

# VMware Cloud Foundation on VxRail Planning and Preparation Guide

## Abstract

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

March 2020

# Revisions

Date	Description
April 2019	Initial release
May 2019	Minor edits to provide more clarity and detail
August 2019	Updates to support VMware Cloud Foundation 3.8
March 2020	Updates to support VMware Cloud Foundation 3.9.1

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

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

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [3/27/2020] [Planning Guide] [H17730]

# Table of contents

1	VMware Cloud Foundation on VxRail Product Overview.....	6
2	VMware Cloud Foundation on VxRail Deployment.....	9
3	Data Center Requirements.....	13
3.1	Physical rack space.....	13
3.2	Physical networking.....	13
3.2.1	Switch port capacity.....	13
3.2.2	Switch port type.....	14
3.2.3	Jumbo frames.....	14
3.2.4	Multicast.....	14
3.2.5	Border Gateway Protocol.....	14
3.2.6	Hardware VTEP (Tunnel endpoint).....	15
3.3	Network services.....	15
4	Use Cases and High-Level Design.....	17
4.1	Hybrid cloud or virtual desktop delivery objective.....	17
4.2	Site locations.....	17
4.3	Application availability.....	18
5	Cloud Foundation on VxRail Workload Planning.....	19
5.1	Determine use cases for Cloud Foundation VI workload domain.....	19
5.2	Deciding on single-site VxRail cluster or stretched cluster.....	19
5.3	Planning the Management workload domain Resources.....	20
5.4	Planning the VI workload domain Resources.....	20
5.5	Sizing the Cloud Foundation Domains.....	21
6	Application Dependencies and Routing Decisions.....	24
7	Cloud Foundation on VxRail Physical Network Planning.....	26
7.1	Select a physical network architecture and topology.....	26
7.2	VxRail Cluster Physical Networking Planning.....	28
7.3	AVN (Application Virtual Network) Physical Network Planning.....	29
7.4	VxRail Stretched Cluster Physical Network Planning.....	32
7.5	Cloud Foundation Domain Physical Network Planning.....	34
7.5.1	Identify the required logical layer 2 networks.....	34
7.5.2	Configure VLANs for each Cloud Foundation domain.....	35
7.5.3	Configure overlay network settings.....	35
7.5.4	Configure routing services.....	35
7.5.5	Identify multicast IP addresses for VXLAN.....	35
7.5.6	Deploy DHCP server for VXLAN Tunnel Endpoints.....	36

8	VxRail Cluster Deployment Preparation.....	37
8.1	Prepare for VxRail cluster initial build.....	37
8.2	Capture and record the VLANs for VxRail cluster.....	37
8.3	Capture and record the network settings for VxRail cluster.....	38
8.4	Capture and record the network settings for VxRail stretched cluster.....	38
8.5	Create Forward and Reverse DNS Entries for VxRail cluster.....	38
8.6	Prepare top-of-rack switches for VxRail cluster.....	38
8.7	Prepare passwords.....	39
9	Prepare for VMware Cloud Foundation Management VI workload domain.....	40
9.1	Provide a temporary IP address for Cloud Builder.....	40
9.2	Global network setting for management VI workload domain.....	40
9.3	Capture and record the network settings for the management VI workload domain.....	41
9.4	Capture and record the DHCP settings for VTEP tunnel endpoints for the management VI workload domain.....	41
9.5	Create Forward and Reverse DNS Entries for the management VI workload domain.....	41
9.6	Select and create the VXLAN VLAN.....	41
9.7	Capture and record the NSX settings for the management VI workload domain.....	41
9.8	Select a multicast IP address range for VXLAN network.....	42
9.9	Select names for resource pools in VI Management Domain.....	42
9.10	Prepare passwords.....	42
9.11	Obtain VMware license keys.....	42
10	Prepare for Cloud Foundation Application Virtual Network.....	43
10.1	Capture the Application Virtual Network Region Settings.....	43
10.2	Capture settings for BGP peering.....	44
10.3	Capture settings for Universal Distributed Logical Router.....	44
10.4	Capture settings for Edge Service Gateways.....	45
10.5	Capture external router settings for the eBGP peering.....	45
10.6	Capture Universal Logical Router settings.....	46
10.7	Capture Region A Management Component Settings.....	47
10.8	Capture 2nd Site settings for Stretched Cluster.....	47
11	Prepare for Cloud Foundation VI workload domain.....	49
11.1	Configure VI workload domain VXLAN.....	49
11.2	Capture settings for VI workload domain.....	50
12	Planning for NSX integration with physical network.....	51
12.1	Capture settings for upstream router in physical network.....	51
12.2	Capture settings for NSX Edge virtual devices.....	51
12.3	Capture settings for NSX logical router.....	51

A	Cloud Foundation on VxRail Footprints for Sizing .....	52
B	Cloud Foundation on VxRail VLANs .....	55
C	VxRail Network Configuration .....	56
D	Cloud Builder and Management VI workload domain Configuration .....	58
E	Application Virtual Network Configuration .....	59
F	VI workload domain Configuration Settings .....	61
G	NSX Routing Settings.....	62
H	Example switch configuration settings for BGP peering .....	63

# 1 VMware Cloud Foundation on VxRail Product Overview

VMware Cloud Foundation on VxRail is a jointly engineered solution between VMware and Dell EMC. The 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.

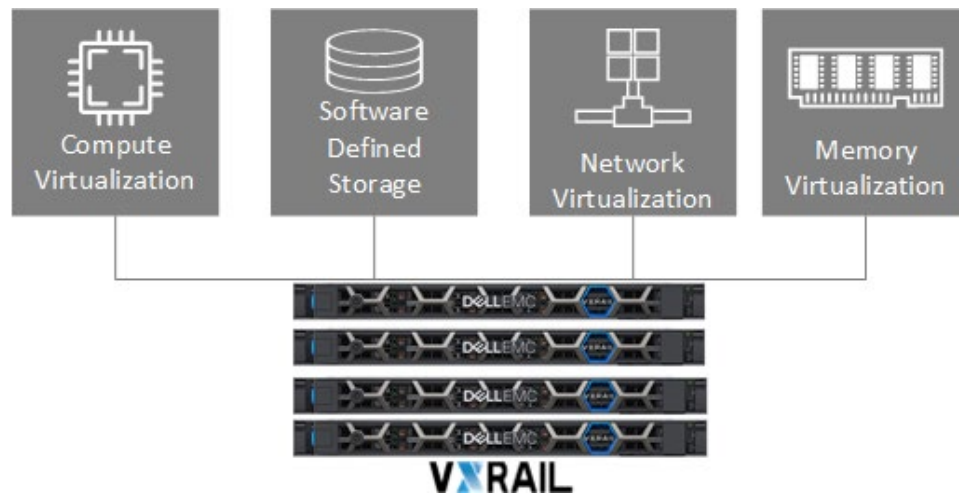


Figure 1 VxRail cluster representing a pool of virtual resources

VxRail integrates software products from VMware with custom software engineered from Dell EMC 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 is 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.

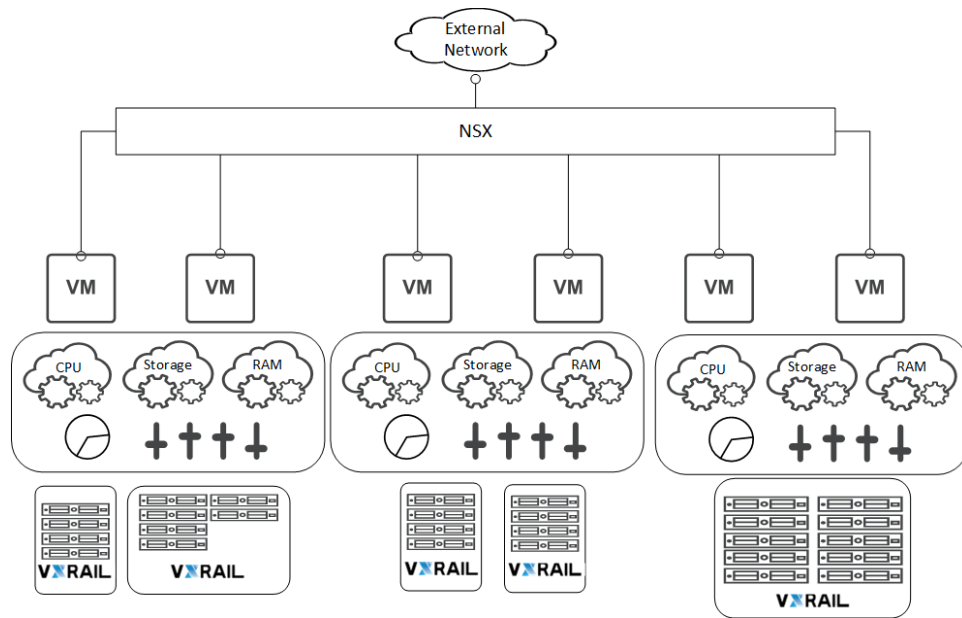


Figure 2 VxRail clusters as building blocks for Cloud Foundation virtual workload consumption

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 are embedded and enabled into each VI workload domain.

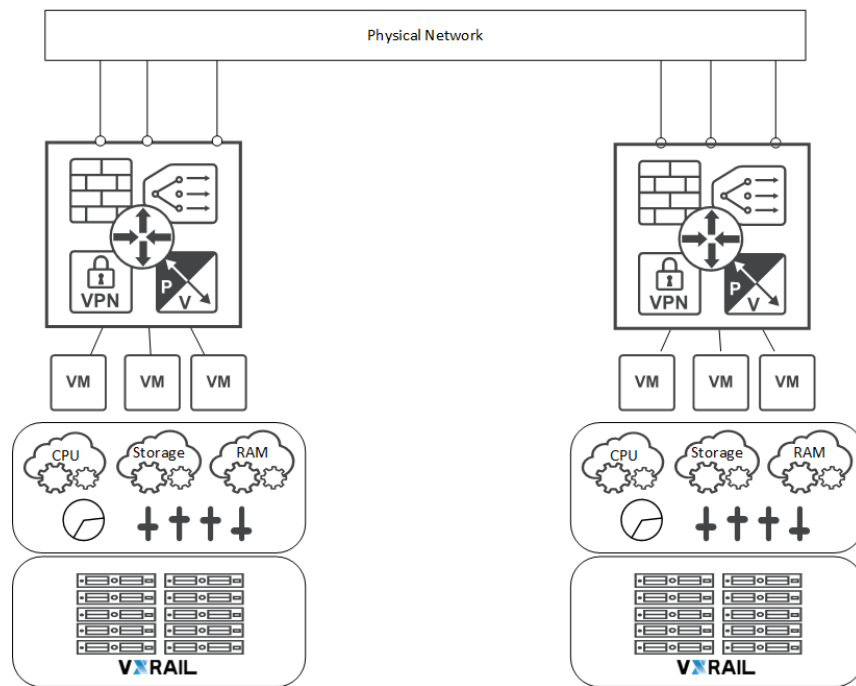


Figure 3 Cloud Foundation VI workload domains with fully virtualized resources

With support for logical routing by NSX, this means virtual machine traffic that previously had to pass through to the physical network can now be pushed down to the virtual network when established on a Cloud Foundation on VxRail VI workload domain.

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

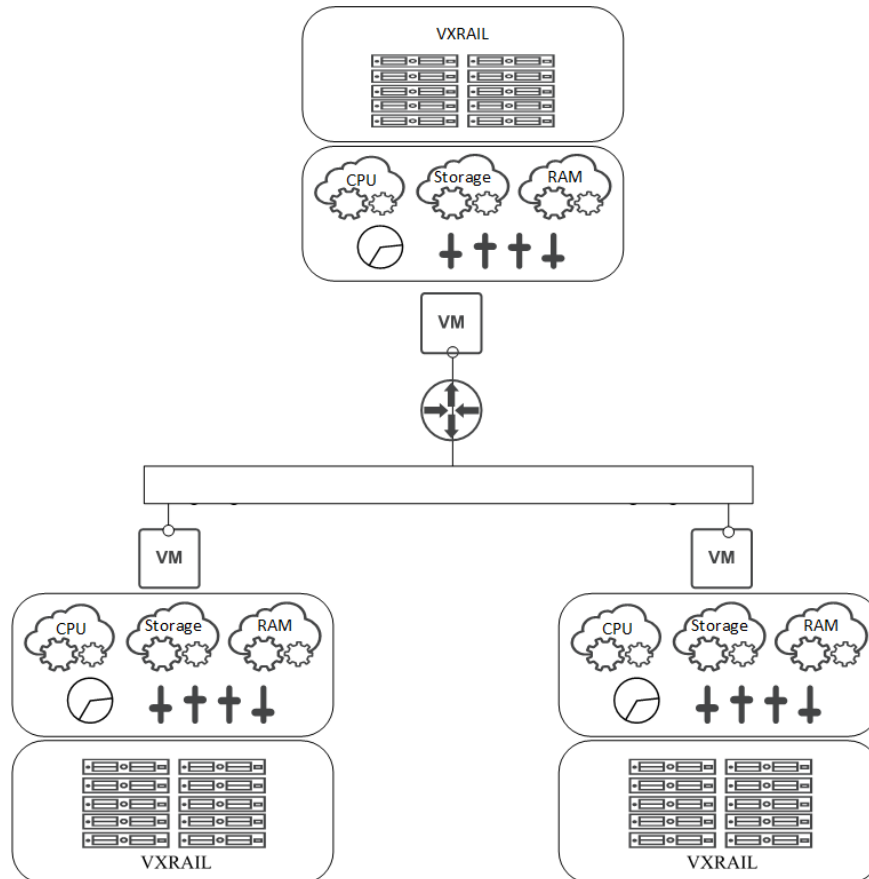


Figure 4 Virtual machines connected to an extended logical switch with routing services

If virtual machine connectivity requires routing, the extended logical switch can use routing services within the virtual network. To support connectivity outside of the virtual network, the virtual routing services will form a peer relationship with existing physical routers in the data center to form a seamless connection between the physical and logical networks.



## 2 VMware Cloud Foundation on VxRail Deployment

The Cloud Foundation on VxRail cloud platform in a data center will have 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 lifecycle 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 can be aligned with their application requirements to propose a high-level design. Dell EMC specialists work jointly with the account team at this stage of the effort.

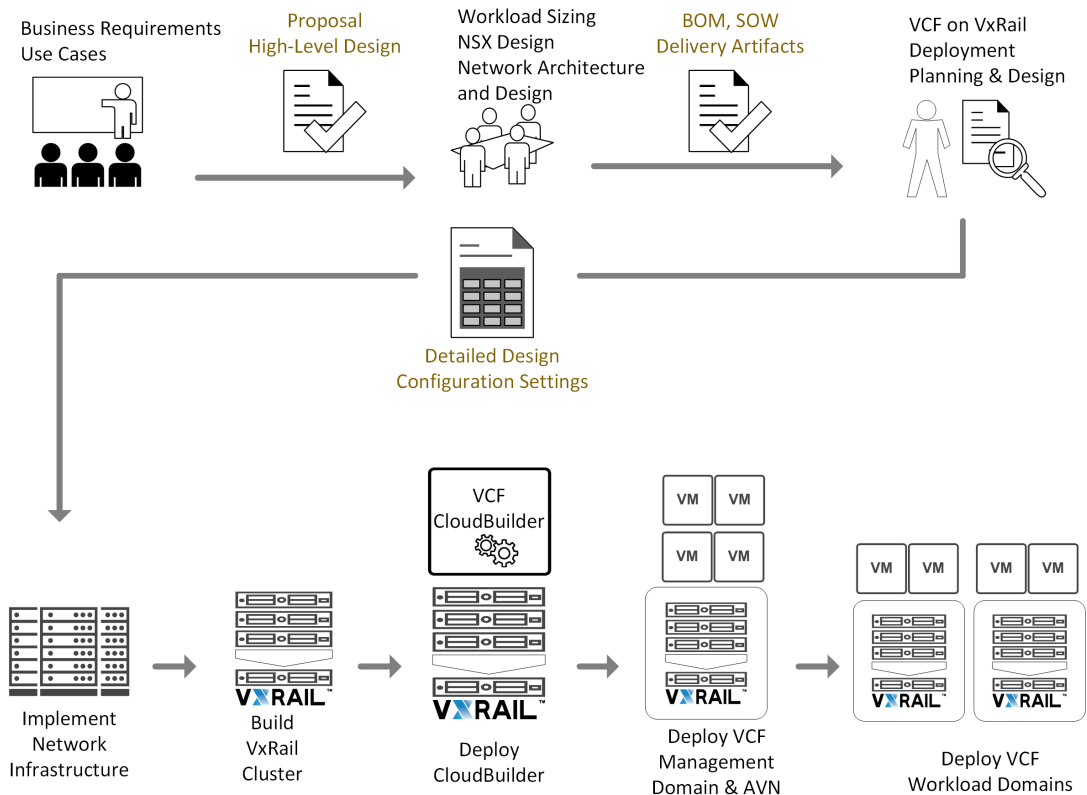


Figure 5 Anatomy of a Cloud Foundation on VxRail deployment experience

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 materials needed to support the planned workload. In addition, the dependencies between the planned sets of applications are analyzed, captured, and used to produce a supporting physical network architecture and design. The application dependency report and physical network design are then used as the baseline in the planning efforts for the virtual networks in the VI workload domains.

During the next phase, the work effort transitions to a professional services engagement. The planning and design phase will commence while awaiting equipment delivery from Dell EMC manufacturing to your site locations. The assigned solutions architect will capture the detailed design and configuration settings for the initial deployment of Cloud Foundation on VxRail. This will include the network settings for the planned VxRail clusters and the Cloud Foundation Cloud Builder virtual appliance. The solutions architect will also perform a validation of the data center environment to make sure all prerequisites are met.

If Dell EMC is responsible for the configuration of the supporting network infrastructure, that service will be performed after the physical hardware is installed and cabled in the data center. The VxRail clusters and VMware Cloud Foundation are dependent on the supporting network infrastructure to be properly configured, and for all required network services to be properly configured before moving to the next phase.

Dell EMC will next deploy the VxRail cluster targeted for the Cloud Foundation management workload domain using the information captured in the planning and design phase, and deploy the Cloud Foundation Cloud Builder virtual appliance on the VxRail cluster.

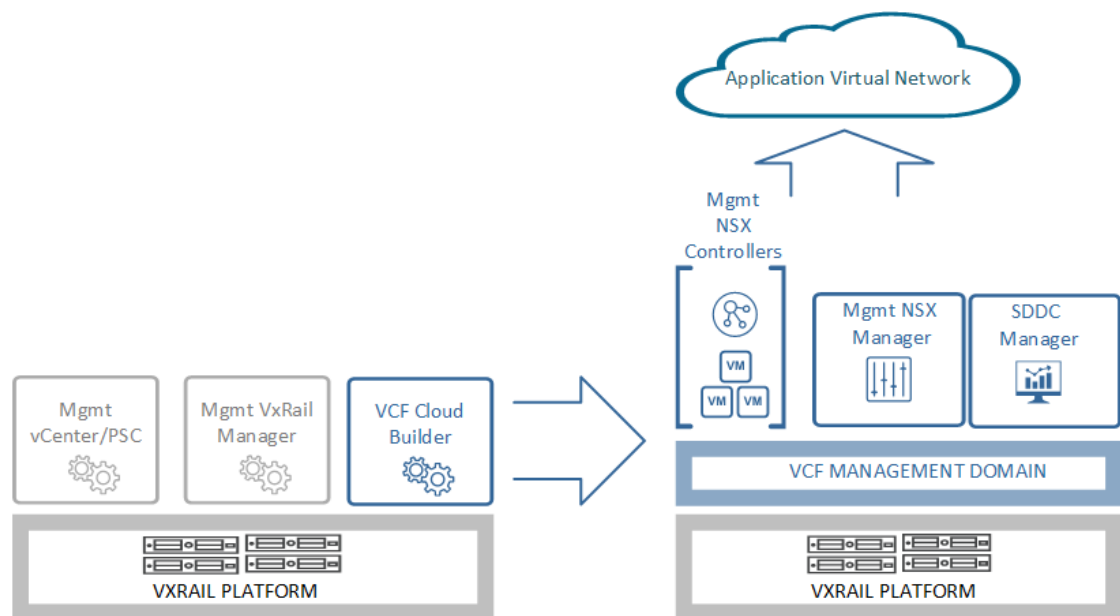


Figure 6 Overview of Cloud Builder automatic deployment of the Cloud Foundation management workload domain and Application Virtual Network (AVN)

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 and creates the Cloud Foundation management workload domain.

After the base deployment of the VI management domain and management components is completed, CloudBuilder will then create a multi-region virtual network environment called the Application Virtual Network (AVN). This configured network is reserved for supporting the vRealize management software suite. CloudBuilder will create two regions within the Application Virtual Network during the deployment process, named 'Region A' and 'xRegion'. The Log Insight virtual components are then deployed into 'Region A', with the deployment of the vRealize management software suite is performed after completion of all CloudBuilder tasks. CloudBuilder will then integrate the Application Virtual Network into the upstream data center network using physical and virtual routing services. Upon completion of all automated deployment tasks, the management workload domain for the Cloud Foundation on VxRail and the base deployment of the Application Virtual Network is complete, and the Cloud Builder appliance is removed.

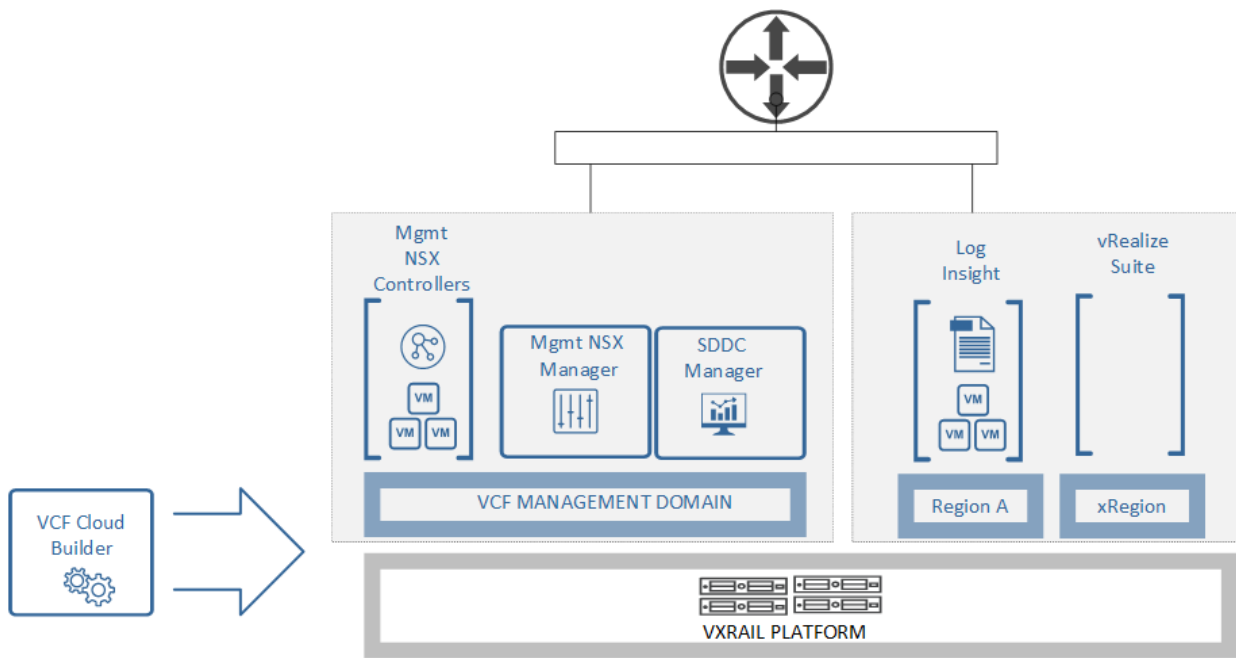


Figure 7 Overview of Cloud Builder automatic deployment of the Application Virtual Network (AVN) regions

The next phase is the deployment of the Cloud Foundation VI workload domains. For this phase, the VI workload domain is initialized using SDDC Manager. The underlying VxRail clusters are deployed and then assigned to the VI workload domain.

The initialization of a VI workload domain using SDDC Manager lays down the basic foundation for the future deployment of virtual machines and their network interconnections. At this stage, a decision must be made as to whether the virtual network supporting the VI workload domain will be based on NSX-V or NSX-T, as the deployment process will differ based on this decision.

Once completed, the virtual network design for the VI workload domain can move from the planning stage to the deployment stage for supporting applications.

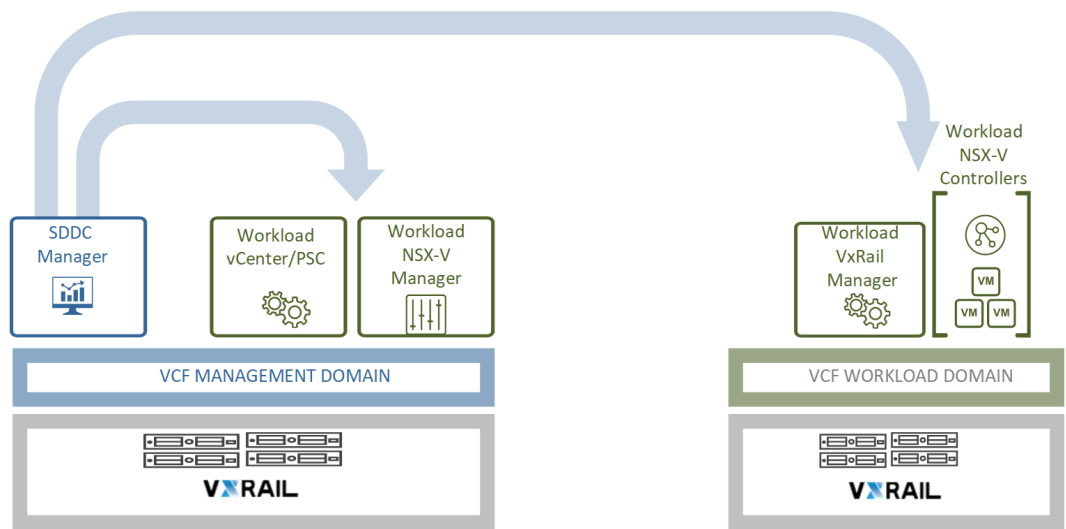


Figure 8 Overview of initial deployment of NSX-V based VI workload by SDDC Manager

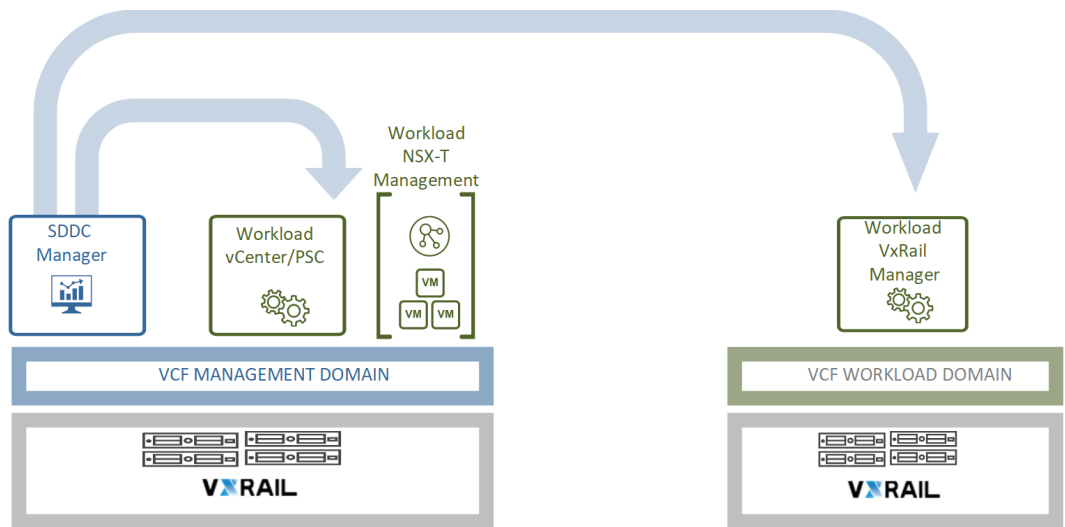


Figure 9 Overview of initial deployment of NSX-T based VI workload by SDDC Manager

If the requirements for Cloud Foundation on VxRail include leveraging the vRealize software to support any workload domains, then the software can download to the Application Virtual Network xRegion after the deployment of the management domain is complete.

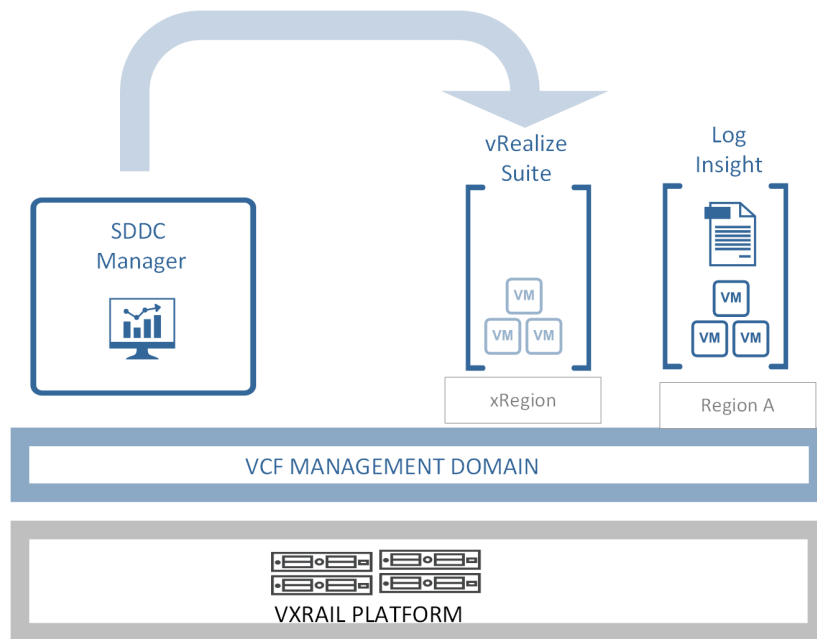


Figure 10 Deploying vRealize management suite into AVN xRegion

## 3 Data Center Requirements

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

### 3.1 Physical rack space

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 is dependent on the model you select to support your Cloud Foundation on VxRail platform. Dell EMC will go through a sizing exercise to produce a bill of materials 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	750W, 1100W	C14
P-Series	2	1100W, 1600W	C14
S-Series	2	1100W, 1600W	C14
V-Series	2	2000W	C20
G-Series	2	2000W, 2400W	C20

The amount of rack space required for the supporting physical network is dependent 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.

### 3.2 Physical networking

You can either bundle a Dell network infrastructure with your Cloud Foundation on VxRail for a single source solution, or select, implement and configure your own supporting network. If you choose a network infrastructure based on Dell network switches, your Dell EMC 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.

#### 3.2.1 Switch port capacity

To support NSX-T in Cloud Foundation workload domains will require four connections per node to your supporting physical network. If your requirements do not include NSX-T, and will depend on NSX-V only, then each node deployed to support Cloud Foundation on VxRail will require a minimum of two connections to your supporting physical network. For fault protection, each node port connects to a separate physical switch.




Figure 11 VxRail node physical network connections

### 3.2.2 Switch port type

VxRail nodes can support either SFP+ or RJ45 network connections. The ports on the physical switch supporting your Cloud Foundation on VxRail cloud platform must match the network type on the VxRail nodes.

### 3.2.3 Jumbo frames

Both NSX-V and NSX-T depend 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. While NSX-V is based on the VXLAN (Virtual Extensible LAN) standard, and NSX-T depends on the GENEVE (GENeric NETwork Virtualization) standard, both depend on 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.

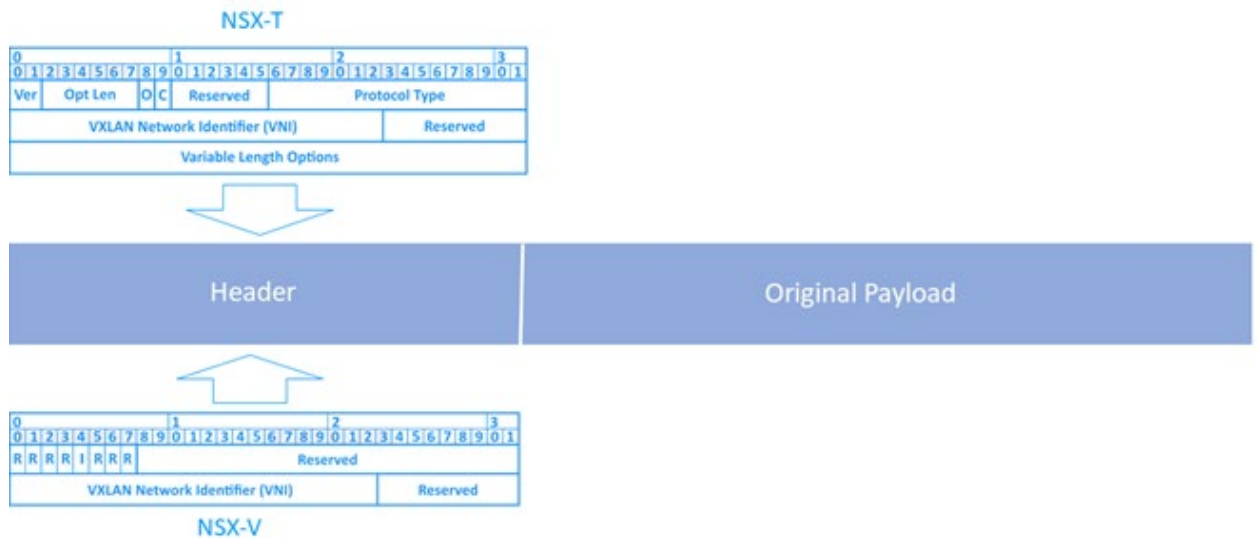


Figure 12 NSX-V and NSX-T extended frames

### 3.2.4 Multicast

NSX-V depends on multicast to support replication requirements for your VXLAN networks. Cloud Foundation on VxRail enables hybrid replication by default to use multicasting only when required. In addition, the physical switch must support IGMP snooping and snooping querier to optimize multicast traffic.

### 3.2.5 Border Gateway Protocol

Cloud Foundation on VxRail leverages a gateway at the edge between the physical and virtual networks to serve as the boundary point. 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 this boundary, the gateway at the edge must be able to peer with an upstream router in the physical network that supports Border Gateway Protocol.

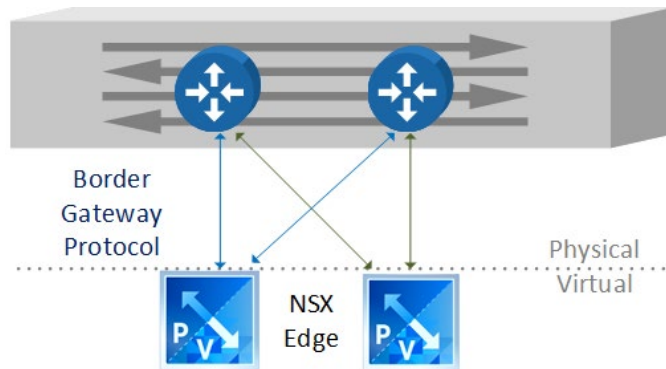


Figure 13 Physical and virtual route peering with Border Gateway Protocol

---

Note that only the network devices that define your layer 2 and layer 3 boundaries with Cloud Foundation on VxRail require support for Border Gateway Protocol.

---

### 3.2.6 Hardware VTEP (Tunnel endpoint)

Dell EMC strongly recommends selecting network switches that support hardware-based VTEP (Virtual Tunnel Endpoints). This feature is beneficial for customers expecting to deploy VxRail over multiple racks. The feature, when enabled on the switch, supports the bridging of Layer 2 network traffic from a logical vSphere switch between physical Ethernet ports on the VxRail nodes through packet encapsulation and decapsulation on a Layer 3 overlay network. While this feature is not required for a single-rack deployment, it is required if you expand into a multi-rack deployment of Cloud Foundation on VxRail.

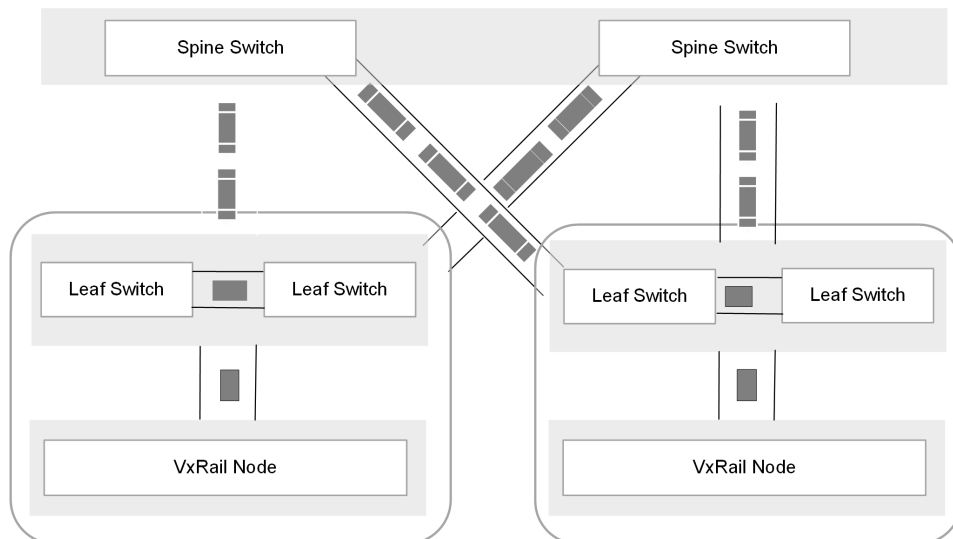


Figure 14 Sample VTEP tunnel network supporting multi-rack deployment

## 3.3 Network services

The network services listed in this section are required in the host data center for your Cloud Foundation on VxRail deployment.

- Active Directory
- Domain Name Service (DNS) – You will need to 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) – For Cloud Foundation on VxRail, DHCP is required for the automated allocation of IP addresses for the VXLAN Tunnel Endpoints. A DHCP server must be deployed in the host data center, and be pre-populated with the IP addresses to be assigned in order to support connectivity on the VXLAN network.

The following network services are optional, but recommended:

- Simple Message Transfer Protocol (SMTP)
- Certificate Authority (CA)



## 4 Use Cases and High-Level Design

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

### 4.1 Hybrid cloud or virtual desktop delivery objective

If your end objective for Cloud Foundation on VxRail is to serve as a platform for a hybrid cloud with the vRealize software suite, Dell EMC and VMware will guide you through this phase. Also, if your end objective is for Cloud Foundation on VxRail to serve as a platform for virtual desktops with the Horizon software suite, Dell EMC and VMware will also guide you through this phase.

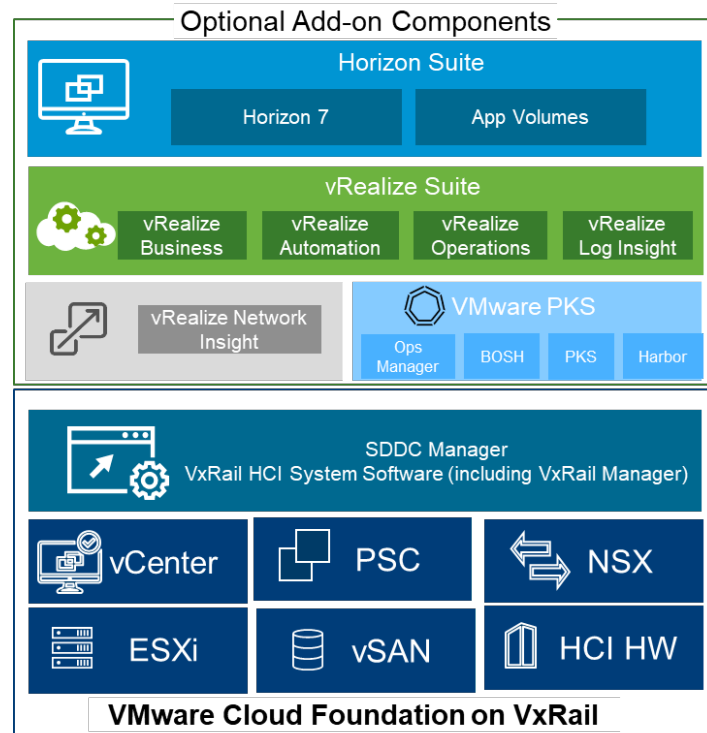


Figure 15 vRealize, PKS and Horizon suites layered on Cloud Foundation on VxRail

### 4.2 Site locations

The requirement for multiple site locations for Cloud Foundation on VxRail deployments will 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.

## 4.3 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 very 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.

## 5 Cloud Foundation on VxRail Workload Planning

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 start with as few as three nodes and 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, refer to: [VxRail 14G Series Specification Sheet](#)

The mixing of different server configurations 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 pre-defined baseline, and adapt and expand as necessary for changing workload requirements.

### 5.1 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 will impact the resources required to support the workload planned for each domain.

If plans include the deployment of any additional layers of software, such as VMware's vRealize product suites, this action will impact the overall size of the management workload domain and associated VI workload domains. The overhead required to support these product suites needs to be taken into consideration with the sizing effort.

### 5.2 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 16 Impact of VxRail stretched cluster on virtual machine workload

In addition, 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.

### 5.3 Planning the Management workload domain Resources

Guest virtual machines cannot be deployed in the Cloud Foundation management workload domain. This domain is reserved for management purposes only, and only virtual machines needed to administer the cloud platform are deployed in this domain. At initial deployment of Cloud Foundation on VxRail, after the first VxRail cluster is built, Dell EMC will deploy a temporary virtual appliance named 'Cloud Foundation Cloud Builder' on the cluster. This tool, when activated, layers the Cloud Foundation management workload domain on the VxRail cluster, and in the same process, deploys an initial core set of virtual machines to support overall Cloud Foundation on VxRail management.

This initial set of core virtual machines from Cloud Builder provides a baseline on the minimum resources needed to bring up the management workload domain and begin deploying VI workload domains. However, for each VI workload domain that is created, a base of management virtual machines is deployed by SDDC Manager and placed in the management workload domain and in the VI workload domain. Therefore, in order to get more accurate sizing guidelines for the management workload domain, the planned use cases for the VI workload domains must be understood, and a resource consumption assessment completed for the VI workload domains, before determining the required size of the management workload domain.

If VMware vRealize functionality is part of the overall plan, then additional resources need to be reserved for those virtual machines. In addition, a workload domain deployed to support Horizon or PKS will also require additional resources. The tables in the [Cloud Foundation on VxRail Footprints for Sizing](#) appendix can be used to provide estimates of the minimum sizing requirements for the management components, based on planned use cases.

### 5.4 Planning the VI workload domain Resources

At least one VI workload domain must be created to support guest virtual machines, 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 will deploy a vCenter and Platform Services Controller virtual machine. For VI workload domains supported by NSX-V, an NSX-V Manager is deployed in the management domain and three NSX-V controllers in the workload domain to support virtual networking. For VI workload domains supported by NSX-T, all workload domains of this type share NSX-T management resources. The first VI workload domain based on NSX-T will deploy three virtual machines in the management domain to be used for overall NSX-T management. These are the minimum requirements just to bring up a VI workload domain.

Depending on the use case, additional virtual machines might need to be deployed to support NSX network traffic services in a VI workload domain. For instance, logical routers are needed to enable networking between virtual machines on different network segments, and networking outside of the Cloud Foundation on VxRail environment. Depending on the NSX services required per use case, additional virtual machines might need to be deployed for this purpose.

Use the tables in the [Cloud Foundation on VxRail Footprints for Sizing](#) appendix for estimating the minimum sizing requirements for the management components, based on planned use cases.

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

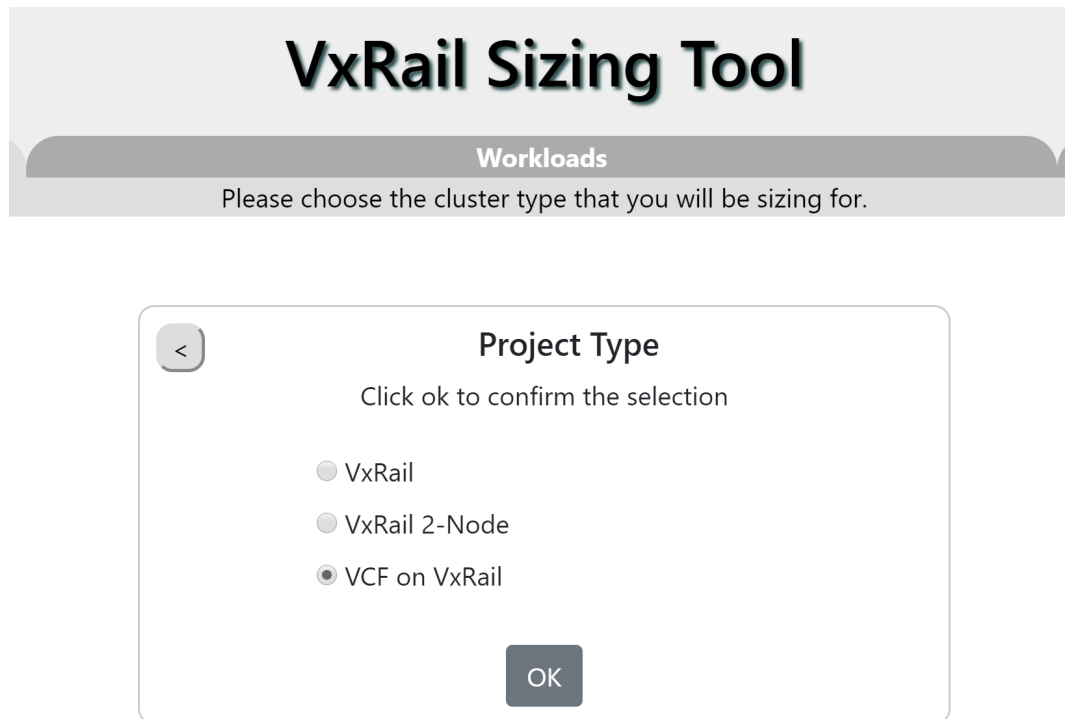


Figure 17 VxRail online sizing tool

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

**Workloads**

Please select all application packages that will be used in this management domain.

VCF Management Domain

NSX	
Edge Service Gateways	1

VCF Virtual Infrastructure	
ESRS	<input checked="" type="checkbox"/>

VCF Operations Management	
vRealize Suite	<input checked="" type="checkbox"/>
vRealize Operations	<input type="checkbox"/>

VCF Cloud Management	
vRealize Automation	<input type="checkbox"/>
vRealize Business	<input type="checkbox"/>
Microsoft SQL Server 2017	<input type="checkbox"/>

VCF Business Continuity	
Site Recovery	<input type="checkbox"/>
vSphere Replication	<input type="checkbox"/>

Figure 18 Selecting options for management domain 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.

## VxRail Sizing Tool

**Workloads**

Please input your requirement. Click + to add more workloads.

VCF Workload Domain

Type	Workload	# of VMs	% Concur	IOPS Per VM	Usable Storage Per VM (GB)	vCPUs Per VM	Memory Per VM (GB)	Read %	Read IO Size (KB)	Write IO Si
Reference	OLTP 4K		100					70	4	4

Start Sizing
Application Packages
Advanced Settings

Figure 19 Sizing 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 EMC's best practice is to use a collector tool for guest virtual machine resource requirements.

Dell EMC 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.

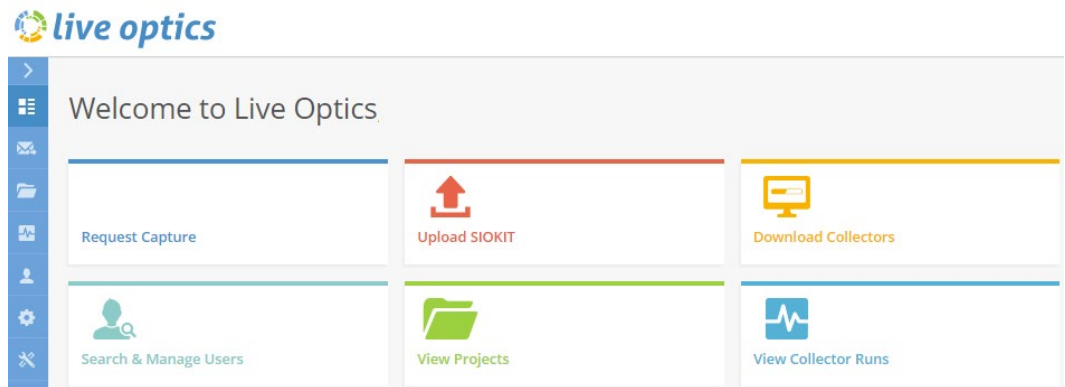


Figure 20 LiveOptics main dashboard

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 EMC 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.

## 6 Application Dependencies and Routing Decisions

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 necessary.

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 utilize VMware NSX, this means the virtual machine network traffic must travel upstream out of the virtual layer, where the routing decisions are made at the physical network layer.

Cloud Foundation on VxRail leverages NSX to enable support for routing in the virtual networks on the VI workload domains. This means 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.

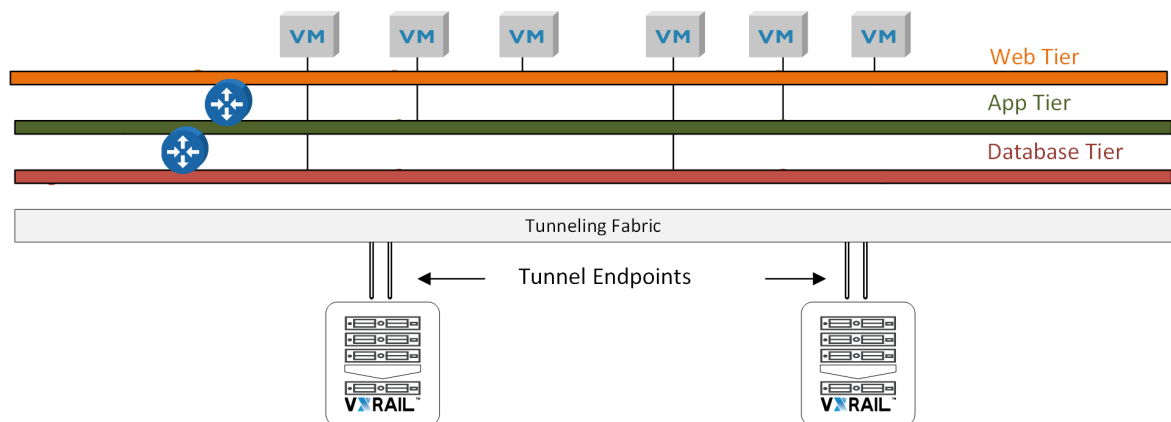


Figure 21 Three-tier application on separate virtual networks connected with virtual routers

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 identifier called a 'virtual LAN' (VLAN), 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 deployed. VXLAN for NSX-V or 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 across the individual virtual switches in the VxRail clusters.



A virtual router is deployed for the applications on one logical switch that need to access an application on another logical switch, such as connecting from the web tier application to the app tier.

The logical router also serves as the gateway to the external network through uplinks. In this instance, the logical router connects to an edge virtual device, which serves as the border between the physical network and virtual network. This virtual device can also serve as a pathway to go further upstream in the physical network, such as connecting to an application on a physical host, or for connecting to the Internet. In addition, the edge device synchronizes with upstream physical routers to form a seamless barrier between the physical and logical networks.

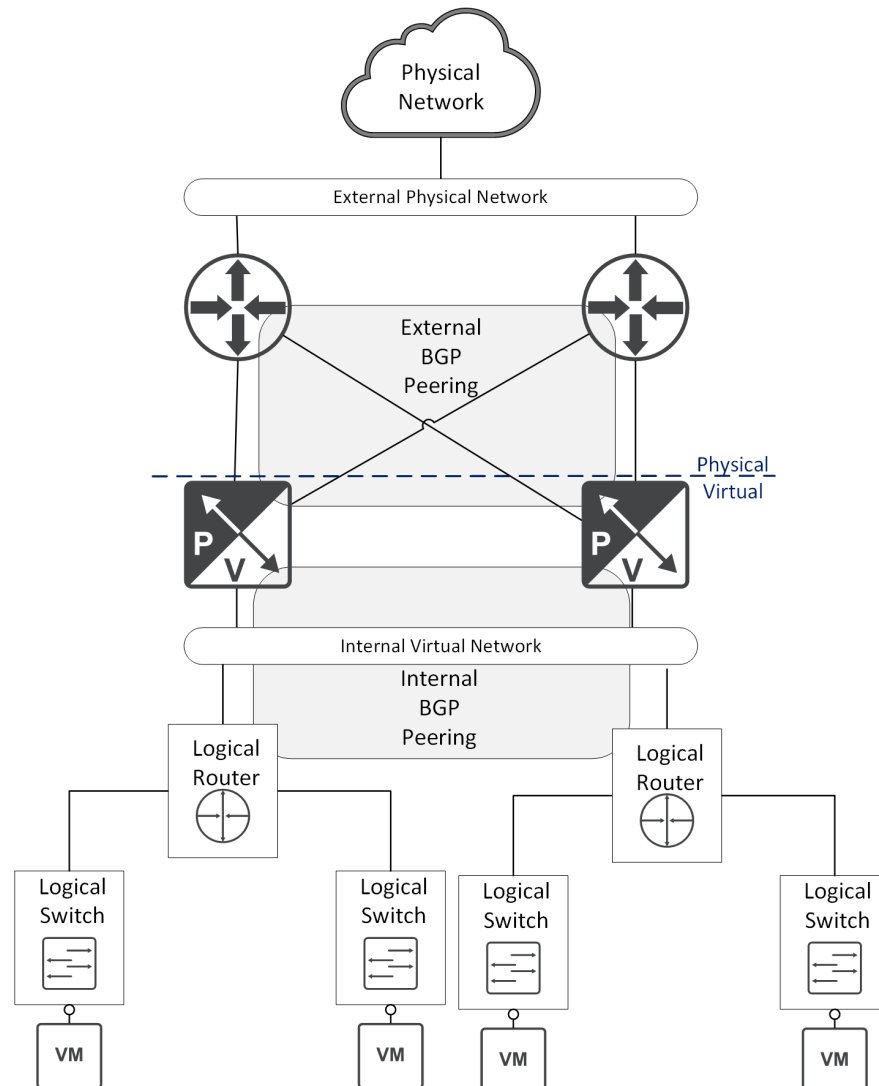


Figure 22 Overview of Physical and Logical Network Routing Relationships

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.

## 7 Cloud Foundation on VxRail Physical Network Planning

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, make sure 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.

### 7.1 Select a physical network architecture and topology

Cloud Foundation on VxRail offers flexibility with regards to the selection of a physical network architecture to support the planned deployment. The most common network topology for Cloud Foundation on VxRail, and is considered a best practice, is a spine-leaf topology. 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 edge 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 NSX edge gateway collects the routing tables from the adjacent routers in the physical network, and passes the updates down to the NSX controllers.

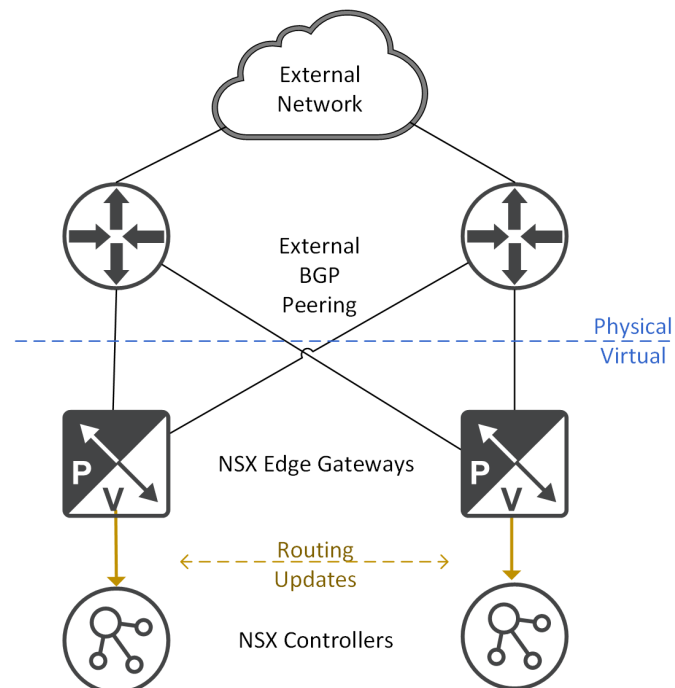


Figure 23 Peering NSX Edge Gateway with Physical Routers

The VLANs used in Cloud Foundation on VxRail to support the guest virtual machine networks terminate at these upstream routing devices. Therefore, using the route mapping for the applications planned for the VI workload domains drives the decisions for the locations for the NSX edge virtual devices in Cloud Foundation domains, and guides the process for selecting the location of 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 choose to deploy a spine-leaf network topology, enabling layer 3 at either the spine layer or the leaf layer is not required. However, this means routing 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 to deploy only a small number of edge virtual devices.




Figure 24 Options for layer 2/layer 3 boundaries in spine-leaf network topology

Establishing the router layer at the spine layer means the uplinks on the leaf layer are trunked ports, and pass through all of 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 of 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.




Figure 25 VxRail cluster nodes extended beyond one physical rack

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

This VLAN limitation can be overcome by establishing the routing at the leaf layer. This will optimize routing traffic, as it requires the least number of hops for the edge virtual devices to peer with an adjacent router. But it requires 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 without additional configuration. Both NSX-V and NSX-T support spanning layer 2 networks between switches using layer 3 services for the virtual machine traffic. To overcome the problem of extending layer 2 networks across racks and switches, best practice is for the leaf switch to support hardware-based 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 edge gateways
  - Multicast & Unicast: Required for VxRail and VXLAN traffic
  - Jumbo Frames: Required for VXLAN
  - Hardware-based tunneling (VTEP): Required to extend layer 2 VM traffic over a layer 3 network at the physical switch layer
2. Decide which physical network layer will support layer 3 routing services.

## 7.2 VxRail Cluster Physical Networking Planning

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, as well as the additional requirements for Cloud Foundation.

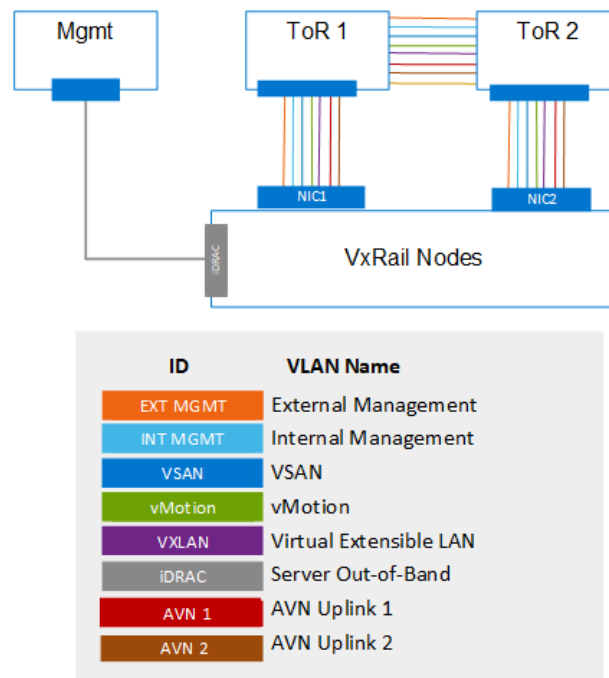


Figure 26 VxRail, VXLAN and AVN Networks

The key points to understand regarding the networking for the VxRail cluster are:

- Each VxRail node requires a minimum of two physical ports to be reserved for Cloud Foundation on VxRail network traffic, and only NSX-V will be used to support the workload domains.
- Each VxRail node requires four physical ports to be reserved if NSX-T is to be deployed for workload domains within Cloud Foundation on VxRail
- Each VxRail node balances the physical connections between the first top-of-rack (ToR) switch and the second top-of-rack (ToR) switch.
- Virtual LANs (VLANs) must be assigned to required VxRail, VXLAN and AVN (Application Virtual Network) networks to isolate the network traffic.
- The assigned VLANs must be configured on the Ethernet switches that directly connect to the VxRail nodes.

- The switch ports directly connected to the ports on the VxRail nodes must be trunked ports and allow passage for all of the required network VLANs.
- The inter-switch link between the switches connected directed to the VxRail nodes must also allow passage of all of these VLANs.
- The uplink ports configured for passage upstream must be configured to allow passage for all VLANs requiring external connectivity.
- The AVN (Application Virtual Network) networks must be configured for passage upstream to BGP routing services
- 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.

The VLAN for the VXLAN overlay network and the VLANs for the Application Virtual Network (AVN) are not VxRail cluster requirements. However, these VLANs are required for Cloud Foundation network operations. The VLANs must be configured on the top-of-rack switches, and also be configured to pass through the trunk ports directly connected to the VxRail nodes.

The following tasks must be performed in the physical top-of-rack switches in order to prepare for a VxRail cluster:

1. Select switches with sufficient open ports for Cloud Foundation on VxRail. If you plan to use NSX-T with any VI workload domains, reserve an additional two Ethernet ports on the switches.
2. Configure at least 1600 MTU to support VXLAN & GENEVE network traffic.
3. Make sure the port type on the switches (RJ45, SFP+) match the port type on the VxRail nodes.
4. Configure each of the VLANs required for VxRail networks on the switches.
5. Configure the switch ports to be directly connected to the VxRail nodes as layer 2 trunk ports.
6. Configure the inter-switch links to allow passage for all VLANs.
7. Configure the uplinks to allow passage for VLANs requiring external access.
8. Configure unicast on the VLAN representing the VSAN network.
9. Configure multicast on the VLAN representing the VxRail Internal Management network.
10. Configure MLD snooping and MLD querier on the VLAN representing the VxRail Internal Management Network (recommended).
11. Configure Spanning Tree on the switch ports to be directly connected to the VxRail nodes as edge ports, or in 'portfast' mode.

For complete details about VxRail cluster network requirements, see the [Dell EMC VxRail Network Planning Guide](#).

## 7.3 AVN (Application Virtual Network) Physical Network Planning

The Application Virtual Network enables linkage for the vRealize suite of management applications and enables connectivity to the upstream external network. The vRealize components, including vRealize Log Insight, vRealize Life Cycle Manager, vRealize Operations Manager and vRealize Automation, connect to the AVN when deployed. The Log Insight components are deployed at the time the Cloud Foundation management domain is created, while the remaining vRealize components can be deployed after the management domain creation process is complete.

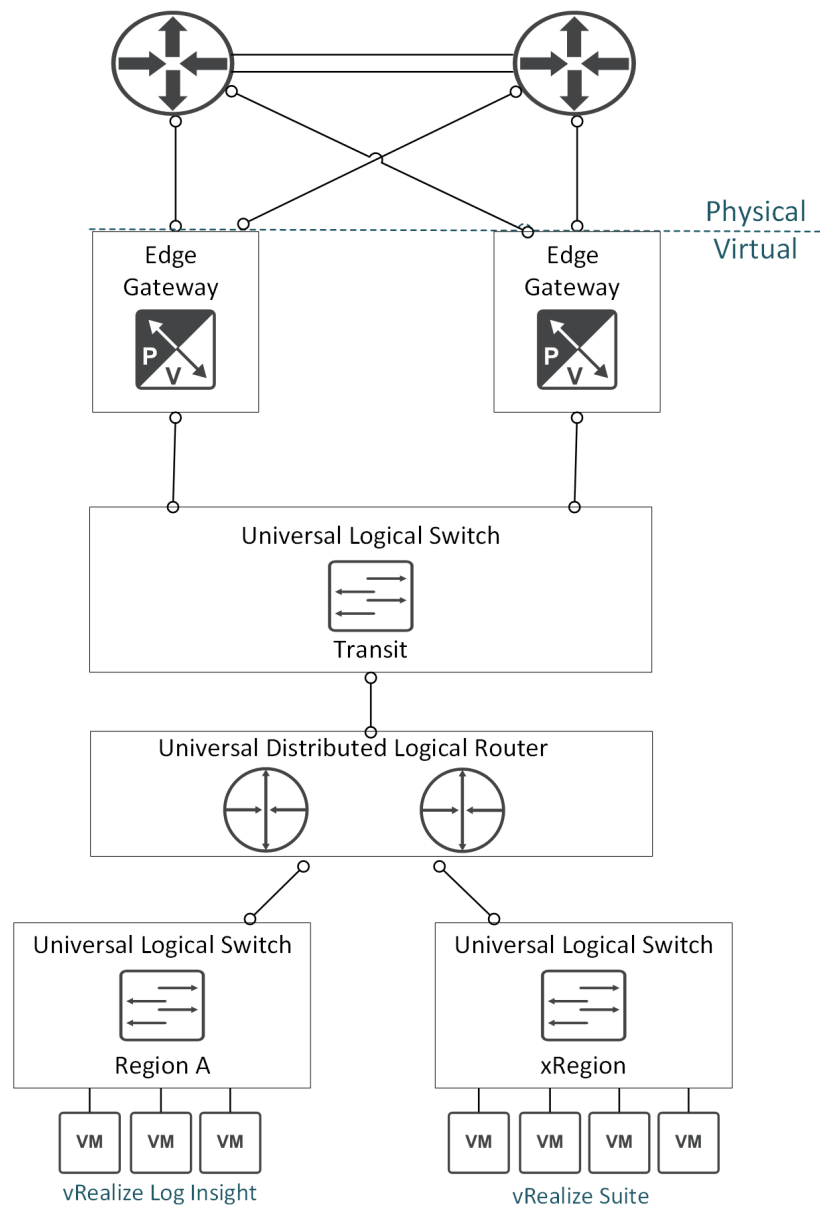


Figure 27 Application Virtual Network (AVN) Overview

During the creation of the Cloud Foundation management domain, the NSX-based logical switches and router needed for AVN are configured, and a pair of NSX edge gateways are also deployed to enable upstream access. The Cloud Foundation build process will perform the following tasks on the edge gateways deployed for the Application Virtual Network:

- Two portgroups on the virtual distributed switch in the VCF management domain are configured for BGP peering
- A VLAN is configured on the first portgroup for establishing an uplink with the first external router
- A VLAN is configured on the second portgroup for establishing an uplink with the second external router
- An IP address is assigned to the first virtual port on each Edge Gateway for BGP peering with the first external router
- An IP address is assigned to the second virtual port on each Edge Gateway for BGP peering with the second external router

- An IP address to connect downstream to the Universal Logical Router is configured on each Edge Gateway
- An ASN (Autonomous System Number) is assigned to the two Edge Gateways and Universal Logical Router
- iBGP (Internal Border Gateway Protocol) is enabled between the Edge Gateways and the Universal Logical Router
- A password to establish peering to each router instance is saved to the Universal Logical Router configuration file
- The ASN for the external router instances is saved to the Universal Logical Router configuration file
- The destination gateway IP address for peering to each external router instance is saved to each Edge Gateway configuration file

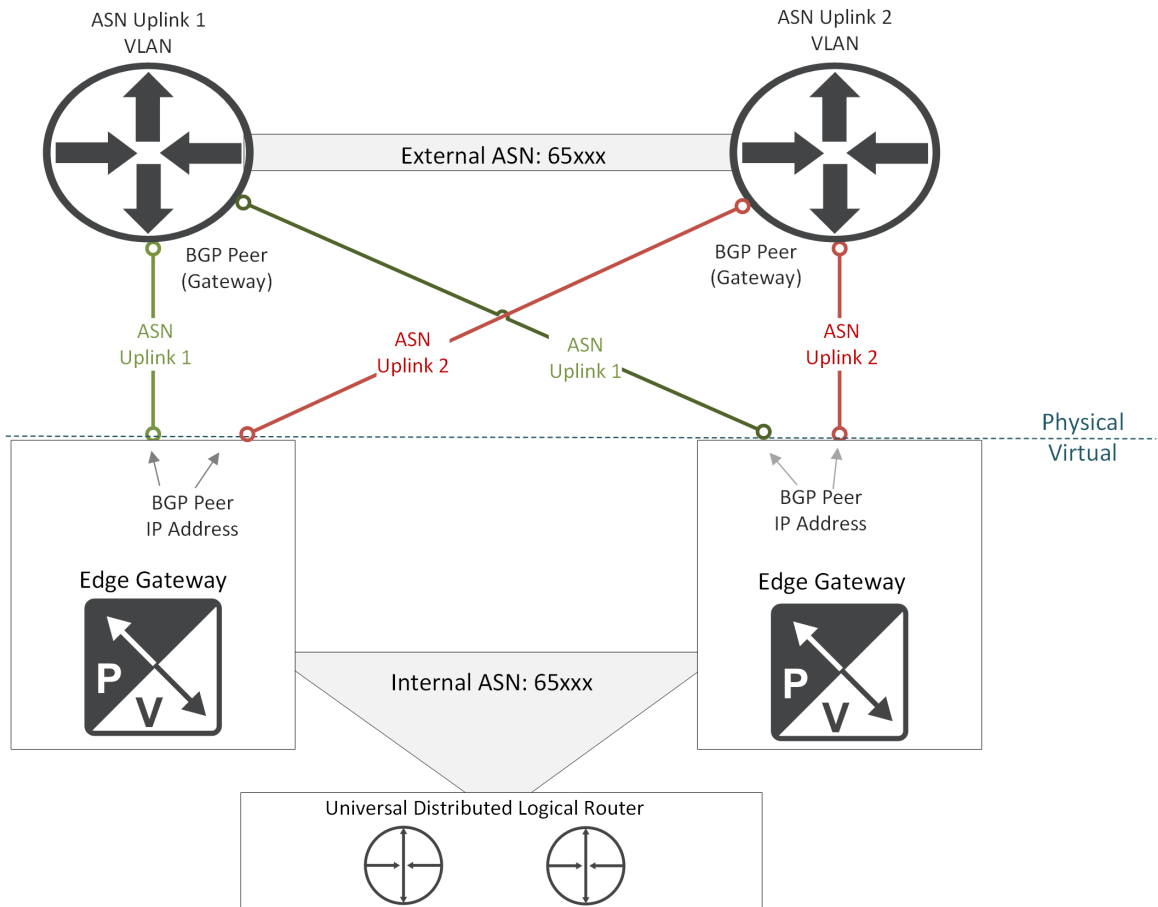


Figure 28 BGP relationship between Edge Gateways and external routers

The edge gateways must be able to configure an eBGP peer relationship with the upstream network as part of the Cloud Foundation deployment process. The following tasks must be completed on the upstream switches peering with the ASN Edge Services Gateways:

1. A VLAN matching the VLANs assigned for the uplinks on the Edge Gateways must be configured on each router instance
2. A gateway IP address matching the IP address saved to each Edge Gateway must be assigned to the VLAN on each router instance
3. BGP is configured on each router instance

- a. The external ASN number configured on each router matches the external ASN value saved to each Edge Gateway
- b. The password configured on each router matches the password saved to each Edge Gateway
- c. The internal ASN value configured on each router matches the internal ASN value configured on the Edge Gateways
- d. The IP addresses configured to establish neighbor relationships on the first router instance matches the IP addresses assigned to the first uplink on Edge Gateway instances
- e. The IP addresses configured to establish neighbor relationships on the second router instance matches the IP addresses assigned to the second uplink on Edge Gateway instances
- f. The timer 'keepalive' value is set to 4
- g. The timer 'holdtime' is set to 12

The example switch configuration syntax displayed in [Example switch configuration settings for BGP peering](#) offers guidance on how to configure an Ethernet switch for peering with a pair of Edge Gateways.

## 7.4 VxRail Stretched Cluster Physical Network Planning

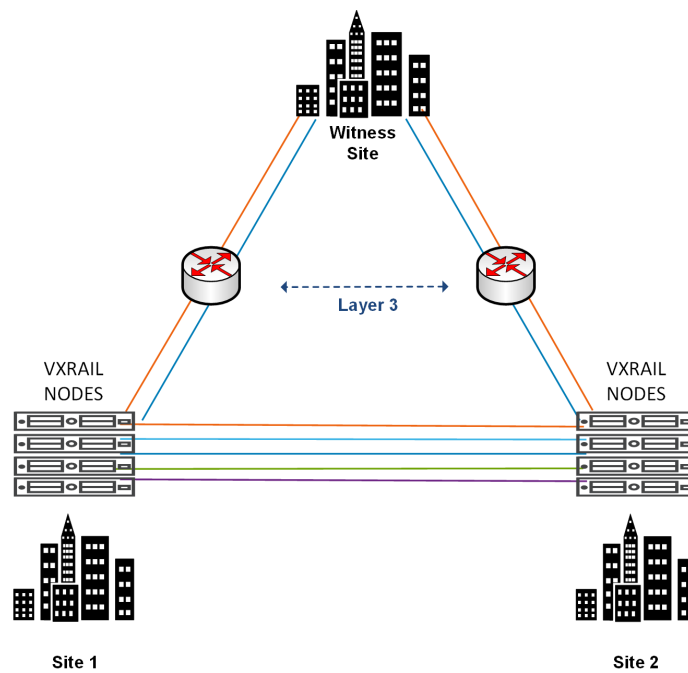
Cloud Foundation on VxRail supports two types of VxRail clusters: one where all nodes are in a single site and a stretched cluster where the nodes are equally distributed between two sites. VxRail stretched cluster is targeted specifically for situations with very high RPO and RTO requirements, and as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

---

For details about VxRail stretched cluster and its requirements, see the '*Dell EMC VxRail vSAN Stretched Cluster Planning Guide*'. It can be found under the 'Deploy' heading on the Dell EMC VxRail product site: [Dell EMC VxRail Stretched Cluster Planning Guide](#). This guide includes details about the strict latency and bandwidth requirements for the supporting physical network.

---





ID	VLAN Name	Network between sites
EXT MGMT	External Management	Layer 2
INT MGMT	Internal Management	Layer 3
VSAN	VSAN	Layer 3
vMotion	vMotion	Layer 2/Layer3
VXLAN	Virtual Extensible LAN	Layer 2 (NSX-V)/Layer 3 (NSX-T)

Figure 29 VxRail stretched cluster network requirements

The guidelines for a VxRail stretched cluster are:

- You must have three physical site locations in a VxRail stretched cluster
- The nodes that comprise a VxRail cluster instance to support the management domain or a workload domain are spread evenly over the two physical sites
- The third site supports the witness that monitors (via 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 virtual machines running in the VI workload domains

If you deploy Cloud Foundation on VxRail with VxRail stretched clusters as the underlying foundation, the cluster supporting the management workload domain must also be deployed as a VxRail 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 deploy the management workload domain as a stretched cluster. Converting an operational single site VxRail cluster instance to a stretched cluster is not supported. The only option is to reset the configured VxRail cluster back to its factory default state and rebuild the VxRail nodes into a supported stretched cluster.




Figure 30 One-to-one mapping of witness to VxRail stretched cluster

In addition to providing a third site to support the stretched cluster witnesses, note that a witness can monitor and support only one VxRail stretched cluster. This means that for every VxRail stretched cluster deployed for Cloud Foundation for VxRail, a witness virtual appliance must be deployed.

To support virtual machine network traffic between Cloud Foundation domains, the MTU size must be set to a minimum of 1600 at each site, and the VXLAN network must be able to pass through each of the sites in a VxRail stretched cluster.

## 7.5 Cloud Foundation Domain Physical Network Planning

The Cloud Foundation software layer is dependent on the VxRail networks and underlying networking infrastructure to be correctly configured and operational. There are additional considerations and steps that must be undertaken to support Cloud Foundation workloads.

---

Note: How the final supporting physical network is deployed will vary depending on individual requirements and networking equipment selection. Guidance provided here is limited to what is needed to be configured on the physical switch infrastructures, providing switch configuration syntax is out of scope for this document.

---

### 7.5.1 Identify the required logical layer 2 networks

Each network configured to bridge layer 2 networks over layer 3 will be given an ID known as a 'segment' ID. The segment range is configured after the Cloud Foundation domain is created. For instance, a segment ID range of 5000 to 7000 will support up to 2000 extended networks. We recommend that you record the NSX segment ID for each extended network for tracking purposes, and capture the properties of the bridged layer 2 network.

## 7.5.2 Configure VLANs for each Cloud Foundation domain

The Cloud Foundation management workload domain and each VI workload domain require VLANs to be configured on the physical switches to support virtual machine traffic. The [Cloud Foundation on VxRail VLANs](#) table in the appendix outlines the required and optional VLANs for the management workload domain and VI workload domain.

## 7.5.3 Configure overlay network settings

SDDC Manager configures a port group on the virtual switch to support VXLAN overlay traffic. Traffic on the overlay network passes through the VTEP on the host up through the physical network layer to its destination. The physical network layer must be configured to support the overlay traffic:

- The VLAN for the overlay network selected for a Cloud Foundation domain must be configured on the physical switches connected to the VxRail nodes supporting the Cloud Foundation domain.
- The VLAN for the overlay network must be configured to pass through the trunk ports connected to the VxRail nodes.
- The DHCP server that will supply IP addresses to the VXLAN tunnel endpoints must be accessible on each planned VXLAN network.
- IGMP snooping must be configured on the switch.
- IGMP snooping querier must be configured on the VLAN.

## 7.5.4 Configure routing services

The NSX edge gateways to support VI workload domains, and the NSX edge gateways to support the vRealize management suite on the AVN depend on synchronization of the routing tables with the upstream routers for network integration.

- Configure routing services at the layer 2/layer 3 boundary in the upstream network with eBGP.
- The following properties are required to peer with an edge gateway:
  - Router IP address
  - Autonomous system ID
  - Password

## 7.5.5 Identify multicast IP addresses for VXLAN

When two virtual machines connected to different hosts need to communicate, VXLAN-encapsulated traffic is exchanged between two VTEPs on the hosts. In VXLAN, all the learning about the virtual machine MAC address and its association with the VXLAN tunnel endpoints (VTEP) is performed through the support of physical network. VXLAN depends on the IP multicast protocol to populate the router forwarding tables in the virtual network.

A unique IP address range is assigned as the multicast group IP address to the VTEP in each VXLAN network. Identify available multicast IP addresses to assign for this role.

4. The multicast address is a Class D address ranging from 224.0.0.0 to 239.255.255.255.
5. Do not use multicast addresses ranging from 224.0.0.1 to 224.0.0.255. These are reserved for routing protocols and other low-level topology discovery or maintenance protocols.
6. From the remaining groups ranging from 224.0.1.0 to 239.255.255.255, use administratively scoped addresses from 239.0.0.0 to 239.255.255.255 if the VXLAN deployment is within a company. This multicast address range is viewed as a private multicast domain similar to the 10.0.0.0/8 range, so it is not used for global Internet traffic.

7. In deployments where VXLAN networks are shared between administrative domains, the addresses should be assigned from the 232.0.0.0/8 (232.0.0.0-232.255.255.255) source-specific block.

### 7.5.6 Deploy DHCP server for VXLAN Tunnel Endpoints

In Cloud Foundation on VxRail, two VXLAN tunnel endpoints are configured on each VxRail node. The endpoints are configured as virtual network adapters and are connected to the port group on the VxRail cluster's virtual distributed switch for the VXLAN network.

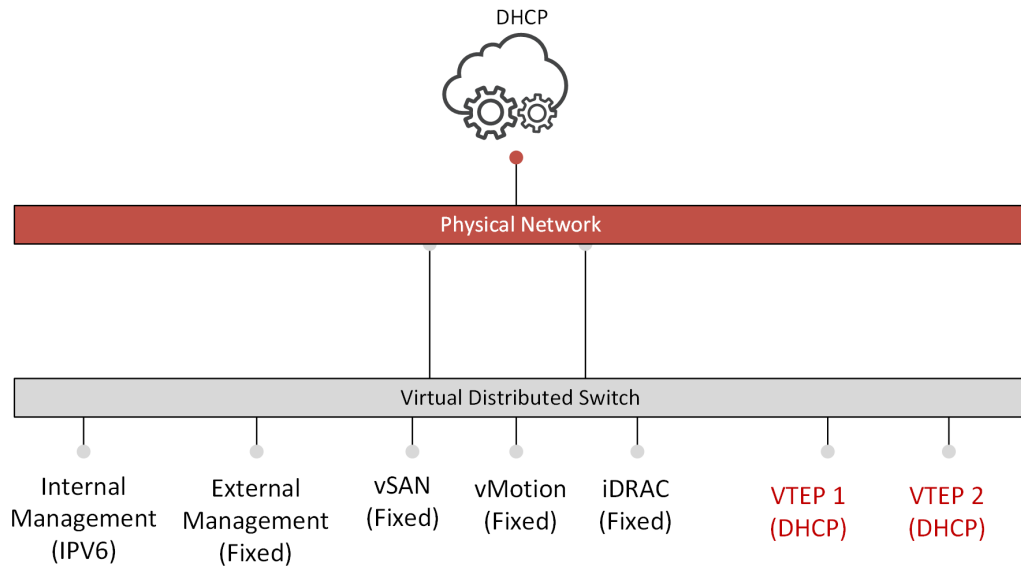


Figure 31 DHCP services for VXLAN Tunnel Endpoints

The IP addresses for the virtual adapters configured for the VxRail clusters are fixed. However, the IP addresses assigned to the two vmkernel interfaces for the tunnel endpoints are obtained from a DHCP server located in the data center. For each overlay network configured, a DHCP server must be accessible on that network to provide IP addresses to the tunnel endpoints.

## 8 VxRail Cluster Deployment Preparation

---

Dell EMC 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 EMC will review a set of prerequisites that must be met for a successful outcome. Dell EMC 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.

### 8.1 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. The VxRail nodes are powered on.
4. VxRail performs self-discovery of the powered-on nodes, and starts VxRail Manager for input of settings to perform automated initial build.

### 8.2 Capture and record the VLANs for VxRail cluster

Refer to the [VxRail Network Configuration](#) appendix for guidance on all of the settings that need to be collected to perform VxRail cluster automated initial build. This table lists the VLANs required for VxRail and for supporting the Application Virtual Network.

- By default, the external VxRail management network is the native VLAN. Since the typical Cloud Foundation on VxRail deployment will have more than one VxRail cluster per top-of-rack switch pair, it is advisable to select a value other than the native VLAN. Dell EMC will change the default external management VLAN on each VxRail node to the recorded value before power-on.
- By default, the internal VxRail management network is 3939. Again, since the typical Cloud Foundation on VxRail deployment will have more than one VxRail cluster per top-of-rack switch pair, it is advisable to set a new VLAN for the second and subsequent VxRail clusters. Dell EMC will change the internal management VLAN on each VxRail node to the recorded value before power-on.
- The vSAN and vMotion VLANs should be unique for each VxRail cluster.
- At least one guest network VLAN is required at initial build. The guest network can be removed after the initial build.

## 8.3 Capture and record the network settings for VxRail cluster

The [VxRail Network Configuration](#) appendix includes the hostnames and network settings required for VxRail management. Note the following:

- Hostnames and IP addresses for VxRail management components are automatically assigned during initial build.
- The IP addresses must be unused and permanent IP addresses, not DHCP-based.
- The IP addresses assigned to VxRail management components must be on the same subnet.
- The IP addresses assigned to VxRail nodes during the initial build are done sequentially, so a range of IP addresses is required.
- All ESXi hostnames in a VxRail cluster are defined by a naming scheme that is comprised of: an ESXi hostname prefix (an alphanumeric string), a separator (“None” or a dash “-“), an iterator (Alpha, Num X, or Num 0X), an offset (empty or numeric), a suffix (empty or alphanumeric string with no .), and a domain.
- Capture the IP addresses of the NTP and DNS servers. Both must be accessible to the VxRail external management network.

## 8.4 Capture and record 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 the [VxRail Network Configuration](#) appendix 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
- Optional VLAN for Witness Traffic Separation

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
- If NSX-T is deployed for the workload domain, then a layer 3 network is required between sites
- If NSX-V is deployed for the workload domain, then only a layer 2 network is supported between sites

## 8.5 Create Forward and Reverse DNS Entries for VxRail cluster

Using the information captured in the [VxRail Network Configuration](#) appendix, create forward and reverse DNS entries for every hostname planned for the VxRail cluster. These include VxRail Manager, vCenter Server, VxRail Platform Service Controller, and each ESXi host in the VxRail cluster.

## 8.6 Prepare top-of-rack switches for VxRail cluster

Using the data collected in the VxRail Network Configuration table, follow these tasks to prepare the top-of-rack switches for VxRail initial build:

- Configure VLANs required for VxRail cluster.
- Configure switch ports connected to VxRail nodes as trunk ports passing through all VLANs.
- Configure switch ports connected to VxRail nodes as STP edge ports.
- Enable Multicast for VxRail Internal Management Network.
- Enable Unicast for VxRail vSAN Network.
- Enable Uplinks to pass external VxRail network traffic.
- Enable Inter-Switch Links.
- Enable MTU size of at least 1600 to support VXLAN/GENEVE
- Configure eBGP for peering with NSX edge gateways

## 8.7 Prepare passwords

For VxRail cluster components, a password is required. The password policy follows VMware standards: Minimum of 8 characters in length and at least one uppercase character, lowercase character, digit and special character (example: @!#\$%^).

## 9 Prepare for VMware Cloud Foundation Management VI workload domain

To configure the management workload domain, Dell EMC will follow these milestones:

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 to the virtual appliance.
4. Activate the Cloud Builder process. This will lay down the SDDC management layer on top of the VxRail cluster using the uploaded settings.

Cloud Builder will perform the following tasks to create the management VI workload domain using the data input into the tool:

1. Deploy NSX-V Manager.
2. Register NSX-V Manager with vCenter & Platform Services Controller.
3. Install License for NSX-V.
4. Deploy 3 NSX-V Controllers.
5. Create Anti-Affinity Rules for NSX-V Controllers.
6. Configure VXLAN Segment ID Range.
7. Configure a new Global Transport Zone.
8. Add management VI workload domain to the new Transport Zone.
9. Deploy and configure the Application Virtual Network
10. Deploy Log Insight in Region A of the Application Virtual Network
11. Deploy SDDC Manager in the management VI workload domain

Refer to the [Cloud Builder and Management VI workload domain Configuration](#) appendix for guidance on all of the settings that need to be collected for Cloud Builder.

### 9.1 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.

### 9.2 Global network setting for management VI workload domain

Capture and record the following global settings:

- The IP addresses for the DNS server(s) to support the management VI workload domain
- The IP addresses for the NTP server(s) to support the management VI workload domain
- The IP address for the DHCP server to support IP address assignment for the VTEP tunnel endpoints for the management VI workload domain
- The site name for single sign-on (SSO). The site name needs to be the same as the site name for the underlying VxRail cluster.
- The domain name and optional sub-domain name



## 9.3 Capture and record the network settings for the management VI workload domain

Capture and record the following management VI workload domain network settings:

- Hostname and IP address settings for SDDC Manager
- Hostname and IP address settings for NSX-V Manager
- Hostname and IP address settings for vSphere Platform Services Controller
- A pool of three contiguous IP addresses for NSX-V Controllers

The IP addresses must be unused and permanent.

## 9.4 Capture and record the DHCP settings for VTEP tunnel endpoints for the management VI workload domain

Capture and record the following DHCP settings to support the VTEP tunnel endpoints:

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

## 9.5 Create Forward and Reverse DNS Entries for the management VI workload domain

Create forward and reverse DNS entries for every hostname planned for the Cloud Foundation management VI workload domain. These include SDDC Manager, NSX-V Manager, NSX-V Controllers, and Log Insight.

## 9.6 Select and create the VXLAN VLAN

A VLAN must be selected for the VXLAN network

- The VXLAN VLAN must be configured on the adjacent top-of-rack switches connected to the VxRail nodes.
- The VXLAN VLAN must pass through 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 VLAN upstream through the uplinks on the adjacent top-of-rack switches.
- Configure IGMP Snooping and an IGMP Querier on the adjacent top-of-rack switches.

## 9.7 Capture and record the NSX settings for the management VI workload domain

Capture and record the following settings for NSX:

- Select the NSX segment range. The values between the starting and ending NSX segment IDs define the number of VXLAN networks that can be supported on the Cloud Foundation on VxRail platform.

- Select the name of the transport zone. VxRail nodes that are members of a transport zone are connected over a logical network, and the virtual machines on those VxRail nodes can communicate over this logical network.

## 9.8 Select a multicast IP address range for VXLAN network

In VXLAN, all the learning about the virtual machine MAC address and its association with VTEP IP address is performed through the support of physical network. Cloud Foundation on VxRail uses IP multicast by default to populate the forwarding tables in the VTEP.

1. Assign a unique IP address range is assigned as the multicast IP address group in the range of 224.0.0.0 to 239.255.255.255. Each address in this range designates a multicast group.
2. Do not use multicast addresses ranging from 224.0.0.1 to 224.0.0.255. These are reserved for routing protocols and other low-level topology discovery or maintenance protocols.
3. Use administratively scoped addresses from 239.0.0.0 to 239.255.255.255 if the VXLAN deployment is within a company. This multicast address range is viewed as a private multicast domain.
4. If VXLAN networks will be shared between administrative domains, the addresses should be assigned from the 232.0.0.0/8 (232.0.0.0-232.255.255.255) block.

## 9.9 Select names for resource pools in VI Management Domain

Four resource pools will be created during the build process for the VI management domain. A resource pool for management components and edge components will be created for SDDC management and SDDC users. The default names provided for these four resource pools can be customized.

## 9.10 Prepare passwords

A password is required for Cloud Foundation management workload domain components. Like VxRail, the password policy follows VMware standards: Minimum of 8 characters in length and at least one uppercase character, lowercase character, digit and special character (example: @!#\$%^).

## 9.11 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. Licenses are required for:

- vSphere
- vSAN
- NSX
- vRealize Log Insight

## 10 Prepare for Cloud Foundation Application Virtual Network

The Application Virtual Network (AVN) is deployed as part of the initial deployment of the Cloud Foundation performed by CloudBuilder. The requirements for the Application Virtual Network phase are input into CloudBuilder and deployed based upon the settings captured in the workbook. The tables in section [Application Virtual Network Configuration](#) represent what needs to be captured to deploy the Application Virtual Network by CloudBuilder.

### 10.1 Capture the Application Virtual Network Region Settings

The network traffic from the Log Insight components in Region A and the vRealize management components in xRegion of the Application Virtual Network must be assigned routable IP addresses that can pass upstream to required data center services (DNS, NTP) and connect to external support sites to enable software lifecycle management.

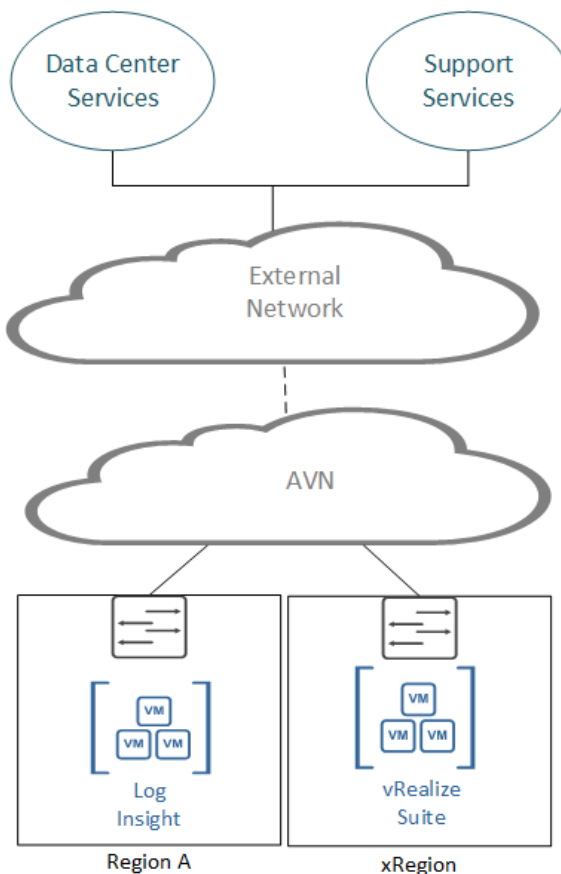


Figure 32 External connectivity requirements for AVN

The table labeled 'Application Virtual Network Regions' in the section [Application Virtual Network Configuration](#) outlines the settings that need to be captured and input into CloudBuilder at the time of initial deployment of Cloud Foundation. For the DNS search domain and zone selected for each region, the networks assigned to each region must be able to route upstream through the AVN and physical networks to reach the required data center services.

## 10.2 Capture settings for BGP peering

The table 'BGP Peering Networks' in the section [Application Virtual Network Configuration](#) captures the VLANs for the eBGP peering for the first uplink and second uplink of each Edge Service Gateway. At the time of deployment, CloudBuilder will configure two virtual switch portgroups, and assign the VLANs to each portgroup respectively. CloudBuilder will then connect the first Edge Service Gateway uplink to the first portgroup, and the second Edge Service Gateway to the second portgroup.

The table 'BGP Peering Networks' also captures the networks to be assigned to enable connectivity for eBGP and iBGP peering

- Capture the IP address ranges for eBGP peering between the external routers and Edge Service Gateways in CIDR format. The range only needs to be large enough to support peering with the external routers
- Capture the IP address range for iBGP peering between the Edge Service Gateways and Universal Distributed Logical Router in CIDR format. The range only needs to be large enough to support peering within the Application Virtual Network.

## 10.3 Capture settings for Universal Distributed Logical Router

The table 'Universal Distributed Logical Router' in the section [Application Virtual Network Configuration](#) outlines the settings needed to configure the Universal Distributed Logical Router (UDLR)

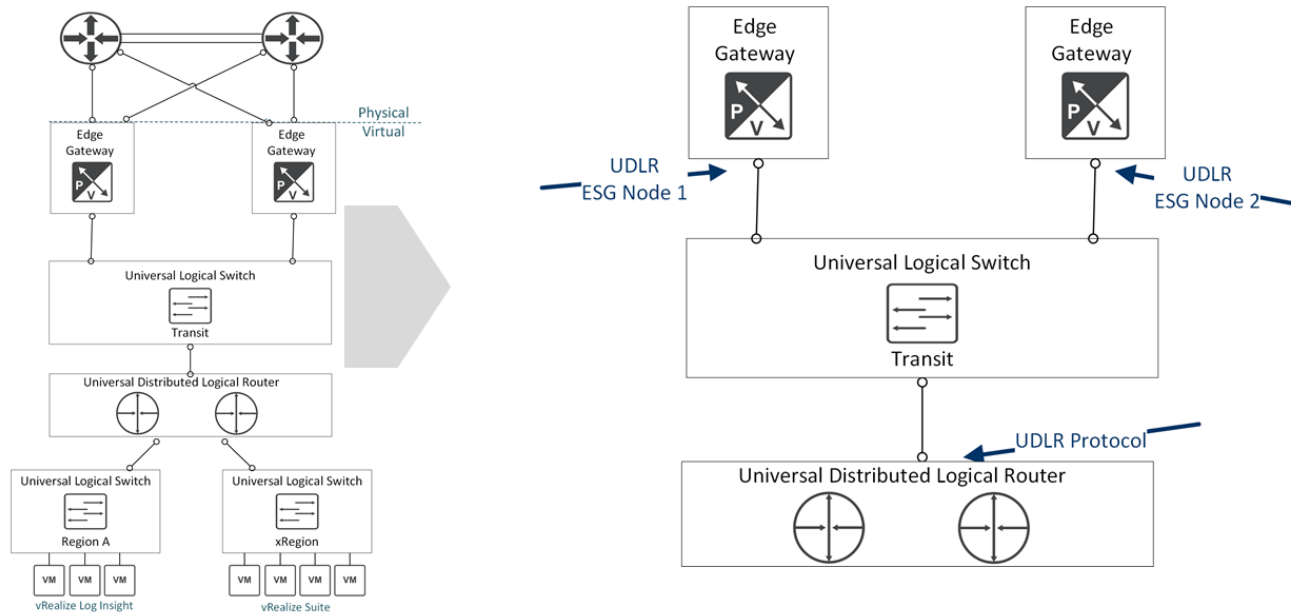


Figure 33 Network settings for UDLR and Edge Gateways

- The ASN selected is what is configured on the virtual router instance by CloudBuilder during initial Cloud Foundation initial deployment. This same ASN value must be configured as the remote ASN in the eBGP configuration settings on each physical router
- An IP address range in CIDR format must be reserved for the UDLR instance. The UDLR uses the IP addresses from this preset range to enable iBGP peering with the Edge Services Gateways. The IP address range must be large enough to assign the required IP addresses to the UDLR, and to assign an IP address to each Edge Service Gateway.

- The ASN selected is what is configured on the UDLR by CloudBuilder during initial Cloud Foundation initial deployment. This same ASN value must be configured as the remote ASN in the eBGP configuration settings on each physical router.
- The password selected is what is configured on the virtual router instance by CloudBuilder during initial Cloud Foundation initial deployment. This same password value must be configured in the eBGP configuration settings on each physical router
- The IP addresses to be assigned to the UDLR instance must be from the IP address range assigned to the UDLR.

## 10.4 Capture settings for Edge Service Gateways

The table 'Edge Service Gateways' in the section [Application Virtual Network Configuration](#) outlines the network settings for each ESG to enable eBGP peering with the upstream routers.

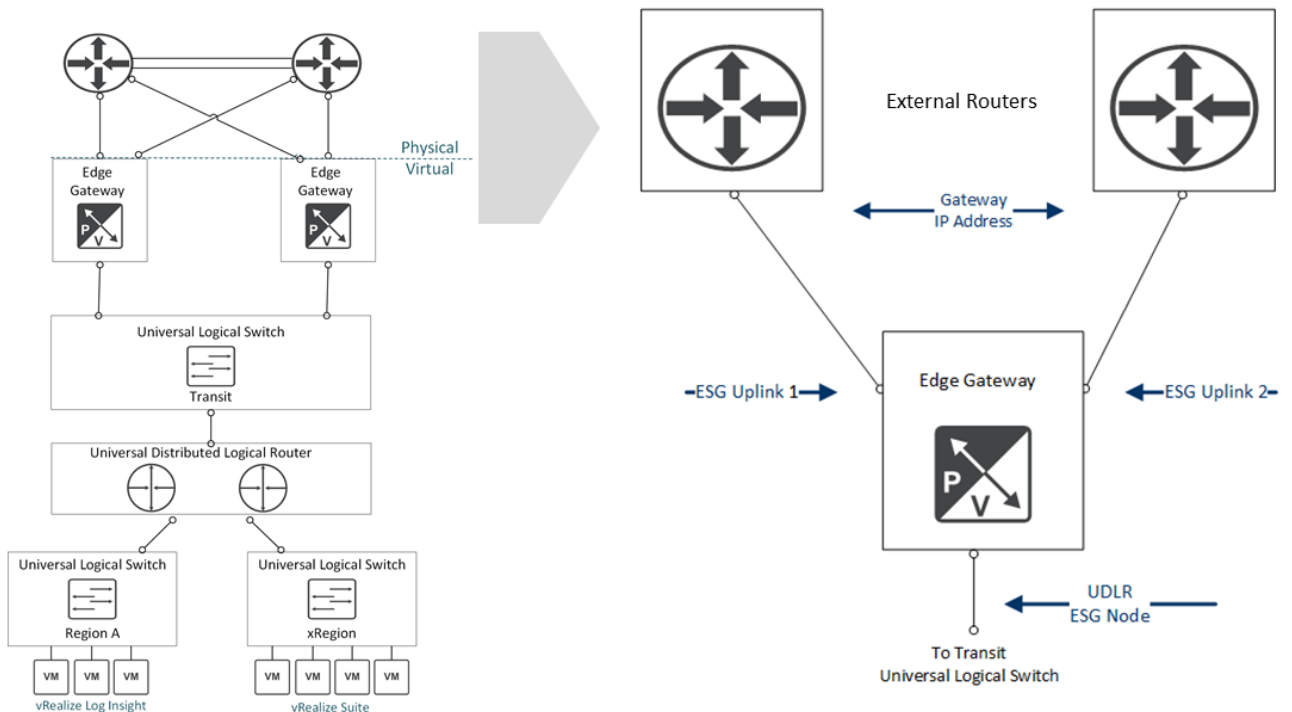


Figure 34 Network settings for AVN Edge Service Gateways

- The individual IP addresses assigned to each uplink must be within the range entered in table 'eBGP Peering Networks'
- The IP address for uplink 1 on each ESG must be from the same subnet
- The IP address assigned to uplink 2 on each ESG must also be in the same subnet
- Each ESG must be able to connect to the Universal Distributed Logical Router to enable iBGP peering. The IP address entered for the UDLR IP address for each must be in the same range as assigned to the Universal Distributed Logical Router.

## 10.5 Capture external router settings for the eBGP peering

The table 'External Routers' in the section [Application Virtual Network Configuration](#) captures what is configured on the upstream routers that will peer with the AVN Edge Service Gateways.

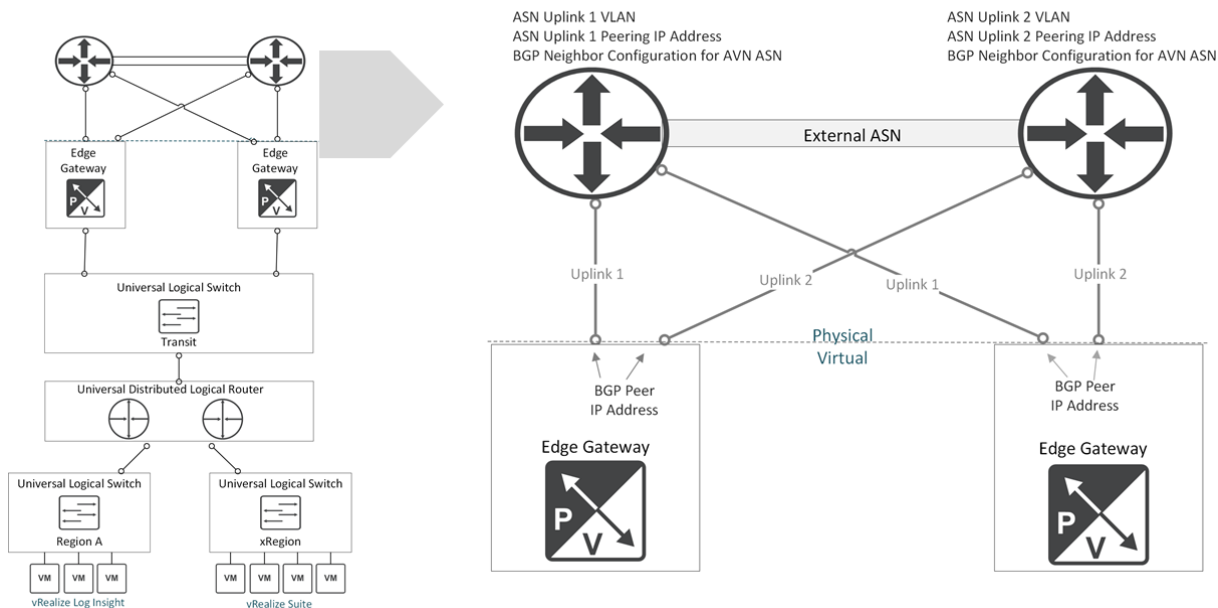


Figure 35 Network settings for external routing services for eBGP peering

- The ASN represents is what is configured for eBGP peering on the external routers
- The password represents what is configured on the external routers for eBGP peering
- For the first physical router:
  - The receiving IP address must be configured on the VLAN entered in the row 'ESG Uplink 1 VLAN' in table 'eBGP Peering Networks'
  - The receiving IP address must be within the range entered in the row 'ESG Uplink 1 Network' in table 'eBGP Peering Networks'
- For the second physical router:
  - The receiving IP address must be configured on the VLAN entered in the row 'ESG Uplink 2 VLAN' in table 'eBGP Peering Networks'
  - The receiving IP address must be within the range entered in the row 'ESG Uplink 2 Network' in table 'eBGP Peering Networks'

## 10.6 Capture Universal Logical Router settings

The table 'Universal Logical Router Settings' in the section [Application Virtual Network Configuration](#) captures the settings and properties for the Universal Distributed Logical Router.

- The segment range are the values between the starting and ending NSX segments. The value defines the number of VXLAN networks that can be supported by the UDLR
- The universal multicast address range needs to map to the networking requirements for the UDLR. The default values should work in most cases unless there is a conflict in the data center.
- The default size of MT

## 10.7 Capture Region A Management Component Settings

The table 'AVN Region A Management Components' in the section [Application Virtual Network Configuration](#) outlines the settings for Log Insight which is deployed by CloudBuilder. The IP addresses assigned to Log Insight must be from within the range entered in the 'Region A VXLAN' section of the table 'Application Virtual Network Regions'.

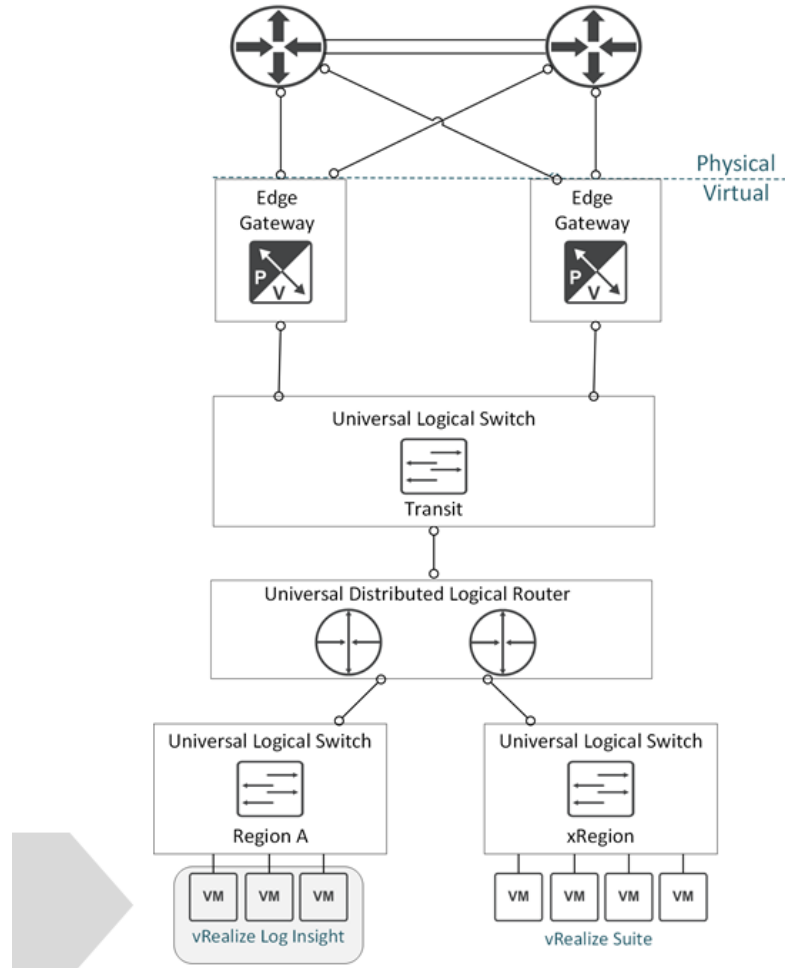


Figure 36 Network settings for vRealize Log Insight

## 10.8 Capture 2nd Site settings for Stretched Cluster

If Cloud Foundation is to be deployed on a VxRail stretched cluster, then network settings must be captured for the second site. The table '2nd Site BGP Peering Networks' and '2nd Site External Routers' in the section [Application Virtual Network Configuration](#) outlines the second set of settings that need to be captured for the second site.

The underlying networks supporting the VxRail cluster and VXLAN are configured across the two sites in the stretched cluster to enable cross-site connectivity. The logical network switches and routers supporting the Application Virtual Network also extend between the two sites using this cross-site network connectivity. To ensure uninterrupted routing services, both sites in the stretched cluster peer with upstream routers, which means an additional pair of Edge Services Gateways need to be deployed at the second site, and the second set of Edge Services Gateways need to form an eBGP peering relationship with routers in the second site.

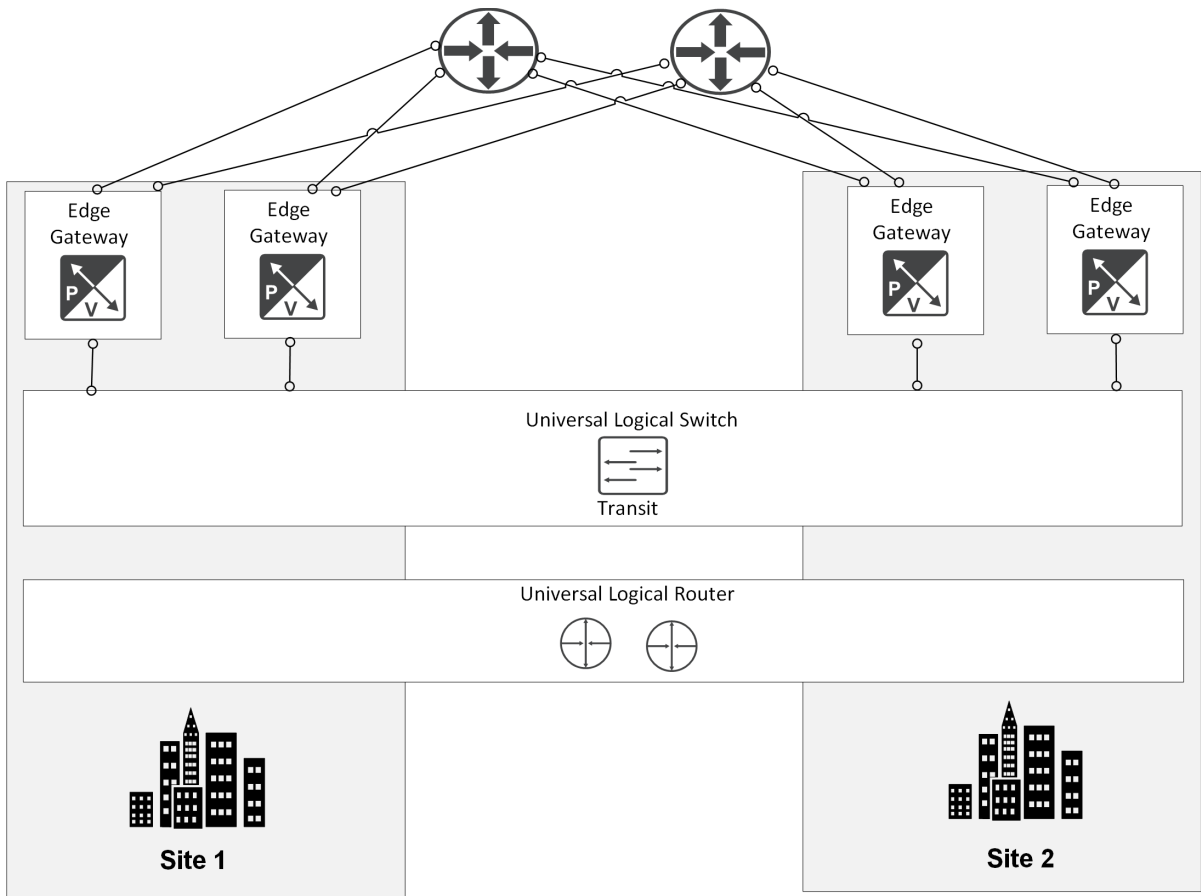


Figure 37 AVN deployed on a VxRail stretched cluster with two pairs of Edge Gateways



# 11 Prepare for Cloud Foundation VI workload domain

Once a Cloud Foundation on VxRail management workload domain is deployed during the Cloud Builder deployment phase, 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.

---

Note: Configuring a Cloud Foundation on VxRail management workload domain assigns the VxRail cluster as the underlying resource pool to support virtual machines, and then deploys the basic management components within the domain. Additional steps are required for the buildout of the NSX-based virtual network in the domain.

---

Complete the following tasks to deploy a Cloud Foundation VI workload domain:

1. Deploy the basic VI workload domain structure from SDDC Manager.
  - This will deploy a vCenter instance in the management workload domain. This vCenter instance is used to manage all VxRail clusters assigned to support the VI workload 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(s) to the VI workload domain using SDDC Manager.
  - For VI workload domains serviced by NSX-V, this action will deploy the NSX-V Manager in the management workload domain and three NSX-V controllers in the VI workload domain.
  - For the first VI workload domains serviced by NSX-T, this action will deploy three NSX-T Managers in the management workload domain. The NSX-T Managers will be used by all subsequent VI workload domains to be serviced by NSX-T.

## 11.1 Configure VI workload domain VXLAN

Each VI workload domain requires at least one VXLAN or overlay VLAN. Each cluster can have a dedicated overlay VLAN. However, each cluster is added to a single Transport Zone.

- If each VxRail cluster in a VI workload domain shares the same VLAN subnet for the VTEPs, use the same VLAN for each cluster added to the VI workload domain.
- Use a unique VLAN if you want to start a dedicated VXLAN overlay network. Networking in this VI workload domain will be limited until more clusters are linked to this VXLAN-based logical switch.
- For a new VXLAN VLAN, configure the VLAN on the adjacent switches providing connectivity to the VxRail nodes.
- Add the VXLAN VLAN to each trunked port on the adjacent switches providing connectivity to the VxRail nodes.
- Ensure connectivity to the DHCP server supplying the IP addresses for the VTEP tunnel endpoints to this VXLAN network.
- Configure IGMP Snooping and an IGMP Querier on the adjacent top-of-rack switches.

## 11.2 Capture settings for VI workload domain

Use the VI workload domain Configuration Settings appendix as a guide for the VI workload domain deployment requirements.

- Each VI workload domain requires a unique name and an organization name.
- Capture the VXLAN overlay VLAN ID.
- Network settings are required for the vCenter instance that is deployed in the management workload domain. Save the network settings for the vCenter instance for the subsequent phase when the VxRail cluster is assigned to the VI workload domain.
- Follow the steps outlined in the VxRail Cluster Deployment Preparation section to capture and record the settings for a VxRail cluster. Use the network settings for the vCenter instance from the previous step for the VxRail cluster.
- Capture the network settings for the NSX components. Select an IP address accessible to the external management network for the NSX management components.

---

Note: Configure a Cloud Foundation on VxRail VI workload domain from SDDC Manager to assign the VxRail cluster as the underlying resource pool to support virtual machines, and then deploy the basic virtual management components on the cluster resources to enable operations. Additional steps are required for the buildout of the NSX-based virtual network in the VI workload domain.

---

## 12 Planning for NSX integration with physical network

The process of deploying a Cloud Foundation on VxRail domain places the pool of resources from VxRail clusters as the foundation for future workload, and initializes the management components to begin the process of virtual machine creation, connectivity and operations. The network connectivity options for the virtual machines at this juncture are the base virtual networks configured for VxRail and the initialization of the Cloud Foundation domain.

Network efficiency and scalability can be realized by using routers in the appropriate locations, which will decrease broadcast domains. With NSX, routing decisions will transition down into the virtual network to further reduce hops for routing purposes and reduce network broadcasts.

The entire breadth of planning the virtual network topology to meet the application dependencies planned for the VI workload domains is out of scope for this guide. The NSX Routing Settings appendix provides guidance on a scenario of capturing the settings for the primary components needed to integrate the supporting physical network with the NSX logical network, and enable logical routing.

### 12.1 Capture settings for upstream router in physical network

The peering of an NSX edge virtual device with an upstream router will enable the passing of routing tables and integrate the routing paths between the physical and virtual network layers. For each upstream router peering candidate, capture the following:

- Router IP address
- Router password
- Autonomous system ID

Ensure that the external Border Gateway Protocol is configured on the upstream router.

### 12.2 Capture settings for NSX Edge virtual devices

The NSX Edge virtual device supports connectivity upstream through uplinks and connectivity to the virtual networks through internal ports. Edge virtual devices can be deployed as needed throughout an NSX network.

- Capture the network settings for the edge gateway uplinks.
- Capture the network settings for the edge gateway internal links.
- If you want to peer an uplink on the edge gateway with an external router, the uplink must connect to a port group that allows external network access.

The best practice is for the edge gateway to peer with two upstream routers for redundancy.

### 12.3 Capture settings for NSX logical router

The logical router performs routing between overlay networks, and between virtual and physical networks. A logical router can peer with multiple edge virtual devices as part of the overall network topology. The router requires the following:

- A forwarding IP address for forwarding data path updates upstream
- A protocol IP address on the same subnet that is used to peer with an upstream router

# A 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	4	16	290
Management	PSC	2	4	60
Management	PSC	2	4	60
Management	NSX-V Manager	4	16	60
Management	NSX-V Controller 1	4	4	28
Management	NSX-V Controller 2	4	4	28
Management	NSX-V Controller 3	4	4	28

Base Virtual Machines for Cloud Foundation Domain Logging\*

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	Log Insight - Master	8	16	1312
Management	Log Insight – Worker #1	8	16	1312
Management	Log Insight – Worker #2	8	16	1312

\* Log Insight can be disabled. Dell EMC support can be engaged to disable Log Insight.

Virtual Machines deployed in Cloud Foundation Management Domain for each Cloud Foundation VI workload domain based on NSX-V or NSX-T

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

Virtual Machines deployed for each Cloud Foundation VI workload domain based on NSX-V

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	NSX-V Manager	4	4	60
Workload	NSX-V Controller 1	4	4	20
Workload	NSX-V Controller 2	4	4	20
Workload	NSX-V Controller 3	4	4	20

(Optional) Virtual Machines deployed for NSX-V Edge Services\*

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	NSX-V Edge Services Gateway	2	1	1

\* Use this as an estimate for a deployment of a single virtual machine instance of an NSX Edge services gateway. The number of NSX-V Edge gateways deployed will be dependent on the final network design.

Virtual Machines deployed for initial Cloud Foundation VI workload domain using NSX-T and shared by subsequent Cloud Foundation VI workload domains using NSX-T

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	NSX-T Manager 1	4	4	20
Management	NSX-T Manager 2	4	4	20
Management	NSX-T Manager 3	4	4	20

(Optional) Virtual Machines deployed for shared edge virtual devices based on NSX-T

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Workload	NSX-T Edge Device 1	8	16	120
Workload	NSX-T Edge Device 2	8	16	120

The following tables are the optional VMware software applications that can be deployed on a VI workload domain.

#### vRealize Cloud Operations Management

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	vRealize Operations Manager - Master Node	8	32	1024
Management	vRealize Operations Manager - Replica Node	8	32	1024
Management	vRealize Operations Manager - Data Node	8	32	1024
Management	vRealize Operations Manager - Remote Collector 1	2	4	275
Management	vRealize Operations Manager - Remote Collector 2	2	4	275
Management	vRealize Lifecycle Manager	2	16	135
Management	vRealize Download Services	2	2	120

#### vRealize Cloud Management

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	vRealize Automation Node 1	4	16	140
Management	vRealize Automation Node 2	4	16	140
Management	vRealize Automation Node 3	4	16	140
Management	vRealize Automation Web Manager 1	2	8	60
Management	vRealize Automation Web Manager 2	2	8	60
Management	vRealize Automation Manager 1	2	8	60
Management	vRealize Automation Manager 2	2	8	60
Management	vRealize Automation DEM Worker 1	4	8	60
Management	vRealize Automation DEM Worker 2	4	8	60

Management	vRealize Automation Proxy Agent 1	2	8	60
Management	vRealize Automation Proxy Agent 2	2	8	60
Management	vRealize Business Cloud Server	4	8	50
Management	vRealize Business Cloud Collector	4	2	50
Management	vRealize IaaS Database	8	16	80

#### Horizon

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	Each Connection Server	2	10	80
Management	Each Composer Server	2	10	80
Management	Each Apps Volumes Server	2	10	80
Management	Each Unified Access Gateway	2	4	20
Management	Each User Environment Manager	2	10	80

#### PKS

Domain	Component	vCPUs	Memory (GB)	Storage (GB)
Management	Operations Manager	1	8	160
Management	BOSH Director	2	8	103
Management	Enterprise PKS	2	8	80
Management	Harbor Container Registry	2	8	170

## B Cloud Foundation on VxRail VLANs

Category	VLAN	Description
Domain	External Management	VxRail cluster and ESXi hosts
	Internal Management	VxRail device discovery
	vMotion	Virtual machine migration
	vSAN	vSphere datastore
NSX	VXLAN	NSX VTEP/Overlay
	Uplink 1	NSX Edge Services
	Uplink 2	NSX Edge Services
Node Management	Out-of-band management	Dell PowerEdge iDRAC network (optional)
Stretched Cluster	Witness Traffic Separation	Segment traffic between witness and two sites (optional)

## C VxRail Network Configuration

Category	Detail	Description
VxRail	External Management	VLAN ID for the management VLAN that passes through 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)
		DNS server(s)
Management	ESXi hostnames and IP addresses	ESXi hostname prefix
		Separator
		Iterator
		Offset
		Suffix
		Domain
		Starting IP address for VxRail node pool
		Ending IP address for VxRail node pool
	vCenter & PSC Server	VxRail vCenter Server hostname
		VxRail vCenter Server IP address
		VxRail Platform Services Controller hostname
		VxRail Platform Services Controller IP address
	VxRail Manager	VxRail hostname
		VxRail IP address
Networking	Subnet mask	
	Gateway	
vMotion		Starting address for IP pool
		Ending address for IP pool
		Subnet mask
		VLAN ID
vSAN		Starting address for IP pool
		Ending address for IP pool
		Subnet mask
		VLAN ID
VM Networks	Minimum 1.	VM Network name and VLAN ID
Dell Node	iDRAC	IP address pool 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
2 <sup>nd</sup> Site	vMotion	Starting address for IP pool



		Ending address for IP pool
		Subnet mask
		VLAN ID
	VSAN	Starting address for IP pool
		Ending address for IP pool
		Subnet mask
	VXLAN	VLAN ID
		Starting address for IP pool
		Ending address for IP pool
		Subnet mask
	VLAN ID	

## D Cloud Builder and Management VI workload domain Configuration

Category	Detail	Description
CloudBuilder	IP Address	Temporary for CloudBuilder virtual appliance
Global	NTP	IP Address
	DNS	IP Address
	SSO Site Name	Same name as used for the VxRail cluster
	SSO Domain	
VXLAN	VLAN	VLAN ID for the VXLAN VLAN that passes through trunk ports connected to VxRail nodes
	DHCP Server	IP address of DHCP server to assign IP addresses to VTEP tunnel endpoints
	DHCP IP address range	Range of IP addresses in DHCP server to be assigned to VTEP tunnel endpoints
SDDC	Manager	Hostname
		IP Address
		Subnet Mask
NSX	Manager	Hostname
		IP Address
	Controller	Starting IP address for NSX Controller Pool
		Ending IP address for NSX Controller Pool
	Segment Range	Starting Segment ID
		Ending Segment ID
	Transport Zone	Name
	Multicast IP Range	Starting IP address for multicast range
		Ending IP address for multicast range

# E Application Virtual Network Configuration

Application Virtual Network Regions		
Region A VXLAN	Name	Name of universal logical switch
	IP Address	IP address range for Region A VXLAN network in CIDR format
	DNS Search Domain	DNS for Region A
	DNS Zone	DNS for Region A
xRegion VXLAN	Name	Name of universal logical switch
	IP Address	IP address range for xRegion VXLAN network in CIDR format
	DNS Search Domain	DNS for xRegion
	DNS Zone	DNS for xRegion

BGP Peering Networks	
ESG Uplink 1 VLAN	VLAN for eBGP peering for first ESG uplink
ESG Uplink 2 VLAN	VLAN for eBGP peering for second ESG uplink
ESG Uplink 1 Network	Network in CIDR format for eBGP peering with first ESG uplink and first router
ESG Uplink 2 Network	Network in CIDR format for eBGP peering with second ESG uplink and second router
UDLR-ESG Network	Network in CIDR format for UDLR assignment and iBGP peering with ESGs

Universal Distributed Logical Router	
Internal ASN	Autonomous System Number for Application Virtual Network
Name	Node name
Password	Neighbor password for BGP peering
Node 1 IP Address	IP address assigned to UDLR node 1
Node 2 IP Address	IP address assigned to UDLR node 2
Forwarding IP Address	Floating IP address used for sending packets
Protocol IP Address	IP address used for peering with Edge Service Gateways

External Routers		
External ASN	ASN Value	Autonomous System Number for external routers
External Router 1	IP Address	IP Address for peering with Edge Service Gateways
	Password	Neighbor password for BGP peering
External Router 2	IP Address	IP Address for peering with Edge Service Gateways
	Password	Neighbor password for BGP peering

Edge Service Gateways		
Edge Border Gateway 1	Name	Name of Edge Service Gateway 1
	Uplink 1 IP Address	IP address for peering with external router 1
	Uplink 2 IP Address	IP address for peering with external router 2
	UDLR 1 IP Address	IP address for iBGP peering with Universal Distributed Logical Router

Edge Border Gateway 2	Name	Name of Edge Service Gateway 2
	Uplink 1 IP Address	IP address for peering with external router 1
	Uplink 2 IP Address	IP address for peering with external router 2
	UDLR 2 IP Address	IP address for iBGP peering with Universal Distributed Logical Router

Universal Distributed Logical Router	
Internal ASN	Autonomous System Number for Application Virtual Network
Name	Node name
Password	Neighbor password for BGP peering
UDLR Network (CIDR)	IP address range reserved for UDLR in CIDR format
Node 1 IP Address	IP address assigned to UDLR node 1
Node 2 IP Address	IP address assigned to UDLR node 2
Forwarding IP Address	Floating IP address used for sending packets
Protocol IP Address	IP address used for peering with Edge Service Gateways

Universal Logical Router Settings	
NSX Universal Segment ID	
NSX Universal Multicast Address Range	
MTU Size	
Appliance Size	

AVN Region A Management Components		
Log Insight	Load Balancer	IP Address shared with all Log Insight nodes
	Master	IP Address of node 1
	Worker #1	IP Address of node 2
	Worker #2	IP Address of node 3

2nd Site BGP Peering Networks	
ESG Uplink 1 VLAN	VLAN for eBGP peering for first ESG uplink
ESG Uplink 2 VLAN	VLAN for eBGP peering for second ESG uplink
ESG Uplink 1 Network	Network in CIDR format for eBGP peering with first ESG uplink and first router
ESG Uplink 2 Network	Network in CIDR format for eBGP peering with second ESG uplink and second router
UDLR-ESG Network	Network in CIDR format for UDLR assignment and iBGP peering with ESGs

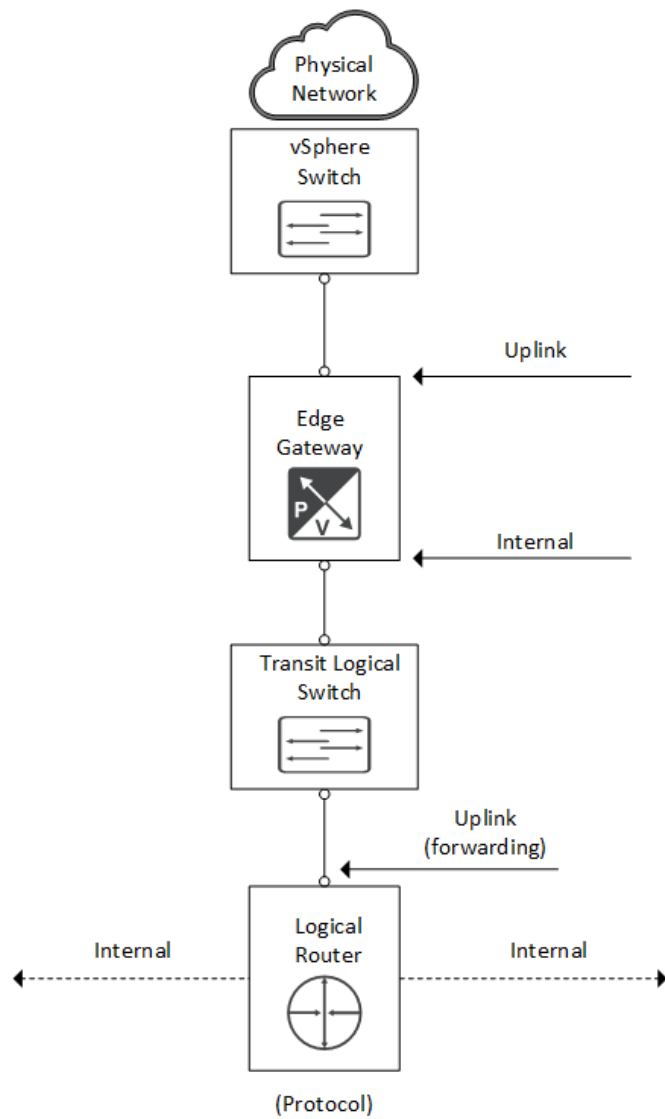
2nd Site External Routers		
External ASN	ASN Value	Autonomous System Number for external routers
External Router 1	IP Address	IP Address for peering with Edge Service Gateways
	Password	Neighbor password for BGP peering
External Router 2	IP Address	IP Address for peering with Edge Service Gateways
	Password	Neighbor password for BGP peering

## F VI workload domain Configuration Settings

Category	Detail	Description
Global	Domain Name	
	Organization Name	
Management	vCenter	Hostname
		IP Address
		Subnet Mask
		Default gateway
		Management Account
VXLAN	VLAN	VLAN ID for the VXLAN VLAN that passes through trunk ports connected to VxRail nodes
	DHCP Server	IP address of DHCP server to assign IP addresses to VTEP tunnel endpoints
	DHCP IP addresses	Range of IP addresses in DHCP server to be assigned to VTEP tunnel endpoints
NSX-V	Manager	Hostname
		IP Address
		Subnet Mask
		Default Gateway
	Controller	IP address for first NSX-V Controller
		IP address for second NSX-V Controller
		IP address for third NSX-V Controller
		Subnet Mask
		Default Gateway
NSX-T	Manager	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

# G NSX Routing Settings

Category	Description
Physical Router	IP address
	Password
	Autonomous System
Edge Service Gateway	Uplink 1 IP Address
	Uplink 2 IP Address
	Internal 1 IP Address
	Internal 2 IP Address
Logical Router	Forwarding IP Address
	Protocol IP Address



## H Example switch configuration settings for BGP peering

This sample syntax is for providing basic guidance on the settings that need to be performed on the top-of-rack switches to support BGP peering for the Application Virtual Network (AVN). The actual code required on the top-of-rack switches is dependent on the existing data center network infrastructure and routing standards.

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 1>
no shutdown
mtu 9216
ip address <Gateway IP address for AVN uplink 1>

interface vlan <VLAN for AVN Uplink 2>
no shutdown
mtu 9216
ip address <Gateway IP address for AVN uplink 2>

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
```