

Architecture and Design

VMware Validated Design 4.0

VMware Validated Design for Remote Office Branch
Office 4.0



vmware®

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

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

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

Copyright © 2016, 2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About VMware Validated Design Architecture and Design for Remote Office and Branch Office 5

1 Remote Office and Branch Office Deployment Models 6

2 Architecture Overview in ROBO 9

- Physical Infrastructure Architecture in ROBO 11
 - Pod Architecture in ROBO 11
 - Pod Types in ROBO 12
 - Physical Network Architecture in ROBO 13
 - Regions, Hubs and Remote and Branch Offices in ROBO 15
- Virtual Infrastructure Architecture in ROBO 16
 - Virtual Infrastructure Overview in ROBO 17
 - Network Virtualization Components in ROBO 19
 - Network Virtualization Services in ROBO 20
- Cloud Management Platform Architecture in ROBO 23
 - Logical Architecture of the Cloud Management Platform 24
 - Cloud Management Platform Logical Architecture 25
- Operations Architecture in ROBO 27
 - Operations Management Architecture in ROBO 28
 - Logging Architecture in ROBO 31
 - Data Protection and Backup Architecture in ROBO 34

3 Detailed Design in ROBO 36

- Physical Infrastructure Design in ROBO 39
 - Physical Design Fundamentals in ROBO 40
 - Physical Networking Design in ROBO 44
 - Physical Storage Design in ROBO 48
- Virtual Infrastructure Design in ROBO 55
 - ESXi Design in ROBO 58
 - vCenter Server Design in ROBO 60
 - Virtualization Network Design in ROBO 70
 - NSX Design 81
 - Shared Storage Design in ROBO 100
- Cloud Management Platform Design in ROBO 117
 - vRealize Automation Design in ROBO 118
- Operations Infrastructure Design in ROBO 135
 - vRealize Operations Manager Design in ROBO 136

[vRealize Log Insight Design in ROBO](#) 142

[vSphere Data Protection Design in ROBO](#) 153

[vSphere Update Manager in ROBO](#) 159

About VMware Validated Design Architecture and Design for Remote Office and Branch Office

VMware Validated Design Architecture for VMware Validated Design™ Remote Office and Branch Office contains a validated model of the Software-Defined Data Center (SDDC) for remote office and branch office (ROBO) locations, and provides a detailed design of each of the management components in the remote office and branch office Software-Defined Data Center (ROBO SDDC) stack.

[Chapter 2 Architecture Overview in ROBO](#) discusses the building blocks and main principles of each ROBO SDDC management layer. [Chapter 3 Detailed Design in ROBO](#) provides the available design options according to the design objective, and a set of design decisions to justify selecting the path for building each ROBO SDDC component.

Required VMware Software

VMware Validated Design Architecture and Design is compliant and validated with certain product versions. See *VMware Validated Design Release Notes* for more information about supported product versions.

Intended Audience

VMware Validated Design Architecture and Design is intended for cloud architects, infrastructure administrators and cloud administrators who are familiar with and want to use VMware software to deploy in a short time and manage an SDDC that meets the requirements for capacity, scalability and backup and restore.

Remote Office and Branch Office Deployment Models

1

The ROBO SDDC provides two deployment models. The first is a centralized model in which the management of ROBO sites is performed from a central hub. The second is a decentralized model in which each ROBO site is managed locally from within itself.

A ROBO location can be a storefront (point of sale), manufacturing facility, medical office, or other IT infrastructure located in a remote, distributed site. These locations are critical to the business but have a smaller number of workloads compared to the corporate datacenter, typically running 100 or less workload virtual machines. Due to the critical nature of the workloads, there should be limited dependencies on a Wide Area Network (WAN) connection. This reduces the impact on workloads should a WAN outage occur.

The VMware Validated Design for ROBO uses the VMware Validated Design for SDDC as the starting point. You must first have the VMware Validated Design for SDDC deployed in either a single or a dual region deployment. The ROBO Validated Design utilizes the provisioning and monitoring components from the SDDC Validated Design. In this VMware Validated Design for ROBO the decentralized management approach is used.

In the centralized model, all management, provisioning, and monitoring components are located in a region deployed as part of the VVD for SDDC. The only components located at the ROBO site are ESXi hosts. The centralized model has the following considerations.

Figure 1-1. Centralized Management for Remote Office and Branch Office

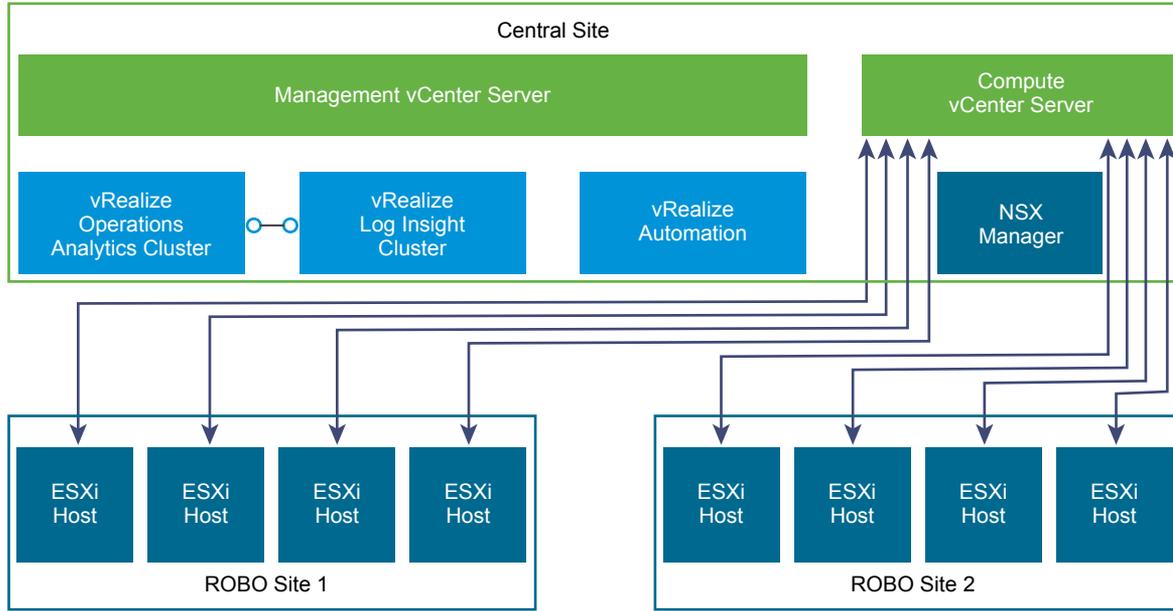


Table 1-1. Centralized Management

Pros	Cons
Single pane of glass management.	Large fault domain.
Centralized lifecycle management operations such as patching and upgrading.	Patching and upgrading involves coordinating downtime windows for management components in all locations.
Less management virtual machine footprint.	Patching and upgrading is a higher risk operation due to the large fault domain.
Rapid and reduced complexity site deployment.	WAN outage leaves the ESXi host disconnected, virtual machine workloads continue to run but can only be managed locally via Host Client or API/CLI access for basic operations. No provisioning via vCenter or CMP is possible. NSX management changes are not possible, however the network data plane continues to function.
	Doesn't allow for the ability to add local disaster recovery.

In the decentralized model, core management components, such as vCenter Server and NSX Manager, are installed locally along with a vRealize Log Insight instance for logging in each ROBO site. This allows for independent management of each ROBO site while leveraging the centralized provisioning and monitoring capabilities of a region deployed as part of the VVD for SDDC. The decentralized model has the following considerations.

Figure 1-2. Decentralized Management for Remote Office and Branch Office

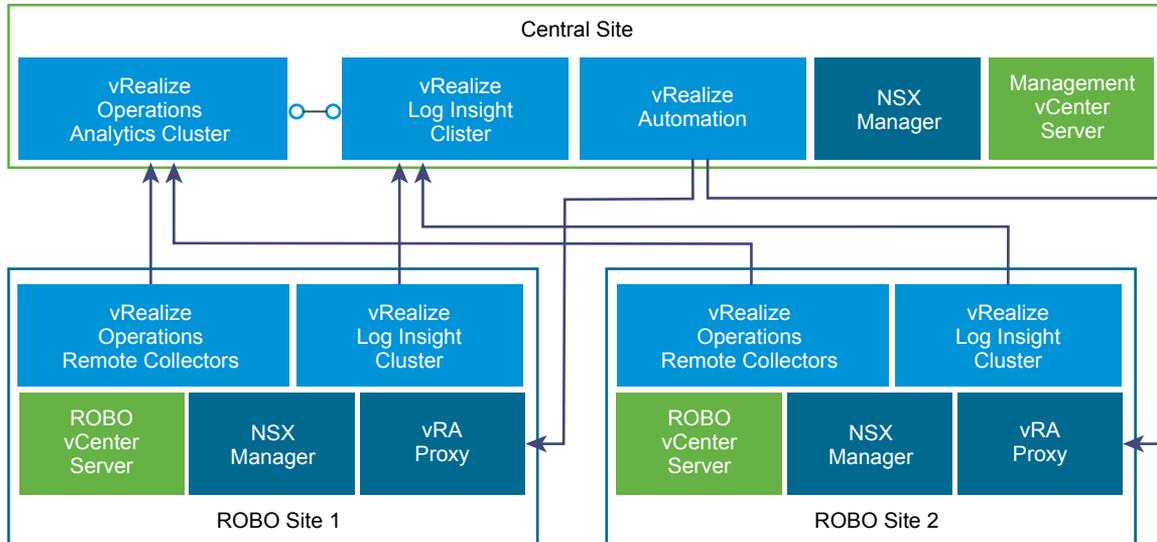


Table 1-2. Decentralized Management

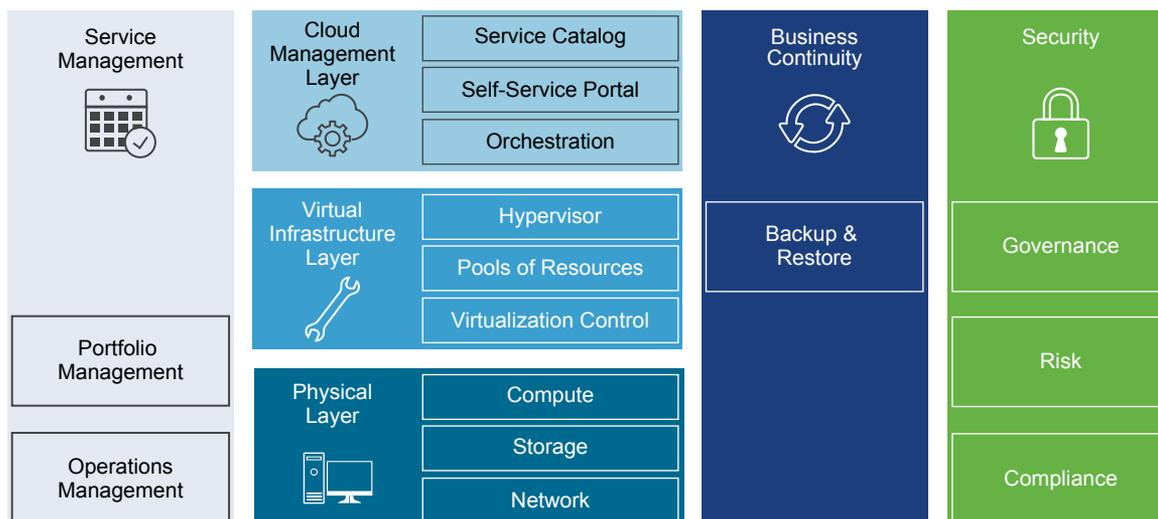
Pros	Cons
A WAN outage does not impact local management or backup operations.	No Single pane of glass for vSphere and NSX management.
Patching and upgrades is less of a risk due to smaller fault domains.	Larger management virtual machine footprint.
Log data is available locally for troubleshooting even when WAN connection is down.	Increased product licensing cost due to deployment at each ROBO.
Utilizes the central provisioning process.	Additional management components to patch and upgrade.
Log data is forwarded to the centralized vRealize Log Insight infrastructure allowing a single pane of glass.	More complex design to deploy and operate.
Event, alert, and utilization monitoring utilizes the central vRealize Operations Manager instance.	
vRealize Operations Manager Remote Collectors are placed in the ROBO so data collection occurs even during a WAN outage.	
Can add local disaster recovery.	

Architecture Overview in ROBO

The VMware Validated Design for Remote Offices and Branch Offices (ROBO) enables IT organizations to automate the provisioning of common, repeatable requests to multiple remote office and branch office locations from a hub. This enables IT to be to respond to business needs with more agility and predictability. Traditionally this has been referred to as Infrastructure as a Service (IAAS), however, the VMware Validated Design for Remote Offices and Branch Offices extends the typical IAAS solution to include a broader and more complete solution.

The VMware Validated Design architecture is based on a number of layers and modules that allows interchangeable components be part of the end solution or outcome, such as the SDDC. If a particular component design does not fit a business or technical requirement for whatever reason, it can be replaced with another, similar component. VMware Validated Designs are one way of putting an architecture together. They are rigorously tested to ensure stability, scalability and compatibility. Ultimately, the system is designed in such a way as to ensure that the desired IT outcome will be achieved.

Figure 2-1. Architecture Overview



Physical Layer

The lowest layer of the solution is the Physical Layer, sometimes referred to as the "core," which consists of three main components: compute, network and storage. Within the compute component there are the x86-based servers that run the management, edge, and tenant compute workloads. There is guidance around the physical capabilities required to run the architecture, however, no recommendations on the type or brand of hardware is given. All components must be supported on the *VMware Hardware Compatibility* guide.

Virtual Infrastructure Layer

Sitting on the Physical Layer components is the Virtual Infrastructure Layer. Within the Virtual Infrastructure Layer, access to the physical underlying infrastructure is controlled and allocated to the management and tenant workloads. The Virtual Infrastructure Layer consists primarily of the physical host's hypervisor and the control of these hypervisors. The management workloads consist of elements in the virtual management layer itself, along with elements in the Cloud Management Layer, Service Management, Business Continuity, and Security areas.

Cloud Management Layer

The Cloud Management Layer is the "top" layer of the stack and is where service consumption occurs. This layer calls for resources and then orchestrates the actions of the lower layers to achieve the request, most commonly by means of a user interface or application programming interface (API). While the SDDC can stand on its own without any additional, ancillary services, for a complete SDDC experience other supporting components are needed. The Service Management, Business Continuity and Security areas complete the architecture by providing this support.

Service Management

When building any type of IT infrastructure, portfolio and operations management play key roles in continued day-to-day service delivery. The Service Management area of this architecture focuses mainly on operations management, in particular monitoring, alerting and log management.

Business Continuity

To ensure that a system is enterprise ready, it must contain elements to support business continuity in the area of data backup, restoration, and disaster recovery. Business continuity ensures that when data loss occurs, the right elements are in place to prevent permanent losses to the business. The design provides comprehensive guidance on how to operate backup and restore functions, along with run books detailing how to fail over components in the event of a disaster.

Security

All systems need to be inherently secure by design. This reduces risk and increases compliance while providing a governance structure. The security area outlines what is needed to ensure that the SDDC is resilient to both internal and external threats.

This chapter includes the following topics:

- [Physical Infrastructure Architecture in ROBO](#)
- [Virtual Infrastructure Architecture in ROBO](#)
- [Cloud Management Platform Architecture in ROBO](#)
- [Operations Architecture in ROBO](#)

Physical Infrastructure Architecture in ROBO

The architecture of the data center physical layer is based on logical hardware pods and a redundant network topology for high availability.

Pod Architecture in ROBO

The VMware Validated Design for Remote Office and Branch Office (ROBO) uses a small set of common building blocks called pods.

Pod Architecture Characteristics

Pods can include different combinations of servers, storage equipment, and network equipment, and can be set up with varying levels of hardware redundancy and varying quality of components. Pods are connected to a network core that distributes data between them. A pod is not defined by any hard, physical properties, as it is a standard unit of connected elements within the SDDC network fabric.

A pod is a logical boundary of functionality for the SDDC platform. While each pod typically spans one rack, it is possible to aggregate multiple pods into a single rack or for a pod to span multiple racks. For all pods, homogeneity and ease of replication are important.

Pod to Rack Mapping

Pods are not mapped one-to-one to 19" data center racks. While a pod is an atomic unit of a repeatable building block, a rack is merely a unit of size. Because pods can have different sizes, how pods are mapped to 19" data center racks depends on the use case.

Note The consolidated pod cannot span racks. This is due to NSX controllers and other virtual machines on a VLAN backed network migrating to a different rack, where that IP subnet is not available due to layer 2 termination at the Top of Rack switch.

Pod Types in ROBO

SDDC ROBO differentiates between different types of pods including the management pod, consolidated pod and storage pod.

Management Pod

The management pod resides in the hub (defined in [About Hubs and Remote Office and Branch Office Sites](#)) and runs the virtual machines that manage the SDDC. The management pod is instantiated as part of the VMware Validated Design for SDDC. The management pod is a pre-requisite for creating consolidated pod's in ROBO locations. For additional information on the management pod refer to the VMware Validated Design for SDDC.

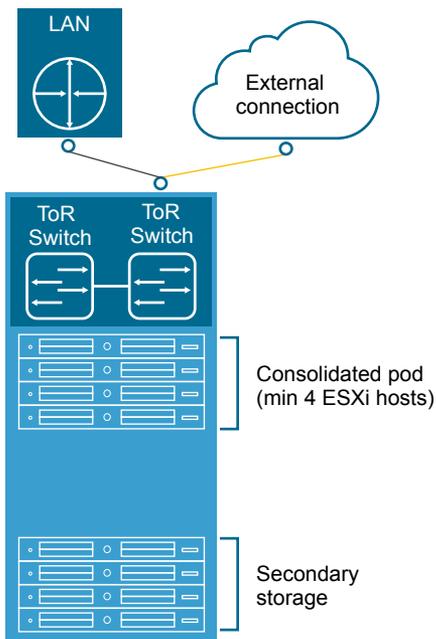
Consolidated Pod

The consolidated pod runs the following services in the ROBO location:

- Virtual machines to manage the SDDC such as vCenter Server, NSX components, Remote Collectors, agents and other shared components.
- Required NSX services to enable north-south routing between the SDDC and the external network, and east-west routing inside the SDDC.
- SDDC tenant virtual machines to support workloads of different Service Level Agreements (SLAs).

Because this pod supports all SDDC, network, and production workloads at a remote office, it is important to ensure highly available physical components such as HVAC, power feeds and power supplies.

Figure 2-2. Consolidated Pod



Storage Pod

Storage pods provide network-accessible storage using NFS or iSCSI. Different types of storage pods can provide different levels of SLA, ranging from just a bunch of disks (JBODs) using IDE drives with minimal to no redundancy, to fully redundant enterprise-class storage arrays. For bandwidth-intense IP-based storage, the bandwidth of these pods can scale dynamically.

Note The VMware Validated Design for ROBO uses VMware vSAN as its primary storage platform, and does not consider block or file storage technology for primary storage. These storage technologies are only referenced for specific use cases such as backups to secondary storage.

Physical Network Architecture in ROBO

The physical network architecture for ROBO has specific requirements. The following sections describe these requirements.

Network Transport in ROBO

You can implement the physical layer switch fabric for a SDDC by offering Layer 2 transport services or Layer 3 transport services to all components. For a scalable and vendor-neutral data center network, use Layer 3 transport.

Benefits and Drawbacks for Layer 2 Transport

In a design that uses Layer 2 transport, leaf switches and spine switches form a switched fabric, effectively acting like one large switch. Using modern data center switching fabric products such as Cisco FabricPath, you can build highly scalable Layer 2 multipath networks without the Spanning Tree Protocol (STP). Such networks are particularly suitable for large virtualization deployments, private clouds, and high-performance computing (HPC) environments.

Layer 2 routing has the following benefits and drawbacks:

- The benefit of this approach is more design freedom. You can span VLANs, which is useful for vSphere vMotion or vSphere Fault Tolerance (FT).
- The drawback is that the size of such a deployment is limited because the fabric elements have to share a limited number of VLANs. In addition, you have to rely on a specialized data center switching fabric product from a single vendor because these products are not designed for interoperability between vendors.

Benefits and Drawbacks for Layer 3 Transport

A design using Layer 3 transport requires these considerations:

- Layer 2 connectivity is limited within the data center rack up to the leaf switch.
- The leaf switch terminates each VLAN and provides default gateway functionality. That is, it has a switch virtual interface (SVI) for each VLAN.

- Uplinks from the leaf switch to the spine layer are routed using point-to-point links. VLAN trunking on the uplinks is not allowed.
- A dynamic routing protocol, such as OSPF, ISIS, or BGP, connects the leaf switches and spine switches. Each leaf switch in the rack advertises a small set of prefixes, typically one per VLAN or subnet. In turn, the leaf switch calculates equal cost paths to the prefixes it received from other leaf switches.

Layer 3 routing has the following benefits and drawbacks:

- The benefit of using Layer 3 transport is that you can choose from a wide array of Layer 3 capable switch products for the physical switching fabric. You can mix switches from different vendors due to general interoperability between implementation of OSPF, ISIS or BGP. This approach is typically more cost effective because it makes use of only the basic functionality of the physical switches.
- A design restriction, and thereby a drawback of Layer 3 routing, is that VLANs are restricted to a single rack. This affects vSphere vMotion, vSphere Fault Tolerance, and storage networks.

Infrastructure Network Architecture in ROBO

A key goal of network virtualization is to provide a virtual-to-physical network abstraction.

To achieve this, the physical fabric must provide a robust IP transport with the following characteristics:

- Simplicity
- Scalability
- High bandwidth
- Fault-tolerant transport
- Support for different levels of quality of service (QoS)

Quality of Service Differentiation in ROBO

Virtualized environments carry different types of traffic, including tenant, storage and management traffic, across the switching infrastructure. Each traffic type has different characteristics and makes different demands on the physical switching infrastructure.

- Management traffic, although typically low in volume, is critical for controlling physical and virtual network state.
- IP storage traffic is typically high in volume and generally stays within a data center.

For virtualized environments, the hypervisor sets the QoS values for the different traffic types. The physical switching infrastructure has to trust the values set by the hypervisor. No reclassification is necessary at the server-facing port of a switch. If there is a congestion point in the physical switching infrastructure, the QoS values determine how the physical network sequences, prioritizes, or potentially drops traffic.

Two types of QoS configuration are supported in the physical switching infrastructure.

- Layer 2 QoS (also referred to as class of service).

- Layer 3 QoS (also referred to as DSCP marking).

vSphere Distributed Switches support both class of service and DSCP marking. You can mark (or label) traffic based on the traffic type or packet classification. When the virtual machines are connected to the VXLAN-based logical switches or networks, the QoS values from the internal packet headers are copied to the VXLAN-encapsulated header. This enables the external physical network to prioritize the traffic based on the tags on the external header.

Server Network Interface Controllers in ROBO

If the server has more than one network interface controller (NIC) of the same speed, use two as uplinks with VLANs trunked to the interfaces.

The vSphere Distributed Switch supports many different NIC Teaming options. Load-based NIC teaming supports optimal use of available bandwidth and supports redundancy in case of a link failure. Use two 10 GbE connections for each server in combination with a pair of top of rack switches. 802.1Q network trunks can support a small number of VLANs. For example, management, storage, VXLAN, and VMware vSphere vMotion traffic.

Regions, Hubs and Remote and Branch Offices in ROBO

Regions support disaster recovery solutions and allow you to place workloads closer to your customer's locations.

This VMware Validated Design uses multiple regions. The full SDDC stack runs in a region which allows for the failover of those components to another region. The remote office and branch office is a separate geographical location, like a region, where a subset of SDDC components run.

About Regions in ROBO

Multiple regions support placing workloads closer to your customers. For example, by operating one region on the US east coast and one region on the US west coast, or operating a region in Europe and another region in the US.

Regions are helpful in several ways.

- Regions let you support disaster recovery solutions. One region is the primary site, and another is the recovery site.
- You can use multiple regions to address data privacy laws and restrictions in certain countries by keeping tenant data within a region in the same country.

The distance between regions can be large. This design uses two example regions: one in San Francisco (SFO) and the other in Los Angeles (LAX). These regions are also referred to as Region A and Region B, respectively.

About Hubs and Remote Office and Branch Office Sites

Operating in multiple geographic locations places workloads closer to end users.

Remote Office and Branch Office (ROBO) sites may be required to support manufacturing, Point of Sale (POS), medical facilities such as hospitals and clinics, or other scenarios. Management of these locations is performed using a hub/spoke method where a single region is typically defined by a geographic location.

- Hubs can act as a central point of management for all connected ROBO sites.
- Remote Office and Branch Office sites can be scaled appropriately as business needs change, such as through mergers and acquisitions.
- Disaster recovery is possible for a hub between regions, ensuring a continuous availability of ROBO site data collection and provisioning.
- Regional SLA dashboards can be created and monitored within a hub to identify investment opportunities in the appropriate ROBO site.

This design uses a hub that can failover between San Francisco (SFO) and Los Angeles (LAX).

Hubs

A hub refers to the centralized provisioning and monitoring components of the SDDC. A hub can be dedicated to ROBO sites (depending on the number of ROBO connections required), or be part of a wider SDDC platform. In either case it can fail over between regions in the event of a disaster, or for disaster avoidance.

Remote and Branch Offices

A remote or branch office is defined as a small location with limited datacenter capabilities providing specific services to the site. Minimal IT investment is typically made to support up to 100 virtual workloads in the SDDC. ROBO sites have lower SLA and connectivity requirements while operating independently from a hub in the event of connectivity failure.

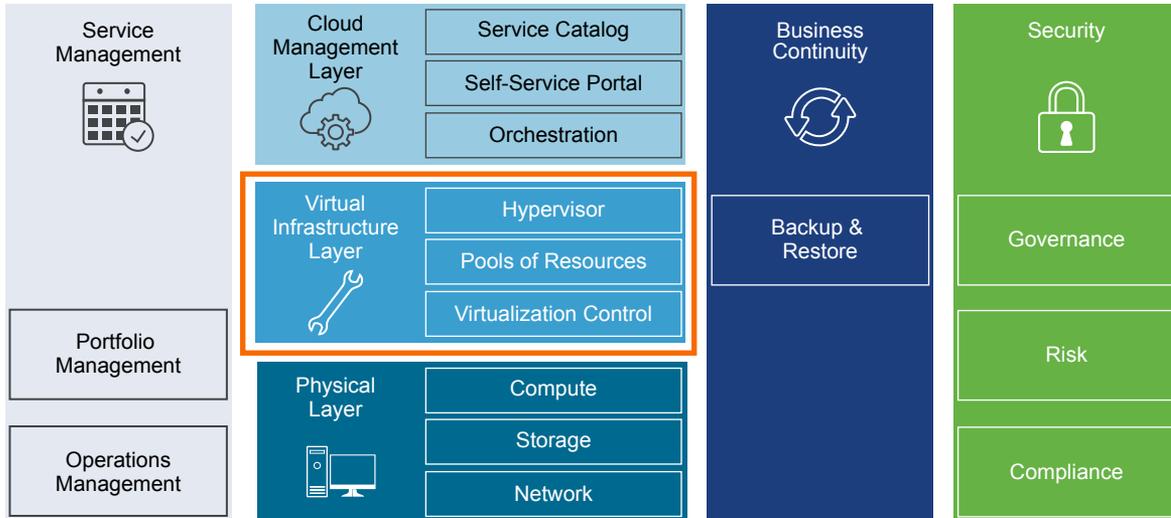
This design uses a single ROBO site located in New York City (NYC).

Virtual Infrastructure Architecture in ROBO

The virtual infrastructure is the foundation of an operational SDDC.

Within the virtual infrastructure layer, access to the underlying physical infrastructure is controlled and allocated to management and tenant workloads. The virtual infrastructure layer consists primarily of the physical hosts' hypervisors, and the control of these hypervisors. The management workloads consist of elements in the virtual management layer itself, along with elements in the cloud management layer and in the service management, business continuity, and security areas.

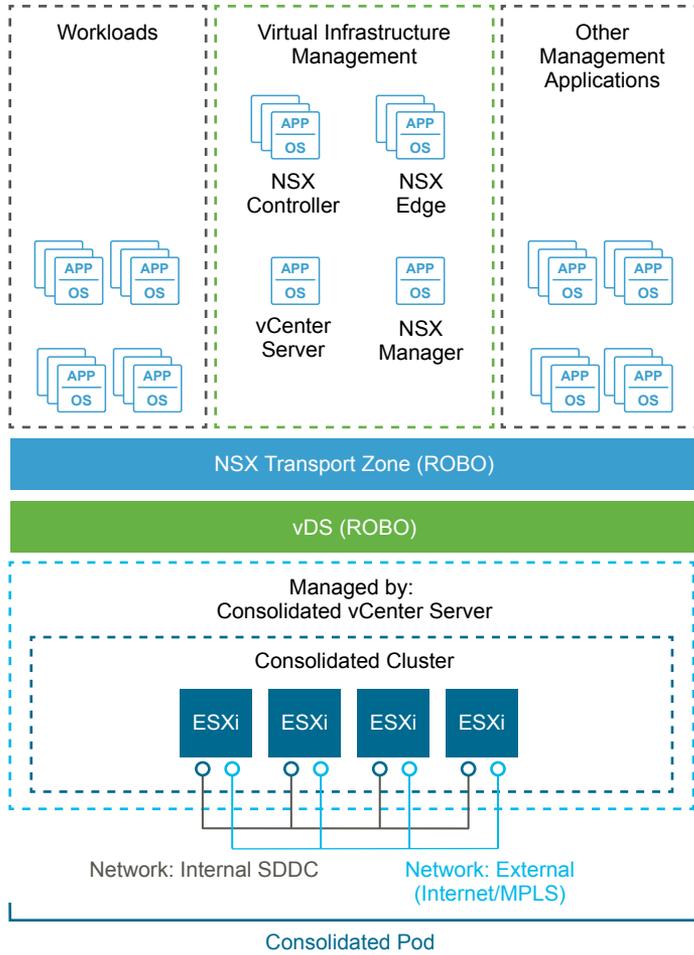
Figure 2-3. Virtual Infrastructure Layer in the Remote Office and Branch Office



Virtual Infrastructure Overview in ROBO

The ROBO SDDC consists of a consolidated pod in a ROBO location, and a hub that is responsible for provisioning workloads and monitoring the SDDC.

Figure 2-4. Remote Office and Branch Office SDDC Logical Design



Management Pod

Management pods run the virtual machines that manage the SDDC in the hub. This includes the vRealize Automation and vRealize Operations components used for provisioning and monitoring all regions. For more information on the management pod see the VMware Validated Design for SDDC.

Consolidated Pod

The consolidated pod runs the required SDDC components to manage ROBO workloads and send data back to the hub for analysis, as well as local backup and recovery. It hosts vCenter Server, NSX components, vRealize Operations Remote Collectors, vRealize Log Insight, vRealize Automation agents and other shared management components.

NSX services enable north-south routing between the SDDC and the external network, and east-west routing inside the SDDC.

The consolidated pod also hosts the SDDC tenant virtual machines (sometimes referred to as workloads or payloads). Workloads run customer business applications supporting varying SLAs. These workloads typically have low latency and security requirements for local deployment.

Note Should your ROBO site need to scale beyond a single pod (100 workloads), consider using the appropriate pod design to properly scale the solution based on the VMware Validated Design for SDDC.

Network Virtualization Components in ROBO

VMware NSX for vSphere, the network virtualization platform, is a key solution in the SDDC architecture. The NSX for vSphere platform consists of several components that are relevant to the network virtualization design.

NSX for vSphere Platform

NSX for vSphere creates a network virtualization layer. All virtual networks are created on top of this layer, which is an abstraction between the physical and virtual networks. Several components are required to create this network virtualization layer:

- vCenter Server
- NSX Manager
- NSX Controller
- NSX Virtual Switch

These components are separated into different planes to create communications boundaries and provide isolation of workload data from system control messages.

Data plane Workload data is contained wholly within the data plane. NSX logical switches segregate unrelated workload data. The data is carried over designated transport networks in the physical network. The NSX Virtual Switch, distributed routing, and the distributed firewall are also implemented in the data plane.

Control plane Network virtualization control messages are located in the control plane. Use secure physical networks (VLANs) that are isolated from the transport networks (which are used by the data plane) to carry control plane communication. Control messages set up networking attributes on NSX Virtual Switch instances, as well as configure and manage disaster recovery and distributed firewall components on each ESXi host.

Management plane Network virtualization orchestration takes place in the management plane. In this layer, cloud management platforms such as VMware vRealize[®] Automation[™] can request, consume, and destroy networking resources for virtual workloads. Communication is directed from the cloud management platform to vCenter Server to create and manage virtual machines, and to NSX Manager to consume networking resources.

Network Virtualization Services in ROBO

Network virtualization services include logical switches, logical routers, logical firewalls, and other components of VMware NSX[®] for vSphere[®].

Logical Switches

NSX for vSphere logical switches create logically abstracted segments to which tenant virtual machines can connect. A single logical switch is mapped to a unique VXLAN segment ID and is distributed across the ESXi hypervisors within a transport zone. This allows line-rate switching in the hypervisor without creating constraints of VLAN sprawl or spanning tree issues.

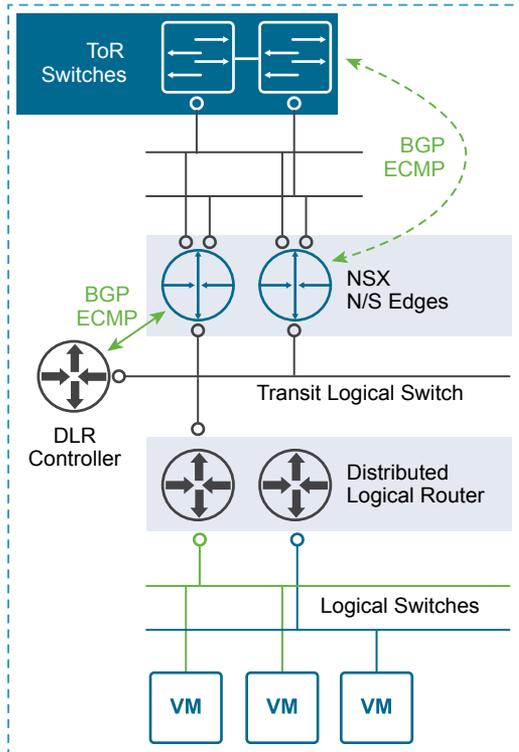
Distributed Logical Router

The NSX for vSphere Distributed Logical Router is optimized for forwarding in the virtualized space (between VMs, on VXLAN- or VLAN-backed port groups). Features include:

- High performance, low overhead first hop routing.
- Scaling the number of hosts.
- Support for up to 1,000 logical interfaces (LIFs) on each distributed logical router.

The Distributed Logical Router is installed in the kernel of every ESXi host, as such it requires a VM to provide the control plane. The distributed logical router Control VM is the control plane component of the routing process, providing communication between NSX Manager and NSX Controller cluster through the User World Agent. NSX Manager sends logical interface information to the Control VM and NSX Controller cluster, and the Control VM sends routing updates to the NSX Controller cluster.

Figure 2-5. NSX for vSphere Distributed Logical Router



Designated Instance

The designated instance is responsible for resolving ARP on a VLAN LIF. There is one designated instance per VLAN LIF. The selection of an ESXi host as a designated instance is performed automatically by the NSX Controller cluster and that information is pushed to all other hosts. Any ARP requests sent by the distributed logical router on the same subnet are handled by the same host. In case of host failure, the controller selects a new host as the designated instance and makes that information available to other hosts.

User World Agent

User World Agent (UWA) is a TCP and SSL client that enables communication between the ESXi hosts and NSX Controller nodes, and the retrieval of information from NSX Manager through interaction with the message bus agent.

Edge Services Gateway

While the Universal Logical Router provides VM to VM or east-west routing, the NSX Edge services gateway provides north-south connectivity, by peering with upstream Top of Rack switches, thereby enabling tenants to access public networks.

Logical Firewall

NSX for vSphere Logical Firewall provides security mechanisms for dynamic virtual data centers.

- The Distributed Firewall allows you to segment virtual data center entities like virtual machines. Segmentation can be based on VM names and attributes, user identity, vCenter objects like data centers, and hosts, or can be based on traditional networking attributes like IP addresses, port groups, and so on.
- The Edge Firewall component helps you meet key perimeter security requirements, such as building DMZs based on IP/VLAN constructs, tenant-to-tenant isolation in multi-tenant virtual data centers, Network Address Translation (NAT), partner (extranet) VPNs, and user-based SSL VPNs.

The Flow Monitoring feature displays network activity between virtual machines at the application protocol level. You can use this information to audit network traffic, define and refine firewall policies, and identify threats to your network.

Logical Virtual Private Networks

NSX Edge supports several types of virtual private networks (VPNs).

- SSL VPN-Plus allows remote users to access private corporate applications.
- IPSec VPN offers site-to-site connectivity between an NSX Edge instance and remote sites.
- L2 VPN allows you to extend your datacenter by allowing virtual machines to retain network connectivity across geographical boundaries.

Logical Load Balancer

The NSX Edge load balancer enables network traffic to follow multiple paths to a specific destination. It distributes incoming service requests evenly among multiple servers in such a way that the load distribution is transparent to users. Load balancing thus helps in achieving optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload. NSX Edge provides load balancing up to Layer 7.

Service Composer

Service Composer helps you provision and assign network and security services to applications in a virtual infrastructure. You map these services to a security group, and the services are applied to the virtual machines in the security group.

Data Security provides visibility into sensitive data that are stored within your organization's virtualized and cloud environments. Based on the violations that are reported by the NSX for vSphere Data Security component, NSX security or enterprise administrators can ensure that sensitive data is adequately protected and assess compliance with regulations around the world.

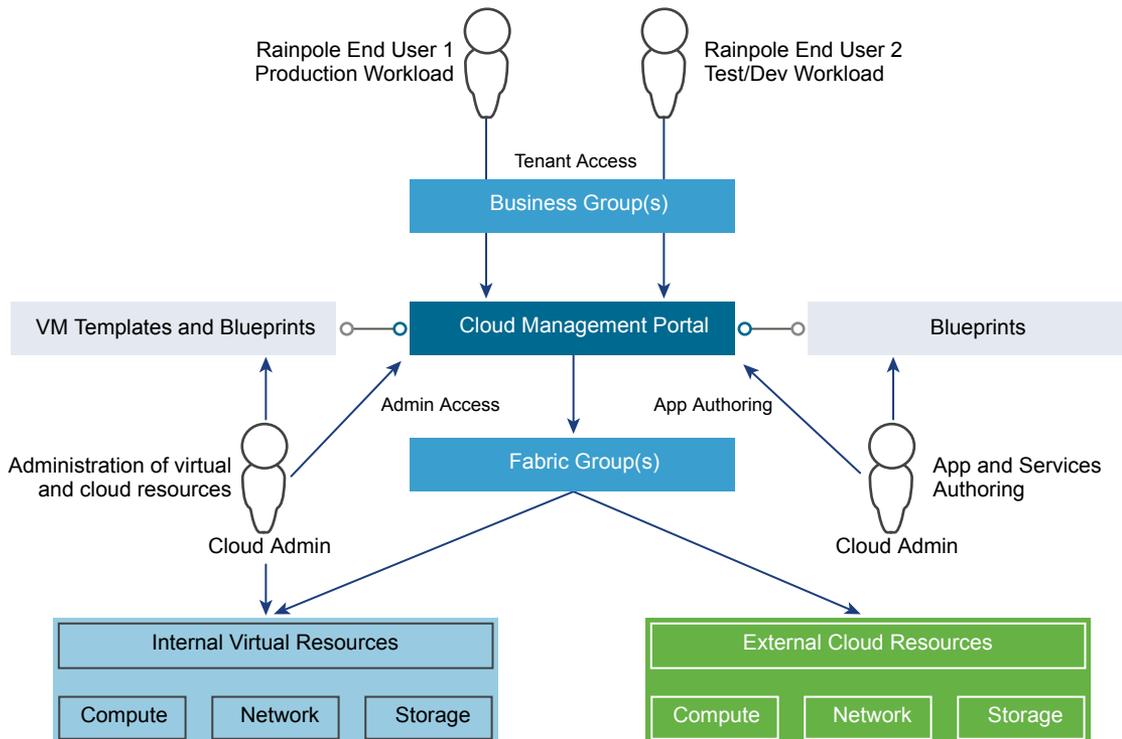
NSX for vSphere Extensibility

VMware partners integrate their solutions with the NSX for vSphere platform to enable an integrated experience across the entire SDDC. Data center operators can provision complex, multi-tier virtual networks in seconds, independent of the underlying network topology or components.

Cloud Management Platform Architecture in ROBO

The Cloud Management Platform (CMP) is the primary consumption portal for the entire Software-Defined Data Center (SDDC). Within the SDDC, users use vRealize Automation to author, administer, and consume VM templates and blueprints.

Figure 2-6. Cloud Management Platform Conceptual Architecture



The Cloud Management Platform consists of the following elements and components.

Table 2-1. Elements and Components of the Cloud Management Platform

Element	Components
Users	<ul style="list-style-type: none"> ■ Cloud administrators. Tenant, group, fabric, infrastructure, service, and other administrators as defined by business policies and organizational structure. ■ Cloud (or tenant) users. Users within an organization who can provision virtual machines, and directly perform operations on them at the level of the operating system.
Tools and supporting infrastructure	VM templates and blueprints are the building blocks that provide the foundation of the cloud. VM templates are used to author the blueprints that tenants (end users) use to provision their cloud workloads.
Provisioning infrastructure	On-premise and off-premise resources which together form a hybrid cloud. <ul style="list-style-type: none"> ■ Internal Virtual Resources. Supported hypervisors and associated management tools. ■ External Cloud Resources. Supported cloud providers and associated APIs.
Cloud management portal	<p>A portal that provides self-service capabilities for users to administer, provision and manage workloads.</p> <ul style="list-style-type: none"> ■ vRealize Automation portal, Admin access. The default root tenant portal URL used to set up and administer tenants and global configuration options. ■ vRealize Automation portal, Tenant access. Refers to a subtenant that is accessed using an appended tenant identifier. <p>Note A tenant portal might refer to the default tenant portal in some configurations. In this case, the URLs match, and the user interface is contextually controlled by the role-based access control permissions that are assigned to the tenant.</p>

Logical Architecture of the Cloud Management Platform

The design of the Cloud Management Platform considers characteristics such as availability, manageability, performance, scalability, and security. To provide this it must deliver a comprehensive set of multi-platform and multi-vendor cloud services.

The Cloud Management Platform layer delivers the following multi-platform and multi-vendor cloud services.

- Comprehensive and purpose-built capabilities to provide standardized resources to global customers in a short time span.
- Multi-platform and multi-vendor delivery methods that integrate with existing enterprise management systems.
- Central user-centric and business-aware governance for all physical, virtual, private, and public cloud services.
- Architecture that meets customer and business needs, and is extensible.

This design considers the following characteristics.

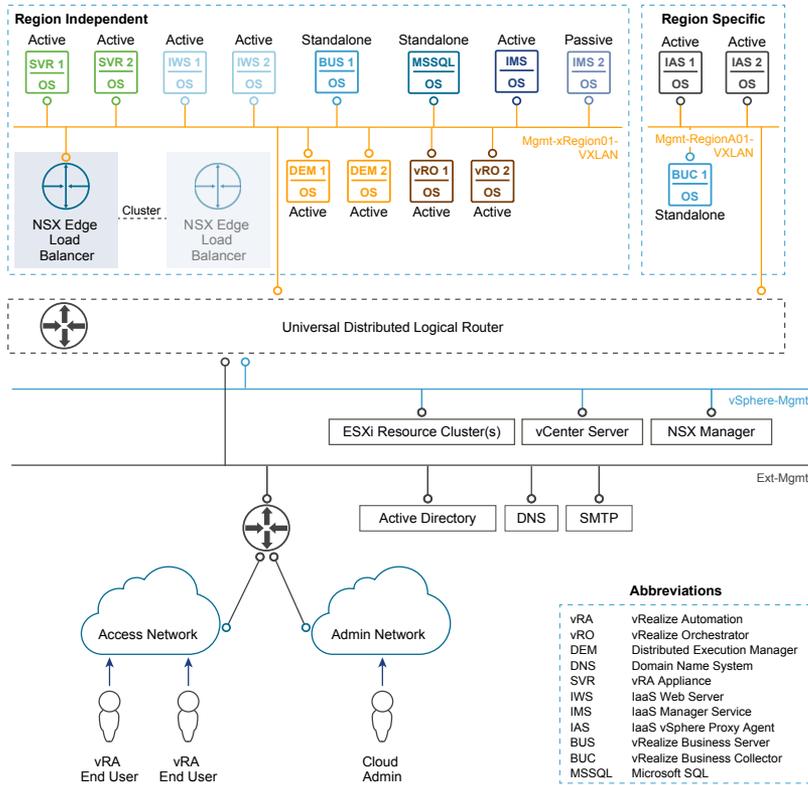
Availability	Indicates the effect a choice has on technology and related infrastructure to provide highly-available operations and sustain operations during system failures. VMware vSphere High Availability will provide the required host redundancy and tolerance of hardware failures where appropriate.
Manageability	Relates to the effect a choice has on overall infrastructure manageability. Key metrics: Accessibility and the lifecycle of the infrastructure being managed.
Performance	Reflects whether the option has a positive or negative impact on overall infrastructure performance. This architecture follows the VMware reference architecture sizing guidelines to provide certain performance characteristics. Key metrics: Performance analysis and tuning of the database, Manager service, Model Manager, portal Web site, and data collection.
Scalability	Determines the ability of the solution to be augmented to achieve better sustained performance within the infrastructure. Key metrics: Web site latency, network traffic, and CPU usage on the database and web servers.
Security	Reflects whether the option has a positive or negative impact on overall infrastructure security. Key metrics: Data confidentiality, integrity, authenticity, and non-repudiation of cloud automation components and the option's integration with supporting and provisioning infrastructures.

Cloud Management Platform Logical Architecture

In this architecture, vRealize Automation and vRealize Orchestrator run in the hub of a VXLAN-backed network that is fronted by an NSX Logical Distributed Router. An NSX Edge services gateway, acting as a load balancer, is deployed to provide load balancing services for the CMP components.

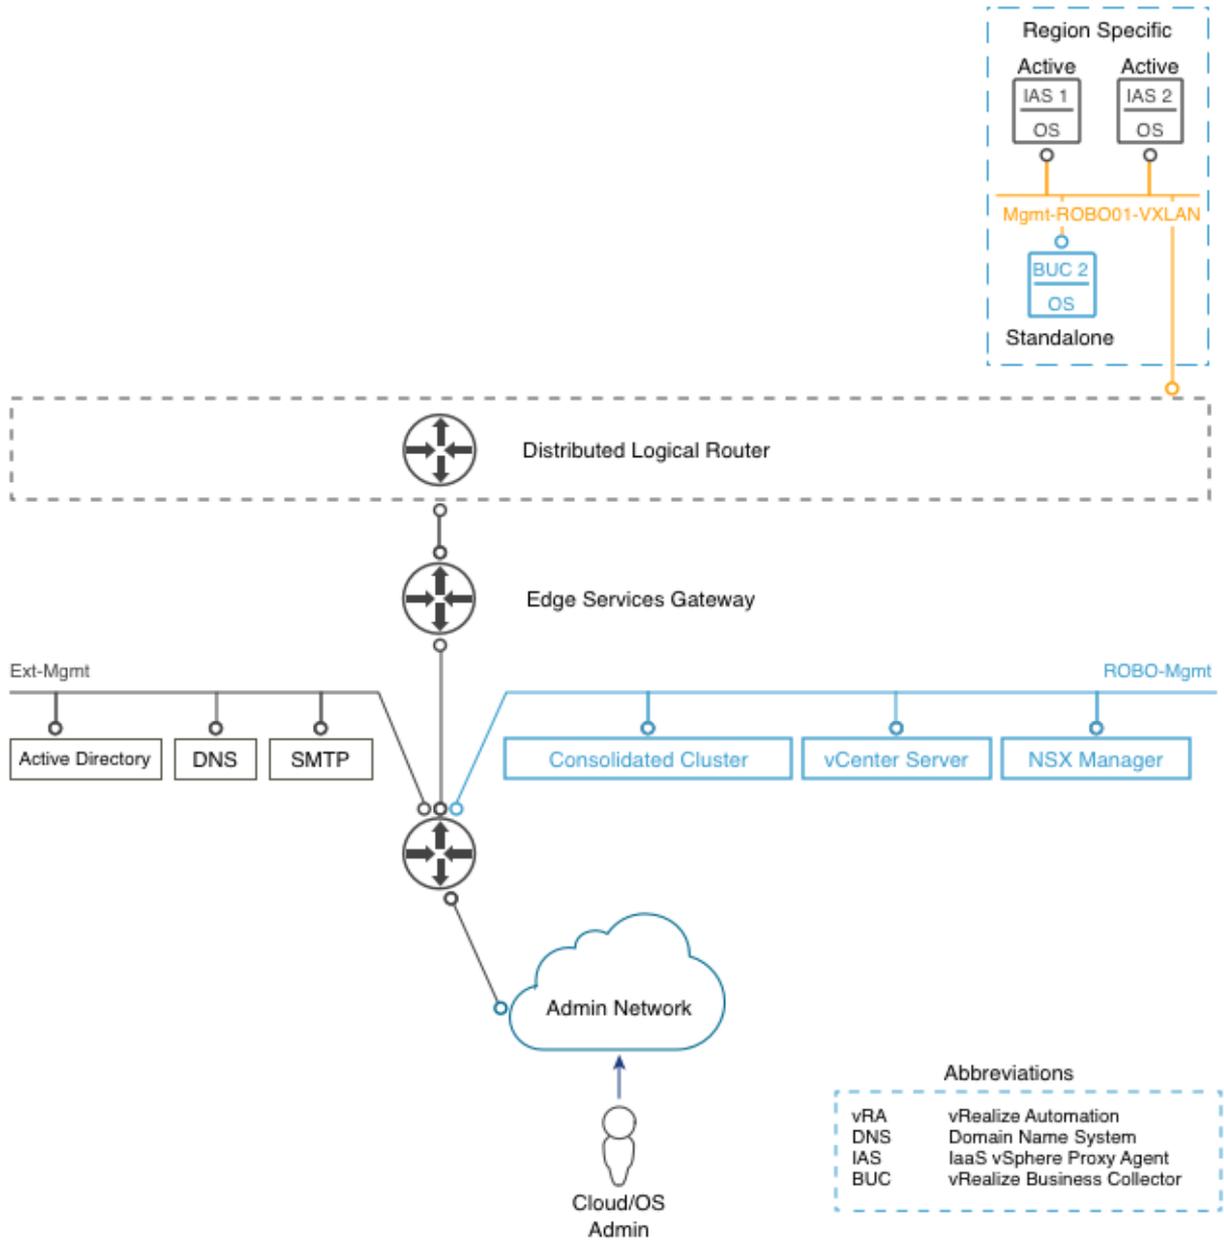
These components run in the hub, providing the primary platform for automating and managing the workloads running in each of the connected remote sites.

Figure 2-7. vRealize Automation Logical Architecture for the Hub



Each ROBO site employs the vRealize Automation Proxy Agent and vRealize Business for Cloud Standard Data Collector as brokers back to the hub CMP. This enables centralized workload management and cost analysis with minimal overhead in each ROBO site.

Figure 2-8. vRealize Automation Logical Architecture for ROBO



Operations Architecture in ROBO

The architecture of the operations management layer includes management components that provide support for the main types of operations in an SDDC. You can perform monitoring, logging, and backup and restore operations.

Operations Management Architecture in ROBO

vRealize Operations Manager tracks and analyzes the operation of multiple data sources within the Software-Defined Data Center (SDDC) using specialized analytics algorithms. These algorithms help vRealize Operations Manager to learn and predict the behavior of every object it monitors. Users access this information using views, reports, and dashboards.

Installation Models

vRealize Operations Manager is available in two different deployment models: a preconfigured virtual appliance, or a Windows or Linux installable package. Select the installation method for your deployment according to the following considerations:

- When you use the vRealize Operations Manager virtual appliance, you deploy the OVF file of the virtual appliance once for each cluster node. You access the product to set up cluster nodes according to their role, and log in to configure the installation.

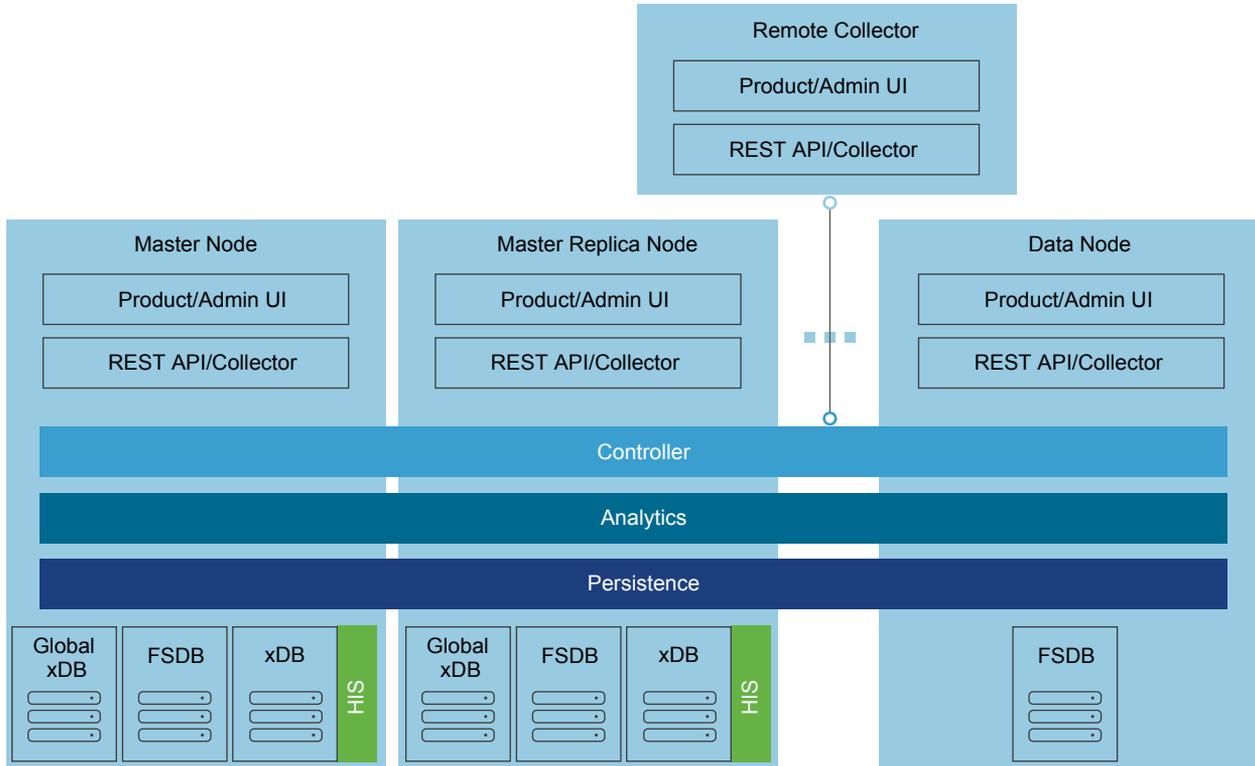
Use the virtual appliance deployment to easily create vRealize Operations Manager nodes with pre-defined, CPU, memory, and disk sizing that is identical among nodes.
- When you use the Windows or Linux installable package, you run the vRealize Operations Manager installation on each cluster node. You access the product to set up cluster nodes according to their role, and log in to configure the installation.

Use the installable package deployment to create vRealize Operations Manager nodes with custom CPU, memory, and disk sizing that must be identical among nodes.

Architecture

vRealize Operations Manager contains functional elements that collaborate for data analysis and storage, and support creating clusters of nodes with different roles.

Figure 2-9. vRealize Operations Manager Architecture



Types of Nodes and Clusters

For high availability and scalability, you can deploy several vRealize Operations Manager instances in a cluster where they can have either of the following roles.

Master Node

Required initial node in the cluster. In large-scale environments, the master node manages all other nodes. In small-scale environments, the master node is the single, standalone vRealize Operations Manager node.

Master Replica Node

(Optional) Enables high availability of the master node.

Data Node

Enables scale-out of vRealize Operations Manager in larger environments. Data nodes have adapters installed to perform collection and analysis. Data nodes also host vRealize Operations Manager management packs.

Larger deployments usually include adapters only on data nodes, not on the master node or replica node

Remote Collector Node

In distributed deployments, enables navigation through firewalls, interfaces with a remote data source, reduces bandwidth across regions, or reduces the load on the vRealize Operations Manager analytics cluster. Remote collector nodes only gather objects for the inventory, and forward collected data to data nodes. Remote collector nodes do not store data or perform analysis. In addition, you can install them on a different operating system than the rest of the cluster nodes.

The master and master replica nodes are data nodes with extended capabilities.

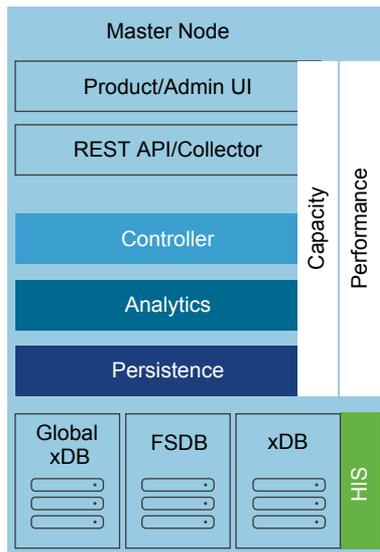
vRealize Operations Manager can form two types of clusters according to the nodes that participate in a cluster.

- Analytics clusters** Tracks, analyzes, and predicts the operation of monitored systems. Consists of a master node, data nodes, and optionally of a master replica node.
- Remote collectors cluster** Only collects diagnostics data without storage or analysis. Consists only of remote collector nodes.

Application Functional Components

The functional components of a vRealize Operations Manager instance interact to provide analysis of diagnostics data from the data center and visualize the result in the Web user interface.

Figure 2-10. vRealize Operations Manager Logical Node Architecture



The components of the vRealize Operations Manager node performs these tasks.

- Admin / Product UI server** The Admin/Product UI server is a Web application that serves as both the user and administrative interface.
- REST API / Collector** The collector collects data from all components in the data center.
- Controller** The controller handles the data flow of the UI server, collector, and analytics engine.
- Analytics** The analytics engine creates associations and correlations between various data sets, handles super metric calculations, performs capacity planning functions, and is responsible for triggering alerts.
- Persistence** The persistence layer handles the read and write operations on the underlying databases across all nodes.

FSDB	The File System Database (FSDB) stores collected metrics in raw format. FSDB is available in all the nodes.
xDB (HIS)	The xDB stores data from the Historical Inventory Service (HIS). This component is available only on the master and master replica nodes.
Global xDB	The Global xDB stores user preferences, alerts, and alarms, and customization that is related to the vRealize Operations Manager. This component is available only on the master and master replica nodes.

Management Packs

Management packs contain extensions and third-party integration software. They add dashboards, alert definitions, policies, reports, additional metrics and other content to the inventory of vRealize Operations Manager. To learn more about management packs and to download them, see *VMware Solutions Exchange*.

Multi-Region vRealize Operations Manager Deployment

The scope of the SDDC design covers multiple regions. Using vRealize Operations Manager across multiple regions requires deploying an analytics cluster in a regional hub that is protected by Site Recovery Manager, and deploying remote collectors in each region.

Logging Architecture in ROBO

vRealize Log Insight provides real-time log management and log analysis with machine learning-based intelligent grouping, high-performance searching, and troubleshooting across physical, virtual, and cloud environments.

Overview

vRealize Log Insight collects data from ESXi hosts using the syslog protocol. It connects to other VMware products, like vCenter Server, to collect events, tasks, and alarms data, and integrates with vRealize Operations Manager to send notification events and enable launch in context. vRealize Log Insight also functions as a collection and analysis point for any system capable of sending syslog data. In addition to syslog data an ingestion agent can be installed on Linux or Windows servers or may come pre-installed on certain VMware products to collect logs. This agent approach is especially useful for custom application logs and operating systems that don't natively support the syslog protocol, such as Windows.

Installation Models

You can deploy vRealize Log Insight as a virtual appliance in one of the following configurations:

- Standalone node
- Highly available cluster of one master and at least two worker nodes using an integrated load balancer (ILB)

The compute and storage resources of the vRealize Log Insight instances can scale-up as growth demands.

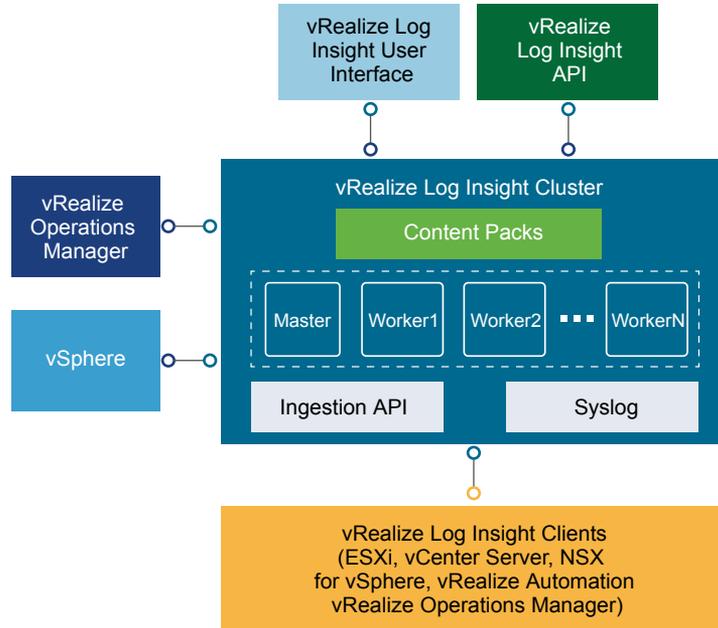
Cluster Nodes

For high availability and scalability, you can deploy several vRealize Log Insight instances in a cluster where they can have either of the following roles:

Master Node	Required initial node in the cluster. In standalone mode, the master node is initially responsible for all activities, including queries and log ingestion; however, after additional nodes have been deployed and configured in a vRealize Log Insight cluster for high availability (HA), these activities are delegated to all available nodes. After HA has been configured for a cluster, the master node still retains responsibility for the lifecycle of a cluster, which includes performing upgrades, as well as adding and removing of Worker nodes. The master node stores logs locally. If master node is down, the logs on the master becomes unavailable.
Worker Node	Enables scale-out in larger environments. A worker node is responsible for queries and log ingestion. A worker node stores logs locally. If a worker node is down, the logs on that worker becomes unavailable. You need at least two worker nodes to form a cluster with the master node.
Integrated Load Balancer (ILB)	Provides high availability. The ILB runs on one of the cluster nodes. If the node that hosts the ILB Virtual IP (VIP) address stops responding, the VIP address is failed over to another node in the cluster. All queries against data are directed to the ILB, in which the query request is delegated to a query master for the duration of the query, and in turn queries all nodes (both master and workers) for data, which then aggregates that data for handoff back to the client. The Web User Interface of the ILB serves as a single pane of glass, presenting data from multiple sources in the cluster in a unified display; while individual nodes can be accessed via their Web User Interfaces, unless you are performing specific administrative activities on the individual nodes, it is advised to use the ILB's Interface.

Architecture of a Cluster

The architecture of vRealize Log Insight enables several channels for HA collection of log messages.

Figure 2-11. Cluster Architecture of vRealize Log Insight

vRealize Log Insight clients connect to ILB VIP address, and use the Web user interface and ingestion (by using syslog or the Ingestion API) to send logs to vRealize Log Insight.

By default, the vRealize Log Insight collects data from vCenter Server systems and ESXi hosts. For forwarding logs from NSX for vSphere, and vRealize Automation, use content packs which contain extensions or provide integration with other systems in the SDDC.

Authentication Models

You can configure vRealize Log Insight user authentication to utilize one or more of the following authentication models:

- Microsoft Active Directory
- Local Accounts
- VMware Identity Manager

Integration with vRealize Operations Manager

The integration with vRealize Operations Manager provides data from multiple sources to a central place for monitoring the SDDC. vRealize Log Insight sends notification events to vRealize Operations Manager. Once integrated, vRealize Operations Manager is able to provide the inventory map of any vSphere objects to vRealize Log Insight, allowing for you to perform launching of log messages from vRealize Log Insight to the vRealize Operations Manager Web user interface, taking you either directly to the object itself or the location of the object within the environment.

Integration with vSphere

The integration with vSphere provides data from multiple sources to a central place for monitoring the SDDC. vRealize Log Insight connects to vCenter Server in two-minute intervals, and collect events, alarms, and tasks data from these vCenter Server systems. Further, this integration can be used to configure the managed ESXi hosts within the vCenter Server to send their logs to the vRealize Log Insight instance.

Archiving

vRealize Log Insight supports data archiving on NFS shared storage that each vRealize Log Insight node can access.

Backup

You back up each vRealize Log Insight cluster using traditional virtual machine backup solutions that use vSphere Storage APIs for Data Protection (VADP) compatible backup software such as vSphere Data Protection.

Multi-Region vRealize Log Insight Deployment

Using vRealize Log Insight in a multi-region design can provide a syslog infrastructure in all regions of the SDDC. Using vRealize Log Insight across multiple regions requires deploying a cluster in each region. vRealize Log Insight supports event forwarding to other vRealize Log Insight deployments across regions in the SDDC. Implementing failover by using vSphere Replication or disaster recovery by using Site Recovery Manager is not necessary. The event forwarding feature adds tags to log message that identify the source region and event filtering prevents looping messages between the regions.

Tagging

vRealize Log Insight provides a capability called tags. When logs are forwarded to a regional hub, each log can be tagged with metadata that allows you to track its origin. With the VMware Validated Design, these tags are used as location identifiers. Configuring these tags as part of the forwarder allows SLA dashboards to be created for high level monitoring and trend analysis at each site, identifying locations with unreliable connectivity or high site resource consumption.

Data Protection and Backup Architecture in ROBO

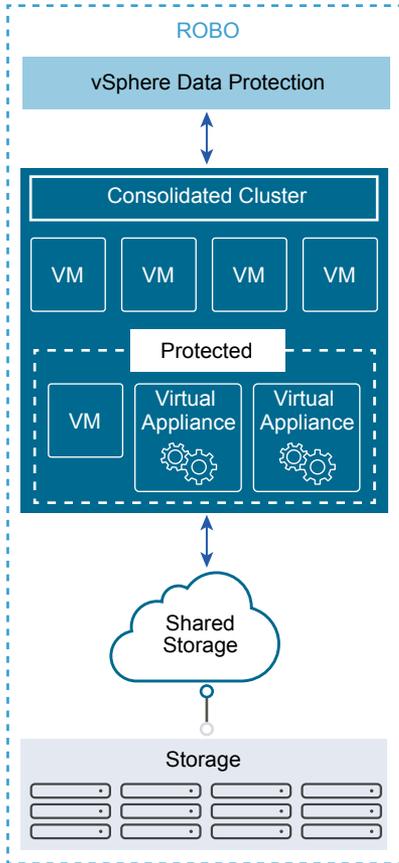
You can use a backup solution, such as vSphere Data Protection, to protect the data of your SDDC management components and of the tenant workloads.

Data protection solutions provide the following functions in the SDDC:

- Backup and restore virtual machines.
- Organization of virtual machines into groups by VMware product.
- Store data according to company retention policies.
- Inform administrators about backup and restore activities through reports.

vSphere Data Protection instances in the ROBO sites provide data protection for the products that implement the management capabilities of the SDDC. vSphere Data Protection stores backups of the management product virtual appliances on a shared storage allocation according to a defined schedule.

Figure 2-12. Data Protection Architecture for Remote office and Branch Office



Detailed Design in ROBO

The ROBO SDDC detailed design considers both physical and virtual infrastructure design. It includes numbered design decisions and the justification and implications of each decision.

The diagrams below illustrate the SDDC hub, a remote office or branch office site, and the connections between them.

Figure 3-1. SDDC Hub with Connections to Remote Office and Branch Office

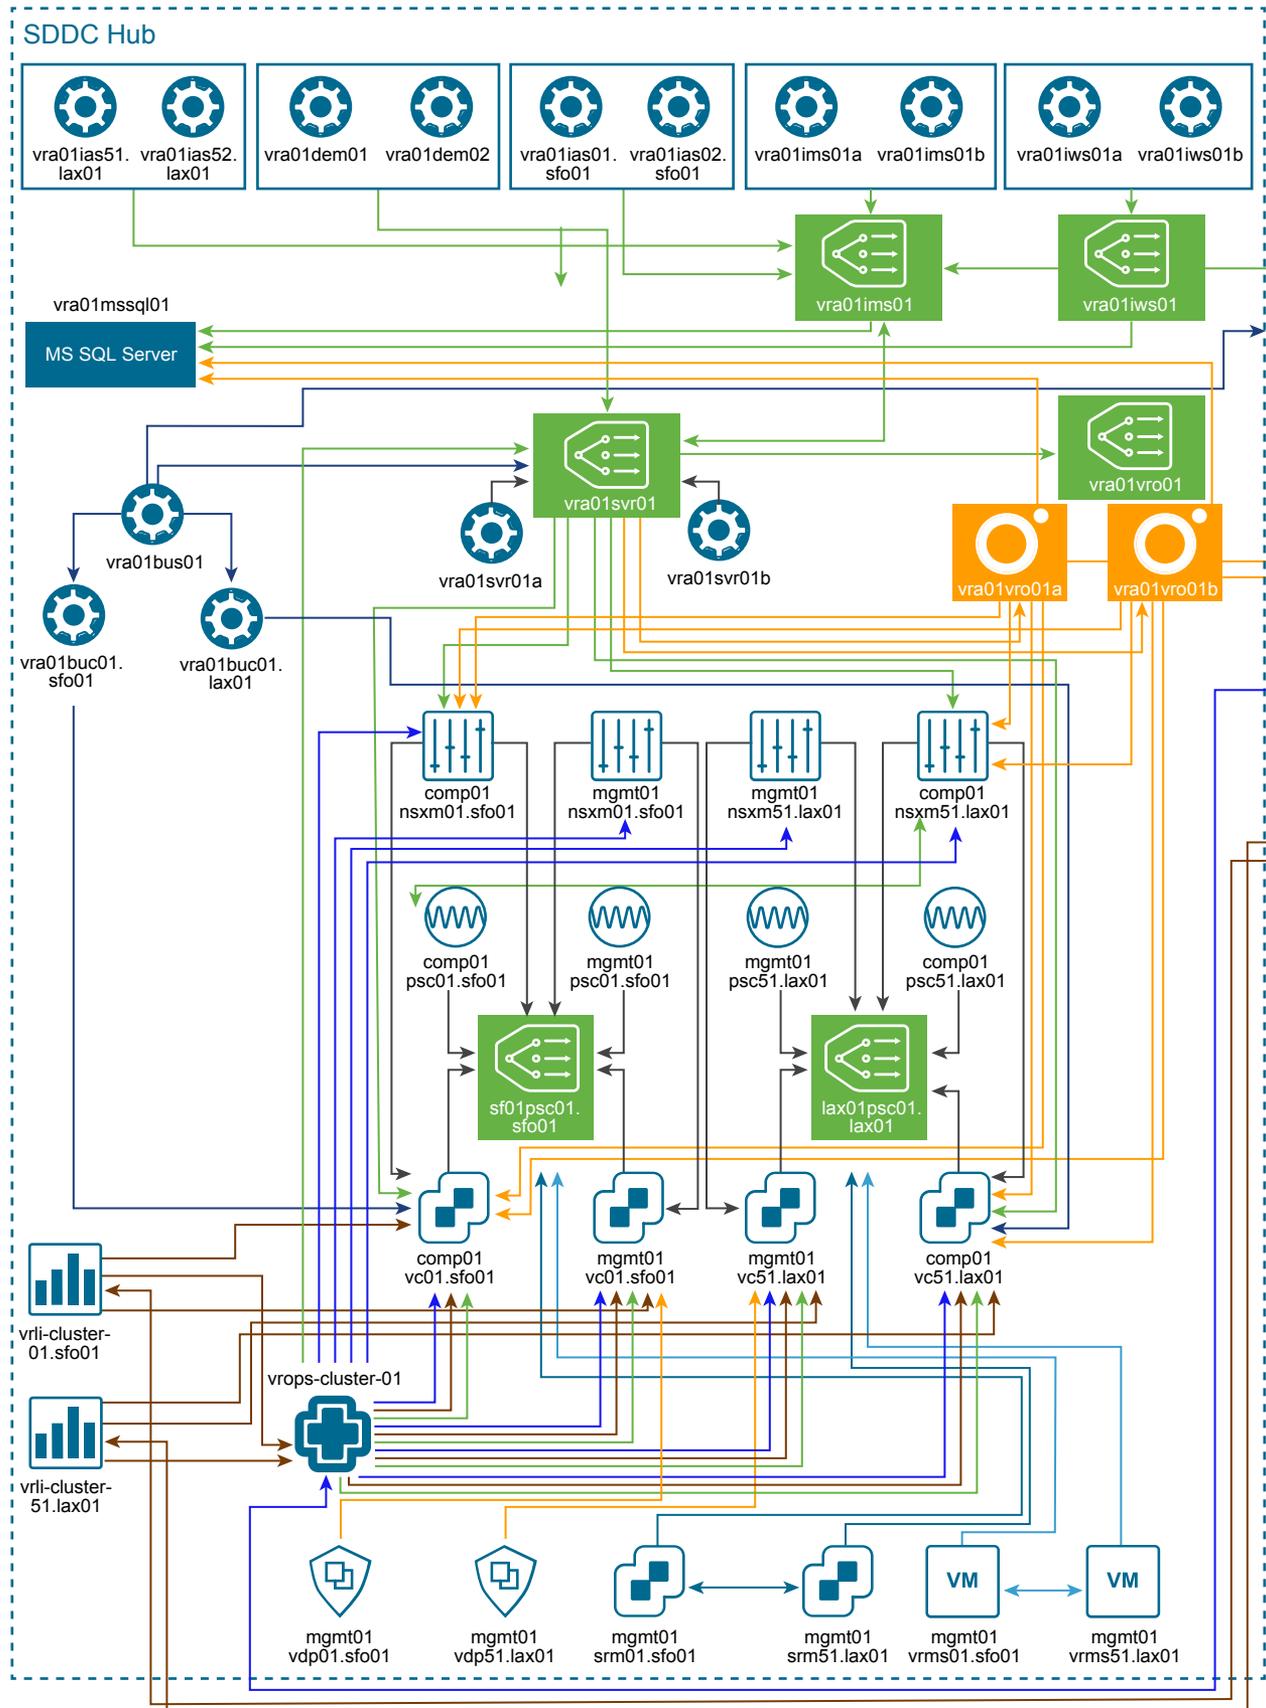
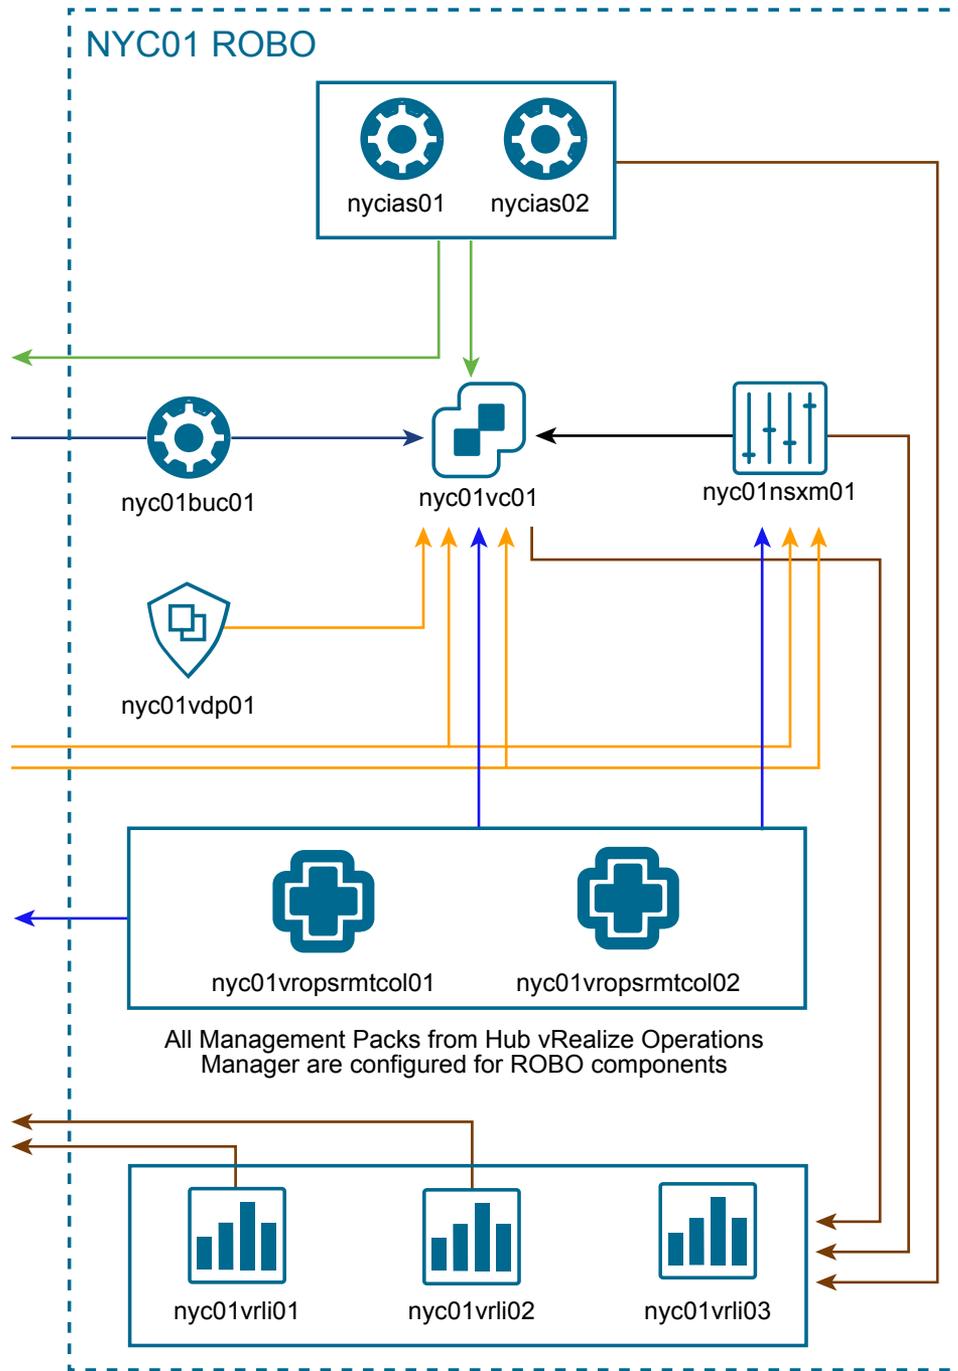


Figure 3-2. Remote Office and Branch Office with Connections to SDDC Hub



Each section also includes detailed discussion and diagrams.

Physical Infrastructure Design	Focuses on the three main pillars of any data center, compute, storage and network. In this section you will find information about regions, hubs, and Remote Office and Branch Office sites. The section also provides details on the rack and pod configuration, and on physical hosts and the associated storage and network configurations.
Virtual Infrastructure Design	Provides details on the core virtualization software configuration. This section has information on the ESXi hypervisor, vCenter Server, the virtual network design including VMware NSX, and on software-defined storage for VMware vSAN. This section also includes details on business continuity (backup and restore).
Cloud Management Platform Design	Contains information on the consumption and orchestration layer of the SDDC stack, which uses vRealize Automation and vRealize Orchestrator. Organizations can use the fully distributed and scalable architecture to streamline their provisioning and decommissioning operations.
Operations Infrastructure Design	Explains how to architect, install, and configure vRealize Operations Manager and vRealize Log Insight. You learn how to ensure that service management within the SDDC is comprehensive. This section ties directly into the <i>Operational Guidance</i> section.

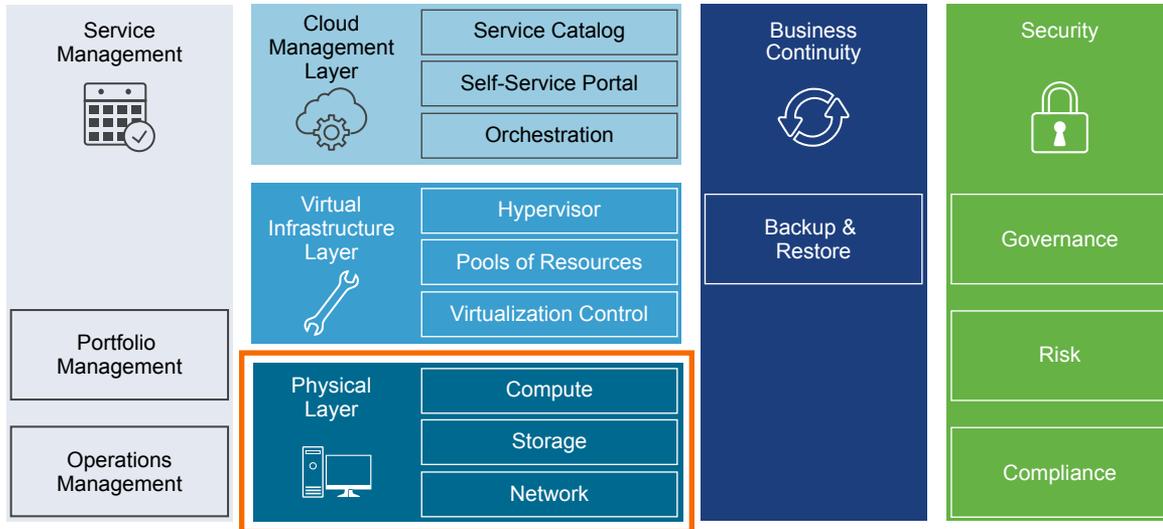
This chapter includes the following topics:

- [Physical Infrastructure Design in ROBO](#)
- [Virtual Infrastructure Design in ROBO](#)
- [Cloud Management Platform Design in ROBO](#)
- [Operations Infrastructure Design in ROBO](#)

Physical Infrastructure Design in ROBO

The physical infrastructure design includes details on decisions for regions and the pod layout within datacenter racks.

Design decisions related to server, networking, and storage hardware are part of the physical infrastructure design.

Figure 3-3. Physical Infrastructure Design

- **Physical Design Fundamentals in ROBO**

Physical design fundamentals include decisions on availability zones and regions, and on pod types, pods, and racks. Included in the physical design fundamentals is the ESXi host physical design.

- **Physical Networking Design in ROBO**

The physical network design for ROBO has specific requirements which are outlined in the following sections.

- **Physical Storage Design in ROBO**

This VMware Validated Design relies on both vSAN and secondary storage.

Physical Design Fundamentals in ROBO

Physical design fundamentals include decisions on availability zones and regions, and on pod types, pods, and racks. Included in the physical design fundamentals is the ESXi host physical design.

Regions, Hubs, and Remote and Branch Offices

Regions, hubs, and ROBO sites are the key elements of the ROBO SDDC.

Regions

Regions provide disaster recovery across different SDDC instances. This design uses two regions. Each region is a separate SDDC instance. The regions have a similar physical layer design and virtual infrastructure design but different naming.

The design uses the following regions. The region identifier uses United Nations Code for Trade and Transport Locations(UN/LOCODE) along with a numeric instance ID.

Table 3-1. Regions

Region	Region Identifier	Region-specific Domain Name	Region Description
A	SFO01	sfo01.rainpole.local	San Francisco, CA, USA based data center
B	LAX01	lax01.rainpole.local	Los Angeles, CA, USA based data center

Note Region Identifiers will vary based on the locations used in your deployment.

Hubs A hub refers to the centralized provisioning and monitoring components of the SDDC. A hub can be dedicated to ROBO sites (depending on the number of ROBO connections required) or part of the VMware Validated Design for SDDC and, in either case, has the capability for fail over between regions in the event of a disaster or for disaster avoidance.

Remote Office and Branch Office A location used to support specific services such as manufacturing, hospitals, or call centers. These locations require minimal workload deployment and have hardware located in space constrained rooms.

The design uses the following identifier. The ROBO identifier uses United Nations Code for Trade and Transport Locations(UN/LOCODE) along with a numeric instance ID.

Table 3-2. Remote Office and Branch Office

ROBO Identifier	ROBO Description
NYC01	New York City, NY, USA based Remote Office and Branch Office

Note ROBO Identifiers will vary based on the locations used in your deployment.

Table 3-3. Regions Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-001	Deploy the hub using two regions.	Supports multi-region failover capabilities while sustaining centralized provisioning and monitoring of ROBO sites.	Having multiple regions will require an increased solution footprint and associated costs.

Pods and Racks in ROBO

The local remote office and branch office functionality is contained in a single consolidated pod designed for use with ROBO.

Figure 3-4. Pod Architecture for Remote Office and Branch Office

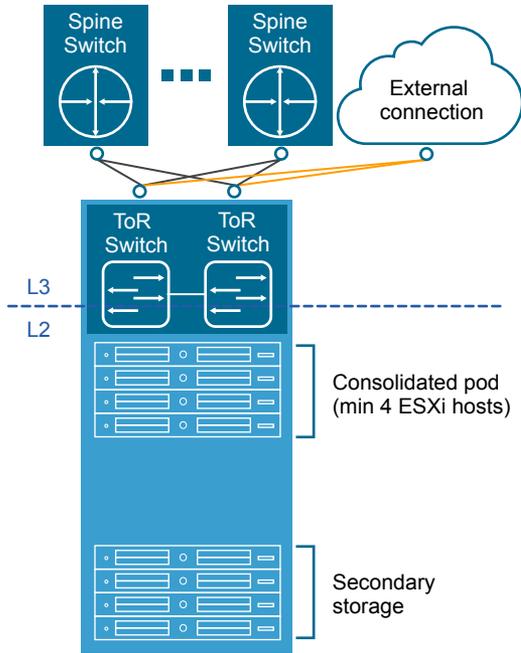


Table 3-4. Required Number of Racks

Pod (Function)	Required Number of Racks (for full scale deployment)	Minimum Number of Racks	Comment
Consolidated Pod	1	1	One rack is sufficient for a consolidated pod. With 1 rack, the consolidated pod requires a minimum of 4 ESXi hosts in order to contain the management components and a minimum set of user workloads. Additional hosts can be added to scale to achieve the supported maximum number of virtual machines in a ROBO configuration. The quantity and performance varies based on the workloads running within the consolidated pod and the actual hardware configuration of the pod hosts.
Storage pods	1	0	Storage that is not vSAN storage is hosted on isolated storage pods.
Total	2 (if secondary storage is not located in the same rack)	1	

Table 3-5. POD and Racks Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-002	The consolidated pod is bound to a physical rack.	The Top of Rack switches are the layer 2 boundary. Any virtual machines on a VLAN backed network, such as vCenter Server and NSX controllers, will lose network access if migrated to a host in a different rack.	The design must include sufficient power and cooling to operate the server equipment. This depends on the selected vendor and products.
ROBO-PHY-003	Each rack has two separate power feeds.	Redundant power feeds increase availability by ensuring that failure of a power feed does not bring down all equipment in a rack. Combined with redundant network connections into a rack and within a rack, redundant power feeds prevent failure of equipment in an entire rack.	All equipment used must support two separate power feeds. The equipment must keep running if one power feed fails.

ESXi Host Physical Design Specifications in ROBO

The physical design specifications of the ESXi host list the characteristics of the hosts that were used during deployment and testing of this VMware Validated Design.

Physical Design Specification Fundamentals

The configuration and assembly process for each system is standardized, with all components installed the same manner on each host. Standardizing the entire physical configuration of the ESXi hosts is critical to providing an easily manageable and supportable infrastructure because standardization eliminates variability. Consistent PCI card slot location, especially for network controllers, is essential for accurate alignment of physical to virtual I/O resources. Deploy ESXi hosts with identical configuration, including identical storage, and networking configurations, across all cluster members. Identical configurations ensure an even balance of virtual machine storage components across storage and compute resources.

Select all ESXi host hardware, including CPUs following the *VMware Compatibility Guide*.

The sizing of the vSAN physical servers for the ESXi hosts are based on the [VMware Virtual SAN Ready Nodes](#) document.

- An average sized VM has two vCPUs with 4 GB of RAM.
- A standard 2U server can host 60 average-sized VMs on a single ESXi host.

Table 3-6. ESXi Host Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-004	Use vSAN Ready Nodes.	Using a vSAN Ready Node ensures seamless compatibility with vSAN during the deployment.	Might limit hardware choices.
ROBO-PHY-005	All nodes must have uniform configurations across a given cluster.	A balanced cluster delivers more predictable performance even during hardware failures. In addition, performance impact during resync/rebuild is minimal when the cluster is balanced.	Vendor sourcing, budgeting and procurement considerations for uniform server nodes will be applied on a per cluster basis.

ESXi Host Memory

The amount of memory required for compute pods will vary depending on the workloads running in the pod. When sizing memory for compute pod hosts it's important to remember the admission control setting (n+1) which reserves one hosts resources for fail over or maintenance.

Note See the *VMware vSAN 6.5 Design and Sizing Guide* for more information about disk groups, including design and sizing guidance. The number of disk groups and disks that an ESXi host manages determines memory requirements. 32 GB of RAM is required to support the maximum number of disk groups.

Table 3-7. Host Memory Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-006	Ensure each ESXi host in the Consolidated pod has a minimum 256 GB RAM.	The management and edge virtual machines require a total of 117 GB RAM from the cluster. The remaining RAM is to support workload virtual machines.	Might limit hardware choices.

Host Boot Device Background Considerations

Minimum boot disk size for ESXi in SCSI-based devices (SAS / SATA / SAN) is greater than 5 GB. ESXi can be deployed using stateful, local SAN SCSI boot devices, or by using vSphere Auto Deploy. The supported boot device depends on the version of vSAN that you use. Consider the following issues when selecting a boot device:

- vSAN does not support stateless vSphere Auto Deploy
- vSAN 5.5 and greater supports USB/SD embedded devices for ESXi boot device (4 GB or greater).
- Beginning with vSAN 6.0, there is an option to use SATADOM as a supported boot device.

To select the boot device option that best fits your hardware, see the *VMware vSAN 6.5 Design and Sizing Guide*.

Physical Networking Design in ROBO

The physical network design for ROBO has specific requirements which are outlined in the following sections.

Switch Types and Network Connectivity in ROBO

Setup of the physical environment requires careful consideration. Follow best practices for physical switches, VLANs and subnets, and access port settings.

Top of Rack Physical Switches

When configuring Top of Rack (ToR) switches, consider the following best practices.

- Configure redundant physical switches to enhance availability.

- Configure switch ports that connect to ESXi hosts manually as trunk ports. Virtual switches are passive devices and do not send or receive trunking protocols, such as Dynamic Trunking Protocol (DTP).
- Modify the Spanning Tree Protocol (STP) on any port that is connected to an ESXi NIC to reduce the time it takes to transition ports over to the forwarding state, for example using the Trunk PortFast feature found in a Cisco physical switch.
- Provide DHCP or DHCP Helper capabilities on all VLANs that are used by Management and VXLAN VMkernel ports. This setup simplifies the configuration by using DHCP to assign IP address based on the IP subnet in use.
- Configure jumbo frames on all switch ports, inter-switch link (ISL) and switched virtual interfaces (SVI's).

Top of Rack Connectivity and Network Settings

Each ESXi host is connected redundantly to the SDDC network fabric ToR switches by means of two 10 GbE ports. Configure the ToR switches to provide all necessary VLANs via an 802.1Q trunk.

VLANs and Subnets

Each ESXi host uses VLANs and corresponding subnets for internal-only traffic, as shown in Sample VLANs and Subnets within a Pod.

Follow these guidelines.

- Use only /24 subnets to reduce confusion and mistakes when dealing with IPv4 subnetting.
- Use the IP address .253 as the (floating) interface with .251 and .252 for Virtual Router Redundancy Protocol (VRRP) or Hot Standby Routing Protocol (HSRP).
- Use the RFC1918 IPv4 address space for these subnets and allocate one octet by region and another octet by function. For example, the mapping *172.regionid.function.0/24* results in the following sample subnets.

Note The following VLANs and IP ranges are meant as samples. Your actual implementation depends on your environment.

Table 3-8. Sample VLANs and Subnets within a Pod

Pod	Function	Sample VLAN	Sample IP range
Consolidated	Management	1811 (Native)	172.18.11.0/24
Consolidated	vMotion	1812	172.18.12.0/24
Consolidated	vSAN	1813	172.18.13.0/24
Consolidated	VXLAN	1814	172.18.14.0/24
Consolidated	Storage	1815	172.18.15.0/24
Consolidated	Uplink 1	1816	172.18.16.0/24
Consolidated	Uplink 2	1817	172.18.17.0/24

Access Port Network Settings

Configure additional network settings on the access ports that connect the switch to the corresponding servers.

Spanning-Tree Protocol (STP)	Designate the access ports as trunk PortFast.
Trunking	Configure the VLANs as members of a 802.1Q trunk with the management VLAN acting as the native VLAN.
MTU	Set MTU for all VLANs and SVIs (Management, vMotion, VXLAN and Storage) to jumbo frames for consistency purposes.
DHCP helper	Configure the VIF of the Management and VXLAN subnet as a DHCP proxy.
Multicast	Configure IGMP snooping on the switches and include an IGMP querier on each VLAN.

Physical Network Design Decisions in ROBO

The physical network design decisions govern the physical layout and use of VLANs. They also include decisions on jumbo frames and on some other network-related requirements such as DNS and NTP.

Physical Network Design Decisions

Routing protocols	Base the selection of the external routing protocol on your current implementation or on available expertise among the IT staff. Take performance requirements into consideration. Possible options are OSPF, BGP and IS-IS.
DHCP proxy	The DHCP proxy must point to a DHCP server by way of its IPv4 address. See the <i>Planning and Preparation</i> documentation for details on the DHCP server.

Table 3-9. Physical Network Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-NET-001	The network infrastructure must support the following requirements: <ul style="list-style-type: none"> ■ BGP support ■ IGMP support ■ 1, 10GbE Port per ESXi host on the ToR switches ■ Minimum 20 GbE throughput from each ToR switch to the upstream physical device 	Two uplinks per ESXi host guarantees availability during a switch failure. Based on predicted traffic 20GbE uplinks guarantee no over subscription. BGP is used as the dynamic routing protocol in this design.	Could limit hardware choice.
ROBO-PHY-NET-002	The Consolidated rack uses two ToR switches. These switches provide connectivity across two 10 GbE links to each server.	This design uses two 10 GbE links to provide redundancy and reduce overall design complexity.	Requires two ToR switches per rack which can increase costs.
ROBO-PHY-NET-003	Use VLANs to segment physical network functions.	Allow for Physical network connectivity without requiring large number of NICs. Segregation is needed for the different network functions that are required in the SDDC. This segregation allows for differentiated services and prioritization of traffic as needed.	Uniform configuration and presentation is required on all the trunks made available to the ESXi hosts.

Additional Design Decisions

Additional design decisions deal with static IP addresses, DNS records, and the required NTP time source.

Table 3-10. Additional Network Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-NET-004	Assign Static IP addresses to all management nodes of the SDDC infrastructure.	Configuration of static IP addresses avoid connection outages due to DHCP availability or misconfiguration.	Accurate IP address management must be in place.
ROBO-PHY-NET-005	Create DNS records for all management nodes to enable forward, reverse, short and FQDN resolution.	Ensures consistent resolution of management nodes using both IP address (reverse lookup) and name resolution.	None
ROBO-PHY-NET-006	Use an NTP time source for all management nodes.	Critical to maintain accurate and synchronized time between management nodes.	None

Jumbo Frames Design Decisions

IP storage throughput can benefit from the configuration of jumbo frames. Increasing the per-frame payload from 1500 bytes to the jumbo frame setting increases the efficiency of data transfer. Jumbo frames must be configured end-to-end, which is easily accomplished in a LAN. When you enable jumbo frames on an ESXi host, you have to select an MTU that matches the MTU of the physical switch ports.

The workload determines whether it makes sense to configure jumbo frames on a virtual machine. If the workload consistently transfers large amounts of network data, configure jumbo frames if possible. In that case, the virtual machine operating systems and the virtual machine NICs must also support jumbo frames.

Using jumbo frames also improves performance of vSphere vMotion.

Note VXLANs need an MTU value of at least 1600 bytes on the switches and routers that carry the transport zone traffic.

Table 3-11. Jumbo Frames Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-NET-007	Configure the MTU size to 9000 bytes (Jumbo Frames) on the portgroups that support the following traffic types. <ul style="list-style-type: none"> ■ vSAN ■ vMotion ■ VXLAN ■ Secondary Storage 	Setting the MTU to 9000 bytes (Jumbo Frames) improves traffic throughput. In order to support VXLAN the MTU setting must be increased to a minimum of 1600 bytes, setting this portgroup also to 9000 bytes has no effect on VXLAN but ensures consistency across portgroups that are adjusted from the default MTU size.	When adjusting the MTU packet size, the entire network path (VMkernel port, distributed switch, physical switches and routers) must also be configured to support the same MTU packet size.

Physical Storage Design in ROBO

This VMware Validated Design relies on both vSAN and secondary storage.

The focus of this section is on physical storage design. For information and on where the SDDC uses vSAN and secondary storage, see [Shared Storage Design in ROBO](#).

vSAN Physical Design in ROBO

Software-defined storage is a key technology in the SDDC. This design uses vSAN to implement software-defined storage.

vSAN is a fully integrated hypervisor-converged storage software. vSAN creates a cluster of server hard disk drives and solid state drives, and presents a flash-optimized, highly resilient, shared storage datastore to hosts and virtual machines. vSAN allows you to control capacity, performance, and availability on a per virtual machine basis through the use of storage policies.

Requirements and Dependencies

The software-defined storage module has the following requirements and options.

- Minimum of 4 ESXi hosts providing storage resources to the ROBO vSAN cluster is required for guarantee vSAN redundancy during maintenance operations.
- vSAN is configured as hybrid storage or all-flash storage.
 - A vSAN hybrid storage configuration requires both magnetic devices and flash caching devices.
 - An All-Flash vSAN configuration requires vSphere 6.0 or later.

- Each ESXi host that provides storage resources to the cluster must meet the following requirements.
 - Minimum of one SSD. The SSD flash cache tier should be at least 10% of the size of the HDD capacity tier.
 - Minimum of two HDDs.
 - RAID controller compatible with vSAN.
 - 10 Gbps network for vSAN traffic with Multicast enabled.
 - vSphere High Availability Isolation Response set to power off virtual machines. With this setting, no possibility of split brain conditions in case of isolation or network partition exists. In a split-brain condition, the virtual machine might be powered on by two hosts by mistake. For more information, see design decision [ROBO-VI-VC-007](#).

Table 3-12. vSAN Physical Storage Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-STO-001	Use one 300 GB SSD and three traditional 1 TB HDDs to create a single disk group.	Allow enough capacity for the management and start point for workload VMs with a minimum of 10% flash-based caching.	When using only a single disk group you limit the amount of striping (performance) capability and increase the size of the fault domain. Disk space must be scaled as necessary to accommodate full 100 workload VMs. Because not all ROBO will be the same, 3TB of disk space provides a good starting point. Disk requirements will likely be higher depending on the workload disk size.

Hybrid Mode and All-Flash Mode

vSphere offers two different vSAN modes of operation, all-flash or hybrid.

Hybrid Mode In a hybrid storage architecture, vSAN pools server-attached capacity devices (in this case magnetic devices) and caching devices, typically SSDs or PCI-e devices to create a distributed shared datastore.

All-Flash Mode vSAN can be deployed as all-flash storage. All-flash storage uses flash-based devices (SSD or PCI-e) only as a write cache while other flash-based devices provide high endurance for capacity and data persistence.

Table 3-13. vSAN Mode Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-STO-002	Configure vSAN in hybrid mode.	Ensures a lower entry point for vSAN. Typically, virtual machines in ROBO do not require the performance of all-flash vSAN. Alternatively all-flash configuration could be used.	vSAN hybrid mode does not provide the potential performance or additional capabilities such as deduplication of an all-flash configuration.

Hardware Considerations in ROBO

You can build your own vSAN cluster or choose from a list of vSAN Ready Nodes.

Build Your Own

Be sure to use hardware from the [VMware Compatibility Guide](#) for the following vSAN components:

- Solid state disks (SSDs)
- Magnetic hard drives (HDDs)
- I/O controllers, including vSAN certified driver/firmware combinations

Use VMware vSAN Ready Nodes

A vSAN Ready Node is a validated server configuration in a tested, certified hardware form factor for vSAN deployment, jointly recommended by the server OEM and VMware. See the [VMware Compatibility Guide](#). The vSAN Ready Node documentation provides examples of standardized configurations, including the numbers of VMs supported and estimated number of 4K IOPS delivered.

As per design decision [ROBO-PHY-005](#), the VMware Validated Design uses vSAN Ready Nodes.

Solid State Disk (SSD) Characteristics in ROBO

In a vSAN configuration, the SSDs are used for the vSAN caching layer for hybrid deployments and for the capacity layer for all flash.

- For a hybrid deployment, the use of the SSD is split between a non-volatile write cache (approximately 30%) and a read buffer (approximately 70%). As a result, the endurance and the number of I/O operations per second that the SSD can sustain are important performance factors.
- For an all-flash model, endurance and performance have the same criteria. However, many more write operations are held by the caching tier, elongating or extending the life of the SSD capacity-tier.

SSD Endurance

This VMware Validated Design uses class D endurance class SSDs for the caching tier.

SDDC Endurance Design Decision Background

For endurance of the SSDs used for vSAN, standard industry write metrics are the primary measurements used to gauge the reliability of the drive. No standard metric exists across all vendors, however, Drive Writes per Day (DWPD) or Petabytes Written (PBW) are the measurements normally used.

For vSphere 5.5, the endurance class was based on Drive Writes Per Day (DWPD). For VMware vSAN 6.0 and later, the endurance class has been updated to use Terabytes Written (TBW), based on the vendor's drive warranty. TBW can be used for VMware vSAN 5.5, VMware vSAN 6.0 and VMware vSAN 6.5 and is reflected in the *VMware Compatibility Guide*.

The reasoning behind using TBW is that VMware provides the flexibility to use larger capacity drives with lower DDPD specifications.

If a SSD vendor uses Drive Writes Per Day as a measurement, you can calculate endurance in Terabytes Written (TBW) with the following equation.

$$TBW \text{ (over 5 years)} = \text{Drive Size} \times DDPD \times 365 \times 5$$

For example, if a vendor specified DDPD = 10 for an 800 GB capacity SSD, you can compute TBW with the following equation.

$$TBW = 0.4TB \times 10DDPD \times 365days \times 5yrs$$

$$TBW = 7300TBW$$

That means the SSD supports 7300TB writes over 5 years (The higher the TBW number, the greater the endurance class.).

For SSDs that are designated for caching and all-flash capacity layers, the following table outlines which endurance class to use for hybrid and for all-flash vSAN.

Endurance Class	TBW	Hybrid Caching Tier	All-Flash Caching Tier	All-Flash Capacity Tier
Class A	>=365	No	No	Yes
Class B	>=1825	Yes	No	Yes
Class C	>=3650	Yes	Yes	Yes
Class D	>=7300	Yes	Yes	Yes

Note This VMware Validated Design does not use All-Flash vSAN.

Table 3-14. SSD Endurance Class Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-STO-003	Use Class D (>=7300TBW) SSDs for the caching tier of the Consolidated cluster.	If a SSD designated for the caching tier fails due to wear-out, the entire vSAN disk group becomes unavailable. The result is potential data loss or operational impact.	SSDs with higher endurance may be more expensive than lower endurance classes.

SSD Performance in ROBO

There is a direct correlation between the SSD performance class and the level of vSAN performance. The highest-performing hardware results in the best performance of the solution. Cost is therefore the determining factor. A lower class of hardware that is more cost effective might be attractive even if the performance or size is not ideal.

For optimal performance of vSAN, select class E or greater SSDs. See the [VMware Compatibility Guide](#) for detail on the different classes.

SSD Performance Design Decision Background

Select a high class of SSD for optimal performance of VMware vSAN. Before selecting a drive size, consider disk groups and sizing as well as expected future growth. VMware defines classes of performance in the [VMware Compatibility Guide](#) as follows.

Table 3-15. SSD Performance Classes

Performance Class	Writes Per Second
Class A	2,500 – 5,000
Class B	5,000 – 10,000
Class C	10,000 – 20,000
Class D	20,000 – 30,000
Class E	30,000 – 100,000
Class F	100,000 +

Select an SSD size that is, at a minimum, 10% of the anticipated size of the consumed HDD storage capacity, before failures to tolerate are considered. For example, select an SSD of at least 100 GB for 1 TB of HDD storage consumed in a 2 TB disk group.

Caching Algorithm

Both hybrid clusters and all-flash configurations adhere to the recommendation that 10% of consumed capacity for the flash cache layer. However, there are differences between the two configurations.

Hybrid vSAN 70% of the available cache is allocated for storing frequently read disk blocks, minimizing accesses to the slower magnetic disks. 30% of available cache is allocated to writes.

All-Flash vSAN All-flash clusters have two types of flash: very fast and durable write cache, and cost-effective capacity flash. Here cache is 100% allocated for writes, as read performance from capacity flash is more than sufficient.

Use Class E SSDs or greater for the highest possible level of performance from the vSAN volume.

Table 3-16. SSD Performance Class Selection

Design Quality	Option 1 Class E	Option 2 Class C	Comments
Availability	o	o	Neither design option impacts availability.
Manageability	o	o	Neither design option impacts manageability.
Performance	↑	↓	The higher the storage class that is used, the better the performance.
Recover-ability	o	o	Neither design option impacts recoverability.
Security	o	o	Neither design option impacts security.

Legend: ↑ = positive impact on quality; ↓ = negative impact on quality; o = no impact on quality.

Table 3-17. SSD Performance Class Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-STO-004	Use Class E SSDs (30,000-100,000 writes per second) for the Consolidated cluster.	The storage I/O performance requirements within the Consolidated cluster dictate the need for at least Class E SSDs.	Class E SSDs might be more expensive than lower class drives.

Magnetic Hard Disk Drives (HDD) Characteristics in ROBO

The HDDs in a vSAN environment have two different purposes, capacity and object stripe width.

Capacity Magnetic disks, or HDDs, unlike caching-tier SSDs, make up the capacity of a vSAN datastore

Stripe Width You can define stripe width at the virtual machine policy layer. vSAN might use additional stripes when making capacity and placement decisions outside a storage policy.

vSAN supports these disk types:

- Serial Attached SCSI (SAS)
- Near Line Serial Attached SCSI (NL-SCSI). NL-SAS can be thought of as enterprise SATA drives but with a SAS interface.
- Serial Advanced Technology Attachment (SATA). Use SATA magnetic disks only in capacity-centric environments where performance is not prioritized.

SAS and NL-SAS get you the best results. This VMware Validated Design uses 10,000 RPM drives to achieve a balance between cost and availability.

HDD Capacity, Cost, and Availability Background Considerations

You can achieve the best results with SAS and NL-SAS.

The VMware vSAN design must consider the number of magnetic disks required for the capacity layer, and how well the capacity layer will perform.

- SATA disks typically provide more capacity per individual drive, and tend to be less expensive than SAS drives. However, the trade-off is performance. Due to lower rotational speeds (typically 7200RPM), SATA performance is not as good as that of SAS .
- In environments where performance is critical, choose SAS magnetic disks instead of SATA magnetic disks.

Consider that failure of a larger capacity drive has operational impact on the availability and recovery of more components.

Rotational Speed (RPM) Background Considerations

HDDs tend to be more reliable, but that comes at a cost. SAS disks can be available up to 15,000 RPM speeds.

Table 3-18. vSAN HDD Environmental Characteristics

Characteristic	Revolutions per Minute (RPM)
Capacity	7,200
Performance	10,000
Additional Performance	15,000

Cache-friendly workloads are less sensitive to disk performance characteristics; however, workloads can change over time. HDDs with 10,000 RPM are the accepted norm when selecting a capacity tier.

For the software-defined storage module, VMware recommends that you use an HDD configuration that is suited to the characteristics of the environment. If there are no specific requirements, selecting 10,000 RPM drives achieves a balance between cost and availability.

Table 3-19. HDD Selection Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-STO-005	Use 10,000 RPM HDDs for the Consolidated cluster.	10,000 RPM HDDs achieve a balance between performance and availability for the VMware vSAN configuration. The performance of 10,000 RPM HDDs avoids disk drain issues. In vSAN hybrid mode, the vSAN periodically flushes uncommitted writes to the capacity tier.	Slower and potentially cheaper HDDs are not available.

I/O Controllers in ROBO

The I/O controllers are as important to a vSAN configuration as the selection of disk drives. vSAN supports SAS, SATA, and SCSI adapters in either pass-through or RAID 0 mode. vSAN supports multiple controllers per host.

- Multiple controllers can improve performance and mitigate a controller or SSD failure to a smaller number of drives or vSAN disk groups.
- With a single controller, all disks are controlled by one device. A controller failure impacts all storage, including the boot media (if configured).

Controller queue depth is possibly the most important aspect for performance. All I/O controllers in the *VMware vSAN Hardware Compatibility Guide* have a minimum queue depth of 256. Consider normal day-to-day operations and increase of I/O due to Virtual Machine deployment operations or re-sync I/O activity as a result of automatic or manual fault remediation.

About SAS Expanders

SAS expanders are a storage technology that lets you maximize the storage capability of your SAS controller card. Like the switches in an Ethernet network, SAS expanders allow you to connect a larger number of devices. That is, more SAS/SATA devices to a single SAS controller. Many SAS controllers support 128 or more hard drives.

Caution VMware has not extensively tested SAS expanders, as a result performance and operational predictability are relatively unknown at this point. For this reason, you should avoid configurations with SAS expanders.

Secondary Storage Design

Secondary storage is recommended for backup data to ensure backups do not reside on primary storage. The consolidated cluster uses vSAN for primary storage, and VMware recommends the use of secondary storage for backup.

Table 3-20. Secondary Storage Design Decisions

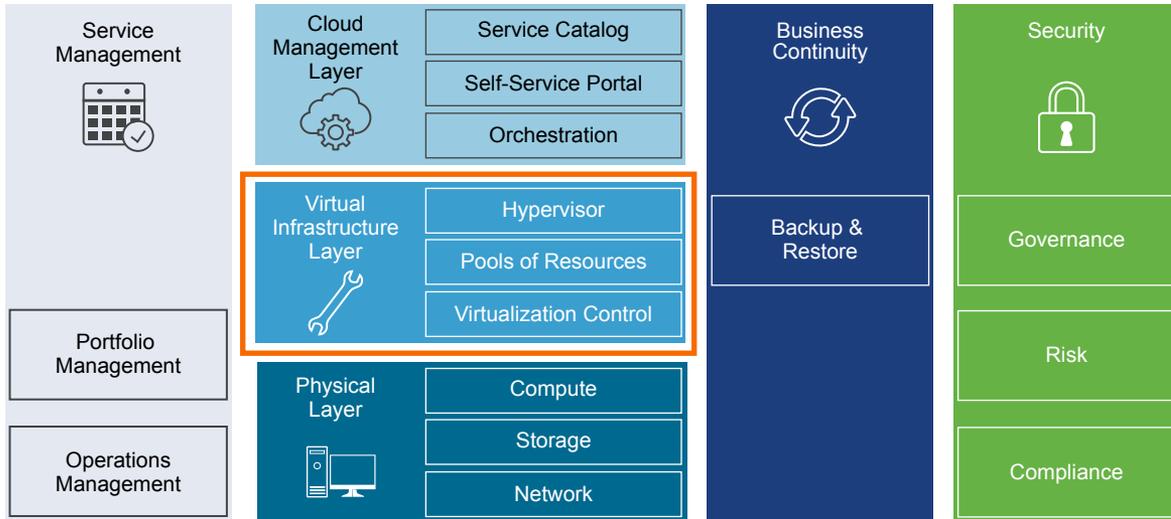
Decision ID	Design Decision	Design Justification	Design Implication
ROBO-PHY-STO-006	Use a secondary storage solution for management and workload backup data.	Separate primary virtual machine storage from backup data in case of primary storage failure.	Secondary storage is required.
ROBO-PHY-STO-007	The secondary storage used must provide adequate size and I/O for backup operations to finish during the scheduled backup window.	The backup and restore process is I/O intensive. The backup retention process is a storage constrained operation.	The secondary storage solution has an impact on the backup and restore SLA.

Virtual Infrastructure Design in ROBO

The virtual infrastructure design includes the software components that make up the virtual infrastructure layer and that support the business continuity of the SDDC.

These components include the software products that provide the virtualization platform hypervisor, virtualization management, storage virtualization, network virtualization, backup and disaster recovery. VMware products in this layer include VMware vSphere, VMware vSAN, VMware NSX, and VMware vSphere Data Protection.

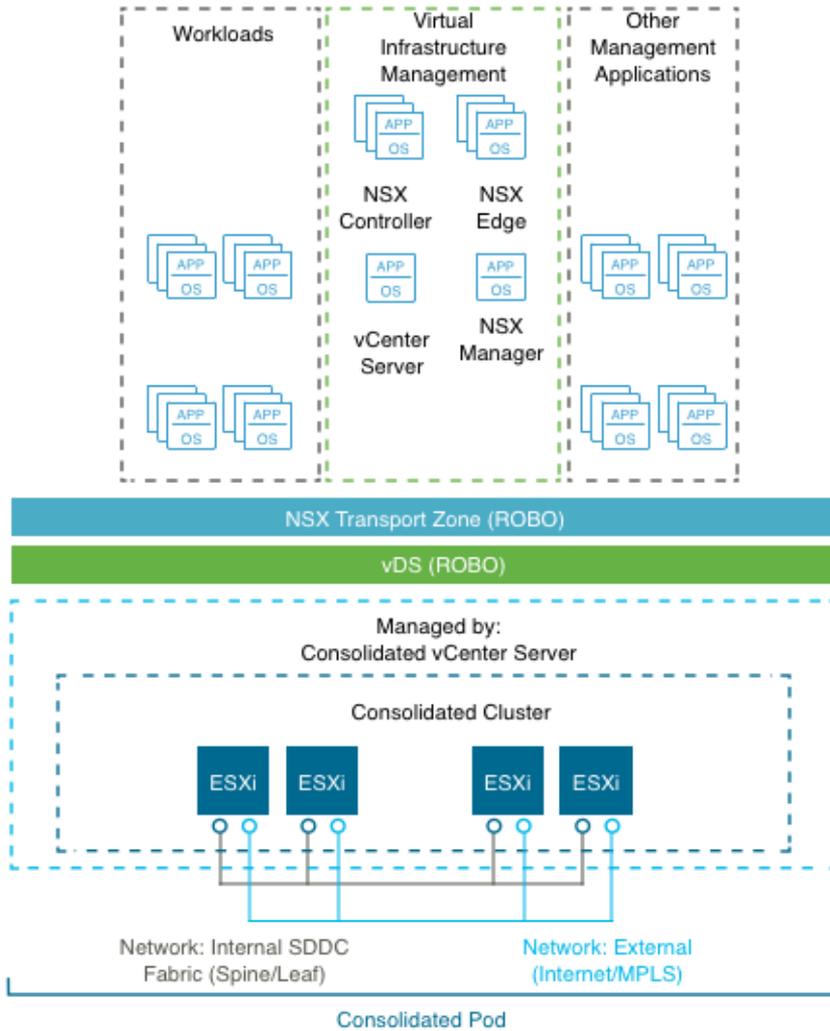
Figure 3-5. Virtual Infrastructure Layer in the SDDC



Virtual Infrastructure Design Overview

The ROBO virtual infrastructure consists of one or more ROBO locations. Each includes a consolidated pod.

Figure 3-6. Remote Office and Branch Office Logical Design



Management Pod

The management pod resides in the hub and runs the virtual machines that manage the SDDC. The management pod is instantiated as part of the VMware Validated Design for the SDDC. The management pod is a pre-requisite for creating consolidated POD's in ROBO locations. For additional information refer to the VMware Validated Design for SDDC.

Consolidated Pod

The consolidated pod runs the following services in the ROBO location:

- Virtual machines to manage the SDDC such as vCenter Server, NSX components, Remote Collectors, agents and other shared components.
- Required NSX services to enable north-south routing between the SDDC and the external network, and east-west routing inside the SDDC.

- SDDC tenant virtual machines to support workloads of different Service Level Agreements (SLAs).

Because this pod supports all SDDC, network, and production workloads at the ROBO site, it is important to ensure highly available physical components such as HVAC, power feeds and power supplies.

ESXi Design in ROBO

The ESXi design includes design decisions for boot options, user access, and the virtual machine swap configuration.

ESXi Hardware Requirements

You can find the ESXi hardware requirements in Physical Design Fundamentals. The following design outlines the design of the ESXi configuration.

ESXi Manual Install and Boot Options

You can install or boot ESXi 6.5 from the following storage systems:

SATA disk drives	SATA disk drives connected behind supported SAS controllers or supported on-board SATA controllers.
Serial-attached SCSI (SAS) disk drives	Supported for installing ESXi.
SAN	Dedicated SAN disk on Fibre Channel or iSCSI.
USB devices	Supported for installing ESXi. 16 GB or larger SD card is recommended.
FCoE	(Software Fibre Channel over Ethernet)

ESXi can boot from a disk larger than 2 TB if the system firmware and the firmware on any add-in card support it. See the vendor documentation.

ESXi Boot Disk and Scratch Configuration

For new installations of ESXi, the installer creates a 4 GB VFAT scratch partition. ESXi uses this scratch partition to store log files persistently. By default, `vm-support output`, which is used by VMware to troubleshoot ESXi host issues, is stored on this scratch partition.

An ESXi installation on USB media does not configure a default scratch partition. VMware recommends that you specify a scratch partition on a shared datastore and configure remote syslog logging for the host.

Table 3-21. ESXi Boot Disk Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-ESXi-001	Install and configure all ESXi hosts to boot using a SD device of 16 GB or greater.	SD cards are an inexpensive and easy to configure option for installing ESXi. Using SD cards allows allocation of all local HDDs to a VMware vSAN storage system.	When you use SD cards ESXi logs are not retained locally.

ESXi Host Access

After installation, ESXi hosts are added to a VMware vCenter Server system and managed through that vCenter Server system.

Direct access to the host console is still available and most commonly used for troubleshooting purposes. You can access ESXi hosts directly using one of these three methods:

Direct Console User Interface (DCUI)	Graphical interface on the console. Allows basic administrative controls and troubleshooting options.
ESXi Shell	A Linux-style bash login on the ESXi console itself.
Secure Shell (SSH) Access	Remote command-line console access.

You can enable or disable each method. By default, the ESXi Shell and SSH are disabled to secure the ESXi host. The DCUI is disabled only if Strict Lockdown Mode is enabled.

ESXi User Access

By default, root is the only user who can log in to an ESXi host directly, however, you can add ESXi hosts to an Active Directory domain. After the host has been added to an Active Directory domain, access can be granted through Active Directory groups. Auditing who has logged into the host also becomes easier.

Table 3-22. ESXi User Access Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-ESXi-002	Add each host to the Active Directory domain.	Using Active Directory membership allows greater flexibility in granting access to ESXi hosts. Ensuring that users log in with a unique user account allows greater visibility for auditing.	Adding hosts to the domain can add some administrative overhead.
ROBO-VI-ESXi-003	Change the default ESX Admins group to the SDDC-Admins Active Directory group. Add ESXi administrators to the SDDC-Admins group following standard access procedures.	Having an SDDC-Admins group is more secure because it removes a known administrative access point. In addition different groups allow for separation of management tasks.	Additional changes to the host's advanced settings are required.

Virtual Machine Swap Configuration

When a virtual machine is powered on, the system creates a VMkernel swap file to serve as a backing store for the virtual machine's RAM contents. The default swap file is stored in the same location as the virtual machine's configuration file. This simplifies configuration, however, it can cause an excess of unnecessary replication traffic.

You can reduce the amount of traffic that is replicated by changing the swap file location to a user-configured location on the host. However, it can take longer to perform VMware vSphere vMotion[®] operations when the swap file has to be recreated.

Table 3-23. Other ESXi Host Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-ESXi-004	Configure all ESXi hosts to synchronize time with the central NTP servers.	Required because deployment of vCenter Server Appliance on an ESXi host might fail if the host is not using NTP.	All firewalls located between the ESXi host and the NTP servers have to allow NTP traffic on the required network ports.

vCenter Server Design in ROBO

The vCenter Server design includes both the design for the vCenter Server instance and the VMware Platform Services Controller instance.

A Platform Services Controller groups a set of infrastructure services including vCenter Single Sign-On, License service, Lookup Service, and VMware Certificate Authority (VMCA). You can deploy the Platform Services Controller and the associated vCenter Server system on the same virtual machine (embedded Platform Services Controller) or on different virtual machines (external Platform Services Controller).

- [vCenter Server Deployment in ROBO](#)

The design decisions for vCenter Server deployment discuss the number of vCenter Server and Platform Services Controller instances, the type of installation, and the topology.

- [vCenter Server Networking in ROBO](#)

As specified in the physical networking design, the vCenter Server system must use static IP address and host name. The IP address must have valid (internal) DNS registration including reverse name resolution.

- [vCenter Server Redundancy in ROBO](#)

Protecting the vCenter Server system is important because it is the central point of management and monitoring for the ROBO SDDC. How you protect vCenter Server depends on maximum downtime tolerated, and on whether failover automation is required.

- [vCenter Server Appliance Sizing in ROBO](#)

The following tables outline minimum hardware requirements for the management vCenter Server appliance and the compute vCenter Server appliance.

- [vSphere Cluster Design in ROBO](#)

The cluster design takes into account the workload that the cluster must handle.

- [vCenter Server Customization in ROBO](#)

vCenter Server supports a rich set of customization options, including monitoring, virtual machine fault tolerance, and so on. For each feature, this VMware Validated Design specifies the design decisions.

■ **Use of Transport Layer Security Certificates in ROBO**

By default, vSphere 6.5 uses TLS/SSL certificates that are signed by the VMware Certificate Authority (VMCA). By default, these certificates are not trusted by end-user devices or browsers. It is a security best practice to replace at least user-facing certificates with certificates that are signed by a third-party or enterprise Certificate Authority (CA). Certificates for machine-to-machine communication can remain as VMCA signed certificates.

vCenter Server Deployment in ROBO

The design decisions for vCenter Server deployment discuss the number of vCenter Server and Platform Services Controller instances, the type of installation, and the topology.

Table 3-24. vCenter Server Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-001	Deploy a single vCenter Server Appliance with Embedded Platform Services Controller in each ROBO.	<ul style="list-style-type: none"> Supports clear separation of roles and responsibilities to ensure that only those administrators with proper authorization can administer certain ROBO sites. Simplifies operations of a large scale ROBO design. 	Requires licenses for each vCenter Server instance.

You can install vCenter Server as a Windows-based system or deploy the Linux-based VMware vCenter Server Appliance. The Linux-based vCenter Server Appliance is preconfigured, enables fast deployment, and potentially results in reduced Microsoft licensing costs.

Table 3-25. vCenter Server Platform Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-002	Deploy ROBO vCenter Server using the vCenter Server Appliance.	Allows for rapid deployment, enables scalability, and reduces Microsoft licensing costs.	Operational staff may need Linux experience to troubleshoot the appliance.

Platform Services Controller Design Decision Background

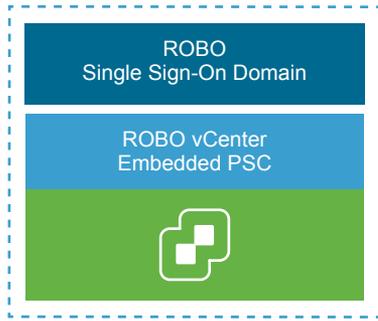
vCenter Server supports installation with an embedded Platform Services Controller (embedded deployment) or with an external Platform Services Controller.

- In an embedded deployment, vCenter Server and the Platform Services Controller run on the same virtual machine. Embedded deployments are recommended for standalone environments with only one vCenter Server system.
- Environments with an external Platform Services Controller can have multiple vCenter Server systems. The vCenter Server systems can use the same Platform Services Controller services. For example, several vCenter Server systems can use the same instance of vCenter Single Sign-On for authentication.
- If there is a need to replicate with other Platform Services Controller instances, or if the solution includes more than one vCenter Single Sign-On instance, you can deploy multiple external Platform Services Controller instances on separate virtual machines.

Table 3-26. Platform Service Controller Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-003	Deploy ROBO vCenter Server with an embedded Platform Services Controller.	Embedded Platform Services Controllers simplifies operations.	Embedded Platform Services Controllers do not support replication.

Figure 3-7. vCenter Server and Platform Services Controller Deployment Model



vCenter Server Networking in ROBO

As specified in the physical networking design, the vCenter Server system must use static IP address and host name. The IP address must have valid (internal) DNS registration including reverse name resolution.

The vCenter Server must maintain network connections to the following components:

- All VMware vSphere Client and vSphere Web Client user interfaces.
- Systems running vCenter Server add-on modules.
- Each ESXi host.

vCenter Server Redundancy in ROBO

Protecting the vCenter Server system is important because it is the central point of management and monitoring for the ROBO SDDC. How you protect vCenter Server depends on maximum downtime tolerated, and on whether failover automation is required.

The following table lists methods available for protecting the vCenter Server system and the vCenter Server Appliance.

Table 3-27. Methods for Protecting vCenter Server System and the vCenter Server Appliance

Redundancy Method	Protects vCenter Server system (Windows)	Protects Platform Services Controller (Windows)	Protects vCenter Server (Appliance)	Protects Platform Services Controller (Appliance)
Automated protection using vSphere HA.	Yes	Yes	Yes	Yes
Manual configuration and manual failover. For example, using a cold standby.	Yes	Yes	Yes	Yes

Table 3-27. Methods for Protecting vCenter Server System and the vCenter Server Appliance (Continued)

Redundancy Method	Protects vCenter Server system (Windows)	Protects Platform Services Controller (Windows)	Protects vCenter Server (Appliance)	Protects Platform Services Controller (Appliance)
HA Cluster with external load balancer	Not Available	Yes	Not Available	Yes
vCenter Server HA	Not Available	Not Available	Yes	Not Available

Table 3-28. vCenter Server Protection Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-004	Protect the vCenter Server appliance by using vSphere HA.	Supports availability objectives for vCenter Server appliance without a required manual intervention during a failure event.	vCenter Server will be unavailable during a vSphere HA failover.

vCenter Server Appliance Sizing in ROBO

The following tables outline minimum hardware requirements for the management vCenter Server appliance and the compute vCenter Server appliance.

Table 3-29. Logical Specification for vCenter Server Appliance

Attribute	Specification
vCenter Server version	6.5 (vCenter Server Appliance)
Physical or virtual system	Virtual (appliance)
Appliance Size	Small (up to 100 hosts / 1,000 VMs)
Platform Services Controller	Embedded
Number of CPUs	4
Memory	16 GB
Disk Space	290 GB

Table 3-30. vCenter Server Appliance Sizing Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-005	Configure the vCenter Server Appliance with the small size setting.	A tiny vCenter Server deployment supports only 100 virtual machines, which does not allow you to fully scale ROBO when taking into account management virtual machines. Based on the number of management and workload virtual machines that a ROBO deployment runs, a vCenter Server Appliance installed with the small size setting is sufficient.	A small vCenter Server deployment can support 100 hosts or 1000 virtual machines, which is more than ROBO is designed for.

vSphere Cluster Design in ROBO

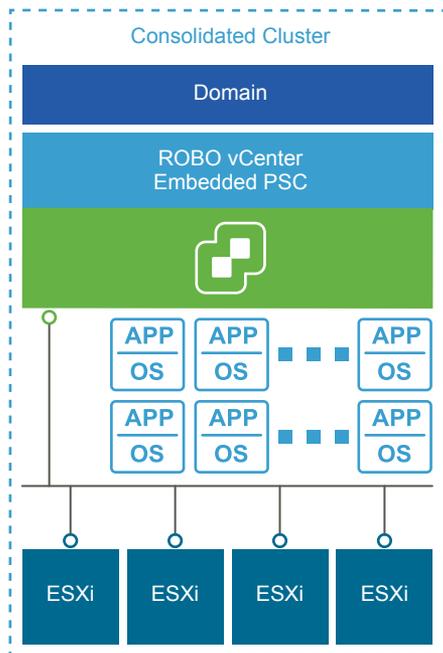
The cluster design takes into account the workload that the cluster must handle.

vSphere Cluster Design Decision Background

The following heuristics help with cluster design decisions.

- Decide to use fewer, larger hosts or more, smaller hosts.
 - A scale-up cluster has fewer, larger hosts.
 - A scale-out cluster has more, smaller hosts.
 - A virtualized server cluster typically has more hosts with fewer virtual machines per host.
- Compare the capital costs of purchasing fewer, larger hosts with the costs of purchasing more, smaller hosts. Costs vary between vendors and models.
- Evaluate the operational costs of managing a few hosts with the costs of managing more hosts.
- Consider the purpose of the cluster.
- Consider the total number of hosts and cluster limits.

Figure 3-8. vSphere Logical Cluster Layout



vSphere High Availability Design in ROBO

VMware vSphere High Availability (vSphere HA) protects your virtual machines in case of host failure by restarting virtual machines on other hosts in the cluster when a host fails.

vSphere HA Design Basics

During configuration of the cluster, the hosts elect a master host. The master host communicates with the vCenter Server system and monitors the virtual machines and secondary hosts in the cluster.

The master hosts detect different types of failure:

- Host failure from an unexpected power failure
- Host network isolation or connectivity failure
- Loss of storage connectivity
- Problems with virtual machine OS availability

Table 3-31. vSphere HA Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-06	Use vSphere HA to protect virtual machines in the Consolidated cluster against host failures.	vSphere HA supports a robust level of protection for virtual machine availability.	Sufficient resources on the remaining hosts are required so that virtual machines can be migrated to those hosts in the event of a host outage.
ROBO-VI-VC-07	Set vSphere HA Host Isolation Response to Power Off.	vSAN requires that the HA Isolation Response be set to Power Off, and to restart VMs on available hosts.	VMs are powered off in case of a false positive and a host is declared isolated incorrectly.

vSphere HA Admission Control Policy Configuration

The vSphere HA Admission Control Policy allows an administrator to configure how the cluster judges available resources. In a smaller vSphere HA cluster, a larger proportion of the cluster resources are reserved to accommodate host failures, based on the selected policy.

The following policies are available:

Host failures the cluster tolerates. vSphere HA ensures that a specified number of hosts can fail and sufficient resources remain in the cluster to fail over all the virtual machines from those hosts.

Percentage of cluster resources reserved. Percentage of cluster resources reserved. vSphere HA ensures that a specified percentage of aggregate CPU and memory resources are reserved for failover.

Specify Failover Hosts. When a host fails, vSphere HA attempts to restart its virtual machines on any of the specified failover hosts. If restart is not possible, for example the failover hosts have insufficient resources or have failed as well, then vSphere HA attempts to restart the virtual machines on other hosts in the cluster.

Consolidated Cluster Design in ROBO

The consolidated cluster design determines the number of hosts and vSphere HA settings for the cluster.

The management virtual machines, NSX controllers and edges, and tenant workloads run on the ESXi hosts in the consolidated cluster.

Figure 3-9. Consolidated Cluster Resource Pools

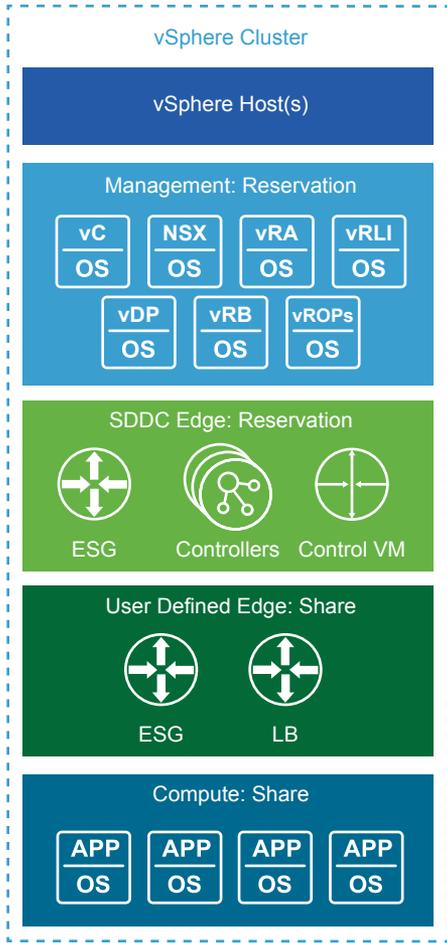


Table 3-32. Management Cluster Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-008	Create a consolidated cluster with 4 hosts.	Three hosts are used to provide n+1 redundancy for the vSAN cluster. The fourth host is used to guarantee n+1 for vSAN redundancy during maintenance operations. The fourth host is also used to ensure adherence of virtual machine DRS anti-affinity rules.	Additional host resources are required for redundancy.
ROBO-VI-VC-009	Configure Admission Control for 1 host failure and percentage based failover capacity.	Using the percentage-based reservation works well in situations where virtual machines have varying and sometime significant CPU or memory reservations. vSphere 6.5 automatically calculates the reserved percentage based on host failures to tolerate and the number of hosts in the cluster.	In a four host cluster only the resources of three hosts are available for use.

Table 3-32. Management Cluster Design Decisions (Continued)

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-010	Create a host profile for the consolidated Cluster.	Utilizing host profiles simplifies configuration of hosts and ensures settings are uniform across the cluster.	Anytime an authorized change to a host is made the host profile must be updated to reflect the change or the status will show non-compliant.
ROBO-VI-VC-011	Set up VLAN-backed port groups for external access and management.	Edge gateways need access to the external network in addition to the management network.	VLAN-backed port groups must be configured with the correct number of ports, or with elastic port allocation.
ROBO-VI-VC-012	Create a resource pool for the required ROBO management virtual machines with a CPU share level of High, a memory share level of normal, and a 102 GB memory reservation.	These virtual machines perform management and monitoring of the ROBO. In a contention situation it is imperative that these virtual machines receive all the resources required.	During contention management components receive more resources than user workloads as such monitoring and capacity management must be a proactive activity.
ROBO-VI-VC-013	Create a resource pool for the required ROBO NSX Controllers and edge appliances with a CPU share level of High, a memory share of normal, and 15 GB memory reservation.	The NSX components control all network traffic in and out of the SDDC as well as update route information for inter-SDDC communication. In a contention situation it is imperative that these virtual machines receive all the resources required.	During contention NSX components receive more resources than user workloads as such monitoring and capacity management must be a proactive activity.
ROBO-VI-VC-014	Create a resource pool for all user NSX Edge devices with a CPU share value of Normal and a memory share value of Normal.	NSX Edges for users, created by vRealize Automation, support functions such as load balancing for user workloads. These Edge devices do not support the entire SDDC, and as such they receive a lower amount of resources during contention.	During contention these NSX Edge devices will receive fewer resources than the SDDC Edge devices. As a result, monitoring and capacity management must be a proactive activity.
ROBO-VI-VC-015	Create a resource pool for all user virtual machines with a CPU share value of Normal and a memory share value of Normal.	Creating virtual machines outside of a resource pool will have a negative impact on all other virtual machines during contention. In a consolidated cluster the SDDC edge devices must be guaranteed resources above all other workloads as to not impact network connectivity. Setting the share values to normal gives the SDDC edges more shares of resources during contention ensuring network traffic is not impacted.	During contention user workload virtual machines could be starved for resources and experience poor performance. It is critical that monitoring and capacity management must be a proactive activity and that capacity is added or a dedicated edge cluster is created before contention occurs. Some workloads cannot be directly deployed to a resource pool, as such additional administrative overhead may be required to move workloads into resource pools.
ROBO-VI-VC-016	Create a DRS VM to Host rule that runs vCenter Server on the first four hosts in the cluster.	In the event of an emergency vCenter Server is easier to find and bring up.	Limits DRS ability to place vCenter Server on any available host in the cluster.

Table 3-33. Consolidated Cluster Attributes

Attribute	Specification
Capacity for host failures per cluster	1
Number of usable hosts per cluster	3
Minimum number of hosts required to support the Consolidated cluster	4

vCenter Server Customization in ROBO

vCenter Server supports a rich set of customization options, including monitoring, virtual machine fault tolerance, and so on. For each feature, this VMware Validated Design specifies the design decisions.

VM and Application Monitoring Service

When VM and Application Monitoring is enabled, the VM and Application Monitoring service, which uses VMware Tools, evaluates whether each virtual machine in the cluster is running. The service checks for regular heartbeats and I/O activity from the VMware Tools process running on guests. If the service receives no heartbeats or I/O activity, it is likely that the guest operating system has failed or that VMware Tools is not being allocated time for heartbeats or I/O activity. In this case, the service determines that the virtual machine has failed and reboots the virtual machine.

Enable Virtual Machine Monitoring for automatic restart of a failed virtual machine. The application or service that is running on the virtual machine must be capable of restarting successfully after a reboot or the VM restart is not sufficient.

Table 3-34. Monitor Virtual Machines Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-017	Enable Virtual Machine Monitoring.	Virtual Machine Monitoring provides adequate in-guest protection for most VM workloads.	There is no downside to enabling Virtual Machine Monitoring.

VMware vSphere Distributed Resource Scheduling (DRS)

vSphere Distributed Resource Scheduling provides load balancing of a cluster by migrating workloads from heavily loaded hosts to less utilized hosts in the cluster. DRS supports manual and automatic modes.

Manual

Recommendations are made but an administrator needs to confirm the changes

Automatic

Automatic management can be set to five different levels. At the lowest setting, workloads are placed automatically at power on and only migrated to fulfill certain criteria, such as entering maintenance mode. At the highest level, any migration that would provide a slight improvement in balancing will be executed.

Table 3-35. vSphere Distributed Resource Scheduling Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-018	Enable DRS and set it to Fully Automated, with the default setting (medium).	The default settings provide the best trade-off between load balancing and excessive migration with vMotion events.	In the event of a vCenter outage, mapping from virtual machines to ESXi hosts might be more difficult to determine.

Enhanced vMotion Compatibility (EVC)

EVC works by masking certain features of newer CPUs to allow migration between hosts containing older CPUs. EVC works only with CPUs from the same manufacturer and there are limits to the version difference gaps between the CPU families.

If you set EVC during cluster creation, you can add hosts with newer CPUs at a later date without disruption. You can use EVC for a rolling upgrade of all hardware with zero downtime.

Set EVC to the highest level possible with the current CPUs in use.

Table 3-36. VMware Enhanced vMotion Compatibility Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-019	Enable Enhanced vMotion Compatibility. Set EVC mode to the highest level supported by all hosts in the cluster.	Allows cluster upgrades without virtual machine downtime.	You can enable EVC only if clusters contain hosts with CPUs from the same vendor.

Use of Transport Layer Security Certificates in ROBO

By default, vSphere 6.5 uses TLS/SSL certificates that are signed by the VMware Certificate Authority (VMCA). By default, these certificates are not trusted by end-user devices or browsers. It is a security best practice to replace at least user-facing certificates with certificates that are signed by a third-party or enterprise Certificate Authority (CA). Certificates for machine-to-machine communication can remain as VMCA signed certificates.

Table 3-37. vCenter Server TLS Certificate Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-VC-020	Replace the vCenter Server machine certificate with a certificate signed by a 3rd party Public Key Infrastructure.	Infrastructure administrators connect to vCenter Server by way of a Web browser to perform configuration, management and troubleshooting activities. Certificate warnings result with the default certificate.	Replacing and managing certificates is an operational overhead.
ROBO-VI-VC-021	Use a SHA-2 or higher algorithm when signing certificates.	The SHA-1 algorithm is considered less secure and has been deprecated.	Not all certificate authorities support SHA-2.

Virtualization Network Design in ROBO

A well-designed network helps the organization meet its business goals. It prevents unauthorized access, and provides timely access to business data.

This network virtualization design uses vSphere and VMware NSX for vSphere to implement virtual networking.

- [Virtual Network Design Guidelines in ROBO](#)

This VMware Validated Design follows high-level network design guidelines and networking best practices.

- [Virtual Switches in ROBO](#)

Virtual switches simplify the configuration process by providing one single pane of glass view for performing virtual network management tasks.

- [NIC Teaming in ROBO](#)

You can use NIC teaming to increase the network bandwidth available in a network path, and to provide the redundancy that supports higher availability.

- [Network I/O Control in ROBO](#)

When Network I/O Control is enabled, the distributed switch allocates bandwidth for the following system traffic types.

- [VXLAN in ROBO](#)

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

- [vMotion TCP/IP Stack in ROBO](#)

Use the vMotion TCP/IP stack to isolate traffic for vMotion and to assign a dedicated default gateway for vMotion traffic.

Virtual Network Design Guidelines in ROBO

This VMware Validated Design follows high-level network design guidelines and networking best practices.

Design Goals

The high-level design goals apply regardless of your environment.

- Meet diverse needs. The network must meet the diverse needs of many different entities in an organization. These entities include applications, services, storage, administrators, and users.
- Reduce costs. Reducing costs is one of the simpler goals to achieve in the vSphere infrastructure. Server consolidation alone reduces network costs by reducing the number of required network ports and NICs, but a more efficient network design is desirable. For example, configuring two 10 GbE NICs with VLANs might be more cost effective than configuring a dozen 1 GbE NICs on separate physical networks.

- Boost performance. You can achieve performance improvement and decrease the time that is required to perform maintenance by providing sufficient bandwidth, which reduces contention and latency.
- Improve availability. A well-designed network improves availability, typically by providing network redundancy.
- Support security. A well-designed network supports an acceptable level of security through controlled access (where required) and isolation (where necessary).
- Enhance infrastructure functionality. You can configure the network to support vSphere features such as vSphere vMotion, vSphere High Availability, and vSphere Fault Tolerance.

Best Practices

Follow networking best practices throughout your environment.

- Separate network services from one another to achieve greater security and better performance.
- Use Network I/O Control and traffic shaping to guarantee bandwidth to critical virtual machines. During network contention these critical virtual machines will receive a higher percentage of the bandwidth.
- Separate network services on a single vSphere Distributed Switch by attaching them to port groups with different VLAN IDs.
- Keep vSphere vMotion traffic on a separate network. When migration with vMotion occurs, the contents of the guest operating system's memory is transmitted over the network. You can put vSphere vMotion on a separate network by using a dedicated vSphere vMotion VLAN.
- When using passthrough devices with a Linux kernel version 2.6.20 or earlier guest OS, avoid MSI and MSI-X modes because these modes have significant performance impact.
- For best performance, use VMXNET3 virtual NICs.
- Ensure that physical network adapters that are connected to the same vSphere Standard Switch or vSphere Distributed Switch are also connected to the same physical network.

Network Segmentation and VLANs

Separating different types of traffic is required to reduce contention and latency. Separate networks are also required for access security.

High latency on any network can negatively affect performance. Some components are more sensitive to high latency than others. For example, reducing latency is important on the IP storage and the vSphere Fault Tolerance logging network because latency on these networks can negatively affect the performance of multiple virtual machines.

Depending on the application or service, high latency on specific virtual machine networks can also negatively affect performance. Use information gathered from the current state analysis and from interviews with key stakeholder and SMEs to determine which workloads and networks are especially sensitive to high latency.

Virtual Networks

Determine the number of networks or VLANs that are required depending on the type of traffic.

- vSphere operational traffic.
 - Management
 - vMotion
 - vSAN
 - NFS Storage
 - VXLAN
- Traffic that supports the organization's services and applications.

Virtual Switches in ROBO

Virtual switches simplify the configuration process by providing one single pane of glass view for performing virtual network management tasks.

Virtual Switch Design Background in ROBO

A vSphere Distributed Switch (distributed switch) offers several enhancements over standard virtual switches.

Centralized management

Because distributed switches are created and managed centrally on a vCenter Server system, they make the switch configuration more consistent across ESXi hosts. Centralized management saves time, reduces mistakes, and lowers operational costs.

Additional features

Distributed switches offer features that are not available on standard virtual switches. Some of these features can be useful to the applications and services that are running in the organization's infrastructure. For example, NetFlow and port mirroring provide monitoring and troubleshooting capabilities to the virtual infrastructure.

Consider the following caveats for distributed switches.

- Distributed switches are not manageable when vCenter Server is unavailable. vCenter Server therefore becomes a tier one application.

Health Check in ROBO

The health check service helps identify and troubleshoot configuration errors in vSphere distributed switches.

Health check helps identify the following common configuration errors.

- Mismatched VLAN trunks between an ESXi host and the physical switches it's connected to.
- Mismatched MTU settings between physical network adapters, distributed switches, and physical switch ports.

- Mismatched virtual switch teaming policies for the physical switch port-channel settings.

Health check monitors VLAN, MTU, and teaming policies.

VLANs Checks whether the VLAN settings on the distributed switch match the trunk port configuration on the connected physical switch ports.

MTU For each VLAN, health check determines whether the physical access switch port's MTU jumbo frame setting matches the distributed switch MTU setting.

Teaming policies Health check determines whether the connected access ports of the physical switch that participate in an EtherChannel are paired with distributed ports whose teaming policy is IP hash.

Health check is limited to the access switch port to which the ESXi hosts' NICs connects.

Design ID	Design Decision	Design Justification	Design Implication
ROBO-VI-NET-001	Enable vSphere Distributed Switch Health Check on the virtual distributed switch.	vSphere Distributed Switch Health Check ensures all VLANS are trunked to all hosts attached to the vSphere Distributed Switch and ensures MTU sizes match the physical network.	You must have a minimum of two physical uplinks to use this feature.

Note For VLAN and MTU checks, at least two physical NICs for the distributed switch are required. For a teaming policy check, at least two physical NICs and two hosts are required when applying the policy.

Type of Virtual Switches in ROBO

Create a single vSphere Distributed Switch.

Table 3-38. Virtual Switch Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-NET-002	Use the vSphere Distributed Switch (VDS).	The vSphere Distributed Switch simplifies management.	Migration from a VSS to a VDS requires a minimum of two physical NICs to maintain redundancy.

Consolidated Cluster Distributed Switch in ROBO

The consolidated cluster uses a single vSphere Distributed Switch with the following configuration settings.

Table 3-39. Virtual Switch for the Consolidated Cluster

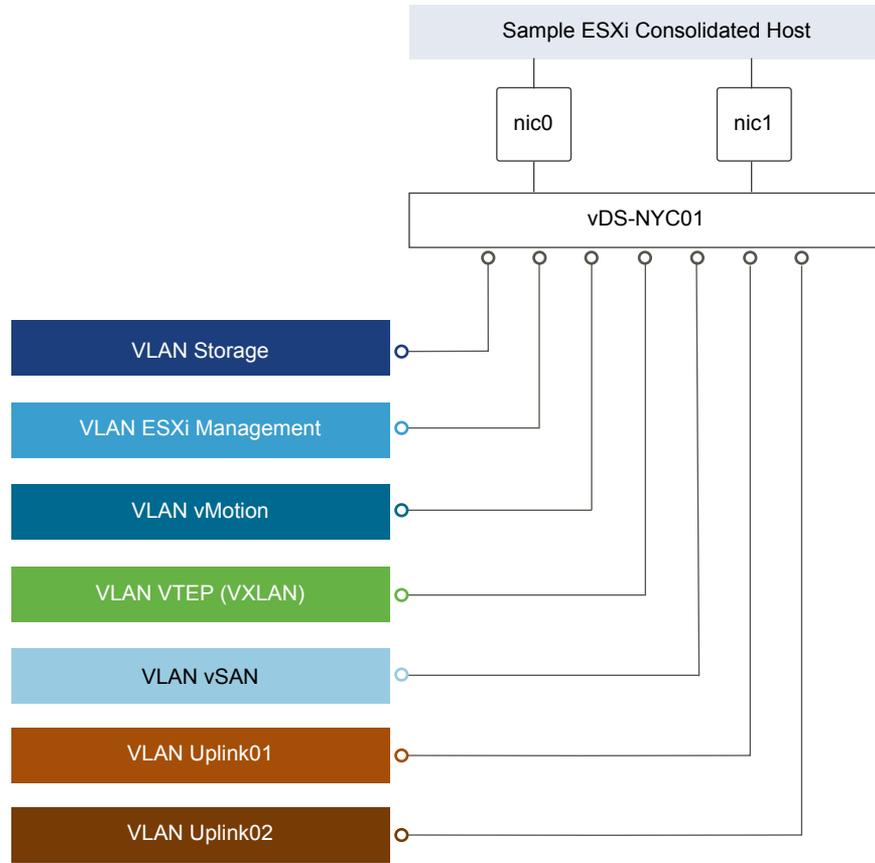
vSphere Distributed Switch Name	Function	Network I/O Control	Number of Physical NIC Ports	MTU
vDS-ROBO01	<ul style="list-style-type: none"> ■ ESXi Management ■ Secondary Storage ■ vSAN ■ vSphere vMotion ■ VXLAN Tunnel Endpoint (VTEP) ■ Uplinks (2) to enable ECMP 	Enabled	2	9000

Note vDS-ROBO01 is a variable that you replace with the name of the vSphere Distributed Switch in your environment. Your vSphere Distributed Switch should be named based on the ROBO location. For example: vDS-NYC01.

Table 3-40. vDS-MgmtPort Group Configuration Settings

Parameter	Setting
Failover detection	Link status only
Notify switches	Enabled
Failback	Yes
Failover order	Active uplinks: Uplink1, Uplink2

Figure 3-10. Network Switch Design for the Consolidated Cluster



This section expands on the logical network design by providing details on the physical NIC layout and physical network attributes.

Table 3-41. Consolidated Virtual Switches by Physical/Virtual NIC

vSphere Distributed Switch	vmnic	Function
vDS-ROBO01	0	Uplink
vDS-ROBO01	1	Uplink

Note The following VLANs are meant as samples. Your actual implementation depends on your environment.

Table 3-42. Consolidated Virtual Switch Port Groups and VLANs

vSphere Distributed Switch	Port Group Name	Teaming Policy	Active Uplinks	VLAN ID
vDS-ROBO01	vDS-ROBO01-Management	Route based on physical NIC load	0, 1	1811
vDS-ROBO01	vDS-ROBO01-vMotion	Route based on physical NIC load	0, 1	1812
vDS-ROBO01	vDS-ROBO01-VSAN	Route based on physical NIC load	0, 1	1813
vDS-ROBO01	Auto Generated (NSX VTEP)	Route based on SRC-ID	0, 1	1814
vDS-ROBO01	vDS-ROBO01-Storage (Optional)	Route based on physical NIC load	0, 1	1815

Table 3-42. Consolidated Virtual Switch Port Groups and VLANs (Continued)

vSphere Distributed Switch	Port Group Name	Teaming Policy	Active Uplinks	VLAN ID
vDS-ROBO01	vDS-ROBO01-Uplink01	Route based on physical NIC load	0, 1	1816
vDS-ROBO01	vDS-ROBO01-Uplink02	Route based on physical NIC load	0, 1	1817

Table 3-43. Management VMkernel Adapter

vSphere Distributed Switch	Network Label	Connected Port Group	Enabled Services	MTU
vDS-ROBO01	Management	vDS-ROBO01-Management	Management Traffic	1500 (Default)
vDS-ROBO01	vMotion	vDS-ROBO01-vMotion	vMotion Traffic	9000
vDS-ROBO01	VSAN	vDS-ROBO01-VSAN	vSAN	9000
vDS-ROBO01	Storage	vDS-ROBO01-Storage (Optional)	-	9000
vDS-ROBO01	VTEP	Auto Generated (NSX VTEP)	-	9000

For more information on the physical network design specifications, see [Physical Networking Design in ROBO](#).

NIC Teaming in ROBO

You can use NIC teaming to increase the network bandwidth available in a network path, and to provide the redundancy that supports higher availability.

NIC teaming helps avoid a single point of failure and provides options for load balancing of traffic. To further reduce the risk of a single point of failure, build NIC teams by using ports from multiple NIC and motherboard interfaces.

Create a single virtual switch with teamed NICs across separate physical switches.

This VMware Validated Design uses an active-active configuration using the route that is based on physical NIC load algorithm for teaming. In this configuration, idle network cards do not wait for a failure to occur, and they aggregate bandwidth.

Benefits and Overview

NIC teaming helps avoid a single point of failure and provides options for load balancing of traffic. To further reduce the risk of a single point of failure, build NIC teams by using ports from multiple NIC and motherboard interfaces.

Create a single virtual switch with teamed NICs across separate physical switches.

This VMware Validated Design uses an active-active configuration using the route that is based on physical NIC load algorithm for teaming. In this configuration, idle network cards do not wait for a failure to occur, and they aggregate bandwidth.

NIC Teaming Design Background

For a predictable level of performance, use multiple network adapters in one of the following configurations.

- An active-passive configuration that uses explicit failover when connected to two separate switches.
- An active-active configuration in which two or more physical NICs in the server are assigned the active role.

This validated design uses an active-active configuration.

Table 3-44. NIC Teaming and Policy

Design Quality	Active-Active	Active-Passive	Comments
Availability	↑	↑	Using teaming regardless of the option increases the availability of the environment.
Manageability	o	o	Neither design option impacts manageability.
Performance	↑	o	An active-active configuration can send traffic across either NIC, thereby increasing the available bandwidth. This configuration provides a benefit if the NICs are being shared among traffic types and Network I/O Control is used.
Recoverability	o	o	Neither design option impacts recoverability.
Security	o	o	Neither design option impacts security.

Legend: ↑ = positive impact on quality; ↓ = negative impact on quality; o = no impact on quality.

Table 3-45. NIC Teaming Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-NET-003	Use the Route based on physical NIC load teaming algorithm for all port groups except for ones that carry VXLAN traffic. VTEP kernel ports and VXLAN traffic will use Route based on SRC-ID.	Reduce complexity of the network design and increase resiliency and performance.	Because NSX does not support Route based on physical NIC load two different algorithms are necessary.

Network I/O Control in ROBO

When Network I/O Control is enabled, the distributed switch allocates bandwidth for the following system traffic types.

- Fault tolerance traffic
- iSCSI traffic
- vSphere vMotion traffic
- Management traffic
- VMware vSphere Replication traffic
- NFS traffic
- vSAN traffic

- vSphere Data Protection backup traffic
- Virtual machine traffic

How Network I/O Control Works

Network I/O Control enforces the share value specified for the different traffic types only when there is network contention. When contention occurs Network I/O Control applies the share values set to each traffic type. As a result, less important traffic, as defined by the share percentage, will be throttled, allowing more important traffic types to gain access to more network resources.

Network I/O Control also allows the reservation of bandwidth for system traffic based on the capacity of the physical adapters on a host, and enables fine-grained resource control at the virtual machine network adapter level. Resource control is similar to the model for vCenter CPU and memory reservations.

Network I/O Control Heuristics

The following heuristics can help with design decisions.

Shares vs. Limits	When you use bandwidth allocation, consider using shares instead of limits. Limits impose hard limits on the amount of bandwidth used by a traffic flow even when network bandwidth is available.
Limits on Certain Resource Pools	Consider imposing limits on a given resource pool. For example, if you put a limit on vSphere vMotion traffic, you can benefit in situations where multiple vSphere vMotion data transfers, initiated on different hosts at the same time, result in oversubscription at the physical network level. By limiting the available bandwidth for vSphere vMotion at the ESXi host level, you can prevent performance degradation for other traffic.
Teaming Policy	When you use Network I/O Control, use Route based on physical NIC load teaming as a distributed switch teaming policy to maximize the networking capacity utilization. With load-based teaming, traffic might move among uplinks, and reordering of packets at the receiver can result occasionally.
Traffic Shaping	Use distributed port groups to apply configuration policies to different traffic types. Traffic shaping can help in situations where multiple vSphere vMotion migrations initiated on different hosts converge on the same destination host. The actual limit and reservation also depend on the traffic shaping policy for the distributed port group where the adapter is connected to.

Network I/O Control Design Decisions

Based on the heuristics, this design has the following decisions.

Table 3-46. Network I/O Control Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-NET-004	Enable Network I/O Control on all distributed switches.	Increase resiliency and performance of the network.	If configured incorrectly Network I/O Control could impact network performance for critical traffic types.
ROBO-VI-NET-005	Set the share value for vMotion traffic to Low.	During times of contention vMotion traffic is not as important as virtual machine or storage traffic.	During times of network contention the vMotion migration of virtual machines from one physical server to another will take longer than usual to complete.
ROBO-VI-NET-006	Set the share value for vSphere Replication traffic to Low.	vSphere Replication is not used in this design therefore it can be set to the lowest priority.	None.
ROBO-VI-NET-007	Set the share value for vSAN to High.	During times of contention vSAN traffic needs guaranteed bandwidth so virtual machine performance does not suffer.	None.
ROBO-VI-NET-008	Set the share value for Management to Normal.	By keeping the default setting of Normal, management traffic is prioritized higher than vMotion traffic, but lower than vSAN traffic. Management traffic is important as it ensures the hosts can still be managed during times of network contention.	None.
ROBO-VI-NET-009	Set the share value for NFS Traffic to Low.	Because NFS can be used for secondary storage, such as backups, it is not as important as vSAN traffic, by prioritizing it lower vSAN is not impacted.	During times of contention services such as backups will be slower than usual.
ROBO-VI-NET-010	Set the share value for vSphere Data Protection Backup traffic to Low.	During times of contention it is more important that primary functions of the ROBO SDDC continue to have access to network resources over backup traffic.	During times of contention VDP backups will be slower than usual.
ROBO-VI-NET-011	Set the share value for virtual machines to High.	Virtual machines are the most important asset in the ROBO SDDC. Leaving the default setting of High ensures that they will always have access to the network resources they need.	None.
ROBO-VI-NET-012	Set the share value for Fault Tolerance to Low.	Fault Tolerance is not used in this design therefore it can be set to the lowest priority.	None.
ROBO-VI-NET-013	Set the share value for iSCSI traffic to Normal.	Because iSCSI can be used for secondary storage, such as backups, it is not as important as vSAN traffic, by prioritizing it lower vSAN is not impacted.	During times of contention services such as backups will be slower than usual.

VXLAN in ROBO

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

The first step in creating these logical networks is to abstract and pool the networking resources. Just as vSphere abstracts compute capacity from the server hardware to create virtual pools of resources that can be consumed as a service, vSphere Distributed Switch and VXLAN abstract the network into a generalized pool of network capacity and separate the consumption of these services from the underlying physical infrastructure. A network capacity pool can span physical boundaries, optimizing compute resource utilization across clusters, pods, and geographically-separated data centers. The unified pool of network capacity can then be optimally segmented into logical networks that are directly attached to specific applications.

VXLAN works by creating Layer 2 logical networks that are encapsulated in standard Layer 3 IP packets. A Segment ID in every frame differentiates the VXLAN logical networks from each other without any need for VLAN tags. As a result, large numbers of isolated Layer 2 VXLAN networks can coexist on a common Layer 3 infrastructure.

In the vSphere architecture, the encapsulation is performed between the virtual NIC of the guest VM and the logical port on the virtual switch, making VXLAN transparent to both the guest virtual machines and the underlying Layer 3 network. Gateway services between VXLAN and non-VXLAN hosts (for example, a physical server or the Internet router) are performed by the NSX Edge Services Gateway appliance. The Edge gateway translates VXLAN segment IDs to VLAN IDs, so that non-VXLAN hosts can communicate with virtual machines on a VXLAN network.

Table 3-47. VXLAN Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI- NET-014	Use NSX for vSphere to introduce VXLANs for the use of virtual application networks and tenant networks.	Simplify the network configuration for each tenant via centralized virtual network management.	Requires additional compute and storage resources to deploy NSX components. Additional training may be needed on NSX.
ROBO-VI- NET-015	Use VXLAN along with NSX Edge gateways, the Distributed Logical Router (DLR) to provide management application and customer/tenant network capabilities.	Create isolated, multi-tenant broadcast domains across data center fabrics to create elastic, logical networks that span physical network boundaries. Leverage benefits of network virtualization.	VXLAN requires an MTU of 1600 bytes or greater.

vMotion TCP/IP Stack in ROBO

Use the vMotion TCP/IP stack to isolate traffic for vMotion and to assign a dedicated default gateway for vMotion traffic.

By using a separate TCP/IP stack, you can manage vMotion and cold migration traffic according to the topology of the network, and as required for your organization.

- Route the traffic for the migration of virtual machines that are powered on or powered off by using a default gateway that is different from the gateway assigned to the default stack on the host.
- Assign a separate set of buffers and sockets.
- Avoid routing table conflicts that might otherwise appear when many features are using a common TCP/IP stack.
- Isolate traffic to improve security.

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI- NET-016	Use the vMotion TCP/IP stack for vMotion traffic.	By leveraging the vMotion TCP/IP stack, vMotion traffic can utilize a default gateway on its own subnet, allowing for vMotion traffic to go over Layer 3 networks.	The vMotion TCP/IP stack is not available in the vDS VMkernel creation wizard, and as such the VMkernel adapter must be created directly on a host.

NSX Design

This design implements software-defined networking by using VMware NSX™ for vSphere®. With NSX for vSphere, virtualization delivers for networking what it has already delivered for compute and storage.

In much the same way that server virtualization programmatically creates, snapshots, deletes, and restores software-based virtual machines (VMs), NSX network virtualization programmatically creates, snapshots, deletes, and restores software-based virtual networks. The result is a transformative approach to networking that not only enables data center managers to achieve orders of magnitude better agility and economics, but also supports a vastly simplified operational model for the underlying physical network. NSX for vSphere is a non-disruptive solution because it can be deployed on any IP network, including existing traditional networking models and next-generation fabric architectures, from any vendor.

When administrators provision workloads, network management is one of the most time-consuming tasks. Most of the time spent provisioning networks is consumed configuring individual components in the physical infrastructure and verifying that network changes do not affect other devices that are using the same networking infrastructure.

The need to pre-provision and configure networks is a major constraint to cloud deployments where speed, agility, and flexibility are critical requirements. Pre-provisioned physical networks can allow for the rapid creation of virtual networks and faster deployment times of workloads utilizing the virtual network. As long as the physical network that you need is already available on the host where the workload is to be deployed, this works well. However, if the network is not available on a given host, you must find a host with the available network and spare capacity to run your workload in your environment.

To get around this bottleneck requires a decoupling of virtual networks from their physical counterparts. This, in turn, requires that you can programmatically recreate all physical networking attributes that are required by workloads in the virtualized environment. Because network virtualization supports the creation of virtual networks without modification of the physical network infrastructure, it allows more rapid network provisioning.

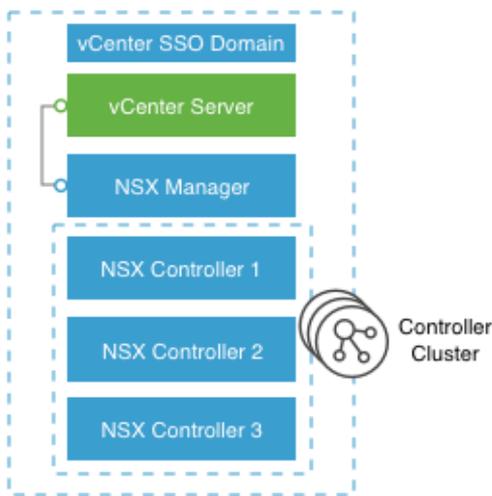
NSX for vSphere Design in ROBO

A NSX instance is tied to a single vCenter Server instance.

Table 3-48. NSX for vSphere Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-001	Use one NSX instance per ROBO.	Software-defined Networking (SDN) capabilities offered by NSX, such as load balancing and firewalls, are crucial to support the cloud management platform operations, and also for the management applications that need these capabilities.	None.

Figure 3-11. Architecture of NSX for vSphere



NSX Components in ROBO

The following sections describe the components in the solution and how they are relevant to the network virtualization design.

Consumption Layer

NSX for vSphere can be consumed by the cloud management platform (CMP), represented by vRealize Automation, by using the NSX REST API and the vSphere Web Client.

Cloud Management Platform

NSX for vSphere is consumed by vRealize Automation. NSX offers self-service provisioning of virtual networks and related features from a service portal. Details of the service requests and their orchestration are outside the scope of this document and can be referenced in the *VMware Validated Design for SDDC* document.

API

NSX for vSphere offers a powerful management interface through its REST API.

- A client can read an object by making an HTTP GET request to the object's resource URL.
- A client can write (create or modify) an object with an HTTP PUT or POST request that includes a new or changed XML document for the object.
- A client can delete an object with an HTTP DELETE request.

vSphere Web Client

The NSX Manager component provides a networking and security plug-in in the vSphere Web Client. This plug-in provides an interface to consuming virtualized networking from the NSX Manager for users that have sufficient privileges.

Table 3-49. Consumption Method Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-002	End user access is accomplished by using vRealize Automation services. Administrators use both the vSphere Web Client and the NSX REST API.	vRealize Automation services are used for the customer-facing portal. The vSphere Web Client consumes NSX for vSphere resources through the Network and Security plug-in. The NSX REST API offers the potential of scripting repeating actions and operations.	Customers typically interact only indirectly with NSX from the vRealize Automation portal. Administrators interact with NSX from the vSphere Web Client and API.

NSX Manager

NSX Manager provides the centralized management plane for NSX for vSphere and has a one-to-one mapping to vCenter Server workloads.

NSX Manager performs the following functions.

- Provides the single point of configuration and the REST API entry-points for NSX in a vSphere environment.
- Deploys NSX Controller clusters, Edge distributed routers, and Edge service gateways in the form of OVF appliances, guest introspection services, and so on.
- Prepares ESXi hosts for NSX by installing VXLAN, distributed routing and firewall kernel modules, and the User World Agent (UWA).
- Communicates with NSX Controller clusters over REST and with hosts over the RabbitMQ message bus. This internal message bus is specific to NSX for vSphere and does not require setup of additional services.
- Generates certificates for the NSX Controller instances and ESXi hosts to secure control plane communications with mutual authentication.

NSX Controller

An NSX Controller performs the following functions.

- Provides the control plane to distribute VXLAN and logical routing information to ESXi hosts.
- Includes nodes that are clustered for scale-out and high availability.
- Slices network information across cluster nodes for redundancy.
- Removes requirement of VXLAN Layer 3 multicast in the physical network.
- Provides ARP suppression of broadcast traffic in VXLAN networks.

NSX control plane communication occurs over the management network.

Table 3-50. NSX Controller Design Decision

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-003	Deploy three NSX Controller instances to provide high availability and scale.	The high availability of NSX Controllers reduce the downtime period in case of failure of one physical host.	None.

NSX Virtual Switch

The NSX data plane consists of the NSX virtual switch. This virtual switch is based on the vSphere Distributed Switch (VDS) with additional components to enable rich services. The add-on NSX components include kernel modules (VIBs) which run within the hypervisor kernel and provide services such as distributed logical router (DLR) and distributed firewall (DFW), and VXLAN capabilities.

The NSX virtual switch abstracts the physical network and provides access-level switching in the hypervisor. It is central to network virtualization because it enables logical networks that are independent of physical constructs such as VLAN. Using an NSX virtual switch includes several benefits.

- Supports overlay networking and centralized network configuration. Overlay networking enables the following capabilities.
 - Creation of a flexible logical Layer 2 overlay over existing IP networks on existing physical infrastructure without the need to re-architect the data center networks.
 - Provisioning of communication (east/west and north/south) while maintaining isolation between tenants.
 - Application workloads and virtual machines that are agnostic of the overlay network and operate as if they were connected to a physical Layer 2 network.
- Facilitates massive scale of hypervisors.
- Because the NSX virtual switch is based on VDS, it provides a comprehensive toolkit for traffic management, monitoring, and troubleshooting within a virtual network through features such as port mirroring, NetFlow/IPFIX, configuration backup and restore, network health check, QoS, and more.

Logical Switching

NSX logical switches create logically abstracted segments to which tenant virtual machines can be connected. A single logical switch is mapped to a unique VXLAN segment and is distributed across the ESXi hypervisors within a transport zone. The logical switch allows line-rate switching in the hypervisor without the constraints of VLAN sprawl or spanning tree issues.

Distributed Logical Router

The NSX distributed logical router (DLR) is optimized for forwarding in the virtualized space, that is, forwarding between VMs on VXLAN- or VLAN-backed port groups. DLR has the following characteristics.

- High performance, low overhead first hop routing
- Scales with number of hosts
- Up to 1,000 Logical Interfaces (LIFs) on each DLR

Distributed Logical Router Control Virtual Machine

The distributed logical router control virtual machine is the control plane component of the routing process, providing communication between NSX Manager and the NSX Controller cluster through the User World Agent (UWA). NSX Manager sends logical interface information to the control virtual machine and the NSX Controller cluster, and the control virtual machine sends routing updates to the NSX Controller cluster.

User World Agent

The User World Agent (UWA) is a TCP (SSL) client that facilitates communication between the ESXi hosts and the NSX Controller instances as well as the retrieval of information from the NSX Manager via interaction with the message bus agent.

VXLAN Tunnel Endpoint

VXLAN Tunnel Endpoints (VTEPs) are instantiated within the vSphere Distributed Switch to which the ESXi hosts that are prepared for NSX for vSphere are connected. VTEPs are responsible for encapsulating VXLAN traffic as frames in UDP packets and for the corresponding decapsulation. VTEPs take the form of one or more VMkernel ports with IP addresses and are used both to exchange packets with other VTEPs and to join IP multicast groups via Internet Group Membership Protocol (IGMP). If you use multiple VTEPs, then you must select a teaming method.

Edge Services Gateway

The NSX Edge services gateways (ESGs) primary function is north/south communication, but it also offers support for Layer 2, Layer 3, perimeter firewall, load balancing and other services such as SSL-VPN and DHCP-relay.

Distributed Firewall

NSX includes a distributed kernel-level firewall known as the distributed firewall. Security enforcement is done at the kernel and VM network adapter level. The security enforcement implementation enables firewall rule enforcement in a highly scalable manner without creating bottlenecks on physical appliances. The distributed firewall has minimal CPU overhead and can perform at line rate.

The flow monitoring feature of the distributed firewall displays network activity between virtual machines at the application protocol level. This information can be used to audit network traffic, define and refine firewall policies, and identify botnets.

Logical Load Balancer

The NSX logical load balancer provides load balancing services up to Layer 7, allowing distribution of traffic across multiple servers to achieve optimal resource utilization and availability. The logical load balancer is a service provided by the NSX Edge service gateway.

NSX for vSphere Requirements in ROBO

NSX for vSphere requirements impact both physical and virtual networks.

Physical Network Requirements

Physical requirements determine the MTU size for networks that carry VLAN traffic, dynamic routing support, type synchronization through an NTP server, and forward and reverse DNS resolution.

Requirement	Comments
Any network that carries VXLAN traffic must have an MTU size of 1600 or greater.	VXLAN packets cannot be fragmented. The MTU size must be large enough to support extra encapsulation overhead. This design uses jumbo frames, MTU size of 9000, for VXLAN traffic.
For the hybrid replication mode, Internet Group Management Protocol (IGMP) snooping must be enabled on the Layer 2 switches to which ESXi hosts that participate in VXLAN are attached. IGMP querier must be enabled on the connected router or Layer 3 switch.	IGMP snooping on Layer 2 switches is a requirement of the hybrid replication mode. Hybrid replication mode is the recommended replication mode for broadcast, unknown unicast, and multicast (BUM) traffic when deploying into an environment with large scale-out potential. The traditional requirement for Protocol Independent Multicast (PIM) is removed.
Dynamic routing support on the upstream Layer 3 data center switches must be enabled.	Enable a dynamic routing protocol supported by NSX on the upstream data center switches to establish dynamic routing adjacency with the ESGs.
NTP server must be available.	The NSX Manager requires NTP settings that synchronize it with the rest of the vSphere environment. Drift can cause problems with authentication. The NSX Manager must be in sync with the vCenter Single Sign-On service on the Platform Services Controller.
Forward and reverse DNS resolution for all management VMs must be established.	The NSX Controller nodes do not require DNS entries.

NSX Component Specifications

The following table lists the components involved in the NSX for vSphere solution and the requirements for installing and running them. The compute and storage requirements have been taken into account when sizing resources to support the NSX for vSphere solution.

Note NSX ESG sizing can vary with tenant requirements, so all options are listed.

VM	vCPU	Memory	Storage	Quantity per ROBO instance
NSX Manager	4	16 GB	60 GB	1
NSX Controller	4	4 GB	20 GB	3
NSX ESG	1 (Compact) 2 (Large) 4 (Quad Large) 6 (X-Large)	512 MB (Compact) 1 GB (Large) 1 GB (Quad Large) 8 GB (X-Large)	512 MB 512 MB 512 MB 4.5 GB (X-Large) (+4 GB with swap)	Optional component. Deployment of the NSX ESG varies per use case.
DLR control VM	1	512 MB	512 MB	Optional component. Varies with use case. Typically 2 per HA pair.
Guest introspection	2	1 GB	4 GB	Optional component. 1 per ESXi host.
NSX data security	1	512 MB	6 GB	Optional component. 1 per ESXi host.

NSX Edge Service Gateway Sizing

The Quad Large model is suitable for high performance firewall abilities and the X-Large is suitable for both high performance load balancing and routing.

You can convert between NSX Edge service gateway sizes upon demand using a non-disruptive upgrade process, so the recommendation is to begin with the Large model and scale up if necessary. A Large NSX Edge service gateway is suitable for medium firewall performance but as detailed later, the NSX Edge service gateway does not perform the majority of firewall functions.

Note Edge service gateway throughput is influenced by characteristics such as uplink speed and WAN circuits. An adaptable approach, that is, converting as necessary, is recommended.

Table 3-51. NSX Edge Service Gateway Sizing Design Decision

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-004	Use large size NSX Edge service gateways.	The large size provides all the performance characteristics needed even in the event of a failure. A larger size would also provide the performance required but at the expense of extra resources that wouldn't be used.	None.

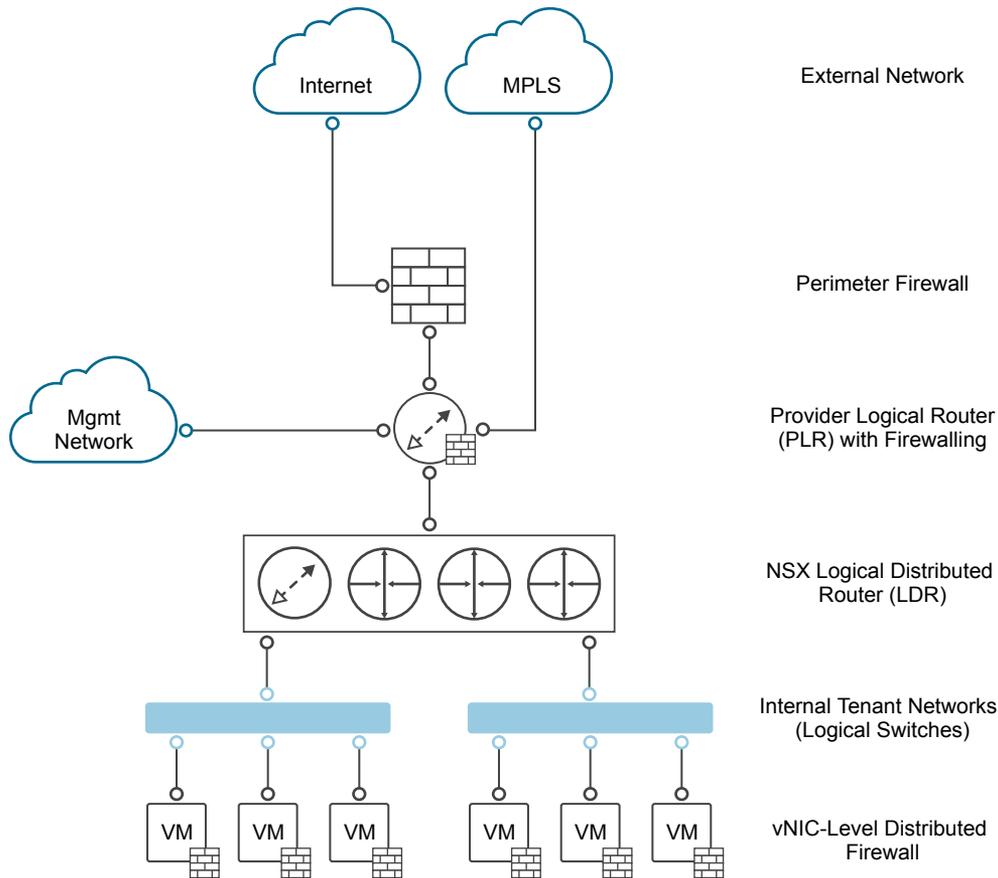
Network Virtualization Conceptual Design in ROBO

This conceptual design provides you with an understanding of the network virtualization design.

The network virtualization conceptual design includes a perimeter firewall, a provider logical router, and the NSX for vSphere Logical Router. It also includes the external network, internal tenant network, and internal non-tenant network.

Note In this document, tenant refers to a tenant of the cloud management platform or to a management application.

Figure 3-12. Conceptual Tenant Overview



The conceptual design has the following key components.

- External Networks** Connectivity to and from external networks is through the perimeter firewall. The main external network is the Internet.
- Perimeter Firewall** The physical firewall exists at the perimeter of the data center. Each tenant receives either a full instance or partition of an instance to filter external traffic.
- Provider Logical Router (PLR)** The PLR exists behind the perimeter firewall and handles north/south traffic that is entering and leaving tenant workloads.
- NSX for vSphere Distributed Logical Router (DLR)** This logical router is optimized for forwarding in the virtualized space, that is, between VMs, on VXLAN port groups or VLAN-backed port groups.

Internal Non-Tenant Network A single management network, which sits behind the perimeter firewall but not behind the PLR. Enables customers to manage the tenant environments.

Internal Tenant Networks Connectivity for the main tenant workload. These networks are connected to a DLR, which sits behind the PLR. These networks take the form of VXLAN-based NSX for vSphere logical switches. Tenant virtual machine workloads will be directly attached to these networks.

Cluster Design for NSX for vSphere in ROBO

Following the vSphere design, the NSX for vSphere design consists of a single consolidated stack providing services for management components and workloads.

Consolidated Stack

In the converted stack, the underlying hosts are prepared for NSX for vSphere. The Consolidated stack has these components.

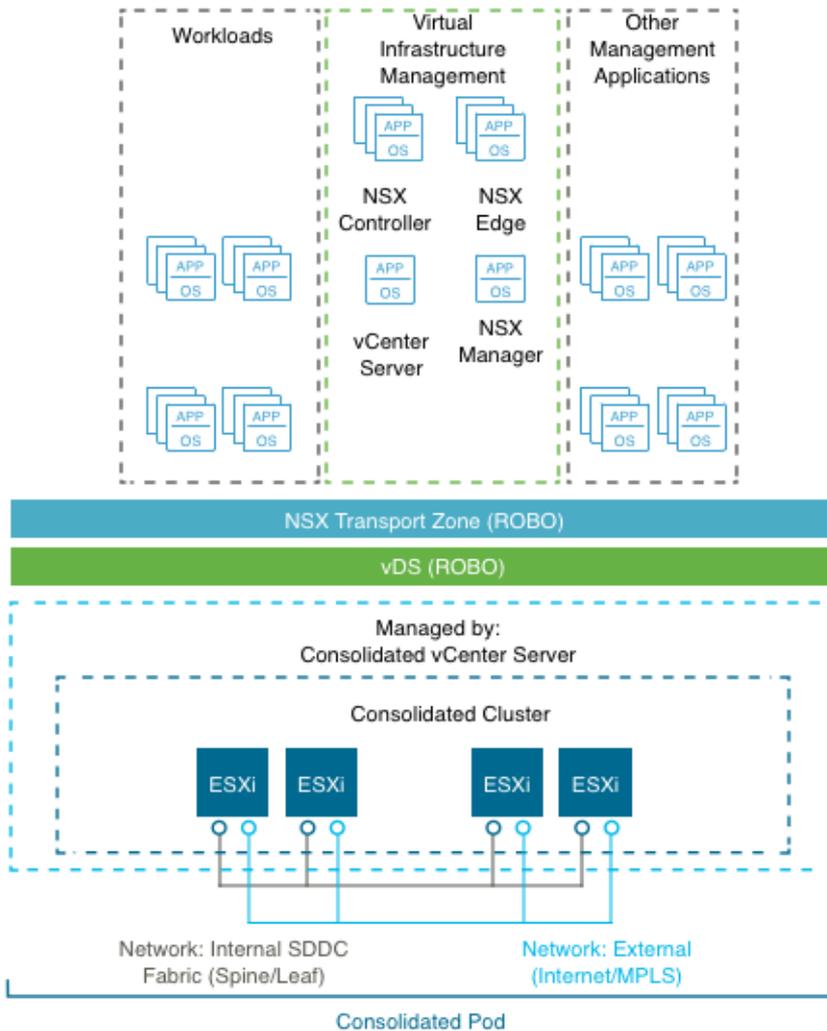
- NSX Manager instance.
- NSX Controller cluster.
- NSX ESG for north/south routing.
- NSX DLR for east/west routing.
- NSX ESG load balancers for workloads, where required.

Table 3-52. NSX for vSphere Cluster Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-005	For the Consolidated stack, do not use a dedicated edge cluster.	Simplifies configuration and minimizes the number of hosts required for initial deployment.	The NSX Controller instances, NSX Edge services gateways, and DLR control VMs are deployed in the converted cluster. The shared nature of the cluster will require the cluster to be scaled out as compute workloads are added so as to not impact network performance.
ROBO-VI-SDN-006	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules to the NSX components.	Using DRS prevents controllers from running on the same ESXi host and thereby risking their high availability capability.	Additional configuration is required to set up anti-affinity rules.

The logical design of NSX considers the vCenter Server clusters and define the place where each NSX component runs.

Figure 3-13. Cluster Design for NSX for vSphere



High Availability of NSX for vSphere Components

vSphere HA protects the NSX Manager instance by ensuring that the NSX Manager VM is restarted on a different host in the event of primary host failure.

The NSX Controller nodes have defined vSphere Distributed Resource Scheduler (DRS) rules to ensure that NSX for vSphere Controller nodes do not run on the same host.

The data plane remains active during outages in the management and control planes although the provisioning and modification of virtual networks is impaired until those planes become available again.

NSX Edge components that are deployed for north/south traffic are configured in equal-cost multi-path (ECMP) mode that supports route failover in seconds. NSX Edge components deployed for load balancing utilize NSX HA. NSX HA provides faster recovery than vSphere HA alone because NSX HA uses an active/passive pair of NSX Edge devices. By default, the passive Edge device becomes active within 15 seconds. All NSX Edge devices are also protected by vSphere HA.

Scalability of NSX Components

A one-to-one mapping between NSX Manager instances and vCenter Server instances exists. If the inventory exceeds the limits supported by a single vCenter Server, then you can deploy a new vCenter Server instance, and must also deploy a new NSX Manager instance. Because ROBO is defined as up to 100 workloads plus the required management components scalability of the products is not an issue.

vSphere Distributed Switch Uplink Configuration in ROBO

Each ESXi host utilizes two physical 10 Gb Ethernet adapters, associated with the uplinks on the vSphere Distributed Switches to which it is connected. Each uplink is connected to a different top-of-rack switch to mitigate the impact of a single top-of-rack switch failure and to provide two paths in and out of the ROBO SDDC.

Table 3-53. VTEP Teaming and Failover Configuration Design Decision

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-007	Set up VXLAN Tunnel Endpoints (VTEPs) to use Route based on SRC-ID for teaming and failover configuration.	Allows for the utilization of the two uplinks of the vDS resulting in better bandwidth utilization and faster recovery from network path failures.	Link aggregation such as LACP between the top-of-rack (ToR) switches and ESXi host must not be configured in order to allow dynamic routing to peer between the ESGs and the upstream switches.

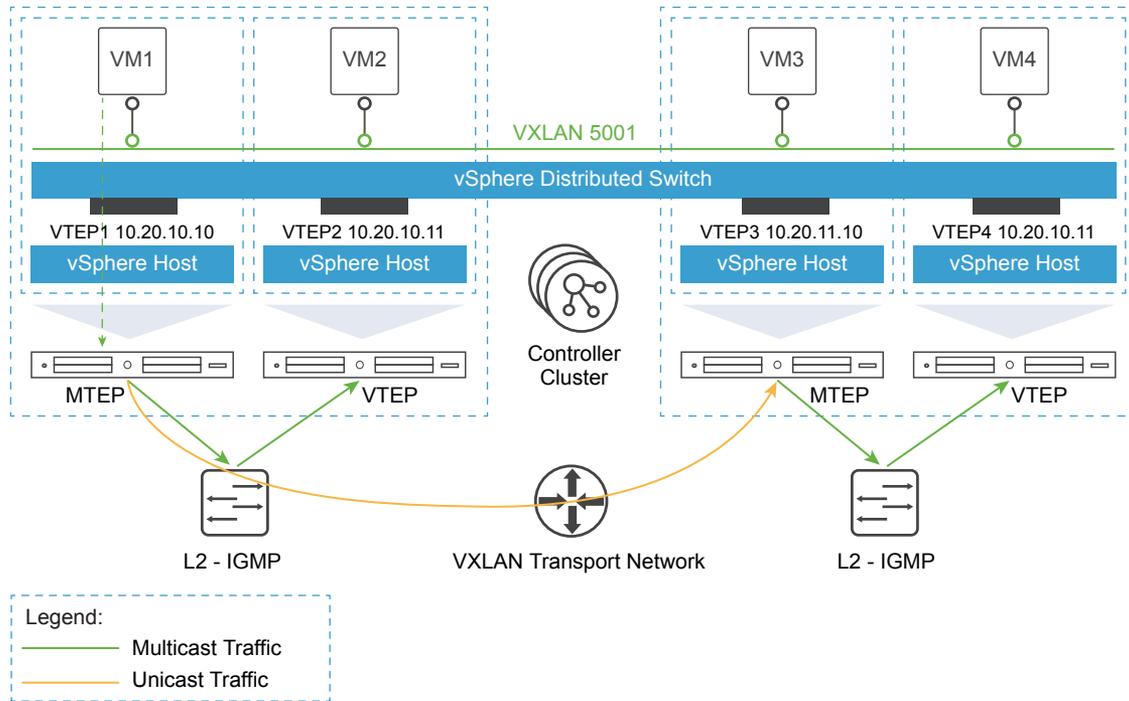
Logical Switch Control Plane Mode Design in ROBO

The control plane decouples NSX for vSphere from the physical network and handles the broadcast, unknown unicast, and multicast (BUM) traffic within the logical switches. The control plane is on top of the transport zone and is inherited by all logical switches that are created within it. It is possible to override aspects of the control plane.

The following options are available.

- Multicast Mode** The control plane uses multicast IP addresses on the physical network. Use multicast mode only when upgrading from existing VXLAN deployments. In this mode, you must configure PIM/IGMP on the physical network.
- Unicast Mode** The control plane is handled by the NSX Controllers and all replication occurs locally on the host. This mode does not require multicast IP addresses or physical network configuration.
- Hybrid Mode** This mode is an optimized version of the unicast mode where local traffic replication for the subnet is offloaded to the physical network. Hybrid mode requires IGMP snooping on the first-hop switch and access to an IGMP querier in each VTEP subnet. Hybrid mode does not require PIM.

Figure 3-14. Logical Switch Control Plane in Hybrid Mode



This design uses hybrid mode for control plane replication.

Table 3-54. Logical Switch Control Plane Mode Design Decision

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-008	Use hybrid mode for control plane replication.	Offloading replication processing to the physical network reduces pressure on VTEPs. In high consolidation ratio environments, hybrid mode is preferable to unicast mode as it saves CPU cycles. Multicast mode is used only when migrating from existing VXLAN solutions.	IGMP snooping must be enabled on the ToR physical switch and an IGMP querier must be available.

Transport Zone Design in ROBO

A transport zone is used to define the scope of a VXLAN overlay network and can span one or more clusters within one vCenter Server domain. One or more transport zones can be configured in an NSX for vSphere solution. A transport zone is not meant to delineate a security boundary.

Table 3-55. Transport Zones Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-009	Use a single global transport zone.	A single transport zone localizes networks and security policies within a ROBO.	None.

Routing Design in ROBO

The routing design considers different levels of routing within the environment from which to define a set of principles for designing a scalable routing solution.

North/south The Provider Logical Router (PLR) handles the north/south traffic to and from a tenant and management applications inside of application virtual networks.

East/west Internal east/west routing at the layer beneath the PLR deals with the application workloads.

Table 3-56. Routing Model Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-011	Deploy NSX Edge Services Gateways in an ECMP configuration for north/south routing.	The NSX ESG is the recommended device for managing north/south traffic. Using ECMP provides multiple paths in and out of the ROBO SDDC. This results in faster failover times than deploying Edge service gateways in HA mode.	ECMP requires 2 VLANS for uplinks which adds an additional VLAN over traditional HA ESG configurations.
ROBO-VI-SDN-012	Deploy a single NSX DLR in HA mode to provide east/west routing.	Using the DLR reduces the hop count between nodes attached to it to 1. This reduces latency and improves performance.	DLRs are limited to 1,000 logical interfaces. When that limit is reached, a new DLR must be deployed.
ROBO-VI-SDN-013	Use BGP as the dynamic routing protocol inside the ROBO SDDC.	Using BGP as opposed to OSPF eases the implementation of dynamic routing. There is no need to plan and design access to OSPF area 0 inside the ROBO SDDC. OSPF area 0 varies based on customer configuration.	BGP requires configuring each ESG and DLR with the remote router that it exchanges routes with.
ROBO-VI-SDN-014	Configure BGP Keep Alive Timer to 1 and Hold Down Timer to 3 between the DLR and all ESGs that provide north/south routing.	With Keep Alive and Hold Timers between the UDLR and ECMP ESGs set low, a failure is detected quicker, and the routing table is updated faster.	If an ESXi host becomes resource constrained, the ESG running on that host might no longer be used even though it is still up.
ROBO-VI-SDN-015	Configure BGP Keep Alive Timer to 4 and Hold Down Timer to 12 between the ToR switches and all ESGs providing north/south routing.	This provides a good balance between failure detection between the ToRs and the ESGs and overburdening the ToRs with keep alive traffic.	By using longer timers to detect when a router is dead, a dead router stays in the routing table for a longer period of time, and continues to send traffic to a dead router.
ROBO-VI-SDN-016	Create one or more static routes on ECMP enabled edges for subnets behind the DLR with a higher administrative cost than the dynamically learned routes.	When the DLR control VM fails over router adjacency is lost and routes from upstream devices such as ToR's to subnets behind the UDLR are lost.	This requires that each ECMP edge device be configured with static routes to the DLR. If any new subnets are added behind the DLR the routes must be updated on the ECMP edges.

Transit Network and Dynamic Routing

Dedicated networks are needed to facilitate traffic between the distributed logical routers and edge gateways, and to facilitate traffic between edge gateways and the top of rack switches. These networks are used for exchanging routing tables and for carrying transit traffic.

Table 3-57. Transit Network Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-017	Create a virtual switch for use as the transit network between the DLR and ESG's. The DLR provides east/west routing in the stack while the ESG's provide north/south routing.	The virtual switch allows the DLR and ESGs to exchange routing information.	A virtual switch for use as a transit network is required.
ROBO-VI-SDN-018	Create two VLANs. Use those VLANs to enable ECMP between the north/south ESGs and the ToR switches. Each ToR has an SVI on one of the two VLANs and each north/south ESG has an interface on both VLANs.	This enables the ESGs to have multiple equal-cost routes which provides more resiliency and better bandwidth utilization in the network.	Extra VLANs are required.

Firewall Logical Design in ROBO

The NSX Distributed Firewall is used to protect all management applications attached to application virtual networks. To secure the ROBO SDDC, only other solutions in the ROBO SDDC and approved administration IPs can directly communicate with individual components. External facing portals are accessible via a load balancer virtual IP (VIP). This simplifies the design by having a single point of administration for all firewall rules. The firewall on individual ESGs is set to allow all traffic. An exception are ESGs that provide ECMP services, which require the firewall to be disabled.

Table 3-58. Firewall Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-019	For all ESGs deployed as load balancers, set the default firewall rule to allow all traffic.	Restricting and granting access is handled by the distributed firewall. The default firewall rule does not have to do it.	Explicit rules to allow or deny access to management applications and tenant workloads must be defined in the distributed firewall.
ROBO-VI-SDN-020	For all ESGs deployed as ECMP north/south routers, disable the firewall.	Use of ECMP on the ESGs is a requirement. Leaving the firewall enabled, even in allow all traffic mode, results in sporadic network connectivity.	Services such as NAT and load balancing cannot be used when the firewall is disabled.
ROBO-VI-SDN-021	Configure the Distributed Firewall to limit access to administrative interfaces on the management virtual applications.	To ensure only authorized administrators can access the administrative interfaces of management applications.	Maintaining firewall rules adds administrative overhead.

Load Balancer Design in ROBO

The ESG implements load balancing within NSX for vSphere.

The ESG has both a Layer 4 and a Layer 7 engine that offer different features, which are summarized in the following table.

Feature	Layer 4 Engine	Layer 7 Engine
Protocols	TCP	TCP HTTP HTTPS (SSL Pass-through) HTTPS (SSL Offload)
Load balancing method	Round Robin Source IP Hash Least Connection	Round Robin Source IP Hash Least Connection URI
Health checks	TCP	TCP HTTP (GET, OPTION, POST) HTTPS (GET, OPTION, POST)
Persistence (keeping client connections to the same back-end server)	TCP: SourceIP	TCP: SourceIP, MSRDP HTTP: SourceIP, Cookie HTTPS: SourceIP, Cookie, ssl_session_id
Connection throttling	No	Client Side: Maximum concurrent connections, Maximum new connections per second Server Side: Maximum concurrent connections
High availability	Yes	Yes
Monitoring	View VIP (Virtual IP), Pool and Server objects and stats via CLI and API View global stats for VIP sessions from the vSphere Web Client	View VIP, Pool and Server objects and statistics by using CLI and API View global statistics about VIP sessions from the vSphere Web Client
Layer 7 manipulation	No	URL block, URL rewrite, content rewrite

Table 3-59. NSX for vSphere Load Balancer Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-022	Use the NSX load balancer in HA mode.	The NSX load balancer can support the needs of the workload applications. Using another load balancer would increase cost and add another component to be managed as part of the ROBO SDDC.	None.

Bridging Physical Workloads in ROBO

NSX for vSphere offers VXLAN to Layer 2 VLAN bridging capabilities with the data path contained entirely in the ESXi hypervisor. The bridge runs on the ESXi host where the DLR control VM is located. Multiple bridges per DLR are supported.

Table 3-60. Virtual to Physical Interface Type Design Decision

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-023	Place all virtual machines, both management and tenant, on VXLAN-backed networks unless you must satisfy an explicit requirement to use VLAN-backed port groups for these virtual machines. If VLAN-backed port groups are required, connect physical workloads that need to communicate to virtualized workloads to routed VLAN LIFs on a DLR.	Bridging and routing are not possible on the same logical switch. As a result, it makes sense to attach a VLAN LIF to a distributed router or ESG and route between the physical and virtual machines. Use bridging only where virtual machines need access only to the physical machines on the same Layer 2.	Access to physical workloads is routed via the DLR or ESG.

Remote Connectivity in ROBO

ROBO sites must be connected to their hub. Connection types could be point-to-point links, MPLS, VPN Tunnels, and other appropriate connection types. The connection types will vary by customer, and is out of scope for this design.

The remote connectivity design must ensure latency is less than 150 ms.

Table 3-61. Remote Connectivity Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-024	Provide a connection between each ROBO site and the hub that is capable of routing.	The vRealize Operations remote collectors, vRealize Log Insight, and vRealize Automation Proxy Agents require connectivity back to the hub.	For provisioning and single pane of glass monitoring the connection between the hub and ROBO site must be up.
ROBO-VI-SDN-025	Ensure that the latency between the ROBO sites and the hub is less than 150 ms.	A latency below 150 ms is required to ensure successful provisioning of workloads in the ROBO site.	None.

Application Virtual Network in ROBO

Management applications, such as VMware vRealize Automation Proxy Agents, VMware vRealize Operations Remote Collectors, and vRealize Log Insight, reside on Application Virtual Networks.

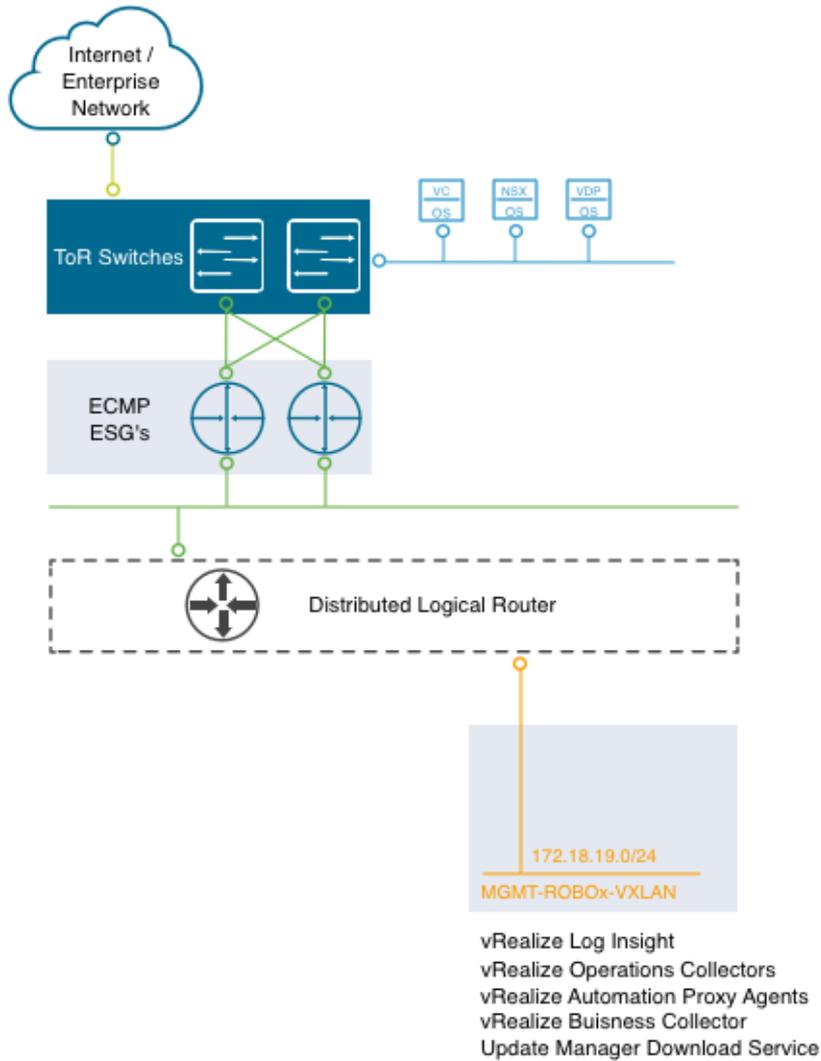
Table 3-62. Isolated Management Applications Design Decisions

Decision ID	Design Decision	Design Justification	Design Implications
ROBO-VI-SDN-026	Place the following management applications on an application virtual network. <ul style="list-style-type: none"> ■ vRealize Automation Proxy Agents ■ vRealize Business collectors ■ vRealize Operations Manager remote collectors ■ vRealize Log Insight ■ Update Manager Download Service 	Access to the management applications is only through published access points.	Direct access to application virtual networks is controlled by distributed firewall rules.
ROBO-VI-SDN-027	Create a single application virtual network.	Using only a single application virtual network simplifies the design by placing all management components on the same virtual wire.	A single /24 subnet is used for the application virtual network. IP management becomes critical to ensure no shortage of IP addresses will appear in the future.

Having software-defined networking based on NSX in the management stack makes all NSX features available to the management applications.

This approach to network virtualization service design improves security and mobility of the management applications, and reduces the integration effort with existing customer networks.

Figure 3-15. Virtual Application Network Components and Design



Certain configuration choices might later facilitate the tenant onboarding process.

- Create the primary NSX ESG to act as the tenant PLR and the logical switch that forms the transit network for use in connecting to the DLR.
- Connect the primary NSX ESG uplinks to the external networks
- Connect the primary NSX ESG internal interface to the transit network.
- Create the NSX DLR to provide routing capabilities for tenant internal networks and connect the DLR uplink to the transit network.

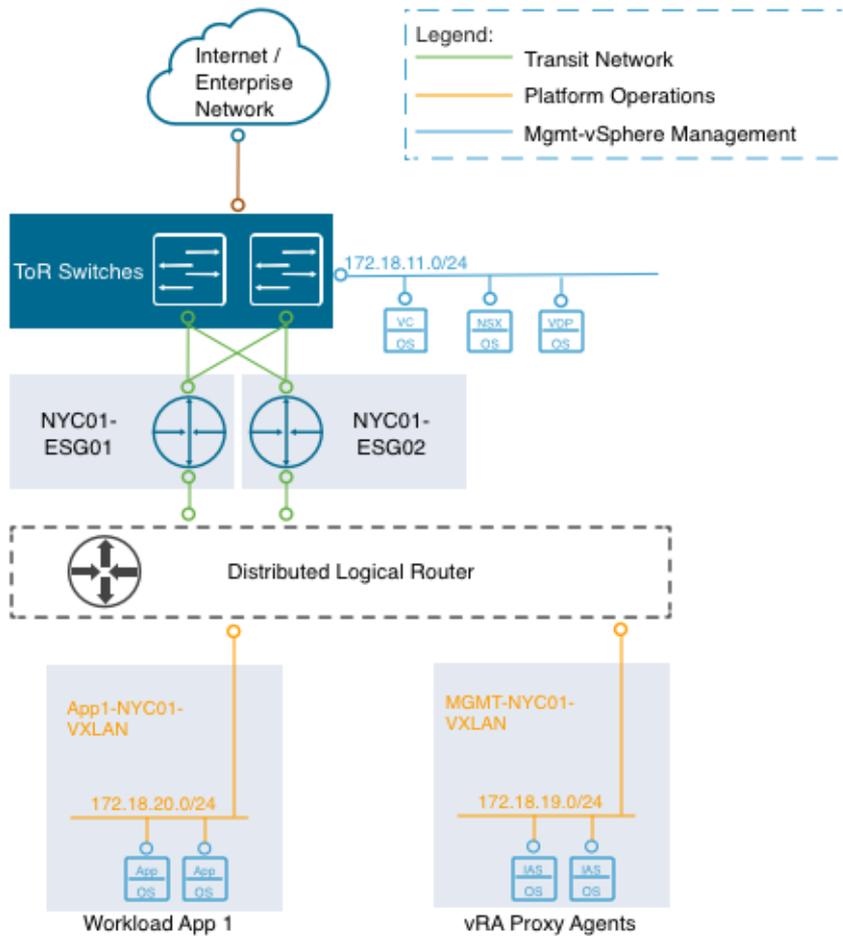
- Create any tenant networks that are known up front and connect them to the DLR.

Virtual Network Design Example in ROBO

The virtual network design example illustrates an implementation for a management application virtual network.

The following figure illustrates how to implement a management application virtual network and a workload virtual network. The example shows vRealize Automation Proxy Agents and a Workload App, but other applications would be implemented similarly.

Figure 3-16. Detailed Example for ROBO Networking



The example is set up as follows.

- You deploy the vRealize Automation Proxy Agents on the application virtual network. This network is provided by a VXLAN virtual wire (orange network).

- The network that is used by vRealize Automation connects to external networks through NSX for vSphere. NSX ESGs and the DLR route traffic between the application virtual networks and the public network.
- Services such as a Web GUI, which must be available to the end users of vRealize Automation, are accessible via the NSX Edge load balancer.

The following table shows an example of a mapping from application virtual networks to IPv4 subnets. The actual mapping depends on the customer environment and is based on available IP subnets.

Note The following IP ranges are an example. Your actual implementation depends on your environment.

Application Virtual Network	Applications	Internal IPv4 Subnet
MGMT-NYC01-VXLAN	vRealize Automation Proxy Agents vRealize Business Collector vRealize Operations Remote Collectors vRealize Log Insight cluster Update Manager Download Service	172.18.19.0/24
App1-NYC01-VXLAN	User workload app 1	172.18.20.0/24
App2-NYC01-VXLAN	User workload app2	172.18.21.0/24

Use of Secure Sockets Layer Certificates in ROBO

By default, NSX Manager uses a self-signed SSL certificate. This certificate is not trusted by end-user devices or browsers. It is a security best practice to replace these certificates with certificates that are signed by a third-party or enterprise Certificate Authority (CA).

Design ID	Design Decision	Design Justification	Design Implication
ROBO-VI-SDN-028	Replace the NSX Manager certificate with a certificate signed by a 3rd party Public Key Infrastructure.	Ensures communication between NSX admins and the NSX Manager are encrypted by a trusted certificate.	Replacing and managing certificates is an operational overhead.

Shared Storage Design in ROBO

The shared storage design includes design decisions for vSAN and secondary storage.

Well-designed shared storage provides the basis for the ROBO SDDC and has the following benefits.

- Prevents unauthorized access to business data
- Protects data from hardware and software failures
- Protects data from malicious or accidental corruption

Follow these guidelines when designing shared storage for your environment.

- Optimize the storage design to meet the diverse needs of applications, services, administrators, and users.

- Strategically align business applications and the storage infrastructure to reduce costs, boost performance, improve availability, provide security, and enhance functionality.
- Provide multiple tiers of storage to match application data access to application requirements.
- Design each tier of storage with different performance, capacity, and availability characteristics. Because not every application requires expensive, high-performance, highly available storage, designing different storage tiers reduces cost.

Shared Storage Platform in ROBO

You can choose between traditional storage, VMware vSphere Virtual Volumes, and vSAN storage.

Storage Types

Traditional Storage	Fibre Channel, NFS, and iSCSI are mature and viable options to support virtual machine needs.
VMware vSAN Storage	vSAN is a software-based distributed storage platform that combines the compute and storage resources of VMware ESXi hosts. When you design and size a vSAN cluster, hardware choices are more limited than for traditional storage.
VMware vSphere Virtual Volumes	This design does not leverage VMware vSphere Virtual Volumes.

Traditional Storage and vSAN Storage

Fibre Channel, NFS, and iSCSI are mature and viable options to support virtual machine needs.

Your decision to implement one technology or another can be based on performance and functionality, and on considerations like the following:

- The organization's current in-house expertise and installation base
- The cost, including both capital and long-term operational expenses
- The organization's current relationship with a storage vendor

vSAN is a software-based distributed storage platform that combines the compute and storage resources of ESXi hosts. It provides a simple storage management experience for the user. This solution makes software-defined storage a reality for VMware customers. However, you must carefully consider supported hardware options when sizing and designing a vSAN cluster.

Storage Type Comparison

ESXi hosts support a variety of storage types. Each storage type supports different vSphere features.

Table 3-63. Network Shared Storage Supported by ESXi Hosts

Technology	Protocols	Transfers	Interface
Fibre Channel	FC/SCSI	Block access of data/LUN	Fibre Channel HBA
Fibre Channel over Ethernet	FCoE/SCSI	Block access of data/LUN	Converged network adapter (hardware FCoE) NIC with FCoE support (software FCoE)
iSCSI	IP/SCSI	Block access of data/LUN	iSCSI HBA or iSCSI enabled NIC (hardware iSCSI) Network Adapter (software iSCSI)
NAS	IP/NFS	File (no direct LUN access)	Network adapter
vSAN	IP	Block access of data	Network adapter

Table 3-64. vSphere Features Supported by Storage Type

Type	vSphere vMotion	Datastore	Raw Device Mapping (RDM)	Application or Block-level Clustering	HA/DRS	Storage APIs Data Protection
Local Storage	Yes	VMFS	No	Yes	No	Yes
Fibre Channel / Fibre Channel over Ethernet	Yes	VMFS	Yes	Yes	Yes	Yes
iSCSI	Yes	VMFS	Yes	Yes	Yes	Yes
NAS over NFS	Yes	NFS	No	No	Yes	Yes
vSAN	Yes	vSAN	No	Yes (via iSCSI Initiator)	Yes	Yes

Shared Storage Logical Design in ROBO

The shared storage design selects the appropriate storage device for a ROBO.

The storage devices for use by each type of cluster are as follows.

- Consolidated clusters use vSAN for primary storage and another technology for secondary storage.

Figure 3-17. Logical Storage Design

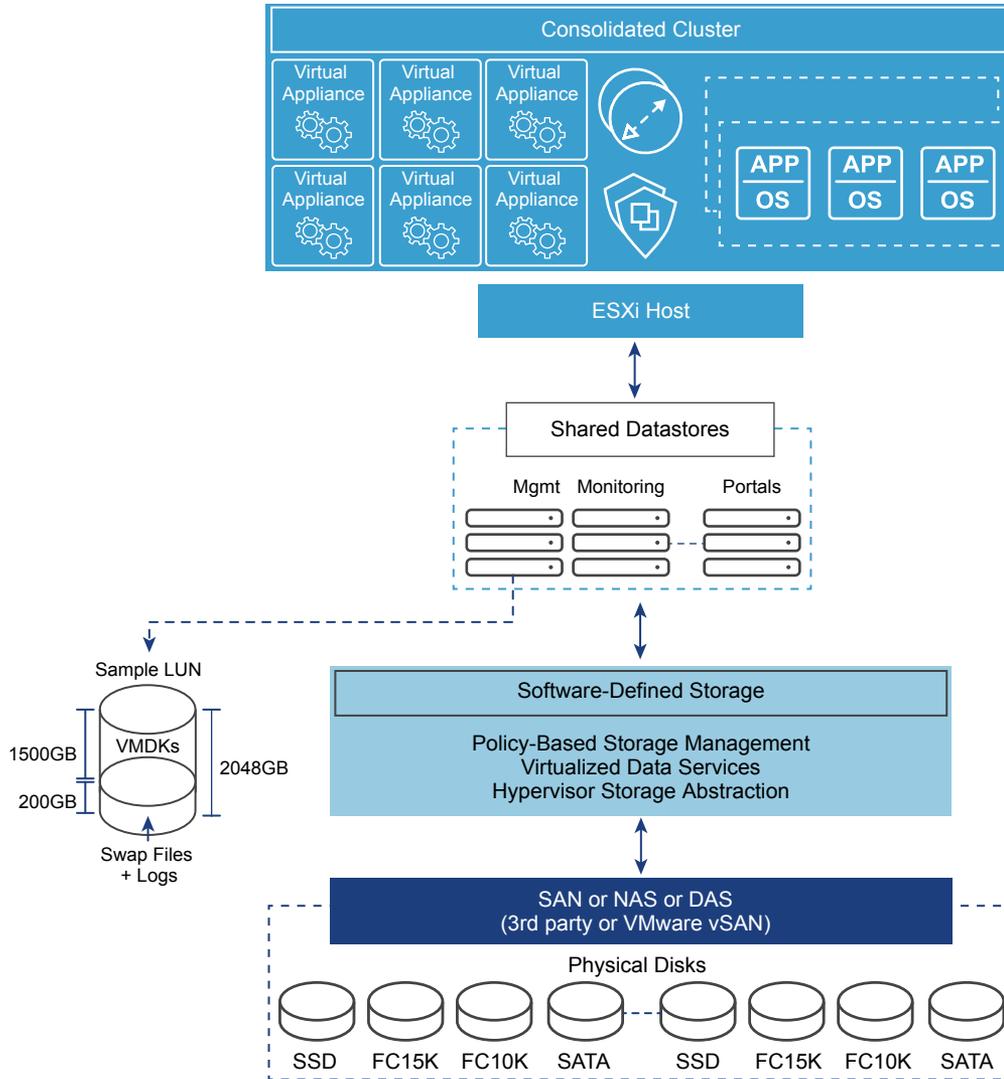


Table 3-65. Storage Type Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-001	<p>In the Consolidated cluster, use vSAN and secondary shared storage:</p> <ul style="list-style-type: none"> Use vSAN as the primary shared storage platform. Use a secondary shared storage platform for backup data. 	<p>vSAN as the primary shared storage solution can take advantage of more cost-effective local storage.</p> <p>Secondary storage is used primarily for archival and the need to maintain historical data.</p>	<p>The use of two different storage technologies increases the complexity and operational overhead.</p>
ROBO-VI-STORAGE-002	<p>Ensure that at least 20% of free space is always available on all non-vSAN datastores.</p>	<p>If the datastore runs out of free space, applications and services within the ROBO SDDC, such as backups will fail. To prevent this, maintain adequate free space.</p>	<p>Monitoring and capacity management are critical, and must be proactively performed.</p>

Storage Tiering in ROBO

Today's enterprise-class storage arrays contain multiple drive types and protection mechanisms. The storage, server, and application administrators face challenges when selecting the correct storage configuration for each application being deployed in the environment. Virtualization can make this problem more challenging by consolidating many different application workloads onto a small number of large devices. Given this challenge, administrators might use single storage type for every type of workload without regard to the needs of the particular workload. However, not all application workloads have the same requirements, and storage tiering allows for these differences by creating multiple levels of storage with varying degrees of performance, reliability and cost, depending on the application workload needs.

The most mission-critical data typically represents the smallest amount of data and offline data represents the largest amount. Details differ for different organizations.

To determine the storage tier for application data, determine the storage characteristics of the application or service.

- I/O operations per second (IOPS) requirements
- Megabytes per second (MBps) requirements
- Capacity requirements
- Availability requirements
- Latency requirements

After you determine the information for each application, you can move the application to the storage tier with matching characteristics.

- Consider any existing service-level agreements (SLAs).
- Move data between storage tiers during the application life cycle as needed.

VMware Hardware Acceleration API/CLI for Storage in ROBO

The VMware Hardware Acceleration API/CLI for storage (previously known as vStorage APIs for Array Integration or VAAI), supports a set of ESXCLI commands for enabling communication between ESXi hosts and storage devices. The APIs define a set of storage primitives that enable the ESXi host to offload certain storage operations to the array. Offloading the operations reduces resource overhead on the ESXi hosts and can significantly improve performance for storage-intensive operations such as storage cloning, zeroing, and so on. The goal of hardware acceleration is to help storage vendors provide hardware assistance to speed up VMware I/O operations that are more efficiently accomplished in the storage hardware.

Without the use of VAAI, cloning or migration of virtual machines by the VMkernel data mover involves software data movement. The data mover issues I/O to read and write blocks to and from the source and destination datastores. With VAAI, the data mover can use the API primitives to offload operations to the array when possible. For example, when you copy a virtual machine disk file (VMDK file) from one

datastore to another inside the same array, the data mover directs the array to make the copy completely inside the array. If you invoke a data movement operation and the corresponding hardware offload operation is enabled, the data mover first attempts to use hardware offload. If the hardware offload operation fails, the data mover reverts to the traditional software method of data movement.

In nearly all cases, hardware data movement performs significantly better than software data movement. It consumes fewer CPU cycles and less bandwidth on the storage fabric. Timing operations that use the VAAI primitives and use `esxtop` to track values such as CMDS/s, READS/s, WRITES/s, MBREAD/s, and MBWRTN/s of storage adapters during the operation show performance improvements.

Table 3-66. vStorage APIs for Array Integration Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-003	When using on premise secondary storage select an array that supports VAAI.	VAAI offloads tasks to the array itself, enabling the ESXi hypervisor to use its resources for application workloads and not become a bottleneck in the storage subsystem. VAAI is required to support the desired number of virtual machine lifecycle operations.	Not all VAAI arrays support VAAI over all protocols.

Virtual Machine Storage Policies in ROBO

You can create a storage policy for a virtual machine to specify which storage capabilities and characteristics are the best match for this virtual machine.

Note vSAN uses storage policies to allow specification of the characteristics of virtual machines, so you can define the policy on an individual disk level rather than at the volume level for vSAN.

You can identify the storage subsystem capabilities by using the VMware vSphere API for Storage Awareness or by using a user-defined storage policy.

VMware vSphere API for Storage Awareness (VASA) With vSphere API for Storage Awareness, storage vendors can publish the capabilities of their storage to VMware vCenter Server, which can display these capabilities in its user interface.

User-defined storage policy Defined by using the VMware Storage Policy SDK or VMware vSphere PowerCL, or from the vSphere Web Client.

You can assign a storage policy to a virtual machine and periodically check for compliance so that the virtual machine continues to run on storage with the correct performance and availability characteristics.

You can associate a virtual machine with a virtual machine storage policy when you create, clone, or migrate that virtual machine. If a virtual machine is associated with a storage policy, the vSphere Web Client shows the datastores that are compatible with the policy. You can select a datastore or datastore cluster. If you select a datastore that does not match the virtual machine storage policy, the vSphere Web Client shows that the virtual machine is using non-compliant storage. See *Creating and Managing vSphere Storage Policies* in the vSphere 6.5 documentation.

Table 3-67. Virtual Machine Storage Policy Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-004	Use the default vSAN storage policy for management virtual machines in the Consolidated cluster.	The default vSAN storage policy is adequate for the management virtual machines.	If workload virtual machines have different storage requirements, additional VM storage policies may be required.

vSphere Storage I/O Control Design in ROBO

VMware vSphere Storage I/O Control allows cluster-wide storage I/O prioritization, which results in better workload consolidation and helps reduce extra costs associated with over provisioning.

vSphere Storage I/O Control extends the constructs of shares and limits to storage I/O resources. You can control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, so that more important virtual machines get preference over less important virtual machines for I/O resource allocation.

When vSphere Storage I/O Control is enabled on a datastore, the ESXi host monitors the device latency when communicating with that datastore. When device latency exceeds a threshold, the datastore is considered to be congested and each virtual machine that accesses that datastore is allocated I/O resources in proportion to their shares. Shares are set on a per-virtual machine basis and can be adjusted.

vSphere Storage I/O Control has several requirements, limitations, and constraints.

- Datastores enabled with vSphere Storage I/O Control must be managed by a single vCenter Server system.
- Storage I/O Control is supported on Fibre Channel-connected, iSCSI-connected, and NFS-connected storage. RDM is not supported.
- Storage I/O Control does not support datastores with multiple extents.
- Before using vSphere Storage I/O Control on datastores with arrays using automated storage tiering capabilities, verify that the storage array has been certified as compatible with vSphere Storage I/O Control. See the *VMware Compatibility Guide*

Table 3-68. Storage I/O Control Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-006	Enable Storage I/O Control with the default values on all non vSAN datastores.	Storage I/O Control ensures that all virtual machines on a datastore receive an equal amount of I/O.	Virtual machines that use more I/O are throttled to allow other virtual machines access to the datastore only when contention occurs on the datastore.

Datastore Cluster Design in ROBO

A datastore cluster is a collection of datastores with shared resources and a shared management interface. Datastore clusters are to datastores what clusters are to ESXi hosts. After you create a datastore cluster, you can use vSphere Storage DRS to manage storage resources.

vSphere datastore clusters group similar datastores into a pool of storage resources. When vSphere Storage DRS is enabled on a datastore cluster, vSphere automates the process of initial virtual machine file placement and balances storage resources across the cluster to avoid bottlenecks. vSphere Storage DRS considers datastore space usage and I/O load when making migration recommendations.

When you add a datastore to a datastore cluster, the datastore's resources become part of the datastore cluster's resources. The following resource management capabilities are also available for each datastore cluster.

Capability	Description
Space utilization load balancing	You can set a threshold for space use. When space use on a datastore exceeds the threshold, vSphere Storage DRS generates recommendations or performs migrations with vSphere Storage vMotion to balance space use across the datastore cluster.
I/O latency load balancing	You can configure the I/O latency threshold to avoid bottlenecks. When I/O latency on a datastore exceeds the threshold, vSphere Storage DRS generates recommendations or performs vSphere Storage vMotion migrations to help alleviate high I/O load.
Anti-affinity rules	You can configure anti-affinity rules for virtual machine disks to ensure that the virtual disks of a virtual machine are kept on different datastores. By default, all virtual disks for a virtual machine are placed on the same datastore.

You can enable vSphere Storage I/O Control or vSphere Storage DRS for a datastore cluster. You can enable the two features separately, even though vSphere Storage I/O control is enabled by default when you enable vSphere Storage DRS.

vSphere Storage DRS Background Information

vSphere Storage DRS supports automating the management of datastores based on latency and storage utilization. When configuring vSphere Storage DRS, verify that all datastores use the same version of VMFS and are on the same storage subsystem. Because vSphere Storage vMotion performs the migration of the virtual machines, confirm that all prerequisites are met.

vSphere Storage DRS provides a way of balancing usage and IOPS among datastores in a storage cluster:

- Initial placement of virtual machines is based on storage capacity.
- vSphere Storage DRS uses vSphere Storage vMotion to migrate virtual machines based on storage capacity.
- vSphere Storage DRS uses vSphere Storage vMotion to migrate virtual machines based on I/O latency.
- You can configure vSphere Storage DRS to run in either manual mode or in fully automated mode.

vSphere vStorage I/O Control and vSphere Storage DRS manage latency differently.

- vSphere Storage I/O Control distributes the resources based on virtual disk share value after a latency threshold is reached.

- vSphere Storage DRS measures latency over a period of time. If the latency threshold of vSphere Storage DRS is met in that time frame, vSphere Storage DRS migrates virtual machines to balance latency across the datastores that are part of the cluster.

When making a vSphere Storage design decision, consider these points:

- Use vSphere Storage DRS where possible.
- vSphere Storage DRS provides a way of balancing usage and IOPS among datastores in a storage cluster:
 - Initial placement of virtual machines is based on storage capacity.
 - vSphere Storage vMotion is used to migrate virtual machines based on storage capacity.
 - vSphere Storage vMotion is used to migrate virtual machines based on I/O latency.
 - vSphere Storage DRS can be configured in either manual or fully automated modes

vSAN Storage Design in ROBO

VMware vSAN Storage design in this VMware Validated Design includes conceptual design, logical design, network design, cluster and disk group design, and policy design.

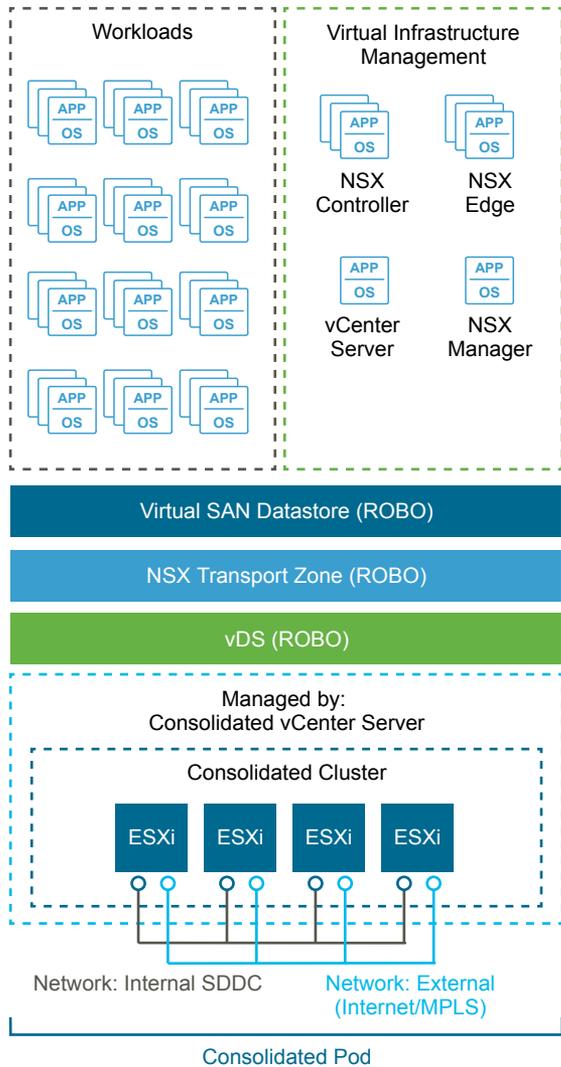
VMware vSAN Conceptual Design and Logical Design in ROBO

The design uses the default storage policy to achieve redundancy and performance within the cluster.

vSAN Conceptual Design

vSAN is used as primary storage for the consolidated cluster.

Figure 3-18. Conceptual vSAN Design



vSAN Logical Design

In a cluster that is managed by vCenter Server, you can manage software-defined storage resources just as you can manage compute resources. Instead of CPU or memory reservations, limits, and shares, you can define storage policies and assign them to virtual machines. The policies specify the characteristics of the storage and can be changed as business requirements change.

vSAN Network Design in ROBO

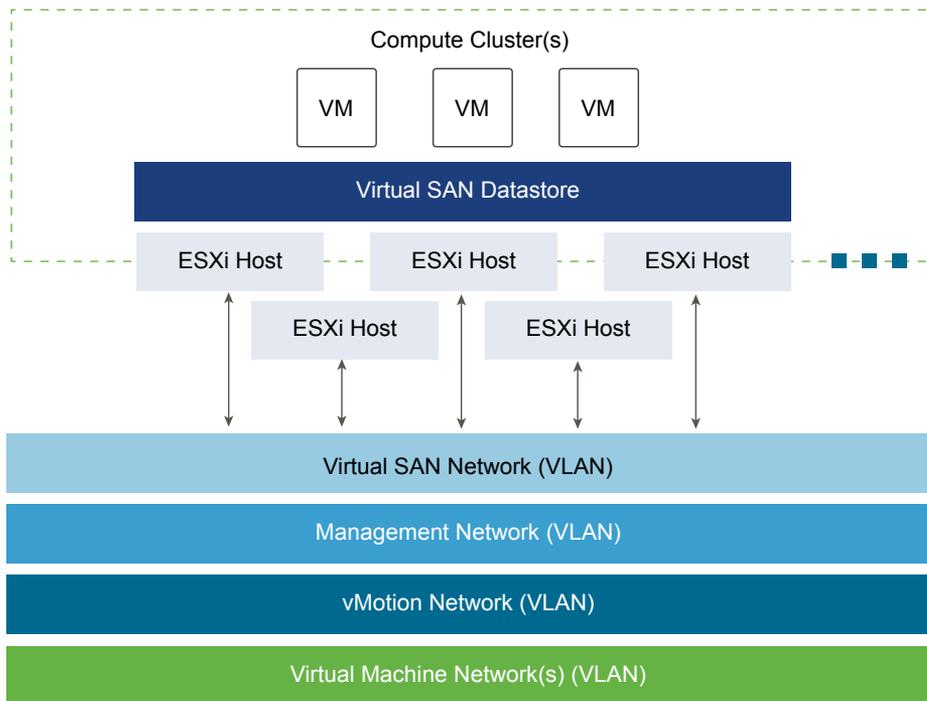
When performing network configuration, you have to consider the traffic and decide how to isolate vSAN traffic.

- Consider how much replication and communication traffic is running between hosts. With VMware vSAN, the amount of traffic depends on the number of VMs that are running in the cluster, and on how write-intensive the I/O is for the applications running in the VMs.
- Isolate vSAN traffic on its own Layer 2 network segment. You can do this with dedicated switches or ports, or by using a VLAN.

The vSAN VMkernel port group is created as part of cluster creation. Configure this port group on all hosts in a cluster, even for hosts that are not contributing storage resources to the cluster.

The following diagram illustrates the logical design of the network.

Figure 3-19. VMware vSAN Conceptual Network



Network Bandwidth Requirements

VMware recommends that solutions use a 10 Gb Ethernet connection for use with vSAN to ensure the best and most predictable performance (IOPS) for the environment. Without it, a significant decrease in array performance results.

Note vSAN all-flash configurations are supported only with 10 GbE.

Table 3-69. Network Speed Selection

Design Quality	1	10	Comments
	Gb	Gb	
Availability	o	o	Neither design option impacts availability.
Manageability	o	o	Neither design option impacts manageability.
Performance	↓	↑	Faster network speeds increase vSAN performance (especially in I/O intensive situations).
Recoverability	↓	↑	Faster network speeds increase the performance of rebuilds and synchronizations in the environment. This ensures that VMs are properly protected from failures.
Security	o	o	Neither design option impacts security.

Legend: ↑ = positive impact on quality; ↓ = negative impact on quality; o = no impact on quality.

Table 3-70. Network Bandwidth Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-001	Use only 10 GbE for VMware vSAN traffic.	Performance with 10 GbE is optimal. Without it, a significant decrease in array performance results.	The physical network must support 10 Gb networking between every host in the vSAN clusters.

VMware vSAN Virtual Switch Type

vSAN supports the use of vSphere Standard Switch or vSphere Distributed Switch. The benefit of using vSphere Distributed Switch is that it supports Network I/O Control which allows for prioritization of bandwidth in case of contention in an environment.

This design uses a vSphere Distributed Switch for the vSAN port group to ensure that priority can be assigned using Network I/O Control to separate and guarantee the bandwidth for vSAN traffic.

Virtual Switch Design Background

Virtual switch type affects performance and security of the environment.

Table 3-71. Virtual Switch Types

Design Quality	vSphere Standard Switch	vSphere Distributed Switch	Comments
Availability	o	o	Neither design option impacts availability.
Manageability	↓	↑	The vSphere Distributed Switch is centrally managed across all hosts, unlike the standard switch which is managed on each host individually.
Performance	↓	↑	The vSphere Distributed Switch has added controls, such as Network I/O Control, which you can use to guarantee performance for vSAN traffic.
Recoverability	↓	↑	The vSphere Distributed Switch configuration can be backed up and restored, the standard switch does not have this functionality.
Security	↓	↑	The vSphere Distributed Switch has added built-in security controls to help protect traffic.

Legend: ↑ = positive impact on quality; ↓ = negative impact on quality; o = no impact on quality.

Table 3-72. Virtual Switch Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-002	Use the existing vSphere Distributed Switch instance.	Provide guaranteed performance for vSAN traffic in case of contention by using existing networking components.	All traffic paths are shared over common uplinks.

Jumbo Frames

VMware vSAN supports jumbo frames for vSAN traffic.

A VMware vSAN design should use jumbo frames 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.

Table 3-73. Jumbo Frames Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-003	Configure jumbo frames on the VLAN dedicated to vSAN traffic.	Jumbo frames are already used to improve performance of vSphere vMotion and support VXLAN traffic.	Every device in the network must support jumbo frames.

VLANs

VMware recommends isolating VMware vSAN traffic 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 prevents interference between clusters and helps with troubleshooting cluster configuration.

Table 3-74. VLAN Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-004	Use a dedicated VLAN for vSAN traffic.	VLANs ensure traffic isolation.	VLANs span only a single pod. A sufficient number of VLANs are available within the consolidated pod and should be used for traffic segregation.

Multicast Requirements

VMware vSAN requires that IP multicast is enabled on the Layer 2 physical network segment that is used for intra-cluster communication. All VMkernel ports on the vSAN network subscribe to a multicast group using Internet Group Management Protocol (IGMP).

A default multicast address is assigned to each vSAN cluster at the time of creation. IGMP (v3) snooping is used to limit Layer 2 multicast traffic to specific port groups. As per the Physical Network Design, IGMP snooping is configured with an IGMP snooping querier to limit the physical switch ports that participate in the multicast group to only vSAN VMkernel port uplinks. In some cases, an IGMP snooping querier can be associated with a specific VLAN. However, vendor implementations might differ.

Cluster and Disk Group Design in ROBO

When considering the cluster and disk group design, you must decide on the vSAN datastore size, number of hosts per cluster, number of disk groups per host, and the vSAN policy.

vSAN Datastore Size

The size of the VMware vSAN datastore depends on the requirements for the datastore. To determine the appropriate size consider cost versus availability.

Table 3-75. VMware vSAN Datastore Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-005	Provide the consolidated cluster with enough vSAN space to run management and workloads virtual machines with FTT=1.	Management virtual machines require a maximum of 6 TB of storage when using FTT=1.	You must provide enough storage for all virtual machines. Monitoring and capacity management are critical and must be performed proactively.
ROBO-VI-STORAGE-SDS-006	On all vSAN datastores, ensure that at least 30% of free space is always available.	When vSAN reaches 80% usage a re-balance task is started which can be resource intensive.	Increases the amount of available storage needed. Monitoring and capacity management are critical and must be performed proactively.

Number of Hosts Per Cluster

The number of hosts in the cluster depends on these factors:

- Amount of available space on the vSAN datastore
- Number of failures you can tolerate in the cluster

For example, if the vSAN cluster has only 3 ESXi hosts, only a single failure is supported. If you require a higher level of availability, you must use additional hosts.

Cluster Size Design Background

Table 3-76. Number of Hosts Per Cluster

Design Quality	4 Hosts	32 Hosts	64 Hosts	Comments
Availability	↓	↑	↑↑	The more hosts that are available in the cluster, the more failures the cluster can tolerate.
Manageability	↓	↑	↑	The more hosts in the cluster, the more virtual machines can be in the VMware vSAN environment.
Performance	↑	↓	↓	Having a larger cluster can impact performance if there is an imbalance of resources. Consider performance as you make your decision.
Recoverability	o	o	o	Neither design option impacts recoverability.
Security	o	o	o	Neither design option impacts security.

Legend: ↑ = positive impact on quality; ↓ = negative impact on quality; o = no impact on quality.

Table 3-77. Cluster Size Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-007	Configure the consolidated cluster with a minimum of 4 ESXi hosts to support vSAN.	Having 4 hosts addresses the availability and sizing requirements, and allows you to take an ESXi host offline for maintenance or upgrades without impacting the overall vSAN cluster health.	Depending on the workloads deployed more hosts may need to be added for both compute and storage resources. Monitoring and capacity management are critical and must be performed proactively.

Number of Disk Groups Per Host

Disk group sizing is an important factor during volume design.

- If more hosts are available in the cluster, more failures are tolerated in the cluster. This capability adds cost because additional hardware for the disk groups is required.
- More available disk groups can increase the recoverability of VMware vSAN during a failure.

Consider these data points when deciding on the number of disk groups per host:

- Amount of available space on the vSAN datastore
- Number of failures you can tolerate in the cluster

The optimal number of disk groups is a balance between hardware and space requirements for the vSAN datastore. More disk groups increase space and provide higher availability. However, adding disk groups can be cost-prohibitive.

Disk Groups Design Background

The number of disk groups can affect availability and performance.

Table 3-78. Number of Disk Groups Per Host

Design Quality	1 Disk Group	3 Disk Groups	5 Disk Groups	Comments
Availability	↓	↑	↑↑	If more hosts are available in the cluster, the cluster tolerates more failures. This capability adds cost because additional hardware for the disk groups is required.
Manageability	o	o	o	If more hosts are in the cluster, more virtual machines can be managed in the vSAN environment.
Performance	o	↑	↑↑	If the flash percentage ratio to storage capacity is large, the vSAN can deliver increased performance and speed.
Recoverability	o	↑	↑↑	More available disk groups can increase the recoverability of vSAN during a failure. Rebuilds complete faster because there are more places to place data and to copy data from.
Security	o	o	o	Neither design option impacts security.

Legend: ↑ = positive impact on quality; ↓ = negative impact on quality; o = no impact on quality.

Table 3-79. Disk Groups Per Host Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-008	Configure VMware vSAN with a minimum of a single disk group per ESXi host.	Single disk group provides the required performance and usable space.	<p>Losing an SSD in a host takes the disk group offline.</p> <p>Using two or more disk groups can increase availability and performance.</p> <p>Depending on the workloads deployed creating a second disk group may be required.</p>

vSAN Policy Design in ROBO

After you enable and configure vSAN, you can create storage policies that define the virtual machine storage characteristics. Storage characteristics specify different levels of service for different virtual machines. The default storage policy tolerates a single failure and has a single disk stripe. Use the default unless your environment requires policies with non-default behavior. If you configure a custom policy, vSAN will guarantee it; however, if vSAN cannot guarantee a policy, you cannot provision a virtual machine that uses the policy unless you enable force provisioning.

VMware vSAN Policy Options

A storage policy includes several attributes, which can be used alone or combined to provide different service levels. Policies can be configured for availability and performance conservatively to balance space consumed and recoverability properties. In many cases, the default system policy is adequate and no additional policies are required. Policies allow any configuration to become as customized as needed for the application's business requirements.

Policy Design Background

Before making design decisions, understand the policies and the objects to which they can be applied. The policy options are listed in the following table.

Table 3-80. VMware vSAN Policy Options

Capability	Use Case	Value	Comments
Number of failures to tolerate	Redundancy	Default 1 Max 3	<p>A standard RAID 1 mirrored configuration that provides redundancy for a virtual machine disk. The higher the value, the more failures can be tolerated. For n failures tolerated, n+1 copies of the disk are created, and 2n+1 hosts contributing storage are required.</p> <p>A higher n value indicates that more replicas of virtual machines are made, which can consume more disk space than expected.</p>
Number of disk stripes per object	Performance	Default 1 Max 12	<p>A standard RAID 0 stripe configuration used to increase performance for a virtual machine disk.</p> <p>This setting defines the number of HDDs on which each replica of a storage object is striped.</p> <p>If the value is higher than 1, increased performance can result. However, an increase in system resource usage might also result.</p>

Table 3-80. VMware vSAN Policy Options (Continued)

Capability	Use Case	Value	Comments
Flash read cache reservation (%)	Performance	Default 0 Max 100%	Flash capacity reserved as read cache for the storage is a percentage of the logical object size that will be reserved for that object. Only use this setting for workloads if you must address read performance issues. The downside of this setting is that other objects cannot use a reserved cache. VMware recommends not using these reservations unless it is absolutely necessary because unreserved flash is shared fairly among all objects.
Object space reservation (%)	Thick provisioning	Default 0 Max 100%	The percentage of the storage object that will be thick provisioned upon VM creation. The remainder of the storage will be thin provisioned. This setting is useful if a predictable amount of storage will always be filled by an object, cutting back on repeatable disk growth operations for all but new or non-predictable storage use.
Force provisioning	Override policy	Default: No	Force provisioning allows for provisioning to occur even if the currently available cluster resources cannot satisfy the current policy. Force provisioning is useful in case of a planned expansion of the vSAN cluster, during which provisioning of VMs must continue. VMware vSAN automatically tries to bring the object into compliance as resources become available.

By default, policies are configured based on application requirements. However, they are applied differently depending on the object.

Table 3-81. Object Policy Defaults

Object	Policy	Comments
Virtual machine namespace	Failures-to-Tolerate: 1	Configurable. Changes are not recommended.
Swap	Failures-to-Tolerate: 1	Configurable. Changes are not recommended.
Virtual disk(s)	User-Configured Storage Policy	Can be any storage policy configured on the system.
Virtual disk snapshot(s)	Uses virtual disk policy	Same as virtual disk policy by default. Changes are not recommended.

Note If you do not specify a user-configured policy, the default system policy of 1 failure to tolerate and 1 disk stripe is used for virtual disk(s) and virtual disk snapshot(s). Policy defaults for the VM namespace and swap are set statically and are not configurable to ensure appropriate protection for these critical virtual machine components. Policies must be configured based on the application's business requirements. Policies give VMware vSAN its power because it can adjust how a disk performs on the fly based on the policies configured.

Policy Design Recommendations

Policy design starts with assessment of business needs and application requirements. Use cases for VMware vSAN must be assessed to determine the necessary policies. Start by assessing the following application requirements:

- I/O performance and profile of your workloads on a per-virtual-disk basis

- Working sets of your workloads
- Hot-add of additional cache (requires repopulation of cache)
- Specific application best practice (such as block size)

After assessment, configure the software-defined storage module policies for availability and performance in a conservative manner so that space consumed and recoverability properties are balanced. In many cases the default system policy is adequate and no additional policies are required unless specific requirements for performance or availability exist.

Table 3-82. Policy Design Decision

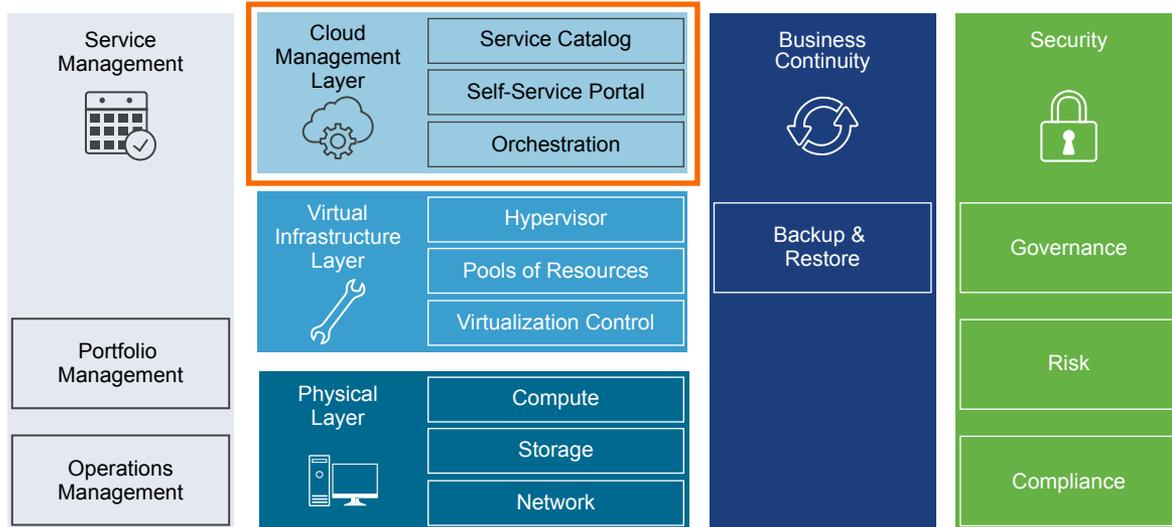
Decision ID	Design Decision	Design Justification	Design Implication
ROBO-VI-STORAGE-SDS-009	Use the default VMware vSAN storage policy.	The default vSAN storage policy provides the level of redundancy that is needed for the management workloads within the consolidated cluster.	Additional policies might be needed for workloads because their performance or availability requirements might differ from what the default VMware vSAN policy supports.
ROBO-VI-STORAGE-SDS-010	Configure the virtual machine swap file as a sparse objects on vSAN .	Enabling this setting creates virtual swap files as a sparse object on the vSAN datastore. Sparse virtual swap files will only consume capacity on vSAN as they are accessed. The result can be significantly less space consumed on the vSAN datastore, provided virtual machines do not experience memory over commitment, requiring use of the virtual swap file.	Administrative overhead to enable the advanced setting on all ESXi hosts running VMware vSAN.

Cloud Management Platform Design in ROBO

The Cloud Management Platform (CMP) layer is the management component of the entire SDDC. The CMP layer allows you to deliver tenants with automated workload provisioning by using a self-service portal.

The CMP layer includes the Service Catalog, which houses the facilities to be deployed, Orchestration which provides the workflows to get the catalog items deployed, and the Self-Service Portal that empowers the end users to take full advantage of the Software Defined Data Center. vRealize Automation provides the Portal and the Catalog, and vRealize Orchestrator takes care of the Orchestration.

Figure 3-20. The Cloud Management Platform Layer Within the Software-Defined Data Center



vRealize Automation Design in ROBO

VMware vRealize Automation provides a service catalog from which tenants can deploy applications, and a portal that lets you deliver a personalized, self-service experience to end users.

For ROBO, vRealize Automation Proxy Agents are deployed in each ROBO location connected to an existing vRealize Automation instance in the hub. Figures of the vRealize Automation deployment are used for reference only. For information on how to deploy vRealize Automation, see the *VMware Validated Design for SDDC* documentation.

vRealize Automation Logical and Physical Design in ROBO

The cloud management layer can deliver multi-platform and multi-vendor cloud services.

The cloud management services in the SDDC provide the following advantages.

- Comprehensive and purpose-built capabilities to provide standardized resources to global customers in a short time span.
- Multi-platform and multi-vendor delivery methods that integrate with existing enterprise management systems.
- Central user-centric and business-aware governance for all physical, virtual, private, and public cloud services.
- Design that meets the customer and business needs and is extensible.

Physical Design in ROBO

The physical design consists of characteristics and decisions that support the logical design.

Deployment Considerations in ROBO

This design uses NSX logical switches to abstract the vRealize Automation application and its supporting services. This abstraction allows the application to be hosted in any given region regardless of the underlying physical infrastructure such as network subnets, compute hardware, or storage types. This design places the core vRealize Automation application and its supporting services in the hub. The same instance of the application manages workloads in all Regions and ROBO locations.

Table 3-83. vRealize Automation Region Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-001	Utilize a single vRealize Automation installation to manage Region A, Region B, and ROBO deployments from a single instance.	vRealize Automation can manage one or more regions as well as ROBO locations. This provides a single consumption portal regardless of region or ROBO location. The abstraction of the vRealize Automation application over virtual networking allows it to be independent from any physical site locations or hardware.	You must size vRealize Automation to accommodate multi-region and ROBO deployments. vRealize Automation supports only ten connected vCenter Servers.

Table 3-84. vRealize Automation Anti-Affinity Rules

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-002	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules to the vRealize Automation components	Using DRS prevents vRealize Automation nodes from residing on the same ESXi host and thereby risking the cluster's high availability capability	Additional configuration is required to set up anti-affinity rules.

Table 3-85. vRealize Automation IaaS AD Requirement

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-003	vRealize Automation IaaS Machines are joined to Active Directory.	This is a hard requirement for vRealize Automation	Active Directory access must be provided using dedicated service accounts

Figure 3-21. vRealize Automation Design Overview for the Hub

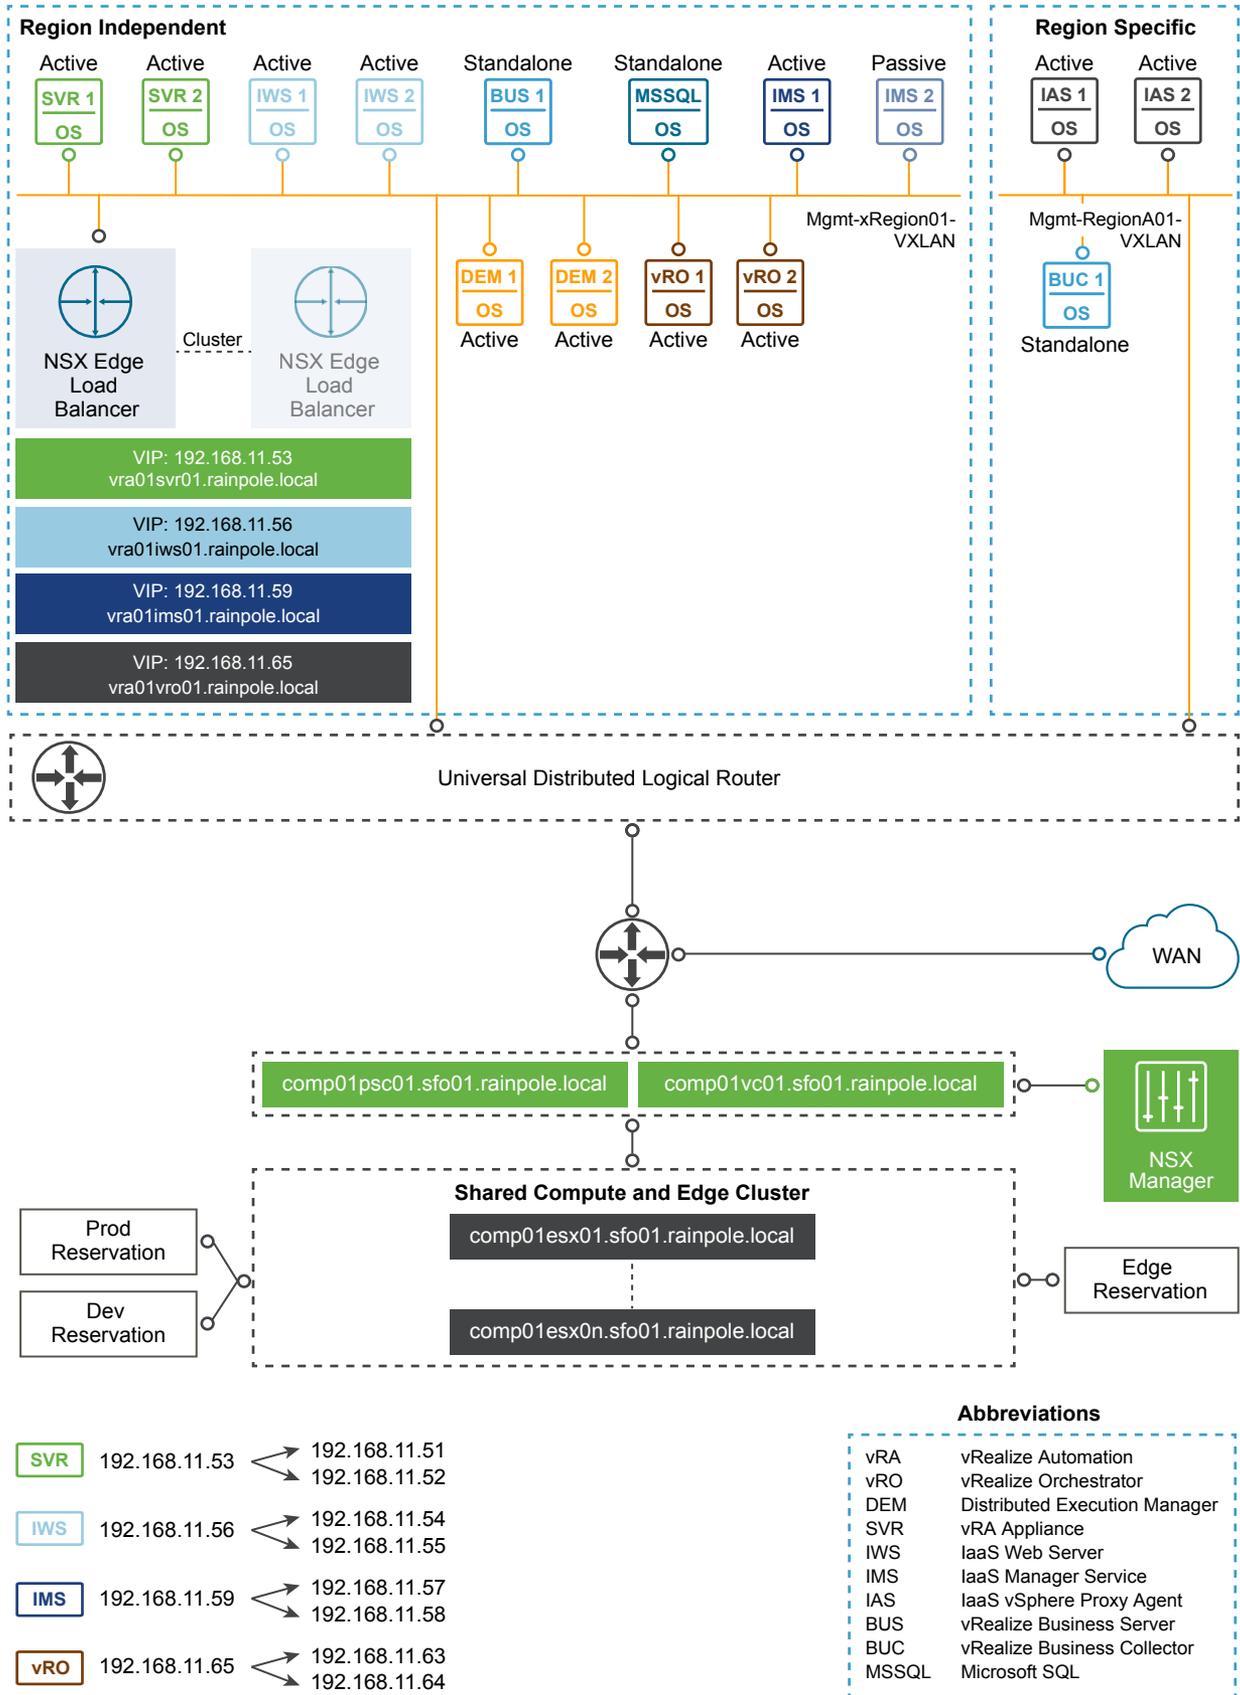
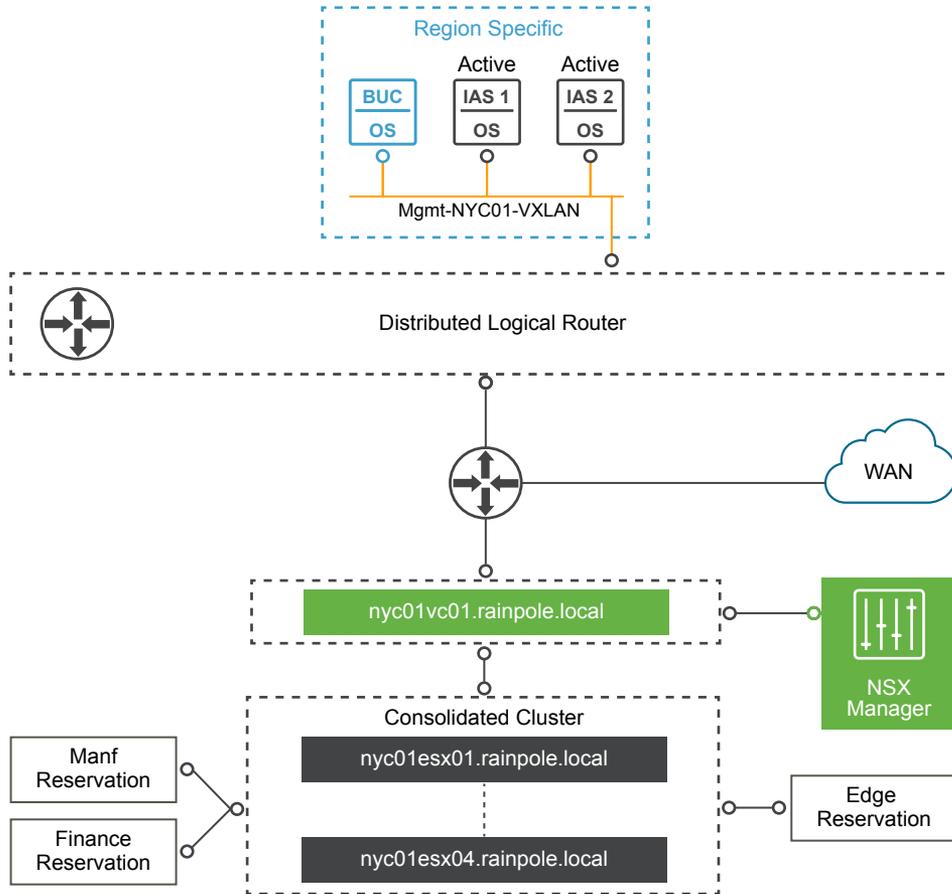


Figure 3-22. vRealize Automation Design Overview for ROBO



vRealize Automation Deployment in ROBO

You deploy vRealize Automation in a single region to manage both Region A and Region B, as well as the ROBO locations.

Centralized management of these sites is performed to reduce operational overhead and simplify automation support for IT services. Proxy Agents are deployed to each ROBO location being managed, these ROBO locations do not require a full vRealize Automation deployment.

This design only applies to the vRealize Automation requirements at the ROBO location.

Table: Cloud Management Platform Elements

Element	Location	Component
vRealize Automation virtual appliance	Hub	<ul style="list-style-type: none"> ■ vRealize Automation Portal Web/Application Server ■ vRealize Automation PostgreSQL Database ■ vRealize Automation service catalog ■ VMware Identity Manager
vRealize Automation IaaS components	Hub	<ul style="list-style-type: none"> ■ vRealize Automation IaaS Web Server ■ vRealize Automation IaaS Manager Services

Element	Location	Component
Distributed execution components	Hub	<ul style="list-style-type: none"> ■ vRealize Automation Distributed Execution Managers. <ul style="list-style-type: none"> ■ Orchestrator ■ Workers
Integration components	Hub ROBO	<ul style="list-style-type: none"> ■ vRealize Automation Agent machines
vRealize Orchestrator components	Hub	<ul style="list-style-type: none"> ■ vRealize Orchestrator virtual appliances
Provisioning infrastructure	Hub ROBO	<ul style="list-style-type: none"> ■ vSphere environment ■ Other supported physical, virtual, or cloud environments.
Costing components	Hub ROBO	<ul style="list-style-type: none"> ■ vRealize Business for Cloud Standard server (Hub only) ■ vRealize Business for Cloud Standard data collector
Supporting infrastructure	Hub ROBO	<ul style="list-style-type: none"> ■ Microsoft SQL database environment (hub only) ■ Active Directory environment ■ DNS ■ SMTP ■ NTP

For more information on the complete design and configuration steps required to stand up the vRealize Automation instance, please see the VMware Validated Design for SDDC.

Note As workloads and the number of ROBO sites increase, it is possible a single vRealize Automation platform will be unable to handle all of the required operations. Additional vRealize Automation instances may be required to support additional ROBO sites. Additional vRealize Automation platforms should be deployed at the appropriate hub as per the VMware Validated Design for SDDC and protected to the recovery region.

vRealize Automation IaaS Proxy Agent in ROBO

The vRealize Automation IaaS Proxy Agent is a windows service used to communicate with specific infrastructure endpoints. In this design, the vSphere Proxy agent is utilized to communicate with vCenter Server.

The IaaS Proxy Agent server provides the following functions:

- vRealize Automation IaaS Proxy Agent can interact with different types of infrastructure components. For this design, only the vSphere Proxy agent is used.
- vRealize Automation does not itself virtualize resources, but works with vSphere to provision and manage the virtual machines. It uses vSphere Proxy agents to send commands to and collect data from vSphere.

Table 3-86. vRealize Automation IaaS Agent Server Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-004	Deploy two vRealize Automation vSphere Proxy Agent virtual machines in each ROBO.	Using two virtual machines provides redundancy for vSphere connectivity.	More resources are required because multiple virtual machines are deployed for this function.
ROBO-CMP-005	Place the vRealize Automation vSphere Proxy Agent virtual machines on the management application virtual network.	Per ROBO-VI-SDN-023 all management virtual machines must reside on VXLAN backed networks unless there is a hard requirement for it to be placed on a VLAN backed network.	All management applications share the same layer 2 broadcast domain.

Table 3-87. vRealize Automation IaaS Proxy Agent Resource Requirements per Virtual Machine

Attribute	Specification
Number of vCPUs	2
Memory	4 GB
Number of vNIC ports	1
Number of local drives	1
vRealize Automation functions	vSphere Proxy agent
Operating system	Microsoft Windows Server 2012 SP2 R2

vRealize Automation Supporting Infrastructure in ROBO

To satisfy the requirements of this ROBO SDDC design, you configure additional components for vRealize Automation.

vRealize Business for Cloud in ROBO

vRealize Business for Cloud provides end-user transparency in the costs that are associated with operating workloads.

vRealize Business integrates with vRealize Automation to display costing during workload request and on an ongoing basis with cost reporting by user, business group or tenant. Additionally, tenant administrators can create a wide range of custom reports to meet the requirements of an organization.

vRealize Business can collect data from geographically distributed endpoints without the need to directly connect the primary instance. The data collector connects to management components such as vCenter Server and pushes data back to the vRealize Business appliance located in the hub.

Table 3-88. vRealize Business for Cloud Standard Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CMP-006	Deploy one vRealize Business remote collector appliance in the management application virtual network.	vRealize Business remote collectors are used to send cost data back to the primary vRealize Business appliance.	The communication with vCenter Server involves an additional Layer 3 hop through an NSX edge device.

Table 3-89. vRealize Business Virtual Appliance Resource Requirements

Attribute	Specification
Number of vCPUs	4
Memory	8 GB for Remote Collector
vRealize Business function	Remote collector

vRealize Automation Cloud Tenant Design in ROBO

A tenant is an organizational unit within a vRealize Automation deployment, and can represent a business unit within an enterprise, or a company that subscribes to cloud services from a service provider. Each tenant has its own dedicated configuration, although some system-level configuration is shared across tenants.

Comparison of Single-Tenant and Multi-Tenant Deployments in ROBO

vRealize Automation supports deployments with a single tenant or multiple tenants. System-wide configuration is always performed using the default tenant, and can then be applied to one or more tenants. For example, system-wide configuration might specify defaults for branding and notification providers.

Infrastructure configuration, including the infrastructure sources that are available for provisioning, can be configured in any tenant and is shared among all tenants. The infrastructure resources, such as cloud or virtual compute resources or physical machines, can be divided into fabric groups managed by fabric administrators. The resources in each fabric group can be allocated to business groups within each tenant by using reservations.

Default-Tenant Deployment

In a default-tenant deployment, all configuration occurs in the default tenant. Tenant administrators can manage users and groups, and configure tenant-specific branding, notifications, business policies, and catalog offerings. All users log in to the vRealize Automation console at the same URL, but the features available to them are determined by their roles.

Single-Tenant Deployment

In a single-tenant deployment, the system administrator creates a single new tenant for the organization that use the same vRealize Automation instance. Tenant users log in to the vRealize Automation console at a URL specific to their tenant. Tenant-level configuration is segregated from the default tenant, although users with system-wide roles can view and manage both configurations. The IaaS administrator for the organization tenant creates fabric groups and appoints fabric administrators. Fabric administrators can create reservations for business groups in the organization tenant.

Multi-Tenant Deployment

In a multi-tenant deployment, the system administrator creates new tenants for each organization that uses the same vRealize Automation instance. Tenant users log in to the vRealize Automation console at a URL specific to their tenant. Tenant-level configuration is segregated from other tenants

and from the default tenant, although users with system-wide roles can view and manage configuration across multiple tenants. The IaaS administrator for each tenant creates fabric groups and appoints fabric administrators to their respective tenants. Although fabric administrators can create reservations for business groups in any tenant, in this scenario they typically create and manage reservations within their own tenants. If the same identity store is configured in multiple tenants, the same users can be designated as IaaS administrators or fabric administrators for each tenant.

Tenant Design in ROBO

This design deploys a single tenant containing two business groups.

- The first business group is designated for manufacturing workloads.
- The second business group is designated for finance workloads.

Tenant administrators manage users and groups, configure tenant-specific branding, notifications, business policies, and catalog offerings. All users log in to the vRealize Automation console at the same URL, but the features available to them are determined by their roles.

The figure below illustrates the ROBO tenant design.

Figure 3-23. Rainpole Cloud Automation Tenant Design for Two ROBO Sites

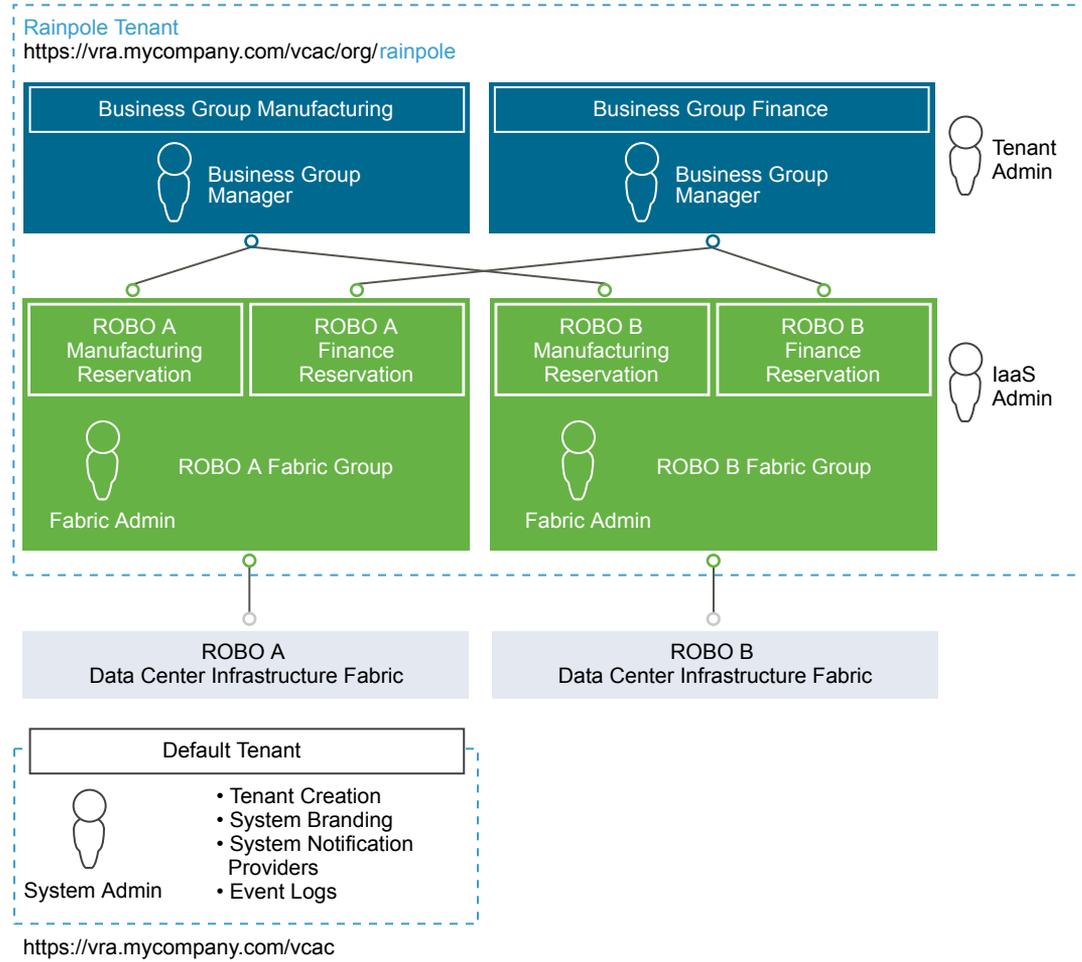


Table 3-90. Tenant Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-007	Utilizes vRealize Automation business groups for separate business units (instead of separate tenants).	Allows transparency across the environments and some level of sharing of resources and services such as blueprints.	Some elements such as build profiles are visible to both business groups. The design does not provide full isolation for security or auditing.
ROBO-CMP-008	Create separate fabric groups for each ROBO. Each fabric group represent ROBO specific data center resources. Each of the business groups have reservations into each of the fabric groups.	Provides future isolation of fabric resources and potential delegation of duty to independent fabric administrators.	Initial deployment will use a single shared fabric that consists of one compute pod.

Table 3-90. Tenant Design Decisions (Continued)

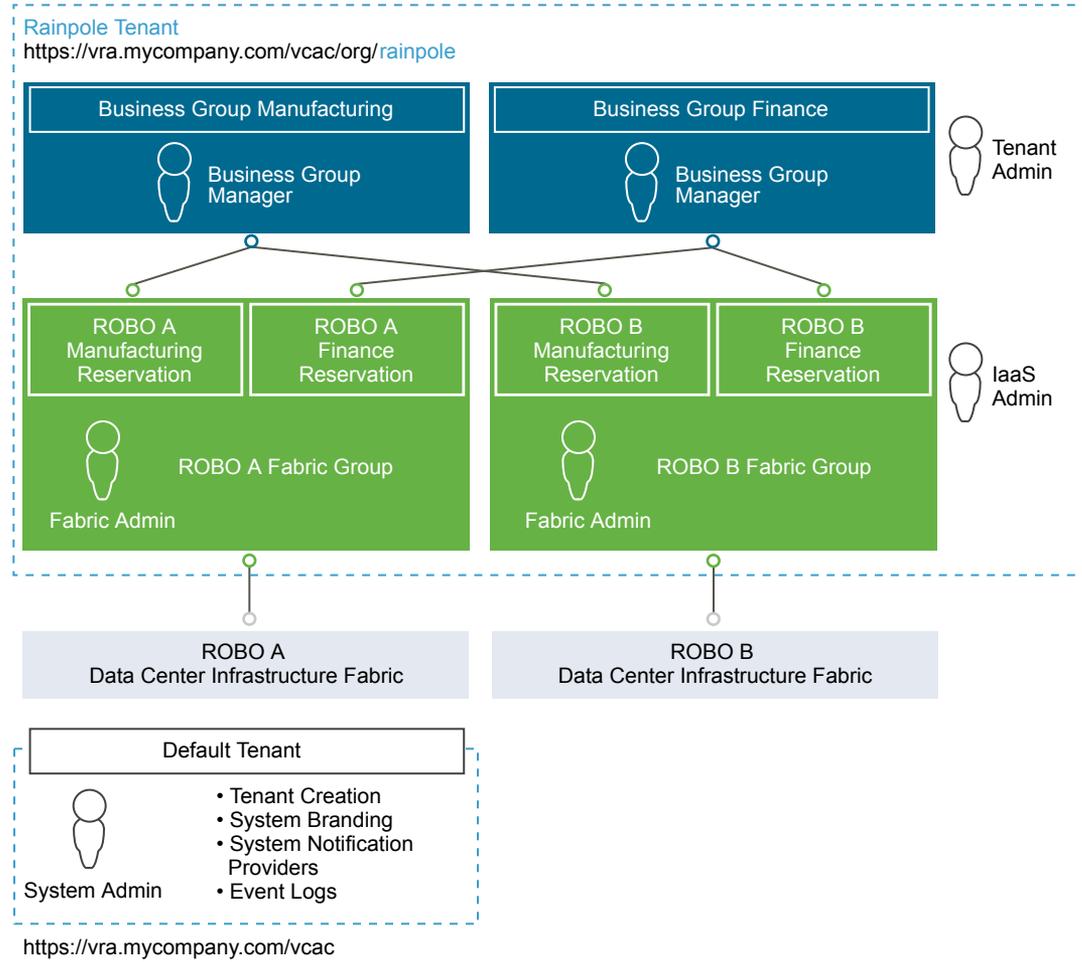
Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-009	Allow access to the default tenant only by the system administrator and for the purposes of managing tenants and modifying system-wide configurations.	Isolates the default tenant from individual tenant configurations.	Each tenant administrator is responsible for managing their own tenant configuration.
ROBO-CMP-010	Evaluate your internal organization structure and workload needs. Configure vRealize Business Groups, Reservations, Service Catalogs, and templates based on your specific organization's needs.	vRealize Automation is designed to integrate with your individual organization's needs. Within this design, guidance for Rainpole is provided as a starting point, but this guidance may not be appropriate for your specific business needs.	Partners and Customers will need to evaluate their specific business needs.

vRealize Automation Infrastructure as a Service Design in ROBO

This topic introduces the integration of vRealize Automation with vSphere resources used to create the Infrastructure as a Service design for use with the SDDC.

vRealize Automation Logical Design illustrates the logical design of the vRealize Automation groups and vSphere resources.

Figure 3-24. vRealize Automation Logical Design



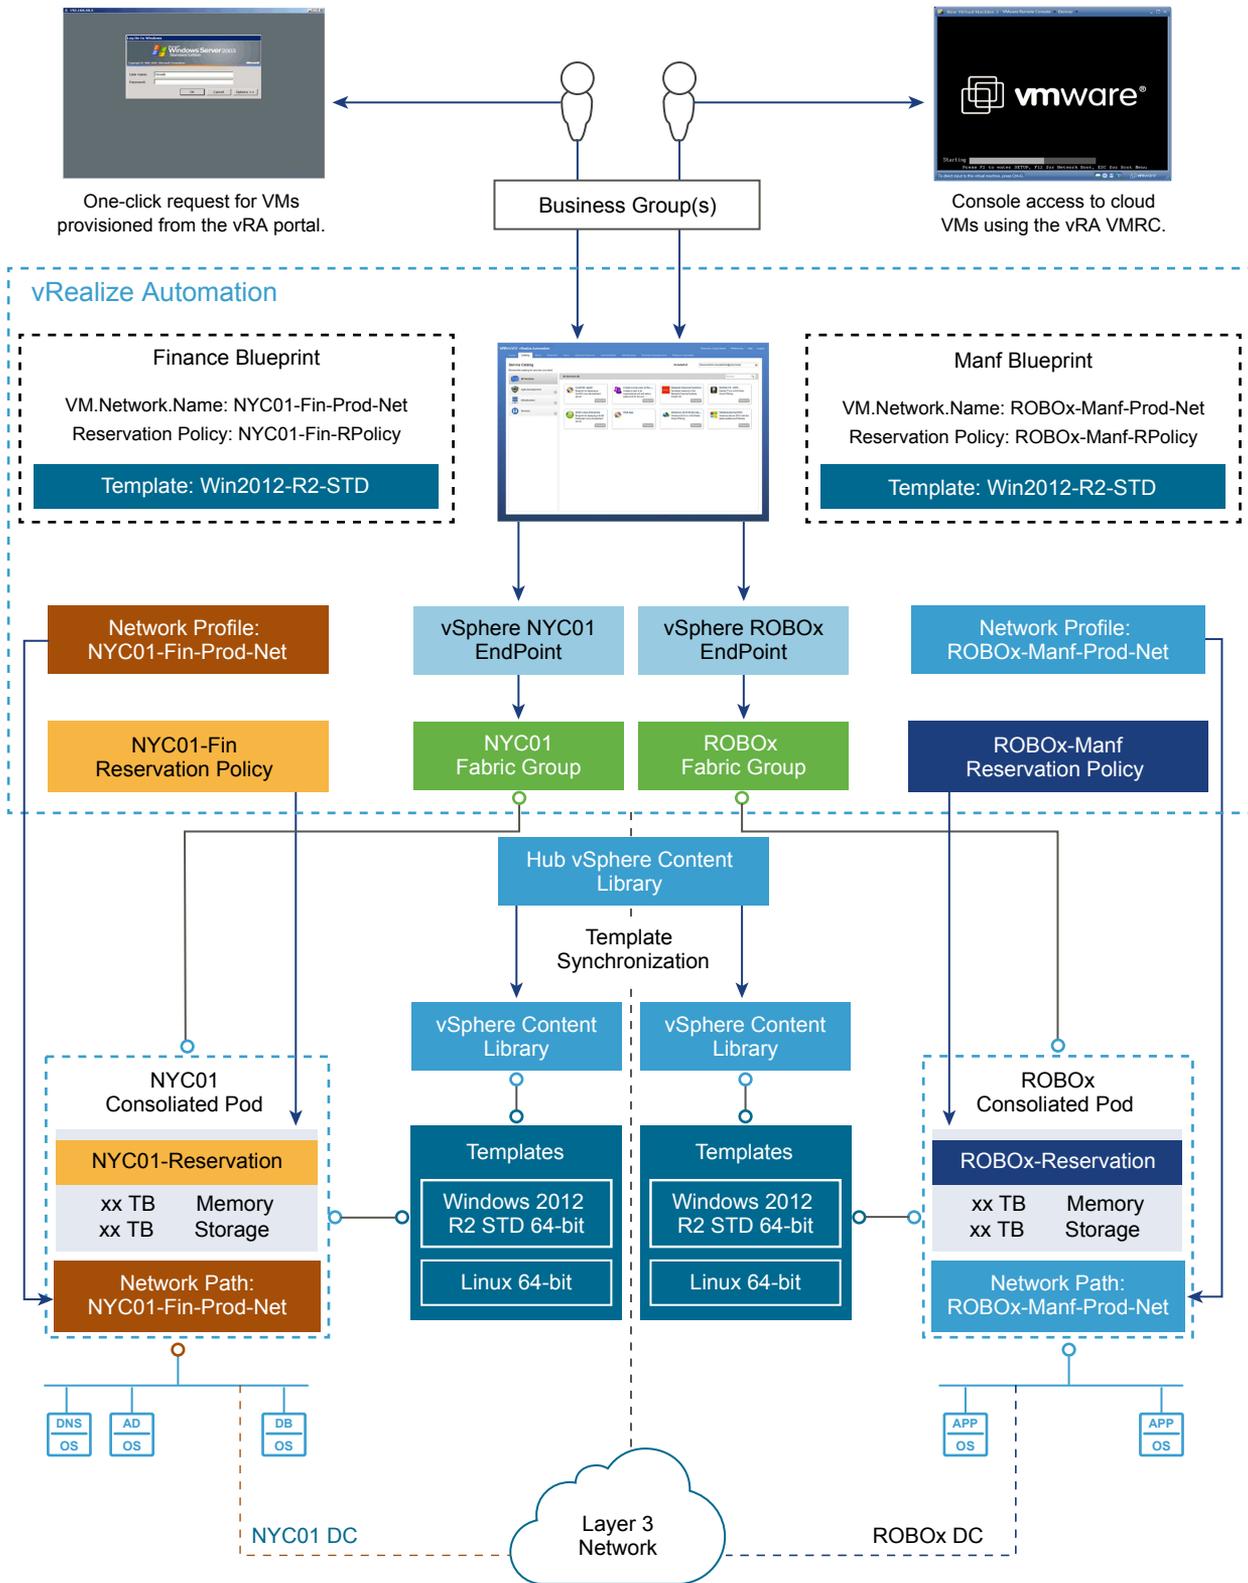
The following terms apply to vRealize Automation when integrated with vSphere. These terms and their meaning may vary from the way they are used when referring only to vSphere.

Term	Definition
vSphere (vCenter Server) endpoint	Provides information required by vRealize Automation IaaS to access vSphere compute resources.
Compute resource	Virtual object within vRealize Automation that represents a vCenter Server cluster or resource pool, and datastores or datastore clusters. Note Compute resources are CPU, memory, storage and networks. Datastores and datastore clusters are part of the overall storage resources.
Fabric groups	vRealize Automation IaaS organizes compute resources into fabric groups.
Fabric administrators	Fabric administrators manage compute resources, which are organized into fabric groups.

Term	Definition
Compute reservation	<p>A share of compute resources (vSphere cluster, resource pool, datastores, or datastore clusters), such as CPU and memory reserved for use by a particular business group for provisioning virtual machines.</p> <p>Note vRealize Automation uses the term reservation to define resources (be they memory, storage or networks) in a cluster. This is different than the use of reservation in vCenter Server, where a share is a percentage of total resources, and reservation is a fixed amount.</p>
Storage reservation	<p>Similar to compute reservation (see above), but pertaining only to a share of the available storage resources. In this context, you specify a storage reservation in terms of gigabytes from an existing LUN or Datastore.</p>
Business groups	<p>A collection of virtual machine consumers, usually corresponding to an organization's business units or departments. Only users in the business group can request virtual machines.</p>
Reservation policy	<p>vRealize Automation IaaS determines its reservation (also called virtual reservation) from which a particular virtual machine is provisioned. The reservation policy is a logical label or a pointer to the original reservation. Each virtual reservation can be added to one reservation policy.</p>
Build profile	<p>A set of user defined properties you apply to a virtual machine when it is provisioned. For example, the operating system used in a blueprint, or the available networks to use for connectivity at the time of provisioning the virtual machine.</p> <p>Build profile properties determine the specification of the virtual machine, the manner in which it is provisioned, operations to perform after it is provisioned, or management information maintained within vRealize Automation.</p>
Blueprint	<p>The complete specification for a virtual machine, determining the machine attributes, the manner in which it is provisioned, and its policy and management settings.</p> <p>Blueprint allows the users of a business group to create virtual machines on a virtual reservation (compute resource) based on the reservation policy, and using platform and cloning types. It also lets you specify or add machine resources and build profiles.</p>

vRealize Automation Integration with vSphere Endpoint shows the logical design constructs discussed in the previous section as they apply to a deployment of vRealize Automation integrated with vSphere in a cross data center provisioning.

Figure 3-25. vRealize Automation Integration with vSphere Endpoint



Infrastructure Source Endpoints in ROBO

An infrastructure source endpoint is a connection to the infrastructure that provides a set (or multiple sets) of resources, which can then be made available by IaaS administrators for consumption by end users. vRealize Automation IaaS regularly collects information about known endpoint resources and the virtual resources provisioned therein. Endpoint resources are referred to as compute resources or compute pods, the terms are often used interchangeably.

Infrastructure data is collected through proxy agents that manage and communicate with the endpoint resources. This information about the compute resources on each infrastructure endpoint and the machines provisioned on each computer resource is collected at regular intervals.

During installation of the vRealize Automation IaaS components, you can configure the proxy agents and define their associated endpoints. Alternatively, you can configure the proxy agents and define their associated endpoints separately after the main vRealize Automation installation is complete.

Table 3-91. Endpoint Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-011	Create a vSphere endpoint for each ROBO.	One vSphere endpoint is required to connect to each vCenter Server instance in each ROBO.	As additional ROBOs are brought online additional vSphere endpoints need to be deployed.

Virtualization Compute Resources in ROBO

A virtualization compute resource is a vRealize Automation object that represents an ESXi host or a cluster of ESXi hosts. When a group member requests a virtual machine, the virtual machine is provisioned on these compute resources. vRealize Automation regularly collects information about known compute resources and the virtual machines provisioned on them through the proxy agents.

Table 3-92. Compute Resource Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-012	Create a compute resources for each ROBO.	Each ROBO has one Consolidated cluster, one compute resource is required for each cluster.	Each ROBO must have a compute resource created.

Note By default, compute resources are provisioned to the root of the compute cluster. In this design, use of vSphere resource pools is mandatory.

Fabric Groups in ROBO

A fabric group is a logical container of several compute resources, and can be managed by fabric administrators.

Table 3-93. Fabric Group Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-013	Create a fabric group for each ROBO and include all the compute resources and edge resources.	To enable ROBO specific provisioning a fabric group in each must be created. This allows different fabric administrators in each location.	A fabric group must be created for each ROBO.

Business Groups in ROBO

A business group is a collection of machine consumers, often corresponding to a line of business, department, or other organizational unit. To request machines, a vRealize Automation user must belong to at least one business group. Each group has access to a set of local blueprints used to request machines.

Business groups have the following characteristics:

- A group must have at least one business group manager, who maintains blueprints for the group and approves machine requests.
- Groups can contain support users, who can request and manage machines on behalf of other group members.
- A vRealize Automation user can be a member of more than one Business group, and can have different roles in each group.

Reservations in ROBO

A reservation is a share of one compute resource's available memory, CPU and storage reserved for use by a particular fabric group. Each reservation is for one fabric group only but the relationship is many-to-many. A fabric group might have multiple reservations on one compute resource, or reservations on multiple compute resources, or both.

Consolidated Cluster and Resource Pools

While reservations provide a method to allocate a portion of the cluster memory or storage within vRealize Automation, reservations do not control how CPU and memory is allocated during periods of contention on the underlying vSphere compute resources. vSphere Resource Pools are utilized to control the allocation of CPU and memory during time of resource contention on the underlying host. To fully utilize this, all VMs must be deployed into one of four resource pools: *ROBO-MGMT*, *ROBO-Edge*, *User-Edge*, and *User-VM*. *ROBO-MGMT* is dedicated for datacenter level management components and should not contain any user workloads. *ROBO-Edge* is dedicated for datacenter level NSX Edge components and should not contain any user workloads. *User-Edge* is dedicated for any statically or dynamically deployed NSX components such as NSX Edges or Load Balancers which serve a specific customer workload. The *User-VM* is dedicated for any statically or dynamically deployed virtual machine, such as Windows, Linux, or databases containing specific customer workloads.

Table 3-94. Reservation Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-014	Create at least one vRealize Automation reservation for each business group.	In our example, each Consolidated cluster will have two reservations, one for manufacturing and one for finance, allowing both manufacturing and finance workloads to be provisioned.	Because manufacturing and finance share the same compute resources, the business groups must be limited to a fixed amount of resources.
ROBO-CMP-015	Create at least one vRealize Automation reservation for edge resources in each ROBO.	An edge reservation in each region allows NSX to create edge services gateways on demand and place them on the edge cluster.	The workload reservation must define the edge reservation in the network settings.
ROBO-CMP-016	Configure all vRealize Automation workloads to utilize dedicated vCenter Resource Pools.	In order to ensure dedicated compute resources of management and NSX networking components, end-user deployed workloads must be assigned to a dedicated end-user workload vCenter Resource Pools. Workloads provisioned at the root resource pool level will receive more resources than resource pools, which would starve those virtual machines in contention situations.	Cloud administrators must ensure all workload reservations are configured with the appropriate resource pool.
ROBO-CMP-017	Configure vRealize Automation reservations for dynamically provisioned NSX Edge components (routed gateway) to utilize dedicated vCenter Resource Pools.	In order to ensure dedicated compute resources of management and NSX networking components, end-user deployed NSX edge components must be assigned to a dedicated end-user network component vCenter Resource Pool. Workloads provisioned at the root resource pool level will receive more resources than resource pools, which would starve those virtual machines in contention situations.	Cloud administrators must ensure all workload reservations are configured with the appropriate resource pool.
ROBO-CMP-018	All vCenter resource pools utilized for Edge or Compute workloads must be created at the "root" level. Nesting of resource pools is not recommended.	Nesting of resource pools can create administratively complex resource calculations that may result in unintended under or over allocation of resources during contention situations.	All resource pools must be created at the root resource pool level.

Reservation Policies in ROBO

You can add each virtual reservation to one reservation policy. The reservation from which a particular virtual machine is provisioned is determined by vRealize Automation based on the reservation policy specified in the blueprint, if any, the priorities and current usage of the fabric group's reservations, and other custom properties.

Table 3-95. Reservation Policy Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-019	Create at least one workload reservation policy for each ROBO.	Reservation policies are used to target a deployment to a specific set of reservations in each ROBO. Reservation policies are also used to target workloads into their appropriate vSphere resource pool.	None
ROBO-CMP-020	Create at least one reservation policy for placement of dynamically created edge services gateways into the User-Edge resource pool.	Required to place the edge devices into their respective edge vSphere resource pool.	None

A storage reservation policy is a set of datastores that can be assigned to a machine blueprint to restrict disk provisioning to only those datastores. Storage reservation policies are created and associated with the appropriate datastores and assigned to reservations.

Table 3-96. Storage Reservation Policy Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-021	Within this design, storage tiers are not used.	The underlying physical storage design does not use storage tiers.	Both business groups will have access to the same storage. For customers who utilize multiple datastores with different storage capabilities will need to evaluate the usage of vRealize Automation Storage Reservation Policies.

Template Synchronization in ROBO

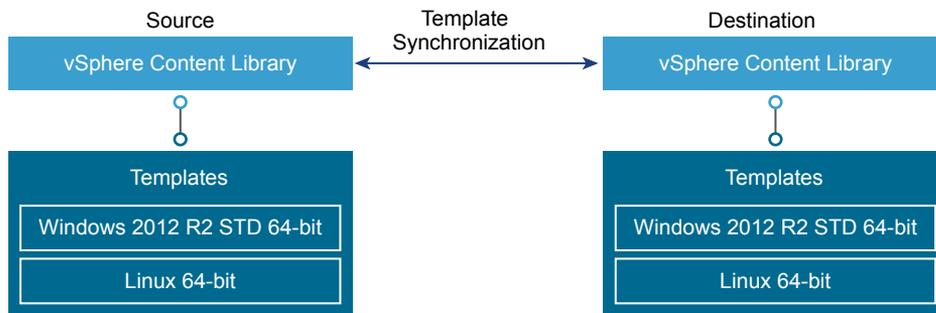
This design supports provisioning workloads across regions/ROBOs from the same portal using the same single-machine blueprints. A synchronization mechanism is required to have consistent templates across regions. There are multiple ways to achieve synchronization, for example, vSphere Content Library or external services like vCloud Connector or vSphere Replication.

Table 3-97. Template Synchronization Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-CMP-022	This design uses vSphere Content Library to synchronize templates across regions.	The vSphere Content Library is built into the version of vSphere being used and meets all the requirements to synchronize templates.	Storage space must be provisioned in each region. vRealize Automation cannot directly consume templates from vSphere Content Library.

When users select this blueprint, vRealize Automation clones a vSphere virtual machine template with preconfigured vCenter customizations.

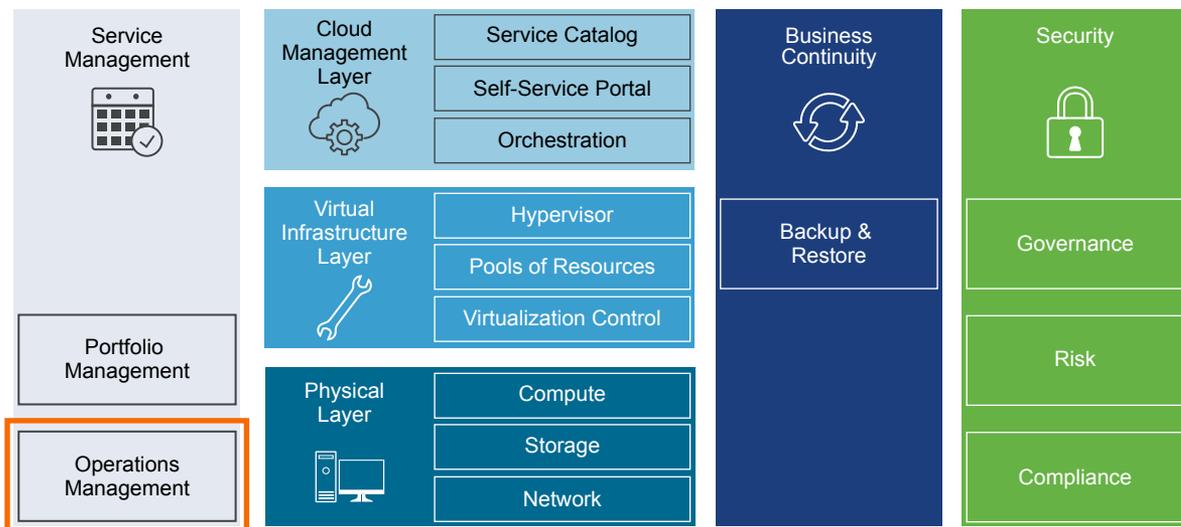
Figure 3-26. Template Synchronization



Operations Infrastructure Design in ROBO

Operations Management is a required element of a software-defined data center. Monitoring operations support in vRealize Operations Manager and vRealize Log Insight provides capabilities for performance and capacity management of related infrastructure and cloud management components.

Figure 3-27. Operations Management Layer for ROBO



- **vRealize Operations Manager Design in ROBO**

The foundation of vRealize Operations Manager is a single instance of a 3-node analytics cluster that is deployed in the hub of the SDDC, and a 2-node remote collector cluster in each ROBO.

- **vRealize Log Insight Design in ROBO**

vRealize Log Insight design enables real-time logging for all components that build up the management capabilities of the SDDC.

- **vSphere Data Protection Design in ROBO**

Design data protection of the management components in your environment to ensure continuous operation of the ROBO SDDC if the data of a management application is damaged.

- [vSphere Update Manager in ROBO](#)

vSphere Update Manager pairs with vCenter Server to enable patch and version management of ESXi hosts and virtual machines.

vRealize Operations Manager Design in ROBO

The foundation of vRealize Operations Manager is a single instance of a 3-node analytics cluster that is deployed in the hub of the SDDC, and a 2-node remote collector cluster in each ROBO.

Logical Design

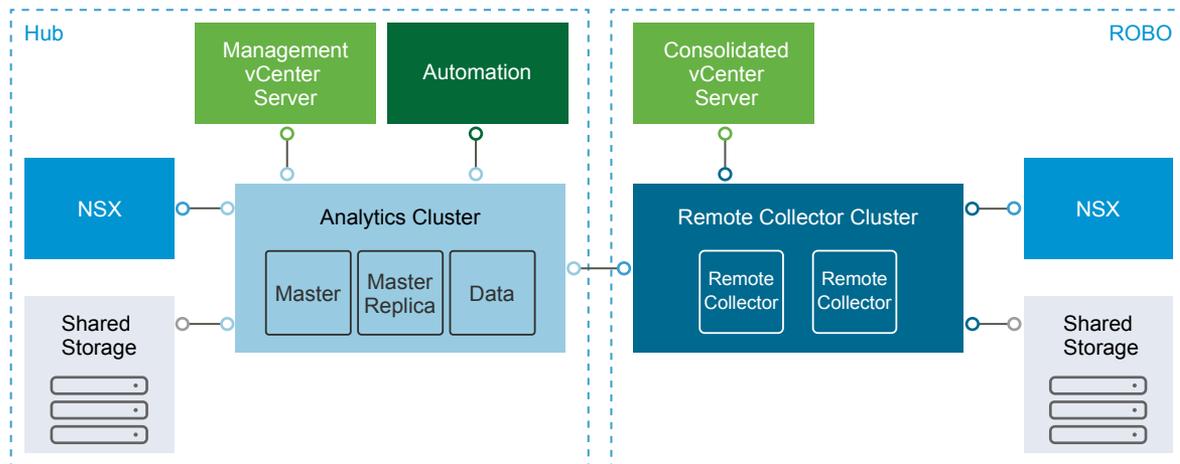
You deploy a vRealize Operations Manager configuration that consists of the following entities.

- 3-node (medium-size) vRealize Operations Manager analytics cluster that is highly available (HA) in the hub. This topology provides high availability, scale-out capacity up to sixteen nodes, and failover.
- 2 remote collector nodes in each ROBO. The remote collectors communicate directly with the data nodes in the vRealize Operations Manager analytics cluster. For high availability and fault tolerance, two remote collectors are deployed in each ROBO.

Each ROBO contains its own pair of remote collectors whose role is to ease scalability by performing the data collection from the applications and periodically sending collected data to the analytics cluster.

The vRealize Operations Manager analytics cluster design is covered in the VMware Validated Design for SDDC.

Figure 3-28. Logical Design of vRealize Operations Manager ROBO Deployment



Physical Design

The vRealize Operations Manager analytics cluster and remote collector nodes run on the management pod in hub and the remote collector nodes run on the Consolidated cluster in the ROBO. For information about the types of pods, see Pod Architecture.

Data Sources

vRealize Operations Manager collects data from the following virtual infrastructure and cloud management components located within the ROBO.

- Consolidated pod
 - vCenter Server
 - ESXi hosts
- NSX for vSphere
 - NSX Manager
 - NSX Controller Instances
 - NSX Edge instances
- vRealize Automation
 - vRealize Agent Servers
- vRealize Log Insight

vRealize Operations Manager Nodes in ROBO

The analytics cluster of the vRealize Operations Manager deployment contains the nodes that analyze and store data from the monitored components. You deploy a configuration of the analytics cluster that satisfies the requirements for monitoring the number of virtual machines according to the design objectives of this VMware Validated Design.

A 3-node vRealize Operations Manager analytics cluster in the cross-region application virtual network is deployed as part of the VMware Validated Design for SDDC. The analytics cluster consists of one master node, one master replica node, and one data node to enable scale out and high availability.

A 2-node vRealize Operations Manager remote collector cluster is deployed in each ROBO.

Table 3-98. Cluster Node Configuration Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-MON-001	Deploy two remote collector nodes per ROBO.	<ul style="list-style-type: none"> ■ Removes the load from the analytics cluster from collecting metrics from applications in ROBO locations. ■ Provides high availability and fault tolerance, where If the Remote Collector running an adapter fails, the adapter is automatically moved to the other Remote Collector in the collector group 	When configuring the monitoring of a solution, you must assign a collector group.

Sizing Compute Resources in ROBO

You size compute resources for vRealize Operations Manager to provide enough resources to run the remote collectors in each ROBO.

Sizing Compute Resources for the Remote Collector Nodes

Unlike the analytics cluster nodes, remote collector nodes have only the collector role. Deploying two remote collector nodes in each ROBO does not increase the capacity for monitored objects.

Table 3-99. Size of a Standard Remote Collector Virtual Appliance for vRealize Operations Manager

Attribute	Specification
Appliance size	Remote Collector - Standard
vCPU	2
Memory	4 GB
Single-node maximum Objects(*)	1,500
Single-Node Maximum Collected Metrics	600,000
Multi-Node Maximum Objects Per Node	N/A
Multi-Node Maximum Collected Metrics Per Node	N/A
Maximum number of End Point Operations Management Agents per Node	250
Maximum Objects for 16-Node Maximum	N/A
Maximum Metrics for 16-Node Configuration	N/A

*The object limit for the remote collector is based on the VMware vCenter adapter.

Table 3-100. Configuration and Resources Requirements of the Remote Collector Nodes Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-MON-002	Deploy the standard-size remote collector virtual appliances.	Enables metric collection for the expected number of objects in the SDDC when at full capacity.	You must provide 4 vCPUs and 8 GB of memory in the Consolidated cluster in each ROBO.

Networking Design in ROBO

You place the vRealize Operations Manager nodes in several network application virtual networks for isolation and failover. The networking design also supports public access to the analytics cluster nodes.

The analytics cluster is depicted only for illustrative purposes. The design for the analytics cluster can be found in the *VMware Validated Design for SDDC* documentation. For secure access, load balancing and portability, the vRealize Operations Manager analytics cluster is deployed in the shared cross-region application isolated network `Mgmt-xRegion01-VXLAN`, and the remote collector clusters in the shared local application isolated networks in each ROBO.

Figure 3-29. Networking Design of vRealize Operations Manager at the SDDC

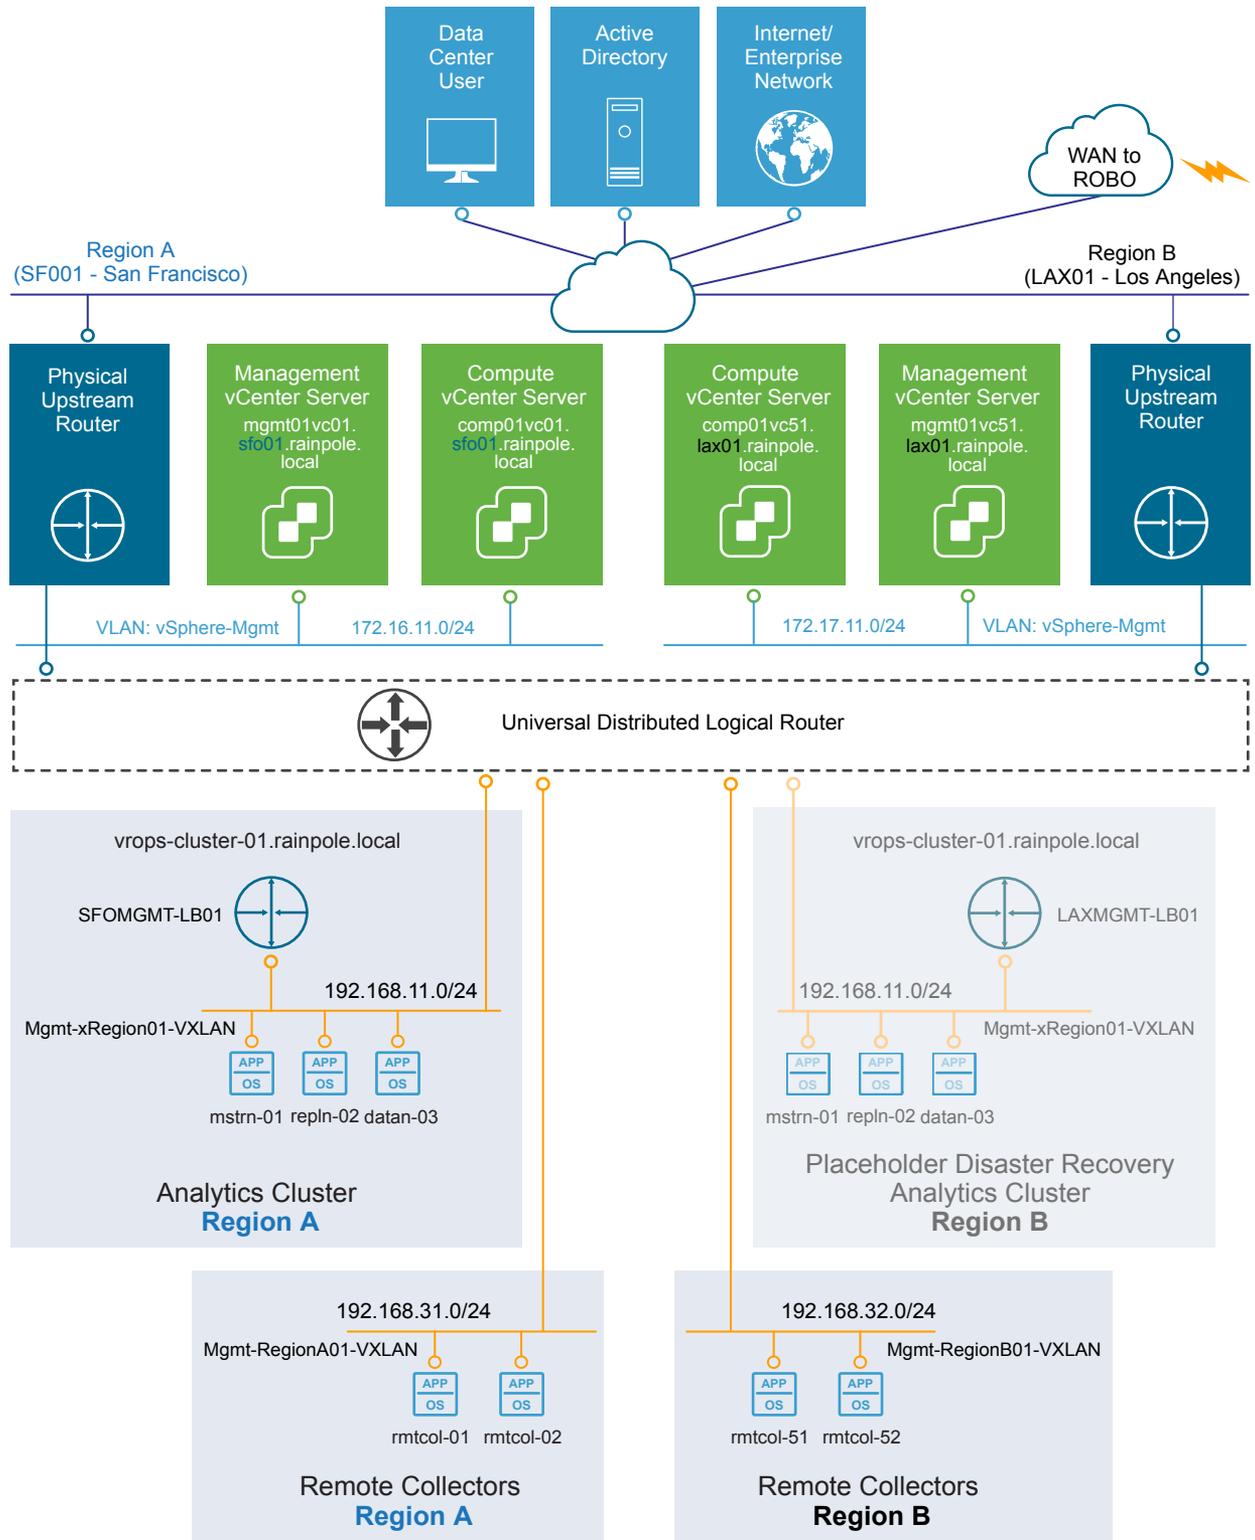
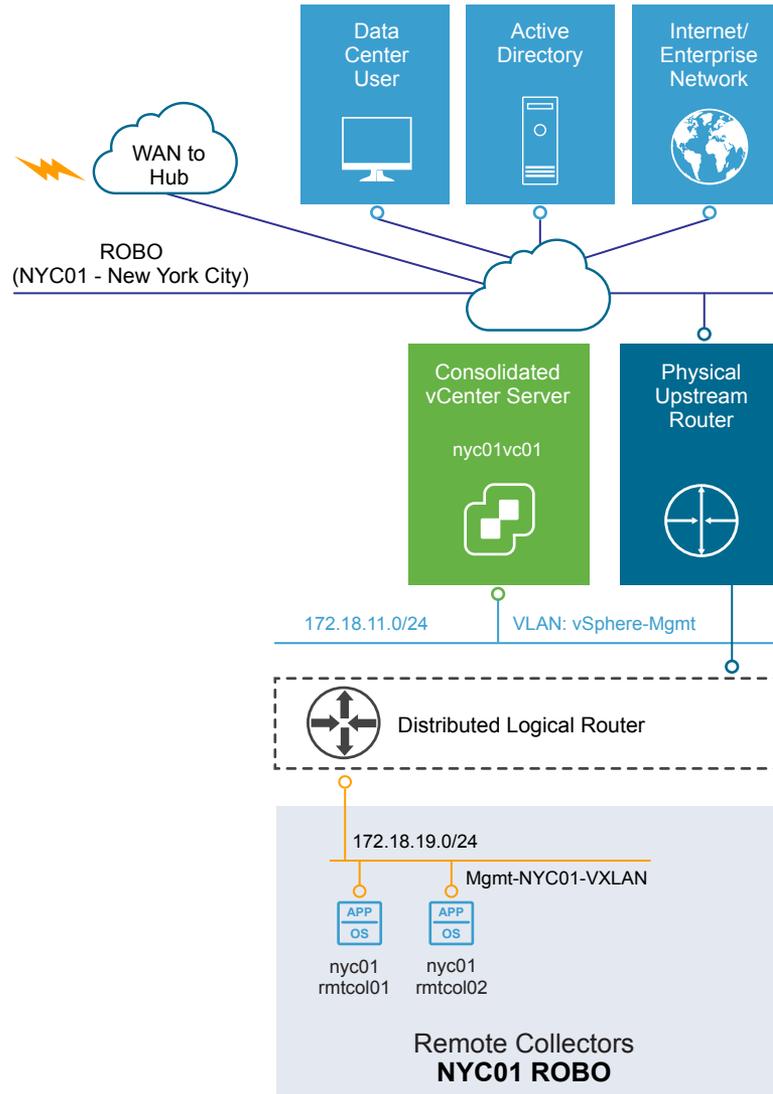


Figure 3-30. Networking Design of vRealize Operations Manager at the Remote office and Branch Office Site



Application Isolated Network Design in ROBO

The vRealize Operations Manager remote collector nodes are installed in their ROBO-specific application virtual network.

This networking design has the following features:

- All nodes have routed access to the vSphere management network through the NSX Distributed Logical Router.
- Routing to the vSphere management network and other external networks is dynamic, and is based on the Border Gateway Protocol (BGP).

For more information about the networking configuration of the application isolated network, see [Virtualization Network Design in ROBO](#) and [NSX Design](#).

Table 3-101. vRealize Operations Manager Isolated Network Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-MON-003	Use the existing ROBO-specific application virtual network for vRealize Operations Manager remote collectors.	Ensures collections of metrics locally per region in the event of a network outage. Additionally, it co-localizes metric collection to the per-ROBO SDDC applications using the virtual networks.	You must implement NSX to support this network configuration.

IP Subnets in ROBO

You can allocate the following example subnets to the vRealize Operations Manager Remote Collector deployment. As each new ROBO deployment is setup, you can increase the second octet of the IP subnet.

Table 3-102. IP Subnets in the Application Virtual Network of vRealize Operations Manager

vRealize Operations Manager Cluster Type	IP Subnet
Analytics cluster in hub (also valid for failover)	192.168.11.0/24
Remote collectors in Region A	192.168.31.0/24
Remote collectors in Region B	192.168.32.0/24
Remote collectors in ROBO NYC01	172.18.19.0/24
Remote collections in ROBO <i>Next01</i>	172.19.19.0/24

Note These IP subnets are only examples. Your implementation will vary based on available IP subnets within your environment.

Table 3-103. IP Subnets Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-MON-04	Allocate separate subnets for each application virtual network.	Placing the remote collectors in the management VXLAN in the ROBO enables all management applications to be on the same IP subnet.	None.

Management Packs in ROBO

The SDDC contains several VMware products for network, storage, and cloud management. You can monitor and perform diagnostics on all of them in vRealize Operations Manager by using management packs.

Table 3-104. Management Packs for vRealize Operations Manager Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-MON-005	Configure the remote collectors in each ROBO to collect data using the management packs for the solutions in that ROBO.	Provides local data collection of metrics from management packs.	The management packs available are dictated by what has been installed in the vRealize Operation Manager Analytics Cluster in the Hub.

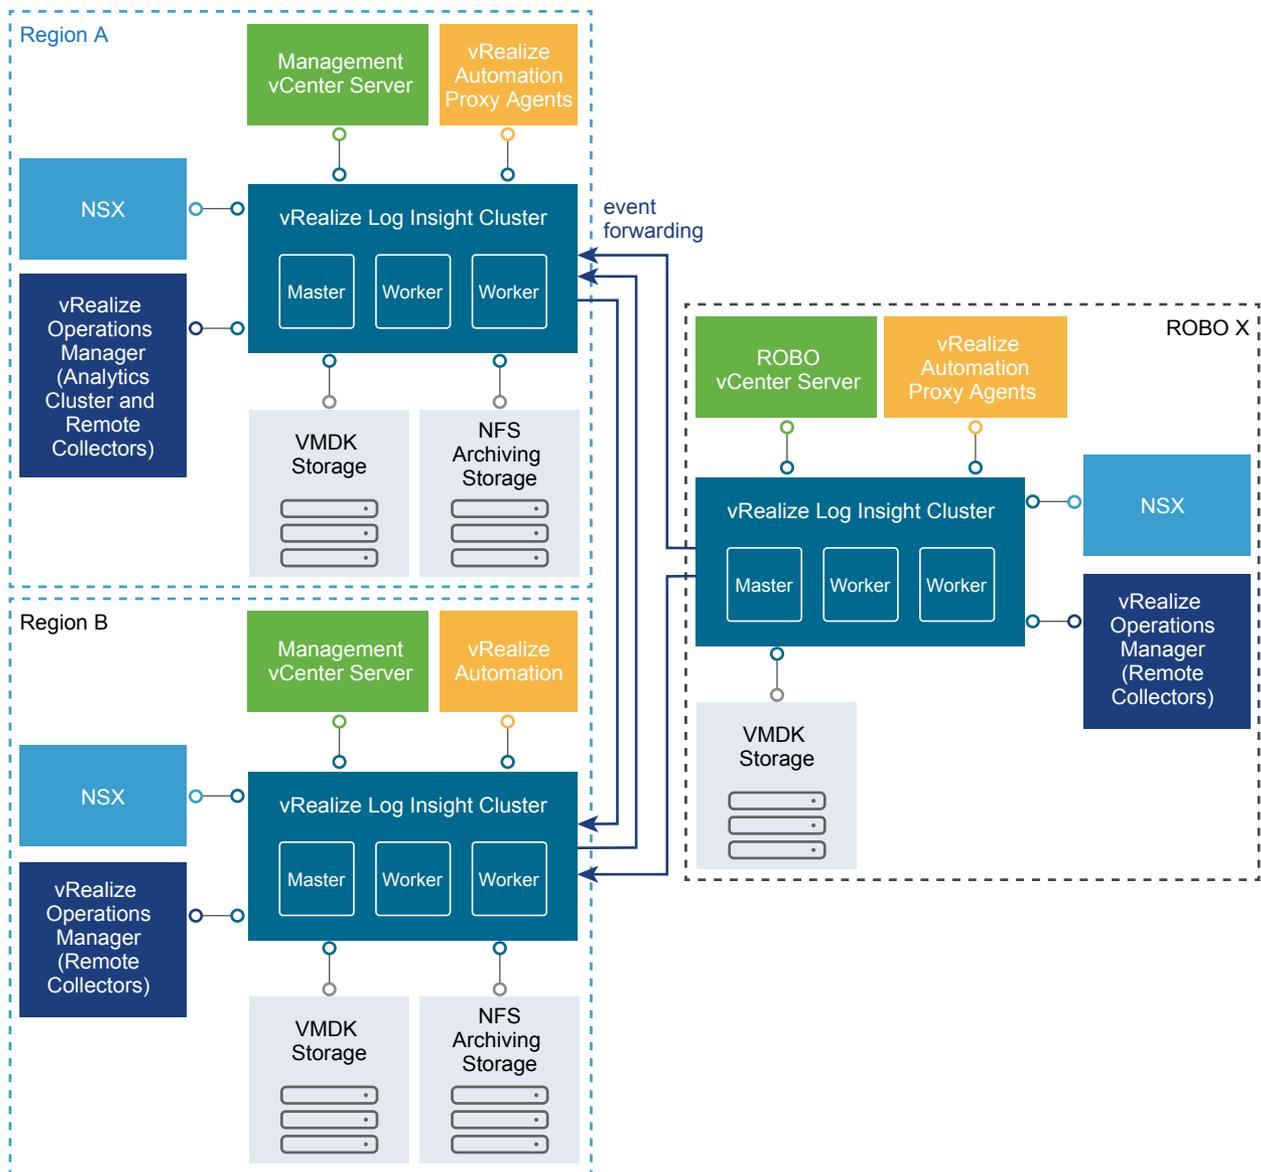
vRealize Log Insight Design in ROBO

vRealize Log Insight design enables real-time logging for all components that build up the management capabilities of the SDDC.

Logical Design

In a multi-region Software Defined Data Center (SDDC) deploy a vRealize Log Insight cluster in each region and in each ROBO consisting of a minimum of three nodes. This configuration allows for continued availability and increased log ingestion rates. The ROBO forwards log data to each Region, this allows a single pane of glass view of all logging data when viewed from either region.

Figure 3-31. Logical Design of vRealize Log Insight



Sources of Log Data

vRealize Log Insight collects logs as to provide monitoring information about the SDDC from a central location.

vRealize Log Insight collects log events from the following virtual infrastructure and cloud management components in the ROBO.

- Consolidated Pod
 - vCenter Server
 - ESXi Hosts
- NSX for vSphere
 - NSX Manager
 - NSX Controller instances
 - NSX Edge instances
- vRealize Automation
 - vRealize Proxy Agent Servers
- vRealize Operations Manager
 - vRealize Operations Manager Remote Collectors

Cluster Nodes in ROBO

The vRealize Log Insight cluster consists of one master node and two worker nodes. You enable the Integrated Load Balancer (ILB) on the cluster to have vRealize Log Insight to balance incoming traffic fairly among available nodes.

vRealize Log Insight clients, using both the Web user interface, and ingestion through syslog or the Ingestion API, connect to vRealize Log Insight that the ILB addresses.

vRealize Log Insight cluster can scale out to 12 nodes, that is, one master and 11 worker nodes.

Table 3-105. Cluster Node Configuration Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-001	Deploy vRealize Log Insight in a cluster configuration of 3 nodes with an integrated load balancer: one master and two worker nodes.	Provides high availability. Using the integrated load balancer simplifies the Log Insight deployment, and prevents from a single point of failure.	<ul style="list-style-type: none"> ■ You must size each node identically. ■ If the capacity requirements for your vRealize Log Insight cluster grow, identical capacity must be added to each node.
ROBO-OPS-LOG-002	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules to the vRealize Log Insight cluster components	Using DRS prevents vRealize Log Insight nodes from on the same ESXi host and thereby risking the cluster's high availability capability.	<ul style="list-style-type: none"> ■ Additional configuration is required to set up anti-affinity rules. ■ Only a single ESXi host in the initial ROBO cluster, of the four ESXi hosts, will be able to be put into maintenance mode at a time.

Sizing Log Insight Nodes in ROBO

To accommodate all log data from the products in the ROBO SDDC, you must size the compute resources and storage for the Log Insight nodes properly.

By default, the vRealize Log Insight virtual appliance uses the predefined values for small configurations, which have 4 vCPUs, 8 GB of virtual memory, and 510 GB of disk space provisioned. vRealize Log Insight uses 100 GB of the disk space to store raw data, index, metadata, and other information.

Sizing Nodes

Select a size for the vRealize Log Insight nodes to collect and store log data from the SDDC management components and tenant workloads according to the objectives of this design.

Table 3-106. Compute Resources for a vRealize Log Insight Small-Size Node

Attribute	Specification
Appliance size	Small
Number of CPUs	4
Memory	8 GB
Disk Capacity	510 GB (490 GB for event storage)
IOPS	500 IOPS
Amount of processed log data when using log ingestion	30 GB/day of processing per node
Number of processed log messages	2,000 event/second of processing per node
Environment	Up to 100 syslog connections per node

Sizing Storage

Sizing is based on IT organization requirements, but this design provides calculations according based on a single region implementation, and is implemented on a per-region basis. This sizing is calculated according to the following node configuration per region:

- Consolidated vCenter Server
- ESXi hosts
- NSX for vSphere
 - NSX Manager
 - NSX Controller instances
 - NSX Edge instances
- vRealize Automation
 - vRealize Agent Servers
- Tenant workloads

These components aggregate to approximately 150 syslog and vRealize Log Insight Agent sources. Assuming that you want to retain 7 days of data, use the following calculations:

For 150 syslog sources at a basal rate of 150 MB of logs ingested per-day per-source over 7 days, you need the following storage space:

```
150 sources * 150 MB of log data ≈ 22.5 GB log data per-day

22.5 GB * 7 days ≈ 157.5 GB log data per vRealize Log Insight node

157.5 GB * 1.7 indexing overhead ≈ 267.75 GB
```

Based on this example, the storage space that is allocated per small-size vRealize Log Insight virtual appliance is enough to monitor the ROBO.

Consider the following approaches when you must increase the Log Insight capacity:

- If you must maintain a log data retention for more than 7 days in your SDDC, you might add more storage per node by adding a new virtual hard disk. vRealize Log Insight supports virtual hard disks of up to 2 TB. If you must add more than 2 TB to a virtual appliance, add another virtual hard disk.

When you add storage to increase the retention period, this must be applied to all available virtual appliances.

Note Do not extend existing retention virtual disks. Once provisioned, do not reduce the size or remove virtual disks to avoid data loss.

- If you must monitor more components by using log ingestion and exceed the number of syslog connections or ingestion limits defined in this design, you can deploy more vRealize Log Insight virtual appliances to scale out your environment. vRealize Log Insight can scale up to 12 nodes in an HA cluster.

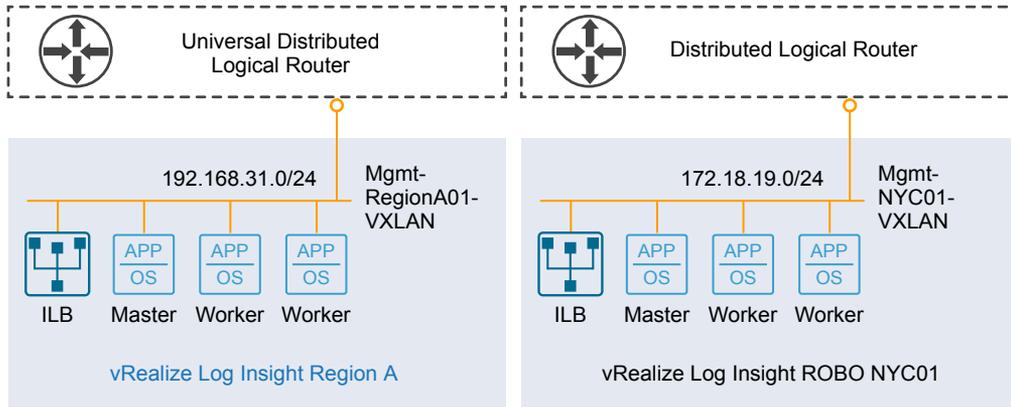
Table 3-107. Compute Resources for the vRealize Log Insight Nodes Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-003	Deploy vRealize Log Insight nodes of small size.	Accommodates the number of expected Syslog and vRealize Log Insight Agent connections from the following. This is approximately 150 syslog and vRealize Log Insight Agent sources. This ensure the storage space for the vRealize Log Insight cluster is sufficient for 7 days of data retention.	You must increase the size of the nodes if you configure Log Insight to monitor additional syslog sources.

vRealize Log Insight Networking Design in ROBO

In all regions and ROBOs, the vRealize Log Insight instances are connected to the region-specific management VXLANs.

Figure 3-32. Networking Design for the vRealize Log Insight Deployment



Application Network Design

This networking design has the following features:

- All nodes have routed access to the vSphere management network through a NSX logical router.
- Routing to the vSphere management network and the external network is dynamic, and is based on the Border Gateway Protocol (BGP).

For more information about the networking configuration of the application isolated networks for vRealize Log Insight, see [Application Virtual Network in ROBO](#) and [Virtual Network Design Example in ROBO](#).

Table 3-108. vRealize Log Insight Network Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-004	Deploy vRealize Log Insight on the region-specific application virtual networks.	<ul style="list-style-type: none"> ■ Ensures centralized access to log data per region if a cross-region network outage occurs. ■ Co-located log collection to the region local SDDC applications using the region-specific application virtual networks. ■ Provides a consistent deployment model for management applications. 	<ul style="list-style-type: none"> ■ Interruption in the network can impact event forwarding between the vRealize Log Insight clusters and cause gaps in log data. ■ You must use NSX to support this network configuration.

IP Subnets

You can allocate the following example subnets to the vRealize Log Insight deployment. As each new ROBO deployment is setup, you can increase the second octet of the IP subnet.

Table 3-109. IP Subnets in the Application Isolated Networks

vRealize Log Insight Cluster	IP Subnet
Region A	192.168.31.0/24
Region B	192.168.32.0/24

Table 3-109. IP Subnets in the Application Isolated Networks (Continued)

vRealize Log Insight Cluster	IP Subnet
ROBO NYC01	172.18.19.0/24
ROBO <i>Next01</i>	172.19.19.0/24

Table 3-110. DNS Names Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-005	Configure forward and reverse DNS records for all vRealize Log Insight nodes and VIPs.	All nodes are accessible by using fully-qualified domain names instead of by using IP addresses only.	You must manually provide a DNS record for each node and VIP.

vRealize Log Insight Retention in ROBO

Configure archive and retention parameters of vRealize Log Insight according to the company policy for compliance and governance.

vRealize Log Insight virtual appliances contain three default virtual disks and can use additional virtual disks for storage.

Table 3-111. Virtual Disk Configuration in the vRealize Log Insight Virtual Appliance

Hard Disk	Size	Usage
Hard disk 1	20 GB	Root file system
Hard disk 2	510 GB for small-size deployment	Contains two partitions: <ul style="list-style-type: none"> ■ /storage/var System logs ■ /storage/core Storage for Collected logs.
Hard disk 3	512 MB	First boot only

Calculate the storage space that is available for log data using the following equation:

$$\text{/storage/core} = \text{hard disk 2 space} - \text{system logs space on hard disk 2}$$

Based on the size of the default disk, the storage core is equal to 490 GB.

$$\begin{aligned} \text{/storage/core} &= 510\text{GB} - 20\text{ GB} = 490\text{ GB} \\ \text{Retention} &= \text{/storage/core} - 3\% * \text{/storage/core} \end{aligned}$$

If /storage/core is 490 GB, vRealize Log Insight can use 475 GB for retention.

$$\text{Retention} = 490\text{ GB} - 3\% * 490 \approx 475\text{ GB}$$

Configure a retention period of 7 days for the small-size vRealize Log Insight appliance.

Table 3-112. Retention Period Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-006	Configure vRealize Log Insight to retain data for 7 days.	Accommodates logs from 300 syslog sources (100 per node).	Disk size allows 750 syslog sources with 7 days retention, although the small appliance size limits each node to 100 connections or 300 total.

If you must support a log retention period that is greater than amount of available storage, additional storage may be added to each virtual appliance. When adding storage to increase the retention, the capacity supplied must be identical across all virtual appliances.

vRealize Log Insight Alerting in ROBO

vRealize Log Insight supports alerts that trigger notifications about its health.

Alert Types

The following types of alerts exist in vRealize Log Insight:

System Alerts	vRealize Log Insight generates notifications when an important system event occurs, for example when the disk space is almost exhausted and vRealize Log Insight must start deleting or archiving old log files.
Content Pack Alerts	Content packs contain default alerts that can be configured to send notifications, these alerts are specific to the content pack and are disabled by default.
User-Defined Alerts	Administrators and users can define their own alerts based on data ingested by vRealize Log Insight. vRealize Log Insight handles alerts in two ways: <ul style="list-style-type: none"> ■ Send an e-mail via SMTP. ■ Send to vRealize Operations Manager. ■ Send a HTTP POST via Webhooks.

SMTP Notification

Enable e-mail notification for alerts in vRealize Log Insight.

Table 3-113. SMTP Alert Notification Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-007	Configure SMTP to send alerts via email.	<ul style="list-style-type: none"> ■ Enables administrators and operators to receive system alerts via email from vRealize Log Insight. ■ Allows for administrators and operations to enable user alerts to send email from vRealize Log Insight 	Requires access to an external SMTP server.

Integration with vRealize Operations Manager

vRealize Log Insight integrates with vRealize Operations Manager to provide a central location for monitoring and diagnostics. You can use the following integration points that you can enable separately:

Notification Events Forward notification events from vRealize Log Insight to vRealize Operations Manager.

Table 3-114. Integration with vRealize Operations Manager Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-008	Forward alerts to vRealize Operations Manager.	Provides monitoring and alerting information that is pushed from vRealize Log Insight to vRealize Operations Manager for centralized administration.	You must install the vRealize Log Insight management pack into vRealize Operation Manager. This management pack is packaged with vRealize Operations Manager 6.0 and later.

vRealize Log Insight Security and Authentication in ROBO

Protect the vRealize Log Insight deployment by providing centralized role-based authentication and secure communication with the other components in the ROBO SDDC.

Authentication

Enable role-based access control in vRealize Log Insight by using the existing rainpole.local Active Directory domain.

Table 3-115. Custom Role-Based User Management Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-009	Use Active Directory for authentication.	Provides fine-grained role and privilege-based access for administrator and operator roles.	<ul style="list-style-type: none"> ■ You must provide access to the Active Directory from all Log Insight nodes. ■ Additional administrative overhead required for maintaining role-based access control configurations between ROBOs as they are not replicated between vRealize Log Insight instances.
ROBO-OPS-LOG-010	Configure Active Directory authentication to specifically use Active Directory Domain Controller(s) located within the ROBO	<ul style="list-style-type: none"> ■ Co-location of Active Directory Domain Controllers to the vRealize Log Insight cluster prevent users from being unable to authenticate in the event of a WAN outage. ■ Co-location of Active Directory Domain Controllers to the vRealize Log Insight cluster ensures the optimal authentication route is taken for ROBO users, providing better authentication performance. 	You must have Active Directory Domain Controller(s) located in your ROBO site.

Encryption

Replace default self-signed certificates with a CA-signed certificate to provide secure access to the vRealize Log Insight Web user interface.

Table 3-116. Custom Certificates Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-011	Replace the default self-signed certificates with a CA-signed certificate.	Configuring a CA-signed certificate ensures that all communication to the externally facing Web UI is encrypted. This allows for encrypted log forwarding.	The administrator must have access to a Public Key Infrastructure (PKI) to acquire certificates.

Configuration for Collecting Logs in ROBO

As part of vRealize Log Insight configuration, you configure syslog and vRealize Log Insight agents.

Client applications can send logs to vRealize Log Insight in one of the following ways:

- Directly to vRealize Log Insight over the syslog protocol
- By using vRealize Log Insight to directly query the vSphere Web Server APIs
- By using a vRealize Log Insight Agent

Table 3-117. Direct Log Communication to vRealize Log Insight Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-012	Configure syslog sources to send log data directly to vRealize Log Insight.	Simplifies the design implementation for log sources that are syslog capable.	You must configure syslog sources to forward logs to the vRealize Log Insight VIP.
ROBO-OPS-LOG-013	Configure the vRealize Log Insight agent for the vRealize Automation Proxy agents.	<ul style="list-style-type: none"> ■ Windows does not natively support syslog. ■ vRealize Automation requires the use of agents to collect all vRealize Automation logs. 	You must manually install and configure the agents on the vRealize Automation Proxy agent nodes to forward logs to the vRealize Log Insight VIP.
ROBO-OPS-LOG-014	Configure the vRealize Log Insight agent for the vRealize Operations Manager Remote Collectors	Simplifies the design implementation for log sources that are pre-installed with vRealize Log Insight agent.	You must manually configure the agents on the vRealize Operations Remote Collectors to forward logs to the vRealize Log Insight VIP.

Table 3-117. Direct Log Communication to vRealize Log Insight Design Decisions (Continued)

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-015	Configure the vCenter Server Appliance as syslog sources to send log data directly to vRealize Log Insight.	Simplifies the design implementation for log sources that are syslog capable.	<ul style="list-style-type: none"> ■ You must manually configure syslog sources to forward logs to the vRealize Log Insight VIP. ■ Certain dashboards within vRealize Log Insight require the use of the vRealize Log Insight Agent for proper ingestion. ■ Not all Operating System-level events are forwarded to vRealize Log Insight.
ROBO-OPS-LOG-016	Configure vRealize Log Insight to ingest events, tasks, and alarms from the vCenter Server.	Ensures that all tasks, events and alarms generated from the vCenter Server instance are captured and analyzed for the administrator.	You must create a service account on vCenter Server to connect vRealize Log Insight for events, tasks, and alarms pulling.

Time Synchronization in ROBO

Time synchronization is critical for the core functionality of vRealize Log Insight. By default, vRealize Log Insight synchronizes time with a predefined list of public NTP servers.

NTP Configuration

Configure consistent NTP sources on all systems that send log data (vCenter Server, ESXi, vRealize Operation Manager). See *Time Synchronization* in the *Planning and Preparation* document.

Table 3-118. Time Synchronization Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-017	Configure consistent NTP sources on all virtual infrastructure and cloud management applications for correct log analysis in vRealize Log Insight.	Guarantees accurate log timestamps.	Requires that all applications synchronize time to the same NTP time source.
ROBO-OPS-LOG-018	Configure NTP source(s) that are located within the ROBO	<ul style="list-style-type: none"> ■ Co-location of NTP source(s) to the vRealize Log Insight cluster guarantees time accuracy across the nodes in the event of a WAN outage. ■ Co-location of NTP source(s) to the vRealize Log Insight cluster ensures time sync between all nodes is as accurate as possible. 	You must have NTP source(s) located in your ROBO site.

Cluster Communication

All vRealize Log Insight cluster nodes must be in the same LAN with no firewall or NAT between the nodes.

External Communication

vRealize Log Insight receives log data over the syslog TCP, syslog TLS/SSL, or syslog UDP protocols. Use the default syslog UDP protocol because security is already designed at the level of the management network.

Table 3-119. Syslog Protocol Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-019	Communicate with the syslog clients, such as ESXi, vCenter Server, NSX for vSphere, on the default UDP syslog port.	Using the default syslog port simplifies configuration for all syslog sources.	<ul style="list-style-type: none"> ■ If the network connection is interrupted, the syslog traffic is lost. ■ UDP syslog traffic is not secure.

vRealize Log Insight Event Forwarding to Regions from ROBO

vRealize Log Insight supports event forwarding to other clusters and standalone instances. While forwarding events, the vRealize Log Insight instance still ingests and stores events locally.

You forward logging data in vRealize Log Insight by using the Ingestion API or a native syslog implementation. In the VMware Validated Design for ROBO, we are using the Ingestion API.

The vRealize Log Insight Ingestion API uses TCP communication. In contrast to syslog streaming, the forwarding module supports the following features for the Ingestion API.

- Sending of encrypted events.
- Both structured and unstructured data, that is, multi-line messages.
- Metadata in the form of tags.
- Client-side compression.
- Configurable disk-backed queue to save events until the server acknowledges the ingestion.

Table 3-120. Protocol for Event Forwarding Across Regions Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-LOG-020	Forward log event to both regions by using the Ingestion API.	Using the forwarding protocol supports structured and unstructured data provides client-side compression, and event throttling to be passed from one vRealize Log Insight cluster to the other. Forwarding ensures that during a disaster recovery situation the administrator has access to all logs from the ROBO although the ROBO may be offline. Forwarding enables a single pane of glass view for Logs across all regions and ROBO sites.	<ul style="list-style-type: none"> ■ You must configure each ROBO to forward log data to both regions. The configuration requires administrative overhead to prevent recursion of logging between regions via exclusion of the tags used. This introduces the potential for a single forwarder to fail from the ROBO site creating a gap in logging data in the region it is forwarding to. If the connectivity is not restored in an adequate amount of time such that the disk cache begins dropping events, the logging data will not be supplemented once the connection to that region is restored due to the exclusion rule. If this occurs, data sets in the regions will not be identical. ■ Log forwarding from ROBO adds more load on each region. You must consider log forwarding in the sizing calculations for the vRealize Log Insight clusters in each region in the VVD for SDDC as each ROBO is deployed.
ROBO-OPS-LOG-021	Configure log forwarding to use SSL.	Ensures that the log forwarding operations from each ROBO to the regions has provable identity and is secure.	The source forwarder must be configured to trust the destination server's signing certificate authority.
ROBO-OPS-LOG-022	Configure disk cache for event forwarding to 2,000 MB (2 GB) for buffering.	Ensures that log forwarding has a buffer for approximately 2 hours if a connectivity outage occurs. The disk cache size is calculated at a base rate of 150 MB per day per syslog source with 105 syslog sources.	<ul style="list-style-type: none"> ■ If the event forwarder of vRealize Log Insight is restarted during the communication outage, messages that reside in the non-persistent cache will be cleared. ■ If a communication outage exceeds 2 hours, the newest local events are dropped and not forwarded to the remote destination even after the connection is restored.

vSphere Data Protection Design in ROBO

Design data protection of the management components in your environment to ensure continuous operation of the ROBO SDDC if the data of a management application is damaged.

Data backup protects the data of your organization against data loss, hardware failure, accidental deletion, or other disaster for each region. For consistent image-level backups, use backup software that is based on the vSphere APIs for Data Protection (VADP). This design uses vSphere Data Protection as an example. You can use any VADP compatible software. Adapt and apply the design decisions to the backup software you use.

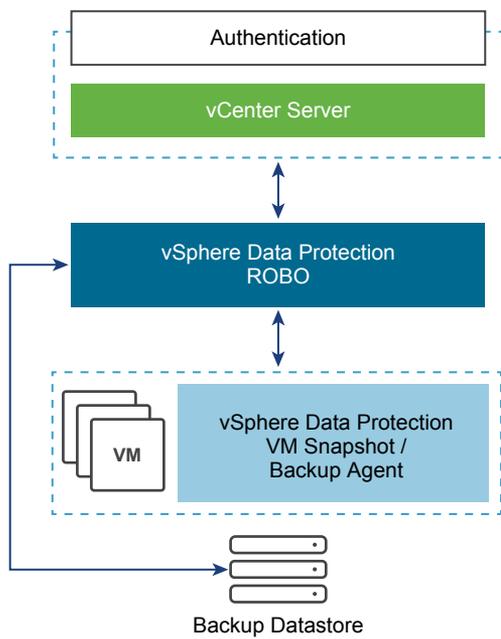
Table 3-121. vSphere Data Protection Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-BKP-001	Use VADP compatible backup software to back up all management components such as vSphere Data Protection.	vSphere Data Protection provides the functionality that is required to back up full image VMs and applications in those VMs, for example, Microsoft SQL Server.	vSphere Data Protection lacks some features that are available in other backup solutions.

Logical Design in ROBO

vSphere Data Protection protects the virtual infrastructure at the VMware vCenter Server layer. Because vSphere Data Protection is connected to the vCenter Server, it can access all ESXi hosts, and can detect the virtual machines that require backups.

Figure 3-33. vSphere Data Protection Logical Design



Backup Datastore in ROBO

The backup datastore stores all the data that is required to recover services according to a Recovery Point Objective (RPO). Determine the target location and make sure that it meets performance requirements.

vSphere Data Protection uses deduplication technology to back up virtual environments at data block level, which enables efficient disk utilization. To optimize backups and leverage the VMware vSphere Storage APIs, all ESXi hosts must have access to the production storage.

Table 3-122.
Options for Backup Storage Location

Option	Benefits	Drawbacks
Store production and backup data on the same storage platform.	<ul style="list-style-type: none"> ■ You do not have to request a new storage configuration from the storage team. ■ You can take full advantage of vSphere capabilities. 	You cannot recover your data if the destination datastore or the production storage is unrecoverable.
Store backup data on dedicated storage.	<ul style="list-style-type: none"> ■ If production storage becomes unavailable, you can recover your data because your backup data is not located on the same shared storage. ■ You separate production and backup virtual machines. ■ The backup schedule does not impact production storage performance because the backup storage is completely separate. 	You might be required to install and configure a dedicated storage volume for backups.

Table 3-123. VMware Backup Store Target Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-BKP-002	Allocate a secondary storage datastore for the vSphere Data Protection appliance and the backup data.	<ul style="list-style-type: none"> ■ vSphere Data Protection emergency restore operations are possible even when the primary VMware vSAN datastore is not available because the vSphere Data Protection storage volume is separate from the primary vSAN datastore. ■ The amount of storage required for backups is greater than the amount of storage available in the vSAN datastore. 	You must provide secondary storage.

Performance in ROBO

vSphere Data Protection generates a significant amount of I/O operations, especially when performing multiple concurrent backups. The storage platform must be able to handle this I/O. If the storage platform does not meet the performance requirements, it might miss backup windows. Backup failures and error messages might occur. Run the vSphere Data Protection performance analysis feature during virtual appliance deployment or after deployment to assess performance.

Table 3-124. VMware vSphere Data Protection Performance

Total Backup Size	Avg Mbps in 4 hours
0.5 TB	306 Mbps
1 TB	611 Mbps
2 TB	1223 Mbps

Volume Sizing in ROBO

vSphere Data Protection can dynamically expand the destination backup store from 2 TB to 8 TB. Using an extended backup storage requires additional memory on the vSphere Data Protection appliance.

Table 3-125. VMware vSphere Data Protection Sizing Guide

Available Backup Storage Capacity	Size On Disk	Minimum Appliance Memory
0.5 TB	0.9 TB	4 GB
1 TB	1.6 TB	4 GB
2 TB	3 TB	6 GB
4 TB	6 TB	8 GB
6 TB	9 TB	10 GB
8 TB	12 TB	12 GB

Table 3-126. VMware Backup Store Size Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-BKP-003	Deploy the vSphere Data Protection virtual appliance initially for 2 TB of available backup storage capacity and (3 TB on-disk size.)	Handles the backup of the Management stack. The management stack currently consumes approximately 1 TB of disk space, uncompressed and without deduplication.	You must provide more secondary storage to accommodate increased disk requirements.

Other Considerations in ROBO

vSphere Data Protection can protect virtual machines that reside on VMware vSAN datastores from host failures. The virtual machine storage policy is not backed up with the virtual machine, but you can restore the storage policy after restoring the virtual machine.

Note The default vSAN storage policy includes `Number Of Failures To Tolerate = 1`, which means that virtual machine data is mirrored.

You use vSphere Data Protection to restore virtual machines that fail or whose data must be reverted to a previous state.

Backup Policies in ROBO

Use vSphere Data Protection backup policies to specify virtual machine backup options, the schedule window, and retention policies.

Virtual Machine Backup Options

vSphere Data Protection provides the following options for a virtual machine backup:

- HotAdd** Provides full image backups of virtual machines, regardless of the guest operating system.
- The virtual machine base disk is attached directly to vSphere Data Protection to back up data. vSphere Data Protection uses Changed Block Tracking to detect and back up blocks that are altered.

- The backup and restore performance is faster because the data flow is through the VMkernel layer instead of over a network connection.
- A quiesced snapshot can be used to redirect the I/O of a virtual machine disk .vmdk file.
- HotAdd does not work in multi-writer disk mode.

Network Block Device (NBD)

Transfers virtual machine data across the network to allow vSphere Data Protection to back up the data.

- The performance of the virtual machine network traffic might be lower.
- NBD takes a quiesced snapshot. As a result, it might interrupt the I/O operations of the virtual machine to swap the .vmdk file or consolidate the data after the backup is complete.
- The time to complete the virtual machine backup might be longer than the backup window.
- NBD does not work in multi-writer disk mode.

vSphere Data Protection Agent Inside Guest OS

Provides backup of certain applications that are running in the guest operating system through an installed backup agent.

- Enables application-consistent backup and recovery with Microsoft SQL Server, Microsoft SharePoint, and Microsoft Exchange support.
- Provides more granularity and flexibility to restore on the file level.

Table 3-127. Virtual Machine Transport Mode Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-BKP-004	Use HotAdd to back up virtual machines.	HotAdd optimizes and speeds up virtual machine backups, and does not impact the vSphere management network.	All ESXi hosts need to have the same visibility of the virtual machine datastores.

Schedule Window

Even though vSphere Data Protection uses the Changed Block Tracking technology to optimize the backup data, to avoid any business impact, do not use a backup window when the production storage is in high demand.

Caution Do not perform any backup or other administrative activities during the vSphere Data Protection maintenance window. You can only perform restore operations. By default, the vSphere Data Protection maintenance window begins at 8 PM local server time and continues uninterrupted until 8 AM or until the backup jobs are complete. Configure maintenance windows according to IT organizational policy requirements.

Table 3-128. Backup Schedule Design Decisions

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-BKP-005	Schedule daily backups.	Allows for the recovery of virtual machines data that is at most a day old	Data that changed since the last backup, 24 hours ago, is lost.
ROBO-OPS-BKP-006	Schedule backups outside the production peak times.	Ensures that backups occur when the system is under the least amount of load. You should verify that backups are completed in the shortest time possible with the smallest risk of errors.	Backups need to be scheduled to start between 8:00 PM and 8:00 AM or until the backup jobs are complete, whichever comes first.

Retention Policies

Retention policies are properties of a backup job. If you group virtual machines by business priority, you can set the retention requirements according to the business priority.

Table 3-129. Retention Policies Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-BKP-007	Retain backups for at least 3 days.	Keeping 3 days of backups enables administrators to restore the management applications to a state within the last 72 hours.	Depending on the rate of change in virtual machines, backup retention policy can increase the storage target size.

Component Backup Jobs in ROBO

You can configure backup for each ROBO SDDC management component separately. For this scenario, no requirement to back up the entire ROBO SDDC exists, and this design does not imply such an operation. Some products can perform internal configuration backups. Use those products in addition to the whole VM component backups as appropriate.

Table 3-130. Component Backup Jobs Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-BKP-008	Use the internal configuration backup features within VMware NSX.	Restoring small configuration files can be a faster and less destructive method to achieve a similar restoration of functionality.	An FTP server is required for the NSX configuration backup.

Backup Jobs in ROBO locations

Create a single backup job for the components of a management application according to the node configuration of the application in Region A.

Table 3-131. VM Backup Jobs in ROBO Locations

Product	Image VM Backup Jobs in ROBO NYC01
ESXi	Backup is not applicable
vCenter Server	<ul style="list-style-type: none"> ■ nyc01vc01.rainpole.local
NSX for vSphere	<ul style="list-style-type: none"> ■ nyc01nsxm01.rainpole.local

Table 3-131. VM Backup Jobs in ROBO Locations (Continued)

Product	Image VM Backup Jobs in ROBO NYC01
vRealize Automation	<ul style="list-style-type: none"> ■ nyc01ias01.rainpole.local ■ nyc01ias02.rainpole.local ■ nyc01buc01.rainpole.local
vRealize Log Insight	<ul style="list-style-type: none"> ■ nyc01vrli01.rainpole.local ■ nyc01vrli02.rainpole.local ■ nyc01vrli03.rainpole.local
vRealize Operations Manager	<ul style="list-style-type: none"> ■ nyc01vroprmtcol01.rainpole.local ■ nyc01vroprmtcol02.rainpole.local
vRealize Business Data Collector	Part of the vRealize Automation backup job

vSphere Update Manager in ROBO

vSphere Update Manager pairs with vCenter Server to enable patch and version management of ESXi hosts and virtual machines.

vSphere Update Manager 6.5 can remediate the following objects:

- VMware Tools and VMware virtual machine hardware upgrade operations for virtual machines running on ESXi 5.0 or later
- ESXi host patching operations for hosts running ESXi 5.0 or later
- ESXi host upgrade operations for hosts running ESXi 5.0 or later

vSphere Update Manager performs host remediation over the network. vSphere Update Manager must be connected to ESXi and vCenter Server.

Physical Design of vSphere Update Manager in ROBO

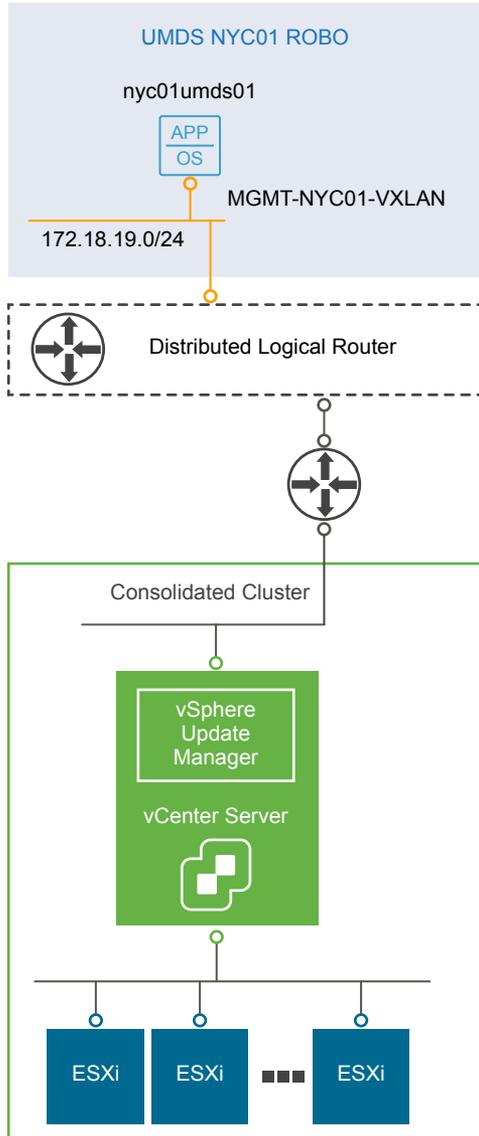
You use the vSphere Update Manager service on the vCenter Server Appliance and deploy a vSphere Update Manager Download Service (UMDS) to download and stage upgrade and patch data.

Networking and Application Design

You can use the vSphere Update Manager 6.5 as a service of the vCenter Server Appliance 6.5. The Update Manager server and client components are a part of the vCenter Server Appliance

You can connect only one vCenter Server instance to a vSphere Update Manager instance. To restrict the access to the external network from vSphere Update Manager and vCenter Server, deploy a vSphere Update Manager Download Service (UMDS) in each ROBO. UMDS downloads upgrades, patch binaries and patch metadata, and stages the downloads on a Web server. The local Update Manager servers download the patches from UMDS.

Figure 3-34. vSphere Update Manager Logical and Networking Design



Deployment Model

vSphere Update Manager is embedded in the vCenter Server Appliance. After you deploy or upgrade the vCenter Server Appliance, the VMware vSphere Update Manager Extension service starts automatically.

In addition to vSphere Update Manager deployment, two models for downloading patches from VMware exist.

Internet-connected model The vSphere Update Manager server is connected to the VMware patch repository to download patches for ESXi 5.x hosts, ESXi 6.x hosts, and virtual appliances. No additional configuration is required, other than scan and remediate the hosts as needed.

Proxied access model vSphere Update Manager has no connection to the Internet and cannot download patch metadata. You install and configure UMDS to download and store patch metadata and binaries to a shared repository. vSphere Update Manager must be configured to use the shared repository as a patch datastore before remediating the ESXi hosts.

Table 3-132. Update Manager Physical Design Decision

Decision ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-VUM-001	Use the vSphere Update Manager service on the vCenter Server Appliance in each ROBO that you configure and use for patch management.	A one-to-one mapping of vCenter Server to vSphere Update Manager is required. Each ROBO vCenter Server needs their own vSphere Update Manager.	All physical design decisions for vCenter Server determine the setup for vSphere Update Manager.
ROBO-OPS-VUM-002	Use the embedded PostgreSQL of the vCenter Server Appliance for vSphere Update Manager.	Reduces both overhead and Microsoft or Oracle licensing costs. Avoids problems with upgrades.	The vCenter Server Appliance has limited database management tools for database administrators.
ROBO-OPS-VUM-003	Use the network settings of the vCenter Server Appliance for vSphere Update Manager.	Simplifies network configuration because of the one-to-one mapping between vCenter Server and vSphere Update Manager. You configure the network settings once for both vCenter Server and vSphere Update Manager.	None.
ROBO-OPS-VUM-004	Deploy and configure vSphere Update Manager Download Service virtual machines in every ROBO.	Limits direct access to the Internet from vSphere Update Manager vCenter Server instances.	None.
ROBO-OPS-VUM-005	Connect the UMDS virtual machines to the ROBO-specific application virtual network.	<ul style="list-style-type: none"> ■ Ensures local storage and access to vSphere Update Manager repository data ■ Provides a consistent deployment model for management applications. 	You must use NSX to support this network configuration.

Logical Design of vSphere Update Manager in ROBO

You configure vSphere Update Manager to apply updates on the management components of the SDDC according to the objectives of this design.

UMDS Virtual Machine Specification

You allocate resources to and configure the virtual machines for UMDS according to the following specification:

Table 3-133. vSphere Update Manager Download Service (UMDS) Virtual Machine Specifications

Attribute	Specification
vSphere Update Manager Download Service	vSphere 6.5
Number of CPUs	2
Memory	2 GB
Disk Space	120 GB
Operating System	Ubuntu 14.04 LTS

ESXi Host and Cluster Settings

When you perform updates by using the vSphere Update Manager, the update operation affects certain cluster and host base settings. You customize these settings according to your business requirements and use cases.

Table 3-134. Host and Cluster Settings That Are Affected by vSphere Update Manager

Settings	Description
Maintenance mode	During remediation, updates might require the host to enter maintenance mode. Virtual machines cannot run when a host is in maintenance mode. For availability during a host update, virtual machines are migrated to other ESXi hosts within a cluster before the host enters maintenance mode. However, putting a host in maintenance mode during update might cause issues with the availability of the cluster.
vSAN	<p>When using vSAN, consider the following factors when you update hosts by using vSphere Update Manager:</p> <ul style="list-style-type: none"> ■ Host remediation might take a significant amount of time to complete because, by design, only one host from a vSAN cluster can be in maintenance mode at any one time. ■ vSphere Update Manager remediates hosts that are a part of a vSAN cluster sequentially, even if you set the option to remediate the hosts in parallel. ■ If the number of failures to tolerate is set to 0 for the vSAN cluster, the host might experience delays when entering maintenance mode. The delay occurs because vSAN copies data between the storage devices in the cluster. <p>To avoid delays, set a vSAN policy where the number failures to tolerate is 1 .The number of failures to tolerate is 1 by default .</p>

You can control the update operation by using a set of host and cluster settings in vSphere Update Manager.

Table 3-135. Host and Cluster Settings for Updates

Level	Description
Host settings	<ul style="list-style-type: none"> VM Power state when entering maintenance mode. You can configure vSphere Update Manager to power off, suspend or do not control virtual machines during remediation. This option applies only if vSphere vMotion is not available for a host. Retry maintenance mode in case of failure. If a host fails to enter maintenance mode before remediation, vSphere Update Manager waits for a retry delay period and retries putting the host into maintenance mode as many times as you indicate. Will attempt to enter maintenance mode with the configured parameter settings, if an initial attempt fails. Allow installation of additional software on PXE-booted hosts. This option is limited to software packages that do not require a host reboot after installation.
Cluster settings	<ul style="list-style-type: none"> Disable vSphere Distributed Power Management (DPM), vSphere High Availability (HA) Admission Control, and Fault Tolerance (FT). Enable parallel remediation of hosts. vSphere Update Manager can remediate multiple hosts. <ul style="list-style-type: none"> Note Parallel remediation is not supported if you use vSAN. Migrate powered-off or suspended virtual machines. vSphere Update Manager migrates the suspended and powered-off virtual machines from hosts that must enter maintenance mode to other hosts in the cluster. The migration is launched on virtual machines that do not prevent the host from entering maintenance mode.

Virtual Machine and Virtual Appliance Update Settings

vSphere Update Manager supports remediation of virtual machines and appliances. You can control the virtual machine and appliance updates by using the following settings:

Table 3-136. vSphere Update Manager Settings for Remediation of Virtual Machines and Appliances

Configuration	Description
Take snapshots before virtual machine remediation	Test before you commit the changes.
Define the window in which a snapshot persists for a remediated virtual machine	Automatically clean up virtual machine snapshots that are taken before remediation.
Enable smart rebooting for VMware vSphere vApps remediation	Start virtual machines post remediation to maintain start-up dependencies no matter if some of the virtual machines are not remediated.

ESXi Image Configuration

You can store full images that you can use to upgrade ESXi hosts. You cannot download such images from the patch repositories and must upload them by using vSphere Update Manager. Import into the repository the ESXi builds that are available in the environment.

By using Image Builder, add the NSX software packages `esx-vdpi`, `esx-vsip` and `esx-vxlan` into the ESXi upgrade image so that you can use the hosts being upgraded in a software-defined networking setup.

Baselines and Groups

vSphere Update Manager baselines and baseline groups are collections of patches that can be assigned to a cluster or host entity in the environment. Depending on the business requirements, the default baselines might not be allowed until patches are tested or verified on development or pre-production hosts. Baselines can be confirmed so that the tested patches are applied to hosts and only updated when appropriate.

Two types of baselines exist:

- Dynamic baselines that can change as items are added to the repository.
- Fixed baselines that remain the same.

vSphere Update Manager contains the following default baselines. Each of these baselines is configured for dynamic selection of new items.

Critical host patches	Upgrades hosts with a collection of critical patches that are high priority as defined by VMware.
Non-critical host patches	Upgrades hosts with patches that are not classified as critical.
VMware Tools Upgrade to Match Host	Upgrades the VMware Tools version to match the host version.
VM Hardware Upgrade to Match Host	Upgrades the virtual machine hardware version to match the host version.
VA Upgrade to Latest	Upgrades a virtual appliance to the latest version available.

vSphere Update Manager Logical Design Decisions

This design applies the following decisions on the logical design of vSphere Update Manager and update policy:

Table 3-137. vSphere Update Manager Logical Design Decisions

Design ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-VUM-006	Use the default patch repositories by VMware.	No additional sources required.	None.
ROBO-OPS-VUM-007	Set the VM power state to Do Not Power Off.	Ensures highest uptime of management components and compute workload virtual machines.	Manual intervention will be required if migration fails.
ROBO-OPS-VUM-008	Enable parallel remediation of hosts assuming that there are enough resources available to support update of multiple hosts at the same time.	Remediation of host patches can occur more quickly.	More resources unavailable at the same time during remediation.
ROBO-OPS-VUM-009	Enable migration of powered-off virtual machines and templates.	Ensures that templates stored on all management hosts are accessible.	Increases the amount of time to start remediation for templates to be migrated.

Table 3-137. vSphere Update Manager Logical Design Decisions (Continued)

Design ID	Design Decision	Design Justification	Design Implication
ROBO-OPS-VUM-010	Use the default critical and non-critical patch baselines for the Consolidated cluster.	No customized baselines required.	All patches are added to the baselines as soon as they are released.
ROBO-OPS-VUM-011	Use the default schedule of a once-per-day check and patch download.	No change required for this engagement.	None.
ROBO-OPS-VUM-012	Remediate hosts, virtual machines, and virtual appliances once a month or per business guidelines.	Ensures up to date hosts, virtual machines and virtual appliances.	Schedule must be aligned to the business policies.
ROBO-OPS-VUM-013	Use Image Builder to add NSX for vSphere software packages to the ESXi upgrade image.	<ul style="list-style-type: none"> ■ Ensures that the ESXi hosts are ready for software-defined networking immediately after the upgrade. ■ Allows for parallel remediation of ESXi hosts. ■ Additional NSX remediation is not required. 	<ul style="list-style-type: none"> ■ You must enable the Image Builder service. ■ NSX for vSphere updates might require new ESXi images updates.
ROBO-OPS-VUM-014	Configure an HTTP Web server on each UMDS service that the connected vSphere Update Manager servers must use to download the patches from.	Without a Web service running, vSphere Update Manager is unable to download patches automatically from UMDS. The alternative is to copy media from one place to another manually.	You must be familiar with third party Web server such as Nginx or Apache.