

# VMware Cloud Foundation 4.4 on VxRail Planning and Preparation Guide

February 2022

## Abstract

This guide is for customers interested in deploying VMware Cloud Foundation on VxRail. This guide covers version 4.4 of Cloud Foundation on VxRail. It outlines the planning and preparation that needs to be undertaken before commencing with the product deployment.

## Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA 02/22 Planning Guide H18416.

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

# Contents

<b>Chapter 1..... Executive Summary</b>	<b>7</b>
Solution description .....	8
Solution version reference .....	8
Document purpose .....	8
Intended audience .....	8
Revisions .....	8
<b>Chapter 2..... VMware Cloud Foundation on VxRail</b>	<b>10</b>
Product overview .....	11
<b>Chapter 3..... VMware Cloud Foundation on VxRail Requirements and Use Cases .... 15</b>	
Introduction.....	16
VMware Cloud Foundation on VxRail workload planning.....	17
VMware Cloud Foundation on VxRail deployment overview .....	18
<b>Chapter 4..... VMware Cloud Foundation on VxRail Workload Domain Planning 20</b>	
Introduction.....	21
NSX-T management options for VI workload domain .....	22
vRealize software option for VI workload domains .....	23
vSphere with Kubernetes workload domains .....	23
<b>Chapter 5..... Data Center and Network Requirements</b>	<b>24</b>
Introduction.....	25
Data center rack space requirements.....	25
Data center networking.....	25
Network switch selection .....	26
Switch port capacity.....	26
Switch port type .....	27
Jumbo Frames.....	27
Multicast .....	27
Border Gateway Protocol .....	27
Hardware VTEP for multi-rack deployments .....	28
Network services .....	28
Storage Services .....	29

<b>Chapter 6..... Cloud Foundation on VxRail High-Level Design</b>	<b>31</b>
Introduction.....	32
Consolidated vs. standard architecture .....	32
Site locations .....	33
Application availability .....	33
VxRail cluster network planning .....	34
<b>Chapter 7..... Cloud Foundation on VxRail Workload Planning</b>	<b>40</b>
Introduction.....	41
Determine use cases for Cloud Foundation VI workload domain.....	41
Deciding on single-site VxRail cluster or stretched cluster.....	41
Planning the management workload domain resource requirements .....	43
Planning the VI workload domain resource requirements .....	44
Sizing the Cloud Foundation domains.....	44
<b>Chapter 8..... Application Dependencies and Routing Decisions</b>	<b>48</b>
Understanding connection dependencies .....	49
<b>Chapter 9..... Cloud Foundation on VxRail Physical Network Planning</b>	<b>52</b>
Introduction.....	53
Select a physical network architecture and topology.....	53
VxRail stretched cluster physical network planning.....	55
Fibre Channel storage network planning.....	58
<b>Chapter 10 ... Cloud Foundation on VxRail Physical Network Preparation</b>	<b>60</b>
Introduction.....	61
Capture configuration settings for Cloud Foundation on VxRail.....	61
VxRail cluster and NSX-T networks leaf switch preparation .....	61
DHCP services for NSX-T host overlay network preparation .....	63
Leaf switch preparation for NSX-T edge .....	65
Layer 3 network preparation.....	65
BGP peering preparation.....	67
<b>Chapter 11 ... VxRail Cluster Deployment Preparation</b>	<b>70</b>
Introduction.....	71
Prepare for VxRail cluster initial build .....	71
Select the external management network subnet .....	71
Select the VxRail cluster VLANs .....	72
Select the network settings for VxRail cluster .....	72
Decide whether to join an existing vCenter SSO domain.....	73
Select the network settings for VxRail stretched cluster.....	73
Create forward and reverse DNS entries for VxRail cluster .....	73

Select passwords .....	73
<b>Chapter 12 ...Prepare for VMware Cloud Foundation Management VI Workload Domain</b>	<b>74</b>
Introduction.....	75
Provide a temporary IP address for Cloud Builder .....	75
Select the settings for the management workload domain.....	75
Provide global settings for management VI workload domain.....	75
Select the settings for NSX-T host overlay network .....	76
Create forward and reverse DNS entries for the management VI workload domain .....	76
Select the NSX-T host overlay VLAN .....	76
Select names for resource pools in VI Management workload domain.....	77
Decide on number of virtual distributed switches .....	77
Prepare passwords.....	78
Obtain VMware license keys .....	78
<b>Chapter 13 ...Prepare for Cloud Foundation VI Workload Domain</b>	<b>79</b>
Introduction.....	80
Cloud Foundation workload domain task outline.....	80
Prepare NSX-T host overlay network.....	80
Capture settings for VI workload domain.....	81
Prepare for vSphere for Tanzu workload domain .....	81
Prepare for multi-region with NSX-T Federation .....	83
<b>Chapter 14 ...Prepare for NSX-T Edge Gateway Services</b>	<b>85</b>
Introduction.....	86
Capture external router settings for eBGP peering .....	86
Capture settings for NSX-T edge gateway uplinks.....	86
Capture NSX-T edge overlay network settings .....	87
Capture second site settings for stretched cluster.....	87
<b>Chapter 15 ...Prepare for Cloud Foundation Application Virtual Network</b>	<b>89</b>
Introduction.....	90
Capture the Application Virtual Network region settings .....	90
<b>Appendixes .92</b>	
Appendix A: Cloud Foundation on VxRail checklist.....	93
Appendix B: Cloud Foundation on VxRail footprints for sizing .....	97
Appendix C: Cloud Foundation on VxRail VLANs .....	99
Appendix D: VxRail network configuration .....	100
Appendix E: Cloud Builder and management VI workload configuration .....	102
Appendix F: VI workload domain configuration settings.....	104

Appendix G: Edge Gateway configuration.....	106
Appendix H: Application Virtual Network configuration .....	107
Appendix I: Sample switch configuration settings .....	108

# Chapter 1 Executive Summary

This chapter presents the following topics:

- Solution description** ..... 8
- Solution version reference** ..... 8
- Document purpose** ..... 8
- Intended audience** ..... 8
- Revisions**..... 8

## Solution description

VMware® Cloud Foundation (VCF) on VxRail™ is a Dell Technologies and VMware jointly engineered integrated solution. It contains features that simplify, streamline, and automate the operations of your entire Software-Defined Datacenter (SDDC) from Day 0 through Day 2. The new platform delivers a set of software-defined services for compute (with vSphere and vCenter), storage (with vSAN), networking (with NSX), security, and cloud management (with vRealize Suite) in both private and public environments, making it the operational hub for your hybrid cloud.

VCF on VxRail provides the simplest path to the hybrid cloud through a fully integrated hybrid cloud platform that leverages native VxRail hardware and software capabilities and other VxRail-unique integrations (such as vCenter plugins and Dell networking). These components work together to deliver a new turnkey hybrid cloud user experience with full-stack integration. Full-stack integration means you get both HCI infrastructure layer and cloud software stack in one complete automated life-cycle turnkey experience.

## Solution version reference

This guide supports the major software release 4.4 of VMware Cloud Foundation, and major software release 7.0 of VxRail. The specific versions of the software stack supported for the major versions of VMware Cloud Foundation on VxRail covered in this guide can be found in the [VMware Cloud Foundation 4.x on VxRail Support Matrix](#).

## Document purpose

This guide provides detailed guidance for the initial deployment of a VCF on VxRail solution in a data center. It outlines the tasks and processes that you should expect and prepare for from the planning and design phase through deployment of the solution. The guide also serves as an aid in helping determine a configuration that meets your business and operational objectives.

## Intended audience

This planning and preparation guide is intended for cloud architects, network architects, and technical sales engineers who are interested in the planning, designing, and deployment of the VCF on VxRail solution to meet business and operational requirements. Readers should be familiar with VMware vSphere, NSX, vSAN, and vRealize product suites in addition to general network architecture concepts.

## Revisions

Date	Description
April 2019	Initial release
August 2019	Updated to support VMware Cloud Foundation 3.8
February 2020	Updated to support VMware Cloud Foundation 3.9.1



Date	Description
May 2020	Updated to support VMware Cloud Foundation 4.0
July 2020	Updated for VxRail stretched cluster requirements post-deployment
August 2020	Updated to support VMware Cloud Foundation 4.0.1
September 2020	Updated to support VMware Cloud Foundation 4.1 on VxRail 7.0.010
October 2020	Removed references to Witness Traffic Separation
November 2020	Updated to support VMware Cloud Foundation 4.1 on VxRail 7.0.100
March 2021	Updated to support VMware Cloud Foundation 4.2 on VxRail 7.0.131
April 2021	<ul style="list-style-type: none"> <li>• Updated for better alignment with Cloud Foundation features</li> <li>• Updated content for stretched cluster networking requirements</li> </ul>
September 2021	Updated to support VMware Cloud Foundation 4.3 on VxRail 7.0.202
October 2021	Updated to support VMware Cloud Foundation 4.3.1 on VxRail 7.0.241
February 2022	Updated to support VMware Cloud Foundation 4.4 on VxRail 7.0.320

## Chapter 2 VMware Cloud Foundation on VxRail


This chapter presents the following topic:

<b>Product overview</b> .....	<b>11</b>
-------------------------------	-----------

## Product overview

The VMware Cloud Foundation on VxRail solution is integrated end-to-end to fully enable a software-defined cloud platform that is designed for the rapid deployment of physical resources into managed consumption pools, and for the provisioning of these resource pools on-demand to meet flexible and resilient workload requirements.

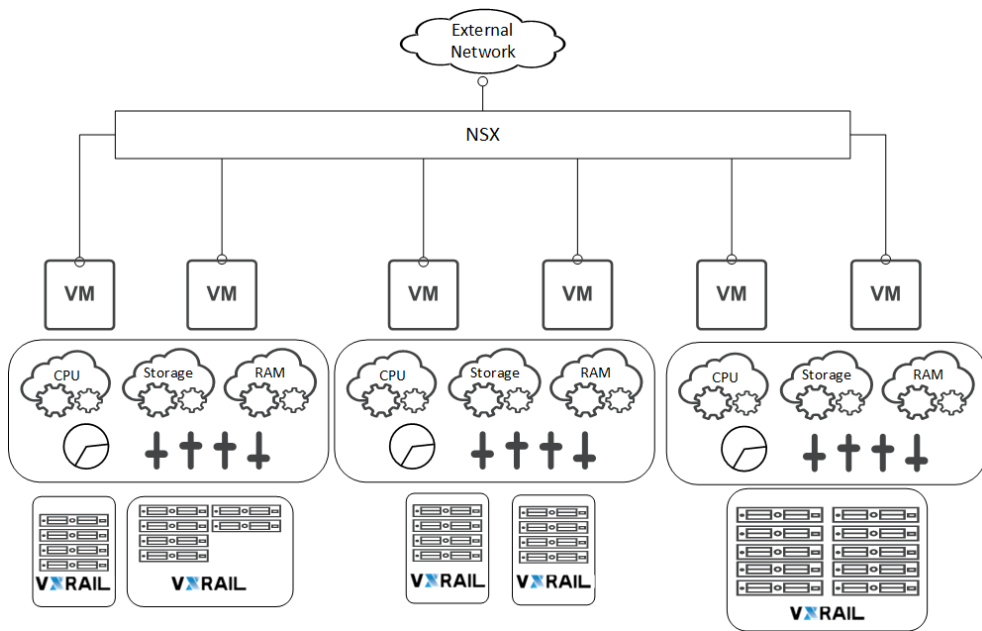
VxRail provides the physical resource foundation for the cloud delivery platform. VxRail is a set of specially engineered and manufactured compute nodes that when logically bound together after initial configuration, represent a single managed cluster for virtual workloads.



**Figure 1. VxRail cluster representing a pool of virtual resources**

VxRail integrates software products from VMware with custom software engineered from Dell Technologies so that the physical compute, memory, network, and storage resources are placed under a virtualization layer to be managed and controlled as an adaptable pool of resources. The physical disk devices on each VxRail node are encapsulated under the virtualization layer to create a single consumable data store for the virtual workloads. In addition, a virtual switch is created during initial configuration and distributed across the entire VxRail cluster. The Ethernet ports on each node are placed under the virtualization layer to enable connectivity between virtual machines on the VxRail cluster, and to enable connectivity to end-users.


When integrated with VMware Cloud Foundation, the VxRail cluster is positioned as an individual building block to supply compute resources for consumption in Cloud Foundation virtual workloads. Cloud Foundation allows users to dynamically allocate and assign VxRail clusters into individual consumption pools, known as Virtual Infrastructure (VI) workload domains. A VI workload domain represents the logical boundary of consumable resources, and all functionality within these boundaries is managed through a single vCenter instance. Under this model, VI workload domains can be planned and deployed to support the distinct requirements of individual organizations or a set of applications.



**Figure 2. VxRail clusters as building blocks for Cloud Foundation virtual workload consumption**

The resources of individual VI workload domains can be expanded through the addition of individual nodes into a VxRail cluster, or through the addition of an entire new VxRail cluster into a VI workload domain. The physical resources are automatically added to the VI workload domain pool upon completion of this event.

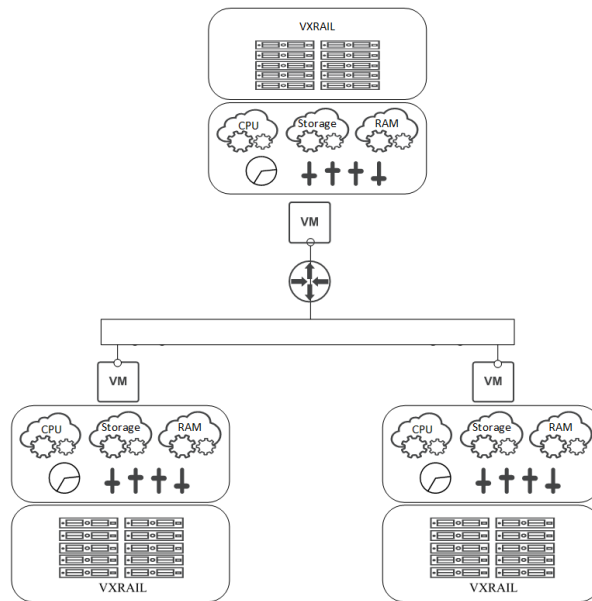
The networking resources for each VI workload domain are also logically segmented, so that the distinct requirements for a set of applications can be individually managed. With the layering of the VMware's Cloud Foundation software stack on VxRail virtual switches, enterprise networking features such as routing, VPN, and security from NSX-T are embedded and enabled into each VI workload domain.



**Figure 3. Cloud Foundation VI workload domains with fully virtualized resources**

With support for NSX-T, virtual machine traffic that previously had to pass upstream through to the physical network for routing purposes can now traverse the virtual network when established on a Cloud Foundation on VxRail VI workload domain.

Virtual machines will connect to the network using a logical switch in a Cloud Foundation domain. Cloud Foundation on VxRail supports the connecting of these virtual switches into an extended logical network, known as a segment. This allows virtual machines in different VI workload domains to connect to each other through this extended switch fabric.



**Figure 4. Virtual machines connected to an extended logical network with logical routing services**

If a virtual machine requires routing services, the extended logical switch, or segment, can use routing services provided by NSX-T within the virtual network. To support connectivity to applications and end-users outside of the virtual network, the NSX-T virtual routing services form a peer relationship with existing upstream physical routers in the data center to share routing information, and form a seamless connection between the physical and logical networks.

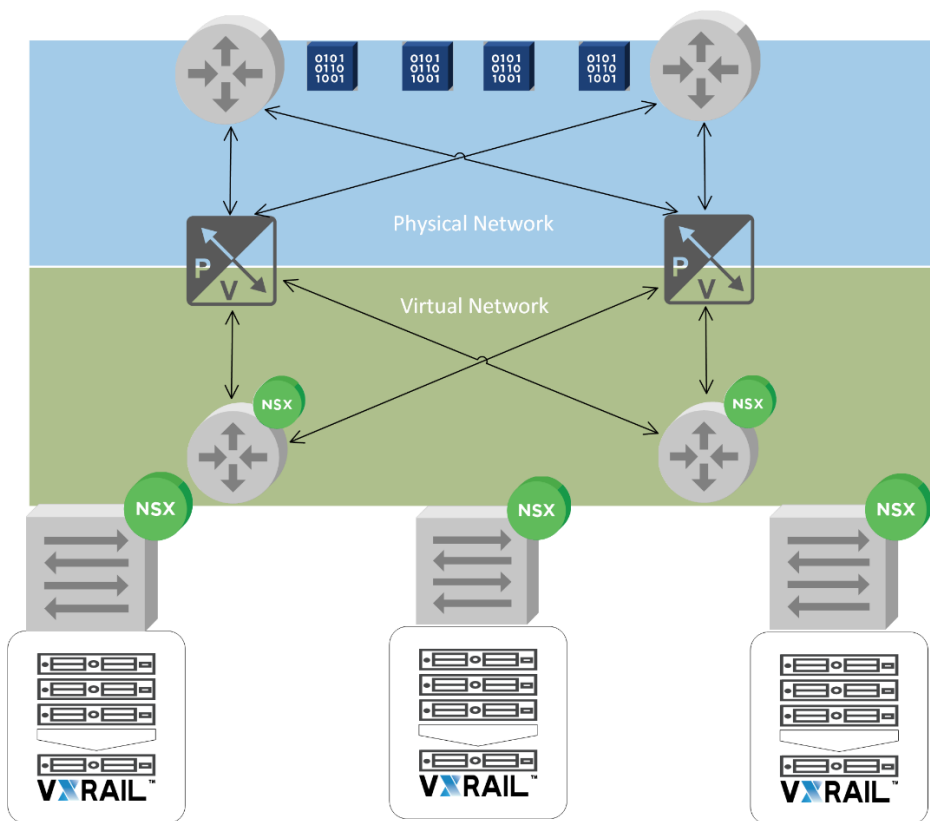


Figure 5. Relationship between physical and virtual networks

## Chapter 3 VMware Cloud Foundation on VxRail Requirements and Use Cases

This chapter presents the following topics:

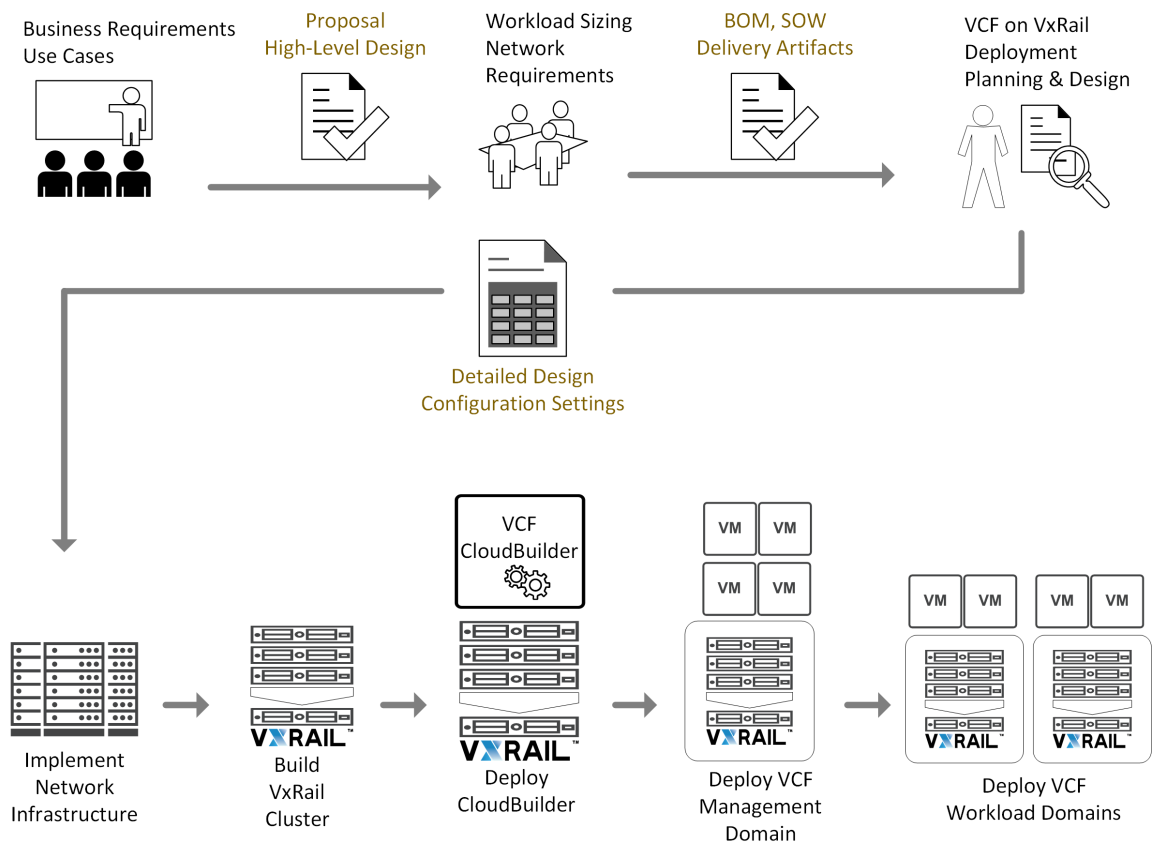
<b>Introduction.....</b>	<b>16</b>
<b>VMware Cloud Foundation on VxRail workload planning .....</b>	<b>17</b>
<b>VMware Cloud Foundation on VxRail deployment overview .....</b>	<b>18</b>

## Introduction

The Cloud Foundation on VxRail cloud platform in a data center has a transformational effect on the way IT resources are delivered to support applications and users. The deployment of a Cloud Foundation on VxRail cloud platform in your environment involves careful and deliberate planning and preparation to ensure an efficient and seamless deployment experience.

The Cloud Foundation on VxRail deployment life cycle starts before a purchase order is issued. In the initial phase, the business and operational requirements are captured and applied toward the overall solution. The requirements process captures the use cases for the planned Cloud Foundation on VxRail deployment. At this stage, decisions can be made about requirements such as site locations and availability. In addition, various organizations and business units are aligned with their application requirements to propose a high-level design. Dell Technologies specialists work jointly with the account team at this stage of the effort.

After acceptance of a high-level design and proposal, technologists and subject matter experts will join the effort. The applications and virtual machines targeted for the Cloud Foundation on VxRail platform are used in a sizing exercise to produce a detailed VxRail infrastructure bill-of-material needed to support the planned workload.



**Figure 6. Anatomy of a Cloud Foundation on VxRail deployment experience**

Also, during this phase, the dependencies between the planned sets of applications for Cloud Foundation are analyzed and used to produce a high-level network design. These



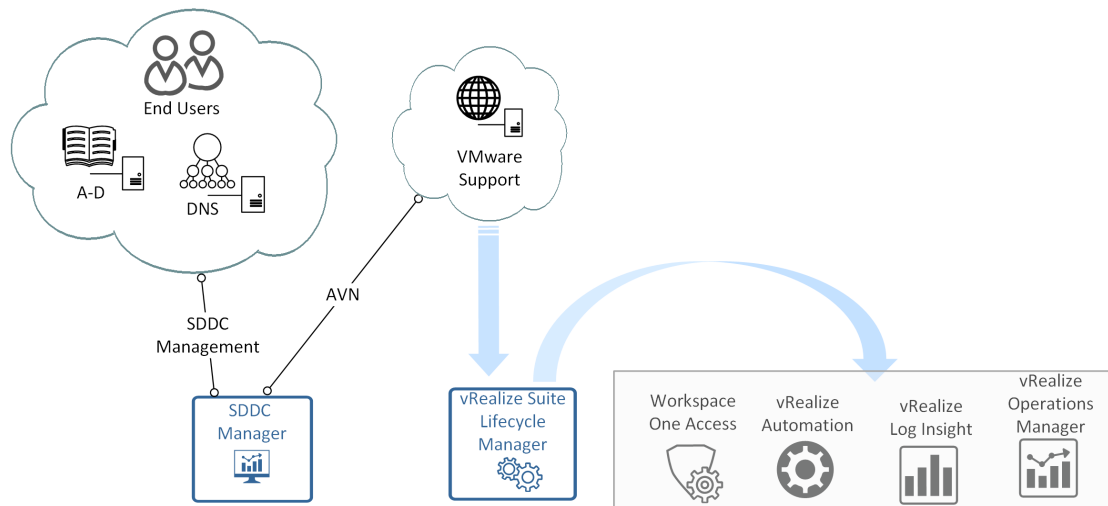
requirements are then used as the baseline in the planning efforts for the upstream external network and for the virtual networks in the management workload domain and the VI workload domains.

## VMware Cloud Foundation on VxRail workload planning

Before moving forward with the initial deployment, decisions must be made on the overall Cloud Foundation architecture based on best practices and use case requirements:

- Resource consumption on the management workload domain depends on decisions made regarding network connectivity for the workload domains. A new workload domain can either leverage existing NSX-T management virtual appliances deployed in an existing workload domain, or new NSX-T management virtual appliances can be deployed to support networking requirements. Capacity must be reserved in instances where additional NSX-T virtual appliances will be deployed.
- The default size of the NSX-T management virtual appliances deployed during Cloud Foundation initial deployment is adequate for most workloads. An expansion of virtual machine workload or the enablement of additional networking services can impact NSX-T management resources, and impose constraints on the management workload domain if there is a lack of resource capacity.
- For smaller, less impactful workload requirements, a consolidated architecture can be considered. A consolidated architecture does not support additional domains beyond the management workload domain, which means the resources for Cloud Foundation management and all application workloads are shared in a single management workload domain. This option should be considered only if the sizing exercises show that a consolidated architecture can be supported, including plans for future growth.
- If a requirement is the deployment of the vRealize software suite to support use case and application requirements, Cloud Foundation on VxRail automates the deployment of vRealize Automation, vRealize Operations Manager, and vRealize Log Insight using vRealize Suite Lifecycle Manager. Selecting this option will require reserving additional resources in the management domain to support this use case.

In addition to the additional resource requirements for vRealize, Cloud Foundation on VxRail will require access to the external network so that SDDC Manager can download vRealize Suite Lifecycle Manager from the VMware support site and perform the installation. To support this, the Application Virtual Network (AVN) needs to be configured at the time of Cloud Foundation on VxRail initial deployment to enable this connectivity. Dell Technologies delivery resources will capture the network integration requirements as part of the planning and design phase.



**Figure 7. SDDC Manager uses AVN to download and deploy vRealize Suite Lifecycle Manager**

## VMware Cloud Foundation on VxRail deployment overview

Once the decisions have been made on workload planning, the work effort transitions to a professional services engagement. The information captured from the requirements and use cases drives the overall sizing effort for the hardware platform. Once the order has been placed and awaiting delivery, Dell Technologies will commence with the planning and design phase. The solutions architect will walk through the deployment process using information gathered in the initial planning phase, and capture the detailed design and configuration settings for the initial deployment of Cloud Foundation on VxRail.

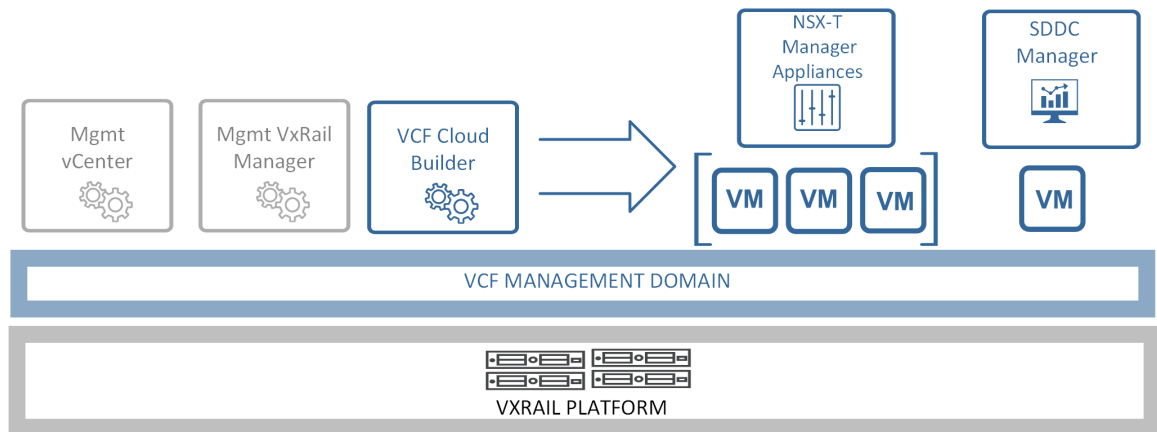
The design and configuration settings that must be captured include specific details about how the Cloud Foundation on VxRail instance will integrate into the data center and network infrastructure. In addition, the settings for the VxRail clusters which will support the Cloud Foundation management and workload domains are captured, including the details for connecting the VxRail nodes into the data center network infrastructure. Then, the settings for the Cloud Foundation Cloud Builder virtual appliance, a tool that is used to drive the initial deployment of Cloud Foundation on VxRail, are captured.

If the deployment of a Cloud Foundation workload domain is part of the scope of work, the settings for the underlying VxRail cluster and initial workload domain properties are also captured.

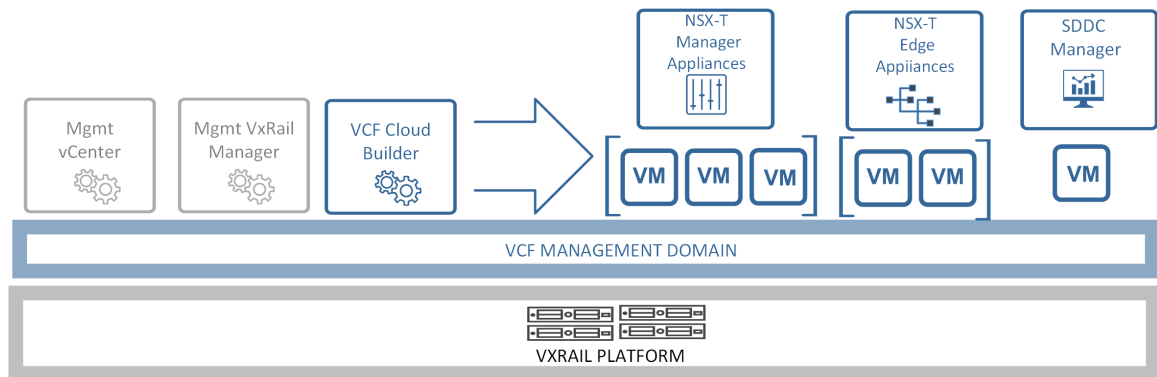
Upon completion of the capture of the required settings for initial deployment, the Dell Technologies solutions architect will perform a validation of the data center environment to ensure that all the prerequisites have been met.

Once the planning and design phase is completed, the next phase is the deployment of the VxRail cluster targeted to support the Cloud Foundation management workload domain, and then the deployment of the Cloud Foundation Cloud Builder virtual appliance on the VxRail cluster. The configuration settings captured from the planning and design phase are fed into the Cloud Builder virtual appliance, which automates the deployment of the Cloud Foundation software onto the VxRail cluster. This creates the Cloud Foundation management workload domain, and deploys the virtual machines required to support the

management workload domain. If a use case requirement is for the future deployment of the vRealize software suite, Cloud Builder deploys three NSX-T Edge virtual appliances into the management workload domain to support upstream connectivity for the Application Virtual Network.



**Figure 8. Overview of Cloud Builder automatic deployment of virtual appliances for a basic Cloud Foundation management workload domain**



**Figure 9. Overview of Cloud Builder automatic deployment of virtual appliances for a Cloud Foundation management workload domain to support the Application Virtual Network**

# Chapter 4 VMware Cloud Foundation on VxRail Workload Domain Planning

This chapter presents the following topics:

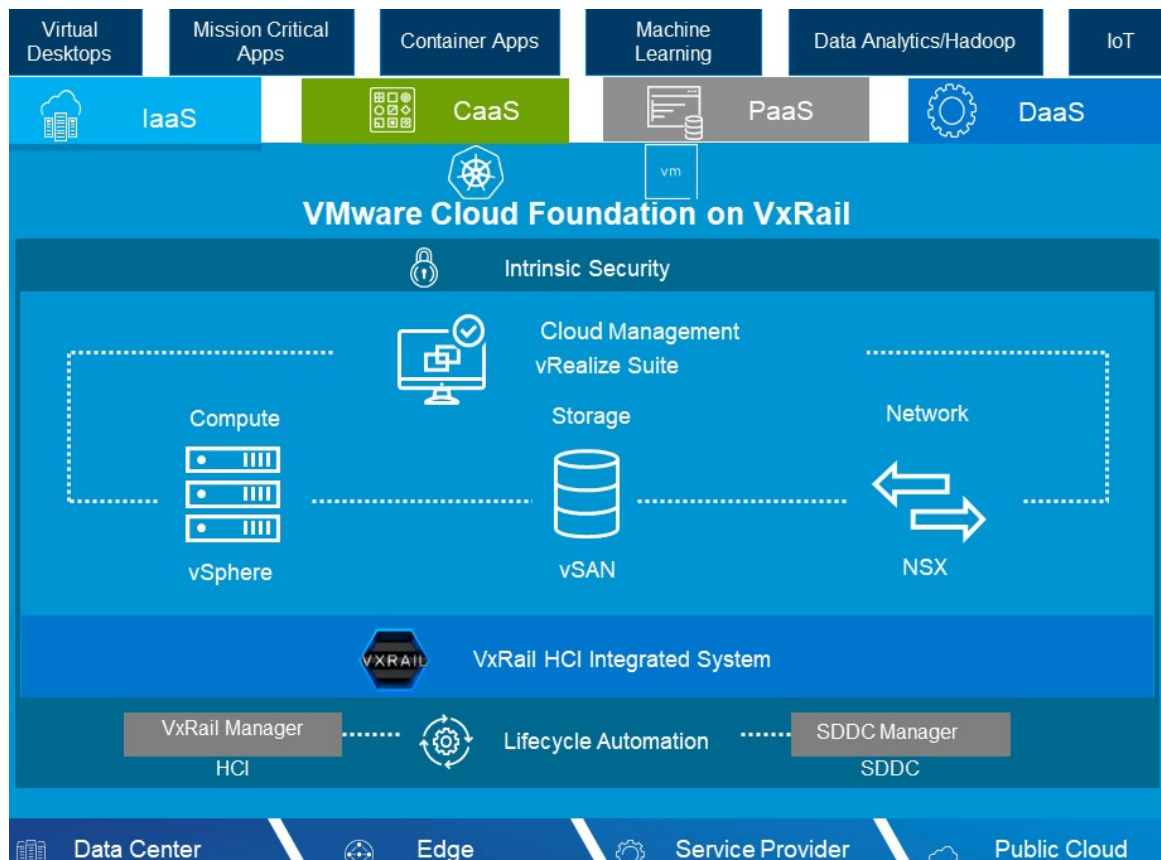
- Introduction..... 21
- NSX-T Management options for VI workload domain ..... 22
- vRealize software option for VI workload domains..... 23
- vSphere with Kubernetes workload domains..... 23

## Introduction

Once Cloud Builder has completed the initial phase of deploying the Cloud Foundation management workload domain, the next phase is the deployment of the Cloud Foundation VI workload domains to support use cases and workload requirements, with the exception being the single-domain consolidated architectures. For this phase, a VI workload domain is initialized using SDDC Manager with initial configuration settings. Next, the underlying VxRail cluster or clusters to support the workload domain are deployed. Then as a final step, the VxRail cluster or clusters are integrated into the VI workload domain.

The initialization of a VI workload domain lays down the basic foundation for the future deployment of virtual machines and their network interconnections. SDDC Manager can also deploy a workload domain to address a specific use case. This action, depending on the option, can alter the pre-requisites and requirements that must be addressed before configuring the workload domain.

This section provides an overview of the workload domains that can be created with Cloud Foundation on VxRail to address specific use cases, and raise awareness of the impact these choices have on the planning and preparation phase. The specific details on how to implement a VI workload domain for a specific use can be found in the Cloud Foundation documentation on the VMware support site: [VMware Cloud Foundation Documentation](#).



**Figure 10. Solutions layers on Cloud Foundation for VxRail**

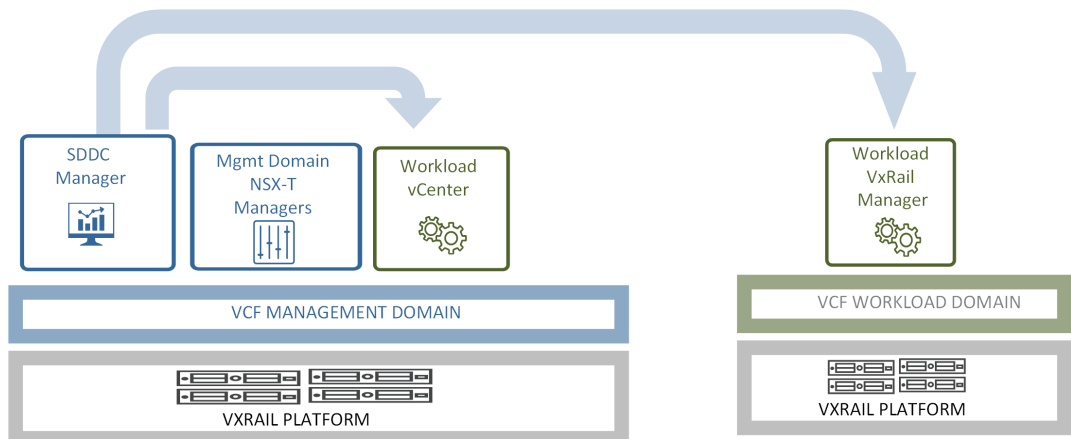
## NSX-T management options for VI workload domain

With each new VI workload domain, there is the option to deploy three new NSX-T management virtual appliances to manage network requirements, or to reuse existing NSX-T managers. The use cases that trigger this decision include:

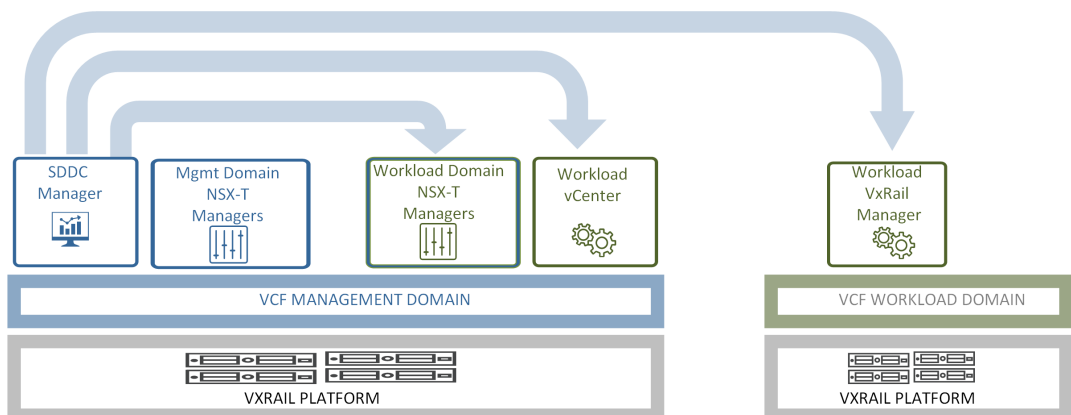
- Test/Development or pre-production workloads that do not use the production network
- Applications in the workload domain have an NSX-T version dependency.
- Applications that require more isolation and security
- Cloud Foundation deployments that have a multitenant workload requirement
- A new VLAN-backed transport zone is required.

If a new set of NSX-T management virtual appliances is planned for the management workload domain, the following points must be considered:

- Each NSX-T manager has a resource reservation of 48 GB of memory.
- NSX-T managers are subject to vSphere HA admission control.
- NSX-T manager virtual appliances can be CPU intensive depending on workload activity.



**Figure 11. VI workload domain deployment using existing NSX-T managers**



**Figure 12. Overview of initial deployment of NSX-T based VI workload by SDDC Manager**

## vRealize software option for VI workload domains

If the use case requirements for any Cloud Foundation VI workload domains include the vRealize software suite, plan for the vRealize management virtual appliances to be deployed in the management workload domain. Additional steps must be undertaken to integrate access into Active Directory and enable vRealize usage.

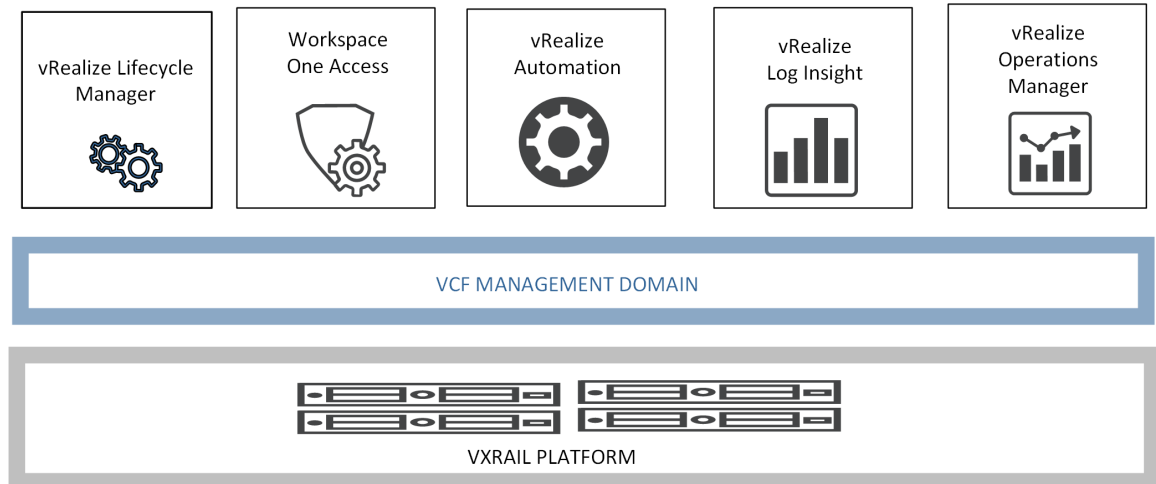


Figure 13. vRealize management virtual appliances in management workload domain

## vSphere with Kubernetes workload domains

SDDC Manager supports VI workload domain configurations that provide the infrastructure foundation required by vSphere with Kubernetes. A VI workload domain configured for Kubernetes transforms the vSphere platform into a platform for running Kubernetes workloads natively on the hypervisor layer. If there is a use case requirement for vSphere with Kubernetes, the following items must be considered:

- All the nodes in the VxRail cluster supporting the VI workload domain must have a vSphere with Kubernetes license.
- An NSX-T edge cluster must be available to the vSphere with Kubernetes workload domain for networking purposes.
- The NSX-T edge cluster requires connectivity upstream using eBGP. Plan on preparing the upstream network for BGP peering and route distribution.
- Additional IP addresses specific for vSphere with Kubernetes will be required when the VI workload domain is configured.
  - Non-routable subnet for pod networking (minimum /22)
  - Non-routable subnet for service IP addresses (minimum /24)
  - Routable subnets for ingress and egress for the NSX-T edge cluster (minimum /27).

# Chapter 5 Data Center and Network Requirements

This chapter presents the following topics:

- Introduction..... 25**
- Data center rack space requirements..... 25**
- Data center networking..... 25**
- Network switch selection..... 26**
- Switch port capacity..... 26**
- Switch port type..... 27**
- Jumbo Frames..... 27**
- Multicast..... 27**
- Border Gateway Protocol..... 27**
- Hardware VTEP for multi-rack deployments..... 28**
- Network services..... 28**
- Storage Services..... 29**



## Introduction

Your data center environment must meet certain requirements to support the deployment of Cloud Foundation on VxRail. Before the product is delivered, Dell Technologies will review these prerequisites with you to ensure compliance.

## Data center rack space requirements

The Cloud Foundation on VxRail platform is a consolidated, self-contained architecture, as there is no requirement for external storage to support workload. Furthermore, there is the expectation that network traffic will migrate onto the virtual network within the VI workload domains, which might free up physical space occupied by excess physical network equipment.

The amount of rack space required for the VxRail nodes depends on the models you select to support your Cloud Foundation on VxRail platform. Dell Technologies will go through a sizing exercise to produce a bill-of-material of the VxRail nodes needed to support your requirements.

**Table 1. VxRail Node Rack Space**

VxRail Model	Rack Units	Power Supply	Plug Type
E-Series	1	750 W, 1100 W	C14
D-Series	1	550 W, 600 W	C14
P-Series	2	1100 W, 1600 W	C14
S-Series	2	1100 W, 1600 W	C14
V-Series	2	2000 W	C20
G-Series	2	2000 W, 2400 W	C20

The amount of rack space required for the supporting physical network depends on the network topology selected for the Cloud Foundation on VxRail deployment. The most common network topology is leaf-spine, so plan on additional rack space for the switches.

## Data center networking

You can either bundle a Dell network infrastructure with your Cloud Foundation on VxRail for a single source solution, or acquire, implement, and configure your own supporting network.


If you choose a network infrastructure based on Dell network switches, your Dell Technologies specialist will work with you to design the network that meets the requirements for your specific Cloud Foundation on VxRail deployment. Regardless of which option you choose, your network infrastructure must support certain requirements for Cloud Foundation on VxRail.

If Dell Technologies is responsible for the configuration of the switches enabling network connectivity for the Cloud Foundation on VxRail infrastructure, those services will be

performed after the networking hardware is installed and cabled in the data center. This work effort includes configuring uplinks to connect to the data center network. The VxRail clusters and VMware Cloud Foundation depend on the supporting network infrastructure to be properly configured before actual deployment.

## Network switch selection

Cloud Foundation on VxRail is supported with either Dell-branded switches or most major enterprise switching products. You should plan on a pair of switches in the top of each data center rack used for Cloud Foundation on VxRail for redundancy purposes. Each VxRail node deployed to support Cloud Foundation on VxRail needs at least one connection into each switch, and possibly more depending on use cases and workload requirements. The first rack which supports the Cloud Foundation on VxRail management domain must integrate with your existing data center infrastructure using Layer 3 network services with a pair of uplinks. If your requirements will expand beyond a single rack, plan on expanding to a leaf-spine network topology. A multi-rack deployment requires additional switch ports be reserved on the top-of-rack switches to connect to the spine layer.



**Figure 14. Single site Cloud Foundation on VxRail deployment across two racks**

## Switch port capacity

If deployed into a 25Gb network, Cloud Foundation for VxRail supports either two or four Ethernet ports per node to be reserved for Cloud Foundation on VxRail networking purposes. Instead, if the nodes are plugged into a 10Gb network, either two, four, or six Ethernet ports per node can be reserved for Cloud Foundation on VxRail networking. The total number of ports required on the adjacent switches is depending on the number of VxRail nodes to be deployed to support the Cloud Foundation working, and the number of ports selected per node to reserve for this purpose.

## Switch port type

VxRail nodes can support either SFP+ or RJ45 for 10 Gb Ethernet network connections, or SFP28 for 25 Gb Ethernet connections. The ports on the physical switch supporting your Cloud Foundation on VxRail cloud platform must match the network type on the VxRail nodes.

## Jumbo Frames

NSX-T depends on extending the standard Ethernet frame beyond the default 1500 bytes to support the tunneling of virtual machine traffic over the physical network in Cloud Foundation on VxRail. NSX-T depends on the GENEVE (GENeric Network Virtualization) standard, which requires an MTU size of 1600 or higher to support the encapsulation of virtual machine traffic and provide the additional required header space. The physical network supporting Cloud Foundation on VxRail must support the ability to increase the MTU size to support tunneling.



**Figure 15. NSX-T extended frames**

## Multicast


VxRail depends on IPV6 multicasting to support device discovery during the cluster build operation and node expansion. The IPV6 is a private network that is isolated within the switches supporting the VxRail cluster in order to limit the impact on the data center network. Enabling MLD snooping and snooping querier is recommended on the physical switch to optimize multicast traffic.

If your network infrastructure does not support IPV6 multicast, or if your network policy does not allow IPV6 multicast, you can choose to have the nodes ingested manually during the cluster build operation and during node expansion. This option is supported with VMware Cloud Foundation 4.2 or later, as this option depends on VxRail version 7.0.130 or later.

## Border Gateway Protocol

Cloud Foundation on VxRail leverages NSX-T edge gateways to serve as the boundary point between the physical and virtual networks. This gateway is the passageway for traffic external to the data center to communicate with the virtual workload running on Cloud Foundation on VxRail. To enable routing between the physical and virtual networks,

the upstream physical network must support Border Gateway Protocol. The NSX-T edge gateways require BGP adjacency to peer with upstream routing services.




**Figure 16. Physical and virtual route peering with Border Gateway Protocol**

Verify that Equal Cost Multi-Path (ECMP) routing is supported on the switches to optimize traffic on these pathways.

## Hardware VTEP for multi-rack deployments

Dell Technologies recommends selecting network switches that support hardware-based Virtual Tunnel Endpoints (VTEP) for multi-rack deployments. This feature is beneficial for customers expecting to deploy VxRail clusters over multiple racks, and do not want to extend the Layer 2 networks across racks. This feature is also beneficial in avoiding a rack being single point of failure, as the virtual machines in the management workload domain can migrate between racks without the need to change the IP address.

The feature supports the bridging of Layer 2 network traffic VxRail nodes in different racks through packet encapsulation and decapsulation on a Layer 3 overlay network. This feature optimizes VxRail network traffic across racks in a multi-rack cluster by eliminating the need to route through upstream routing services.



**Figure 17. VTEP tunnel network supporting multi-rack deployment**

## Network services

The network services listed in this section are required in the host data center for your Cloud Foundation on VxRail deployment. These services must be enabled in the data center planned for the Cloud Foundation on VxRail deployment, and configured with the settings required for your specific deployment.

- Domain Name Service (DNS) – You must enter forward and reverse DNS entries for every VxRail node. In addition, the virtual components used for the management of the VxRail clusters and the Cloud Foundation domains also require forward and reverse DNS entries.
- Network Time Protocol (NTP)
- Dynamic Host Configuration Protocol (DHCP) – IP addresses are assigned to each host in the VxRail cluster to serve as the endpoints for NSX-T inbound/outbound traffic at the edge. The IP addresses can either be assigned manually or dynamically using DHCP. A DHCP server must be deployed in the host data center, and be pre-populated with the IP addresses to be assigned to the host endpoints.

The following network services are optional, but recommended:

- Simple Message Transfer Protocol (SMTP)
- Certificate Authority (CA) – The Certificate Authority must be able to ingest a Certificate Signing Request from the SDDC components, and issue a signed certificate. Cloud Foundation on VxRail supports Microsoft Windows Enterprise Certificate Authority. The domain controller must be configured with the Certificate Authority Service and the Certificate Authority Web Enrollment roles.
- SFTP Server – The SFTP server supports backups of NSX-T Data Center instances and SDDC Manager.

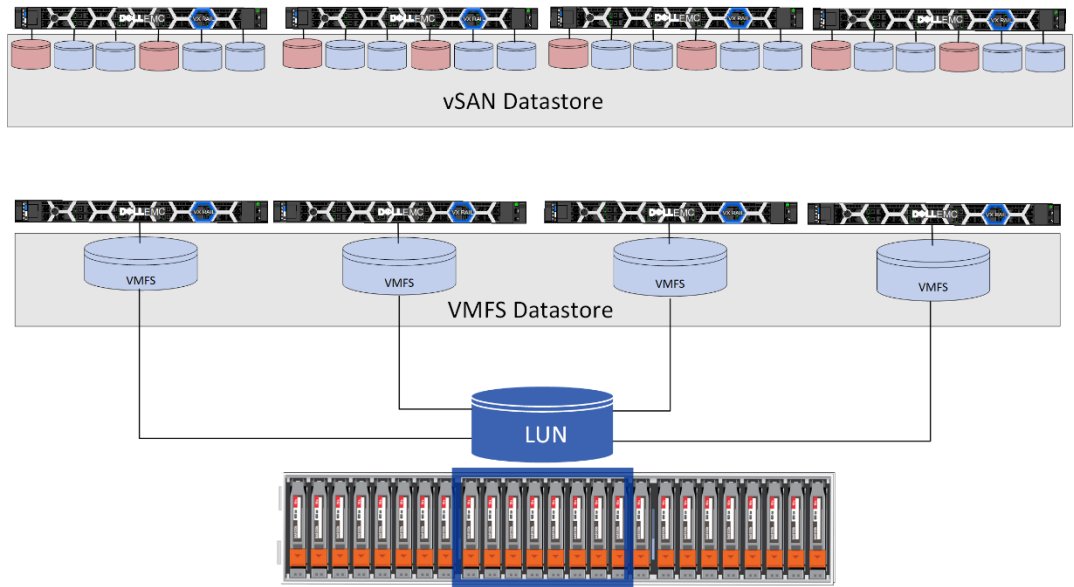
The following network services may be required, depending on the use cases targeted for Cloud Foundation on VxRail VI workload domains:

- Active Directory - Cloud Foundation on VxRail uses Active Directory service accounts for application-to-application communications.

## Storage Services

For the VI workload domains you deploy with your Cloud Foundation on VxRail instance, the VxRail clusters support two options for provisioning primary storage resources to support virtual machine operations.

- One option is to purchase VxRail nodes with installed disk drives, and then use the disk drives on the VxRail nodes for the formation of a vSAN datastore.
- The second option is to purchase VxRail nodes with no installed disk drives, then use a storage array in your data center to provide a LUN to the nodes to form a VMFS datastore.



**Figure 18. vSAN and VMFS datastore options for supporting VI workload domains**

If you choose to deploy one or more clusters for your Cloud Foundation on VxRail with no installed disk drives, then your data center must be prepared to provide storage resources to these nodes from a Fibre Channel storage array over a Fibre Channel network. Consult with your Dell Technologies account team for the storage arrays that are compatible with Cloud Foundation on VxRail.

# Chapter 6 Cloud Foundation on VxRail High-Level Design

This chapter presents the following topics:

**Introduction..... 32**

**Consolidated vs. standard architecture ..... 32**

**Site locations ..... 33**

**Application availability..... 33**

**VxRail Cluster Network Planning..... 34**

## Introduction

There are many factors to consider when deciding on the use cases for a Cloud Foundation on the VxRail platform. Because of the adaptive architecture designed into Cloud Foundation on VxRail, the business and operational requirements vary from one situation to another.

## Consolidated vs. standard architecture

In a standard deployment, the Cloud Foundation management workload domain consists of workloads supporting the virtual infrastructure, cloud operations, cloud automation, business continuity, and security and compliance components for the SDDC. Using SDDC Manager, separate workload domains are allocated to tenant or containerized workloads. In a consolidated architecture, the Cloud Foundation management workload domain runs both the management workloads and tenant workloads.

There are limitations to the consolidated architecture model that must be considered that will impact this decision-making process.


- The decision to deploy a consolidated architecture must be made at the time of deployment, as a consolidated architecture cannot be converted to a standard architecture.
- Use cases that require a VI workload domain to be configured to meet specific application requirements cannot run on a consolidated architecture. The singular management workload domain cannot be tailored to support management functionality and these use cases. If your plans include applications that require a specialized VI workload domain, plan to deploy a standard architecture.
- Life-cycle management can be applied to individual VI workload domains in a standard architecture. If the applications targeted for Cloud Foundation on VxRail have strict dependencies on the underlying platform, consolidated architecture is not an option.
- Autonomous licensing can be used in a standard architecture, where licensing can be applied to individual VI workload domains. In a consolidated architecture, this is not an option.
- Scalability in a consolidated architecture has less flexibility than a standard architecture. Expansion is limited to the underlying VxRail cluster or clusters supporting the single management workload domain in a consolidated architecture, as all resources are shared. The minimum node count is eight for a standard architecture. Dell-Technologies recommends that any workload requirements that will require eight or more nodes should plan for a deployment using a standard architecture.
- If a VxRail cluster was built using two Ethernet ports, consolidating VxRail traffic and NSX-T traffic, additional nodes added to a cluster are limited to two Ethernet ports being used for Cloud Foundation for VxRail.
- VxRail nodes of differing network port speeds cannot be mixed in a VxRail cluster. If the workload is constrained due to network bandwidth and/or throughput, expanding the cluster with nodes that support higher port speeds is not an option.



## Site locations

The requirement for multiple site locations for Cloud Foundation on VxRail deployments affect the overall high-level design. Factors such as distance and the quality of the network between sites must be considered.

- The management of multiple sites from a single management instance has network guidelines that must be met for supportability.
- The migration of virtual machines between sites using vMotion has network guidelines that must be considered for supportability.
- Both private and public WAN connections are supported for Cloud Foundation on VxRail provided compatibility requirements are met.
- Considerations for SD-WAN are merited for deployments spanning across multiple remote sites.



**Figure 19. Support for remote sites over Internet and private networks**

Primary and secondary active WAN links are strongly recommended. Without network redundancy, you might encounter conditions that can lead to failure states such as a two-failure state condition, which can result in unrecoverable virtual machines and application failure.

## Application availability

Assess and categorize the availability requirements for the sets of application planned for deployment on Cloud Foundation on VxRail.

- If a primary objective is protection from a site-level failure with no loss of service, stretching the vSAN datastore in the VxRail cluster between sites is an option. In this configuration, synchronous I/O is supported for the virtual machines operating in the Cloud Foundation domains where a VxRail vSAN stretched cluster is present. However, there are strict latency requirements for the network between the sites, and this option requires a third site for a 'witness' to monitor the stretched vSAN datastore.
- Determine the operational recovery and disaster recovery objectives of the application sets planned for Cloud Foundation on VxRail. Certain application sets can then be placed in VI workload domains configured to support these objectives.

## VxRail cluster network planning

There are options and design decisions to be considered on the integration of the VxRail cluster physical and virtual networks with your data center networks. The decisions made regarding VxRail networking cannot be easily modified after the cluster is deployed and supporting Cloud Foundation, and should be decided before actual VxRail cluster deployment.

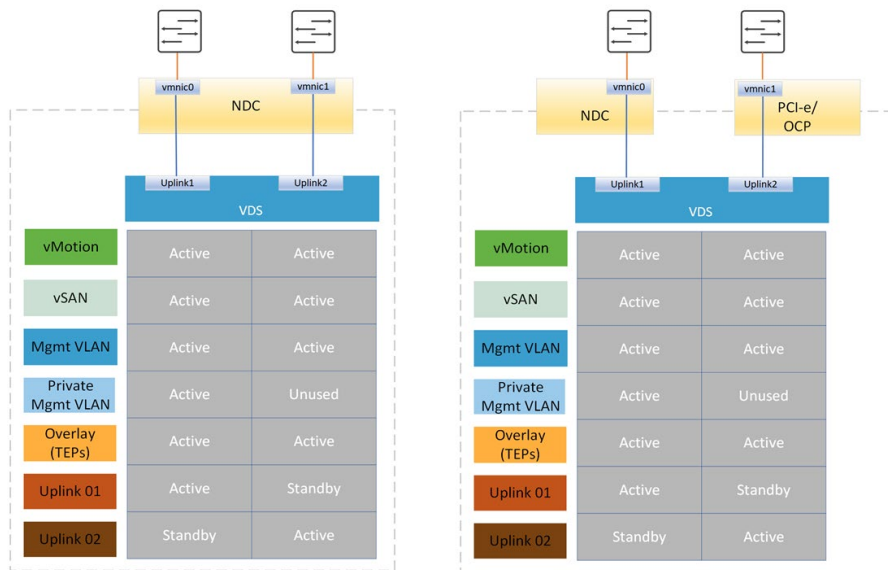
### VxRail network profile selection

Each VxRail node has an on-board, integrated network card. Depending on the VxRail models selected and supported options, the on-board Ethernet ports can be configured as either 2x10Gb, 4x10Gb, 2x25Gb, or 4x10Gb. You can choose to support your Cloud Foundation on VxRail workload using only the on-board Ethernet ports, or deploy with both on-board Ethernet ports and with Ethernet ports from expansion adapter cards. If NIC-level redundancy is a business requirement, a decision can be made to install optional Ethernet adapter cards into each VxRail node for this purpose. Depending on the VxRail nodes selected for the cluster, the adapter cards can support 10 Gb, 25 Gb and 100 expansion ports.

VxRail supports both predefined network profiles and custom network profiles when deploying the cluster to support Cloud Foundation VI workload domains. The best practice is to select the network profile that aligns with the number of on-board ports and expansion ports being selected per node to support Cloud Foundation on VxRail networking. This ensures that VxRail and CloudBuilder will configure the supporting virtual networks by following the guidance designed into these network profiles.

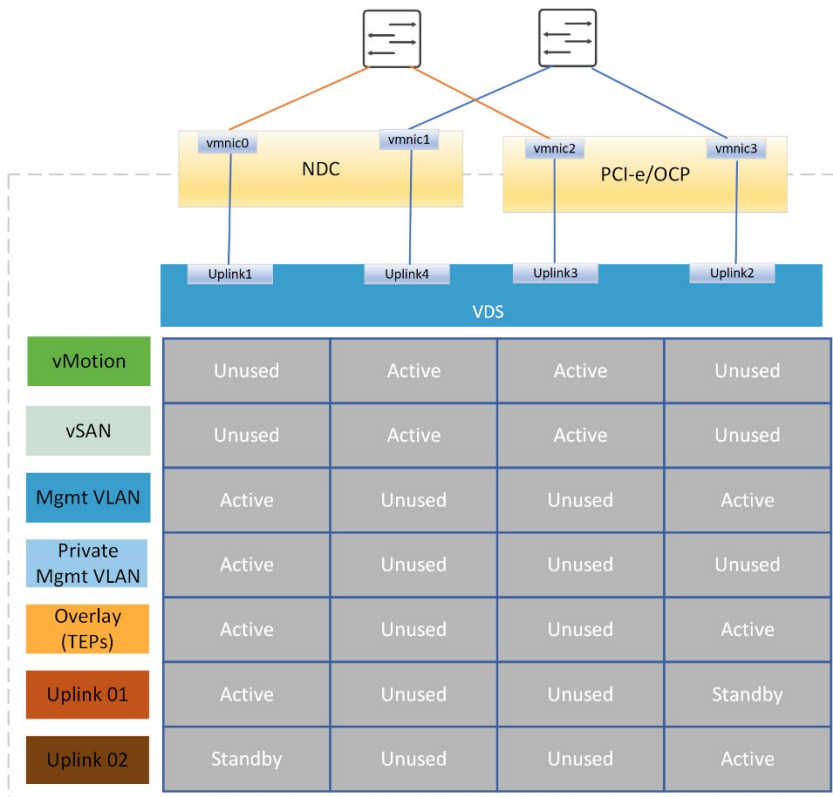
If VCF on VxRail networking will be configured on two node ports, a decision must be made whether to add an expansion card into each VxRail node to eliminate the on-board port as a single point of failure. If only the on-board ports are present, the first two ports on each VxRail node will be reserved to support VCF on VxRail networking using a predefined network profile. If both on-board and expansion ports are present, a custom network profile can be configured, with the option to select an on-board port and expansion port to reserve for VCF on VxRail networking.

In both two-port instances, the NSX and VxRail networks are configured to share the bandwidth capacity of the two Ethernet ports.



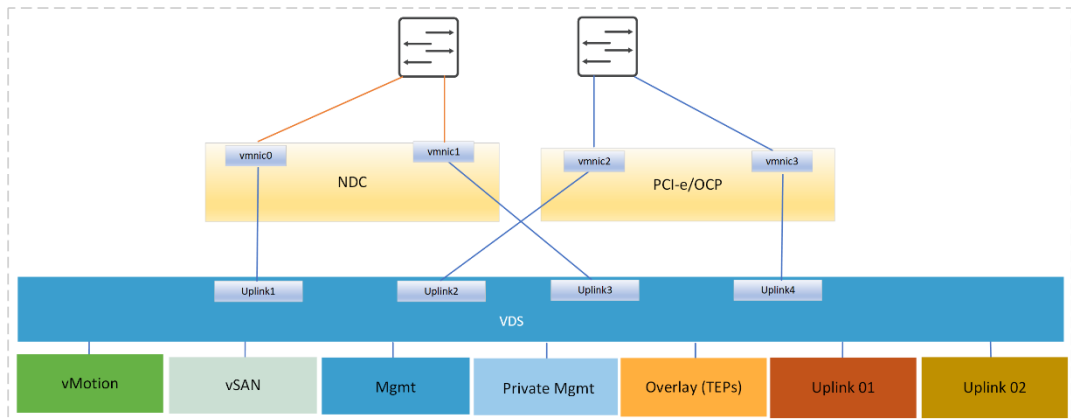
**Figure 20. 2 ports reserved for VCF on VxRail networking**

If VCF on VxRail networking will be configured with four ports, either all four on-board ports can be used, or the workload can be spread across on-board and expansion ports. The option to use only on-board ports uses a predefined network profile, with automatic assignment of the VMnics to the uplinks. Configuring with both on-board and expansion ports is preferred because it enables resiliency across the node devices and across the pair of switches.



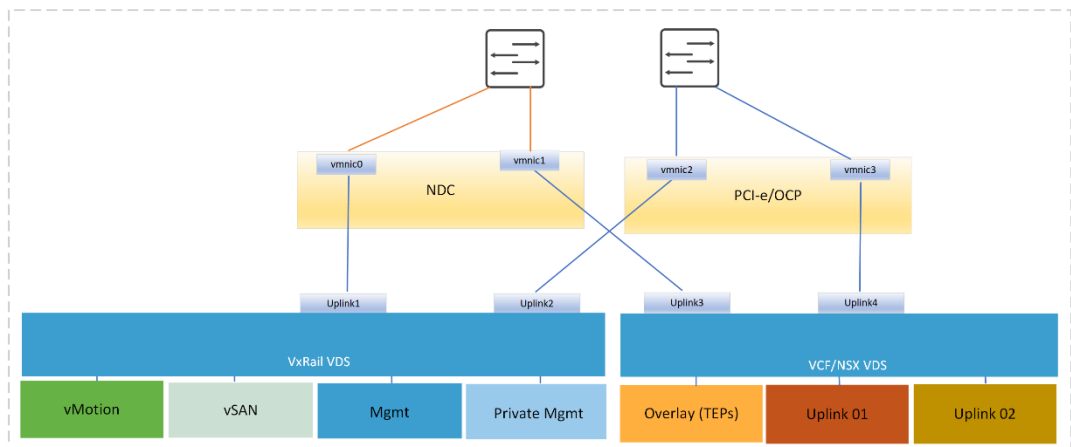
**Figure 21. 4 ports reserved for VCF on VxRail networking in pre-defined network profile**

Deploying the VxRail cluster using NDC-based ports and PCIe/OCP-based ports with a custom network profile offers flexibility for network assignments. The best practice with a custom network profile to enable network resiliency is to plug the NDC-based ports into one switch, and then plug the PCIe/OCP-based ports into the second switch. Then, use the custom network profile feature to map the uplinks from the virtual distributed switch to the VMnics to spread the workload across both switches and also protect from a single point of failure.



**Figure 22. 4 ports and 1 VDS for VCF on VxRail networking using a custom network profile**

Using a custom network profile also eases the assignment of uplinks if the Cloud Foundation on VxRail instance is deployed with separate virtual distributed switches to segment VxRail network traffic and VCF/NSX network traffic.




**Figure 23. 4 ports and 2 VDS for VCF on VxRail networking using a custom network profile**

## VxRail virtual network planning

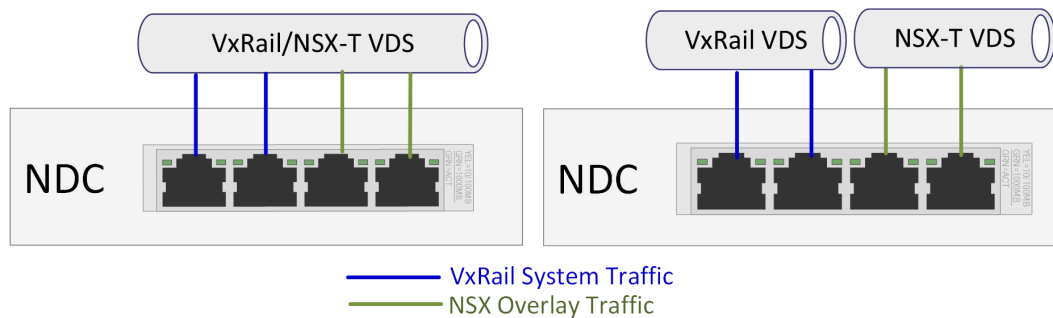
Cloud Foundation on VxRail supports the physical segmentation of VxRail and Cloud Foundation network traffic onto dedicated network ports, and onto separate, dedicated virtual distributed switches.

If the VxRail nodes are configured with two Ethernet ports, all the VxRail network traffic and Cloud Foundation/NSX-T traffic is consolidated onto the two ports. A second virtual distributed switch is not supported for the two-port connectivity option, so all VxRail and Cloud Foundation/NSX-T traffic flows through a single virtual distributed switch.



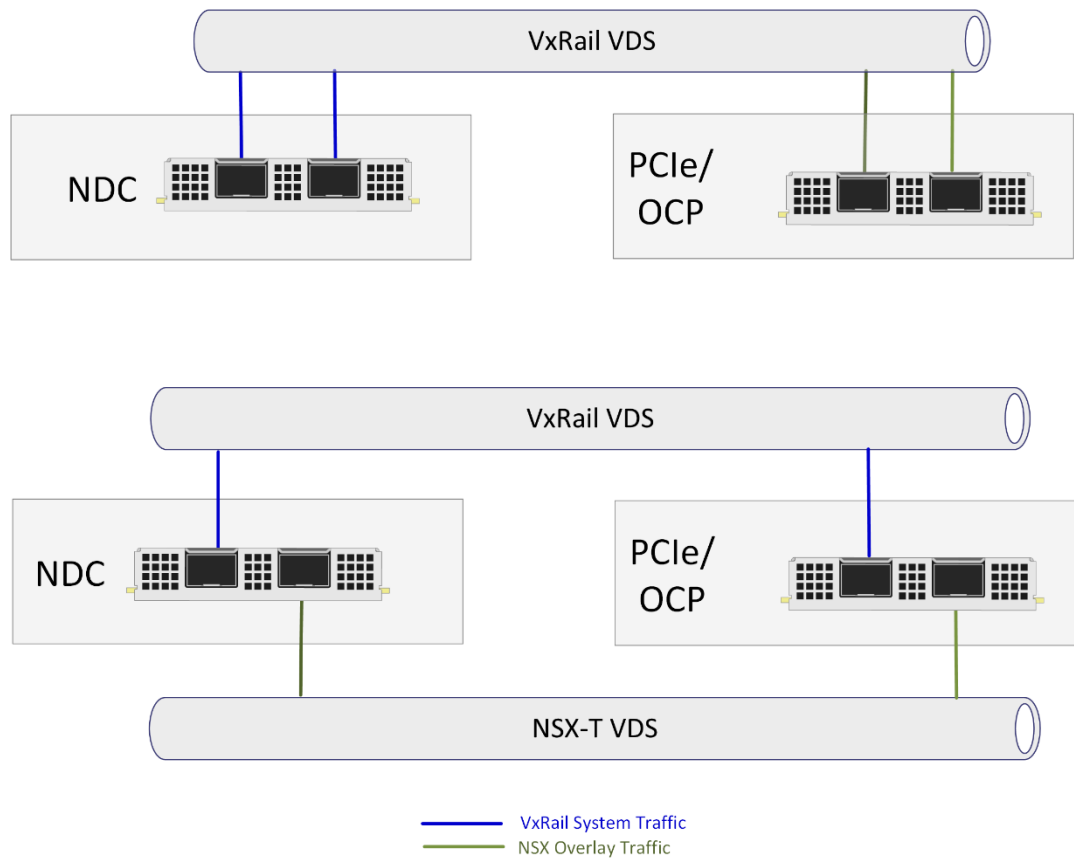
**Figure 24. Connectivity options for VxRail nodes with two on-board RJ45 ports or two on-board SFP+ ports**

With the option of deploying four on-board ports, the vMotion and vSAN network traffic supporting VxRail are positioned on the second port, and the Cloud Foundation/NSX-T traffic is assigned to the last two ports. With this network profile, a second virtual distributed switch can be deployed to isolate the VxRail network traffic on the first virtual distributed switch, and the Cloud Foundation/NSX-T traffic on the second virtual distributed switch.



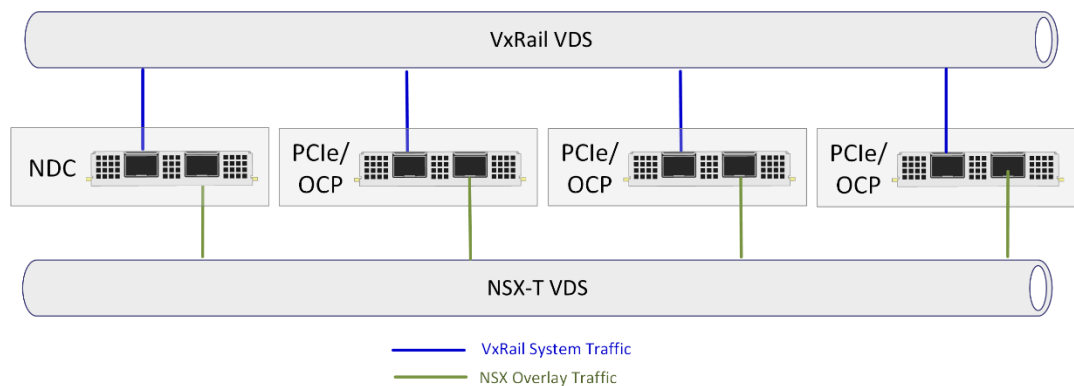
**Figure 25. Connectivity options for VxRail nodes with four on-board ports**

For the four-port option using both on-board ports and expansion ports from each VxRail node, a decision can be made to direct all Cloud Foundation on VxRail traffic onto a single virtual distributed switch, or to redirect the NSX-T network traffic onto a separate virtual distributed switch.



**Figure 26. Connectivity options for VxRail nodes with two on-board ports and two expansion ports**

For planned workloads that have very high bandwidth requirements, up to eight Ethernet ports can be used across the on-board and expansion cards. The VxRail network traffic is spread across four ports, and Cloud Foundation/NSX-T network traffic is spread across the other four ports.



**Figure 27. Sample connectivity option for VxRail nodes with two on-board ports and 6 expansion ports**

The reservation and assignment of the physical ports on the VxRail nodes to support Cloud Foundation on VxRail networking is performed during the initial deployment the VxRail cluster. Dell Technologies recommends that careful consideration be taken to

ensure that sufficient network capacity is built into the overall design to support planned workloads. If possible, Dell Technologies recommends an overcapacity of physical networking resources to support future workload growth.

# Chapter 7 Cloud Foundation on VxRail Workload Planning

This chapter presents the following topics:

- Introduction.....41**
- Determine use cases for Cloud Foundation VI workload domain ..... 41**
- Deciding on single-site VxRail cluster or stretched cluster ..... 41**
- Planning the Management workload domain resource requirements..... 43**
- Planning the VI workload domain resource requirements ..... 44**
- Sizing the Cloud Foundation domains ..... 44**



## Introduction

The primary building block for compute resources for Cloud Foundation on VxRail is the server node. VxRail leverages the Dell PowerEdge server products as the foundation for a cluster. A VxRail cluster can scale to a maximum of 64 nodes. The first VxRail cluster deployed is always used to support the management workload domain, which requires a minimum of four nodes. VxRail supports a wide variety of server physical configurations, with flexibility on CPU model, CPU quantity and speed, RAM capacity, physical storage capacity, and network port quantity and speed.

For more details, see the [VxRail Series Specification Sheet](#).

The mixing of different server models in a single cluster is supported because VxRail views the individual server node as a static pool of compute resources. This offers additional flexibility to start the initial configuration to meet a predefined baseline, and adapt and expand as necessary for changing workload requirements.

## Determine use cases for Cloud Foundation VI workload domain

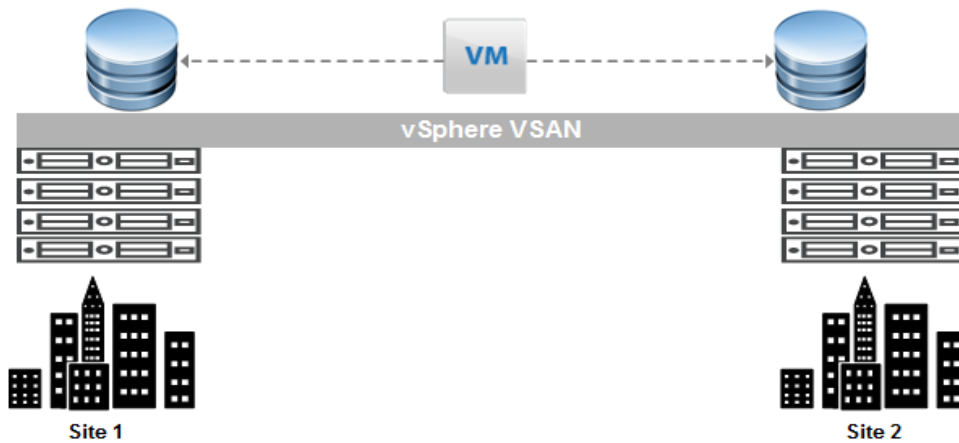
Before performing an overall sizing effort for Cloud Foundation on VxRail, decisions must be made on the rules and criteria for the creation of VI workload domains within your business. The criteria can be for a range of reasons, for instance:

- Logical grouping of applications or application sets for streamlined interconnectivity
- Ease of assigning and controlling a pool of IT resources to internal organizations
- Managing multiple sites from a single management entity

Each of these criteria impact the resources that are required to support the workload planned for each domain. If plans include special use cases, such as the deployment of the vRealize suite, the overhead required to support these product suites must be considered with the sizing effort.

## Deciding on single-site VxRail cluster or stretched cluster

If you deploy VxRail stretched clusters instead of a single-site cluster in order to meet availability requirements, be aware of the impact this decision has in planning the workload. A VxRail stretched cluster requires double the number of VxRail nodes to support any planned workload, as each site must be able to support that workload in the event of a site failure.



**Figure 28. Impact of VxRail stretched cluster on virtual machine workload**


The reason for this is because every write operation by a virtual machine is performed on the vSAN datastore at both sites. Therefore, the size of the physical storage positioned to support any planned workload must be doubled.

In addition, there are a couple of other items of importance to consider with using stretched cluster to support your Cloud Foundation workloads:

- You must use DHCP to assign IP addresses to each VxRail node to support the overlay network required by NSX-T if your plans include VxRail stretched clusters to support the Cloud Foundation workload domains. Assigning static IP addresses to support the overlay network is not supported with stretched clusters.
- A feature in vSAN called 'HCI Mesh' support the sharing of vSAN datastore resources between VxRail clusters. This is accomplished by a VxRail cluster mounting a remote vSAN datastore and allocating those resources to local virtual machines. The 'HCI Mesh' feature is not supported with stretched clusters.

### Deciding on local vs. remote workload domains

If you require support for workloads at multiple locations, you can choose a single-site deployment with a single point of management at each location, or have a centralized site designated to manage workloads at the local site and at remote sites.



**Figure 29. Remote workload domain WAN requirements**

The decision whether to pursue a centralized site model must meet the following considerations:

- The number of nodes at the remote site to support the VI workload domain is limited to four nodes.
- The minimum network bandwidth between sites is 10 Mb per second.
- The maximum latency between sites is 50-millisecond round-trip time.
- Redundant WAN links between sites are strongly recommended to eliminate the WAN as a single point of failure.

### Deciding on single region vs. multi-region

If your plans include expansion of Cloud Foundation on VxRail instances across regions using NSX-T Federation, consider both the latency and bandwidth of the physical network between regions. You can begin by deploying Cloud Foundation on VxRail in a single region, and decide to scale out into a multi-region architecture as a future expansion.

The maximum round-trip latency supported for NSX-T Federation is 150 milliseconds. Ideally, your latency between regions should fall within the performance threshold requirements set for the applications that are interconnected across regions. In addition, the network between sites must be configured to support NSX-T. The MTU size of the connecting network must be 1600 or larger.

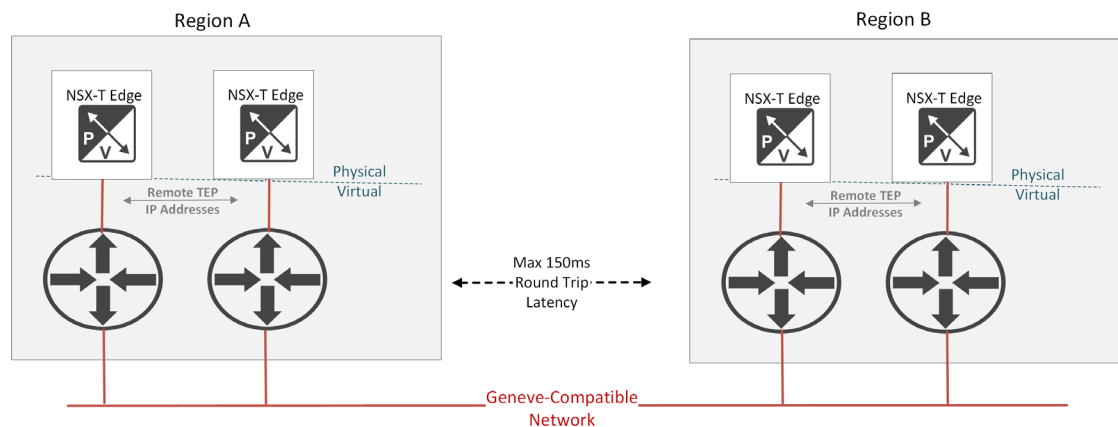


Figure 30. RTT latency for NSX-T Federation with Cloud Foundation on VxRail

## Planning the management workload domain resource requirements

At initial deployment of Cloud Foundation on VxRail, Cloud Builder deploys an initial core set of virtual appliances on the management domain to support Cloud Foundation on VxRail management.

This initial set of core virtual machines from Cloud Builder provides a baseline on the resources required to bring up the management workload domain and to begin deploying VI workload domains on a standard architecture. With a consolidated architecture, guest virtual machines are deployed on the single workload domain, sharing resources with the Cloud Foundation on VxRail management components. With a standard architecture, guest virtual machines can only be deployed on separate VI workload domains. For each

VI workload domain that is created, additional virtual machines are deployed for management purposes. Therefore, in order to get more accurate sizing guidelines for the management workload domain and VI workload domains in a standard architecture, the planned use cases for the VI workload domains must be understood before initial deployment. Then, a resource consumption assessment is completed for both the planned management workload domain and the VI workload domains, which are used to determine the required size of the underlying VxRail platform for all the planned domains.

The tables in [Appendix A: Cloud Foundation on VxRail checklist](#) can be used to provide estimates of the minimum sizing requirements for the management components, based on planned use cases.

## Planning the VI workload domain resource requirements

At least one VI workload domain must be created to support guest virtual machines for a standard architecture, and at least one VxRail cluster of any supported size and configuration must be used as the resource foundation for a VI workload domain. A VxRail cluster that is assigned to support the workload of a Cloud Foundation domain is dedicated to that domain, and its resources cannot be shared with other Cloud Foundation domains.

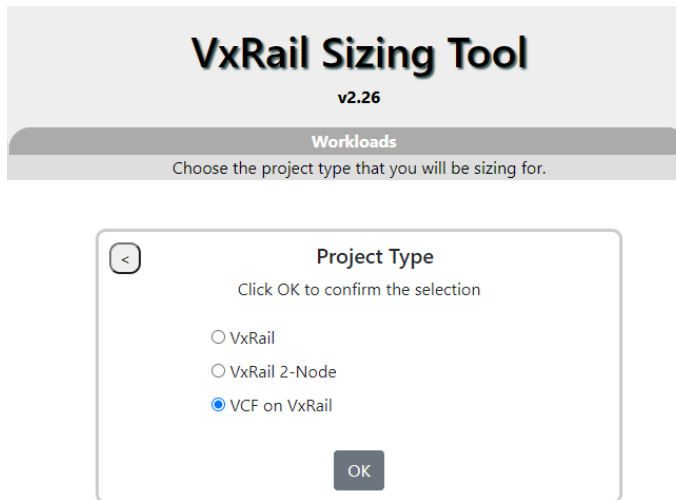
For each VI workload domain that is created, SDDC Manager deploys a vCenter virtual machine in the management workload domain. For VI workload domains that are not sharing existing NSX-T management resources, a new set of NSX-T management virtual appliances are deployed in the management workload domain. Depending on the use case, additional virtual machines might need to be deployed to support those specific applications.

Use the tables in [Appendix B: Cloud Foundation on VxRail footprints for sizing](#) for an understanding of the baseline sizing at the creation of the management workload domain, and for estimating the sizing requirements for additional components based on planned use cases.

## Sizing the Cloud Foundation domains

The best practice for the resource sizing effort of the Cloud Foundation domains is to consider initial baseline of resources required for overall management based on use cases, and then calculate the additional resources needed for guest virtual machines.

Dell Technologies uses a sizing tool to calculate the workload resource requirements for the Cloud Foundation domains. Dell Technologies will conduct a sizing exercise to determine the pools of resources required to satisfy VI workload domain demands and their service level objectives at an optimal cost.



**VxRail Sizing Tool**  
v2.26

**Workloads**  
Choose the project type that you will be sizing for.

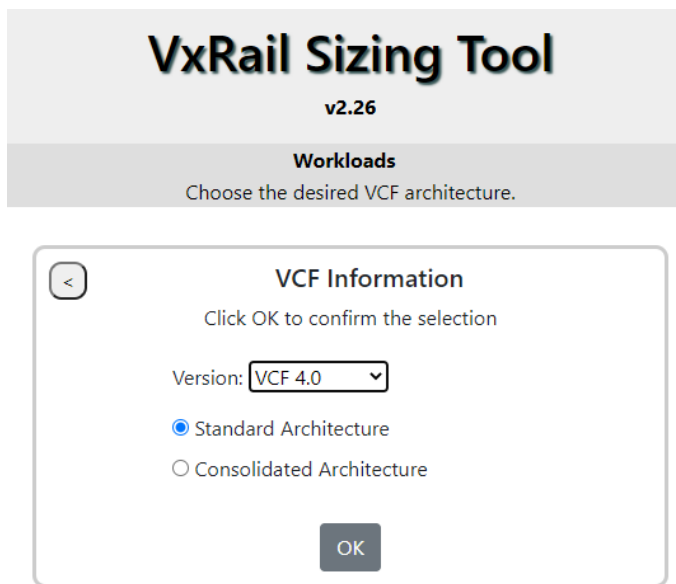
**Project Type**  
Click OK to confirm the selection

☐ VxRail  
☐ VxRail 2-Node  
☒ VCF on VxRail

OK

**Figure 31. VxRail online sizing tool**

The VxRail sizing tool performs calculations on one Cloud Foundation domain at a time. Therefore, the resources overhead required for management of each of the VI workload domains can be factored into the sizing effort for the management workload domain.



**VxRail Sizing Tool**  
v2.26

**Workloads**  
Choose the desired VCF architecture.

**VCF Information**  
Click OK to confirm the selection

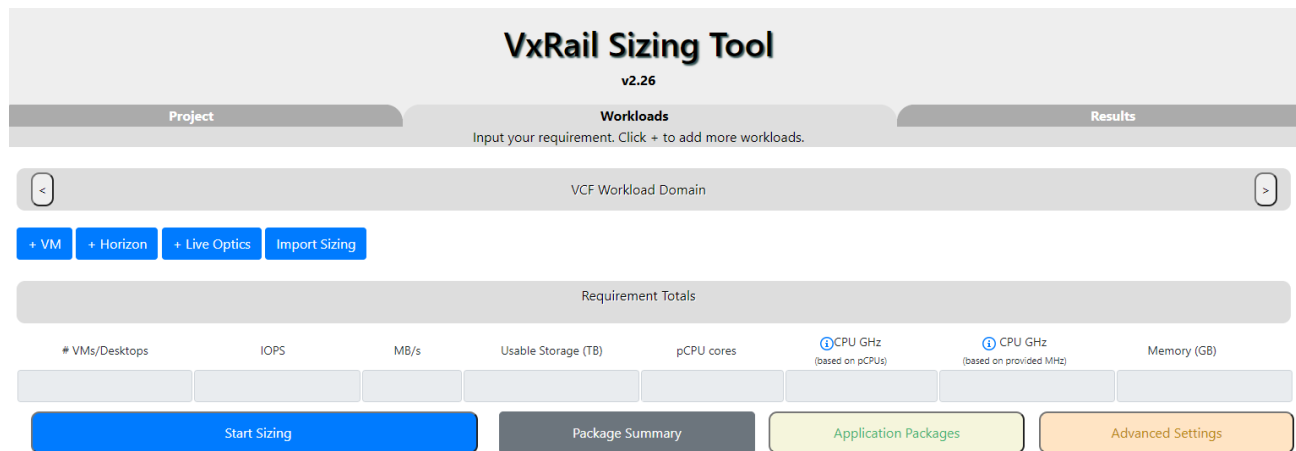
Version:

☒ Standard Architecture  
☐ Consolidated Architecture

OK

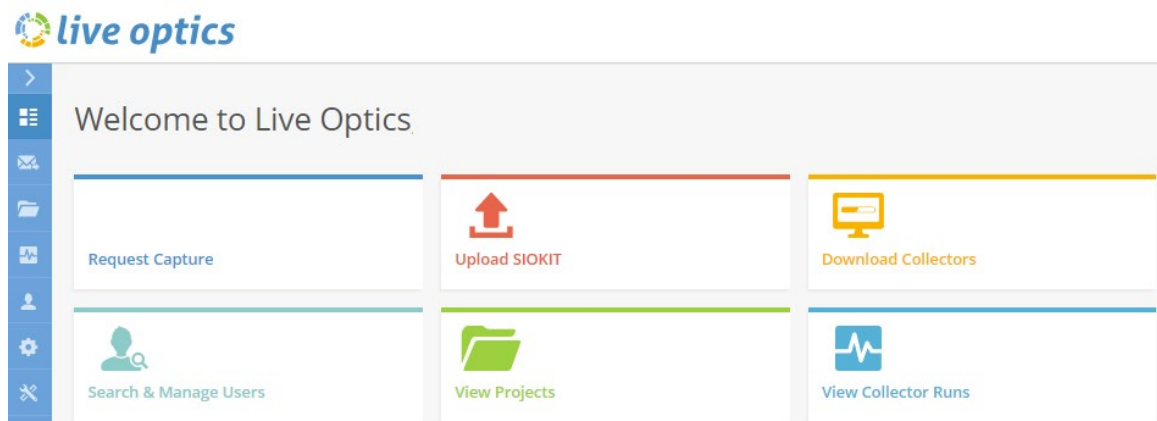
**Figure 32. Selecting options for architecture in VxRail sizing tool**

At least one workload domain is also included in the sizing effort. It is important to understand the applications planned for each respective workload domain to ensure accurate sizing.



**Figure 33. Sizing VI workload domain in VxRail sizing tool**

For calculating resource requirements for guest virtual machines, the VxRail sizing tool accepts sizing data either through manual entry or by downloading metrics from a collector tool. For the most accurate sizing calculations, Dell Technologies' best practice is to use a collector tool for guest virtual machine resource requirements. Dell Technologies uses [LiveOptics](#) data collection for this purpose. The capture from the data collector can then be input directly into the VxRail sizing tool to produce the sizing report for each VxRail cluster.



**Figure 34. LiveOptics main dashboard**

The VxRail sizing tool also supports reference workloads. A reference workload is a synthetic workload that attempts to represent real-life workloads. Select the reference workloads that best represent what is planned for a given VI workload domain to enable proper sizing.

The VxRail sizing tool performs its calculations using virtual machine profiles, and number of virtual machines that fit for each profile. Note that more than one profile can be defined for the same sizing exercise.

For best results, define the following metrics for each virtual machine profile:

- A reference workload
- The expected I/O activity per VM

- The usable storage capacity per VM
- The number of vCPUs or the amount of CPU in MHz per VM
- The amount of memory per VM

Dell Technologies will include the sizing metrics entered for each virtual infrastructure domain, and then perform the sizing analysis. When the settings are finalized, the resulting report from the VxRail sizing tool will show the required node count for a VxRail model and the HW characteristics for each node to meet the overall workload requirements.

## Chapter 8 Application Dependencies and Routing Decisions

This chapter presents the following topic:

**Understanding connection dependencies .....49**



## Understanding connection dependencies

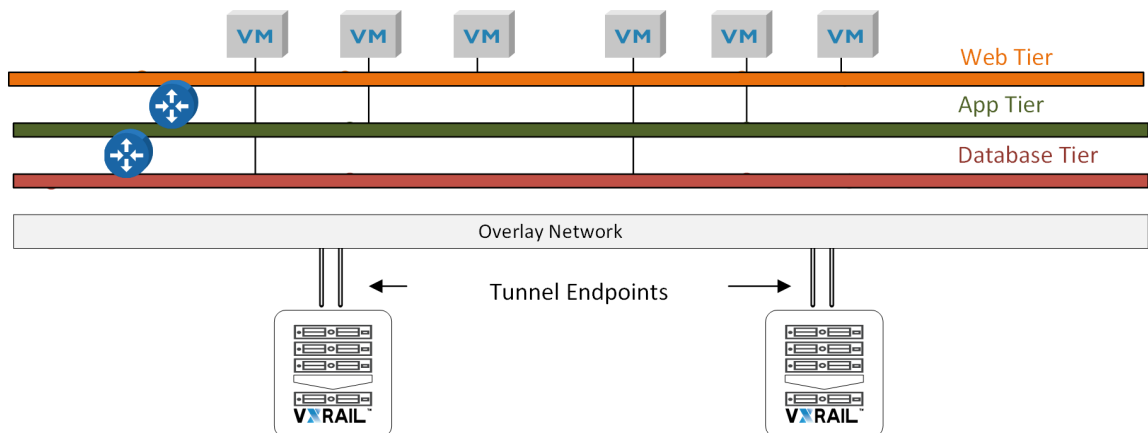
Understanding the connection dependencies between the applications planned for Cloud Foundation on VxRail will streamline the high-level network design process and improve its effectiveness. It will also simplify the final decisions to be made on the placement of application sets on specific VI workload domains. To reduce the routing workload at the physical network layer and optimize the efficiency of the virtual networks, an assessment of the routing maps and dependencies for the sets of applications targeted for Cloud Foundation on VxRail is recommended.

When applications running on different subnets need to connect with each other, the network traffic is directed to a router, which then decides the path the network traffic takes to communicate. For environments that do not use VMware NSX-T, this means that the virtual machine network traffic must travel upstream out of the virtual network layer, where the routing decisions are made at the physical network layer.

Cloud Foundation on VxRail leverages NSX-T to enable support for routing in the virtual networks on the VI workload domains. This means that the defined network paths can be in different locations:

- Between applications within a Cloud Foundation VI workload domain
- Between applications in different Cloud Foundation VI workload domains
- Connected to external applications outside of a Cloud Foundation VI workload domain

Your sets of applications are likely separated by factors such as function or end-user accessibility (such as web tiers and database tiers). Within Cloud Foundation on VxRail VI workload domains, you might want to segment those application sets for network isolation, so that end-users can only access, for instance, the web tier network. You might also want flexibility so that the applications in each isolated network are not tied to a static pool of resources, or a static location.

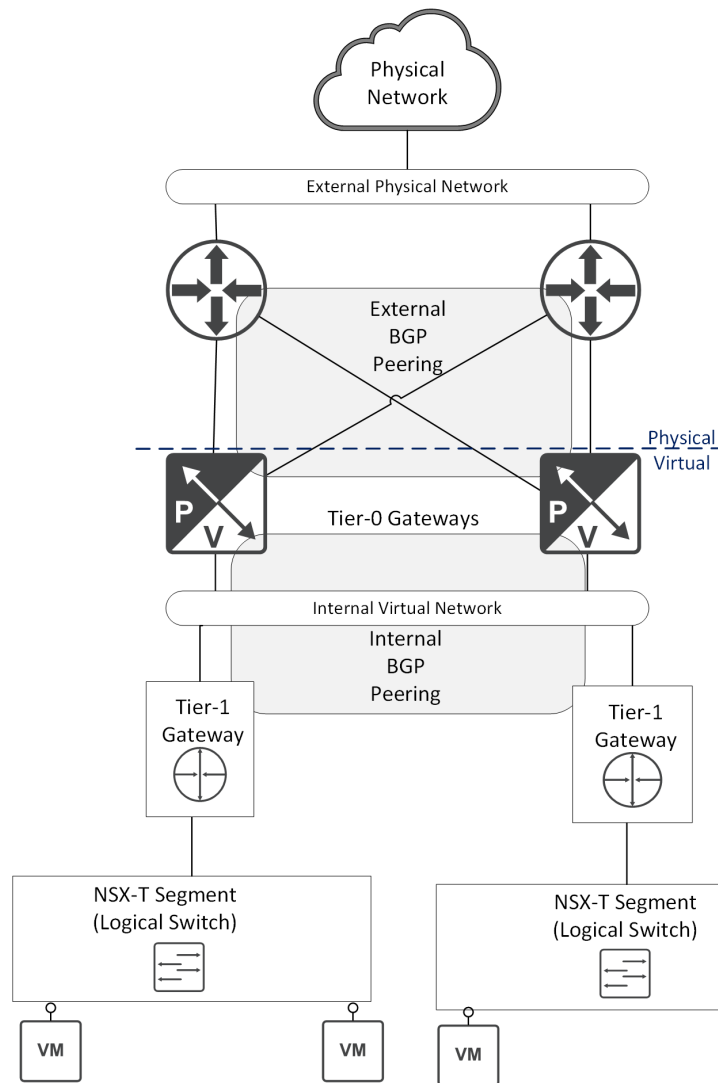


**Figure 35. Three-tier application on separate virtual networks connected with virtual routers**

Virtual machines deployed in a Cloud Foundation VI workload domain connect to port groups on a virtual switch in their Cloud Foundation VI workload domain. Each port group on a virtual switch is assigned a unique 'virtual LAN' (VLAN) identifier, and the traffic on a VLAN is logically isolated from the network traffic on another VLAN. If a virtual machine needs to connect with another virtual machine on the same VLAN but not on the same virtual switch, an extended network is used. Using GENEVE for NSX-T supports extending the non-routable VLAN-based network over a routable network. The traffic from one virtual machine flows through the virtual switch on a host and up a Tunnel Endpoint, over the physical network, and down through the Tunnel Endpoint on the second host, and is delivered to the second virtual machine.

An extended physical network supports accessibility between the virtual networks in the VI workload domains. This configuration forms an extended logical switch, or segment within NSX-T, across the individual virtual switches in the VxRail clusters. A virtual router, known as a Tier-1 gateway, is deployed within NSX-T for the applications on one segment that need to access an application on another segment, such as connecting from the web tier application to the app tier.

The Tier-1 gateways are positioned in the NSX-T network adjacent to edge devices, or Tier-0 gateways, which serve as the ingress and egress point with the external network. The Tier-1 gateways peer with the Tier-0 gateways in the NSX-T network using BGP for sharing routing information. The Tier-0 gateway, represented by NSX-T edge-virtual devices, peer with upstream external routers for the purposes of sharing routing information. This enables a pathway for traffic from a virtual machine to connect to an application on an external host, or connect to external networking services. These peering relationships form a seamless barrier between the physical network and the NSX-T virtual networks.



**Figure 36. Overview of Physical and Logical Network Routing Relationships**

Documenting the interdependencies between the applications will guide the high-level network design to support the application connectivity dependencies, and serve as the basis for the planning process of the placement of the virtual machines into VI workload domains.

# Chapter 9 Cloud Foundation on VxRail Physical Network Planning

This chapter presents the following topics:

- Introduction..... 53
- Select a physical network architecture and topology ..... 53
- VxRail stretched cluster physical network planning ..... 55
- Fibre Channel storage network planning ..... 58

## Introduction

A complete planning phase of the physical and logical networking is critical for a successful deployment of Cloud Foundation on VxRail and the ongoing operations of the Cloud Foundation management and VI workload domains. VxRail clusters are dependent on a set of physical Ethernet switches to serve as the backplane for all networking communications. The Cloud Foundation management and VI workload domains are also dependent on the supporting physical network layer to enable virtual machine connectivity within a domain, between domains, and to the external network. The supporting physical network for VxRail must be properly configured before building the cluster, and the same interconnected network must also meet the requirements for VMware Cloud Foundation before attempting initial deployment. Before moving to a planning and design phase, ensure that the key requirements for Cloud Foundation on VxRail are understood. As a starting point, have a good understanding of the interdependencies between the applications targeted for the Cloud Foundation VI workload domains.


## Select a physical network architecture and topology

Cloud Foundation on VxRail offers flexibility regarding the selection of a physical network architecture to support the planned deployment. A spine-leaf topology is the most common network topology for Cloud Foundation on VxRail and is considered a best practice. In this model, the VxRail nodes connect directly to the leaf-layer switches, and multiple VxRail clusters can be supported on a single pair of leaf-layer switches. The spine layer is positioned primarily for aggregating upstream traffic, providing connectivity to external resources and enabling VTEP tunneling between racks.

Decisions must be made regarding the location of the Layer 2 and Layer 3 boundaries to support Cloud Foundation on VxRail networking. The NSX-T Tier-0 gateways depend on peering with a router upstream in the physical network using External Border Gateway Protocol (eBGP) to update routing tables in the virtual network.

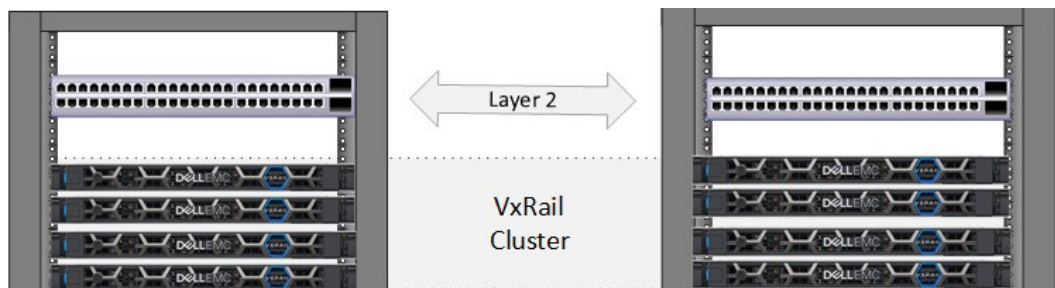
The VLANs used in Cloud Foundation on VxRail to support the guest virtual machine networks terminate at these upstream routers in the physical network. Therefore, using the route mapping for the applications planned for the VI workload domains drives the decisions for the peering of the NSX-T edge virtual devices in Cloud Foundation, and guides the process for enabling and configuring the adjacent routers in the physical network.

In most cases, routing outside of the virtual network is positioned in either the spine layer or leaf layer. If you deploy a spine-leaf network topology, enabling Layer 3 at either the spine layer or the leaf layer is not required. However, this means Layer 2 traffic must pass through both the leaf and the spine layers to reach the routers. This option is more suitable for small-scale deployments, and it is easy to deploy and configure. It is appealing for sites that have low routing requirements, or the plan is for a small workload.



**Figure 37. Options for Layer 2/ Layer 3 boundaries in spine-leaf network topology**

Establishing the router layer at the spine layer means that the uplinks on the leaf layer are trunked ports, and pass through all the required VLANs to the routing services on 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 networks that extend beyond one rack because the switches at the leaf layer do not need to support Layer 3 services, and enabling VTEP tunneling between the switches in different racks is not necessary.



**Figure 38. VxRail cluster nodes extended beyond one physical rack**

A major drawback to this topology is scalability. Ethernet standards enforce a limitation of addressable VLANs to 4094, which can be a constraint in a shared switch layer fabric. Do not select this topology option if your deployment might breach this threshold.

Enabling routing services at the leaf layer is preferred for Cloud Foundation on VxRail deployments. This option overcomes the VLAN limitation imposed by establishing the routing at the spine layer. This option optimizes routing traffic, as it requires the least number of hops for the NSX-T edge virtual devices to peer with an adjacent upstream router. A caveat is that this option does require Layer 3 services to be licensed and configured at the leaf layer. In addition, since Layer 2 networks now terminate at the leaf layer, they cannot span leaf switches. If there is a requirement to extend Layer 2 networks across switches in multiple racks, the best practice is to enable hardware-based (VTEP) tunneling.

The key points to consider for the decisions regarding the network architecture and topology are:

1. Select Ethernet switches that support the features required for Cloud Foundation on VxRail:
  - Border Gateway Protocol: Required for peering with NSX-T edge gateways
  - Unicast: Required for VxRail traffic
  - Multicast: Required for device discovery. Not required if selecting manual device discovery option instead.

- Jumbo Frames: Required for GENEVE
  - Hardware-based tunneling (VTEP): Required to extend Layer 2 traffic over a Layer 3 network at the physical switch layer
  - DHCP 'helper': Switches that support the DHCP 'helper' functionality ease connecting DHCP services in the data center to the Cloud Foundation on VxRail environment.
2. Decide which physical network layer will support Layer 3 routing services.

## VxRail stretched cluster physical network planning

Cloud Foundation on VxRail supports two types of VxRail clusters: one where all nodes are in a single site and a stretched cluster where the nodes are equally distributed between two sites. VxRail stretched cluster is targeted specifically for situations with very high RPO and RTO requirements, and as such includes additional requirements to those for a cluster in a single location.


---

**Note:** For full details about VxRail stretched cluster and its requirements, see the [Dell EMC VxRail Stretched Cluster Planning Guide](#).

---

The foundation for VxRail stretched cluster is based on vSphere vSAN stretched cluster. The basic guidelines for a vSphere vSAN stretched cluster are:

- You must have three physical site locations.
- The VxRail nodes that consist of the stretched cluster instance are spread evenly over two physical sites.
- The third site supports the witness that monitors (using heartbeat) the health of the vSAN datastore that is positioned between the two sites. The required witness is a VMware virtual appliance, so the third site must have a vSphere platform at a supported VCF on VxRail version to support the witness.
- The network between the sites must meet strict latency and bandwidth requirements since it must support synchronous I/O to vSAN for the running virtual machines in the stretched cluster.
  - 5 millisecond RTT between data node sites
  - 200 millisecond RTT between data node sites and the Witness site



**Figure 39. VxRail stretched cluster network requirements**

Cloud Foundation on VxRail is supported on a vSphere vSAN stretched cluster as the underlying foundation, and the basic tenets for vSphere stretched cluster are applicable. There are additional networking requirements specific to Cloud Foundation on VxRail.

To ensure connectivity in the event of a site outage, the Cloud Foundation on VxRail networks must be extended across the two sites, and must adhere to specific connectivity requirements.

Cloud Foundation on VxRail Networks	Site to Site Connectivity
External Management	Layer 2
vSAN	Layer 3
vMotion	Layer 2 or Layer 3
NSX-T host overlay network	Layer 3
NSX-T edge overlay network	Layer 2
NSX-T edge node uplinks	Layer 2


- Only Layer 2 is supported for the external management network between sites to prevent the need to re-IP the management components at the surviving site.
- Reserve a routable IP subnet for the vSAN network.
- vMotion can be either Layer 2 or Layer 3. A routable IP subnet is recommended for expansive virtual machine mobility.
- The NSX-T host overlay network must be routed between sites over a Layer 3 network.



- If the Application Virtual Network is deployed during Cloud Foundation on VxRail, the NSX-T edge overlay network and NSX-T node uplinks must be configured between sites. Layer 2 is supported for these networks.
- For Cloud Foundation on VxRail domains, Layer 3 is required between the sites housing the VxRail nodes and the witness site.

Cloud Foundation on VxRail Networks	Site to Witness Connectivity
External Management	Layer 3
vSAN	Layer 3

The witness site should be geographically dispersed from the stretched cluster sites in a separate failure domain, and therefore require routing services to enable connectivity. Having the witness in a separate failure domain enables it to distinguish a site failure from a network interruption between the stretched cluster sites. There might be limited instances where the stretched cluster is confined to a campus, but this model does not offer the same level of protection as a separate failure domain.



**Figure 40. Mapping of witnesses to VxRail stretched cluster sites**


To support virtual machine network traffic between Cloud Foundation domains, the MTU size must be set to a minimum of 1600 at each site. The MTU size selected must also be configured for traffic destined for the witness site.

If you deploy any Cloud Foundation on VxRail workload domains with stretched clusters as the underlying foundation, the VxRail cluster supporting the management workload domain must also be configured as a stretched cluster. If at some point in the future there is the possibility of a VI workload domain being configured with VxRail stretched clusters, it is best practice to configure the management workload domain on a VxRail stretched cluster at initial deployment. Converting an operational single site VxRail cluster instance to a stretched cluster requires additional planning and preparation as outlined in this section, as well as additional deployment work. The milestones for consideration include:

- Identify a second data center site to host the VxRail nodes to support the stretched cluster.
- Identify a third data center location for the stretched cluster witness.
- Configure the supporting networks in all sites to support VxRail stretched cluster requirements.
- Deploy the witness at the third data center site.
- Deploy additional VxRail nodes in the second data center site to balance with the VxRail nodes in the first data center site.
- Add the additional nodes into the existing VxRail cluster supporting the management workload domain.
- Convert the single site VxRail cluster to a stretched cluster.

## Fibre Channel storage network planning

If your plans for any VI workload domains include using external storage resources in your data center to support virtual machine operations, then you must prepare your data center to provide at least one LUN to the nodes that are members of a VxRail cluster supporting a VI workload domain.



**Figure 41. Fibre channel network providing storage resources for VI workload domain**

Before initial deployment, ensure that these pre-requisites are met:

- Your VI workload domain will be deployed on a standard VxRail cluster, and not on a VxRail stretched cluster. vSAN is the only storage option with VI workload domains supported by VxRail stretched clusters.
- Your planned deployment is not based on a Cloud Foundation on VxRail consolidated architecture, where the management domain and VI workload domains share VxRail cluster resources. vSAN is the only storage option for the management domain, and therefore external storage is not supported for consolidated architectures.
- Ensure that your Fibre Channel storage array and Fibre Channel switches are compatible with the Cloud Foundation on VxRail version you select to support VI workload domains. Your Dell Technologies account team can help identify if your data center resources are compatible.
- The VxRail nodes ordered to support the VI workload domains are configured at the factory with compatible Fibre Channel adapter cards. Your Dell Technologies account team will guide you through the ordering and selection process.
- Ensure that there are sufficient open ports on the Fibre Channel switches in your data center to support connecting the LUNs to the nodes. Plan on a minimum of two Fibre Channel connections to each VxRail node.
- Ensure you have sufficient unused storage capacity on the storage array to support the planned VI workload domains. Dell Technologies resources will perform a sizing exercise to determine the size of the LUN or LUNs required to support the planned workload.
- The VxRail cluster will depend on the LUNs for primary storage. Plan on creating at least one LUN of a minimum of 800 GB on the storage array. Also plan on masking any LUNs, and zoning the VxRail nodes to the storage array, before any VxRail cluster initial deployment activity.

# Chapter 10 Cloud Foundation on VxRail Physical Network Preparation

This chapter presents the following topics:

- Introduction..... 61**
- Capture configuration settings for Cloud Foundation on VxRail ..... 61**
- VxRail cluster and NSX-T networks leaf switch preparation..... 61**
- DHCP services for NSX-T host overlay network preparation..... 63**
- Application Virtual Network leaf switch preparation ..... 65**
- Layer 3 network preparation ..... 65**
- BGP peering preparation ..... 67**

## Introduction

This section outlines the tasks that must be undertaken to prepare the data center network for the deployment of Cloud Foundation on VxRail.

## Capture configuration settings for Cloud Foundation on VxRail

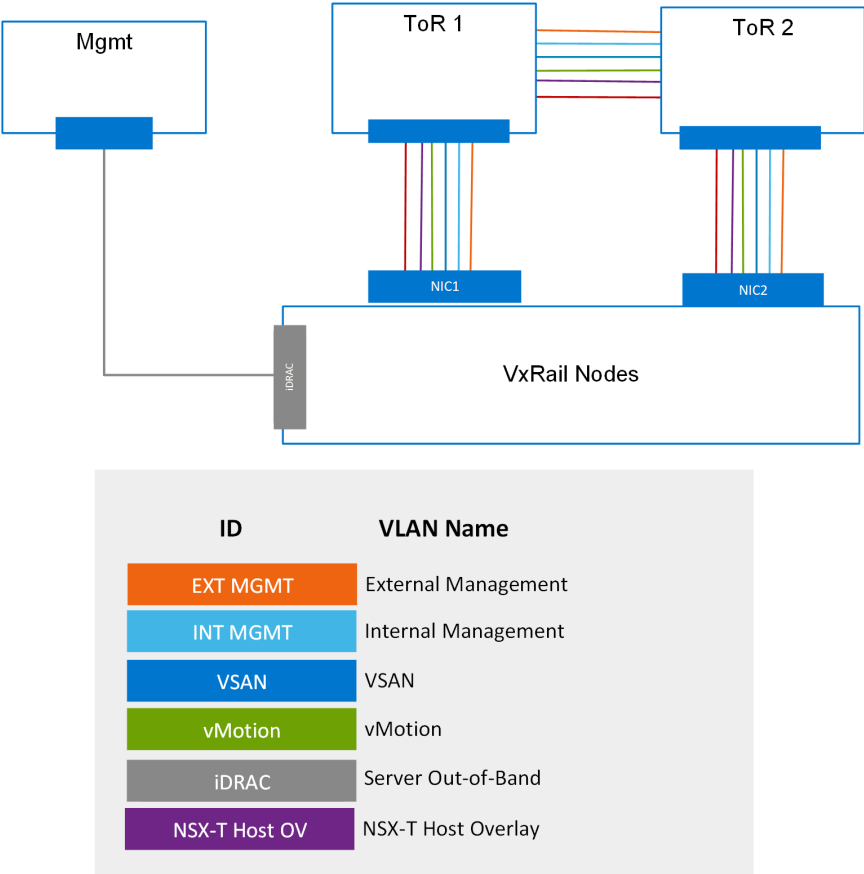
After the data center network architecture is defined and core network requirements are addressed, the next step is to prepare the data center network for the deployment of Cloud Foundation on VxRail. The first step is to capture and record the configuration settings for the deployment of Cloud Foundation on VxRail. Dell Technologies professional services will work with key stakeholders to capture the required information before moving forward to the network preparation phase.

- The table in [Appendix C: Cloud Foundation on VxRail VLANs](#) describes the core VLANs that are required for the initial deployment of Cloud Foundation on VxRail.
- The table in [Appendix D: VxRail network configuration](#) describes the configuration settings to deploy a VxRail cluster by VxRail Manager.
- The table in [Appendix E: Cloud Builder and management VI workload configuration](#) describes the configuration settings required for an initial deployment of the Cloud Foundation management workload domain.
- The table in [Appendix F: VI workload domain configuration settings](#) describes the configuration settings required to deploy a standard VI workload domain.
- The table in [Appendix G: Edge Gateway configuration](#) describes the settings required to enable connectivity to the upstream physical network with the NSX-T edge gateways.
- The table in [Appendix H: Application Virtual Network configuration](#) describes the settings required for Cloud Builder to auto-configure the Application Virtual Network after the initial deployment of the Cloud Foundation management workload domain.

The remaining tasks to prepare the data center network for initial deployment of the Cloud Foundation on VxRail, and move forward with the rest of the workflows, depends on this activity being completed.

## VxRail cluster and NSX-T networks leaf switch preparation

While the networking requirements for VxRail and Cloud Foundation differ, there is overlap in the sense that Cloud Foundation domains depend on the networking resources enabled by VxRail for connectivity. Therefore, the supporting physical network must be properly designed and configured to support VxRail cluster network traffic, and the additional requirements for Cloud Foundation.



**Figure 42. VxRail and NSX-T Overlay Networks**

A leaf switch is at the lowest tier in a multi-tier architecture, and often referred to as a ‘top-of-rack’ switch. The VxRail nodes will only connect with a leaf switches in a single rack, with the upper tier switches, known as spine switches, enable multi-rack interconnectivity.

The number of Ethernet ports from each VxRail node you reserve for Cloud Foundation on VxRail networking will drive the configuration process for each switch port connected to a VxRail node port. Starting with Cloud Foundation on VxRail version 4.0.1, up to six ports on each node can be reserved for Cloud Foundation on VxRail networking.

- a. In a 2-port configuration, VxRail network traffic and Cloud Foundation network traffic flow through the same pair of Ethernet ports
- b. In a 4-port configuration, the VxRail network traffic flows through the first two ports on the on-board network daughter card (NDC). The Cloud Foundation traffic flows either through the other two ports on the NDC, or on two ports on the PCIe/OCP expansion card.
- c. In a 6-port configuration, the VxRail network traffic flows through the four ports on the NDC, and the Cloud Foundation traffic flows through the two ports on the PCIe/OCP expansion card.

If the VxRail network traffic and Cloud Foundation network traffic will be physically separated between the nodes and the leaf switches, the VLANs for VxRail and Cloud Foundation only need to be assigned only to the required switch ports.

The following tasks must be performed in the top-of-rack switches in order to prepare for a VxRail cluster deployment and to prepare to support NSX-T:

1. Select switches with sufficient open ports capacity to connect all the VxRail nodes, connect the interswitch links between the leaf switches, and connect upstream to the adjacent network layer.
2. Configure at least 1600 MTU to support host overlay network traffic (9000 preferred). A minimum MTU size of 1600 (9000 preferred) must be configured on the leaf switches.
3. Ensure that the port type on the switches (RJ45, SFP+) match the port type on the VxRail nodes.
4. Configure each of the VLANs required for the VxRail clusters on the switches.
5. Configure VLAN for the NSX-T host overlay network on each switch. If you plan to use DHCP to supply IP addresses for the host overlay network, configure this network so that it can reach the DHCP server.
6. For each switch ports to be directly connected to the VxRail nodes:
  - a. Configure each port as Layer 2 trunk port.
  - b. Configure the VLANs required for the VxRail cluster on each port.
  - c. Configure the VLAN for the NSX-T host overlay network on each port.
  - d. Configure Spanning Tree on the switch ports to be directly connected to the VxRail nodes as edge ports, or in 'portfast' mode.
7. Configure unicast on the VLAN representing the vSAN network.
8. If opting for VxRail automatic device discovery, configure IPv6 multicast on the VLAN representing the VxRail Internal Management network.
9. Configure MLD snooping and MLD querier on the VLAN representing the VxRail Internal Management Network (recommended).
10. Configure the interswitch links to allow network traffic to pass between the two switches.

Each VxRail node has a separate Ethernet port for out-of-band server management called 'Integrated Dell Remote Access Controller' (iDRAC). A separate Ethernet switch is recommended to provide connectivity for server maintenance. The server maintenance traffic can also be redirected through the existing network infrastructure. For complete details about VxRail cluster network requirements, see the *Dell EMC VxRail Network Planning Guide*.

The table in [Appendix C: Cloud Foundation on VxRail VLANs](#) lists the individual VLANs that must be configured on the top-of-rack switches. The example switch configuration syntax displayed in [Appendix I: Sample switch configuration settings](#) offers guidance on how to configure an Ethernet switch with sample VLANs and a sample switch port configuration.


## DHCP services for NSX-T host overlay network preparation

## NSX-T host overlay network preparation

Two additional virtual NICs are created on each VxRail node to support the host overlay network during the initial deployment of Cloud Foundation on VxRail. This overlay network uses encapsulation to enable the passage of Layer 2 traffic across transport zones in the VI workload domains with a Layer 3 network. Each virtual NIC supporting the overlay network requires a routable IP address.


You can have these IP addresses input into Cloud Builder and assigned to the NICs during Cloud Foundation on VxRail deployment, or leverage DHCP services for IP address assignment. If your requirements include deploying VxRail stretched clusters to support Cloud Foundation, DHCP is the only supported method for IP address assignment.

1. Prepare a pool of IP addresses that is at least equal to double the number of VxRail nodes planned for the deployment.
2. Deploy a DHCP server in the data center, if necessary.
3. Configure the pool of IP addresses for DHCP services to support the NSX-T host overlay network.
4. Configure the data center network so that NSX-T host overlay network can reach the DHCP server.



**Figure 43. NSX-T host overlay uplinks IP addresses assigned by DHCP**

**Note:** If the VLAN cannot be extended out to the DHCP server, enabling 'DHCP helper' services on the host overlay VLAN is recommended, if supported on the leaf switches.



**Figure 44. DHCP helper connects DHCP services to host overlay VLAN**



## Leaf switch preparation for NSX-T edge

The NSX-T edge gateways enable the connectivity between the upstream physical network and NSX-T, the product which enables virtual networking in the VI workload domains. If your requirements include using NSX-T to support workloads on Cloud Foundation on VxRail, or you plan to deploy the vRealize Suite on your Cloud Foundation instance, you must prepare the leaf switches supporting the Cloud Foundation on VxRail environment to connect to the NSX-T edge gateways.

Follow these steps to prepare the leaf switches:

1. Configure a unique VLAN on each switch to support the NSX-T edge uplinks.
2. Configure a common VLAN on both switches to support the NSX-T edge overlay network.
3. The NSX-T edge uplink network and the NSX-T edge overlay network require static IP addresses. Apply Layer 3 settings to the VLANs on the leaf switches representing the NSX-T edge uplink network and the NSX-T edge overlay network.
4. For each switch ports to be directly connected to the VxRail nodes:
  - a. Configure the VLAN for the NSX-T edge uplinks.
  - b. Configure the VLAN for the NSX-T edge overlay network.




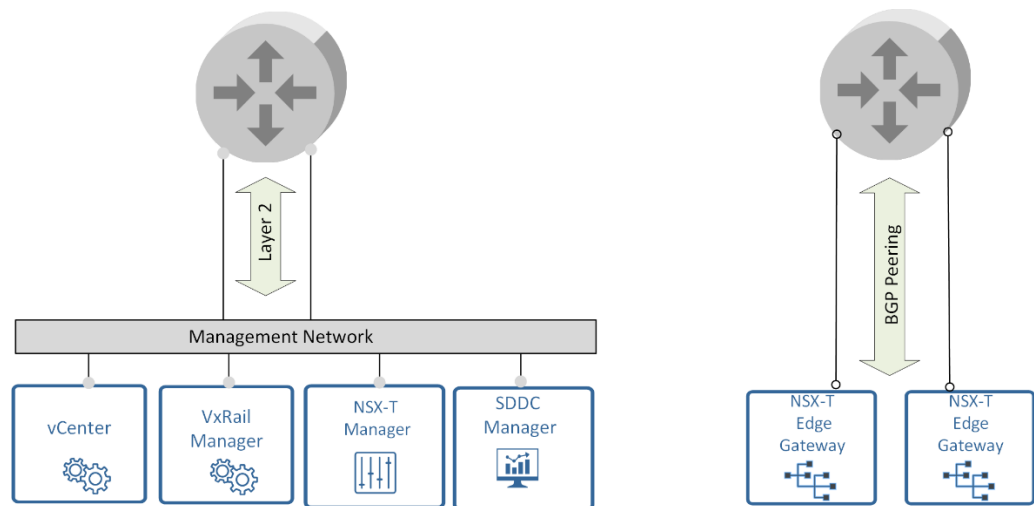
Figure 45. NSX-T Edge and Overlay Networks

Use the table in [Appendix G: Edge Gateway configuration](#) as a reference for the required settings.

## Layer 3 network preparation

The preparation for network routing services can be understood to follow these rules:

- The Application Virtual Network in the management domain, and any future workload domains planned for NSX-T will require peering with upstream BGP neighbors to sync routing tables and enable external access.
- There is a single management network shared by both VxRail and Cloud Foundation. This network enables connectivity to key management components such as SDDC Manager, NSX-T Manager and VxRail Manager. This VLAN-backed network must be configured to pass upstream to data center services such as DNS, and NTP, and end-users at the Layer 2/3 boundary.



**Figure 46. Comparison of upstream connectivity for Cloud Foundation on VxRail**

- The VxRail vSAN network can be configured with either a public IP subnet or a private IP subnet. A public IP subnet will support the extension of a VxRail cluster vSAN datastore across racks using Layer 3 services, and is the recommended option.
- The VxRail vMotion network can also be configured with either a public IP subnet or a private IP subnet. Like the vSAN network, a public IP subnet will expand support for VM migrations across racks using Layer 3 services.
- The VxRail nodes supporting a workload domain participate in the same overlay network, known as the 'NSX-T Host Overlay Network', which enables the virtual machines on these different nodes to use the overlay network for communication. The VxRail nodes connected to this overlay network must be able to communicate with each other, and this network must be routable to the 'NSX-T Edge Overlay Network'.
- The NSX-T edge nodes supporting gateway services connect to the 'NSX-T Edge Overlay Network', and this network must be routable. The NSX-T host overlay network and NSX-T edge overlay network are routed to each other.

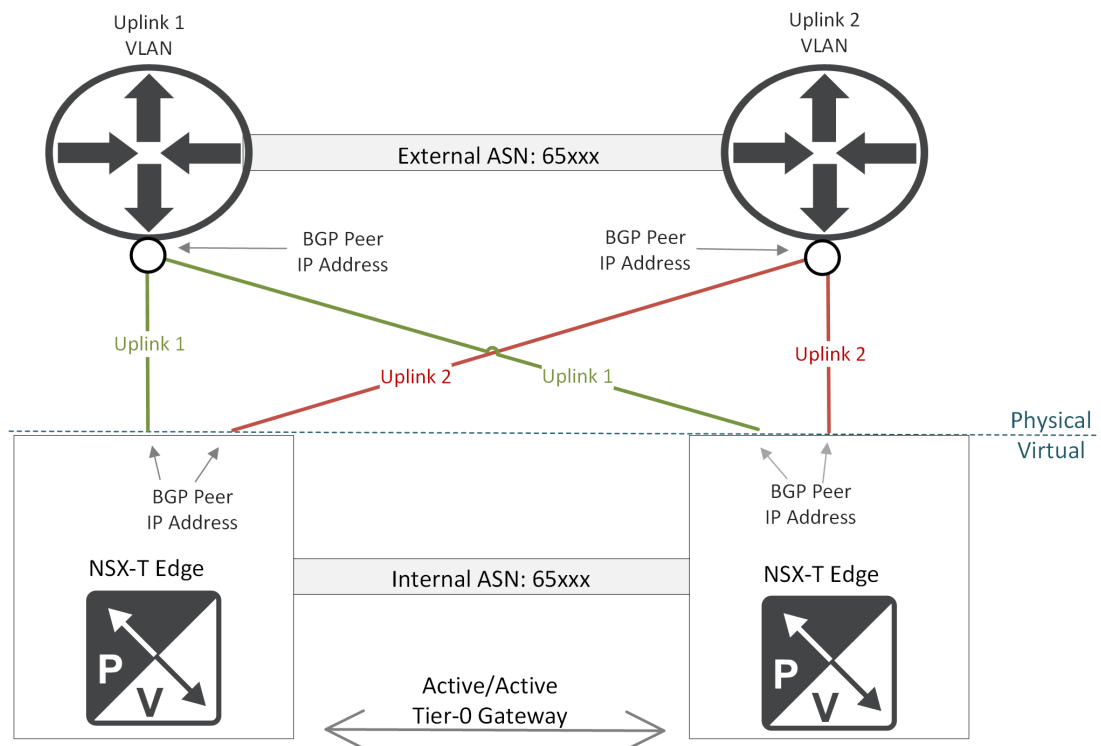
After the initial deployment of the management domain, the VI workload domains that can be constructed in a standard architecture that require NSX-T will require the configuration of at least one NSX-T edge cluster for upstream connectivity. This edge cluster will include two additional NSX-T edge nodes that must peer with upstream routers running

BGP to enable external network access. Any additional VI workload domains can share an existing edge cluster, or a new edge cluster can be deployed to support NSX-T networking.

## BGP peering preparation

Border Gateway Protocol (BGP) services on the Layer 2-3 network boundary should be configured before the initial deployment of Cloud Foundation on VxRail. Neighbor relationships upstream should be established to enable connectivity to required data center services and end-users, and to the external Dell Technologies and VMware support sites.

The tables in [Appendix G: Edge Gateway configuration](#) provide guidance on the settings that must be captured to enable BGP peering with the NSX-T Tier-0 gateways for the Application Virtual Network.



**Figure 47. BGP relationship between NSX-T Edge Gateways and external routers**

During the Cloud Builder deployment process, the NSX-T edge devices required for AVN must be able to establish an eBGP peer relationship with upstream routing services. The following tasks must be completed on the upstream switches to enable peering with the NSX-T Edge Tier-0 gateways:

1. BGP is configured on each router instance.
  - a. Configure BGP with a common Autonomous System Number (ASN) on the network devices targeted for peering with the NSX-T edge gateways.
  - b. Configure the IP prefix list to allow passage of all networks between the physical and virtual networks.

2. Configure two eBGP neighbors on first router instance.
  - a. Configure the IP address assigned to the first uplink on the first NSX-T edge device for peering. This IP address is assigned to the first NSX-T edge uplink VLAN.
  - b. Configure the IP address assigned to the first uplink on the second NSX-T edge device for peering. This uplink will also be assigned to the first NSX-T edge uplink VLAN.
  - c. Configure the timer 'keepalive' value is to 4 and the timer 'holdtime' value to 12.
  - d. Configure a password on the neighbor instance. This password is captured and configured on the adjacent NSX-T Tier-0 gateways.
  - e. Configure the internal ASN value assigned to the NSX-T edge devices.
3. Configure two eBGP neighbors on second router instance.
  - a. Configure the IP address assigned to the second uplink on the first NSX-T edge device for peering. This IP address is assigned to the second NSX-T edge uplink VLAN.
  - b. Configure the IP address assigned to the second uplink on the second NSX-T edge device for peering. This uplink will also be assigned to the second NSX-T edge uplink VLAN.
  - c. Configure the timer 'keepalive' value is to 4 and the timer 'holdtime' value to 12.
  - d. Configure a password on the neighbor instance. This password is captured and configured on the adjacent NSX-T Tier-0 gateways.
  - e. Configure the internal ASN value assigned to the NSX-T edge devices.
4. Configure a VLAN to match the VLAN assigned to the uplinks on the NSX-T Tier-0 Gateways on each router instance.
5. Configure a gateway IP address for the VLAN assigned to the uplinks on the NSX-T Tier-0 Gateways on each router instance.

If the AVN option is selected during the deployment of the Cloud Foundation management workload domain, the Cloud Builder process performs the following tasks:

- If a separate virtual distributed switch is planned to support NSX-T traffic, it is configured on the VxRail cluster.
  - The first uplink is assigned to the first VMnic reserved for NSX-T traffic.
  - The second uplink is assigned to the second VMnic reserved for NSX-T traffic.
- Two portgroups for the NSX-T edge nodes are configured on the virtual distributed switch supporting NSX-T in the management workload domain.
  - The first uplink is active on the first portgroup, and the second uplink is on standby.
  - The second uplink is active on the first portgroup, and the first uplink is on standby.

- Two NSX-T edge node virtual appliances are configured in the management domain with three virtual network adapters, forming the edge cluster.
  - One virtual network adapter connects to the management network
  - The other two virtual network adapters are connected to the two portgroups configured for NSX-T edge nodes respectively.
  - Two IP addresses are assigned to each node for the tunnel endpoints (TEP) to enable connectivity to the NSX-T edge overlay network.
- Two NSX-T edge gateway instances are configured in the edge cluster.
  - An Autonomous System Number (ASN) is assigned to the NSX-T edge gateways.
  - The BGP timers and passwords (if applicable) are configured.
  - An IP address is assigned to each instance to enable BGP peering with the first external router using the first NSX-T edge uplink network.
  - An IP address is assigned to each instance to enable BGP peering with the second external router using the second NSX-T edge uplink network.
- Internal Border Gateway Protocol (iBGP) services are enabled on the NSX-T edge gateways for connectivity with downstream NSX-T logical routing services.

The sample switch configuration syntax displayed in [Appendix I: Sample switch configuration settings](#) provides guidance on how to configure an Ethernet switch for peering with a pair of Edge Gateways.

## Chapter 11 VxRail Cluster Deployment Preparation

This chapter presents the following topics:

<b>Introduction.....</b>	<b>71</b>
<b>Prepare for VxRail cluster initial build.....</b>	<b>71</b>
<b>Select the external management network subnet .....</b>	<b>71</b>
<b>Select the VxRail cluster VLANs .....</b>	<b>72</b>
<b>Select the network settings for VxRail cluster .....</b>	<b>72</b>
<b>Decide whether to join an existing vCenter SSO domain.....</b>	<b>73</b>
<b>Select the network settings for VxRail stretched cluster .....</b>	<b>73</b>
<b>Create forward and reverse DNS entries for VxRail cluster .....</b>	<b>73</b>
<b>Select passwords .....</b>	<b>73</b>

## Introduction

---

**Note:** Dell Technologies Professional Services will be responsible for the deployment of Cloud Foundation on VxRail per the agreed-upon statement of work. This section provides guidance on what to prepare for during that phase.

---

In preparation for the deployment of Cloud Foundation on VxRail, Dell Technologies will review a set of prerequisites that must be met for a successful outcome. Dell Technologies will also capture and record the settings and properties required for full deployment of Cloud Foundation on VxRail in your data center.

The data capture process is performed in the following phases:

1. The initial phase focuses on the VxRail clusters that form the resource building blocks for Cloud Foundation. Each Cloud Foundation domain requires at least one VxRail cluster.
2. The next phase captures the settings and properties for layering the Cloud Foundation management workload domain on the first VxRail cluster.
3. The next phase captures the settings and properties for the planning Cloud Foundation VI workload domains.
4. The final phase focuses on the deployment of the NSX virtual networks in the respective Cloud Foundation VI workload domains.

[Appendix D: VxRail network configuration](#) describes the settings required for VxRail cluster formation and management.

## Prepare for VxRail cluster initial build

The initial build operation of a VxRail cluster transforms the physical nodes into a single, managed vSphere cluster with a single vSAN datastore and a single virtual distributed switch. The initial build operation occurs after the following steps are completed:

1. The adjacent top-of-rack switches are configured per VxRail requirements.
2. The VxRail nodes are installed in the racks and cabled to power and network sources.
3. If FC LUNs instead of vSAN will serve as the primary storage resource for a VxRail cluster, perform the pre-requisites to configure the LUNs on the storage array and configure the storage area network to present the LUNs to the VxRail nodes.
4. The VxRail nodes are powered on.
5. VxRail performs self-discovery of the powered-on nodes, and starts VxRail Manager for input of settings to perform automated initial build.

## Select the external management network subnet

There is one overall management network for Cloud Foundation on VxRail. The management components for VxRail, Cloud Foundation and NSX-T all share this same subnet. Ensure the subnet range selected is of sufficient size.

The 'VxRail' category in the table in [Appendix D: VxRail network configuration](#) captures this network in CIDR format.

## Select the VxRail cluster VLANs

See [Appendix D: VxRail network configuration](#) for guidance on all the settings that must be collected to perform VxRail cluster automated initial build.

- The external VxRail management network, which represents the overall management for VxRail, Cloud Foundation and NSX-T, is the native VLAN by default. Since the typical Cloud Foundation on VxRail deployment has more than one VxRail cluster, it is advisable to select a value other than the native VLAN.
- You can choose to have VxRail Manager discover the nodes to form a cluster either through automatic discovery, or by manually ingesting the nodes into VxRail Manager. With automatic discovery, the internal VxRail management network VLAN is 3939 by default. Since the typical Cloud Foundation on VxRail deployment has more than one VxRail cluster per leaf switch pair, it is possible that an existing VxRail Manager will detect the new powered-on nodes in the rack on this same network. To bypass this, a new VLAN can be set for the second and subsequent VxRail clusters. Dell Technologies will change the internal management VLAN on each VxRail node to the recorded value before cluster formation.
- The vSAN and vMotion VLANs should be unique for each VxRail cluster. Using a unique VLAN for vSAN and vMotion removes the possibility of conflicts when multiple VxRail clusters are sharing a set of switches.

## Select the network settings for VxRail cluster

The rows in the table in [Appendix D: VxRail network configuration](#) describe the settings required for VxRail cluster formation and management. The following should be adhered to when selecting these settings:

- The IP addresses assigned to the VxRail management components must be within the range selected for the external management network.
- The IP addresses must be unused and permanent IP addresses. They cannot be assigned by DHCP.
- The NTP and DNS servers must be accessible to this external management network.
- While using contiguous IP addresses is recommended for the VxRail management networks, it is not required.
- The network range selected for vMotion and vSAN can be public or private. If you want to ensure that these networks can expand across racks with Layer 3, use a routable IP address range.



- Hostnames and IP addresses selected for VxRail management components are automatically assigned during initial build. You can choose to customize the hostnames or follow a preset format for hostnames.

## Decide whether to join an existing vCenter SSO domain

If a Cloud Foundation for VxRail instance is already running in the data center, you can decide if you want this new instance to join the existing vCenter SSO domain, or create a new vCenter SSO domain. It should be noted that when you join an existing SSO domain, it reduces the number of workload domains the Cloud Foundation on VxRail instance can support. This is due to the limitation of 15 vCenter instances supported in enhanced linked mode.

## Select the network settings for VxRail stretched cluster

If your plans include the deployment of a VxRail stretched cluster, additional network settings must be captured. The second table in [Appendix D: VxRail network configuration](#) includes the additional network settings required for VxRail stretched clusters.

Capture the settings to deploy the witness virtual appliance at the third site, which is required before performing initial build of the VxRail stretched cluster.

- Settings for the witness management network
- Settings for the witness vSAN network
- IP address of the vSphere host supporting the witness virtual appliance

Additional network settings need to be captured for the second site. Depending on the decisions made for the stretched cluster deployment, additional network configuration must be performed at the second site:

- Layer 3 network services are required for the VSAN network between sites.
- The vMotion network supports either Layer 2 or Layer 3 networks between the sites.
- A Layer 3 network is required between sites to support NSX-T.

## Create forward and reverse DNS entries for VxRail cluster

Using the information captured in [Appendix D: VxRail network configuration](#), create forward and reverse DNS entries for every hostname planned for the VxRail cluster. These include VxRail Manager, vCenter Server, and each ESXi host in the VxRail cluster.

## Select passwords

For VxRail cluster components, a password is required. The password policy follows VMware standards, as described in the 'vSphere Security' guide at [VMware vSphere Documentation](#).

# Chapter 12 Prepare for VMware Cloud Foundation Management VI Workload Domain

This chapter presents the following topics:

- Introduction.....75
- Provide a temporary IP address for Cloud Builder .....75
- Select the settings for the management workload domain .....75
- Provide global settings for management VI workload domain .....75
- Select the settings for NSX-T host overlay network .....76
- Create forward and reverse DNS entries for the management VI workload domain.....76
- Select the NSX-T host overlay VLAN.....76
- Select names for resource pools in VI Management workload domain.....77
- Decide on number of virtual distributed switches .....77
- Prepare passwords .....78
- Obtain VMware license keys .....78

## Introduction

To configure the management workload domain, Dell Technologies will follow these steps:

1. Capture and record the settings specific to the Cloud Foundation management VI workload domain and Application Virtual Network in a workbook.
2. Download and deploy the Cloud Foundation Cloud Builder virtual appliance on the VxRail cluster.
3. Upload the settings captured in the workbook for the Cloud Foundation on VxRail to the virtual appliance.
4. Activate the Cloud Builder process. This lays down the management workload domain on top of the VxRail cluster using the uploaded settings.

See [Appendix E: Cloud Builder and management VI workload configuration](#) for guidance on all the settings that must be collected for Cloud Builder.

## Provide a temporary IP address for Cloud Builder

Deploying the Cloud Builder virtual appliance requires an IP address to be accessible from the VxRail external management network.

## Select the settings for the management workload domain

The IP address range for the management workload domain is captured in the 'Management Network' row in the table in [Appendix D: VxRail network configuration](#). The IP addresses assigned to the components in the management VI workload domain:

- Must be on the same subnet as the VxRail, Cloud Foundation and NSX-T management components
- Must not be in use by another component
- Cannot be assigned by DHCP

## Provide global settings for management VI workload domain

Provide the following global settings:

- The IP addresses for the DNS server(s) to support the management VI workload domain. The required forward and reverse entries to support the Cloud Foundation management workload domain must have already been completed.
- The IP addresses for the NTP server(s) to support the management VI workload domain
- The site name for single sign-on (SSO). The site name must be the same as the site name used for the underlying VxRail cluster, with the default of 'Default-First-Site'.
- The domain name and sub-domain name
- Hostname and IP address for the SDDC Manager

- Hostnames and IP addresses for the three NSX-T Management nodes, and the single virtual IP address for NSX-T Management
- Hostnames and IP addresses for the NSX-T edge nodes, if AVN is part of the deployment plan

## Select the settings for NSX-T host overlay network

If you choose to have the IP addresses for the NSX-T host uplinks assigned by CloudBuilder during Cloud Foundation on VxRail deployment, capture and record the following settings:

- A name for this static IP address pool to provide to Cloud Builder
- The IP address range to reserve for the NSX-T host uplinks in CIDR format. The range should be at least double the number of VxRail nodes planned for initial deployment.
- The starting and ending IP addresses for Cloud Builder to assign to the NSX-T host uplinks
- The gateway for the reserved IP address range

If you choose to have the IP addresses assigned dynamically during initial deployment, capture and record the following DHCP settings to support the VTEP tunnel endpoints:

- The IP address of the DHCP server that will supply two IP addresses to each VxRail node for enable connectivity to the host overlay network
- The pool of IP addresses that were configured in the DHCP server for assignment to each VxRail node. Each node requires two IP addresses.

## Create forward and reverse DNS entries for the management VI workload domain

Create forward and reverse DNS entries for the Cloud Foundation management VI workload domain. These are the SDDC Manager and all the components in the NSX-T Management cluster. If Application Virtual Network is enabled, include the forward and reverse entries for the NSX-T edge nodes.

## Select the NSX-T host overlay VLAN

- The NSX-T host overlay VLAN must be configured on the adjacent top-of-rack switches connected to the VxRail nodes.
- The NSX-T host overlay VLAN must be configured on the trunk ports connected to the VxRail nodes.
- Depending on where the Layer 2/ Layer 3 boundary is in the supporting physical network, pass the NSX-T host overlay VLAN upstream through the uplinks on the adjacent top-of-rack switches to the spine switch layer.
- Gateway IP addresses must be assigned to the switches at the Layer 2/ Layer 3 boundary to enable routing.

- If DHCP is used to supply the IP addresses, configuring 'dhcp helper' is recommended on the NSX-T host overlay VLAN, if supported by the switches.


## Select names for resource pools in VI Management workload domain

For a consolidated architecture, four resource pools will be created during the build process for the VI management workload domain. Resource pools for management components and edge components will be created. The default names provided for these four resource pools can be customized.

## Decide on number of virtual distributed switches

There are four supported options to choose from for the deployment of virtual distributed switches to support the management and VI workload domains. The options provide a way to segment traffic based on network type based on the number of NICs configured to support the Cloud Foundation on VxRail instance.

If four Ethernet ports per node are configured to support Cloud Foundation on VxRail networking, then you can choose to either have a single virtual distributed switch be deployed for both VxRail networking and NSX-T networking, or deploy separate virtual distributed switches and move the resource-intensive networks to a separate virtual distributed switch.

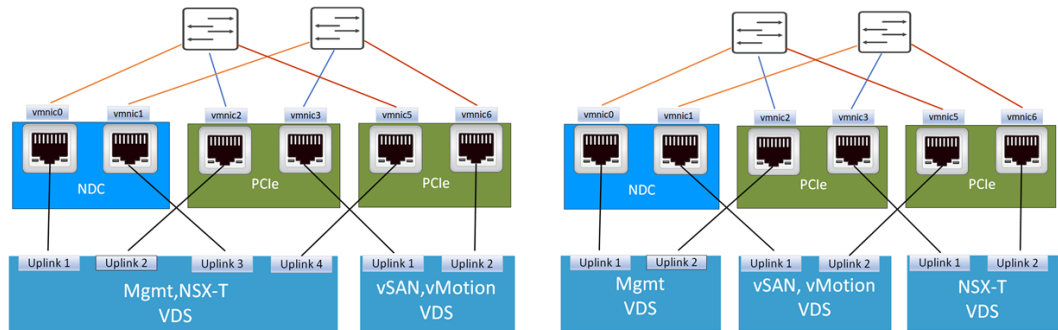


**Figure 48. VDS options with 4 reserved NICs**

With a second virtual distributed switch, the vSAN network, which is the most resource-intensive network required by Cloud Foundation on VxRail, can be isolated on its own virtual distributed switch with its own dedicated uplinks. Additionally, the vMotion network can either remain on the first virtual distributed switch and share bandwidth with the management and NSX-T networks, or migrate over the second virtual distributed switch.

For Cloud Foundation on VxRail workloads with very high resource requirements, you can elect to reserve eight Ethernet ports per node, and configure four uplinks on each virtual distributed switch.

If six Ethernet ports per node are configured to support Cloud Foundation on VxRail networking, then you can choose to deploy as many as three virtual distributed switches to support VxRail networking and NSX-T networking.



**Figure 49. VDS options with 6 reserved NICs**

If you elect to deploy two virtual distributed switches, then the resource-intensive networks can be moved to the second virtual distributed switch, and two uplinks on the first virtual distributed switch can be reserved for NSX-T traffic. Alternatively, you can choose to deploy a third virtual distributed switch to isolate the NSX-T traffic.

## Prepare passwords

A password is required for Cloud Foundation management workload domain components. Like VxRail, the password policy follows VMware standards: See the 'vSphere Security' guide at VMware's documentation site for vSphere: [VMware vSphere Documentation](https://docs.vmware.com/en/VMware-vSphere/7.0/vSphere-Security/vSphere-Security.html).

## Obtain VMware license keys

Cloud Foundation on VxRail is deployed with temporary license keys. Permanent license keys must be entered before the expiration of the grace period. The following licenses are required, depending on your final configuration:

- vCenter Server Standard
- ESXi Enterprise Plus (for Management and VI Workload Domains)
- vSphere Add-on for Kubernetes (Required for VI Workload Domains for Kubernetes)
- vSAN Advanced or higher
- NSX-T Data Center Advanced or higher (required for AVN)
- vRealize Suite (if vRealize deployed)
- SDDC Manager

## Chapter 13 Prepare for Cloud Foundation VI Workload Domain

This chapter presents the following topics:

<b>Introduction.....</b>	<b>80</b>
<b>Cloud Foundation workload domain task outline.....</b>	<b>80</b>
<b>Prepare NSX-T host overlay network .....</b>	<b>80</b>
<b>Capture settings for VI workload domain .....</b>	<b>81</b>
<b>Prepare for vSphere for Tanzu workload domain .....</b>	<b>81</b>
<b>Prepare for multi-region with NSX-T Federation .....</b>	<b>83</b>

## Introduction

A consolidated architecture combines management and workload into a single domain instance running on a common pool of resources supplied by VxRail. A standard architecture instead has a separate pool of resources for each domain, each with a separate pool of underlying VxRail clusters.

With a standard architecture, once a Cloud Foundation on VxRail management workload domain is deployed, the cloud platform is ready for the deployment of Cloud Foundation on VxRail VI workload domains. This milestone must be completed before moving forward with any other activities with VI workload domains.

This section provides a basic outline of how to configure a VI workload domain using SDDC Manager, and describes the key steps for a deployment of Cloud Foundation with VxRail as the underlying platform.

## Cloud Foundation workload domain task outline

Configuring a Cloud Foundation on VxRail workload domain requires at least one fully deployed VxRail cluster to serve as the underlying resource pool. The following outline describes the steps to deploy a Cloud Foundation VI workload domain:

1. Deploy the VI workload domain logical structure from SDDC Manager.  
This task creates the workload domain logical construct in SDDC Manager and deploys a vCenter instance in the management workload domain. This vCenter instance is used to manage all VxRail clusters that are assigned to support the VI workload domain. The IP address assigned to this vCenter instance must be in the management network subnet that is reserved for the Cloud Foundation management domain.
2. Deploy at least one VxRail cluster.  
Use the VI workload domain vCenter as the point of management during the initial build of the VxRail cluster.
3. Add the VxRail cluster to the VI workload domain using SDDC Manager.
4. Decide on deploying new NSX-T management resources or use existing NSX-T management resources.
  - For the first VI workload domain, three NSX-T Managers are deployed in the management domain and configured with a virtual IP (VIP).
  - Subsequent VI workload domains can share an existing set of NSX-T Managers, or a new set of NSX-T Managers can be deployed.

## Prepare NSX-T host overlay network



Each VI workload domain requires at least one VLAN for the NSX-T host overlay network. The VLAN drives the Transport Zone selection the VI workload domain can be a member of.

- Best practice is to configure a new VLAN for each VI workload domain to support the NSX-T host overlay network.
- Applying the same VLAN to each VxRail cluster added to the VI workload domain ensures that all VxRail nodes are members of the host overlay network.
- Follow these tasks for a new VLAN for the NSX-T host overlay network:
  - Configure the VLAN on the adjacent leaf switches providing connectivity to the VxRail nodes.
  - Add the VLAN to each trunked port on the adjacent leaf switches providing connectivity to the VxRail nodes.
  - If using DHCP to supply IP addresses for the host overlay network, ensure connectivity to a DHCP server supplying the IP addresses

## Capture settings for VI workload domain

Use the tables in [Appendix F: VI workload domain configuration settings](#) as a guide for the VI workload domain deployment requirements.

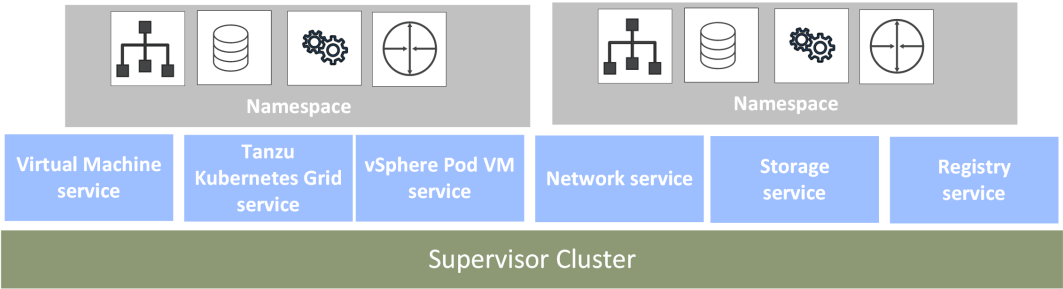
- Each VI workload domain requires a unique name.
- A data center name is required. The VxRail cluster is placed under this data center in vCenter.
- Network settings are required for the vCenter instance that is deployed in the management workload domain.
- Credentials for a management account for the vCenter instance

Follow the steps to capture and record the settings for a VxRail cluster. Use the network settings for the vCenter instance configured for the workload domain as the management target for the VxRail cluster.

## Prepare for vSphere for Tanzu workload domain

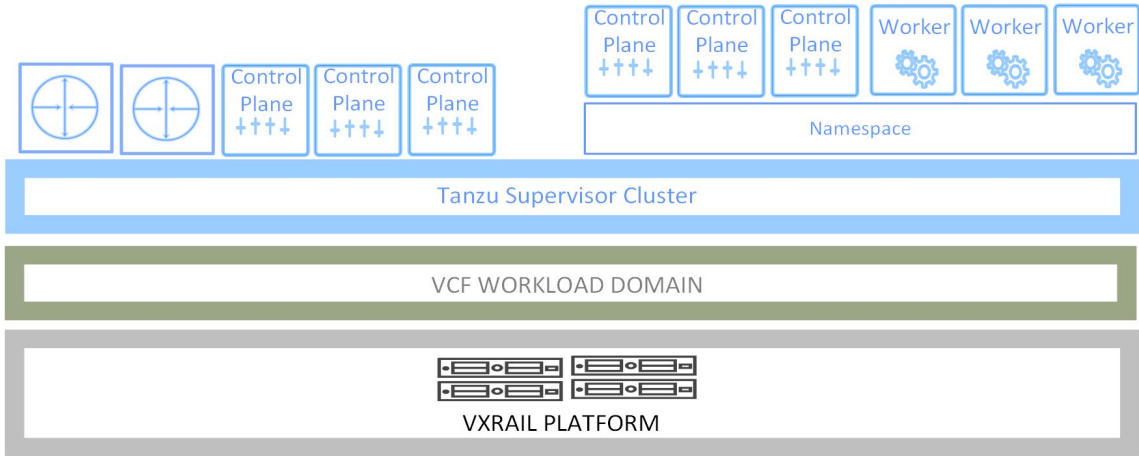
You can use the SDDC Manager in Cloud Foundation on VxRail to deploy a workload domain that will support Tanzu. Tanzu is a full distribution of the open-source Kubernetes container orchestration software that is packaged, signed, and supported by VMware. SDDC Manager will perform the configuration of the workload domain to support a Kubernetes supervisor cluster, and enable all the underlying services to support namespaces on the workload domain resources.

Under this environment, the supervisor cluster uses the services enabled in vSphere to support Kubernetes, and uses the resources provided by the ESXi hosts as worker nodes instead of Linux hosts.




**Figure 50. Tanzu Kubernetes services on VI Workload Domain**

To prepare for the deployment of a vSphere for Tanzu workload domain using Cloud Foundation on VxRail, ensure that there are enough resources on the planned workload domain to support the planned workload. The Tanzu Kubernetes Grid Service deploys a baseline of virtual appliances on the supervisor cluster to spurt management activities from a vCenter perspective, which include the creation of namespaces for DevOps. It will also deploy a pair of NSX-T edge appliances to enable connectivity upstream to the NSX-T tier-0 gateway. In addition, each time a namespace is configured by the vSphere administrator, a set of control plane virtual appliances are deployed to enable management access. The table in [Appendix B: Cloud Foundation on VxRail footprints for sizing](#) should be used to reserve resources in the supervisor cluster to support management overhead.



**Figure 51. vSphere for Tanzu workload domain management components**

As part of the deployment process, SDDC Manager will configure a workload network to support connectivity to the Tanzu supervisor cluster, deploy NSX-T load balancers to separate the external and internal networks within the cluster, and deploy an NSX-T tier-1 gateway for ingress and egress access. NAT rules will also be established in NSX-T to enforce the separation the public and private networks.



**Figure 52. Rules for supervisor cluster networks**

The routable management network connects the management components in the supervisor cluster to vCenter, while the workload network uses NSX-T to support traffic to the Kubernetes APIs and to the pods created within the namespaces.

A set of IP address ranges must be reserved for usage by the vSphere for Tanzu workload domain.

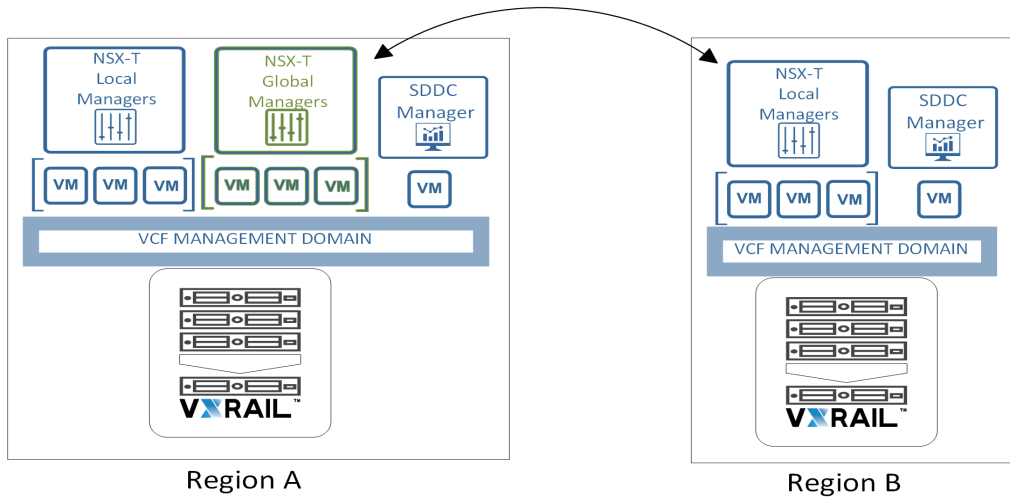
- A private IP address pool used to support pods and workloads in a namespace.
- A private IP address pool used by Kubernetes applications for service exposed within the namespace. This pool is assigned to the east-west load balancer within the supervisor cluster.
- A public IP address pool for exposing services outside of the supervisor cluster through the namespace load balancer. Each namespace will get an IP assigned to be used for a NAT rule for external access.
- A public IP address pool for NAT to use for traffic outside of the supervisor cluster

## Prepare for multi-region with NSX-T Federation

If your business requirements including supporting Cloud Foundation on VxRail across multiple regions, there are additional preparation steps that are required.

### Plan for NSX-T global managers

In a multi-region deployment with NSX-T Federation, one region is selected as the global manager for the NSX-T Federation. The NSX-T Local Managers in each region connect with the NSX-T global manager in the selected region. The global manager provides the administration support for NSX-T global objects across the regions.



**Figure 53. Global and Local Managers in NSX-T Federation**

You can also configure two regions to support global management to support active/standby failover configuration, but at least one must be selected for this purpose. The size of the virtual machines that must be deployed in the management domain selected for global management depends on the size of the multi-region NSX-T federation. See [Appendix B: Cloud Foundation on VxRail footprints for sizing](#) for guidance on sizing the virtual machines in the VCF management domain.

## Chapter 14 Prepare for NSX-T Edge Gateway Services

This chapter presents the following topics.

<b>Introduction .....</b>	<b>85</b>
<b>Capture external router settings for eBGP peering .....</b>	<b>86</b>
<b>Capture settings for NSX-T edge gateway uplinks .....</b>	<b>86</b>
<b>Capture NSX-T edge overlay network settings .....</b>	<b>87</b>
<b>Capture second site settings for stretched cluster .....</b>	<b>87</b>

## Introduction

NSX-T edge gateways peer with upstream physical network routers to enable north-south connectivity between the application and services on the Cloud Foundation on VxRail platform and external applications and services, and end users. You can choose to deploy NSX-T edge gateways and have those gateways provide service for one or more VI workload domains. With a new VI workload domain, you can use existing NSX-T edge gateways for upstream network connectivity, or deploy a new NSX-T edge gateway cluster to support the new VI workload domain.

If your requirements include using the vRealize Suite for workloads on the Cloud Foundation on VxRail platform, the installation and management of the vRealize suite is dependent on the Application Virtual Network (AVN), which is covered in [Prepare for Cloud Foundation Application Virtual Network](#). An NSX-T edge gateway cluster is a prerequisite for supporting the AVN.

## Capture external router settings for eBGP peering

The *External Routers* table in [Appendix G: Edge Gateway configuration](#) captures what is configured on the upstream routers that will peer with the NSX-T edge gateways.


- The IP addresses to be assigned to the VLANs configured on the upstream adjacent switches configured for BGP peering
- The ASN value configured on the switches
- The passwords, if any, for peering with the upstream BGP routers

## Capture settings for NSX-T edge gateway uplinks

The *NSX-T Edge Gateways* table in [Appendix G: Edge Gateway configuration](#) outlines the edge gateway settings required to support upstream connectivity.

- Assign an ASN for the workload domain edge cluster.
- Reserve hostnames for the two NSX-T edge nodes.
- Reserve four unused IP addresses from the management network subnet for the three NSX-T edge nodes and the virtual IP address to serve as the connection point for NSX-T management.
- Reserve two new VLANs to support the two NSX-T edge uplink networks for the new edge cluster.
  - The VLANs must be configured on the switch trunk ports connected to the VxRail nodes.
  - One VLAN is assigned to first of the two upstream router instances.
  - One VLAN is assigned to the second of the two upstream router instances.
- Reserve three new IP address subnet ranges
  - One subnet range for the first NSX-T edge uplink network
  - One subnet range for the second NSX-T edge uplink network

- One subnet range for the NSX-T edge overlay network
- Assign an IP address to each edge node for the first NSX-T edge uplink network.
- Assign an IP address to each edge node for the second NSX-T edge uplink network.
- Assign two IP addresses to each edge node for the NSX-T edge overlay network.
- Assign two VLANs to enable BGP peering through the two NSX-T edge uplink networks.



**Figure 54. NSX-T edge gateway peering relationship**

## Capture NSX-T edge overlay network settings


The NSX-T edge nodes must be able to pass network traffic over an edge overlay network. This edge overlay network is separate from the NSX-T host overlay network in that it enables communications between the ESX-T edge nodes.

- The selected VLAN must be configured on each leaf switch to enable connectivity over the edge overlay network.
- An IP address is assigned to each NSX-T edge node to enable connectivity over the NSX-T edge overlay network.

## Capture second site settings for stretched cluster

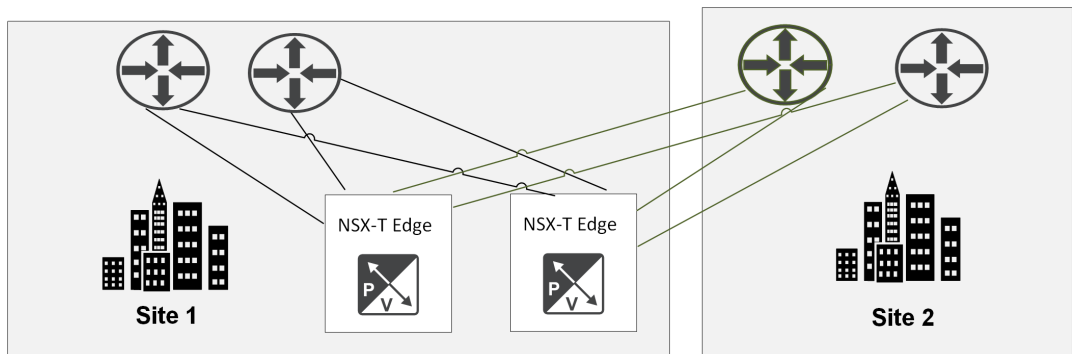
If Cloud Foundation is to be deployed on a VxRail stretched cluster, network settings must be captured for the second site. The table for the second site in [Appendix G: Edge Gateway configuration](#) outlines the second set of settings that must be selected.

The underlying networks supporting the VxRail cluster and NSX-T are configured across the two sites in the stretched cluster to enable cross-site connectivity. This enables a seamless failover to the surviving site in the event of a single site failure.



**Figure 55. Transition to surviving site for NSX-T edge nodes support Tier-0 gateway**

To ensure uninterrupted routing services, the NSX-T edge devices configured as Tier-0 gateways peer with upstream routers in both locations in a stretched cluster deployment.



**Figure 56. VxRail stretched cluster with two pairs of NSX-T Tier-0 Gateways**



## Chapter 15 Prepare for Cloud Foundation Application Virtual Network

This chapter presents the following topics.

<b>Introduction.....</b>	<b>90</b>
<b>Capture the Application Virtual Network region settings .....</b>	<b>90</b>

## Introduction

The Application Virtual Network (AVN) is required if the vRealize Suite product suite will be deployed on the Cloud Foundation on VxRail platform post-deployment. The Application Virtual Network is dependent on NSX-T edge gateways being deployed in the management domain as a pre-requisite, which is described in [Prepare for NSX-T Edge Gateway Services](#). The tables in [Appendix H: Application Virtual Network configuration](#) represent what needs to be captured to deploy the Application Virtual Network.

## Capture the Application Virtual Network region settings

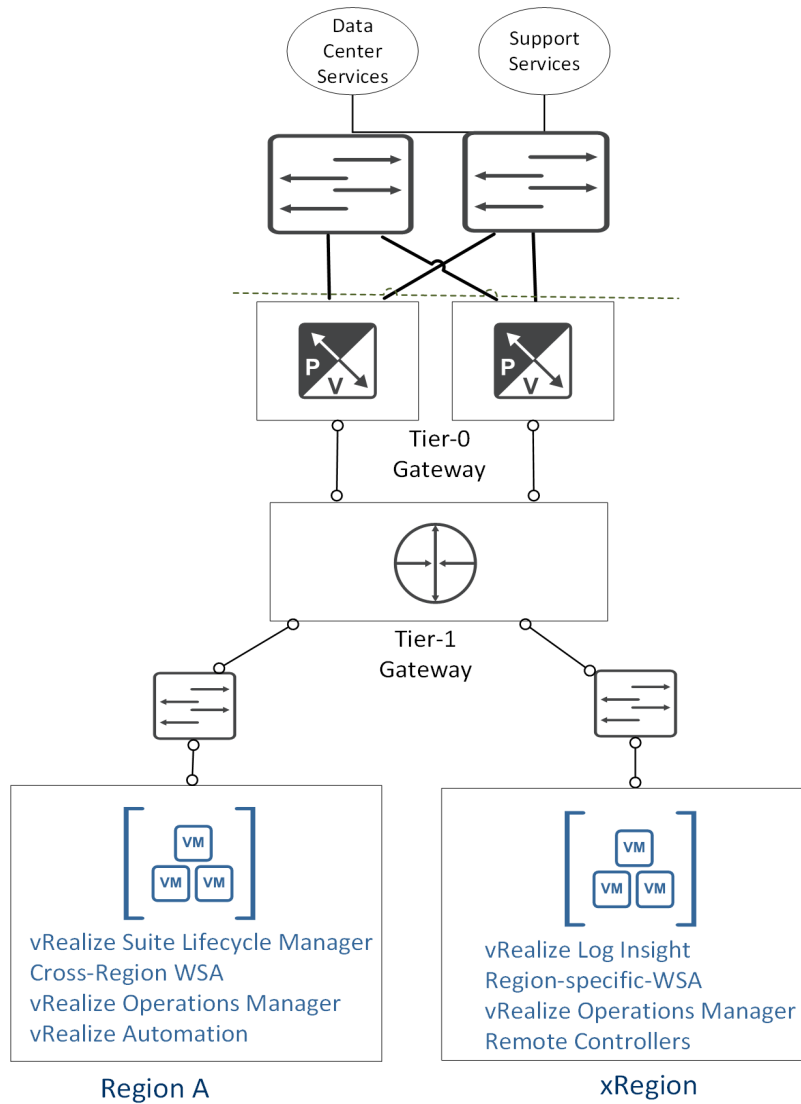
The segment 'Region A' is for the management components assigned to a specific region, whereas the segment 'xRegion' is for the vRealize management components that require mobility between regions.

Each region is connected to the NSX-T Tier-1 gateway configured for the Application Virtual Network. The Tier-1 gateway is a logical router used for managing network traffic within the NSX-T network. The NSX-T Tier-1 gateway is positioned just below the NSX-T Tier-0 gateway that peers with external routing services, enabling upstream access for the vRealize management components.

Each region must be given an IP address range or VLAN that is used to assign to the management components running in each respective region. If using VLAN-backed segments, VLANs are required for both segments. If dual regions are being planned with NSX-T Federation, the VLAN option is not supported. The table in [Appendix H: Application Virtual Network configuration](#) outlines the settings that must be captured to enable AVN:

- The name of the region-specific (Region A) logical segment
- The IP address range or VLAN reserved for the region-specific (Region A) logical segment
- The name of the cross-region x(Region) logical segment
- The IP address range or VLAN reserved for the cross-region (xRegion) logical segment

Specific details for the deployment and configuration of the vRealize software product are out of scope for this guide. See the VMware documentation [VMware Validated Design Documentation](#) for more information.



**Figure 57. AVN Regions and Logical Segments**

# Appendixes

This appendix presents the following topics:

- Appendix A: Cloud Foundation on VxRail checklist ..... 99**
- Appendix B: Cloud Foundation on VxRail footprints for sizing..... 99**
- Appendix C: Cloud Foundation on VxRail VLANs ..... 99**
- Appendix D: VxRail network configuration..... 100**
- Appendix E: Cloud Builder and management VI workload configuration .. 102**
- Appendix F: VI workload domain configuration settings ..... 104**
- Appendix G: Edge Gateway configuration..... 106**
- Appendix H: Application Virtual Network configuration ..... 107**
- Appendix I: Sample switch configuration settings ..... 108**

## Appendix A: Cloud Foundation on VxRail checklist

Workload Requirements	
Use Cases	<ul style="list-style-type: none"> <li>Determination of use cases planned for VCF on VxRail integrated platform</li> <li>Determination of application availability requirements for VCF on VxRail integrated platform</li> </ul>
Workload Planning	<ul style="list-style-type: none"> <li>Captured performance metrics from applications targeted for VCF on VxRail integrated platform</li> <li>Completed sizing exercise with Dell Technologies VCF on VxRail sizing tool</li> <li>Converted sizing report into top-level architecture for VCF on VxRail integrated platform</li> </ul>
Data Center Requirements	
Rack Space	<ul style="list-style-type: none"> <li>Calculated data center rack space and power requirements for VCF on VxRail integrated platform</li> <li>Enable redundant power sources for each data center rack.</li> </ul>
Data Center Infrastructure	<ul style="list-style-type: none"> <li>Ethernet switch ports compatible with VxRail node ports</li> <li>Sufficient open ports for VxRail nodes</li> <li>Jumbo frames enabled on data center network</li> <li>Ethernet switches supporting VCF on VxRail integrated platform support Unicast and, if applicable, Multicast</li> <li>Ethernet switches supporting VCF on VxRail integrated platform support Border Gateway Protocol</li> <li>Ethernet switches supporting VCF on VxRail integrated platform support hardware-based VTEP</li> <li>If planning to use FC storage to support VI workload, Fibre Channel infrastructure compatible with VxRail</li> </ul>
Data Center Services	<ul style="list-style-type: none"> <li>Domain Name Services (DNS) deployed in data center planned for VCF on VxRail integrated platform</li> <li>Network Time Protocol (NTP) services deployed in data center planned for VCF on VxRail integrated platform</li> <li>Active Directory (A-D) configured in data center planned for VCF on VxRail integrated platform (required for certain use cases)</li> <li>Dynamic Host Configuration Protocol (DHCP) services configured in data center planned for VCF on VxRail integrated platform. This is not required if using static IP addresses for the host overlay networks.</li> <li>SFTP server for backups for NSX-T and SDDC Manager instances configured in data center planned for VCF on VxRail integrated platform</li> <li>Certificate generation utility (required for certain use cases)</li> <li>Syslog server (optional)</li> </ul>

Remote Sites (if applicable)	
WAN	<ul style="list-style-type: none"> <li>• Minimum 10 Mb bandwidth between the central management site and any planned remote workload sites</li> <li>• Maximum 50-millisecond latency between the central management site and any planned remote workload sites</li> <li>• WAN redundancy enabled between the central management site and any planned remote workload sites</li> <li>• Geneve-compatible network between VCF on VxRail management domain instances across regions if NSX-T Federation is planned.</li> </ul>
Licensing	
Licenses	<ul style="list-style-type: none"> <li>• vCenter Server Standard</li> <li>• ESXi Enterprise Plus (Management and VI Workload Domains)</li> <li>• ESXi Enterprise Plus for Kubernetes (VI Workload Domains for Kubernetes)</li> <li>• vSAN Advanced or higher</li> <li>• NSX-T Data Center</li> <li>• vRealize Suite (minimum 2019)</li> <li>• SDDC Manager</li> </ul>
Credentials	
	<ul style="list-style-type: none"> <li>• Login credentials for Dell-Technologies support site</li> <li>• Login credentials for VMware support site</li> </ul>
VCF on VxRail Configuration Settings	
Reserve VLANs	<ul style="list-style-type: none"> <li>• One external management VLAN for VxRail Manager, vCenter Server, ESXi, SDDC Manager and other components deployed in management workload domain requiring external access</li> <li>• One internal management VLAN with IPV6 multicast for VxRail node auto-discovery and device management. The default is 3939. (This VLAN reservation can be bypassed if opting for manual node discovery.)</li> <li>• One VLAN with IPv4 unicast for vSAN traffic, unless planning for FC storage</li> <li>• One VLAN for vSphere vMotion</li> <li>• One VLAN for NSX-T Host Overlay network</li> <li>• One VLAN for the first NSX-T edge uplink (for NSX-T edge services)</li> <li>• One VLAN for the second NSX-T edge uplink (for NSX-T edge services)</li> <li>• One VLAN for the NSX-T edge overlay network (for NSX-T edge services)</li> <li>• One VLAN for iDRAC management of the VxRail nodes</li> </ul>

VCF on VxRail Configuration Settings	
	<ul style="list-style-type: none"> <li>• Determine default gateway and subnet mask.</li> <li>• Reserve four or more IP addresses for VxRail nodes for each VxRail cluster.</li> <li>• Reserve one IP address for vCenter Server.</li> <li>• Reserve one IP address for VxRail Manager.</li> <li>• Decide whether you want to use the default TCP-IP stack for vMotion, or a separate IP addressing scheme for the dedicated vMotion TCP-IP stack.</li> <li>• Reserve four or more IP addresses and a subnet mask for vSphere vMotion.</li> <li>• Select the gateway for either the default TCP-IP stack, or the dedicated vMotion TCP-IP stack.</li> <li>• Reserve four or more IP addresses and a subnet mask for vSAN, unless using external storage for VI workload.</li> <li>• Reserve IP address for SDDC Manager.</li> <li>• Reserve IP addresses for NSX-T Management VIP and appliance nodes.</li> <li>• Reserve IP addresses for the first NSX-T edge uplink (for NSX-T edge services).</li> <li>• Reserve IP addresses for the second NSX-T edge uplink (for NSX-T edge services).</li> <li>• Reserve IP addresses for the NSX-T edge overlay network (for NSX-T edge services).</li> <li>• Reserve IP addresses for the NSX-T host overlay network (unless using DHCP).</li> <li>• If witness is required for stretched cluster, reserve one IP address for the management network and one IP address for the vSAN network.</li> <li>• If NSX-T Federation is a requirement, reserve IP addresses for the remote TEPs on the edge gateways in each region.</li> </ul>
Reserve Hostnames	<ul style="list-style-type: none"> <li>• Determine parent and child DNS domains.</li> <li>• Decide on your VxRail host naming scheme. The naming scheme is applied to all VxRail hosts.</li> <li>• Reserve hostname for vCenter Server</li> <li>• Reserve hostname for VxRail Manager</li> <li>• Reserve hostname for SDDC Manager</li> <li>• Reserve hostnames for NSX-T Management VIP and appliance nodes.</li> </ul>
Passwords	<ul style="list-style-type: none"> <li>• Determine password structure following VMware password policy.</li> <li>• Select passwords for VxRail management components.</li> <li>• Select passwords for NSX-T Data Center.</li> <li>• Select passwords for SDDC Manager.</li> </ul>
Prepare Data Center Services	
Prepare DNS	<ul style="list-style-type: none"> <li>• Configure forward and reverse DNS records for VxRail Manager.</li> <li>• Configure forward and reverse DNS records for vCenter Server.</li> <li>• Configure forward and reverse DNS records for all VxRail nodes.</li> <li>• Configure forward and reverse DNS records for SDDC Manager.</li> <li>• Configure forward and reverse DNS records for NSX-T Management Cluster.</li> </ul>
Prepare DHCP	<ul style="list-style-type: none"> <li>• Configure IP address scope for NSX-T host overlay network (unless using static IP addresses).</li> </ul>

VCF on VxRail Configuration Settings	
Prepare Active Directory	<ul style="list-style-type: none"> <li>• If a use case for Cloud Foundation on VxRail include vRealize Suite to support a future VI workload domain, Active Directory must be deployed in the data center to support this requirement.</li> </ul>
Prepare Leaf Switches	<ul style="list-style-type: none"> <li>• Configure at least 1600 MTU (9000 preferred).</li> <li>• Configure the required VLANs on the top-of-rack switches.</li> <li>• Configure Layer 3 settings on VxRail external management network VLAN.</li> <li>• Configure Layer 3 settings on NSX-T host overlay network.</li> <li>• Configure Layer 3 settings on NSX-T edge overlay network (for NSX-T edge services).</li> <li>• Configure the switch ports to be directly connected to the VxRail nodes as Layer 2 trunk ports.</li> <li>• Configure unicast on the vSAN network.</li> <li>• Configure multicast on the VxRail internal management network, unless manually assigning VxRail management IP addresses.</li> <li>• Configure MLD snooping and MLD querier on the VxRail internal management network, unless manually assigning VxRail management IP addresses.</li> <li>• Configure Spanning Tree on the switch ports supporting VxRail nodes as edge ports, or in 'portfast' mode.</li> <li>• Configure inter-switch links on switches below the Layer 2/3 boundary.</li> </ul>
Prepare Routing Services	<ul style="list-style-type: none"> <li>• Configure Border Gateway Protocol at the Layer 2/3 network boundary.</li> <li>• Configure BGP peering with NSX-T Tier-0 Gateway (for NSX-T edge services)</li> </ul>
External Storage (if applicable)	<ul style="list-style-type: none"> <li>• Configure LUN or LUNs on FC storage array</li> <li>• Perform zoning and masking to present LUNs to VxRail nodes</li> </ul>



## Appendix B: Cloud Foundation on VxRail footprints for sizing

Use these tables to obtain footprint estimates of the resources for Cloud Foundation on VxRail.

Base Virtual Machines for every Cloud Foundation Management workload domain

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	SDDC Manager	4	16	800
Management	vCenter	8	24	500
Management	NSX-T Manager 1	6	24	200
Management	NSX-T Manager 2	6	24	200
Management	NSX-T Manager 3	6	24	200

**Note:** NSX-T Manager and NSX-T edge devices can be deployed in three sizes: Small, Medium, Large. Cloud Builder deploys the 'Medium' sized NSX-T Manager virtual appliances into the management workload domain.

Virtual Machines deployed in Cloud Foundation Management workload domain for each Cloud Foundation VI workload domain

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	vCenter	8	24	500

Virtual Machines deployed in Cloud Foundation VI Workload Domain to support the Application Virtual Network. The default size is 'Medium'.

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Workload	NSX-T Edge 1	4	8	120
Workload	NSX-T Edge 2	4	8	120

Virtual Machines deployed in Cloud Foundation Management workload domain for each Cloud Foundation VI workload domain which does not use a shared NSX-T management instance. The default size is 'Large'.

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	NSX-T Manager 1	8	32	200
Management	NSX-T Manager 2	8	32	200
Management	NSX-T Manager 3	8	32	200

Virtual Machines deployed in Cloud Foundation Management workload domain for each Cloud Foundation VI workload domain for Kubernetes which does not use a shared NSX-T management instance. The default size is 'Large'.

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	NSX-T Manager 1	8	32	200
Management	NSX-T Manager 2	8	32	200
Management	NSX-T Manager 3	8	32	200

Virtual Machines deployed in one of the Cloud Foundation Management workload domains to support NSX-T Federation across regions. The size of the virtual machines depends on the size of the federation under management, either Medium or Large.

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	NSX-T Global Manager 1	6/12	24/48	300
Management	NSX-T Global Manager 2	6/12	24/48	300
Management	NSX-T Global Manager 3	6/12	24/48	300

The following table lists the sizing to prepare for a vRealize Suite Lifecycle Manager download and deployment.

---

**Note:** Cloud Foundation on VxRail does not automate the deployment or the life cycle management of the other vRealize Suite components. vRealize Suite Lifecycle Manager is used to deploy and manage those components.

---

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	vRealize Suite Lifecycle Manager	4	16	135

The following table lists the sizing to prepare for the deployment of vSphere with Tanzu workload domain.

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Workload	Supervisor Cluster control plane	12	48	200
Workload	Registry Service	7	7	200
Workload	NSX-T Edge 1	8	32	200
Workload	NSX-T Edge 2	8	32	200
Workload	Tanzu Kubernetes Cluster control plane	6	12	48
Workload	Tanzu Kubernetes Cluster worker nodes	6	12	48

## Appendix C: Cloud Foundation on VxRail VLANs

Following are the core VLANs that must be configured on the data center switches supporting the Cloud Foundation on VxRail platform Cloud Foundation on VxRail VLANs.

Category	VLAN	Description
VxRail	External Management	VxRail and Cloud Foundation components
	Internal Management	VxRail device discovery
	vMotion	Virtual machine migration
	vSAN	vSphere datastore
NSX-T	Host Overlay	NSX-T VTEP. Must be able to reach DHCP server to assign IP addresses for host overlay network, if using DHCP for this purpose.
Node Management	Out-of-band management	Dell PowerEdge iDRAC network (optional)

## Appendix D: VxRail network configuration

The following table lists the configuration settings required by VxRail Manager to deploy a VxRail cluster.

Category	Detail	Description
VxRail	Management Network	VxRail and Cloud Foundation management network subnet. Must be large enough for all VxRail and Cloud Foundation management components.
	External Management	VLAN ID for the management network that passes upstream from the top-of-rack switches
	Internal Management	VLAN ID for VxRail device discovery. This network stays isolated on the top-of-rack switches. The default VLAN ID is 3939.
System	Global settings	Time zone
		NTP server(s) IP Address
		DNS server(s) IP Address
		Syslog Server
Management	ESXi hostnames and IP addresses	ESXi hostname prefix
		Separator
		Iterator
		Offset
		Suffix
		Domain
		IP address pool for VxRail nodes
	vCenter Server	VxRail vCenter Server hostname
		VxRail vCenter Server IP address
	VxRail Manager	VxRail hostname
		VxRail IP address
	Networking	Subnet mask
		Gateway
vMotion		IP address pool for vMotion
		Gateway
		Subnet mask
		VLAN ID
vSAN		IP address pool for vSAN
		Subnet mask
		VLAN ID
Dell Node	iDRAC	IP address for iDRAC port on each VxRail node

The following table applies to stretched clusters only.

Category	Detail	Description
Witness	Management	Hostname
		IP Address
		Subnet Mask
		Gateway
	VSAN	IP Address
		Subnet Mask
		Gateway
Witness Site	vSphere Host	IP Address
Network	Witness Traffic Separation	Optional VLAN ID to manage traffic between two sites hosting VxRail nodes and witness site
Second Site	vMotion	IP address pool for vMotion
		Subnet mask
		VLAN ID
	vSAN	IP address pool for vSAN
		Subnet mask
		VLAN ID
	VxLAN	IP address pool for VxLAN
		Subnet mask
		VLAN ID

## Appendix E: Cloud Builder and management VI workload configuration

This table lists the configuration settings that are required by Cloud Builder to deploy the Cloud Foundation management workload domain on the VxRail cluster platform.

Category	Detail	Description
Cloud Builder	IP Address	Temporary for Cloud Builder virtual appliance
Global	NTP	IP Address
	DNS	IP Address
	SSO Site Name	Must be the same site name as used for the VxRail cluster.
	SSO Domain	
	DNS Zone Name	
SDDC	Manager	Hostname
		IP Address
		Domain Name
NSX-T	Manager (VIP)	Hostname
		IP Address
	Manager Node 1	Hostname
		IP Address
	Manager Node 2	Hostname
		IP Address
	Manager Node 3	Hostname
		IP Address
	Appliance Size	(Small, Medium, Large)
NSX-T Host Overlay Network	Static IP Assignment Method	Name of static IP address pool
		IP address range in CIDR format
		Starting IP address to be assigned for host overlay network
		Ending IP address to be assigned for host overlay network
		Gateway
	Dynamic IP Assignment Method	IP address of DHCP server to assign IP addresses to VTEP tunnel endpoints for host overlay network
		Range of IP addresses in DHCP server to be assigned to VTEP tunnel endpoints for host overlay network
vSphere Objects	Data Center Name	Must match value in VxRail cluster
	Cluster Name	Must match value in VxRail cluster

Category	Detail	Description
	VxRail Distributed Switch Name(s)	Must match values used in VxRail cluster
	NSX-T Distributed Switch Name	If deploying separate VDS for NSX-T networking. VDS name must be unique in VxRail cluster.
	vSAN Datastore Name	Must match value used in VxRail cluster
vSphere Resource Pools	SDDC Management	Required for consolidated architecture
	SDDC Edge	Required for consolidated architecture
	User Edge	Required for consolidated architecture
	User VM	Required for consolidated architecture

## Appendix F: VI workload domain configuration settings

This table lists the configuration settings required to support the configuration of a standard VI workload domain by SDDC Manager.

Category	Detail	Description
Global	Domain Name	
	Datacenter Name	Name of vSphere data center to be configured in vCenter instance supporting VI workload domain
Management	vCenter	Hostname
		IP Address
		Subnet Mask
		Default gateway
		Management Account
NSX-T Host Overlay Network	Static IP Assignment Method	Name of static IP address pool
		IP address range in CIDR format
		Starting IP address to be assigned for host overlay network
		Ending IP address to be assigned for host overlay network
		Gateway
	Dynamic IP Assignment Method	IP address of DHCP server to assign IP addresses to VTEP tunnel endpoints for host overlay network
		Range of IP addresses in DHCP server to be assigned to VTEP tunnel endpoints for host overlay network

Use this table only if a new NSX-T edge cluster is deployed as part of the VI workload domain.

Category	Detail	Description
NSX-T ASN	Autonomous System Number	ASN for the workload domain edge cluster
NSX-T Manager	NSX-T management cluster	IP address of NSX-T Manager VIP
		IP address for first NSX-T Manager
		IP address for second NSX-T Manager
		IP address for third NSX-T Manager
		Subnet Mask
		Default Gateway
NSX-T Edge Node 1	Name	Hostname of virtual appliance
	Management IP Address	Must be within workload domain management network subnet
	Uplink 1 IP Address	IP address for BGP peering on first NSX-T edge uplink VLAN for this workload domain



Category	Detail	Description
	Uplink 2 IP Address	IP address for BGP peering on second NSX-T edge uplink VLAN for this workload domain
	Overlay IP Address	IP address for overlay network between edge nodes in this workload domain
NSX-T Edge Node 2	Name	Hostname of virtual appliance
	Management IP Address	Must be within workload domain management network subnet
	Uplink 1 IP Address	IP address for BGP peering on first NSX-T edge uplink VLAN for this workload domain
	Uplink 2 IP Address	IP address for BGP peering on second NSX-T edge uplink VLAN for this workload domain
	Overlay IP Address	IP address for overlay network between edge nodes in this workload domain
NSX-T Edge Uplink 1	VLAN	Used for BGP peering with upstream routing services for this workload domain
NSX-T Edge Uplink 2	VLAN	Used for BGP peering with upstream routing services for this workload domain
NSX-T Edge Overlay Network	VLAN	Used for edge overlay network connecting NSX-T edge nodes in this workload domain

If a new edge cluster is deployed, use this table to capture the BGP neighbors for the NSX-T edge gateway.

Category	Detail	Description
External ASN	ASN Value	Autonomous System Number for external routers
External Router 1	IP Address	IP Address for peering with NSX-T edge gateway on first NSX-T edge uplink VLAN
	Password	Neighbor password for BGP peering
External Router 2	IP Address	IP Address for peering with NSX-T edge gateway on second NSX-T edge uplink VLAN
	Password	Neighbor password for BGP peering

Use this table only if a vSphere for Tanzu supervisor cluster will be configured on the VI workload domain.

Detail	Type	Description
Pod CIDRs	Internal	Used by Kubernetes pods that run in the cluster
Service CIDRs	Internal	Used by Kubernetes applications that need a service IP address
Ingress CIDRs	External	Used for load balancing
Egress CIDRs	External	Used for NAT endpoint use

## Appendix G: Edge Gateway configuration

This table lists the configuration settings required to support deployment of the NSX-T edge gateways

External Routers		
External ASN	ASN Value	Autonomous System Number for external routers
External Router 1	IP Address	IP Address for peering with NSX-T edge gateway on first NSX-T edge uplink VLAN
	Password	Neighbor password for BGP peering
External Router 2	IP Address	IP Address for peering with NSX-T edge gateway on second NSX-T edge uplink VLAN
	Password	Neighbor password for BGP peering
NSX-T Edge Gateways		
Cluster	Name	Name of NSX-T Edge Cluster
Internal ASN	ASN Value	BGP Autonomous System Number for NSX-T Edge Gateways
Edge Node 1	Name	Hostname of virtual appliance
	Management IP Address	Must be within management network subnet range
	Uplink 1 IP Address	IP address for BGP peering on first NSX-T edge uplink VLAN
	Uplink 2 IP Address	IP address for BGP peering on second NSX-T edge uplink VLAN
	Overlay IP Address	IP address for overlay network between edge nodes
Edge Node 2	Name	Hostname of virtual appliance
	Management IP Address	Must be within management network subnet range
	Uplink 1 IP Address	IP address for BGP peering on first NSX-T edge uplink VLAN
	Uplink 2 IP Address	IP address for BGP peering on second NSX-T edge uplink VLAN
	Overlay IP Address	IP address for overlay network between edge nodes
Edge Gateway VLANs	Uplink 1 VLAN	First NSX-T edge uplink
	Uplink 2 VLAN	Second NSX-T edge uplink
	Edge Overlay VLAN	Used for edge overlay network connecting NSX-T edge nodes
Second Site - External Routers		
External ASN	ASN Value	Autonomous System Number for external routers
External Router 1	IP Address	IP Address for peering with NSX-T edge gateway on first NSX-T edge uplink VLAN
	Password	Neighbor password for BGP peering
External Router 2	IP Address	IP Address for peering with NSX-T edge gateway on second NSX-T edge uplink VLAN
	Password	Neighbor password for BGP peering

## Appendix H: Application Virtual Network configuration

This table lists the configuration settings required to support deployment of the optional Application Virtual Network.

Application Virtual Network Regions		
Region A	Logical Segment	Name of Region A logical segment
	VLAN	VLAN for Region A network
	IP Addresses	IP address range for Region A network
xRegion	Logical Segment	Name of xRegion logical segment
	VLAN	VLAN for xRegion network
	IP Address	IP address range for xRegion network

## Appendix I: Sample switch configuration settings

This set of sample syntax is for providing basic guidance on the settings that must be performed on the top-of-rack switches to configure VLANs and a switch port for a Cloud Foundation on VxRail deployment, and configuring a switch support BGP peering for the Application Virtual Network (AVN). The actual code that is required on the top-of-rack switches depends on the existing data center network infrastructure, switch operating system and routing standards.

The sample syntax highlights the following required items:

- VLAN for VxRail external management
- VLAN for VxRail internal management – node discovery
- VLAN for NSX-T host overlay network – with DHCP helper
- Switch port configuration for VxRail node

```
interface vlan <VxRail External Management>
no shutdown
ip address <gateway>/24
vrrp-group <id>
priority <priority>
virtual-address <virtual gateway>

interface vlan <VxRail Internal Management>
no shutdown
ipv6 mld snooping querier

interface <Host Overlay>
description
no shutdown
mtu 9216
ip address <gateway>/24
ip helper-address <DHCP server IP Address>
vrrp-group <id>
priority <priority>
virtual-address <virtual gateway>

interface ethernet <port>
no shutdown
switchport mode trunk
switchport access vlan <native vlan>
```

The sample syntax highlights the following required items:

- VLANs for AVN uplinks with assigned IP addresses
- Prefix list to permit route filtering for routes from the AVN
- BGP External ASN setting
- BGP peering with the pair of AVN Edge Service Gateways

```
interface vlan <VLAN for AVN Uplink>
no shutdown
mtu 9216
ip address <Gateway IP address for AVN uplink>

ip prefix-list <Router-ESGs route map name> permit <IP address range parameters>

router bgp <External ASN>
maximum-paths ebgp 4
router-id <External router ID>
address-family ipv4 unicast
redistribute connected route-map <Router-ESG route map name>

template external-router-to-ESG
advertisement-interval <value>
password <password saved to Edge Gateways>
timers 4 12

neighbor <IP address assigned to first Edge Gateway>
inherit template external-router-to-ESG
remote-as <ASN assigned to Edge Gateways>
no shutdown

neighbor <IP address assigned to second Edge Gateway>
inherit template external-router-to-ESG
remote-as <ASN assigned to Edge Gateways>
no shutdown
```