellite node
Network Configuration Table ✓ Rows 82 and 83	Enter the management IP address and hostname for the second satellite node
Network Configuration Table ✓ Row 84 and 85	Enter the management IP address and hostname for the third satellite node

Assign passwords for VxRail satellite nodes

Since ESXi is installed into every VxRail satellite node at the factory, a password must be assigned for the 'root' account. Use the table in [Appendix B: VxRail Passwords](#) to capture the passwords for each satellite node.

The password rules for the 'root' account for satellite nodes is the same as for VxRail management components, and 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.

Chapter 9 Configure the Network for VxRail

This chapter presents the following topic:

- Introduction..... 94**
- Setting up the network switch for VxRail connectivity 94**
- Setting up the upstream network for VxRail connectivity..... 99**
- Confirm your data center network 101**
- Confirm your data center environment 102**

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 the VxRail nodes and power 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 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 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

Note: If you do not plan to use vSAN as the primary storage resource on the VxRail cluster, this task can be skipped.

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 vCenter Server Management VLAN (If different from VxRail External Management VLAN)
- 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 – in cases where vSAN is the primary storage resource. Ensure that unicast is enabled.
- vSphere vMotion VLAN
- VM Networks VLAN (can be configured after VxRail initial deployment)

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.

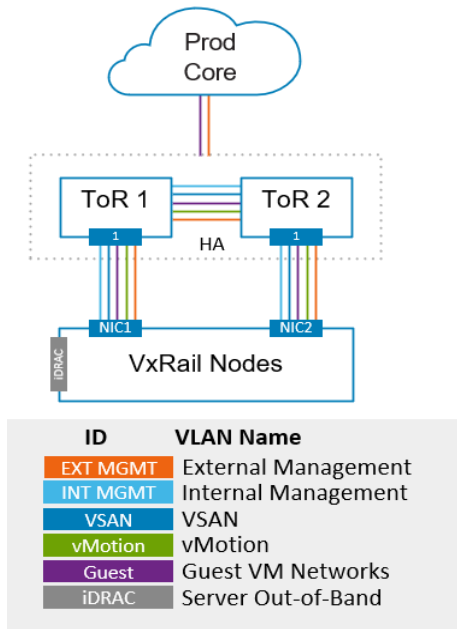


Figure 54. VxRail Logical Networks: Version 4.7 and later

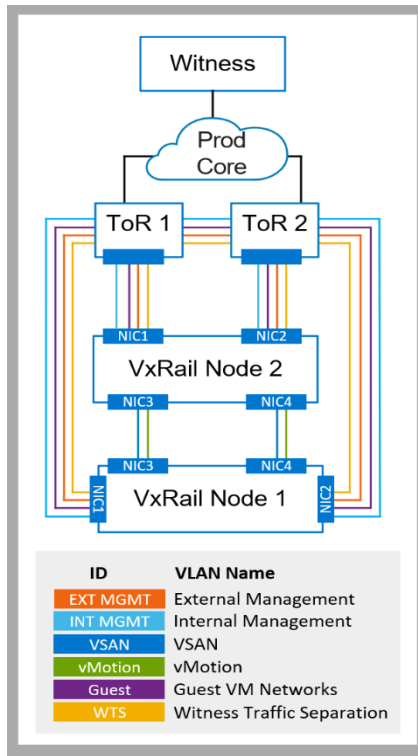


Figure 55. VxRail Logical Networks: 2-Node Cluster with Witness

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 to support device discovery.
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. If you choose to create a separate subnet for the vCenter Server Network, configure the **vCenter Server Network VLAN (Row 7)**, configure the VLAN on the switches.
8. Configure the optional **VxRail Witness Traffic Separation VLAN (Row 74)** on the switches ports if required.
9. 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)**, **vSAN VLAN (Row 4)** and **vCenter Server Network VLAN (Row 7)**. If a vSAN witness is required for the VxRail cluster, include the **VxRail Witness Traffic Separation VLAN (Row 74)** 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 **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 two or four 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 configures a virtual-distributed switch on the cluster, and then configures a portgroup on that virtual-distributed switch for each VxRail management network.




Figure 56. 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, link aggregation can be configured to support that network traffic. These can include any unused ports on the NDC/OCP 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 **vCenter Server Network VLAN (Row 7)** on the switch ports.
12. Configure the optional **VxRail Witness Traffic Separation VLAN (Row 74)** 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)**, any **VM Network VLANs (Row 6)**, and the optional **vCenter Server Network VLAN (Row 7)**. If a vSAN witness is required for the VxRail cluster, include the **VxRail Witness Traffic Separation VLAN (Row 74)** 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.

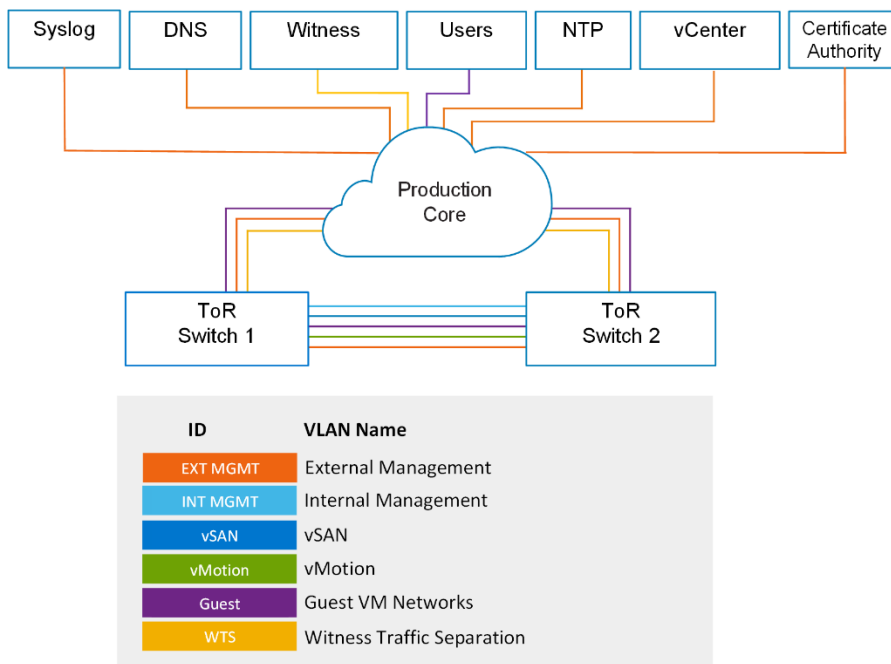


Figure 57. Logical networks connecting to upstream elements

If your Layer 2/Layer 3 boundary is at the lowest network tier (top-of-rack switch), 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), 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 'Ether Channels', to allow upstream passage of external VxRail networks.

Configure network to support RoCE

Note: This section is only relevant if you are planning to deploy a VxRail cluster with vSAN using RoCE-compliant Ethernet adapters.

If the VxRail nodes targeted for deployment are configured with Ethernet adapters that support RDMA over Converged Ethernet (RoCE), then the switches and supporting network need to be configured to enable a 'lossless' network for the vSAN traffic.

Converting a cluster to enable RoCE on the vSAN datastore is performed after the VxRail initial build. Consult the technical reference guides from your switch vendor for the specific steps to configure on the physical network. The basic step to ensure a 'lossless' network are as follows:

- vSAN with RDMA supports NIC failover, but does not support link aggregation or NIC teaming based on IP hash.
- Data Center Bridging must be enabled on the switches supporting the VxRail cluster
- Control traffic flow on the switches using a mechanism such as Priority Flow Control (PFC). Set the RoCE network to a higher priority as outlined in your vendor documentation.
- Configure RoCE on vSAN on a priority-enabled VLAN
- If the vSAN traffic will be on a routed Layer 3 network, then the 'lossless' network settings must be preserved when routed across network devices using a feature such as the Differentiated Serviced Code Point (DSCP) QoS setting.

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 performs 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 uses 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 VxRail can communicate with your syslog server if you plan to capture logging.
4. Verify that your IT administrators can communicate with the VxRail management system.
5. If you plan to use a customer-supplied vCenter, verify open communication between the vCenter instance and the VxRail managed hosts.
6. 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.
7. If you plan to deploy a separate network for ESXi host management (iDRAC), verify that your IT administrators can communicate with the iDRAC network.
8. 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. Confirm the vCenter Server management gateway IP is accessible, if configured.
6. 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.
7. If you have configured NTP servers, confirm that you can reach them from your configured VxRail external management network.
8. If you have configured a third-party syslog server for logging, confirm that you can reach it from the network supporting VxRail Manager.
9. If you plan to use a customer-supplied vCenter, confirm that it is accessible from the network supporting VxRail Manager.
10. If you plan to use a local certificate authority for certificate renewal on VxRail, verify that it is accessible from network supporting VxRail Manager.
11. If you plan to use Secure Connect Gateways to enable connectivity to the back-end Customer Support centers, verify that the gateways are accessible from network supporting VxRail Manager.
12. 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.
13. 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 10 Preparing to Build the VxRail Cluster

This chapter provides the following topics:

Introduction	105
Complete pre-requisites for dynamic clusters	105
Configure nodes for tagged VxRail management VLAN	105
Configure a jump host or laptop for VxRail initialization	105
Perform initialization to create a VxRail cluster	107

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.

Complete pre-requisites for dynamic clusters

Note: This section is only relevant if you are planning to deploy a dynamic cluster.

A dynamic cluster does not support a local vSAN datastore, and is therefore dependent on external storage resources to support cluster implementation and operations. A dynamic cluster will not fully complete implementation unless an external storage resource is made available to support virtual machine workload.

If you plan to deploy a dynamic cluster where the primary storage will be an FC-based VMFS datastore, confirm you have completed all the steps on your Fibre Channel network and storage array to present a LUN to the VxRail nodes.

If you plan to deploy a dynamic cluster where the primary storage resource is a remote vSAN datastore, confirm that your data center network has been properly configured so that the vSAN network on the cluster serving the vSAN datastore and the vSAN network on the dynamic cluster can connect.

If you plan to deploy a dynamic cluster with either iSCSI-based storage or NFS-based storage, confirm that your data center network is properly configured so the VxRail nodes can connect to the supporting storage array.

Configure nodes for tagged VxRail management VLAN

The VLAN pre-configured for the VxRail external management network on the nodes during the manufacturing process is zero, or the “native” VLAN. This “native” VLAN is ‘untagged’, which means all network traffic is allowed passage. If you choose to use a ‘tagged’ VLAN for the VxRail external management network, which restricts network passage, then the following steps must be taken:

- The VLAN must be configured on the trunked port as a ‘tagged’ VLAN.
- The default VLAN for the VxRail external management network must be changed on the virtual switch each of the VxRail nodes to the ‘tagged’ VLAN.

If you choose to use a ‘tagged’ VLAN, Dell professional services will change the VLAN on the VxRail nodes before performing cluster initialization.

Configure a jump host or laptop for VxRail initialization

The initialization process requires a desktop with an operating system such as Windows that can reach the VxRail external management network, and a supported web browser that can access VxRail management interface.

There are two options that can be used for this purpose:

- A workstation or laptop that can plug into an open Ethernet port on one of the top-of-rack switches
- A jump host ([Jump Server Description](#)) in your data center that can reach the VxRail external management network

If you choose the method to plug a workstation or laptop into an open Ethernet port on a switch, and you choose a VLAN other than the default 'native' VLAN for the VxRail external management network, first configure the port to enable connectivity to the VxRail management interface.

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 either the temporary, preconfigured VxRail initial IP address: 192.168.10.200/24, or have Dell services apply the permanent IP address to VxRail Manager. This temporary IP address will automatically change during VxRail initialization to your desired permanent address, and assigned to VxRail Manager during cluster formation.

Your workstation/laptop must be able to reach both the temporary VxRail initial IP address and the permanent VxRail Manager IP address (Row 15 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.

The following table is an example of how to set the IP addresses on the Ethernet port. The IP address settings on the jump host or laptop should be in the same subnet range as the temporary IP address and permanent IP address planned for VxRail Manager.

Example Configuration	VxRail	Jump host/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

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: Log in to a jump host, or 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 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 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.** If your planned deployment is a dynamic cluster, ensure that the necessary preparations for the selected external storage resource are complete.
- Step 5.** 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 6.** Attach Ethernet cables between the ports on the VxRail nodes and switch ports that are configured to support VxRail network traffic.
- Step 7.** 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 8.** Connect a workstation/laptop or jump host configured for VxRail initialization to access the VxRail external management network on your selected VLAN. It must be either plugged into the top-of-rack switch, or able to logically reach the VxRail external management network from elsewhere on your network.
- Step 9.** Open a browser to the VxRail initial IP address to begin the VxRail initialization process. This is either the default IP address assigned to VxRail Manager at the factory, or the permanent IP address set by Dell services.

- Step 10.** 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 11.** If you have enabled Dell SmartFabric Services, VxRail will automatically configure the switches that are connected to VxRail nodes using the information populated on the input screens.
- Step 12.** VxRail performs the verification process, using the information input into the menus.
- Step 13.** After validation is successful, the initialization process will begin to build a new VxRail cluster.
- Step 14.** The new permanent IP address for VxRail Manager will be displayed.
- If the permanent IP address was set on VxRail Manager before initialization, the browser will work seamlessly through the process.
 - 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 reenter 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 15.** Progress is shown as the VxRail cluster is built. The process takes about 25-40 minutes.
- Step 16.** 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 17.** Connect to VxRail Manager using either the VxRail Manager IP address (**Row 15**) or the fully qualified domain name (FQDN) (**Row 14**) that you configured on your DNS server. This will lead you to the vCenter instance.

Chapter 11 VxRail Network Considerations After Implementation

This chapter provides the following topics:

Introduction	110
Using unassigned physical ports for VxRail networks	110
Configure link aggregation on VxRail networks	111

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 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.

Using unassigned physical ports for VxRail networks

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 58. VxRail node with NDC/OCP ports and ports from optional PCIe adapter card

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

Starting configuration	Ending configuration
2 NDC/OCP ports	1 NDC/OCP port and 1 PCIe port
2 NDC/OCP Ports	2 NDC/OCP ports and 2 PCIe ports
4 NDC/OCP Ports	2 NDC/OCP ports and 2 PCIe ports
4 NDC/OCP Ports	1 NDC/OCP port and 1 PCIe port

The following rules apply for migrating VxRail networks from NDC/OCP-only ports to mixed NDC/OCP-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 this first reserved port.
- The switch ports enabling connectivity to the PCIe-based ports are properly configured to support VxRail network traffic.

- All 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/OCP ports can be reconfigured to run on one NDC/OCP 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 59. Enabling link aggregation for load-balancing

The following rules are applicable for enabling link aggregation:

- A minimum of four Ethernet ports per VxRail node is required for the VxRail networks.
 - Two ports configured to support VxRail management networks (Management/vCenter/VxRail discovery)
 - Two ports configured to support vSAN and vMotion.
- Can be all NDC/OCP ports, all PCIe ports, or a mixture of NDC/OCP and PCIe ports
 - If all your VxRail networks are on NDC/OCP-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.
- The top-of-rack switches support dynamic link aggregation with LACP.

Dynamic link aggregation requires:

- Dynamic port channels are configured on the adjacent top-of-rack switches.
- LACP is enabled on the adjacent top-of-rack switches.
- 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.

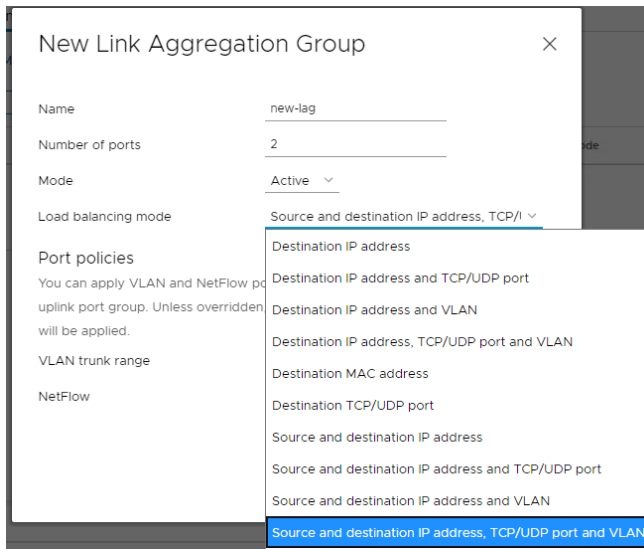


Figure 60. LACP Policy configuration on virtual-distributed- switch

Appendixes

This section presents the following topics:

- Appendix A: VxRail Network Configuration Table 115**
- Appendix B: VxRail Passwords 119**
- Appendix C: VxRail Setup Checklist 120**
- Appendix D: VxRail Open Ports Requirements 122**
- Appendix E: Virtual Distributed Switch Portgroup Default Settings 124**
- Appendix F: Physical Network Switch Examples 127**

Appendix A: VxRail Network Configuration Table

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

Row	Topic	Category	Description	
1	VxRail Networks	External Management VLAN	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 VLAN	This network traffic should stay isolated on the top-of-rack switches. The default VLAN ID is 3939.	
3		vMotion VLAN		
4		vSAN VLAN		
5		Guest Networks VLAN	Network Name	
6			VLAN	
7		vCenter Server Management VLAN	vCenter Server network on same VLAN as external management network by default. Can also be assigned a unique VLAN	
8	External Management Network Settings	Subnet Mask	Subnet mask for VxRail External Management Network	
9		Default Gateway	Default gateway for VxRail External Management Network	
10	vCenter Server Management Network	Subnet Mask	vCenter Server network on same subnet as external management network by default. Can optionally be assigned a new subnet mask.	
11		Gateway	vCenter Server network on same subnet as external management network by default. Can optionally be assigned a new gateway.	
12	System	Global Settings	Time zone	
13			NTP servers	
14			DNS servers	
15			Top Level Domain	
16	VxRail Manager	Hostname		
17		IP Address		
18	ESXi Hostnames	VxRail auto-assign method	Prefix	
19			Separator	
20			Iterator	

Appendixes

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

Row	Topic	Category	Description
51	vMotion	VxRail auto-assign method	Starting address for IP pool
52			Ending address for IP pool
53		Customer-supplied method	vMotion IP Address 1
54			vMotion IP Address 2
55			vMotion IP Address 3
56			vMotion IP Address 4
57		Subnet Mask	
58		Gateway	Default or vMotion stack to enable routing
59	vSAN	VxRail auto-assign method	Starting address for IP pool
60			Ending address for IP pool
61		Customer-supplied method	vSAN IP Address 1
62			vSAN IP Address 2
63			vSAN IP Address 3
64			vSAN IP Address 4
65		Subnet Mask	
66		Gateway	Default or Custom for routable network
67	Logging	Log Insight	vRealize Log Insight hostname
68			vRealize Log Insight IP address
69		Syslog Server	Syslog server IP address
70	SmartFabric	Switch out-of-band management	Out-of-band management IP address for switch 1
71			Out-of-band management IP address for switch 2
72		Dell OMNI plug-in	IP address
73			Subnet Mask
74			Gateway
75	Witness Site	Management IP Address	Witness management network IP address
76		vSAN IP Address	Witness vSAN network IP address
77	Witness Traffic Separation	WTS VLAN	Optional to enable Witness traffic separation on stretched cluster or 2-Node Cluster
78	2-Node Cluster	Node 1 WTS IP address	Must be routable to Witness
79		Node 2 WTS IP address	Must be routable to Witness
80	Satellite Nodes	Node 1	ESXi hostname
81			Management IP address

Appendixes

Row	Topic	Category	Description
82		Node 2	ESXi hostname
83			Management IP address
84		Node 3	ESXi hostname
85			Management IP address

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	

Item	Account	Password
Satellite Node #1	Root	
Satellite Node #2	Root	
Satellite Node #3	Root	

Appendix C: VxRail Cluster Setup Checklist

<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 switches: Ensure that your switches support VxRail requirements and provides the connectivity option that you chose for your VxRail nodes. Verify cable requirements.</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>	
Reserve VLANs	<ul style="list-style-type: none"> ✓ One external management VLAN ✓ 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 (before v4.5.0) for vSAN traffic ✓ One VLAN for vSphere vMotion ✓ One or more VLANs for your VM Guest Networks ✓ One VLAN for vCenter Server Network (if applicable) ✓ 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 a vCenter provided by VxRail. ✓ VxRail vCenter Server: Reserve IP addresses for vCenter Server and PSC (if supplied by VxRail). ✓ 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.

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, if using vSAN for primary storage
Logging	<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 Switches	<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, vCenter Server Network 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 or point-to-points links to carry networks requiring external connectivity upstream. ✓ 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 UI/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.

Note: VxRail manages the 'VxRail Customer Firewall Rules' interactive workbook. Access to the workbook requires Dell Technologies customer credentials. If you do not have Dell Technologies login credentials, contact your account team to download the tool for you.

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 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 configures one or two virtual-distributed switches on the selected vCenter instance using best practices for VxRail.

Default standard settings

For each VxRail network portgroup, VxRail initialization applies 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 initialization configures a teaming and failover policy for the port groups on the virtual-distributed switches. Overriding the default load-balancing policy is supported on VxRail.

Setting	Value
Load balancing (active-passive)	Route based on originating virtual port
Load balancing (active-active)	Route based on physical NIC load
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 virtual-distributed switches, and configure custom Network I-O Control (NIOC) settings for the following network traffic types. The settings depend 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 supports customizing the assignment of uplinks to VxRail networks, and also supports setting either an 'active/active' or 'active/standby' failover policy for the VxRail network portgroups. The following tables contain the default active/standby settings that are applied for the four predefined network traffic types that are required for initialization: Management, VxRail Management, vCenter Server Network, vMotion, and vSAN.

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	Unused	Active	Unused
VxRail Management	Active	Unused	Standby	Unused

Appendix F: Physical Network Switch Examples

VxRail enforces a predefined 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 predefined 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/OCP. 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, additional slots on the VxRail nodes must be reserved for PCIe adapter cards. If this is a current requirement or potential future requirement, 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 NDC/OCP

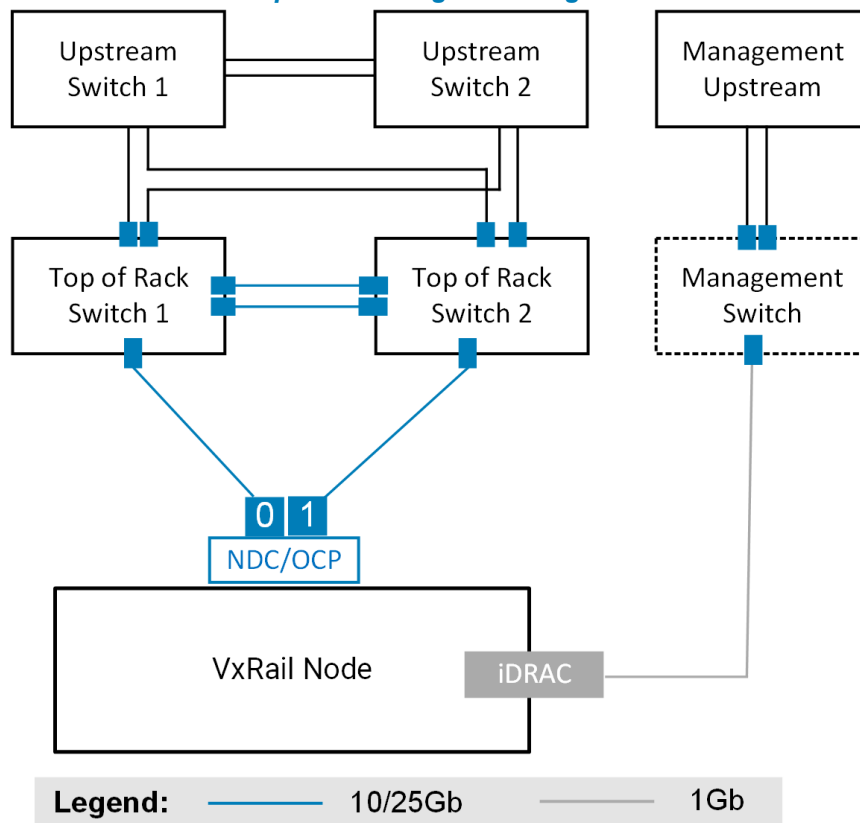


Figure 61. VxRail nodes with two 10Gb or 25Gb NDC/OCP ports connected to two TOR switches, and one optional connection to management switch for iDRAC

VxRail selects the two ports on the NDC/OCP to support VxRail networking. If the NDC/OCP on the VxRail nodes is shipped with four Ethernet ports, 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/OCP as a potential single point of failure.

Predefined network profile: 4x10gb or 4x25gb NDC/OCP

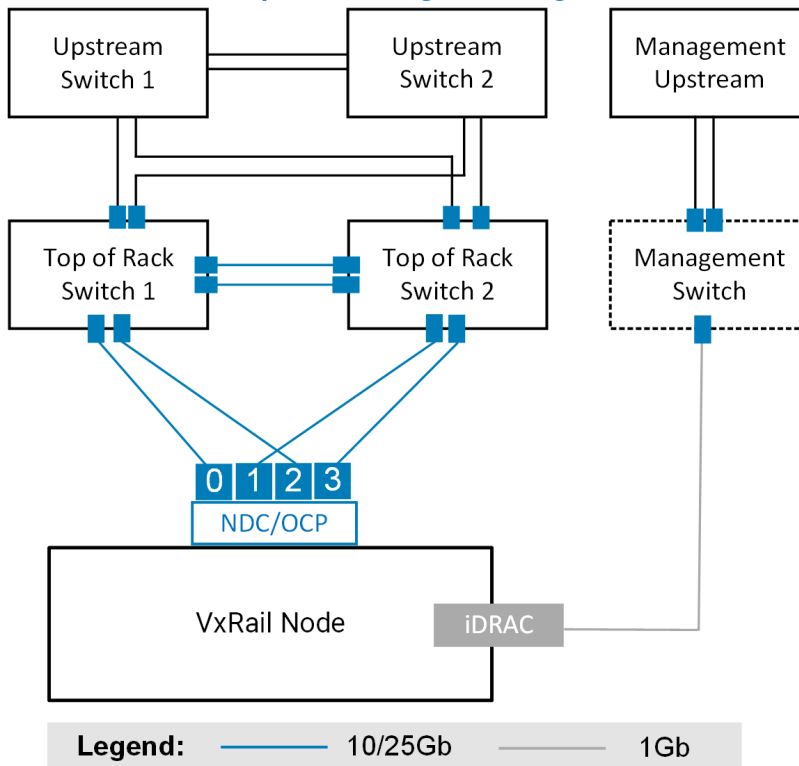


Figure 62. VxRail nodes with four 10gb NDC/OCP ports connected to 2 TOR switches, and one optional connection to management switch for iDRAC

VxRail selects all four ports on the NDC/OCP to support VxRail networking. The same number of cable connections should be made to each switch.

Predefined network profile: 2x10/25Gb NDC/OCF & 2x10/25Gb PCIe

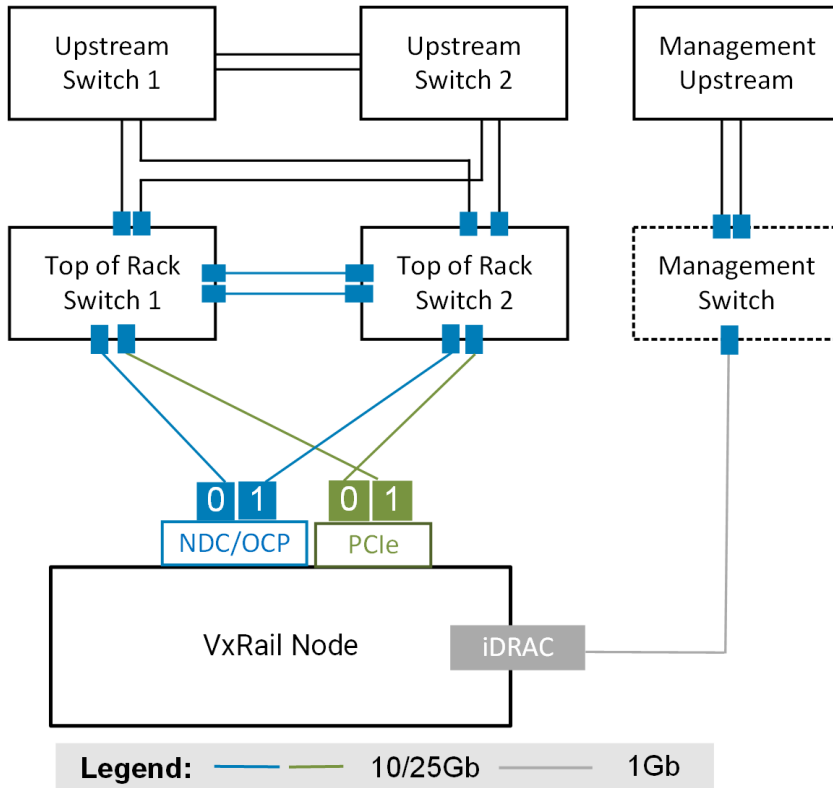


Figure 63. VxRail nodes with two 10/25Gb NDC/OCF ports and two 10/25Gb PCIe ports connected to 2 TOR switches, and one optional connection to management switch for iDRAC

In this option, two NDC/OCF 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/OCF 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/OCF or PCIe adapter card.

Custom option: Any NDC/OCP ports paired with PCIe ports

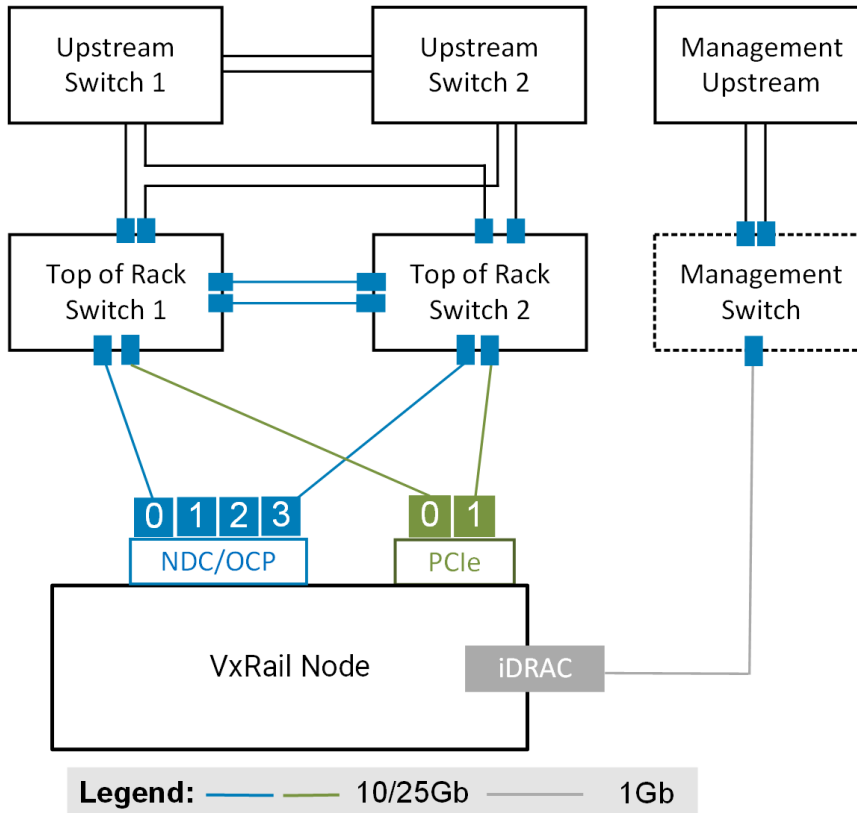


Figure 64. VxRail nodes with any two 10/25gb NDC/OCP ports and two 10/25gb PCIe ports connected to two TOR switch, and one optional connection to management switch for iDRAC

This is an example of an optional cabling setup for 2 NDC/OCP ports and 2 PCIe ports connected to a pair of 10gb switches or 25gb switches. Any NDC/OCP 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/OCF ports paired with PCIe ports other than the first slot

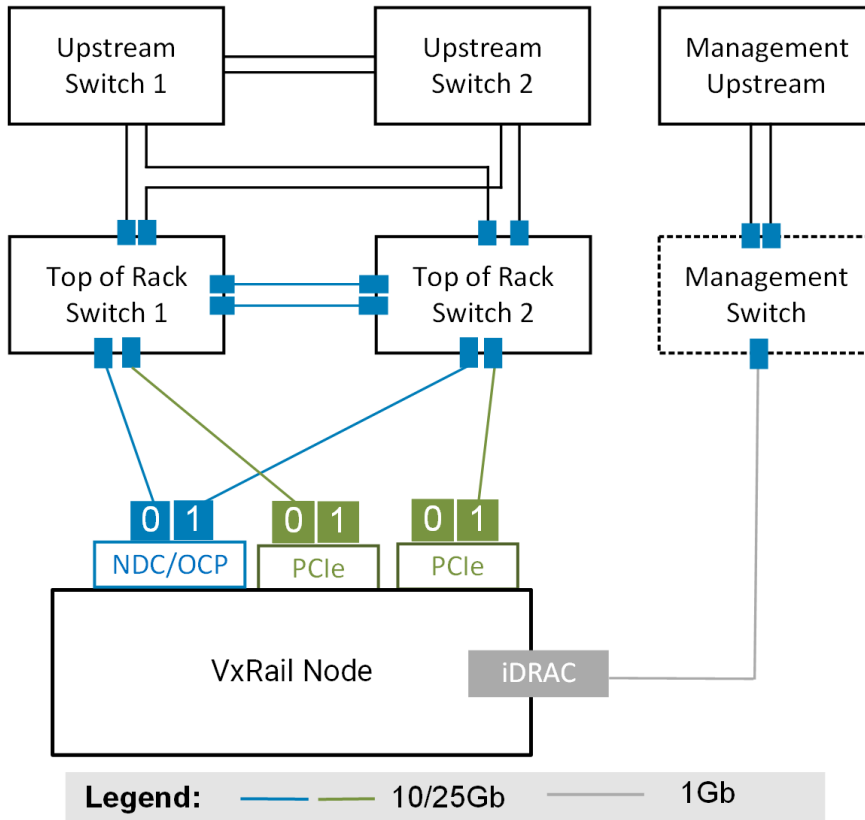


Figure 65. VxRail nodes with any two 10/25gb NDC/OCF ports and any two 10/25gb PCIe ports connected to 2 TOR switches, and one optional connection to management switch for iDRAC

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. The ports must be of the same type and running at the same speed.

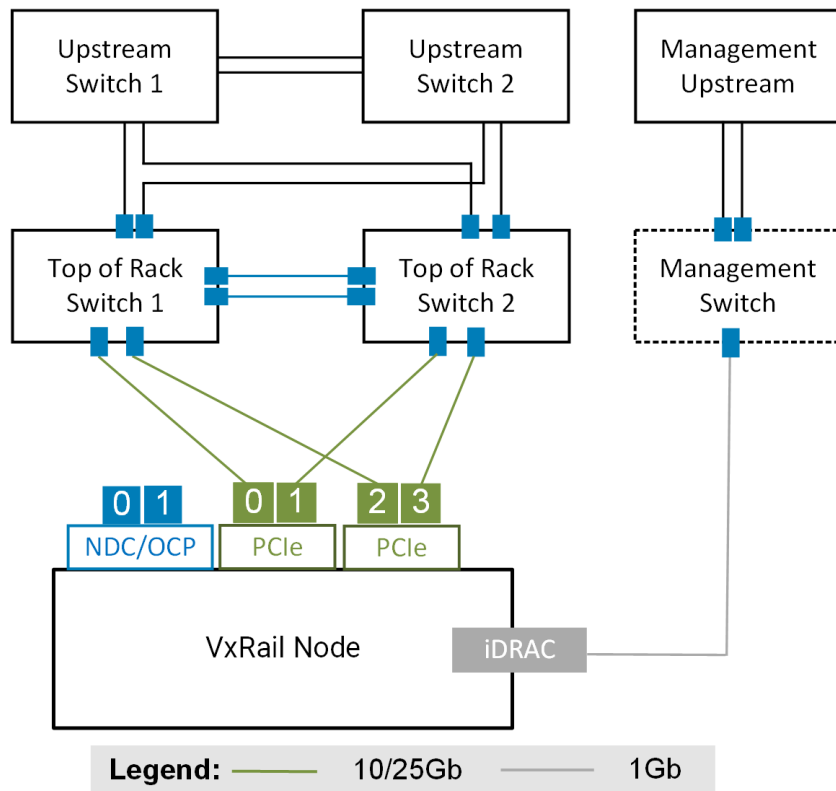


Figure 66. VxRail nodes with two or four PCIe ports connected to a pair of TOR switch, and one optional connection to management switch for iDRAC

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

Custom option: 6 ports

VxRail is most commonly deployed with 2 or 4 ports. For more network-intense workload requirements, VxRail can be deployed with 6 or even 8 network ports. This option supports spreading the VxRail networking between NDC/OCP ports and PCIe ports, and between ports on two different PCIe adapter cards.

In this topology, resource-intense workloads such as vSAN and vMotion can each have a dedicated Ethernet port instead of shared Ethernet ports. This prevents the possibility of saturation of shared Ethernet port resources.

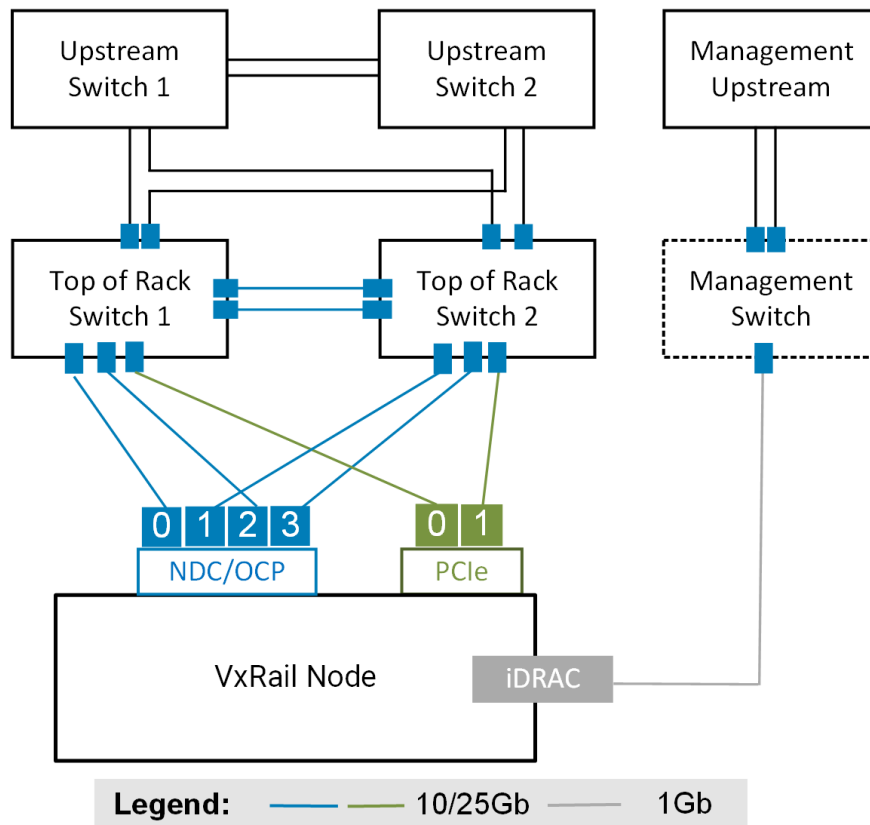


Figure 67. VxRail nodes with four NDC/OCP ports and a pair of PCIe ports connected to a pair of TOR switch, and one optional connection to management switch for iDRAC

With the six-port option, you can use more of the PCIe ports as opposed to the NDC/OCP ports. If your nodes have two PCIe slots occupied with network adapter cards, then this offers the flexibility to spread the VxRail networking workload across three network adapter cards.

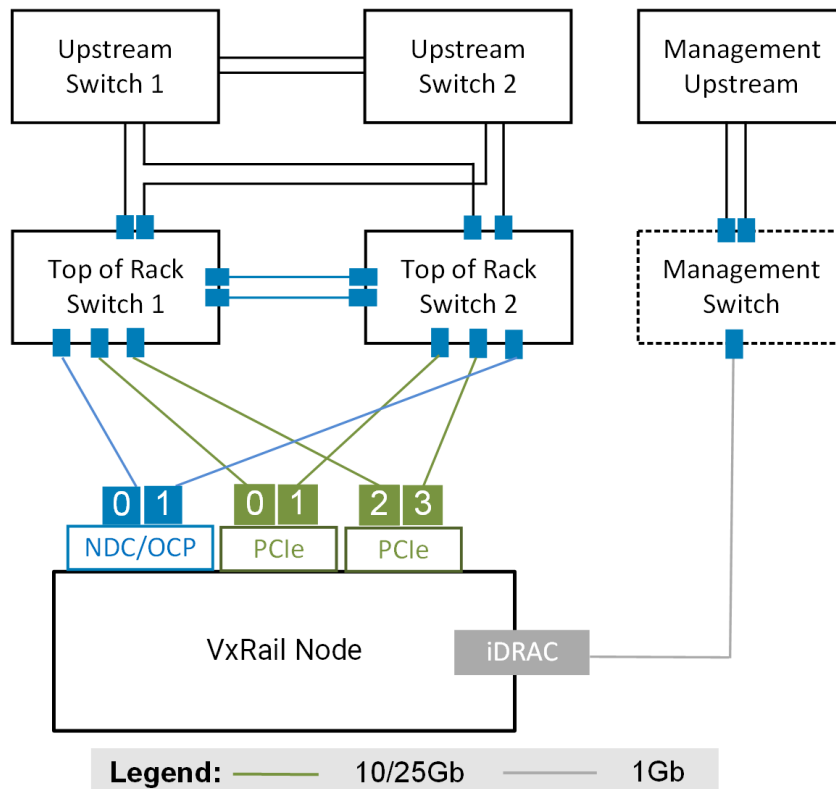


Figure 68. VxRail nodes with two NDC/OCP ports and ports from a pair of PCIe adapter cards connected to a pair of TOR switches, and one optional connection to management switch for iDRAC

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.

This option offers more flexibility in that it can support adapters at different speeds. The less resource-intensive networks can run at lower speeds, and more resource-intensive networks run at a higher speed. For example, the two NDC/OCP ports can be connected to 10gb switches to support the management networks, and the two ports on the PCIe can connect to a pair of 25gb switches to support non-management networks.

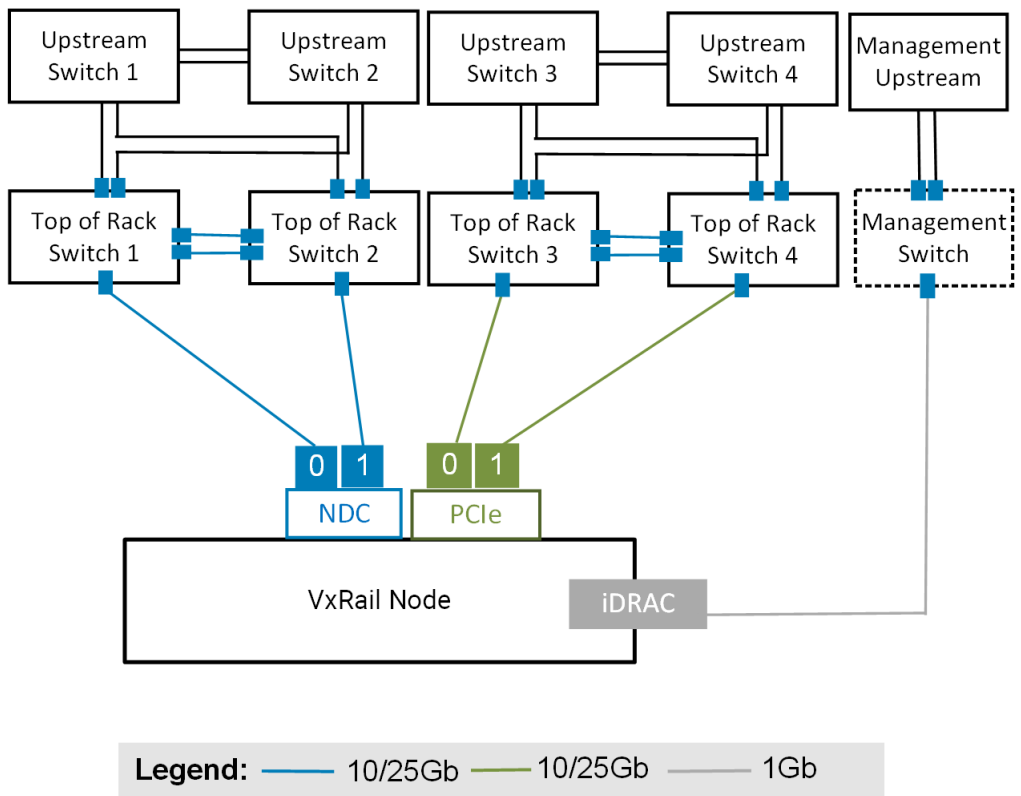


Figure 69. VxRail nodes with four ports connected to four TOR switches, and one optional connection to management switch for iDRAC