

Dell VxRail Network Planning Guide

H15300.31

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Revision history.....	7
Chapter 1: Introduction.....	9
Chapter 2: Plan the VxRail network.....	10
VxRail clusters and nodes.....	10
Network switch.....	12
VxRail relationship with the Ethernet switch	12
VxRail node discovery and the Ethernet switch	13
Basic switch requirements.....	13
Switch performance considerations.....	13
Storage network considerations.....	14
Network redundancy and performance considerations.....	14
Data center network.....	15
Routing services.....	16
VxRail networking port options.....	16
VxRail Ethernet adapter options.....	20
VxRail FC adapter options.....	20
VxRail RoCE adapter options.....	21
VxRail DPU options.....	21
VxRail node connectivity options.....	22
VxRail networking rules and restrictions.....	23
Topology and connections.....	24
Chapter 3: VxRail Cluster Types.....	26
Original vSAN architecture.....	26
Express vSAN architecture.....	27
Dynamic cluster.....	27
FC storage option.....	28
Remote VMware vSAN data store option.....	28
PowerFlex storage option.....	29
IP-based storage options.....	30
NVMe option.....	31
VMware vSAN stretched cluster.....	31
2-node cluster.....	32
Satellite nodes.....	35
Chapter 4: VxRail feature-driven decision points.....	36
Software-defined data center.....	36
VMware vSphere with Kubernetes on VxRail.....	37
Chapter 5: VxRail hardware and switch selection decision points.....	39
Chapter 6: Prepare the data center to implement VxRail	41
Prepare external network connectivity for VxRail.....	41

Prepare service connectivity for VxRail.....	42
Prepare for VMware vSphere+ subscription licensing.....	42
Prepare data center routing services.....	43
Prepare for multirack VxRail cluster.....	44
Prepare for VMware vSAN HCI mesh topology.....	46
Prepare external FC storage for dynamic clusters.....	48
Prepare for VxRail custom uplink assignments.....	49
Prepare data center network MTU.....	51
Prepare for LAG of VxRail networks.....	52
Identify switch ports to be configured for LAG.....	55
Plan LAG on switch port pairs.....	55
Prepare certificate authority server for VxRail.....	56
Identify isolation IP addresses for VMware vSphere High Availability.....	56
Chapter 7: Plan the VxRail cluster implementation.....	58
Decide on VxRail single point of management.....	58
Decide on VxRail network traffic segmentation.....	59
Decide on teaming and failover policies for VxRail networks.....	60
Plan the VxRail logical network.....	61
IP address considerations for VxRail networks.....	62
VLAN considerations for VxRail networks	63
Plan network settings for VxRail management components.....	65
Plan network settings for VMware vCenter Server management network.....	66
Identify IP addresses for VxRail management components.....	66
Select hostnames for VxRail management components.....	67
Select a top-level domain.....	67
Select a hostname for VxRail Manager	67
Select the ESXi hostnames.....	68
Select a hostname for the VxRail-managed VMware vCenter Server.....	69
Identify external applications and settings for VxRail.....	69
Set the DNS for VxRail management components.....	69
Prepare the customer-managed VMware vCenter Server.....	70
Prepare a customer-managed virtual-distributed switch.....	71
Prepare LAG on a customer-managed virtual-distributed switch.....	73
Reserve IP addresses for VxRail-managed VMware vSphere vMotion network.....	74
Reserve IP addresses for VxRail vSAN network.....	75
Decide on VxRail logging solution.....	75
Assign passwords for VxRail management.....	76
Chapter 8: Planning vSAN Witness Networking.....	77
Overview of vSAN Witness for VxRail stretched cluster.....	77
Overview of vSAN Witness for VxRail 2-node cluster.....	78
Networking rules for vSAN witness and VxRail.....	78
Reserve network settings for vSAN witness networks.....	80
Chapter 9: Planning VxRail Satellite Nodes Networking.....	82
Plan networking to support VxRail satellite node management.....	82
Assign network settings to VxRail satellite nodes.....	83
Assign passwords for VxRail satellite nodes.....	83

Chapter 10: Configure the Network for VxRail.....	84
Setting up the network switch for VxRail connectivity.....	84
Configure multicast for VxRail internal management network.....	84
Configure unicast for VxRail vSAN network.....	84
Configure VLANs for the VxRail networks.....	85
Configure the inter-switch links.....	87
Configure switch port mode.....	87
Configure LAG.....	87
Limit spanning tree protocol.....	88
Enable flow control.....	88
Set up the network switch ports for VxRail connectivity.....	88
Set up the upstream network for VxRail connectivity.....	89
Configure your network to support RoCE.....	90
Confirm your data center network.....	90
Confirm your firewall settings.....	91
Confirm your data center environment.....	91
Chapter 11: Prepare to build the VxRail cluster.....	92
Complete prerequisites for dynamic clusters.....	92
Configure nodes for tagged VxRail management VLAN.....	92
Configure a jump host or laptop for VxRail initialization.....	92
Perform initialization to create a VxRail cluster.....	93
Chapter 12: VxRail network considerations after implementation.....	95
Configure LAG on VxRail networks.....	95
Appendix A: Appendix A: VxRail Network Configuration Table.....	98
Appendix B: Appendix B: VxRail Passwords.....	101
Appendix C: Appendix C: VxRail cluster setup checklist.....	102
Appendix D: Appendix D: VxRail Open Ports Requirements.....	104
Appendix E: Appendix E: VMware VDS port group default settings.....	106
Default standard settings.....	106
Default teaming and failover policy.....	106
Default network I-O control (NIOC).....	107
Default failover order policy.....	107
Appendix F: Appendix F: Physical Network Switch Examples.....	109
Pre-defined network profile: 2x10gb or 2x25gb from a single NDC/OCP.....	110
Predefined network profile: 4x10gb or 4x25gb NDC/OCP.....	111
Predefined network profile: 2 x 10/25 GB NDC/OCP and 2 x 10/25 GB PCIe.....	112
Custom option: Any NDC/OCP ports paired with PCIe ports.....	113
Custom option: Two NDC/OCP ports paired with PCIe ports other than the first slot.....	114
Custom option: PCIe ports only.....	114

Custom option: Six ports.....	115
Custom option: Eight ports.....	117
Custom option: Eight ports connected to four Ethernet switches.....	118

Revision history

Date	Revision	Description
October 2023	H15300.31	<ul style="list-style-type: none"> Updated section on service connectivity. Added support for 10GbE NICs on P570 nodes. Added support for 10GbE for VMware vSAN ESA
August 2023	H15300.30	Added support for VxRail VE-660 and VP-760 models.
June 2023	H15300.29	Added content: <ul style="list-style-type: none"> Support for vSAN HCI mesh with vSAN Express Storage Architecture (ESA) Support for HCI mesh with VxRail stretched cluster
May 2023	H15300.28	New content for LAG is supported on VxRail-managed VMware VDS.
April 2023	H15300.27	Updates for hardware branding and removed SmartFabric content.
March 2023	H15300.26	<ul style="list-style-type: none"> Included content for release of VxRail 15G VD-Series hardware platform. Added sections to provide networking preparation guidance for connecting to Dell and VMware external sites. Added chapter for networking requirements for vSAN witness. Removed references to Platform Services Controller. Updated SmartFabric content based on new VxRail rules.
January 2023	H15300.25	Support for DPUs and ESXio in VxRail.
December 2022	H15300.24	<ul style="list-style-type: none"> Support for vSAN Express Architecture. Support for SmartFabric Services in integrated mode and decoupled mode.
November 2022	H15300.23	Support VxRail-managed VMware vCenter Server on a 2-node cluster.
August 2022	H15300.22	Support for new features in 7.0.400.
March 2022	H15300.21	Support for new features in 7.0.350.
January 2022	H15300.20	Support for new 15G VxRail models.
December 2021	H15300.19	<ul style="list-style-type: none"> Update dynamic cluster content with link to Dell published guide. Update content for VxRail Manager network exclusions.
November 2021	H15300.18	Support for PowerFlex as external storage for dynamic cluster.
October 2021	H15300.17	Support for satellite nodes.
August 2021	H15300.16	Support for new features in 7.0.240.
June 2021	H15300.15	<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) are unsupported for VxRail networking.
April 2021	H15300.14	<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.
February 2021	H15300.13	Support for new features in 7.0.131.
November 2020	H15300.12	Removed requirement for VxRail guest network during initial configuration.

Date	Revision	Description
October 2020	H15300.11	Support for new features in VxRail 7.0.100 and removed references to VMware Log Insight.
September 2020	H15300.10	Outlined best practices for link aggregation on non-VxRail ports.
August 2020	H15300.9	Updated requirement for NIC redundancy enablement.
July 2020	H15300.8	Support for new features in VxRail 7.0.010.
June 2020	H15300.7	Updated networking requirements for multirack VxRail clusters.
May 2020	H15300.6	Updated switch requirements for VxRail IPv6 multicast.
April 2020	H15300.5	Support for: <ul style="list-style-type: none"> ● VxRail SmartFabric multirack switch network. ● Optional 100 GbE Ethernet and FC network ports on VxRail nodes.
March 2020	H15300.4	Support for new functionality in VMware vSphere 7.0.
February 2020	H15300.3	Support for new features in VxRail 4.7.410.
August 2019	H15300.2	Support for VxRail 4.7.300 with Layer 3 VxRail networks.
June 2019	H15300.1	Support for VxRail 4.7.210 and updates to 25 GbE networking.
April 2019	H15300	<ul style="list-style-type: none"> ● First inclusion of this version history table. ● Support of VMware Cloud Foundation on VxRail.

Introduction

This guide provides the network details for VxRail deployment planning that include best practices recommendations, and requirements for both physical and virtual network environments.

VxRail is an HCI solution that consolidates compute, storage, and network into a single and unified system. With careful planning, you can deploy VxRail into an existing data center environment to immediately deploy applications and services.

VxRail 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 VxRail are managed as a single shared pool. They are allocated to applications and services that are based on defined business and operational requirements.

VxRail scale-out architecture leverages VMware vSphere and VMware vSAN to provide server virtualization and software-defined storage, with simplified deployment, upgrades, and maintenance through VxRail Manager.

Network connectivity is fundamental to the VxRail clustered architecture. Through logical and physical networks, 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:

- VxRail Manager
- VMware vCenter Server
- VMware vSAN
- VMware vSphere

Audience

This document has been prepared for anyone who is involved in planning, installing, and maintaining VxRail, including Dell 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 service representative to perform the installation.

Licenses

The VxRail includes the following Dell RecoverPoint for Virtual Machines licenses that you can download, install, and configure:

- Five full VM licenses per single node
- Fifteen full VM licenses for the G Series chassis

Temporary evaluation licenses are provided for VMware vSphere and VMware vSAN. The VMware vSphere licenses can be purchased from Dell, VMware, or your preferred VMware reseller partner.

The Deployment Guide contains additional information about licenses.

Plan the VxRail network

Take the time to plan out your data center network for VxRail. The network considerations for VxRail are the same as that of any enterprise IT infrastructure: availability, performance, and extensibility.

VxRail is manufactured in the factory per your purchase order, and delivered to your data center ready for deployment. The VxRail nodes can connect to any compatible network infrastructure to enable operations. Follow all the guidance and decision points that are described in this document. If there are separate teams for network and servers in your data center, work together to design the network and configure the switches.

The following planning recommendations ensure a proper deployment and functioning of your VxRail:

- Select the VxRail hardware and physical network infrastructure that best aligns with your business and operational objectives.
- Plan and prepare for VxRail implementation in your data center before product delivery.
- Set up the network switch infrastructure in your data center for VxRail before product delivery.
- Prepare for physical installation and VxRail initialization into the final product.

The VxRail nodes connect to one or more network switches to form a VxRail cluster. VxRail communicates with the physical data center network through one or more VMware VDS that is deployed in the VxRail cluster. The VMware VDS integrates with the physical network infrastructure to 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 VMware vSAN data store. It also provides connectivity to applications and end-users outside of the VxRail cluster.

The following are the physical components and selection criteria for VxRail clusters:

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

VxRail clusters and nodes

VxRail consists of server nodes that are designed and engineered for VxRail. The VxRail physical node is a PowerEdge server that goes through VxRail product engineering specifications to produce a VxRail node ready for shipment. The set of components that are manufactured into VxRail nodes is based on the 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 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, capacity, and network connectivity.

VxRail models

The naming standards for VxRail models are structured to target specific objectives

- E Series: Balanced compute and storage.
- V Series: Virtual desktop enablement with support for GPUs.
- P Series: High performance.
- S Series: Dense storage.

Newer VxRail new models use a naming standard that is more closely aligned with PowerEdge naming standards. This new naming standard now uses a two-character prefix instead of a single character, and consolidates three target objectives from the previous standard under a single model series.

- VE-Series: Balance compute and storage requirements.
- VP-Series: Support multiple business objectives, including high performance, storage density, and GPUs
- VD-Series: Ruggedized, short-depth, sled-based models for space optimization

Each VxRail model series offers choices for network connectivity. Every VxRail node includes support for base connectivity through onboard Ethernet ports. Network expansion is supported on VxRail nodes using available PCIe slots for Ethernet adapter cards. The number and type of PCIe slots that can be used for networking purposes depends on the model that is selected and the hardware configuration ordered.

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

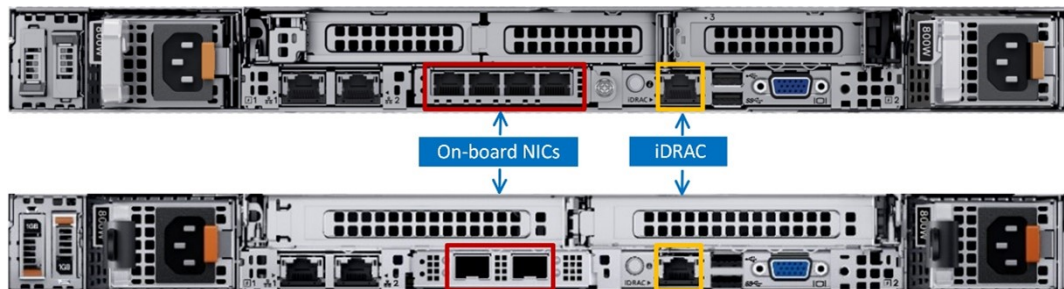


Figure 1. Back view of VxRail E-Series VE-Series nodes

With the 1U rack mount models, such as the VE-Series and E-Series models, PCIe slots can be used for network expansion. Certain models support both low-profile PCIe slots and full-height PCIe slots. The number and type of slots on each node vary depending on the hardware options that are configured, such as whether to include support for GPUs.




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

The 2U models offer a higher number of PCIe slots in contrast to the 1U models. However, the 2U rackmount models can be configured at the time of ordering to address a wider range of use cases, which can reduce the number of available PCIe slots. For instance, network expansion options would be reduced in cases where those slots instead are populated with storage devices or GPUs.




Figure 3. Back view of VxRail VD-Series rackable and stackable chassis options

The VxRail VD-Series support sleds of 1U in size and 2U in size, and both sled options are supported in either a rackable chassis or stackable chassis. Only the 2U sleds support network expansion outside of the onboard Ethernet ports through PCIe slots.

Network switch

A VxRail cluster depends on adjacent ToR Ethernet switches to support cluster operations.

VxRail is compatible with most Ethernet switches. For best results, select a switch platform that meets the operational and performance criteria for your planned use cases.

Data center network

VxRail is dependent of specific data center services to implement the cluster and for day-to-day operations. Configure the ToR switches to the upstream network to enable connectivity to these data center services, and to enable connectivity to the end-user community.

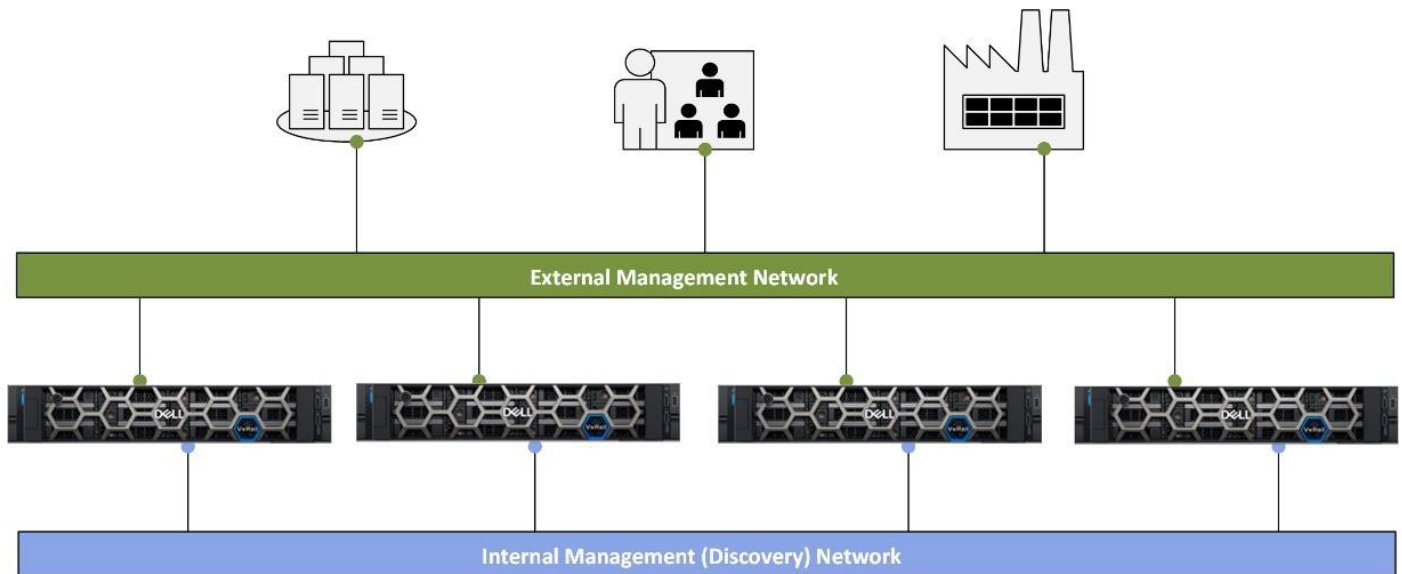
VxRail relationship with the Ethernet switch

The VxRail product does not have a backplane, so the adjacent ToR 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 ToR 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 developed to enable efficiency with the physical ToR 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.



The network that is visible only to the VxRail nodes depends on IPv6 multi casting services configured on the adjacent ToR 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 Layer3 services. You can enable routing services further upstream on the network infrastructure, or enable routing services at this ToR 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.

Here are some use cases:

- If your plans include deploying all flash storage on your VxRail cluster, 10 GbE network switches are the minimum requirement for this feature. Use 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

Consider the additional feature requirements while selecting Ethernet switches for your VxRail cluster depending on interoperability requirements for different types of storage resources.

If your VxRail cluster includes adapters that support RoCE (RDMA over Converged Ethernet) for VMware 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 DCB must support bandwidth allocation based on priority settings, which are 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

Determine whether 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. Two or more switches are required to support sustained performance, high availability, and fail over in production environments.

VxRail is a software-defined data center which relies on the physical ToR switching for network communications. It is engineered to enable full redundancy and failure protection across the cluster. For 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 LAG 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 7.0.130 and later. These logical port pairings can peer with a pair of ports on the adjacent switches to enable the load-balancing of demanding VxRail networks.

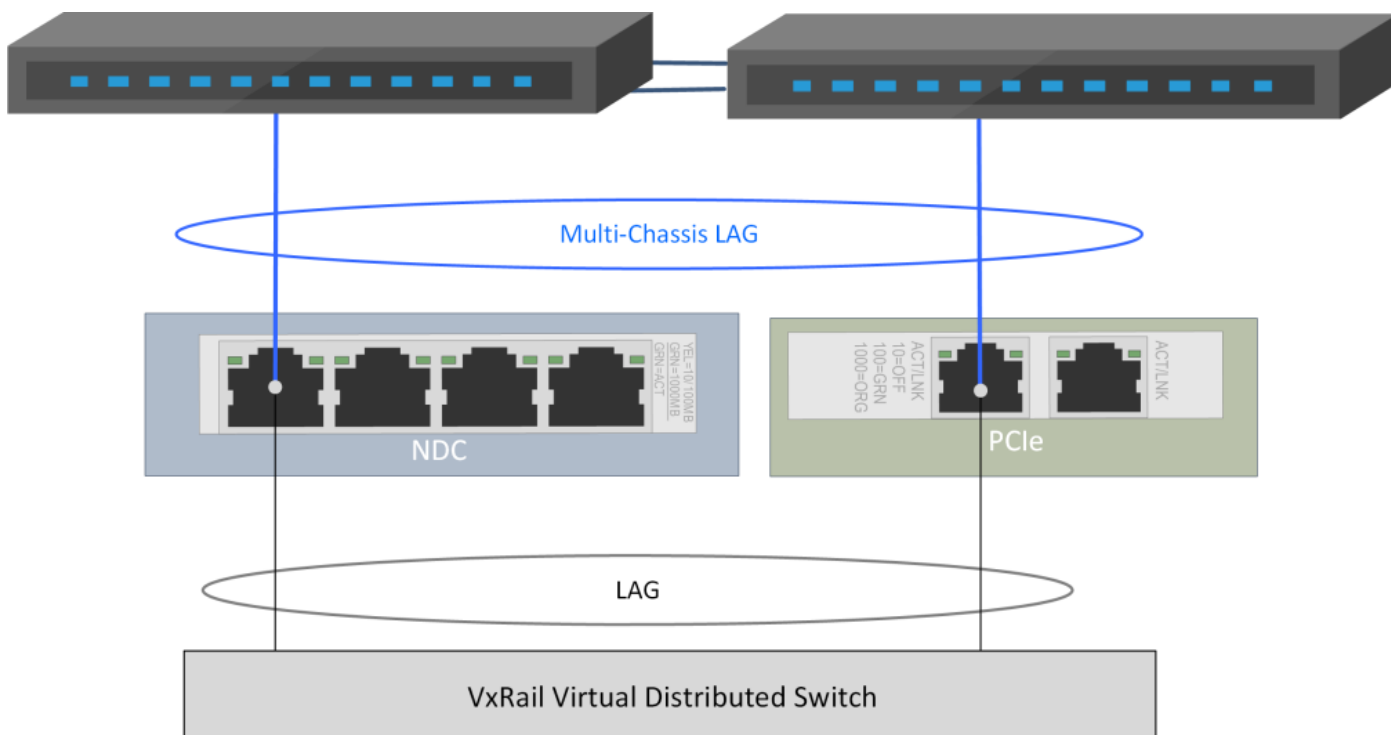


Figure 4. Multi-chassis LAG across two switches

For network-intensive workloads that require high availability, consider switches that support multi-chassis LAG, such as Cisco Virtual Port Channel or Dell 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 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 ToR 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

Requirements and options for data center services include:

- 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 requires synchronization of the clock settings on the various VxRail components to ensure proper operations. Dell Technologies recommends using a reliable global timing service for VxRail such as Network Time Protocol (NTP) for this purpose.
- Syslog service is supported with VxRail, but is not required.
- VxRail depends on VMware vCenter Server for cluster management and operations. You can use either the embedded VMware vCenter Server instance that is included with VxRail, or a customer-managed VMware vCenter Server instance in your data center.

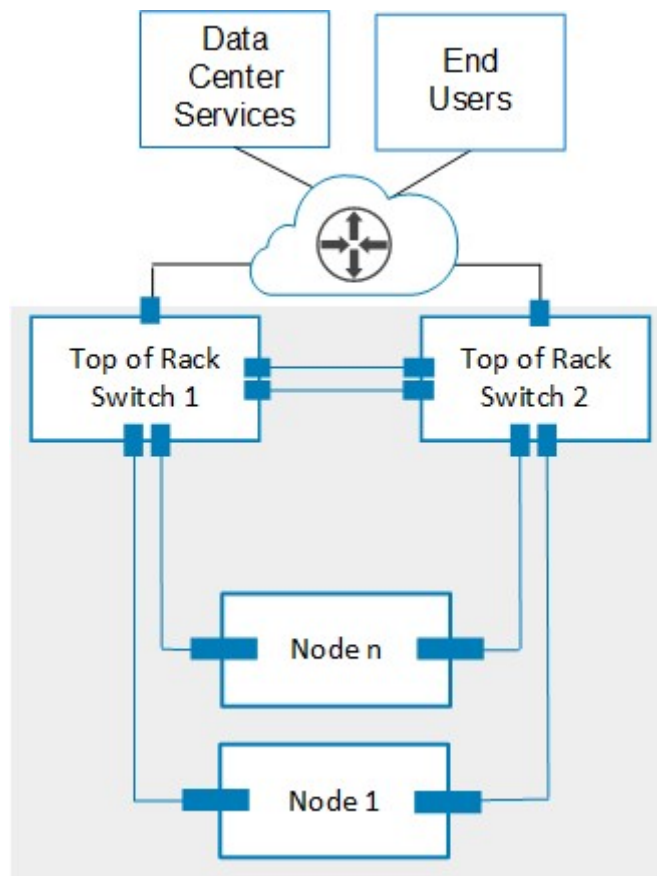


Figure 5. Connecting data center services with VxRail cluster

Routing services

VxRail cluster operations depend on a set of networks that run on both the virtual network inside the cluster and on the adjoining physical network switches.

Some of these networks, specifically for VxRail management and for end-user access must be passed to the upstream network, while other VxRail networks can stay isolated on the adjoining network switches.

It is best practice to reserve a set of 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/Layer3layer. 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 supported Ethernet and FC adapter card models, the number of slots available for expansion, and the maximum number of networking ports that are supported per node is driven by factors such as the VxRail series selected for the cluster, and the number of CPUs installed per node.

The following networking connectivity rules apply to VxRail nodes:

- The integrated 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.

The integrated and expansion network connectivity options are displayed for the following VxRail nodes based on the PowerEdge:

- 16th generation Intel platform
- 15th generation Intel platform
- 14th generation Intel platform
- Platform with AMD processors

The following figures show the network connectivity options that are supported for base connectivity on the integrated adapter cards for each VxRail model series, and the available PCIe-based adapter options for the expansion slots on the nodes.

VxRail 16G Intel-Based Node Integrated Connectivity Options								
VxRail Series	Storage Type	1x1 iDRAC	2x10 RJ45	2x10 SFP+	4x10 RJ45	4x10 SFP+	2x25 SFP28	4x25 SFP28
VE-Series	H/F							
VP-Series	H/F							
Storage Types: H=Hybrid F=All-Flash N=All-NVMe								

Figure 6. Integrated connectivity options for 16th generation VxRail models with Intel CPUs

VxRail 16G Intel-Based Node Expansion Connectivity Options								
VxRail Series	Storage Type	2x10 RJ45	2x10 SFP+	4x10 RJ45	4x10 SFP+	2x25 SFP28	4x25 SFP28	2x100 QSFP/28/56
VE-Series	H/F							
VP-Series	H/F							

Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 7. Expansion connectivity options for 16th generation VxRail models with Intel CPUs

VxRail 15G Intel-Based Node Integrated Connectivity Options								
VxRail Series	Storage Type	1x1 iDRAC	2x10 RJ45	2x10 SFP+	4x10 RJ45	4x10 SFP+	2x25 SFP28	4x25 SFP28
E-Series	H/F							
E-Series	N							
P-Series	F							
P-Series	N							
S-Series	H							
V-Series	F							
VD-Series	N							

Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 8. Integrated connectivity options for 15th generation VxRail models with Intel CPUs

VxRail 15G-Based Intel Node Expansion Connectivity Options										
VxRail Series	Storage Type	4x1 RJ45	2x10 RJ45	2x10 SFP+	4x10 RJ45	4x10 SFP+	2x25 SFP28	4x25 SFP28	2x100 QSFP	2x16/32 FC
E-Series	H/F	Red	Green	Green	Green	Green	Green	Green	Green	Green
E-Series	N	Red	Green	Green	Green	Green	Green	Red	Green	Green
P-Series	F	Red	Green	Green	Green	Green	Green	Green	Green	Green
P-Series	N	Red	Green	Green	Green	Green	Green	Green	Green	Green
S-Series	H	Red	Green	Green	Green	Green	Green	Green	Green	Green
V-Series	F	Red	Green	Green	Green	Green	Green	Green	Green	Green
VD-Series	N	Green	Green	Green	Red	Red	Green	Red	Red	Red

Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 9. Expansion connectivity options for 15th generation VxRail models with Intel CPUs

VxRail 14G Intel-Based Node Integrated Connectivity Options									
VxRail Series	Storage Type	1x1 iDRAC	2x10 RJ45	2x10 SFP+	4x10 RJ45	4x10 SFP+	2x25 SFP28	4x25 SFP28	2x100 QSFP
D-Series	H/F	Green	Green	Green	Red	Red	Green	Red	Red
E-Series	H/F	Green	Green	Green	Green	Green	Green	Red	Red
E-Series	N	Green	Green	Green	Green	Green	Green	Red	Green
G-Series	H/F	Green	Red	Green	Red	Red	Green	Red	Red
P-Series	H/F	Green	Green	Green	Green	Green	Green	Red	Green
P-Series	N	Green	Green	Green	Green	Green	Green	Red	Red
S-Series	H	Green	Green	Green	Green	Green	Green	Red	Red
V-Series	F	Green	Green	Green	Green	Green	Green	Red	Red

Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 10. Integrated connectivity options for 14th generation VxRail models with Intel CPUs

VxRail 14G Intel-Based Node Expansion Connectivity Options									
VxRail Series	Storage Type	2x10 RJ45	2x10 SFP+	4x10 RJ45	4x10 SFP+	2x25 SFP28	4x25 SFP28	2x100 QSFP	2x16/32 FC
D-Series	H/F	Green	Green	Green	Red	Green	Red	Green	Red
E-Series	H/F	Green	Green	Green	Red	Green	Red	Green	Green
E-Series	N	Green	Green	Green	Red	Green	Red	Green	Green
G-Series	H/F	Green	Green	Green	Red	Green	Red	Green	Red
P-Series	H/F	Green	Green	Green	Green	Green	Red	Green	Green
P-Series	N	Green	Green	Green	Green	Green	Red	Green	Green
S-Series	H	Green	Green	Green	Green	Green	Red	Red	Green
V-Series	F	Green	Green	Green	Green	Green	Red	Green	Green

Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 11. Expansion connectivity options for 14th generation VxRail models with Intel CPUs

VxRail AMD-Based Node Built-In Connectivity Options								
VxRail Series	Storage Type	1x1 iDRAC	2x10 RJ45	2x10 SFP+	4x10 RJ45	4x10 SFP+	2x25 SFP28	4x25 SFP28
E-Series	H/F	Green	Green	Green	Red	Red	Green	Red
E-Series	N	Green	Green	Green	Red	Red	Green	Red
P-Series	F	Green	Green	Green	Red	Red	Green	Red
P-Series	N	Green	Green	Green	Red	Red	Green	Red

Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 12. Built-in connectivity options for VxRail models with AMD CPUs

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

Storage Types: H=Hybrid F=All-Flash N=All-NVMe

Figure 13. Expansion 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 that is selected for support with VxRail must pass tests to qualify.

Table 1. Networking products that pass qualification and are supported for VxRail

Port speed	Vendor
10 GbE	<ul style="list-style-type: none"> • Intel • Broadcom • QLogic
25 GbE	<ul style="list-style-type: none"> • Broadcom • Intel • Mellanox • QLogic
100 GbE	Mellanox

Follow the guidelines to drive port adapter selection:

- When a VxRail cluster is initially built, it is recommended, but not required, that all the NIC 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.
- Use 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 works as expected. Consult the Dell cable and optics support matrix before attempting to mix vendor equipment in a VxRail cluster.
- The feature sets supported from NIC card suppliers do not always match. There is a dependency on the firmware and driver in the adapter card to support certain features. If a feature is needed to meet a business requirement, consult with a sales specialist to verify that the needed feature is supported for a specific vendor.

VxRail FC adapter options

FC 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.

The following FC adapter cards are supported for use with VxRail:

Table 2. FC adapter cards

FC speed	Vendor
16 GB	<ul style="list-style-type: none">• Emulex• QLogic
32 GB	<ul style="list-style-type: none">• Emulex• QLogic

VxRail RoCE adapter options

The RoCE adapter supports RDMA connectivity over converged networks. RDMA is a technology to send data over a network without involving the CPU in the transfer.

Enabling RDMA over a Converged Ethernet network infrastructure provides faster data transfer for network-intensive applications through lower I-O latencies on this network. IOPS performance is improved in comparison to other network-based storage connectivity options. The port speed for the Mellanox adapter is 25 GbE.

The following guidelines apply for RoCE-supported adapters:

- Verify that the VxRail cluster is configured with a minimum version of VxRail 7.0.200.
- Configure all nodes in the cluster that are connected to the common VMware vSAN data store with RoCE-supported adapter cards from the same vendor with the same model. Variations in adapters from different vendors or different models can disrupt I/O on the VMware vSAN data store.
- The Ethernet ports on the RoCE-supported adapters are reserved for VMware vSAN traffic only.
- The physical network supporting VMware vSAN is configured as a lossless network. To determine which Ethernet adapters support RDMA, see the appropriate *VxRail Support Matrix* on the [Dell Support site](#).

VxRail DPU options

DPUs support the offloading of network and security services from the main CPUs on the VxRail nodes.

ADPU is installed in the factory on the VxRail node instead of a standard Ethernet adapter card. The DPU differs from a traditional Ethernet adapter card because ESXi is installed on the device, and the DPU is configured with its own CPU.




Figure 14. Integrated DPU with ESXi image and CPU

Adhere to the following guidelines for VxRail clusters with DPUs:

- The VxRail cluster must be running a minimum version of VxRail 8.0.010.
- The VxRail nodes are configured in the factory with qualified and supported DPUs.
- The data center network supporting the VxRail cluster can support the physical network speed requirements of the DPU.

Table 3. DPU port speeds

Port speed	Vendor
25 GbE	<ul style="list-style-type: none"> • NVIDIA • Pensando
100 GbE	Pensando

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 or OCP. Starting with version 7.0.130, for workloads that require a higher level of failure protection, VxRail supports spreading the networks across NDC or 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 that are supported in VxRail, or you can select the Ethernet ports for 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 or OCP, mix the ports from the NDC or OCP and a PCIe adapter card, or select only ports from PCIe adapter cards. For more details, see [Appendix F: Physical Network Switch Examples](#) to understand the most common node connectivity options.




Figure 15. Mixing NDC/OCP and PCIe ports to support a VxRail cluster

If change the topology by migrating the VxRail networks onto other uplinks on the VxRail nodes, you can perform this activity after the cluster is built, as long as the VxRail cluster is at VxRail 7.0.010 or later.

VxRail networking rules and restrictions

The following network rules and restrictions apply.

- The Ethernet ports selected during the VxRail initial build to support the VxRail cluster are reserved 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 that are 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:
 - Reserve four ports on each node 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 1 GbE:
 - The most common topology is to configure the cluster with either two ports or four ports per node to support VxRail networking traffic.
 - Starting with VxRail 7.0.400, six ports or eight ports per node can be selected to support VxRail networking. This option is best used for deployments supporting demanding and network-intense workloads.
 - Adding Ethernet ports beyond the ports that are 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 VxRail 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 demanding and network-intense workloads.
- Custom Ethernet port configurations are supported with restrictions:
 - Before VxRail 7.0.130, all Ethernet ports on the VxRail nodes that are selected for a VxRail cluster must be the same port type and run at the same speed.
 - For VxRail version 7.0.130 and later, the ports on the NDC or OCP and PCIe adapter cards that are configured in the VxRail nodes can run at different speeds. The NDC or OCP ports can run at 10 GbE and the ports on the PCIe adapter cards can run at 25 GbE.
 - Any ports that are assigned to the same VxRail network, whether based on NDC/OCP or PCIe, must run 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. 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.
 - You should not mix Ethernet port types such as RJ45 and SFP+ to support VxRail cluster network operations. Mixing different Ethernet port types invites complexity regarding firmware, drivers, and cabling with the data center network.




Figure 16. 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 that span across data centers. Simpler workloads can be satisfied with the nodes and adjacent switches that are confined to a single rack, with routing services configured further upstream.

Before cabling and powering on VxRail nodes and performing an initial build of the VxRail cluster, create a site diagram showing the proposed network components and connectivity.

Select the network architecture to support the VxRail cluster and the protocols to connect to data center services and end users. For VxRail clusters managing production workloads, VLANs are configured to support the VxRail networks. Determine which network tier the VxRail networking VLANs terminate and which tier to configure routing services.

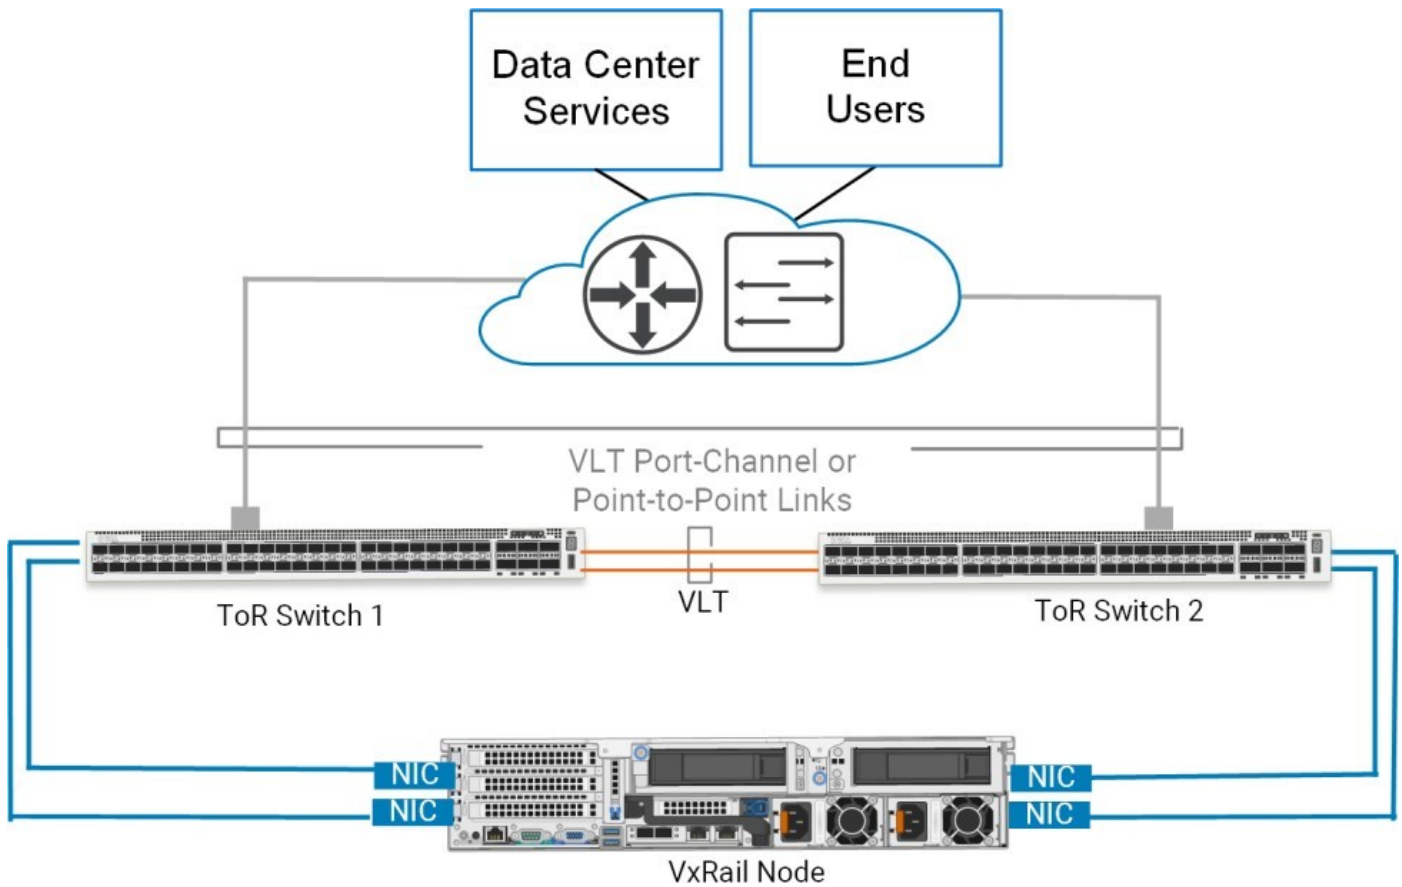


Figure 17. High-level network topology with Layer 2 and Layer 3 options

VxRail cluster operations require several ports on each switch. To determine the base number of ports, multiply the number of Ethernet ports on each VxRail node to support VxRail networking by the number of nodes to be configured into the cluster. For a dual-switch configuration, reserve ports on each switch to form an interswitch link for network traffic passage. 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 at a data center that you cannot easily access, set up an OOB management switch to facilitate direct communication with each node.

To use OOB management, connect the iDRAC port to a separate switch to provide physical network separation. Default values, capabilities, and recommendations for OOB management are provided with server hardware information. Reserve an IP address for each iDRAC in your VxRail cluster (one per node).

VxRail Cluster Types

The primary building block for VxRail is the individual node. 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 data store 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.

Cluster with VMware 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 VMware 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 VMware vSAN data store during the cluster initialization process.

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

- With a local VMware vSAN data store, 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.

VxRail supports two different vSAN architectures. The first architecture is known as the Original vSAN Architecture and the second architecture is referred to as the Express vSAN Architecture.

Original vSAN architecture

The Original vSAN Architecture is a two-tier model that is built on the foundation of disk groups.

A disk group consists of a single cache drive that is partnered with one or more capacity drives. A collection of disk groups is used to form a vSAN data store.




Figure 18. Local vSAN OSA datastore deployed on a VxRail cluster

The cluster initialization process performs an inventory of the disk drives on the nodes, and uses that discovery process to identify the number of disk groups on each node to be used as the foundation for the vSAN data store. The high-endurance SSD drives discovered on each node serve as a cache for virtual machine I-O operations in each disk group, while the high-capacity drives discovered are the primary permanent storage resource for the virtual machines for each disk group. The vSAN build process partners the discovered cache drives with one or more capacity drives to form disk groups, with the resulting vSAN data store consisting of this collection of disk groups.

VxRail clusters with vSAN datastores based on the Original vSAN Architecture support solid-state and NVMe drives for both cache and capacity, and solid-state and hard drives for capacity only. This architecture supports network speeds of 10 GbE, 25 GbE, and 100 GbE.

Express vSAN architecture

The Express vSAN Architecture is a single-tier model that takes advantage of the adoption of high-performance NVMe drives and the continuous increase in CPU cores to support a data store based on a single pool of storage instead of disk groups.

With this architecture, every drive can serve both cache and capacity requirements.




Figure 19. Local vSAN ESA datastore deployed on a VxRail cluster

During the inventory process, VxRail queries the drive slots and verifies whether a datastore can be built with the Express vSAN Architecture. The data store will only be built with this architecture if all the drive slots contain compatible NVMe drives. If you select this vSAN architecture option, the network that is configured to support the data store must be running at either 10 GbE, 25 GbE, or 100 GbE.

Dynamic cluster

Dynamic clusters differentiate themselves from other VxRail cluster types with the resource that is selected for primary storage. With other cluster types, there is a dependency on the local vSAN data store as the primary storage resource. With a dynamic cluster, the nodes that are 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. 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 that are 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:

- Verify that the storage resource in the data center you plan for VxRail dynamic clusters is supported. See the support matrix for dynamic nodes on the [VxRail Product Support](#) to verify compatibility.
- 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](#). Use the technical documentation provided for the selected external storage to prepare the storage for a VxRail dynamic cluster.
- 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 data store.

FC storage option

You can configure a compatible FC storage array which can be configured to supply a single VMFS data store or multiple data stores to a VxRail dynamic cluster.




Figure 20. Dynamic cluster using FC-connected VMFS for primary storage

This option has the following prerequisites:

- Verify that the storage array in your data center to support the dynamic cluster is supported with VxRail. See 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 FC adapter cards. You should have two FC 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 FC switches to accommodate the connections required from each VxRail node.

Remote VMware vSAN data store option

For primary storage with dynamic clusters, you can use an existing VMware vSphere or VxRail cluster with a local vSAN data store in your data center.

The VMs running on the dynamic cluster use the free storage resource on the VMware vSAN data store, and the compute resources on the local nodes.




Figure 21. Dynamic cluster using a remote VMware vSAN datastore for primary storage

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 data store and any dynamic cluster that is connected to this remote vSAN data store must be configured on the same VMware vCenter instance under a common data center object.
- If you already have a cluster that is sharing its vSAN data store to other clusters, ensure that you have not reached the maximum of five clusters that are already mounted to this vSAN data store.
- 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 data store 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 data store and dynamic cluster nodes.

PowerFlex storage option

PowerFlex provides IP-based storage to VxRail dynamic clusters that can be configured as primary storage for virtual machine workload.




Figure 22. VxRail dynamic cluster storage provided by PowerFlex virtual volume

PowerFlex 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, PowerEdge servers serving as the hardware foundation for block storage capacity.

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

- See [Dell PowerFlex Networking Best Practices and Considerations](#) to ensure the supporting network infrastructure is properly planned and configured.
- See [How to configure Dell PowerFlex Storage with VxRail Dynamic Nodes](#) for more information about provisioning PowerFlex storage.
- 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, configure a separate VMware VDS with new port groups in the VMware vCenter Server to support connectivity to the PowerFlex front-end system.

IP-based storage options

If the data center does not support FC storage, shared VMware vSAN resources or a deployed PowerFlex storage array, you can have storage over an IP network.

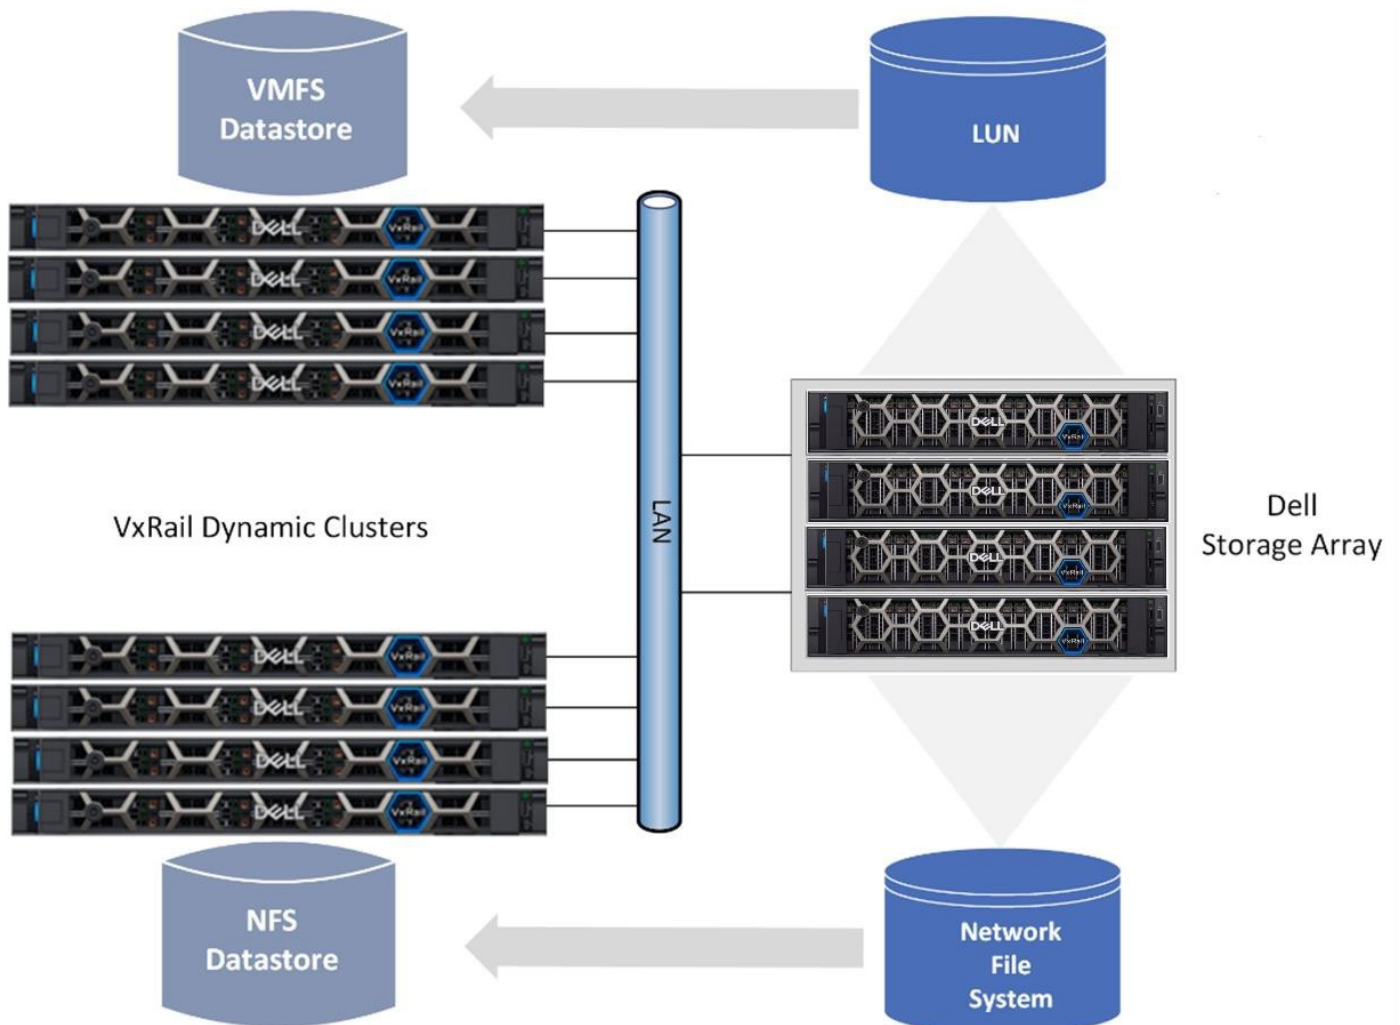


Figure 23. IP-based external storage supporting VxRail dynamic clusters

The supported storage resources can be either block-based or file system based. With block-level storage option, the LUN is presented to the VxRail cluster nodes over an IP network. VMware vSphere is used to configure a VMFS data store from the LUN.

iSCSI is a standard feature in VMware vSphere supported with VxRail. iSCSI is enabled by configuring a software adapter on a NIC on the VxRail nodes. The adapter then serves as an iSCSI initiator by targeting external storage arrays to present LUNs back to the initiators.

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 data store.

NVMe option

You can leverage NVMe (Non-Volatile Memory Express) to connect to block-based storage.

Leveraging NVMe can support local storage devices over PCIe, and storage devices over an FC network or an IP network. NVMe can serve as an alternative to FC or iSCSI storage for demanding workloads, since it is designed for usage with faster storage devices that are enabled with non-volatile memory.

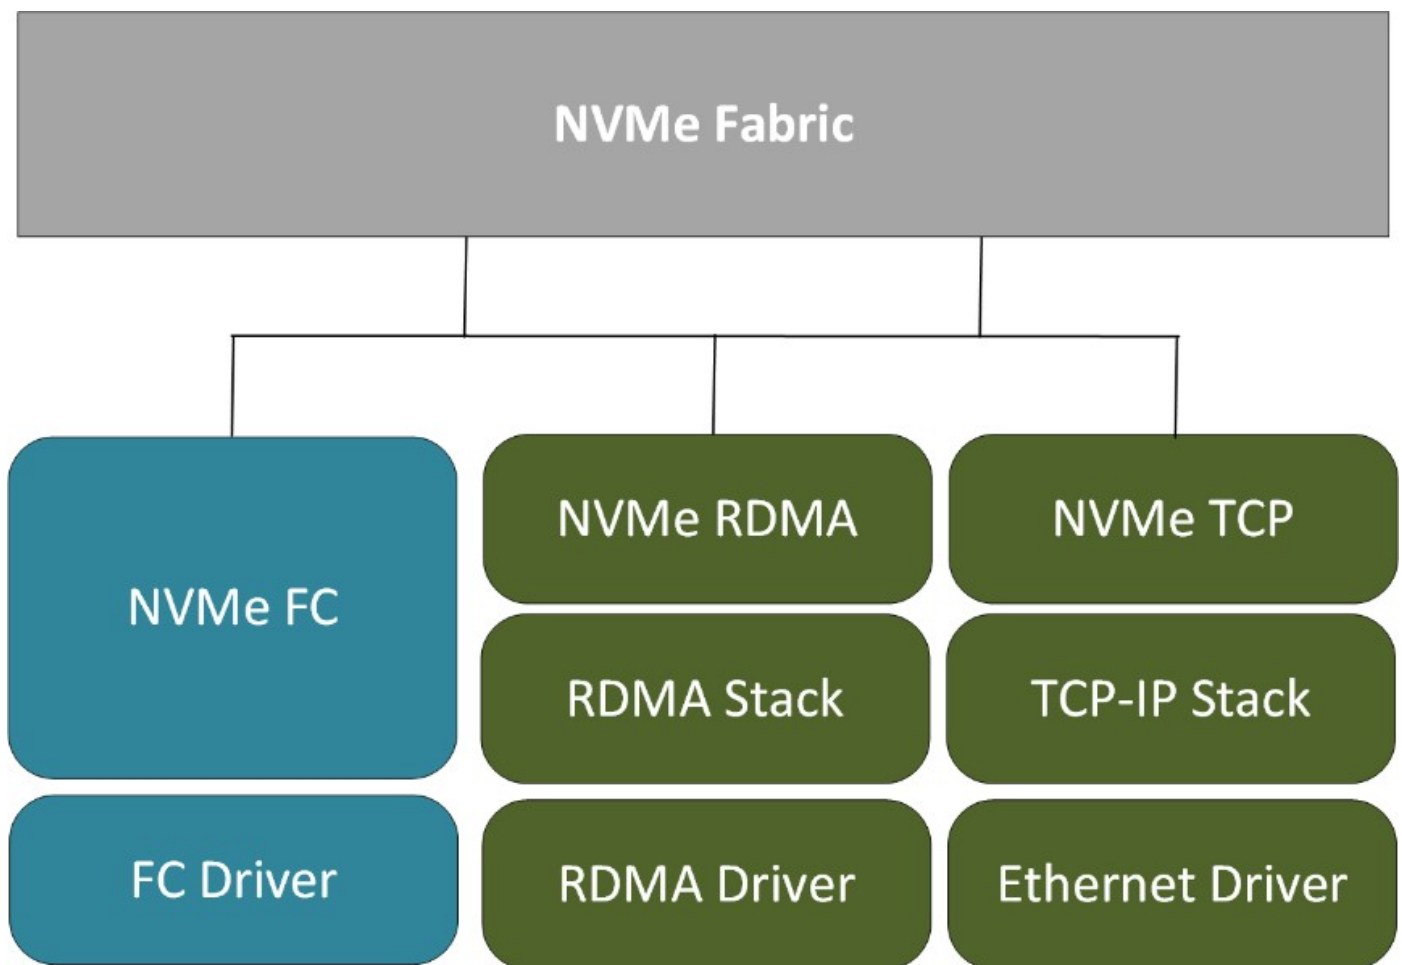


Figure 24. NVMe software stack in VMware vSphere

With all the storage options for dynamic clusters, verify that the storage resource in the data center you plan for VxRail dynamic clusters is supported. See the [VxRail E-Lab Navigator](#) to verify compatibility.

VMware vSAN stretched cluster

VMware vSAN stretched cluster supports synchronous I/O on a local vSAN data store on two sites that are separated geographically.

The VMware vSAN stretched cluster enables site-level failure protection with no loss of service or loss of data.

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

- The VxRail nodes must be populated with the compatible disk drives so that a VMware vSAN data store can be configured.
- The compute and storage resources that are planned for the cluster should be doubled. It is best practice to reserve 50 percent of resource capacity in the cluster for failure protection.
- This cluster type requires a witness. A witness is a small virtual appliance that monitors the health of the stretched cluster.
 - Network connectivity is required between the stretched cluster and the witness.
 - A VMware ESXi instance is required at the Witness site.
- Three data center sites are required for this solution: a primary and secondary data center site to host the VxRail infrastructure, and a third site to support a witness to monitor the stretched cluster.
- A minimum of one ToR switch for the VxRail nodes in the primary site and in the secondary site.

The VMware vSAN stretched cluster feature has strict networking guidelines for the WAN to support the solution.

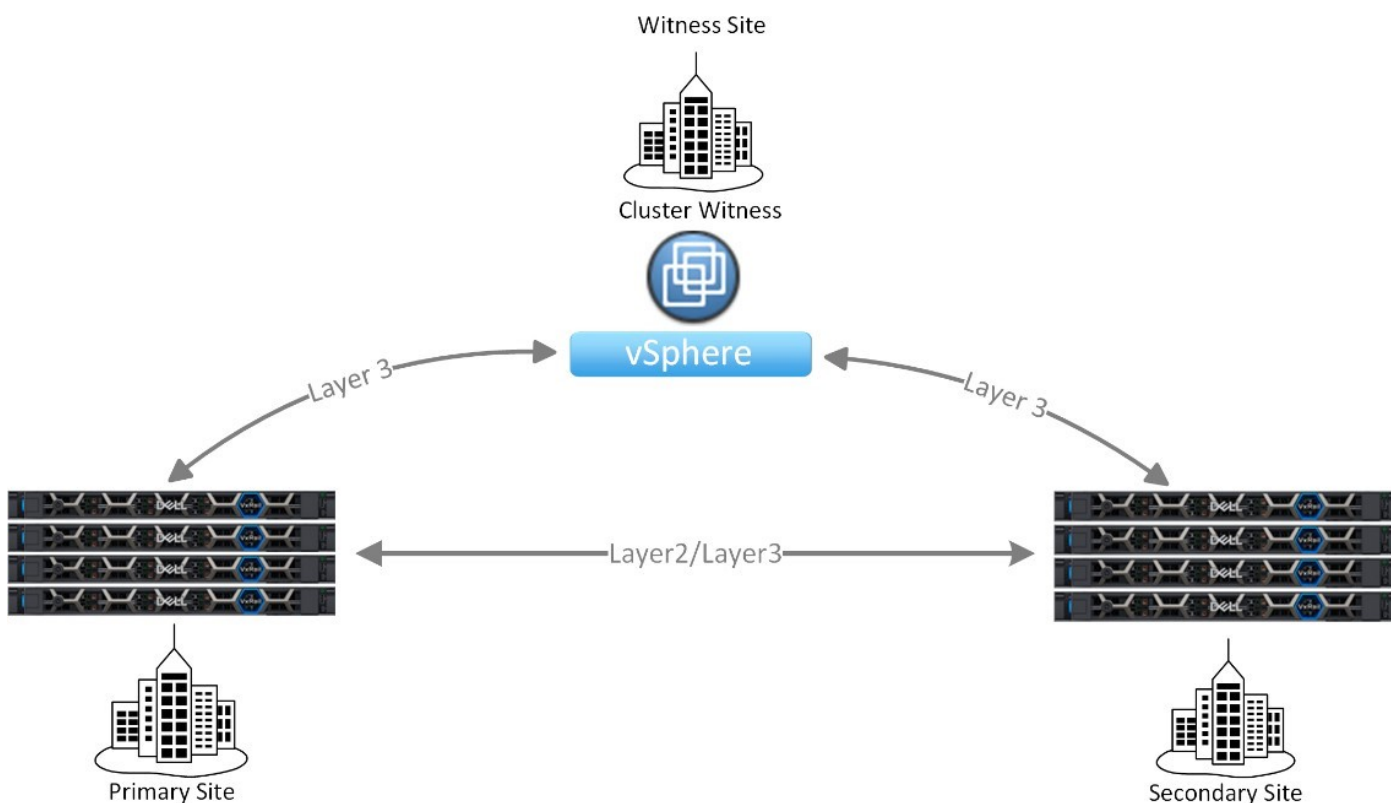


Figure 25. VMware vSAN stretched cluster topology

See [Dell VxRail vSAN Stretched Cluster Planning Guide](#) for detailed information about vSAN stretched cluster and the networking requirements.

2-node cluster

A 2-node cluster supports small-scale deployments with reduced workload and availability requirements, such as those in a remote office setting.

The solution is limited to two VxRail nodes only, and like the stretched cluster solution, requires a witness.

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.
- A single ToR switch is supported.
- Four Ethernet ports per node are required to deploy a 2-node cluster. Supported networking options include:
 - Four 10 GbE NICs.
 - Four 25 GbE NICs.
 - Two 10 GbE NICs and two 25 GbE NICs.

- Two 1 GbE NICs and two 10 GbE NICs.
- Two 1 GbE NICs and two 25 GbE NICs.
- In a mixed network configuration, the vSAN network must be placed on the pair of Ethernet ports running at the higher speed
- A 2-node cluster can be deployed with one of two network topologies:
 - Two Ethernet ports connect to create two links between the physical VxRail nodes, and the other two Ethernet ports connect to the ToR switch.




Figure 26. Direct connect option for a 2-node cluster

- All four Ethernet ports connect to the ToR switch or switches.




Figure 27. Switch-only option for a 2-node cluster

The following guidelines must be followed for the required witness:

- Network connectivity is required between the 2-node cluster and the witness.
- Witness can be deployed at the same physical site as the VxRail nodes, but cannot be deployed on the 2-node cluster.
- If there is more than one 2-node clusters deployed at the site, the witness can reside on a 2-node cluster it is not monitoring.
- For a 2-node cluster based on the VxRail XR-Series hardware, targeted for edge use cases:
 - The witness sled integrated into the XR-Series chassis can support only one VxRail 2-node cluster.
 - The Ethernet ports on the witness sled supporting the 2-node cluster only support 1 GbE.




Figure 28. 2-node cluster on VD-Series compute sleds and a witness sled

- The VMware vCenter Server instance supporting the 2-node cluster can be either a VxRail-managed VMware vCenter Server deployed on the cluster, or a customer-managed VMware vCenter Server external to the 2-node cluster. A minimum VxRail version of 7.0.410 is required for the VxRail supplied VMware vCenter option.




Figure 29. Two VMware vCenter Server placement options for a 2-node cluster

Like the VMware vSAN stretched cluster feature, this small-scale solution has strict networking guidelines 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 [Planning Guide—vSAN 2-Node Cluster on VxRail](#).

Satellite nodes

For cost efficiency and economies of scale, you can position a VxRail cluster at a central location to monitor and manage pools of VxRail satellite nodes, deployed locally and at remote locations.




Figure 30. VxRail satellite nodes managed by VxRail cluster

The PowerEdge models used as the hardware foundation for the other VxRail cluster types are the same for satellite nodes. Satellite nodes 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.

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.

VxRail feature-driven decision points

Applications, software stacks, and product features that are supported on VxRail can impact the architecture, deployment, and operations of the cluster. If you plan to include any of the feature sets or software stacks that are listed in this section, make note of the requirements that may impact plans for VxRail.

Software-defined data center

If you plan to include the transformation of your current data center with disparate technologies and processes towards an SDDC, 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 cloud management and virtual desktop software solutions. These resources can be managed as a logical pool for end-user consumption. By using VxRail clusters as the underlying foundation, the SDDC can be designed and deployed to meet specific business and operational requirements.

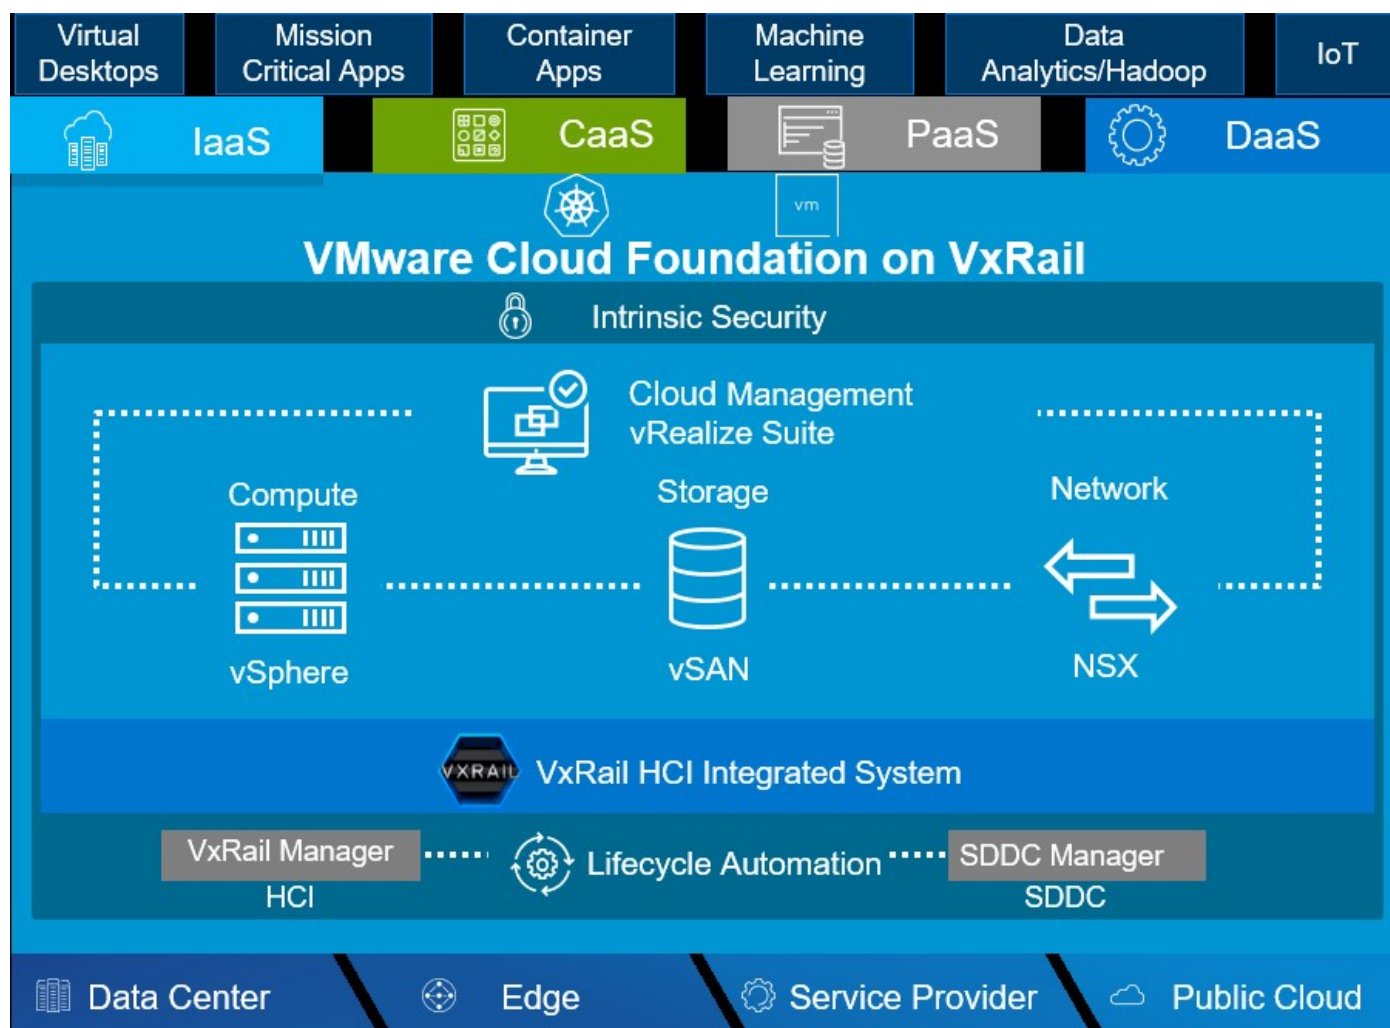


Figure 31. VxRail as the foundation for the SDDC

The path starts with a structured discovery and planning process that focuses on business use cases and strategic goals. These goals drive the selection of software layers that consist of the SDDC. Software layers are implemented 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 VCF 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 and deployment of VCF on VxRail, see [Dell VxRail technical Guides](#).

VMware vSphere with Kubernetes on VxRail

If you require workload management using Kubernetes, you can configure a VxRail cluster as a supervisor cluster for Kubernetes.

Kubernetes is a portable, extensible, API-driven platform for the management of containerized workload and services. The VMware Tanzu feature enables the conversion of a VxRail cluster with a VMware vSphere foundation into a platform for running Kubernetes workloads inside dedicated resource pools. A VxRail cluster that is enabled for VMware 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 VMware vSphere with Tanzu:

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

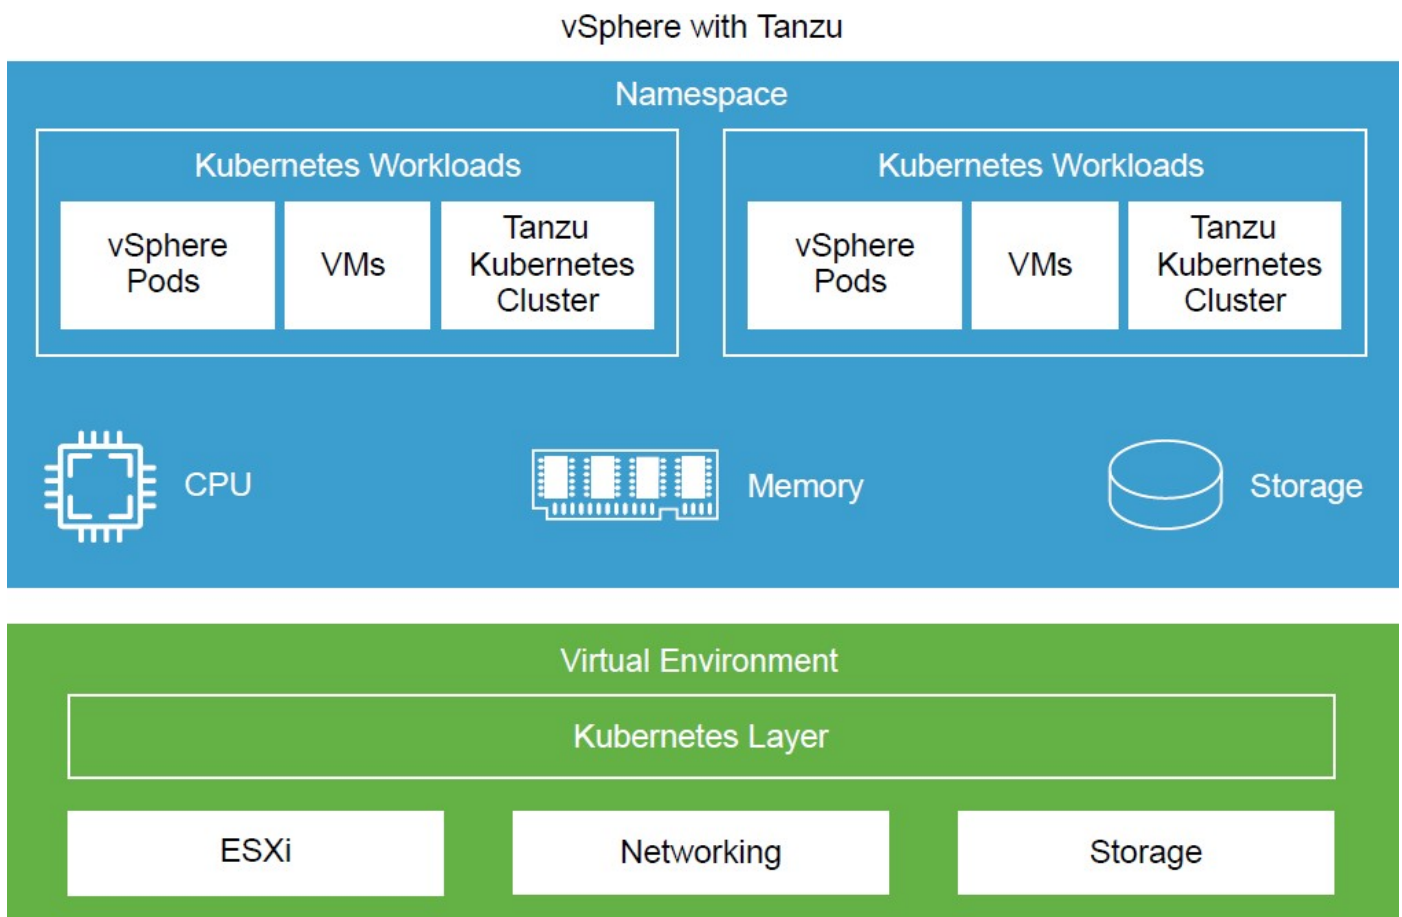


Figure 32. VMware vSphere with Tanzu on a VxRail cluster

As a VxRail administrator using VMware vSphere management capabilities, you can create name spaces on the Supervisor Cluster, and configure them with a specified amount of memory, CPU, and storage. Within the name spaces, 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 VMware vSphere Enterprise Plus license with an add-on license for Kubernetes.
- This feature requires port groups 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 also be configured on the VxRail cluster to enable connectivity from client network to workloads running in the name spaces.
- 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 VMware vSphere with Tanzu, see the [vSphere with Tanzu Configuration and Management Guide](#).

VxRail hardware and switch selection decision points

Step-by-step networking decision points are described.

Steps

1. Assess your requirements and perform a sizing exercise to determine the quantity and characteristics of the VxRail nodes you require to meet planned workload and targeted use cases.
2. Determine the number of physical racks required to support the quantity and footprint of VxRail nodes to meet workload requirements, including the ToR switches. Verify that the data center has sufficient floor space, power, and cooling.
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](#) for guidance on the options supported for VxRail cluster operations.
4. Based on the sizing exercise, determine the number of Ethernet ports on each VxRail node you want to reserve for VxRail networking.
 - a. Two ports might be sufficient in cases where the resource consumption on the cluster is low and will not exceed available bandwidth, or if the network infrastructure supports a high enough bandwidth that two ports are sufficient to support planned future growth.
 - b. Workloads with a high resource requirement or with a high potential for growth benefits from a 4-port deployment. Resource-intensive networks, such as the vSAN and VMware vSphere vMotion networks, benefit from the 4-port option because two ports can be reserved just for those demanding networks.
 - c. 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.
 - d. 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.

NOTE: You can reserve more than four ports per node for VxRail networking for cases where it is desirable for certain individual VxRail networks to not share any bandwidth with other VxRail networks.
5. Determine the optimal VxRail adapter and Ethernet port types to meet planned workload and availability requirements.
 - a. VxRail supports 1 GbE, 10 GbE, 25 GbE, and 100 GbE connectivity options to build the initial cluster.
 - b. Starting with VxRail 7.0.130, you have the flexibility to reserve and use the following Ethernet adapter types:
 - Only ports on the NDC or OCP for VxRail cluster networking.
 - Both NDC or OCP-based and PCIe-based ports for VxRail cluster networking.
 - Only PCIe-based ports for VxRail cluster networking.
 - c. If your performance and availability requirements might change later, you can reserve and use just NDC or OCP ports to build the initial cluster, and then migrate certain VxRail networks to PCIe-based ports.
 - d. If your requirements include using FC storage to support VxRail workload, you can select either 16 GB or 32 GB connectivity to your FC network.

NOTE: The VxRail cluster must be at version 7.0.010 or later to migrate VxRail networks to PCIe-based ports.
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 the type of adapter cards 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 or OCP. Optics that are rated for standard thermal specifications can be used on the expansion PCIe network ports supporting SFP28 connectivity.

7. Determine the additional ports and port speed on the switches for the uplinks to your core network infrastructure and interswitch links for dual switch topologies. Select a switch or switches that provide sufficient port capacity and characteristics.
8. Determine whether to enable OOB management. Dell iDRAC functionality is built into each VxRail node, and requires a 1GbE connection. Deploy a dedicated 1 GbE switch for this purpose. If this is not practical, you can also use open ports on the ToR switches.
9. Determine 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 ToR 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.

Prepare the data center to implement VxRail

VxRail is an entire SDDC in a cluster form factor. Administrative activities, including initial implementation and initialization, configuration, capacity expansion, online upgrades, and maintenance and support are handled within VxRail management.

When the VxRail 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 that are made in addressing these topics and performing these tasks drive the capability and functionality of the VxRail cluster

Be attentive to data center network planning and preparation topics to ensure that the VxRail cluster when deployed meets your business and operational requirements.

Prepare external network connectivity for VxRail

VxRail depends on access to web services outside the customer data center for certain features and functionality to work.

VxRail can be deployed and operated in a data center where access to the external network is blocked, but certain features are inoperable and a workaround might be needed.

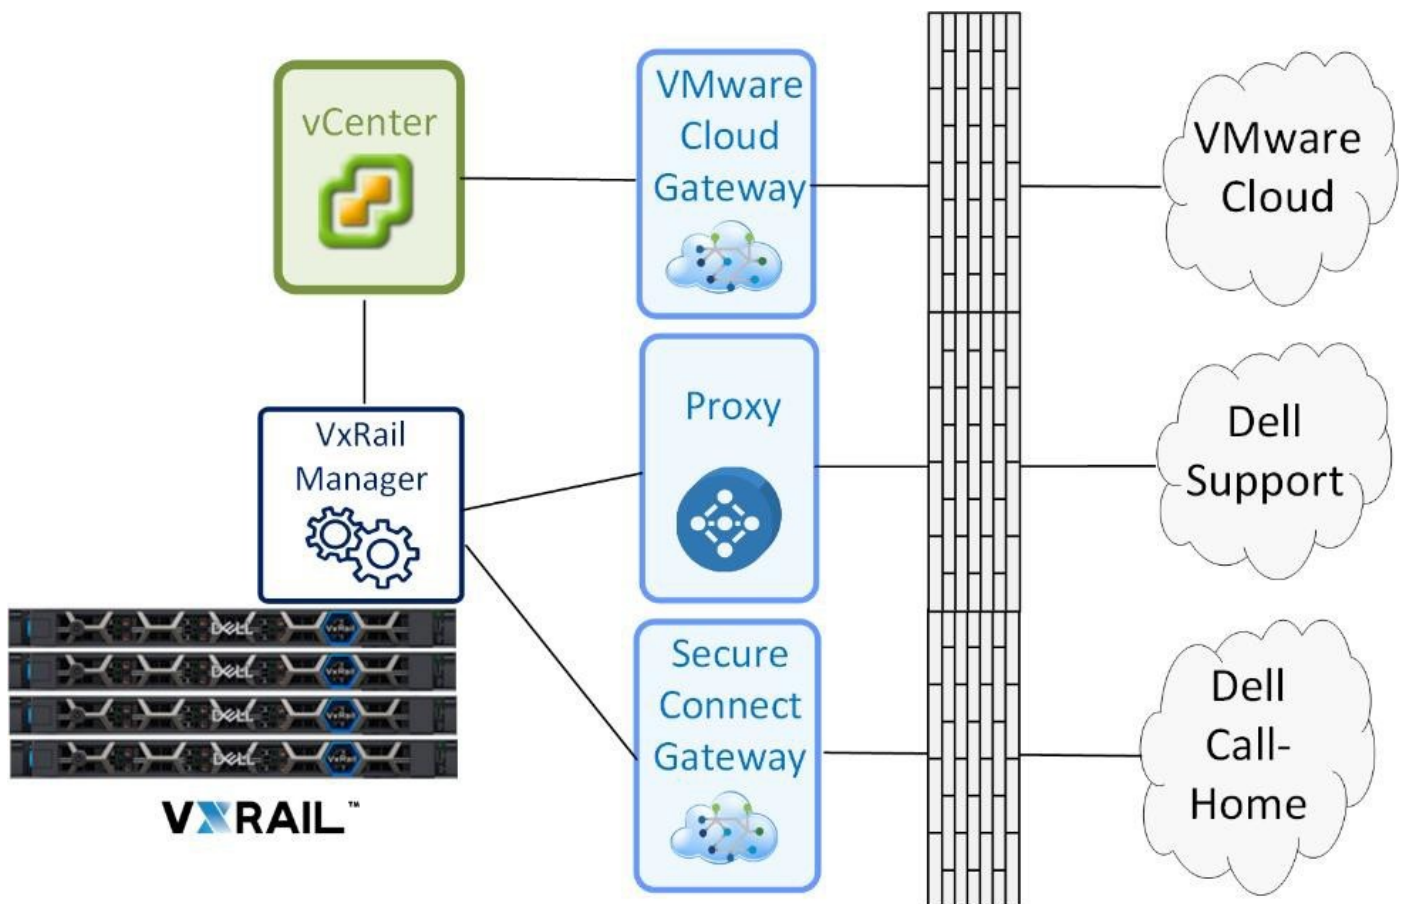


Figure 33. VxRail external network dependencies

- The VMware-branded software that is deployed on VxRail requires a VMware license. The licenses can be perpetual or based on a subscription licensing model. A perpetual license requires a key value that is retrieved from the VMware support site

to be entered into VMware vCenter. The subscription model requires access to the VMware cloud through a cloud gateway virtual appliance for license activation.

- You can configure VxRail to register directly with Dell support through VxRail Manager to streamline access to support services and product updates. You can also configure VxRail Manager to directly connect with Dell support services or to use a proxy server to align with customer networking security policies.
- Dell call-home enables VxRail Manager to directly contact support services in the event of a serious or critical failure on the cluster, thereby expediting resolution with Dell Support. The secure connect gateway virtual appliance can be deployed in the customer data center to consolidate call-home connectivity for Dell products in the data center.

If external network access to these sites is not allowed for VxRail:

- Work with VMware or your VMware partner to understand the licensing options.
- Download software updates without the aid of VxRail Manager and then transport to the data center for upload to VxRail Manager.
- Dell support personnel are not able to configure a remote session to troubleshoot VxRail.

The data center network must be configured to enable VxRail to connect to these external sites for the described functionality and services to work. For more information about the firewall rules for these services, see [Appendix D: VxRail Open Ports Requirements](#).

Prepare service connectivity for VxRail

VxRail supports network connectivity to Dell Customer Support, enables health monitoring, real-time analytics, event notification, and remote support when necessary.

If you plan to enable VxRail to connect to Dell back-end Customer Support centers, you can connect VxRail Manager directly to the back-end services sites or use a centralized instance or pool of secure connect gateways.

If you already have a secure connect gateway that is deployed in your data center to support this service-known as call-home for other Dell products, this same pool of gateway servers can be used for VxRail. If this is the first Dell product that you plan to connect to this service, be sure to first complete the prerequisites:

- Create and validate a Dell Support account.
- Configure at least one site ID for identity purposes in the Customer Support databases.
- Register the VxRail product with a Site ID.

When service connectivity is enabled, hosts can send health data to the VxRail Support team and provide secure automated access between Support and your VxRail. VxRail uses the secure connect gateway for connectivity. See [KB 000196945](#) and the [Secure Remote Services 3.52 Upgrade to Secure Connect Gateway Supplement Documentation](#) for more information. Verify that the minimum secure connect gateway version is 5.00.07.10. Enable service connectivity in direct connection mode or through an external secure connect gateway in VxRail Manager.

Prepare for VMware vSphere+ subscription licensing

This section is only relevant if your VMware licensing is going to be subscription-based and require connectivity to the VMware cloud for these services.

If VMware licenses are subscription-based, then the VMware vCenter Server instance supporting your VxRail cluster requires connectivity to the VMware cloud through the VMware vCenter Cloud Gateway.

Deployment and administration of the VMware vCenter Cloud Gateway virtual appliance is out of scope for this guide. See [VMware vSphere+ Documentation](#) for more details.

Table 4. Open ports to connect VMware vCenter Server to the VMware Cloud

Source	Destination	Port
Web browser	VMware vCenter Server Cloud Gateway	5480
VMware vCenter Server	VMware Cloud	433
VMware vCenter Server	VMware vCenter Server Cloud Gateway	5480
VMware vCenter Server	VMware vCenter Server Cloud Gateway	5484
VMware vCenter Server	VMware vCenter Server Cloud Gateway	7444

Table 4. Open ports to connect VMware vCenter Server to the VMware Cloud (continued)

Source	Destination	Port
VMware vCenter Server	VMware vCenter Server Cloud Gateway	433
VMware vCenter Server	VMware vCenter Server Cloud Gateway	5010-5019
VMware vCenter Server	VMware vCenter Server Cloud Gateway	2020

Prepare data center routing services

The VxRail external management network and any external-facing networks that are configured for VxRail must have routing services that support connectivity to external services and applications, and end users.

A leaf-spine network topology is the most common use case for VxRail clusters. A single VxRail cluster can start on a single pair of switches in a single rack. When workload requirements expand beyond a single rack, expansion racks can be deployed to support the additional VxRail nodes and switches. The ToR switches which are positioned as a leaf layer, can be connected using switches at the adjacent upper layer or spine layer. If you use spine-leaf network topology to support the VxRail clusters in your data center, you can enable Layer 3 routing services at either the spine layer or the leaf layer.




Figure 34. 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, you can extend Layer 2 networks between switches in physical racks over a Layer 3 overlay network to support a multirack VxRail cluster.

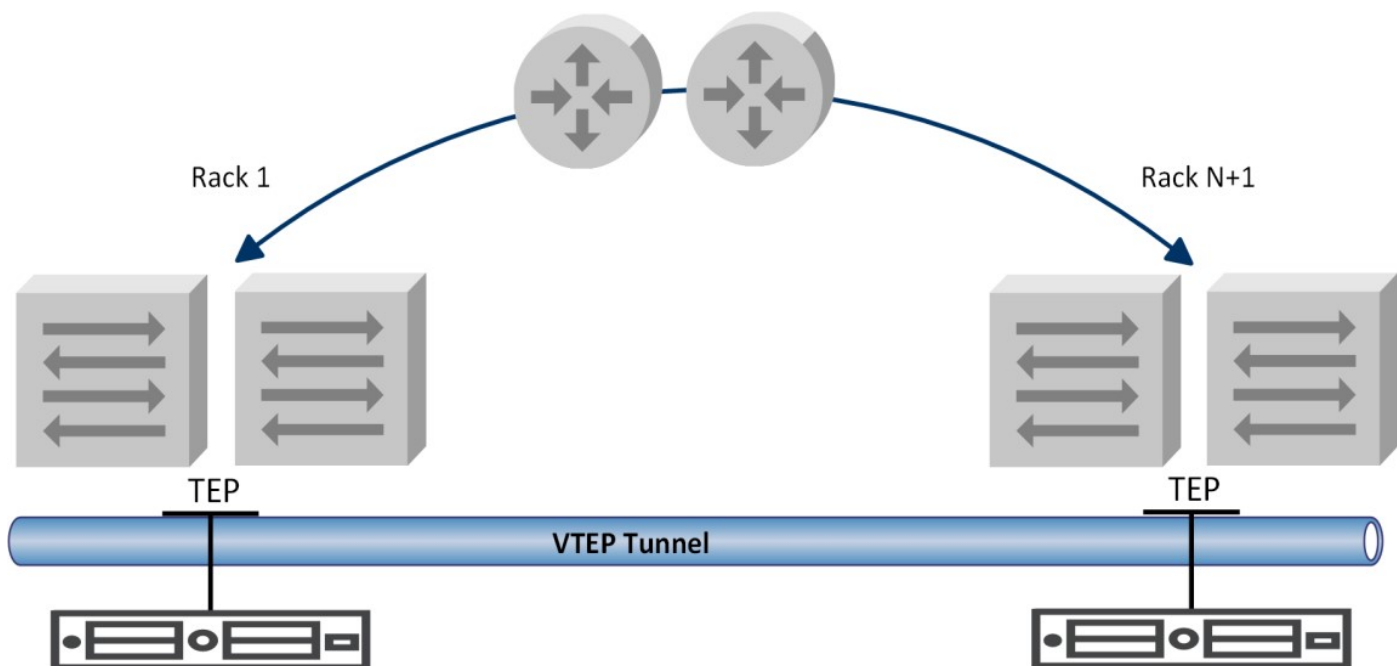


Figure 35. VTEP tunneling between leaf switches across racks

Prepare for multirack 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 in integrated mode 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.)

The following two conditions must be met to expand a cluster beyond a single rack with a new subnet range:

- Your VxRail cluster must be running VxRail 4.7.300 or later.
- You are not deploying VCF on the VxRail cluster.

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 terminate at the router layer, and routing services must be configured to enable connectivity between the racks.




Figure 36. Multirack VxRail sharing the same subnet




Figure 37. Multirack VxRail with different subnets

Prepare for VMware vSAN HCI mesh topology

This section is relevant only in situations for sharing vSAN data store resources over the network.

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 VMware vCenter Server instance.

With a vSAN HCI Mesh network, the VxRail cluster that leverages a local vSAN datastore for primary storage can also leverage the capacity on a remote vSAN datastore for application workload. If dynamic clusters are to 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 38. Storage resource sharing between VxRail clusters with VMware vSAN HCI mesh

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. This method is preferable if the client cluster does not have a local vSAN datastore. If a common VLAN is assigned, and the VxRail clusters are deployed against different sets of ToR switches, the VLAN must be configured to stretch between the set of switches.
- If a unique VLAN is assigned to the vSAN network on each cluster, then connectivity can be enabled using Layer 3 routing services. If this option is selected, be sure to assign routable IP addresses to the vSAN network on each participating VxRail cluster.

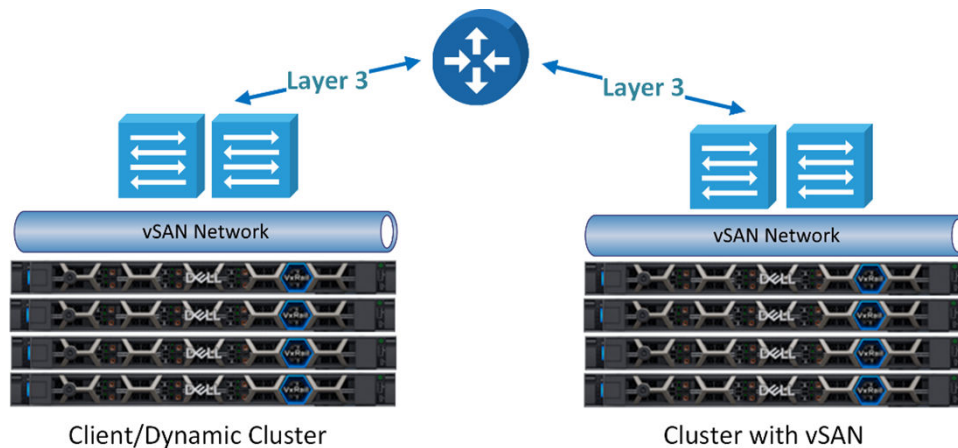


Figure 39. Enabling vSAN HCI mesh connectivity over a Layer 3 network

To share storage resources between VxRail clusters both with VMware vSAN datastores in a VMware vSAN HCI mesh, prepare your data center to meet the following prerequisites:

- A VMware vCenter Server instance at a version that supports VxRail 7.0.100 or later.
- A VMware vSAN Enterprise license for each VxRail cluster using VMware vSAN HCI mesh topology.

If your plans include sharing the VMware vSAN resources of a VxRail cluster with one or more VxRail dynamic clusters, meet the following prerequisites:

- A VMware vCenter Server instance at a version that supports VxRail 7.0.240 or later.
- A VMware vSAN Enterprise license is needed only for the VxRail cluster that is sharing its VMware vSAN storage resources. This license is not needed on a dynamic cluster because it does not have a local VMware vSAN datastore.

If your plans include sharing the VMware vSAN resources from a VxRail stretched cluster, meet the following prerequisites:

- The VxRail cluster participating in the vSAN HCI mesh topology must be running VxRail 8.0.100 or later.
- The VMware vCenter Server instance that supports the VxRail clusters must be running at a version that supports VxRail 8.0.100.
- The VMware vSAN datastore being shared is based in the OSA.
-

The VxRail client cluster can consist of:

- Dynamic nodes with no local VMware vSAN datastore.
- VxRail cluster with a local VMware vSAN datastore using the remote stretched VMware vSAN datastore as secondary storage.
- VxRail cluster with a stretched VMware vSAN datastore using the remote stretched VMware vSAN datastore as secondary storage.




Figure 40. VxRail stretched cluster sharing VMware vSAN storage with a client cluster

Prepare external FC storage for dynamic clusters

Perform this task if you plan to use FC storage as the primary storage resource for a VxRail dynamic cluster.

Preconfigure a VMFS datastore on each VxRail node before the initial cluster build.




Figure 41. Enable FC connectivity for dynamic clusters

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

- See the [VxRail 8.0 Support Matrix](#) to verify if the FC array is compatible with dynamic clusters.

- Include FC adapter cards with the VxRail order. Each node that is a member of a dynamic cluster requires FC connectivity. Install at least one dual-port FC adapter card for each VxRail node.
- Deploy a pair of FC switches for redundancy purposes to support the network topology.
- The minimum LUN size that is supported for dynamic clusters is 800 GB. Verify that you have sufficient capacity on your storage array.
- If multiple VMFS datastores are detected, the largest free capacity on the primary host is selected.
- Before the initial cluster build, perform LUN masking to all dynamic nodes.
- Format LUNs as a VMFS datastore to be supported for dynamic clusters.

Prepare for VxRail custom uplink assignments

During the initial build of the cluster, you have flexibility in how to assign uplinks to the VxRail networks.

Deploy VxRail using the predefined uplink assignment templates or select uplinks on each node you want to assign to a given VxRail network. A VxRail cluster where the networking profile is predefined follows 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.

- To deploy VxRail with a predefined network profile, each VxRail node port that is used to support VxRail networking must be running at the same speed.
- To create a custom profile option, the following general rules are applicable:
 - You can configure the VxRail nodes with Ethernet ports running at different speeds. For instance, you can have 10 GbE ports on the NDC or OCP, and 25 GbE ports on a PCIe adapter card.
 - The Ethernet ports that 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 VMware vSAN and VMware 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 or 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 VMware vMotion and VMware vSAN networks are assigned the other two uplinks.




Figure 42. Default network profiles for a 2-port and a 4-port VxRail network

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




Figure 43. 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 VMware vMotion network on the same pair of ports as reserved for the VxRail management networks, and isolate the VMware vSAN network on a pair of Ethernet ports.

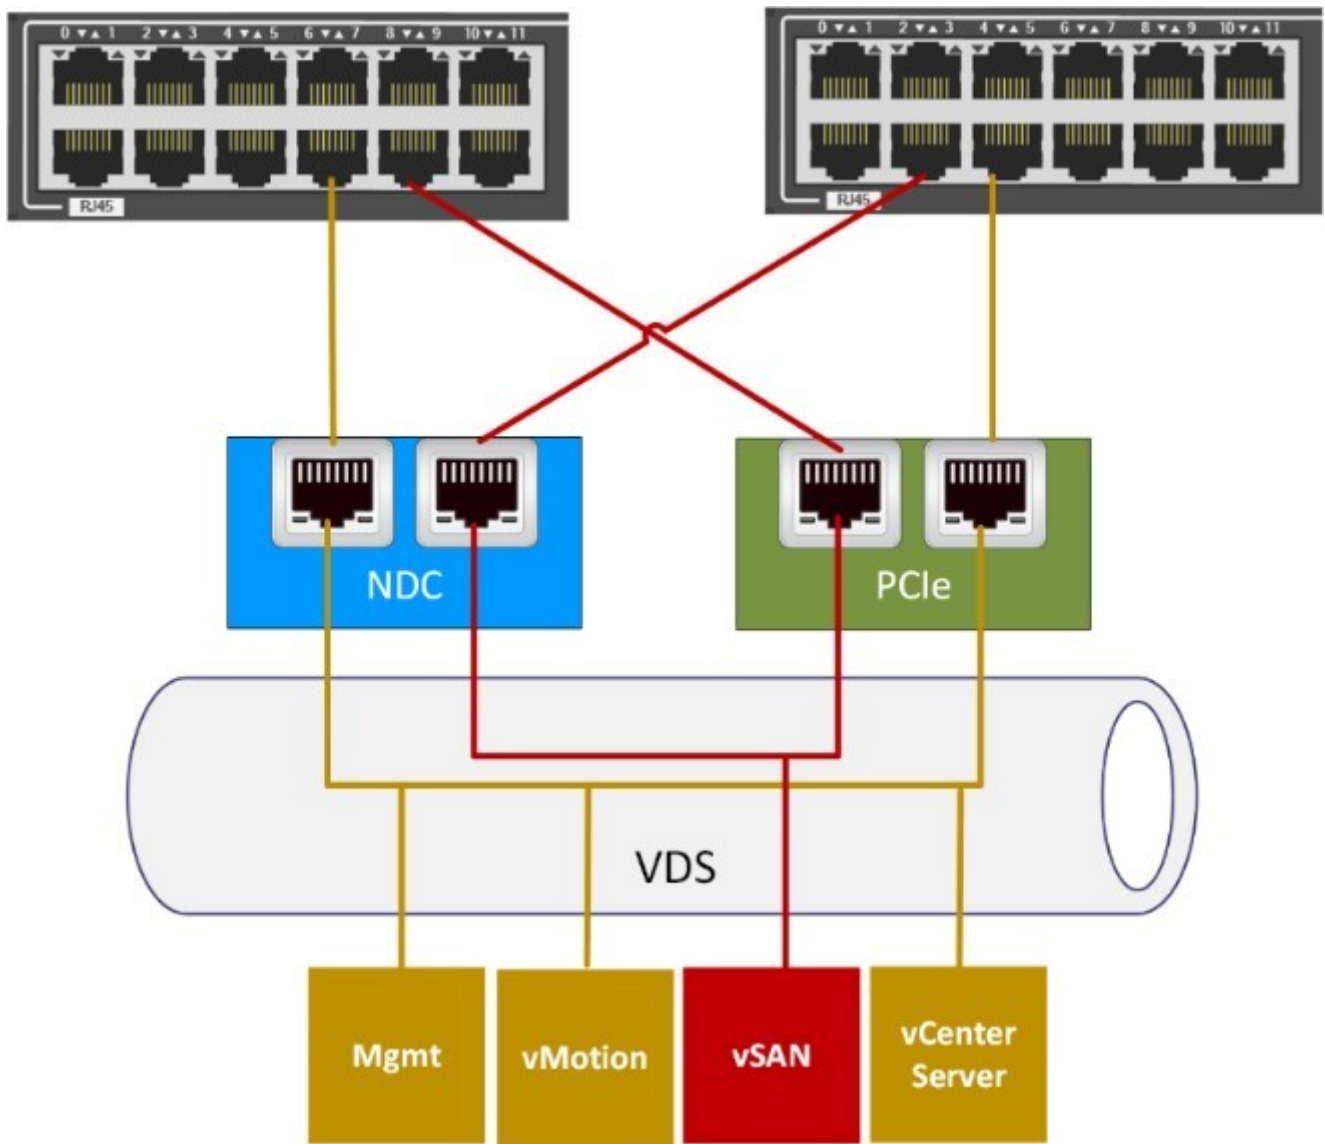


Figure 44. Custom uplink assignment with VMware vSAN network isolated on two Ethernet ports

Customizing the uplink assignments can impact the ToR 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 VMware 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 or OCP ports and PCIe Ethernet adapter cards, you can migrate certain VxRail networks off the NDC or 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

During the initial build of the cluster, select the MTU size to assign to the VMware VDS. This allows you to configure the switch ports supporting the VxRail cluster to match the MTU setting on the VMware VDS, and reduce the potential for network fragmentation.

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

Prepare for LAG of VxRail networks

LAG for specific VxRail networks is supported starting with VxRail 7.0.130.

NIC teaming in VxRail is the foundation to support LAG, which is the bundling of two physical network links to form a single logical channel. LAG allows ports on a VxRail node to peer with a matching pair of ports on the ToR switches to support load-balancing and optimize network traffic distribution across the physical ports. VxRail networks with heavy resource requirements, such as VMware vSAN and perhaps VMware vSphere vMotion, benefit most from network traffic optimization.

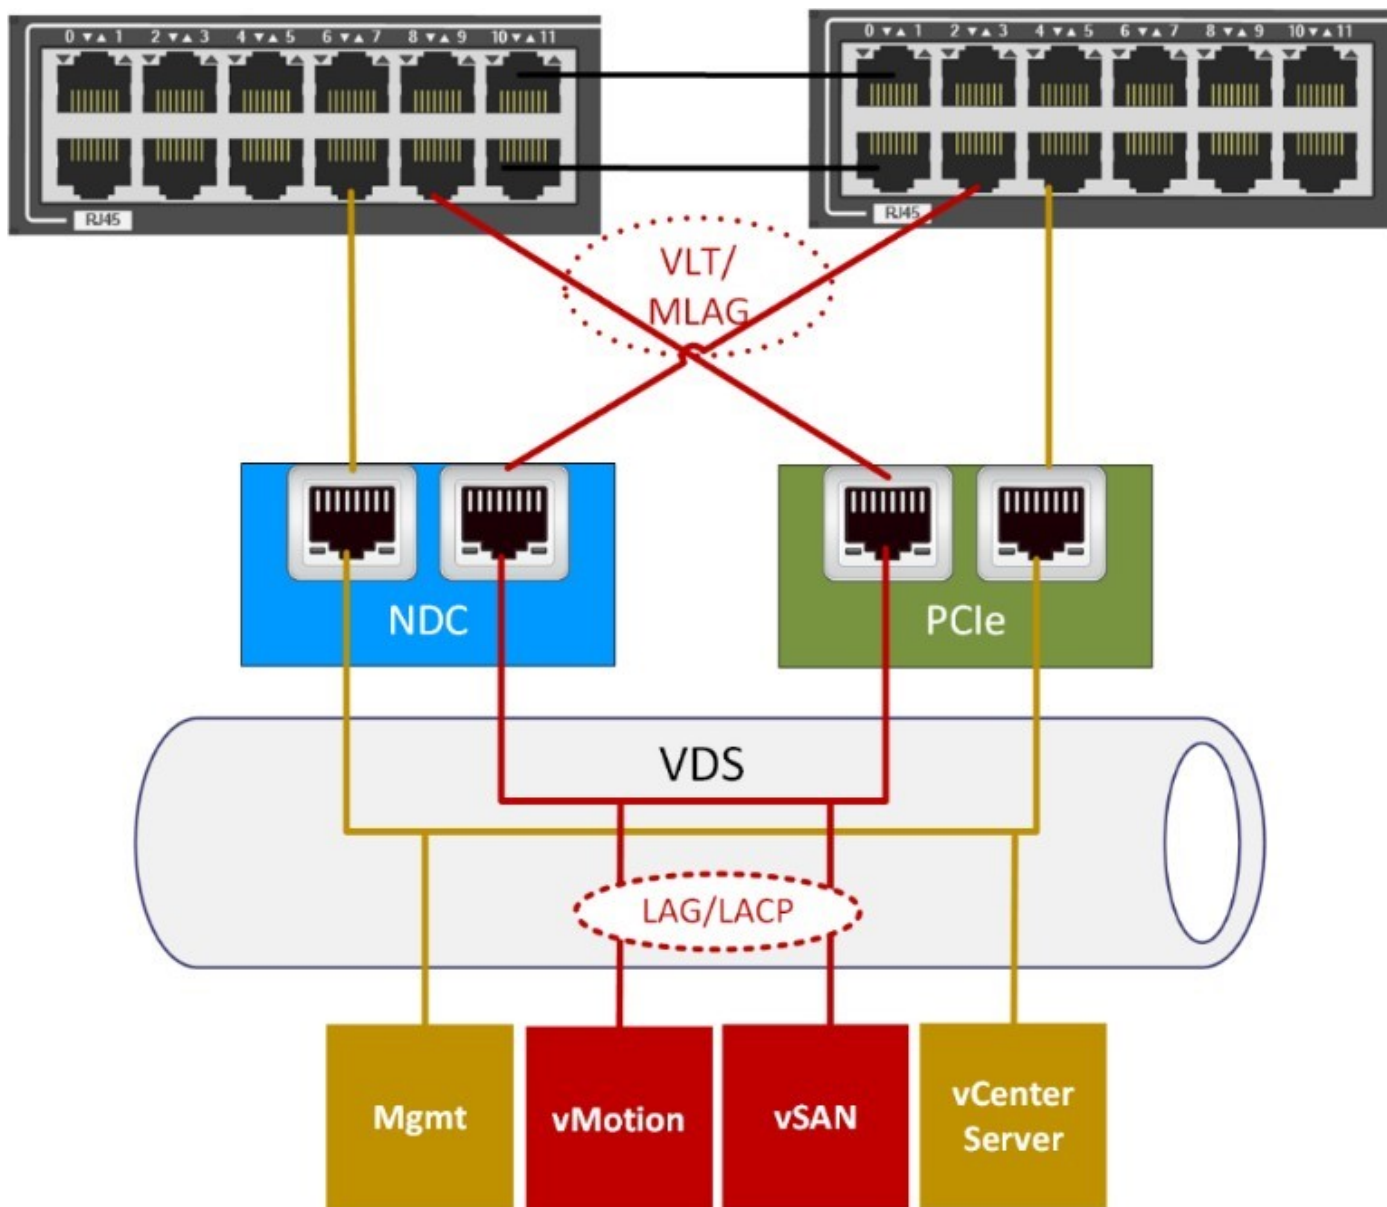


Figure 45. Dependent LAG features

Each VxRail network is assigned two uplinks by default during the initial implementation operation. Even if the virtual distributed switch port group is assigned a teaming and fail over policy to enable better distribution across the two uplinks, true load balancing is not achievable without LAG. Enabling LAG allows the VxRail network to better use the bandwidth of both uplinks, with the traffic flow coordinated based on the load-balancing hash algorithm that is configured on the virtual distributed switch and the ToR switches.

The following functionality dependencies must be understood if considering LAG with VxRail:

- The switches that are targeted for the VxRail cluster must support LACP. LACP is the protocol that dynamically forms a peering relationship between ports on two separate switches. Dynamic LAG requires an LACP policy to be configured on the VMware VDS to enable this peering to be established with the adjacent ToR switches.

- LACP is considered the best practice for LAG because it offers support for more load-balancing hashing algorithms and superior management capabilities.
- For network topologies using two or more switches to support VxRail networking:
 - The switch operating system must support a feature to allow LAG to span the two switches. This is known as Multi-Chassis Link Aggregation (MLAG) with Dell networking products.
 - The switch operating system must support a feature that enables the networks in a port channel to pass through the two switches at one end of a LAG group. This is known as Virtual Link Trunking with Dell networking products.
- LAG is not supported if the Ethernet adapters supporting the VMware vSAN network are enabled for RDMA over Converged Ethernet (RoCE).

LAG is not supported on the VxRail management networks prior to VxRail 7.0.450. The following minimum VxRail versions must be used when using LAG with VxRail:

- Customer-managed VMware VDS: VxRail 7.0.130
- VxRail-managed VMware VDS: VxRail 7.0.450
- All VxRail networks: VxRail 7.0.450
-

The following VxRail network guidelines must be adhered to enable LAG:

- VxRail supports a single LAG group per cluster.
- Two Ethernet ports can only be configured into a LAG group.
- The Ethernet ports that are selected for LAG can be all NDC/OCP-based, all PCIe-based, or a mixture of NDC/OCP and PCIe Ethernet ports.
- All ports that are configured for LAG must be running at the same speed.

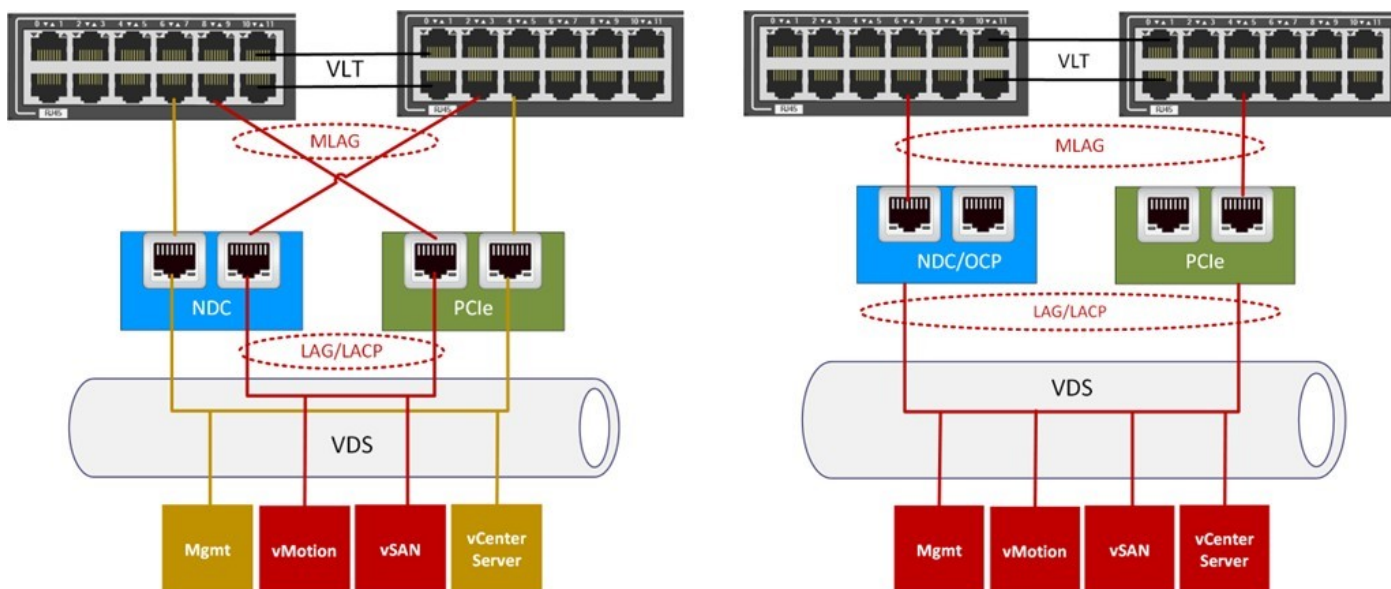


Figure 46. Sample LAG options for VxRail networking

The following guidelines must be followed to enable LAG on a customer-managed VMware VDS:

- You must supply a compatible VMware vCenter Server instance to serve as the target for the VxRail cluster.
- You must supply a compatible and pre-configured virtual distributed switch to support the VxRail networking requirements.
- The LACP policy to support LAG and load-balancing must be pre-configured on the virtual distributed switch.

Follow these guidelines to enable LAG on a VxRail-managed VMware VDS:

- The default mode for switch ports configured as port channels is that the ports shut down if they do not form a pairing with another pair of ports on another switch.
- VxRail requires those switch ports to remain open and active during the VxRail initial implementation process.
- The switches that are targeted for the VxRail cluster must support enabling the individual ports in the port channel to stay open and active if they do not form a LAG partnership with other switch port pairs within the configured timeout setting. (On Dell-branded switches running OS10, this is known as the LACP individual' feature.)

On Cisco-branded switches, the feature is LACP suspend-individual. This feature should be disabled on the switch ports in an EtherChannel to prevent the ports from shutting down. Check the documentation for switches from other vendors for the proper feature description.

Verify that the switches support LAG

Support for LACP, the selection of load-balancing hashing algorithms and the formation of LAG on the physical switches depends on the switch vendor and operating system. These features are 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 that are planned for the VxRail cluster supports this feature.

Verify support for multi-chassis LAG

To deploy a pair of switches to support the VxRail cluster, and enable load-balancing across both switches, the switches must support the ability for the networks in a LAG group to logically flow across both switches. The switch operating system must support the multi-chassis LAG feature, such as Cisco Virtual Port Channel. See the guides provided by your switch vendor for the steps to complete this task.

Verify support for LACP individual or similar feature

If you plan to deploy VxRail with a VxRail-supplied virtual distributed switch, then the switches must support the LACP individual or compatible feature. By default, switch ports that are configured for LAG are set to an inactive state until such time it has exchanged LACP PDUs with a LAG on another adjacent switch. When these PDUs are exchanged, the two LAGs can then sync into a partnership, and the sets of ports become active.




Figure 47. Enable connectivity for VxRail initial implementation with LAG

An individual VxRail node connects to the adjacent ToR with a standard virtual switch at power-on, and virtual switches do not support LAG. The LACP policy that is configured by VxRail on the VxRail-managed VMware VDS during the initial implementation process cannot exchange LACP PDUs at the power-on stage. This peering relationship does not occur until VxRail begins the virtual network formation stage later in the initial implementation process. Using this feature enables a switch port that is configured for LAG to be set to an active state in order to enable VxRail connectivity and allow initial implementation to proceed.

Identify switch ports to be configured for LAG

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 ToR switches.




Figure 48. Plug into equivalent switch ports for LAG

If you are enabling LAG 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. Create a table to map each VxRail node port to a corresponding switch port. Identify which ports on each VxRail to be enabled for LAG.

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

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 49. Sample port mapping

Assuming that you are using the second port on the NDC or 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 LAG. Create a table mapping the VxRail ports to the switch ports as part of the planning and design phase.

Plan LAG on switch port pairs

For each pair of ports on each node that is supporting a VxRail network to be enabled for LAG, configure the corresponding pair of switch ports they are connected to for LAG.

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 these tasks.

The tasks that are applicable to enable LAG include:

- Configure virtual link trunk or similar functionality between the switches for data center networks if more than one switch is planned to support the VxRail cluster.
- Configure port channels on each switch.
- Configure the VLAN or VLANs targeted for LAG in each port channel.
- Configure individual LACP or similar features if your plans include implementation with a VxRail-managed VMware VDS.
- Assign the respective port channels to the switch ports.

Prepare certificate authority server for VxRail

VxRail Manager is deployed with a self-signed certificate by default that is likely to be triggered as untrusted by most web browsers, depending on security settings

Replace the default certificate with one that is signed by your organization's root Certification Authority (CA) for production VxRail environments.

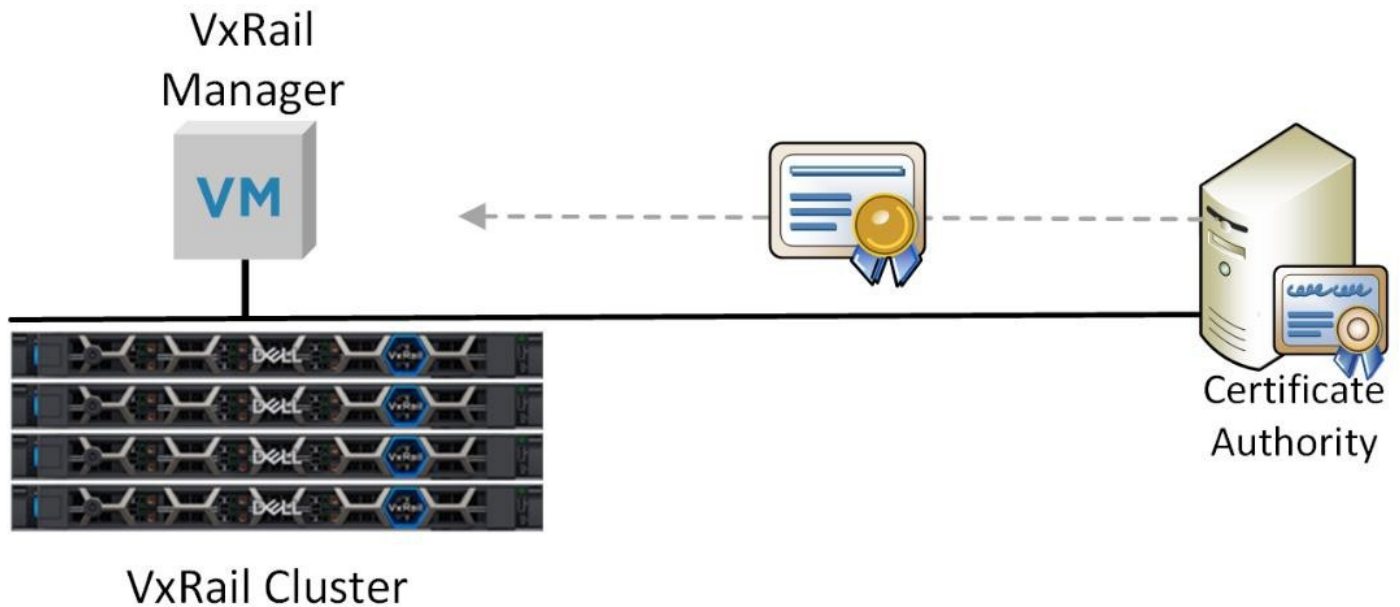


Figure 50. Trusted certificate authority supporting VxRail deployment

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

Identify isolation IP addresses for VMware vSphere High Availability

The VMware vSphere High Availability feature is configured by default when the VxRail cluster is built.

The VMware vSphere agents on each node monitor the state of the VxRail cluster using heartbeat signals on the VxRail external management network. If a node in question believes it has become separated by network from the rest of the cluster, it can perform an additional heartbeat check to the isolation IP addresses to confirm the problem.

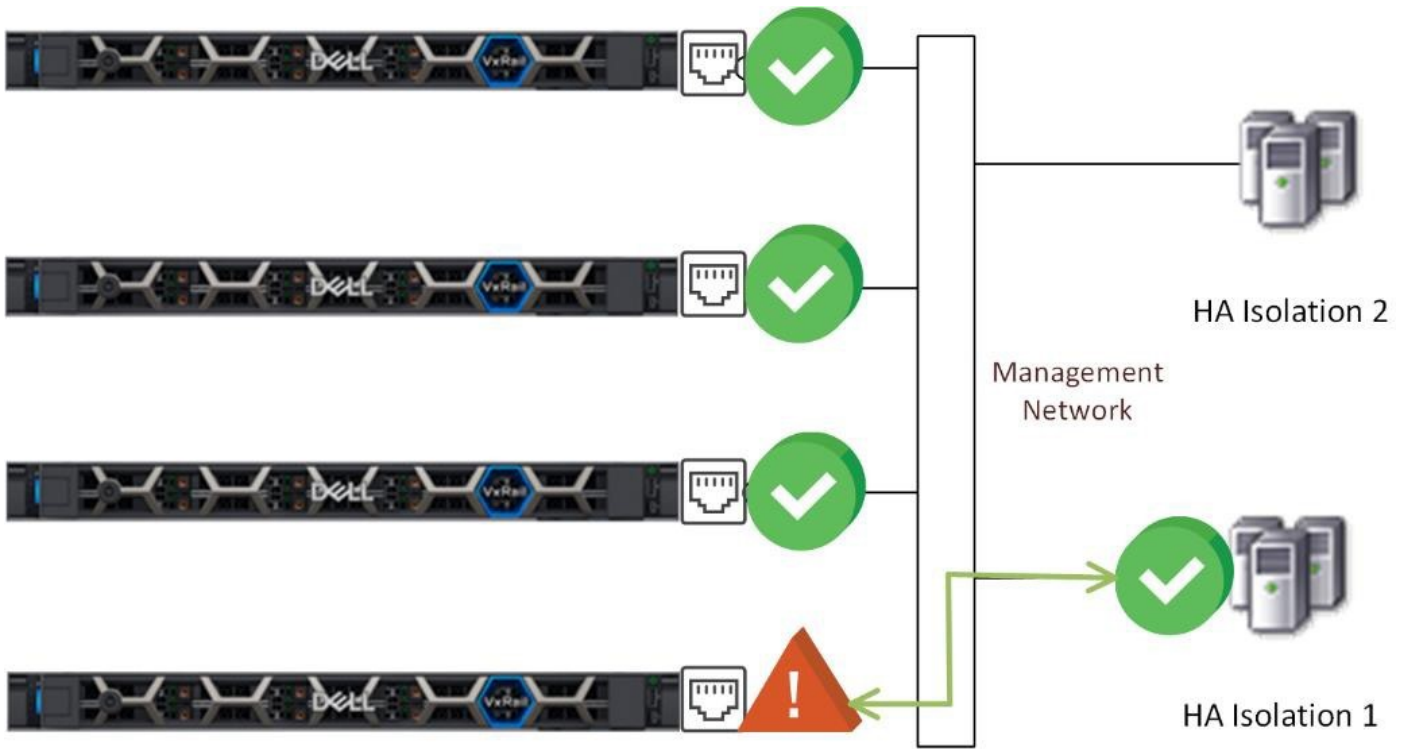


Figure 51. Confirming separation with VMware vSphere HA isolation IP addresses

The isolation IP addresses can be configured as part of the automated VxRail initial build process, or configured after the cluster is built. A minimum of two IP addresses reachable from the VxRail external management network should be selected for this purpose.

Plan the VxRail cluster implementation

VxRail is an SDDC in a cluster 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 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 SDDC based on your settings and input.

Before getting to this phase, perform planning and preparation steps to ensure a seamless integration of the final product into your data center environment.

Use the [Appendix C: VxRail Cluster 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.

Once you set up the VxRail cluster and complete the initial initialization phase to produce the final product, the configuration cannot easily be changed. Select the configurations that work most effectively for your organization during this planning and preparation phase.

Decide on VxRail single point of management

VxRail virtual infrastructure is defined and managed as a VMware vSphere cluster under a single instance of VMware vCenter Server.

Select whether to use one of the following:

- VxRail-managed VMware vCenter Server, which is deployed in the cluster during the initialization process.
- Customer-managed VMware vCenter Server, which is external to the cluster.

During the VxRail initialization process, select to deploy the VxRail-managed VMware vCenter Server on the cluster or deploy the cluster on an customer-managed VMware vCenter Server. Once the initialization process is complete, migrating to a new VMware vCenter Server single point of management requires professional services assistance and is difficult to change. Consider the ramifications during this planning and preparation phase and decide on the single point of management option that works most effectively for your organization.

Select customer-managed or VxRail-managed VMware vCenter Server

Consider the following requirements for selecting the VxRail vCenter Server:

- A VMware vCenter Server Standard license is included with VxRail cluster with vSAN and does not require a separate license. This license cannot be transferred to another VMware vCenter Server instance.
- VxRail Lifecycle Management supports the upgrade of the VxRail-managed VMware vCenter Server. VxRail does not support upgrading a customer-managed VMware vCenter Server using VxRail Lifecycle Management.
- DNS services are required for VxRail. With the VxRail-managed VMware vCenter Server, you can use the internal DNS supported within the VxRail cluster, or leveraging external DNS in your data center. The internal DNS option only supports naming services on which the VxRail cluster it is deployed. This option cannot support naming services outside of that cluster.

For a customer-managed VMware vCenter Server, consider the following requirements:

- The VMware vCenter Server Standard license included with VxRail cannot be transferred to a VMware vCenter Server instance outside of the VxRail cluster.
- Multiple VxRail clusters can be configured on a single customer-managed VMware vCenter Server, limiting the points of management.
- External DNS must be configured to support the VxRail cluster.
- Ensuring version compatibility of the customer-managed VMware vCenter Server with VxRail is the responsibility of the customer.
- You can configure the VMware VDS to support the VxRail cluster, or have VxRail deploy a VMware VDS and perform the configuration. This provides better control and manageability of the virtual networking in your data center, and consolidates the number of VMware VDS in your VMware vCenter Server instance.

- To deploy one or more dynamic VxRail clusters where the primary storage resource is going to be a remote VMware vSAN datastore, a customer-managed VMware vCenter Server is preferable. All clusters participating in sharing and receiving VMware vSAN datastore resources must reside in a common VMware vCenter Server instance.

The option to use the internal DNS or to deploy the VxRail cluster against a preconfigured VMware VDS requires VxRail 7.0.010 or later. For more details on the planning steps for a customer-managed VMware vCenter Server, 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 VMware VDS level. You can configure the required VxRail networks to a single VMware VDS, or you can deploy a second VMware VDS to isolate the VxRail management network traffic and the VxRail non-management network traffic.

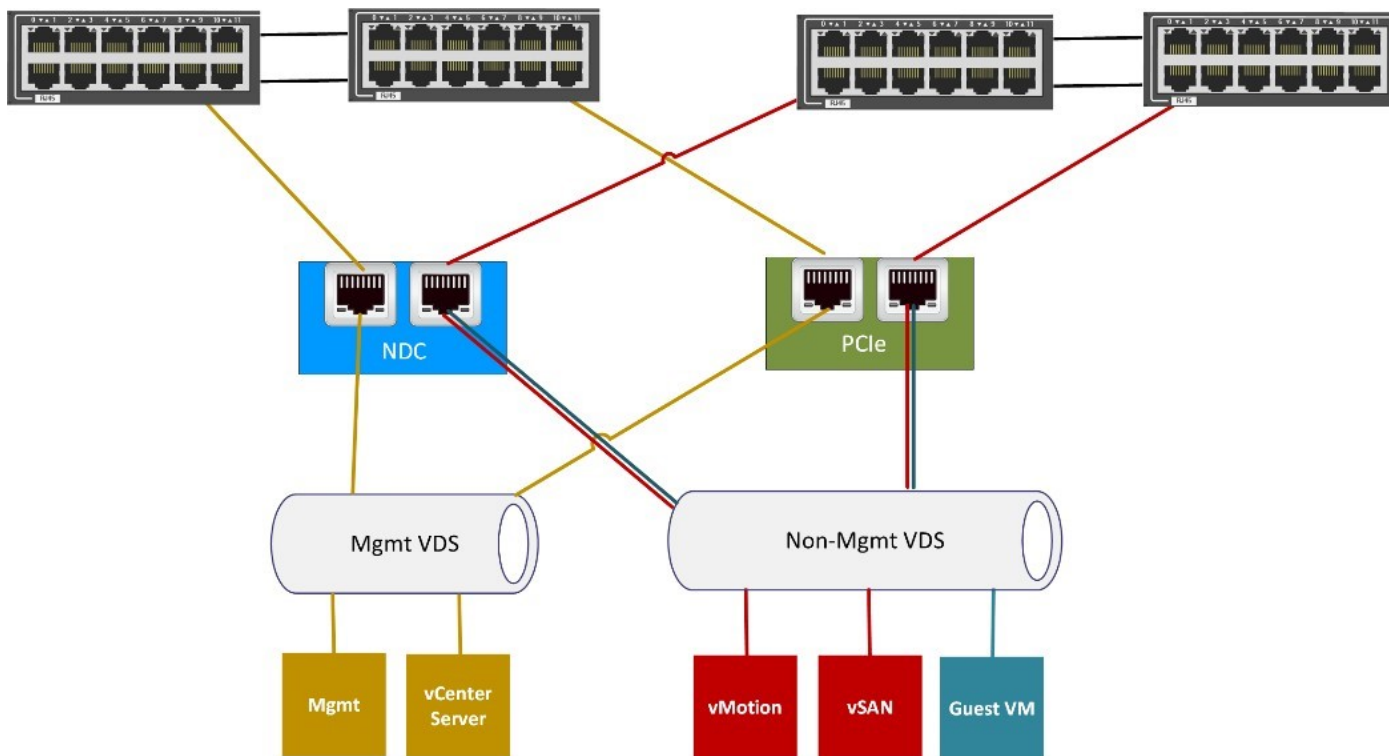


Figure 52. VxRail network segmentation with two VMware VDS

If your company or organization has stringent security policies regarding network separation, splitting the VxRail networks between two VMware VDS enables 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 VMware VDS.
- Use two VMware VDS to segment the VxRail management network traffic from the VxRail non-management traffic and guest VM network traffic.
- Deploy a separate VMware VDS to support guest virtual machine network traffic.




Figure 53. VxRail network segmentation with two VMware VDS

VxRail supports either a single VMware VDS or two VMware VDS as part of the initial implementation process. If your security posture changes after the VxRail cluster initial implementation has completed, a second VMware VDS can still be deployed, and the VxRail network traffic can be redirected to that second VMware VDS. Any additional VMware VDS 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 can select and assign the teaming and failover policy on each port group for each required VxRail network.

The following load-balancing policies are supported for VxRail clusters:

- **Route based on the originating virtual port:** After the VMware VDS selects an uplink for a VM or VMkernel adapter, it always forwards traffic through the same uplink. This option makes a simple selection based on the available physical uplinks. 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 VM based on the VM MAC address. This option requires more resources than the originating virtual port, but 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. 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 teaming policy settings:

- **Active/Active:** The default load balance policy set by VxRail is **Route based on physical NIC load**.
- **Active/Standby:** The default load balance policy set by VxRail is **Route based on originating virtual port**.

If the Ethernet adapters targeted to support VMware vSAN support RDMA, 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 VMware VDS to enable communication within the cluster, and out to IT management and the application user community.

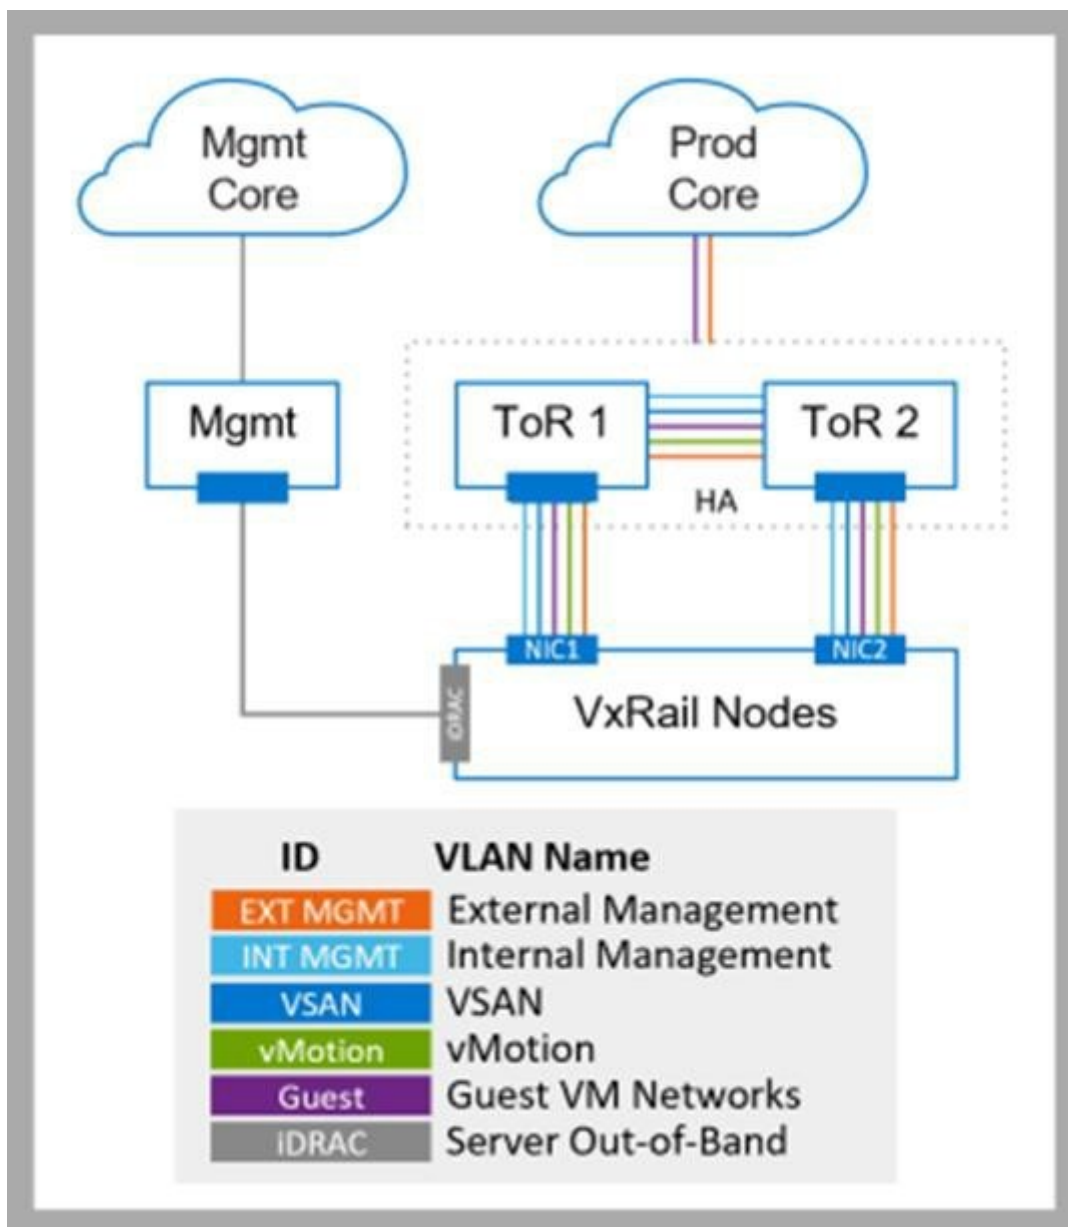


Figure 54. VxRail Logical Network Topology

VxRail has predefined logical networks to manage and control traffic within and outside of the cluster. Make VxRail logical networks 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 VMware vSAN datastore is required unless you plan to use FC storage as the primary storage resource with a dynamic cluster. Configure a network to support VMware vSphere vMotion, which is used to dynamically migrate virtual machines between VxRail nodes to balance workload. You also need an internal management network by VxRail for device discovery. You can skip the internal management network if you plan to use manual device discovery. All PowerEdge servers that serve as the foundation for VxRail nodes include a separate Ethernet port that enables OOB connectivity to the platform to perform hardware-based maintenance and troubleshooting tasks.

A separate network to support management access to the PowerEdge servers is recommended, but not required.

IP address considerations for VxRail networks

Assign IP addresses to the VxRail external management network, the VMware vMotion network, and any guest networks. If your cluster is going to use VMware vSAN as primary storage, IP addresses are required for the vSAN network. You can segment the external management network to separate subnets for the physical and logical components. Determine the IP address ranges reserved for each VxRail network.

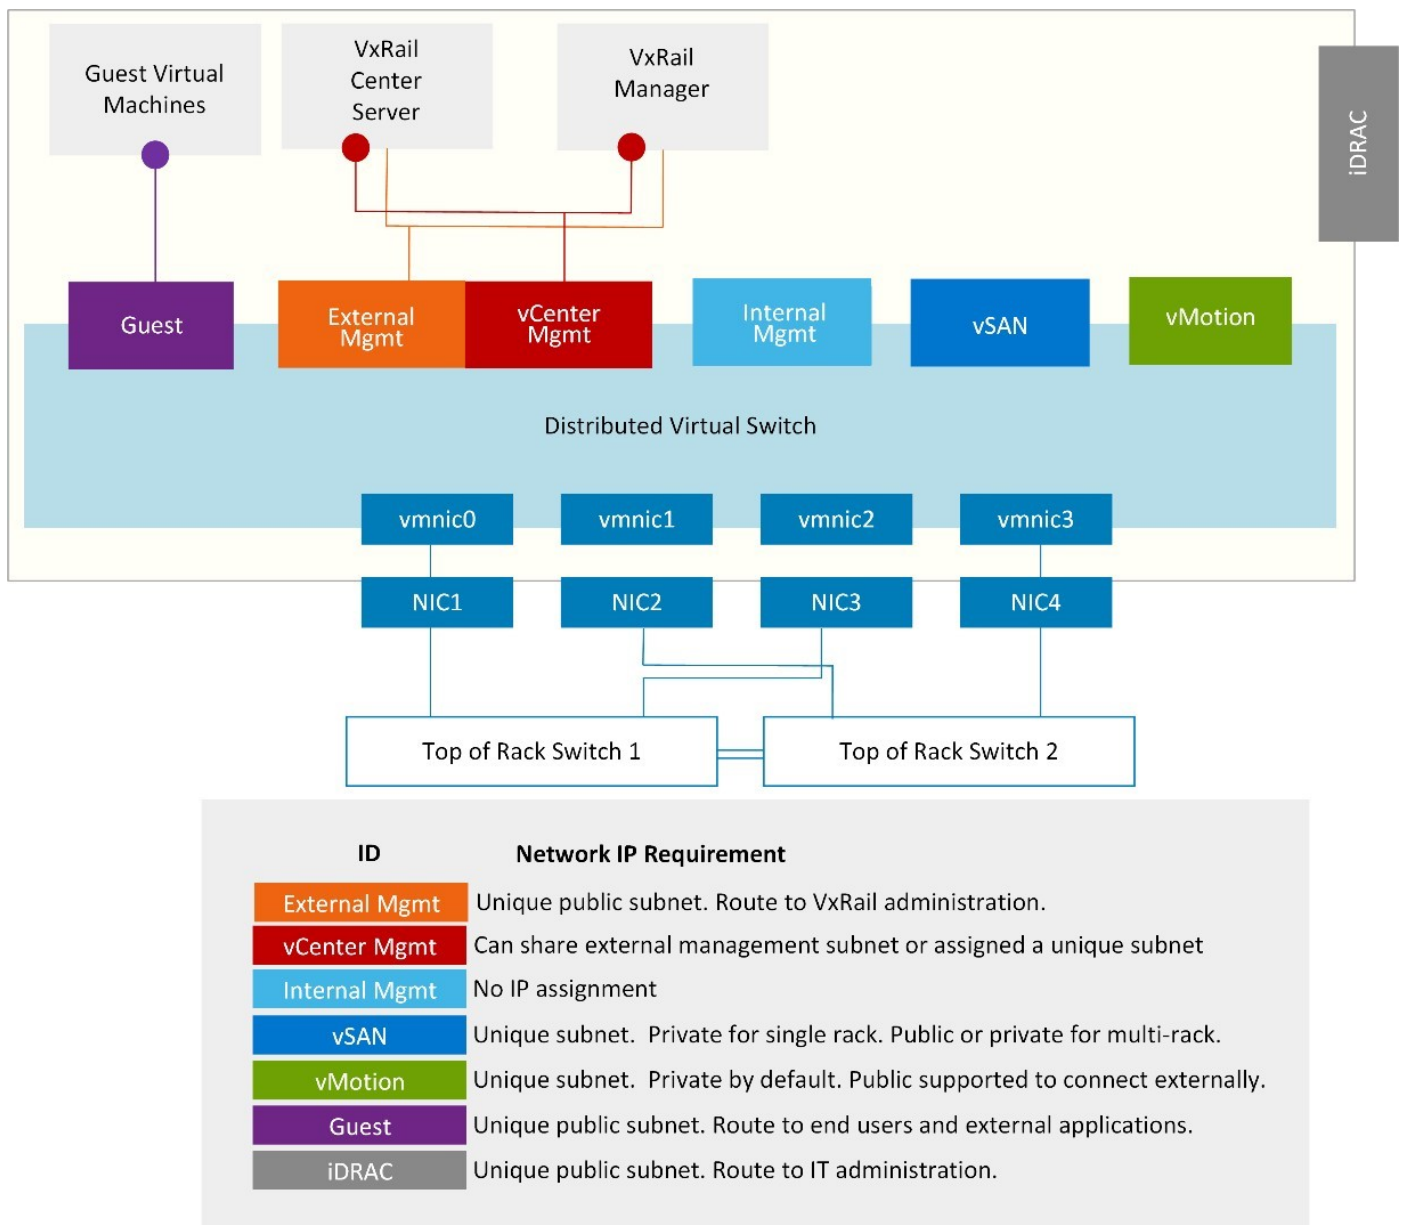


Figure 55. VxRail Core Network IP Requirements

- 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 VMware vCenter Server management network share the same non-private, routable IP address range. The VMware vCenter Server management network provides connectivity to VxRail Manager and the embedded VMware vCenter Server instance. Starting with 7.0.350, if you do not want VxRail Manager and the VMware vCenter Server instance to share the same subnet with the VxRail (ESXi) nodes, 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, VxRail does 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 multi-rack 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 multi-rack expansion that will use a different subnet range on any expansion racks, reserve a routable IP address range.

VLAN considerations for VxRail networks

VLANs define the VxRail logical networks within the cluster and the method used to control the paths that a logical network can pass through. A VLAN is a numeric ID that is assigned to a VxRail logical network. The same VLAN ID is also configured on the individual ports on your ToR switches, and on the virtual ports in the VMware VDS during the automated implementation process.

When an application or service in the VxRail cluster sends a network packet on the VMware VDS, the VLAN ID for the logical network is attached to the packet. The packet can only be able to pass through the ports on the ToR switch and the VMware VDS where there is a match in VLAN IDs. You should isolate the VxRail logical network traffic using separate VLANs. This is recommended, but not required. A flat network is recommended only for test, non-production purposes.

Meet with the network team and virtualization team to plan the VxRail network architecture.

- The virtualization team discuss with the application owners 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 required to support traffic from non-management virtual machines.
- If you plan to have multiple independent VxRail clusters, use different VLAN IDs across multiple VxRail clusters to reduce network traffic congestion.
- The network team must plan the following:
 - Define the pool of VLAN IDs needed to support the VxRail logical networks, and determine which VLANs restrict traffic to the cluster, and which VLANs can pass through the switch up to the core network.
 - Plan to configure the VLANs on the upstream network, and on the switches attached to the VxRail nodes.
 - 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
- Virtual Machine

VxRail assigns the settings that you specify for each logical networks during initialization.

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

External Management network

External Management network supports communications to the ESXi hosts, and has common network settings with the VMware vCenter Server Management Network. All VxRail external management traffic is untagged by default and should be able to go over the native VLAN on your ToR 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 are deployed on a single set of ToR switches. To support using a tagged VLAN for the VxRail external management network, configure the VLAN on the ToR switches. Configure trunking for every switch port that is connected to a VxRail node to tag the external management traffic.

vCenter Management network


The vCenter Management Network hosts the VxRail Manager and the VxRail-managed VMware 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.

Internal 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 ToR 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 fails 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.

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

To leverage vSAN for VxRail cluster storage resources, configure a VLAN for the **vSAN** network and the **VMware vSphere vMotion** network. Configure a VLAN for each network on the ToR switches, and include the VLANs on the trunked switch ports that are connected to VxRail nodes.

Virtual Machine

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 or afterward using the VMware vClient after initial configuration is complete. Dedicated VLANs are preferred to divide VM 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. When you create VMs in the VMware vSphere Web Client to run your applications and services, you can assign the VM 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 5. Network configuration

Network configuration table	Action
Row 1	Enter the external management VLAN ID for VxRail management network (VxRail Manager, ESXi, VMware vCenter Server, 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 .

Table 5. Network configuration (continued)

Network configuration table	Action
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.
Row 3	Enter a VLAN ID for VMware vSphere vMotion (enter 0 in the VLAN ID field for untagged traffic).
Row 4	Enter a VLAN ID for vSAN, if applicable (enter 0 in the VLAN ID field for untagged traffic).
Row 5-6	Enter a Name and VLAN ID pair for each VM guest network that 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).
Row 7	Enter the VMware vCenter Server Network VLAN ID (if different from the external management 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 VMware 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 must 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 are going to use the VMware vCenter Server instance that is deployed in the VxRail cluster or use an external VMware vCenter Server already operational in your data center.
 - If you use the VMware vCenter Server instance that is supplied by VxRail, you must reserve an IP address for VMware vCenter Server.
- Decide if you are going to use VMware vSphere Log Insight that can be deployed in the VxRail cluster.
 - For VxRail version 7.0 and earlier, if you use the VMware vCenter Server instance that is deployed in the VxRail cluster, you have the option to deploy VMware 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 VMware vSphere Log Insight in the VxRail cluster, you must reserve one IP address.
 - VMware vRealize Log Insight is not an option for deployment during the initial VxRail configuration process starting with version 7.0.010.
 - If you use an external VMware vCenter Server already operational in your data center for VxRail, VMware vSphere Log Insight cannot be deployed.

Use the following table to determine the number of public IP addresses required for external connectivity:

Table 6. IP addresses required for external connectivity

Component	Condition
VxRail Node	One per VxRail Node
VxRail Manager	One

Table 6. IP addresses required for external connectivity (continued)

Component	Condition
VMware vCenter Server	<ul style="list-style-type: none">• If you are supplying VMware vCenter Server for VxRail: 0• If VxRail is supplying VMware vCenter Server: 1
VMware vSphere Log Insight	<ul style="list-style-type: none">• If you are supplying VMware 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 VMware vSphere 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.

Table 7. Network configuration table

Network configuration table	Action
Row 8	Enter the subnet mask for the VxRail External Management network.
Row 9	Enter the default gateway for the VxRail External Management network.

Plan network settings for VMware vCenter Server management network

Before VxRail 7.0.350, the VxRail-managed VMware vCenter Server and VxRail Manager were placed on the same VxRail external management network with the ESXi hosts, and a shared VLAN and subnet. From VxRail 7.0.350 and later, these virtual appliances can 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 that is separate from the ESXi hosts, make a note of the network properties.

Table 8. Network configuration table

Network configuration table	Action
Row 10	Enter the VMware vCenter Server network subnet mask.
Row 11	Enter the VMware 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 pool of unused IP addresses.

Record the IP address range for the ESXi hosts.

Table 9. Record the IP address range for the ESXi hosts

Network configuration table	Action
Row 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.

Table 10. Record the IP address of each ESXi host

Network configuration table	Action
Row 29 and 32	Enter the IP addresses for the ESXi hosts.

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

Table 11. Record the permanent IP address for VxRail Manager

Network configuration table	Action
Row 17	Enter the permanent IP address for VxRail Manager.

If you are going to deploy the embedded VMware vCenter Server on the VxRail cluster provided with VxRail, record the permanent IP address for VMware vCenter Server. Leave these entries blank if you are going to provide an external VMware vCenter Server for VxRail.

Table 12. Record the permanent IP address for VMware vCenter Server

Network configuration table	Action
Row 34	Enter the IP address for VxRail vCenter.

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

Select a top-level domain

DNS is a requirement for VxRail, so select a domain where the naming services can support that domain.

Begin the process by selecting the domain to use for VxRail and assign to the fully qualified hostnames.

Table 13. Select a top-level domain

Network configuration table	Action
Row 15	Enter the top-level domain.

Select a hostname for VxRail Manager

Assign a hostname for VxRail Manager. The domain is also automatically applied to the chosen hostname.

Follow the naming format that is selected for the ESXi hosts to simplify cluster management.

Table 14. Select a hostname for VxRail Manager

Network configuration table	Action
Row 16	Enter the hostname for VxRail Manager.

Select the ESXi hostnames

All VxRail nodes in a cluster require hostnames.

From VxRail 7.0.010 and later, you can use any host naming convention that you want, if 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, ensure that you follow the rules that are stated here.

All ESXi hostnames in a VxRail cluster are defined by a naming convention that consists of:

- 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 .)
- 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 iteration 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. For example: `my-vcenter.local`.

Table 15. ESXi hostname examples

Parameter	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 that is the chosen method.

Table 16. To auto-assign the ESXi hostnames

Network configuration table	Action
Row 18–22	Enter an example of your chosen ESXi host-naming scheme. Be sure to show your chosen prefix, separator, iterator, offset, suffix, and domain.

If you are going to assign the ESXi hostnames manually, capture the name for each ESXi host that is planned for the VxRail initial build operation.

Table 17. To manually assign the ESXi hostnames

Network configuration table	Action
Row 23–26	Enter the reserved hostname for each ESXi host.

Select a hostname for the VxRail-managed VMware vCenter Server

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

If you want to deploy a new VMware vCenter Server on the VxRail cluster, you must specify a hostname for the VxRail-managed VMware vCenter Server. The domain is also automatically applied to the chosen hostname. Follow the naming format that is selected for the ESXi hosts to simplify cluster management.

Table 18. VMware vCenter Server hostname

Network configuration table	Action
Row 33	Enter an alphanumeric string for the new VMware vCenter Server hostname. The specified domain is 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 the time zone and NTP server

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

An NTP server is not required, but it is recommended. If you provide an NTP server, the VMware vCenter Server is configured using 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).

From VxRail 7.0.400 and later, an outside service such as the ones offered by pool.ntp.org or similar websites are supported. Ensure that the NTP hostname or IP address is accessible from the VxRail External Management Network which the VxRail nodes will be connected to and is functioning properly.

Table 19. Set the time zone and NTP server

Network configuration table	Action
Row 12	Enter your time zone.
Row 13	Enter the hostnames or IP addresses of your NTP servers.

Set the DNS for VxRail management components

From VxRail 7.0.010 and later, you can either use an internal DNS included with VxRail-managed VMware vCenter Server or use an external DNS in your data center.

If you choose to use the internal DNS method, you can skip the steps that are mentioned here to set up DNS.

If the internal DNS option is not selected, one or more external, customer-managed 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, VMware vCenter Server, and so on). Ensure that the DNS IP address is accessible from the network to which VxRail Manager is connected and is functioning properly.

Table 20. IP addresses for DNS servers

Network configuration table	Action
Row 14	Enter the IP addresses for your DNS servers.

You must create lookup records in your selected DNS for every VxRail management component that you are deploying in the cluster and are assigning a hostname and IP address. These components can include VxRail Manager, VxRail vCenter Server, Log Insight, and each ESXi host in the VxRail cluster. The DNS entries must support both forward and reverse lookup.








 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 56. 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 57. 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. VMware vSphere vMotion and vSAN IP addresses do not require hostnames, so no entries are required in the DNS server.

Prepare the customer-managed VMware vCenter Server

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

You must complete certain prerequisites before VxRail initial implementation if you use a customer-managed VMware vCenter Server as the VxRail cluster management platform. During the VxRail initialization process, it connects to your customer-managed VMware vCenter Server to perform necessary validation steps, and perform configuration steps, to deploy the VxRail cluster on your VMware vCenter Server instance.

- Determine if your customer-managed VMware vCenter Server is compatible with your VxRail version. See the Knowledge Base article [VxRail: VxRail and External VMware vCenter Server Interoperability Matrix](#) on the [Dell VxRail Support site](#) for the latest support matrix.
- Enter the FQDN of your selected, compatible customer-managed VMware vCenter Server in the [Appendix A: VxRail Network Configuration Table](#).

Table 21. Enter the FQDN

Network configuration table	Action
Row 35	Enter the FQDN of the customer-supplied vCenter Server.

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

Table 22. Enter the SSO domain

Network configuration table	Action
Row 36	Enter the single sign-on(SSO) domain for the customer-managed VMware vCenter Server (for example, vsphere.local).

- The VxRail initialization process requires login credentials to your customer-managed VMware vCenter Server. 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 set of credentials in your VMware vCenter Server for this purpose. Two new roles are created and assigned to this user by your Dell delivery services.

Table 23. Enter the administrative username and password

Network configuration table	Action
Row 37	Enter the administrative username/password for the customer-managed VMware vCenter Server, or the VxRail non-admin username/password you create on the customer-managed VMware vCenter Server.

- You must create a set of credentials in the customer-supplied vCenter for VxRail management with no permissions and no assigned roles. These credentials are assigned to a role with limited privileges during the VxRail initialization process, and then assigned to VxRail to enable connectivity to the customer-managed VMware vCenter Server after initialization completes.
 - If this is the first VxRail cluster on the customer-managed VMware vCenter Server, enter the credentials for the customer-managed VMware vCenter Server.
 - If you already have an account for a previous VxRail cluster in the customer-managed VMware vCenter Server, enter those credentials.

Table 24. Enter the VxRail management username and password

Network configuration table	Action
Row 38	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 data center, or select an existing Data center on the customer-managed VMware vCenter Server.

Table 25. Enter the name of a data center

Network configuration table	Action
Row 39	Enter the name of a data center on the customer-managed VMware 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-managed VMware vCenter Server.

Table 26. Enter the name of the cluster

Network configuration table	Action
Row 40	Enter the name of the cluster that you are going to use for VxRail.

Prepare a customer-managed 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-managed virtual-distributed switches.

Before VxRail 7.0.010, if you deployed the VxRail cluster on an external, customer-managed VMware vCenter Server, a virtual-distributed switch was configured on the VMware vCenter Server 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 VMware vCenter Server instance, and then attached the VxRail networks to the port groups on the virtual-distributed switch. Depending on the target version planned for your VxRail cluster, you can pre-configure one or two virtual-distributed switches on your external VMware vCenter Server instance to support VxRail networking.

- From VxRail 7.0.010 and later, you have the choice of configuring a single virtual distributed switch to the external VMware vCenter Server before the initial cluster build process.
- From VxRail 7.0.130 and later, you have the choice of configuring one or two virtual-distributed switches to the external VMware vCenter Server instance before the initial cluster build process.

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

- Unless your data center already has a VMware vCenter Server instance compatible with VxRail, deploy a VMware vCenter Server instance that serves 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 VMware vCenter Server instance.
- Configure a port group for each of the required VxRail networks. Use naming standards that clearly identify the VxRail network traffic type.
- Configure the VLAN assigned to each required VxRail network on the respective port group. The VLANs for each VxRail network traffic type are in the VxRail Networks section in [Appendix A: VxRail Network Configuration Table](#).
- Configure at least two uplinks on the virtual-distributed switch 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 VxRail networks, configure the LACP policy on the virtual-distributed switch, and apply the policy to the appropriate port group or port groups.

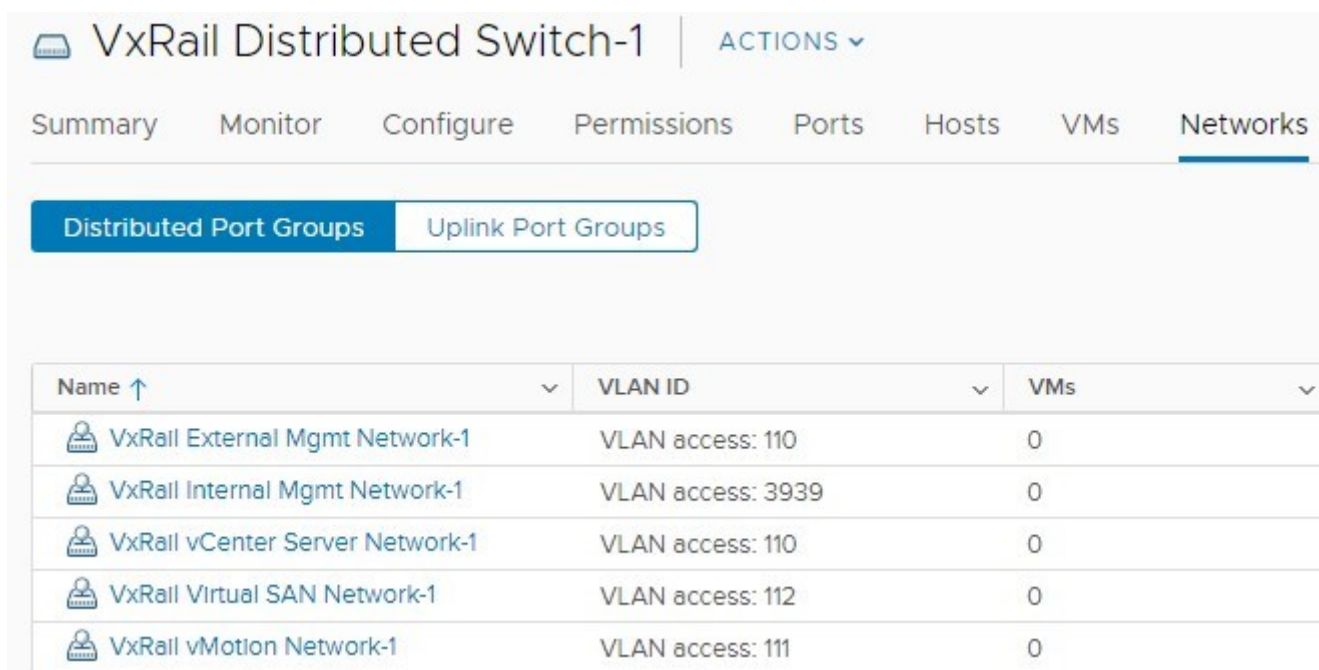


Figure 58. Sample port groups on customer-managed virtual distributed switch

Refer the configuration settings that are 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-managed virtual-distributed switch. The settings that are used by the automated initial build process can be found in [Appendix E: Virtual Distributed Switch Portgroup Default Settings](#).

Table 27. Network Configuration Table

Network configuration table	Action
Row 41	Enter the name of the virtual-distributed switch that will support the VxRail cluster networking.
Row 42	If a decision is made to configure two virtual-distributed switches, enter the name of the second virtual-distributed switch.
Row 43	Enter the name of the port group that will enable connectivity for the VxRail external management network.
Row 44	Enter the name of the port group that will enable connectivity for the VxRail-managed VMware vCenter Server network.
Row 45	Enter the name of the port group that will enable connectivity for the VxRail internal management network.
Row 46	Enter the name of the port group that will enable connectivity for the VMware vSphere vMotion network.
Row 47	Enter the name of the port group that will enable connectivity for the vSAN network.

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

Configuring port groups on the virtual distributed switch for any guest networks you want is not required for the VxRail initial build process. These port groups can be configured after the VxRail initial build process is complete. Follow a distinctive naming standard for these distributed port groups.

Prepare LAG on a customer-managed 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 LAG against one or more VxRail networks on the customer-managed virtual-distributed switch.

From VxRail 7.0.130 and later, you can configure LAG against the VxRail non-management networks, which include the vSAN network and VMware vSphere vMotion network.

Ensure that you meet the following prerequisites to enable LAG on the VxRail non-management networks:

- Configure four ports from each VxRail node to support VxRail networking.
 - Two ports to support VxRail management networks. LAG is not supported on these networks.
 - Two ports to support VxRail non-management networks. LAG is supported on these networks.
- LAG can be configured on the vSAN network, VMware vSphere vMotion network, or both.
- If you are not configuring LAG on the VMware vSphere vMotion network, assign this network to the same uplinks supporting the VxRail management networks.
- Configure the adjacent ToR switches to support LAG. See the guides provided by your switch vendor to perform this task.

Complete the following tasks on the virtual-distributed switch on your customer-managed VMware vCenter Server instance to support LAG:

- Configure an LACP policy on the virtual distributed switch.

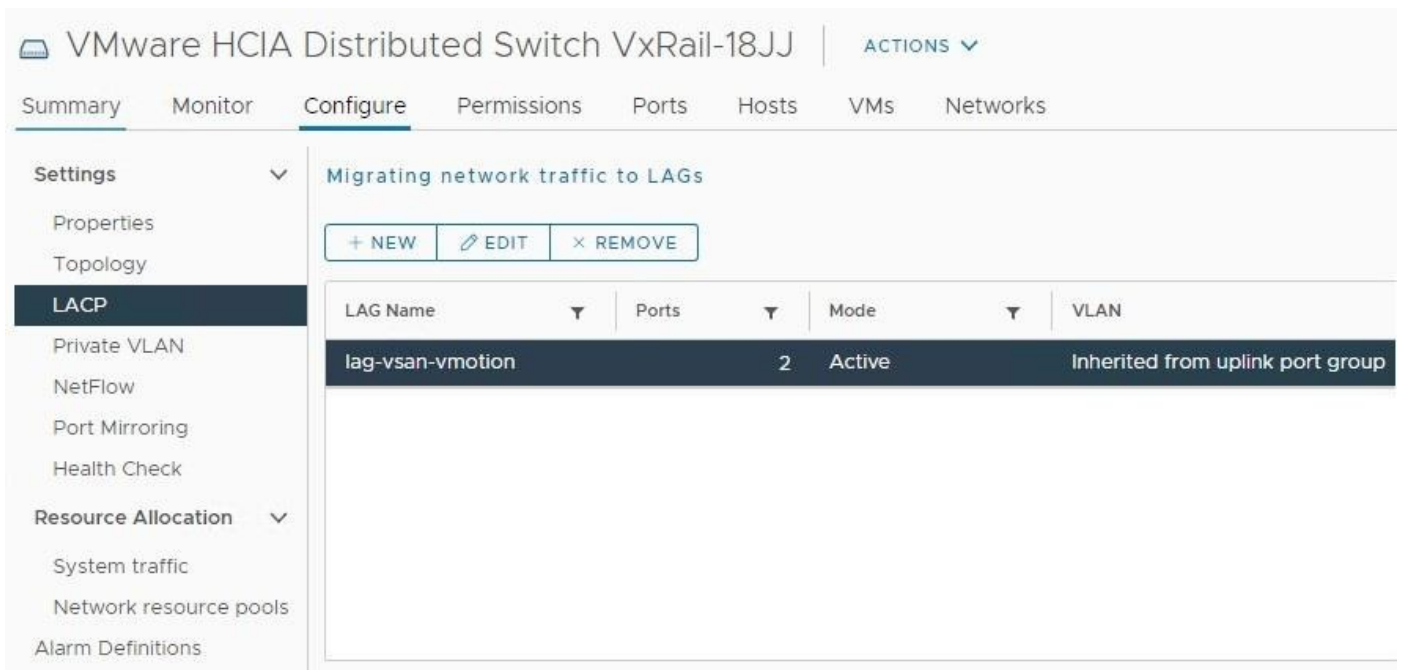


Figure 59. Sample LACP policy configured on virtual distributed switch

- Configure the teaming and failover policy on the port groups that are targeted for LAG.

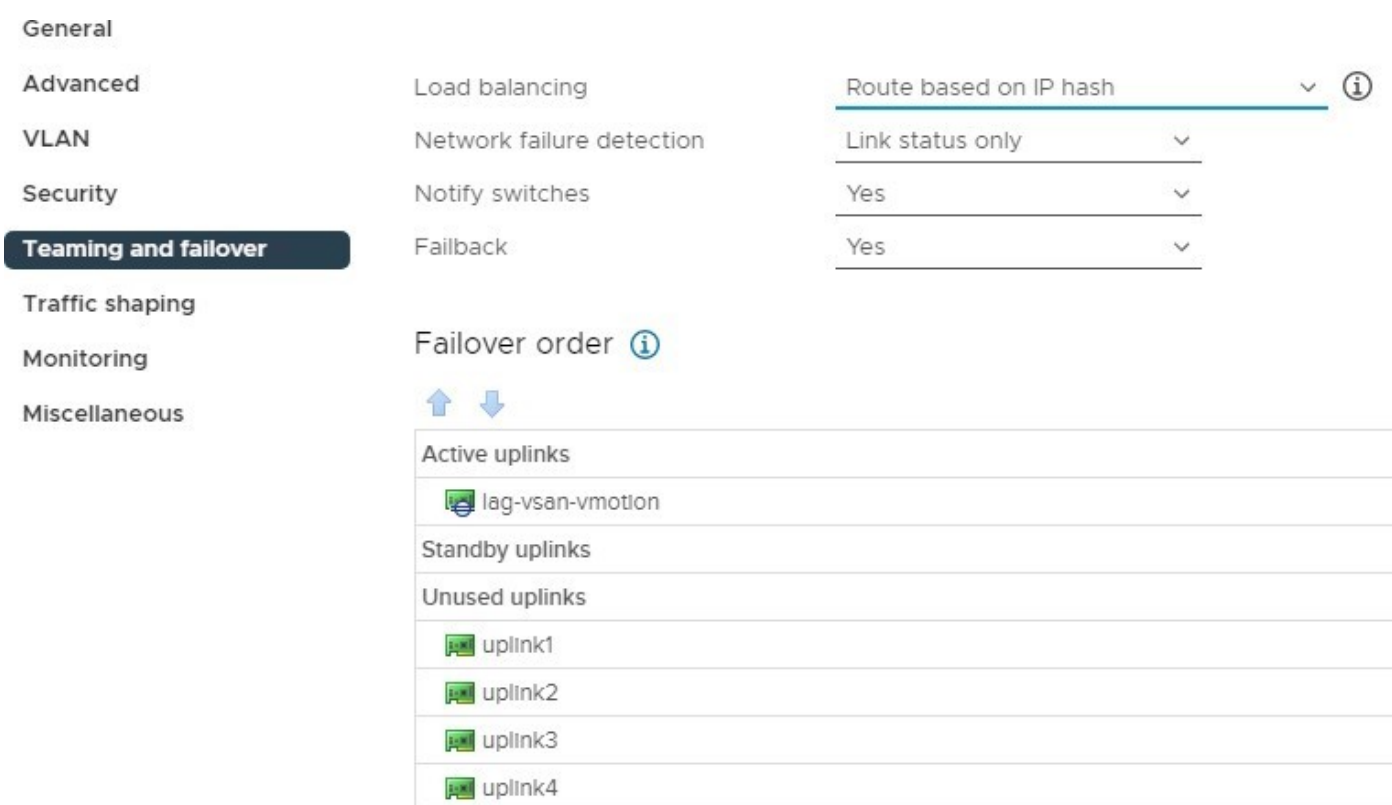


Figure 60. Sample LACP policy configured as active uplink in teaming and failover policy

Reserve IP addresses for VxRail-managed VMware vSphere vMotion network

An IP address is required for the VMware vSphere vMotion network for each ESXi host in the VxRail cluster.

A private address range is acceptable if you decide the VMware vSphere vMotion network will not be routable. If your plans include the ability to migrate VMs outside of the VxRail cluster, consider it while selecting the IP address scheme.

From VxRail 7.0.010 and later, you can choose to have the IP addresses assigned automatically during the VxRail initial build, or manually select the IP addresses for each ESXi host. If the VxRail version is earlier than 7.0.010, the 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.

Table 28. Record the IP address range for auto-assignment method

Network configuration table	Action
Rows 48-49	Enter the starting and ending IP addresses for VMware vSphere vMotion.

For the manual assignment method, record the IP addresses.

Table 29. Record the IP addresses for the manual assignment method

Network configuration table	Action
Rows 50-53	Enter the IP addresses for VMware vSphere vMotion.

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

Table 30. Record the subnet mask and gateway

Network configuration table	Action
Row 54	Enter the subnet mask for VMware vSphere vMotion.
Row 55	Enter the gateway for VMware vSphere 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 that is 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.

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

Table 31. Record the IP address range for auto-assignment method

Network configuration table	Action
Rows 56-57	Enter the starting and ending IP addresses for vSAN.

For the manual assignment method, record the IP addresses.

Table 32. Record the IP addresses for the manual assignment method

Network configuration table	Action
Rows 58-61	Enter the IP addresses for vSAN.

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

Table 33. Record the subnet mask and gateway

Network configuration table	Action
Row 62	Enter the subnet mask for vSAN.
Row 63	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 VMware vRealize Log Insight solution, or no logging solution.

You can only select the VMware vRealize Log Insight option if:

- You can deploy the VMware vCenter Server instance that 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 VMware vRealize Log Insight virtual appliance onto the cluster during the initial build.
- It is possible to deploy VMware vRealize Log Insight after the VxRail initial build operation is complete.

If you choose the VMware vRealize Log Insight option, the IP address that is assigned to VMware vRealize Log Insight must be on the same subnet as the VxRail management network.

Table 34. Hostname and IP address for VMware vRealize Log Insight

Network configuration table	Action
Row 64	Enter the hostname for VMware vRealize Log Insight.
Row 65	Enter the IP address for VMware vRealize Log Insight.

If a syslog server is already deployed in your data center and you are going to use it as a logging solution, capture the IP address.

Table 35. IP address of the syslog server

Network configuration table	Action
Row 66	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 service representative needs 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, assign passwords to the **root** account. You can use one password for each ESXi host or apply the same password to each host.
- For VxRail Manager, assign a password to the **root** account (Row 1). This credential is for access to the console.
- To access the VxRail Manager web interface, use the **administrator@<SSO Domain>** credentials.
 - If you deploy the VxRail-managed VMware vCenter Server, VxRail Manager and VMware vCenter Server, share the same default administrator login **administrator@vsphere.local**. Enter the password that you want to use (Row 2).
 - If you use a customer-managed VMware vCenter Server, VxRail Manager uses the same **administrator@<SSO Domain>** login credentials that you use to access the customer-managed VMware vCenter Server.
- If you deploy the VxRail-managed VMware vCenter Server:
 - Enter the **root** password for the VxRail-managed VMware vCenter Server (Row 3).
 - Enter a password for management for the VxRail-managed VMware vCenter Server (Row 4).
- If you deploy VMware vRealize Log Insight:
 - Enter a password for **root** (Row 6).
 - Enter a password for **admin** (Row 7).

Passwords must comply with the following 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 VMware vSphere password and VMware vCenter Server password documentation.

Planning vSAN Witness Networking

You can skip this section if your plans do not include deploying a VxRail stretched cluster or a VxRail 2-node cluster.

VxRail supports two types of vSAN clusters that require a witness to serve as a quorum if there is a failure. They are the VxRail stretched cluster and the VxRail 2-node cluster. If your plans include deploying one of these VxRail cluster types, planning for a vSAN witness must be considered first before building the VxRail cluster. Full details on the preparation and planning steps for these two VxRail cluster types can be found in the [Dell VxRail vSAN Stretched Cluster Planning Guide](#) or the [Dell vSAN 2- Node Cluster Planning and Preparation Guide](#). This chapter provides an overview of the networking planning steps for the required vSAN witness.

Overview of vSAN Witness for VxRail stretched cluster

For a stretched cluster with sites across geographic distance, plan for the following network requirements:

- The vSAN witness must be positioned in a separate failure domain to ensure resiliency.
- The VxRail External Management Network must be stretched across the primary and secondary sites on a Layer 2 network to ensure proper connectivity for the management components if there is a failure.
- The networks supporting the vSAN witness must be able to route between both VxRail cluster sites and the witness site.

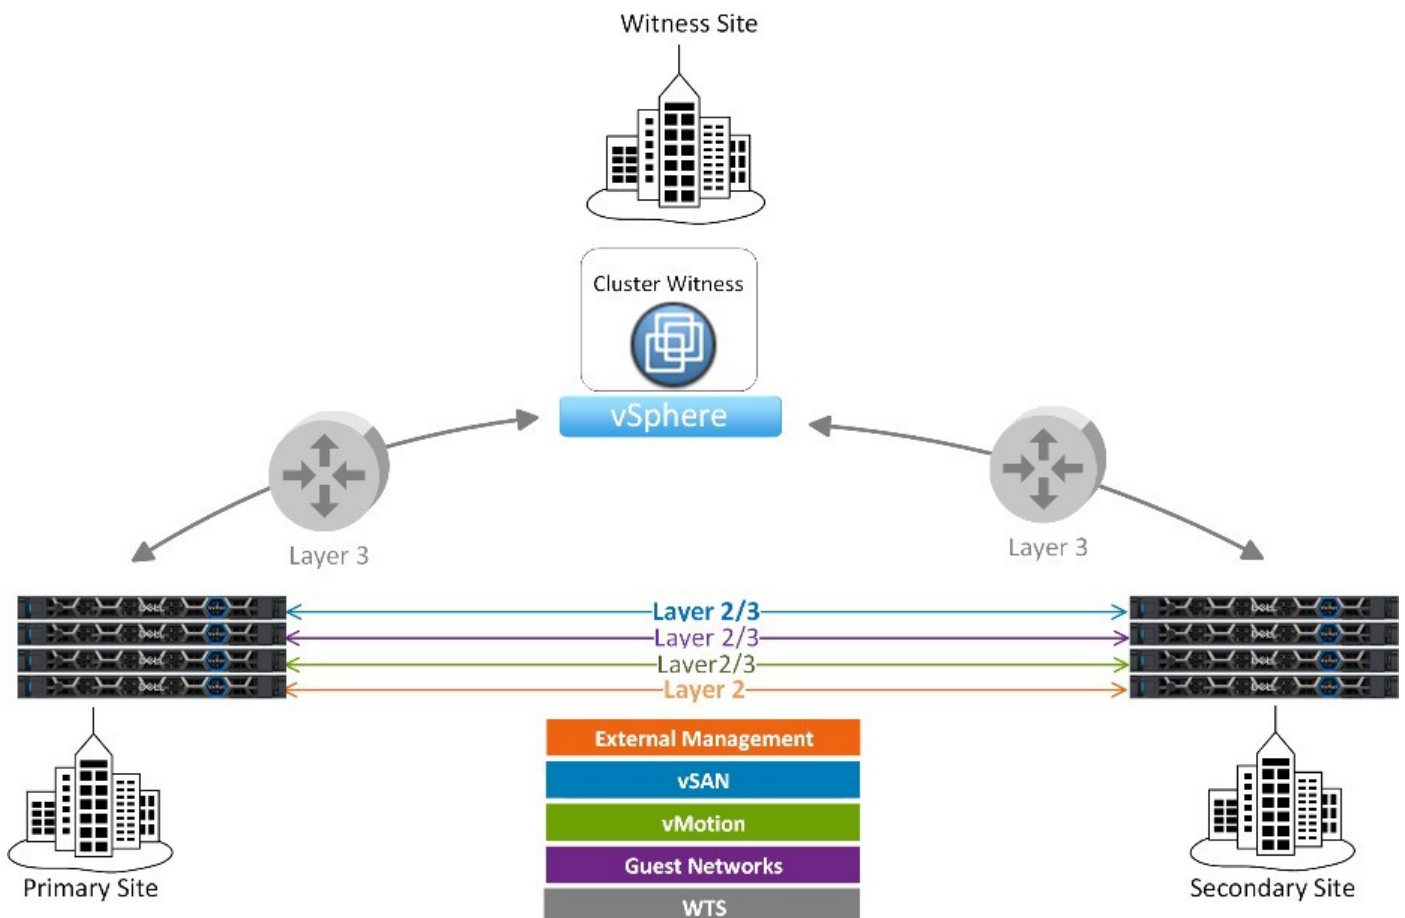


Figure 61. Stretched cluster with vSAN witness at third site

Overview of vSAN Witness for VxRail 2-node cluster

For a VxRail 2 node cluster, it is not necessary to place the nodes at separate sites across geographic distance. This type of cluster can use a compatible vSAN witness at a third site, or a vSAN witness deployed local to the VxRail cluster.

As there are only two nodes in this cluster type, the VxRail management, vSAN, and VMware vSphere vMotion networks can be placed on Layer 2 networks between the nodes. If the vSAN witness is to be deployed local to the two VxRail nodes, Layer 2 can be used for the Witness Traffic Separation network. A vSAN witness at a remote site can also be considered to support a VxRail 2- node cluster.

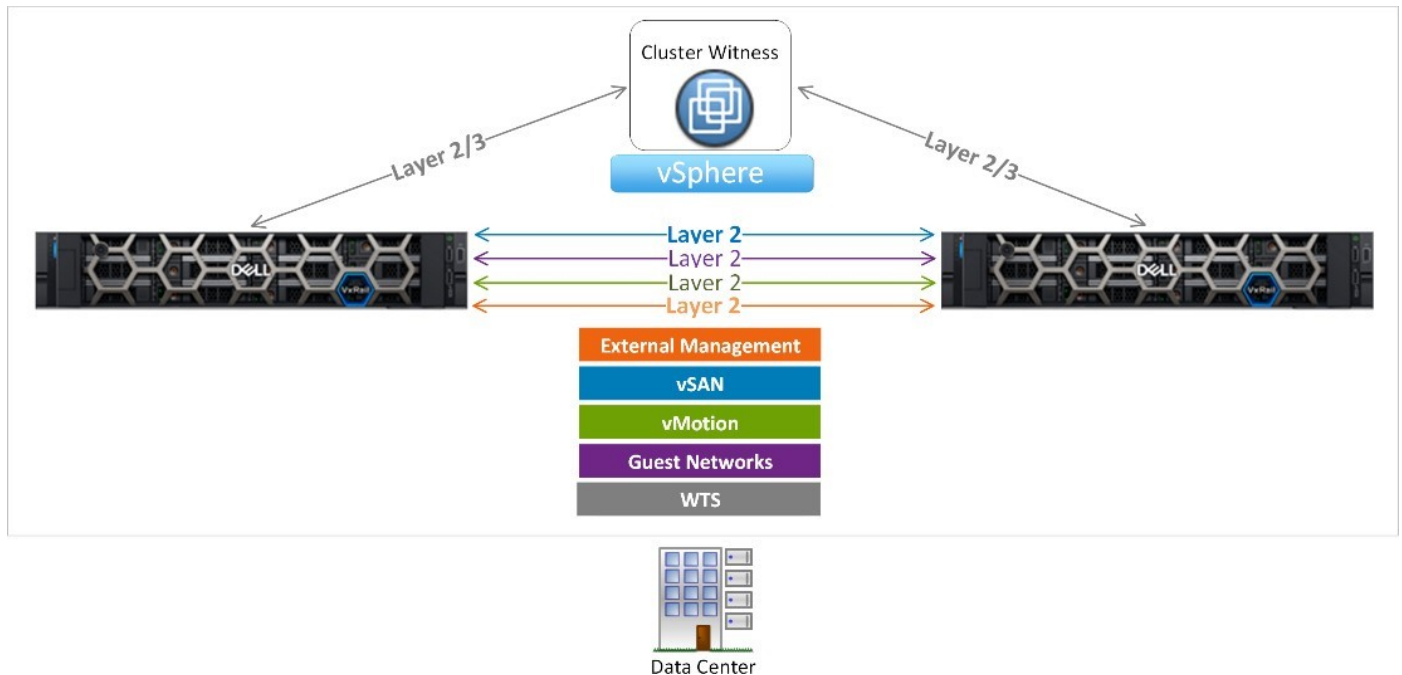


Figure 62. 2-node cluster with local vSAN witness

Networking rules for vSAN witness and VxRail

For both the VxRail stretched cluster and the VxRail 2-node cluster, the networks on the vSAN witness virtual appliance must be properly configured to ensure proper cluster operations.

You must follow these guidelines:

- The vSAN witness virtual appliance must be configured with two separate networks during deployment:
 - One for management called the **Management Network**.
 - One to support vSAN witness traffic called the **Secondary Network**.
- The IP address that is assigned for vSAN witness management on the Management Network must be reachable by the VMware vCenter Server instance supporting the VxRail cluster.
 - The vSAN witness must be added to the VMware vCenter Server inventory as an ESXi host to support vSAN witness traffic with the VxRail cluster.
- Each VxRail node and the vSAN witness must be able to connect over the vSAN witness traffic network.
 - An IP address is assigned to each VxRail node in the cluster to support vSAN witness network traffic.
 - An IP address is assigned to the vSAN witness virtual appliance also to support vSAN witness network traffic on the **Secondary Network**.

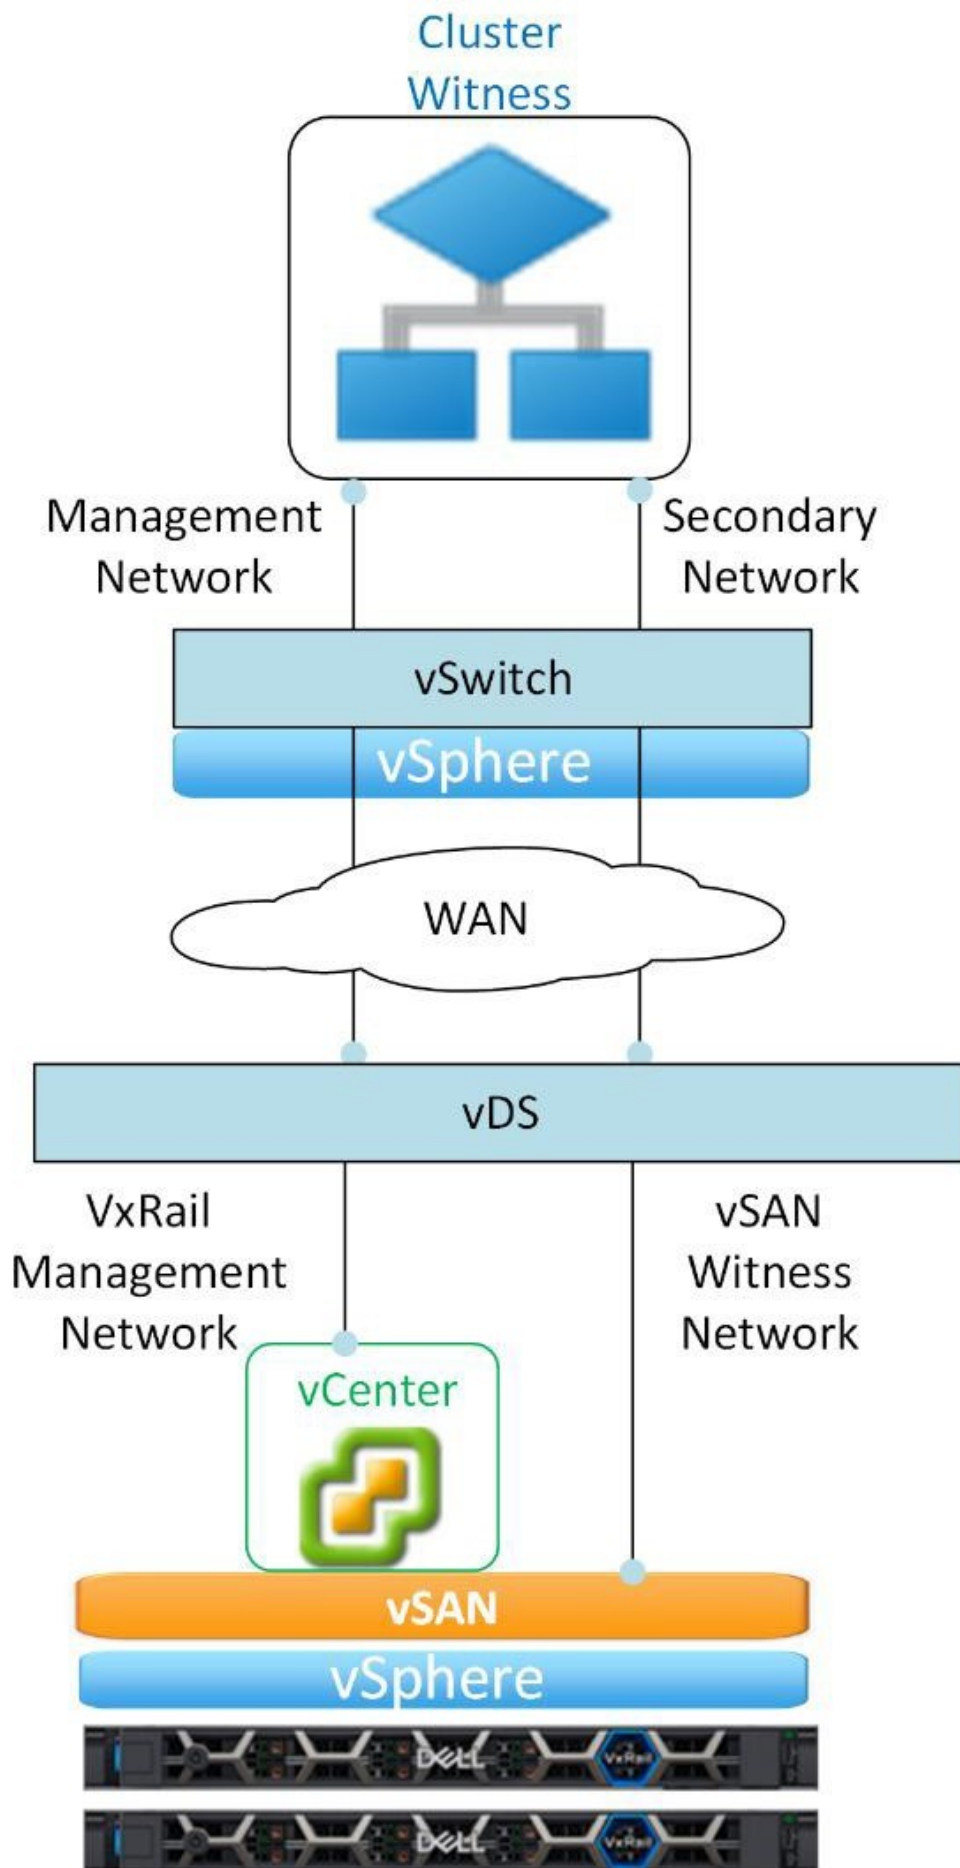


Figure 63. Network relationships for vSAN witness and VxRail

- If it is a VxRail-managed VMware vCenter Server, then the vSAN management traffic on the **Management Network** must route to the VxRail-managed VMware Server network. VxRail Manager configures this network during the initial build process.
- For a 2-node cluster planned with only Layer 2 networking, and a VxRail-managed VMware vCenter Server, the IP address that is assigned for vSAN witness management must be in the VxRail-managed VMware vCenter Server network subnet range.
- If it is a customer-managed VMware vCenter Server instance, then the vSAN management traffic on the **Management Network** must route to this VMware vCenter Server instance.

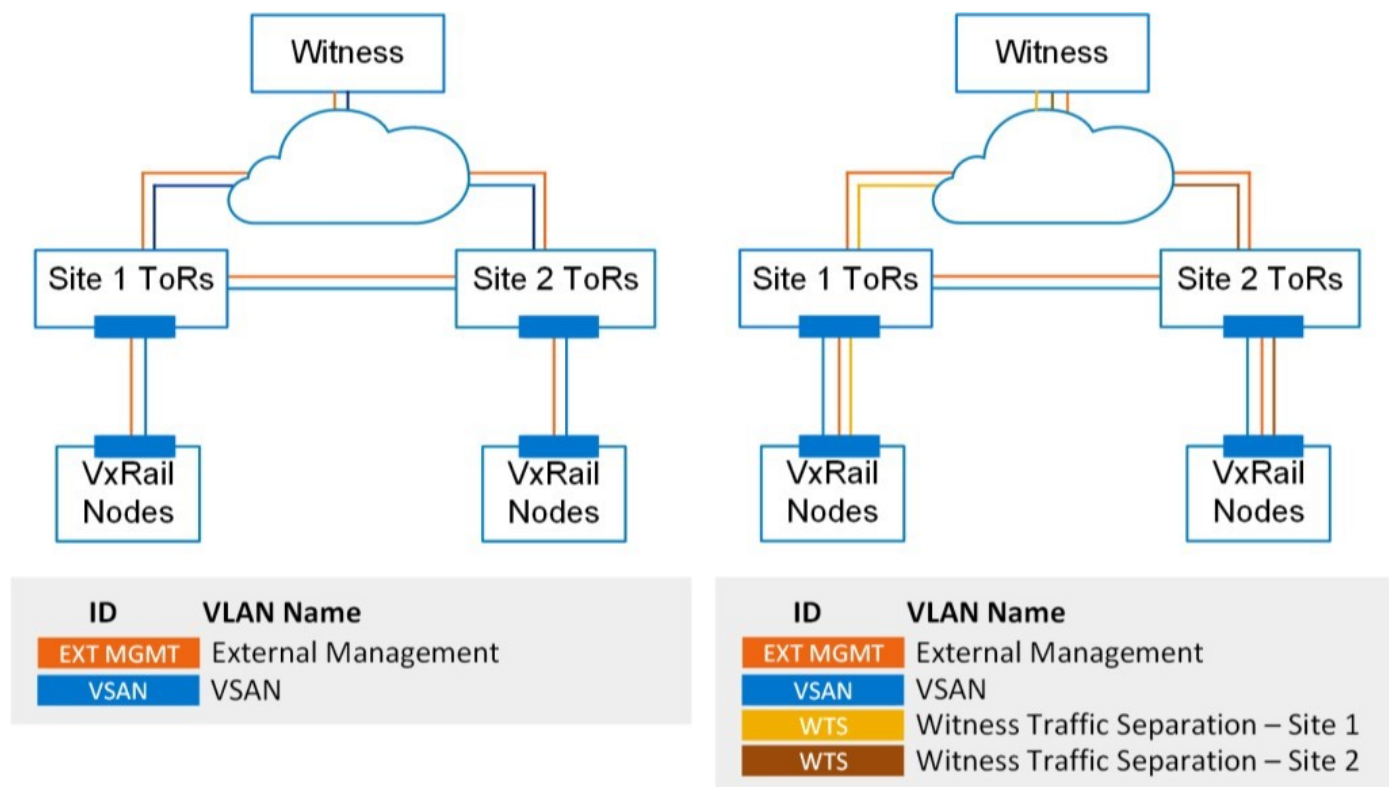


Figure 64. With or without Witness Traffic Separation network

- If a Witness Traffic Separation (WTS) network is being planned for the VxRail cluster, then the vSAN witness traffic on the **Secondary Network** must be able to route to this network.
- If a Witness Traffic Separation (WTS) network is not being considered for the VxRail cluster, then the vSAN witness traffic on the **Secondary Network** must be able to route to the vSAN network planned for the VxRail cluster.

Reserve network settings for vSAN witness networks

The vSAN witness can be VxRail-managed or customer-managed.

- For a VxRail stretched cluster, a customer-managed vSAN witness is the only option.
 - Configuring connectivity between the VxRail cluster networks and the vSAN witness networks occurs after VxRail Manager completes the initial build of the cluster at the primary site.
- For a VxRail 2-node cluster, the vSAN witness can be VxRail-managed or customer-managed.
 - For a VxRail-managed vSAN witness, VxRail Manager configures the vSAN witness with the network settings that are provided during the initial build process.
 - For a customer-managed vSAN witness, VxRail Manager validates network connectivity as part of the validation process. Validation must pass before continuing to the initial build phase.

Record the two IP addresses for the vSAN witness virtual appliance. The Witness Management Network and the Witness vSAN Network cannot share the same subnet.

Table 36. Network configuration table

Network configuration table	Action
Row 67	Enter IP address for Witness Management Network.
Row 73	Enter IP address for Witness vSAN Network.

Decide whether to use the vSAN network or the Witness Traffic Separation (WTS) Network for monitoring.

- For a VxRail 2-node cluster, the WTS network is required.
- For a VxRail stretched cluster, a WTS network is recommended but not required.

Record the VLAN for the WTS network using the network configuration table: Row 74 and enter the VLAN for the WTS network.

For a 2-node cluster deployment, record the IP addresses for each node that is required for vSAN witness traffic.

Table 37. Network configuration table

Network configuration table	Action
Row 75	Enter the IP address for the first of the two nodes in the 2-node cluster.
Row 76	Enter the IP address for the second of the two nodes in the 2-node cluster.

Planning VxRail Satellite Nodes Networking

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.
- VMware vMotion networks for VM 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 that are 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.

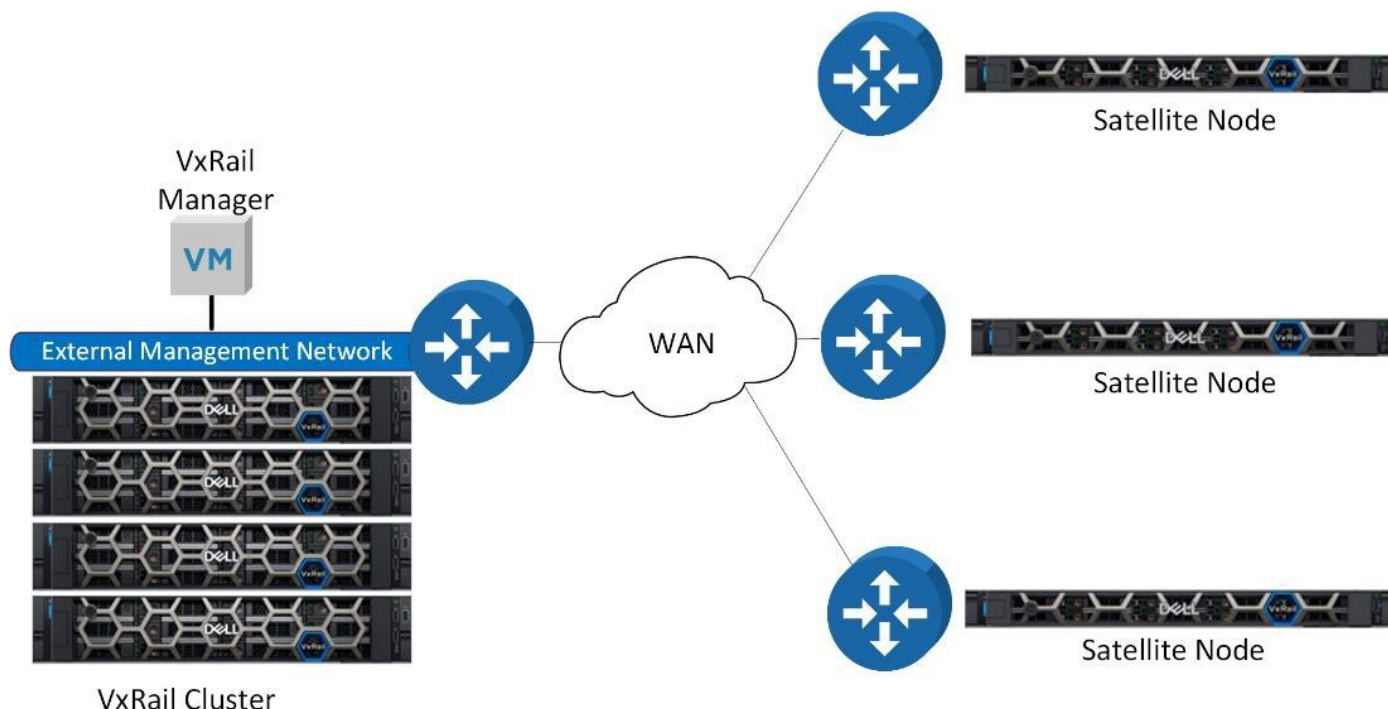


Figure 65. Routing between VxRail Manager and VxRail satellite nodes

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 a connection with a VxRail Manager instance is established. Plan your data center

network to allow the VxRail external management network to route externally. 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.

Table 38. Network Configuration Table

Network configuration table	Action
Row 77 and 78	Enter the management IP address and hostname for the first satellite node.
Row 79 and 80	Enter the management IP address and hostname for the second satellite node.
Row 81 and 82	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 are 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 VMware vSphere password and VMware vCenter Server password documentation.

Configure the Network for VxRail

For the VxRail initialization process to pass validation and build the cluster, you must configure the adjacent ToR 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 Cluster Setup Checklist](#) for guidance. Be sure to follow your vendor documentation for specific switch configuration activities and for best practices for performance and availability.

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

If you do not plan to use the autodiscover method due to multicast restrictions, and will use the manual method instead for selecting nodes for the cluster build operation, you can skip this task.

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 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 VMware Loudmouth service.

VMware 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, only on the ports connected to VxRail nodes.

VxRail creates little traffic through multicasting for autodiscovery 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.

Configure unicast for VxRail vSAN network

You can skip this task if you do not plan to use vSAN as the primary storage resource on the VxRail cluster.

For early versions of VxRail, multicast was required for the vSAN VLAN. One or more network switches that connect to VxRail must allow the pass-through of the multicast traffic on the vSAN VLAN. From VxRail 4.5.x onwards, all vSAN traffic replaces multicast with unicast. This change helps reduce network configuration complexity and simplifies switch configuration with unicast. Unicast is a common protocol that is enabled by default on most enterprise Ethernet switches. If you are required to configure multicast, VxRail multicast traffic for vSAN is 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. If your switch supports, enable both IGMP Snooping, and IGMP Querier and 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 takes care of the following:

- Sends out IGMP group membership queries at a fixed interval.
- Retrieves IGMP membership reports from active members.
- Allows updates to group membership tables.

By default, most switches enable IGMP Snooping but disable IGMP Querier. If so, you must change the settings. 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

Prerequisites

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

- 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. This is not required if you are going to use manual node discovery instead of automatic node discovery.
- vSAN VLAN: In cases where vSAN is the primary storage resource. Ensure that unicast is enabled. This is not required for VxRail dynamic clusters.
- VMware vSphere vMotion VLAN
- VM Networks VLAN: It 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 required only if you are deploying a VxRail stretched cluster or 2-Node cluster.

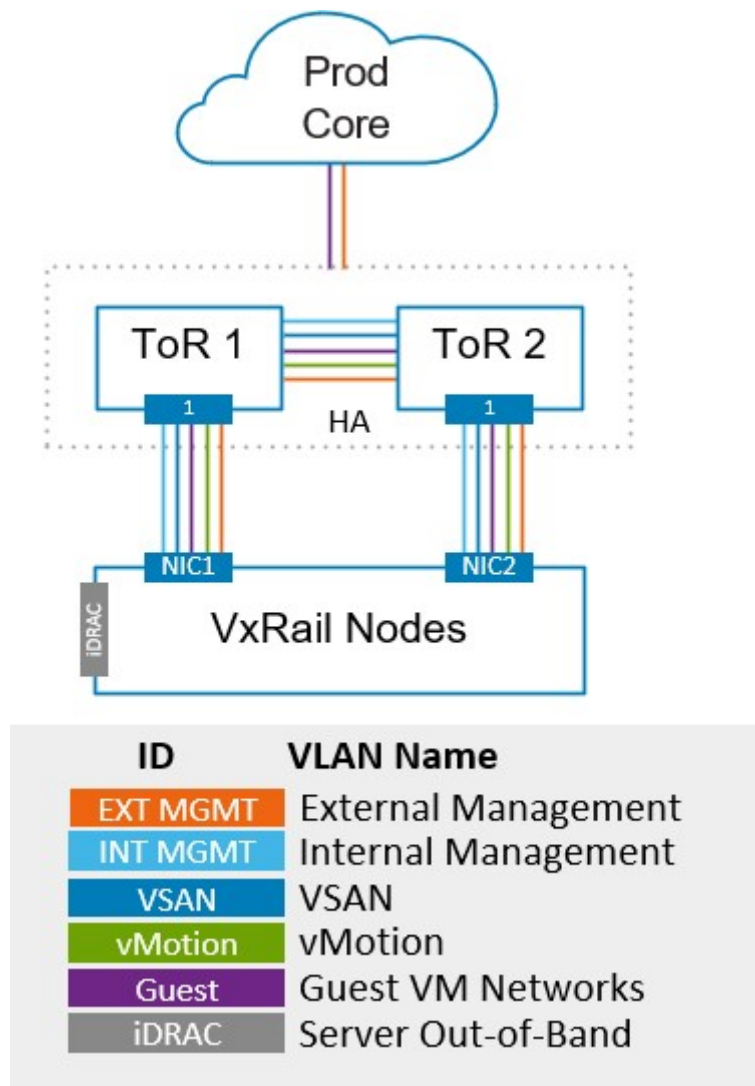


Figure 66. VxRail logical networks: VxRail cluster with vSAN

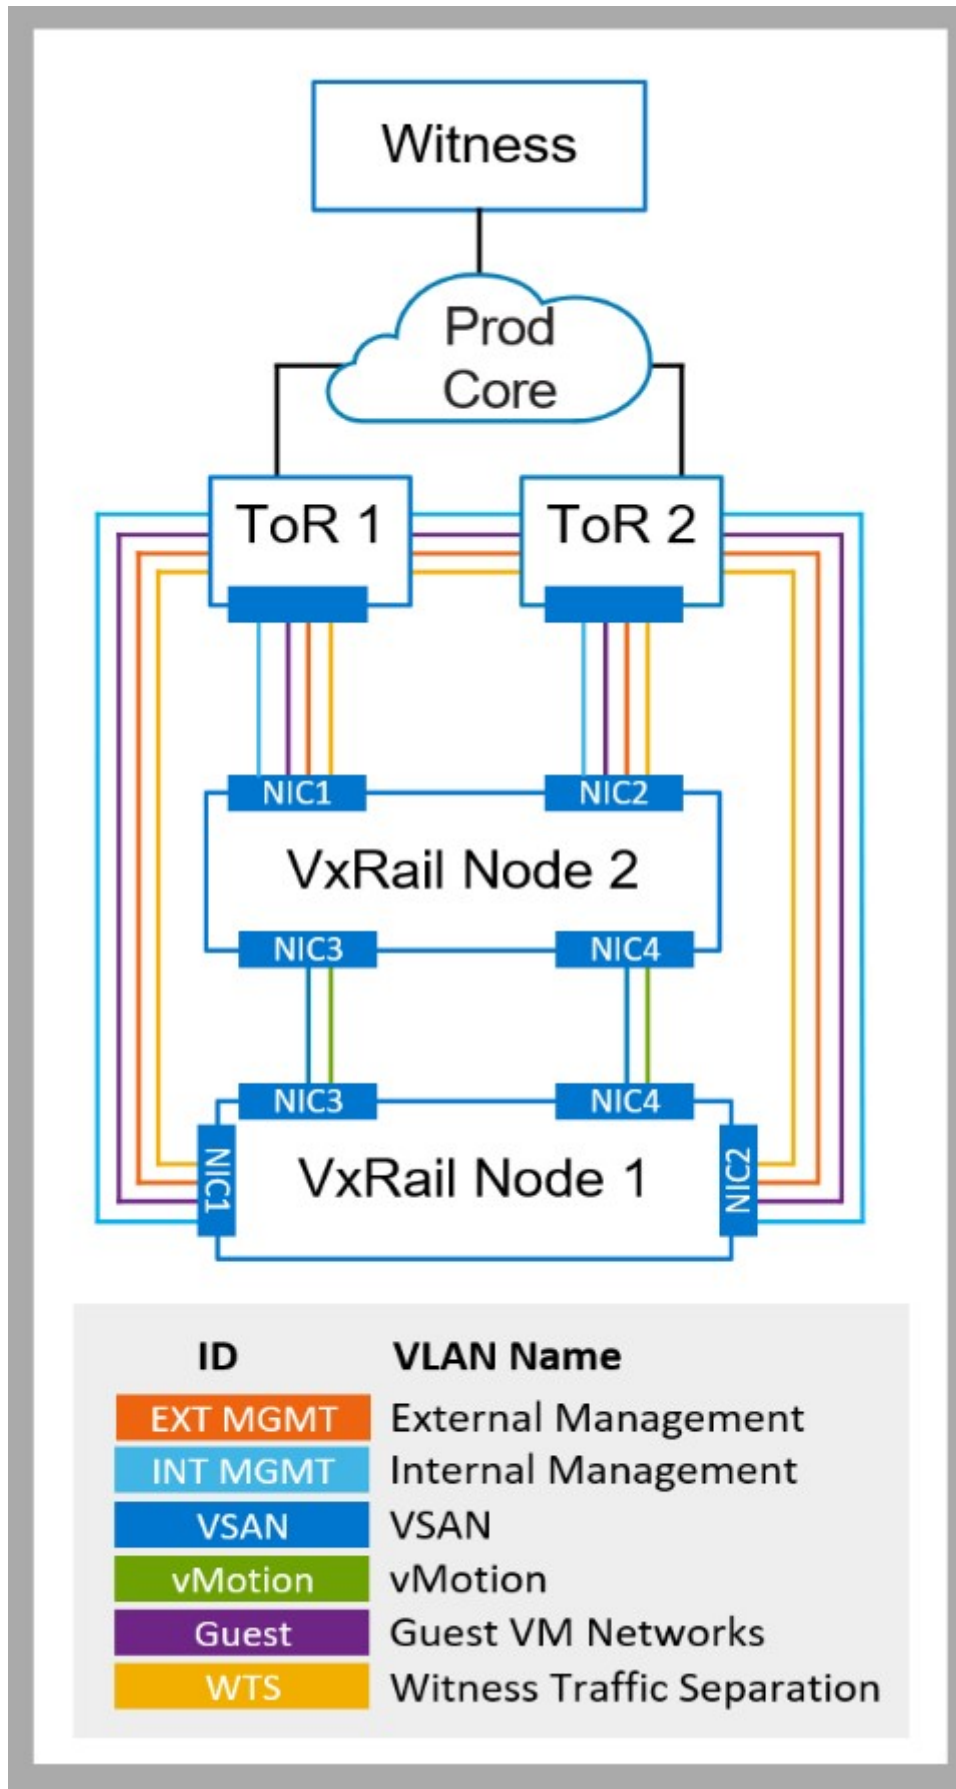


Figure 67. VxRail logical networks: 2-node cluster with Witness

About this task

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

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 4.7.x 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 VMware 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 VxRail 4.5.x 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 69) 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 69) on the uplinks.

Configure the inter-switch links

If more than one ToR 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 port mode

Configure the port mode on your switch based on the plan for the VxRail logical networks, and whether you are going to use VLANs 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.

Configure LAG

LAG is supported for the VxRail initial implementation process only under the following conditions:

- The nodes are running on VxRail 7.0.130 or earlier.
- LAG is being applied only to non-management VxRail networks.
- VxRail is being deployed with customer-supplied virtual distributed switches.

If these conditions are not applicable to your plans, do not enable LAG, including protocols such as LACP and Ether Channel, on any switch ports that are connected to VxRail node ports before initial implementation. Doing so results in VxRail initial implementation to fail. When the initial implementation process completes, you can configure LAG on the operational VxRail cluster, as described in the section [Configure LAG on VxRail networks](#).

If your plans meet these conditions for supporting LAG during the VxRail initial implementation process, then perform these action items before starting:

- For data center networks with more than one switch planned to support the VxRail cluster, configure virtual link trunking between the switches.
- Configure a port channel on each switch for each VxRail node.

- Configure the VLAN or VLANs targeted for LAG in each port channel.
- If your plans include implementation with a VxRail-supplied virtual distributed switch, then configure the port channel so that the Ethernet ports are active before a LAG partnership is formed. This feature is commonly known as **lacp individual** on Dell-branded switches, but may be described differently with switches from other vendors.
- Depending on your switch operating system, other network characteristics, such as MTU setting and spanning tree protocol settings, can be configured in a port channel. These settings transfer to the switch port when LAG is active. Configure these additional network settings in the port channel according to guidance provided by your switch vendor.

Limit spanning tree protocol

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 VMware 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 is **receive on** and **transmit off**.

Set up the network switch ports for VxRail connectivity

Once the switch base settings are complete, the next step is to set up the switch ports.

About this task

Perform the following steps for each switch port that has to be connected to a VxRail node:

Steps

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 VMware 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 69) on the switch ports, if required.

If link aggregation is configured on the switchport, some settings are configured on the port channel, and some are configured on the interface:

- Configure the appropriate port channel or equivalent, depending on the switch operating system on the switch port.
- Configure any additional network port settings that do not transfer from the port channel on the switch port.

Set 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](#) as a 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.

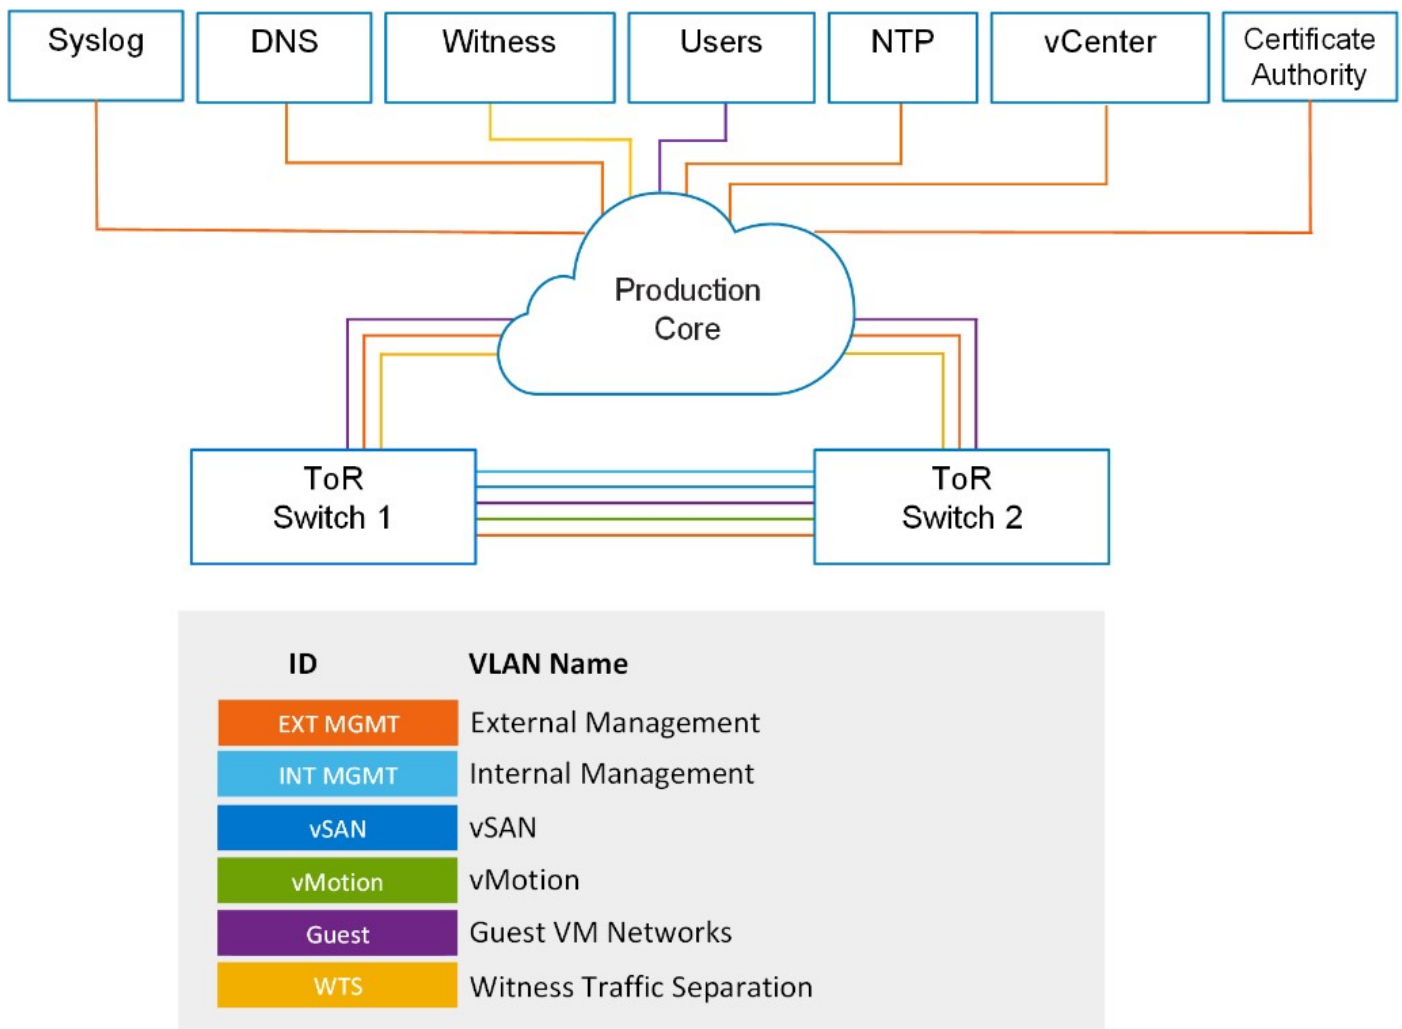


Figure 68. Logical networks connecting to upstream elements

If your Layer 2/Layer3 boundary is at the lowest network tier (ToR switch), perform the following tasks:

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

If your Layer 2/Layer3 boundary is upstream from at the lowest network tier (ToR switch), perform the following tasks:

- Connect ports on the adjacent upstream switch to the uplinks on the ToR switches.
- Configure logical pairings of the ports on the adjacent upstream switch and the ToR switch.
- Configure the logical port pairings, commonly known as **port channels** or **Ether Channels**, to allow upstream passage of external VxRail networks.

Configure your network to support RoCE

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 that are targeted for deployment are configured with Ethernet adapters that support RDMA over Converged Ethernet (RoCE), the switches and supporting network must 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 steps to ensure a lossless network are as follows:

- vSAN with RDMA supports NIC failover, but does not support LAG 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, the lossless network settings must be preserved when routed across network devices using a feature such as the Differentiated Service 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.

Prerequisites

Before forming the VxRail cluster, the VxRail initialization process performs several verification steps, including:

- Verifying switch and data center environment support.
- 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 cause the validation to fail, and the VxRail cluster is not formed. When validation fails, the data center settings and switch configurations must undergo troubleshooting to resolve the problems reported.

About this task

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

Steps

1. External management traffic is untagged on the native VLAN by default. If a tagged VLAN is used instead, customize the switches with the new VLAN.
2. Internal device discovery network traffic uses the default VLAN of 3939. If this has changed, customize all ESXi hosts with the new VLAN, or device discovery will not work.
3. Confirm that the switch ports that attach to the 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 that an inter-switch link is configured between them to support passage of the VxRail network VLANs.

Confirm your firewall settings

Prerequisites

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.

Steps

1. Verify that VxRail can communicate with your DNS server.
2. Verify that VxRail can communicate with your NTP server, if planned for clock synchronization.
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 connect gateway in your data center instead of the secure connection deployed in the VxRail cluster, verify the open communications between VxRail management and the secure connect gateway.
9. If you are planning to use VMware subscription licenses with VxRail, confirm connectivity to the VMware Cloud from the VMware vCenter Cloud Gateway.
See [Appendix D: VxRail Open Ports Requirements](#) for information about VxRail port requirements.

Confirm your data center environment

Steps

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 that the VMware 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 VMware vSphere 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 the 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, confirm that there is a routable path between the witness and this network for the management traffic and vSAN witness traffic.
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.

Prepare to build the VxRail cluster

Dell professional services is responsible for performing the following steps to build the cluster during the delivery engagement.

Complete prerequisites for dynamic clusters

Perform these steps only if you intend to deploy a dynamic cluster.

A dynamic cluster depends on external storage resources to support cluster implementation and operations. A dynamic cluster does not support a local VMware vSAN datastore. A dynamic cluster cannot not fully complete implementation unless an external storage resource is available to support a virtual machine workload.

If you plan to deploy a dynamic cluster over an FC network, complete all the steps on your FC switches and storage array to present LUNs to the VxRail nodes. To deploy a dynamic cluster where the primary storage resource is a remote vSAN datastore, confirm that your data center network is 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 using storage over an IP network, confirm that your data center network is configured so that the VxRail nodes can connect to the supporting storage array.

Configure nodes for tagged VxRail management VLAN

The VLAN preconfigured 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 use a tagged VLAN for the VxRail external management network, which restricts network passage, perform the following steps:

- Configure the VLAN on the trunked port as a tagged VLAN.
- Change the default VLAN for the VxRail external management network on the VMware VDS for each of the VxRail nodes to the tagged VLAN.

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

Configure a jump host or laptop for VxRail initialization

Use a desktop with an operating system such as Windows to reach the VxRail external management network for initialization. The desktop must include a supported web browser that can access the VxRail management interface.

You can either use a workstation or laptop to plug into an open Ethernet port on one of the ToR switches or a jump host (jump server description) in your data center that can reach the VxRail external management network.

If you 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. When the VxRail initialization process is complete, the switch port or jump host is no longer required to manage VxRail.

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 switch, and the workstation/laptop must be logically configured to reach the necessary networks.

A supported web browser is required to access the VxRail management interface. The latest versions of Firefox, Chrome, and Internet Explorer10+ 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, pre-configured 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 changes during VxRail initialization to your desired permanent address, and assigned to VxRail Manager during cluster formation.

Your workstation or 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 reminds you that you might need to reconfigure your workstation or 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, the Dell 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.

Table 39. Set the IP addresses on the Ethernet port

Example configuration	VxRail IP address/netmask	Jump host/Laptop		
		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

Log in to a jump host, or connect the laptop or workstation to a switch port, to perform the final step of VxRail initialization.

About this task

Dell service representatives perform the steps in this section. The steps are included here to help you understand the complete process.

Steps

1. Before coming on-site, the Dell service representative contacts you to capture and record the information that is described in [Appendix A: VxRail Network Configuration Table](#) and walk through [Appendix C: VxRail Cluster Setup Checklist](#).
2. If your planned VxRail deployment requires a witness at a remote data center location, the witness virtual appliance is deployed.
3. If your deployment includes 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.
4. If your planned deployment is a dynamic cluster, complete the necessary preparations for the selected external storage resource.
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.
6. Attach Ethernet cables between the ports on the VxRail nodes and switch ports that are configured to support VxRail network traffic.
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.
8. Connect a workstation or laptop or jump host that is configured for VxRail initialization to access the VxRail external management network on your selected VLAN. Plug into the ToR switch or logically reach the VxRail external management network from elsewhere on your network.
9. Open a browser to the VxRail IP address to begin the VxRail initialization. This is either the default IP address that is assigned to VxRail Manager at the factory or the permanent IP address set by Dell services.
10. The Dell service representative populates the input screens on the menu with the data that is collected from the customer during the planning and design process.
11. VxRail performs the verification process using the information input into the menus.
12. After validation is successful, the initialization process begins to build a new VxRail cluster. The new permanent IP address for VxRail Manager is displayed.

13. If the permanent IP address was set on VxRail Manager before initialization, the browser works through the process.
 - If you configured the workstation or laptop to enable connectivity to both the temporary VxRail IP address and the new permanent IP address, the browser session makes the switch automatically. If not, you must manually change the IP settings on your workstation or 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 receive a message to fix your network and try again. If you are unable to connect to the new IP address after 20 minutes, VxRail reverts to its unconfigured state, and you must reenter your configuration at the temporary VxRail IP address.
 - After the build process starts, if you close your browser, you must browse to the new, permanent VxRail IP address.
14. Progress is shown as the VxRail cluster is built. The process takes about 25-40 minutes.
15. When VxRail initialization is complete and a new VxRail cluster is built, click **Manage VxRail** to continue to VxRail management. You should also bookmark this IP address in your browser for future use.
16. 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 leads you to the vCenter instance.

VxRail network considerations after implementation

Other actions are provided to optimize VxRail cluster operations. The cluster version, the configuration settings, and supported features on your data center network determine the options available to modify the VxRail networking. If the VxRail cluster was originally configured using only the integrated NDC/OCP ports, you can configure the ports on the optional PCIe adapter cards to support VxRail network traffic. 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 starting and ending network reconfiguration that VxRail supports:

Table 40. NDC/OCP port configurations

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.
- Reserve the first port configured for VxRail networking, which is known as **vmnic0** or **vmnic1**, 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 running on two NDC/OCP ports can be reconfigured to run on one NDC/OCP port and one PCIe port.

Follow the official instructions or procedures from VMware and Dell for these operations.

The supported operations include:

- Create a VMware Standard Switch and connect to unused ports.
- Connect unused ports to new port groups on the default VMware VDS.
- Create a VMware VDS and add VxRail nodes. Connect their unused network ports to the VMware VDS.
- Create VMkernel adapters, and enable services of IP storage and VMware vSphere replication.
- Create VM Networks and assign them to new port groups.

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

- You cannot migrate or move VxRail traffic to the optional ports. VxRail traffic includes the management, vSAN, VMware vCenter Server, and VMware vSphere vMotion Networks.
- You cannot migrate VxRail traffic to other port groups.
- You cannot migrate VxRail traffic to another VMware VDS.

 **CAUTION: Unsupported operations impact the stability and operations of the VxRail cluster and cause a failure.**

Configure LAG on VxRail networks

Starting with VxRail version 7.0.130, configure NIC teaming on VxRail non-management networks with both a customer-supplied and VxRail-supplied VMware VDS after initial cluster implementation. NIC teaming enables the formation of a LAG, which is a logical port that represents a pair of physical ports on a VxRail node. If the ports on the ToR switches connected to these logical

ports are also configured into a LAG, peering between the two LAGs enables an active/active port configuration and support load-balancing for network optimization.




Figure 69. Enabling link aggregation for load balancing

The following rules are applicable to enable LAG on VxRail networks:

- No more than two Ethernet ports can be configured in a LAG. The Ethernet ports can be the following:
 - All NDC/OCP ports
 - All PCIe ports, or
 - A mixture of NDC/OCP and PCIe ports
- Configure all VxRail node ports to run at the same speed for LAG.
- In VxRail versions earlier than 7.0.450, LAG can only be configured on the non-management VxRail networks.
- The adjacent ToR switches support LACP.

Dynamic LAG requires:

- Dynamic port channels are configured on the adjacent ToR switches.
- LACP is enabled on the adjacent ToR switches.
- An LACP policy is configured on the VMware VDS.
- The load-balancing setting on the LACP policy is compatible with the supported load-balancing hash algorithms on the adjacent ToR switches.

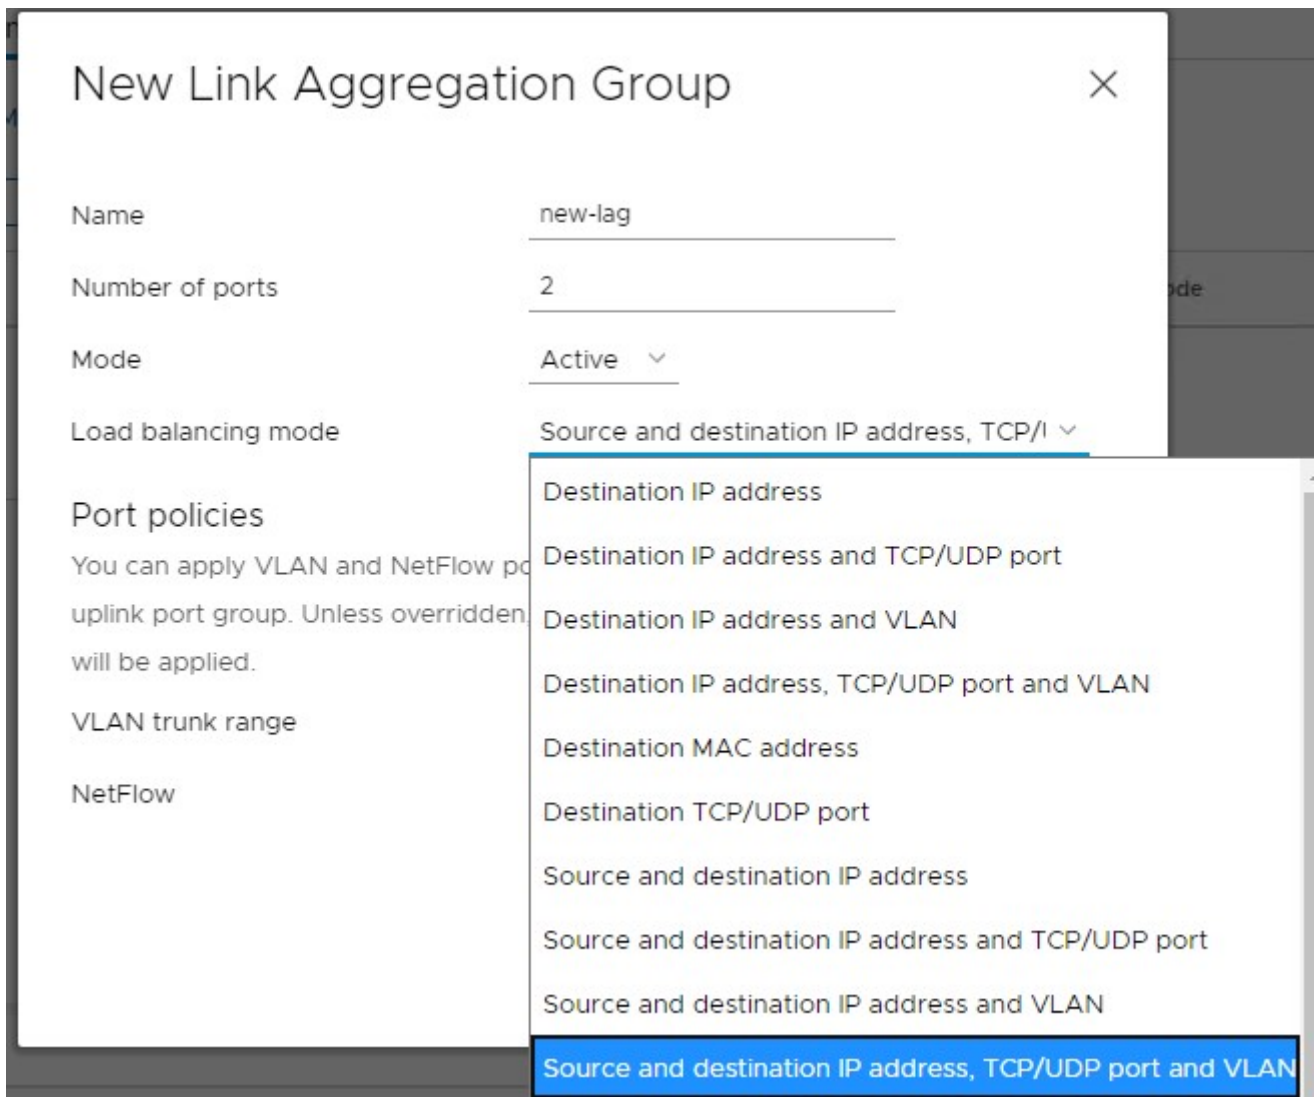


Figure 70. LACP policy configuration on virtual distributed switch

You can configure any unused network ports on the VxRail nodes that were not configured for VxRail network traffic use cases to support LAG. These can include any unused ports on the NDC/OCP or on the optional PCIe adapter cards. Updates to support these new networks can be configured on the VMware VDS deployed during VxRail initial build, or you can configure a new VMware VDS. Since the initial VMware VDS is under the management and control of VxRail, configure a separate VMware VDS on the VMware vCenter Server instance to support these networking use cases.

Appendix A: VxRail Network Configuration Table

The Dell 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:

Table 41. Data collection workbook content

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 ToR switches. The default VLAN ID is 3939.	
3		vMotion VLAN	VMware vSphere vMotion VLAN	
4		vSAN VLAN	VMware vSAN VLAN	
5		Guest Networks VLAN	Network name	
6			VLAN	
7		VMware vCenter Server Management VLAN	VMware vCenter Server network on the 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	VMware vCenter Server Management Network	Subnet Mask	VMware vCenter Server network on the same subnet as external management network by default. Can optionally be assigned a new subnet mask.	
11		Gateway	VMware vCenter Server network on the 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 autoassign method	Prefix	
19			Separator	
20			Iterator	

Table 41. Data collection workbook content (continued)

Row	Topic	Category	Description
21		Customer-supplied method	Offset
22			Suffix
23			ESXi hostname 1
24			ESXi hostname 2
25			ESXi hostname 3
26			ESXi hostname 4
27	ESXi IP Addresses	VxRail autoassign method	Starting IP address
28			Ending IP address
29		Customer-supplied method	VMware ESXi IP Address 1
30			VMware ESXi IP Address 2
31			VMware ESXi IP Address 3
32			VMware ESXi IP Address 4
33	VMware vCenter Server	VxRail VMware vCenter Server	VMware vCenter Server Hostname
34			VMware vCenter Server IP Address
35			VMware vCenter Server Hostname (FQDN)
36			VMware vCenter Server SSO Domain
37			Admin username/password or the VxRail nonadmin username and password
38			New VxRail management username and password
39			VMware vCenter Data Center Name
40			VMware vCenter Cluster Name
41	VMware VDS	Customer-supplied Switch Names	Name of first VMware VDS
42			Name of second VMware VDS
43		Customer-supplied VMware VDS port groups	Name of VMware VDS port group supporting VxRail external management network
44			Name of VMware VDS port group supporting VxRail VMware vCenter Server network
45			Name of VMware VDS port group supporting VxRail internal management network
46			Name of VMware VDS port group supporting VxRail vMotion network
47			Name of VMware VDS port group supporting VxRail vSAN network
48	vMotion	VxRail autoassign method	Starting address for IP pool
49			Ending address for IP pool
50		Customer-supplied method	VMware vSphere vMotion IP Address 1
51			VMware vSphere vMotion IP Address 2
			VMware vSphere vMotion IP Address 3
52			VMware vSphere vMotion IP Address 4
55		Subnet Mask	
	Gateway	Default or vMotion stack to enable routing	

Table 41. Data collection workbook content (continued)

Row	Topic	Category	Description
56	vSAN	VxRail autoassign method	Starting address for IP pool
57			Ending address for IP pool
58		Customer-supplied method	VMware vSAN IP Address 1
59			VMware vSAN IP Address 2
60			VMware vSAN IP Address 3
61			VMware vSAN IP Address 4
62		Subnet Mask	
63		Gateway	Default or Custom for routable network
64	Logging	Log Insight	VMware vRealize Log Insight hostname
65			VMware vRealize Log Insight IP address
66		Syslog Server	Syslog server IP address
67	Witness site	Management IP Address	Witness management network IP address
68		vSAN IP Address	Witness vSAN network IP address
69	Witness Traffic Separation	WTS VLAN	Optional to enable Witness traffic separation on a stretched cluster or 2-node cluster
70	2-Node Cluster	Node 1 WTS IP address	Must be routable to Witness.
71		Node 2 WTS IP address	Must be routable to Witness.
72	Satellite Nodes	Node 1	VMware ESXi hostname
73			Management IP address
74		Node 2	VMware ESXi hostname
75			Management IP address
76		Node 3	VMware ESXi hostname
77			Management IP address

Appendix B: VxRail Passwords

Record VxRail passwords.

Table 42. VxRail Manager and VMware passwords

Item	Account	Password
VxRail Manager	root	
	mystic	
	service	
VxRail VMware vCenter Server	administrator@<SSO Domain>	
	root	
	Management	
VMware vRealize Log Insight	root	
	admin	

Table 43. VMware ESXi hosts

Item	Account	Password
VMware ESXi Host 1	root	
	Management	
VMware ESXi Host 2	root	
	Management	
VMware ESXi Host 3	root	
	Management	
VMware ESXi Host 4	root	
	Management	

Table 44. Satellite node

Item	Account	Password
Satellite Node 1	root	
Satellite Node 2	root	
Satellite Node 3	root	

Appendix C: VxRail cluster setup checklist

Follow the checklist to set up the cluster.

- VxRail cluster: Plan for additional nodes beyond the initial three (or four)-node cluster. You can have up to 64 nodes in a VxRail cluster.
- VxRail ports: Determine how many ports to configure per VxRail node, what port type, and what network speed.
- Network switches: Ensure that your switches support VxRail requirements and provides the connectivity option that you chose for your VxRail nodes. Verify cable requirements.
- Data center: Verify that the required external applications for VxRail are accessible over the network and correctly configured.
- Topology: If you are deploying VxRail over more than one rack, be sure that network connectivity is set up between the racks. Determine the L2/L3 boundary in the planned network topology.
- Workstation/laptop: Any operating system with a browser to access the VxRail Manager. The latest versions of Firefox, Chrome, and Internet Explorer 10+ are all supported.
- Out-of-band Management (optional): One available port that supports 1 Gb for each VxRail node.

Component	Description
Reserve VLANs	<p>One external management VLA</p> <p>One internal management VLAN with multicast for autodiscovery and device management. The default is 3939.</p> <p>One VLAN with unicast enabled for VMware vSAN traffic</p> <p>One VLAN for VMware vSphere vMotion</p> <p>One or more VLANs for your VM Guest Networks</p> <p>One VLAN for VMware vCenter Server Network (if applicable)</p> <p>If you are enabling witness traffic separation, reserve one VLAN for the VxRail witness traffic separation network.</p>
System	<p>Select the timezone.</p> <p>Select the top-level domain.</p> <p>Hostname or IP address of the NTP servers on your network (recommended)</p> <p>IP address of the DNS servers on your network (if external DNS)</p> <p>Forward and reverse DNS records for VxRail management components (if external DNS).</p>
Management	<p>Decide on your VxRail host naming scheme. The naming scheme applies to all VxRail management components.</p> <p>Reserve IP addresses for VMware ESXi hosts.</p> <p>Reserve one IP address for VxRail Manager</p> <p>Determine default gateway and subnet mask.</p> <p>Select passwords for VxRail management components.</p>
vCenter	<p>Determine whether to use a VMware vCenter Server that is customer-supplied or a VMware vCenter provided by VxRail.</p> <p>VxRail vCenter Server: Reserve IP addresses for VMware vCenter Server (if supplied by VxRail).</p> <p>Customer-managed VMware vCenter Server: Determine hostname and IP address for vCenter, administration user, and name of VMware vSphere data center. Create a VxRail management user in VMware vCenter. Select a unique VxRail cluster name. (Optional) Create a VxRail nonadmin user.</p>

Virtual Distributed Switch	<p>Preconfigure a customer-managed VMware VDS or have VxRail deploy a VMware VDS in your VMware vCenter instance.</p> <p>Customer-managed VMware VDS: Configure target port groups for required VxRail networks.</p>
vMotion	<p>Decide whether you want to use the default TCP-IP stack for VMware vMotion, or a separate IP addressing scheme for the dedicated VMware vMotion TCP-IP stack.</p> <p>Reserve IP addresses and a subnet mask for VMware vSphere vMotion.</p> <p>Select the gateway for either the default TCP-IP stack, or the dedicated VMware vMotion TCP-IP stack.</p>
vSAN	Reserve IP addresses and a subnet mask for vSAN, if using vSAN for primary storage.
Logging	<p>To use VMware vRealize Log Insight: Reserve one IP address.</p> <p>To use an existing syslog server: Get the hostname or IP address of your third-party syslog server.</p>
Witness Site	If a witness is required, reserve one IP address for the management network and one IP address for the vSAN network.
Workstation	<p>Configure your workstation/laptop to reach the VxRail initial IP address.</p> <p>Ensure you know how to configure the laptop to reach the VxRail Manager IP address after configuration.</p>
Set up Switches	<p>Configure your selected external management VLAN (default is untagged/native).</p> <p>Configure your internal management VLAN.</p> <p>Confirm unicast is enabled for device discovery.</p> <p>Configure your selected VLANs for VMware vSAN, VMware vSphere vMotion, VMware vCenter Server Network and VM Guest Networks.</p> <p>If applicable, configure your witness traffic separation VLAN.</p> <p>In dual-switch environments, configure the interswitch links to carry traffic between switches.</p> <p>Configure uplinks or point-to-points links to carry networks requiring external connectivity upstream.</p> <p>Configure one port as an access port for laptop/workstation to connect to VxRail Manager for initial configuration.</p> <p>Confirm configuration and network access.</p>
Workstation/Laptop	<p>Configure your workstation/laptop to reach the VxRail Manager initial IP address.</p> <p>Configure the laptop to reach the VxRail Manager IP address after permanent IP address assignment.</p>

Appendix D: VxRail Open Ports Requirements

Firewall settings specific for the deployment of a VxRail cluster are provided. Use the links that are provided after the tables for firewall rules that are driven by product feature and use case. The VxRail cluster must 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.

Table 45. Data center 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	DNS Servers Host ESXi Management Interface, VMware vRealize Log Insight	Syslog Server	TCP	514
LDAP	VMware vCenter Servers	LDAP Server	TCP	389, 636
SMTP	Secure connect gateway VMs, VMware vRealize Log Insight.	SMTP Servers	TCP	25

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

Table 46. 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, VMware vRealize Log Insight	TCP	80, 443
Dell server management	Administrators	Dell iDRAC	TCP	623, 5900, 5901
SSH and SCP	Administrators	Host ESXi Management, vCenter Server, Dell iDRAC port, VxRail Manager Console	TCP	22

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

Table 47. VMware vCenter Server and VMware vSphere Client

Description	Source Devices	Destination Devices	Protocol	Ports
VMware vSphere Clients to VMware vCenter Server	VMware vSphere Clients	VMware vCenter Server	TCP	5480, 8443, 9443, 10080, 10443
Managed Hosts to VMware vCenter Server	Host ESXi Management	VMware vCenter Server	TCP	443, 902, 5988, 5989, 6500, 8000, 8001
Managed Hosts to VMware vCenter Server Heartbeat	Host ESXi Management	VMware vCenter Server	TCP	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. VxRail manages the VxRail Customer Firewall Rules interactive workbook. Access to the

workbook requires Dell customer credentials. If you do not have Dell login credentials, contact your account team to download the tool for you.

Table 48. VxRail Customer Firewall Rules interactive workbook

Description	Reference
VMware Ports and Protocols	VMware Ports and Protocols
Network port diagram for VMware 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 Connect Gateway Documentation	Dell Secure Connect Gateway Documentation
VMware vCenter Cloud Gateway Requirements	VMware Cloud Gateway for vSphere+ Requirements

Appendix E: VMware VDS port group default settings

If you configure an external VMware VDS in your customer-managed VMware vCenter Server for the VxRail cluster, the VxRail initial build process configures one or two VMware VDS on the selected VMware vCenter Server instance using best practices.

Default standard settings

Default settings are applied for each VxRail network port group.

VxRail initialization applies the following standard settings.

Table 49. 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 VMware VDS.

Overriding the default load-balancing policy is supported on VxRail.

Table 50. Teaming and failover policy

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
Network failure detection	Yes
Failback	Yes

Default network I-O control (NIOC)

VxRail enables network I-O control on the VMware VDS and configures 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 two Ethernet ports per node that is reserved for the VxRail cluster, or if four Ethernet ports were reserved for the VxRail cluster:

Table 51. Traffic type settings

Traffic type	NIOC shares	
	Four ports	Two ports
Management traffic	40	20
VMware vSphere vMotion traffic	50	50
vSAN traffic	100	100
VM 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 port groups.

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, VMware vCenter Server Network, VMware vSphere vMotion, and vSAN.

Table 52. 4 x 10 GbE traffic configuration

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

Table 53. 2 x 10 GbE or 2 x 25 GbE traffic configuration

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

Table 54. 2 x 10 GbE or 2 x 25 GbE traffic configuration

Traffic type	Uplink1 VMNIC0	Uplink2 VMNIC1	Uplink3 VMNIC2	Uplink4 VMNIC3
Management	Active	Unused	Standby	Unused

Table 54. 2 x 10 GbE or 2 x 25 GbE traffic configuration (continued)

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

Appendix F: Physical Network Switch Examples

The diagrams in this appendix show various options for physical wiring between VxRail nodes and the adjacent, ToR switches. They are provided as illustrative examples to help with the planning and design process for physical network connectivity. All VxRail nodes are manufactured with Ethernet ports that are built into the chassis. VxRail nodes support NDC or adapters that are supported by the OCP specifications for the integrated Ethernet ports, depending on VxRail model. Optional PCIe adapter cards can be installed in the VxRail nodes to provide additional Ethernet ports for redundancy purposes and increased bandwidth. These examples display options starting with connectivity from a single NDC/OCP adapter to network topologies that address requirements for failure protection and optimal bandwidth distribution.

If additional Ethernet connectivity is required to support other use cases outside of VxRail networking, such as additional guest networks or external storage, additional slots on the VxRail nodes should 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. The examples include topologies for predefined network profiles and custom network profiles. With a predefined network profile, VxRail Manager automatically selects the Ethernet ports for assignment to VxRail networks. With a custom network profile, the network topology is manually configured using VxRail Manager. Custom network profiles are supported with VxRail 7.0.130 or later.

Network profiles using six ports and eight ports for VxRail networking are supported with VxRail 7.0.400 or later.

Pre-defined network profile: 2x10gb or 2x25gb from a single NDC/OCP

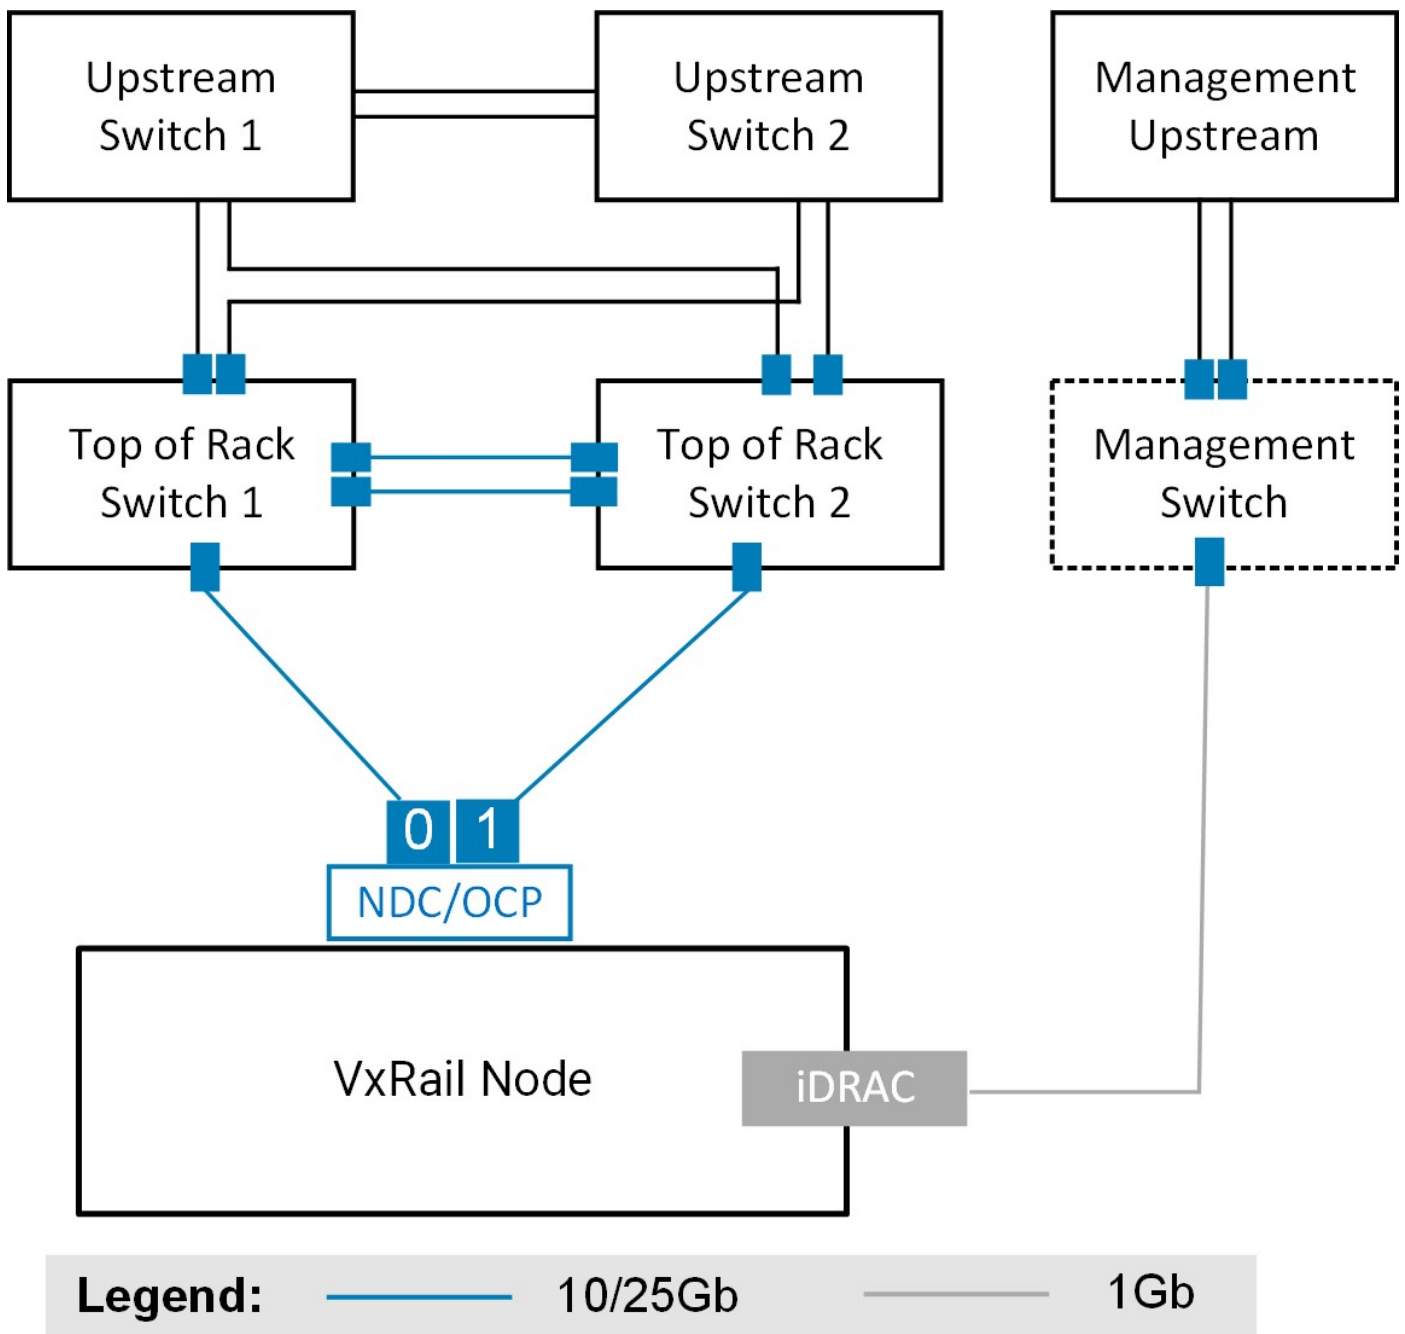


Figure 71. VxRail nodes with two 10 Gb or 25 Gb NDC/OCP ports connected to two ToR switches, and one optional connection to management switch for iDRAC

With this predefined profile, VxRail selects the two ports on the NDC/OCP to support VxRail networking. If the NDC/OCP adapter on the VxRail nodes is shipped with four Ethernet ports, the two left most 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 loss of the NDC/OCP adapter as a single point of failure.

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

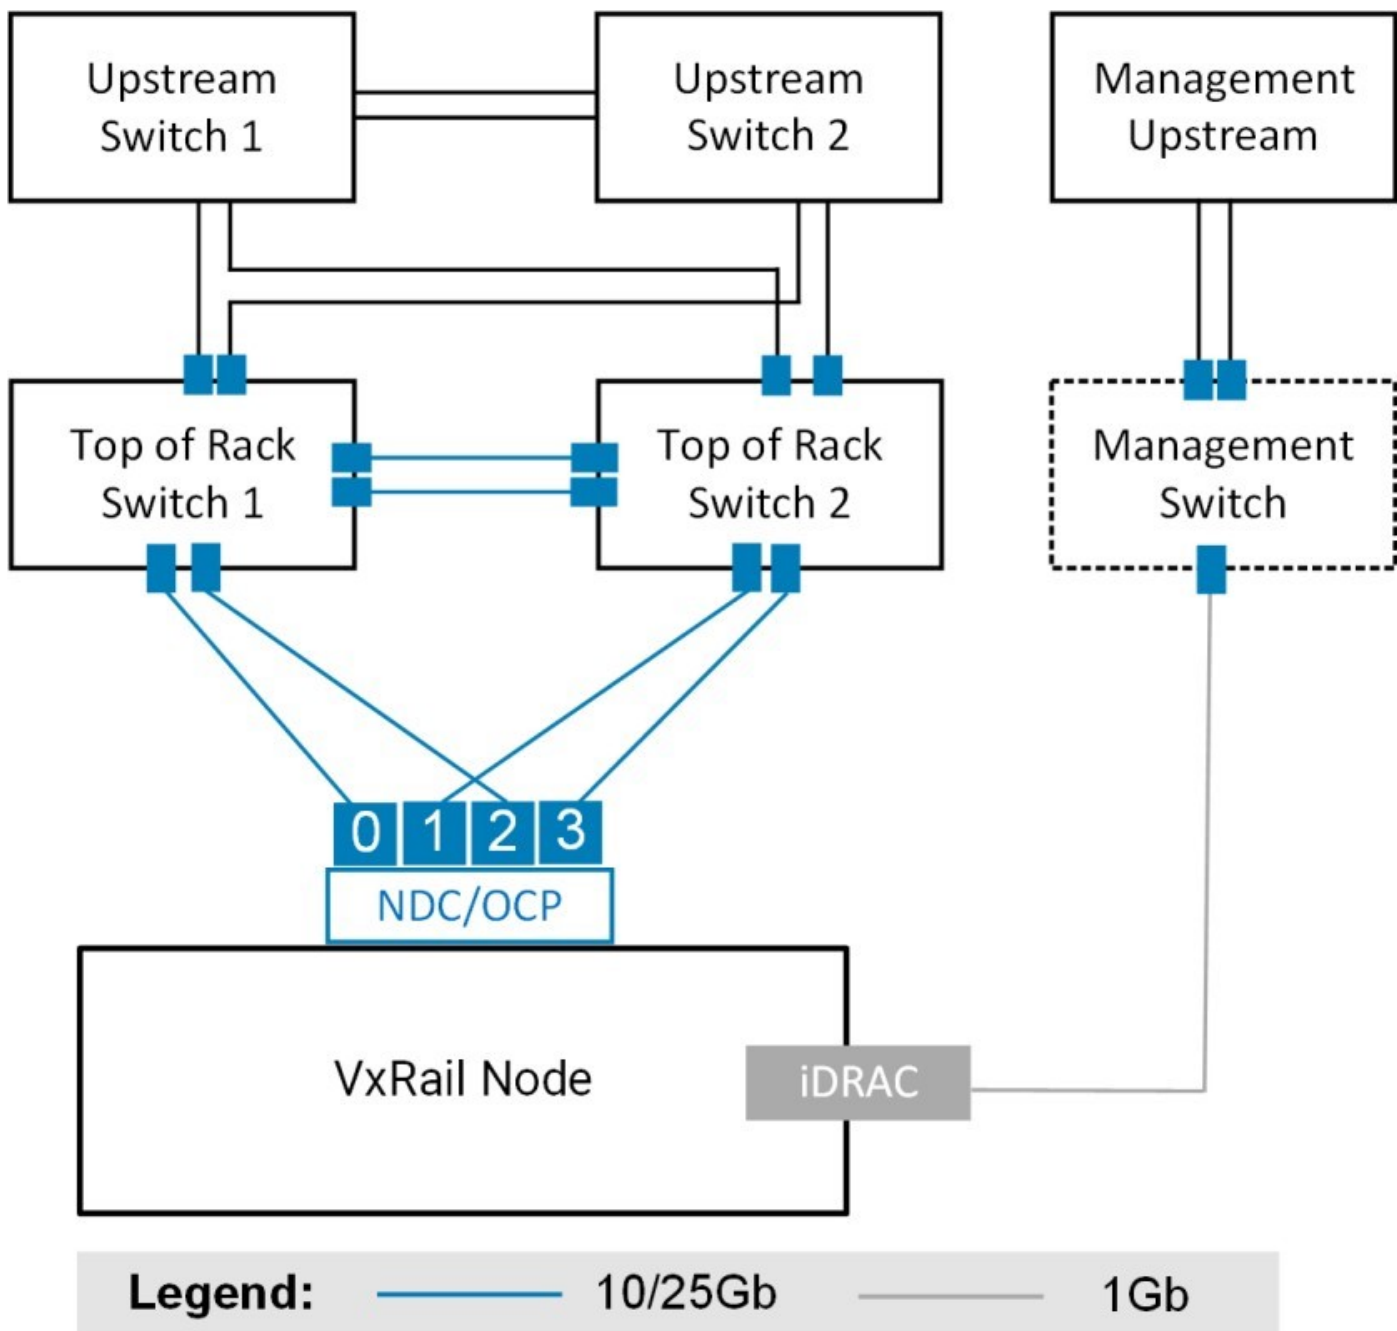


Figure 72. VxRail nodes with four 10 Gb NDC/OCP ports connected to two ToR switches, and one optional connection to management switch for iDRAC

In this predefined network profile, VxRail selects all four ports on the NDC/OCP to support VxRail networking instead of two. The same number of cable connections should be made to each switch. This topology provides additional bandwidth over the two-port option, but provides no protection resulting from a failure with the network adapter card.

Predefined network profile: 2 x 10/25 GB NDC/OCP and 2 x 10/25 GB PCIe

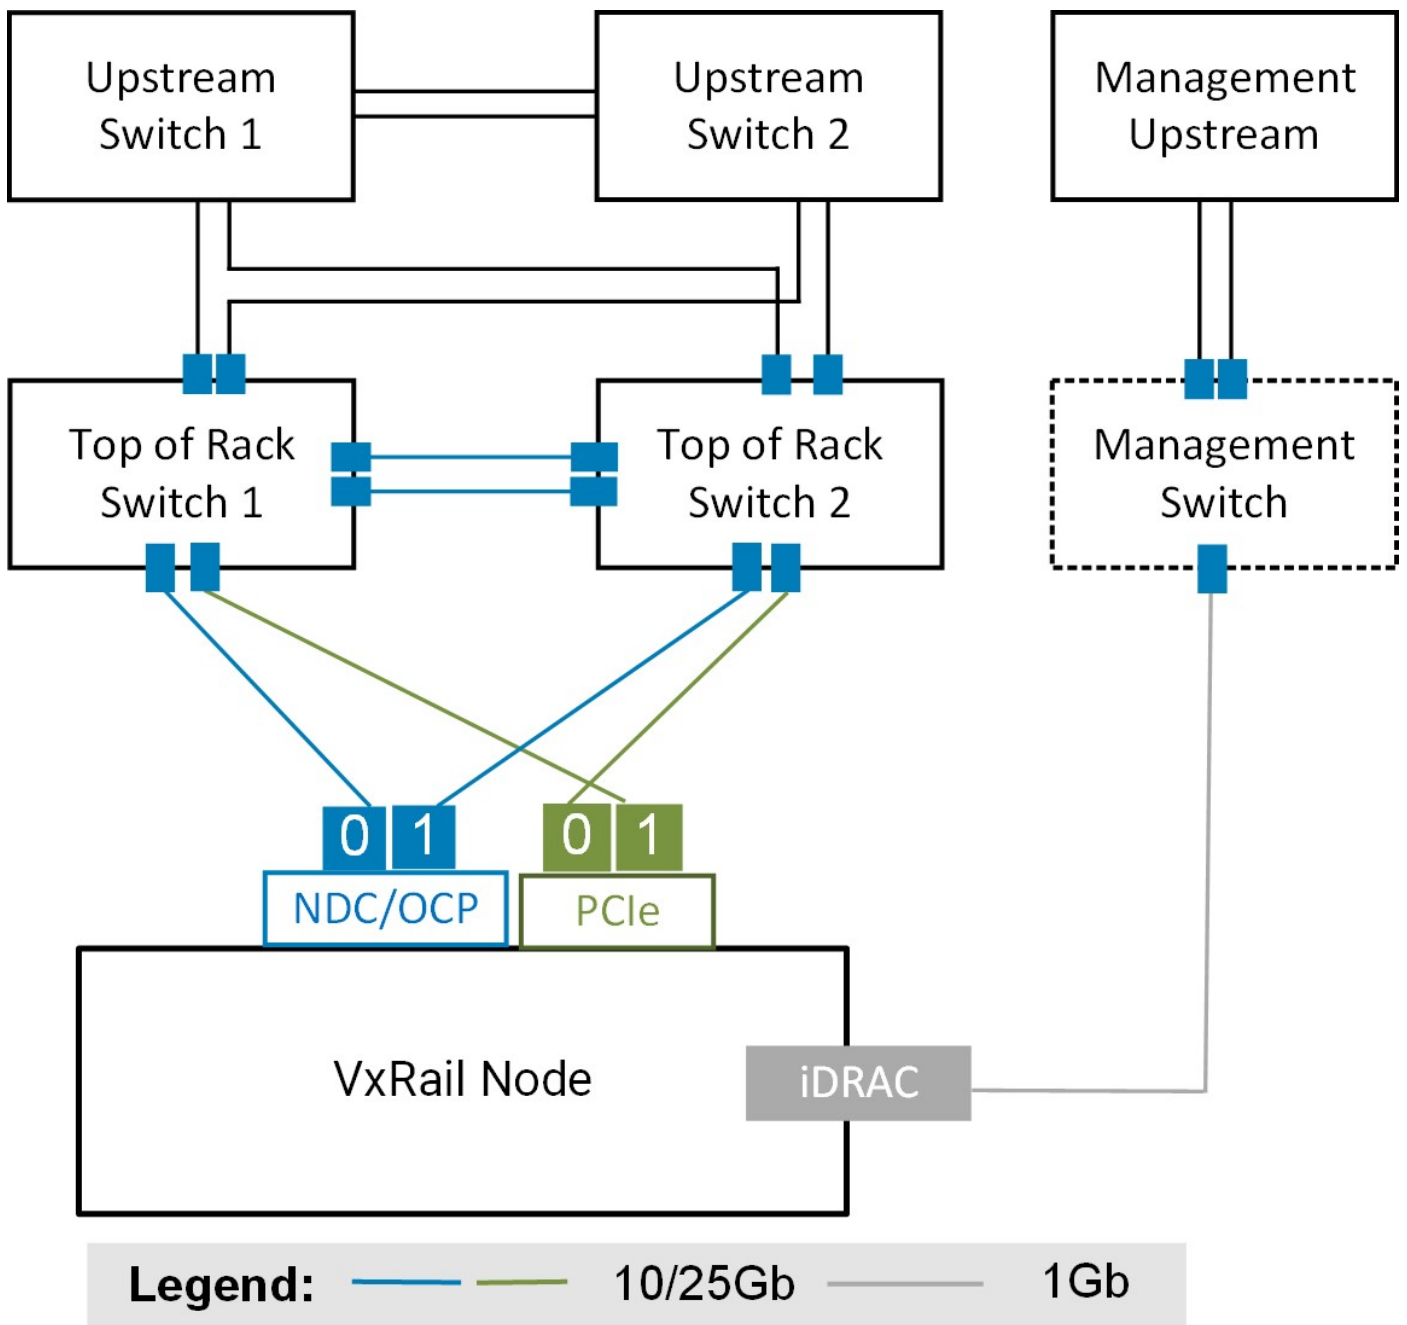


Figure 73. VxRail nodes with two 10/25 Gb NDC/OCP ports and two 10/25 Gb PCIe ports connected to two ToR switches, and one optional connection to management switch for iDRAC

In this predefined network profile option, two NDC/OCP 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/OCP 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, and also with a failure in either the NDC/OCP or PCIe adapter card.

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

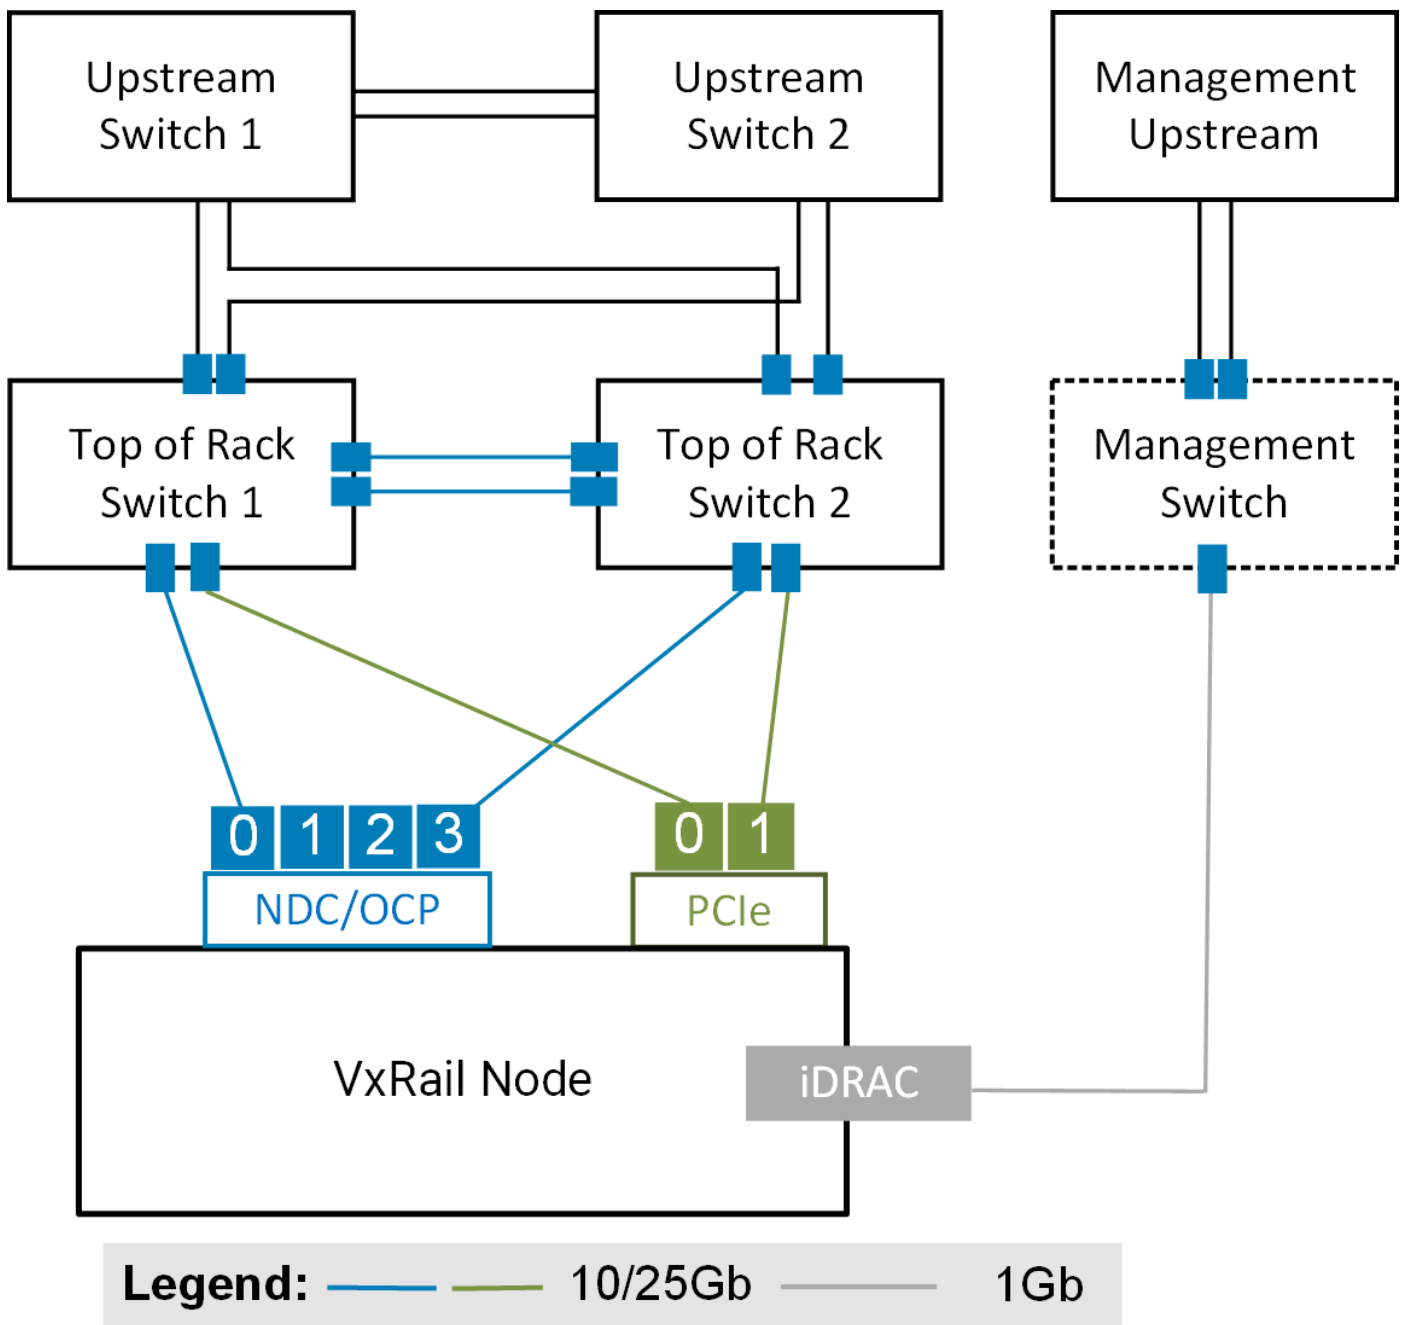


Figure 74. VxRail nodes with any two 10/25 Gb NDC/OCP ports and two 10/25 Gb PCIe ports connected to two ToR switches, and one optional connection to management switch for iDRAC

This is an example of a custom cabling setup with two NDC/OCP ports and 2 PCIe ports that are connected to a pair of 10gb switches or 25gb switches. Any NDC/OCP port and any PCIe port can be selected to support VxRail networking. However, the two NICs assigned to support a specific VxRail network must be of the same type and running at the same speed.

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

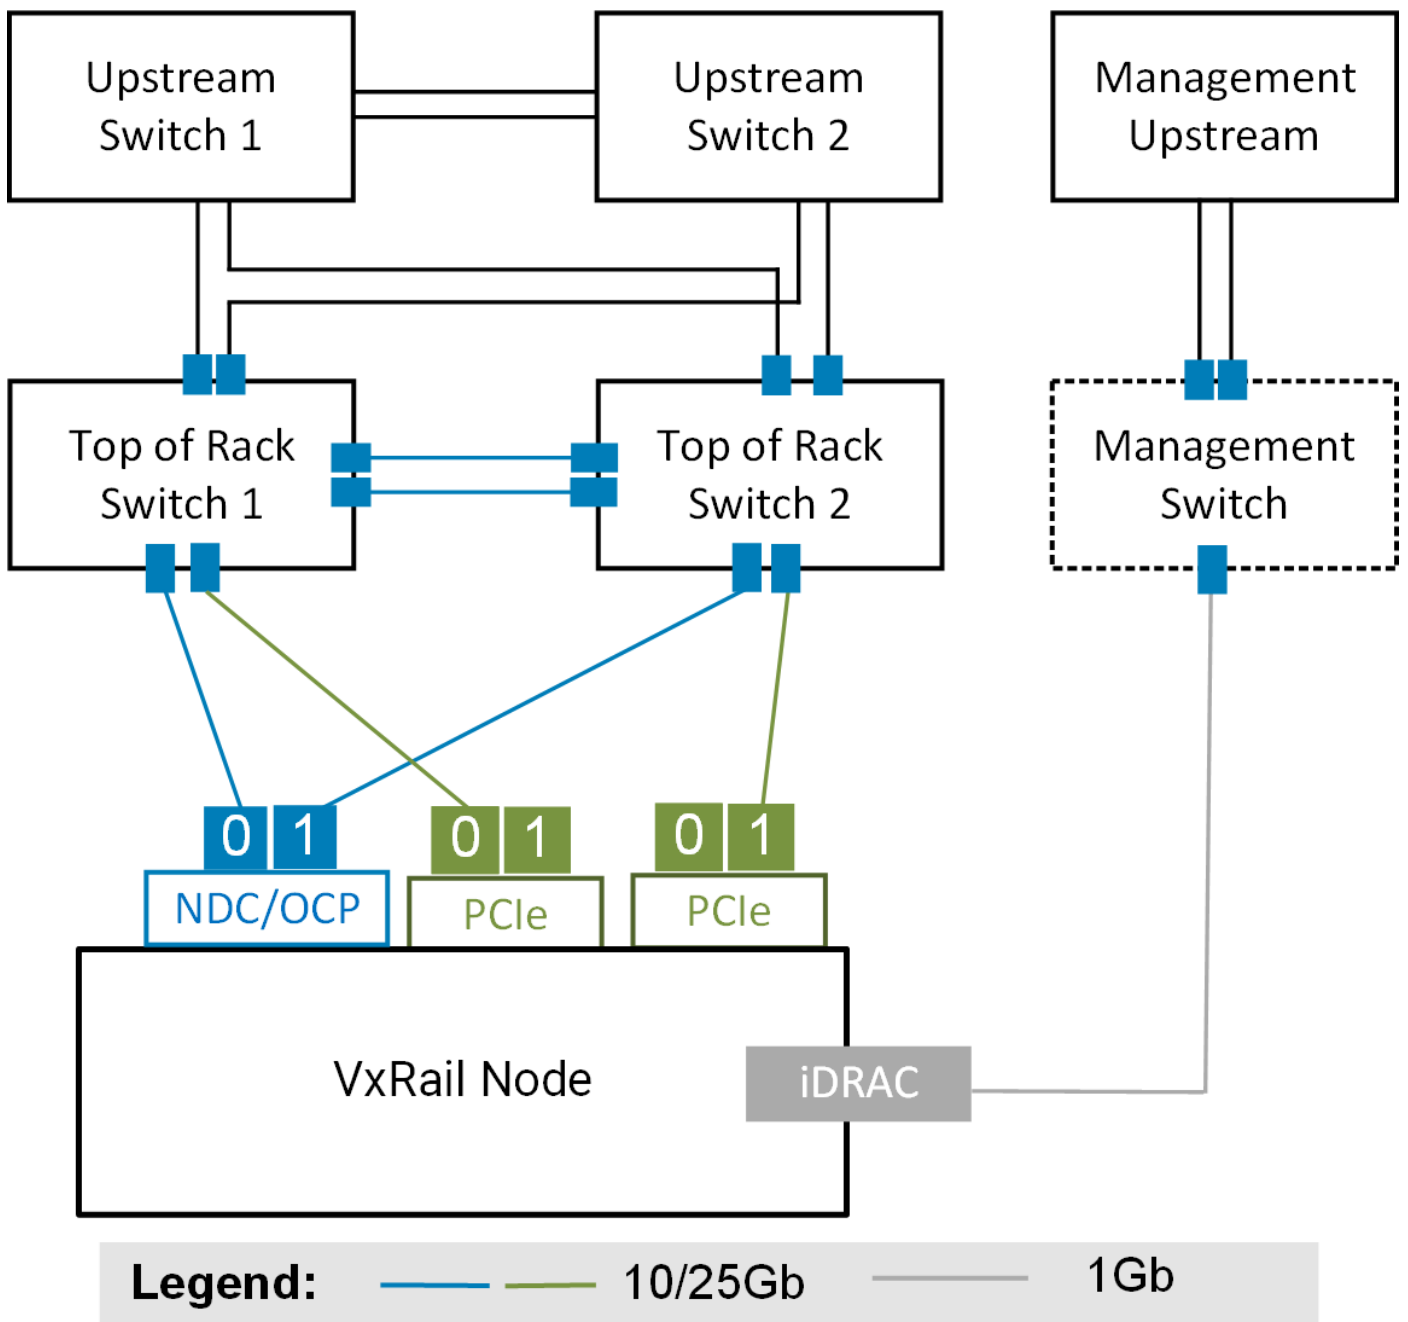


Figure 75. VxRail nodes with any two 10/25 Gb NDC/OCP ports and any two 10/25 Gb PCIe ports connected to two ToR switches, and one optional connection to management switch for iDRAC

With the custom option, there is no restriction about the ports that are selected for VxRail networking reside on the PCIe adapter card in the first slot. If there is more than one PCIe adapter card, ports can be selected from either card.

Custom option: PCIe ports only

In this outlier use case where there is a specific business or operational requirement for this topology, 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.

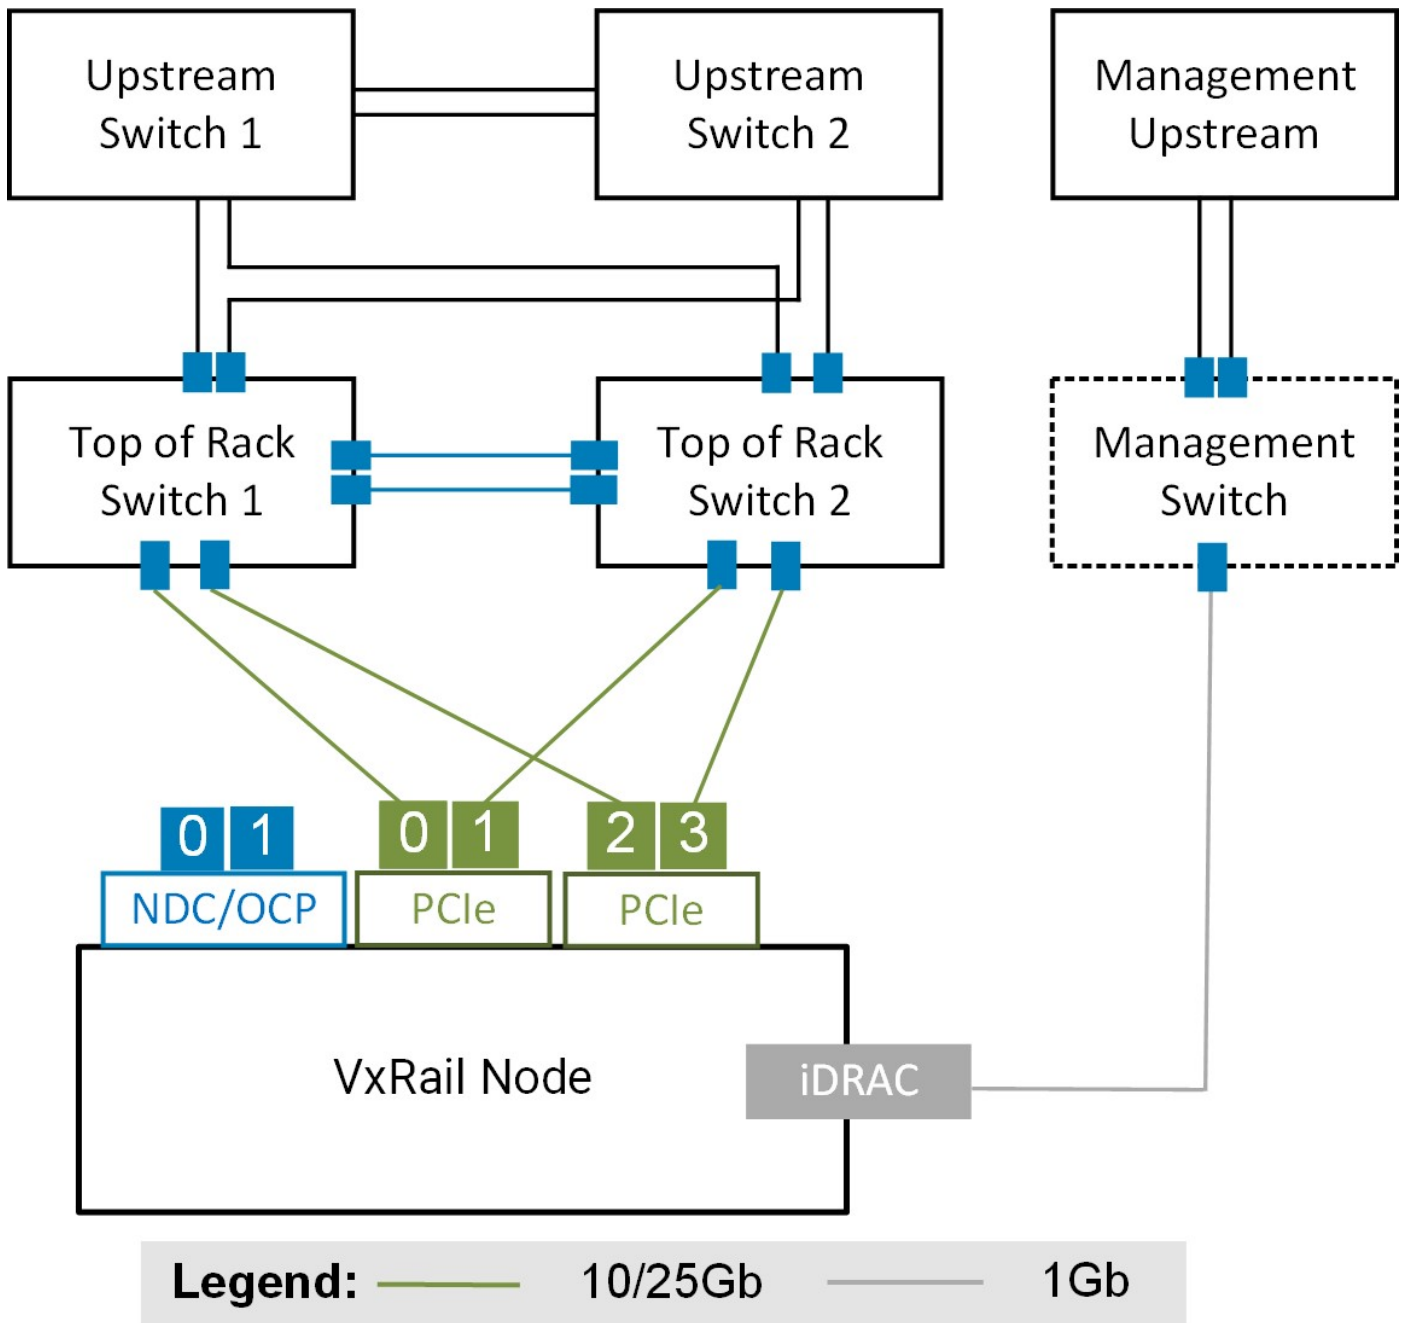


Figure 76. VxRail nodes with two or four PCIe ports connected to a pair of ToR switches, and one optional connection to management switch for iDRAC

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

Custom option: Six ports

VxRail is most commonly deployed with two or four ports. For more network-intensive workload requirements, VxRail can be deployed with six or eight 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-intensive 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.

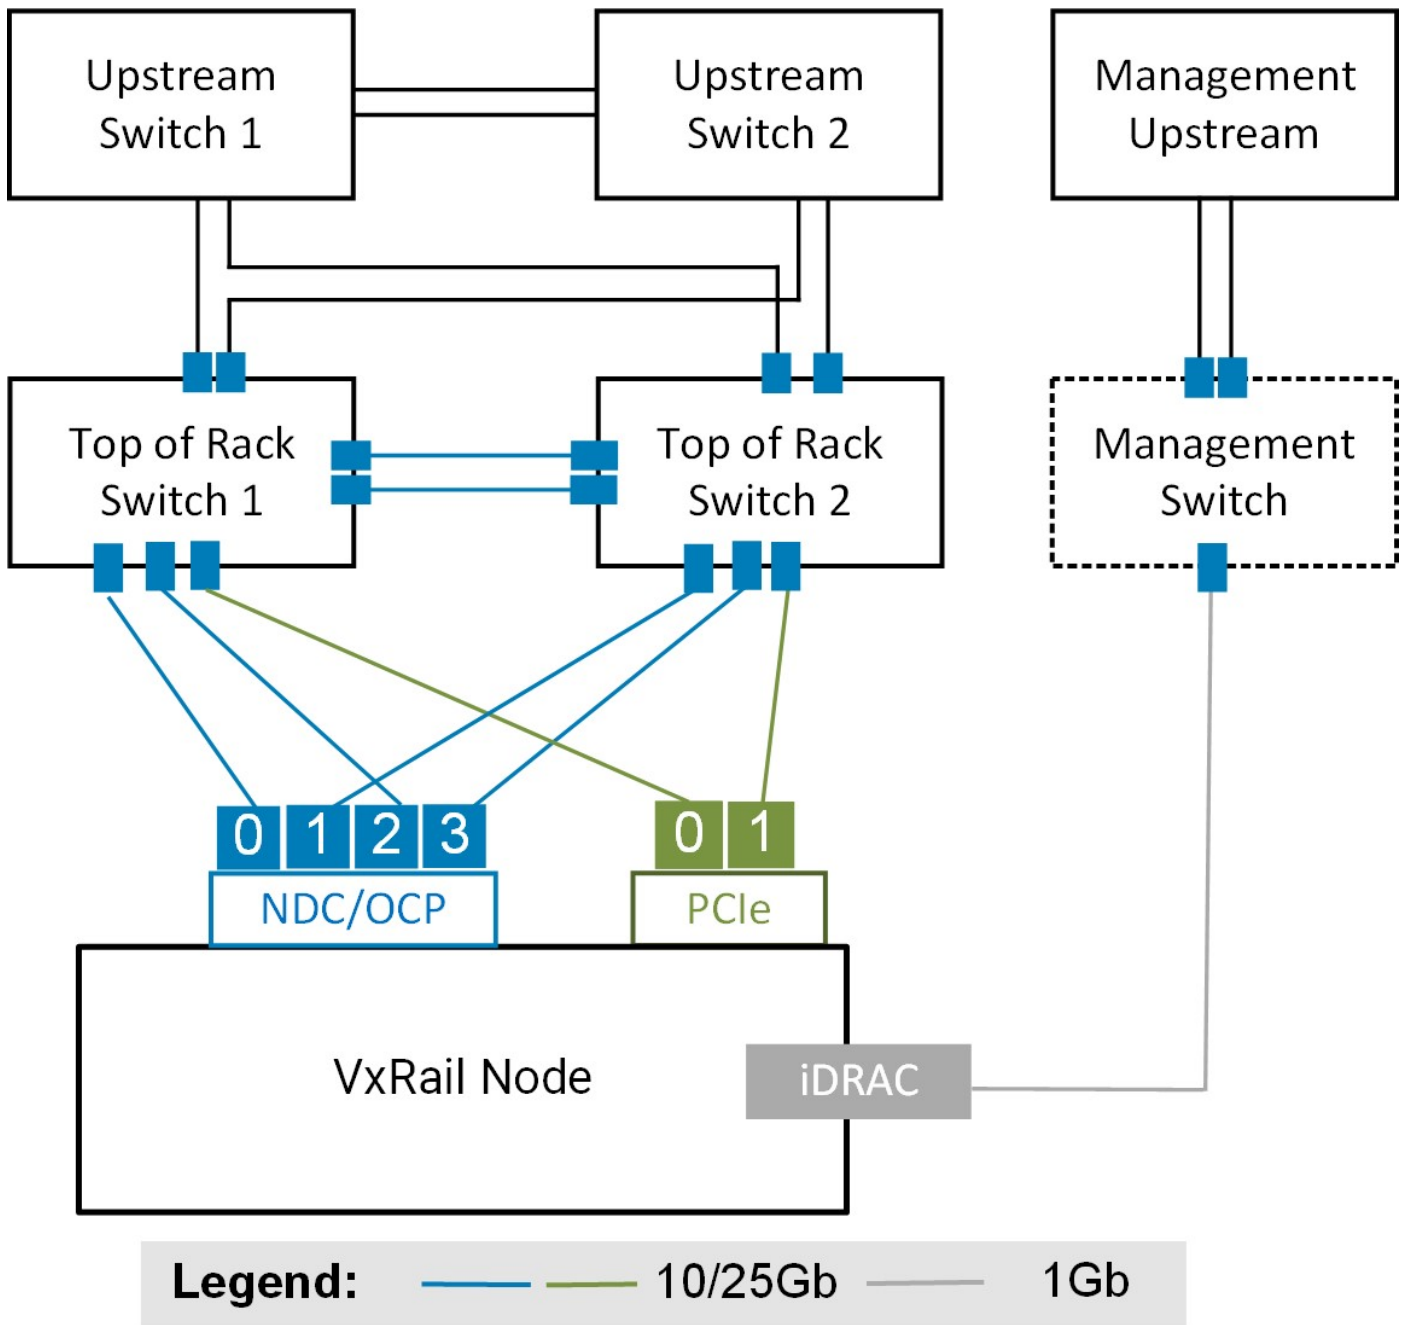


Figure 77. VxRail nodes with four NDC/OCP ports and a pair of PCIe ports connected to a pair of ToR switches, and one optional connection to management switch for iDRAC

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 that are occupied with network adapter cards, this offers the flexibility to spread the VxRail networking workload across three network adapter cards.

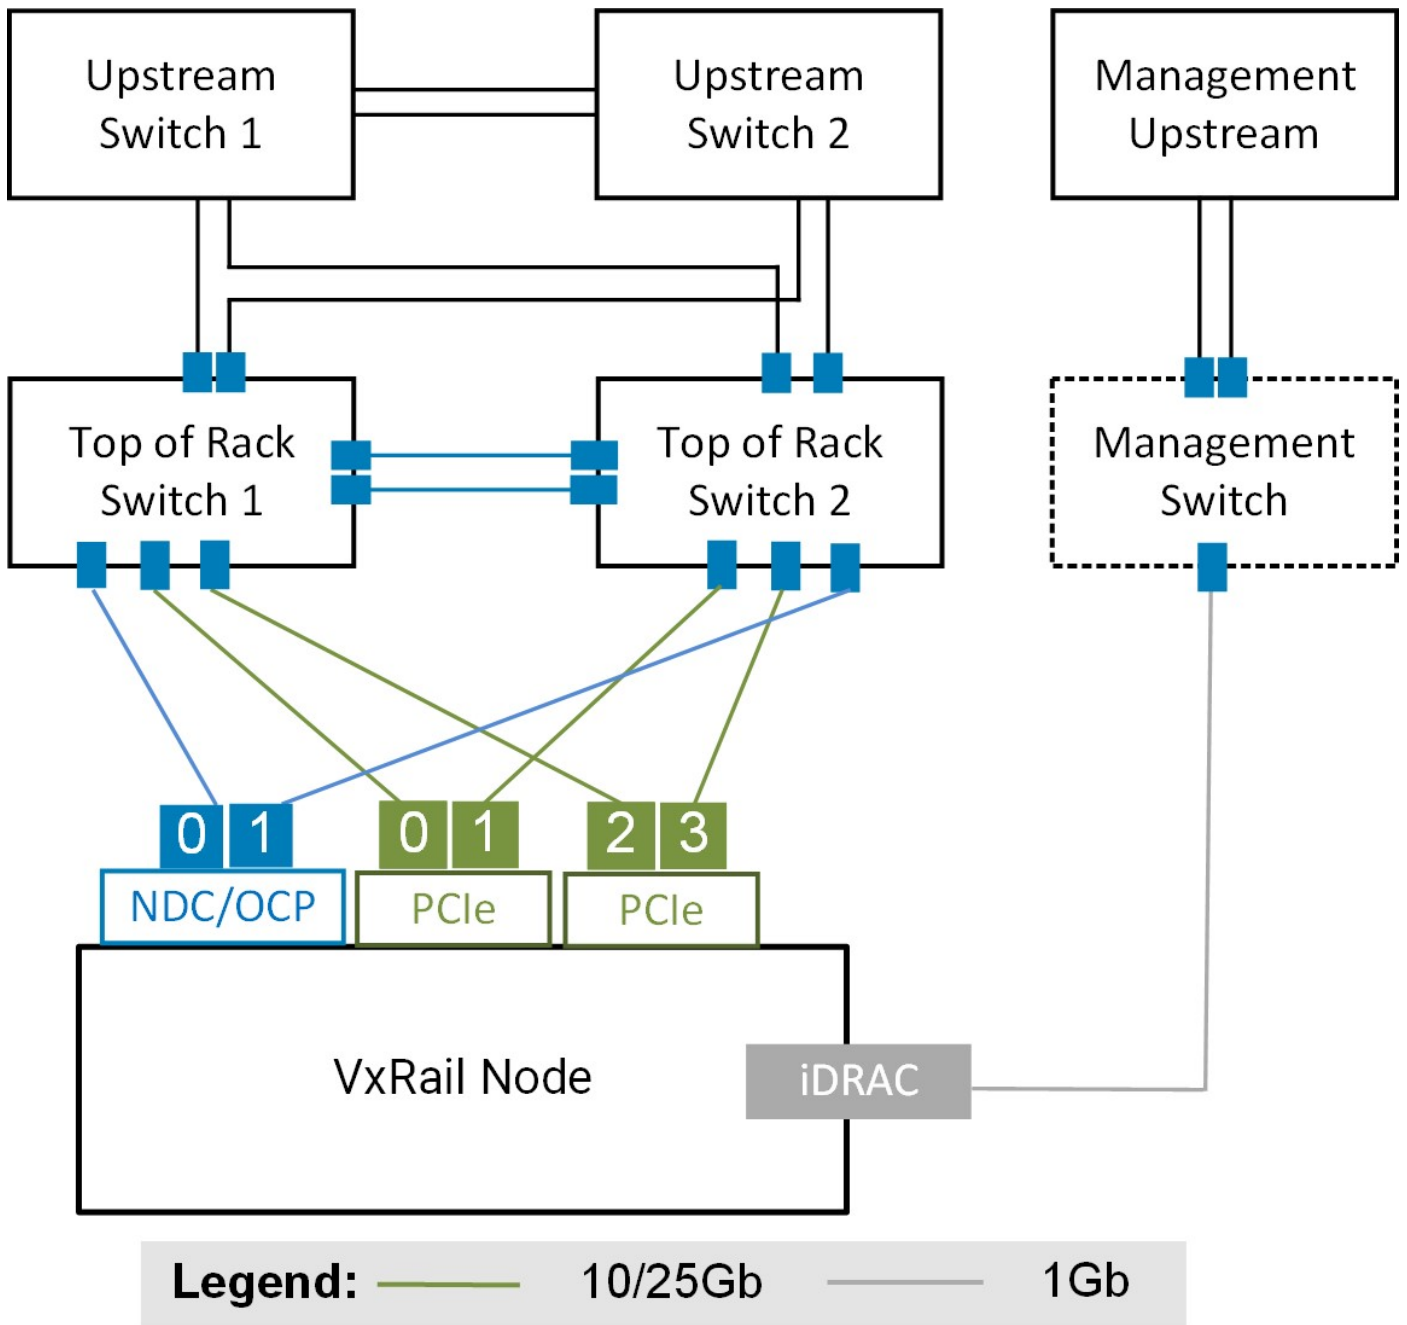


Figure 78. VxRail nodes with two NDC/OCP ports and ports from a pair of PCIe adapter cards connected to a pair of ToR switches, and one optional connection to management switch for iDRAC

Custom option: Eight ports

For workload use cases with extreme availability, scalability, and performance requirements, up to eight ports can be selected to support VxRail networking. This option is advantageous if it is desirable to have resource-intensive networks such as vSAN or VMware vSphere vMotion have dedicated Ethernet ports. This option can also be useful if you want VxRail to configure a guest network with dedicated Ethernet ports as part of the initial build process.

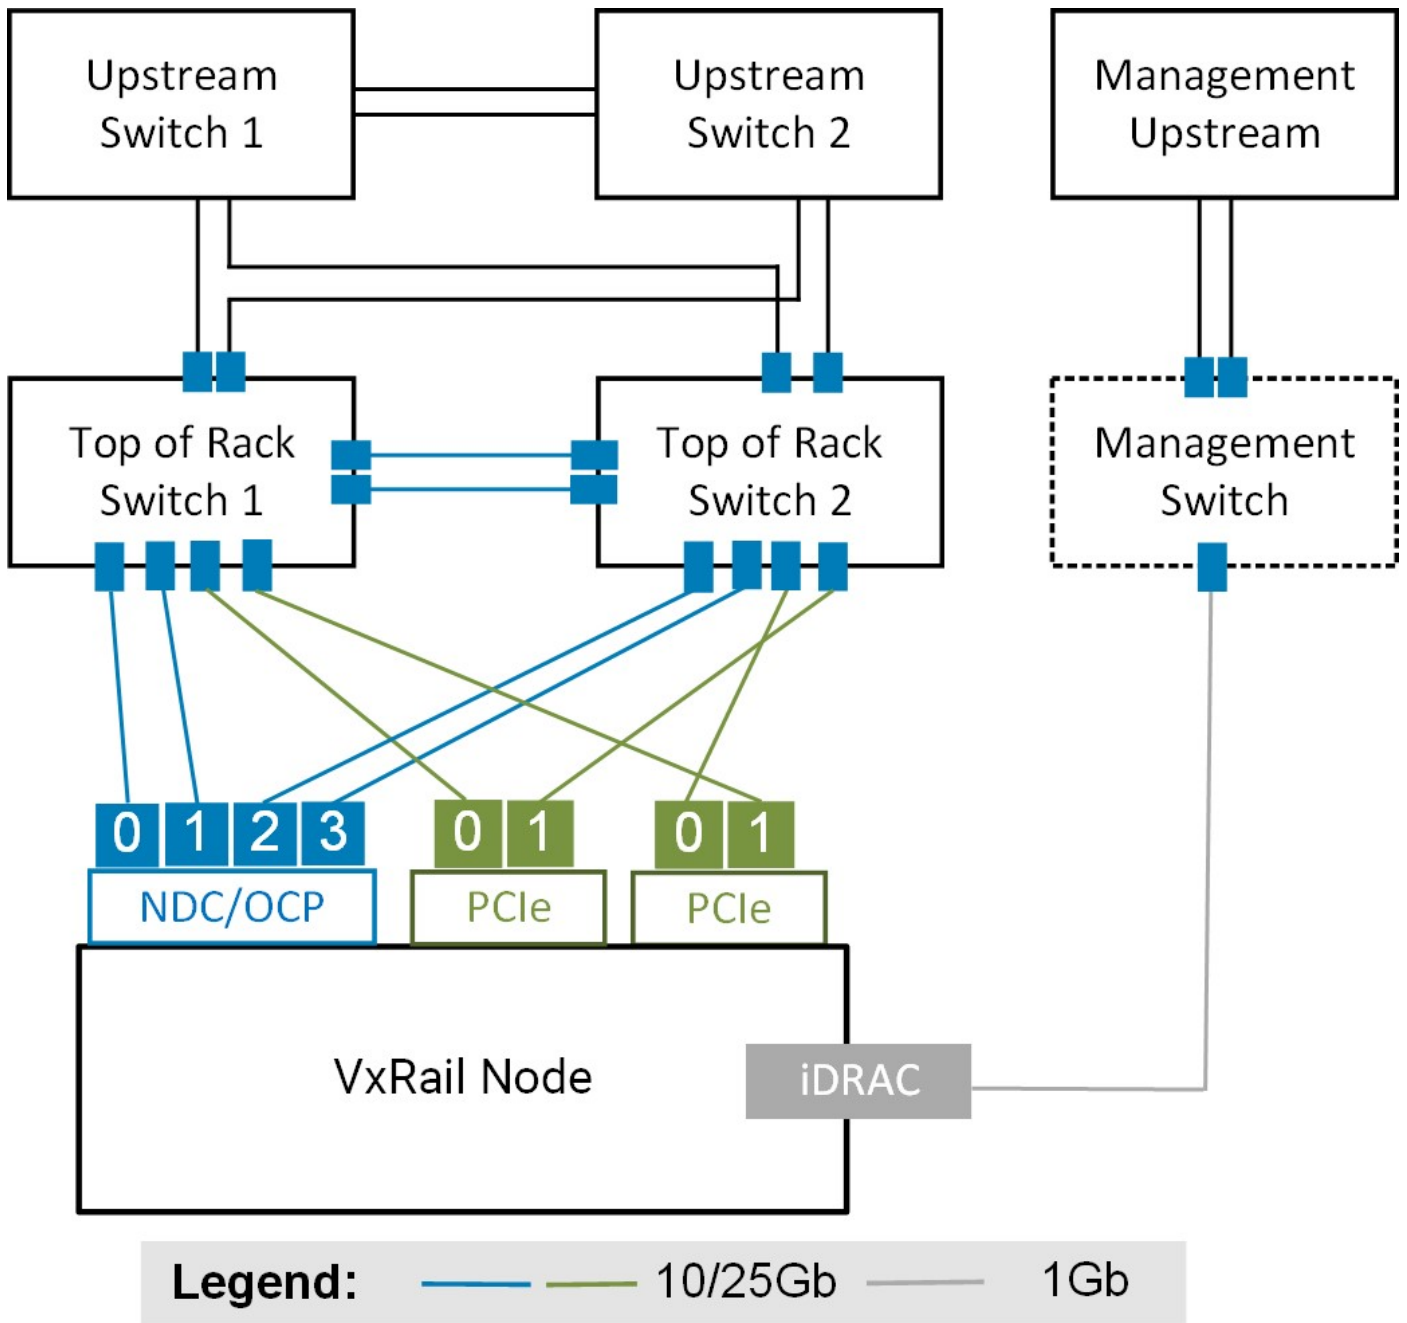
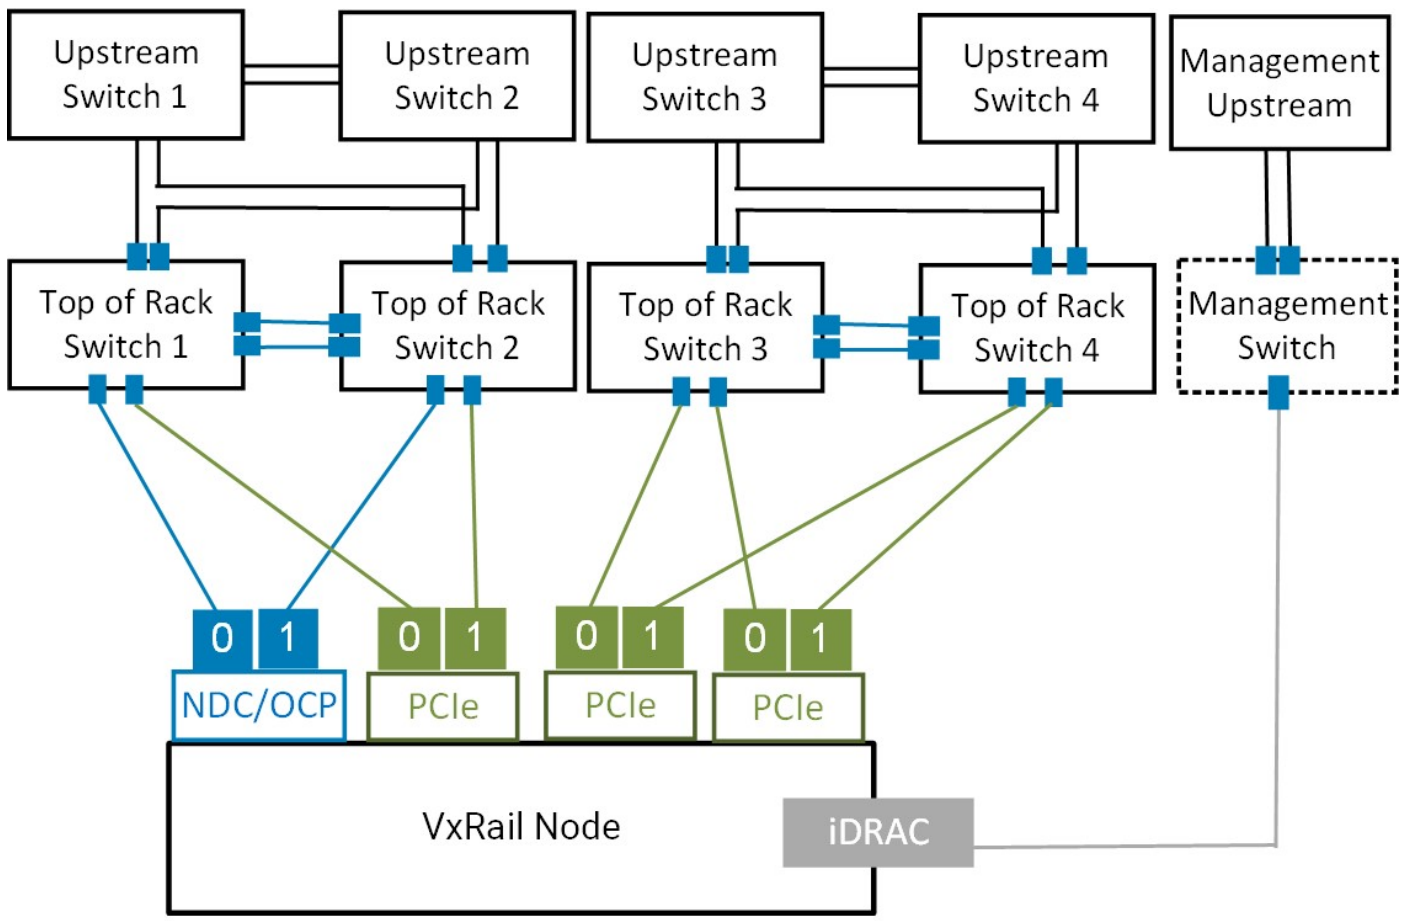


Figure 79. VxRail nodes with four NDC/OCP ports and two ports from a pair of PCIe adapter cards connected to a pair of ToR switches, and one optional connection to management switch for iDRAC

Custom option: Eight ports connected to four Ethernet switches

This topology also addresses the use case of physical network separation to meet specific security policies or governance requirements, with four Ethernet switches that are positioned to support VxRail networking. For instance, the networks that are 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 topology is also applicable for workload use cases with extreme availability, scalability, and performance requirements. For instance, each VxRail network can be configured for redundancy at the switch level and network adapter level if the nodes are installed with four adapter cards containing two ports each.



Legend: — 10/25Gb — 10/25Gb — 1Gb

Figure 80. VxRail nodes with eight ports connected to four Ethernet switches, and one optional connection to management switch for iDRAC