VMware Cloud Foundation Design Guide

07 NOV 2023

VMware Cloud Foundation 5.1

You can find the most up-to-date technical documentation on the VMware by Broadcom website at:

https://docs.vmware.com/

VMware by Broadcom
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About VMware Cloud Foundation Design Guide

The *VMware Cloud Foundation Design Guide* contains a design model for VMware Cloud Foundation (also called VCF) that is based on industry best practices for SDDC implementation.

The *VMware Cloud Foundation Design Guide* provides the supported design options for VMware Cloud Foundation, and a set of decision points, justifications, implications, and considerations for building each component.

Intended Audience

This *VMware Cloud Foundation Design Guide* is intended for cloud architects who are familiar with and want to use VMware Cloud Foundation to deploy and manage an SDDC that meets the requirements for capacity, scalability, backup and restore, and extensibility for disaster recovery support.

Before You Apply This Guidance

The sequence of the VMware Cloud Foundation documentation follows the stages for implementing and maintaining an SDDC.

To apply this *VMware Cloud Foundation Design Guide*, you must be acquainted with the *Getting Started with VMware Cloud Foundation* documentation and with the *VMware Cloud Foundation Release Notes*. See VMware Cloud Foundation documentation.

For performance best practices for vSphere, see Performance Best Practices for VMware vSphere 8.0 Update 1.

Design Elements

This *VMware Cloud Foundation Design Guide* contains requirements and recommendations for the design of each component of the SDDC. In situations where a configuration choice exists, requirements and recommendations are available for each choice. Implement only those that are relevant to your target configuration.

Design Element	Description
Requirement	Required for the operation of VMware Cloud Foundation. Deviations are not permitted.
Recommendation	Recommended as a best practice. Deviations are permitted.

VMware Cloud Foundation Deployment Options in This Design

This design guidance is for the all architecture models of VMware Cloud Foundation. By following the guidance, you can examine the design for these deployment options:

- Single VMware Cloud Foundation instance.
- Single VMware Cloud Foundation instance with multiple availability zones (also known as stretched deployment). The default vSphere cluster of the workload domain is stretched between two availability zones by using Chapter 6 vSAN Design for VMware Cloud Foundation and configuring vSphere Cluster Design Requirements and Recommendations for VMware Cloud Foundation and BGP Routing Design for VMware Cloud Foundation accordingly.
- Multiple VMware Cloud Foundation instances. You deploy several instances of VMware Cloud Foundation to address requirements for scale and co-location of users and resources.

  For disaster recovery, workload mobility, or propagation of common configuration to multiple VMware Cloud Foundation instances, you can deploy Chapter 8 NSX Design for VMware Cloud Foundation for the SDDC management and workload components.
- Multiple VMware Cloud Foundation instances with multiple availability zones. You apply the configuration for stretched clusters for a single VMware Cloud Foundation instance to one or more additional VMware Cloud Foundation instances in your environment.

vCenter Single Sign-On Options in This Design

This design guidance covers the topology with a single vCenter Single Sign-On domain in a VMware Cloud Foundation instance and the topology with several isolated vCenter Single Sign-On domains in a single instance. See vCenter Single Sign-On Design Requirements for VMware Cloud Foundation.

VMware Cloud Foundation Design Blueprints

You can follow design blueprints for selected architecture models and topologies that list the applicable design elements. See VMware Cloud Foundation Design Blueprints.

VMware Cloud Foundation Glossary

See the VMware Cloud Foundation Glossary for constructs, operations, and other terms specific to VMware Cloud Foundation. It is important to understand these constructs before continuing with this design guidance.

VMware Cloud Foundation Concepts

1

To design a VMware Cloud Foundation deployment, you need to understand certain *VMware Cloud Foundation* concepts.

Read the following topics next:

- Architecture Models and Workload Domain Types in VMware Cloud Foundation
- VMware Cloud Foundation Topologies
- VMware Cloud Foundation Design Blueprints

Architecture Models and Workload Domain Types in VMware Cloud Foundation

When you design a VMware Cloud Foundation deployment, you decide what architecture model, that is, standard or consolidated, and what workload domain types, for example, consolidated, isolated, or standard, to implement according to the requirements for hardware, expected number of workloads and workload domains, co-location of management and customer workloads, identity isolation, and other.

Architecture Models

Decide on a model according to your organization's requirements and your environment's resource capabilities. Implement a standard architecture for workload provisioning and mobility across VMware Cloud Foundation instances according to production best practices. If you plan to deploy a small-scale environment, or if you are working on an SDDC proof-of-concept, implement a consolidated architecture.

Figure 1-1. Choosing a VMware Cloud Foundation Architecture Model

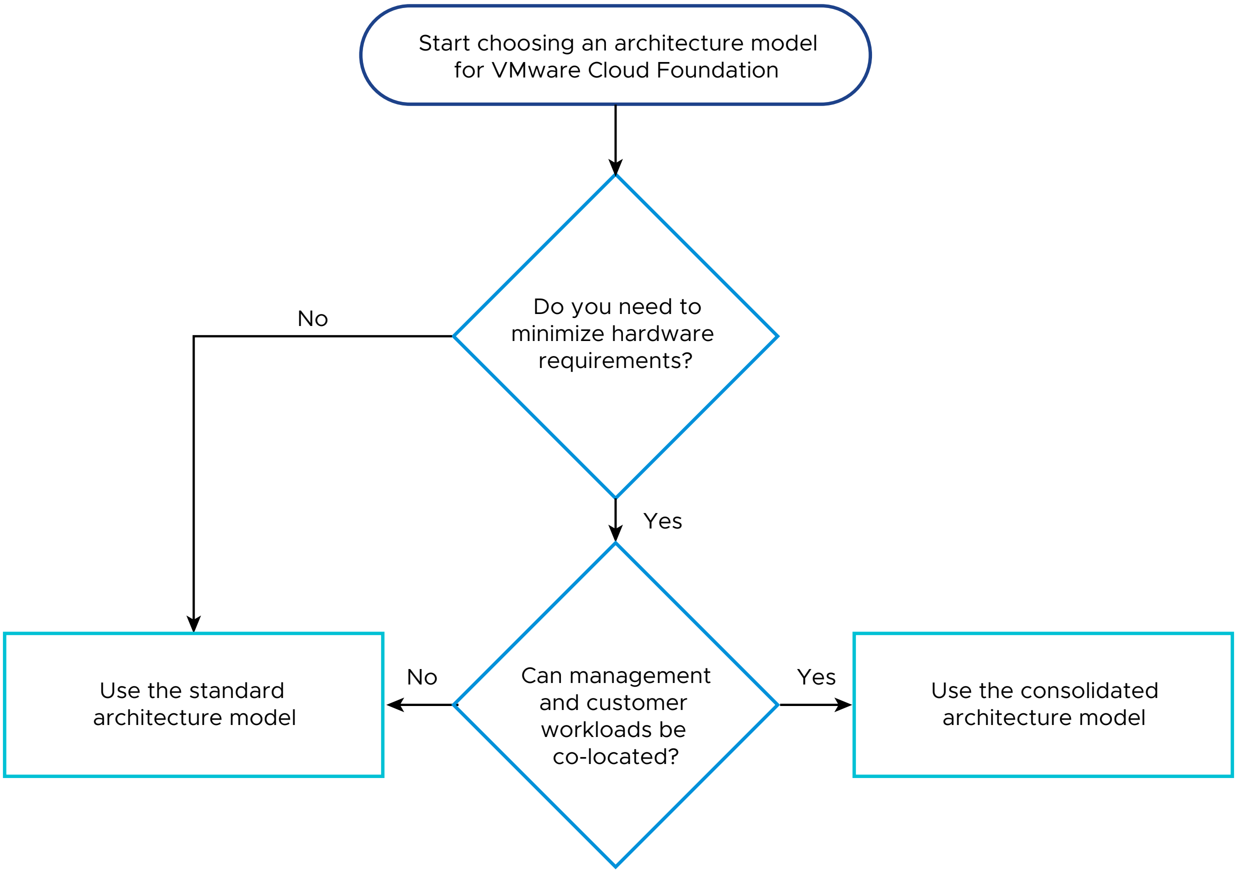
Start choosing an architecture model for VMware Cloud Foundation
Do you need to minimize hardware requirements?
No
Yes
Use the standard architecture model
No
Can management and customer workloads be co-located?
Yes
Use the consolidated architecture model

Table 1-1. Architecture Model Recommendations for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-ARCH-RCMD-CFG-001	Use the standard architecture model of VMware Cloud Foundation.	■ Aligns with the VMware best practice of separating management workloads from customer workloads. ■ Provides better long-term flexibility and expansion options.	Requires additional hardware.

Workload Domain Types

A workload domain represents a logical unit of application-ready infrastructure that groups ESXi hosts managed by a vCenter Server instance with specific characteristics according to VMware recommended practices. A workload domain can consist of one or more vSphere clusters, provisioned by SDDC Manager.

Table 1-2. Workload Domain Types

Workload Domain Type	Description	Benefits	Drawbacks
Management domain	■ First domain deployed. ■ Contains the following management appliances for all workload domains: ■ vCenter Server ■ NSX Manager ■ SDDC Manager ■ Optional. VMware Aria Suite components ■ Optional. Management domain NSX Edge nodes ■ Has dedicated ESXi hosts ■ First domain to upgrade.	■ Guaranteed sufficient resources for management components	■ You must carefully size the domain to accommodate planned deployment of VI workload domains and additional management components. ■ Hardware might not be fully utilized until full-scale deployment has been reached.
Consolidated domain	■ Represents a management domain which also runs customer workloads. ■ Uses resource pools to ensure sufficient resources for management components.	■ Considers the minimum possible initial hardware and management component footprint. ■ Can be scaled to a standard architecture model.	■ Management components and customer workloads are not isolated. ■ You must constantly monitor it to ensure sufficient resources for management components. ■ Migrating customer workloads to dedicated VI workloads domains is more complex.

Table 1-2. Workload Domain Types (continued)

Workload Domain Type	Description	Benefits	Drawbacks
VI workload domain	■ Represents an additional workload domain for running customer workloads. ■ Shares a vCenter Single Sign-On domain with the management domain. ■ Shares identity provider configuration with the management domain. ■ Has dedicated ESXi hosts.	■ Can share an NSX Manager instance with other VI workload domains. ■ All workload domains can be managed through a single pane of glass. ■ Minimizes password management overhead. ■ Allows for independent life cycle management.	This workload domain type cannot provide distinct vCenter Single Sign-On domains for customer workloads.
Isolated VI workload domain	■ Represents an additional workload domain for running customer workloads. ■ Has a distinct vCenter Single Sign-On domain. ■ Has a distinct identity provider configuration. ■ Has dedicated ESXi hosts.	■ Can provide distinct vCenter Single Sign-On domains for customer workloads. ■ Supports a scale beyond 14 VI workload domains. ■ Allows for independent life cycle management.	■ Workload domains of this type cannot share an NSX Manager instance with other VI workload domains. ■ Workload domain vCenter Server instances are managed through different panes of glass. ■ Additional password management overhead exists for administrators of VMware Cloud Foundation. ■ You can scale up to 24 VI workload domains per VMware Cloud Foundation instance.

Figure 1-2. Choosing a VMware Cloud Foundation Workload Domain Type for Customer Workloads

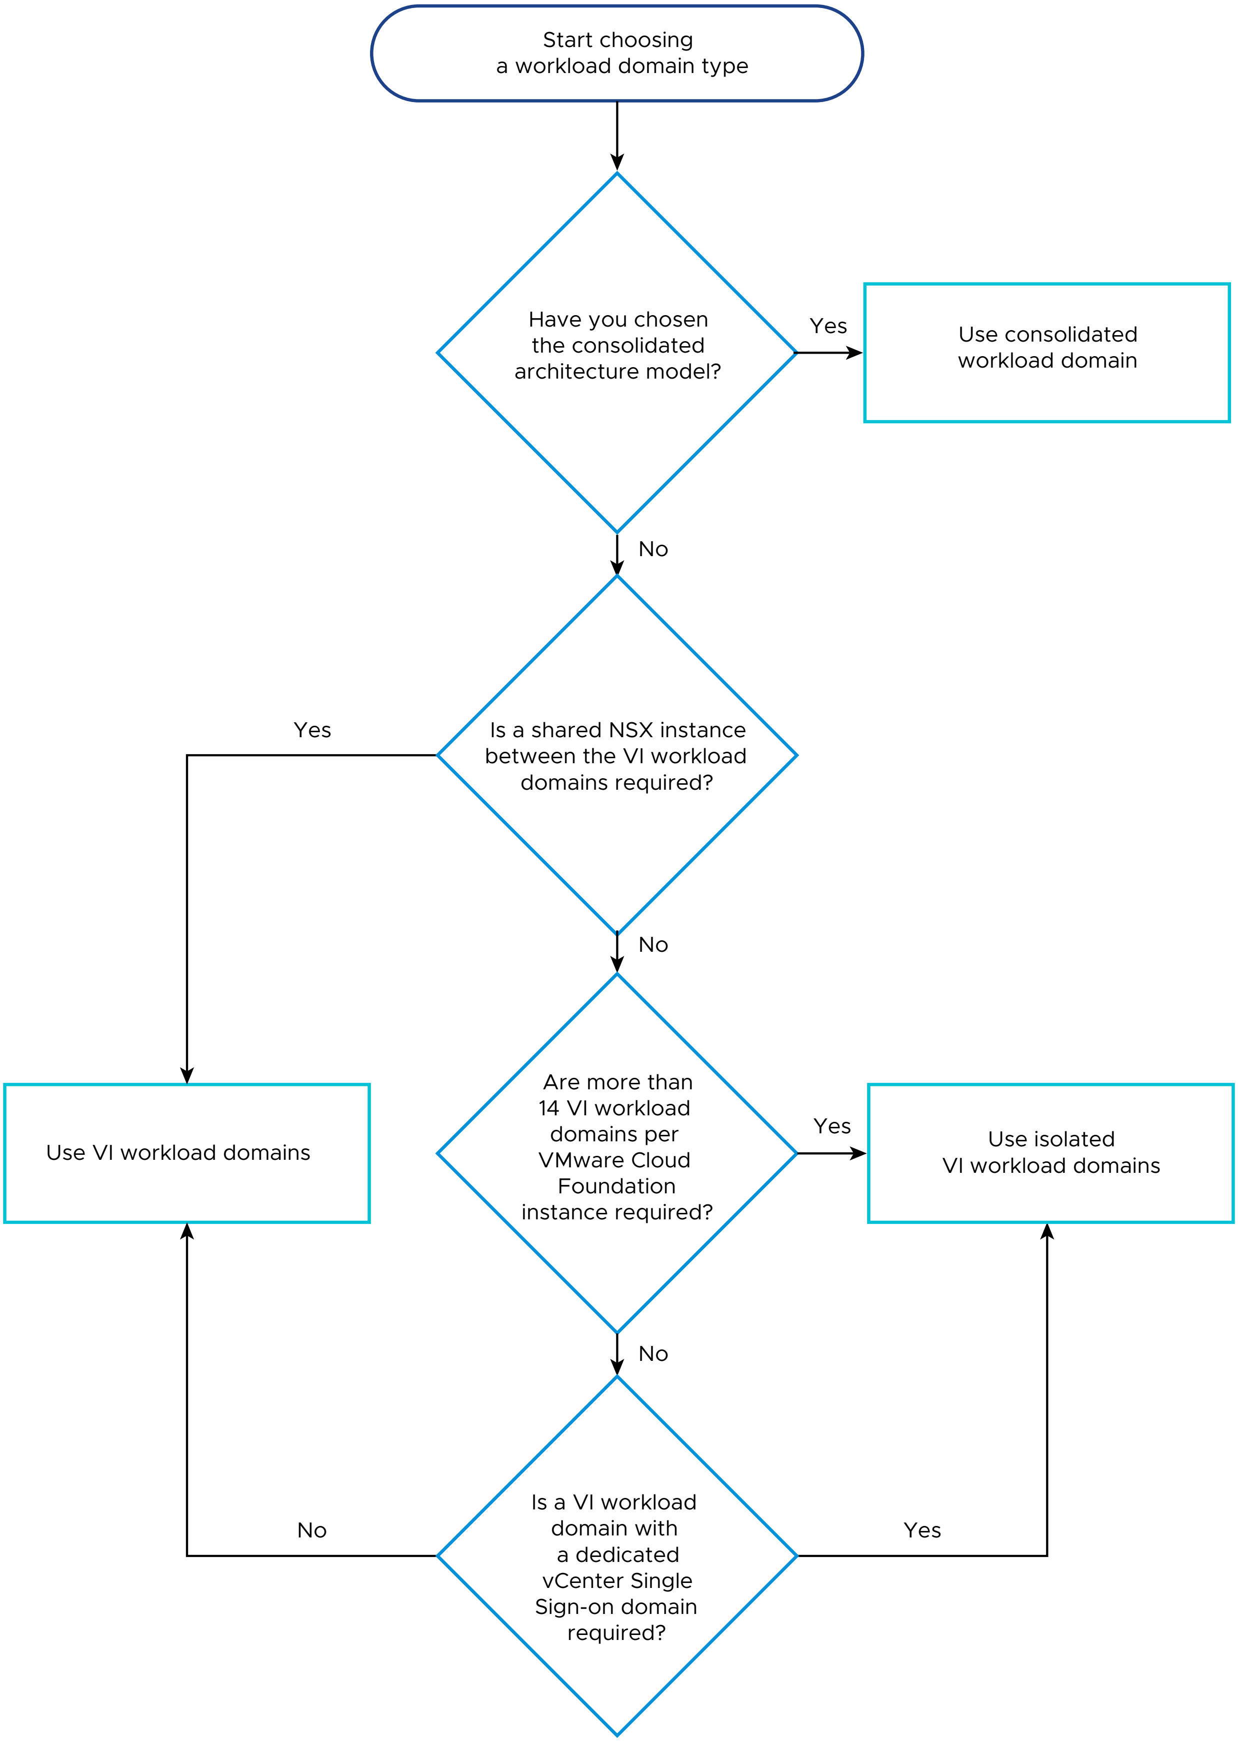

Table 1-3. Workload Domain Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WLD-RCMD-CFG-001	Use VI workload domains or isolated VI workload domains for customer workloads.	■ Aligns with the VMware best practice of separating management workloads from customer workloads. ■ Provides better long term flexibility and expansion options.	Requires additional hardware.

VMware Cloud Foundation Topologies

VMware Cloud Foundation supports multiple topologies that provide different levels of availability and scale.

Availability Zones and VMware Cloud Foundation Instances

Availability zone

An availability zone is a fault domain at the SDDC level.

You create multiple availability zones for the purpose of creating vSAN stretched clusters. Using multiple availability zones can improve availability of management components and workloads running within the SDDC, minimize downtime of services, and improve SLAs.

Availability zones are typically located either within the same data center, but in different racks, chassis, rooms, or in different data centers with low-latency high-speed links connecting them. One availability zone can contain several fault domains.

Note Only stretched clusters created by using the Stretch Cluster API, and are therefore vSAN storage based, are considered by and treated as stretched clusters by VMware Cloud Foundation.

VMware Cloud Foundation Instance

Each VMware Cloud Foundation instance is a separate VMware Cloud Foundation deployment and might contain one or two availability zones. VMware Cloud Foundation instances may be geographically separate.

VMware Cloud Foundation Topologies

Several topologies of VMware Cloud Foundation exist according to the number of availability zones and VMware Cloud Foundation instances.

Table 1-4. VMware Cloud Foundation Topologies

Topology	Description
Single Instance - Single Availability Zone	Workload domains are deployed in a single availability zone.
Single Instance - Multiple Availability Zones	Workload domains might be stretched between two availability zones.
Multiple Instances - Single Availability Zone per VMware Cloud Foundation instance	Workload domains in each instance are deployed in a single availability zone.
Multiple Instances - Multiple Availability Zones per VMware Cloud Foundation instance	Workload domains in each instance might be stretched between two availability zones.

Figure 1-3. Choosing a VMware Cloud Foundation Topology

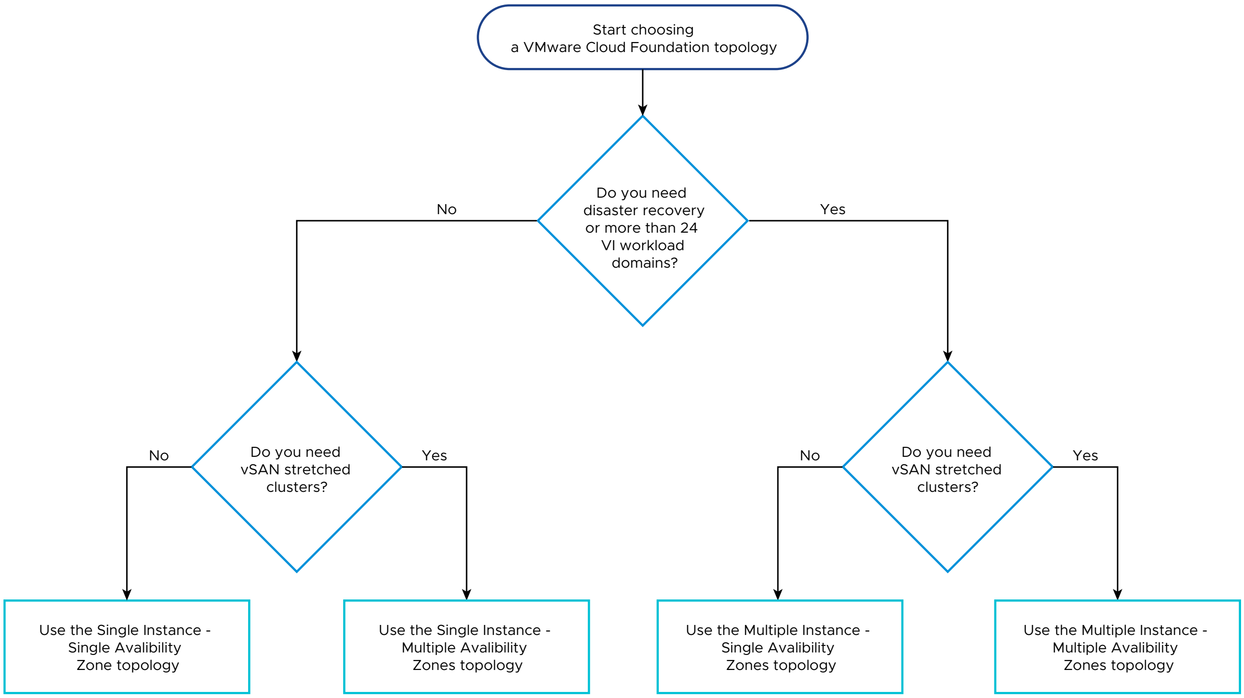
Start choosing
a VMware Cloud Foundation topology
Do you need disaster recovery or more than 24 VI workload domains?
No
Yes
Do you need vSAN stretched clusters?
No
Yes
Do you need vSAN stretched clusters?
No
Yes
Use the Single Instance - Single Avalibility Zone topology
Use the Single Instance - Multiple Avalibility Zones topology
Use the Multiple Instance - Single Avalibility Zones topology
Use the Multiple Instance - Multiple Avalibility Zones topology

What to read next

- Single Instance - Single Availability Zone

  Single Instance - Single Availability Zone is the simplest VMware Cloud Foundation topology where workload domains are deployed in a single availability zone.

- Single Instance - Multiple Availability Zones

  You protect your VMware Cloud Foundation environment against a failure of a single hardware fault domain by implementing multiple availability zones.

- Multiple Instances - Single Availability Zone per Instance

  You protect against a failure of a single VMware Cloud Foundation instance by implementing multiple VMware Cloud Foundation instances.

- Multiple Instances - Multiple Availability Zones per Instance

  You protect against a failure of a single VMware Cloud Foundation instance by implementing multiple VMware Cloud Foundation instances. Implementing multiple availability zones in an instance protects against a failure of a single hardware fault domain.

Single Instance - Single Availability Zone

Single Instance - Single Availability Zone is the simplest VMware Cloud Foundation topology where workload domains are deployed in a single availability zone.

The Single Instance - Single Availability Zone topology relies on vSphere HA to protect against host failures.

Figure 1-4. Single VMware Cloud Foundation Instance with a Single Availability Zone

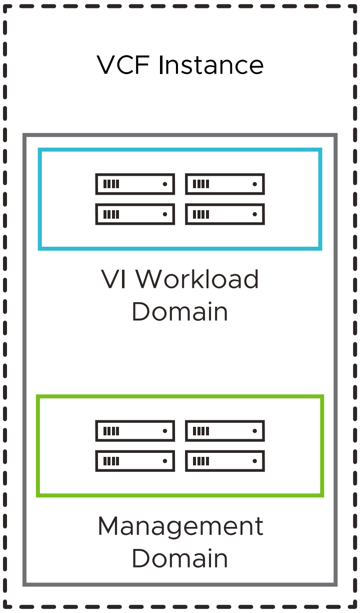

Table 1-5. Single Instance - Single Availability Zone Attributes

Attributes	Detail
Data centers	Single data center
Workload domain cluster rack mappings	■ Workload domain cluster in a single rack ■ Workload domain cluster spanning multiple racks
Scale	■ Up to 25 workload domains Up to 15 workload domains in a single vCenter Single Sign-On domain
Resilience	vSphere HA provides protection against host failures

Single Instance - Multiple Availability Zones

You protect your VMware Cloud Foundation environment against a failure of a single hardware fault domain by implementing multiple availability zones.

Incorporating multiple availability zones in your design can help reduce the blast radius of a failure and can increase application availability. You usually deploy multiple availability zones across two independent data centers.

Figure 1-5. Multiple Availability Zones in VMware Cloud Foundation

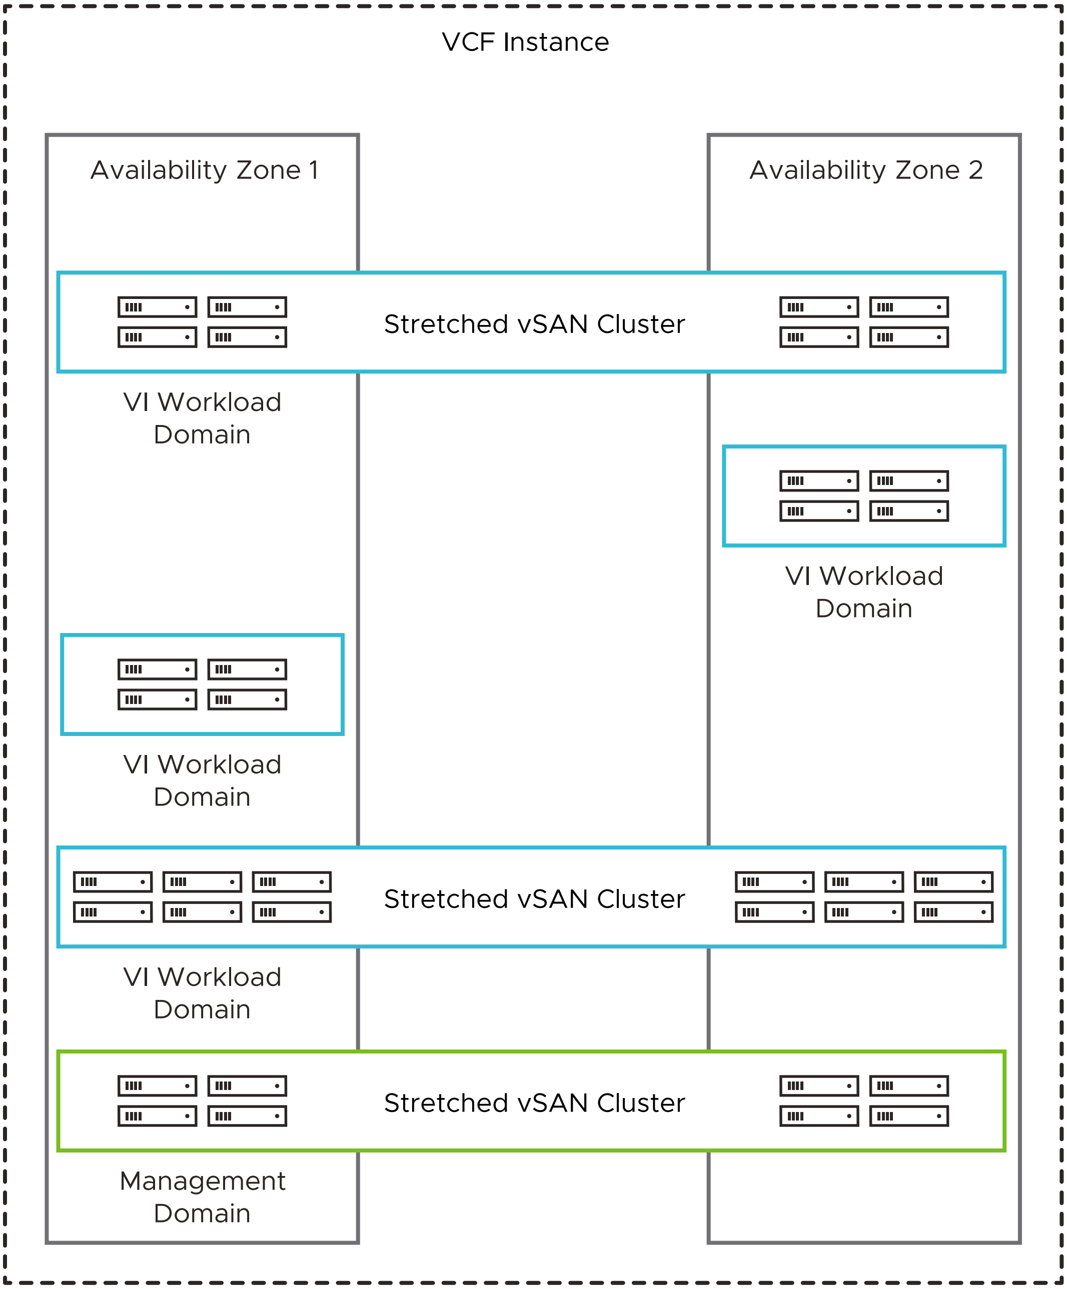

Table 1-6. Single Instance - Multiple Availability Zone Attributes

Attributes	Detail
Workload domain cluster rack mappings	■ Workload domain cluster in a single rack ■ Workload domain cluster spanning multiple racks ■ Workload domain cluster with multiple availability zones, each zone in a single rack ■ Workload domain cluster with multiple availability zones, each zone spanning multiple racks
Stretched cluster	■ Because availability zones use VMware vSAN™ stretched clusters, the bandwidth between the zones must be at least 10 Gbps and the round-trip latency must be less than 5 ms. ■ Having the management domain on a vSAN stretched cluster is a prerequisite to configure and implement vSAN stretched clusters in your VI workload domains. ■ You can have up to two availability zones.
Scale	■ Up to 25 workload domains. Up to 15 workload domains in a single vCenter Single Sign-On domain
Resilience	■ vSphere HA provides protection against host failures. ■ Multiple availability zones protect against data center failures.

Multiple Instances - Single Availability Zone per Instance

You protect against a failure of a single VMware Cloud Foundation instance by implementing multiple VMware Cloud Foundation instances.

Incorporating multiple VMware Cloud Foundation instances in your design can help reduce the blast radius of a failure and can increase application availability across larger geographical distances than cannot be achieved by using multiple availability zones. You usually deploy this topology in the same data center for scale or across independent data centers for resilience.

Figure 1-6. Multiple Instance - Single Availability Zone Topology for VMware Cloud Foundation

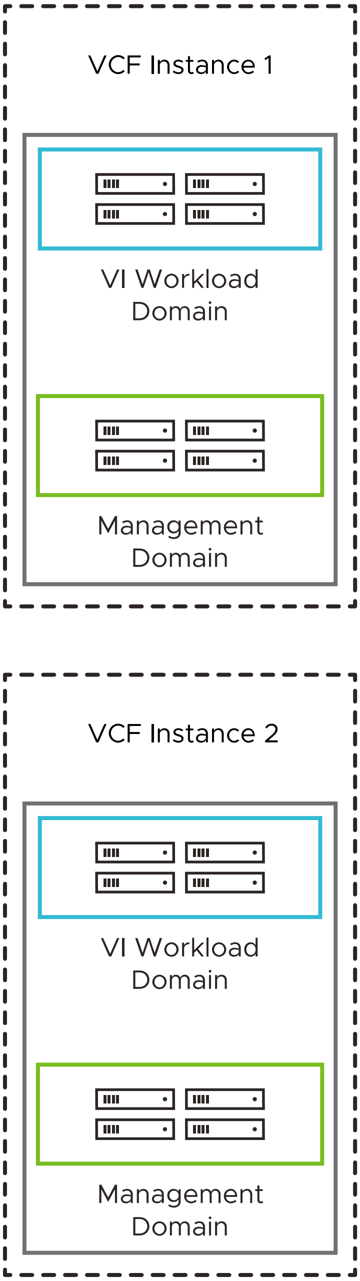

Table 1-7. Multiple Instance - Single Availability Zone Attributes

Attributes	Detail
Workload domain cluster rack mapping	■ Workload domain cluster in a single rack ■ Workload domain cluster spanning multiple racks
Multiple instances	Using multiple VMware Cloud Foundation instances can facilitate the following use cases: ■ Disaster recovery across different VMware Cloud Foundation instances at longer distances ■ Scale beyond the maximums of a single VMware Cloud Foundation instance. ■ Co-location of end users and resources If you plan to use NSX Federation between VMware Cloud Foundation instances, the following considerations exist: ■ Maximum of four locations when using medium-size NSX Global Managers ■ Up to 16 locations when using large-size NSX Global Managers ■ Maximum of four locations per cross-instance Tier-0 gateway ■ Life cycle management must be planned carefully
Scale	■ Up to 25 workload domains per VMware Cloud Foundation instance Up to 15 workload domains in a single vCenter Single Sign-on domain per instance
Resilience	■ vSphere HA provides protection against host failures. ■ Deploying multiple instances can protect against natural disasters by providing recovery locations at greater geographical distances.

Multiple Instances - Multiple Availability Zones per Instance

You protect against a failure of a single VMware Cloud Foundation instance by implementing multiple VMware Cloud Foundation instances. Implementing multiple availability zones in an instance protects against a failure of a single hardware fault domain.

Incorporating multiple VMware Cloud Foundation instances into your design can help reduce the blast radius of a failure and can increase application availability across larger geographical distances that cannot be achieved using multiple availability zones.

Figure 1-7. Multiple Instance - Multiple Availability Zones Topology for VMware Cloud Foundation

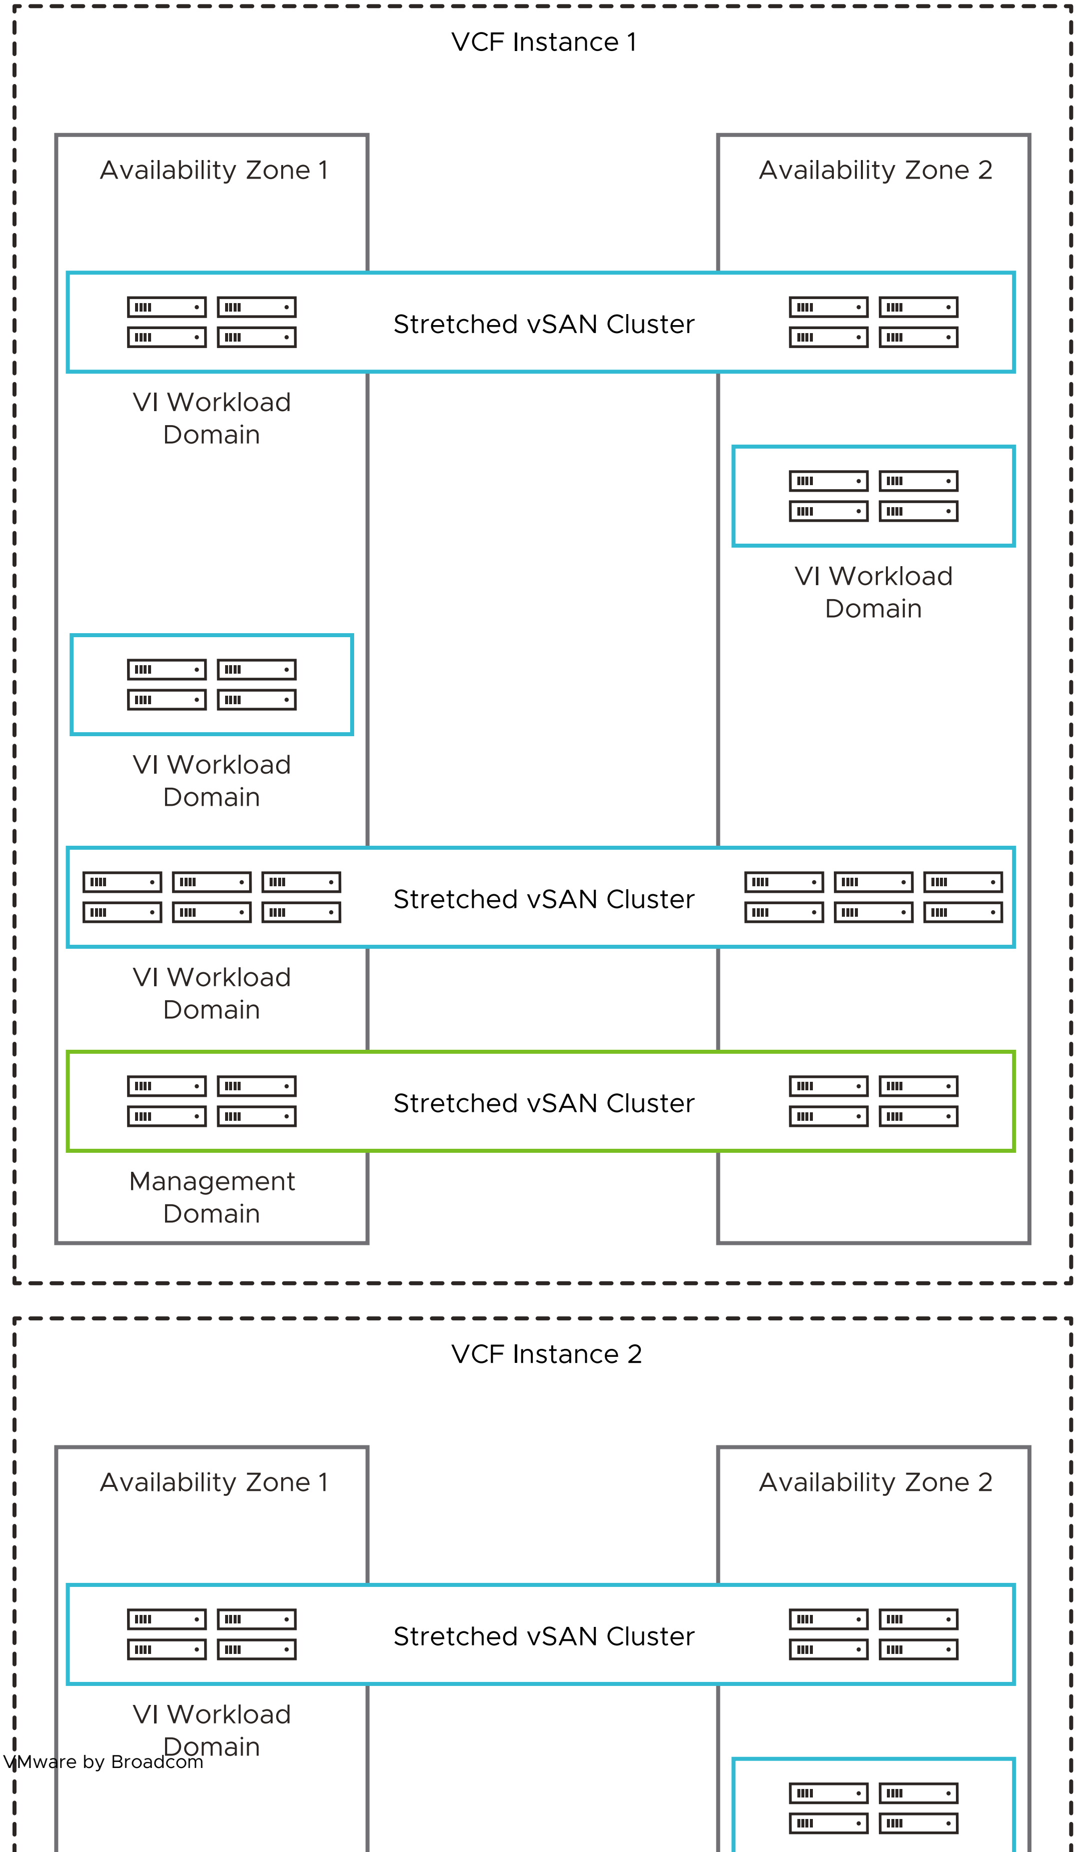

Table 1-8. Multiple Instance - Multiple Availability Zone Attributes

Attributes	Detail
Workload domain cluster rack mapping	■ Workload domain cluster in a single rack ■ Workload domain cluster spanning multiple racks ■ Workload domain cluster with multiple availability zones, each zone in a single rack ■ Workload domain cluster with multiple availability zones, each zone spanning multiple racks
Multiple instances	Using multiple VMware Cloud Foundation instances can facilitate the following: ■ Disaster recovery across different VMware Cloud Foundation instances at longer distances ■ Scale beyond the maximums of a single VMware Cloud Foundation instance ■ Co-location of end users and resources If you plan to use NSX Federation between instances, VMware Cloud Foundation consider the following: ■ Maximum of 4 locations when using meduim size global managers. ■ Up to 16 locations when using large size global managers. ■ Maximum of 4 locations per Stretched Tier-0 Gateway. ■ Lifecycle management will need to be carefully planned.
Stretched cluster	■ Because availability zones use VMware vSAN™ stretched clusters, the bandwidth between the zones must be at least 10 Gbps and the round-trip latency must be less than 5 ms. ■ You can have up to two availability zones. ■ Having the management domain on a vSAN stretched cluster is a prerequisite to configure and implement vSAN stretched clusters in your VI workload domains.
Scale	■ Up to 25 workload domains per VMware Cloud Foundation instance Up to 15 workload domains in a single vCenter Single Sign-On domain per instance
Resilience	■ vSphere HA provides protection against host failures. ■ Multiple availability zones protect against data center failures. ■ Multiple instances can protect against natural disasters by providing recovery locations at greater geographical distances.

VMware Cloud Foundation Design Blueprints

A *VMware Cloud Foundation* design blueprint is a collection of design requirements and recommendations based on a chosen architecture model, workload domain type, and topology. It can be used as a full end-to-end design for a *VMware Cloud Foundation* deployment.

Design Blueprint One: Multiple Instance - Multiple Availability Zone

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology that includes multiple VMware Cloud Foundation instances, each instance containing multiple availability zones, for an organization called Rainpole.

Design Choices for Design Blueprint One

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 1-9. Design Choices for Design Blueprint One

Design Aspect	Choice Made
Architecture Models and Workload Domain Types in VMware Cloud Foundation	Standard
Workload Domain Types	Management domain and VI workload domains
Multiple Instances - Single Availability Zone per Instance	Multiple Instances - Multiple Availability Zones per instance
Leaf-Spine Physical Network Design Requirements and Recommendations for VMware Cloud Foundation	Leaf-Spine
Routing Design for VMware Cloud Foundation	BGP
Chapter 6 vSAN Design for VMware Cloud Foundation	vSAN
Chapter 10 VMware Aria Suite Lifecycle Design for VMware Cloud Foundation	Included
Chapter 11 Workspace ONE Access Design for VMware Cloud Foundation	Standard Workspace ONE Access

Design Elements for Design Blueprint One

Table 1-10. External Services Design Elements

Design Area	Applicable Design Elements
External Services Design Elements for VMware Cloud Foundation	External Services Design Requirements

Table 1-11. Physical Network Design Elements

Design Area	Applicable Design Elements
Physical Network Design Elements for VMware Cloud Foundation	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Requirements for NSX Federation
	Leaf-Spine Physical Network Design Recommendations
	Leaf-Spine Physical Network Design Recommendations for Stretched Clusters
	Leaf-Spine Physical Network Design Recommendations for NSX Federation

Table 1-12. Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN Design Elements for VMware Cloud Foundation	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi Design Elements for VMware Cloud Foundation	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server Design Elements	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On Design Elements	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere Cluster Design Elements for VMware Cloud Foundation	vSphere Cluster Design Requirements
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere Networking Design Elements for VMware Cloud Foundation	vSphere Networking Design Recommendations
NSX Manager Design Elements	NSX Manager Design Requirements
	NSX Manager Design Recommendation

Table 1-12. Management Domain Design Elements (continued)

Design Area	Applicable Design Elements
	NSX Manager Design Recommendations for Stretched Clusters
NSX Global Manager Design Elements	NSX Global Manager Design Requirements for NSX Federation
	NSX Global Manager Design Recommendations for NSX Federation
	NSX Global Manager Design Recommendations for Stretched Clusters
NSX Edge Design Elements	NSX Edge Design Requirements
	NSX Edge Design Requirements for NSX Federation
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
BGP Routing Design Elements for VMware Cloud Foundation	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Requirements for NSX Federation
	BGP Routing Design Recommendations
	BGP Routing Design Recommendations for NSX Federation
Overlay Design Elements for VMware Cloud Foundation	Overlay Design Requirements
	Overlay Design Recommendations
Application Virtual Network Design Elements for VMware Cloud Foundation	Application Virtual Network Design Requirements
	Application Virtual Network Design Requirements for NSX Federation
Load Balancing Design Elements for VMware Cloud Foundation	Load Balancing Design Requirements
	Load Balancing Design Requirements for NSX Federation
SDDC Manager Design Elements for VMware Cloud Foundation	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 1-13. VI Workload Domain Design Elements

Design Area	Applicable Design Elements
vSAN Design Elements for VMware Cloud Foundation	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters

Table 1-13. VI Workload Domain Design Elements (continued)

Design Area	Applicable Design Elements
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi Design Elements for VMware Cloud Foundation	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server Design Elements	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On Design Elements	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single SSO Domain Topology
vSphere Cluster Design Elements for VMware Cloud Foundation	vSphere Cluster Design Requirements VMware Cloud Foundation
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere Networking Design Elements for VMware Cloud Foundation	vSphere Networking Design Recommendations
NSX Manager Design Elements	NSX Manager Design Requirements
	NSX Manager Design Recommendations
	NSX Manager Design Recommendations for Stretched Clusters
NSX Global Manager Design Elements	NSX Global Manager Design Requirements for NSX Federation
	NSX Global Manager Design Recommendations for NSX Federation
	NSX Global Manager Design Recommendations for Stretched Clusters
NSX Edge Design Elements	NSX Edge Design Requirements
	NSX Edge Design Requirements for NSX Federation
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
BGP Routing Design Elements for VMware Cloud Foundation	BGP Routing Design Requirements

Table 1-13. VI Workload Domain Design Elements (continued)

Design Area	Applicable Design Elements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Requirements for NSX Federation
	BGP Routing Design Recommendations
	BGP Routing Design Recommendations for NSX Federation
Overlay Design Elements for VMware Cloud Foundation	Overlay Design Requirements
	Overlay Design Recommendations

Table 1-14. VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle Design Elements for VMware Cloud Foundation	VMware Aria Suite Lifecycle Design Requirements
	VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters
	VMware Aria Suite Lifecycle Design Requirements for NSX Federation
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access Design Elements for VMware Cloud Foundation	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Requirements for Stretched Clusters
	Workspace ONE Access Design Requirements for NSX Federation
	Workspace ONE Access Design Recommendations

Table 1-15. Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life Cycle Management Design Elements for VMware Cloud Foundation	Life Cycle Management Design Requirements

Table 1-16. Account and Password Management Design Elements

Design Area	Applicable Design Elements
Information Security Design Elements for VMware Cloud Foundation	Account and Password Management Design Recommendations

Table 1-17. Certificate Management Design Elements

Design Area	Applicable Design Elements
Information Security Design Elements for VMware Cloud Foundation	Certificate Management Design Recommendations

Design Blueprint Two: Single Instance - Multiple Availability Zones

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology that includes one VMware Cloud Foundation instance with multiple availability zones for an organization called Rainpole.

Design Choices for Design Blueprint Two

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 1-18. Design Choices for Design Blueprint Two

Design Aspect	Choice Made
Architecture Models and Workload Domain Types in VMware Cloud Foundation	Standard
Workload Domain Types	Management domain and VI workload domains
Multiple Instances - Single Availability Zone per Instance	Single Instance - Multiple Availability Zones per instance
Leaf-Spine Physical Network Design Requirements and Recommendations for VMware Cloud Foundation	Leaf-Spine
Routing Design for VMware Cloud Foundation	BGP
Chapter 6 vSAN Design for VMware Cloud Foundation	vSAN
Chapter 10 VMware Aria Suite Lifecycle Design for VMware Cloud Foundation	Included
Chapter 11 Workspace ONE Access Design for VMware Cloud Foundation	Standard Workspace ONE Access

Design Elements for Design Blueprint Two

Table 1-19. External Services Design Elements

Design Area	Applicable Design Elements
External Services Design Elements for VMware Cloud Foundation	External Services Design Requirements

Table 1-20. Physical Network Design Elements

Design Area	Applicable Design Elements
Physical Network Design Elements for VMware Cloud Foundation	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Recommendations
	Leaf-Spine Physical Network Design Recommendations for Stretched Clusters

Table 1-21. Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN Design Elements for VMware Cloud Foundation	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi Design Elements for VMware Cloud Foundation	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server Design Elements	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On Design Elements	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere Cluster Design Elements for VMware Cloud Foundation	vSphere Cluster Design Requirements
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere Networking Design Elements for VMware Cloud Foundation	vSphere Networking Design Recommendations
NSX Manager Design Elements	NSX Manager Design Requirements
	NSX Manager Design Recommendation
	NSX Manager Design Recommendations for Stretched Clusters
NSX Edge Design Elements	NSX Edge Design Requirements
	NSX Edge Design Recommendations

Table 1-21. Management Domain Design Elements (continued)

Design Area	Applicable Design Elements
	NSX Edge Design Recommendations for Stretched Clusters
BGP Routing Design Elements for VMware Cloud Foundation	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Recommendations
Overlay Design Elements for VMware Cloud Foundation	Overlay Design Requirements
	Overlay Design Recommendations
Application Virtual Network Design Elements for VMware Cloud Foundation	Application Virtual Network Design Requirements
Load Balancing Design Elements for VMware Cloud Foundation	Load Balancing Design Requirements
SDDC Manager Design Elements for VMware Cloud Foundation	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 1-22. VI Workload Domain Design Elements

Design Area	Applicable Design Elements
vSAN Design Elements for VMware Cloud Foundation	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi Design Elements for VMware Cloud Foundation	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server Design Elements	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On Design Elements	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single SSO Domain Topology
vSphere Cluster Design Elements for VMware Cloud Foundation	vSphere Cluster Design Requirements VMware Cloud Foundation
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations

Table 1-22. VI Workload Domain Design Elements (continued)

Design Area	Applicable Design Elements
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere Networking Design Elements for VMware Cloud Foundation	vSphere Networking Design Recommendations
NSX Manager Design Elements	NSX Manager Design Requirements
	NSX Manager Design Recommendations
	NSX Manager Design Recommendations for Stretched Clusters
NSX Edge Design Elements	NSX Edge Design Requirements
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
BGP Routing Design Elements for VMware Cloud Foundation	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Recommendations
Overlay Design Elements for VMware Cloud Foundation	Overlay Design Requirements
	Overlay Design Recommendations

Table 1-23. VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle Design Elements for VMware Cloud Foundation	VMware Aria Suite Lifecycle Design Requirements
	VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access Design Elements for VMware Cloud Foundation	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Requirements for Stretched Clusters
	Workspace ONE Access Design Recommendations

Table 1-24. Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life Cycle Management Design Elements for VMware Cloud Foundation	Life Cycle Management Design Requirements

Table 1-25. Account and Password Management Design Elements

Design Area	Applicable Design Elements
Information Security Design Elements for VMware Cloud Foundation	Account and Password Management Design Recommendations

Table 1-26. Certificate Management Design Elements

Design Area	Applicable Design Elements
Information Security Design Elements for VMware Cloud Foundation	Certificate Management Design Recommendations

Design Blueprint Three: Single Instance - Consolidated

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology for an organization called Rainpole which includes one VMware Cloud Foundation instance where the management domain runs both management and customer workloads in a single availability zone.

Design Choices for Design Blueprint Three

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 1-27. Design Choices for Design Blueprint Three

Design Aspect	Choice Made
Architecture Models and Workload Domain Types in VMware Cloud Foundation	Consolidated
Workload Domain Types	Consolidated
Multiple Instances - Single Availability Zone per Instance	Consolidated topology
Leaf-Spine Physical Network Design Requirements and Recommendations for VMware Cloud Foundation	Leaf-Spine
Routing Design for VMware Cloud Foundation	BGP
Chapter 6 vSAN Design for VMware Cloud Foundation	vSAN
Chapter 10 VMware Aria Suite Lifecycle Design for VMware Cloud Foundation	Included
Chapter 11 Workspace ONE Access Design for VMware Cloud Foundation	Standard Workspace ONE Access

Design Elements for Design Blueprint Three

Table 1-28. External Services Design Elements

Design Area	Applicable Design Elements
External Services Design Elements for VMware Cloud Foundation	External Services Design Requirements

Table 1-29. Physical Network Design Elements

Design Area	Applicable Design Elements
Physical Network Design Elements for VMware Cloud Foundation	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Recommendations

Table 1-30. Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN Design Elements for VMware Cloud Foundation	vSAN Design Requirements
	vSAN Design Recommendations
ESXi Design Elements for VMware Cloud Foundation	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server Design Elements	vCenter Server Design Requirements
	vCenter Server Design Recommendations
vCenter Single Sign-On Design Elements	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere Cluster Design Elements for VMware Cloud Foundation	vSphere Cluster Design Requirements
	vSphere Cluster Design Recommendations
vSphere Networking Design Elements for VMware Cloud Foundation	vSphere Networking Design Recommendations
NSX Manager Design Elements	NSX Manager Design Requirements
	NSX Manager Design Recommendation
NSX Edge Design Elements	NSX Edge Design Requirements
	NSX Edge Design Recommendations
BGP Routing Design Elements for VMware Cloud Foundation	BGP Routing Design Requirements
	BGP Routing Design Recommendations
Overlay Design Elements for VMware Cloud Foundation	Overlay Design Requirements
	Overlay Design Recommendations

Table 1-30. Management Domain Design Elements (continued)

Design Area	Applicable Design Elements
Application Virtual Network Design Elements for VMware Cloud Foundation	Application Virtual Network Design Requirements
Load Balancing Design Elements for VMware Cloud Foundation	Load Balancing Design Requirements
SDDC Manager Design Elements for VMware Cloud Foundation	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 1-31. VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle Design Elements for VMware Cloud Foundation	VMware Aria Suite Lifecycle Design Requirements
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access Design Elements for VMware Cloud Foundation	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Recommendations

Table 1-32. Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life Cycle Management Design Elements for VMware Cloud Foundation	Life Cycle Management Design Requirements

Table 1-33. Account and Password Management Design Elements

Design Area	Applicable Design Elements
Information Security Design Elements for VMware Cloud Foundation	Account and Password Management Design Recommendations

Table 1-34. Certificate Management Design Elements

Design Area	Applicable Design Elements
Information Security Design Elements for VMware Cloud Foundation	Certificate Management Design Recommendations

Workload Domain Cluster to Rack Mapping in VMware Cloud Foundation

2

VMware Cloud Foundation distributes the functionality of the SDDC across multiple workload domains and vSphere clusters. A workload domain, whether it is the management workload domain or a VI workload domain, is a logical abstraction of compute, storage, and network capacity, and consists of one or more clusters. Each cluster can exist vertically in a single rack or be spanned horizontally across multiple racks.

The relationship between workload domain clusters and data center racks in VMware Cloud Foundation is not one-to-one. While a workload domain cluster is an atomic unit of repeatable building blocks, a rack is a unit of size. Because workload domain clusters can have different sizes, you map workload domain clusters to data center racks according to your requirements and physical infrastructure constraints. You determine the total number of racks for each cluster type according to your scalability needs.

Table 2-1. Workload Domain Cluster to Rack Configuration Options

Workload Domain Cluster to Rack Configuration	Description
Workload domain cluster in a single rack	The workload domain cluster occupies a single rack.
Workload domain cluster spanning multiple racks	■ The management domain can span multiple racks if the data center fabric can provide Layer 2 adjacency, such as BGP EVPN, between racks. If the Layer 3 fabric does not support this requirement, then the management cluster should be mapped to a single rack. ■ A VI workload domain can span multiple racks. If you are using a Layer 3 network fabric, NSX Edge clusters cannot be hosted on clusters that span racks.

Table 2-1. Workload Domain Cluster to Rack Configuration Options (continued)

Workload Domain Cluster to Rack Configuration	Description
Workload domain cluster with multiple availability zones, each zone in a single rack	To span multiple availability zones, the network fabric must support stretched Layer 2 networks and Layer 3 routed networks between the availability zones.
Workload domain cluster with multiple availability zones, each zone spanning multiple racks	■ A VI workload domain cluster with customer workloads and no NSX Edge clusters can span racks by using Layer 3 network fabric without Layer 2 adjacency between racks. If you are using a Layer 3 network fabric, NSX Edge clusters cannot be hosted on clusters that span racks. ■ To span multiple racks, the network fabric must support stretched Layer 2 networks between these racks if NSX Edge clusters are deployed on the vSphere cluster. ■ To span multiple availability zones, the network fabric must support stretched Layer 2 networks and Layer 3 routed networks between availability zones.

Figure 2-1. Workload Domains in a Single Rack

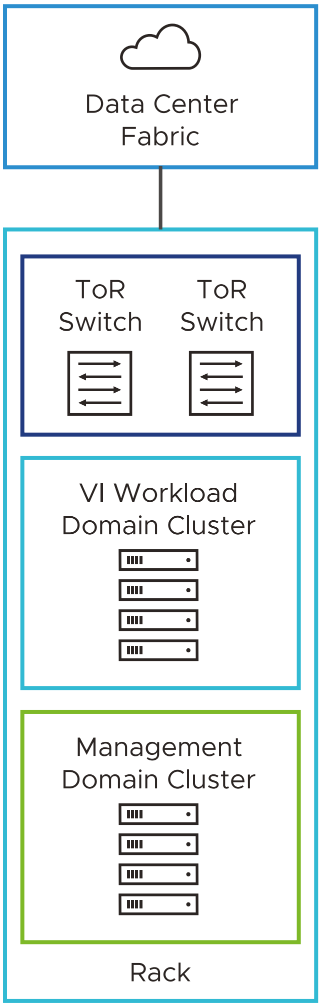

Figure 2-2. Workload Domain Spanning Multiple Racks

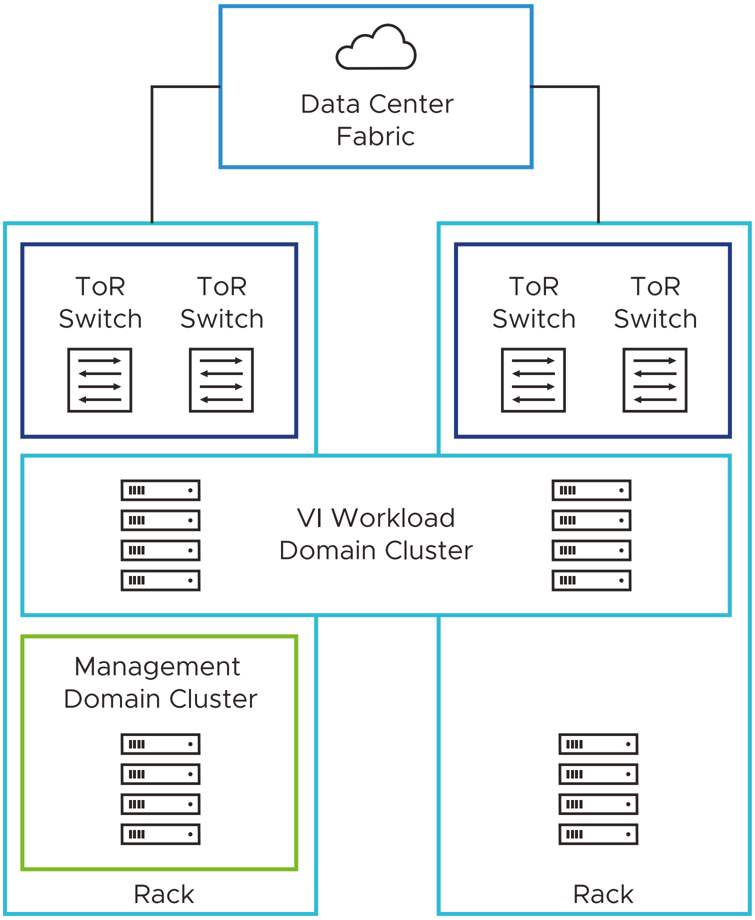

Figure 2-3. Workload Domains with Multiple Availability Zones, Each Zone in One Rack

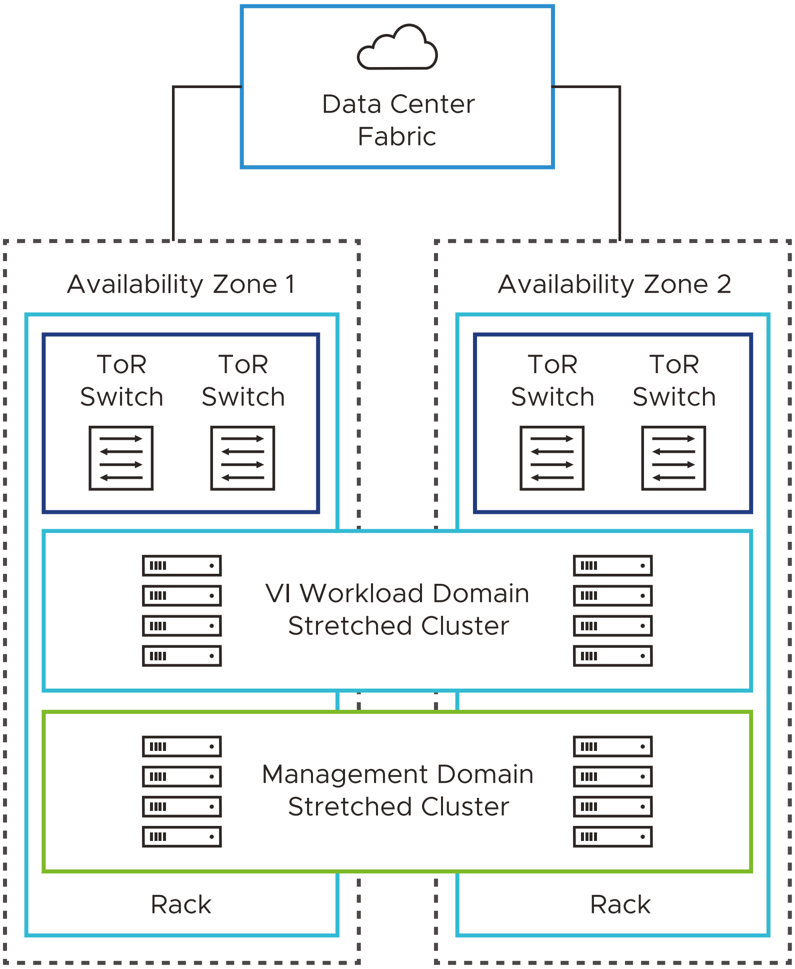

Supported Storage Types for VMware Cloud Foundation

3

Storage design for VMware Cloud Foundation includes the design for principal and supplemental storage.

Principal storage is used during the creation of a workload domain and is capable of running workloads. Supplemental storage can be added after the creation of a workload domain and can be capable of running workloads or be used for data at rest storage such as virtual machine templates, backup data, and ISO images.

Special considerations apply if you plan to add clusters to the management domain, for example, to separate additional management components that require specific hardware resources or might impact the performance of the main management components in the default cluster, or, in the case of the consolidated architecture of VMware Cloud Foundation, to separate customer workloads from the management components.

VMware Cloud Foundation supports the following principal and supplemental storage combinations:

Table 3-1. Supported Storage Types in VMware Cloud Foundation

Storage Type	Management Domain	VI Workload Domain
vSAN Original Storage Architecture (OSA)	Principal	Principal
vSAN Express Storage Architecture (ESA)	Principal	Principal
VMware vSAN Max™	Not Supported	Not Supported
Cross-cluster capacity sharing (HCI Mesh)	Supplemental	■ Principal (additional clusters only) ■ Supplemental
VMware vSphere® Virtual Volumes™ (FC, iSCSI, or NFS)	Supplemental	■ Principal ■ Supplemental
VMFS on FC	Supplemental	■ Principal ■ Supplemental
NFS	Supplemental (NFS 3 and NFS 4.1)	■ Principal (NFS 3) ■ Supplemental (NFS 3 and NFS 4.1)
iSCSI	Supplemental	Supplemental

Table 3-1. Supported Storage Types in VMware Cloud Foundation (continued)

Storage Type	Management Domain	VI Workload Domain
NVMe/TCP	Supplemental	Supplemental
NVMe/FC	Supplemental	Supplemental

Note For a consolidated VMware Cloud Foundation architecture model, the storage types that are supported for the management domain apply.

External Services Design for VMware Cloud Foundation

4

IP addressing scheme, name resolution, and time synchronization must support the requirements for VMware Cloud Foundation deployments.

Table 4-1. External Services Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-EXT-REQD-NET-001	Allocate statically assigned IP addresses and host names for all workload domain components.	Ensures stability across the VMware Cloud Foundation instance, and makes it simpler to maintain, track, and implement a DNS configuration.	You must provide precise IP address management.
VCF-EXT-REQD-NET-002	Configure forward and reverse DNS records for all workload domain components.	Ensures that all components are accessible by using a fully qualified domain name instead of by using IP addresses only. It is easier to remember and connect to components across the VMware Cloud Foundation instance.	You must provide DNS records for each component.
VCF-EXT-REQD-NET-003	Configure time synchronization by using an internal NTP time source for all workload domain components.	Ensures that all components are synchronized with a valid time source.	An operational NTP service must be available in the environment.
VCF-EXT-REQD-NET-004	Set the NTP service for all workload domain components to start automatically.	Ensures that the NTP service remains synchronized after you restart a component.	None.

Physical Network Infrastructure Design for VMware Cloud Foundation

5

Design of the physical data center network includes defining the network topology for connecting physical switches and ESXi hosts, determining switch port settings for VLANs and link aggregation, and designing routing.

A software-defined network (SDN) both integrates with and uses components of the physical data center. SDN integrates with your physical network to support east-west transit in the data center and north-south transit to and from the SDDC networks.

Several typical data center network deployment topologies exist:

- Core-Aggregation-Access
- Leaf-Spine
- Hardware SDN

Note Leaf-Spine is the default data center network deployment topology used for VMware Cloud Foundation.

Read the following topics next:

- VLANs and Subnets for VMware Cloud Foundation
- Leaf-Spine Physical Network Design Requirements and Recommendations for VMware Cloud Foundation

VLANs and Subnets for VMware Cloud Foundation

Configure your VLANs and subnets according to the guidelines and requirements for VMware Cloud Foundation.

When designing the VLAN and subnet configuration for your VMware Cloud Foundation deployment, consider the following guidelines:

Table 5-1. VLAN and Subnet Guidelines for VMware Cloud Foundation

All Deployment Topologies	Multiple Availability Zones	NSX Federation Between Multiple VMware Cloud Foundation Instances
■ Ensure your subnets are scaled appropriately to allow for expansion as expanding at a later time can be disruptive. ■ Use the IP address of the floating interface for Virtual Router Redundancy Protocol (VRPP) or Hot Standby Routing Protocol (HSRP) as the gateway. ■ Use the RFC 1918 IPv4 address space for these subnets and allocate one octet by VMware Cloud Foundation instance and another octet by function.	■ For network segments which are stretched between availability zones, the VLAN ID must meet the following requirements: ■ Be the same in both availability zones with the same Layer 3 network segments. ■ Have a Layer 3 gateway at the first hop that is highly available such that it tolerates the failure of an entire availability zone. ■ For network segments of the same type which are not stretched between availability zones, the VLAN ID can be the same or different between the zones.	■ An RTEP network segment should have a VLAN ID and Layer 3 range that are specific to the VMware Cloud Foundation instance. ■ In a VMware Cloud Foundation instance with multiple availability zones, the RTEP network segment must be stretched between the zones and assigned the same VLAN ID and IP range. ■ All Edge RTEP networks must reach each other.

When deploying VLANs and subnets for VMware Cloud Foundation, they must conform to the following requirements according to the VMware Cloud Foundation topology:

Figure 5-1. Choosing a VLAN Model for Host and Management VM Traffic

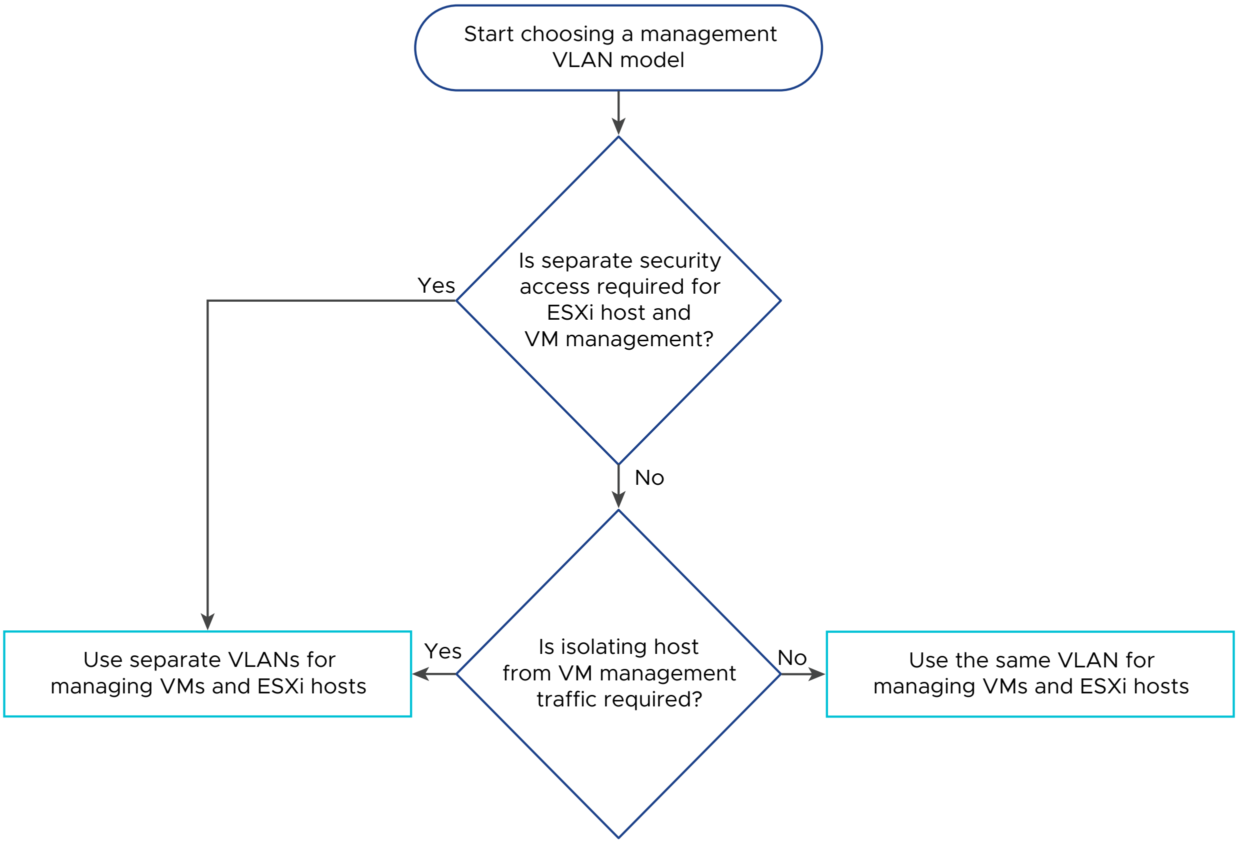
Start choosing a management VLAN model
Is separate security access required for ESXi host and VM management?
Yes
No
Is isolating host from VM management traffic required?
Yes
No
Use separate VLANs for managing VMs and ESXi hosts
Use the same VLAN for managing VMs and ESXi hosts

Table 5-2. VLANs and Subnets for VMware Cloud Foundation

Function	VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
VM management	■ Required ■ Highly available gateway within the instance	■ Required ■ Must be stretched within the instance ■ Highly available gateway across availability zones within the instance
Host management - first availability zone	■ Required ■ Highly available gateway within the instance	■ Required ■ Highly available gateway across availability zones within the instance
vSphere vMotion - first availability zone	■ Required ■ Highly available gateway within the instance	■ Required ■ Highly available gateway in first availability zone within the instance
vSAN - first availability zone	■ Required ■ Highly available gateway within the instance	■ Required ■ Highly available gateway in first availability zone within the instance

Table 5-2. VLANs and Subnets for VMware Cloud Foundation (continued)

Function	VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
Host overlay - first availability zone	■ Required ■ Highly available gateway within the instance	■ Required ■ Highly available gateway in first availability zone within the instance
Uplink01	■ Required ■ Gateway optional	■ Required ■ Gateway optional ■ Must be stretched within the instance
Uplink02	■ Required ■ Gateway optional	■ Required ■ Gateway optional ■ Must be stretched within the instance
Edge overlay	■ Required ■ Highly available gateway within the instance	■ Required ■ Must be stretched within the instance ■ Highly available gateway across availability zones within the instance
Host management - second availability zone	■ Not required	■ Required ■ Highly available gateway in second availability zone within the instance
vSphere vMotion - second availability zone	■ Not required	■ Required ■ Highly available gateway in second availability zone within the instance
vSAN - second availability zone	■ Not required	■ Required ■ Highly available gateway in second availability zone within the instance
Host overlay - second availability zone	■ Not required	■ Required ■ Highly available gateway in second availability zone within the instance
Edge RTEP	■ Required for NSX Federation only ■ Highly available gateway within the instance	■ Required for NSX Federation only ■ Must be stretched within the instance ■ Highly available gateway across availability zones within the instance
Management and Witness - witness appliance at a third location	■ Not required	■ Required ■ Highly available gateway at the witness location

Leaf-Spine Physical Network Design Requirements and Recommendations for VMware Cloud Foundation

Leaf-Spine is the default data center network deployment topology used for VMware Cloud Foundation. Consider network bandwidth, trunk port configuration, jumbo frames and routing configuration for NSX in a deployment with a single or multiple VMware Cloud Foundation instances.

Leaf-Spine Physical Network Logical Design

Each ESXi host is connected redundantly to the top-of-rack (ToR) switches of the SDDC network fabric by two 25-GbE ports. The ToR switches are configured to provide all necessary VLANs using an 802.1Q trunk. These redundant connections use features in vSphere Distributed Switch and NSX to guarantee that no physical interface is overrun and available redundant paths are used.

Figure 5-2. Leaf-Spine Physical Network Logical Design

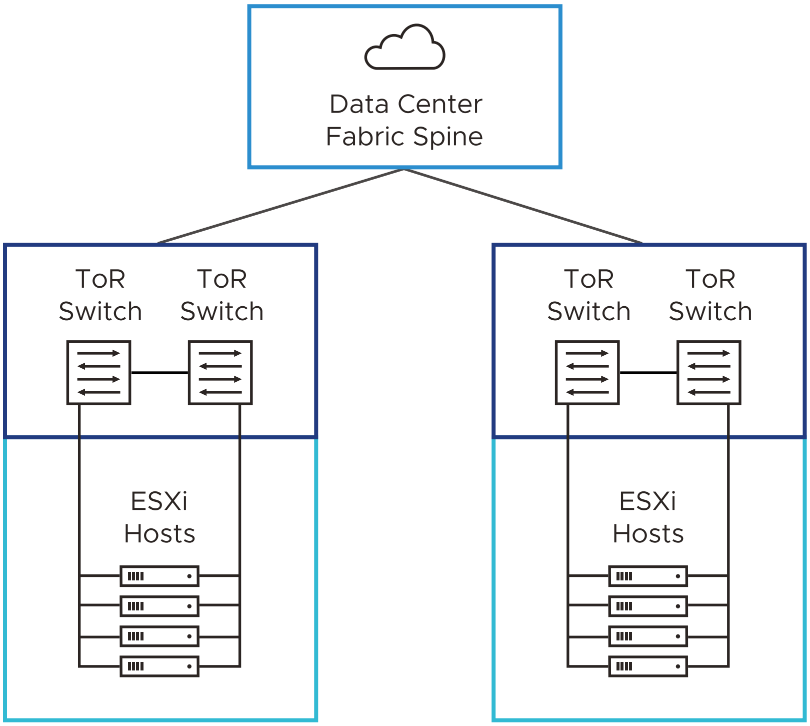

Leaf-Spine Physical Network Design Requirements and Recommendations

The requirements and recommendations for the leaf-spine network configuration determine the physical layout and use of VLANs. They also include requirements and recommendations on jumbo frames, and on network-related requirements such as DNS and NTP.

Table 5-3. Leaf-Spine Physical Network Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-REQD-CFG-001	Do not use EtherChannel (LAG, LACP, or vPC) configuration for ESXi host uplinks.	■ Simplifies configuration of top of rack switches. ■ Teaming options available with vSphere Distributed Switch provide load balancing and failover. ■ EtherChannel implementations might have vendor-specific limitations.	None.
VCF-NET-REQD-CFG-002	Use VLANs to separate physical network functions.	■ Supports physical network connectivity without requiring many NICs. ■ Isolates the different network functions in the SDDC so that you can have differentiated services and prioritized traffic as needed.	Requires uniform configuration and presentation on all the trunks that are made available to the ESXi hosts.
VCF-NET-REQD-CFG-003	Configure the VLANs as members of a 802.1Q trunk.	All VLANs become available on the same physical network adapters on the ESXi hosts.	Optionally, the management VLAN can act as the native VLAN.
VCF-NET-REQD-CFG-004	Set the MTU size to at least 1,700 bytes (recommended 9,000 bytes for jumbo frames) on the physical switch ports, vSphere Distributed Switches, vSphere Distributed Switch port groups, and N-VDS switches that support the following traffic types: ■ Overlay (Geneve) ■ vSAN ■ vSphere vMotion	■ Improves traffic throughput. ■ Supports Geneve by increasing the MTU size to a minimum of 1,600 bytes. ■ Geneve is an extensible protocol. The MTU size might increase with future capabilities. While 1,600 bytes is sufficient, an MTU size of 1,700 bytes provides more room for increasing the Geneve MTU size without the need to change the MTU size of the physical infrastructure.	When adjusting the MTU packet size, you must also configure the entire network path (VMkernel network adapters, virtual switches, physical switches, and routers) to support the same MTU packet size. In an environment with multiple availability zones, the MTU must be configured on the entire network path between the zones.

Table 5-4. Leaf-Spine Physical Network Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-REQD-CFG-005	Set the MTU size to at least 1,500 bytes (1,700 bytes preferred; 9,000 bytes recommended for jumbo frames) on the components of the physical network between the VMware Cloud Foundation instances for the following traffic types. ■ NSX Edge RTEP	■ Jumbo frames are not required between VMware Cloud Foundation instances. However, increased MTU improves traffic throughput. ■ Increasing the RTEP MTU to 1,700 bytes minimizes fragmentation for standard-size workload packets between VMware Cloud Foundation instances.	When adjusting the MTU packet size, you must also configure the entire network path, that is, virtual interfaces, virtual switches, physical switches, and routers to support the same MTU packet size.
VCF-NET-REQD-CFG-006	Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 500 ms.	A latency lower than 500 ms is required for NSX Federation.	None.
VCF-NET-REQD-CFG-007	Provide a routed connection between the NSX Manager clusters in VMware Cloud Foundation instances that are connected in an NSX Federation.	Configuring NSX Federation requires connectivity between the NSX Global Manager instances, NSX Local Manager instances, and NSX Edge clusters.	You must assign unique routable IP addresses for each fault domain.

Table 5-5. Leaf-Spine Physical Network Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-001	Use two ToR switches for each rack.	Supports the use of two 10-GbE (25-GbE or greater recommended) links to each server, provides redundancy and reduces the overall design complexity.	Requires two ToR switches per rack which might increase costs.
VCF-NET-RCMD-CFG-002	Implement the following physical network architecture: ■ One 25-GbE (10-GbE minimum) port on each ToR switch for ESXi host uplinks (Host to ToR). ■ Layer 3 device that supports BGP.	■ Provides availability during a switch failure. ■ Provides support for BGP dynamic routing protocol	■ Might limit the hardware choices. ■ Requires dynamic routing protocol configuration in the physical network.
VCF-NET-RCMD-CFG-003	Use a physical network that is configured for BGP routing adjacency.	■ Supports design flexibility for routing multi-site and multi-tenancy workloads. ■ BGP is the only dynamic routing protocol that is supported for NSX Federation. ■ Supports failover between ECMP Edge uplinks.	Requires BGP configuration in the physical network.
VCF-NET-RCMD-CFG-004	Assign persistent IP configurations for NSX tunnel endpoints (TEPs) that use static IP pools instead of dynamic IP pool addressing.	■ Ensures that endpoints have a persistent TEP IP address. ■ In VMware Cloud Foundation, TEP IP assignment by using static IP pools is recommended for all topologies. ■ This configuration removes any requirement for external DHCP services.	If you add more hosts to the cluster, expanding the static IP pools might be required.

Table 5-5. Leaf-Spine Physical Network Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-005	Configure the trunk ports connected to ESXi NICs as trunk PortFast.	Reduces the time to transition ports over to the forwarding state.	Although this design does not use the STP, switches usually have STP configured by default.
VCF-NET-RCMD-CFG-006	Configure VRRP, HSRP, or another Layer 3 gateway availability method for these networks. ■ Management ■ Edge overlay	Ensures that the VLANs that are stretched between availability zones are connected to a highly- available gateway. Otherwise, a failure in the Layer 3 gateway will cause disruption in the traffic in the SDN setup.	Requires configuration of a high availability technology for the Layer 3 gateways in the data center.

Table 5-6. Leaf-Spine Physical Network Design Recommendations for NSX Federation in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-007	Provide BGP routing between all VMware Cloud Foundation instances that are connected in an NSX Federation setup.	BGP is the supported routing protocol for NSX Federation.	None.
VCF-NET-RCMD-CFG-008	Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 150 ms for workload mobility.	A latency lower than 150 ms is required for the following features: ■ Cross vCenter Server vMotion	None.

vSAN Design for VMware Cloud Foundation

6

VMware Cloud Foundation uses VMware vSAN as the principal storage type for the management domain and is recommended for use as principal storage in VI workload domains. You must determine the size of the compute and storage resources for the vSAN storage, and the configuration of the network carrying vSAN traffic. For multiple availability zones, you extend the resource size and determine the configuration of the vSAN witness host.

Read the following topics next:

- Logical Design for vSAN for VMware Cloud Foundation
- Hardware Configuration for vSAN for VMware Cloud Foundation
- Network Design for vSAN for VMware Cloud Foundation
- vSAN Witness Design for VMware Cloud Foundation
- vSAN Design Requirements and Recommendations for VMware Cloud Foundation

Logical Design for vSAN for VMware Cloud Foundation

vSAN is a cost-efficient storage technology that provides a simple storage management user experience, and permits a fully automated initial deployment of VMware Cloud Foundation. It also provides support for future storage expansion and implementation of vSAN stretched clusters in a workload domain.

Figure 6-1. vSAN Logical Design for VMware Cloud Foundation

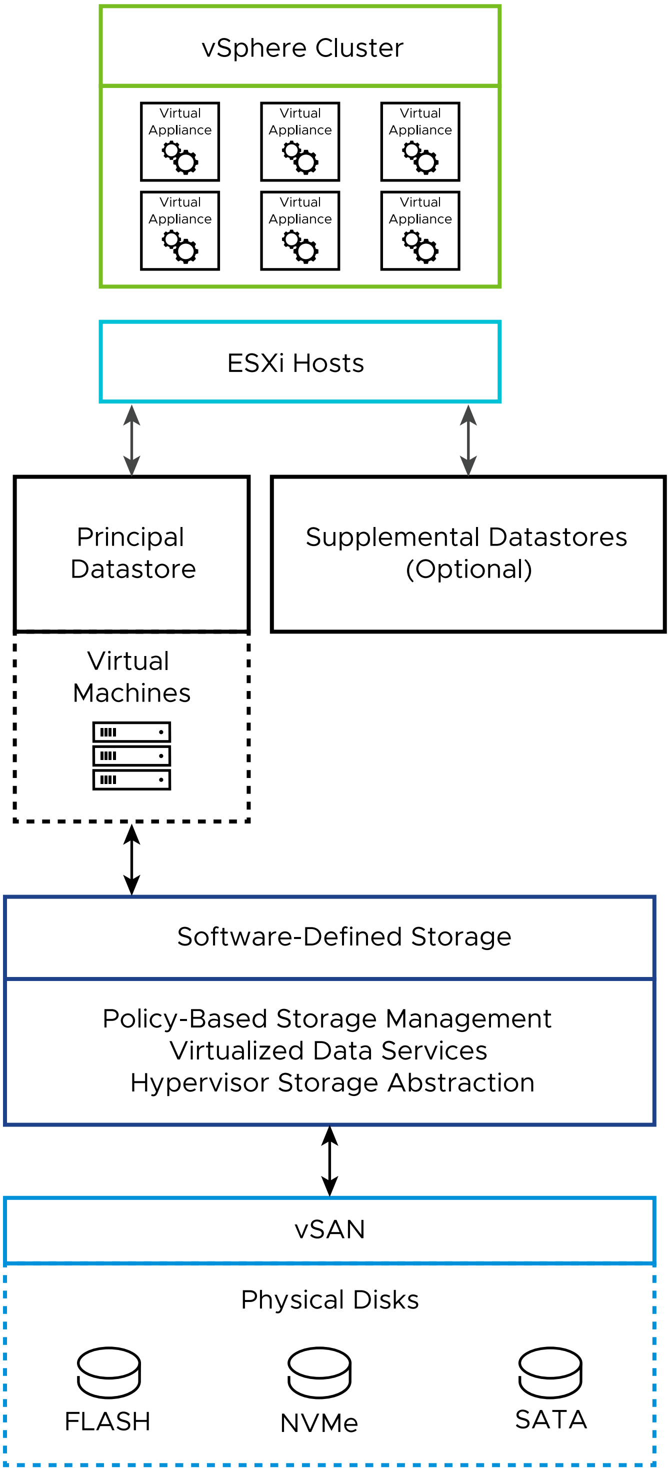

Table 6-1. vSAN Logical Design

Workload Domain Type	VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud FoundationInstances with Multiple Availability Zones
Management domain (default cluster)	Four nodes minimum	■ Must be stretched first ■ 8 node minimum, equally distributed across availability zones ■ vSAN witness appliance in a third fault domain
Management domain (additional clusters)	■ Three nodes minimum ■ Four nodes minimum is recommended for higher availability	■ Six nodes minimum, equally distributed across availability zones ■ Eight nodes minimum is recommended for higher availability ■ vSAN witness appliance in a third fault domain
VI workload domain (all clusters)	■ Three nodes minimum ■ Four nodes minimum is recommended for higher availability	■ Six nodes minimum, equally distributed across availability zones ■ Eight nodes minimum is recommended for higher availability ■ vSAN witness appliance in a third fault domain

Hardware Configuration for vSAN for VMware Cloud Foundation

Determine the vSAN architecture and the storage controllers for performance and stability according to the requirements of the management components of VMware Cloud Foundation.

Figure 6-2. Choosing a vSAN Architecture Model

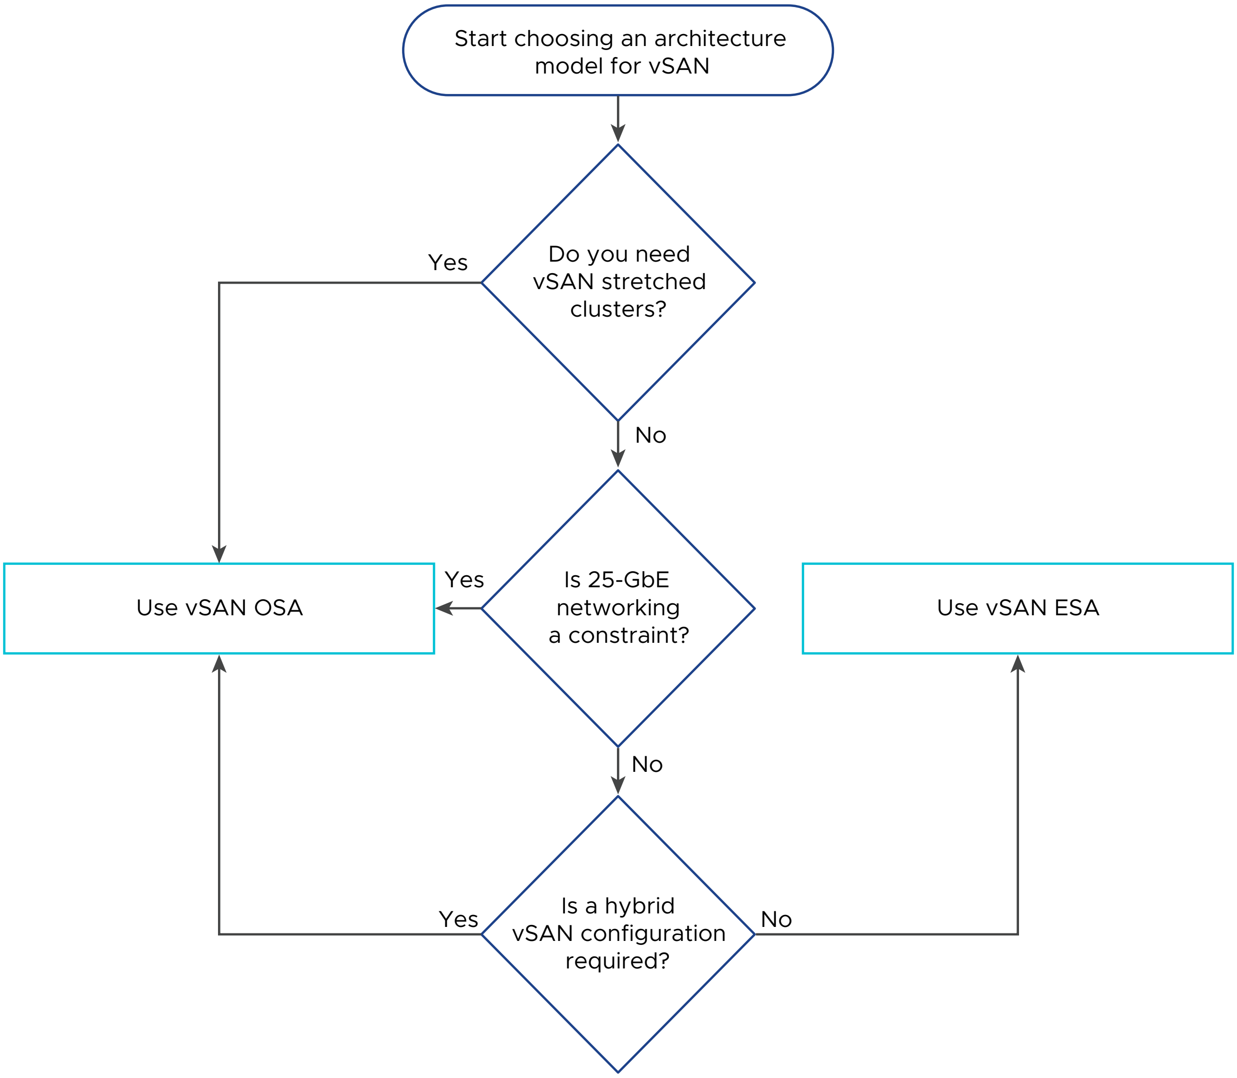
Start choosing an architecture model for vSAN
Do you need vSAN stretched clusters?
Yes
No
Is 25-GbE networking a constraint?
Yes
No
Use vSAN OSA
Use vSAN ESA
Is a hybrid vSAN configuration required?
Yes
No

vSAN Physical Requirements and Dependencies

vSAN has the following requirements and options:

- vSAN Original Storage Architecture (OSA) as hybrid storage or all-flash storage.
  - A vSAN hybrid storage configuration requires both magnetic devices and flash caching devices. The cache tier must be at least 10% of the size of the capacity tier.
  - An all-flash vSAN configuration requires flash devices for both the caching and capacity tiers.
  - VMware vSAN ReadyNodes or hardware from the VMware Compatibility Guide to build your own.

- vSAN Express Storage Architecture (ESA)
  - All storage devices claimed by vSAN contribute to capacity and performance. Each host's storage devices claimed by vSAN form a storage pool. The storage pool represents the amount of caching and capacity provided by the host to the vSAN datastore.
  - ESXi hosts must be on the vSAN ESA Ready Node HCL with a minimum of 512 GB RAM per host.

Note vSAN ESA stretched clusters are not supported by VMware Cloud Foundation.

For best practices, capacity considerations, and general recommendations about designing and sizing a vSAN cluster, see the VMware vSAN Design and Sizing Guide.

Network Design for vSAN for VMware Cloud Foundation

In the network design for vSAN in VMware Cloud Foundation, you determine the network configuration for vSAN traffic.

Consider the overall traffic bandwidth and decide how to isolate storage traffic.

- Consider how much vSAN data traffic is running between ESXi hosts.
- The amount of storage traffic depends on the number of VMs that are running in the cluster, and on how write-intensive the I/O process is for the applications running in the VMs.

For information on the physical network setup for vSAN traffic, and other system traffic, see Chapter 5 Physical Network Infrastructure Design for VMware Cloud Foundation.

For information on the virtual network setup for vSAN traffic, and other system traffic, see Logical vSphere Networking Design for VMware Cloud Foundation.

The vSAN network design includes these components.

Table 6-2. Components of vSAN Network Design

Design Component	Description
Physical NIC speed	For best and predictable performance (IOPS) of the environment, this design uses a minimum of a 10-GbE connection, with 25-GbE recommended, for use with vSAN OSA all-flash configurations. For vSAN ESA, 25-GbE connection is recommended.
VMkernel network adapters for vSAN	The vSAN VMkernel network adapter on each ESXi host is created when you enable vSAN on the cluster. Connect the vSAN VMkernel network adapters on all ESXi hosts in a cluster to a dedicated distributed port group, including ESXi hosts that are not contributing storage resources to the cluster.

Table 6-2. Components of vSAN Network Design (continued)

Design Component	Description
VLAN	All storage traffic should be isolated on its own VLAN. When a design uses multiple vSAN clusters, each cluster should use a dedicated VLAN or segment for its traffic. This approach increases security, prevents interference between clusters, and helps with troubleshooting cluster configuration. If a cluster spans a rack, the vSAN VLAN can be allocated per rack to enable Layer 3 multi-rack deployments.
Jumbo frames	vSAN traffic can be handled by using jumbo frames. Use jumbo frames for vSAN traffic only if the physical environment is already configured to support them, they are part of the existing design, or if the underlying configuration does not create a significant amount of added complexity to the design.

What to read next

vSAN Witness Design for VMware Cloud Foundation

The vSAN witness appliance is a specialized ESXi installation that provides quorum and tiebreaker services for stretched clusters in VMware Cloud Foundation.

vSAN Witness Deployment Specification

You must deploy a witness ESXi host when using vSAN in a stretched cluster configuration. This appliance must be deployed in a third location that is not local to the ESXi hosts on either side of the stretched cluster.

Table 6-3. vSAN Witness Appliance Sizing Considerations

Appliance Size	Maximum Number of Supported Virtual Machines	Maximum Number of Supported Witness Components
Tiny	10	750
Medium	500	21, 000
Large	More than 500	45,000
Extra Large	More than 500	64, 000

vSAN Witness Network Design

When using two availability zones, connect the vSAN witness appliance to the workload domain vCenter Server so that you can perform the initial setup of the stretched cluster and have workloads failover between the zones.

VMware Cloud Foundation uses vSAN witness traffic separation where you can use a VMkernel adapter for vSAN witness traffic that is different from the adapter for vSAN data traffic. In this design, you configure vSAN witness traffic in the following way:

- On each ESXi host in both availability zones, place the vSAN witness traffic on the management VMkernel adapter.
- On the vSAN witness appliance, use the same VMkernel adapter for both management and witness traffic.

For information about vSAN witness traffic separation, see vSAN Stretched Cluster Guide on VMware Cloud Platform Tech Zone.

Management network

Routed to the management networks in both availability zones. Connect the first VMkernel adapter of the vSAN witness appliance to this network. The second VMkernel adapter on the vSAN witness appliance is not used.

Place the following traffic on this network:

- Management traffic
- vSAN witness traffic

Figure 6-3. vSAN Witness Network Design

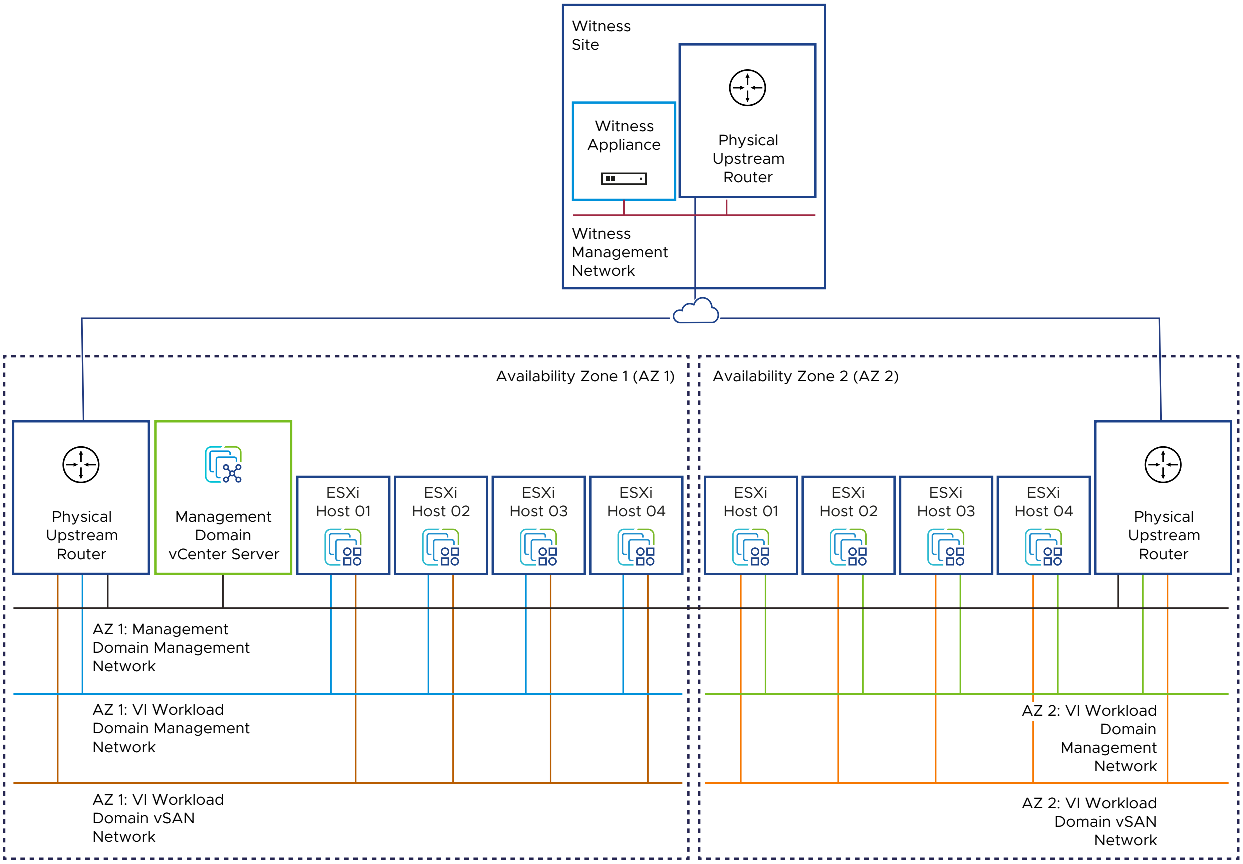

Witness Site
Witness Appliance
Physical Upstream Router
Witness Management Network
Availability Zone 1 (AZ 1)
Availability Zone 2 (AZ 2)
Physical Upstream Router
Management Domain vCenter Server
ESXi Host 01
ESXi Host 02
ESXi Host 03
ESXi Host 04
ESXi Host 01
ESXi Host 02
ESXi Host 03
ESXi Host 04
Physical Upstream Router
AZ 1: Management Domain Management Network
AZ 1: VI Workload Domain Management Network
AZ 1: VI Workload Domain vSAN Network
AZ 2: VI Workload Domain Management Network
AZ 2: VI Workload Domain vSAN Network

vSAN Design Requirements and Recommendations for VMware Cloud Foundation

Consider the requirements for using vSAN storage for standard and stretched clusters in VMware Cloud Foundation, such as required capacity, number of hosts, storage policies, and the similar best practices for having vSAN operate in an optimal way.

For related vSphere cluster requirements and recommendations, see vSphere Cluster Design Requirements and Recommendations for VMware Cloud Foundation.

vSAN Design Requirements

You must meet the following design requirements for standard and stretched clusters in your vSAN design for VMware Cloud Foundation.

Table 6-4. vSAN Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-001	Provide sufficient raw capacity to meet the initial needs of the workload domain cluster.	Ensures that sufficient resources are present to create the workload domain cluster.	None.
VCF-VSAN-REQD-CFG-002	Provide at least the required minimum number of hosts according to the cluster type.	Satisfies the requirements for storage availability.	None.

Table 6-5. vSAN ESA Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-003	Verify the hardware components used in your vSAN deployment are on the vSAN Hardware Compatibility List.	Prevents hardware-related failures during workload deployment	Limits the number of compatible hardware configurations that can be used.

Table 6-6. vSAN Design Requirements for Stretched Clusters with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-004	Add the following setting to the default vSAN storage policy: Site disaster tolerance = Site mirroring - stretched cluster	Provides the necessary protection for virtual machines in each availability zone, with the ability to recover from an availability zone outage.	You might need additional policies if third-party virtual machines are to be hosted in these clusters because their performance or availability requirements might differ from what the default VMware vSAN policy supports.
VCF-VSAN-REQD-CFG-005	Configure two fault domains, one for each availability zone. Assign each host to their respective availability zone fault domain.	Fault domains are mapped to availability zones to provide logical host separation and ensure a copy of vSAN data is always available even when an availability zone goes offline.	You must provide additional raw storage when the site mirroring - stretched cluster option is selected, and fault domains are enabled.
VCF-VSAN-REQD-CFG-006	Use vSAN OSA to create a stretched cluster.	Stretched clusters on top of vSAN ESA are not supported by VMware Cloud Foundation	None.
VCF-VSAN-REQD-CFG-007	Configure an individual vSAN storage policy for each stretched cluster.	The vSAN storage policy of a stretched cluster cannot be shared with other clusters.	You must configure additional vSAN storage policies.
VCF-VSAN-WTN-REQD-CFG-001	Deploy a vSAN witness appliance in a location that is not local to the ESXi hosts in any of the availability zones.	Ensures availability of vSAN witness components in the event of a failure of one of the availability zones.	You must provide a third physically separate location that runs a vSphere environment. You might use a VMware Cloud Foundation instance in a separate physical location.
VCF-VSAN-WTN-REQD-CFG-002	Deploy a witness appliance that corresponds to the required cluster capacity.	Ensures the witness appliance is sized to support the projected workload storage consumption.	The vSphere environment at the witness location must satisfy the resource requirements of the witness appliance.
VCF-VSAN-WTN-REQD-CFG-003	Connect the first VMkernel adapter of the vSAN witness appliance to the management network in the witness site.	Enables connecting the witness appliance to the workload domain vCenter Server.	The management networks in both availability zones must be routed to the management network in the witness site.
VCF-VSAN-WTN-REQD-CFG-004	Allocate a statically assigned IP address and host name to the management adapter of the vSAN witness appliance.	Simplifies maintenance and tracking, and implements a DNS configuration.	Requires precise IP address management.

Table 6-6. vSAN Design Requirements for Stretched Clusters with VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-WTN-REQD-CFG-005	Configure forward and reverse DNS records for the vSAN witness appliance for the VMware Cloud Foundation instance.	Enables connecting the vSAN witness appliance to the workload domain vCenter Server by FQDN instead of IP address.	You must provide DNS records for the vSAN witness appliance.
VCF-VSAN-WTN-REQD-CFG-006	Configure time synchronization by using an internal NTP time for the vSAN witness appliance.	Prevents any failures in the stretched cluster configuration that are caused by time mismatch between the vSAN witness appliance and the ESXi hosts in both availability zones and workload domain vCenter Server.	■ An operational NTP service must be available in the environment. ■ All firewalls between the vSAN witness appliance and the NTP servers must allow NTP traffic on the required network ports.

vSAN Design Recommendations

In your vSAN design for VMware Cloud Foundation, you can apply certain best practices for standard and stretched clusters.

Table 6-7. vSAN Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-001	Provide sufficient raw capacity to meet the planned needs of the workload domain cluster.	Ensures that sufficient resources are present in the workload domain cluster, preventing the need to expand the vSAN datastore in the future.	None.
VCF-VSAN-RCMD-CFG-002	Ensure that at least 30% of free space is always available on the vSAN datastore,.	This reserved capacity is set aside for host maintenance mode data evacuation, component rebuilds, rebalancing operations, and VM snapshots.	Increases the amount of available storage needed.

Table 6-7. vSAN Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-003	Use the default VMware vSAN storage policy.	■ Provides the level of redundancy that is needed in the workload domain cluster. ■ Provides the level of performance that is enough for the individual workloads.	You might need additional policies for third-party virtual machines hosted in these clusters because their performance or availability requirements might differ from what the default VMware vSAN policy supports.
VCF-VSAN-RCMD-CFG-004	Leave the default virtual machine swap file as a sparse object on vSAN.	Sparse virtual swap files consume capacity on vSAN only as they are accessed. As a result, you can reduce the consumption on the vSAN datastore if virtual machines do not experience memory over-commitment, which would require the use of the virtual swap file.	None.
VCF-VSAN-RCMD-CFG-005	Use the existing vSphere Distributed Switch instance for the workload domain cluster.	■ Reduces the complexity of the network design. ■ Reduces the number of physical NICs required.	All traffic types can be shared over common uplinks.

Table 6-7. vSAN Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-006	Configure jumbo frames on the VLAN for vSAN traffic.	■ Simplifies configuration because jumbo frames are also used to improve the performance of vSphere vMotion and NFS storage traffic. ■ Reduces the CPU overhead, resulting in high network usage.	Every device in the network must support jumbo frames.
VCF-VSAN-RCMD-CFG-007	Configure vSAN in an all-flash configuration in the default workload domain cluster.	Meets the performance needs of the default workload domain cluster.	All vSAN disks must be flash disks, which might cost more than magnetic disks.

Table 6-8. vSAN OSA Design Recommendations for with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-008	Ensure that the storage I/O controller has a minimum queue depth of 256 set.	Storage controllers with lower queue depths can cause performance and stability problems when running vSAN. vSAN ReadyNode servers are configured with the correct queue depths for vSAN.	Limits the number of compatible I/O controllers that can be used for storage.
VCF-VSAN-RCMD-CFG-009	Do not use the storage I/O controllers that are running vSAN disk groups for another purpose.	Running non-vSAN disks, for example, VMFS, on a storage I/O controller that is running a vSAN disk group can impact vSAN performance.	If non-vSAN disks are required in ESXi hosts, you must have an additional storage I/O controller in the host.

Table 6-8. vSAN OSA Design Recommendations for with VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-010	Configure vSAN with a minimum of two disk groups per ESXi host.	Reduces the size of the fault domain and spreads the I/O load over more disks for better performance.	Using multiple disk groups requires more disks in each ESXi host.
VCF-VSAN-RCMD-CFG-011	For the cache tier in each disk group, use a flash-based drive that is at least 600 GB large.	Provides enough cache for both hybrid or all-flash vSAN configurations to buffer I/O and ensure disk group performance. Additional space in the cache tier does not increase performance.	Using larger flash disks can increase the initial host cost.

Table 6-9. vSAN ESA Design Recommendations for with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-012	Activate auto-policy management.	Configures optimized storage policies based on the cluster type and the number of hosts in the cluster inventory. Changes to the number of hosts in the cluster or Host Rebuild Reserve will prompt you to make a suggested adjustment to the optimized storage policy.	You must activate auto-policy management manually.
VCF-VSAN-RCMD-CFG-013	Activate vSAN ESA compression.	Activated by default, it also improves performance.	PostgreSQL databases and other applications might use their own compression capabilities. In these cases, using a storage policy with the compression capability turned off will save CPU cycles. You can disable vSAN ESA compressions for such workloads through the use of the Storage Policy Based Management (SPBM) framework.
VCF-VSAN-RCMD-CFG-014	Use NICs with a minimum 25-GbE capacity.	10-GbE NICs will limit the scale and performance of a vSAN ESA cluster because usually performance requirements increase over the lifespan of the cluster.	Requires 25-GbE or faster network fabric.

Table 6-10. vSAN Design Recommendations for Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-WTN-RCMD-CFG-001	Configure the vSAN witness appliance to use the first VMkernel adapter, that is the management interface, for vSAN witness traffic.	Removes the requirement to have static routes on the witness appliance as witness traffic is routed over the management network.	The management networks in both availability zones must be routed to the management network in the witness site.
VCF-VSAN-WTN-RCMD-CFG-002	Place witness traffic on the management VMkernel adapter of all the ESXi hosts in the workload domain.	Separates the witness traffic from the vSAN data traffic. Witness traffic separation provides the following benefits: ■ Removes the requirement to have static routes from the vSAN networks in both availability zones to the witness site. ■ Removes the requirement to have jumbo frames enabled on the path between each availability zone and the witness site because witness traffic can use a regular MTU size of 1500 bytes.	The management networks in both availability zones must be routed to the management network in the witness site.

vSphere Design for VMware Cloud Foundation

7

The vSphere design includes determining the configuration of the vCenter Server instances, ESXi hosts, vSphere clusters,and vSphere networking for a VMware Cloud Foundation environment.

Read the following topics next:

- ESXi Design for VMware Cloud Foundation
- vCenter Server Design for VMware Cloud Foundation
- vSphere Cluster Design for VMware Cloud Foundation
- vSphere Networking Design for VMware Cloud Foundation

ESXi Design for VMware Cloud Foundation

In the design of the ESXi host configuration for your VMware Cloud Foundation environment, consider the resources, networking, and security policies that are required to support the virtual machines in each workload domain cluster.

- Logical Design for ESXi for VMware Cloud Foundation

  In the logical design for ESXi, you determine the high-level integration of the ESXi hosts with the other components of the VMware Cloud Foundation instance for providing virtual infrastructure to management and workload components.

- Sizing Considerations for ESXi for VMware Cloud Foundation

  You decide on the number of ESXi hosts per cluster and the number of physical disks per ESXi host.

- ESXi Design Requirements and Recommendations for VMware Cloud Foundation

  The requirements for the ESXi hosts in a workload domain in VMware Cloud Foundation are related to the system requirements of the workloads hosted in the domain. The ESXi requirements include number, server configuration, amount of hardware resources, networking, and certificate management. Similar best practices help you design optimal environment operation.

Logical Design for ESXi for VMware Cloud Foundation

In the logical design for ESXi, you determine the high-level integration of the ESXi hosts with the other components of the VMware Cloud Foundation instance for providing virtual infrastructure to management and workload components.

To provide the resources required to run the management and workload components of the VMware Cloud Foundation instance, each ESXi host consists of the following elements:

- CPU and memory
- Storage devices
- Out of band management interface
- Network interfaces

Figure 7-1. ESXi Logical Design for VMware Cloud Foundation

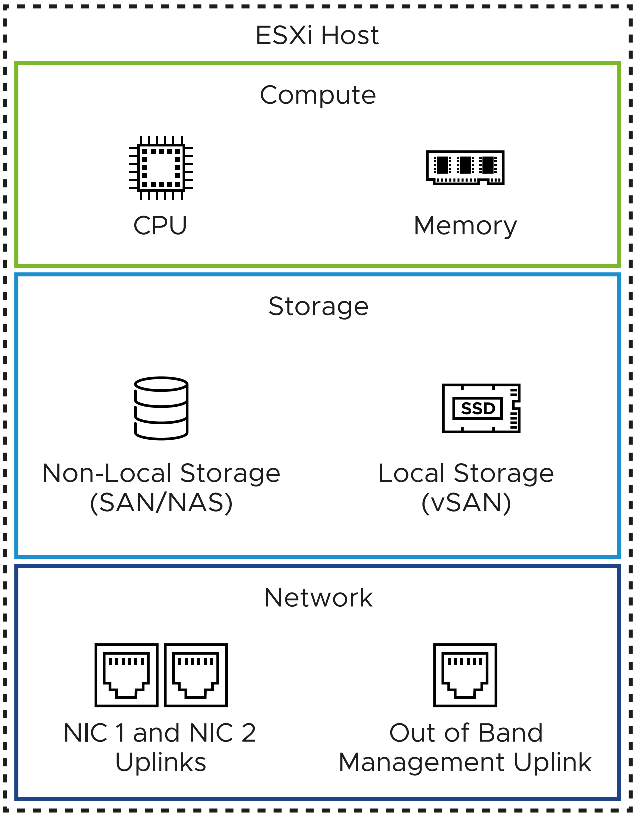

Sizing Considerations for ESXi for VMware Cloud Foundation

You decide on the number of ESXi hosts per cluster and the number of physical disks per ESXi host.

For detailed sizing based on the overall profile of the VMware Cloud Foundation instance you plan to deploy, see VMware Cloud Foundation Planning and Preparation Workbook.

The configuration and assembly process for each system should be standardized, with all components installed in the same manner on each ESXi host. Because standardization of the physical configuration of the ESXi hosts removes variability, the infrastructure is easily managed and supported. ESXi hosts are deployed with identical configuration across all cluster members, including storage and networking configurations. For example, consistent PCIe card slot placement, especially for network interface controllers, is essential for accurate mapping of physical network interface controllers to virtual network resources. By using identical configurations, you have an even balance of virtual machine storage components across storage and compute resources.

Table 7-1. ESXi Server Sizing Considerations by Hardware Element

Hardware Element	Considerations
CPU	■ Total CPU requirements for the workloads that are running in the cluster. ■ Host failure and maintenance scenarios. Keep the overcommitment ratio vCPU-to-pCPU less than or equal to 2:1 for the management domain and less than or equal to 8:1 for VI workload domains . ■ Additional third-party management components. ■ Number of physical cores, not logical cores. Simultaneous multithreading (SMT) technologies in CPUs, such as hyper-threading in Intel CPUs, improve CPU performance by allowing multiple threads to run in parallel on the same CPU core. Although a single CPU core can be viewed as two logical cores, the performance enhancement will not be equivalent to 100% more CPU power. It will also differ from one environment to another.
Memory	■ Total memory requirements for the workloads that are running in the cluster. ■ When sizing memory for the ESXi hosts in a cluster, to reserve the resources of one host for failover or maintenance, set the admission control setting to N+1, which reserves the resources of one host for failover or maintenance. ■ vSAN OSA. Number of vSAN disk groups and disks on an ESXi host. To support the maximum number of disk groups, you must provide 32 GB of RAM. For more information about disk groups, including design and sizing guidance, see Administering VMware vSAN in the vSphere documentation. ■ vSAN ESA. You must provide 512 GB of RAM.
Storage	■ Use high-endurance device such as a hard drive or SSD for boot device ■ Use 128-GB boot device to maximize the space available for ESX-OS Data ■ vSAN OSA ■ Provide at least one 600-GB cache disk. ■ Use a minimum of two capacity disks. ■ vSAN ESA. Use a minimum of two NVME devices. ■ Use hosts with homogeneous configuration.

ESXi Design Requirements and Recommendations for VMware Cloud Foundation

The requirements for the ESXi hosts in a workload domain in VMware Cloud Foundation are related to the system requirements of the workloads hosted in the domain. The ESXi requirements include number, server configuration, amount of hardware resources, networking, and certificate management. Similar best practices help you design optimal environment operation.

ESXi Server Design Requirements

You must meet the following design requirements for the ESXi hosts in a workload domain in a VMware Cloud Foundation deployment.

Table 7-2. Design Requirements for ESXi Server Hardware

Requirement ID	Design Requirement	Requirement Justification	Requirement Implication
VCF-ESX-REQD-CFG-001	Install no less than the minimum number of ESXi hosts required for the cluster type being deployed.	■ Ensures availability requirements are met. ■ If one of the hosts is not available because of a failure or maintenance event, the CPU overcommitment ratio becomes 2:1.	None.
VCF-ESX-REQD-CFG-002	Ensure each ESXi host matches the required CPU, memory and storage specification.	■ Ensures workloads will run without contention even during failure and maintenance conditions.	Assemble the server specification and number according to the sizing in VMware Cloud Foundation Planning and Preparation Workbook which is based on projected deployment size.
VCF-ESX-REQD-SEC-001	Regenerate the certificate of each ESXi host after assigning the host an FQDN.	Establishes a secure connection with VMware Cloud Builder during the deployment of a workload domain and prevents man-in-the-middle (MiTM) attacks.	You must manually regenerate the certificates of the ESXi hosts before the deployment of a workload domain.

ESXi Server Design Recommendations

In your ESXi host design for VMware Cloud Foundation, you can apply certain best practices.

Table 7-3. Design Recommendations for ESXi Server Hardware

Recommendation ID	Recommendation	Justification	Implication
VCF-ESX-RCMD-CFG-001	Use vSAN ReadyNodes with vSAN storage for each ESXi host in the management domain.	Your management domain is fully compatible with vSAN at deployment. For information about the models of physical servers that are vSAN-ready, see vSAN Compatibility Guide for vSAN ReadyNodes.	Hardware choices might be limited. If you plan to use a server configuration that is not a vSAN ReadyNode, your CPU, disks and I/O modules must be listed on the VMware Compatibility Guide under *CPU Series* and *vSAN Compatibility List* aligned to the ESXi version specified in VMware Cloud Foundation 5.1 Release Notes.
VCF-ESX-RCMD-CFG-002	Allocate hosts with uniform configuration across the default management vSphere cluster.	A balanced cluster has these advantages: ■ Predictable performance even during hardware failures ■ Minimal impact of resynchronization or rebuild operations on performance	You must apply vendor sourcing, budgeting, and procurement considerations for uniform server nodes on a per cluster basis.
VCF-ESX-RCMD-CFG-003	When sizing CPU, do not consider multithreading technology and associated performance gains.	Although multithreading technologies increase CPU performance, the performance gain depends on running workloads and differs from one case to another.	Because you must provide more physical CPU cores, costs increase and hardware choices become limited.
VCF-ESX-RCMD-CFG-004	Install and configure all ESXi hosts in the default management cluster to boot using a 128-GB device or larger.	Provides hosts that have large memory, that is, greater than 512 GB, with enough space for the scratch partition when using vSAN.	None.
VCF-ESX-RCMD-CFG-005	Use the default configuration for the scratch partition on all ESXi hosts in the default management cluster.	■ If a failure in the vSAN cluster occurs, the ESXi hosts remain responsive and log information is still accessible. ■ It is not possible to use vSAN datastore for the scratch partition.	None.

Table 7-3. Design Recommendations for ESXi Server Hardware (continued)

Recommendation ID	Recommendation	Justification	Implication
VCF-ESX-RCMD-CFG-006	For workloads running in the default management cluster, save the virtual machine swap file at the default location.	Simplifies the configuration process.	Increases the amount of replication traffic for management workloads that are recovered as part of the disaster recovery process.
VCF-ESX-RCMD-NET-001	Place the ESXi hosts in each management domain cluster on a host management network that is separate from the VM management network.	Enables the separation of the physical VLAN between ESXi hosts and the other management components for security reasons.	Increases the number of VLANs required.
VCF-ESX-RCMD-NET-002	Place the ESXi hosts in each VI workload domain on a separate host management VLAN-backed network.	Enables the separation of the physical VLAN between the ESXi hosts in different VI workload domains for security reasons.	Increases the number of VLANs required. For each VI workload domain, you must allocate a separate management subnet.
VCF-ESX-RCMD-SEC-001	Deactivate SSH access on all ESXi hosts in the management domain by having the SSH service stopped and using the default SSH service policy `Start and stop manually`.	Ensures compliance with the *vSphere Security Configuration Guide* and with security best practices. Disabling SSH access reduces the risk of security attacks on the ESXi hosts through the SSH interface.	You must activate SSH access manually for troubleshooting or support activities as VMware Cloud Foundation deactivates SSH on ESXi hosts after workload domain deployment.
VCF-ESX-RCMD-SEC-002	Set the advanced setting `UserVars.SuppressShellWarning` to `0` across all ESXi hosts in the management domain.	■ Ensures compliance with the *vSphere Security Configuration Guide* and with security best practices ■ Enables the warning message that appears in the vSphere Client every time SSH access is activated on an ESXi host.	You must turn off SSH enablement warning messages manually when performing troubleshooting or support activities.

vCenter Server Design for VMware Cloud Foundation

vCenter Server design considers the location, size, high availability, and identity domain isolation of the vCenter Server instances for the workload domains in a VMware Cloud Foundation environment.

- Logical Design for vCenter Server for VMware Cloud Foundation

  Each workload domain has a dedicated vCenter Server that manages the ESXi hosts running NSX Edge nodes and customer workloads. All vCenter Server instances run in the management domain.

- Sizing Considerations for vCenter Server for VMware Cloud Foundation

  You select an appropriate vCenter Server appliance size according to the scale of your environment.

- High Availability Design for vCenter Server for VMware Cloud Foundation

  Protecting vCenter Server is important because it is the central point of management and monitoring for each workload domain.

- vCenter Server Design Requirements and Recommendations for VMware Cloud Foundation

  Each workload domain in VMware Cloud Foundation is managed by a single vCenter Server instance. You determine the size of this vCenter Server instance and its storage requirements according to the number of ESXi hosts per cluster and the number of virtual machines you plan to run on these clusters.

- vCenter Single Sign-On Design Requirements for VMware Cloud Foundation

  vCenter Server instances for the VI workload domains in a VMware Cloud Foundation deployment can be either joined to the vCenter Single Sign-On domain of the vCenter Server instance for the management domain or deployed in isolated vCenter Single Sign-On domains.

Logical Design for vCenter Server for VMware Cloud Foundation

Each workload domain has a dedicated vCenter Server that manages the ESXi hosts running NSX Edge nodes and customer workloads. All vCenter Server instances run in the management domain.

Figure 7-2. Design of vCenter Server for VMware Cloud Foundation

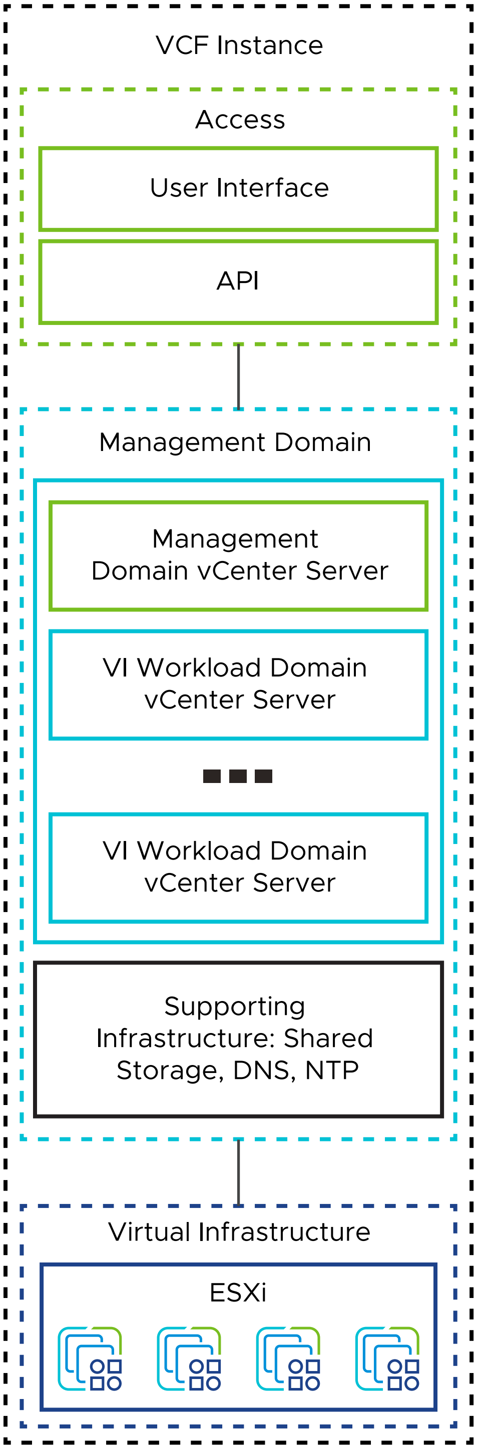
VCF Instance
Access
User Interface
API
Management Domain
Management Domain vCenter Server
VI Workload Domain vCenter Server
VI Workload Domain vCenter Server
Supporting Infrastructure: Shared Storage, DNS, NTP
Virtual Infrastructure
ESXi

Table 7-4. vCenter Server Layout

VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
■ One vCenter Server instance for the management domain that manages the management components of the SDDC, such as the vCenter Server instances for the VI workload domains, NSX Manager cluster nodes, SDDC Manager, and other solutions. ■ Optionally, additional vCenter Server instances for the VI workload domains to support customer workloads. ■ vSphere HA protecting all vCenter Server appliances.	■ One vCenter Server instance for the management domain that manages the management components of the SDDC, such as vCenter Server instances for the VI workload domains, NSX Manager cluster nodes, SDDC Manager, and other solutions. ■ Optionally, additional vCenter Server instances for the VI workload domains to support customer workloads. ■ vSphere HA protecting all vCenter Server appliances. ■ A should-run-on-host-in-group VM-Host affinity rule in vSphere DRS specifying that the vCenter Server appliances should run in the primary availability zone unless an outage in this zone occurs.

Sizing Considerations for vCenter Server for VMware Cloud Foundation

You select an appropriate vCenter Server appliance size according to the scale of your environment.

When you deploy a workload domain, you select a vCenter Server appliance size that is suitable for the scale of your environment. The option that you select determines the number of CPUs and the amount of memory of the appliance. For detailed sizing according to a collective profile of the VMware Cloud Foundation instance you plan to deploy, refer to the VMware Cloud Foundation Planning and Preparation Workbook .

Table 7-5. Sizing Considerations for vCenter Server

vCenter Server Appliance Size	Management Capacity
Tiny	Up to 10 hosts or 100 virtual machines
Small *	Up to 100 hosts or 1,000 virtual machines
Medium **	Up to 400 hosts or 4,000 virtual machines
Large	Up to 1,000 hosts or 10,000 virtual machines
X-Large	Up to 2,000 hosts or 35,000 virtual machines

* Default for the management domain vCenter Server

** Default for VI workload domain vCenter Server instances

High Availability Design for vCenter Server for VMware Cloud Foundation

Protecting vCenter Server is important because it is the central point of management and monitoring for each workload domain.

VMware Cloud Foundation supports only vSphere HA as a high availability method for vCenter Server.

Table 7-6. Methods for Protecting the vCenter Server Appliance

High Availability Method	Supported in VMware Cloud Foundation	Considerations
vSphere High Availability	Yes	-
vCenter High Availability (vCenter HA)	No	■ vCenter Server services must fail over to the passive node so there is no continuous availability. ■ Recovery time can be up to 30 mins. ■ You must meet additional networking requirements for the private network. ■ vCenter HA requires additional resources for the Passive and Witness nodes. ■ Life cycle management is complicated because you must manually delete and recreate the standby virtual machines during a life cycle management operation.
vSphere Fault Tolerance (vSphere FT)	No	■ The vCPU limit of vSphere FT vCPU would limit vCenter Server appliance size to medium. ■ You must provide a dedicated network.

vCenter Server Design Requirements and Recommendations for VMware Cloud Foundation

Each workload domain in VMware Cloud Foundation is managed by a single vCenter Server instance. You determine the size of this vCenter Server instance and its storage requirements according to the number of ESXi hosts per cluster and the number of virtual machines you plan to run on these clusters.

vCenter Server Design Requirements for VMware Cloud Foundation

You allocate vCenter Server appliances according to the requirements for workload isolation, scalability, and resilience to failures.

Table 7-7. vCenter Server Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-CFG-001	Deploy a dedicated vCenter Server appliance for the management domain of the VMware Cloud Foundation instance.	■ Isolates vCenter Server failures to management or customer workloads. ■ Isolates vCenter Server operations between management and customers. ■ Supports a scalable cluster design where you can reuse the management components as more customer workloads are added to the SDDC. ■ Simplifies capacity planning for customer workloads because you do not consider management workloads for the VI workload domain vCenter Server. ■ Improves the ability to upgrade the vSphere environment and related components by enabling for explicit separation of maintenance windows: ■ Management workloads remain available while you are upgrading the tenant workloads ■ Customer workloads remain available while you are upgrading the management nodes ■ Supports clear separation of roles and responsibilities to ensure that only administrators with granted authorization can control the management workloads. ■ Facilitates quicker troubleshooting and problem resolution. ■ Simplifies disaster recovery operations by supporting a clear separation between recovery of the management components and tenant workloads. ■ Provides isolation of potential network issues by introducing network separation of the clusters in the SDDC.	Requires a separate license for the vCenter Server instance in the management domain
VCF-VCS-REQD-NET-001	Place all workload domain vCenters Server	■ Simplifies IP addressing for management VMs by using the	None.

Table 7-7. vCenter Server Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
	appliances on the VM management network in the management domain.	same VLAN and subnet. - Provides simplified secure access to management VMs in the same VLAN network.	

vCenter Server Design Recommendations

In your vCenter Server design for VMware Cloud Foundation, you can apply certain best practices for sizing and high availability.

Table 7-8. vCenter Server Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VCS-RCMD-CFG-001	Deploy an appropriately sized vCenter Server appliance for each workload domain.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is Small and for VI workload domains is Medium. To override these values, you must use the Cloud Builder API and the SDDC Manager API.
VCF-VCS-RCMD-CFG-002	Deploy a vCenter Server appliance with the appropriate storage size.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is **Small** and for VI Workload Domains is **Medium**. To override these values, you must use the API.
VCF-VCS-RCMD-CFG-003	Protect workload domain vCenter Server appliances by using vSphere HA.	vSphere HA is the only supported method to protect vCenter Server availability in VMware Cloud Foundation.	vCenter Server becomes unavailable during a vSphere HA failover.
VCF-VCS-RCMD-CFG-004	In vSphere HA, set the restart priority policy for the vCenter Server appliance to high.	vCenter Server is the management and control plane for physical and virtual infrastructure. In a vSphere HA event, to ensure the rest of the SDDC management stack comes up faultlessly, the workload domain vCenter Server must be available first, before the other management components come online.	If the restart priority for another virtual machine is set to highest, the connectivity delay for the management components will be longer.

vCenter Server Design Recommendations for Stretched Clusters with VMware Cloud Foundation

The following additional design recommendations apply when using stretched clusters.

Table 7-9. vCenter Server Design Recommendations for vSAN Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VCS-RCMD-CFG-005	Add the vCenter Server appliance to the virtual machine group for the first availability zone.	Ensures that, by default, the vCenter Server appliance is powered on a host in the first availability zone.	None.

vCenter Single Sign-On Design Requirements for VMware Cloud Foundation

vCenter Server instances for the VI workload domains in a VMware Cloud Foundation deployment can be either joined to the vCenter Single Sign-On domain of the vCenter Server instance for the management domain or deployed in isolated vCenter Single Sign-On domains.

You select the vCenter Single Sign-On topology according to the needs and design objectives of your deployment.

Table 7-10. vCenter Single Sign-On Topologies for VMware Cloud Foundation

VMware Cloud Foundation Topology	vCenter Single Sign-On Domain Topology	Benefits	Drawbacks
Single vCenter Server Instance - Single vCenter Single Sign-On Domain	One vCenter Single Sign-On domain with the management domain vCenter Server instance only.	Enables a small environment where customer workloads run in the same cluster as the management domain components.	-
Multiple vCenter Server Instances - Single vCenter Single Sign-On Domain	One vCenter Single Sign-On domain with the management domain and all VI workload domain vCenter Server instances in enhanced linked mode (ELM) using a ring topology.	Enables sharing of vCenter Server roles, tags and licenses between all workload domain instances.	Limited to 15 workload domains per VMware Cloud Foundation instance including the management domain.
Multiple vCenter Server Instances - Multiple vCenter Single Sign-On Domains	■ One vCenter Single Sign-On domain with at least the management domain vCenter Server instance ■ Additional VI workload domains, each with their own isolated vCenter Single Sign-On domain.	■ Enables isolation at the vCenter Single Sign-On domain layer for increased security separation. ■ Supports up to 25 workload domains per VMware Cloud Foundation instance.	Additional password management overhead per vCenter Single Sign-On domain.

Figure 7-3. Single vCenter Server Instance - Single vCenter Single Sign-On Domain

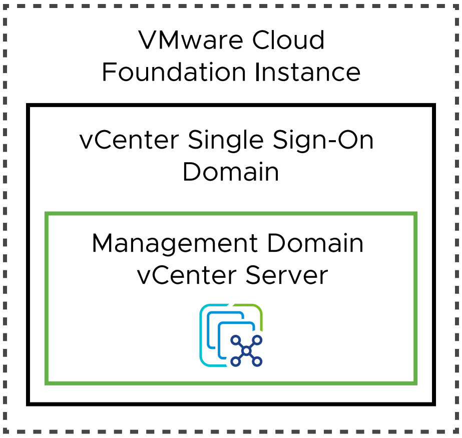
VMware Cloud Foundation Instance
vCenter Single Sign-On Domain
Management Domain vCenter Server

Because the Single vCenter Server Instance - Single vCenter Single Sign-On Domain topology contains a single vCenter Server instance by definition, no relevant design requirements or recommendations for vCenter Single Sign-On are needed.

Figure 7-4. Multiple vCenter Server Instances - Single vCenter Single Sign-On Domain

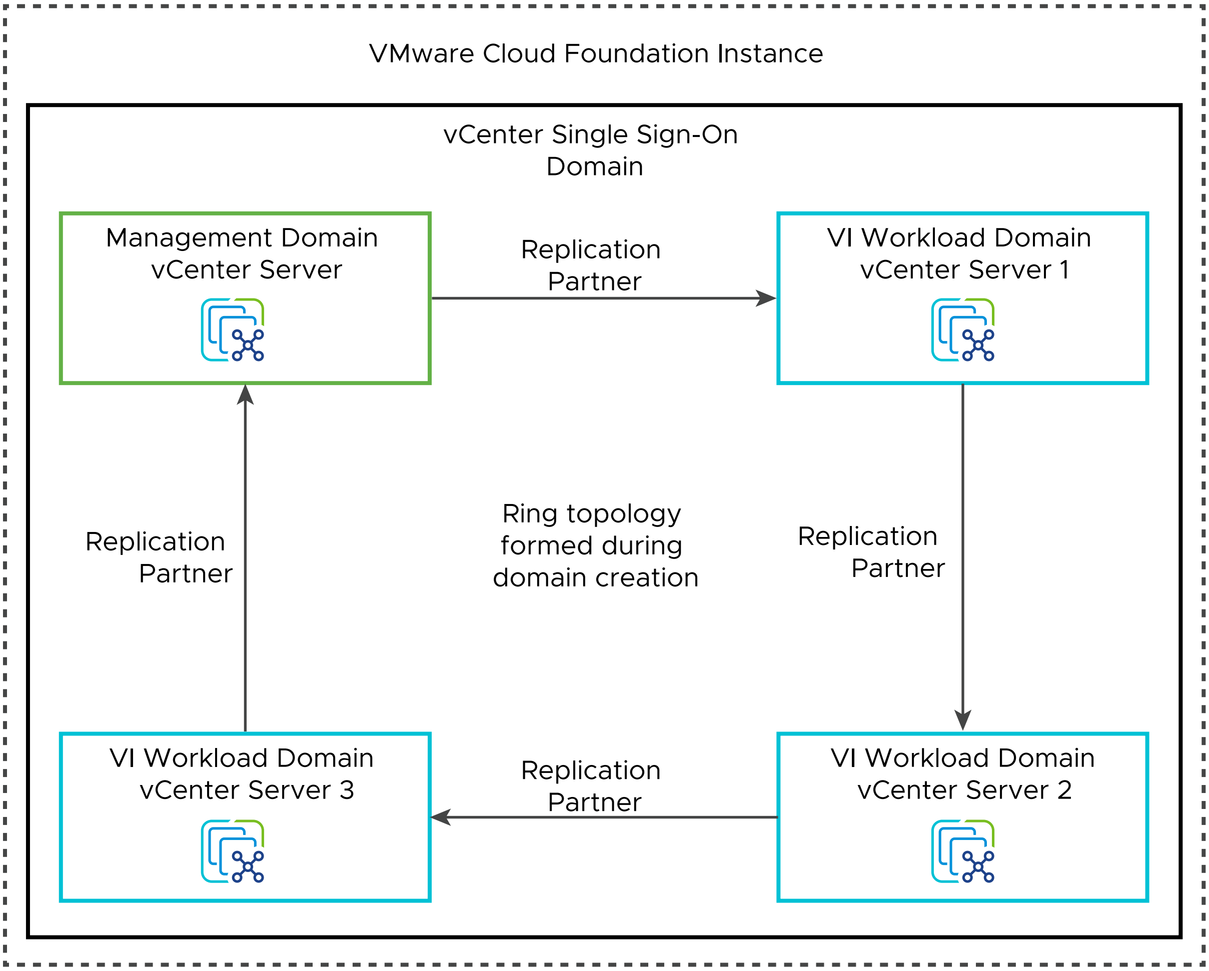
VMware Cloud Foundation Instance
vCenter Single Sign-On Domain
Management Domain vCenter Server
Replication Partner
VI Workload Domain vCenter Server 1
Replication Partner
Ring topology formed during domain creation
Replication Partner
VI Workload Domain vCenter Server 3
Replication Partner
VI Workload Domain vCenter Server 2

Table 7-11. Design Requirements for the Multiple vCenter Server Instance - Single vCenter Single Sign-on Domain Topology for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-SSO-STD-001	Join all vCenter Server instances within aVMware Cloud Foundation instance to a single vCenter Single Sign-On domain.	When all vCenter Server instances are in the same vCenter Single Sign-On domain, they can share authentication and license data across all components.	■ Only one vCenter Single Sign-On domain exists. ■ The number of linked vCenter Server instances in the same vCenter Single Sign-On domain is limited to 15 instances. Because each workload domain uses a dedicated vCenter Server instance, you can deploy up to 15 domains within each VMware Cloud Foundation instance.
VCF-VCS-REQD-SSO-STD-002	Create a ring topology between the vCenter Server instances within the VMware Cloud Foundation instance.	By default, one vCenter Server instance replicates only with another vCenter Server instance. This setup creates a single point of failure for replication. A ring topology ensures that each vCenter Server instance has two replication partners and removes any single point of failure.	None.

Figure 7-5. Multiple vCenter Server Instances - Multiple vCenter Single Sign-On Domain

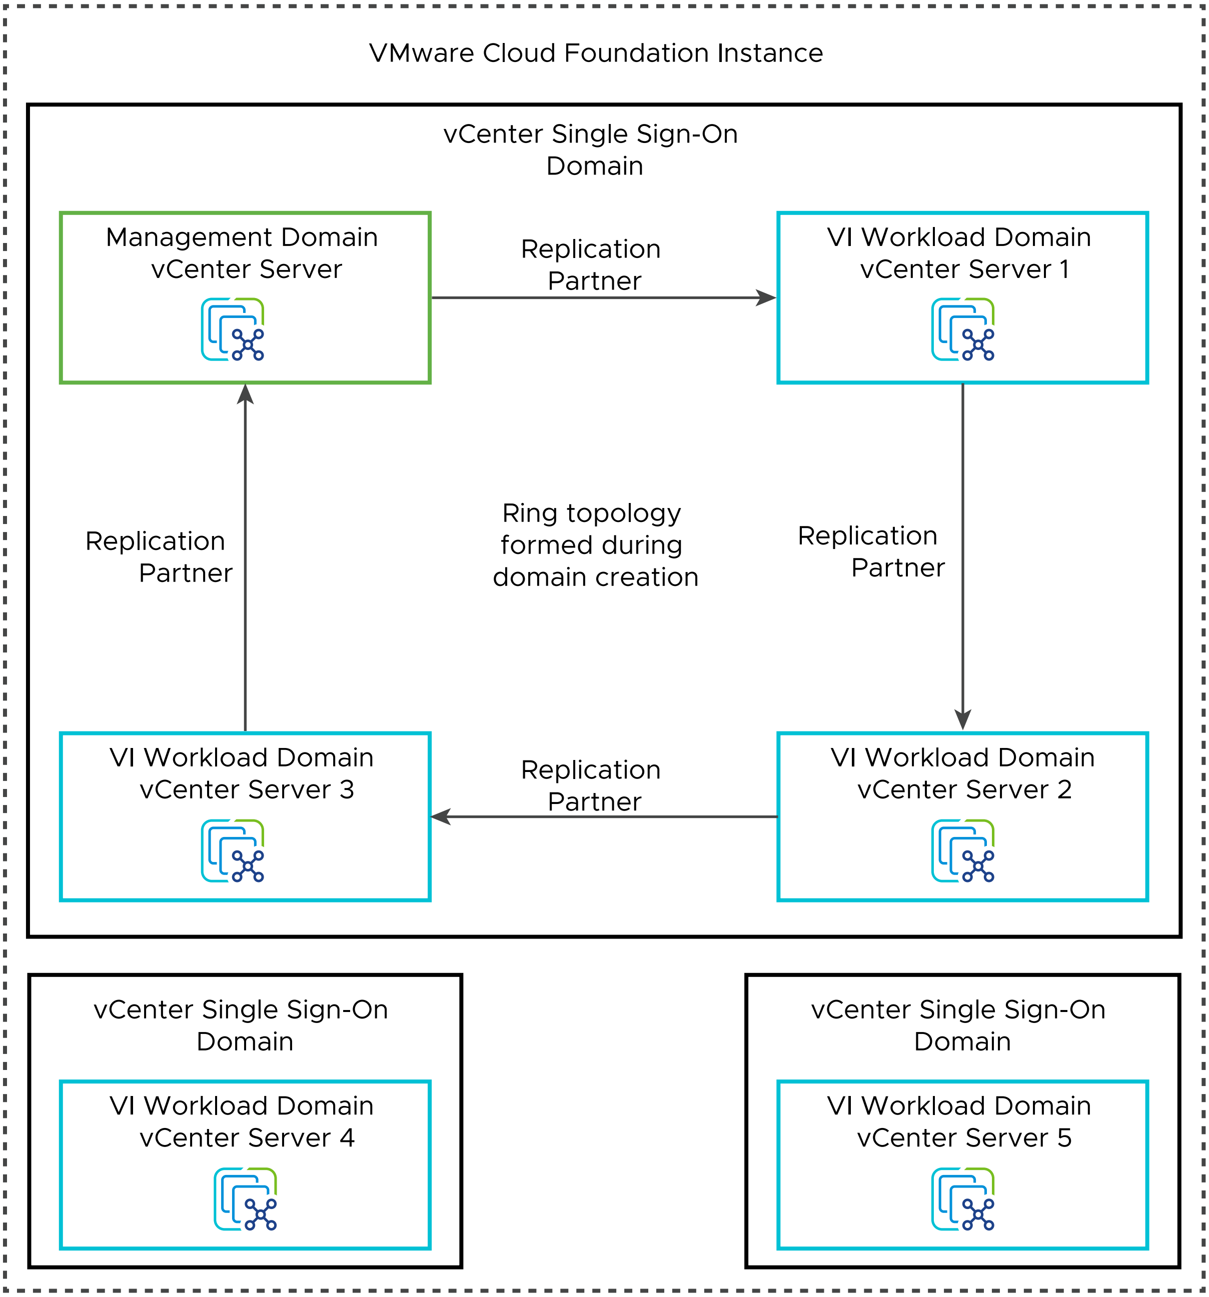
VMware Cloud Foundation Instance
vCenter Single Sign-On Domain
Management Domain vCenter Server
Replication Partner
VI Workload Domain vCenter Server 1
Replication Partner
Ring topology formed during domain creation
Replication Partner
VI Workload Domain vCenter Server 3
Replication Partner
VI Workload Domain vCenter Server 2
vCenter Single Sign-On Domain
VI Workload Domain vCenter Server 4
vCenter Single Sign-On Domain
VI Workload Domain vCenter Server 5

Table 7-12. Design Requirements for Multiple vCenter Server Instance - Multiple vCenter Single Sign-On Domain Topology for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-SSO-ISO-001	Create all vCenter Server instances within a VMware Cloud Foundation instance in their own unique vCenter Single Sign-On domains.	■ Enables isolation at the vCenter Single Sign-On domain layer for increased security separation. ■ Supports up to 25 workload domains.	■ Each vCenter server instance is managed through its own pane of glass using a different set of administrative credentials. ■ You must manage password rotation for each vCenter Single Sign-On domain separately.

vSphere Cluster Design for VMware Cloud Foundation

vSphere cluster design must consider the requirements for standard, stretched and remote clusters and for the life cycle management of the ESXi hosts in the clusters according to the characteristics of the workloads.

- Logical vSphere Cluster Design for VMware Cloud Foundation

  The cluster design must consider the characteristics of the workloads that are deployed in the cluster.

- vSphere Cluster Life Cycle Method Design for VMware Cloud Foundation

  vSphere Lifecycle Manager is used to manage the vSphere clusters in each VI workload domain.

- vSphere Cluster Design Requirements and Recommendations for VMware Cloud Foundation

  The design of a vSphere cluster is a subject to a minimum number of hosts, design requirements, and design recommendations.

Logical vSphere Cluster Design for VMware Cloud Foundation

The cluster design must consider the characteristics of the workloads that are deployed in the cluster.

When you design the cluster layout in vSphere, consider the following guidelines:

- Compare the capital costs of purchasing fewer, larger ESXi hosts with the costs of purchasing more, smaller ESXi hosts. Costs vary between vendors and models. Evaluate the risk of losing one larger host in a scaled-up cluster and the impact on the business with the higher chance of losing one or more smaller hosts in a scale-out cluster.
- Evaluate the operational costs of managing a few ESXi hosts with the costs of managing more ESXi hosts.
- Consider the purpose of the cluster.

- Consider the total number of ESXi hosts and cluster limits.

Figure 7-6. Logical vSphere Cluster Layout with a Single Availability Zone for VMware Cloud Foundation

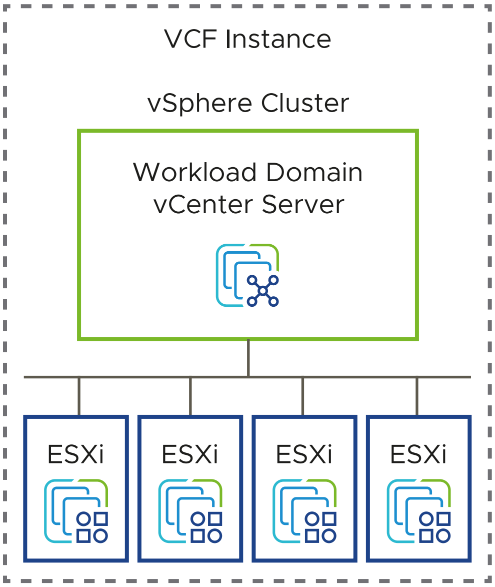

Figure 7-7. Logical vSphere Cluster Layout for Multiple Availability Zones for VMware Cloud Foundation

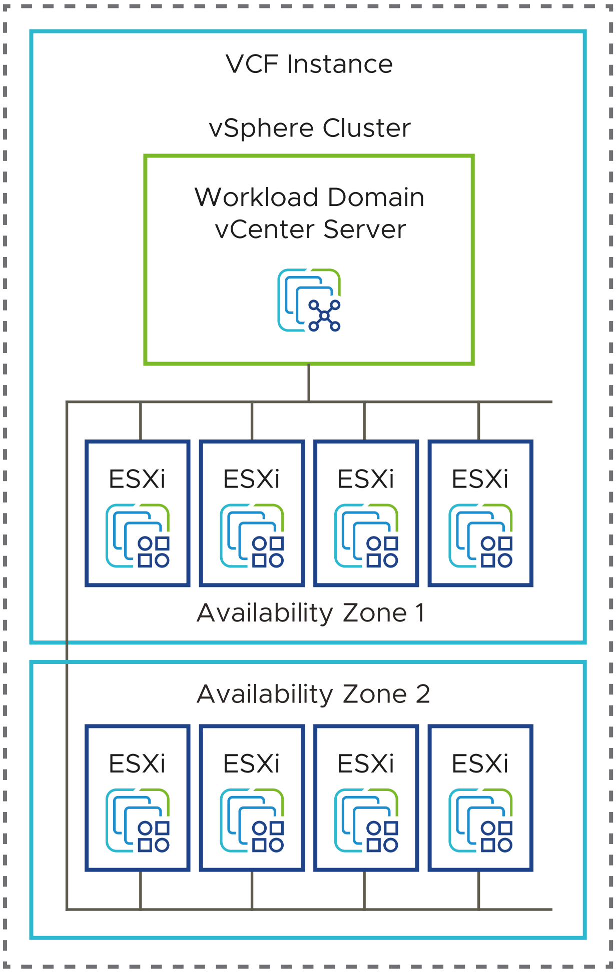

Remote Cluster Design Considerations

Remote clusters are managed by the management infrastructure at the central site.

Table 7-13. Remote Cluster Design Considerations

Remote Cluster Attribute	Consideration
Number of hosts per remote cluster	■ Minimum: 3 ■ Maximum: 16
Number of remote clusters per VMware Cloud Foundation instance	■ Maximum: 8
Number of remote clusters per VI workload domain	1
Cluster types per VI workload domain	A VI workload domain can include either local clusters or a remote cluster.
Latency between the central site and the remote site	■ Maximum: 100 ms
Bandwidth between the central site and the remote site	■ Minimum: 10 Mbps

vSphere Cluster Life Cycle Method Design for VMware Cloud Foundation

vSphere Lifecycle Manager is used to manage the vSphere clusters in each VI workload domain.

When you deploy a workload domain, you choose a vSphere cluster life cycle method according to your requirements.

Table 7-14. vSphere Lifecycle Manager choices

Cluster Life Cycle Method	Description	Benefits	Drawbacks
vSphere Lifecycle Manager images	vSphere Lifecycle Manager images contain base images, vendor add-ons, firmware, and drivers.	■ Supports vSAN stretched clusters. ■ Supports VI workload domains with vSphere with Tanzu. ■ Supports NVIDIA GPU-enabled clusters. ■ Supports 2-node NFS, FC, or vVols clusters.	■ An initial cluster image is required during workload domain or cluster deployment.
vSphere Lifecycle Manager baselines	An upgrade baseline contains the ESXi image and a patch baseline contains the respective patches for ESXi host.	■ Supports vSAN stretched clusters. ■ Supports VI workload domains with vSphere with Tanzu.	■ Not supported for NVIDIA GPU-enabled clusters. ■ Not supported for 2-node NFS, FC, or vVols clusters.

vSphere Cluster Design Requirements and Recommendations for VMware Cloud Foundation

The design of a vSphere cluster is a subject to a minimum number of hosts, design requirements, and design recommendations.

For vSAN design requirements and recommendations, see vSAN Design Requirements and Recommendations for VMware Cloud Foundation.

The requirements for the ESXi hosts in a workload domain in VMware Cloud Foundation are related to the system requirements of the workloads hosted in the domain. The ESXi requirements include number, server configuration, amount of hardware resources, networking, and certificate management. Similar best practices help you design optimal environment operation

vSphere Cluster Design Considerations

You consider different number of hosts per cluster according to the storage type and specific resource requirements for standard and stretched vSAN clusters.

Table 7-15. Host-Related Design Considerations per Cluster

Attribute	Specification	Management Domain (Default Cluster)	Management Domain (Additional Clusters) or VI Workload Domain (All Clusters)
Minimum number of ESXi hosts	vSAN (single availability zone)	4	3
	vSAN (two availability zones)	8	6
	NFS, FC, or vVols	Not supported	■ 2 ■ VI workload domain only ■ Requires vSphere Lifecycle Manager images ■ 3 ■ Additional management clusters
Reserved capacity for handling ESXi host failures per cluster	Single availability zone	■ 25% CPU and memory ■ Tolerates one host failure	■ 33% CPU and memory ■ Tolerates one host failure
	Two availability zones	■ 50% CPU and memory ■ Tolerates one availability zone failure	■ 50% CPU and memory ■ Tolerates one availability zone failure

vSphere Cluster Design Requirements VMware Cloud Foundation

You must meet the following design requirements for standard and stretched clusters in your vSphere cluster design for VMware Cloud Foundation.The cluster design considers the storage type for the cluster, the architecture model of the environment, and the life cycle management method .

Table 7-16. vSphere Cluster Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-001	Create a cluster in each workload domain for the initial set of ESXi hosts.	■ Simplifies configuration by isolating management from customer workloads. ■ Ensures that customer workloads have no impact on the management stack.	Management of multiple clusters and vCenter Server instances increases operational overhead.
VCF-CLS-REQD-CFG-002	Allocate a minimum number of ESXi hosts according to the cluster type being deployed.	■ Ensures correct level of redundancy to protect against host failure in the cluster.	To support redundancy, you must allocate additional ESXi host resources.

Table 7-16. vSphere Cluster Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-003	If using a consolidated workload domain, configure the following vSphere resource pools to control resource usage by management and customer workloads. ■ *cluster-name*-rp-sddc-mgmt ■ *cluster-name*-rp-sddc-edge ■ *cluster-name*-rp-user-edge ■ *cluster-name*-rp-user-vm	■ Ensures sufficient resources for the management components.	You must manage the vSphere resource pool settings over time.
VCF-CLS-REQD-CFG-004	Configure the vSAN network gateway IP address as the isolation address for the cluster.	Allows vSphere HA to validate if a host is isolated from the vSAN network.	None.
VCF-CLS-REQD-CFG-005	Set the advanced cluster setting `das.usedefaultisolationaddress` to false.	Ensures that vSphere HA uses the manual isolation addresses instead of the default management network gateway address.	None.

Table 7-17. vSphere Cluster Design Requirements for vSAN Stretched Clusters with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-006	Configure the vSAN network gateway IP addresses for the second availability zone as an additional isolation addresses for the cluster.	Allows vSphere HA to validate if a host is isolated from the vSAN network for hosts in both availability zones.	None.
VCF-CLS-REQD-CFG-007	Enable the **Override default gateway for this adapter** setting on the vSAN VMkernel adapters on all ESXi hosts.	Enables routing the vSAN data traffic through the vSAN network gateway rather than through the management gateway.	vSAN networks across availability zones must have a route to each other.
VCF-CLS-REQD-CFG-008	Create a host group for each availability zone and add the ESXi hosts in the zone to the respective group.	Makes it easier to manage which virtual machines run in which availability zone.	You must create and maintain VM-Host DRS group rules.

vSphere Cluster Design Recommendations for VMware Cloud Foundation

In your vSphere cluster design, you can apply certain best practices for standard and stretched clusters .

Table 7-18. vSphere Cluster Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-001	Use vSphere HA to protect all virtual machines against failures.	vSphere HA supports a robust level of protection for both ESXi host and virtual machine availability.	You must provide sufficient resources on the remaining hosts so that virtual machines can be restarted on those hosts in the event of a host outage.
VCF-CLS-RCMD-CFG-002	Set host isolation response to Power Off and restart VMs in vSphere HA.	vSAN requires that the host isolation response be set to Power Off and to restart virtual machines on available ESXi hosts.	If a false positive event occurs, virtual machines are powered off and an ESXi host is declared isolated incorrectly.
VCF-CLS-RCMD-CFG-003	Configure admission control for 1 ESXi host failure and percentage-based failover capacity.	Using the percentage-based reservation works well in situations where virtual machines have varying and sometimes significant CPU or memory reservations. vSphere automatically calculates the reserved percentage according to the number of ESXi host failures to tolerate and the number of ESXi hosts in the cluster.	In a cluster of 4 ESXi hosts, the resources of only 3 ESXi hosts are available for use.
VCF-CLS-RCMD-CFG-004	Enable VM Monitoring for each cluster.	VM Monitoring provides in-guest protection for most VM workloads. The application or service running on the virtual machine must be capable of restarting successfully after a reboot or the virtual machine restart is not sufficient.	None.
VCF-CLS-RCMD-CFG-005	Set the advanced cluster setting `das.iostatsinterval` to 0 to deactivate monitoring the storage and network I/O activities of the management appliances.	Enables triggering a restart of a management appliance when an OS failure occurs and heartbeats are not received from VMware Tools instead of waiting additionally for the I/O check to complete.	If you want to specifically enable I/O monitoring, you must configure the das.iostatsinterval advanced setting.

Table 7-18. vSphere Cluster Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-006	Enable vSphere DRS on all clusters, using the default fully automated mode with medium threshold.	Provides the best trade-off between load balancing and unnecessary migrations with vSphere vMotion.	If a vCenter Server outage occurs, the mapping from virtual machines to ESXi hosts might be difficult to determine.
VCF-CLS-RCMD-CFG-007	Enable Enhanced vMotion Compatibility (EVC) on all clusters in the management domain.	Supports cluster upgrades without virtual machine downtime.	You must enable EVC only if the clusters contain hosts with CPUs from the same vendor. You must enable EVC on the default management domain cluster during bringup.
VCF-CLS-RCMD-CFG-008	Set the cluster EVC mode to the highest available baseline that is supported for the lowest CPU architecture on the hosts in the cluster.	Supports cluster upgrades without virtual machine downtime.	None.
VCF-CLS-RCMD-LCM-001	Use images as the life cycle management method for VI workload domains.	vSphere Lifecycle Manager images simplify the management of firmware and vendor add-ons manually.	An initial cluster image is required during workload domain or cluster deployment.

Table 7-19. vSphere Cluster Design Recommendations for vSAN Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-009	Increase admission control percentage to half of the ESXi hosts in the cluster.	Allocating only half of a stretched cluster ensures that all VMs have enough resources if an availability zone outage occurs.	In a cluster of 8 ESXi hosts, the resources of only 4 ESXi hosts are available for use. If you add more ESXi hosts to the default management cluster, add them in pairs, one per availability zone.
VCF-CLS-RCMD-CFG-010	Create a virtual machine group for each availability zone and add the VMs in the zone to the respective group.	Ensures that virtual machines are located only in the assigned availability zone to avoid unnecessary vSphere vMotion migrations.	You must add virtual machines to the allocated group manually.
VCF-CLS-RCMD-CFG-011	Create a should-run-on-hosts-in-group VM-Host affinity rule to run each group of virtual machines on the respective group of hosts in the same availability zone.	Ensures that virtual machines are located only in the assigned availability zone to avoid unnecessary vSphere vMotion migrations.	You must manually create the rules.

vSphere Networking Design for VMware Cloud Foundation

VMware Cloud Foundation uses vSphere Distributed Switch for virtual networking.

Logical vSphere Networking Design for VMware Cloud Foundation

When you design vSphere networking, consider the configuration of the vSphere Distributed Switches, distributed port groups, and VMkernel adapters in the VMware Cloud Foundation environment.

vSphere Distributed Switch Design

The default cluster in a workload domain uses a single vSphere Distributed Switch with a configuration for system traffic types, NIC teaming, and MTU size.

VMware Cloud Foundation supports NSX Overlay traffic over a single vSphere Distributed Switch per cluster. Additional distributed switches are supported for other traffic types.

When using vSAN ReadyNodes, you must define the number of vSphere Distributed Switches at workload domain deployment time. You cannot add additional vSphere Distributed Switches post deployment.

Table 7-20. Configuration Options for vSphere Distributed Switch for VMware Cloud Foundation

vSphere Distributed Switch Configuration	Management Domain Options	VI Workload Domain Options	Benefits	Drawbacks
Single vSphere Distributed Switch for hosts with two physical NICs	■ One vSphere Distributed Switch for each cluster with all traffic using two uplinks.	■ One vSphere Distributed Switch for each cluster with all traffic using two uplinks.	Requires the least number of physical NICs and switch ports.	All traffic shares the same two uplinks.
Single vSphere Distributed Switch for hosts with four or six physical NICs	■ One vSphere Distributed Switch for each cluster with four uplinks by using the predefined profiles in the *Deployment Parameters Workbook* in VMware Cloud Builder to deploy the default management cluster. ■ One vSphere Distributed Switch for each cluster with four or six uplinks by using the VMware Cloud Builder API to deploy the default management cluster.	■ One vSphere Distributed Switch for each cluster with four or six uplinks.	■ Provides support for traffic separation across different uplinks.	■ You must provide additional physical NICs and switch ports.
Multiple vSphere Distributed Switches	■ Maximum two vSphere Distributed Switches by using the predefined profiles in the *Deployment Parameters Workbook* in VMware Cloud Builder to deploy the default management cluster.	■ Maximum 16 vSphere Distributed Switches per cluster. ■ You can use only one of the vSphere Distributed Switches for NSX overlay traffic.	■ Provides support for traffic separation across different uplinks or vSphere Distributed Switches. ■ Provides support for traffic separation onto different physical network fabrics.	■ You must provide additional physical NICs and switch ports. ■ More complex with additional configuration and management overhead.

Table 7-20. Configuration Options for vSphere Distributed Switch for VMware Cloud Foundation (continued)

vSphere Distributed Switch Configuration	Management Domain Options	VI Workload Domain Options	Benefits	Drawbacks
	■ Maximum 16 vSphere Distributed Switches per cluster. You use the VMware Cloud Builder API to deploy the default management cluster using combinations of vSphere Distributed Switches and physical NIC configurations that are not available as predefined profiles in the *Deployment Parameters Workbook* ■ You can use only one of the vSphere Distributed Switches for NSX overlay traffic.			

Distributed Port Group Design

VMware Cloud Foundation requires several port groups on the vSphere Distributed Switch for a workload domain. The VMkernel adapters for the host TEPs are connected to the host overlay network, but does not require a dedicated port group on the distributed switch. The VMkernel network adapter for host TEP is automatically created VMware Cloud Foundation configures the ESXi host as a transport node.

Table 7-21. Distributed Port Group Configuration for VMware Cloud Foundation

Function	Teaming Policy	Management Domain	VI Workload Domain
■ Management ■ vSphere vMotion ■ vSAN	Route based on physical NIC load	Required.	Recommended.
	■ Failover Detection: Link status only ■ Failback: Yes Occurs only on saturation of the active uplink. ■ Notify Switches: Yes	Recommended.	Recommended.
■ Host Overlay	Not applicable.	Not applicable.	Not applicable.
■ Edge Uplinks and Overlay	Use explicit failover order.	Required.	Required.
■ Edge RTEP (NSX Federation Only)	Not applicable.	Not applicable.	Not applicable.

VMkernel Network Adapter Design

The VMkernel networking layer provides connectivity to hosts and handles the system traffic for management, vSphere vMotion, vSphere HA, vSAN, and others.

Table 7-22. Default VMkernel Adapters for a Workload Domain per Availability Zone

VMkernel Adapter Service	Connected Port Group	Activated Services	Recommended MTU Size (Bytes)
Management	Management Port Group	Management Traffic	1500 (Default)
vMotion	vMotion Port Group	vMotion Traffic	9000
vSAN	vSAN Port Group	vSAN	9000
Host TEPs	Not applicable	Not applicable	9000

vSphere Networking Design Recommendations for VMware Cloud Foundation

Consider the recommendations for vSphere networking in VMware Cloud Foundation, such as MTU size, port binding, teaming policy and traffic-specific network shares.

Table 7-23. vSphere Networking Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VDS-RCMD-CFG-001	Use a single vSphere Distributed Switch per cluster.	■ Reduces the complexity of the network design. ■ Reduces the size of the fault domain.	Increases the number of vSphere Distributed Switches that must be managed.
VCF-VDS-RCMD-CFG-002	Configure the MTU size of the vSphere Distributed Switch to 9000 for jumbo frames.	■ Supports the MTU size required by system traffic types. ■ Improves traffic throughput.	When adjusting the MTU packet size, you must also configure the entire network path (VMkernel ports, virtual switches, physical switches, and routers) to support the same MTU packet size.
VCF-VDS-RCMD-DPG-001	Use ephemeral port binding for the Management VM port group.	Using ephemeral port binding provides the option for recovery of the vCenter Server instance that is managing the distributed switch.	Port-level permissions and controls are lost across power cycles, and no historical context is saved.
VCF-VDS-RCMD-DPG-002	Use static port binding for all non-management port groups.	Static binding ensures a virtual machine connects to the same port on the vSphere Distributed Switch. This allows for historical data and port level monitoring.	None.
VCF-VDS-RCMD-DPG-003	Use the `Route based on physical NIC load` teaming algorithm for the VM management port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-004	Use the `Route based on physical NIC load` teaming algorithm for the ESXi management port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-005	Use the `Route based on physical NIC load` teaming algorithm for the vSphere vMotion port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-006	Use the `Route based on physical NIC load` teaming algorithm for the vSAN port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.

Table 7-23. vSphere Networking Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VDS-RCMD-NIO-001	Enable Network I/O Control on vSphere Distributed Switch of the management domain cluster.	Increases resiliency and performance of the network.	Network I/O Control might impact network performance for critical traffic types if misconfigured.
VCF-VDS-RCMD-NIO-002	Set the share value for management traffic to Normal.	By keeping the default setting of Normal, management traffic is prioritized higher than vSphere vMotion but lower than vSAN traffic. Management traffic is important because it ensures that the hosts can still be managed during times of network contention.	None.
VCF-VDS-RCMD-NIO-003	Set the share value for vSphere vMotion traffic to Low.	During times of network contention, vSphere vMotion traffic is not as important as virtual machine or storage traffic.	During times of network contention, vMotion takes longer than usual to complete.
VCF-VDS-RCMD-NIO-004	Set the share value for virtual machines to High.	Virtual machines are the most important asset in the SDDC. Leaving the default setting of High ensures that they always have access to the network resources they need.	None.
VCF-VDS-RCMD-NIO-005	Set the share value for vSAN traffic to High.	During times of network contention, vSAN traffic needs guaranteed bandwidth to support virtual machine performance.	None.
VCF-VDS-RCMD-NIO-006	Set the share value for other traffic types to Low.	By default, VVMware Cloud Foundation does not use other traffic types, like vSphere FT traffic. Hence, these traffic types can be set the lowest priority.	None.

NSX Design for VMware Cloud Foundation

8

In VMware Cloud Foundation, you use NSX for connecting management and customer virtual machines by using virtual network segments and routing. You also create constructs for solutions that are deployed for a single VMware Cloud Foundation instance or are available across multiple VMware Cloud Foundation instances. These constructs provide routing to the data center and load balancing.

Table 8-1. NSX Logical Concepts and Components

Component	Description
NSX Manager	■ Provides the user interface and the REST API for creating, configuring, and monitoring NSX components, such as segments, and Tier-0 and Tier-1 gateways. ■ In a deployment with NSX Federation, NSX Manager is called NSX Local Manager.
NSX Edge nodes	■ Is a special type of transport node which contains service router components. ■ Provides north-south traffic connectivity between the physical data center networks and the NSX SDN networks. Each NSX Edge node has multiple interfaces where traffic flows. ■ Can provide east-west traffic flow between virtualized workloads. They provide stateful services such as load balancers and DHCP. In a deployment with multiple VMware Cloud Foundation instances, east-west traffic between the VMware Cloud Foundation instances flows through the NSX Edge nodes too.
NSX Federation (optional design extension)	■ Propagates configurations that span multiple NSX instances in a single VMware Cloud Foundation instance or across multiple VMware Cloud Foundation instances. You can stretch overlay segments, activate failover of segment ingress and egress traffic between VMware Cloud Foundation instances, and implement a unified firewall configuration. ■ In a deployment with multiple VMware Cloud Foundation instances, you use NSX to provide cross-instance services to SDDC management components that do not have native support for availability at several locations, such as VMware Aria Automation and VMware Aria Operations. ■ Connect only workload domains of matching types (management domain to management domain or VI workload domain to VI workload domain).

Table 8-1. NSX Logical Concepts and Components (continued)

Component	Description
NSX Global Manager (Federation only)	■ Is part of deployments with multiple VMware Cloud Foundation instances where NSX Federation is required. NSX Global Manager can connect multiple NSX Local Manager instances under a single global management plane. ■ Provides the user interface and the REST API for creating, configuring, and monitoring NSX global objects, such as global virtual network segments, and global Tier-0 and Tier-1 gateways. ■ Connected NSX Local Manager instances create the global objects on the underlying software-defined network that you define from NSX Global Manager. An NSX Local Manager instance directly communicates with other NSX Local Manager instances to synchronize configuration and state needed to implement a global policy. ■ NSX Global Manager is a deployment-time role that you assign to an NSX Manager appliance.
NSX Manager instance shared between VI workload domains	■ An NSX Manager instance can be shared between up to 14 VI workload domains that are part of the same vCenter Single Sign-On domain. ■ VI workload domains sharing an NSX Manager instance must use the same vSphere cluster life cycle method. ■ Using a shared NSX Manager instance reduces resource requirements for the management domain. ■ A single transport zone is shared across all clusters in all VI workload domains that share the NSX Manager instance. ■ The management domain NSX instance cannot be shared. ■ Isolated workload domain NSX instances cannot be shared.

Read the following topics next:

- Logical Design for NSX for VMware Cloud Foundation
- NSX Manager Design for VMware Cloud Foundation
- NSX Edge Node Design for VMware Cloud Foundation
- Routing Design for VMware Cloud Foundation
- Overlay Design for VMware Cloud Foundation
- Application Virtual Network Design for VMware Cloud Foundation
- Load Balancing Design for VMware Cloud Foundation

Logical Design for NSX for VMware Cloud Foundation

NSX provides networking services to workloads in VMware Cloud Foundation such as load balancing, routing and virtual networking.

Table 8-2. NSX Logical Design

Component	VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
NSX Manager Cluster	■ Three appropriately sized nodes with a virtual IP (VIP) address with an anti-affinity rule to keep them on different hosts. ■ vSphere HA protects the cluster nodes applying high restart priority	■ Three appropriately sized nodes with a VIP address with an anti-affinity rule to keep them on different hosts. ■ vSphere HA protects the cluster nodes applying high restart priority ■ vSphere DRS rule should-run-on-hosts-in-group keeps the NSX Manager VMs in the first availability zone.
NSX Global Manager Cluster (Conditional)	■ Manually deployed three appropriately sized nodes with a VIP address with an anti-affinity rule to run them on different hosts. ■ One active and one standby cluster. ■ vSphere HA protects the cluster nodes applying high restart priority.	■ Manually deployed three appropriately sized nodes with a VIP address with an anti-affinity rule to run them on different hosts. ■ One active and one standby cluster. ■ vSphere HA protects the cluster nodes applying high restart priority. ■ vSphere DRS rule should-run-on-hosts-in-group keeps the NSX Global Manager VMs in the first availability zone.
NSX Edge Cluster	■ Two appropriately sized NSX Edge nodes with an anti-affinity rule to separate them on different hosts. ■ vSphere HA protects the cluster nodes applying high restart priority.	■ Two appropriately sized NSX Edge nodes in the first availability zone with an anti-affinity rule to separate them on different hosts. ■ vSphere HA protects the cluster nodes applying high restart priority. ■ vSphere DRS rule should-run-on-hosts-in-group keeps the NSX Edge VMs in the first availability zone.
Transport Nodes	■ Each ESXi host acts as a host transport node. ■ Two edge transport nodes.	■ Each ESXi host acts as a host transport node. ■ Two edge transport nodes in the first availability zone.

Table 8-2. NSX Logical Design (continued)

Component	VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
Transport zones	■ One VLAN transport zone for north-south traffic. ■ Maximum one overlay transport zone for overlay segments per NSX instance. ■ One VLAN tranpsort zone for VLAN-backed segments.	■ One VLAN transport zone for north-south traffic. ■ Maximum one overlay transport zone for overlay segments per NSX instance. ■ One or more VLAN tranpsort zones for VLAN-backed segments.
VLANs and IP subnets allocated to NSX For information about the networks for virtual infrastructure management, see Distributed Port Group Design.	See VLANs and Subnets for VMware Cloud Foundation.	See VLANs and Subnets for VMware Cloud Foundation.
Routing configuration	■ BGP for a single VMware Cloud Foundation instance. ■ In a VMware Cloud Foundation deployment with NSX Federation, BGP with ingress and egress traffic to the first VMware Cloud Foundation instance during normal operating conditions.	■ BGP with path prepend to control ingress traffic and local preference to control egress traffic through the first availability zone during normal operating condition. ■ In a VMware Cloud Foundation deployment with NSX Federation, BGP with ingress and egress traffic to the first instance during normal operating conditions.

For a description of the NSX logical component in this design, see Table 8-1. NSX Logical Concepts and Components.

Single Instance - Single Availability Zone

The NSX design for the Single Instance - Single Availability Zone topology consists of the following components:

Figure 8-1. NSX Logical Design for a Single Instance - Single Availability Zone Topology

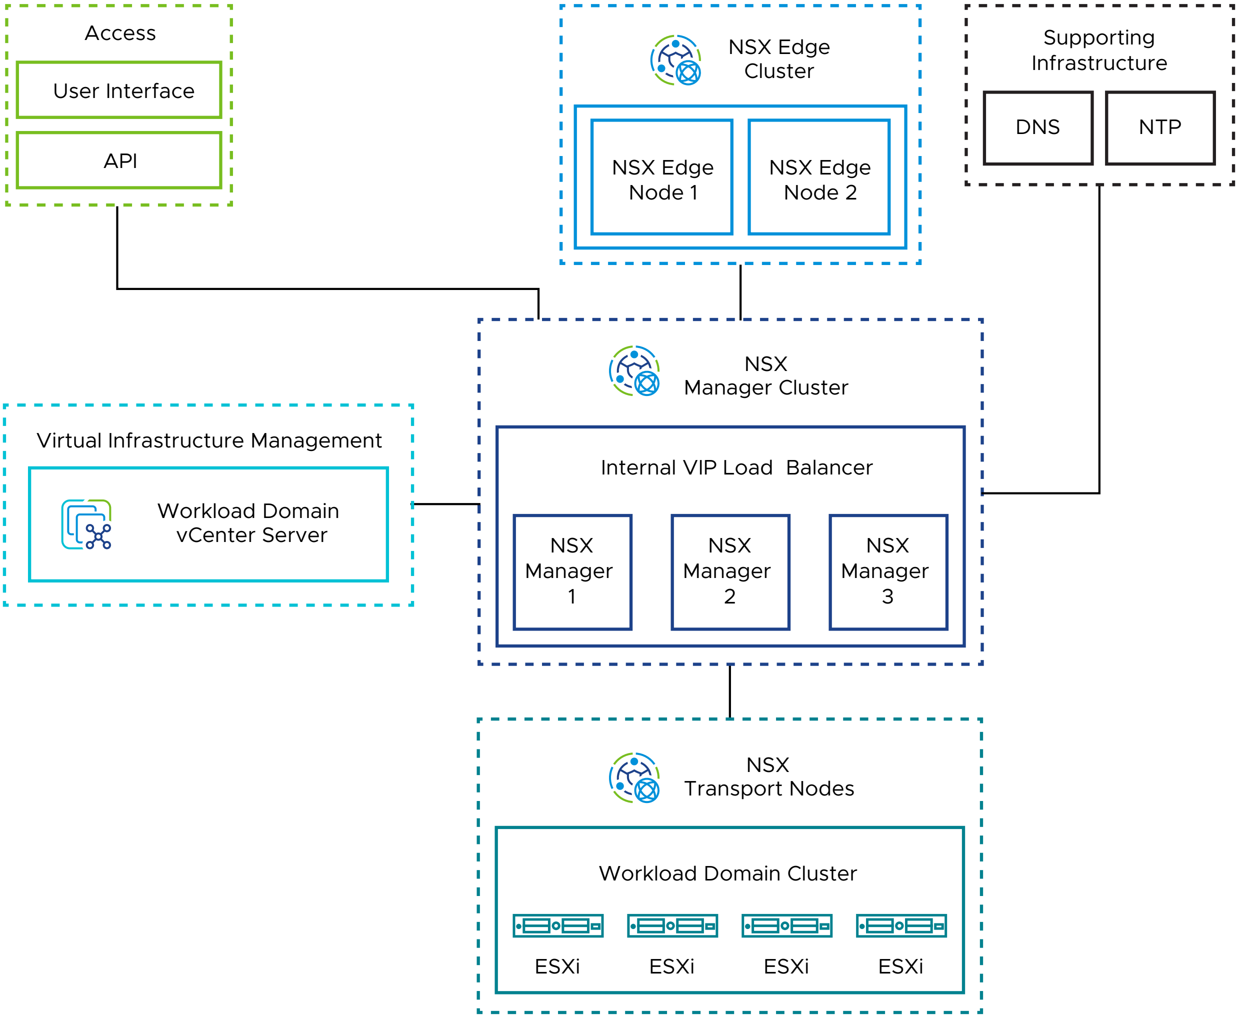

- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes in the workload domain that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts in the workload domain that are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.

Single Instance - Multiple Availability Zones

The NSX design for a Single Instance - Multiple Availability Zone topology consists of the following components:

Figure 8-2. NSX Logical Design for a Single Instance - Multiple Availability Zone Topology

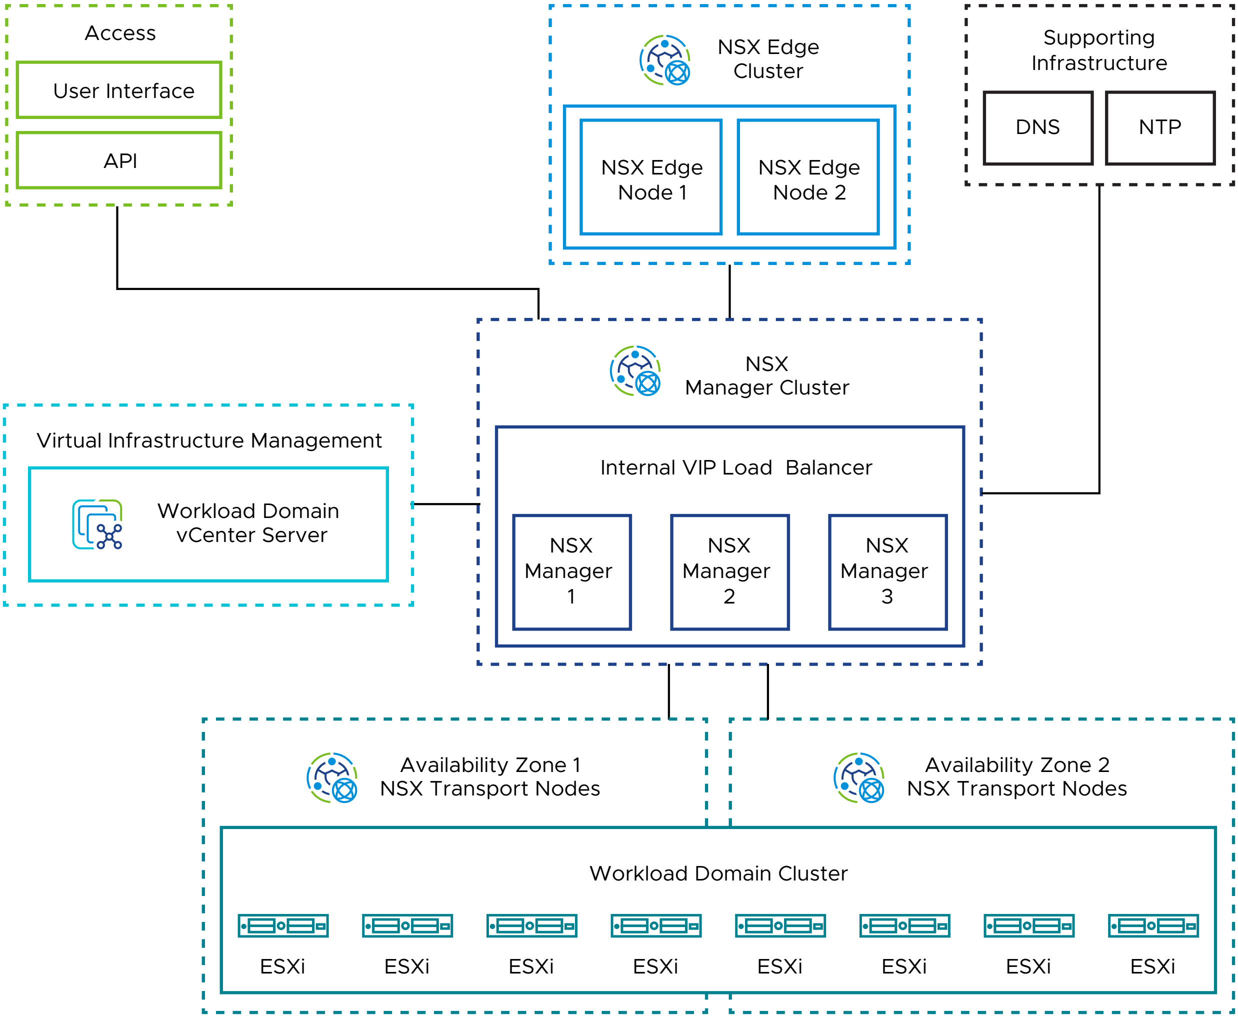
Access
User Interface
API
NSX Edge Cluster
NSX Edge Node 1
NSX Edge Node 2
Supporting Infrastructure
DNS
NTP
NSX Manager Cluster
Internal VIP Load Balancer
NSX Manager 1
NSX Manager 2
NSX Manager 3
Virtual Infrastructure Management
Workload Domain vCenter Server
Availability Zone 1 NSX Transport Nodes
Availability Zone 2 NSX Transport Nodes
Workload Domain Cluster
ESXi
ESXi
ESXi
ESXi
ESXi
ESXi
ESXi
ESXi

- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts that are distributed evenly across availability zones in the workload domain and are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.

Multiple Instances - Single Availability Zone

The NSX design for a Multiple Instance - Single Availability Zone topology consists of the following components:

Figure 8-3. NSX Logical Design for a Multiple Instance - Single Availability Zone Topology

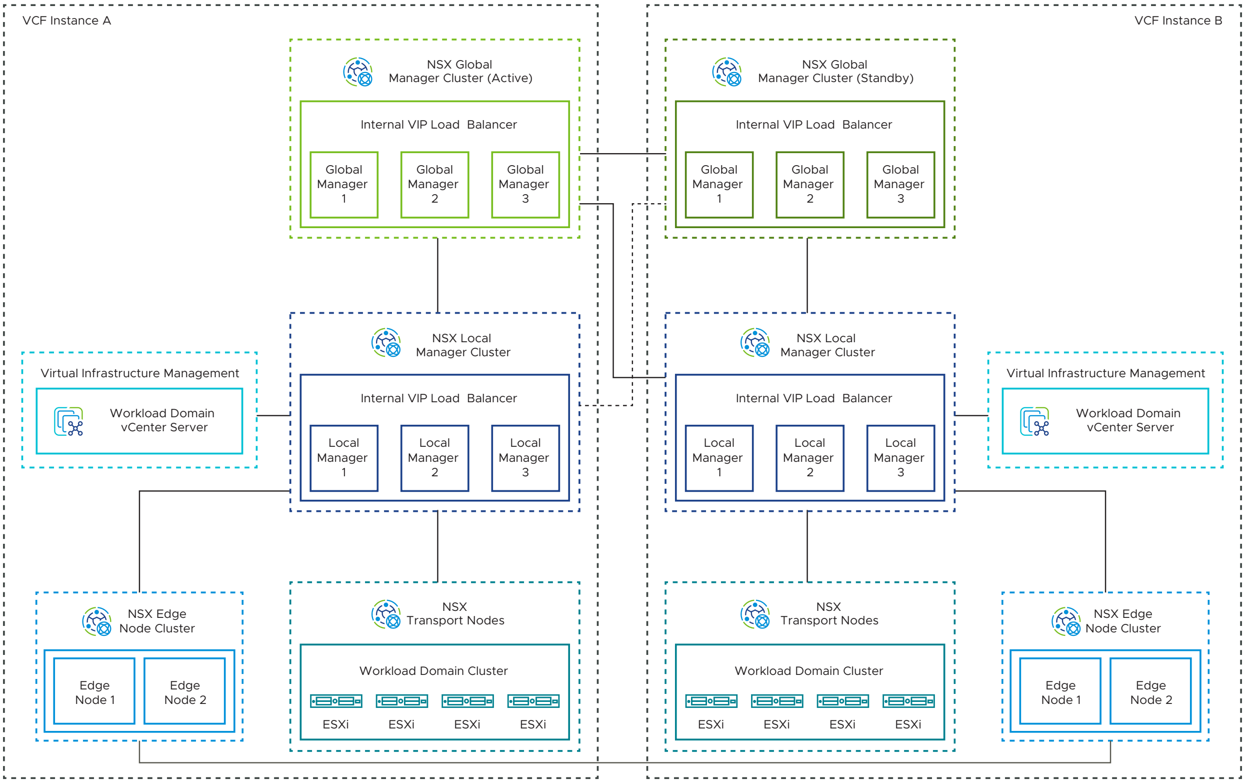

- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts in the workload domain that are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.
- NSX Global Manager cluster in each of the first two VMware Cloud Foundation instances.

  You deploy the NSX Global Manager cluster in each VMware Cloud Foundation instance so that you can use NSX Federation for global management of networking and security services.
- An additional infrastructure VLAN in each VMware Cloud Foundation instance to carry instance-to-instance traffic (RTEP).

Multiple Instances - Multiple Availability Zones

The NSX design for a Multiple Instance - Multiple Availability Zone topology consists of the following components:

Figure 8-4. NSX Logical Design for Multiple Instance - Multiple Availability Zone Topology

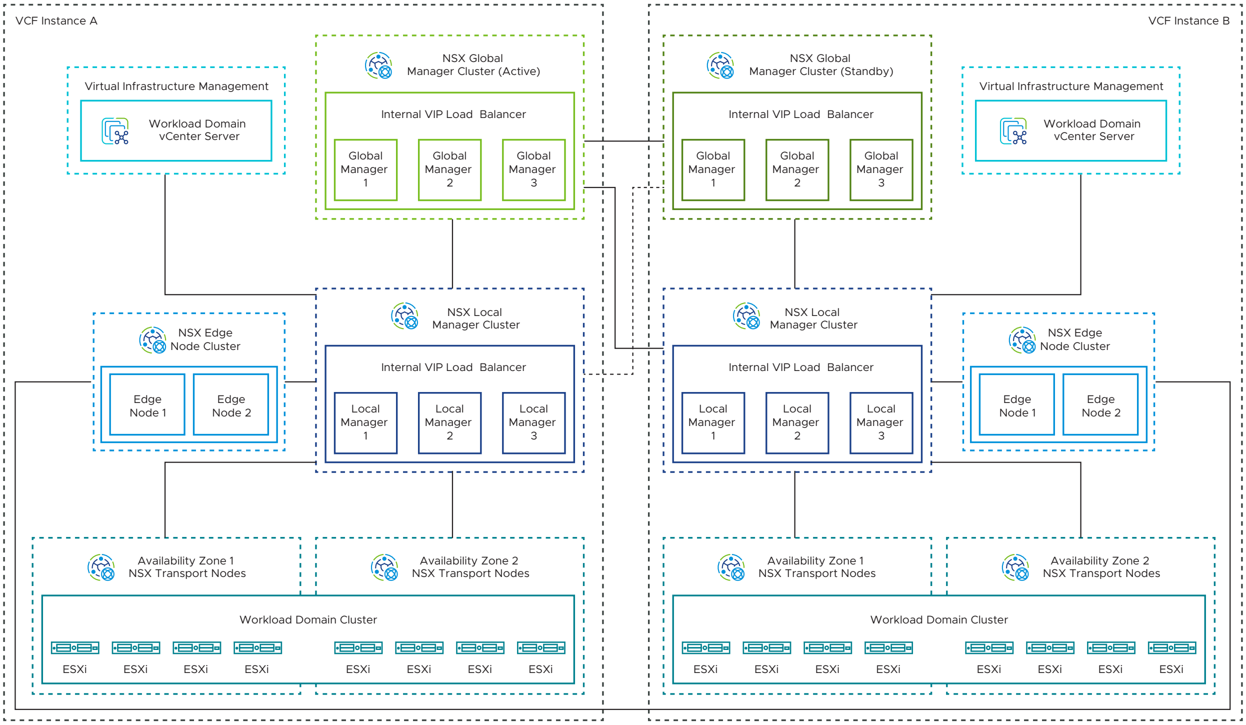
VCF Instance A
VCF Instance B
NSX Global Manager Cluster (Active)
NSX Global Manager Cluster (Standby)
Virtual Infrastructure Management
Workload Domain vCenter Server
Internal VIP Load Balancer
Global Manager 1
Global Manager 2
Global Manager 3
NSX Local Manager Cluster
Local Manager 1
Local Manager 2
Local Manager 3
NSX Edge Node Cluster
Edge Node 1
Edge Node 2
Availability Zone 1 NSX Transport Nodes
Availability Zone 2 NSX Transport Nodes
Workload Domain Cluster
ESXi

- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts that are distributed evenly across availability zones in the workload domain in a VMware Cloud Foundation instance, and are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.
- NSX Global Manager cluster in each of the first two VMware Cloud Foundation instances.

  You deploy the NSX Global Manager cluster in each VMware Cloud Foundation instance so that you can use NSX Federation for global management of networking and security services.
- An additional infrastructure VLAN in each VMware Cloud Foundation instance to carry instance-to-instance traffic (RTEP).

NSX Manager Design for VMware Cloud Foundation

Following the principles of this design and of each product, you determine the size of, deploy and configure NSX Manager as part of your VMware Cloud Foundation deployment.

Sizing Considerations for NSX Manager for VMware Cloud Foundation

You select an appropriate NSX Manager appliance size that is suitable for the scale of your environment.

When you deploy NSX Manager appliances, either with a local or global scope, you select to deploy the appliance with a size that is suitable for the scale of your environment. The option that you select determines the number of CPUs and the amount of memory of the appliance. For detailed sizing according to the overall profile of the VMware Cloud Foundation instance you plan to deploy, see VMware Cloud Foundation Planning and Preparation Workbook.

Table 8-3. Sizing Considerations for NSX Manager

NSX Manager Appliance Size	Scale
Extra-Small	Cloud Service Manager only
Small	Proof of concept
Medium Default for the management domain	Up to 128 ESXi hosts
Large Default for VI workload domains	Up to 1,024 ESXi hosts

Note To deploy an NSX Manager appliance in the VI workload domain with a size different from the default one, you must use the API.

NSX Manager Design Requirements and Recommendations for VMware Cloud Foundation

Consider the placement requirements for using NSX Manager in VMware Cloud Foundation, and the best practices for having an NSX Manager cluster operate in an optimal way, such as number and size of the nodes, and high availability, on a standard or stretched management cluster.

NSX Manager Design Requirements for VMware Cloud Foundation

You must meet the following design requirements for in your NSX Manager design for VMware Cloud Foundation.

Table 8-4. NSX Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LM-REQD-CFG-001	Place the appliances of the NSX Manager cluster on the VM management network in the management domain.	■ Simplifies IP addressing for management VMs by using the same VLAN and subnet. ■ Provides simplified secure access to management VMs in the same VLAN network.	None.
VCF-NSX-LM-REQD-CFG-002	Deploy three NSX Manager nodes in the default vSphere cluster in the management domain for configuring and managing the network services for the workload domain.	Supports high availability of the NSX manager cluster.	You must have sufficient resources in the default cluster of the management domain to run three NSX Manager nodes.

NSX Manager Design Recommendations for VMware Cloud Foundation

In your NSX Manager design for VMware Cloud Foundation, you can apply certain best practices for standard and stretched clusters.

Table 8-5. NSX Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-LM-RCMD-CFG-001	Deploy appropriately sized nodes in the NSX Manager cluster for the workload domain.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is Medium, and for VI workload domains is Large.
VCF-NSX-LM-RCMD-CFG-002	Create a virtual IP (VIP) address for the NSX Manager cluster for the workload domain.	Provides high availability of the user interface and API of NSX Manager.	■ The VIP address feature provides high availability only. It does not load-balance requests across the cluster. ■ When using the VIP address feature, all NSX Manager nodes must be deployed on the same Layer 2 network.

Table 8-5. NSX Manager Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-LM-RCMD-CFG-003	Apply VM-VM anti-affinity rules in vSphere Distributed Resource Scheduler (vSphere DRS) to the NSX Manager appliances.	Keeps the NSX Manager appliances running on different ESXi hosts for high availability.	You must allocate at least four physical hosts so that the three NSX Manager appliances continue running if an ESXi host failure occurs.
VCF-NSX-LM-RCMD-CFG-004	In vSphere HA, set the restart priority policy for each NSX Manager appliance to high.	■ NSX Manager implements the control plane for virtual network segments. vSphere HA restarts the NSX Manager appliances first so that other virtual machines that are being powered on or migrated by using vSphere vMotion while the control plane is offline lose connectivity only until the control plane quorum is re-established. ■ Setting the restart priority to high reserves the highest priority for flexibility for adding services that must be started before NSX Manager.	If the restart priority for another management appliance is set to highest, the connectivity delay for management appliances will be longer.

Table 8-6. NSX Manager Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-LM-RCMD-CFG-006	Add the NSX Manager appliances to the virtual machine group for the first availability zone.	Ensures that, by default, the NSX Manager appliances are powered on a host in the primary availability zone.	None.

NSX Global Manager Design Requirements and Recommendations for VMware Cloud Foundation

For a deployment with multiple VMware Cloud Foundation instances, you use NSX Federation, which requires the manual deployment of NSX Global Manager nodes in the first two instances. Consider the placement requirements for using NSX Global Manager in VMware Cloud Foundation, and the best practices for having an NSX Global Manager cluster operate in an optimal way, such as the number and size of the nodes, high availability, on a standard or stretched management cluster.

NSX Global Manager Design Requirements

You must meet the following design requirements in your NSX Global Manager design for VMware Cloud Foundation.

Table 8-7. NSX Global Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-GM-REQD-CFG-001	Place the appliances of the NSX Global Manager cluster on the Management VM network in each VMware Cloud Foundation instance.	■ Simplifies IP addressing for management VMs. ■ Provides simplified secure access to all management VMs in the same VLAN network.	None.

NSX Global Manager Design Recommendations

In your NSX Global Manager design for VMware Cloud Foundation, you can apply certain best practices for standard and stretched clusters.

Table 8-8. NSX Global Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-001	Deploy three NSX Global Manager nodes for the workload domain to support NSX Federation across VMware Cloud Foundation instances.	Provides high availability for the NSX Global Manager cluster.	You must have sufficient resources in the default cluster of the management domain to run three NSX Global Manager nodes.
VCF-NSX-GM-RCMD-CFG-002	Deploy appropriately sized nodes in the NSX Global Manager cluster for the workload domain.	Ensures resource availability and usage efficiency per workload domain.	The recommended size for a management domain is Medium and for VI workload domains is Large.
VCF-NSX-GM-RCMD-CFG-003	Create a virtual IP (VIP) address for the NSX Global Manager cluster for the workload domain.	Provides high availability of the user interface and API of NSX Global Manager.	■ The VIP address feature provides high availability only. It does not load-balance requests across the cluster. ■ When using the VIP address feature, all NSX Global Manager nodes must be deployed on the same Layer 2 network.

Table 8-8. NSX Global Manager Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-004	Apply VM-VM anti-affinity rules in vSphere DRS to the NSX Global Manager appliances.	Keeps the NSX Global Manager appliances running on different ESXi hosts for high availability.	You must allocate at least four physical hosts so that the three NSX Manager appliances continue running if an ESXi host failure occurs.
VCF-NSX-GM-RCMD-CFG-005	In vSphere HA, set the restart priority policy for each NSX Global Manager appliance to medium.	■ NSX Global Manager implements the management plane for global segments and firewalls. NSX Global Manager is not required for control plane and data plane connectivity. ■ Setting the restart priority to medium reserves the high priority for services that impact the NSX control or data planes.	■ Management of NSX global components will be unavailable until the NSX Global Manager virtual machines restart. ■ The NSX Global Manager cluster is deployed in the management domain, where the total number of virtual machines is limited and where it competes with other management components for restart priority.
VCF-NSX-GM-RCMD-CFG-006	Deploy an additional NSX Global Manager Cluster in the second VMware Cloud Foundation instance.	Enables recoverability of NSX Global Manager in the second VMware Cloud Foundation instance if a failure in the first VMware Cloud Foundation instance occurs.	Requires additional NSX Global Manager nodes in the second VMware Cloud Foundation instance.

Table 8-8. NSX Global Manager Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-007	Set the NSX Global Manager cluster in the second VMware Cloud Foundation instance as standby for the workload domain.	Enables recoverability of NSX Global Manager in the second VMware Cloud Foundation instance if a failure in the first instance occurs.	Must be done manually.
VCF-NSX-GM-RCMD-SEC-001	Establish an operational practice to capture and update the thumbprint of the NSX Local Manager certificate on NSX Global Manager every time the certificate is updated by using SDDC Manager.	Ensures secured connectivity between the NSX Manager instances. Each certificate has its own unique thumbprint. NSX Global Manager stores the unique thumbprint of the NSX Local Manager instances for enhanced security. If an authentication failure between NSX Global Manager and NSX Local Manager occurs, objects that are created from NSX Global Manager will not be propagated on to the SDN.	The administrator must establish and follow an operational practice by using a runbook or automated process to ensure that the thumbprint is up-to-date.

Table 8-9. NSX Global Manager Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-008	Add the NSX Global Manager appliances to the virtual machine group for the first availability zone.	Ensures that, by default, the NSX Global Manager appliances are powered on a host in the primary availability zone.	Done automatically by VMware Cloud Foundation when stretching a cluster.

NSX Edge Node Design for VMware Cloud Foundation

Following the principles of this design and of each product, you deploy, configure, and connect the NSX Edge nodes to support networks in the NSX instances in your VMware Cloud Foundation deployment.

Deployment Model for the NSX Edge Nodes for VMware Cloud Foundation

For NSX Edge nodes, you determine the form factor, number of nodes and placement according to the requirements for network services in a VMware Cloud Foundation workload domain.

An NSX Edge node is an appliance that provides centralized networking services which cannot be distributed to hypervisors, such as load balancing, NAT, VPN, and physical network uplinks. Some services, such as Tier-0 gateways, are limited to a single instance per NSX Edge node. However, most services can coexist in these nodes.

NSX Edge nodes are grouped in one or more edge clusters, representing a pool of capacity for NSX services.

An NSX Edge node can be deployed as a virtual appliance, or installed on bare-metal hardware. The edge node on bare-metal hardware can have better performance capabilities at the expense of more difficult deployment and limited deployment topology use cases. For details on the trade-offs of using virtual or bare-metal NSX Edges, see the *NSX documentation*.

Table 8-10. NSX Edge Deployment Model Considerations

Deployment Model	Benefits	Drawbacks
NSX Edge virtual appliance deployed by using SDDC Manager	■ Deployment and life cycle management by using SDDC Manager workflows that call NSX Manager ■ Automated password management by using SDDC Manager ■ Benefits from vSphere HA recovery ■ Can be used across availability zones ■ Easy to scale up by modifying the specification of the virtual appliance	■ Might not provide best performance in individual customer scenarios
NSX Edge virtual appliance deployed by using NSX Manager	■ Benefits from vSphere HA recovery ■ Can be used across availability zones ■ Easy to scale up by modifying the specification of the virtual appliance	■ Might not provide best performance in individual customer scenarios ■ Manually deployed by using NSX Manager ■ Manual password Management by using NSX Manager ■ Cannot be used to support Application Virtual Networks (AVNs) in the management domain
Bare-metal NSX Edge appliance	■ Might provide better performance in individual customer scenarios	■ Has hardware compatibility requirements ■ Requires individual hardware life cycle management and monitoring of failures, firmware and drivers ■ Manual password management ■ Must be manually deployed and connected to the environment ■ Requires manual recovery after hardware failure ■ Requires deploying a bare-metal NSX Edge appliance in each availability zone for network failover ■ Deploying a bare-metal edge in each availability zone requires considering asymmetric routing ■ Requires edge fault domains if more than one edge is deployed in each availability zone for Active/Standby Tier-0 or Tier-1 gateways

Table 8-10. NSX Edge Deployment Model Considerations (continued)

Deployment Model	Benefits	Drawbacks
		■ Requires redeployment to new host to achieve scale-up ■ Cannot be used to support AVNs in the management domain

Sizing Considerations for NSX Edges for VMware Cloud Foundation

When you deploy NSX Edge appliances, you select a size according to the scale of your environment. The option that you select determines the number of CPUs and the amount of memory of the appliance.

For detailed sizing according to the overall profile of the VMware Cloud Foundation instance you plan to deploy, see VMware Cloud Foundation Planning and Preparation Workbook.

Table 8-11. Sizing Considerations for NSX Edges

NSX Edge Appliance Size	Scale
Small	Proof of concept
Medium	Suitable when only Layer 2 through Layer 4 features such as NAT, routing, Layer 4 firewall, Layer 4 load balancer are required and the total throughput requirement is less than 2 Gbps.
Large	Suitable when only Layer 2 through Layer 4 features such as NAT, routing, Layer 4 firewall, Layer 4 load balancer are required and the total throughput is 2 ~ 10 Gbps. It is also suitable when Layer 7 load balancer, for example, SSL offload is required.
Extra Large	Suitable when the total throughput required is multiple Gbps for Layer 7 load balancer and VPN.

Network Design for the NSX Edge Nodes for VMware Cloud Foundation

In each VMware Cloud Foundation instance, you implement an NSX Edge configuration with a single N-VDS. You connect the uplink network interfaces of the edge appliance to VLAN trunk port groups that are connected to particular physical NICs on the host.

NSX Edge Network Configuration

The NSX Edge node contains a virtual switch, called an N-VDS, that is managed by NSX. This internal N-VDS is used to define traffic flow through the interfaces of the edge node. An N-VDS can be connected to one or more interfaces. Interfaces cannot be shared between N-VDS instances.

If you plan to deploy multiple VMware Cloud Foundation instances, apply the same network design to the NSX Edge cluster in the second and other additional VMware Cloud Foundation instances.

Figure 8-5. NSX Edge Network Configuration

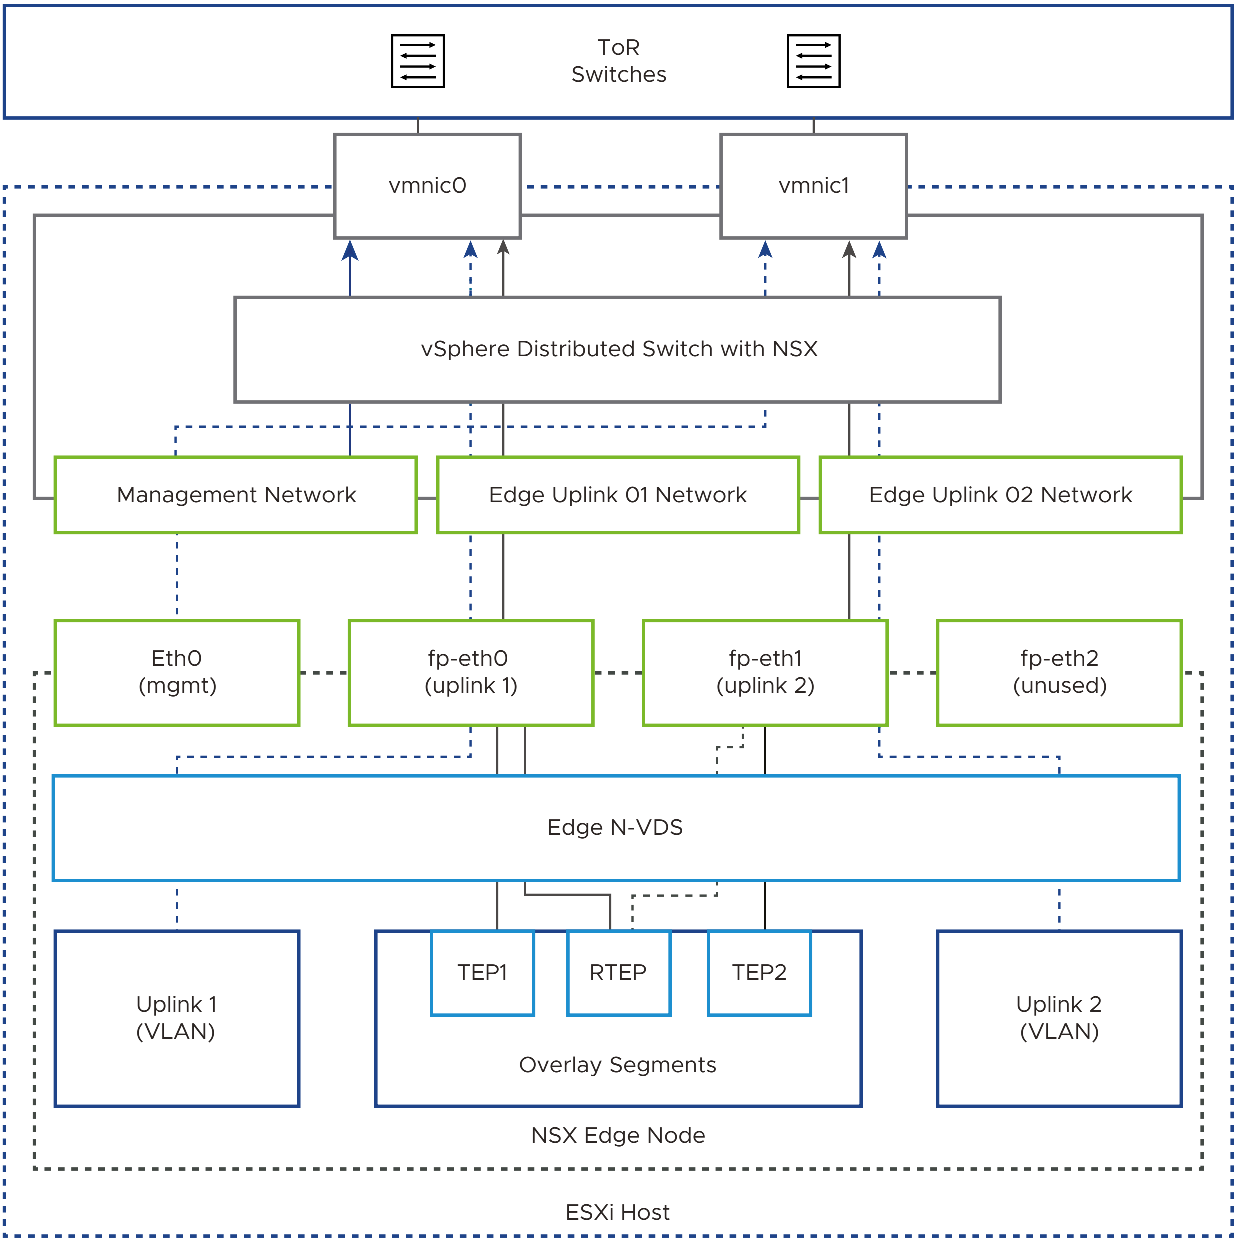

Uplink Policy Design for the NSX Edge Nodes for VMware Cloud Foundation

A transport node can participate in an overlay and VLAN network. Uplink profiles define policies for the links from the NSX Edge transport nodes to top of rack switches. Uplink profiles are containers for the properties or capabilities for the network adapters. Uplink profiles are applied to the N-VDS of the edge node.

Uplink profiles can use either load balance source or failover order teaming. If using load balance source, multiple uplinks can be active. If using failover order, only a single uplink can be active.

Teaming can be configured by using the default teaming policy or a user-defined named teaming policy. You can use named teaming policies to pin traffic segments to designated edge uplinks.

NSX Edge Node Requirements and Recommendations for VMware Cloud Foundation

Consider the network, N-VDS configuration and uplink policy requirements for using NSX Edge nodes in VMware Cloud Foundation, and the best practices for having NSX Edge nodes operate in an optimal way, such as number and size of the nodes, high availability, and N-VDS architecture, on a standard or stretched cluster.

NSX Edge Design Requirements

You must meet the following design requirements for standard and stretched clusters in your NSX Edge design for VMware Cloud Foundation.

Table 8-12. NSX Edge Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-001	Connect the management interface of each NSX Edge node to the VM management network.	Provides connection from the NSX Manager cluster to the NSX Edge.	None.
VCF-NSX-EDGE-REQD-CFG-002	■ Connect the `fp-eth0` interface of each NSX Edge appliance to a VLAN trunk port group pinned to physical NIC 0 of the host, with the ability to failover to physical NIC 1. ■ Connect the `fp-eth1` interface of each NSX Edge appliance to a VLAN trunk port group pinned to physical NIC 1 of the host, with the ability to failover to physical NIC 0. ■ Leave the `fp-eth2` interface of each NSX Edge appliance unused.	■ Because VLAN trunk port groups pass traffic for all VLANs, VLAN tagging can occur in the NSX Edge node itself for easy post-deployment configuration. ■ By using two separate VLAN trunk port groups, you can direct traffic from the edge node to a particular host network interface and top of rack switch as needed. ■ In the event of failure of the top of rack switch, the VLAN trunk port group will failover to the other physical NIC and to ensure both `fp-eth0` and `fp-eth1` are available.	None.

Table 8-12. NSX Edge Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-003	Use a dedicated VLAN for edge overlay that is different from the host overlay VLAN.	A dedicated edge overlay network provides support for edge mobility in support of advanced deployments such as multiple availability zones or multi-rack clusters.	■ You must have routing between the VLANs for edge overlay and host overlay. ■ You must allocate another VLAN in the data center infrastructure for edge overlay.
VCF-NSX-EDGE-REQD-CFG-004	Create one uplink profile for the edge nodes with three teaming policies. ■ Default teaming policy of load balance source with both active uplinks `uplink1` and `uplink2`. ■ Named teaming policy of failover order with a single active uplink `uplink1` without standby uplinks. ■ Named teaming policy of failover order with a single active uplink `uplink2` without standby uplinks.	■ An NSX Edge node that uses a single N-VDS can have only one uplink profile. ■ For increased resiliency and performance, supports the concurrent use of both edge uplinks through both physical NICs on the ESXi hosts. ■ The default teaming policy increases overlay performance and availability by using multiple TEPs, and balancing of overlay traffic. ■ By using named teaming policies, you can connect an edge uplink to a specific host uplink and from there to a specific top of rack switch in the data center. ■ Enables ECMP because the NSX Edge nodes can uplink to the physical network over two different VLANs.	None.

Table 8-13. NSX Edge Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-005	Allocate a separate VLAN for edge RTEP overlay that is different from the edge overlay VLAN.	The RTEP network must be on a VLAN that is different from the edge overlay VLAN. This is an NSX requirement that provides support for configuring different MTU size per network.	You must allocate another VLAN in the data center infrastructure.

NSX Edge Design Recommendations

In your NSX Edge design for VMware Cloud Foundation, you can apply certain best practices for standard and stretched clusters.

Table 8-14. NSX Edge Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-001	Use appropriately sized NSX Edge virtual appliances.	Ensures resource availability and usage efficiency per workload domain.	You must provide sufficient compute resources to support the chosen appliance size.
VCF-NSX-EDGE-RCMD-CFG-002	Deploy the NSX Edge virtual appliances to the default vSphere cluster of the workload domain, sharing the cluster between the workloads and the edge appliances.	Simplifies the configuration and minimizes the number of ESXi hosts required for initial deployment.	Workloads and NSX Edges share the same compute resources.
VCF-NSX-EDGE-RCMD-CFG-003	Deploy two NSX Edge appliances in an edge cluster in the default vSphere cluster of the workload domain.	Creates the minimum size NSX Edge cluster while satisfying the requirements for availability.	For a VI workload domain, additional edge appliances might be required to satisfy increased bandwidth requirements.
VCF-NSX-EDGE-RCMD-CFG-004	Apply VM-VM anti-affinity rules for vSphere DRS to the virtual machines of the NSX Edge cluster.	Keeps the NSX Edge nodes running on different ESXi hosts for high availability.	None.

Table 8-14. NSX Edge Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-005	In vSphere HA, set the restart priority policy for each NSX Edge appliance to high.	■ The NSX Edge nodes are part of the north-south data path for overlay segments. vSphere HA restarts the NSX Edge appliances first to minimise the time an edge VM is offline. ■ Setting the restart priority to high reserves highest for future needs.	If the restart priority for another VM in the cluster is set to highest, the connectivity delays for edge appliances will be longer.
VCF-NSX-EDGE-RCMD-CFG-006	Create an NSX Edge cluster with the default Bidirectional Forwarding Detection (BFD) configuration between the NSX Edge nodes in the cluster.	■ Satisfies the availability requirements by default. ■ Edge nodes must remain available to create services such as NAT, routing to physical networks, and load balancing.	None.
VCF-NSX-EDGE-RCMD-CFG-007	Use a single N-VDS in the NSX Edge nodes.	■ Simplifies deployment of the edge nodes. ■ The same N-VDS switch design can be used regardless of edge form factor. ■ Supports multiple TEP interfaces in the edge node. ■ vSphere Distributed Switch is not supported in the edge node.	None.

Table 8-15. NSX Edge Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-008	Add the NSX Edge appliances to the virtual machine group for the first availability zone.	Ensures that, by default, the NSX Edge appliances are powered on upon a host in the primary availability zone.	None.

Routing Design for VMware Cloud Foundation

NSX Edge clusters in VMware Cloud Foundation provide pools of capacity for service router functions in NSX.

Routing Options for VMware Cloud Foundation

VMware Cloud Foundation supports the following routing options:

Routing Type	Description	Benefits	Drawbacks
Static routing	■ The administrator manages the routing information, adding routing information to the routing table. ■ If any change occurs in the network, the administrator has to update the related information in the routing table.	■ No dynamic routing protocol required on the ToR switches. ■ In some cases, no additional license needed on the ToR switches to implement dynamic routing.	■ You must manually create static routes in NSX Manager on the Tier-0 gateway. ■ If required, you must manually create an HA VIP in the NSX Manager on the Tier-0 gateway to provide redundancy across ToR switches. ■ Not supported with vSAN stretched clusters.
OSPF	■ The routing protocol automatically adds and manages the routing information in the routing table. If any change occurs in the network, the routing protocol automatically updates the related information in the routing table. If any new segments or subnets are added in NSX, they are automatically added to the routing table.	■ If the physical fabric is running OSPF routing protocol, using OSPF at the virtual layer might be a simpler approach for the network administrator.	■ Needs additional manual configuration. See VMware Knowledge Base article 85916. ■ Not supported with vSAN stretched clusters. ■ Not supported with NSX Federation. ■ Combined use of BGP and OSPF on a single Tier-0 gateway not supported.
BGP	■ BGP is known as an exterior gateway protocol. It is designed to share routing information between disparate networks, known as autonomous systems (ASes). ■ When multiple BGP-derived paths exist, the protocol chooses a path to send traffic based on certain criteria. ■ The routing protocol automatically adds and manages the routing information in the routing table. If any new segments or	■ Fully supported by the automated edge workflows in VMware Cloud Foundation. ■ Fully supported for all VMware Cloud Foundation topologies.	■ None.

Routing Type	Description	Benefits	Drawbacks
	subnets are added in NSX, they are automatically added to the routing table.		

BGP routing is the routing option recommended for VMware Cloud Foundation.

BGP Routing Design for VMware Cloud Foundation

Determine the number, networking, and high-availability configuration of the Tier-0 and Tier-1 gateways in NSX for VMware Cloud Foundation workload domains. Identify the BGP configuration for a single availability zone and two availability zones in the environment.

Table 8-16. Routing Direction Definitions

Routing Direction	Description
North-south	Traffic leaving or entering the NSX domain, for example, a virtual machine on an overlay network communicating with an end-user device on the corporate network.
East-west	Traffic that remains in the NSX domain, for example, two virtual machines on the same or different segments communicating with each other.

North-South Routing

The routing design considers different levels of routing in the environment, such as number and type of gateways in NSX, dynamic routing protocol, and others.

The following models for north-south traffic exist:

Table 8-17. Considerations for the Operating Model for North-South Service Routers

North-South Service Router Operating Model	Description	Benefits	Drawbacks
Active-Active	■ Bandwidth independent of the Tier-0 gateway failover model. ■ Configured in active-active equal-cost multi-path (ECMP) mode. ■ Failover takes approximately 2 seconds for virtual edges and is sub-second for bare-metal edges.	■ The active-active mode can support up to 8 NSX Edge nodes per northbound service router (SR). ■ Availability can be as high as N+7, with up to 8 active-active NSX Edge nodes. ■ Supports ECMP north-south routing on all nodes in the NSX Edge cluster.	■ Cannot provide some stateful services, such as SNAT or DNAT.
Active-Standby	■ Bandwidth independent of the Tier-0 gateway failover model. ■ Failover takes approximately 2 seconds for virtual edges and is sub-second for bare-metal edges.	■ Can provide stateful services such as NAT.	■ The active-standby mode is limited to a single node. ■ Availability limited to N+1.

BGP North-South Routing for a Single or Multiple Availability Zones

For multiple availability zones, plan for failover of the NSX Edge nodes by configuring BGP so that traffic from the top of rack switches is directed to the first availability zone unless a failure in this zone occurs.

Figure 8-6. BGP North-South Routing for VMware Cloud Foundation Instances with a Single Availability Zone

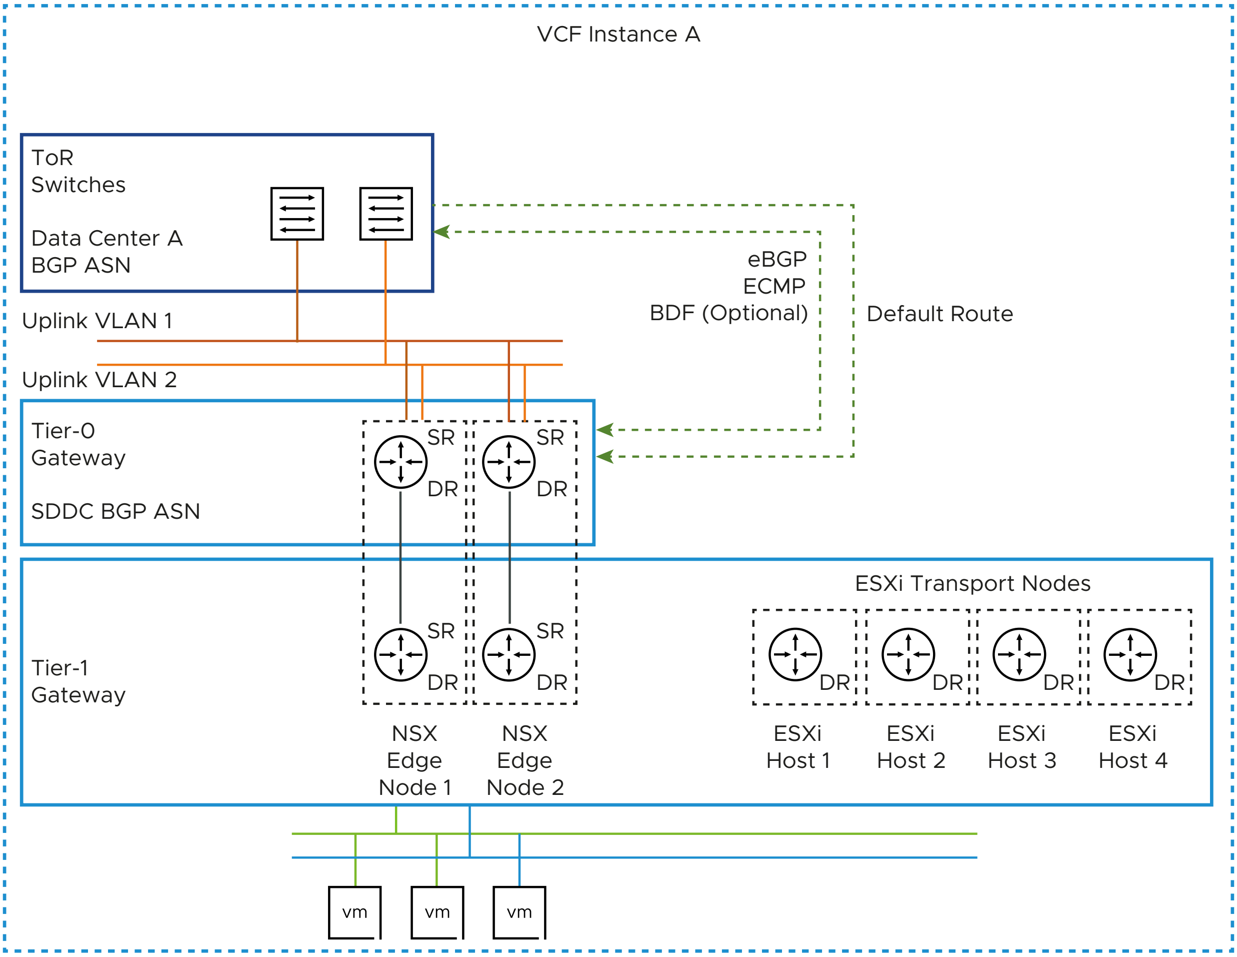

Figure 8-7. BGP North-South Routing for VMware Cloud Foundation Instances with Multiple Availability Zones

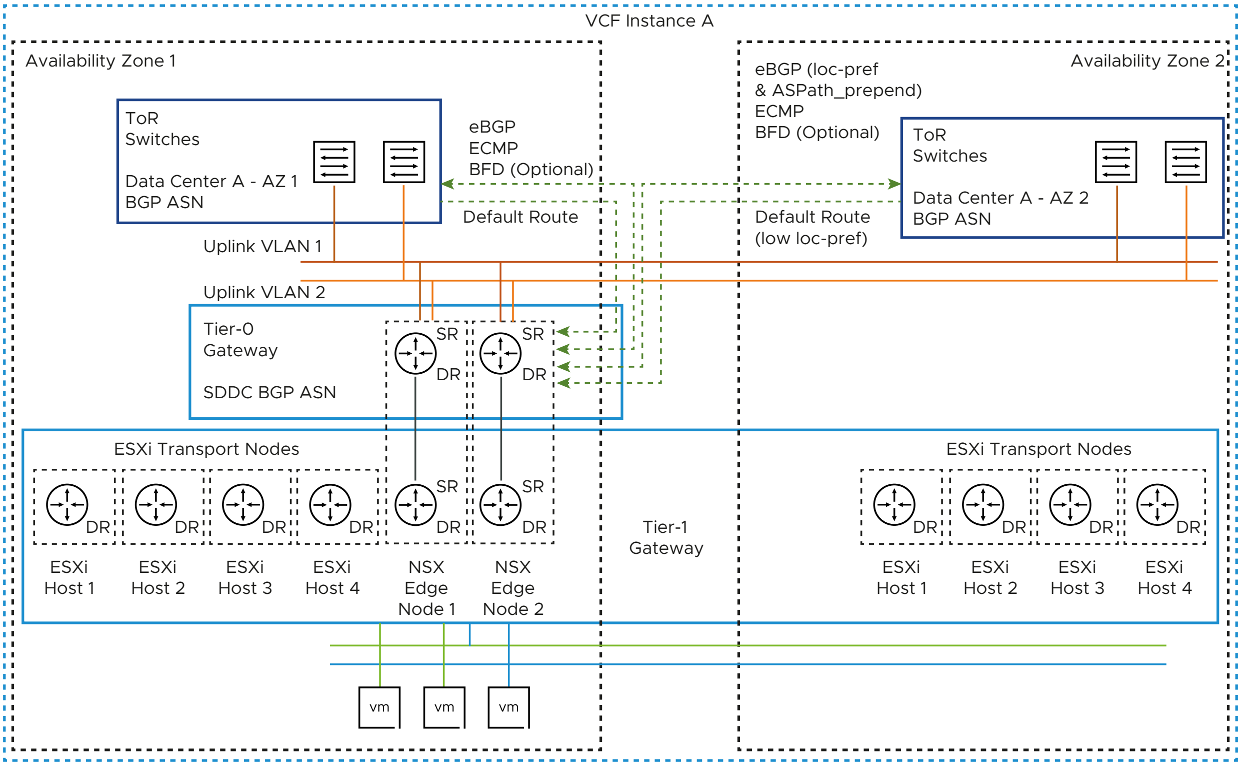

BGP North-South Routing Design for NSX Federation

In a routing design for an environment with VMware Cloud Foundation instances that use NSX Federation, you identify the instances that an SDN network must span and at which physical location ingress and egress traffic should occur.

Local egress allows traffic to leave any location which the network spans. The use of local-egress would require controlling local-ingress to prevent asymmetrical routing. This design does not use local-egress. Instead, this design uses a preferred and failover VMware Cloud Foundation instances for all networks.

Figure 8-8. BGP North-South Routing for VMware Cloud Foundation Instances with NSX Federation

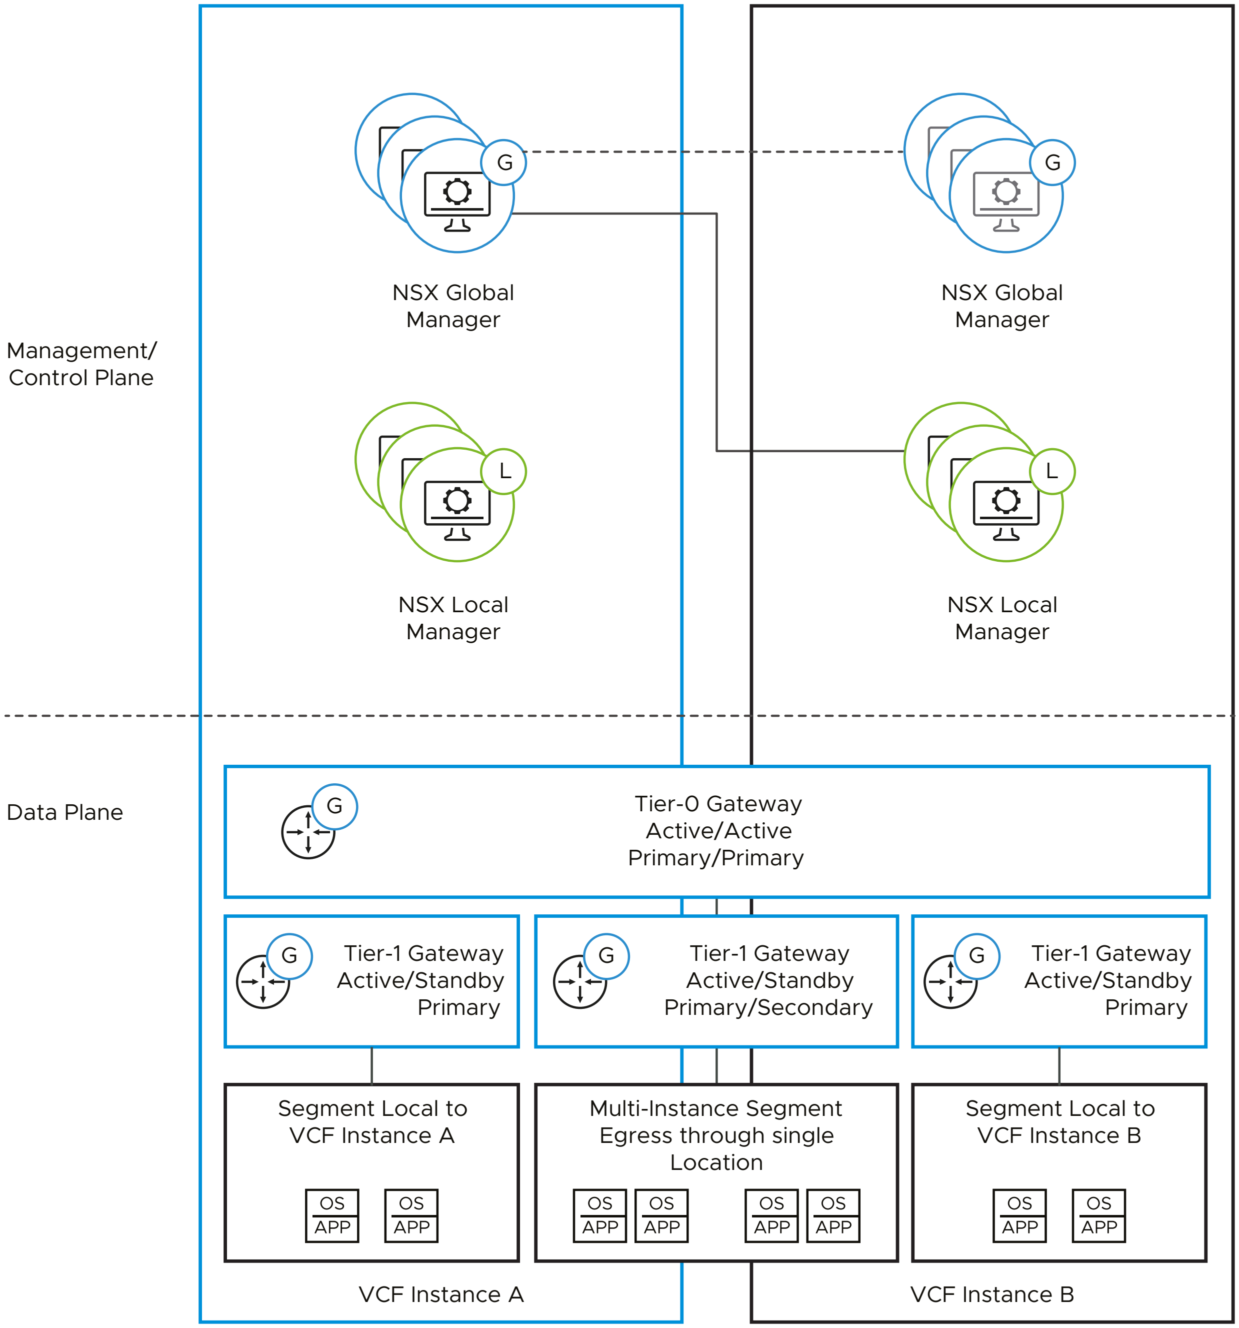

Tier-0 Gateways with NSX Federation

In NSX Federation, a Tier-0 gateway can span multiple VMware Cloud Foundation instances.

Each VMware Cloud Foundation instance that is in the scope of a Tier-0 gateway can be configured as primary or secondary. A primary instance passes traffic for any other SDN service such as Tier-0 logical segments or Tier-1 gateways. A secondary instance routes traffic locally but does not egress traffic outside the SDN or advertise networks in the data center.

When deploying an additional VMware Cloud Foundation instance, the Tier-0 gateway in the first instance is extended to the new instance.

In this design, the Tier-0 gateway in each VMware Cloud Foundation instance is configured as primary. Although the Tier-0 gateway technically supports local-egress, the design does not recommend the use of local-egress. Ingress and egress traffic is controlled at the Tier-1 gateway level.

Each VMware Cloud Foundation instance has its own NSX Edge cluster with associated uplink VLANs for north-south traffic flow for that instance. The Tier-0 gateway in each instance peers with the top of rack switches over eBGP.

Figure 8-9. BGP Peering to Top of Rack Switches for VMware Cloud Foundation Instances with NSX Federation

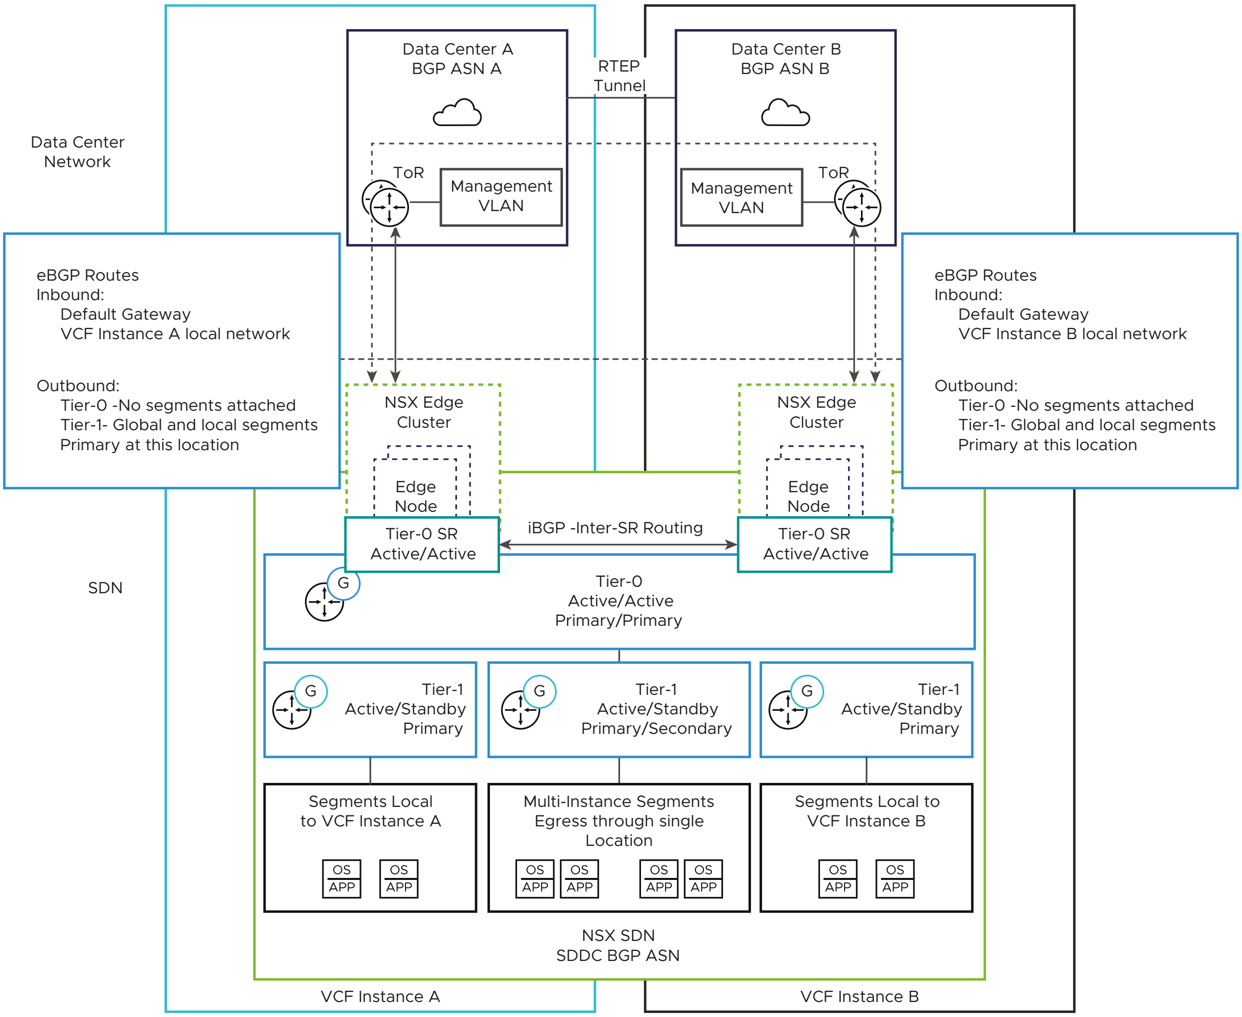

Tier-1 Gateways with NSX Federation

A Tier-1 gateway can span several VMware Cloud Foundation instances. As with a Tier-0 gateway, you can configure an instance's location as primary or secondary for the Tier-1 gateway. The gateway then passes ingress and egress traffic for the logical segments connected to it.

Any logical segments connected to the Tier-1 gateway follow the span of the Tier-1 gateway. If the Tier-1 gateway spans several VMware Cloud Foundation instances, any segments connected to that gateway become available in both instances.

Using a Tier-1 gateway enables more granular control on logical segments in the first and second VVMware Cloud Foundation instances. You use three Tier-1 gateways - one in each VMware Cloud Foundation instance for segments that are local to the instance, and one for segments which span the two instances.

Table 8-18. Location Configuration of the Tier-1 Gateways for Multiple VMware Cloud Foundation Instances

Tier-1 Gateway	First VMware Cloud Foundation Instance	Second VMware Cloud Foundation Instance	Ingress and Egress Traffic
Connected to both VMware Cloud Foundation instances	Primary	Secondary	First VMware Cloud Foundation instance Second VMware Cloud Foundation instance
Local to the firstVMware Cloud Foundation instance	Primary	-	First VMware Cloud Foundation instance only
Local to the second VMware Cloud Foundation instance	-	Primary	Second VMware Cloud Foundation instance only

The Tier-1 gateway advertises its networks to the connected local-instance unit of the Tier-0 gateway. In the case of primary-secondary location configuration, the Tier-1 gateway advertises its networks only to the Tier-0 gateway unit in the location where the Tier-1 gateway is primary. The Tier-0 gateway unit then re-advertises those networks to the data center in the sites where that Tier-1 gateway is primary. During failover of the components in the first VMware Cloud Foundation instance, an administrator must manually set the Tier-1 gateway in the second VMware Cloud Foundation instance as primary. Then, networks become advertised through the Tier-1 gateway unit in the second instance.

In a Multiple Instance-Multiple Availability Zone topology, the same Tier-0 and Tier-1 gateway architecture applies. The ESXi transport nodes from the second availability zone are also attached to the Tier-1 gateway as per the Figure 8-7. BGP North-South Routing for VMware Cloud Foundation Instances with Multiple Availability Zones design.

BGP Routing Design Requirements and Recommendations for VMware Cloud Foundation

Consider the requirements for the configuration of Tier-0 and Tier-1 gateways for implementing BGP routing in VMware Cloud Foundation, and the best practices for having optimal traffic routing on a standard or stretched cluster in a environment with a single or multiple VMware Cloud Foundation instances.

BGP Routing

The BGP routing design has the following characteristics:

- Enables dynamic routing by using NSX.
- Offers increased scale and flexibility.
- Is a proven protocol that is designed for peering between networks under independent administrative control - data center networks and the NSX SDN.

Note These design recommendations do not include BFD. However, if faster convergence than BGP timers is required, you must enable BFD on the physical network and also on the NSX Tier-0 gateway.

BGP Routing Design Requirements

You must meet the following design requirements for standard and stretched clusters in your routing design for a single VMware Cloud Foundation instance. For NSX Federation, additional requirements exist.

Table 8-19. BGP Routing Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-001	To enable ECMP between the Tier-0 gateway and the Layer 3 devices (ToR switches or upstream devices), create two VLANs. The ToR switches or upstream Layer 3 devices have an SVI on one of the two VLANS, and each Edge node in the cluster has an interface on each VLAN.	Supports multiple equal-cost routes on the Tier-0 gateway and provides more resiliency and better bandwidth use in the network.	Additional VLANs are required.
VCF-NSX-BGP-REQD-CFG-002	Assign a named teaming policy to the VLAN segments to the Layer 3 device pair.	Pins the VLAN traffic on each segment to its target edge node interface. From there, the traffic is directed to the host physical NIC that is connected to the target top of rack switch.	None.

Table 8-19. BGP Routing Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-003	Create a VLAN transport zone for edge uplink traffic.	Enables the configuration of VLAN segments on the N-VDS in the edge nodes.	Additional VLAN transport zones might be required if the edge nodes are not connected to the same top of rack switch pair.
VCF-NSX-BGP-REQD-CFG-004	Deploy a Tier-1 gateway and connect it to the Tier-0 gateway.	Creates a two-tier routing architecture. Abstracts the NSX logical components which interact with the physical data center from the logical components which provide SDN services.	A Tier-1 gateway can only be connected to a single Tier-0 gateway. In cases where multiple Tier-0 gateways are required, you must create multiple Tier-1 gateways.
VCF-NSX-BGP-REQD-CFG-005	Deploy a Tier-1 gateway to the NSX Edge cluster.	Enables stateful services, such as load balancers and NAT, for SDDC management components. Because a Tier-1 gateway always works in active-standby mode, the gateway supports stateful services.	None.

Table 8-20. BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-006	Extend the uplink VLANs to the top of rack switches so that the VLANs are stretched between both availability zones.	Because the NSX Edge nodes will fail over between the availability zones, ensures uplink connectivity to the top of rack switches in both availability zones regardless of the zone the NSX Edge nodes are presently in.	You must configure a stretched Layer 2 network between the availability zones by using physical network infrastructure.
VCF-NSX-BGP-REQD-CFG-007	Provide this SVI configuration on the top of the rack switches. ■ In the second availability zone, configure the top of rack switches or upstream Layer 3 devices with an SVI on each of the two uplink VLANs. ■ Make the top of rack switch SVI in both availability zones part of a common stretched Layer 2 network between the availability zones.	Enables the communication of the NSX Edge nodes to the top of rack switches in both availability zones over the same uplink VLANs.	You must configure a stretched Layer 2 network between the availability zones by using the physical network infrastructure.
VCF-NSX-BGP-REQD-CFG-008	Provide this VLAN configuration: ■ Use two VLANs to enable ECMP between the Tier-0 gateway and the Layer 3 devices (top of rack switches or Leaf switches). ■ The ToR switches or upstream Layer 3 devices have an SVI to one of the two VLANS and each NSX Edge node has an interface to each VLAN.	Supports multiple equal-cost routes on the Tier-0 gateway, and provides more resiliency and better bandwidth use in the network.	■ Extra VLANs are required. ■ Requires stretching uplink VLANs between availability zones

Table 8-20. BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-009	Create an IP prefix list that permits access to route advertisement by `any` network instead of using the default IP prefix list.	Used in a route map to prepend a path to one or more autonomous system (AS-path prepend) for BGP neighbors in the second availability zone.	You must manually create an IP prefix list that is identical to the default one.
VCF-NSX-BGP-REQD-CFG-010	Create a route map-out that contains the custom IP prefix list and an AS-path prepend value set to the Tier-0 local AS added twice.	■ Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. ■ Ensures that all ingress traffic passes through the first availability zone.	You must manually create the route map. The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.
VCF-NSX-BGP-REQD-CFG-011	Create an IP prefix list that permits access to route advertisement by network `0.0.0.0/0` instead of using the default IP prefix list.	Used in a route map to configure local-reference on learned default-route for BGP neighbors in the second availability zone.	You must manually create an IP prefix list that is identical to the default one.

Table 8-20. BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-012	Apply a route map-in that contains the IP prefix list for the default route `0.0.0.0/0` and assign a lower local-preference , for example, `80`, to the learned default route and a lower local-preference, for example, `90` any routes learned.	■ Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. ■ Ensures that all egress traffic passes through the first availability zone.	You must manually create the route map. The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.
VCF-NSX-BGP-REQD-CFG-013	Configure the neighbors of the second availability zone to use the route maps as In and Out filters respectively.	Makes the path in and out of the second availability zone less preferred because the AS path is longer and the local preference is lower. As a result, all traffic passes through the first zone.	The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.

Table 8-21. BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-014	Extend the Tier-0 gateway to the second VMware Cloud Foundation instance.	■ Supports ECMP north-south routing on all nodes in the NSX Edge cluster. ■ Enables support for cross-instance Tier-1 gateways and cross-instance network segments.	The Tier-0 gateway deployed in the second instance is removed.
VCF-NSX-BGP-REQD-CFG-015	Set the Tier-0 gateway as primary for all VMware Cloud Foundation instances.	■ In NSX Federation, a Tier-0 gateway lets egress traffic from connected Tier-1 gateways only in its primary locations. ■ Local ingress and egress traffic is controlled independently at the Tier-1 level. No segments are provisioned directly to the Tier-0 gateway. ■ A mixture of network spans (local to a VMware Cloud Foundation instance or spanning multiple instances) is enabled without requiring additional Tier-0 gateways and hence edge nodes. ■ If a failure in a VMware Cloud Foundation instance occurs, the local-instance networking in the other instance remains available without manual intervention.	None.

Table 8-21. BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-016	From the global Tier-0 gateway, establish BGP neighbor peering to the ToR switches connected to the second VMware Cloud Foundation instance.	■ Enables the learning and advertising of routes in the second VMware Cloud Foundation instance. ■ Facilitates a potential automated failover of networks from the first to the second VMware Cloud Foundation instance.	None.
VCF-NSX-BGP-REQD-CFG-017	Use a stretched Tier-1 gateway and connect it to the Tier-0 gateway for cross-instance networking.	■ Enables network span between the VMware Cloud Foundation instances because NSX network segments follow the span of the gateway they are attached to. ■ Creates a two-tier routing architecture.	None.

Table 8-21. BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-018	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the stretched Tier-1 gateway. Set the first VMware Cloud Foundation instance as primary and the second instance as secondary.	■ Enables cross-instance network span between the first and second VMware Cloud Foundation instances. ■ Enables deterministic ingress and egress traffic for the cross-instance network. ■ If a VMware Cloud Foundation instance failure occurs, enables deterministic failover of the Tier-1 traffic flow. ■ During the recovery of the inaccessible VMware Cloud Foundation instance, enables deterministic failback of the Tier-1 traffic flow, preventing unintended asymmetrical routing. ■ Eliminates the need to use BGP attributes in the first and second VMware Cloud Foundation instances to influence location preference and failover.	You must manually fail over and fail back the cross-instance network from the standby NSX Global Manager.
VCF-NSX-BGP-REQD-CFG-019	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the local Tier-1 gateway for that VMware Cloud Foundation instance.	■ Enables instance-specific networks to be isolated to their specific instances. ■ Enables deterministic flow of ingress and egress traffic for the instance-specific networks.	You can use the service router that is created for the Tier-1 gateway for networking services. However, such configuration is not required for network connectivity.
VCF-NSX-BGP-REQD-CFG-020	Set each local Tier-1 gateway only as primary in that instance. Avoid setting the gateway as secondary in the other instances.	Prevents the need to use BGP attributes in primary and secondary instances to influence the instance ingress-egress preference.	None.

BGP Routing Design Recommendations

In your routing design for a single VMware Cloud Foundation instance, you can apply certain best practices for standard and stretched clusters. For NSX Federation, additional recommendations are available.

Table 8-22. BGP Routing Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
VCF-NSX-BGP-RCMD-CFG-001	Deploy an active-active Tier-0 gateway.	Supports ECMP north-south routing on all Edge nodes in the NSX Edge cluster.	Active-active Tier-0 gateways cannot provide stateful services such as NAT.
VCF-NSX-BGP-RCMD-CFG-002	Configure the BGP Keep Alive Timer to 4 and Hold Down Timer to 12 or lower between the top of tack switches and the Tier-0 gateway.	Provides a balance between failure detection between the top of rack switches and the Tier-0 gateway, and overburdening the top of rack switches with keep-alive traffic.	By using longer timers to detect if a router is not responding, the data about such a router remains in the routing table longer. As a result, the active router continues to send traffic to a router that is down. These timers must be aligned with the data center fabric design of your organization.
VCF-NSX-BGP-RCMD-CFG-003	Do not enable Graceful Restart between BGP neighbors.	Avoids loss of traffic. On the Tier-0 gateway, BGP peers from all the gateways are always active. On a failover, the Graceful Restart capability increases the time a remote neighbor takes to select an alternate Tier-0 gateway. As a result, BFD-based convergence is delayed.	None.
VCF-NSX-BGP-RCMD-CFG-004	Enable helper mode for Graceful Restart mode between BGP neighbors.	Avoids loss of traffic. During a router restart, helper mode works with the graceful restart capability of upstream routers to maintain the forwarding table which in turn will forward packets to a down neighbor even after the BGP timers have expired causing loss of traffic.	None.

Table 8-22. BGP Routing Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
VCF-NSX-BGP-RCMD-CFG-005	Enable Inter-SR iBGP routing.	In the event that an edge node has all of its northbound eBGP sessions down, north-south traffic will continue to flow by routing traffic to a different edge node.	None.
VCF-NSX-BGP-RCMD-CFG-006	Deploy a Tier-1 gateway in non-preemptive failover mode.	Ensures that after a failed NSX Edge transport node is back online, it does not take over the gateway services thus preventing a short service outage.	None.
VCF-NSX-BGP-RCMD-CFG-007	Enable standby relocation of the Tier-1 gateway.	Ensures that if an edge failure occurs, a standby Tier-1 gateway is created on another edge node.	None.

Table 8-23. BGP Routing Design Recommendations for NSX Federation in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-BGP-RCMD-CFG-008	Use Tier-1 gateways to control the span of networks and ingress and egress traffic in the VMware Cloud Foundation instances.	Enables a mixture of network spans (isolated to a VMware Cloud Foundation instance or spanning multiple instances) without requiring additional Tier-0 gateways and hence edge nodes.	To control location span, a Tier-1 gateway must be assigned to an edge cluster and hence has the Tier-1 SR component. East-west traffic between Tier-1 gateways with SRs need to physically traverse an edge node.
VCF-NSX-BGP-RCMD-CFG-009	Allocate a Tier-1 gateway in each instance for instance-specific networks and connect it to the stretched Tier-0 gateway.	■ Creates a two-tier routing architecture. ■ Enables local-instance networks that are not to span between the VMware Cloud Foundation instances. ■ Guarantees that local-instance networks remain available if a failure occurs in another VMware Cloud Foundation instance.	None.

Overlay Design for VMware Cloud Foundation

As part of the overlay design, you determine the NSX configuration for handling traffic between workloads, management or customer, in VMware Cloud Foundation. You determine the NSX segments and the transport zones.

Logical Overlay Design for VMware Cloud Foundation

This conceptual design provides the network virtualization design of the logical components that handle the data to and from the workloads in the environment. For an environment with multiple VMware Cloud Foundation instances, you replicate the design of the first VMware Cloud Foundation instance to the additional VMware Cloud Foundation instances.

ESXi Host Transport Nodes

A transport node in NSX is a node that is capable of participating in an NSX data plane. The workload domains contain multiple ESXi hosts in a vSphere cluster to support management or customer workloads. You register these ESXi hosts as transport nodes so that networks and workloads on that host can use the capabilities of NSX. During the preparation process, the native vSphere Distributed Switch for the workload domain is extended with NSX capabilities.

Virtual Segments

Geneve provides the overlay capability to create isolated, multi-tenant broadcast domains in NSX across data center fabrics, and enables customers to create elastic, logical networks that span physical network boundaries, and physical locations.

Transport Zones

A transport zone identifies the type of traffic, VLAN or overlay, and the vSphere Distributed Switch name. You can configure one or more VLAN transport zones and a single overlay transport zone per virtual switch. A transport zone does not represent a security boundary. VMware Cloud Foundation supports a single overlay transport zone per NSX Instance. All vSphere clusters, within and across workload domains that share the same NSX instance subsequently share the same overlay transport zone.

Figure 8-10. Transport Zone Design

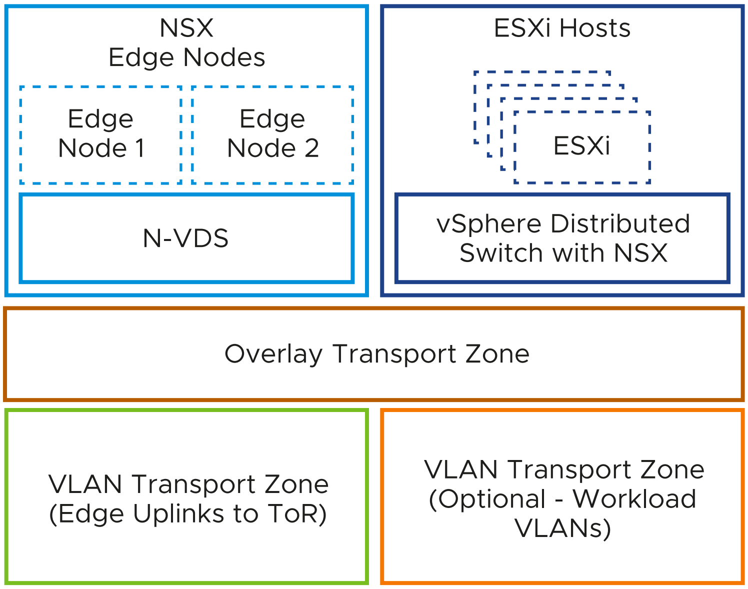

Uplink Policy for ESXi Host Transport Nodes

Uplink profiles define policies for the links from ESXi hosts to NSX segments or from NSX Edge appliances to top of rack switches. By using uplink profiles, you can apply consistent configuration of capabilities for network adapters across multiple ESXi hosts or NSX Edge nodes.

Uplink profiles can use either load balance source or failover order teaming. If using load balance source, multiple uplinks can be active. If using failover order, only a single uplink can be active.

Replication Mode of Segments

The control plane decouples NSX from the physical network. The control plane handles the broadcast, unknown unicast, and multicast (BUM) traffic in the virtual segments.

The following options are available for BUM replication on segments.

Table 8-24. BUM Replication Modes of NSX Segments

BUM Replication Mode	Description
Hierarchical Two-Tier	The ESXi host transport nodes are grouped according to their TEP IP subnet. One ESXi host in each subnet is responsible for replication to an ESXi host in another subnet. The receiving ESXi host replicates the traffic to the ESXi hosts in its local subnet. The source ESXi host transport node knows about the groups based on information it has received from the control plane. The system can select an arbitrary ESXi host transport node as the mediator for the source subnet if the remote mediator ESXi host node is available.
Head-End	The ESXi host transport node at the origin of the frame to be flooded on a segment sends a copy to every other ESXi host transport node that is connected to this segment.

Overlay Design Requirements and Recommendations for VMware Cloud Foundation

Consider the requirements for the configuration of the ESXi hosts in a workload domain as NSX transport nodes, transport zone layout, uplink teaming policies, and the best practices for IP allocation and BUM replication mode in a VMware Cloud Foundation deployment.

Overlay Design Requirements

You must meet the following design requirements in your overlay design.

Table 8-25. Overlay Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-OVERLAY-REQD-CFG-001	Configure all ESXi hosts in the workload domain as transport nodes in NSX.	Enables distributed routing, logical segments, and distributed firewall.	None.
VCF-NSX-OVERLAY-REQD-CFG-002	Configure each ESXi host as a transport node using transport node profiles.	■ Enables the participation of ESXi hosts and the virtual machines running on them in NSX overlay and VLAN networks. ■ Transport node profiles can only be applied at the cluster level.	None.

Table 8-25. Overlay Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-OVERLAY-REQD-CFG-003	To provide virtualized network capabilities to workloads, use overlay networks with NSX Edge nodes and distributed routing.	■ Creates isolated, multi-tenant broadcast domains across data center fabrics to deploy elastic, logical networks that span physical network boundaries. ■ Enables advanced deployment topologies by introducing Layer 2 abstraction from the data center networks.	Requires configuring transport networks with an MTU size of at least 1,600 bytes.
VCF-NSX-OVERLAY-REQD-CFG-004	Create a single overlay transport zone in the NSX instance for all overlay traffic across the host and NSX Edge transport nodes of the workload domain.	■ Ensures that overlay segments are connected to an NSX Edge node for services and north-south routing. ■ Ensures that all segments are available to all ESXi hosts and NSX Edge nodes configured as transport nodes.	All clusters in all workload domains that share the same NSX Manager share the same transport zone.
VCF-NSX-OVERLAY-REQD-CFG-005	Create an uplink profile with a load balance source teaming policy with two active uplinks for ESXi hosts.	For increased resiliency and performance, supports the concurrent use of both physical NICs on the ESXi hosts that are configured as transport nodes.	None.

Overlay Design Recommendations

In your overlay design for VMware Cloud Foundation, you can apply certain best practices.

Table 8-26. Overlay Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-OVERLAY-RCMD-CFG-001	Use static IP pools to assign IP addresses to the host TEP interfaces.	■ Removes the need for an external DHCP server for the host overlay VLANs. ■ You can use NSX Manager to verify static IP pool configurations.	None.
VCF-NSX-OVERLAY-RCMD-CFG-002	Use hierarchical two-tier replication on all overlay segments.	Hierarchical two-tier replication is more efficient because it reduced the number of ESXi hosts the source ESXi host must replicate traffic to if hosts have different TEP subnets. This is typically the case with more than one cluster and will improve performance in that scenario.	None.

Table 8-27. Overlay Design Recommendations for Stretched Clusters inVMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-OVERLAY-RCMD-CFG-003	Configure an NSX sub-transport node profile.	■ You can use static IP pools for the host TEPs in each availability zone. ■ The NSX transport node profile can remain attached when using two separate VLANs for host TEPs at each availability zone as required for clusters that are based on vSphere Lifecycle Manager images. ■ Using an external DHCP server for the host overlay VLANs in both availability zones is not required.	Changes to the host transport node configuration are done at the vSphere cluster level.

Application Virtual Network Design for VMware Cloud Foundation

In VMware Cloud Foundation, you place VMware Aria Suite components on a pre-defined configuration of NSX segments (known as application virtual networks or AVNs) for dynamic routing and load balancing.

Logical Application Virtual Network Design VMware Cloud Foundation

NSX segments provide flexibility for workload placement by removing the dependence on traditional physical data center networks. This approach also improves security and mobility of the management applications, and reduces the integration effort with existing customer network.

Table 8-28. Comparing Application Virtual Network Types

Design Component	Overlay-Based NSX Segments	VLAN-Backed NSX Segments
Benefits	■ Supports IP mobility with dynamic routing. ■ Limits the number of VLANs needed in the data center fabric. ■ In an environment with multiple availability zones, limits the number of VLANs needed to expand from an architecture with one availability zone to an architecture with two availability zones.	Uses the data center fabric for the network segment and the next-hop gateway.
Requirement	Requires routing between the data center fabric and the NSX Edge nodes.	

Figure 8-11. Application Virtual Networks in VMware Cloud Foundation

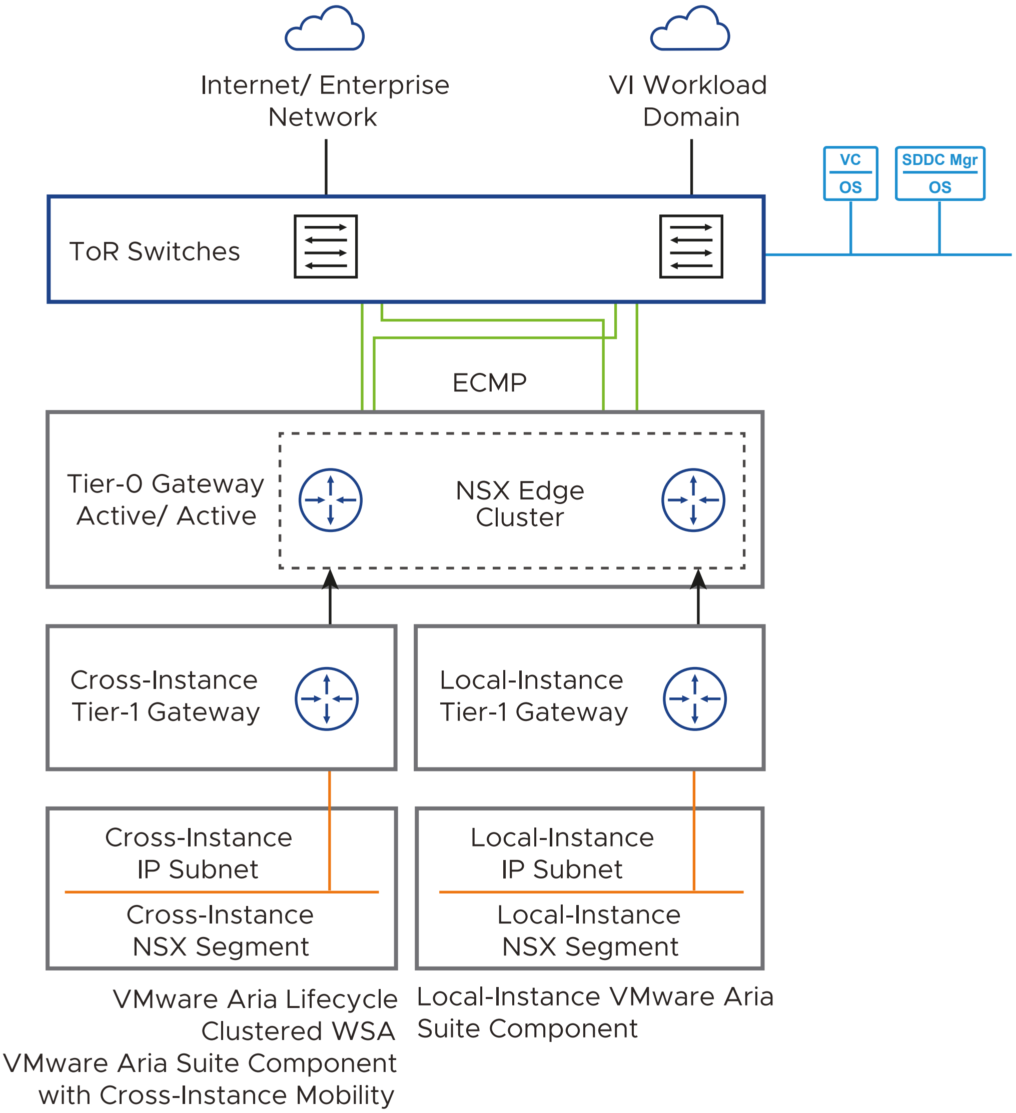
Internet/ Enterprise Network
VI Workload Domain
VC
OS
SDDC Mgr
OS
ToR Switches
ECMP
Tier-0 Gateway Active/ Active
NSX Edge Cluster
Cross-Instance Tier-1 Gateway
Local-Instance Tier-1 Gateway
Cross-Instance IP Subnet
Cross-Instance NSX Segment
Local-Instance IP Subnet
Local-Instance NSX Segment
VMware Aria Lifecycle
Clustered WSA
VMware Aria Suite Component with Cross-Instance Mobility
Local-Instance VMware Aria Suite Component

For the design for specific VMware Aria Suite components, see this design and VMware Validated Solutions. For identity and access management design for NSX, see Identity and Access Management for VMware Cloud Foundation.

Important If you plan to use NSX Federation in the management domain, create the AVNs before you enable the federation. Creating AVNs in an environment where NSX Federation is already active is not supported.

With NSX Federation, an NSX segment can span multiple instances of NSX and VMware Cloud Foundation. A single network segment can be available in different physical locations over the NSX SDN. In an environment with multiple VMware Cloud Foundation instances, the cross-instance NSX network in the management domain is extended between the first two instances. This configuration provides IP mobility for management components which fail over from the first to the second instance.

Application Virtual Network Design Requirements and Recommendations forVMware Cloud Foundation

Consider the requirements and best practices for the configuration of the NSX segments for using the Application Virtual Networks in VMware Cloud Foundation for a single VMware Cloud Foundation or multiple VMware Cloud Foundation instances.

Application Virtual Network Design Requirements

You must meet the following design requirements in your Application Virtual Network design for a single VMware Cloud Foundation instance and for multiple VMware Cloud Foundation instances.

Table 8-29. Application Virtual Network Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-AVN-REQD-CFG-001	Create one cross-instance NSX segment for the components of a VMware Aria Suite application or another solution that requires mobility between VMware Cloud Foundation instances.	Prepares the environment for the deployment of solutions on top of VMware Cloud Foundation, such as VMware Aria Suite, without a complex physical network configuration. The components of the VMware Aria Suite application must be easily portable between VMware Cloud Foundation instances without requiring reconfiguration.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-002	Create one or more local-instance NSX segments for the components of a VMware Aria Suite application or another solution that are assigned to a specific VMware Cloud Foundation instance.	Prepares the environment for the deployment of solutions on top of VMware Cloud Foundation, such as VMware Aria Suite, without a complex physical network configuration.	Each NSX segment requires a unique IP address space.

Table 8-30. Application Virtual Network Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-AVN-REQD-CFG-003	Extend the cross-instance NSX segment to the second VMware Cloud Foundation instance.	Enables workload mobility without a complex physical network configuration. The components of a VMware Aria Suite application must be easily portable between VMware Cloud Foundation instances without requiring reconfiguration.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-004	In each VMware Cloud Foundation instance, create additional local-instance NSX segments.	Enables workload mobility within a VMware Cloud Foundation instance without complex physical network configuration. Each VMware Cloud Foundation instance should have network segments to support workloads which are isolated to that VMware Cloud Foundation instance.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-005	In each VMware Cloud Foundation instance, connect or migrate the local-instance NSX segments to the corresponding local-instance Tier-1 gateway.	Configures local-instance NSX segments at required sites only.	Requires an individual Tier-1 gateway for local-instance segments.

Application Virtual Network Design Recommendations

In your Application Virual Network design for VMware Cloud Foundation, you can apply certain best practices.

Table 8-31. Application Virtual Network Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-AVN-RCMD-CFG-001	Use overlay-backed NSX segments.	■ Supports expansion to deployment topologies for multiple VMware Cloud Foundation instances. ■ Limits the number of VLANs required for the data center fabric.	Using overlay-backed NSX segments requires routing, eBGP recommended, between the data center fabric and edge nodes.

Load Balancing Design for VMware Cloud Foundation

Following the principles of this design and of each product, you deploy and configure NSX load balancing services to support VMware Aria Suite and Workspace ONE Access components.

Logical Load Balancing Design for VMware Cloud Foundation

The logical load balancer capability in NSX offers a high-availability service for applications in VMware Cloud Foundation and distributes the network traffic load among multiple servers.

A standalone Tier-1 gateway is created to provide load balancing services with a service interface on the cross-instance application virtual network.

Figure 8-12. NSX Logical Load Balancing Design for VMware Cloud Foundation

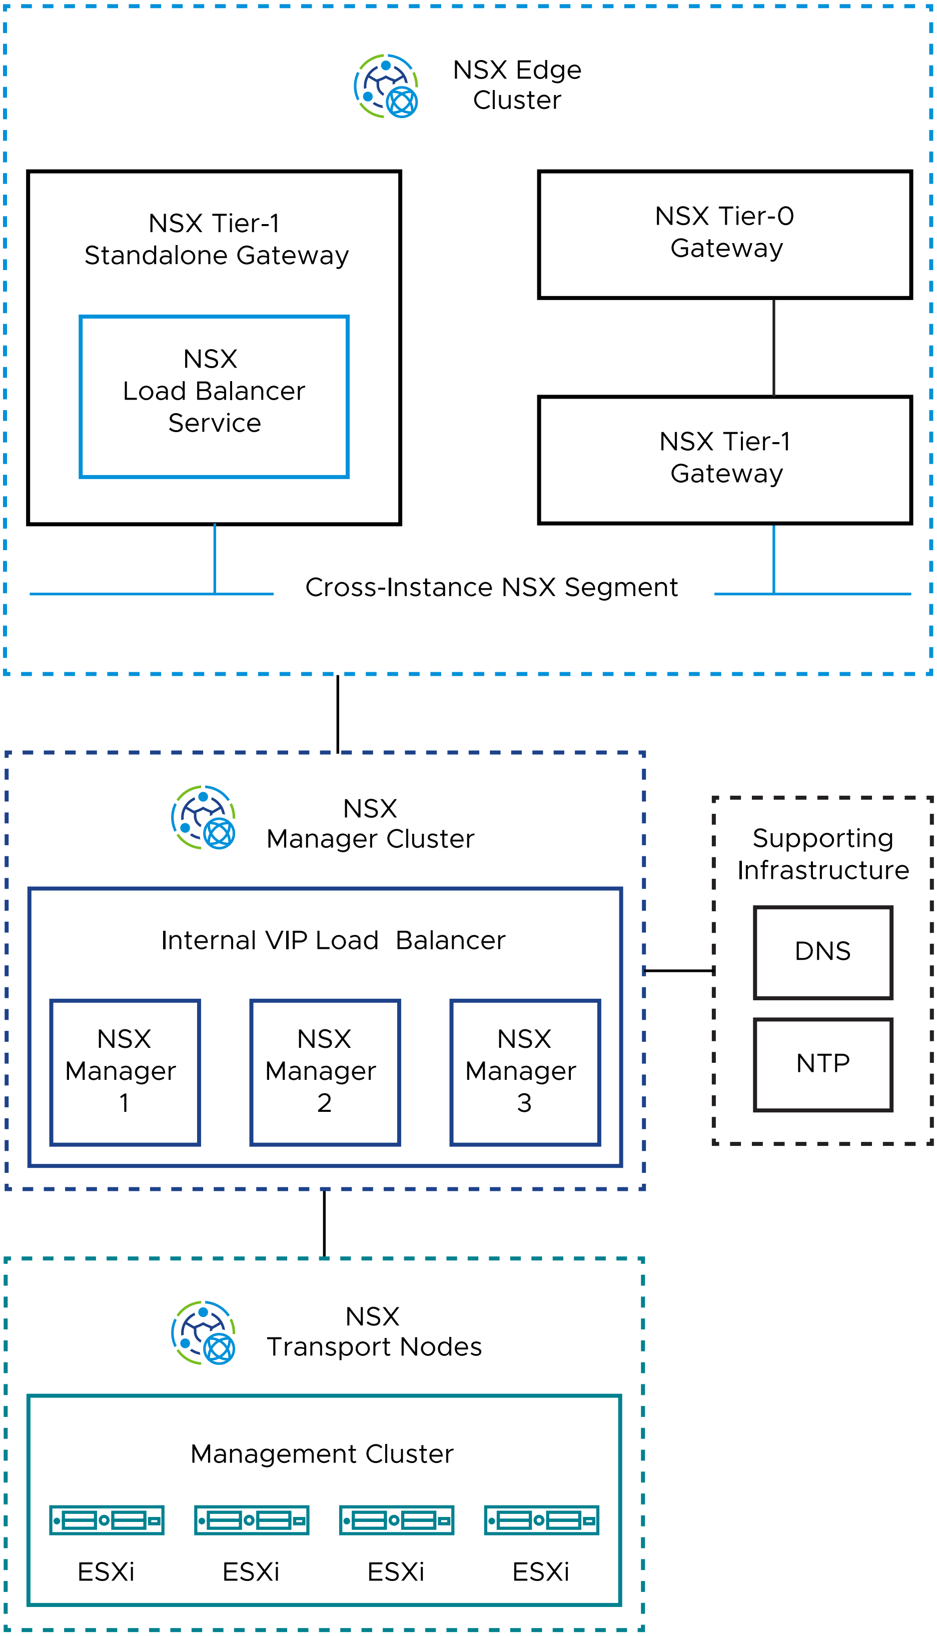

Load Balancing Design Requirements for VMware Cloud Foundation

Consider the requirements for running a load balancing service including creating a standalone Tier-1 gateway and connecting it to the client applications. Separate requirements exist for a single VMware Cloud Foundation instance and for multiple VMware Cloud Foundation instances.

Table 8-32. Load Balancing Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-001	Deploy a standalone Tier-1 gateway to support advanced stateful services such as load balancing for other management components.	Provides independence between north-south Tier-1 gateways to support advanced deployment scenarios.	You must add a separate Tier-1 gateway.
VCF-NSX-LB-REQD-CFG-002	When creating load balancing services for Application Virtual Networks, connect the standalone Tier-1 gateway to the cross-instance NSX segments.	Provides load balancing to applications connected to the cross-instance network.	You must connect the gateway to each network that requires load balancing.
VCF-NSX-LB-REQD-CFG-003	Configure a default static route on the standalone Tier-1 gateway with a next hop the Tier-1 gateway for the segment to provide connectivity to the load balancer.	Because the Tier-1 gateway is standalone, it does not auto-configure its routes.	None.

Table 8-33. Load Balancing Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-004	Deploy a standalone Tier-1 gateway in the second VMware Cloud Foundation instance.	Provides a cold-standby non-global service router instance for the second VMware Cloud Foundation instance to support services on the cross-instance network which require advanced services not currently supported as NSX global objects.	■ You must add a separate Tier-1 gateway. ■ You must manually configure any services and synchronize them between the non-global service router instances in the first and second VMware Cloud Foundation instances. ■ To avoid a network conflict between the two VMware Cloud Foundation instances, make sure that the primary and standby networking services are not both active at the same time.
VCF-NSX-LB-REQD-CFG-005	Connect the standalone Tier-1 gateway in the second VMware Cloud Foundationinstance to the cross-instance NSX segment.	Provides load balancing to applications connected to the cross-instance network in the second VMware Cloud Foundation instance.	You must connect the gateway to each network that requires load balancing.

Table 8-33. Load Balancing Design Requirements for NSX Federation in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-006	Configure a default static route on the standalone Tier-1 gateway in the second VMware Cloud Foundation instance with a next hop as the Tier-1 gateway for the segment it connects with to provide connectivity to the load balancers.	Because the Tier-1 gateway is standalone, it does not autoconfigure its routes.	None.
VCF-NSX-LB-REQD-CFG-007	Establish a process to ensure any changes made on to the load balancer instance in the first VMware Cloud Foundationinstance are manually applied to the disconnected load balancer in the second instance.	Keeps the network service in the failover load balancer instance ready for activation if a failure in the first VMware Cloud Foundation instance occurs. Because network services are not supported as global objects, you must configure them manually in each VMware Cloud Foundation instance. The load balancer service in one instance must be connected and active, while the service in the other instance must be disconnected and inactive.	■ Because of incorrect configuration between the VMware Cloud Foundation instances, the load balancer service in the second instance might come online with an invalid or incomplete configuration. ■ If both VMware Cloud Foundation instances are online and active at the same time, a conflict between services could occur resulting in a potential outage. ■ The administrator must establish and follow an operational practice by using a runbook or automated process to ensure that configuration changes are reproduced in each VMware Cloud Foundation instance.

SDDC Manager Design for VMware Cloud Foundation

9

In VMware Cloud Foundation, operational day-to-day efficiencies are delivered through SDDC Manager. These efficiencies include full life cycle management tasks such as deployment, configuration, patching and upgrades.

Read the following topics next:

- Logical Design for SDDC Manager
- SDDC Manager Design Requirements and Recommendations for VMware Cloud Foundation

Logical Design for SDDC Manager

You deploy an SDDC Manager appliance in the management domain for creating VI workload domains, provisioning additional virtual infrastructure, and life cycle management of the SDDC management components.

Figure 9-1. Logical Design of SDDC Manager

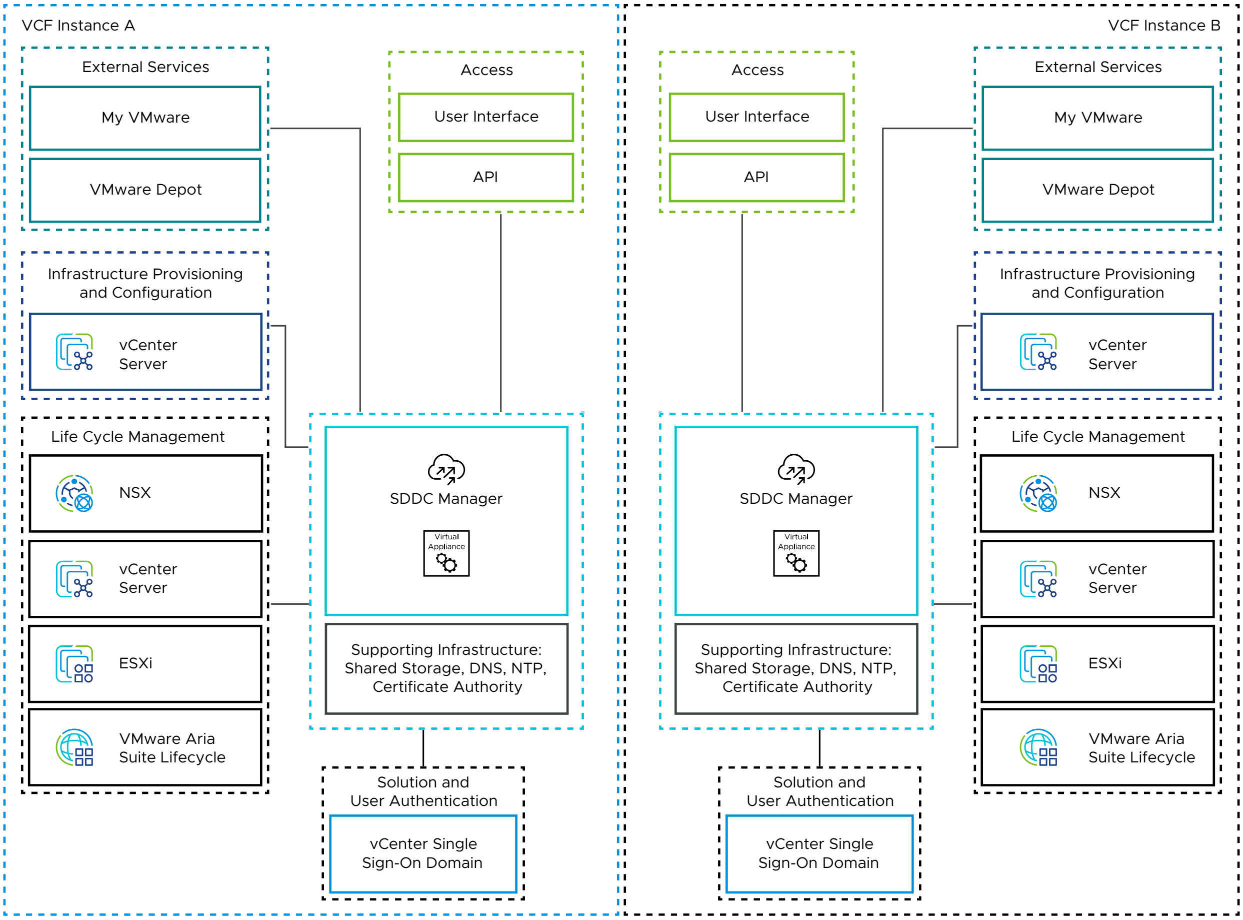
VCF Instance A
External Services
My VMware
VMware Depot
Access
User Interface
API
Infrastructure Provisioning and Configuration
vCenter Server
Life Cycle Management
NSX
vCenter Server
ESXi
VMware Aria Suite Lifecycle
SDDC Manager
Virtual Appliance
Supporting Infrastructure: Shared Storage, DNS, NTP, Certificate Authority
Solution and User Authentication
vCenter Single Sign-On Domain
VCF Instance B
Access
User Interface
API
External Services
My VMware
VMware Depot
Infrastructure Provisioning and Configuration
vCenter Server
SDDC Manager
Virtual Appliance
Supporting Infrastructure: Shared Storage, DNS, NTP, Certificate Authority
Life Cycle Management
NSX
vCenter Server
ESXi
VMware Aria Suite Lifecycle
Solution and User Authentication
vCenter Single Sign-On Domain

You use SDDC Manager to perform the following operations:

- Commissioning or decommissioning ESXi hosts
- Deployment of VI workload domains
- Deployment of VMware Aria Suite Lifecycle
- Deployment of NSX Edge clusters in workload domains
- Adding and extending clusters in workload domains
- Life cycle management of the virtual infrastructure components in all workload domains and of VMware Aria Suite Lifecycle
- Storage management for vVOL VASA providers
- Identity provider management
- Composable infrastructure management
- Creation of network pools for host configuration workload domains
- Product licenses storage
- Certificate management

- Password management and rotation
- Backup configuration

Table 9-1. SDDC Manager Logical Components

VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
■ A single SDDC Manager appliance is deployed on the management network. ■ vSphere HA protects the SDDC Manager appliance.	■ A single SDDC Manager appliance is deployed on the management network. ■ vSphere HA protects the SDDC Manager appliance. ■ A vSphere DRS rule specifies that the SDDC Manager appliance should run on an ESXi host in the first availability zone.

SDDC Manager Design Requirements and Recommendations for VMware Cloud Foundation

Consider the placement and network design requirements for SDDC Manager, and the best practices for configuring the access to install and upgrade software bundles.

SDDC Manager Design Requirements

You must meet the following design requirements for in your SDDC Manager design.

Table 9-2. SDDC Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-SDDCMGR-REQD-CFG-001	Deploy an SDDC Manager system in the first availability zone of the management domain.	SDDC Manager is required to perform VMware Cloud Foundation capabilities, such as provisioning VI workload domains, deploying solutions, patching, upgrading, and others.	None.
VCF-SDDCMGR-REQD-CFG-002	Deploy SDDC Manager with its default configuration.	The configuration of SDDC Manager is not configurable and should not be changed from its defaults.	None.
VCF-SDDCMGR-REQD-CFG-003	Place the SDDC Manager appliance on the VM management network.	■ Simplifies IP addressing for management VMs by using the same VLAN and subnet. ■ Provides simplified secure access to management VMs in the same VLAN network.	None.

SDDC Manager Design Recommendations

In your SDDC Manager design, you can apply certain best practices.

Table 9-3. SDDC Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-SDDCMGR-RCMD-CFG-001	Connect SDDC Manager to the Internet for downloading software bundles.	SDDC Manager must be able to download install and upgrade software bundles for deployment of VI workload domains and solutions, and for upgrade from a repository.	The rules of your organization might not permit direct access to the Internet. In this case, you must download software bundles for SDDC Manager manually.
VCF-SDDCMGR-RCMD-CFG-002	Configure a network proxy to connect SDDC Manager to the Internet.	To protect SDDC Manager against external attacks from the Internet.	The proxy must not use authentication because SDDC Manager does not support proxy with authentication.
VCF-SDDCMGR-RCMD-CFG-003	Configure SDDC Manager with a VMware Customer Connect account with VMware Cloud Foundation entitlement to check for and download software bundles.	Software bundles for VMware Cloud Foundation are stored in a repository that is secured with access controls.	Requires the use of a VMware Customer Connect user account with access to VMware Cloud Foundation licensing. Sites without an internet connection can use local upload option instead.
VCF-SDDCMGR-RCMD-CFG-004	Configure SDDC Manager with an external certificate authority that is responsible for providing signed certificates.	Provides increased security by implementing signed certificate generation and replacement across the management components.	An external certificate authority, such as Microsoft CA, must be locally available.

VMware Aria Suite Lifecycle Design for VMware Cloud Foundation

10

In VMware Cloud Foundation, VMware Aria Suite Lifecycle provides life cycle management capabilities for VMware Aria Suite components and Workspace ONE Access, including automated deployment, configuration, patching, and upgrade, and content management across VMware Aria Suite products.

You deploy VMware Aria Suite Lifecycle by using SDDC Manager. SDDC Manager deploys VMware Aria Suite Lifecycle in VMware Cloud Foundation mode. In this mode, VMware Aria Suite Lifecycle is integrated with SDDC Manager, providing the following benefits:

- Integration with the SDDC Manager inventory to retrieve infrastructure details when creating environments for Workspace ONE Access and VMware Aria Suite components, such as NSX segments and vCenter Server details.
- Automation of the NSX load balancer configuration when deploying Workspace ONE Access, VMware Aria Operations, and VMware Aria Automation.
- Deployment details for VMware Aria Suite Lifecycle environments are populated in the SDDC Manager inventory and can be queried using the SDDC Manager API.
- Day-two workflows in SDDC Manager to connect VMware Aria Operations for Logs and VMware Aria Operations to workload domains.
- The ability to manage password life cycle for Workspace ONE Access and VMware Aria Suite components.

For information about deploying VMware Aria Suite components, see VMware Validated Solutions.

Read the following topics next:

Logical Design for VMware Aria Suite Lifecycle for VMware Cloud Foundation

You deploy VMware Aria Suite Lifecycle to provide life cycle management capabilities for VMware Aria Suite components and a Workspace ONE Access cluster.

Logical Design

In a VMware Cloud Foundation environment, you use VMware Aria Suite Lifecycle in VMware Cloud Foundation mode. In this mode, VMware Aria Suite Lifecycle is integrated with VMware Cloud Foundation in the following way:

- SDDC Manager deploys the VMware Aria Suite Lifecycle appliance. Then, you deploy the VMware Aria Suite products that are supported by VMware Cloud Foundation by using VMware Aria Suite Lifecycle.
- Supported versions are controlled by the VMware Aria Suite Lifecycle appliance and Product Support Packs. See the *VMware Interoperability Matrix*.
- To orchestrate the deployment, patching, and upgrade of Workspace ONE Access and the VMware Aria Suite products, VMware Aria Suite Lifecycle communicates with SDDC Manager and the management domain vCenter Server in the environment.
- SDDC Manager configures the load balancer for Workspace ONE Access, VMware Aria Operations, and VMware Aria Automation.

Figure 10-1. Logical Design of VMware Aria Suite Lifecycle

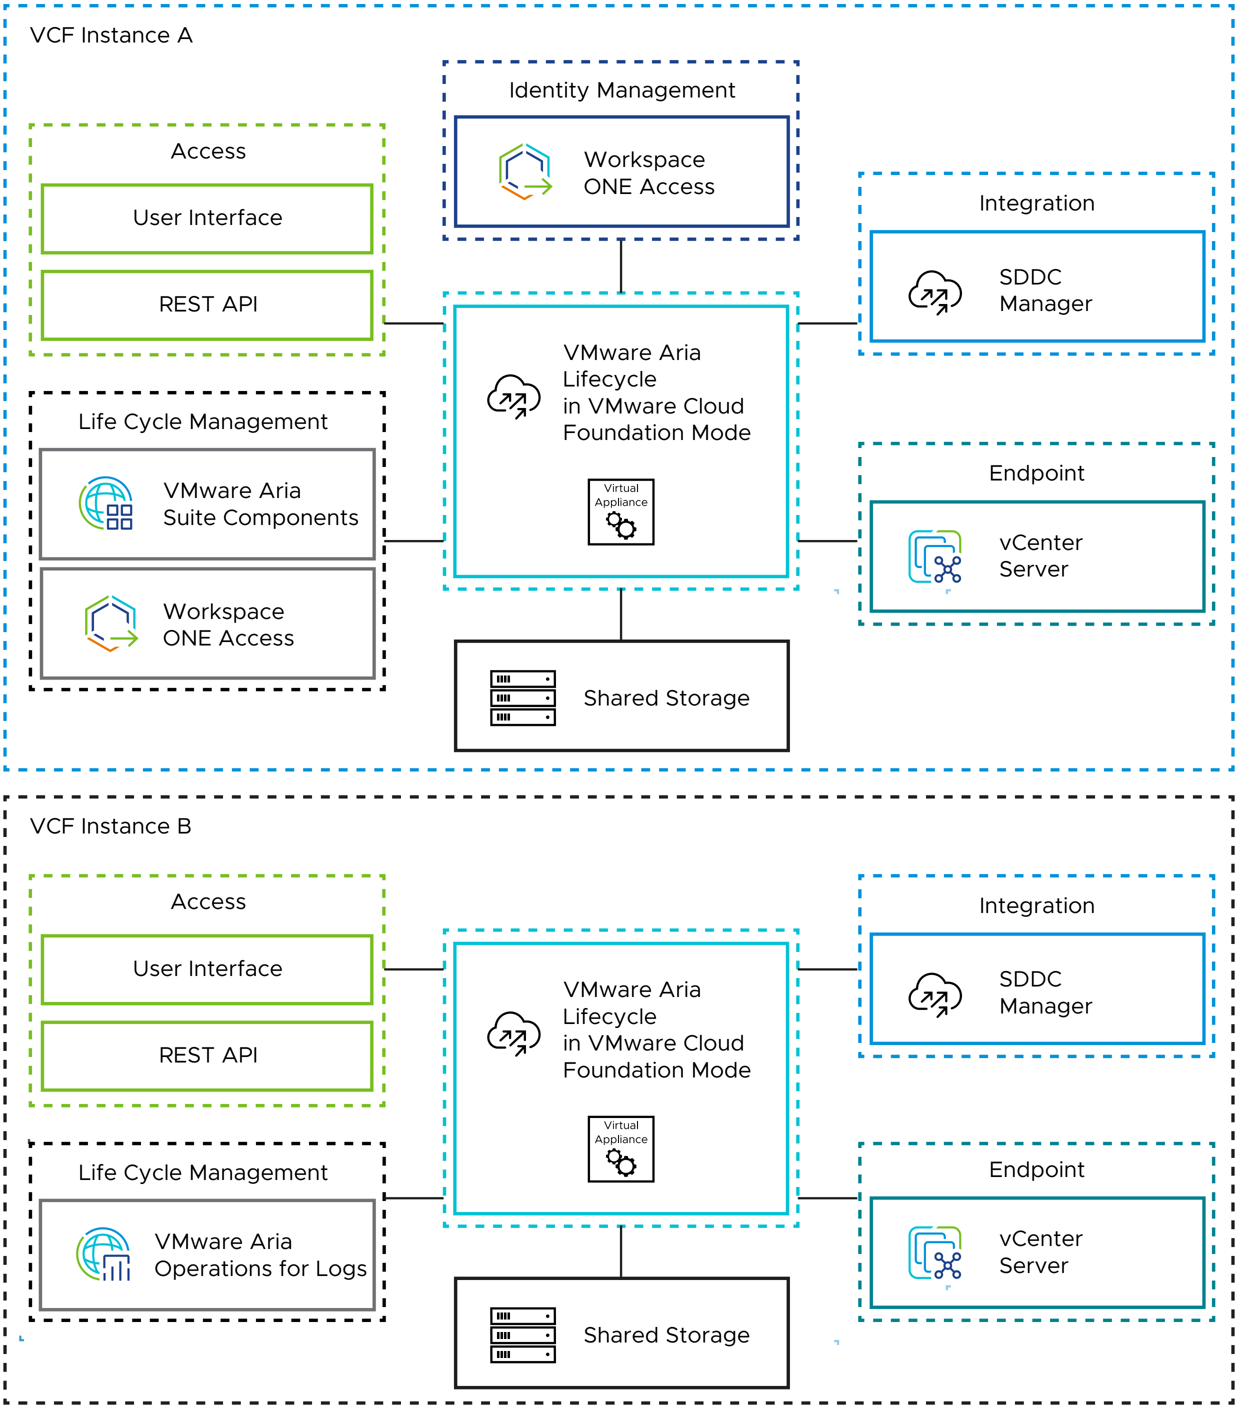

According to the VMware Cloud Foundation topology deployed, VMware Aria Suite Lifecycle is deployed in one or more locations and is responsible for the life cycle of the VMware Aria Suite components in one or more VMware Cloud Foundation instances.

VMware Cloud Foundation instances might be connected for the following reasons:

- Disaster recovery of the VMware Aria Suite components.

- Over-arching management of those instances from the same VMware Aria Suite deployments.

Table 10-1. VMware Aria Suite Lifecycle Component Layout

VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones	Connected VMware Cloud Foundation Instances
■ A single VMware Aria Suite Lifecycle appliance deployed on the cross-instance NSX segment. ■ vSphere HA protects the VMware Aria Suite Lifecycle appliance. Life cycle management for: ■ Workspace ONE Access ■ VMware Aria Suite	■ A single VMware Aria Suite Lifecycle appliance deployed on the cross-instance NSX segment. ■ vSphere HA protects the VMware Aria Suite Lifecycle appliance. ■ A should-run vSphere DRS rule specifies that the VMware Aria Suite Lifecycle appliance should run on an ESXi host in the first availability zone. Life cycle management for: ■ Workspace ONE Access ■ VMware Aria Suite	The VMware Aria Suite Lifecycle instance in the first VMware Cloud Foundation instance provides life cycle management for: ■ Workspace ONE Access ■ VMware Aria Suite VMware Aria Suite Lifecycle in each additional VMware Cloud Foundation instance provides life cycle management for: ■ VMware Aria Operations for Logs

Network Design for VMware Aria Suite Lifecycle

For secure access to the UI and API, you place the VMware Aria Suite Lifecycle appliance on an overlay-backed (recommended) or VLAN-backed Application Virtual Network.

VMware Aria Suite Lifecycle must have routed access to the management VLAN through the Tier-0 gateway in the NSX instance for the management domain.

Figure 10-2. Network Design for VMware Aria Suite Lifecycle

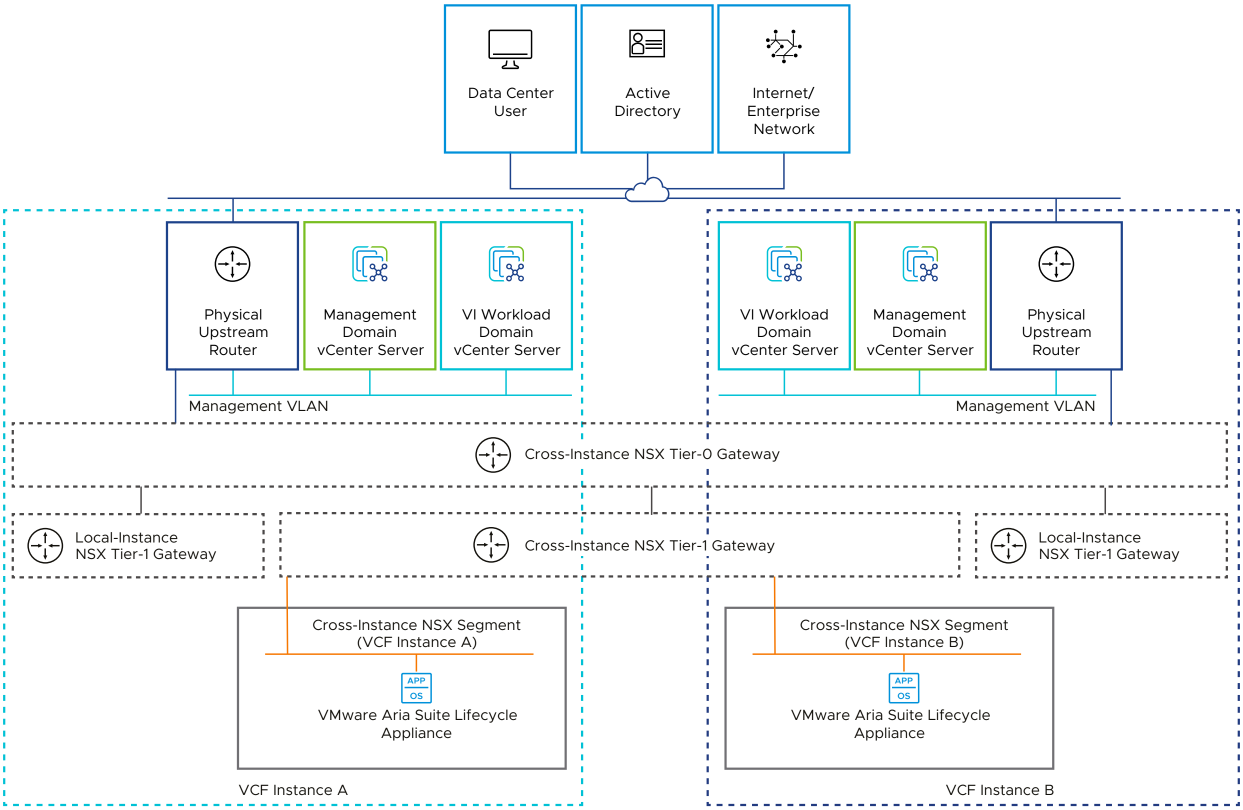
Data Center User
Active Directory
Internet/ Enterprise Network
Physical Upstream Router
Management Domain vCenter Server
VI Workload Domain vCenter Server
VI Workload Domain vCenter Server
Management Domain vCenter Server
Physical Upstream Router
Management VLAN
Management VLAN
Cross-Instance NSX Tier-0 Gateway
Local-Instance NSX Tier-1 Gateway
Cross-Instance NSX Tier-1 Gateway
Local-Instance NSX Tier-1 Gateway
Cross-Instance NSX Segment (VCF Instance A)
APP
OS
VMware Aria Suite Lifecycle Appliance
Cross-Instance NSX Segment (VCF Instance B)
APP
OS
VMware Aria Suite Lifecycle Appliance
VCF Instance A
VCF Instance B

Data Center and Environment Design for VMware Aria Suite Lifecycle

To deploy VMware Aria Suite products by using VMware Aria Suite Lifecycle, you configure product support, data centers, environment structures, and product specifications.

Product Support

VMware Aria Suite Lifecycle provides several methods to obtain and store product binaries for the install, patch, and upgrade of the VMware Aria Suite products.

Table 10-2. Methods for Obtaining and Storing Product Binaries

Method	Description
Product Upload	■ You can upload and discover product binaries to the VMware Aria Suite Lifecycle appliance.
VMware Customer Connect	■ You can integrate vVMware Aria Suite Lifecycle with VMware Customer Connect to access and download VMware Aria Suite product entitlements from an online depot over the Internet. This method simplifies, automates, and organizes the repository.

Data Centers and Environments

VMware Aria Suite Lifecycle supports the deployment and upgrade of VMware Aria Suite products in a logical environment grouping.

You create data centers and environments in VMware Aria Suite Lifecycle to manage the life cycle operations on the VMware Aria Suite products and to support the growth of the SDDC.

Table 10-3. VMware Aria Suite Lifecycle Logical Constructs

Construct	Definition
Datacenter	Represents a geographical or logical location for an organization. Management domain vCenter Server instances are added to specific data centers.
Environment	Is mapped to a data center object. Each environment can contain only one instance of a VMware Aria Suite product.

Table 10-4. Logical Datacenter to vCenter Server Mappings in VMware Aria Suite Lifecycle

Logical Datacenter	vCenter Server Type	Description
Cross-instance	■ Management domain vCenter Server for the local VMware Cloud Foundation instance. ■ Management domain vCenter Server for an additional VMware Cloud Foundation instance.	Supports the deployment of cross-instance components, such as Workspace ONE Access, VMware Aria Operations, and VMware Aria Automation, including any per-instance collector components.
Local-instance	Management domain vCenter Server for the local VMware Cloud Foundation instance.	Supports the deployment of VMware Aria Operations for Logs.

Table 10-5. VMware Aria Suite Lifecycle Environment Types

Environment Type	Description
Global Environment	Contains the Workspace ONE Access instance that is required before you can deploy VMware Aria Automation.
VMware Cloud Foundation Mode	■ Infrastructure details for the deployed products, including vCenter Server, networking, DNS and NTP information are retrieved from the SDDC Manager inventory. ■ Successful deployment details are synced back to the SDDC Manager inventory. ■ Limited to one instance of each VMware Aria Suite product.
Standalone Mode	■ Infrastructure details for the deployed products are entered manually. ■ Successful deployment details are not synced back to the SDDC Manager inventory. ■ Supports deployment of more than one instance of a VMware Aria Suite product.

Note You can deploy new VMware Aria Suite products to the SDDC environment or import existing product deployments.

Table 10-6. Environment Topologies

Environment Name	VMware Cloud Foundation Mode	Logical Datacenter	Product Components
Global Environment	Enabled	Cross-instance	Workspace ONE Access
Cross-instance	Enabled	Cross-instance	■ VMware Aria Operations analytics nodes ■ VMware Aria Operations remote collectors ■ VMware Aria Automation cluster
Each instance	Enabled	Local-instance	VMware Aria Operations for Logs cluster nodes

Locker Design for VMware Aria Suite Lifecycle

The VMware Aria Suite Lifecycle Locker allows you to secure and manage passwords, certificates, and licenses for VMware Aria Suite product solutions and integrations.

Passwords

VMware Aria Suite Lifecycle stores passwords in the locker repository which are referenced during life cycle operations on data centers, environments, products, and integrations.

Table 10-7. Life Cycle Operations Use of Locker Passwords in VMware Aria Suite Lifecycle

Life Cycle Operations Element	Password Use
Datacenters	vCenter Server credentials for aVMware Aria Suite Lifecycle-to-vSphere integration user.
Environments	■ Global environment default configuration administrator,**configadmin**. ■ Environment password, for example, for product default **admin** or **root** password.
Products	■ Product administrator password, for example, the **admin** password for an individual product. ■ Product appliance password, for example, the **root** password for an individual product.

Certificates

VMware Aria Suite Lifecycle stores certificates in the Locker repository which can be referenced during product life cycle operations. Externally provided certificates, such as Certificate Authority-signed certificates, can be imported or certificates can be generated by the VMware Aria Suite Lifecycle appliance.

Licenses

VMware Aria Suite Lifecycle stores licenses in the Locker repository which can be referenced during product life cycle operations. Licenses can be validated and added to the repository directory or imported through an integration with VMware Customer Connect.

VMware Aria Suite Lifecycle Design Requirements and Recommendations for VMware Cloud Foundation

Consider the placement, networking, sizing and high availability requirements for using VMware Aria Suite Lifecycle for deployment and life cycle management of VMware Aria Suite components in VMware Cloud Foundation. Apply similar best practices for having VMware Aria Suite Lifecycle operate in an optimal way.

VMware Aria Suite Lifecycle Design Requirements

You must meet the following design requirements for standard and stretched clusters in your VMware Aria Suite Lifecycle design for VMware Cloud Foundation. For NSX Federation, additional requirements exist.

Table 10-8. VMware Aria Suite Lifecycle Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-001	Deploy a VMware Aria Suite Lifecycle instance in the management domain of each VMware Cloud Foundation instance to provide life cycle management for VMware Aria Suite and Workspace ONE Access.	Provides life cycle management operations for VMware Aria Suite applications and Workspace ONE Access.	You must ensure that the required resources are available.
VCF-VASL-REQD-CFG-002	Deploy VMware Aria Suite Lifecycle by using SDDC Manager.	■ Deploys VMware Aria Suite Lifecycle in VMware Cloud Foundation mode, which enables the integration with the SDDC Manager inventory for product deployment and life cycle management of VMware Aria Suite components. ■ Automatically configures the standalone Tier-1 gateway required for load balancing the clustered Workspace ONE Access and VMware Aria Suite components.	None.
VCF-VASL-REQD-CFG-003	Allocate extra 100 GB of storage to the VMware Aria Suite Lifecycle appliance for VMware Aria Suite product binaries.	■ Provides support for VMware Aria Suite product binaries (install, upgrade, and patch) and content management. ■ SDDC Manager automates the creation of storage.	None.
VCF-VASL-REQD-CFG-004	Place the VMware Aria Suite Lifecycle appliance on an overlay-backed (recommended) or VLAN-backed NSX network segment.	Provides a consistent deployment model for management applications.	You must use an implementation in NSX to support this networking configuration.
VCF-VASL-REQD-CFG-005	Import VMware Aria Suite product licenses to the Locker repository for product life cycle operations.	■ You can review the validity, details, and deployment usage for the license across the VMware Aria Suite products. ■ You can reference and use licenses during product life cycle operations, such as deployment and license replacement.	When using the API, you must specify the Locker ID for the license to be used in the JSON payload.

Table 10-8. VMware Aria Suite Lifecycle Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-ENV-001	Configure datacenter objects in VMware Aria Suite Lifecycle for local and cross-instance VMware Aria Suite deployments and assign the management domain vCenter Server instance to each data center.	You can deploy and manage the integrated VMware Aria Suite components across the SDDC as a group.	You must manage a separate datacenter object for the products that are specific to each instance.
VCF-VASL-REQD-ENV-002	If deploying VMware Aria Operations for Logs, create a local-instance environment in VMware Aria Suite Lifecycle.	Supports the deployment of an instance of VMware Aria Operations for Logs.	None.
VCF-VASL-REQD-ENV-003	If deploying VMware Aria Operations or VMware Aria Automation, create a cross-instance environment in VMware Aria Suite Lifecycle	■ Supports deployment and management of the integrated VMware Aria Suite products across VMware Cloud Foundation instances as a group. ■ Enables the deployment of instance-specific components, such as VMware Aria Operations remote collectors. In VMware Aria Suite Lifecycle, you can deploy and manage VMware Aria Operations remote collector objects only in an environment that contains the associated cross-instance components.	You can manage instance-specific components, such as remote collectors, only in an environment that is cross-instance.

Table 10-8. VMware Aria Suite Lifecycle Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-SEC-001	Use the custom vCenter Server role for VMware Aria Suite Lifecycle that has the minimum privileges required to support the deployment and upgrade of VMware Aria Suite products.	VMware Aria Suite Lifecycle accesses vSphere with the minimum set of permissions that are required to support the deployment and upgrade of VMware Aria Suite products. SDDC Manager automates the creation of the custom role.	You must maintain the permissions required by the custom role.
VCF-VASL-REQD-SEC-002	Use the service account in vCenter Server for application-to-application communication from VMware Aria Suite Lifecycle to vSphere. Assign global permissions using the custom role.	■ Provides the following access control features: ■ VMware Aria Suite Lifecycle accesses vSphere with the minimum set of required permissions. ■ You can introduce improved accountability in tracking request-response interactions between the components of the SDDC. ■ SDDC Manager automates the creation of the service account.	■ You must maintain the life cycle and availability of the service account outside of SDDC manager password rotation.

Table 10-9. VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-006	For multiple availability zones, add the VMware Aria Suite Lifecycle appliance to the VM group for the first availability zone.	Ensures that, by default, the VMware Aria Suite Lifecycle appliance is powered on a host in the first availability zone.	If VMware Aria Suite Lifecycle is deployed after the creation of the stretched management cluster, you must add the VMware Aria Suite Lifecycle appliance to the VM group manually.

Table 10-10. VMware Aria Suite Lifecycle Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-007	Configure the DNS settings for the VMware Aria Suite Lifecycle appliance to use DNS servers in each instance.	Improves resiliency in the event of an outage of external services for a VMware Cloud Foundation instance.	As you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the DNS settings of the VMware Aria Suite Lifecycle appliance must be updated.
VCF-VASL-REQD-CFG-008	Configure the NTP settings for the VMware Aria Suite Lifecycle appliance to use NTP servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	As you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the NTP settings on the VMware Aria Suite Lifecycle appliance must be updated.
VCF-VASL-REQD-ENV-004	Assign the management domain vCenter Server instance in the additional VMware Cloud Foundation instance to the cross-instance data center.	Supports the deployment of VMware Aria Operations remote collectors in an additional VMware Cloud Foundation instance.	None.

VMware Aria Suite Lifecycle Design Recommendations

In your VMware Aria Suite Lifecycle design for VMware Cloud Foundation, you can apply certain best practices .

Table 10-11. VMware Aria Suite Lifecycle Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VASL-RCMD-CFG-001	Protect VMware Aria Suite Lifecycle by using vSphere HA.	Supports the availability objectives for VMware Aria Suite Lifecycle without requiring manual intervention during a failure event.	None.
VCF-VASL-RCMD-LCM-001	Obtain product binaries for install, patch, and upgrade in VMware Aria Suite Lifecycle from VMware Customer Connect.	■ You can upgrade VMware Aria Suite products based on their general availability and endpoint interoperability rather than being listed as part of VMware Cloud Foundation bill of materials (BOM). ■ You can deploy and manage binaries in an environment that does not allow access to the Internet or are dark sites.	The site must have an Internet connection to use VMware Customer Connect. Sites without an Internet connection should use the local upload option instead.
VCF-VASL-RCMD-LCM-002	Use support packs (PSPAKS) for VMware Aria Suite Lifecycle to enable upgrading to later versions of VMware Aria Suite products.	Enables the upgrade of an existing VMware Aria Suite Lifecycle to permit later versions of VMware Aria Suite products without an associated VMware Cloud Foundation upgrade. See VMware Knowledge Base article 88829	None.
VCF-VASL-RCMD-SEC-001	Enable integration between VMware Aria Suite Lifecycle and your corporate identity source by using the Workspace ONE Access instance.	■ Enables authentication to VMware Aria Suite Lifecycle by using your corporate identity source. ■ Enables authorization through the assignment of organization and cloud services roles to enterprise users and groups defined in your corporate identity source.	You must deploy and configure Workspace ONE Access to establish the integration between VMware Aria Suite Lifecycle and your corporate identity sources.
VCF-VASL-RCMD-SEC-002	Create corresponding security groups in your corporate directory services for VMware Aria Suite Lifecycle roles: ■ **VCF** ■ **Content Release Manager** ■ **Content Developer**	Streamlines the management of VMware Aria Suite Lifecycle roles for users.	■ You must create the security groups outside of the SDDC stack. ■ You must set the desired directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.

Workspace ONE Access Design for VMware Cloud Foundation

11

Workspace ONE Access in VMware Cloud Foundation mode provides identity and access management services to specific components in the SDDC, such as VMware Aria Suite.

Workspace ONE Access provides the following capabilities:

- Directory integration to authenticate users against an identity provider (IdP), such as Active Directory or LDAP.
- Multiple authentication methods.
- Access policies that consist of rules to specify criteria that users must meet to authenticate.

The Workspace ONE Access instance that is integrated with VMware Aria Suite Lifecycle provides identity and access management services to VMware Aria Suite solutions that either run in a VMware Cloud Foundation instance or must be available across VMware Cloud Foundation instances.

For identity management design for a VMware Aria Suite product, see VMware Cloud Foundation Validated Solutions .

For identity and access management for components other than VMware Aria Suite, such as NSX, you can deploy a standalone Workspace ONE Access instance. See Identity and Access Management for VMware Cloud Foundation.

Read the following topics next:

Logical Design for Workspace ONE Access

To provide identity and access management services to supported SDDC components, such as VMware Aria Suite components, this design uses a Workspace ONE Access instance that is deployed on an NSX network segment.

Figure 11-1. Logical Design for Standard Workspace ONE Access

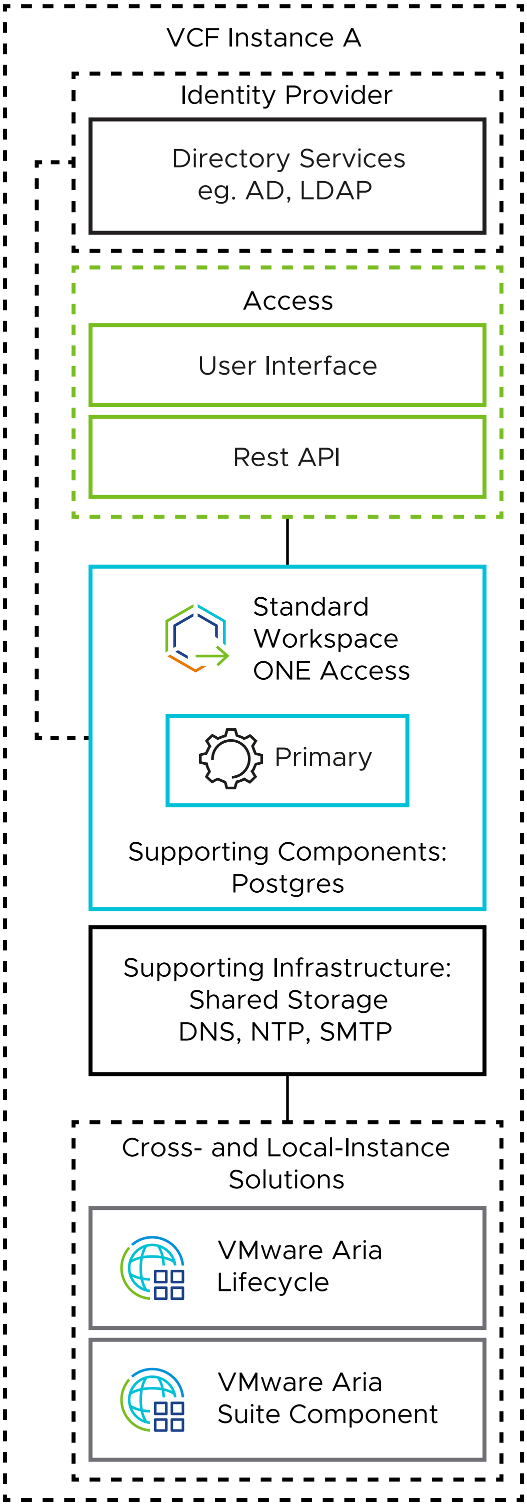

Table 11-1. Logical Components of Standard Workspace ONE Access

VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
■ A single-node Workspace ONE Access instance deployed on an overlay-backed (recommended) or VLAN-backed NSX segment. ■ SDDC solutions that are portable across VMware Cloud Foundation instances are integrated with the Workspace ONE Access instance in the first VMware Cloud Foundation instance.	■ A single-node Workspace ONE Access instance deployed on an overlay-backed (recommended) or VLAN-backed NSX segment. ■ SDDC solutions that are portable across VMware Cloud Foundation instances are integrated with the Workspace ONE Access instance in the first VMware Cloud Foundation instance. ■ A should-run-on-hosts-in-group vSphere DRS rule ensures that, under normal operating conditions, the Workspace ONE Access node runs on a management ESXi host in the first availability zone.

Figure 11-2. Logical Design for Clustered Workspace ONE Access

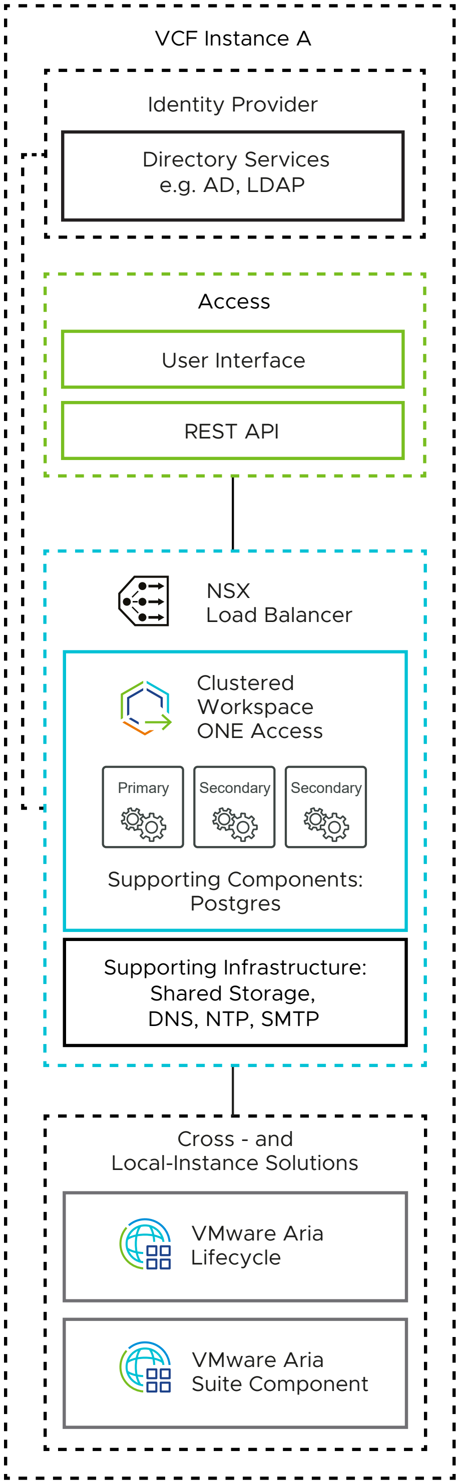

Table 11-2. Logical Components of Clustered Workspace ONE Access

VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
■ A three-node Workspace ONE Access cluster behind an NSX load balancer and deployed on an overlay-backed (recommended) or VLAN-backed NSX segment is deployed in the first VMware Cloud Foundation instance. ■ All Workspace ONE Access services and databases are configured for high availability using a native cluster configuration. SDDC solutions that are portable across VMware Cloud Foundation instances are integrated with this Workspace ONE Access cluster. ■ Each node of the three node cluster is configured as a connector to any relevant identity providers ■ vSphere HA protects the Workspace ONE Access nodes. ■ vSphere DRS anti-affinity rules ensure that the Workspace ONE Access nodes run on different ESXi hosts. ■ Additional single-node Workspace ONE Access instance is deployed on an overlay-backed (recommended) or VLAN-backed NSX segment in all other VMware Cloud Foundation instances.	■ A three-node Workspace ONE Access cluster behind an NSX load balancer and deployed on an overlay-backed (recommended) or VLAN-backed NSX segment. ■ All Workspace ONE Access services and databases are configured for high availability using a native cluster configuration. SDDC solutions that are portable across VMware Cloud Foundation instances are integrated with this Workspace ONE Access cluster. ■ Each node of the three-node cluster is configured as a connector to any relevant identity providers ■ vSphere HA protects the Workspace ONE Access nodes. ■ A vSphere DRS anti-affinity rule ensures that the Workspace ONE Access nodes run on different ESXi hosts. ■ A should-run-on-hosts-in-group vSphere DRS rule ensures that, under normal operating conditions, the Workspace ONE Access nodes run on management ESXi hosts in the first availability zone. ■ Additional single-node Workspace ONE Access instance is deployed on an overlay-backed (recommended) or VLAN-backed NSX segment in all other VMware Cloud Foundation instances.

Sizing Considerations for Workspace ONE Access for VMware Cloud Foundation

When you deploy Workspace ONE Access, you select to deploy the appliance with a size that is suitable for the scale of your environment. The option that you select determines the number of CPUs and the amount of memory of the appliance.

For detailed sizing based on the overall profile of the VMware Cloud Foundation instance you plan to deploy, see VMware Cloud Foundation Planning and Preparation Workbook.

Table 11-3. Sizing Considerations for Workspace ONE Access

Workspace ONE Access Appliance Size	Supported Limits
Extra Small	■ 3,000 users ■ 30 groups
Small	■ 5,000 users ■ 50 groups
Medium	■ 10,000 Users ■ 100 groups A minimum requirement for VMware Aria Automation

Table 11-3. Sizing Considerations for Workspace ONE Access (continued)

Workspace ONE Access Appliance Size	Supported Limits
Large	▪ 25,000 users ▪ 250 groups
Extra Large	▪ 50,000 users ▪ 500 groups
Extra Extra Large	▪ 100,000 users ▪ 1,000 groups

Network Design for Workspace ONE Access

For secure access to the UI and API of Workspace ONE Access, you deploy the nodes on an overlay-backed or VLAN-backed NSX network segment.

Network Segment

This network design has the following features:

- All Workspace ONE Access components have routed access to the management VLAN through the Tier-0 gateway in the NSX instance for the management domain.
- Routing to the management network and other external networks is dynamic and is based on the Border Gateway Protocol (BGP).

Figure 11-3. Network Design for Standard Workspace ONE Access

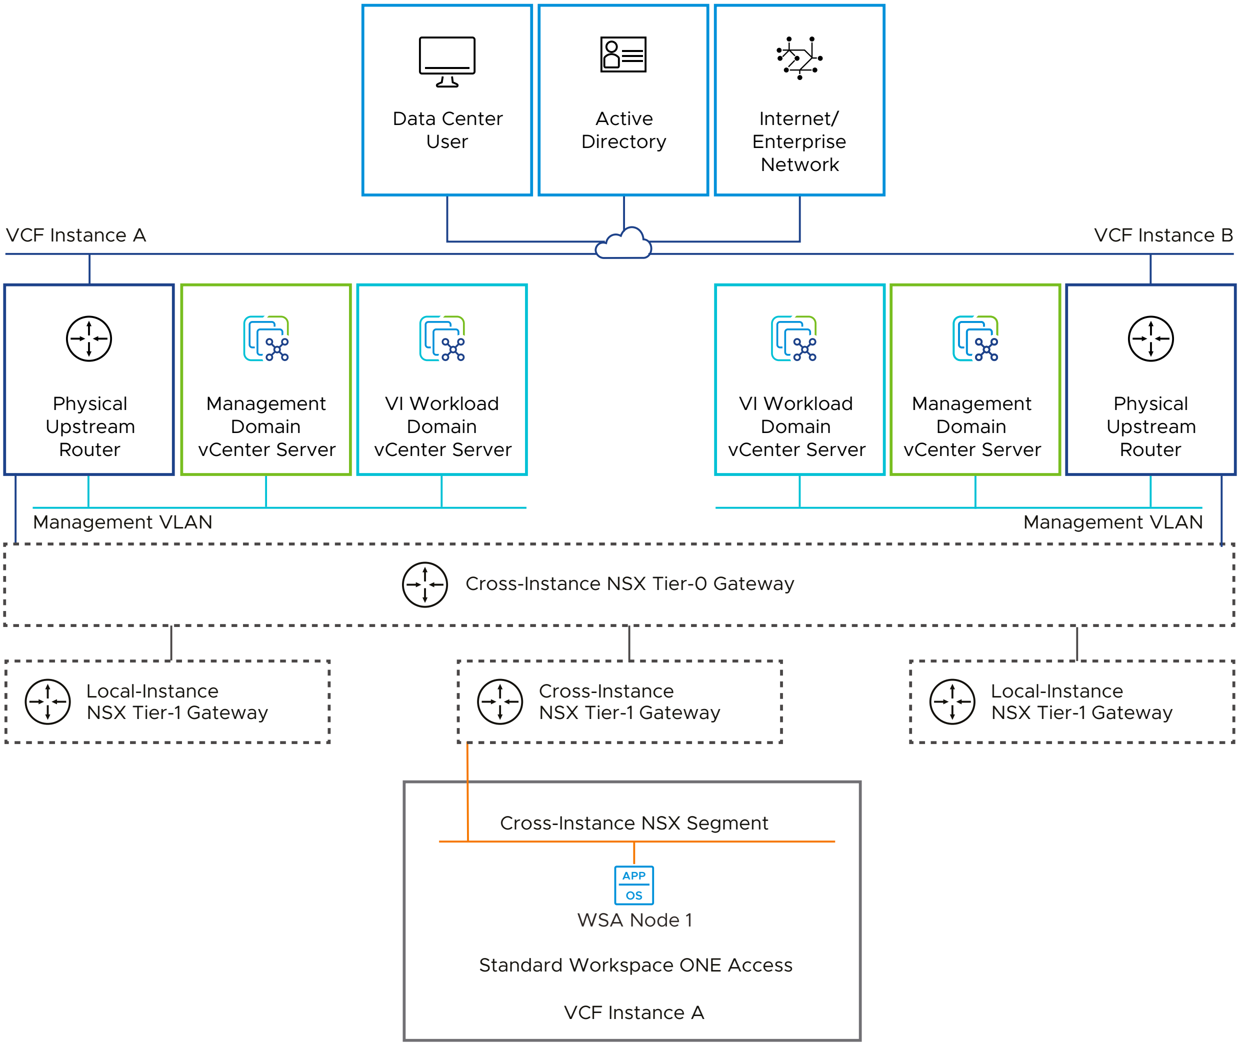
Data Center User
Active Directory
Internet/ Enterprise Network
VCF Instance A
VCF Instance B
Physical Upstream Router
Management Domain vCenter Server
VI Workload Domain vCenter Server
VI Workload Domain vCenter Server
Management Domain vCenter Server
Physical Upstream Router
Management VLAN
Management VLAN
Cross-Instance NSX Tier-0 Gateway
Local-Instance NSX Tier-1 Gateway
Cross-Instance NSX Tier-1 Gateway
Local-Instance NSX Tier-1 Gateway
Cross-Instance NSX Segment
APP
OS
WSA Node 1
Standard Workspace ONE Access
VCF Instance A

Figure 11-4. Network Design for Clustered Workspace ONE Access

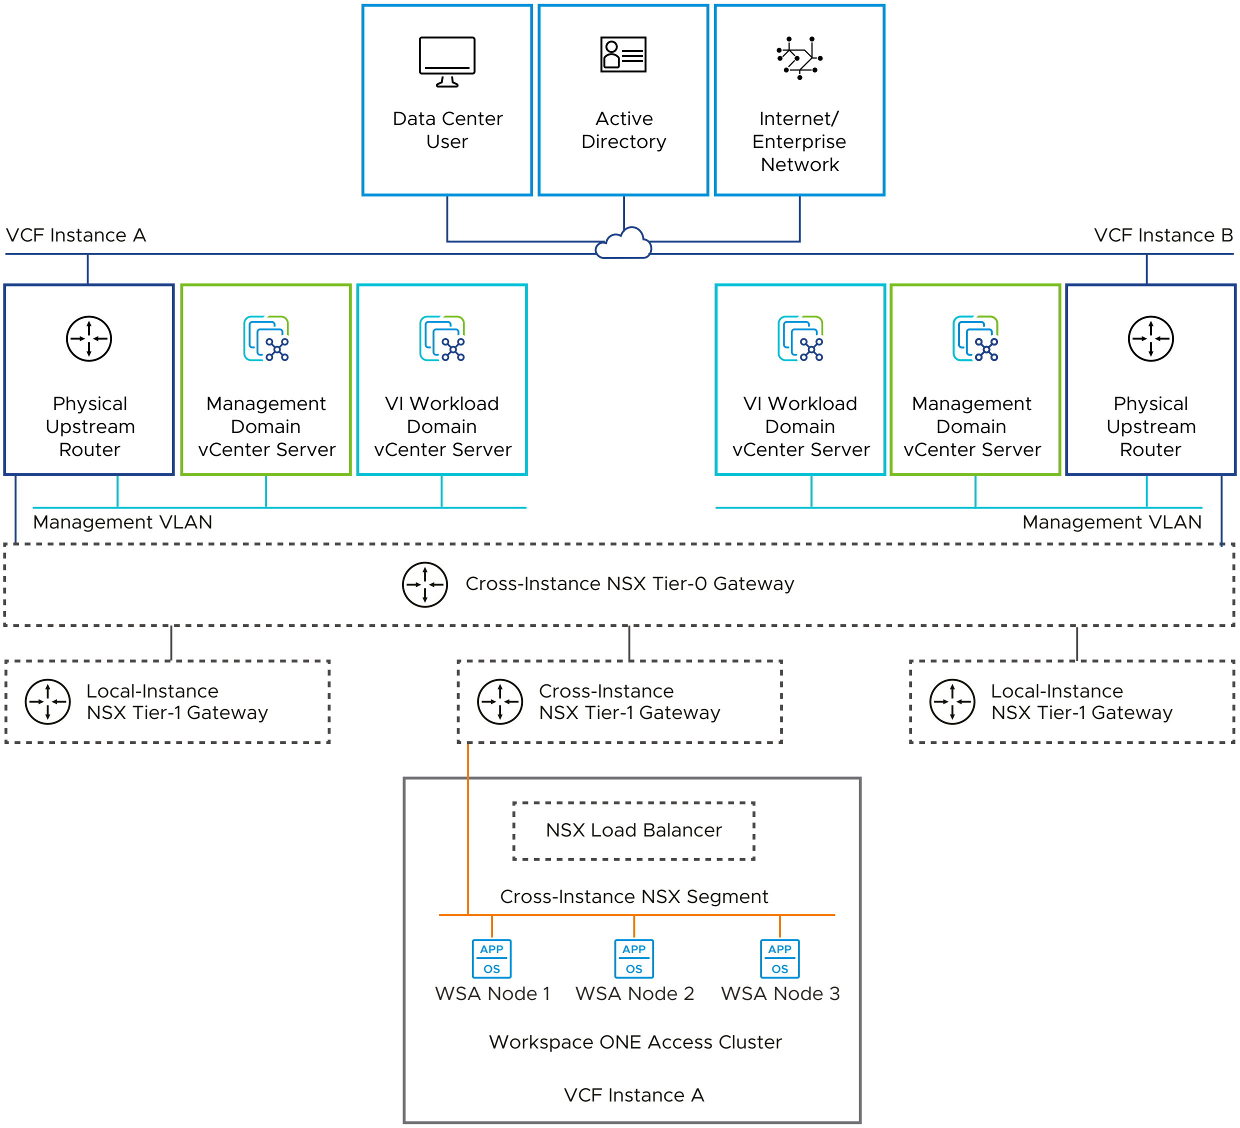
Data Center User
Active Directory
Internet/ Enterprise Network
VCF Instance A
VCF Instance B
Physical Upstream Router
Management Domain vCenter Server
VI Workload Domain vCenter Server
VI Workload Domain vCenter Server
Management Domain vCenter Server
Physical Upstream Router
Management VLAN
Management VLAN
Cross-Instance NSX Tier-0 Gateway
Local-Instance NSX Tier-1 Gateway
Cross-Instance NSX Tier-1 Gateway
Local-Instance NSX Tier-1 Gateway
NSX Load Balancer
Cross-Instance NSX Segment
APP OS
APP OS
APP OS
WSA Node 1
WSA Node 2
WSA Node 3
Workspace ONE Access Cluster
VCF Instance A

Load Balancing

A Workspace ONE Access cluster deployment requires a load balancer to manage connections to the Workspace ONE Access services.

Load-balancing services are provided by NSX. During the deployment of the Workspace ONE Access cluster or scale-out of a standard deployment, VMware Aria Suite Lifecycle and SDDC Manager coordinate to automate the configuration of the NSX load balancer. The load balancer is configured with the following settings:

Table 11-4. Clustered Workspace ONE Access Load Balancer Configuration

Load Balancer Element	Settings
Service Monitor	■ Use the default intervals and timeouts: ■ Monitoring interval: 3 seconds ■ Idle timeout period: 10 seconds ■ Rise/Fall: 3 seconds ■ HTTP request: ■ HTTP method: Get ■ HTTP request version: 1.1 ■ Request URL: /SAAS/API/1.0/REST/system/health/heartbeat. ■ HTTP response: ■ HTTP response code: 200 ■ HTTP response body: OK ■ SSL configuration: ■ Server SSL: Enabled ■ Client certificate: Cross-instance Workspace ONE Access cluster certificate ■ SSL profile: default-balanced-server-ssl-profil.
Server Pool	■ LEAST_CONNECTION algorithm. ■ Set the SNAT translation mode to Auto Map for the pool. ■ Static members: ■ Name: host name IP: ■ IP address ■ Port: 443 ■ Weight: 1 ■ State: Enabled ■ Set the above service monitor.
HTTP Application Profile	■ Timeout ■ 3,600 seconds (60 minutes). ■ X-Forwarded-For ■ Insert

Table 11-4. Clustered Workspace ONE Access Load Balancer Configuration (continued)

Load Balancer Element	Settings
Cookie Persistence Profile	■ Cookie name ■ JSESSIONID. ■ Cookie mode ■ Rewrite
Virtual Server	■ HTTP type ■ L7 ■ Port ■ 443 ■ IP ■ Workspace ONE Access cluster IP ■ Persistence ■ Above Cookie Persistence Profile. ■ Application profile ■ Above HTTP application profile. ■ Server pool ■ Above server pool

Integration Design for Workspace ONE Access with VMware Cloud Foundation

You integrate supported SDDC components with the Workspace ONE Access cluster to enable authentication through the identity and access management services.

After the integration, information security and access control configurations for the integrated SDDC products can be configured.

Table 11-5. Workspace ONE Access SDDC Integration

SDDC Component	Integration	Considerations
vCenter Server	Not Supported	For directory services you must connect vCenter Server directly to Active Directory. See Identity and Access Management for VMware Cloud Foundation.
SDDC Manager	Not Supported	SDDC Manager uses vCenter Single Sign-On. For directory services, you must connect vCenter Server directly to Active Directory

Table 11-5. Workspace ONE Access SDDC Integration (continued)

SDDC Component	Integration	Considerations
NSX	Supported	If you intend to scale out to an environment with multiple VMware Cloud Foundation instances, for example, for disaster recovery, you must deploy an additional standard instance of Workspace ONE Access in each VMware Cloud Foundation instance. The Workspace ONE Access instance that is leveraged by components protected across VMware Cloud Foundation instances might fail over between physical locations which will impact the authentication to NSX in the first VMware Cloud Foundation instance. See Identity and Access Management for VMware Cloud Foundation.
VMware Aria Suite Lifecycle	Supported	None.

See VMware Cloud Foundation Validated Solutions for the design for specific VMware Aria Suite components including identity management.

Deployment Model for Workspace ONE Access

Workspace ONE Access is distributed as a virtual appliance in OVA format that you can deploy and manage from VMware Aria Suite Lifecycle together with other VMware Aria Suite products. The Workspace ONE Access appliance includes identity and access management services.

Deployment Type

You consider the deployment type, standard or cluster, according to the design objectives for the availability and number of users that the system and integrated SDDC solutions must support. You deploy Workspace ONE Access on the default management vSphere cluster.

Table 11-6. Topology Attributes of Workspace ONE Access

Deployment Type	Description	Benefit	Drawbacks
Standard (Recommended)	■ Single node ■ NSX load balancer automatically deployed.	■ Can be scaled out to a 3-node cluster behind an NSX load balancer ■ Can leverage vSphere HA for recovery after a failure occurs. ■ Consumes less resources.	■ Does not provide high availability for Identity Provider connectors.
Cluster	■ Three node clustered deployment using internal PostgreSQL database. ■ NSX load balancer automatically deployed.	■ Provides high availability for Identity Provider connectors.	■ May require manual intervention after a failure occurs. ■ Consumes additional resources.

Workspace ONE Access Design Requirements and Recommendations for VMware Cloud Foundation

Consider the placement, networking, sizing and high availability requirements for using Workspace ONE Access for identity and access management of SDDC solutions on a standard or stretched management cluster in VMware Cloud Foundation. Apply similar best practices for having Workspace ONE Access operate in an optimal way.

Workspace ONE Access Design Requirements

You must meet the following design requirements in your Workspace ONE Access design for VMware Cloud Foundation, considering standard or stretched clusters. For NSX Federation, additional requirements exist.

Table 11-7. Workspace ONE Access Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-ENV-001	Create a global environment in VMware Aria Suite Lifecycle to support the deployment of Workspace ONE Access.	A global environment is required by VMware Aria Suite Lifecycle to deploy Workspace ONE Access.	None.
VCF-WSA-REQD-SEC-001	Import certificate authority-signed certificates to the Locker repository for Workspace ONE Access product life cycle operations.	■ You can reference and use certificate authority-signed certificates during product life cycle operations, such as deployment and certificate replacement.	When using the API, you must specify the Locker ID for the certificate to be used in the JSON payload.

Table 11-7. Workspace ONE Access Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-001	Deploy an appropriately sized Workspace ONE Access instance according to the deployment model you have selected by using VMware Aria Suite Lifecycle in VMware Cloud Foundation mode.	The Workspace ONE Access instance is managed by VMware Aria Suite Lifecycle and imported into the SDDC Manager inventory.	None.
VCF-WSA-REQD-CFG-002	Place the Workspace ONE Access appliances on an overlay-backed or VLAN-backed NSX network segment.	Provides a consistent deployment model for management applications in an environment with a single or multiple VMware Cloud Foundation instances.	You must use an implementation in NSX to support this network configuration.
VCF-WSA-REQD-CFG-003	Use the embedded PostgreSQL database with Workspace ONE Access.	Removes the need for external database services.	None.
VCF-WSA-REQD-CFG-004	Add a VM group for Workspace ONE Access and set VM rules to restart the Workspace ONE Access VM group before any of the VMs that depend on it for authentication.	You can define the startup order of virtual machines regarding the service dependency. The startup order ensures that vSphere HA powers on the Workspace ONE Access virtual machines in an order that respects product dependencies.	None.
VCF-WSA-REQD-CFG-005	Connect the Workspace ONE Access instance to a supported upstream Identity Provider.	You can integrate your enterprise directory with Workspace ONE Access to synchronize users and groups to the Workspace ONE Access identity and access management services.	None.

Table 11-7. Workspace ONE Access Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-006	If using clustered Workspace ONE Access, configure second and third native connectors that correspond to the second and third Workspace ONE Access cluster nodes to support the high availability of directory services access.	Adding the additional native connectors provides redundancy and improves performance by load-balancing authentication requests.	Each of the Workspace ONE Access cluster nodes must be joined to the Active Directory domain to use Active Directory with Integrated Windows Authentication with the native connector.
VCF-WSA-REQD-CFG-007	If using clustered Workspace ONE Access, use the NSX load balancer that is configured by SDDC Manager on a dedicated Tier-1 gateway.	■ During the deployment of Workspace ONE Access by using VMware Aria Suite Lifecycle, SDDC Manager automates the configuration of an NSX load balancer for Workspace ONE Access to facilitate scale-out.	You must use the load balancer that is configured by SDDC Manager and the integration with VMware Aria Suite Lifecycle.

Table 11-8. Workspace ONE Access Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-008	Add the Workspace ONE Access appliances to the VM group for the first availability zone.	Ensures that, by default, the Workspace ONE Access cluster nodes are powered on a host in the first availability zone.	■ If the Workspace ONE Access instance is deployed after the creation of the stretched management cluster, you must add the appliances to the VM group manually. ■ ClusteredWorkspace ONE Access might require manual intervention after a failure of the active availability zone occurs.

Table 11-9. Workspace ONE Access Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-009	Configure the DNS settings for Workspace ONE Access to use DNS servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	None.
VCF-WSA-REQD-CFG-010	Configure the NTP settings on Workspace ONE Access cluster nodes to use NTP servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	If you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the NTP settings on Workspace ONE Access must be updated.

Workspace ONE Access Design Recommendations

In your Workspace ONE Access design for VMware Cloud Foundation, you can apply certain best practices.

Table 11-10. Workspace ONE Access Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-CFG-001	Protect all Workspace ONE Access nodes using vSphere HA.	Supports high availability for Workspace ONE Access.	None for standard deployments. Clustered Workspace ONE Access deployments might require intervention if an ESXi host failure occurs.
VCF-WSA-RCMD-CFG-002	When using Active Directory as an Identity Provider, use Active Directory over LDAP as the Directory Service connection option.	The native (embedded) Workspace ONE Access connector binds to Active Directory over LDAP using a standard bind authentication.	■ In a multi-domain forest, where the Workspace ONE Access instance connects to a child domain, Active Directory security groups must have global scope. Therefore, members added to the Active Directory global security group must reside within the same Active Directory domain. ■ If authentication to more than one Active Directory domain is required, additional Workspace ONE Access directories are required.
VCF-WSA-RCMD-CFG-003	When using Active Directory as an Identity Provider, use an Active Directory user account with a minimum of read-only access to Base DNs for users and groups as the service account for the Active Directory bind.	Provides the following access control features: ■ Workspace ONE Access connects to the Active Directory with the minimum set of required permissions to bind and query the directory. ■ You can introduce improved accountability in tracking request-response interactions between the Workspace ONE Access and Active Directory.	■ You must manage the password life cycle of this account. ■ If authentication to more than one Active Directory domain is required, additional accounts are required for the Workspace ONE Access connector to bind to each Active Directory domain over LDAP.

Table 11-10. Workspace ONE Access Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-CFG-004	Configure the directory synchronization to synchronize only groups required for the integrated SDDC solutions.	■ Limits the number of replicated groups required for each product. ■ Reduces the replication interval for group information.	You must manage the groups from your enterprise directory selected for synchronization to Workspace ONE Access.
VCF-WSA-RCMD-CFG-005	Activate the synchronization of enterprise directory group members when a group is added to the Workspace ONE Access directory.	When activated, members of the enterprise directory groups are synchronized to the Workspace ONE Access directory when groups are added. When deactivated, group names are synchronized to the directory, but members of the group are not synchronized until the group is entitled to an application or the group name is added to an access policy.	None.
VCF-WSA-RCMD-CFG-006	Enable Workspace ONE Access to synchronize nested group members by default.	Allows Workspace ONE Access to update and cache the membership of groups without querying your enterprise directory.	Changes to group membership are not reflected until the next synchronization event.
VCF-WSA-RCMD-CFG-007	Add a filter to the Workspace ONE Access directory settings to exclude users from the directory replication.	Limits the number of replicated users for Workspace ONE Access within the maximum scale.	To ensure that replicated user accounts are managed within the maximums, you must define a filtering schema that works for your organization based on your directory attributes.

Table 11-10. Workspace ONE Access Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-CFG-008	Configure the mapped attributes included when a user is added to the Workspace ONE Access directory.	You can configure the minimum required and extended user attributes to synchronize directory user accounts for the Workspace ONE Access to be used as an authentication source for cross-instance VMware Aria Suite solutions.	User accounts in your organization's enterprise directory must have the following required attributes mapped: ■ `firstname`, for example, `givenname` for Active Directory ■ `lastName`, for example, `sn` for Active Directory ■ `email`, for example, `mail` for Active Directory ■ `userName`, for example,`sAMAccountName` for Active Directory ■ If you require users to sign in with an alternate unique identifier, for example, `userPrincipalName`, you must map the attribute and update the identity and access management preferences.
VCF-WSA-RCMD-CFG-009	Configure the Workspace ONE Access directory synchronization frequency to a reoccurring schedule, for example, 15 minutes.	Ensures that any changes to group memberships in the corporate directory are available for integrated solutions in a timely manner.	Schedule the synchronization interval to be longer than the time to synchronize from the enterprise directory. If users and groups are being synchronized to Workspace ONE Access when the next synchronization is scheduled, the new synchronization starts immediately after the end of the previous iteration. With this schedule, the process is continuous.

Table 11-10. Workspace ONE Access Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-SEC-001	Create corresponding security groups in your corporate directory services for these Workspace ONE Access roles: - **Super Admin** - **Directory Admins** - **ReadOnly Admin**	Streamlines the management of Workspace ONE Access roles to users.	- You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period. - You must create the security group outside of the SDDC stack.
VCF-WSA-RCMD-SEC-002	Configure a password policy for Workspace ONE Access local directory users, **admin** and **configadmin**.	You can set a policy for Workspace ONE Access local directory users that addresses your corporate policies and regulatory standards. The password policy is applicable only to the local directory users and does not impact your organization directory.	You must set the policy in accordance with your organization policies and regulatory standards, as applicable. You must apply the password policy on the Workspace ONE Access cluster nodes.

Life Cycle Management Design for VMware Cloud Foundation

12

In a VMware Cloud Foundation instance, you use SDDC Manager for life cycle management of the management components in the entire instance except for NSX Global Manager and VMware Aria Suite Lifecycle. VMware Aria Suite Lifecycle manages the life cycle of the components that it deploys.

Life cycle management of a VMware Cloud Foundation instance is the process of performing patch updates or upgrades to the underlying management components.

Table 12-1. Life Cycle Management for VMware Cloud Foundation

<table>
<tr><th>Component</th><th>Management Domain</th><th>VI Workload Domain</th></tr>
<tr><td>SDDC Manager</td><td>SDDC Manager performs its own life cycle management.</td><td>Not applicable</td></tr>
<tr><td>NSX Local Manager</td><td colspan="2">SDDC Manager uses the NSX upgrade coordinator service in the NSX Local Manager.</td></tr>
<tr><td>NSX Edges</td><td colspan="2">SDDC Manager uses the NSX upgrade coordinator service in NSX Manager.</td></tr>
<tr><td>NSX Global Manager</td><td colspan="2">You manually use the NSX upgrade coordinator service in the NSX Global Manager.</td></tr>
<tr><td>vCenter Server</td><td colspan="2">You use SDDC Manager for life cycle management of all vCenter Server instances.</td></tr>
<tr><td>ESXi</td><td>■ SDDC Manager uses either vSphere Lifecycle Manager baselines and baseline groups or vSphere Lifecycle Manager images to update and upgrade the ESXi hosts.
■ Custom vendor ISOs are supported and might be required depending on the ESXi hardware in use.</td><td>■ SDDC Manager uses either vSphere Lifecycle Manager baselines and baseline groups or vSphere Lifecycle Manager images to update and upgrade the ESXi hosts.
■ Custom vendor ISOs are supported and might be required depending on the ESXi hardware in use.</td></tr>
<tr><td>VMware Aria Suite Lifecycle</td><td>VMware Aria Suite Lifecycle performs its own life cycle management.</td><td>Not applicable</td></tr>
</table>

VMware Cloud Foundation Life Cycle Management Requirements

Consider the design requirements for automated and centralized life cycle management in the context of the entire VMware Cloud Foundation environment.

Table 12-2. Life Cycle Management Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-001	Use SDDC Manager to perform the life cycle management of the following components: - SDDC Manager - NSX Manager - NSX Edges - vCenter Server - ESXi	Because the deployment scope of SDDC Manager covers the full VMware Cloud Foundation stack, SDDC Manager performs patching, update, or upgrade of these components across all workload domains.	The operations team must understand and be aware of the impact of a patch, update, or upgrade operation by using SDDC Manager.
VCF-LCM-REQD-002	Use VMware Aria Suite Lifecycle to manage the life cycle of the following components: - VMware Aria Suite Lifecycle - Workspace ONE Access	VMware Aria Suite Lifecycle automates the life cycle of VMware Aria Suite Lifecycle and Workspace ONE Access.	- You must deploy VMware Aria Suite Lifecycle by using SDDC Manager. - You must manually apply Workspace ONE Access patches, updates, and hotfixes. Patches, updates, and hotfixes for Workspace ONE Access are not generally managed by VMware Aria Suite Lifecycle.
VCF-LCM-RCMD-001	Use vSphere Lifecycle Manager images to manage the life cycle of vSphere clusters.	- With vSphere Lifecycle Manager images, firmware updates are carried out through firmware and driver add-ons, which you add to the image you use to manage a cluster. - You can check the hardware compatibility of the hosts in a cluster against the VMware Compatibility Guide. - You can validate a vSphere Lifecycle Manager image to check if it applies to all hosts in the cluster. You can also perform a remediation pre-check.	- Updating the firmware with images requires an OEM-provided hardware support manager plug-in, which integrates with vSphere Lifecycle Manager. - An updated vSAN Hardware Compatibility List (vSAN HCL) is required during bring-up.

Table 12-3. Life Cycle Management Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-003	Use the upgrade coordinator in NSX to perform life cycle management on the NSX Global Manager appliances.	The version of SDDC Manager in this design is not currently capable of life cycle operations (patching, update, or upgrade) for NSX Global Manager.	■ You must explicitly plan upgrades of the NSX Global Manager nodes. An upgrade of the NSX Global Manager nodes might require a cascading upgrade of the NSX Local Manager nodes and underlying SDDC Manager infrastructure before upgrading the NSX Global Manager nodes. ■ You must always align the version of the NSX Global Manager nodes with the rest of the SDDC stack in VMware Cloud Foundation.
VCF-LCM-REQD-004	Establish an operations practice to ensure that prior to the upgrade of any workload domain, the impact of any version upgrades is evaluated in relation to the need to upgrade NSX Global Manager.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.
VCF-LCM-REQD-005	Establish an operations practice to ensure that prior to the upgrade of the NSX Global Manager, the impact of any version change is evaluated against the existing NSX Local Manager nodes and workload domains.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.

Logging and Monitoring Design for VMware Cloud Foundation

13

By using VMware or third-party components, collect log data from all SDDC management components in your VMware Cloud Foundation environment in a central place. You can use VMware Aria Operations for Logs as the central platform because of its native integration with VMware Aria Suite Lifecycle.

After you deploy VMware Aria Operations for Logs by using VMware Aria Suite Lifecycle in VMware Cloud Foundation mode, SDDC Manager configures VMware Aria Suite Lifecycle logging to VMware Aria Operations for Logs over the log ingestion API. For information about on-premises VMware Aria Operations for Logs in VMware Cloud Foundation, see Intelligent Logging and Analytics for VMware Clod Foundation.

Information Security Design for VMware Cloud Foundation

14

You design management of access controls, certificates and accounts for VMware Cloud Foundation according to the requirements of your organization.

Read the following topics next:

Access Management for VMware Cloud Foundation

You design access management for VMware Cloud Foundation according to industry standards and the requirements of your organization.

Component	Access Method	Additional Information
SDDC Manager	■ UI ■ API ■ SSH	SSH is active by default. **root** user access is deactivated.
NSX Local Manager	■ UI ■ API ■ SSH	SSH is deactivated by default.
NSX Edges	■ API ■ SSH	SSH is deactivated by default.
NSX Global Manager	■ UI ■ API ■ SSH	SSH setting is defined during deployment.
vCenter Server	■ UI ■ API ■ SSH ■ VAMI	SSH is active by default.

Component	Access Method	Additional Information
ESXi	■ Direct Console User Interface (DCUI) ■ ESXi Shell ■ SSH ■ VMware Host Client	SSH and ESXi shell are deactivated by default.
VMware Aria Suite Lifecycle	■ UI ■ API ■ SSH	SSH is active by default.
Workspace ONE Access	■ UI ■ API ■ SSH	SSH is active by default.

Account Management Design for VMware Cloud Foundation

You design account management for VMware Cloud Foundation according to industry standards and the requirements of your organization.

Password Management Methods

SDDC Manager manages the life cycle of passwords for the components that are part of the VMware Cloud Foundation instance. Multiple methods for managing password life cycle are supported.

Table 14-1. Password Management Methods in VMware Cloud Foundation

Method	Description
Rotate	Update one or more accounts with an auto-generated password
Update	Update password for a single account with a manually entered password
Remediate	Reconcile a single account with a password that has been set manually at the component.
Schedule	Schedule auto-rotation for one or more selected accounts.
Manual	Update a password manually directly in the component.

Account and Password Management

VMware Cloud Foundation comprises multiple types of interactive, local, and service accounts. Each account has different attributes and can be managed in the following ways:

For more information on password complexity, account lockout or integration with additional Identity Providers, refer to the Identity and Access Management for VMware Cloud Foundation.

Table 14-2. Account and Password Management in VMware Cloud Foundation

Component	User Account	Password Management	Additional Information
SDDC Manager	admin@local	■ Manual by using the SDDC Manager API ■ Default Expiry: Never	■ Local appliance account ■ API access (break-glass account)
	vcf	■ Manual by using the OS ■ Default Expiry: 365 days	■ Local appliance account ■ OS level access
	root	■ Manual by using the OS ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access
	backup	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 365 days	■ Local appliance account ■ OS level access
	administrator@vsphere.local	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ vCenter Single Sign-On account. ■ Application and API access. ■ Additional VMware Cloud Foundation**Admin** account required to perform manual password rotation.
NSX Local Manager	admin	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level, API, and application access
	root	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access
	audit	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access ■ Read-only application level access

Table 14-2. Account and Password Management in VMware Cloud Foundation (continued)

Component	User Account	Password Management	Additional Information
NSX Edges	admin	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level, API, and application access
	root	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access
	audit	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access ■ Read-only application level access
NSX Global Manager	admin	■ Manual by using the NSX Global Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level, API, and application access
	root	■ Manual by using each NSX Global Manager appliance ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access
	audit	■ Manual by using the NSX Global Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access ■ Read-only application level access
vCenter Server	root	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ Local appliance account ■ OS level access ■ VAMI access
	administrator@isolatedsso.local	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 90 days	■ vCenter Single Sign-On account. ■ Application and API access. ■ Relevant to isolated workload domain

Table 14-2. Account and Password Management in VMware Cloud Foundation (continued)

Component	User Account	Password Management	Additional Information
	svc-*sddc-manager-hostname-vcenter-server-hostname*@vsphere.local	■ System managed. ■ Automatically rotated every 30 days by default ■ Default Expiry: None	Service account between SDDC Manager and vCenter Server
	svc-*nsx-manager-hostname-vcenter-server-hostname*@vsphere.local	■ System managed. ■ Automatically rotated every 30 days by default ■ Default Expiry: None	Service account between NSX Manager and vCenter Server
	svc-*vrslcm-hostname-vcenter-server-hostname*@vsphere.local	■ System managed ■ Automatically rotated every 30 days by default ■ Default Expiry: None	Service account between VMware Aria Suite Lifecycle and vCenter Server
ESXi	root	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 99999 (never)	Manual
	svc-vcf-*esxi-hostname*	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 99999 (never)	Service account between SDDC Manager and the ESXi host
VMware Aria Suite Lifecycle	vcfadmin@local	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: Never	API and application access
	root	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 365 days	■ Local appliance account ■ OS level access

Table 14-2. Account and Password Management in VMware Cloud Foundation (continued)

Component	User Account	Password Management	Additional Information
Workspace ONE Access	root	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: 60 days	■ Local appliance account ■ OS level access
	sshuser	■ Managed by VMware Aria Suite Lifecycle ■ Default Expiry: 60 days	■ Local appliance account ■ OS level access
	admin (port 8443)	Managed by VMware Aria Suite Lifecycle	System Admin
	Admin (port 443)	■ Rotate, update,remediate or schedule by using the SDDC Manager UI or API ■ Default Expiry: Never	Default application administrator
	configadmin	■ You must use both Workspace ONE Access and VMware Aria Suite Lifecycle to manage the password rotation schedule of the **configadmin** user. ■ Default Expiry: Never	Application configuration administrator

Account Management Design Recommendations

In your account management design, you can apply certain best practices.

Table 14-3. Design Requirements for Account and Password Management for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-ACTMGT-REQD-SEC-001	Enable scheduled password rotation in SDDC Manager for all accounts supporting scheduled rotation.	■ Increases the security posture of your SDDC. ■ Simplifies password management across your SDDC management components.	You must retrieve new passwords by using the API if you must use accounts interactively.
VCF-ACTMGT-REQD-SEC-003	Establish operational practice to rotate passwords using SDDC Manager on components that do not support scheduled rotation in SDDC Manager.	Rotates passwords and automatically remediates SDDC Manager databases for those user accounts.	None.
VCF-ACTMGT-REQD-SEC-003	Establish operational practice to manually rotate passwords on components that cannot be rotated by SDDC Manager.	Maintains password policies across components not handled by SDDC Manager password management.	None.

Certificate Management for VMware Cloud Foundation

You design certificate management for VMware Cloud Foundation according to industry standards and the requirements of your organization.

Access to all management component interfaces must be over a Secure Socket Layer (SSL) connection. During deployment, each component is assigned a certificate from a default signing CA. To provide secure access to each component, replace the default certificate with a trusted enterprise CA-signed certificate.

Table 14-4. Certificate Management in VMware Cloud Foundation

Component	Default Signing CA	Life cycle for Enterprise CA-Signed Certificates
SDDC Manager	Management domain VMCA	Using SDDC Manager
NSX Local Manager	Management domain VMCA	Using SDDC Manager
NSX Edges	Not applicable	Not applicable
NSX Global Manager	Self Signed	Manual
vCenter Server	Local workload domain VMCA	Using SDDC Manager

Table 14-4. Certificate Management in VMware Cloud Foundation (continued)

Component	Default Signing CA	Life cycle for Enterprise CA-Signed Certificates
ESXi	Local workload domain VMCA	Manual*
VMware Aria Suite Lifecycle	Management domain VMCA	Using SDDC Manager

Note * To use enterprise CA-Signed certificates with ESXi, the initial deployment of VMware Cloud Foundation must be done using the API providing the Trusted Root certificate.

Table 14-5. Certificate Management Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-SDDC-RCMD-SEC-001	Replace the default VMCA or signed certificates on all management virtual appliances with a certificate that is signed by an internal certificate authority.	Ensures that the communication to all management components is secure.	Replacing the default certificates with trusted CA-signed certificates from a certificate authority might increase the deployment preparation time because you must generate and submit certificate requests.
VCF-SDDC-RCMD-SEC-002	Use a SHA-2 algorithm or higher for signed certificates.	The SHA-1 algorithm is considered less secure and has been deprecated.	Not all certificate authorities support SHA-2 or higher.
VCF-SDDC-RCMD-SEC-003	Perform SSL certificate life cycle management for all management appliances by using SDDC Manager.	SDDC Manager supports automated SSL certificate lifecycle management rather than requiring a series of manual steps.	Certificate management for NSX Global Manager instances must be done manually.

Appendix: Design Elements for VMware Cloud Foundation

15

The appendix aggregates all design requirements and recommendations in the design guidance for VMware Cloud Foundation. You can use this list for reference related to the end state of your platform and potentially to track your level of adherence to the design and any justification for deviation.

Read the following topics next:

- Architecture Design Elements for VMware Cloud Foundation
- Workload Domain Design Elements for VMware Cloud Foundation
- External Services Design Elements for VMware Cloud Foundation
- Physical Network Design Elements for VMware Cloud Foundation
- vSAN Design Elements for VMware Cloud Foundation
- ESXi Design Elements for VMware Cloud Foundation
- vCenter Server Design Elements for VMware Cloud Foundation
- vSphere Cluster Design Elements for VMware Cloud Foundation
- vSphere Networking Design Elements for VMware Cloud Foundation
- NSX Design Elements for VMware Cloud Foundation
- SDDC Manager Design Elements for VMware Cloud Foundation
- VMware Aria Suite Lifecycle Design Elements for VMware Cloud Foundation
- Workspace ONE Access Design Elements for VMware Cloud Foundation
- Life Cycle Management Design Elements for VMware Cloud Foundation
- Information Security Design Elements for VMware Cloud Foundation

Architecture Design Elements for VMware Cloud Foundation

Use this list of requirements for reference related to using the standard or consolidated architecture of VMware Cloud Foundation.

For full design details, see Architecture Models and Workload Domain Types in VMware Cloud Foundation.

Table 15-1. Architecture Model Recommendations for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-ARCH-RCMD-CFG-001	Use the standard architecture model of VMware Cloud Foundation.	■ Aligns with the VMware best practice of separating management workloads from customer workloads. ■ Provides better long-term flexibility and expansion options.	Requires additional hardware.

Workload Domain Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the types of virtual infrastructure (VI) workload domains in a VMware Cloud Foundation environment.

For full design details, see Workload Domain Types.

Table 15-2. Workload Domain Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WLD-RCMD-CFG-001	Use VI workload domains or isolated VI workload domains for customer workloads.	■ Aligns with the VMware best practice of separating management workloads from customer workloads. ■ Provides better long term flexibility and expansion options.	Requires additional hardware.

External Services Design Elements for VMware Cloud Foundation

Use this list of requirements for reference related to the configuration of external infrastructure services in an environment with a single or multiple VMware Cloud Foundation instances. The requirements define IP address allocation, name resolution, and time synchronization.

For full design details, see Chapter 4 External Services Design for VMware Cloud Foundation.

Table 15-3. External Services Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-EXT-REQD-NET-001	Allocate statically assigned IP addresses and host names for all workload domain components.	Ensures stability across the VMware Cloud Foundation instance, and makes it simpler to maintain, track, and implement a DNS configuration.	You must provide precise IP address management.
VCF-EXT-REQD-NET-002	Configure forward and reverse DNS records for all workload domain components.	Ensures that all components are accessible by using a fully qualified domain name instead of by using IP addresses only. It is easier to remember and connect to components across the VMware Cloud Foundation instance.	You must provide DNS records for each component.
VCF-EXT-REQD-NET-003	Configure time synchronization by using an internal NTP time source for all workload domain components.	Ensures that all components are synchronized with a valid time source.	An operational NTP service must be available in the environment.
VCF-EXT-REQD-NET-004	Set the NTP service for all workload domain components to start automatically.	Ensures that the NTP service remains synchronized after you restart a component.	None.

Physical Network Design Elements for VMware Cloud Foundation

Use this design decision list for reference related to the configuration of the physical network in an environment with a single or multiple VMware Cloud Foundation instances. The design also considers if an instance contains a single or multiple availability zones.

For full design details, see Chapter 5 Physical Network Infrastructure Design for VMware Cloud Foundation.

Table 15-4. Leaf-Spine Physical Network Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-REQD-CFG-001	Do not use EtherChannel (LAG, LACP, or vPC) configuration for ESXi host uplinks.	■ Simplifies configuration of top of rack switches. ■ Teaming options available with vSphere Distributed Switch provide load balancing and failover. ■ EtherChannel implementations might have vendor-specific limitations.	None.
VCF-NET-REQD-CFG-002	Use VLANs to separate physical network functions.	■ Supports physical network connectivity without requiring many NICs. ■ Isolates the different network functions in the SDDC so that you can have differentiated services and prioritized traffic as needed.	Requires uniform configuration and presentation on all the trunks that are made available to the ESXi hosts.
VCF-NET-REQD-CFG-003	Configure the VLANs as members of a 802.1Q trunk.	All VLANs become available on the same physical network adapters on the ESXi hosts.	Optionally, the management VLAN can act as the native VLAN.
VCF-NET-REQD-CFG-004	Set the MTU size to at least 1,700 bytes (recommended 9,000 bytes for jumbo frames) on the physical switch ports, vSphere Distributed Switches, vSphere Distributed Switch port groups, and N-VDS switches that support the following traffic types: ■ Overlay (Geneve) ■ vSAN ■ vSphere vMotion	■ Improves traffic throughput. ■ Supports Geneve by increasing the MTU size to a minimum of 1,600 bytes. ■ Geneve is an extensible protocol. The MTU size might increase with future capabilities. While 1,600 bytes is sufficient, an MTU size of 1,700 bytes provides more room for increasing the Geneve MTU size without the need to change the MTU size of the physical infrastructure.	When adjusting the MTU packet size, you must also configure the entire network path (VMkernel network adapters, virtual switches, physical switches, and routers) to support the same MTU packet size. In an environment with multiple availability zones, the MTU must be configured on the entire network path between the zones.

Table 15-5. Leaf-Spine Physical Network Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-REQD-CFG-005	Set the MTU size to at least 1,500 bytes (1,700 bytes preferred; 9,000 bytes recommended for jumbo frames) on the components of the physical network between the VMware Cloud Foundation instances for the following traffic types. ■ NSX Edge RTEP	■ Jumbo frames are not required between VMware Cloud Foundation instances. However, increased MTU improves traffic throughput. ■ Increasing the RTEP MTU to 1,700 bytes minimizes fragmentation for standard-size workload packets between VMware Cloud Foundation instances.	When adjusting the MTU packet size, you must also configure the entire network path, that is, virtual interfaces, virtual switches, physical switches, and routers to support the same MTU packet size.
VCF-NET-REQD-CFG-006	Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 500 ms.	A latency lower than 500 ms is required for NSX Federation.	None.
VCF-NET-REQD-CFG-007	Provide a routed connection between the NSX Manager clusters in VMware Cloud Foundation instances that are connected in an NSX Federation.	Configuring NSX Federation requires connectivity between the NSX Global Manager instances, NSX Local Manager instances, and NSX Edge clusters.	You must assign unique routable IP addresses for each fault domain.

Table 15-6. Leaf-Spine Physical Network Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-001	Use two ToR switches for each rack.	Supports the use of two 10-GbE (25-GbE or greater recommended) links to each server, provides redundancy and reduces the overall design complexity.	Requires two ToR switches per rack which might increase costs.
VCF-NET-RCMD-CFG-002	Implement the following physical network architecture: ■ One 25-GbE (10-GbE minimum) port on each ToR switch for ESXi host uplinks (Host to ToR). ■ Layer 3 device that supports BGP.	■ Provides availability during a switch failure. ■ Provides support for BGP dynamic routing protocol	■ Might limit the hardware choices. ■ Requires dynamic routing protocol configuration in the physical network.
VCF-NET-RCMD-CFG-003	Use a physical network that is configured for BGP routing adjacency.	■ Supports design flexibility for routing multi-site and multi-tenancy workloads. ■ BGP is the only dynamic routing protocol that is supported for NSX Federation. ■ Supports failover between ECMP Edge uplinks.	Requires BGP configuration in the physical network.
VCF-NET-RCMD-CFG-004	Assign persistent IP configurations for NSX tunnel endpoints (TEPs) that use static IP pools instead of dynamic IP pool addressing.	■ Ensures that endpoints have a persistent TEP IP address. ■ In VMware Cloud Foundation, TEP IP assignment by using static IP pools is recommended for all topologies. ■ This configuration removes any requirement for external DHCP services.	If you add more hosts to the cluster, expanding the static IP pools might be required.

Table 15-6. Leaf-Spine Physical Network Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-005	Configure the trunk ports connected to ESXi NICs as trunk PortFast.	Reduces the time to transition ports over to the forwarding state.	Although this design does not use the STP, switches usually have STP configured by default.
VCF-NET-RCMD-CFG-006	Configure VRRP, HSRP, or another Layer 3 gateway availability method for these networks. ■ Management ■ Edge overlay	Ensures that the VLANs that are stretched between availability zones are connected to a highly- available gateway. Otherwise, a failure in the Layer 3 gateway will cause disruption in the traffic in the SDN setup.	Requires configuration of a high availability technology for the Layer 3 gateways in the data center.

Table 15-7. Leaf-Spine Physical Network Design Recommendations for NSX Federation in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-007	Provide BGP routing between all VMware Cloud Foundation instances that are connected in an NSX Federation setup.	BGP is the supported routing protocol for NSX Federation.	None.
VCF-NET-RCMD-CFG-008	Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 150 ms for workload mobility.	A latency lower than 150 ms is required for the following features: ■ Cross vCenter Server vMotion	None.

vSAN Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to shared storage, vSAN principal storage, and NFS supplemental storage in an environment with a single or multiple VMware Cloud Foundation instances. The design also considers whether an instance contains a single or multiple availability zones.

After you set up the physical storage infrastructure, the configuration tasks for most design decisions are automated in VMware Cloud Foundation. You must perform the configuration manually only for a limited number of design elements as noted in the design implication.

For full design details, see Chapter 6 vSAN Design for VMware Cloud Foundation.

Table 15-8. vSAN Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-001	Provide sufficient raw capacity to meet the initial needs of the workload domain cluster.	Ensures that sufficient resources are present to create the workload domain cluster.	None.
VCF-VSAN-REQD-CFG-002	Provide at least the required minimum number of hosts according to the cluster type.	Satisfies the requirements for storage availability.	None.

Table 15-9. vSAN ESA Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-003	Verify the hardware components used in your vSAN deployment are on the vSAN Hardware Compatibility List.	Prevents hardware-related failures during workload deployment	Limits the number of compatible hardware configurations that can be used.

Table 15-10. vSAN Design Requirements for Stretched Clusters with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-004	Add the following setting to the default vSAN storage policy: Site disaster tolerance = Site mirroring - stretched cluster	Provides the necessary protection for virtual machines in each availability zone, with the ability to recover from an availability zone outage.	You might need additional policies if third-party virtual machines are to be hosted in these clusters because their performance or availability requirements might differ from what the default VMware vSAN policy supports.
VCF-VSAN-REQD-CFG-005	Configure two fault domains, one for each availability zone. Assign each host to their respective availability zone fault domain.	Fault domains are mapped to availability zones to provide logical host separation and ensure a copy of vSAN data is always available even when an availability zone goes offline.	You must provide additional raw storage when the site mirroring - stretched cluster option is selected, and fault domains are enabled.
VCF-VSAN-REQD-CFG-006	Use vSAN OSA to create a stretched cluster.	Stretched clusters on top of vSAN ESA are not supported by VMware Cloud Foundation	None.
VCF-VSAN-REQD-CFG-007	Configure an individual vSAN storage policy for each stretched cluster.	The vSAN storage policy of a stretched cluster cannot be shared with other clusters.	You must configure additional vSAN storage policies.

Table 15-10. vSAN Design Requirements for Stretched Clusters with VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-WTN-REQD-CFG-001	Deploy a vSAN witness appliance in a location that is not local to the ESXi hosts in any of the availability zones.	Ensures availability of vSAN witness components in the event of a failure of one of the availability zones.	You must provide a third physically separate location that runs a vSphere environment. You might use a VMware Cloud Foundation instance in a separate physical location.
VCF-VSAN-WTN-REQD-CFG-002	Deploy a witness appliance that corresponds to the required cluster capacity.	Ensures the witness appliance is sized to support the projected workload storage consumption.	The vSphere environment at the witness location must satisfy the resource requirements of the witness appliance.
VCF-VSAN-WTN-REQD-CFG-003	Connect the first VMkernel adapter of the vSAN witness appliance to the management network in the witness site.	Enables connecting the witness appliance to the workload domain vCenter Server.	The management networks in both availability zones must be routed to the management network in the witness site.
VCF-VSAN-WTN-REQD-CFG-004	Allocate a statically assigned IP address and host name to the management adapter of the vSAN witness appliance.	Simplifies maintenance and tracking, and implements a DNS configuration.	Requires precise IP address management.
VCF-VSAN-WTN-REQD-CFG-005	Configure forward and reverse DNS records for the vSAN witness appliance for the VMware Cloud Foundation instance.	Enables connecting the vSAN witness appliance to the workload domain vCenter Server by FQDN instead of IP address.	You must provide DNS records for the vSAN witness appliance.
VCF-VSAN-WTN-REQD-CFG-006	Configure time synchronization by using an internal NTP time for the vSAN witness appliance.	Prevents any failures in the stretched cluster configuration that are caused by time mismatch between the vSAN witness appliance and the ESXi hosts in both availability zones and workload domain vCenter Server.	■ An operational NTP service must be available in the environment. ■ All firewalls between the vSAN witness appliance and the NTP servers must allow NTP traffic on the required network ports.

Table 15-11. vSAN Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-001	Provide sufficient raw capacity to meet the planned needs of the workload domain cluster.	Ensures that sufficient resources are present in the workload domain cluster, preventing the need to expand the vSAN datastore in the future.	None.
VCF-VSAN-RCMD-CFG-002	Ensure that at least 30% of free space is always available on the vSAN datastore,.	This reserved capacity is set aside for host maintenance mode data evacuation, component rebuilds, rebalancing operations, and VM snapshots.	Increases the amount of available storage needed.
VCF-VSAN-RCMD-CFG-003	Use the default VMware vSAN storage policy.	■ Provides the level of redundancy that is needed in the workload domain cluster. ■ Provides the level of performance that is enough for the individual workloads.	You might need additional policies for third-party virtual machines hosted in these clusters because their performance or availability requirements might differ from what the default VMware vSAN policy supports.
VCF-VSAN-RCMD-CFG-004	Leave the default virtual machine swap file as a sparse object on vSAN.	Sparse virtual swap files consume capacity on vSAN only as they are accessed. As a result, you can reduce the consumption on the vSAN datastore if virtual machines do not experience memory over-commitment, which would require the use of the virtual swap file.	None.
VCF-VSAN-RCMD-CFG-005	Use the existing vSphere Distributed Switch instance for the workload domain cluster.	■ Reduces the complexity of the network design. ■ Reduces the number of physical NICs required.	All traffic types can be shared over common uplinks.

Table 15-11. vSAN Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-006	Configure jumbo frames on the VLAN for vSAN traffic.	■ Simplifies configuration because jumbo frames are also used to improve the performance of vSphere vMotion and NFS storage traffic. ■ Reduces the CPU overhead, resulting in high network usage.	Every device in the network must support jumbo frames.
VCF-VSAN-RCMD-CFG-007	Configure vSAN in an all-flash configuration in the default workload domain cluster.	Meets the performance needs of the default workload domain cluster.	All vSAN disks must be flash disks, which might cost more than magnetic disks.

Table 15-12. vSAN OSA Design Recommendations for with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-008	Ensure that the storage I/O controller has a minimum queue depth of 256 set.	Storage controllers with lower queue depths can cause performance and stability problems when running vSAN. vSAN ReadyNode servers are configured with the correct queue depths for vSAN.	Limits the number of compatible I/O controllers that can be used for storage.
VCF-VSAN-RCMD-CFG-009	Do not use the storage I/O controllers that are running vSAN disk groups for another purpose.	Running non-vSAN disks, for example, VMFS, on a storage I/O controller that is running a vSAN disk group can impact vSAN performance.	If non-vSAN disks are required in ESXi hosts, you must have an additional storage I/O controller in the host.

Table 15-12. vSAN OSA Design Recommendations for with VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-010	Configure vSAN with a minimum of two disk groups per ESXi host.	Reduces the size of the fault domain and spreads the I/O load over more disks for better performance.	Using multiple disk groups requires more disks in each ESXi host.
VCF-VSAN-RCMD-CFG-011	For the cache tier in each disk group, use a flash-based drive that is at least 600 GB large.	Provides enough cache for both hybrid or all-flash vSAN configurations to buffer I/O and ensure disk group performance. Additional space in the cache tier does not increase performance.	Using larger flash disks can increase the initial host cost.

Table 15-13. vSAN ESA Design Recommendations for with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-012	Activate auto-policy management.	Configures optimized storage policies based on the cluster type and the number of hosts in the cluster inventory. Changes to the number of hosts in the cluster or Host Rebuild Reserve will prompt you to make a suggested adjustment to the optimized storage policy.	You must activate auto-policy management manually.
VCF-VSAN-RCMD-CFG-013	Activate vSAN ESA compression.	Activated by default, it also improves performance.	PostgreSQL databases and other applications might use their own compression capabilities. In these cases, using a storage policy with the compression capability turned off will save CPU cycles. You can disable vSAN ESA compressions for such workloads through the use of the Storage Policy Based Management (SPBM) framework.
VCF-VSAN-RCMD-CFG-014	Use NICs with a minimum 25-GbE capacity.	10-GbE NICs will limit the scale and performance of a vSAN ESA cluster because usually performance requirements increase over the lifespan of the cluster.	Requires 25-GbE or faster network fabric.

Table 15-14. vSAN Design Recommendations for Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-WTN-RCMD-CFG-001	Configure the vSAN witness appliance to use the first VMkernel adapter, that is the management interface, for vSAN witness traffic.	Removes the requirement to have static routes on the witness appliance as witness traffic is routed over the management network.	The management networks in both availability zones must be routed to the management network in the witness site.
VCF-VSAN-WTN-RCMD-CFG-002	Place witness traffic on the management VMkernel adapter of all the ESXi hosts in the workload domain.	Separates the witness traffic from the vSAN data traffic. Witness traffic separation provides the following benefits: ■ Removes the requirement to have static routes from the vSAN networks in both availability zones to the witness site. ■ Removes the requirement to have jumbo frames enabled on the path between each availability zone and the witness site because witness traffic can use a regular MTU size of 1500 bytes.	The management networks in both availability zones must be routed to the management network in the witness site.

ESXi Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the ESXi host configuration in an environment with a single or multiple VMware Cloud Foundation instances. The design elements determine the ESXi hardware configuration, networking, life cycle management and remote access.

The configuration tasks for most design requirements and recommendations are automated in VMware Cloud Foundation. You must perform the configuration manually only for a limited number of decisions as noted in the design implications.

For full design details, see ESXi Design for VMware Cloud Foundation.

Table 15-15. Design Requirements for ESXi Server Hardware

Requirement ID	Design Requirement	Requirement Justification	Requirement Implication
VCF-ESX-REQD-CFG-001	Install no less than the minimum number of ESXi hosts required for the cluster type being deployed.	■ Ensures availability requirements are met. ■ If one of the hosts is not available because of a failure or maintenance event, the CPU overcommitment ratio becomes 2:1.	None.
VCF-ESX-REQD-CFG-002	Ensure each ESXi host matches the required CPU, memory and storage specification.	■ Ensures workloads will run without contention even during failure and maintenance conditions.	Assemble the server specification and number according to the sizing in VMware Cloud Foundation Planning and Preparation Workbook which is based on projected deployment size.
VCF-ESX-REQD-SEC-001	Regenerate the certificate of each ESXi host after assigning the host an FQDN.	Establishes a secure connection with VMware Cloud Builder during the deployment of a workload domain and prevents man-in-the-middle (MiTM) attacks.	You must manually regenerate the certificates of the ESXi hosts before the deployment of a workload domain.

Table 15-16. Design Recommendations for ESXi Server Hardware

Recommendation ID	Recommendation	Justification	Implication
VCF-ESX-RCMD-CFG-001	Use vSAN ReadyNodes with vSAN storage for each ESXi host in the management domain.	Your management domain is fully compatible with vSAN at deployment. For information about the models of physical servers that are vSAN-ready, see vSAN Compatibility Guide for vSAN ReadyNodes.	Hardware choices might be limited. If you plan to use a server configuration that is not a vSAN ReadyNode, your CPU, disks and I/O modules must be listed on the VMware Compatibility Guide under *CPU Series* and *vSAN Compatibility List* aligned to the ESXi version specified in VMware Cloud Foundation 5.1 Release Notes.
VCF-ESX-RCMD-CFG-002	Allocate hosts with uniform configuration across the default management vSphere cluster.	A balanced cluster has these advantages: ■ Predictable performance even during hardware failures ■ Minimal impact of resynchronization or rebuild operations on performance	You must apply vendor sourcing, budgeting, and procurement considerations for uniform server nodes on a per cluster basis.
VCF-ESX-RCMD-CFG-003	When sizing CPU, do not consider multithreading technology and associated performance gains.	Although multithreading technologies increase CPU performance, the performance gain depends on running workloads and differs from one case to another.	Because you must provide more physical CPU cores, costs increase and hardware choices become limited.
VCF-ESX-RCMD-CFG-004	Install and configure all ESXi hosts in the default management cluster to boot using a 128-GB device or larger.	Provides hosts that have large memory, that is, greater than 512 GB, with enough space for the scratch partition when using vSAN.	None.
VCF-ESX-RCMD-CFG-005	Use the default configuration for the scratch partition on all ESXi hosts in the default management cluster.	■ If a failure in the vSAN cluster occurs, the ESXi hosts remain responsive and log information is still accessible. ■ It is not possible to use vSAN datastore for the scratch partition.	None.

Table 15-16. Design Recommendations for ESXi Server Hardware (continued)

Recommendation ID	Recommendation	Justification	Implication
VCF-ESX-RCMD-CFG-006	For workloads running in the default management cluster, save the virtual machine swap file at the default location.	Simplifies the configuration process.	Increases the amount of replication traffic for management workloads that are recovered as part of the disaster recovery process.
VCF-ESX-RCMD-NET-001	Place the ESXi hosts in each management domain cluster on a host management network that is separate from the VM management network.	Enables the separation of the physical VLAN between ESXi hosts and the other management components for security reasons.	Increases the number of VLANs required.
VCF-ESX-RCMD-NET-002	Place the ESXi hosts in each VI workload domain on a separate host management VLAN-backed network.	Enables the separation of the physical VLAN between the ESXi hosts in different VI workload domains for security reasons.	Increases the number of VLANs required. For each VI workload domain, you must allocate a separate management subnet.
VCF-ESX-RCMD-SEC-001	Deactivate SSH access on all ESXi hosts in the management domain by having the SSH service stopped and using the default SSH service policy `Start and stop manually`.	Ensures compliance with the *vSphere Security Configuration Guide* and with security best practices. Disabling SSH access reduces the risk of security attacks on the ESXi hosts through the SSH interface.	You must activate SSH access manually for troubleshooting or support activities as VMware Cloud Foundation deactivates SSH on ESXi hosts after workload domain deployment.
VCF-ESX-RCMD-SEC-002	Set the advanced setting `UserVars.SuppressShellWarning` to `0` across all ESXi hosts in the management domain.	■ Ensures compliance with the *vSphere Security Configuration Guide* and with security best practices ■ Enables the warning message that appears in the vSphere Client every time SSH access is activated on an ESXi host.	You must turn off SSH enablement warning messages manually when performing troubleshooting or support activities.

vCenter Server Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the vCenter Server configuration in an environment with a single or multiple VMware Cloud Foundation instances. The design elements also consider if an instance contains a single or multiple availability zones. The vCenter Server design also includes the configuration of the default management cluster.

The configuration tasks for most design requirements and recommendations are automated in VMware Cloud Foundation. You must perform the configuration manually only for a limited number of decisions as noted in the design implications.

For full design details, see vCenter Server Design for VMware Cloud Foundation.

vCenter Server Design Elements

Table 15-17. vCenter Server Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-CFG-001	Deploy a dedicated vCenter Server appliance for the management domain of the VMware Cloud Foundation instance.	■ Isolates vCenter Server failures to management or customer workloads. ■ Isolates vCenter Server operations between management and customers. ■ Supports a scalable cluster design where you can reuse the management components as more customer workloads are added to the SDDC. ■ Simplifies capacity planning for customer workloads because you do not consider management workloads for the VI workload domain vCenter Server. ■ Improves the ability to upgrade the vSphere environment and related components by enabling for explicit separation of maintenance windows: ■ Management workloads remain available while you are upgrading the tenant workloads ■ Customer workloads remain available while you are upgrading the management nodes ■ Supports clear separation of roles and responsibilities to ensure that only administrators with granted authorization can control the management workloads. ■ Facilitates quicker troubleshooting and problem resolution. ■ Simplifies disaster recovery operations by supporting a clear separation between recovery of the management components and tenant workloads.	Requires a separate license for the vCenter Server instance in the management domain

Table 15-17. vCenter Server Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
		■ Provides isolation of potential network issues by introducing network separation of the clusters in the SDDC.	
VCF-VCS-REQD-NET-001	Place all workload domain vCenters Server appliances on the VM management network in the management domain.	■ Simplifies IP addressing for management VMs by using the same VLAN and subnet. ■ Provides simplified secure access to management VMs in the same VLAN network.	None.

Table 15-18. vCenter Server Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VCS-RCMD-CFG-001	Deploy an appropriately sized vCenter Server appliance for each workload domain.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is Small and for VI workload domains is Medium. To override these values, you must use the Cloud Builder API and the SDDC Manager API.
VCF-VCS-RCMD-CFG-002	Deploy a vCenter Server appliance with the appropriate storage size.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is **Small** and for VI Workload Domains is **Medium**. To override these values, you must use the API.
VCF-VCS-RCMD-CFG-003	Protect workload domain vCenter Server appliances by using vSphere HA.	vSphere HA is the only supported method to protect vCenter Server availability in VMware Cloud Foundation.	vCenter Server becomes unavailable during a vSphere HA failover.
VCF-VCS-RCMD-CFG-004	In vSphere HA, set the restart priority policy for the vCenter Server appliance to high.	vCenter Server is the management and control plane for physical and virtual infrastructure. In a vSphere HA event, to ensure the rest of the SDDC management stack comes up faultlessly, the workload domain vCenter Server must be available first, before the other management components come online.	If the restart priority for another virtual machine is set to highest, the connectivity delay for the management components will be longer.

Table 15-19. vCenter Server Design Recommendations for vSAN Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VCS-RCMD-CFG-005	Add the vCenter Server appliance to the virtual machine group for the first availability zone.	Ensures that, by default, the vCenter Server appliance is powered on a host in the first availability zone.	None.

vCenter Single Sign-On Design Elements

Table 15-20. Design Requirements for the Multiple vCenter Server Instance - Single vCenter Single Sign-on Domain Topology for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-SSO-STD-001	Join all vCenter Server instances within aVMware Cloud Foundation instance to a single vCenter Single Sign-On domain.	When all vCenter Server instances are in the same vCenter Single Sign-On domain, they can share authentication and license data across all components.	■ Only one vCenter Single Sign-On domain exists. ■ The number of linked vCenter Server instances in the same vCenter Single Sign-On domain is limited to 15 instances. Because each workload domain uses a dedicated vCenter Server instance, you can deploy up to 15 domains within each VMware Cloud Foundation instance.
VCF-VCS-REQD-SSO-STD-002	Create a ring topology between the vCenter Server instances within the VMware Cloud Foundation instance.	By default, one vCenter Server instance replicates only with another vCenter Server instance. This setup creates a single point of failure for replication. A ring topology ensures that each vCenter Server instance has two replication partners and removes any single point of failure.	None.

Table 15-21. Design Requirements for Multiple vCenter Server Instance - Multiple vCenter Single Sign-On Domain Topology for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-SSO-ISO-001	Create all vCenter Server instances within a VMware Cloud Foundation instance in their own unique vCenter Single Sign-On domains.	■ Enables isolation at the vCenter Single Sign-On domain layer for increased security separation. ■ Supports up to 25 workload domains.	■ Each vCenter server instance is managed through its own pane of glass using a different set of administrative credentials. ■ You must manage password rotation for each vCenter Single Sign-On domain separately.

vSphere Cluster Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the vSphere cluster configuration in an environment with a single or multiple VMware Cloud Foundation instances. The design elements also consider if an instance contains a single or multiple availability zones.

For full design details, see Logical vSphere Cluster Design for VMware Cloud Foundation.

Table 15-22. vSphere Cluster Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-001	Create a cluster in each workload domain for the initial set of ESXi hosts.	■ Simplifies configuration by isolating management from customer workloads. ■ Ensures that customer workloads have no impact on the management stack.	Management of multiple clusters and vCenter Server instances increases operational overhead.
VCF-CLS-REQD-CFG-002	Allocate a minimum number of ESXi hosts according to the cluster type being deployed.	■ Ensures correct level of redundancy to protect against host failure in the cluster.	To support redundancy, you must allocate additional ESXi host resources.

Table 15-22. vSphere Cluster Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-003	If using a consolidated workload domain, configure the following vSphere resource pools to control resource usage by management and customer workloads. ■ *cluster-name*-rp-sddc-mgmt ■ *cluster-name*-rp-sddc-edge ■ *cluster-name*-rp-user-edge ■ *cluster-name*-rp-user-vm	■ Ensures sufficient resources for the management components.	You must manage the vSphere resource pool settings over time.
VCF-CLS-REQD-CFG-004	Configure the vSAN network gateway IP address as the isolation address for the cluster.	Allows vSphere HA to validate if a host is isolated from the vSAN network.	None.
VCF-CLS-REQD-CFG-005	Set the advanced cluster setting `das.usedefaultisolationaddress` to false.	Ensures that vSphere HA uses the manual isolation addresses instead of the default management network gateway address.	None.

Table 15-23. vSphere Cluster Design Requirements for vSAN Stretched Clusters with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-006	Configure the vSAN network gateway IP addresses for the second availability zone as an additional isolation addresses for the cluster.	Allows vSphere HA to validate if a host is isolated from the vSAN network for hosts in both availability zones.	None.
VCF-CLS-REQD-CFG-007	Enable the **Override default gateway for this adapter** setting on the vSAN VMkernel adapters on all ESXi hosts.	Enables routing the vSAN data traffic through the vSAN network gateway rather than through the management gateway.	vSAN networks across availability zones must have a route to each other.
VCF-CLS-REQD-CFG-008	Create a host group for each availability zone and add the ESXi hosts in the zone to the respective group.	Makes it easier to manage which virtual machines run in which availability zone.	You must create and maintain VM-Host DRS group rules.

Table 15-24. vSphere Cluster Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-001	Use vSphere HA to protect all virtual machines against failures.	vSphere HA supports a robust level of protection for both ESXi host and virtual machine availability.	You must provide sufficient resources on the remaining hosts so that virtual machines can be restarted on those hosts in the event of a host outage.
VCF-CLS-RCMD-CFG-002	Set host isolation response to Power Off and restart VMs in vSphere HA.	vSAN requires that the host isolation response be set to Power Off and to restart virtual machines on available ESXi hosts.	If a false positive event occurs, virtual machines are powered off and an ESXi host is declared isolated incorrectly.
VCF-CLS-RCMD-CFG-003	Configure admission control for 1 ESXi host failure and percentage-based failover capacity.	Using the percentage-based reservation works well in situations where virtual machines have varying and sometimes significant CPU or memory reservations. vSphere automatically calculates the reserved percentage according to the number of ESXi host failures to tolerate and the number of ESXi hosts in the cluster.	In a cluster of 4 ESXi hosts, the resources of only 3 ESXi hosts are available for use.
VCF-CLS-RCMD-CFG-004	Enable VM Monitoring for each cluster.	VM Monitoring provides in-guest protection for most VM workloads. The application or service running on the virtual machine must be capable of restarting successfully after a reboot or the virtual machine restart is not sufficient.	None.
VCF-CLS-RCMD-CFG-005	Set the advanced cluster setting `das.iostatsinterval` to 0 to deactivate monitoring the storage and network I/O activities of the management appliances.	Enables triggering a restart of a management appliance when an OS failure occurs and heartbeats are not received from VMware Tools instead of waiting additionally for the I/O check to complete.	If you want to specifically enable I/O monitoring, you must configure the das.iostatsinterval advanced setting.
VCF-CLS-RCMD-CFG-006	Enable vSphere DRS on all clusters, using the default fully automated mode with medium threshold.	Provides the best trade-off between load balancing and unnecessary migrations with vSphere vMotion.	If a vCenter Server outage occurs, the mapping from virtual machines to ESXi hosts might be difficult to determine.

Table 15-24. vSphere Cluster Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-007	Enable Enhanced vMotion Compatibility (EVC) on all clusters in the management domain.	Supports cluster upgrades without virtual machine downtime.	You must enable EVC only if the clusters contain hosts with CPUs from the same vendor. You must enable EVC on the default management domain cluster during bringup.
VCF-CLS-RCMD-CFG-008	Set the cluster EVC mode to the highest available baseline that is supported for the lowest CPU architecture on the hosts in the cluster.	Supports cluster upgrades without virtual machine downtime.	None.
VCF-CLS-RCMD-LCM-001	Use images as the life cycle management method for VI workload domains.	vSphere Lifecycle Manager images simplify the management of firmware and vendor add-ons manually.	An initial cluster image is required during workload domain or cluster deployment.

Table 15-25. vSphere Cluster Design Recommendations for vSAN Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-009	Increase admission control percentage to half of the ESXi hosts in the cluster.	Allocating only half of a stretched cluster ensures that all VMs have enough resources if an availability zone outage occurs.	In a cluster of 8 ESXi hosts, the resources of only 4 ESXi hosts are available for use. If you add more ESXi hosts to the default management cluster, add them in pairs, one per availability zone.
VCF-CLS-RCMD-CFG-010	Create a virtual machine group for each availability zone and add the VMs in the zone to the respective group.	Ensures that virtual machines are located only in the assigned availability zone to avoid unnecessary vSphere vMotion migrations.	You must add virtual machines to the allocated group manually.
VCF-CLS-RCMD-CFG-011	Create a should-run-on-hosts-in-group VM-Host affinity rule to run each group of virtual machines on the respective group of hosts in the same availability zone.	Ensures that virtual machines are located only in the assigned availability zone to avoid unnecessary vSphere vMotion migrations.	You must manually create the rules.

vSphere Networking Design Elements for VMware Cloud Foundation

Use this list of recommendations for reference related to the configuration of the vSphere Distributed Switch instances and VMkernel adapters in a VMware Cloud Foundation environment.

For full design details, see vSphere Networking Design for VMware Cloud Foundation.

Table 15-26. vSphere Networking Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VDS-RCMD-CFG-001	Use a single vSphere Distributed Switch per cluster.	■ Reduces the complexity of the network design. ■ Reduces the size of the fault domain.	Increases the number of vSphere Distributed Switches that must be managed.
VCF-VDS-RCMD-CFG-002	Configure the MTU size of the vSphere Distributed Switch to 9000 for jumbo frames.	■ Supports the MTU size required by system traffic types. ■ Improves traffic throughput.	When adjusting the MTU packet size, you must also configure the entire network path (VMkernel ports, virtual switches, physical switches, and routers) to support the same MTU packet size.
VCF-VDS-RCMD-DPG-001	Use ephemeral port binding for the Management VM port group.	Using ephemeral port binding provides the option for recovery of the vCenter Server instance that is managing the distributed switch.	Port-level permissions and controls are lost across power cycles, and no historical context is saved.
VCF-VDS-RCMD-DPG-002	Use static port binding for all non-management port groups.	Static binding ensures a virtual machine connects to the same port on the vSphere Distributed Switch. This allows for historical data and port level monitoring.	None.
VCF-VDS-RCMD-DPG-003	Use the `Route based on physical NIC load` teaming algorithm for the VM management port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-004	Use the `Route based on physical NIC load` teaming algorithm for the ESXi management port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.

Table 15-26. vSphere Networking Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VDS-RCMD-DPG-005	Use the `Route based on physical NIC load` teaming algorithm for the vSphere vMotion port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-006	Use the `Route based on physical NIC load` teaming algorithm for the vSAN port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-NIO-001	Enable Network I/O Control on vSphere Distributed Switch of the management domain cluster.	Increases resiliency and performance of the network.	Network I/O Control might impact network performance for critical traffic types if misconfigured.
VCF-VDS-RCMD-NIO-002	Set the share value for management traffic to Normal.	By keeping the default setting of Normal, management traffic is prioritized higher than vSphere vMotion but lower than vSAN traffic. Management traffic is important because it ensures that the hosts can still be managed during times of network contention.	None.
VCF-VDS-RCMD-NIO-003	Set the share value for vSphere vMotion traffic to Low.	During times of network contention, vSphere vMotion traffic is not as important as virtual machine or storage traffic.	During times of network contention, vMotion takes longer than usual to complete.
VCF-VDS-RCMD-NIO-004	Set the share value for virtual machines to High.	Virtual machines are the most important asset in the SDDC. Leaving the default setting of High ensures that they always have access to the network resources they need.	None.

Table 15-26. vSphere Networking Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VDS-RCMD-NIO-005	Set the share value for vSAN traffic to High.	During times of network contention, vSAN traffic needs guaranteed bandwidth to support virtual machine performance.	None.
VCF-VDS-RCMD-NIO-006	Set the share value for other traffic types to Low.	By default, VVMware Cloud Foundation does not use other traffic types, like vSphere FT traffic. Hence, these traffic types can be set the lowest priority.	None.

NSX Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the configuration of NSX in an environment with a single or multiple VMware Cloud Foundation instances. The design also considers if an instance contains a single or multiple availability zones.

For full design details, see Chapter 8 NSX Design for VMware Cloud Foundation.

NSX Manager Design Elements

Table 15-27. NSX Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LM-REQD-CFG-001	Place the appliances of the NSX Manager cluster on the VM management network in the management domain.	■ Simplifies IP addressing for management VMs by using the same VLAN and subnet. ■ Provides simplified secure access to management VMs in the same VLAN network.	None.
VCF-NSX-LM-REQD-CFG-002	Deploy three NSX Manager nodes in the default vSphere cluster in the management domain for configuring and managing the network services for the workload domain.	Supports high availability of the NSX manager cluster.	You must have sufficient resources in the default cluster of the management domain to run three NSX Manager nodes.

Table 15-28. NSX Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-LM-RCMD-CFG-001	Deploy appropriately sized nodes in the NSX Manager cluster for the workload domain.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is Medium, and for VI workload domains is Large.
VCF-NSX-LM-RCMD-CFG-002	Create a virtual IP (VIP) address for the NSX Manager cluster for the workload domain.	Provides high availability of the user interface and API of NSX Manager.	■ The VIP address feature provides high availability only. It does not load-balance requests across the cluster. ■ When using the VIP address feature, all NSX Manager nodes must be deployed on the same Layer 2 network.
VCF-NSX-LM-RCMD-CFG-003	Apply VM-VM anti-affinity rules in vSphere Distributed Resource Scheduler (vSphere DRS) to the NSX Manager appliances.	Keeps the NSX Manager appliances running on different ESXi hosts for high availability.	You must allocate at least four physical hosts so that the three NSX Manager appliances continue running if an ESXi host failure occurs.
VCF-NSX-LM-RCMD-CFG-004	In vSphere HA, set the restart priority policy for each NSX Manager appliance to high.	■ NSX Manager implements the control plane for virtual network segments. vSphere HA restarts the NSX Manager appliances first so that other virtual machines that are being powered on or migrated by using vSphere vMotion while the control plane is offline lose connectivity only until the control plane quorum is re-established. ■ Setting the restart priority to high reserves the highest priority for flexibility for adding services that must be started before NSX Manager.	If the restart priority for another management appliance is set to highest, the connectivity delay for management appliances will be longer.

Table 15-29. NSX Manager Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-LM-RCMD-CFG-006	Add the NSX Manager appliances to the virtual machine group for the first availability zone.	Ensures that, by default, the NSX Manager appliances are powered on a host in the primary availability zone.	None.

NSX Global Manager Design Elements

You must perform manually the configuration tasks for the design requirements and recommendations for NSX Global Manager.

Table 15-30. NSX Global Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-GM-REQD-CFG-001	Place the appliances of the NSX Global Manager cluster on the Management VM network in each VMware Cloud Foundation instance.	■ Simplifies IP addressing for management VMs. ■ Provides simplified secure access to all management VMs in the same VLAN network.	None.

Table 15-31. NSX Global Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-001	Deploy three NSX Global Manager nodes for the workload domain to support NSX Federation across VMware Cloud Foundation instances.	Provides high availability for the NSX Global Manager cluster.	You must have sufficient resources in the default cluster of the management domain to run three NSX Global Manager nodes.
VCF-NSX-GM-RCMD-CFG-002	Deploy appropriately sized nodes in the NSX Global Manager cluster for the workload domain.	Ensures resource availability and usage efficiency per workload domain.	The recommended size for a management domain is Medium and for VI workload domains is Large.
VCF-NSX-GM-RCMD-CFG-003	Create a virtual IP (VIP) address for the NSX Global Manager cluster for the workload domain.	Provides high availability of the user interface and API of NSX Global Manager.	■ The VIP address feature provides high availability only. It does not load-balance requests across the cluster. ■ When using the VIP address feature, all NSX Global Manager nodes must be deployed on the same Layer 2 network.

Table 15-31. NSX Global Manager Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-004	Apply VM-VM anti-affinity rules in vSphere DRS to the NSX Global Manager appliances.	Keeps the NSX Global Manager appliances running on different ESXi hosts for high availability.	You must allocate at least four physical hosts so that the three NSX Manager appliances continue running if an ESXi host failure occurs.
VCF-NSX-GM-RCMD-CFG-005	In vSphere HA, set the restart priority policy for each NSX Global Manager appliance to medium.	■ NSX Global Manager implements the management plane for global segments and firewalls. NSX Global Manager is not required for control plane and data plane connectivity. ■ Setting the restart priority to medium reserves the high priority for services that impact the NSX control or data planes.	■ Management of NSX global components will be unavailable until the NSX Global Manager virtual machines restart. ■ The NSX Global Manager cluster is deployed in the management domain, where the total number of virtual machines is limited and where it competes with other management components for restart priority.
VCF-NSX-GM-RCMD-CFG-006	Deploy an additional NSX Global Manager Cluster in the second VMware Cloud Foundation instance.	Enables recoverability of NSX Global Manager in the second VMware Cloud Foundation instance if a failure in the first VMware Cloud Foundation instance occurs.	Requires additional NSX Global Manager nodes in the second VMware Cloud Foundation instance.

Table 15-31. NSX Global Manager Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-007	Set the NSX Global Manager cluster in the second VMware Cloud Foundation instance as standby for the workload domain.	Enables recoverability of NSX Global Manager in the second VMware Cloud Foundation instance if a failure in the first instance occurs.	Must be done manually.
VCF-NSX-GM-RCMD-SEC-001	Establish an operational practice to capture and update the thumbprint of the NSX Local Manager certificate on NSX Global Manager every time the certificate is updated by using SDDC Manager.	Ensures secured connectivity between the NSX Manager instances. Each certificate has its own unique thumbprint. NSX Global Manager stores the unique thumbprint of the NSX Local Manager instances for enhanced security. If an authentication failure between NSX Global Manager and NSX Local Manager occurs, objects that are created from NSX Global Manager will not be propagated on to the SDN.	The administrator must establish and follow an operational practice by using a runbook or automated process to ensure that the thumbprint is up-to-date.

Table 15-32. NSX Global Manager Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-008	Add the NSX Global Manager appliances to the virtual machine group for the first availability zone.	Ensures that, by default, the NSX Global Manager appliances are powered on a host in the primary availability zone.	Done automatically by VMware Cloud Foundation when stretching a cluster.

NSX Edge Design Elements

Table 15-33. NSX Edge Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-001	Connect the management interface of each NSX Edge node to the VM management network.	Provides connection from the NSX Manager cluster to the NSX Edge.	None.
VCF-NSX-EDGE-REQD-CFG-002	■ Connect the `fp-eth0` interface of each NSX Edge appliance to a VLAN trunk port group pinned to physical NIC 0 of the host, with the ability to failover to physical NIC 1. ■ Connect the `fp-eth1` interface of each NSX Edge appliance to a VLAN trunk port group pinned to physical NIC 1 of the host, with the ability to failover to physical NIC 0. ■ Leave the `fp-eth2` interface of each NSX Edge appliance unused.	■ Because VLAN trunk port groups pass traffic for all VLANs, VLAN tagging can occur in the NSX Edge node itself for easy post-deployment configuration. ■ By using two separate VLAN trunk port groups, you can direct traffic from the edge node to a particular host network interface and top of rack switch as needed. ■ In the event of failure of the top of rack switch, the VLAN trunk port group will failover to the other physical NIC and to ensure both `fp-eth0` and `fp-eth1` are available.	None.

Table 15-33. NSX Edge Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-003	Use a dedicated VLAN for edge overlay that is different from the host overlay VLAN.	A dedicated edge overlay network provides support for edge mobility in support of advanced deployments such as multiple availability zones or multi-rack clusters.	■ You must have routing between the VLANs for edge overlay and host overlay. ■ You must allocate another VLAN in the data center infrastructure for edge overlay.
VCF-NSX-EDGE-REQD-CFG-004	Create one uplink profile for the edge nodes with three teaming policies. ■ Default teaming policy of load balance source with both active uplinks `uplink1` and `uplink2`. ■ Named teaming policy of failover order with a single active uplink `uplink1` without standby uplinks. ■ Named teaming policy of failover order with a single active uplink `uplink2` without standby uplinks.	■ An NSX Edge node that uses a single N-VDS can have only one uplink profile. ■ For increased resiliency and performance, supports the concurrent use of both edge uplinks through both physical NICs on the ESXi hosts. ■ The default teaming policy increases overlay performance and availability by using multiple TEPs, and balancing of overlay traffic. ■ By using named teaming policies, you can connect an edge uplink to a specific host uplink and from there to a specific top of rack switch in the data center. ■ Enables ECMP because the NSX Edge nodes can uplink to the physical network over two different VLANs.	None.

Table 15-34. NSX Edge Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-005	Allocate a separate VLAN for edge RTEP overlay that is different from the edge overlay VLAN.	The RTEP network must be on a VLAN that is different from the edge overlay VLAN. This is an NSX requirement that provides support for configuring different MTU size per network.	You must allocate another VLAN in the data center infrastructure.

Table 15-35. NSX Edge Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-001	Use appropriately sized NSX Edge virtual appliances.	Ensures resource availability and usage efficiency per workload domain.	You must provide sufficient compute resources to support the chosen appliance size.
VCF-NSX-EDGE-RCMD-CFG-002	Deploy the NSX Edge virtual appliances to the default vSphere cluster of the workload domain, sharing the cluster between the workloads and the edge appliances.	Simplifies the configuration and minimizes the number of ESXi hosts required for initial deployment.	Workloads and NSX Edges share the same compute resources.
VCF-NSX-EDGE-RCMD-CFG-003	Deploy two NSX Edge appliances in an edge cluster in the default vSphere cluster of the workload domain.	Creates the minimum size NSX Edge cluster while satisfying the requirements for availability.	For a VI workload domain, additional edge appliances might be required to satisfy increased bandwidth requirements.
VCF-NSX-EDGE-RCMD-CFG-004	Apply VM-VM anti-affinity rules for vSphere DRS to the virtual machines of the NSX Edge cluster.	Keeps the NSX Edge nodes running on different ESXi hosts for high availability.	None.
VCF-NSX-EDGE-RCMD-CFG-005	In vSphere HA, set the restart priority policy for each NSX Edge appliance to high.	■ The NSX Edge nodes are part of the north-south data path for overlay segments. vSphere HA restarts the NSX Edge appliances first to minimise the time an edge VM is offline. ■ Setting the restart priority to high reserves highest for future needs.	If the restart priority for another VM in the cluster is set to highest, the connectivity delays for edge appliances will be longer.

Table 15-35. NSX Edge Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-006	Create an NSX Edge cluster with the default Bidirectional Forwarding Detection (BFD) configuration between the NSX Edge nodes in the cluster.	■ Satisfies the availability requirements by default. ■ Edge nodes must remain available to create services such as NAT, routing to physical networks, and load balancing.	None.
VCF-NSX-EDGE-RCMD-CFG-007	Use a single N-VDS in the NSX Edge nodes.	■ Simplifies deployment of the edge nodes. ■ The same N-VDS switch design can be used regardless of edge form factor. ■ Supports multiple TEP interfaces in the edge node. ■ vSphere Distributed Switch is not supported in the edge node.	None.

Table 15-36. NSX Edge Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-008	Add the NSX Edge appliances to the virtual machine group for the first availability zone.	Ensures that, by default, the NSX Edge appliances are powered on upon a host in the primary availability zone.	None.

BGP Routing Design Elements for VMware Cloud Foundation

Table 15-37. BGP Routing Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-001	To enable ECMP between the Tier-0 gateway and the Layer 3 devices (ToR switches or upstream devices), create two VLANs. The ToR switches or upstream Layer 3 devices have an SVI on one of the two VLANS, and each Edge node in the cluster has an interface on each VLAN.	Supports multiple equal-cost routes on the Tier-0 gateway and provides more resiliency and better bandwidth use in the network.	Additional VLANs are required.
VCF-NSX-BGP-REQD-CFG-002	Assign a named teaming policy to the VLAN segments to the Layer 3 device pair.	Pins the VLAN traffic on each segment to its target edge node interface. From there, the traffic is directed to the host physical NIC that is connected to the target top of rack switch.	None.
VCF-NSX-BGP-REQD-CFG-003	Create a VLAN transport zone for edge uplink traffic.	Enables the configuration of VLAN segments on the N-VDS in the edge nodes.	Additional VLAN transport zones might be required if the edge nodes are not connected to the same top of rack switch pair.
VCF-NSX-BGP-REQD-CFG-004	Deploy a Tier-1 gateway and connect it to the Tier-0 gateway.	Creates a two-tier routing architecture. Abstracts the NSX logical components which interact with the physical data center from the logical components which provide SDN services.	A Tier-1 gateway can only be connected to a single Tier-0 gateway. In cases where multiple Tier-0 gateways are required, you must create multiple Tier-1 gateways.
VCF-NSX-BGP-REQD-CFG-005	Deploy a Tier-1 gateway to the NSX Edge cluster.	Enables stateful services, such as load balancers and NAT, for SDDC management components. Because a Tier-1 gateway always works in active-standby mode, the gateway supports stateful services.	None.

Table 15-38. BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-006	Extend the uplink VLANs to the top of rack switches so that the VLANs are stretched between both availability zones.	Because the NSX Edge nodes will fail over between the availability zones, ensures uplink connectivity to the top of rack switches in both availability zones regardless of the zone the NSX Edge nodes are presently in.	You must configure a stretched Layer 2 network between the availability zones by using physical network infrastructure.
VCF-NSX-BGP-REQD-CFG-007	Provide this SVI configuration on the top of the rack switches. ■ In the second availability zone, configure the top of rack switches or upstream Layer 3 devices with an SVI on each of the two uplink VLANs. ■ Make the top of rack switch SVI in both availability zones part of a common stretched Layer 2 network between the availability zones.	Enables the communication of the NSX Edge nodes to the top of rack switches in both availability zones over the same uplink VLANs.	You must configure a stretched Layer 2 network between the availability zones by using the physical network infrastructure.
VCF-NSX-BGP-REQD-CFG-008	Provide this VLAN configuration: ■ Use two VLANs to enable ECMP between the Tier-0 gateway and the Layer 3 devices (top of rack switches or Leaf switches). ■ The ToR switches or upstream Layer 3 devices have an SVI to one of the two VLANS and each NSX Edge node has an interface to each VLAN.	Supports multiple equal-cost routes on the Tier-0 gateway, and provides more resiliency and better bandwidth use in the network.	■ Extra VLANs are required. ■ Requires stretching uplink VLANs between availability zones

Table 15-38. BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-009	Create an IP prefix list that permits access to route advertisement by `any` network instead of using the default IP prefix list.	Used in a route map to prepend a path to one or more autonomous system (AS-path prepend) for BGP neighbors in the second availability zone.	You must manually create an IP prefix list that is identical to the default one.
VCF-NSX-BGP-REQD-CFG-010	Create a route map-out that contains the custom IP prefix list and an AS-path prepend value set to the Tier-0 local AS added twice.	■ Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. ■ Ensures that all ingress traffic passes through the first availability zone.	You must manually create the route map. The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.
VCF-NSX-BGP-REQD-CFG-011	Create an IP prefix list that permits access to route advertisement by network `0.0.0.0/0` instead of using the default IP prefix list.	Used in a route map to configure local-reference on learned default-route for BGP neighbors in the second availability zone.	You must manually create an IP prefix list that is identical to the default one.

Table 15-38. BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-012	Apply a route map-in that contains the IP prefix list for the default route `0.0.0.0/0` and assign a lower local-preference , for example, `80`, to the learned default route and a lower local-preference, for example, `90` any routes learned.	■ Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. ■ Ensures that all egress traffic passes through the first availability zone.	You must manually create the route map. The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.
VCF-NSX-BGP-REQD-CFG-013	Configure the neighbors of the second availability zone to use the route maps as In and Out filters respectively.	Makes the path in and out of the second availability zone less preferred because the AS path is longer and the local preference is lower. As a result, all traffic passes through the first zone.	The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.

Table 15-39. BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-014	Extend the Tier-0 gateway to the second VMware Cloud Foundation instance.	■ Supports ECMP north-south routing on all nodes in the NSX Edge cluster. ■ Enables support for cross-instance Tier-1 gateways and cross-instance network segments.	The Tier-0 gateway deployed in the second instance is removed.
VCF-NSX-BGP-REQD-CFG-015	Set the Tier-0 gateway as primary for all VMware Cloud Foundation instances.	■ In NSX Federation, a Tier-0 gateway lets egress traffic from connected Tier-1 gateways only in its primary locations. ■ Local ingress and egress traffic is controlled independently at the Tier-1 level. No segments are provisioned directly to the Tier-0 gateway. ■ A mixture of network spans (local to a VMware Cloud Foundation instance or spanning multiple instances) is enabled without requiring additional Tier-0 gateways and hence edge nodes. ■ If a failure in a VMware Cloud Foundation instance occurs, the local-instance networking in the other instance remains available without manual intervention.	None.

Table 15-39. BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-016	From the global Tier-0 gateway, establish BGP neighbor peering to the ToR switches connected to the second VMware Cloud Foundation instance.	■ Enables the learning and advertising of routes in the second VMware Cloud Foundation instance. ■ Facilitates a potential automated failover of networks from the first to the second VMware Cloud Foundation instance.	None.
VCF-NSX-BGP-REQD-CFG-017	Use a stretched Tier-1 gateway and connect it to the Tier-0 gateway for cross-instance networking.	■ Enables network span between the VMware Cloud Foundation instances because NSX network segments follow the span of the gateway they are attached to. ■ Creates a two-tier routing architecture.	None.

Table 15-39. BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-018	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the stretched Tier-1 gateway. Set the first VMware Cloud Foundation instance as primary and the second instance as secondary.	■ Enables cross-instance network span between the first and second VMware Cloud Foundation instances. ■ Enables deterministic ingress and egress traffic for the cross-instance network. ■ If a VMware Cloud Foundation instance failure occurs, enables deterministic failover of the Tier-1 traffic flow. ■ During the recovery of the inaccessible VMware Cloud Foundation instance, enables deterministic failback of the Tier-1 traffic flow, preventing unintended asymmetrical routing. ■ Eliminates the need to use BGP attributes in the first and second VMware Cloud Foundation instances to influence location preference and failover.	You must manually fail over and fail back the cross-instance network from the standby NSX Global Manager.
VCF-NSX-BGP-REQD-CFG-019	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the local Tier-1 gateway for that VMware Cloud Foundation instance.	■ Enables instance-specific networks to be isolated to their specific instances. ■ Enables deterministic flow of ingress and egress traffic for the instance-specific networks.	You can use the service router that is created for the Tier-1 gateway for networking services. However, such configuration is not required for network connectivity.
VCF-NSX-BGP-REQD-CFG-020	Set each local Tier-1 gateway only as primary in that instance. Avoid setting the gateway as secondary in the other instances.	Prevents the need to use BGP attributes in primary and secondary instances to influence the instance ingress-egress preference.	None.

Table 15-40. BGP Routing Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
VCF-NSX-BGP-RCMD-CFG-001	Deploy an active-active Tier-0 gateway.	Supports ECMP north-south routing on all Edge nodes in the NSX Edge cluster.	Active-active Tier-0 gateways cannot provide stateful services such as NAT.
VCF-NSX-BGP-RCMD-CFG-002	Configure the BGP Keep Alive Timer to 4 and Hold Down Timer to 12 or lower between the top of tack switches and the Tier-0 gateway.	Provides a balance between failure detection between the top of rack switches and the Tier-0 gateway, and overburdening the top of rack switches with keep-alive traffic.	By using longer timers to detect if a router is not responding, the data about such a router remains in the routing table longer. As a result, the active router continues to send traffic to a router that is down. These timers must be aligned with the data center fabric design of your organization.
VCF-NSX-BGP-RCMD-CFG-003	Do not enable Graceful Restart between BGP neighbors.	Avoids loss of traffic. On the Tier-0 gateway, BGP peers from all the gateways are always active. On a failover, the Graceful Restart capability increases the time a remote neighbor takes to select an alternate Tier-0 gateway. As a result, BFD-based convergence is delayed.	None.
VCF-NSX-BGP-RCMD-CFG-004	Enable helper mode for Graceful Restart mode between BGP neighbors.	Avoids loss of traffic. During a router restart, helper mode works with the graceful restart capability of upstream routers to maintain the forwarding table which in turn will forward packets to a down neighbor even after the BGP timers have expired causing loss of traffic.	None.
VCF-NSX-BGP-RCMD-CFG-005	Enable Inter-SR iBGP routing.	In the event that an edge node has all of its northbound eBGP sessions down, north-south traffic will continue to flow by routing traffic to a different edge node.	None.

Table 15-40. BGP Routing Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
VCF-NSX-BGP-RCMD-CFG-006	Deploy a Tier-1 gateway in non-preemptive failover mode.	Ensures that after a failed NSX Edge transport node is back online, it does not take over the gateway services thus preventing a short service outage.	None.
VCF-NSX-BGP-RCMD-CFG-007	Enable standby relocation of the Tier-1 gateway.	Ensures that if an edge failure occurs, a standby Tier-1 gateway is created on another edge node.	None.

Table 15-41. BGP Routing Design Recommendations for NSX Federation in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-BGP-RCMD-CFG-008	Use Tier-1 gateways to control the span of networks and ingress and egress traffic in the VMware Cloud Foundation instances.	Enables a mixture of network spans (isolated to a VMware Cloud Foundation instance or spanning multiple instances) without requiring additional Tier-0 gateways and hence edge nodes.	To control location span, a Tier-1 gateway must be assigned to an edge cluster and hence has the Tier-1 SR component. East-west traffic between Tier-1 gateways with SRs need to physically traverse an edge node.
VCF-NSX-BGP-RCMD-CFG-009	Allocate a Tier-1 gateway in each instance for instance-specific networks and connect it to the stretched Tier-0 gateway.	■ Creates a two-tier routing architecture. ■ Enables local-instance networks that are not to span between the VMware Cloud Foundation instances. ■ Guarantees that local-instance networks remain available if a failure occurs in another VMware Cloud Foundation instance.	None.

Overlay Design Elements for VMware Cloud Foundation

Table 15-42. Overlay Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-OVERLAY-REQD-CFG-001	Configure all ESXi hosts in the workload domain as transport nodes in NSX.	Enables distributed routing, logical segments, and distributed firewall.	None.
VCF-NSX-OVERLAY-REQD-CFG-002	Configure each ESXi host as a transport node using transport node profiles.	■ Enables the participation of ESXi hosts and the virtual machines running on them in NSX overlay and VLAN networks. ■ Transport node profiles can only be applied at the cluster level.	None.
VCF-NSX-OVERLAY-REQD-CFG-003	To provide virtualized network capabilities to workloads, use overlay networks with NSX Edge nodes and distributed routing.	■ Creates isolated, multi-tenant broadcast domains across data center fabrics to deploy elastic, logical networks that span physical network boundaries. ■ Enables advanced deployment topologies by introducing Layer 2 abstraction from the data center networks.	Requires configuring transport networks with an MTU size of at least 1,600 bytes.
VCF-NSX-OVERLAY-REQD-CFG-004	Create a single overlay transport zone in the NSX instance for all overlay traffic across the host and NSX Edge transport nodes of the workload domain.	■ Ensures that overlay segments are connected to an NSX Edge node for services and north-south routing. ■ Ensures that all segments are available to all ESXi hosts and NSX Edge nodes configured as transport nodes.	All clusters in all workload domains that share the same NSX Manager share the same transport zone.
VCF-NSX-OVERLAY-REQD-CFG-005	Create an uplink profile with a load balance source teaming policy with two active uplinks for ESXi hosts.	For increased resiliency and performance, supports the concurrent use of both physical NICs on the ESXi hosts that are configured as transport nodes.	None.

Table 15-43. Overlay Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-OVERLAY-RCMD-CFG-001	Use static IP pools to assign IP addresses to the host TEP interfaces.	■ Removes the need for an external DHCP server for the host overlay VLANs. ■ You can use NSX Manager to verify static IP pool configurations.	None.
VCF-NSX-OVERLAY-RCMD-CFG-002	Use hierarchical two-tier replication on all overlay segments.	Hierarchical two-tier replication is more efficient because it reduced the number of ESXi hosts the source ESXi host must replicate traffic to if hosts have different TEP subnets. This is typically the case with more than one cluster and will improve performance in that scenario.	None.

Table 15-44. Overlay Design Recommendations for Stretched Clusters inVMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-OVERLAY-RCMD-CFG-003	Configure an NSX sub-transport node profile.	■ You can use static IP pools for the host TEPs in each availability zone. ■ The NSX transport node profile can remain attached when using two separate VLANs for host TEPs at each availability zone as required for clusters that are based on vSphere Lifecycle Manager images. ■ Using an external DHCP server for the host overlay VLANs in both availability zones is not required.	Changes to the host transport node configuration are done at the vSphere cluster level.

Application Virtual Network Design Elements for VMware Cloud Foundation

Table 15-45. Application Virtual Network Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-AVN-REQD-CFG-001	Create one cross-instance NSX segment for the components of a VMware Aria Suite application or another solution that requires mobility between VMware Cloud Foundation instances.	Prepares the environment for the deployment of solutions on top of VMware Cloud Foundation, such as VMware Aria Suite, without a complex physical network configuration. The components of the VMware Aria Suite application must be easily portable between VMware Cloud Foundation instances without requiring reconfiguration.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-002	Create one or more local-instance NSX segments for the components of a VMware Aria Suite application or another solution that are assigned to a specific VMware Cloud Foundation instance.	Prepares the environment for the deployment of solutions on top of VMware Cloud Foundation, such as VMware Aria Suite, without a complex physical network configuration.	Each NSX segment requires a unique IP address space.

Table 15-46. Application Virtual Network Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-AVN-REQD-CFG-003	Extend the cross-instance NSX segment to the second VMware Cloud Foundation instance.	Enables workload mobility without a complex physical network configuration. The components of a VMware Aria Suite application must be easily portable between VMware Cloud Foundation instances without requiring reconfiguration.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-004	In each VMware Cloud Foundation instance, create additional local-instance NSX segments.	Enables workload mobility within a VMware Cloud Foundation instance without complex physical network configuration. Each VMware Cloud Foundation instance should have network segments to support workloads which are isolated to that VMware Cloud Foundation instance.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-005	In each VMware Cloud Foundation instance, connect or migrate the local-instance NSX segments to the corresponding local-instance Tier-1 gateway.	Configures local-instance NSX segments at required sites only.	Requires an individual Tier-1 gateway for local-instance segments.

Table 15-47. Application Virtual Network Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-AVN-RCMD-CFG-001	Use overlay-backed NSX segments.	■ Supports expansion to deployment topologies for multiple VMware Cloud Foundation instances. ■ Limits the number of VLANs required for the data center fabric.	Using overlay-backed NSX segments requires routing, eBGP recommended, between the data center fabric and edge nodes.

Load Balancing Design Elements for VMware Cloud Foundation

Table 15-48. Load Balancing Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-001	Deploy a standalone Tier-1 gateway to support advanced stateful services such as load balancing for other management components.	Provides independence between north-south Tier-1 gateways to support advanced deployment scenarios.	You must add a separate Tier-1 gateway.
VCF-NSX-LB-REQD-CFG-002	When creating load balancing services for Application Virtual Networks, connect the standalone Tier-1 gateway to the cross-instance NSX segments.	Provides load balancing to applications connected to the cross-instance network.	You must connect the gateway to each network that requires load balancing.
VCF-NSX-LB-REQD-CFG-003	Configure a default static route on the standalone Tier-1 gateway with a next hop the Tier-1 gateway for the segment to provide connectivity to the load balancer.	Because the Tier-1 gateway is standalone, it does not auto-configure its routes.	None.

Table 15-49. Load Balancing Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-004	Deploy a standalone Tier-1 gateway in the second VMware Cloud Foundation instance.	Provides a cold-standby non-global service router instance for the second VMware Cloud Foundation instance to support services on the cross-instance network which require advanced services not currently supported as NSX global objects.	■ You must add a separate Tier-1 gateway. ■ You must manually configure any services and synchronize them between the non-global service router instances in the first and second VMware Cloud Foundation instances. ■ To avoid a network conflict between the two VMware Cloud Foundation instances, make sure that the primary and standby networking services are not both active at the same time.
VCF-NSX-LB-REQD-CFG-005	Connect the standalone Tier-1 gateway in the second VMware Cloud Foundationinstance to the cross-instance NSX segment.	Provides load balancing to applications connected to the cross-instance network in the second VMware Cloud Foundation instance.	You must connect the gateway to each network that requires load balancing.

Table 15-49. Load Balancing Design Requirements for NSX Federation in VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-006	Configure a default static route on the standalone Tier-1 gateway in the second VMware Cloud Foundation instance with a next hop as the Tier-1 gateway for the segment it connects with to provide connectivity to the load balancers.	Because the Tier-1 gateway is standalone, it does not autoconfigure its routes.	None.
VCF-NSX-LB-REQD-CFG-007	Establish a process to ensure any changes made on to the load balancer instance in the first VMware Cloud Foundationinstance are manually applied to the disconnected load balancer in the second instance.	Keeps the network service in the failover load balancer instance ready for activation if a failure in the first VMware Cloud Foundation instance occurs. Because network services are not supported as global objects, you must configure them manually in each VMware Cloud Foundation instance. The load balancer service in one instance must be connected and active, while the service in the other instance must be disconnected and inactive.	■ Because of incorrect configuration between the VMware Cloud Foundation instances, the load balancer service in the second instance might come online with an invalid or incomplete configuration. ■ If both VMware Cloud Foundation instances are online and active at the same time, a conflict between services could occur resulting in a potential outage. ■ The administrator must establish and follow an operational practice by using a runbook or automated process to ensure that configuration changes are reproduced in each VMware Cloud Foundation instance.

SDDC Manager Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to SDDC Manager in an environment with a single or multiple VMware Cloud Foundation instances.

For full design details, see Chapter 9 SDDC Manager Design for VMware Cloud Foundation.

Table 15-50. SDDC Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-SDDCMGR-REQD-CFG-001	Deploy an SDDC Manager system in the first availability zone of the management domain.	SDDC Manager is required to perform VMware Cloud Foundation capabilities, such as provisioning VI workload domains, deploying solutions, patching, upgrading, and others.	None.
VCF-SDDCMGR-REQD-CFG-002	Deploy SDDC Manager with its default configuration.	The configuration of SDDC Manager is not configurable and should not be changed from its defaults.	None.
VCF-SDDCMGR-REQD-CFG-003	Place the SDDC Manager appliance on the VM management network.	■ Simplifies IP addressing for management VMs by using the same VLAN and subnet. ■ Provides simplified secure access to management VMs in the same VLAN network.	None.

Table 15-51. SDDC Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-SDDCMGR-RCMD-CFG-001	Connect SDDC Manager to the Internet for downloading software bundles.	SDDC Manager must be able to download install and upgrade software bundles for deployment of VI workload domains and solutions, and for upgrade from a repository.	The rules of your organization might not permit direct access to the Internet. In this case, you must download software bundles for SDDC Manager manually.
VCF-SDDCMGR-RCMD-CFG-002	Configure a network proxy to connect SDDC Manager to the Internet.	To protect SDDC Manager against external attacks from the Internet.	The proxy must not use authentication because SDDC Manager does not support proxy with authentication.

Table 15-51. SDDC Manager Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-SDDCMGR-RCMD-CFG-003	Configure SDDC Manager with a VMware Customer Connect account with VMware Cloud Foundation entitlement to check for and download software bundles.	Software bundles for VMware Cloud Foundation are stored in a repository that is secured with access controls.	Requires the use of a VMware Customer Connect user account with access to VMware Cloud Foundation licensing. Sites without an internet connection can use local upload option instead.
VCF-SDDCMGR-RCMD-CFG-004	Configure SDDC Manager with an external certificate authority that is responsible for providing signed certificates.	Provides increased security by implementing signed certificate generation and replacement across the management components.	An external certificate authority, such as Microsoft CA, must be locally available.

VMware Aria Suite Lifecycle Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to VMware Aria Suite Lifecycle in an environment with a single or multiple VMware Cloud Foundation instances.

For full design details, see Chapter 10 VMware Aria Suite Lifecycle Design for VMware Cloud Foundation.

Table 15-52. VMware Aria Suite Lifecycle Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-001	Deploy a VMware Aria Suite Lifecycle instance in the management domain of each VMware Cloud Foundation instance to provide life cycle management for VMware Aria Suite and Workspace ONE Access.	Provides life cycle management operations for VMware Aria Suite applications and Workspace ONE Access.	You must ensure that the required resources are available.
VCF-VASL-REQD-CFG-002	Deploy VMware Aria Suite Lifecycle by using SDDC Manager.	■ Deploys VMware Aria Suite Lifecycle in VMware Cloud Foundation mode, which enables the integration with the SDDC Manager inventory for product deployment and life cycle management of VMware Aria Suite components. ■ Automatically configures the standalone Tier-1 gateway required for load balancing the clustered Workspace ONE Access and VMware Aria Suite components.	None.
VCF-VASL-REQD-CFG-003	Allocate extra 100 GB of storage to the VMware Aria Suite Lifecycle appliance for VMware Aria Suite product binaries.	■ Provides support for VMware Aria Suite product binaries (install, upgrade, and patch) and content management. ■ SDDC Manager automates the creation of storage.	None.
VCF-VASL-REQD-CFG-004	Place the VMware Aria Suite Lifecycle appliance on an overlay-backed (recommended) or VLAN-backed NSX network segment.	Provides a consistent deployment model for management applications.	You must use an implementation in NSX to support this networking configuration.
VCF-VASL-REQD-CFG-005	Import VMware Aria Suite product licenses to the Locker repository for product life cycle operations.	■ You can review the validity, details, and deployment usage for the license across the VMware Aria Suite products. ■ You can reference and use licenses during product life cycle operations, such as deployment and license replacement.	When using the API, you must specify the Locker ID for the license to be used in the JSON payload.

Table 15-52. VMware Aria Suite Lifecycle Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-ENV-001	Configure datacenter objects in VMware Aria Suite Lifecycle for local and cross-instance VMware Aria Suite deployments and assign the management domain vCenter Server instance to each data center.	You can deploy and manage the integrated VMware Aria Suite components across the SDDC as a group.	You must manage a separate datacenter object for the products that are specific to each instance.
VCF-VASL-REQD-ENV-002	If deploying VMware Aria Operations for Logs, create a local-instance environment in VMware Aria Suite Lifecycle.	Supports the deployment of an instance of VMware Aria Operations for Logs.	None.
VCF-VASL-REQD-ENV-003	If deploying VMware Aria Operations or VMware Aria Automation, create a cross-instance environment in VMware Aria Suite Lifecycle	■ Supports deployment and management of the integrated VMware Aria Suite products across VMware Cloud Foundation instances as a group. ■ Enables the deployment of instance-specific components, such as VMware Aria Operations remote collectors. In VMware Aria Suite Lifecycle, you can deploy and manage VMware Aria Operations remote collector objects only in an environment that contains the associated cross-instance components.	You can manage instance-specific components, such as remote collectors, only in an environment that is cross-instance.

Table 15-52. VMware Aria Suite Lifecycle Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-SEC-001	Use the custom vCenter Server role for VMware Aria Suite Lifecycle that has the minimum privileges required to support the deployment and upgrade of VMware Aria Suite products.	VMware Aria Suite Lifecycle accesses vSphere with the minimum set of permissions that are required to support the deployment and upgrade of VMware Aria Suite products. SDDC Manager automates the creation of the custom role.	You must maintain the permissions required by the custom role.
VCF-VASL-REQD-SEC-002	Use the service account in vCenter Server for application-to-application communication from VMware Aria Suite Lifecycle to vSphere. Assign global permissions using the custom role.	■ Provides the following access control features: ■ VMware Aria Suite Lifecycle accesses vSphere with the minimum set of required permissions. ■ You can introduce improved accountability in tracking request-response interactions between the components of the SDDC. ■ SDDC Manager automates the creation of the service account.	■ You must maintain the life cycle and availability of the service account outside of SDDC manager password rotation.

Table 15-53. VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-006	For multiple availability zones, add the VMware Aria Suite Lifecycle appliance to the VM group for the first availability zone.	Ensures that, by default, the VMware Aria Suite Lifecycle appliance is powered on a host in the first availability zone.	If VMware Aria Suite Lifecycle is deployed after the creation of the stretched management cluster, you must add the VMware Aria Suite Lifecycle appliance to the VM group manually.

Table 15-54. VMware Aria Suite Lifecycle Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-007	Configure the DNS settings for the VMware Aria Suite Lifecycle appliance to use DNS servers in each instance.	Improves resiliency in the event of an outage of external services for a VMware Cloud Foundation instance.	As you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the DNS settings of the VMware Aria Suite Lifecycle appliance must be updated.
VCF-VASL-REQD-CFG-008	Configure the NTP settings for the VMware Aria Suite Lifecycle appliance to use NTP servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	As you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the NTP settings on the VMware Aria Suite Lifecycle appliance must be updated.
VCF-VASL-REQD-ENV-004	Assign the management domain vCenter Server instance in the additional VMware Cloud Foundation instance to the cross-instance data center.	Supports the deployment of VMware Aria Operations remote collectors in an additional VMware Cloud Foundation instance.	None.

Table 15-55. VMware Aria Suite Lifecycle Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VASL-RCMD-CFG-001	Protect VMware Aria Suite Lifecycle by using vSphere HA.	Supports the availability objectives for VMware Aria Suite Lifecycle without requiring manual intervention during a failure event.	None.
VCF-VASL-RCMD-LCM-001	Obtain product binaries for install, patch, and upgrade in VMware Aria Suite Lifecycle from VMware Customer Connect.	■ You can upgrade VMware Aria Suite products based on their general availability and endpoint interoperability rather than being listed as part of VMware Cloud Foundation bill of materials (BOM). ■ You can deploy and manage binaries in an environment that does not allow access to the Internet or are dark sites.	The site must have an Internet connection to use VMware Customer Connect. Sites without an Internet connection should use the local upload option instead.

Table 15-55. VMware Aria Suite Lifecycle Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VASL-RCMD-LCM-002	Use support packs (PSPAKS) for VMware Aria Suite Lifecycle to enable upgrading to later versions of VMware Aria Suite products.	Enables the upgrade of an existing VMware Aria Suite Lifecycle to permit later versions of VMware Aria Suite products without an associated VMware Cloud Foundation upgrade. See VMware Knowledge Base article 88829	None.
VCF-VASL-RCMD-SEC-001	Enable integration between VMware Aria Suite Lifecycle and your corporate identity source by using the Workspace ONE Access instance.	■ Enables authentication to VMware Aria Suite Lifecycle by using your corporate identity source. ■ Enables authorization through the assignment of organization and cloud services roles to enterprise users and groups defined in your corporate identity source.	You must deploy and configure Workspace ONE Access to establish the integration between VMware Aria Suite Lifecycle and your corporate identity sources.
VCF-VASL-RCMD-SEC-002	Create corresponding security groups in your corporate directory services for VMware Aria Suite Lifecycle roles: ■ **VCF** ■ **Content Release Manager** ■ **Content Developer**	Streamlines the management of VMware Aria Suite Lifecycle roles for users.	■ You must create the security groups outside of the SDDC stack. ■ You must set the desired directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.

Workspace ONE Access Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related toWorkspace ONE Access in an environment with a single or multiple VMware Cloud Foundation instances. The design elements also considers whether the management domain has a single or multiple availability zones.

For full design details, see Chapter 11 Workspace ONE Access Design for VMware Cloud Foundation.

Table 15-56. Workspace ONE Access Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-ENV-001	Create a global environment in VMware Aria Suite Lifecycle to support the deployment of Workspace ONE Access.	A global environment is required by VMware Aria Suite Lifecycle to deploy Workspace ONE Access.	None.
VCF-WSA-REQD-SEC-001	Import certificate authority-signed certificates to the Locker repository for Workspace ONE Access product life cycle operations.	■ You can reference and use certificate authority-signed certificates during product life cycle operations, such as deployment and certificate replacement.	When using the API, you must specify the Locker ID for the certificate to be used in the JSON payload.
VCF-WSA-REQD-CFG-001	Deploy an appropriately sized Workspace ONE Access instance according to the deployment model you have selected by using VMware Aria Suite Lifecycle in VMware Cloud Foundation mode.	The Workspace ONE Access instance is managed by VMware Aria Suite Lifecycle and imported into the SDDC Manager inventory.	None.
VCF-WSA-REQD-CFG-002	Place the Workspace ONE Access appliances on an overlay-backed or VLAN-backed NSX network segment.	Provides a consistent deployment model for management applications in an environment with a single or multiple VMware Cloud Foundation instances.	You must use an implementation in NSX to support this network configuration.
VCF-WSA-REQD-CFG-003	Use the embedded PostgreSQL database with Workspace ONE Access.	Removes the need for external database services.	None.
VCF-WSA-REQD-CFG-004	Add a VM group for Workspace ONE Access and set VM rules to restart the Workspace ONE Access VM group before any of the VMs that depend on it for authentication.	You can define the startup order of virtual machines regarding the service dependency. The startup order ensures that vSphere HA powers on the Workspace ONE Access virtual machines in an order that respects product dependencies.	None.

Table 15-56. Workspace ONE Access Design Requirements for VMware Cloud Foundation (continued)

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-005	Connect the Workspace ONE Access instance to a supported upstream Identity Provider.	You can integrate your enterprise directory with Workspace ONE Access to synchronize users and groups to the Workspace ONE Access identity and access management services.	None.
VCF-WSA-REQD-CFG-006	If using clustered Workspace ONE Access, configure second and third native connectors that correspond to the second and third Workspace ONE Access cluster nodes to support the high availability of directory services access.	Adding the additional native connectors provides redundancy and improves performance by load-balancing authentication requests.	Each of the Workspace ONE Access cluster nodes must be joined to the Active Directory domain to use Active Directory with Integrated Windows Authentication with the native connector.
VCF-WSA-REQD-CFG-007	If using clustered Workspace ONE Access, use the NSX load balancer that is configured by SDDC Manager on a dedicated Tier-1 gateway.	■ During the deployment of Workspace ONE Access by using VMware Aria Suite Lifecycle, SDDC Manager automates the configuration of an NSX load balancer for Workspace ONE Access to facilitate scale-out.	You must use the load balancer that is configured by SDDC Manager and the integration with VMware Aria Suite Lifecycle.

Table 15-57. Workspace ONE Access Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-008	Add the Workspace ONE Access appliances to the VM group for the first availability zone.	Ensures that, by default, the Workspace ONE Access cluster nodes are powered on a host in the first availability zone.	■ If the Workspace ONE Access instance is deployed after the creation of the stretched management cluster, you must add the appliances to the VM group manually. ■ ClusteredWorkspace ONE Access might require manual intervention after a failure of the active availability zone occurs.

Table 15-58. Workspace ONE Access Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-009	Configure the DNS settings for Workspace ONE Access to use DNS servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	None.
VCF-WSA-REQD-CFG-010	Configure the NTP settings on Workspace ONE Access cluster nodes to use NTP servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	If you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the NTP settings on Workspace ONE Access must be updated.

Table 15-59. Workspace ONE Access Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-CFG-001	Protect all Workspace ONE Access nodes using vSphere HA.	Supports high availability for Workspace ONE Access.	None for standard deployments. Clustered Workspace ONE Access deployments might require intervention if an ESXi host failure occurs.
VCF-WSA-RCMD-CFG-002	When using Active Directory as an Identity Provider, use Active Directory over LDAP as the Directory Service connection option.	The native (embedded) Workspace ONE Access connector binds to Active Directory over LDAP using a standard bind authentication.	■ In a multi-domain forest, where the Workspace ONE Access instance connects to a child domain, Active Directory security groups must have global scope. Therefore, members added to the Active Directory global security group must reside within the same Active Directory domain. ■ If authentication to more than one Active Directory domain is required, additional Workspace ONE Access directories are required.
VCF-WSA-RCMD-CFG-003	When using Active Directory as an Identity Provider, use an Active Directory user account with a minimum of read-only access to Base DNs for users and groups as the service account for the Active Directory bind.	Provides the following access control features: ■ Workspace ONE Access connects to the Active Directory with the minimum set of required permissions to bind and query the directory. ■ You can introduce improved accountability in tracking request-response interactions between the Workspace ONE Access and Active Directory.	■ You must manage the password life cycle of this account. ■ If authentication to more than one Active Directory domain is required, additional accounts are required for the Workspace ONE Access connector to bind to each Active Directory domain over LDAP.

Table 15-59. Workspace ONE Access Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-CFG-004	Configure the directory synchronization to synchronize only groups required for the integrated SDDC solutions.	■ Limits the number of replicated groups required for each product. ■ Reduces the replication interval for group information.	You must manage the groups from your enterprise directory selected for synchronization to Workspace ONE Access.
VCF-WSA-RCMD-CFG-005	Activate the synchronization of enterprise directory group members when a group is added to the Workspace ONE Access directory.	When activated, members of the enterprise directory groups are synchronized to the Workspace ONE Access directory when groups are added. When deactivated, group names are synchronized to the directory, but members of the group are not synchronized until the group is entitled to an application or the group name is added to an access policy.	None.
VCF-WSA-RCMD-CFG-006	Enable Workspace ONE Access to synchronize nested group members by default.	Allows Workspace ONE Access to update and cache the membership of groups without querying your enterprise directory.	Changes to group membership are not reflected until the next synchronization event.
VCF-WSA-RCMD-CFG-007	Add a filter to the Workspace ONE Access directory settings to exclude users from the directory replication.	Limits the number of replicated users for Workspace ONE Access within the maximum scale.	To ensure that replicated user accounts are managed within the maximums, you must define a filtering schema that works for your organization based on your directory attributes.

Table 15-59. Workspace ONE Access Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-CFG-008	Configure the mapped attributes included when a user is added to the Workspace ONE Access directory.	You can configure the minimum required and extended user attributes to synchronize directory user accounts for the Workspace ONE Access to be used as an authentication source for cross-instance VMware Aria Suite solutions.	User accounts in your organization's enterprise directory must have the following required attributes mapped: - `firstname`, for example, `givenname` for Active Directory - `lastName`, for example, `sn` for Active Directory - `email`, for example, `mail` for Active Directory - `userName`, for example,`sAMAccountName` for Active Directory - If you require users to sign in with an alternate unique identifier, for example, `userPrincipalName`, you must map the attribute and update the identity and access management preferences.
VCF-WSA-RCMD-CFG-009	Configure the Workspace ONE Access directory synchronization frequency to a reoccurring schedule, for example, 15 minutes.	Ensures that any changes to group memberships in the corporate directory are available for integrated solutions in a timely manner.	Schedule the synchronization interval to be longer than the time to synchronize from the enterprise directory. If users and groups are being synchronized to Workspace ONE Access when the next synchronization is scheduled, the new synchronization starts immediately after the end of the previous iteration. With this schedule, the process is continuous.

Table 15-59. Workspace ONE Access Design Recommendations for VMware Cloud Foundation (continued)

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-SEC-001	Create corresponding security groups in your corporate directory services for these Workspace ONE Access roles: - **Super Admin** - **Directory Admins** - **ReadOnly Admin**	Streamlines the management of Workspace ONE Access roles to users.	- You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period. - You must create the security group outside of the SDDC stack.
VCF-WSA-RCMD-SEC-002	Configure a password policy for Workspace ONE Access local directory users, **admin** and **configadmin**.	You can set a policy for Workspace ONE Access local directory users that addresses your corporate policies and regulatory standards. The password policy is applicable only to the local directory users and does not impact your organization directory.	You must set the policy in accordance with your organization policies and regulatory standards, as applicable. You must apply the password policy on the Workspace ONE Access cluster nodes.

Life Cycle Management Design Elements for VMware Cloud Foundation

Use this list of requirements for reference related to life cycle management in a VMware Cloud Foundation environment.

For full design details, see Chapter 12 Life Cycle Management Design for VMware Cloud Foundation.

Table 15-60. Life Cycle Management Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-001	Use SDDC Manager to perform the life cycle management of the following components: ■ SDDC Manager ■ NSX Manager ■ NSX Edges ■ vCenter Server ■ ESXi	Because the deployment scope of SDDC Manager covers the full VMware Cloud Foundation stack, SDDC Manager performs patching, update, or upgrade of these components across all workload domains.	The operations team must understand and be aware of the impact of a patch, update, or upgrade operation by using SDDC Manager.
VCF-LCM-REQD-002	Use VMware Aria Suite Lifecycle to manage the life cycle of the following components: ■ VMware Aria Suite Lifecycle ■ Workspace ONE Access	VMware Aria Suite Lifecycle automates the life cycle of VMware Aria Suite Lifecycle and Workspace ONE Access.	■ You must deploy VMware Aria Suite Lifecycle by using SDDC Manager. ■ You must manually apply Workspace ONE Access patches, updates, and hotfixes. Patches, updates, and hotfixes for Workspace ONE Access are not generally managed by VMware Aria Suite Lifecycle.
VCF-LCM-RCMD-001	Use vSphere Lifecycle Manager images to manage the life cycle of vSphere clusters.	■ With vSphere Lifecycle Manager images, firmware updates are carried out through firmware and driver add-ons, which you add to the image you use to manage a cluster. ■ You can check the hardware compatibility of the hosts in a cluster against the VMware Compatibility Guide. ■ You can validate a vSphere Lifecycle Manager image to check if it applies to all hosts in the cluster. You can also perform a remediation pre-check.	■ Updating the firmware with images requires an OEM-provided hardware support manager plug-in, which integrates with vSphere Lifecycle Manager. ■ An updated vSAN Hardware Compatibility List (vSAN HCL) is required during bring-up.

Table 15-61. Life Cycle Management Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-003	Use the upgrade coordinator in NSX to perform life cycle management on the NSX Global Manager appliances.	The version of SDDC Manager in this design is not currently capable of life cycle operations (patching, update, or upgrade) for NSX Global Manager.	■ You must explicitly plan upgrades of the NSX Global Manager nodes. An upgrade of the NSX Global Manager nodes might require a cascading upgrade of the NSX Local Manager nodes and underlying SDDC Manager infrastructure before upgrading the NSX Global Manager nodes. ■ You must always align the version of the NSX Global Manager nodes with the rest of the SDDC stack in VMware Cloud Foundation.
VCF-LCM-REQD-004	Establish an operations practice to ensure that prior to the upgrade of any workload domain, the impact of any version upgrades is evaluated in relation to the need to upgrade NSX Global Manager.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.
VCF-LCM-REQD-005	Establish an operations practice to ensure that prior to the upgrade of the NSX Global Manager, the impact of any version change is evaluated against the existing NSX Local Manager nodes and workload domains.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.

Information Security Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the management of access controls, certificates and accounts in a VMware Cloud Foundation environment.

For full design details, see Chapter 14 Information Security Design for VMware Cloud Foundation.

Table 15-62. Design Requirements for Account and Password Management for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-ACTMGT-REQD-SEC-001	Enable scheduled password rotation in SDDC Manager for all accounts supporting scheduled rotation.	■ Increases the security posture of your SDDC. ■ Simplifies password management across your SDDC management components.	You must retrieve new passwords by using the API if you must use accounts interactively.
VCF-ACTMGT-REQD-SEC-003	Establish operational practice to rotate passwords using SDDC Manager on components that do not support scheduled rotation in SDDC Manager.	Rotates passwords and automatically remediates SDDC Manager databases for those user accounts.	None.
VCF-ACTMGT-REQD-SEC-003	Establish operational practice to manually rotate passwords on components that cannot be rotated by SDDC Manager.	Maintains password policies across components not handled by SDDC Manager password management.	None.

Table 15-63. Certificate Management Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-SDDC-RCMD-SEC-001	Replace the default VMCA or signed certificates on all management virtual appliances with a certificate that is signed by an internal certificate authority.	Ensures that the communication to all management components is secure.	Replacing the default certificates with trusted CA-signed certificates from a certificate authority might increase the deployment preparation time because you must generate and submit certificate requests.
VCF-SDDC-RCMD-SEC-002	Use a SHA-2 algorithm or higher for signed certificates.	The SHA-1 algorithm is considered less secure and has been deprecated.	Not all certificate authorities support SHA-2 or higher.
VCF-SDDC-RCMD-SEC-003	Perform SSL certificate life cycle management for all management appliances by using SDDC Manager.	SDDC Manager supports automated SSL certificate lifecycle management rather than requiring a series of manual steps.	Certificate management for NSX Global Manager instances must be done manually.