

Architecture and Design for vRealize Suite 2019 on VMware Validated Design 5.1.1 or VMware Cloud Foundation 3.9.1

Early Access

27 Jan 2020

VMware Validated Design 5.1.1

VMware Cloud Foundation 3.9.1

vRealize Suite 2019



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 © 2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About Architecture and Design for VMware vRealize Suite 2019 on VMware Validated Design 5.1.1 or VMware Cloud Foundation 3.9.1 4

1 Architecture Overview 7

Identity and Access Management Architecture 7

Cloud Automation Architecture 7

VMware Cloud Assembly 8

vRealize Orchestrator Architecture 9

VMware Service Broker 9

VMware Code Stream 10

Cloud Operations Architecture 11

Monitoring Architecture 11

Logging Architecture 14

Life Cycle Architecture 17

2 Detailed Design 20

Identity and Access Management Design 20

Workspace ONE Access Design 20

Cloud Automation Design 48

vRealize Automation Design 48

Cloud Operations Design 111

vRealize Suite Lifecycle Manager Design 111

vRealize Operations Manager Design 135

vRealize Log Insight Design 168

About Architecture and Design for VMware vRealize Suite 2019 on VMware Validated Design 5.1.1 or VMware Cloud Foundation 3.9.1

The *Architecture and Design for VMware vRealize Suite 2019 on VMware Validated Design 5.1.1 or VMware Cloud Foundation 3.9.1* document contains a detailed design for adding and connecting the vRealize® Suite 2019 products to a Software-Defined Data Center (SDDC) deployment of VMware Validated Design™ 5.1.1 or VMware Cloud Foundation™ 3.9.1.

The bills of materials of VMware Validated Design 5.1.1 and VMware Cloud Foundation 3.9.1 include vRealize software products that are of versions earlier than the versions of the vRealize Suite 2019 products. This design substitutes the vRealize software products in the bills of materials with the vRealize Suite 2019 products.

Intended Audience

This design is intended for cloud architects and administrators who want to deploy and use vRealize Suite 2019 on an SDDC that is deployed according to VMware Validated Design or by using VMware Cloud Foundation.

Required VMware Software

Architecture and Design for VMware vRealize Suite 2019 on VMware Validated Design 5.1.1 or VMware Cloud Foundation 3.9.1 is compliant and validated with certain vRealize Suite 2019 products and components.

Table 1-1. vRealize Suite 2019 Products and Components

Product Group and Edition	Product	Version
VMware vRealize® Suite Lifecycle Manager™	vRealize Suite Lifecycle Manager	8.0.1
VMware vRealize® Operations Manager™ Advanced or higher	vRealize Operations Manager	8.0.1
	VMware SDDC Health Monitoring Solution™	8.0
	VMware vRealize® Operations Management Pack for NSX for vSphere	3.6
	VMware vRealize® Operations Management Pack for NSX-T™	2.3

Table 1-1. vRealize Suite 2019 Products and Components (continued)

Product Group and Edition	Product	Version
	VMware vRealize [®] Operations Management Pack for VMware Identity Manager [™] (Workspace ONE Access)	1.1
	VMware vRealize [®] Operations Management Pack for Storage Devices	8.0
	VMware vRealize [®] Operations Management Pack for Site Recovery Manager 8.1	8.2.0.2
VMware vRealize [®] Log Insight [™]	vRealize Log Insight	8.0
	VMware vRealize [®] Log Insight [™] Content Pack for NSX Data Center for vSphere	4.0
	VMware vRealize [®] Log Insight [™] Content Pack for NSX-T Data Center	3.8.2
	VMware vRealize [®] Log Insight [™] Content Pack for Linux	2.1
	VMware vRealize [®] Log Insight [™] Content Pack for Workspace ONE Access (Identity Manager)	2.0
	VMware vRealize [®] Log Insight [™] Content Pack for Site Recovery Manager 8.1+	2.1
VMware vRealize [®] Automation [™] Advanced or higher	vRealize Automation	8.0.1
VMware Workspace ONE [®] Access [™]	Workspace ONE Access	3.3.1

Before You Apply This Guidance

To use *Architecture and Design for VMware vRealize Suite 2019 on VMware Validated Design 5.1.1 or VMware Cloud Foundation 3.9.1*, you must have a VMware Validated Design 5.1.1 or VMware Cloud Foundation 3.9.1 SDDC deployment with the following requirements:

- Single or multi-region architecture
- One or more availability zones
- A management domain and at least one virtual infrastructure (VI) workload domain in each region
- Each workload domain is composed of the SDDC virtual infrastructure components only

Table 1-2. SDDC Virtual Infrastructure Components

Product	Management Domain	VI Workload Domain
SDDC Manager (only for VMware Cloud Foundation)	✓	x
VMware vSphere [®]	✓	✓

Table 1-2. SDDC Virtual Infrastructure Components (continued)

Product	Management Domain	VI Workload Domain
VMware vSAN™	✓	Optional. Supports also NFS and FC.
VMware NSX® Data Center for vSphere®	✓	✓
VMware NSX-T™ Data Center	x	✓
A backup solution that is compatible with VMware vSphere Storage APIs – Data Protection (VADP)	✓	x
Site Recovery Manager and vSphere Replication (for a multi- region SDDC)	✓	x

For information about the versions of the SDDC virtual infrastructure components, see *VMware Validated Design 5.1.1 Release Notes* or *VMware Cloud Foundation 3.9.1 Release Notes*.

For information about deploying an SDDC according to VMware Validated Design 5.1.1, see:

- *VMware Validated Design Deployment of Region A*
- *Optionally, VMware Validated Design Deployment of Region B*
- *Optionally, VMware Validated Design Deployment of Multiple Availability Zones*

See .

For information about deploying an SDDC by using VMware Cloud Foundation 3.9.1, see *VMware Cloud Foundation Architecture and Deployment Guide*.

See [VMware Validated Design Documentation](#) and [VMware Cloud Foundation Documentation](#).

Architecture Overview

1

By implementing VMware Validated Design for VMware Cloud Automation Services you can automate provisioning and life cycle management of workloads across the Software Defined Data Center and public clouds by using a policy-driven, extensible, as-a-service cloud automation platform.

This chapter includes the following topics:

- [Identity and Access Management Architecture](#)
- [Cloud Automation Architecture](#)
- [Cloud Operations Architecture](#)

Identity and Access Management Architecture

VMware Workspace ONE Access, formerly VMware Identity Manager™, provides identity and access management to end users.

Workspace ONE Access implements the Zero Trust Access Control model by providing users continuous access to their applications and data based on many factors like their device, location, how they are authenticated, intelligence, and risk signals. Workspace ONE Access ensures that users do not have access to applications that they must not access. Workspace ONE Access provides a common experience for accessing on-premises or SaaS applications, while also providing administrators visibility into user application accessibility, who uses what or when, and the frequency of access. Workspace ONE Access works together with your primary identity providers while acting as a broker into the Software-Defined Data Center and End User Computing platforms.

In the context of the SDDC, Workspace ONE Access is the broker between existing authentication providers in your data center, for example, Active Directory and LDAP directory, and SDDC solutions, such as the vRealize Suite products. Workspace ONE Access provides identity and access management services to each SDDC solution and ensures that the SAML is valid across solutions and regions in the SDDC.

Cloud Automation Architecture

VMware vRealize Automation streamlines multi-cloud infrastructure and application delivery, enhances visibility and cross-functional collaboration, and provides continuous delivery and release automation.

VMware vRealize Automation is a bundled offering of Cloud Assembly, Service Broker, and Code Stream.

- **VMware Cloud Assembly**

VMware Cloud Assembly can automate the delivery of cloud services across multiple clouds.

- **vRealize Orchestrator Architecture**

VMware vRealize Orchestrator contains a workflow library and a workflow engine to allow you to create and run workflows that automate orchestration processes. You run workflows on objects of different technologies that vRealize Orchestrator accesses through a series of plug-ins.

- **VMware Service Broker**

VMware Service Broker aggregates native content from multiple clouds and platforms into a single catalog with role-based policies.

- **VMware Code Stream**

VMware Code Stream accelerates the software delivery and streamlines the troubleshooting through release pipelines automation and analytics.

VMware Cloud Assembly

VMware Cloud Assembly can automate the delivery of cloud services across multiple clouds.

By using the VMware Cloud Assembly service, you can create and deploy virtual machines, applications, and services to your cloud infrastructure. With Cloud Assembly you can:

- Curate content - define and configure what content is available to projects.
- Design and deploy - iteratively build and deploy blueprints for infrastructure and applications.

As a cloud administrator, you configure the cloud vendor infrastructure to support cloud-agnostic blueprint development and deployment for multiple clouds. You set up projects, add users, and enable access to resources in cloud accounts/regions. You import or develop blueprints, or delegate development to the project administrators and members.

As a project member, you use Cloud Assembly to develop and deploy blueprints with a declarative and iterative approach. You deploy blueprints to the cloud accounts/regions, which are configured by the cloud administrator as part of the project, and manage their resources throughout the development life cycle. You the integrate resource delivery into continuous integration and continuous delivery (CI/CD) processes.

Some of the Cloud Assembly capabilities include the following:

Cloud-Agnostic Blueprints	Build blueprints on a set of building blocks that can be deployed on any supported cloud.
Infrastructure as Code	Define blueprints in YAML to facilitate deployment, configuration, definition repeatability, version control, and collaboration.
Policy-Based Resource Delivery	Establish governance that includes access definition to resources and which clouds can support specific activities or teams.

Development Model Create environment definitions by writing code, drawing, or both in an intuitive canvas.

Declarative Approach Modify and iterate to reach a desired end state definition.

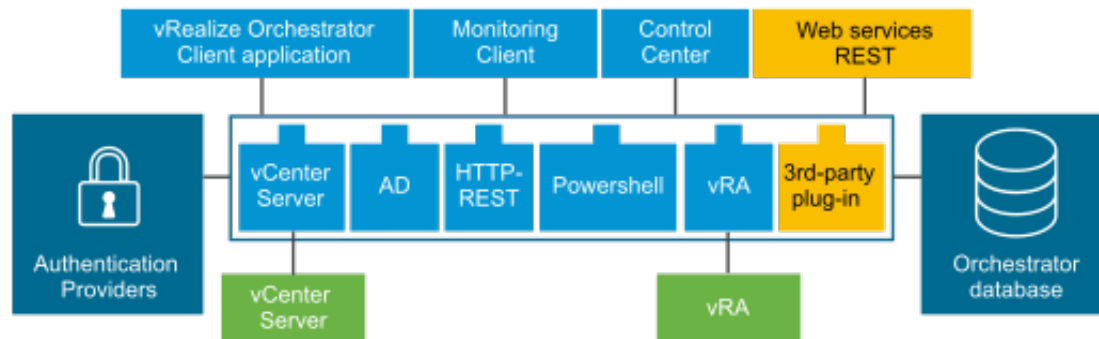
vRealize Orchestrator Architecture

VMware vRealize Orchestrator contains a workflow library and a workflow engine to allow you to create and run workflows that automate orchestration processes. You run workflows on objects of different technologies that vRealize Orchestrator accesses through a series of plug-ins.

VMware vRealize Orchestrator provides a standard set of plug-ins, including a plug-in for vCenter Server, to allow you to orchestrate tasks in the different environments that the plug-ins expose.

vRealize Orchestrator also presents an open architecture for plugging in external third-party applications to the orchestration platform. You can run workflows on the objects of plugged-in technologies that you define yourself. vRealize Orchestrator connects to an authentication provider to manage user accounts and to a preconfigured PostgreSQL database to store information from the workflows that it runs. You can access vRealize Orchestrator, the objects it exposes, and the vRealize Orchestrator workflows through the vRealize Orchestrator Client service, or through Web services. Monitoring and configuration of vRealize Orchestrator workflows and services is done through the vRealize Orchestrator Client and Control Center.

Figure 1-1. vRealize Orchestrator Architecture



VMware Service Broker

VMware Service Broker aggregates native content from multiple clouds and platforms into a single catalog with role-based policies.

VMware Service Broker is a cloud service that aggregates content in native formats from multiple clouds and platforms into a simplified and efficient catalog. You use the catalog to manage the available catalog items, as well as how and where they are deployed in cloud accounts/regions.

As a cloud administrator, Service Broker is the portal that you provide to your users, such as, operations and development teams. You import content such as Cloud Assembly blueprints, AWS CloudFormation templates, and extensibility actions, and configure governance in the form of projects to control accessibility of resources and deployment location.

As a user, you request and monitor the provisioning process. After deployment, you manage the deployed catalog items throughout the deployment lifecycle.

Some of the Service Broker capabilities include the following:

- **Self Service**

A portal for users that provides access to both infrastructure and application level services.

- **Governance**

Policy based management that provide access control over resources and deployment locations.

- **Definition Abstraction**

Support for integrating end user services, that have been designed by using a range of definition tools, such as VMware Cloud Assembly.

- **Multi-cloud Support**

Unified delivery of predefined services, that run on different cloud environments, including VMware based private and hybrid clouds, as well as native public clouds.

VMware Code Stream

VMware Code Stream accelerates the software delivery and streamlines the troubleshooting through release pipelines automation and analytics.

VMware Code Stream is a cloud service that provides continuous integration and continuous delivery (CI/CD) that enables you to deliver software rapidly and reliably. Code Stream simplifies the ability to build, test, and deploy your applications, and increases productivity as you release source code from the development repository, through testing, to production. Code Stream integrates your release with custom and common developer tools and objects, such as Cloud Assembly blueprints.

You create a pipeline that runs actions to build, deploy, test, and release your software. Code Stream runs your software through each stage of the pipeline until it is ready to be released. You integrate the pipeline with one or more DevOps tools, which provide data for the pipeline to run. For example, when a developer checks in code to a Git repository, Code Stream can trigger the pipeline and automate the build, test, and deployment of an application.

You can integrate Code Stream with other VMware Cloud Service. For example, you can publish your Code Stream pipeline to Service Broker as a catalog item that can be requested and deployed on cloud accounts/regions. You can also deploy a Cloud Assembly blueprint and use the parameter values that the blueprint exposes.

Some of the Code Stream capabilities include the following:

- Artifacts**

Ensure that correct versions are used across all stages of the development life cycle.

- Dashboards**

Facilitate collaboration by providing dashboards and reports for release pipelines KPIs.

Integration	Use existing investments in build, test, provisioning, deployment, and monitoring tools.
Multi-cloud Support	Run pipelines on different cloud environments including VMware based private and hybrid clouds, as well as native public clouds.

Important VMware Code Stream is out of scope for this VMware Validated Design.

Cloud Operations Architecture

The architecture of the products of the cloud operations layer supports centralized monitoring of and logging data about the other solutions in the SDDC. You use this architecture to deliver core operational procedures in the data center.

In the operations management layer, the physical infrastructure, virtual infrastructure, and tenant workloads are monitored in real time, collecting the following information for intelligent and dynamic operational management:

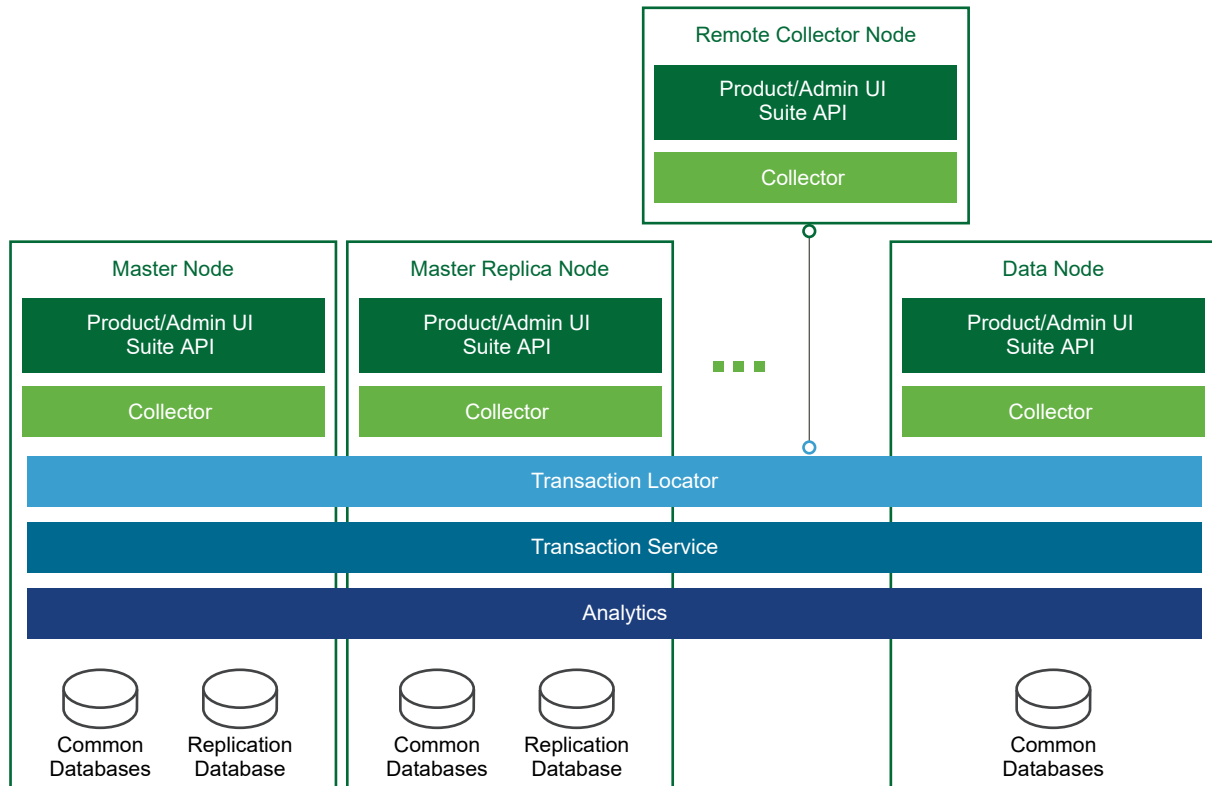
- Monitoring data, such as structured (metrics) and unstructured (logs) data
- Topology data, such as physical and virtual compute, networking, and storage objects

Monitoring Architecture

vRealize Operations Manager tracks and analyzes the operation of multiple data sources in the SDDC by using specialized analytic algorithms. These algorithms help vRealize Operations Manager learn and predict the behavior of every object it monitors. Users access this information by using views, reports, and dashboards.

Architecture Overview

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

Figure 1-2. vRealize Operations Manager Architecture

Types of Nodes

For high availability and scalability, you can deploy several vRealize Operations Manager instances in a cluster to track, analyze, and predict the operation of monitored systems. Cluster nodes can have either of the following roles:

Master Node

Required initial node in the cluster. In large-scale environments, 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

Optional. 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.

Remote Collector Node

Overcomes data collection issues across the enterprise network, such as limited network performance. Remote collector nodes gather statistics about inventory objects and forward collected data to the data nodes. Remote collector nodes do not store data or perform analysis.

Types of Node Groups

Analytics Cluster

Tracks, analyzes, and predicts the operation of monitored systems. Consists of a master node, data nodes, and optionally of a master replica node.

Remote Collector Group

Consists of remote collector nodes. Only collects diagnostics data without storage or analysis. A vRealize Operations Manager deployment can contain several collector groups.

Use collector groups to achieve adapter resiliency in cases where the collector experiences network interruption or becomes unavailable.

Deployment

vRealize Operations Manager is available as a preconfigured virtual appliance in an OVA template. By using the virtual appliance, you can create vRealize Operations Manager nodes with predefined identical sizes.

You deploy the OVA template once for each node. After the node deployment, you access the product to set up cluster nodes according to their role and log in to configure the installation.

Deployment Models

You can deploy vRealize Operations Manager in one of the following configurations:

- A standalone node
- A cluster of one master, at least one data node, and optionally a group of remote collector nodes.

You can establish high availability by using an external load balancer.

The compute and storage resources of the vRealize Operations Manager instances can scale up as growth demands.

Authentication Sources

You can configure the vRealize Operations Manager user authentication to use one or more of the following authentication sources:

- VMware vCenter[®] Single Sign-On
- Workspace ONE Access
- OpenLDAP via LDAP
- Active Directory via LDAP

Management Packs

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

Backup

You back up each vRealize Operations Manager node by using traditional virtual machine backup solutions that are compatible with VMware vSphere Storage APIs – Data Protection (VADP).

Multi-Region vRealize Operations Manager Deployment

The scope of this design can cover both multiple regions and availability zones.

This design implements a large-scale vRealize Operations Manager deployment across multiple regions by using the following configuration:

- Load-balanced analytics cluster that runs multiple nodes is protected by Site Recovery Manager to fail over across regions.
- Multiple remote collector nodes that are assigned to a remote collector group in each region to handle data coming from management solutions.

In a multi-availability zone implementation, which is a super-set of the multi-region design, vRealize Operations Manager continues to provide monitoring of the solutions in all regions of the SDDC. All components of vRealize Operations Manager reside in Availability Zone 1 in Region A. If this zone becomes compromised, all nodes are brought up in Availability Zone 2.

Logging Architecture

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

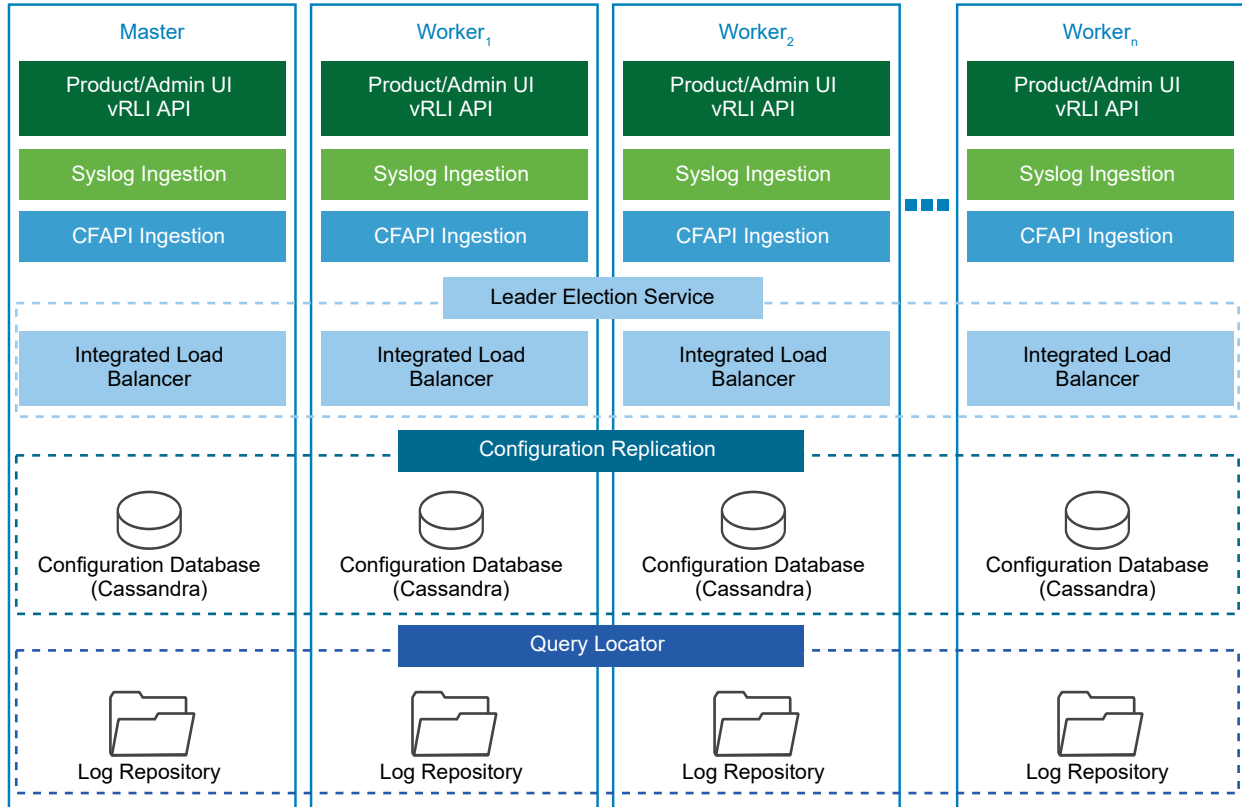
vRealize Log Insight collects data from ESXi hosts by using the syslog protocol. vRealize Log Insight has the following capabilities:

- Connects to other VMware products, such as vCenter Server, to collect events, tasks, and alarm data.
- Integrates with vRealize Operations Manager to send notification events and enable launch in context.
- Functions as a collection and analysis point for any system that is capable of sending syslog data.

To collect additional logs, you can install an ingestion agent on Linux or Windows servers, or you can use the preinstalled agent on certain VMware products. Preinstalled agents are useful for custom application logs and operating systems that do not natively support the syslog protocol, such as Windows.

Architecture

The architecture of vRealize Log Insight in the SDDC enables several channels for the collection of log messages.

Figure 1-3. Architecture of vRealize Log Insight

vRealize Log Insight clients connect to the Load Balancer (ILB) FQDN, and use the syslog or the Ingestion API via the vRealize Log Insight agent to send logs to vRealize Log Insight. Users and administrators interact with the ingested logs by using the user interface or the API.

By default, vRealize Log Insight collects data from vCenter Server systems and ESXi hosts. For analyzing forwarded logs from other components, such as NSX Data Center, use content packs. Content packs contain extensions or provide integration with other systems in the SDDC.

Types of Nodes

For functionality, high availability, and scalability, vRealize Log Insight supports the following types of nodes which have inherent roles:

Master Node

Required initial node in the cluster. In standalone mode, the master node is responsible for all activities, including queries and log ingestion. The master node also handles operations that are related to the life cycle of a cluster, such as performing upgrades and addition or removal of worker nodes. In a scaled-out and highly available environment, the master node still performs life cycle operations, such as addition or removal of worker nodes. However, it functions as a generic worker about queries and log ingestion activities.

The master node stores logs locally. If the master node is down, the logs stored on it become unavailable.

Worker Node

Optional. This component enables a scale-out growth in larger environments. As you add and configure more worker nodes in a vRealize Log Insight cluster for high availability (HA), queries and log ingestion activities are delegated to all available nodes. You must have at least two worker nodes to form a cluster with the master node. The worker node stores logs locally. If any of the worker nodes is down, the logs on the worker become unavailable.

Integrated Load Balancer (ILB)

In cluster mode, the ILB is the centralized entry point which ensures that vRealize Log Insight accepts incoming ingestion traffic. As nodes are added to the vRealize Log Insight instance to form a cluster, the ILB feature simplifies the configuration for high availability. The ILB balances the incoming traffic fairly among the available vRealize Log Insight nodes.

The ILB runs on one of the cluster nodes at all times. In environments that contain several nodes, an election process determines the leader of the cluster. Periodically, the ILB performs a health check to determine whether re-election is required. 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 using an election process.

All queries against data are directed to the ILB. The ILB delegates queries to a query master for the duration of the query. The query master queries all nodes, both master and worker nodes, for data and then sends the aggregated data back to the client.

Use the ILB for administrative activities unless you are performing administrative activities on individual nodes. The Web user interface of the ILB presents data from the master and from the worker nodes in a scaled-out cluster in a unified display (single pane of glass).

Multi-Region vRealize Log Insight Deployment

The scope of this validated design can cover both multiple regions and availability zones.

In a multi-region implementation, vRealize Log Insight provides a separate logging infrastructure in each region of the SDDC. Using vRealize Log Insight across multiple regions requires the following configuration:

- Cluster in each region.
- Event forwarding to other vRealize Log Insight deployments across regions in the SDDC.

In a multi-availability zone implementation, which is a sub-set of the multi-region design, vRealize Log insight continues to provide a logging infrastructure in all regions of the SDDC. All components of the vRealize Log Insight cluster reside in Availability Zone 1 within Region A. If this zone becomes compromised, all nodes are brought up in the Availability Zone 2.

Failover by using vSphere Replication or disaster recovery by using Site Recovery Manager is not necessary. The event forwarding feature adds tags to log messages that identify the source region. Event filtering prevents looping messages between the regions.

Backup

You back up vRealize Log Insight by using traditional virtual machine backup solutions that are compatible with VMware vSphere Storage APIs – Data Protection (VADP).

Life Cycle Architecture

VMware vRealize Suite Lifecycle Manager is used to automate the deployment, upgrade, and patching of the VMware vRealize products in this design.

Overview

In this design, the vRealize Suite Lifecycle Manager solution supports the deployment, upgrade, and patching of the following vRealize products :

- Workspace ONE Access
- vRealize Log Insight
- vRealize Operations
- vRealize Automation

vRealize Suite Lifecycle Manager is a pre-configured appliance distributed in an Open Virtual Appliance (.ova) format. After the appliance deployment, you can access the vRealize Suite Lifecycle Manager services by using both the browser application user interface and the API.

After you deploy vRealize Suite Lifecycle Manager, you register one or more Management Domain vCenter Server instances with it.

An administrator can automate life cycle operations for vRealize products. vRealize Suite Lifecycle Manager provides the following features for management of vRealize products:

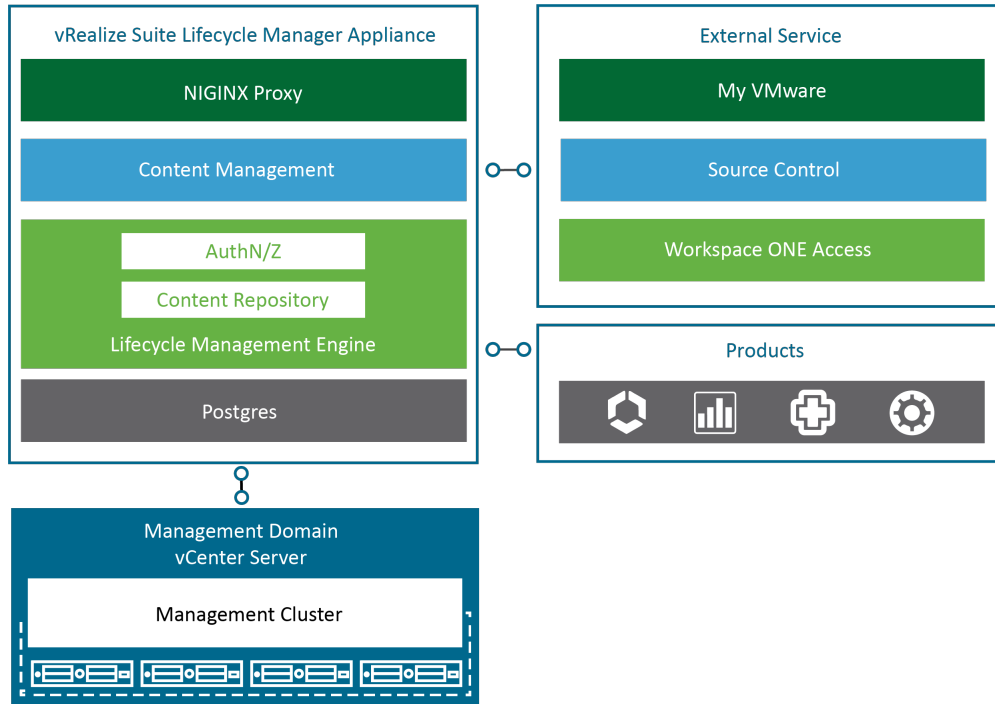
- Manage a product install, upgrade, and patch repository.
- Deploy products using supported topologies.
- Patch and upgrade product deployments.
- Scale-out product deployments.
- Support the import of existing product deployments.
- Organize product deployments in logical environments.
- Manage product certificates, licenses, and passwords.

- Manage and deploy Marketplace content across vRealize solutions.

Architecture

vRealize Suite Lifecycle Manager contains the functional elements that collaborate to orchestrate the life cycle management operations of the vRealize products in this design.

Figure 1-4. Architecture of vRealize Suite Lifecycle Manager



vRealize Suite Lifecycle Manager contains modules for installation, upgrade, and patching of vRealize products in a vSphere environment. vRealize Suite Lifecycle Manager manages product binaries, downloads product content from VMware Marketplace, and integrates with Workspace ONE Access for a centralized identity and access management.

Authentication Models

You can configure the vRealize Suite Lifecycle Manager user authentication to use the following authentication models:

- Local administrator account
- Workspace ONE Access

Marketplace Integration

By using vRealize Suite Lifecycle Manager, you can deploy additional vRealize Operations management packs, vRealize Log Insight content packs, and vRealize Automation blueprints and OVA files directly from the VMware Marketplace.

Multi-Region Deployment of vRealize Suite Lifecycle Manager

The scope of this design can cover both a single region and multiple regions, and availability zones.

In a multi-region implementation, the design implements a vRealize Suite Lifecycle Manager setup in multiple regions by using the following configuration:

- A single vRealize Suite Lifecycle Manager appliance is replicated by vSphere Replication and recovered by Site Recovery Manager. You can fail over the vRealize Suite Lifecycle Manager appliance across regions when there is a planned migration or disaster recovery event.
- The vRealize Suite Lifecycle Manager instance manages the deployment, upgrade, and patching of the vRealize products across all regions.

In a multi-availability zone implementation, vRealize Suite Lifecycle Manager continues to provide life cycle operations services for the vRealize product deployments in all regions of the SDDC. The vRealize Suite Lifecycle Manager virtual appliance resides in Availability Zone 1 in Region A. If this zone becomes compromised, the appliance instance is brought back online in Availability Zone 2.

Backup

You back up vRealize Suite Lifecycle Manager by using traditional virtual machine backup solutions that are compatible with VMware vSphere Storage APIs – Data Protection (VADP).

Detailed Design

2

VMware Validated Design for VMware Cloud Automation Services considers both virtual infrastructure and services design. It includes numbered design decisions, and the justification and implications of each decision.

This chapter includes the following topics:

- [Identity and Access Management Design](#)
- [Cloud Automation Design](#)
- [Cloud Operations Design](#)

Identity and Access Management Design

Identity and Access Management is a key foundational part of the design required to control access to the components making up the Software-Defined Data Center. In this solution, VMware Workspace ONE Access provides this functionality.

Workspace ONE Access Design

VMware Workspace ONE Access provides identity and access management services to several components within the Software-Defined Data Center. Some of these components are region-dependent (or specific) and some are region-independent. The design aligns with the design objectives, constraints, and use cases of these components – in terms of the number of users, availability requirements, and so on.

Logical Design for Workspace ONE Access

The logical design dictates two separate instances modes for Workspace ONE Access - one serving the region-specific components and another serving the region-independent components in the SDDC.

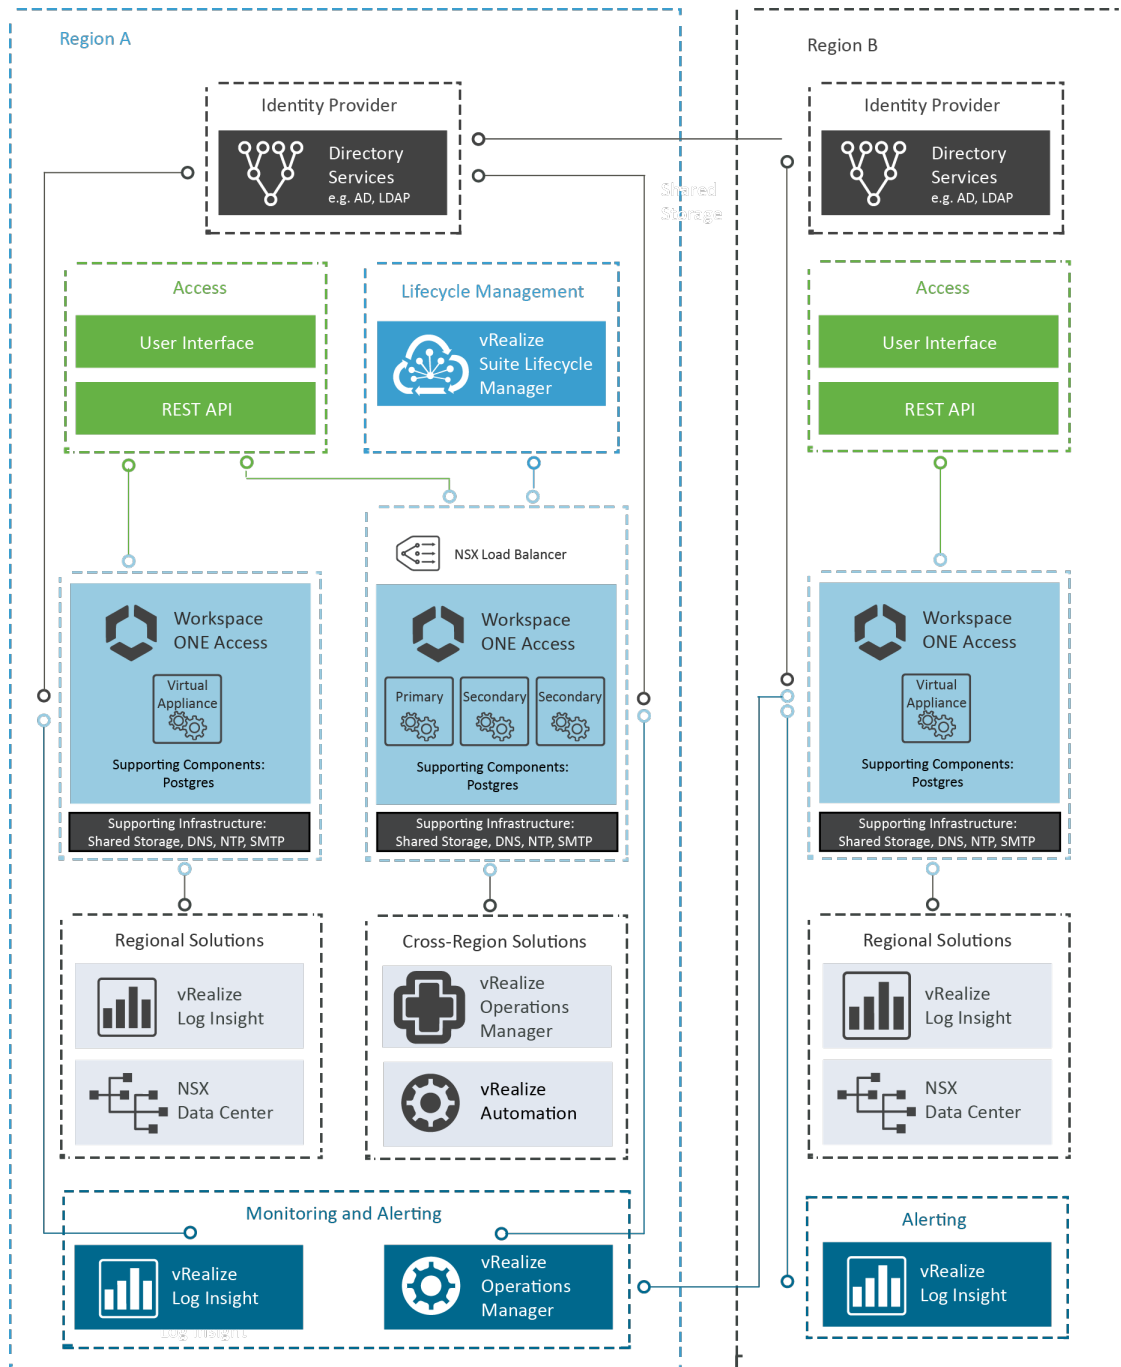
The design addresses the separation of duties, capacity constraints, scalability, sizing, and deployment type to suit the needs of the components a particular Workspace ONE Access deployment serves.

Each mode provides:

- Directory integration to authenticate users against existing directories, such as Active Directory or LDAP.

- Addition of a two-factor authentication through integration with third-party software, such as RSA SecurID, Entrust, and others.

Figure 2-1. Logical Design of Workspace ONE Access in a Multi-Region Deployment



Region-Specific Mode

A single node Workspace ONE Access instance is deployed on the region-specific management virtual network in each region to provide identity and access management for the following SDDC solutions:

- vRealize Log Insight: Workspace ONE Access provides identity and access management services for vRealize Log Insight. Workspace ONE Access is also integrated with vRealize Log Insight as its logging service provider to aggregate and manage its logs.

Region-Independent Mode

A three-node Workspace ONE Access cluster is deployed on a common virtualized network that spans between the regions, that is, cross-region. The Workspace ONE Access cluster is deployed and managed by vRealize Suite Lifecycle Manager and is integrated with the following cross-region SDDC solutions:

- vRealize Operations Manager: Workspace ONE Access provides an authentication source for vRealize Operations which provides authentication services to the enterprise directory or directories. Workspace ONE Access is also integrated with vRealize Operations by using a vRealize Operations Management Pack that provides the monitoring and alerting capabilities, such as health, risk, and efficiency.
- vRealize Automation: Workspace ONE Access provides the authentication services for vRealize Automation which provides authentication services to the enterprise directory or directories.

Supporting Infrastructure

All instances of Workspace ONE Access in this design integrate with the following supporting infrastructure:

- NTP for time synchronization
- DNS for name resolution
- Active Directory (or LDAP) directories

Important Workspace ONE Access does not replace Active Directory or LDAP. Workspace ONE Access integrates with Directory or LDAP for authentication and solution authorization.

Configuration Design for Workspace ONE Access

The configuration design consists of characteristics and decisions that support the logical design. The design objective is to deploy a fully functional identity and access management solution with high availability and the ability to provision workloads in a multi-region SDDC.

Deployment Model of Workspace ONE Access

You deploy a standalone Workspace ONE Access node in each region. You deploy a Workspace ONE Access cluster in Region A.

To accomplish this design objective, you deploy or reuse the following components to deploy the identity and access management layer for the SDDC.

- SDDC Manager

- vRealize Suite Lifecycle Manager
- NSX for vSphere Application Virtual Networks
- NSX for vSphere Load Balancer

Workspace ONE Access is distributed as a virtual appliance in OVA format. The Workspace ONE Access appliance includes identity and access management services. You consider the deployment type - standard or cluster - according to the design objectives for the availability and number of users that the system and integrated SDDC solutions must support. Workspace ONE Access is deployed to the first vSphere cluster in the management domain.

In this design, you deploy Workspace ONE Access in the following topology:

Table 2-1. Workspace ONE Access Topology Attributes

Topology	Use	User Scale	Description
Standard deployment	Regional	1,000 users	You deploy a Workspace ONE Access cluster instance - a single virtual appliance - on the first cluster in the management domain of each region.
Cluster deployment	Cross-region	10,000 users	<p>You deploy a Workspace ONE Access cluster instance - three virtual appliances and a load balancer - in the management domain of Region A.</p> <p>All Workspace ONE Access services and databases are configured for high availability using the underlying appliance configuration. Portable, cross-region SDDC solutions are integrated with this Workspace ONE Access cluster instance.</p> <p>vSphere High Availability protects Workspace ONE Access by restarting each virtual appliance on an alternate ESXi host if a primary ESXi host failure occurs.</p> <p>vSphere Distributed Resource Scheduler anti-affinity rules ensure that Workspace ONE Access virtual appliances in the cluster must be running on different ESXi hosts in the first vSphere cluster in the management domain.</p>

Sizing Compute and Storage Resources for Workspace ONE Access

A Workspace ONE Access standard and cluster deployment models have the following resource requirements:

Table 2-2. Workspace ONE Access CPU, Memory, and Storage Resources

Attribute	Standard Deployment	Cluster Deployment
Number of appliances	1	3 + Load Balancer
CPU	2 vCPUs	6 vCPUs
Memory	6 GB	18 GB
Storage	4.7 GB (thin provisioned)	14.1 GB (thin provisioned)
	60.2 GB (thick provisioned)	180.6 GB (thick provisioned)

Table 2-3. Design Decisions on Deployment and Configuration of Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-001	Deploy a standalone Workspace ONE Access instance on the first vSphere cluster in the management domain in each region.	Each appliance provides an identity and access management service to the regional SDDC solutions.	None
SDDC-IAM-002	Deploy a separate Workspace ONE Access cluster in the first vSphere cluster in the management domain in Region A.	It provides an identity and access management service to cross-region SDDC solutions.	You must use a Workspace ONE Access deployment type that accommodates multi-region deployments.
SDDC-IAM-003	Deploy the Workspace ONE Access instance (standalone) using the standard deployment type to provide identity and access management services to regional SDDC solutions.	Deploying the standard configuration that includes the single-node appliance architecture satisfies the design objectives in the scope for the design allowing Workspace ONE Access to scale to a higher number of consuming users for vRealize Log Insight.	The region-specific Workspace ONE Access instance is not managed by vRealize Suite Lifecycle Manager. Availability is managed by vSphere High Availability only.

Table 2-3. Design Decisions on Deployment and Configuration of Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-004	Deploy the Workspace ONE Access cluster instance through SDDC Manager using the cluster deployment type to provide identity and access management services to cross-region SDDC solutions.	Deploying the cluster configuration that includes the three-node appliance architecture satisfies the design objectives in scope for the design allowing Workspace ONE Access to scale to a higher number of consuming users for vRealize Automation and vRealize Operations authentication. The cross-region Workspace ONE Access cluster is managed by vRealize Suite Lifecycle Manager.	None
SDDC-IAM-005	Protect all Workspace ONE Access virtual appliances by using vSphere High Availability.	Supports the availability objectives for Workspace ONE Access without a required manual intervention during a failure event.	The Workspace ONE Access instance for region-specific SDDC solutions becomes unavailable during a vSphere HA failover.
SDDC-IAM-006	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules for the Workspace ONE Access virtual appliances in the cluster for cross-region SDDC solutions.	Using vSphere DRS prevents Workspace ONE Access virtual appliances from residing on the same ESXi host and risking the high availability of the deployment.	You can only place a single ESXi host at a time into maintenance mode for a management cluster of four ESXi hosts. Requires at least four physical hosts to guarantee the three Workspace ONE Access virtual appliances continue to run in the cluster if an ESXi host failure occurs.
SDDC-IAM-007	Add the VM groups for the cross-region Workspace ONE Access cluster virtual appliances and set VM rules to restart the Workspace ONE Access VM group before the vRealize Automation VM group.	Allows you to define the startup order of virtual appliances regarding service dependency. The startup order ensures that vSphere HA powers on the virtual machines for Workspace ONE Access in the correct order.	None

Table 2-3. Design Decisions on Deployment and Configuration of Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-008	When using two availability zones in Region A, add the Workspace ONE Access virtual appliances to the primary availability zone VM group, for example, <code>sfo01-m01-primary-az-vm-group</code> .	Ensures that the Workspace ONE Access virtual appliances are powered on within the primary availability zone hosts group by default.	If Workspace ONE Access is deployed after the creation of the stretched clusters for management domain availability zones, the VM group for the primary availability zone virtual machines must be updated to include the Workspace ONE Access virtual appliances.
SDDC-IAM-009	Place all region-specific Workspace ONE Access virtual appliances in a dedicated virtual machine folder in each region, for example, <code>sfo01-m01fd-wsa</code> and <code>lax01-m01fd-wsa</code> .	Provides the organization of region-specific Workspace ONE Access virtual appliances in the management domain inventory.	None
SDDC-IAM-010	Place all cross-region Workspace ONE Access virtual appliances in a dedicated virtual machine folder in Region A, for example, <code>xregion-sfo01-lax01-m01fd-wsa</code> .	Provides the organization of cross-region Workspace ONE Access virtual appliances in the management domain inventory and preparation for Site Recovery Manager folder mappings for disaster recovery.	A corresponding virtual machine folder in Region B must be created in preparation for Site Recovery Manager folder mapping, for example, <code>xregion-lax01-sfo01-m01fd-wsa</code> .

Logging Design for Workspace ONE Access

You integrate Workspace ONE Access with vRealize Log Insight to provide operational visibility.

The integration to vRealize Log Insight from Workspace ONE Access provides the ability to send logs from the service containers for aggregation and analysis, as needed.

Logging is enabled through the vRealize Log Insight ingestion API. It is established by installing the vRealize Log Insight Agent and configuring the Workspace ONE Access agent groups for each vRealize Log Insight cluster.

Table 2-4. Design Decisions on Logging for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-011	Do not configure the Workspace ONE Access virtual appliances to use the Syslog protocol for logs.	Workspace ONE Access virtual appliances are configured to use the vRealize Log Insight ingestion API.	None

For more information, see [Integration of vRealize Log Insight with vRealize Operations Manager and Workspace ONE Access](#).

Network Design for Workspace ONE Access

For secure access to the UI and API of Workspace ONE Access, and for failover of the cross-region Workspace ONE Access cluster, you deploy the virtual appliances on application virtual networks.

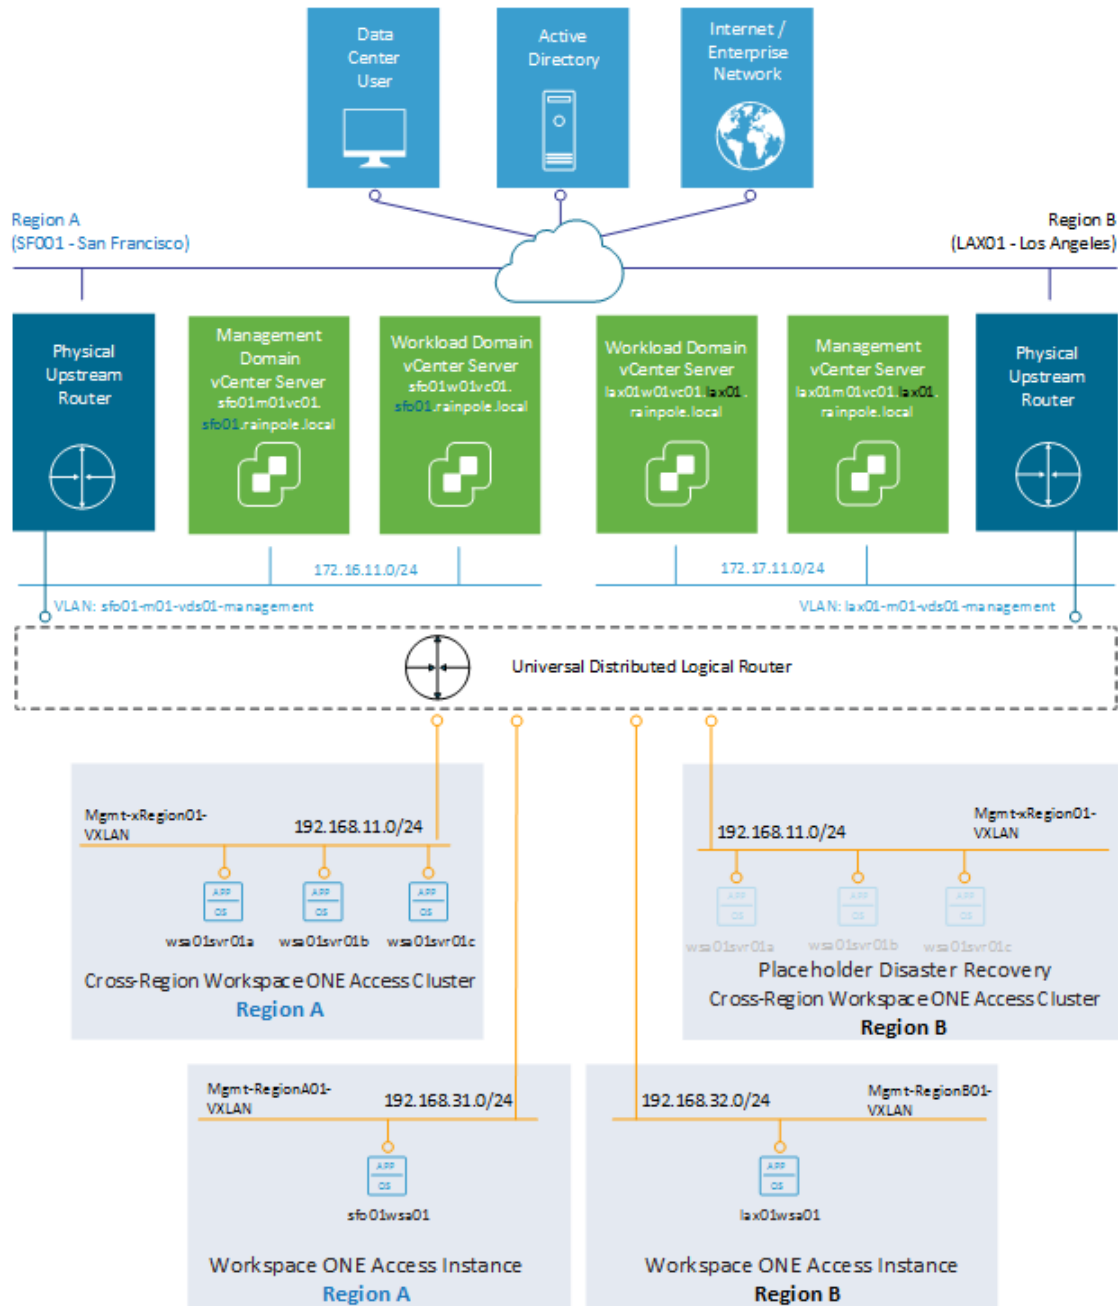
Application Virtual Network

This design uses NSX for vSphereVXLAN backed application virtual networks to abstract Workspace ONE Access and its supporting services from the underlying physical infrastructure.

This networking design has the following features:

- A Workspace ONE Access node for the region-specific SDDC solutions is deployed on a region-specific application virtual network in each region.
- Workspace ONE Access cluster nodes for cross-region SDDC solutions are deployed together on the same cross-region application virtual network in a designated region. With this configuration, you can fail over Workspace ONE Access between regions after expanding to a multi-region SDDC design.
- All Workspace ONE Access components have routed access to the VLAN-backed management network through the NSX Universal Distributed Logical Router.
- Routing to the VLAN-backed management network and other external networks is dynamic and is based on the Border Gateway Protocol (BGP).

The Workspace ONE Access cluster nodes are connected to the virtual segments based on their role to provide secure access to the UI and API, or failover support.

Figure 2-2. Networking Design of the Workspace ONE Access Deployment

As part of this design, use the application virtual network configuration to connect Workspace ONE Access with the other management solutions in the SDDC. Regional Workspace ONE Access nodes are connected to the region-specific application virtual networks, for example, Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN. The cross-region Workspace ONE Access cluster is connected to the cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN, and uses the load balancer on the NSX for vSphere edge for high availability and balancing user access across the Workspace ONE Access cluster.

Table 2-5. Design Decisions on the Application Virtual Network for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-012	Place the Workspace ONE Access virtual appliances for regional SDDC solutions on the existing region-specific application virtual network, for example, Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN.	Authentication and authorization can sustain operations in the event of service interruption in the cross-region network. Ensures a consistent deployment model for management applications.	You must use an implementation in NSX for vSphere to support this network configuration.
SDDC-IAM-013	Place the Workspace ONE Access virtual appliances for cross-region SDDC solutions on the existing cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN.	Authentication and authorization can sustain operations in the event of a planned failover or disaster recovery between regions. Supports disaster recovery by isolating the Workspace ONE Access on the NSX for vSphere network segment, for example, Mgmt-xRegion01-VXLAN, in the management domain.	You must use an implementation in NSX for vSphere to support this network configuration.

IP Addressing Scheme

You allocate the following example subnets in the management domains to the Workspace ONE Access deployments.

Table 2-6. Example IP Subnets for Workspace ONE Access

Solution	IP Subnet	Gateway	NSX for vSphere Application Virtual Network
Workspace ONE Access for Region A	192.168.31.0/24	192.168.31.1	Mgmt-RegionA01-VXLAN
Workspace ONE Access for Region B	192.168.32.0/24	192.168.32.1	Mgmt-RegionB01-VXLAN
Workspace ONE Access for cross-region	192.168.11.0/24	192.168.11.1	Mgmt-xRegion01-VXLAN

Name Resolution

The Workspace ONE Access nodes follow a specific domain name resolution.

The Workspace ONE Access components in each region have the following characteristics:

- The IP addresses of the cross-region Workspace ONE Access cluster and nodes are associated with a fully qualified name whose suffix is set to the root domain, for example, `rainpole.local`.
- The IP addresses of the regional Workspace ONE Access nodes are associated with a fully qualified name whose suffix is set to the child domain, for example, `sfo01.rainpole.local` and `lax01.rainpole.local`.

Table 2-7. Example FQDNs and IP Addresses for Workspace ONE Access

FQDN	IP Address	Description	Region	Failed over to Region B
sfo01wsa01.sfo01.rainpole.local	192.168.31.60	Workspace ONE Access Virtual Appliance for Region A	Region A	No
lax01wsa01.lax01.rainpole.local	192.168.32.60	Workspace ONE Access Virtual Appliance for Region B	Region B	No
wsa01svr01.rainpole.local	192.168.11.60	Workspace ONE Access Cluster NSX for vSphere Load-Balancer Virtual Server for Cross-Region	Region A	Yes
wsa01svr01a.rainpole.local	192.168.11.61	Workspace ONE Access Virtual Appliance A for Cross-Region	Region A	Yes
wsa01svr01b.rainpole.local	192.168.11.62	Workspace ONE Access Virtual Appliance B for Cross-Region	Region A	Yes
wsa01svr01c.rainpole.local	192.168.11.63	Workspace ONE Access Virtual Appliance C for Cross-Region	Region A	Yes

Note The design uses an Active Directory forest with two regional Active Directory child domains, so the examples use a hierarchical DNS name space. However, the design supports the use of a flat DNS name space.

Table 2-8. Design Decisions on DNS for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-014	Configure forward and reverse DNS records for each Workspace ONE Access appliance IP address for each regional instance.	Workspace ONE Access is accessible by using a fully qualified domain name instead of by using IP addresses only.	You must provide DNS records for each Workspace ONE Access appliance IP address.
SDDC-IAM-015	Configure forward and reverse DNS records for each Workspace ONE Access appliance IP address and the load-balancer virtual IP address for the cross-region instance.	Workspace ONE Access is accessible by using a fully qualified domain name instead of by using IP addresses only.	You must provide DNS records for each Workspace ONE Access appliance and the load-balancer virtual IP address.
SDDC-IAM-016	In a multi-region SDDC deployment, configure the DNS settings for the appliances in the cross-region Workspace ONE Access cluster to use DNS servers in each region.	Workspace ONE Access appliances can resolve DNS from regional DNS servers during a planned migration or disaster recovery between regions.	As you scale from a single-region to multi-region SDDC deployment, the DNS settings on each Workspace ONE Access appliance must be updated.

Time Synchronization

Workspace ONE Access is dependent on time synchronization for all appliances.

Table 2-9. Design Decisions on NTP for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-017	Configure NTP for each Workspace ONE Access virtual appliance.	Workspace ONE Access is dependent on time synchronization for all virtual appliances.	All firewalls located between the Workspace ONE Access virtual appliances and the NTP servers must allow NTP traffic.
SDDC-IAM-018	In a multi-region SDDC deployment, configure the NTP settings on the appliances in the cross-region Workspace ONE Access cluster to use NTP servers in each region.	Workspace ONE Access appliances can query NTP from regional NTP servers to synchronize time during a planned migration or disaster recovery between regions.	As you scale from a single region to multi-region SDDC deployment, the NTP settings on each Workspace ONE Access virtual appliance must be updated.

Load Balancing

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

Table 2-10. Design Decisions on Load Balancing for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-019	Add an NSX for vSphere edge in the management domain to load balance connections across cross-region Workspace ONE Access cluster members.	Required to deploy Workspace ONE Access as a cluster deployment type, enabling it to handle a greater load and obtain a higher level of availability for vRealize Automation and vRealize Operations which also share this load balancer.	You must use an implementation in NSX for vSphere to support this network configuration.
SDDC-IAM-020	<ul style="list-style-type: none"> ■ Add an NSX for vSphere load balancer service monitor, for example, <code>wsa-https-monitor</code>, for the cross-region Workspace ONE Access cluster with an active HTTPS monitor on monitoring port 443. ■ Use the following intervals and timeouts for the monitor: <ul style="list-style-type: none"> ■ Interval: 5 ■ Timeout: 10 seconds ■ Max retries: 3 seconds. ■ Set the HTTP request for the monitor: <ul style="list-style-type: none"> ■ HTTP Method: Get ■ Request URL: <code>/SAAS/API/1.0/REST/system/health/heartbeat</code>. ■ Set the HTTP Response for the monitor: <ul style="list-style-type: none"> ■ Expected: 200 ■ Receive: OK. ■ Set the SSL Configuration for the monitor: <ul style="list-style-type: none"> ■ Server SSL: Enabled ■ Client Certificate: Cross-Region Workspace ONE Access Cluster Certificate ■ SSL Profile: <code>default-balanced-server-ssl-profile</code>. 	<ul style="list-style-type: none"> ■ The active monitor uses HTTPS requests to monitor the application health reported by Workspace ONE Access. ■ Ensures that connections to unhealthy cross-region Workspace ONE Access members in the pool are disabled until a subsequent periodic health check finds the members to be healthy. 	<ul style="list-style-type: none"> ■ You must manage the lifecycle of the certificate used on the load-balancer for the cross-region Workspace ONE Access cluster. ■ If a higher-level SSL cipher profile is required, set the SSL Configuration to use the <code>default-high-security-server-ssl-profile</code> SSL profile.

Table 2-10. Design Decisions on Load Balancing for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-021	<ul style="list-style-type: none"> ■ Add an NSX for vSphere load balancer server pool, for example, <code>wsa-server-pool</code>, for cross-region Workspace ONE Access to use the LEASTCONN algorithm. ■ Set the static members for the pool: <ul style="list-style-type: none"> ■ Name: hostname ■ IP: IP Address ■ Port: 443 ■ Weight: 1 ■ State: Enabled 	<ul style="list-style-type: none"> ■ LEASTCONN distributes requests to members based on the number of current connections. New connections are sent to the pool member with the fewest connections. ■ Workspace ONE Access services respond on TCP 443. 	None
SDDC-IAM-022	<ul style="list-style-type: none"> ■ Add an NSX for vSphere load balancer Cookie persistence Application Profile, for example, <code>wsa-cookie-persistence-profile</code>, for the cross-region Workspace ONE Access. ■ Set the Application Profile Type to HTTPS End-To-End. ■ Set the Persistence to Source IP. ■ Set Expires in to 3600 seconds. ■ Set X-Forwarded-For HTTP header to Enable. 	The cross-region Workspace ONE Access cluster requires cookie session persistence for the Workspace ONE Access UI.	None

Table 2-10. Design Decisions on Load Balancing for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-023	<ul style="list-style-type: none"> ■ Configure the NSX for vSphere load balancer application profile, for example, <code>wsa-cookie-persistence-profile</code>, for the cross-region Workspace ONE Access cluster to use an SSL Configuration. ■ Set the SSL Configuration for the Client SSL: <ul style="list-style-type: none"> ■ Client Certificate: Cross-Region Workspace ONE Access Cluster Certificate ■ Cipher: default. ■ Set the SSL Configuration for the Server SSL: <ul style="list-style-type: none"> ■ Client Certificate: Cross-Region Workspace ONE Access Cluster Certificate ■ SSL Profile: default. 	End-to-end SSL is required to support load balancing for the cross-region Workspace ONE Access cluster deployment type.	<ul style="list-style-type: none"> ■ You must manage the life cycle of the certificate used on the load balancer for the cross-region Workspace ONE Access cluster.
SDDC-IAM-024	<ul style="list-style-type: none"> ■ Add an NSX for vSphere load balancer virtual server, for example, <code>wsa-https</code>, for the cross-region Workspace ONE Access cluster to use the L7 HTTP type and port 443. ■ Set Acceleration to Disabled. ■ Set the IP for the Load Balancer. ■ Set the application profile, for example, <code>wsa-cookie-persistence-profile</code>. ■ Set the server pool to use the cross-region Workspace ONE Access cluster server pool, for example, <code>wsa-server-pool</code>. 	<ul style="list-style-type: none"> ■ The virtual server receives all the client connections and distributes them among the pool members based on the state of the pool members. ■ The cross-region Workspace ONE Access cluster requires cookie session persistence. 	None

Table 2-10. Design Decisions on Load Balancing for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-025	<ul style="list-style-type: none"> ■ Add another NSX for vSphere load balancer HTTP application profile, for example, <code>wsa-http-profile-redirect</code>, for the cross-region Workspace ONE Access to redirect HTTP to HTTPS. ■ Set the Timeout to 3600 seconds (60 minutes). ■ Set Redirection to HTTP to HTTPS Redirect. 	<ul style="list-style-type: none"> ■ Ensures that connections to non-secure HTTP are automatically redirected to HTTPS for the cross-region Workspace ONE Access cluster. ■ The cross-region Workspace ONE Access cluster requires a longer timeout. 	None
SDDC-IAM-026	<ul style="list-style-type: none"> ■ Add another NSX for vSphere load balancer virtual server, for example, <code>wsa-http-redirect</code> for the cross-region Workspace ONE Access cluster HTTP to HTTPS redirection to use the L7 HTTP type and Port 80. ■ Set Acceleration to Disable. ■ Set the IP Address for the Load Balancer to the same IP Address used for the HTTPS virtual server, for example, <code>wsa-https</code>. ■ Set the application profile to the HTTP to HTTPS Redirect profile, for example, <code>wsa-http-app-profile-redirect</code>. 	Ensures that connections to non-secure HTTP are automatically redirected to HTTPS for the cross-region Workspace ONE Access cluster.	None

Information Security and Access Control Design for Workspace ONE Access

You manage access to your Workspace ONE Access deployments by assigning users and groups to Workspace ONE Access roles.

Authentication and Authorization for Workspace ONE Access

In Workspace ONE Access, you can assign users three types of role-based access.

Table 2-11. Workspace ONE Access Roles

Role	Description	Example Enterprise Group
Super Admins	A role with the privileges to administer all Workspace ONE Access services and settings.	rainpole.local\ug-wsa-admins
Directory Admins	A role with the privileges to administer Workspace ONE Access users, groups, and directory management.	rainpole.local\ug-wsa-directory-admins
ReadOnly Admins	A role with read-only privileges to Workspace ONE Access.	rainpole.local\ug-wsa-readonly

For more information on roles and their permissions, refer to the Workspace ONE Access documentation.

As the cloud administrator for Workspace ONE Access, you establish an integration with your corporate directories which allows you to use your organization identity source for authentication. You can also set up a multi-factor authentication as part of access policy settings.

Identity and Access Management allows you to control authorization to your SDDC solutions - vRealize Suite Lifecycle Manager, vRealize Log Insight, vRealize Operations, and vRealize Automation - by assigning roles to your organization directory groups, such as Active Directory security groups.

Assigning roles to groups is more efficient than assigning roles to individual users. As a cloud administrator, you determine the members that make up your groups and what roles they are assigned. Groups in the connected directories are available for use Workspace ONE Access. In this design, enterprise groups are used to assign roles in Workspace ONE Access.

Table 2-12. Design Decisions on Authentication and Authorization for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-027	Rotate the appliance root user password on a schedule post-deployment.	The password for the root user account expires 60 days after the initial deployment.	<p>You must manage the password rotation schedule for the root user account in accordance with your organization policies and regulatory standards, as applicable.</p> <p>You must manage the password rotation schedule on the region-specific Workspace ONE Access instances and the cross-region Workspace ONE Access cluster.</p>
SDDC-IAM-028	Rotate the appliance sshuser user password on a schedule post-deployment.	The password for the appliance sshuser user account expires 60 days after the initial deployment.	<p>You must manage the password rotation schedule for the appliance sshuser user account in accordance with your organization policies and regulatory standards, as applicable.</p> <p>You must manage the password rotation schedule on the region-specific Workspace ONE Access instances and the cross-region Workspace ONE Access cluster.</p>
SDDC-IAM-029	Rotate the admin application user password on a schedule post-deployment.	The password for the default administrator application user account does not expire after the initial deployment.	<p>You must manage the password rotation schedule for the admin application user account in accordance with your organization policies and regulatory standards, as applicable.</p> <p>You must manage the password rotation schedule on the region-specific Workspace ONE Access instances and the cross-region Workspace ONE Access cluster.</p> <p>You must use the API to manage the Workspace ONE Access local directory user password changes.</p>

Table 2-12. Design Decisions on Authentication and Authorization for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-030	Rotate the password of the configuration administrator application user, for example, configadmin , on a schedule post-deployment for the cross-region Workspace ONE Access cluster.	The password for the configuration administrator application user account does not expire after the initial deployment.	<p>You must manage the password rotation schedule for the configuration administrator application user account in accordance with your organization policies and regulatory standards, as applicable.</p> <p>You must manage the password rotation schedule on the cross-region Workspace ONE Access cluster.</p> <p>You must use the API to manage the Workspace ONE Access local directory user password changes.</p>
SDDC-IAM-031	Configure a password policy for Workspace ONE Access local directory users, for example admin and configadmin .	<p>Allows you to set a policy for Workspace ONE Access local directory users that addressed your organization policies and regulatory standards.</p> <p>Note The password policy is applicable only to the local directory users and does not impact your organization directory.</p>	<p>You must set the policy in accordance with your organization policies and regulatory standards, as applicable.</p> <p>You must apply the password policy on the region-specific Workspace ONE Access instances and the cross-region Workspace ONE Access instance.</p>
SDDC-IAM-032	Assign roles to groups, synchronized from your corporate identity source for Workspace ONE Access.	It allows access management and administration of Workspace ONE Access by using corporate security group membership.	You must define and manage security groups, group membership and, security controls in your corporate identity source for Workspace ONE Access administrative consumption.
SDDC-IAM-033	Create a security group in your organization directory services for the Super Admin role and synchronize the group in the Workspace ONE Access configuration.	Allows you to streamline the management of Workspace ONE Access roles to users.	<p>You must create the security group outside of the SDDC stack.</p> <p>You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.</p>

Table 2-12. Design Decisions on Authentication and Authorization for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-034	Assign the enterprise group for administrators, for example, rainpole.local\ug-wsa-admins , the Super Admins Workspace ONE Access role.	Provides the following access control features: <ul style="list-style-type: none"> ■ Access to Workspace ONE Access services is granted to a managed set of individuals that are members of the security group. ■ Improved accountability and tracking access to Workspace ONE Access. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-IAM-035	Create a security group in your organization directory services for the Directory Admin role and synchronize the group in the Workspace ONE Access configuration.	Allows you to streamline the management of Workspace ONE Access roles to users.	You must create the security group outside of the SDDC stack. You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-IAM-036	Assign the enterprise group for directory administrator users, for example, rainpole.local\ug-wsa-directory-admins , the Directory Admins Workspace ONE Access role.	Provides the following access control features: <ul style="list-style-type: none"> ■ Access to Workspace ONE Access services is granted to a managed set of individuals that are members of the security group. ■ Improved accountability and tracking access to Workspace ONE Access. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-12. Design Decisions on Authentication and Authorization for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-037	Create a security group in your organization directory services for the ReadOnly Admin role and synchronize the group in the Workspace ONE Access configuration.	Allows you to streamline the management of Workspace ONE Access roles to users.	<p>You must create the security group outside of the SDDC stack.</p> <p>You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.</p>
SDDC-IAM-038	Assign the enterprise group for read-only users, for example, rainpole.local\ug-wsa-readonly , the ReadOnly Admin Workspace ONE Access role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to Workspace ONE Access services is granted to a managed set of individuals that are members of the security group. ■ Improved accountability and tracking access to Workspace ONE Access. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Note In an Active Directory forest, consider using a security group with a universal scope. Add security groups with a global scope that includes service accounts and users from the domains in the Active Directory forest.

Encryption Design for Workspace ONE Access

The Workspace ONE Access user interface and API endpoint use an HTTPS connection.. By default, Workspace ONE Access uses a self-signed certificate. To provide secure access to the Workspace ONE Access user interface and API, replace the default self-signed certificates with a CA-signed certificate.

Table 2-13. Design Decisions for Workspace ONE Access Encryption

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-039	Replace the default self-signed certificates with a Certificate Authority-signed certificate during deployment.	Ensures that all communications to the externally facing Workspace ONE Access browser-based UI, API, and between the components are encrypted.	Replacing the default certificates with trusted CA-signed certificates from a certificate authority increases the deployment preparation time as certificates requests are generated and delivered. You must manage the life cycle of the certificate replacement. You must use a multi-SAN certificate for the cross-region Workspace ONE Access cluster instance.
SDDC-IAM-040	Import the certificate for the Root Certificate Authority to each Workspace ONE Access instance.	Ensures that the certificate authority is trusted by each Workspace ONE Access instance.	None

Branding Design for Workspace ONE Access

You can change the appearance of the Workspace ONE Access browser-based user interface to meet minimal branding guidelines of an organization.

You can change the logo, the background and text colors, or the information in the header and footer.

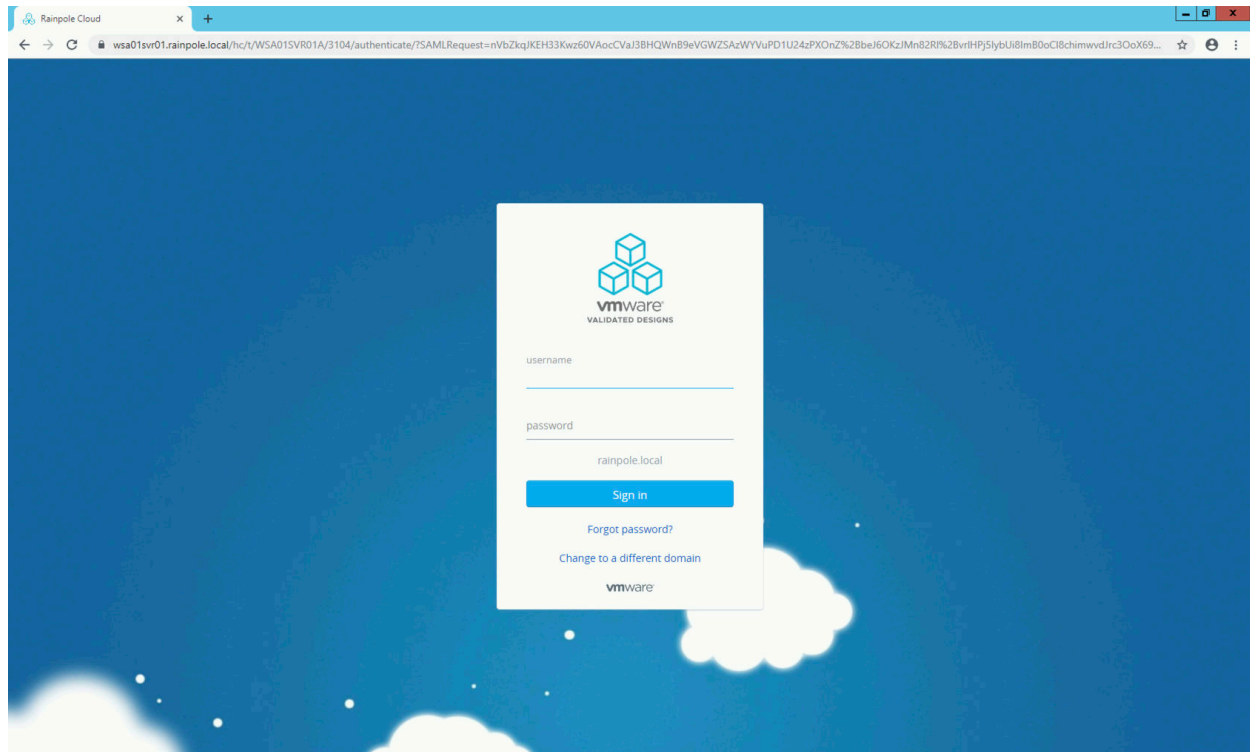
Figure 2-3. Example of Branding

Table 2-14. Design Decisions on Branding for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-041	Apply branding customizations for the Workspace ONE Access user interface that is presented to users when logging into integrated SDDC solutions.	Provides minimal corporate branding to the user interface consumed by end users. <ul style="list-style-type: none"> ■ Company Name ■ Product ■ Favorite Icon ■ Logo Image ■ Background Image ■ Colors 	You must provide an icon, logo, and background image icon that meets the minimum the correct size and resolution.

Integration Design for Workspace ONE Access

You integrate the SDDC solutions with Workspace ONE Access.

In this design, you configure the SDDC solutions with Workspace ONE Access to enable authentication through the identity and access management services. SDDC Solutions are integrated with a Workspace ONE Access instance that matches its location - region-specific or cross-region. Once enabled, information security and access control configurations for the integrated SDDC products can be configured.

Table 2-15. Design Decisions on Integrations in Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-042	Configure the region-specific SDDC solutions with the region-specific Workspace ONE Access instance as the authentication provider: <ul style="list-style-type: none"> ■ vRealize Log Insight Cluster 	Enables authentication through Workspace ONE Access identity and access management services for region-specific enabled SDDC solutions. Allows users to authenticate to a region-specific SDDC solution in the event of connectivity loss between regions.	None
SDDC-IAM-043	Configure the cross-region SDDC solutions with the cross-region Workspace ONE Access instance as the authentication provider: <ul style="list-style-type: none"> ■ vRealize Operations ■ vRealize Automation 	Enables authentication through Workspace ONE Access identity and access management services for cross-region enabled SDDC solutions. Required for vRealize Automation authentication.	The cross-region Workspace ONE Access cluster must be online and operational before you can authenticate to vRealize Automation.

Table 2-15. Design Decisions on Integrations in Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-044	Add Web Application Links to Workspace ONE Access catalog for the SDDC solutions: <ul style="list-style-type: none"> ■ vRealize Suite Lifecycle Manager Instance ■ vRealize Log Insight clusters ■ vRealize Operations cluster ■ vRealize Automation cluster 	Provides a catalog of integrated SDDC solutions from the Workspace ONE Access user portal to authenticated users.	None
SDDC-IAM-045	Assign Web Application Links to ALL USERS and an Automatic deployment type.	Provides a web application link to each Workspace ONE Access enabled SDDC solution within Workspace ONE Access for authenticated users.	None

Directories and Identity Provider Design for Workspace ONE Access

You integrate your enterprise directory with Workspace ONE Access to synchronize users and groups to the Workspace ONE Access identity and access management services. You configure the Workspace ONE Access identity provider connector to perform the synchronization of users and groups from your organization directory.

Directories

Workspace ONE Access has its own concept of a directory, corresponding to Active Directory or LDAP directories in your environment. This internal Workspace ONE Access directory uses attributes to define users and groups. You create one or more directories in the identity and access management service and then synchronize each directory with your corresponding Active Directory or LDAP directory. Workspace ONE Access integrates with the following types of directories:

Table 2-16. Supported External Directories in Workspace ONE Access

Directory Type	Description
Active Directory over LDAP	You create this directory type if you plan to connect to a single Active Directory domain environment. The connector binds to Active Directory using simple bind authentication. If you have more than one domain in a forest, you create a directory for each domain.
Active Directory over Integrated Windows Authentication	You create this directory type if you plan to connect to a multi-domain or multi-forest Active Directory environment. The connector binds to Active Directory using Integrated Windows Authentication. The type and number of directories that you create vary depending on your Active Directory environment, such as single domain or multi-domain, and on the type of trust used between domains. In most environments, you create a single directory.
LDAP directory	Create the LDAP directory to integrate your enterprise LDAP directory with VMware Identity Manager. You can only integrate a single-domain LDAP directory. VMware Identity Manager supports only those OpenLDAP implementations that support paged search queries.

Integration of Workspace ONE Access requires the following:

- Specify the attributes for users required in the Workspace ONE Access service.
- Add a directory in the Workspace ONE Access for the directory type for your organization.
- Map user attributes between your organization directory and Workspace ONE Access.
- Specify and synchronize directory users and groups.
- Establish a synchronization schedule or synchronize on-demand.

Table 2-17. Design Decisions on Directories for Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-046	Configure a directory service connection, for example, <code>rainpole.local</code> , for the Workspace ONE Access instance for each region and the cross-region Workspace ONE Access cluster.	Allows you to integrate your corporate directory with Workspace ONE Access to synchronize users and groups to the Workspace ONE Access identity and access management services.	None
SDDC-IAM-047	Use Active Directory with Integrated Windows Authentication as the Directory Service connection option.	Integrated Windows Authentication supports establishing trust relationships in a multi-domain or multi-forest Active Directory environment.	The Workspace ONE Access appliances must be joined to the Active Directory domain. Refer to Active Directory Environments in the Workspace ONE Access for more information on integration with differing Active Directory Forest and Domain models.

Table 2-17. Design Decisions on Directories for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-048	Configure the directory synchronization to only synchronize groups required for the integrated SDDC solutions.	<p>Allows you to limit the number of replicated groups required for each product.</p> <p>Reduces the replication interval for group information. Refer to the Information Security and Access Control section of the design for each integrated product.</p> <ul style="list-style-type: none"> ■ Workspace ONE Access ■ vRealize Log Insight ■ vRealize Operations ■ vRealize Automation. 	You must manage the groups from your organization directory selected for synchronization to the Workspace ONE Access directory.
SDDC-IAM-049	Enable the synchronization of group members to the directory when a group is added to the Workspace ONE Access directory.	It allows members of the groups to be synced to Workspace ONE Access when groups are added from the corporate directory. When disabled, group names are synced to the directory, but members of the group are not synced until the group is entitled to an application or the group name is added to an access policy.	None
SDDC-IAM-050	Enable Workspace ONE Access to synchronize nested group members by default.	Allows Workspace ONE Access to update and cache the membership of groups without querying your organization directory.	Changes to a group membership are not reflected until the next synchronization event.
SDDC-IAM-051	Add a filter to the directory settings to exclude users from the directory replication.	<p>Allows you to limit the number of replicated users for each Workspace ONE Access within the maximum scale.</p> <ul style="list-style-type: none"> ■ Region-specific instance: 1,000 users accounts (each) ■ Cross-region cluster: 10,000 user accounts 	You must define a filtering schema that works for your organization based on your directory attributes to ensure that replicated user accounts are managed within the maximums.

Table 2-17. Design Decisions on Directories for Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-052	Configure the mapped attributes included when a user is added to the Workspace ONE Access directory.	Allows you to configure the minimum required and extended user attributes to synchronize directory user account for Workspace ONE Access to be used as an authentication source for SDDC Solutions.	<p>User accounts in your organization must have the following required attributes mapped:</p> <ul style="list-style-type: none"> ■ <code>firstname</code>, for example, <code>givenname</code> for Active Directory ■ <code>lastName</code>, for example, <code>sn</code> for Active Directory ■ <code>email</code>, for example, <code>mail</code> for Active Directory ■ <code>userName</code>, for example, <code>sAMAccountName</code> for Active Directory ■ If you require users to sign in with an alternate unique identifier, for example, <code>userPrincipalName</code>, you must map the attribute and update the Identity and access management preferences.
SDDC-IAM-053	Configure the directory synchronization frequency to a reoccurring schedule, for example, 15 minutes.	It ensures that any changes to group memberships in the corporate directory are available for integrated SDDC solutions in a timely manner.	Schedule the synchronization interval to be longer than the time to synchronize from the corporate directory. If users and groups are being synchronized to Workspace ONE Access when the next synchronization is scheduled, the new synchronization starts immediately after the end of the previous iteration. With this schedule, the process is continuous.

Identity Providers and Connectors

Workspace ONE Access synchronizes with corporate directories, for example, Active Directory `rainpole.local`, by using the connector component. Any required users and groups that are provided access to the SDDC solutions are synchronized into Workspace ONE Access. In addition, the connector is the default identity provider and authenticates users to the identity and access management service. Authentication uses your organization directory, but searches are made against the local Workspace ONE Access directory mirror. You can configure high availability for directory synchronization by associating the directory with multiple connector instances.

Table 2-18. Design Decisions on Identity Providers and Connectors in Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-IAM-054	For the cross-region Workspace ONE Access cluster, configure second and third connectors that correspond to the second and third Workspace ONE Access appliances to support the high-availability of directory services access.	This design supports high availability by installing three Workspace ONE Access appliances load-balanced by an NSX for vSphere load-balancer instance. Adding the additional connectors provides redundancy and improves performance by load balancing authentication requests.	None

Monitoring and Alerting Design for Workspace ONE Access

You integrate vRealize Operations with Workspace ONE Access to provide operational visibility.

The integration to Workspace ONE Access from vRealize Operations provides the ability to monitor the health, efficiency, and capacity risks associated with the identity and access management services.

For more information, see the vRealize Operations section of the design.

Data Protection Design for Workspace ONE Access

To preserve the identity management and access services functionality when data or system loss occurs, the design supports the use of data protection.

Workspace ONE Access supports data protection through the creation of consistent image-level backups, using backup software that is based on the vSphere Storage APIs - Data Protection (VADP).

Disaster Recovery Design for Workspace ONE Access

To preserve the identity management and access services functionality when a disaster occurs, the design supports the failover of the Workspace ONE Access cluster from cross-region SDDC solutions between regions.

You place Workspace ONE Access on the cross-region virtual network in the management domain. As a result, after recovery, you continue using the same IP addresses, DNS records, and routing configuration. Workspace ONE Access also uses this network for its cross-region failover capabilities.

If a planned migration or disaster occurs, you use Site Recovery Manager and vSphere Replication for an orchestrated recovery of the Workspace ONE Access cluster. After the recovery, Workspace ONE Access continues to provide identity and access management services to cross-region SDDC solutions.

Cloud Automation Design

The VMware Cloud Automation design includes Cloud Assembly and Service Broker design. The design provides guidance on configuration, organization, and consumption of the services to enable declarative blueprint orchestration in a multi-cloud environment.

■ [vRealize Automation Design](#)

The Cloud Automation layer using vRealize Automation provides the consumption model for the SDDC. All the IT services that the vRealize Automation self-service portal offers allow end users to provision workloads in an automated way while taking full advantage of the powerful compute, storage, security, and networking features and capabilities offered by the SDDC.

vRealize Automation Design

The Cloud Automation layer using vRealize Automation provides the consumption model for the SDDC. All the IT services that the vRealize Automation self-service portal offers allow end users to provision workloads in an automated way while taking full advantage of the powerful compute, storage, security, and networking features and capabilities offered by the SDDC.

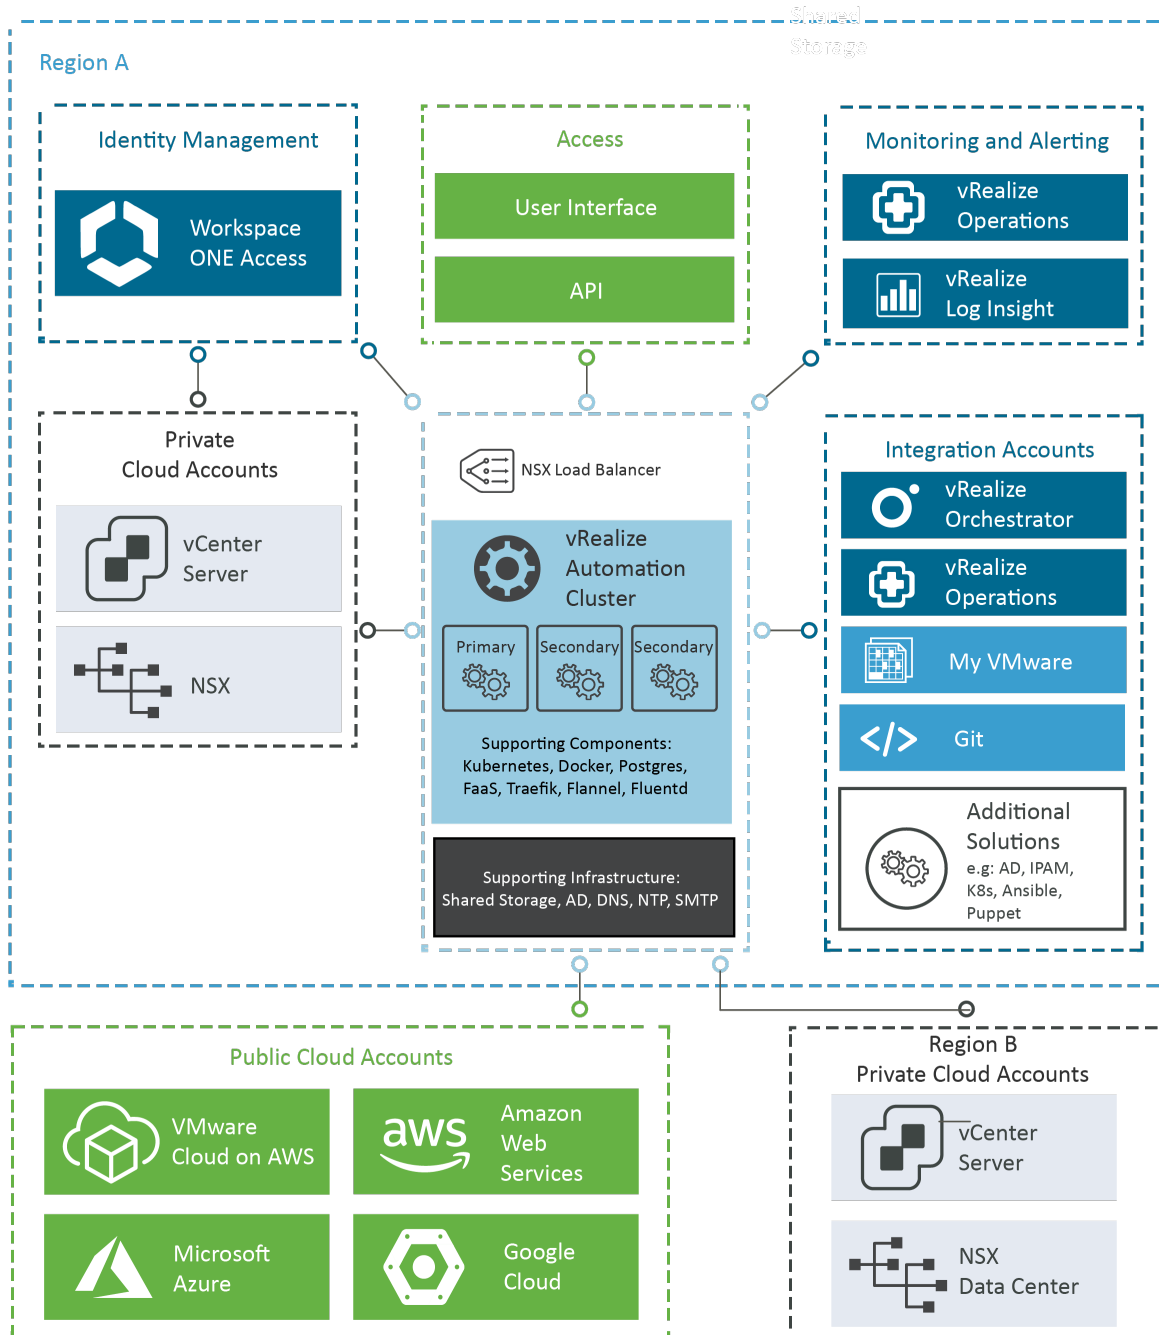
Owing to the services offered, the design objectives of the Cloud Automation layer include:

- Accessibility to all layers of the SDDC
- Resiliency and high availability across a dual region deployment
- Scalability to meet the provisioning requirements of multiple tenants
- Ease of use for consumers
- Ability to provision to multiple clouds
- Provisioning of complex workloads, for example, three-tier applications involving many components native and external to the SDDC
- Simplicity for easier life cycle management of the products in this layer

While some of the design objectives are native to the vRealize Automation product, a set of design and deployment decisions ensures that all Cloud Automation design objectives are met.

Logical Design for vRealize Automation

The logical design of vRealize Automation includes all the integrations of vRealize Automation with other components of the SDDC, supporting infrastructure, public cloud solutions, and external entities, such as Git, IPAM, and Puppet. The networking, identity and access management, product configurations, and secure access must be carefully designed for seamless integrations between vRealize Automation and other components.

Figure 2-4. Logical Design

User Access

vRealize Automation provides a UI and RESTful API that for initiating vRealize Automation services.

Cloud Accounts

vRealize Automation can simplify the multi-cloud experience by managing and working with resources in private cloud and public cloud infrastructures. Each supported type of infrastructure is represented by a cloud account.

■ Private Cloud

- Public Cloud

Integrations

- My VMware
- vRealize Operations Manager
- vRealize Log Insight

The vRealize Operations Management Pack for vRealize Automation provides performance and capacity metrics of tenant business groups and underlying cloud infrastructure.

Supporting Infrastructure

vRealize Automation integrates with the following supporting infrastructure: You can assign users two types of role-based access:

- DNS for providing name resolution for the vRealize Automation components
- NTP servers for providing time synchronization for the vRealize Automation components
- Workspace ONE Access, eventually with Active Directory and LDAP, for tenant user authentication and authorization
- SMTP for sending and receiving notification emails for various actions that can be run in the vRealize Automation console

NSX-T Data Center

The integration of vRealize Automation with NSX-T supports designing and authoring blueprints by using the networking and security features of NSX-T. You can use all NSX-T network constructs, such as logical switches, distributed logical routing, and distributed firewalls.

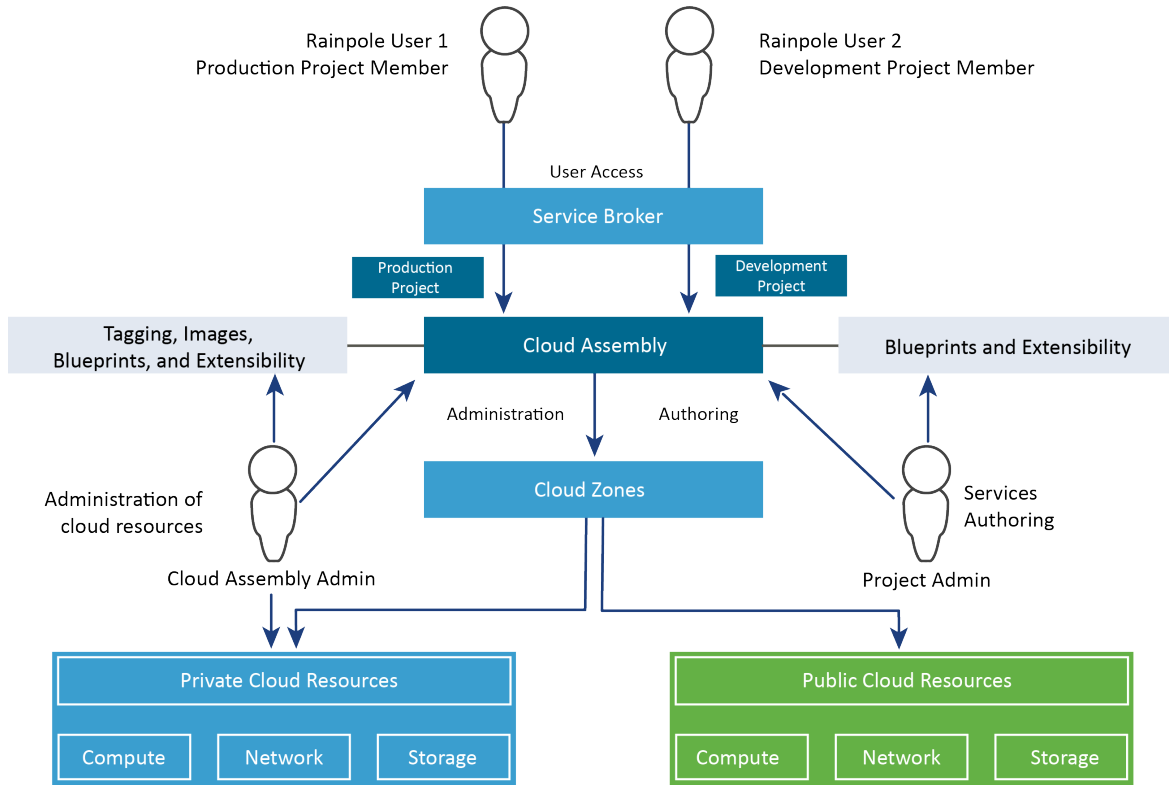
In a blueprint, you can place an on-demand load balancer, NAT network, routed network, and security groups. When a user requests the blueprint, vRealize Automation provisions these constructs in NSX-T.

You can configure automated network provisioning as a part of the blueprint design instead of as a separate operation outside vRealize Automation.

Usage Model

vRealize Automation enables a usage model that includes interaction between the cloud automation services, the supporting infrastructure, and the provisioning infrastructure. The usage model of vRealize Automation contains the following elements and components in them:

Figure 2-5. vRealize Automation Usage Model



Users

Cloud, tenant, group, fabric, infrastructure, service, and other administrators as defined by business policies and organization structure. Cloud (or tenant) users in an organization can provision virtual machines and directly perform operations on them at the level of the operating system.

Blueprints

VM templates and blueprints. VM templates are used to author the blueprints that tenants (business users) use to request workloads.

Images and Flavors

Image and flavor mappings simplify the blueprint creation while adding greater flexibility and customization.

An image mapping groups a set of predefined target OS specifications for a specific cloud account/region in vRealize Automation Cloud Assembly by using natural language naming.

A flavor mapping groups a set of target deployment sizings for a specific cloud account/region in vRealize Automation Cloud Assembly by using natural language naming.

Provisioning infrastructure

Private and public cloud resources which together form a hybrid cloud.

Private cloud resources are supported hypervisors and associated management tools.

Public cloud resources are supported cloud providers and associated APIs.

Cloud Assembly

Self-service capabilities for users to administer, provision, and manage workloads.

The default tenant uses the vRealize Automation administrator portal to set up and administer tenants and global configuration options.

A custom tenant uses the vRealize Automation tenant portal, which you access by appending a tenant identifier.

vRealize Orchestrator

VMware vRealize Orchestrator provides a standard set of plug-ins, including a plug-in for vCenter Server, with which you can orchestrate tasks in the different environments that the plug-ins expose.

Service Broker

Aggregates native content from multiple clouds and platforms into a single catalog with role-based policies.

Configuration Design for vRealize Automation

The configuration design consists of characteristics and decisions that support the logical design and satisfy the design objectives listed in the Logical Design section.

To accomplish this, you require the following components to deploy the cloud automation solution for the SDDC.

- SDDC Manager
- vRealize Suite Lifecycle Manager
- NSX for vSphere Load Balancer for Workspace ONE Access
- Workspace ONE Access cluster
- Supporting infrastructure services, such as Active Directory, DNS, NTP, and SMTP

Deployment Model for vRealize Automation

vRealize Automation is distributed as a virtual appliance in OVA format. The vRealize Automation appliance includes the cloud automation services, including an embedded vRealize Orchestrator, and database services.

You consider the deployment type - standard or cluster - according to the design objectives for the availability and number of workloads that the system must support. vRealize Automation is deployed to the first cluster in the management domain and it uses the vRealize Suite Lifecycle Manager instance as its lifecycle management engine.

In this design you deploy a vRealize Automation cluster - three (3) virtual appliances and load balancer - and the supporting Workspace ONE Access service into the first cluster in the management domain of Region A. vSphere High Availability protects vRealize Automation by restarting each virtual appliance on an alternate ESXi host if a primary ESXi host failure occurs. vSphere Distributed Resource Scheduler anti-affinity rules prevent vRealize Automation virtual appliances in the vRealize Automation cluster from running on the same ESXi host.

All vRealize Automation services and databases are configured for high availability using the underlying Kubernetes cluster and pods within the appliances. The vRealize Automation cluster manages the workloads in each region.

Sizing Compute and Storage Resources for vRealize Automation

A vRealize Automation cluster deployment has the following resource requirements.

Table 2-19. Design Decisions for vRealize Automation Deployment

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-001	Deploy a single vRealize Automation instance in the first cluster in the management domain in Region A to provide cloud automation services to all regions.	<ul style="list-style-type: none"> ■ vRealize Automation can manage one or more regions and provides a single cloud automation service, regardless of region. ■ vRealize Automation can also manage VMware Cloud on AWS and public cloud instances. ■ Because of the abstraction of the vRealize Automation over virtual networking, it is independent of any physical site locations or hardware. 	You must use a vRealize Automation deployment type that accommodates multi-region deployments.
SDDC-CAS-002	Deploy vRealize Automation through SDDC Manager using the cluster deployment type.	<ul style="list-style-type: none"> ■ Deploying the cluster configuration that includes the three-node appliance architecture satisfies the design objectives in scope for the design. ■ This design enables the future growth of virtual machines after you expand the cloud infrastructure. 	None
SDDC-CAS-003	Protect all vRealize Automation virtual appliances by using vSphere High Availability.	Supports the availability objectives for vRealize Automation without a required manual intervention during a failure event.	None

Table 2-19. Design Decisions for vRealize Automation Deployment (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-004	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules for the vRealize Automation virtual appliances.	Using vSphere DRS prevents vRealize Automation virtual appliance from residing on the same ESXi host and risking the high availability of the deployment.	<ul style="list-style-type: none"> ■ You can only place a single ESXi host at a time into maintenance mode for a management cluster of four ESXi hosts. ■ Requires at least four physical hosts to guarantee the three vRealize Automation virtual appliances continue to run if an ESXi host failure occurs.
SDDC-CAS-005	Add the VM groups for the vRealize Automation virtual appliances and set VM rules to restart the Workspace ONE Access VM group before the vRealize Automation VM group.	Defines the startup order of virtual appliances regarding service dependency. The startup order ensures that vSphere HA powers on the virtual machines for vRealize Automation in the correct order.	None
SDDC-CAS-006	When using two availability zones in Region A, add the vRealize Automation virtual appliances to the primary availability zone VM group, for example, sfo01-m01-mgmt01-primary-az-vm-group.	It ensures the vRealize Automation virtual appliances are powered on within the primary availability zone hosts group by default.	If vRealize Automation is deployed after the creation of the stretched clusters for management domain availability zones, the VM Group for the primary availability zone virtual machines must be updated to include the vRealize Automation virtual appliances.
SDDC-CAS-007	Place all cross-region vRealize Automation virtual appliances in a dedicated virtual machine folder in Region A, for example, xregion-sfo01-lax01-m01fd-vra.	It provides the organization of cross-region vRealize Automation virtual appliances in the management domain inventory and preparation for Site Recovery Manager folder mappings for disaster recovery.	A corresponding virtual machine folder in Region B must be created in preparation for Site Recovery Manager folder mapping, for example, xregion-lax01-sfo01-m01fd-vra.

Logging Design for vRealize Automation

You integrate vRealize Automation with vRealize Log Insight to provide operational visibility.

The native integration to vRealize Log Insight from vRealize Automation enables sending logs from the service containers for aggregation and analysis, as needed.

Logging to a vRealize Log Insight instance through the ingestion API is established using the vRealize Automation Command-line Interface.

```
root@vra01svr01a [ ~ ]# vracli vrli set -k -e cross-region-production http://
sfo01vrli01.sfo01.rainpole.local:9000 root@vra01svr01a [ ~ ]# vracli vrli { "agentId": "0",
"environment": "cross-region-production", "host": "sfo01vrli01.sfo01.rainpole.local", "port": 9000,
"scheme": "http", "sslVerify": false
```

Table 2-20. Design Decisions on Logging for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-008	Configure vRealize Automation to send logs to the vRealize Log Insight.	It allows logs from vRealize Automation services to be forwarded from Fluentd to vRealize Log Insight.	A vRealize Log Insight content pack for vRealize Automation is not available.
SDDC-CAS-009	Communicate with the vRealize Log Insight by using the default ingestion API, cfapi, port 9000, and a non-default SSL number.	Supports disaster recovery of vRealize Automation in the SDDC. During the failover, the DNS records for vRealize Log Insight in Region A are update to redirect to the instance in Region B to ensure the log collection remains operational. The <code>ssl=no</code> setting must be used when there is a certificate mismatch post failover.	Transmission traffic for logs is not secure.

For more information, refer to the vRealize Log Insight section of the design.

Network Design for vRealize Automation

For secure access to the UI and API and for failover of vRealize Automation, you deploy the virtual appliance in the cross-region application virtual network.

Application Virtual Network

The vRealize Automation cluster virtual appliances are connected to the cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN, for secure access to the UI and API, and for failover support.

As part of this design, use the application virtual network configuration to connect vRealize Automation with the other management solutions in the SDDC. Use the load balancer in the NSX for vSphere edge for the cross-region application virtual network for high availability and balancing user access across the vRealize Automation cluster.

This design uses NSX for vSphere application virtual networks to abstract vRealize Automation and its supporting services. You can place them in a single designated region, for example, Region A by default, regardless of the underlying physical infrastructure, such as network subnets, compute hardware, or storage types.

This networking design has the following features:

- The vRealize Automation appliances are deployed together on the same application virtual network in a designated region. This configuration provides the ability to fail over vRealize Automation between regions after expanding to a multi-region SDDC design.
- All vRealize Automation components have routed access to the VLAN-backed management network through the NSX Universal Distributed Logical Router.
- Routing to the VLAN-backed management network and other external networks is dynamic and is based on the Border Gateway Protocol (BGP).

Figure 2-6. Networking Design of the vRealize Automation Deployment

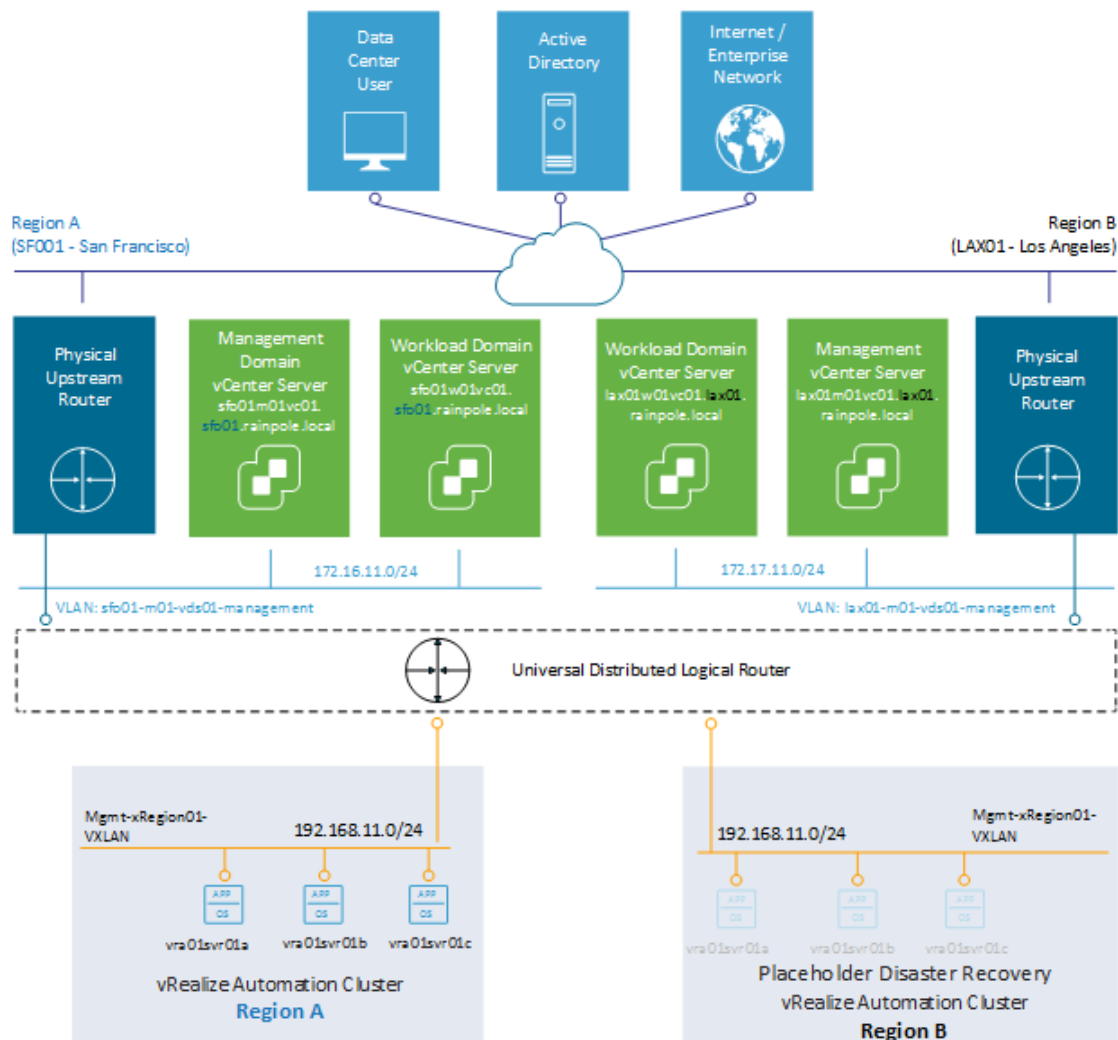


Table 2-21. Design Decisions on the Application Virtual Network for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-010	Place the vRealize Automation virtual appliances on the cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN.	Supports disaster recovery by deploying the vRealize Automation virtual appliances on the NSX for vSphere application virtual network in the management domain.	You must use an implementation in NSX for vSphere to support this network configuration.

IP Addressing Scheme

You allocate a subnet for the cross-region network segment in the management domain and use it for the vRealize Automation deployment.

Table 2-22. Example IP Subnet for vRealize Automation

Solution	IP Subnet	NSX for vSphere Application Virtual Network
vRealize Automation in Region A	192.168.11.0/24	Mgmt-xRegion01-VXLAN

Important The following network ranges are reserved for intra-service communication. vRealize Automation cannot be deployed with an IP address in this range, nor can vRealize Automation access external services with IP addresses in these ranges.

- 10.244.0.0/22
- 10.244.4.0/22

Name Resolution

The vRealize Automation appliances are resolvable by using domain name resolution.

The IP addresses of the vRealize Automation appliances are associated with a fully qualified name suffix in a designated domain name space, for example, `rainpole.local`.

Table 2-23. Example FQDNs and IP Addresses for vRealize Automation

Fully Qualified Domain Name	IP Address	Description	Region	Failed Over to Region B
vra01svr01.rainpole.local	192.168.11.50	vRealize Automation Cluster NSX Load-Balancer Virtual Server	Region A	✓
vra01svr01a.rainpole.local	192.168.11.51	vRealize Automation Virtual Appliance A	Region A	✓
vra01svr01b.rainpole.local	192.168.11.52	vRealize Automation Virtual Appliance B	Region A	✓
vra01svr01c.rainpole.local	192.168.11.53	vRealize Automation Virtual Appliance C	Region A	✓

Table 2-24. Design Decisions on DNS for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-011	Configure forward and reverse DNS records for each vRealize Automation appliance IP address and the load-balancer virtual IP address.	vRealize Automation is accessible by using a fully qualified domain name instead of by using IP addresses only.	<ul style="list-style-type: none"> ■ You must provide DNS records for each vRealize Automation appliance and the load-balancer virtual IP address. ■ All firewalls located between the vRealize Automation virtual appliances and the DNS servers must allow DNS traffic.
SDDC-CAS-012	In a multi-region SDDC deployment, configure the DNS settings for each vRealize Automation appliance to use DNS servers in each region.	vRealize Automation appliances can resolve DNS from regional DNS servers during a planned migration or disaster recovery between regions.	As you scale from a single region to multi-region SDDC deployment, the DNS settings on each vRealize Automation appliance must be updated.

Time Synchronization

vRealize Automation is dependent on time synchronization for all appliances.

Table 2-25. Design Decisions on NTP for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-013	Configure NTP for each vRealize Automation virtual appliance.	vRealize Automation is dependent on time synchronization for all virtual appliances.	All firewalls located between the vRealize Automation virtual appliances and the NTP servers must allow NTP traffic.
SDDC-CAS-014	In a multi-region SDDC deployment, configure the NTP settings for each vRealize Automation appliance to use NTP servers in each region.	vRealize Automation appliances can query NTP from regional NTP servers to synchronize time during a planned migration or disaster recovery between regions.	As you scale from a single region to multi-region SDDC deployment, the NTP settings on each vRealize Automation appliance must be updated.

Load Balancing

A vRealize Automation cluster deployment requires a load balancer to manage connections to vRealize Automation services. The design uses load balancing services provided by NSX for vSphere in the management domain.

Table 2-26. Design Decisions on Load Balancing for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-015	Configure the NSX load balancer that was created in NSX for vSphere in the management domain to load balance Workspace ONE Access to load also balance connections across vRealize Automation cluster members.	Required to deploy vRealize Automation as a cluster deployment type, enabling it to handle a greater load and obtain a higher level of availability.	You must use an implementation in NSX for vSphere to support this network configuration.
SDDC-CAS-016	<ul style="list-style-type: none"> ■ Add an NSX load balancer service monitor, for example, <code>vra-https-monitor</code>, for vRealize Automation with an Active HTTP monitor on Monitoring Port 8008. ■ Set the intervals and timeouts for the monitor: <ul style="list-style-type: none"> ■ Monitoring Interval: 3 seconds ■ Timeout Period: 10 seconds ■ Max Retries : 3 seconds ■ Set the HTTP Request for the monitor: <ul style="list-style-type: none"> ■ HTTP Method: Get ■ Request URL: <code>/health</code> ■ Set the HTTP Response for the monitor: Expected: 200. 	<ul style="list-style-type: none"> ■ The vRealize Automation health check is provided over HTTP on port 8008. ■ The Active Monitor uses HTTP requests to monitor the application health reported by vRealize Automation. ■ Ensures that connections to unhealthy vRealize Automation members in the pool are disabled until a subsequent periodic health check finds the members to be healthy. 	None
SDDC-CAS-017	<ul style="list-style-type: none"> ■ Add an NSX load balancer server pool, for example, <code>vra-server-pool</code>, for vRealize Automation to use the LEASTCON algorithm. ■ Set the static members for the pool: <ul style="list-style-type: none"> ■ Name: <i>Hostname</i> ■ IP: <i>IP Address</i> ■ Port: 443 ■ Weight: 1 ■ State: Enabled. 	<ul style="list-style-type: none"> ■ Least Connection distributes requests to members based on the number of current connections. New connections are sent to the pool member with the fewest connections. ■ vRealize Automation services respond on TCP 443. 	None

Table 2-26. Design Decisions on Load Balancing for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-018	<ul style="list-style-type: none"> ■ Add an NSX load balancer Fast TCP application profile, for example, <code>vra-tcp-app-profile</code>, for vRealize Automation. ■ Set Application Profile Type to SSL Passthrough. ■ Set persistence to None ■ Set the Timeout to 1800 seconds (30 minutes). 	An application profile is required to set the required timeout for HTTPS requests to vRealize Automation.	None
SDDC-CAS-019	<ul style="list-style-type: none"> ■ Add an NSX load balancer virtual server, for example, <code>vra-https</code>, for vRealize Automation to use the L4 TCP type and port 443. ■ Set the acceleration to Enable. ■ Set the IP for the load balancer. ■ Set the application profile, for example, <code>vra-tcp-app-profile</code>. ■ Set the server pool to use the vRealize Automation server pool, for example, <code>vra-server-pool</code>. 	The virtual server receives all the client connections and distributes them among the pool members based on the state of the pool members.	None

Table 2-26. Design Decisions on Load Balancing for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-020	<ul style="list-style-type: none"> ■ Add an NSX load balancer HTTP application profile, for example, <code>vra-http-app-profile-redirect</code>, for vRealize Automation to redirect HTTP to HTTPS. ■ Set the Timeout to 1800 seconds (30 minutes). ■ Set Redirection to HTTP to HTTPS Redirect. 	Ensures that connections to non-secure HTTP are automatically redirected to HTTPS for vRealize Automation.	None
SDDC-CAS-021	<ul style="list-style-type: none"> ■ Add another NSX load balancer virtual server, for example, <code>vra-http-redirect</code>, for vRealize Automation HTTP to HTTPS redirection to use the L7 HTTP type and port 80. ■ Set the acceleration to Disable. ■ Set the IP address for the load balancer to the same IP address used for the HTTPS virtual server, for example, <code>vra-https</code>. ■ Set the application profile to the HTTP to HTTPS redirect profile, for example, <code>vra-http-app-profile-redirect</code>. 	Ensures that connections to non-secure HTTP are automatically redirected to HTTPS for vRealize Automation.	None

Information Security and Access Control Design for vRealize Automation

You manage access to your vRealize Automation organization by assigning users and groups, synchronized to Workspace ONE Access, to organization and service roles.

As an organization owner, you add users to your organization and provide access to the vRealize Automation services associated with it.

You can assign users two types of role-based access:

Organization Role A role within the vRealize Automation organization - owner or member.

Service Role A role within the vRealize Automation for the services within the cloud automation platform.

For more information, on organization roles and their permissions, refer to the vRealize Automation documentation.

As the cloud administrator for vRealize Automation, you establish an integration with your corporate directories which allows you to use your organization identity source for vRealize Automation authentication. You can also set up a multi-factor authentication as part of access policy settings.

After the integration, you can control authorization to your vRealize Automation organization and services by assigning organization and service roles to your enterprise directory groups, such as Active Directory security groups.

Assigning roles to groups is more efficient than assigning the roles to individual users. As an organization owner, you determine the members that make up your groups and what roles they are assigned. IN vRealize Automation, enterprise groups are groups that are derived from your Workspace ONE Access (formerly known as Identity Manager) connected directories and available for use in your organization. As an organization owner, you can add and change the role assignment for an enterprise group. In this design, enterprise groups are used to assign organization and service roles.

Authentication and Authorization

The organization owner role allows you add users and enterprise groups to your organization and provide access to the vRealize Automation services.

Table 2-27. Example Organization Owner Assignment

Role	Description	Enterprise Group
Organization Owner	Used to assign other enterprise groups to vRealize Automation services.	rainpole.local\ug-vra-org-owners

See the following sections for service roles for the vRealize Automation services in this design.

- Cloud Assembly Service Roles
- vRealize Orchestrator Service Roles
- Service Broker Service Roles

Table 2-28. Design Decisions on Authentication and Authorization for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-022	Rotate the root password on a schedule post-deployment.	The password for the root user account does not expire after the initial deployment.	You must manage the password rotation schedule for the root user account in accordance with your organization policies and regulatory standards, as applicable.
SDDC-CAS-023	Enable the vRealize Automation integration with your corporate identity source using Workspace ONE Access.	Allows authentication, including multi-factor, to vRealize Automation using your corporate identity source. Allows authorization through the assignment of organization and cloud services roles to enterprise users and groups defined in your corporate identity source.	You must deploy and configure the Workspace ONE Access to establish the integration between vRealize Automation and your corporate identity sources.
SDDC-CAS-024	Assign organization and service roles to designated enterprise groups, synchronized from your corporate identity source through Workspace ONE Access.	Allows access management and administration to vRealize Automation services by using corporate security group membership.	You must define and manage security groups, group membership and, security controls in your corporate identity source for vRealize Automation consumption.
SDDC-CAS-025	Create a security group in your organization directory services for the Organization Owner organization role and synchronize the group in the Workspace ONE Access configuration for vRealize Automation.	Allows you to grant a managed set of individuals the ability to assign enterprise groups to vRealize Automation service roles.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-CAS-026	Assign the enterprise group for organization owners, for example, rainpole.local\ug-vra-org-owners , the Organization Owner organization role. Service roles for vRealize Automation are not assigned to this group.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to vRealize Automation services, such as Cloud Assembly, are granted to a managed set of individuals that are members of a security group. ■ Improved accountability and tracking organization owner access to vRealize Automation. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-28. Design Decisions on Authentication and Authorization for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-027	Create a security group in your organization directory services for the Cloud Assembly Administrator service role and synchronize the group in the Workspace ONE Access configuration for vRealize Automation.	Allows you to streamline the management of vRealize Automation organization and service roles to users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-CAS-028	Assign the enterprise group for Cloud Assembly administrators, for example, rainpole.local\ug-vra-cloud-assembly-admins , the Organization Member organization role and Cloud Assembly Administrator service role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to the vRealize Automation Cloud Assembly service is granted to a managed set of individuals that are members of the security group. ■ Improved accountability and tracking access to the vRealize Automation service. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-CAS-029	Create a security group in your organization directory services for the Cloud Assembly User service role and synchronize the group in the Workspace ONE Access configuration for vRealize Automation.	Allows you to streamline the management of vRealize Automation organization and service roles to users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-CAS-030	Assign the enterprise group for Cloud Assembly users, for example, rainpole.local\ug-vra-cloud-assembly-users , the Organization Member organization role and Cloud Assembly User service role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to the vRealize Automation Cloud Assembly service is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking access to the vRealize Automation service. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-28. Design Decisions on Authentication and Authorization for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-031	Create a security group in your organization directory services for the Service Broker Administrator service role and synchronize the group in the Workspace ONE Access configuration for vRealize Automation.	It allows you to streamline the management of vRealize Automation organization and service roles to users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-CAS-032	Assign the enterprise group for Service Broker administrators, for example, rainpole.local\ug-vra-service-broker-admins the Organization Member organization role and Service Broker Administrator service role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to the vRealize Automation Service Broker service is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking access to the vRealize Automation service. 	<p>You must maintain the life cycle and availability of the security group outside of the SDDC stack.</p> <p>Important Project Administrators must be granted the Service Broker Administrator to perform customizations to blueprint icons and forms. However, members of this role are also entitled to manage cloud accounts, cloud zones, and integrations created by a Cloud Assembly administrator.</p>
SDDC-CAS-033	Create a security group in your organization directory services for the Service Broker User service role and synchronize the group in the Workspace ONE Access configuration for vRealize Automation.	Allows you to streamline the management of vRealize Automation organization and service roles to users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.

Table 2-28. Design Decisions on Authentication and Authorization for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-034	Assign the enterprise group for Service Broker users, for example, rainpole.local\ug-vra-service-broker-users the Organization Member organization role and Service Broker User service role.	Provides the following access control features: <ul style="list-style-type: none"> ■ Access to the vRealize Automation Service Broker service is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking access to the vRealize Automation service. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-CAS-035	Create a security group in your organization directory services for the Orchestrator Administrator service role and synchronize the group in the Workspace ONE Access configuration for vRealize Automation.	It allows you to streamline the management of vRealize Automation organization and service roles to users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-CAS-036	Assign the enterprise group for vRealize Orchestrator administrators, for example, rainpole.local\ug-vra-orchestrator-admins the Organization Member organization role and Orchestrator Administrator service role.	Provides the following access control features: <ul style="list-style-type: none"> ■ Access to the vRealize Automation Orchestrator service is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking access to the vRealize Automation service. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-28. Design Decisions on Authentication and Authorization for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-037	Create a security group in your organization directory services for the Orchestrator Workflow Designer service role and synchronize the group in the Workspace ONE Access configuration for vRealize Automation.	It allows you to streamline the management of vRealize Automation organization and service roles to users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-CAS-038	Assign the enterprise group for vRealize Orchestrator workflow designers, for example, rainpole.local\ug-vra-orchestrator-designers , the Organization Member organization role and Orchestrator Workflow Designer service role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to the vRealize Automation Orchestrator service is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking access to the vRealize Automation service. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Important In an Active Directory forest, consider using a security group with a universal scope. Add security groups with a global scope that includes service accounts and users from the domains in the Active Directory forest.

Figure 2-7. Example Service Role Assignment in vRealize Automation

The screenshot shows the 'Identity & Access Management' section of the vRealize Automation console. Under the 'Enterprise Groups' tab, there is a table listing various groups and their assigned roles. The table has columns for checkboxes, Enterprise Group Name, Domain, Member Count, Organization Role, and Service Roles. The groups listed include 'ug-vra-org-owners', 'ug-vra-orchestrator-designers', 'ug-vra-code-stream-admins', 'ug-vra-code-stream-users', 'ug-vra-service-broker-users', 'ug-vra-code-stream-executors', 'ug-vra-code-stream-viewers', 'ug-vra-service-broker-admins', 'ug-vra-cloud-assembly-users', 'ug-vra-orchestrator-admins', and 'ug-vra-cloud-assembly-admins'. The 'ug-vra-orchestrator-admins' group is expanded, showing two roles: 'Orchestrator Workflow Designer' and 'Orchestrator Administrator'. A 'SUPPORT' button is visible on the right side of the console.

<input type="checkbox"/>	Enterprise Group Name	Domain	Member Count	Organization Role	Service Roles
<input type="checkbox"/>	ug-vra-org-owners@rainpole.local	rainpole.local	3	Organization Owner	
<input type="checkbox"/>	ug-vra-orchestrator-designers@rainpole.local	rainpole.local	1	Organization Member	Orchestrator Workflow Designer
<input type="checkbox"/>	ug-vra-code-stream-admins@rainpole.local	rainpole.local	2	Organization Member	Code Stream Administrator
<input type="checkbox"/>	ug-vra-code-stream-users@rainpole.local	rainpole.local	1	Organization Member	Code Stream User
<input type="checkbox"/>	ug-vra-service-broker-users@rainpole.local	rainpole.local	1	Organization Member	Service Broker User
<input type="checkbox"/>	ug-vra-code-stream-executors@rainpole.local	rainpole.local	2	Organization Member	Code Stream Executor
<input type="checkbox"/>	ug-vra-code-stream-viewers@rainpole.local	rainpole.local	1	Organization Member	Code Stream Viewer
<input type="checkbox"/>	ug-vra-service-broker-admins@rainpole.local	rainpole.local	3	Organization Member	Service Broker Administrator
<input type="checkbox"/>	ug-vra-cloud-assembly-users@rainpole.local	rainpole.local	2	Organization Member	Cloud Assembly User
<input checked="" type="checkbox"/>	ug-vra-orchestrator-admins@rainpole.local	rainpole.local	2	Organization Member	Orchestrator Workflow Designer Orchestrator Administrator
<input type="checkbox"/>	ug-vra-cloud-assembly-admins@rainpole.local	rainpole.local	2	Organization Member	Cloud Assembly Administrator

1 - 10 of 11 groups | < 1 / 2 >

Encryption

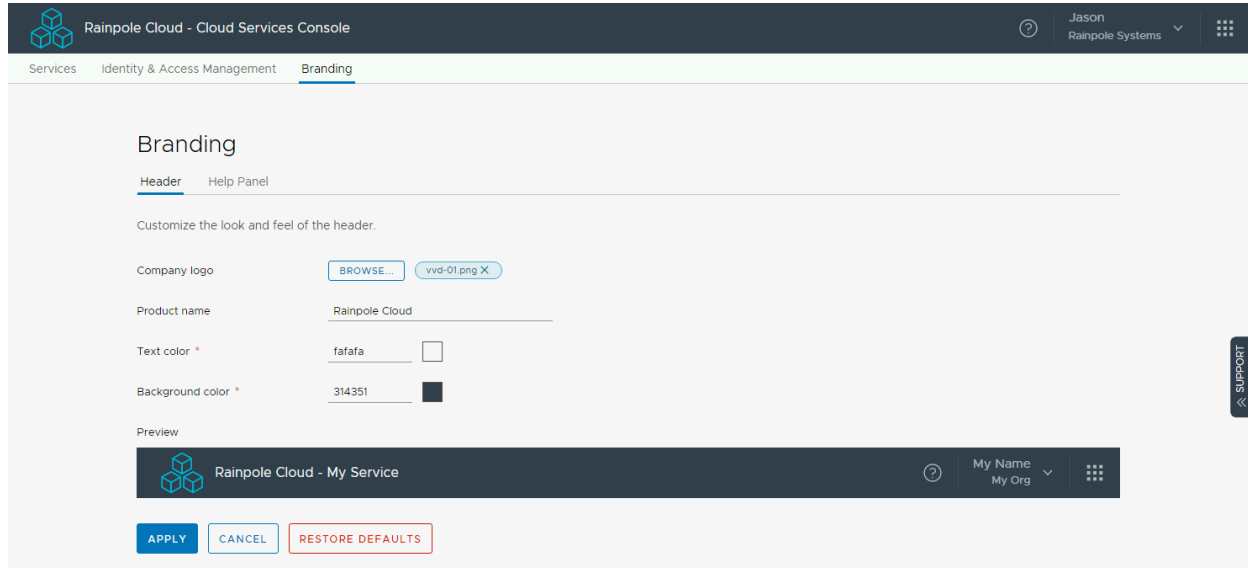
The vRealize Automation user interface and API endpoint use an HTTPS connection. By default, vRealize Automation uses a self-signed certificate. To provide secure access to the vRealize Automation user interface and API, replace the default self-signed certificates with a CA-signed certificate.

Table 2-29. Design Decisions for vRealize Automation Encryption

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-039	Replace the default self-signed certificates with a CA-signed certificate during deployment.	Ensures that all communications to the externally facing vRealize Automation browser-based UI, API, and between the components are encrypted.	<ul style="list-style-type: none"> Replacing the default certificates with trusted CA-signed certificates from a certificate authority increases the deployment preparation time as certificates requests are generated and delivered. You must manage the life cycle of the certificate replacement.

Branding Design for vRealize Automation

You can change the appearance of the vRealize Automation browser-based user interface to meet minimal branding guidelines of an organization by changing the logo, the background and text color, or information in the header and footer.

Figure 2-8. Example of Branding**Table 2-30. Design Decisions for vRealize Automation Encryption**

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-040	Apply branding customizations for the vRealize Automation user interface.	<ul style="list-style-type: none"> Provides minimal corporate branding to the user interface consumed by end users. Allows you to provide an additional help text for end users, which can provide useful information for new users. 	You must provide a logo or icon as a transparent image with the correct size and resolution.
SDDC-CAS-041	Set the organization name in the organization settings.	Allows you to customize the organization name to a user-friendly name instead of the host name from the Workspace ONE Access cluster.	None

Cloud Assembly Design for vRealize Automation

VMware Cloud Assembly orchestrates and expedites infrastructure and application delivery in line with DevOps principles.

Service Roles Design for Cloud Assembly in vRealize Automation

You manage access to Cloud Assembly by assigning enterprise groups to service roles in your organization.

Cloud Assembly has two service roles assigned from the organization identity and access management. You assign the service roles to designated enterprise groups, synchronized from your corporate identity source through Workspace ONE Access.

Table 2-31. Example Service Roles Assignments for Cloud Assembly in vRealize Automation

Role	Description	Enterprise Group
Cloud Assembly Administrator	<ul style="list-style-type: none"> ■ Read and write access to the entire Cloud Assembly user interface and API. ■ Configure cloud accounts, integrations, cloud zones, and Kubernetes zones. ■ Create and manage projects, including project membership. 	rainpole.local\ug-vra-cloud-assembly-admins
Cloud Assembly User	<ul style="list-style-type: none"> ■ Access to the Cloud Assembly user interface and API. ■ The access type is based on the project membership - project administrator or project member. 	rainpole.local\ug-vra-cloud-assembly-users

For information about the service role design decisions for the vRealize Automation Cloud Assembly service, see [Information Security and Access Control Design for vRealize Automation](#).

Tagging Design for Cloud Assembly in vRealize Automation

Tags express capabilities and constraints that determine how and where resources are allocated to workloads during the provisioning process.

Overview

Tagging serves as the foundation for workload placement in Cloud Assembly. Tags are labels that you apply to Cloud Assembly constructs that enable policy-driven placement by directing how and where Cloud Assembly uses resources and infrastructure to build deployable services across private and public clouds. Structurally, tags must follow the *key:value* pair convention, for example, `region:sfo`, but their construction is largely open. Tags also enable the search and identification of compute, storage, and network resources, as well as provisioned machines, using logical and natural language context.

Before you create and use tags in Cloud Assembly, you must establish a well-defined and adaptive tagging strategy and taxonomy. Such a strategy ensures that users who create and use tags understand what they mean, how they must be used, and where and when they must be applied. For example, which tags must be discovered, for example, vSphere tags, and which must be user-defined and managed through Cloud Assembly.

Some practices for an effective tagging strategy:

Plan and Communicate

Create, communicate, and execute a plan for tagging that relates to the structure of your organization. Your plan must support your deployment needs, use a clear natural language, and be understandable to all applicable users.

Simple and Adaptive

Use simple, clear, and meaningful names and values for tags. Users can easily understand capabilities and constraints when using tags in blueprints or reviewing tag assignments for a resource.

In terms of origination, tags can be external and internal:

- External tags are discovered and imported from vSphere, NSX Data Center, and VMware Cloud on AWS, as well as from public clouds like Amazon Web Services and Microsoft Azure.

When imported, these external tags are available for use in the same manner as user-defined tags. External tags are visible in the originating cloud account and from within Cloud Assembly.

- Internal tags are defined in Cloud Assembly and are only visible from within Cloud Assembly.

Tags can be further divided into standard and user-defined:

- Standard tags are applied automatically during provisioning on vSphere, Amazon Web Services, and Microsoft Azure deployments.

Unlike other tags, users cannot use standard tags during deployment configuration, and no constraints are applied. Standard tags are stored as system custom properties and are added to deployments after provisioning.

- User-defined tags defined by a Cloud Assembly user.

In terms of use, tags can be divided into capability and constraints.

Capability Tags

Capability tags are tags that define the capabilities of an object and enable you to define placement logic for deployment. These tags define the required connectivity, functionality, and capabilities for deployments.

You can create capability tags on resources, such as cloud zones, storage and storage profiles, networks and network profiles. Capability tags on storage or network components affect only the components on which they are applied. Cloud Assembly matches capability tags with constraints from cloud zones and on blueprints at deployment time.

Constraint Tags

Constraint tags on blueprints and components match capabilities defined on resources, cloud zones, network and storage profiles to generate deployments with the required configuration.

Constraint tags are applied to two main constructs.

- Project and image configuration
- Blueprint deployment

Constraints applied in both areas are merged in blueprints to form a set of deployment requirements.

When configuring Cloud Assembly, you apply constraint tags on projects which provide governance directly at the project level. All constraints added at this level are applied to all blueprints, requested for the applicable project. If a tag on a project conflicts with a tag on a blueprint, the project tag takes precedence, allowing you to enforce governance rules.

On blueprints, you add constraint tags as YAML code to match the appropriate capability tags that your cloud administrator created on resources, cloud zones, and storage and network profiles. In addition, there are other more complex options for implementing constraint tags. For example, you can use a variable to populate one or more tags on a request. Using a variable enables you to specify one or more of the tags at request time.

Create constraint tags by using the tag label in the blueprint YAML code. Constraint tags from projects are added to the constraint tags created in blueprints.

In the following example, the blueprint constraint attempts to deploy on objects with the `cloud:private` capability tag applied.

```
constraints:
  - tag:
      cloud:private
```

In this example, you pass a blueprint expression for a user selection with a blueprint input.

```
inputs:
  targetCloud:
    type:string
    enum:
      - private
      - vmc
      - aws
      - azure
      - gcp
      - .....
  constraints:
    - tag: '${"cloud:" + to_lower(input.targetCloud)}'
```

If the input of `private` is selected, the constraint tag can be `cloud:private`.

Constraints are typically defined in a blueprint in the format - `[!] tag_key[: tag_value] [:hard|:soft]`.

Consider the following formats when you configure constraints:

Table 2-32. Constraint Formats

Constraint Formats	Description
<code>key:value</code> or <code>key:value:hard</code>	Use this tag when a blueprint must be provisioned on resources with the matching capability tag. The deployment process fails when no matching tag is found.
<code>key:value:soft</code>	Use this tag when you prefer a matching resource. The deployment process proceeds without failing and accepts resources where there is no matching tag.
<code>!key:value</code>	Use this tag, with the hard or soft value, when you want the deployment process to avoid resources with a matching tag.

Table 2-33. Comparison of Capability and Constraint Tags in Cloud Assembly

Object Type	Object	Capability	Constraint
General	Cloud Accounts	✓	x
	Integrations	✓	x
	Cloud Zones	✓	x
	Projects	x	✓
Mappings	Image Mappings	x	x
	Flavor Mappings	x	✓
Profiles	Storage Profiles	✓	x
	Network Profiles	✓	x
Compute	Clusters	✓	x
	Resource Pools	✓	x
	Availability Zones	✓	x
Storage	Storage Policies	x	x
	Datastore / Clusters	x	x
Network	Network Profiles	✓	x
	IP Ranges	x	x
	Load Balancers	✓	x
	Network Domains	x	x
Machines	Machines	x	x
Volumes	Volumes	x	x
Kubernetes	Kubernetes	x	x
Security	Security Groups	✓	x
Blueprints	Blueprint	x	✓

Sequencing and Simulation

The following list summarizes the high-level operations and sequence of capability and constraint tag processing:

- 1 Cloud zones are filtered by several criteria, including availability and profiles. Tags in profiles for the zone are matched.
- 2 Cloud zone and compute capability tags are used to filter the remaining cloud zones by hard constraints.
- 3 Provisioning priority is used to select a cloud zone from the remaining filtered cloud zones. If there are several cloud zones with the same provisioning priority, they are sorted by matching soft constraints, using a combination of the cloud zone and compute capabilities.
- 4 After a cloud zone is selected, a host is selected by matching a series of filters, including hard and soft constraints as expressed in blueprints.

You can simulate a provisioning request to validate your configurations. Based on the provided values, the request goes through the projects, cloud zones, and profiles configurations without executing the provisioning.

Note The design decisions regarding the use of tags on cloud zones, projects, integrations, and profiles are provided within the specific architecture topics.

Table 2-34. Design Decisions on Tagging for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-001	Establish and publish a well-defined strategy and taxonomy for the tagging of cloud resources.	Capability and constraint tags enable you to organize and activate cloud resources and profiles for resource consumption by using the declarative nature of blueprints to define deployment configuration.	Your strategy must account for external tags, for example, vSphere and NSX Data Center tags, and internal (user-defined) tags, managed through Cloud Assembly.
SDDC-CAS-CA-002	Apply constraint tags to blueprints in the YAML code.	During a provisioning operation, capabilities are matched with constraints, each expressed as tags, in blueprints and images to determine the deployment configuration.	You must manage the capability tags on your cloud resources, such as cloud zones, storage and storage profiles, networks and network profiles.

Cloud Accounts Design for Cloud Assembly in vRealize Automation

Cloud accounts in Cloud Assembly provide a centralized authentication mechanism to cloud resources. You configure the necessary permissions to collect data and deploy blueprints to SDDC regions or data centers.

You create cloud accounts for the projects in which your organization team members work. Resource information, such as network and security, compute, storage, and tags data, is collected from your cloud accounts.

For organizations that are distributed across multiple geographic regions, vRealize Automation connects to the cloud accounts directly by using HTTPS.

vRealize Automation connects to on-premises cloud accounts and integrations, such as a workload domain vCenter Server, NSX-T Data Center Manager, and vRealize Orchestrator. You provide the details, such as the endpoint IP/FQDN/URL, user name, password, name, description, and capability tags.

vRealize Automation uses a custom role in vSphere with permissions for a designated service account to perform vRealize Automation operations on cloud accounts in the Software-Defined Data Center. Dedicated service accounts are assigned a custom role for communication between vRealize Automation and the vCenter Server instances in the environment.

You can also configure Cloud Assembly cloud accounts to public cloud services. This configuration allows Cloud Assembly to collect infrastructure data from the public cloud and enables you to deploy blueprints to one or more of the account regions in the public-cloud-backed cloud account.

Table 2-35. Design Decisions for Cloud Accounts for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-003	Define a custom vCenter Server role, for example, vRealize Automation to vSphere Integration , for vRealize Automation that has the minimum privileges required to support a vCenter Server-based cloud account.	vRealize Automation interacts with vSphere by using the minimum set of permissions that are required to support the cloud account.	You must maintain the permissions required by the custom role.
SDDC-CAS-CA-004	Configure a service account, for example, svc-vra-vsphere@rainpole.local , in vCenter Server for application-to-application communication from vRealize Automation to vSphere.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ vRealize Automation services, such as Cloud Assembly, accesses vSphere with the minimum set of required permissions. ■ If there is a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the vRealize Automation and the SDDC cloud account. 	You must maintain the life cycle and availability of the service account outside of the SDDC stack.
SDDC-CAS-CA-005	Assign global permissions for the vRealize Automation-to-vSphere service account, for example, svc-vra-vsphere@rainpole.local .	<p>vRealize Automation access designated workload domains with the minimum set of permissions that are required to support vCenter Server-backed cloud accounts in the design.</p> <p>See the vRealize Automation documentation for the required minimum permissions.</p>	<p>All vCenter Server instances must be in the same vSphere domain.</p> <p>You must set the role for the management domain vCenter Server instances to No Access to ensure that the account cannot communicate with the management domain.</p>

Table 2-35. Design Decisions for Cloud Accounts for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-006	Configure a service account in vCenter Server for application-to-application communication, for example, svc-vra-nsx-v , from vRealize Automation to NSX for vSphere.	Provides the following access control features: <ul style="list-style-type: none"> ■ vRealize Automation accesses NSX with the minimum set of permissions, required to perform life cycle management of virtual networking objects. ■ If an account becomes compromised, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between VMware Cloud Services and the SDDC cloud account. 	You must maintain the life cycle of the service account outside of the SDDC stack to ensure its availability.
SDDC-CAS-CA-007	Use global permissions when you create the vRealize Automation-to-NSX for vSphere service account, for example, vc-vra-nsx-v , in vCenter Server.	<ul style="list-style-type: none"> ■ Simplifies and standardizes the deployment of the service account across all vCenter Server instances in the same vSphere domain. ■ Provides a consistent authorization layer. 	All vCenter Server instances must be in the same vSphere domain.
SDDC-CAS-CA-008	Use the admin account with NSX-T Data Center Enterprise Administrator role for application-to-application communication from vRealize Automation to NSX-T Data Center.	Although NSX-T Data Center supports the use of Workspace ONE Access as an authentication source and access control, it is not supported in vRealize Automation at the time of this design writing. The default admin account is used.	You must control access to the default admin account in NSX-T Data Center for use with vRealize Automation cloud accounts.
SDDC-CAS-CA-009	Add a cloud account for the vCenter Server instance for each workload domain in each Software-Defined Data Center region.	You can integrate on-premises vSphere-backed workload domains with vRealize Automation for workload provisioning.	<ul style="list-style-type: none"> ■ You must manage the cloud account credentials with the life cycle management of the service account. ■ You must manage capability tags for the cloud account.

Table 2-35. Design Decisions for Cloud Accounts for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-010	Add a cloud account for the NSX-T Data Center instance for each workload domain in each Software-Defined Data Center region.	You can integrate on-premises NSX-T Data Center workload domains with vRealize Automation for workload provisioning.	<ul style="list-style-type: none"> ■ You must manage the cloud account credentials with the life cycle management of the service account. ■ You must manage capability tags for the cloud account. ■ Workload domains can share an NSX-T Data Center instance (e.g. many:1:many). You must map the vCenter Server cloud accounts that share an NSX-T Data Center instance to the shared NSX-T Data Center cloud account.
SDDC-CAS-CA-011	Add a cloud account for VMware Cloud on AWS Software-Defined Data Center instances.	You can integrate the vSphere and NSX Data Center resources provided through VMware Cloud on AWS with vRealize Automation for workload provisioning.	<ul style="list-style-type: none"> ■ The VMware Cloud on AWS token used during the cloud account registration must have the expiration set. It must be updated and re-authenticated before the token expiration. ■ You must manage capability tags for the cloud account.

Cloud Zones Design for Cloud Assembly in vRealize Automation

Cloud zones in Cloud Assembly are a method of partitioning the resources that cloud accounts provide to projects.

You create cloud zones to define a set of resources in a cloud account. Each cloud zone is associated with a Cloud Assembly project. Cloud zones can be shared across multiple groups of users and are not limited by a one-to-one relationship.

You can add or remove tags to filter the compute resources that are used in the cloud zone. For example, you can have multiple clusters within a workload domain provided through a vCenter Server-backed cloud account. The clusters can serve different purposes due to specific functional or non-functional constraints or classifications.

Cloud zones do not control how CPU and memory are allocated during periods of resource contention. This function is performed by the use of vSphere resource pools on vSphere-backed cloud accounts. To ensure resource availability, all virtual machines must be deployed to a resource pool on the first cluster (shared edge and workload cluster) for the workload domain.

Table 2-36. Resource Pools in the Default Cluster of a Workload Domain

Resource Pool	Object Types
sfo01-w01rp-sddc-edge	NSX-T Data Center Edge components. Do not place user workloads in this resource pool.
sfo01-w01rp-user-vm	Statically or dynamically deployed virtual machines that contain organization workloads.

For additional clusters in a workload domain to be used by vRealize Automation, apply tags to ensure that the resources are available to the cloud zone.

By default, on vCenter Server-backed cloud zones, workloads are placed on random hosts. Optionally, one of the following strategies can be applied to a cloud zone:

Table 2-37. Placement Policies

Placement Policy	Description
Default	Uses random hosts for workload placement.
Binpack	The most loaded host with enough resources to run the workload is selected for workload placement.
Spread	Evenly spreads workloads across hosts during workload placement.
Advanced	Uses vRealize Operations recommendations, if integrated, for workload placement.

By default, workloads are placed in the default data center folder. Workload placement can be set by defining a relative path within the data center in the blueprint YAML code.

For example, blueprint expressions can evaluate configuration options to return a value. Due to restrictive permissions, in VMware Cloud on AWS, workloads must be deployed in the **Workloads** folder (or subfolder), but for an on-premises cloud account, you might want to deploy the workload to an alternate folder. You can accomplish this with a blueprint expression that evaluates end-state values. If a user selects the `targetCloud` value of `vmc`, the `folderName` property value is set to `Workload`. The `else` option can apply an alternative value to `folderName` - a blank value (""), or another input value `{input.myFolderName}`. See the following sample YAML blueprint:

```
resources:Cloud_vSphere_Machine_1:type:Cloud.vSphere.Machine properties:image:'$
{input.operatingSystem}'flavor:'${input.nodeSize}'count:'${input.nodeCount}'FolderName:'${input.targetCloud ==
"vmc" ? "Workload" : ${input.targetEnvironment}}' networks:- network:'$
{resource.Cloud_NSX_Network_1.id}'constraints:- tag:'${"cloud:" + to_lower(input.targetCloud)}'
```

You can add capability tags to match blueprint constraints to a cloud zone. Capability tags are optional because the tags on compute resources, such as vSphere clusters and resource pools, are also used as capability tags for a cloud zone. For example: `cloud:private`, `region:sfo`, `region:lax`, `region:us-west-1`, `region:us-west-2`, `env:prod`, `env:dev`, `env:dmz`.

Table 2-38. Design Decisions on Cloud Zones for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-012	Create a cloud zone for each Software-Defined Data Center region.	Enables region-specific workload provisioning.	None
SDDC-CAS-CA-013	Add tags to compute resources instead of the cloud zone.	Ensures that the scope of the workload tagging is managed with the compute resources.	<ul style="list-style-type: none"> ■ You must establish and manage an effective tagging strategy. ■ You must ensure that constraint tags are included in the blueprint YAML file.
SDDC-CAS-CA-014	Only add tags to the *-user-vm resource pool in the default shared edge and compute cluster for cloud zones that contain NSX-T Data Center-backed cloud accounts.	Ensures that virtual machines, using NSX-T Data Center, are deployed to the designated resource pool in the default cluster of the workload domain.	You must ensure that constraint tags are included in the blueprint YAML.
SDDC-CAS-CA-015	Add tags to the vSphere cluster for additional clusters added to the workload domain.	Ensures that, as you add new vSphere clusters to a workload domain, you can enable workload provisioning readiness by adding the appropriate organization tags.	<ul style="list-style-type: none"> ■ You must ensure that you manage the tagging of compute resources as you add clusters to a workload domain. ■ You must ensure that constraint tags are included in the blueprint YAML.

Table 2-38. Design Decisions on Cloud Zones for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-016	Add a Workloads folder in the vCenter Server data center for each workload domain, for example, VMs and Templates > sfo01m01vc01.sfo01.rainpole.local > sfo01-w01dc > Workloads > .	<ul style="list-style-type: none"> ■ Ensures that blueprints, which do not include the <code>folderName</code> value, are deployed in a default Workloads folder. ■ Ensures that workloads deployed to VMware Cloud on AWS are provisioned to the default folder due to the restrictive security mode. 	<ul style="list-style-type: none"> ■ You must add the Workloads default folder to all workload domains. ■ To override the default folder, you must add a logic in your blueprint YAML code to deploy blueprints in alternative folders. <p>Note The destination folder, where the blueprints are deployed, must exist and cannot be created by Cloud Assembly without extensibility.</p>
SDDC-CAS-CA-017	Use the Default placement policy for each cloud zone.	<ul style="list-style-type: none"> ■ All vSphere clusters have the vSphere Distributed Resource Scheduler enabled to optimize initial and ongoing workload placement within a cluster. ■ The Advanced option for vRealize Operations does not support workload placement on resource pools in vCenter Server or vSAN datastores for workload placement. 	You must ensure that the vSphere Distributed Resource Scheduler is enabled for all workload domain clusters.

Integrations Design for Cloud Assembly in vRealize Automation

Integrations in Cloud Assembly enable you to support other systems and services, allowing for extended automation and life cycle management capabilities.

Integrations create a connection between Cloud Assembly and the target service. You configure integrations with a direct connection to a target service.

Table 2-39. Cloud Assembly Integrations

Integration	Description
vRealize Orchestrator	<p>You can use vRealize Orchestrator workflows in extensibility subscriptions.</p> <p>With a vRealize Orchestrator integration to Cloud Assembly, you can use vRealize Orchestrator workflows in Extensibility subscriptions. Cloud Assembly supports the integration with the default embedded or external vRealize Orchestrator instances.</p> <p>You can use Cloud Assembly capability tags to manage the placement logic of your vRealize Orchestrator integrations. For information about capability tags, see Tagging Design for Cloud Assembly in vRealize Automation.</p>
vRealize Operations	You can use vRealize Operations to direct vSphere-based workload placement and display metrics post-placement.
My VMware	You can use VMware Marketplace to consume existing blueprints.
Ansible	You can manage deployments for configuration and drift.
Puppet	You can manage deployments for configuration and drift.
GitHub	You can use GitHub (SaaS) repositories to manage blueprints and action scripts under source control, allowing you to use modern development tools, for example, Visual Studio Code, for blueprints and actions authoring.
GitLab	You can use GitLab (SaaS or self-managed) repositories to manage blueprints and action scripts under source control, allowing you to use modern development tools, for example, Visual Studio Code, for blueprints and actions authoring.
VMware Enterprise PKS	You can use the Kubernetes integration with Enterprise PKS to create Kubernetes clusters using blueprints.
IPAM	You can use an IPAM provider for IP Address Management.
Active Directory	You can use Active Directory to manage workload computer accounts.

Table 2-40. Design Decisions on Integrations for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-018	Add an integration in Cloud Assembly with My VMware.	Provides the ability to download curated blueprints and images from VMware Marketplace to Cloud Assembly.	You must consider using a dedicated My VMware account for the integration with Cloud Assembly, for example, svc-vra-marketplace@rainpole.local , with the minimum My VMware permissions to access the VMware Marketplace.
SDDC-CAS-CA-019	Use the default integration to the vRealize Orchestrator instance that is embedded in vRealize Automation.	<p>The use of embedded vRealize Orchestrator instances has the following advantages:</p> <ul style="list-style-type: none"> ■ Operates in cluster mode. ■ Provides a faster time to value. ■ Reduces the number of appliances to manage. ■ Provides an easier upgrade path and better supportability. ■ Improves performance. ■ Removes the requirement for an external database. <p>Note Using the embedded instance of vRealize Orchestrator is applicable in most use cases. However, refer to the product documentation to be aware of the cases where using the external vRealize Orchestrator is applicable.</p>	Less efficient to execute a workflow in a multi-region deployment from a centralized instance. However, the embedded vRealize Orchestrator instance is easier to manage.

Table 2-40. Design Decisions on Integrations for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-020	Configure a service account, for example, svc-vra-vrops@rainpole.local , in vRealize Operations for application-to-application communication from vRealize Automation-to-vRealize Operations Manager.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ vRealize Automation services access vRealize Operations with the minimum set of required permissions. ■ If there is a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the vRealize Automation and the SDDC cloud account. 	<p>You must maintain the life cycle and availability of the service account outside of the SDDC stack.</p> <p>You must maintain the synchronization and availability of the service account in Workspace ONE Access and vRealize Operations Manager.</p> <p>You must use the format of <i>user@domain@source</i> when configuring the integration to use a service account backed by Workspace ONE Access, for example, svc-vra-vrops@rainpole.local@Workspace ONE.</p>
SDDC-CAS-CA-021	Add an integration in Cloud Assembly with vRealize Operations Manager.	<ul style="list-style-type: none"> ■ You can use data from vRealize Operations Manager in vRealize Automation to display live vSphere-based virtual machine metrics for CPU, memory, storage IOPS, and network MBps after placement. Metrics for the past day, week, or month are available. 	<ul style="list-style-type: none"> ■ You must use a service account, created for application-to-application integration for vRealize Automation to vRealize Operations when configuring the integration. ■ You must manage the service account password in the integration configuration when the service account password is updated. ■ You must configure vRealize Operations Manager with vRealize Automation to ensure that both applications are set to the same time zone. vRealize Automation uses only UTC.

Table 2-40. Design Decisions on Integrations for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
		<ul style="list-style-type: none"> You can use data from vRealize Operations Manager in vRealize Automation to display a cost estimation at the time of deployment and over time. 	
		<p>Important</p> <ul style="list-style-type: none"> The vRealize Operations Manager integration is not used for workload placement in this design. The integration does not support resource pools in vCenter Server or vSAN datastores for workload placement. Project costs include only the costs for private cloud workloads. If a project contains deployments that belong to public clouds, the costs for these deployments are not included in the project cost. 	

Projects Design for Cloud Assembly in vRealize Automation

In Cloud Assembly, projects control who has access to blueprints and where blueprints can be deployed.

A project in Cloud Assembly determines the cloud zones to which a set of users or groups can deploy, a priority value, the maximum number of virtual machine instances to deploy, and a maximum amount of memory that the deployment can use. A project is typically defined by using an organization structure, such as a cost-center, or a specific business group or purpose.

By using projects, you can organize and govern what business users can do and on which cloud zones they can deploy blueprints in your cloud infrastructure. You create a project for which you manage membership - project administrators and project members - and cloud zones, so that the project members can deploy their blueprints to the associated cloud zones. Project administrators use the infrastructure that the cloud administrators created to ensure that the project members have the resources they need. Project members are users who create, iterate, and deploy blueprints.

For example, as a Cloud Assembly administrator, you can first create a project for a development team and add the users and groups as project members or project administrators. Then you can add only the cloud zones that are designated for these workloads.

Projects provide a context that you can assign blueprints to, and moving beyond Cloud Assembly they also bind integration endpoints, such as Git-based repositories in Code Stream.

When you assign a cloud zone within a project, you can include additional configurations that determine the project behavior.

Provisioning Priority	The priority order for a cloud zone use in a project when all other constraints have been met.
Instances Limit	The maximum number of instances that can be provisioned in the cloud zone for the project.

Resource tags can be created on cloud resources during provisioning. For vSphere-backed cloud accounts, the tags are associated with the workloads inside the workload domain.

Project constraints can be used as a governance definition to restrict or enable resource consumption. These include network constraints matched with capability tags on network profiles and subnets, storage constraints matched with capability tags on storage profiles, and extensibility constraints matched with capabilities tags on vRealize Orchestrator integrations. Project constraints define what resources the deployment request consumes or avoids in the project cloud zones.

During the deployment process, the blueprint requirements of a project are evaluated against the cloud zones associated with the project to determine the best deployment location.

You can also add a custom project property to trigger and populate extensibility actions or workflows, override blueprint level properties, or for reporting.

Figure 2-9. Tagging Example

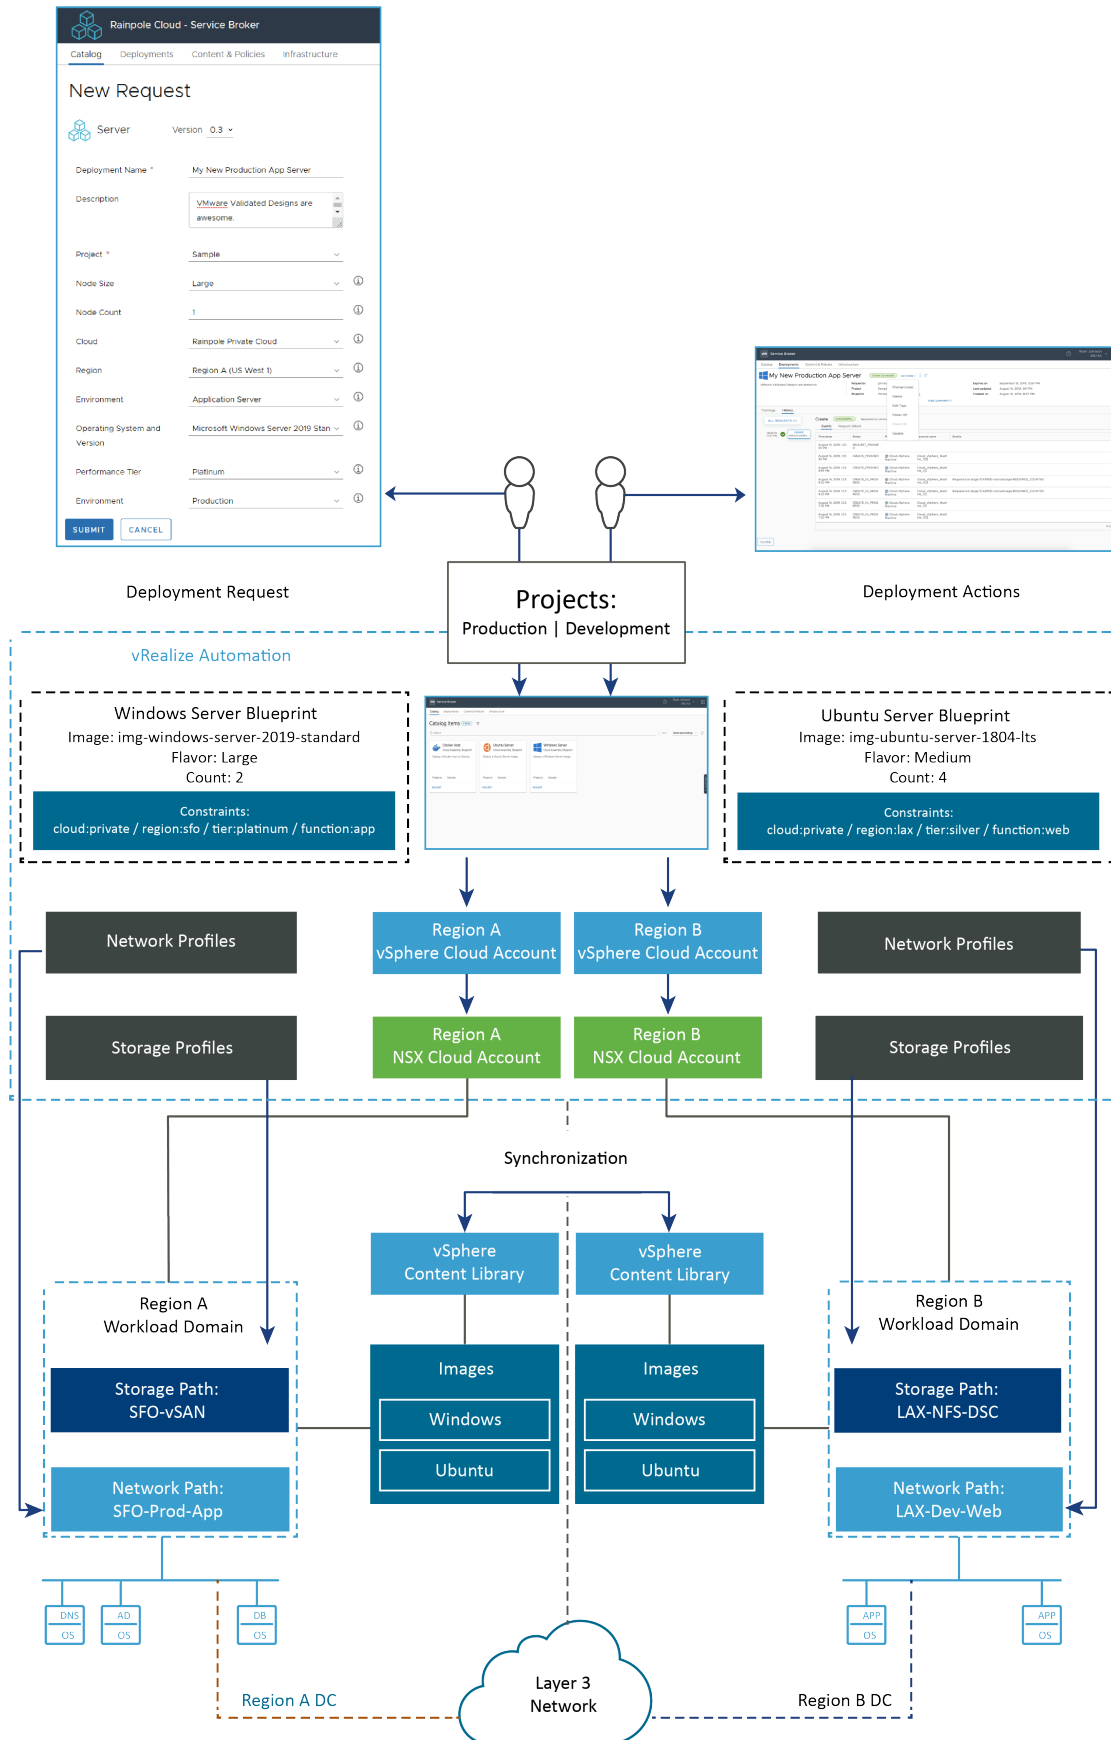


Table 2-41. Design Decisions on Projects for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-022	For each project, add one or more enterprise groups to the Project Administrator and Project Member roles rather than assigning individual users.	An enterprise group is a collection of directory users, such as an Active Directory security group. You can use enterprise groups to specify the role permissions for a collection of users, which makes it easier to manage project membership.	You must maintain the membership and availability of the security groups outside of the SDDC stack.
SDDC-CAS-CA-023	For each project, add one or more Cloud Zones based on the project requirements and allowed cloud resources.	Allows you to provide one or more cloud zones and their resources for project consumption.	None
SDDC-CAS-CA-024	For each project, set a provisioning priority for each cloud zone based on your deployment prioritization.	Enables you to prioritize one cloud zone over another within a project. Note The default priority is 0 (highest priority).	You must manage the provisioning priority for each cloud zones in each project.
SDDC-CAS-CA-025	For each project, set an instance limit for the project cloud zones.	Enables you to set the maximum number of workloads that can be provisioned in this cloud zone for the project. Note The default instance limit is 0 (unlimited instances).	If a value greater than 0 (unlimited) is used for an instance limit, you must manage the limit for each cloud zone in each project when requirements change.
SDDC-CAS-CA-026	For each project, specify a network, storage, and extensibility constraints that must be applied to all requests in the project.	Ensures proper placement of the workloads in a project and its cloud zones.	If the same constraint or the same constraint category is specified in both the project, for example, <code>region:sfo</code> , and the blueprint, for example, <code>region:lax</code> , the constraint specified in the project takes precedence.

Table 2-41. Design Decisions on Projects for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-027	For each project, add one or more custom properties, for example, <code>project:foo</code> .	Custom properties can be used for provisioning or to capture additional metadata. For example, for reporting or extensibility actions.	If the same custom property is specified in both the project, for example, <code>project:foo</code> , and the blueprint, for example, <code>project:bar</code> , the property value specified in the project takes precedence, for example, <code>project:foo</code> .
SDDC-CAS-CA-028	For each project, add a custom naming template to be used for virtual machine names provisioned in the project.	The template provides a custom virtual machine name and does not affect the host name of the virtual machine.	The template substitutes auto-generated virtual machine names by using available properties, such as resource properties, custom properties, endpoint properties, project properties, and/or a random number with a specified number of digits. Important You must ensure that the template generates unique names for this project and between other projects.

Important Project administrators must be granted the **Service Broker Administrator** role to perform customizations to blueprint icons and forms. However, members of this role are also entitled to manage cloud accounts, cloud zones, and integrations created by a **Cloud Assembly Administrator**.

Mappings Design for Cloud Assembly in vRealize Automation

Cloud Assembly mappings use a natural language terminology to define compute resource sizes and virtual machine images for a specific cloud account/region.

Flavor Mappings

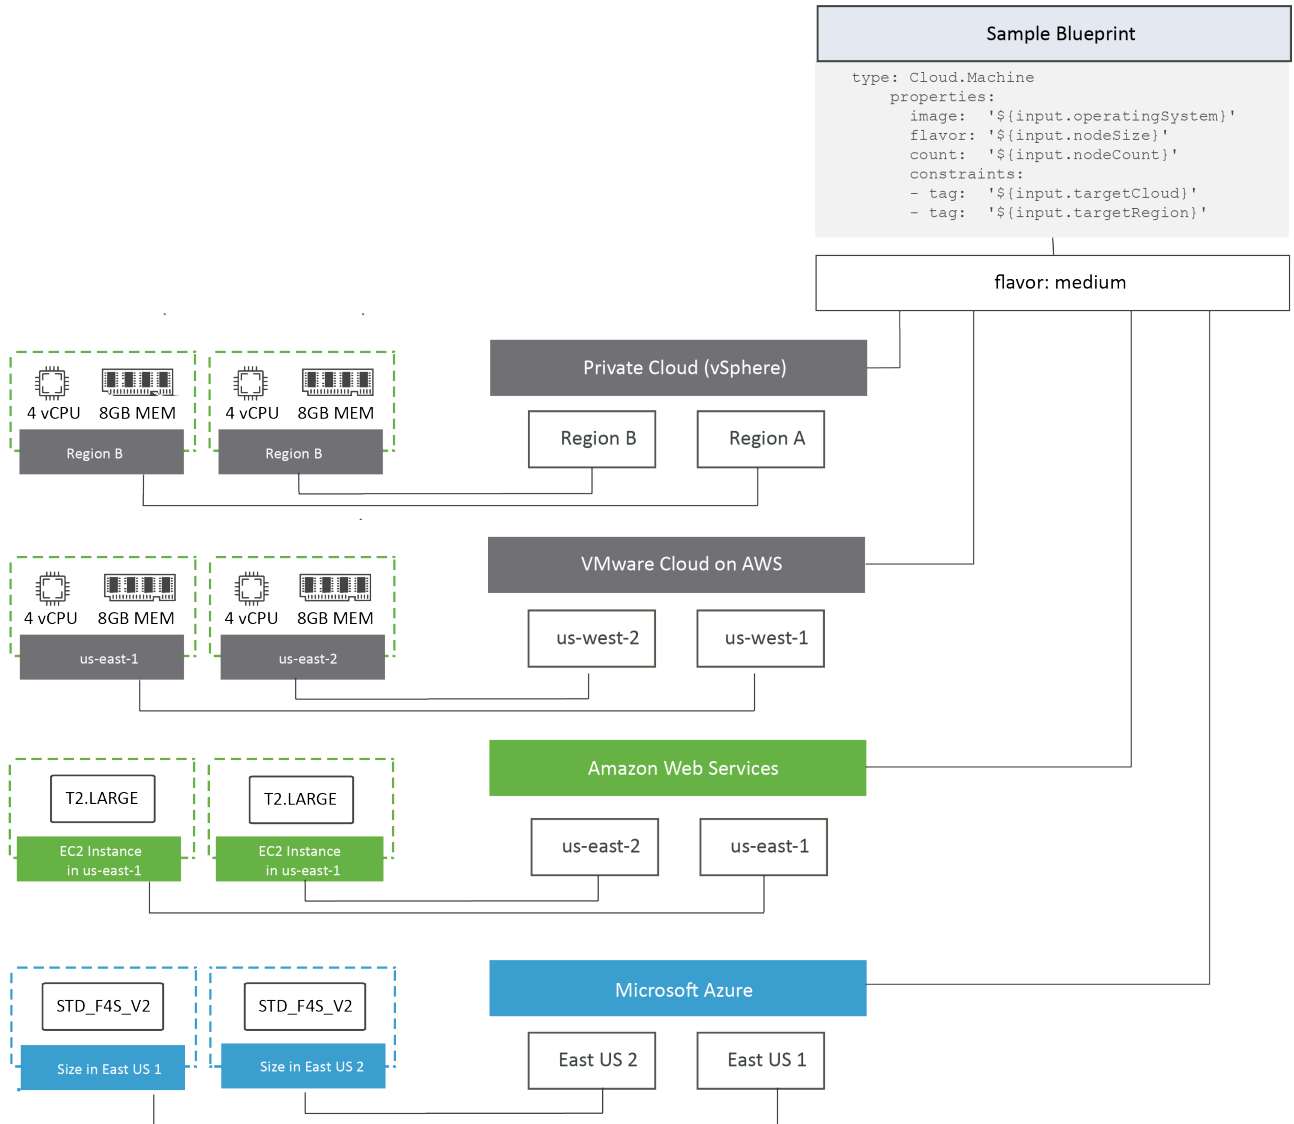
Cloud Assembly flavor mappings allow you to use a natural language naming to define a group of deployment size specifications in a specified cloud account/region.

A flavor mapping associates a defined cloud region with one or more resource sizing options or instance types. A mapping defines a common term that is mapped to the specific constructs in each of your platform environments. A flavor can equate to a specific number of CPUs and memory allocation, allowing you to scale your resources to the requirements of your target workload.

This is commonly expressed as t-shirt sizes. For example, `x-small` can represent 1 CPU and 512 MB memory, while `medium` can represent 4 CPUs and 16 GB memory for a vCenter Server-backed cloud account in a named data center. The same resources can be mapped to `t2.nano` and `t2.large` for an Amazon Web Services account in a named region.

You can create a flavor-mapping schema across your cloud accounts and regions. For example, a flavor map named x-large can contain a similar flavor sizing for some or all available accounts/regions in your project. When you build a blueprint, you pick an available flavor that fits your needs.

Figure 2-10. Flavor Mapping and Blueprint Consumption Across Clouds



To simplify the blueprint creation, you can select a preconfigured option when you add a new cloud account. Using the preconfiguration option, selects the most popular flavor mapping and image mapping for the specified region in your organization.

Table 2-42. Design Decisions on Flavor Mappings for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-029	Create standardized flavor mappings based on a common taxonomy and deployment intent. For example, you can use "t-shirt sizes".	Provides a simple, natural language naming to define common deployment size specifications.	You must publish and communicate the updates to blueprint developers.
SDDC-CAS-CA-030	For each flavor mapping, add all applicable cloud zones.	It provides a simple, natural language naming to define common deployment size specifications when used in a specific cloud account/region.	You must maintain the mapping for any image mapping create or update operation.

Image Mappings

Cloud Assembly image mappings allow you to use a natural language naming to define a group of virtual machine operating system image specifications when used in a cloud account or region.

Image mappings provide a method to define the virtual machine image that an environment can consume. An image mapping associates an image name with a virtual machine template in a region. You can create one or more image names and map to a metadata file that contain pre-defined value sets.

For example, if you want to deploy an Ubuntu virtual machine, you can map `ubuntu-server-1804-lts` to a specific virtual machine template or an image from the vSphere Content Library in a vSphere-backed cloud region or account, or to an AMI file in Amazon Web Services, or a specific VHD in Microsoft Azure.

A mapping links a common term to the specific construct in each of your platform environments. In addition, an image can map to a virtual machine image that contains pre-populated cost or region specifications to import into a blueprint.

Cloud accounts backed by vCenter Server and NSX-based environments, such as on-premises and VMware Cloud on AWS Software-Defined Data Center environments, use image mappings to group a set of target deployment conditions together, including virtual machine configuration settings and guest operating system. Cloud accounts, such as Amazon Web Services and Microsoft Azure, also use a similar grouping mechanism to define a set of deployment conditions.

Cloud Assembly can consume images for vCenter Server and NSX-based environments by using traditional virtual machine templates, or Open Virtual Format (OVF) images, packaged in the Open Virtual Appliance (OVA) format in the vSphere Content Library or a source URL, such as a GIT-based repository. In this design, the vSphere Content Library is used to manage and distribute virtual machine images across workload domains in the Software-Defined Data Center instances.

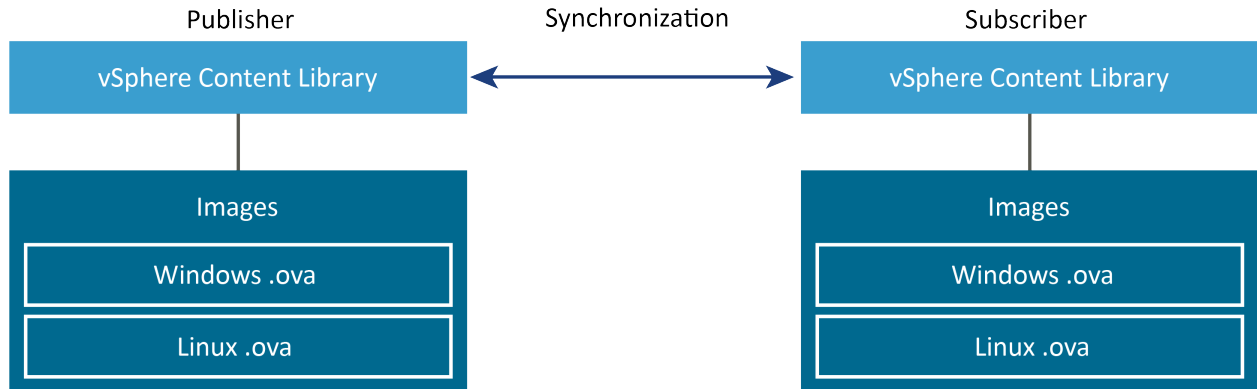
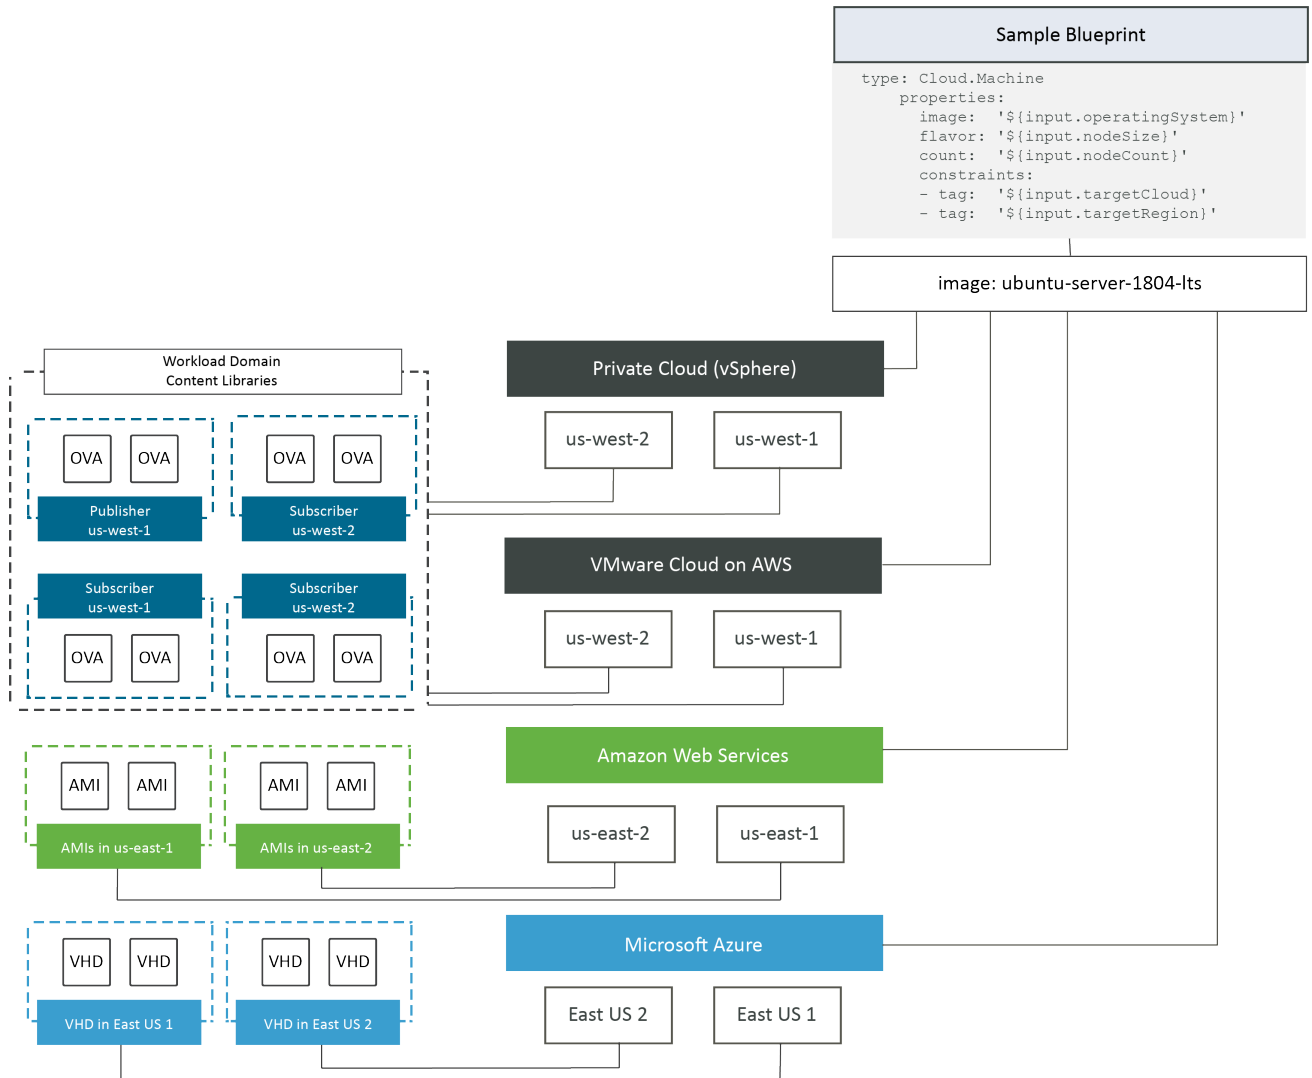
Figure 2-11. Image Distribution Using the vSphere Content Library

Image mappings can consume both private images and public cloud images, such as Amazon Machine Images (AMIs) on Amazon Web Services, or Azure Images (VHD/VHDX) on Microsoft Azure. Cloud Assembly automatically performs private image collection every 24 hours or you can manually trigger a synchronization. Synchronization is disabled during an image enumeration and for 10 minutes after the last image enumeration has completed.

Figure 2-12. Image Mapping and Blueprint Consumption Across Clouds

When you build, deploy, and iterate a blueprint, you pick an available image that best fits your requirements. To simplify the blueprint creation, you can select a preconfiguration option when you add a new cloud account. Using the preconfiguration option, selects the most popular flavor mapping and image mapping for the specified region in your organization.

Table 2-43. Design Decisions on Image Mappings for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-031	Use the Open Virtual Format (OVF) images packaged in the Open Virtual Appliance (OVA) format for vCenter Server-backed Cloud Accounts.	It provides a simple, standards-based, and portable file format for all virtual machine images that can be replicated by the vSphere Content Library and natively consumed by vRealize Automation.	You must use OVFTool or PowerCLI to export virtual machines to the OVA format. If necessary, native vSphere templates can be used instead of OVA-based images. templates can be imported to the vSphere Content Library for distribution.
SDDC-CAS-CA-032	Use the vSphere Content Library to synchronize virtual machine images across workload domains and regions for vSphere-backed cloud accounts. The vSphere Content Library can also be used by Software-Defined Data Center instances in VMware Cloud on AWS as a subscriber to Region A.	<ul style="list-style-type: none"> ■ The vSphere Content Library is built into vSphere and meets all the requirements to synchronize virtual machine images across multiple regions in a Software-Defined Data Center, including VMware Cloud on AWS. ■ vRealize Automation can consume Open Virtual Format-based images from the vSphere Content Library for a cloud account. 	<ul style="list-style-type: none"> ■ The vSphere Content Library permissions are connected to global permissions in the permissions hierarchy. To allow Cloud Assembly to sync and use images in the Content Library, the user and role must be applied at the global permissions level. ■ You must provide storage space in each region for images. ■ You must ensure that the service account used for the Cloud Account for vCenter Server has the minimum required permissions to consume images from the vSphere Content Library. ■ You must ensure that the number, size, and structure of the OVA images are kept within the vSphere Content Library configuration maximums. ■ You must ensure HTTPS communications between each workload domain vCenter Server instance.

Table 2-43. Design Decisions on Image Mappings for vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-033	Create standardized image mappings based on similar operating systems, functional deployment intent, and cloud zone availability.	It allows you to create a simple taxonomy to map images to blueprints.	You must publish and communicate the image-mapping standards and updates to blueprint developers.
SDDC-CAS-CA-034	For each image in an image mapping, add a constraint tag, as applicable.	Refines the image selection in an image mapping by matching constraints. For example, you can specify a provider as a constraint tag <code>provider:rainpole</code> on one image and <code>provider:marketplace</code> on another image within the same image mapping. During blueprint provisioning, you can pass the <code>provider:rainpole</code> constraint tag to ensure it selects a corporate image.	You must manage multiple images in each region based on the use of constraint tags in your organization.

Profiles Design for Cloud Assembly in vRealize Automation

Cloud Assembly profiles allow you to define network and storage profiles for a specific cloud account or region.

Network Profiles

A network profile in Cloud Assembly describes the behavior of the cloud-specific network to be deployed. For example, a network might require specific network services or access.

Network profile capability tags define a group of networks and workload-specific network characteristics that are available for a cloud account in a particular region. Network profile capability tags are matched to blueprint constraint tags during provisioning. Capability tags are applied to all networks in the network profile.

Network profiles can match the *key:value* constraint pair in a blueprint network configuration. If a network profile is enabled for a public IP address, only networks that match `networkType:public` are matched in the deployment process.

Network policies within a network profile allow you to define network settings, such as private or public networks, on-demand networks, and security groups generation. For example, you can create an on-demand network for each deployment by using an NSX transport zone or a public cloud network domain. The existing public and private networks within the network profile are used as the underlying or upstream networks. You can match network type tags with blueprint constraints to place the deployed workload in a specific network.

Table 2-44. Example Network Profiles

Account / Region	Cloud Account Type	Network Profile Name	Network Characteristics	Capability Tags
vmw-rainpole / sfo01dc01	VMware vSphere	net-existing-sfo01dc01	<ul style="list-style-type: none"> Name: subnet-ks81jg0 CIDR: 172.11.10.0/24 	<ul style="list-style-type: none"> cloud:private region:sfo network:web env:prod
			<ul style="list-style-type: none"> Name: subnet-j47fj47 CIDR: 172.11.20.0/24 	<ul style="list-style-type: none"> cloud:private region:sfo network:app env:prod
			<ul style="list-style-type: none"> Name: subnet-19h4ud4 CIDR: 172.11.30.0/24 	<ul style="list-style-type: none"> cloud:private region:sfo network:db env:prod
vmw-rainpole / lax01dc01	VMware vSphere	net-existing-lax01dc01	<ul style="list-style-type: none"> Name: subnet-j58d73kk CIDR: 172.21.10.0/24 	<ul style="list-style-type: none"> cloud:private region:lax network:web env:prod
			<ul style="list-style-type: none"> Name: subnet-89d738j5 CIDR: 172.21.20.0/24 	<ul style="list-style-type: none"> cloud:private region:lax network:app env:prod
			<ul style="list-style-type: none"> Name: subnet-19jd8yt4 CIDR: 172.21.30.0/24 	<ul style="list-style-type: none"> cloud:private region:lax network:db env:prod
vmc-rainpole / SDDC-Datacenter	VMware Cloud on AWS	net-existing-vmc01sddc01	<ul style="list-style-type: none"> Name: vmc01-routed-01 CIDR: 192.168.18.0/24 	<ul style="list-style-type: none"> cloud:vmc region:us-east-1 network:web env:dev
aws-rainpole / us-east-1	Amazon Web Services	aws-us-east-1	<ul style="list-style-type: none"> Name: subnet-5da92b01 Zone: us-east-1a Network Domain: vpc-7685580c Subnet: 172.31.31.0/20 	<ul style="list-style-type: none"> cloud:aws region:us-east network:api env:prod

Table 2-44. Example Network Profiles (continued)

Account / Region	Cloud Account Type	Network Profile Name	Network Characteristics	Capability Tags
			<ul style="list-style-type: none"> ■ Name: subnet-fa90ce9d ■ Zone: us-east-1b ■ Network Domain: vpc-ee462c94 ■ Subnet: 172.31.32.0/20 	<ul style="list-style-type: none"> ■ cloud:aws ■ region:us-east ■ network:db ■ env:prod
azure-rainpole / east us	Microsoft Azure	azure-us-east	<ul style="list-style-type: none"> ■ Name: default ■ Zone: eastus ■ Network Domain: vnet-eff82c97 ■ Subnet: 10.0.0.0/24 	<ul style="list-style-type: none"> ■ cloud:azure ■ region:us-east ■ network:sbx ■ env:sbx

Note The scope of this design includes VMware-enabled clouds. Examples include both private and public cloud network profiles for Cloud Assembly in a multi-cloud context.

Table 2-45. Design Decisions on Network Profiles for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-035	For each account/region, add one or more network profiles based on network characteristics available for consumption.	Allows you to add networks with pre-defined characteristics that can be consumed during a deployment process.	You must manage network profiles for each account/region.
SDDC-CAS-CA-036	For each network in a network profile, add one or more capability tags.	You use capability tags to manage the workload network placement logic during the deployment process.	You must manage capability tagging on each network profile for workflow placement selection during a deployment process.

Storage Profiles

A storage profile in Cloud Assembly describes a cloud-specific storage on which workloads can be deployed based on a set of characteristics.

Storage profiles are organized under cloud-specific regions. A public cloud account, such as Amazon Web Services, can have multiple regions, for example, AWS `us-east-1` and `us-west-1`, with multiple storage profiles for each. However, an on-premises vSphere-backed cloud account can have a single region, for example, `sfo`, with multiple storage profiles.

Storage profiles include cloud-specific configurations and a means to identify the type of storage by capability tags. During the deployment process, tags are matched against the provisioning service request constraints to create the required storage. You can also specify if a profile must be used as the default for an account/region.

Table 2-46. Example Storage Profiles

Account / Region	Cloud Account Type	Storage Profile Name	Storage Characteristics	Capability Tags
vmw-rainpole / sfo01dc01	VMware vSphere	sfo-platinum	<ul style="list-style-type: none"> ■ StoragePolicy: vSAN Default Storage Policy ■ Datastore/ Datastore Cluster: sfo01-w01-vsan ■ Disk Mode: Dependent ■ Supports Encryption: Enabled ■ Preferred Storage: Enabled 	<ul style="list-style-type: none"> ■ cloud:private ■ region:sfo ■ tier:platinum ■ Storage:vsan
vmw-rainpole / sfo01dc01	VMware vSphere	sfo-silver	<ul style="list-style-type: none"> ■ StoragePolicy: Datastore Default ■ Datastore/ Datastore Cluster: sfo01-w01-nfs-dscluster-01 ■ Disk Mode: Dependent ■ Supports Encryption: Disabled ■ Preferred Storage: Disabled 	<ul style="list-style-type: none"> ■ cloud:private ■ region:sfo ■ tier:silver ■ Storage:nfs
vmw-rainpole / lax01dc01	VMware vSphere	lax-platinum	<ul style="list-style-type: none"> ■ StoragePolicy: vSAN Default Storage Policy ■ Datastore/ Datastore Cluster: lax01-w01-vsan ■ Disk Mode: Dependent ■ Supports Encryption: Enabled ■ Preferred Storage: Enabled 	<ul style="list-style-type: none"> ■ cloud:private ■ region:lax ■ tier:platinum ■ Storage:vsan

Table 2-46. Example Storage Profiles (continued)

Account / Region	Cloud Account Type	Storage Profile Name	Storage Characteristics	Capability Tags
vmw-rainpole / lax01dc01	VMware vSphere	lax-silver	<ul style="list-style-type: none"> ■ StoragePolicy: Datastore Default ■ Datastore/ Datastore Cluster: lax01-w01-nfs-dscluster-01 ■ Disk Mode: Dependent ■ Supports Encryption: Disabled ■ Preferred Storage: Disabled 	<ul style="list-style-type: none"> ■ cloud:private ■ region:lax ■ tier:silver ■ Storage:nfs
aws-rainpole / us-east-1	Amazon Web Services	us-east-platinum	<ul style="list-style-type: none"> ■ Device Type: EBS ■ Volume Type: Provisioned IOPS SSD ■ Supports Encryption: Disabled ■ Preferred Storage: Enabled 	<ul style="list-style-type: none"> ■ cloud:aws ■ region:us-east ■ tier:platinum
azure-rainpole / east us	Microsoft Azure	azure-us-east	<ul style="list-style-type: none"> ■ Storage Type: Managed Disks ■ Disk Type: Premium SSD ■ Preferred Storage: Enabled 	<ul style="list-style-type: none"> ■ cloud:azure ■ region:us-east ■ tier:platinum

Note The scope of this design includes VMware enabled clouds. Examples include both private and public cloud storage profiles for Cloud Assembly in a multi-cloud context.

Cloud-independent placement is possible. For example, different cloud storage might have different performance characteristics but still be considered for deployment based on capability tags. For example, a platform can have two different vendor accounts and a region. A storage profile for each region is tagged with a `tier:platinum` constraint key:value pair. During provisioning, a request containing a `tier:platinum:hard` tag for a hard constraint looks for a matching capability, regardless of the cloud vendor. A match then applies the settings from the associated storage profile during the creation of the deployed storage item.

Table 2-47. Design Decisions on Storage Profiles for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-037	For each account/region, add one or more storage profiles based on storage characteristics available for consumption.	Allows you to add storage with defined characteristics that can be consumed during a deployment process.	You must manage storage profiles for each account/region as storage is added, removed, and updated in a cloud environment.
SDDC-CAS-CA-038	For each storage profile, add one or more capability tags.	Allows you to manage the workload storage placement logic during the deployment process.	You must manage the capability tagging on each storage profile for the workflow placement logic during a deployment process.

Extensibility Design for Cloud Assembly in vRealize Automation

Cloud Assembly extensibility enables you to assign an extensibility action or workflow to an event by using subscriptions. When the specified event occurs, the subscription initiates the action or workflow to execute.

Workflow Extensibility

You can use vRealize Orchestrator workflows with Cloud Assembly to extend the application life cycle. Using this integration enables the use of your existing on-premises workflows with extensibility subscriptions.

Cloud Assembly uses vRealize Orchestrator integration to import and link workflows to a subscription. You can create, modify, and delete workflows using the vRealize Orchestrator Client. Workflows are maintained on the vRealize Orchestrator server instance.

Cloud Assembly supports the integration of multiple vRealize Orchestrator instances that can be used in workflow subscriptions. You can manage which vRealize Orchestrator integration is used in workflow subscriptions by using project extensibility constraints and capability tags.

Both soft or hard extensibility constraints on the project enable you to manage which vRealize Orchestrator integrations are used in blueprint provisioning for a project. Capability tags are also used to manage which vRealize Orchestrator integrations are used in workflow subscriptions.

For example, when you deploy a blueprint, Cloud Assembly uses the capability tags, associated with a cloud account, for example, `cloud:private` and `region:sfo`, to manage what vRealize Orchestrator integrations are used in workflow subscriptions.

Workflow and event subscription usage is beyond the scope of this design. The design for vRealize Orchestrator is in scope for the design.

For information about integrating vRealize Orchestrator workflows into the provisioning life cycle, see the vRealize Automation documentation.

Actions Extensibility

You can use lightweight action-based extensibility (ABX) code within Cloud Assembly to automate extensibility actions. You can assign an extensibility action to a Cloud Assembly subscription to extend your application life cycle.

Action-based extensibility provides a lightweight and flexible runtime where you can define small scriptable actions and configure them to initiate during particular events provided through the Event Broker Service (EBS). ABX supports both Node.js and Python runtime environments.

You can create or import extensibility action scripts in Cloud Assembly and assign them to subscriptions. Similarly to workflows, the extensibility action script triggers when an event included in an extensibility subscription occurs. Extensibility action scripts are used for more lightweight integrations and customizations. ABX is a functions-as-a-service (FaaS) natively through an on-premises provider vRealize Automation or a public cloud provider, such as Amazon Web Services Lambda or Microsoft Azure Functions. This is an alternative method to workflows that are hosted on-premises using a vRealize Orchestrator.

For example, the following Python script can be added as an action in Cloud Assembly.

```
import requests
import json

def handler(context, inputs):
    name = inputs['resourceNames'][0]
    os = inputs['customProperties']['image']
    reqid = inputs['requestId']
    print(name)
    data = {"text": "Workload "+name+" has been successfully deployed.\n Hack your workload with an
API call to request ID "+reqid}
    url = "https://hooks.slack.com/services/my-secret-webhook-url"
    r = requests.post(url, data = json.dumps(data), verify=False)
    if r.status_code == 200:
        print("success")
        outputs = {"statusCode":r.status_code}
        return outputs
    else:
        print("Failure")
        outputs = {"statusCode":r.status_code}
        return outputs
```

You can create extensibility actions by either writing a user-defined action script code or importing a predefined script code as a .ZIP package.

Action-Based Extensibility and the use of event subscription usage are beyond the scope of the design. Refer to the vRealize Automation documentation on integrating ABX into the provisioning lifecycle.

Table 2-48. Design Decisions on Action-Based Extensibility for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-CA-039	Use the embedded, on-premises functions-as-a-service provider for action-based extensibility.	You can use the native functions-as-a-service provider vRealize Automation for the execution of lightweight actions through event subscriptions without the requirement for a public cloud account provider, for example, Amazon Web Services Lambda or Microsoft Azure Functions.	<p>The use of action-based extensibility requires the vRealize Automation instance to have outbound access to the Internet to pull container images from publicly available Internet repositories. For example, to resolve Python and Node.js dependencies included in actions.</p> <p>If vRealize Automation is deployed on an isolated network that does not allow outbound traffic to the Internet, an HTTP proxy must be configured and applied using the <code>vraccli proxy</code> command option.</p>

vRealize Orchestrator Design for vRealize Automation

VMware vRealize Orchestrator contains a workflow library and a workflow engine with which you to create and run workflows that automate orchestration processes. You run workflows on objects of different technologies that vRealize Orchestrator accesses through a series of plug-ins.

Service Roles Design for vRealize Orchestrator in vRealize Automation

You manage access to vRealize Orchestrator by assigning enterprise groups to service roles in your organization.

vRealize Orchestrator has two service roles assigned from the organization identity and access management. You assign the service roles to designated enterprise groups, synchronized from your corporate identity source through Workspace ONE Access.

Table 2-49. Example Service Roles and Groups for vRealize Orchestrator in vRealize Automation

Role	Description	Enterprise Group
Orchestrator Administrator	Read and write access to the vRealize Orchestrator user interface and API to manage services and all objects.	rainpole.local\ug-vra-orchestrator-admins
Orchestrator Workflow Designer	Access to the vRealize Orchestrator user interface to design workflows.	rainpole.local\ug-vra-orchestrator-designers

For information about the decisions on the vRealize Orchestrator service roles, [Information Security and Access Control Design for vRealize Automation](#).

Authentication and Authorization Design for vRealize Orchestrator in vRealize Automation

You use service accounts, security groups, and roles to manage authentication and authorization. You establish secure communication with the vCenter Server instances by using CA-signed certificates.

Configure service accounts, security groups, and vCenter Server roles to control and manage the vRealize Orchestrator-to-workload domain vCenter Server endpoint instances. You define a service account with only the minimum set of permissions to perform necessary operations on the workload domain instances defined in the SDDC.

Table 2-50. Design Decisions on vRealize Orchestrator Authentication and Authorization

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-OR-001	Configure a service account in vCenter Server for application-to-application communication from the vRealize Orchestrator that is embedded in vRealize Automation to the workload domain vCenter Server instances in each region, for example, svc-vro-vsphere@rainpole.local .	Provides the following access control features: <ul style="list-style-type: none"> ■ The vRealize Orchestrator instance accesses the workload domain vCenter Server instance with the minimum set of required permissions. ■ If there is a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce improved accountability in tracking request-response interactions between the components of the SDDC. 	You must maintain the life cycle and availability of the service account outside of the SDDC stack.
SDDC-CAS-OR-002	Add the service account for application-to-application communication from the vRealize Orchestrator instance to a directory services security group, for example, rainpole.local\ug-vra-orchestrator-admins . Add any named administrative groups or users to the security group. Note In an Active Directory forest, consider using a security group with a universal scope. Add security groups with a global scope that includes service accounts and users from the domains in the Active Directory forest.	Only groups and users defined in the directory services security group can administer the vRealize Orchestrator instance.	You must maintain the security group membership, add or remove users, outside of the SDDC stack to ensure its membership.

Table 2-50. Design Decisions on vRealize Orchestrator Authentication and Authorization (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-OR-003	Create and apply a custom vCenter Server role for the service accounts used to add workload domain vCenter Server instances to the vRealize Orchestrator configuration, for example, vRealize Orchestrator to vSphere Integration .	You can limit the privileges for application-to-application integration.	<ul style="list-style-type: none"> ■ You must provide administrator privileges for the custom role to register vRealize Orchestrator with vCenter Server. Currently, your organization cannot restrict access based on the privileges required by the organization workflows on the workload domain vCenter Server instance. ■ You cannot integrate vRealize Orchestrator with VMware Cloud on AWS SDDC instances. Due to the restrictive access mode on the vCenter Server instance in VMware Cloud on AWS, the requirement for administrative-level permissions for vRealize Orchestrator inhibits the vCenter Server instance registration.
SDDC-CAS-OR-004	Assign global permissions for the vRealize Orchestrator-to-vSphere service account, for example, svc-vro-vsphere@rainpole.local .	Ensures that only the workload domain vCenter Server instance is accessible from vRealize Orchestrator.	<ul style="list-style-type: none"> ■ All vCenter Server instances must be in the same vSphere domain. ■ You must set the role for the management domain vCenter Server instances to No Access to ensure that the account cannot communicate with the management domain.

Integration Design for vRealize Orchestrator in vRealize Automation

The vRealize Orchestrator design includes guidance on client configuration, database configuration, SSL certificates, and plug-ins.

vRealize Orchestrator Client

With the vRealize Orchestrator user interface in vRealize Automation, you can import packages, create, run, schedule workflows, manage tags, and manage user permissions.

The vRealize Orchestrator user interface is available as a browser-based HTML5 client from **My Services** in the vRealize Automation, https://vra_cluster_fqdn/orchestrator-ui, or from the vRealize Orchestrator start page at https://vra_cluster_fqdn/vco.

Trusted Certificates

The vRealize Orchestrator user interface and API endpoint use a secure connection to communicate with workload domain vCenter Server instances, database systems, LDAP, and other servers. You can import an SSL certificate from a URL or from a PEM file to replace the SSL certificates that the embedded vRealize Orchestrator instance must trust, for example, an Enterprise Certificate Authority.

You can import workload domain vCenter Server instance SSL certificates from **Certificates > Trust Certificate** in the vRealize Orchestrator HTML5-based ControlCenter UI at https://vra_cluster_fqdn/vco-controlcenter.

Table 2-51. Design Decisions on vRealize Orchestrator Trusted Certificates

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-OR-005	Import the certificate for each workload domain vCenter Server instance in a region to the embedded vRealize Orchestrator instance in vRealize Automation.	Ensures that the certificate for each workload domain vCenter Server instance in a region is trusted by the embedded vRealize Orchestrator instance in vRealize Automation.	As workload domains are added or removed, you must add or remove the vCenter Server instance SSL certificate trust from the embedded vRealize Orchestrator instance in vRealize Automation.

vRealize Orchestrator Plug-Ins

You use vRealize Orchestrator plug-ins to access and control external services and applications. The external technologies that you can access by using plug-ins include visualization management tools, email systems, databases, directory services, and remote control interfaces. vRealize Orchestrator provides a set of standard plug-ins for technologies, such as, as the vCenter Server API.

vCenter Server Plug-In Design for vRealize Orchestrator in vRealize Automation

You can use the VMware vCenter Server[®] Plug-in for vRealize[®] Orchestrator[™] to manage multiple workload domain vCenter Server instances. You can create workflows that use the vCenter Server plug-in API to automate tasks in your workload domain environment.

The vCenter Server plug-in maps the vCenter Server API to JavaScript code that you can use in workflows. The plug-in also provides actions that perform individual vCenter Server tasks that you can include in workflows.

The vCenter Server plug-in provides a library of standard workflows that automate vCenter Server operations. For example, you can run workflows that create, clone, migrate, or delete virtual machines. Before managing and running workflows on the objects in your vSphere inventory, you must configure the vCenter Server plug-in and connect vRealize Orchestrator to the workload domain vCenter Server instances that you want to orchestrate.

Table 2-52. Design Decisions on vCenter Server Plug-In for vRealize Orchestrator in vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-OR-006	Register each workload domain vCenter Server instance in each region with the embedded vRealize Orchestrator instance. Do not use per session authentication.	Required for communication from the embedded vRealize Orchestrator to the workload domain vCenter Server instances.	<ul style="list-style-type: none"> ■ You cannot use per-session authentication for vRealize Orchestrator communication to vSphere as vCenter Server does not accept the token from vRealize Automation and the registration. ■ As workload domains are added or removed, you must add or remove the vCenter Server instance in vRealize Orchestrator. ■ You must execute the Update a vCenter Server Instance workflow with new login properties when the service account password changes during its life cycle.

Service Broker Design for vRealize Automation

Service Broker aggregates content in native formats from multiple clouds and platforms into a common catalog with the ability to apply governance based on roles and projects.

Service Roles Design for Service Broker in vRealize Automation

You manage access to Service Broker by assigning enterprise groups to service roles in your organization.

Service Broker has two service roles assigned from the organization identity and access management. You assign the service roles to designated enterprise groups, synchronized from your corporate identity source through Workspace ONE Access.

Table 2-53. Example Service Roles and Groups for Service Broker in vRealize Automation

Role	Description	Enterprise Group
Service Broker Administrator	<ul style="list-style-type: none"> ■ Read and write access to the Service Broker user interface and API resources. ■ Configure content sources and sharing, and customizations. ■ Configure policies. <p>Important Project Administrators must be granted the Service Broker Administrator role to perform customizations to blueprint icons and forms. However, members of this role are also entitled to manage cloud accounts, cloud zones, and integrations created by a Cloud Assembly Administrator.</p>	rainpole.local\ug-vra-service-broker-admins
Service Broker User	Request services from projects.	rainpole.local\ug-vra-service-broker-users

See the Identity and Access Management section for service role design decisions the vRealize Automation Service Broker service.

Catalog Content Design for Service Broker in vRealize Automation

The availability of the catalog items in Service Broker is determined by project membership. Projects link users, catalog items, and cloud resources.

Catalog content sources can include blueprints, actions, and workflows. Content sources are entitled to projects. Projects link a set of users, project members, with one or more target cloud zone regions or datastores.

When users request a catalog item, the deployment location is dependent on the project. Projects might have one or more cloud zones.

Content Source Design for Service Broker in vRealize Automation

Before you can release content, such as Cloud Assembly blueprints and actions, to the Service Broker catalog for project members, you must add a content source based on the content type and project.

You create content sources in Service Broker to import content types to the catalog.

Table 2-54. Content Types in Service Broker

Content Source Type	Description
Blueprints	Cloud Assembly versioned and released blueprints, created and managed in a project. Can be created directly in Cloud Assembly or imported from a Git repository.
Actions	<p>Cloud Assembly versioned and released extensibility actions, created and managed in a project. Can be either created directly in Cloud Assembly or synchronized with a Git repository.</p> <p>Important ABX using Amazon Web Services or Microsoft Azure is not applicable to this design.</p>
Workflows	<p>vRealize Orchestrator workflows:</p> <ul style="list-style-type: none"> Versioned workflows created and managed in vRealize Orchestrator or imported Available from the vRealize Orchestrator integration with the instance embedded in the vRealize Automation cluster
CloudFormation Templates	<p>Amazon CloudFormation templates of specifications for services or applications that you can deploy to an Amazon Web Services cloud account.</p> <p>Important Not applicable to this design.</p>

By default, content items are refreshed every six hours. Any released changes in an item from the source are reflected in the catalog after the refresh. Content can also be refreshed outside the standard cycle by initiating an on-demand import operation.

Table 2-55. Design Decisions on Content Sources in vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-SB-001	Add a blueprint content source for each Cloud Assembly project where blueprints are authored and released.	Provides the ability to share released blueprints with project members or other projects.	None
SDDC-CAS-SB-002	Add an actions content source for each Cloud Assembly project where actions are authored and released.	Provides the ability to share released actions with project members.	None
SDDC-CAS-SB-003	Add a vRealize Orchestrator workflows content source for each Cloud Assembly project, as required.	Provides the ability to share specific workflows with project members.	None

Content Customization Design for Service Broker in vRealize Automation

You customize icon and request forms for catalog items in Service Broker to enhance the user experience and capture additional workload inputs.

You use the content list to view the import source and entitled projects for each item. For each item, you can customize the catalog item icon and request form that is presented to the project members during a request. Icon customization allows you to use non-default icons to represent the catalog item and provide a better user experience to project members.

Form customizations allow you to add and configure elements, as well as data validation on a custom request form. By using input parameters, you can create useful forms and design how the information appears during a request, how the parameter values are populated and add any specialized constraints. Custom forms can be enabled or disabled on a per-catalog item basis.

Table 2-56. Design Decisions on Content Customizations in vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-SB-004	For each shared content item, customize the icon based on the catalog item.	Allows you to provide a meaningful visual indicator on the catalog item or type to enhance the user experience for the project members.	<ul style="list-style-type: none"> ■ Images must not exceed the dimension and file size maximums. ■ Images in PNG format are not supported.
SDDC-CAS-SB-005	For each shared content item, customize the form based on the catalog item and user experience requirements.	Allows you to create an intuitive user experience by using simple and discoverable forms that capture additional user inputs and in-form validations.	Requires customization of request forms per catalog item.

Content Sharing Design for Service Broker in vRealize Automation

You publish imported content and make it available in the Service Broker catalog for project members.

After you create a content source and import content items, you can share the items based on the release scope. When creating a blueprint in Cloud Assembly, you set the scope for sharing in Service Broker. You can restrict the blueprint to be shared only within its own project or make them available to any project in the organization.

For each project, you add released content items that are shared with project members. Items can be added by using two methods:

Content Source

A dynamic method to share all content items from the content source. As new items are added to a content source, they are made available in the catalog for the members of a target project.

All Content

A static method for sharing specific content items within a project. As new items are added to the content source, you must make them available in the catalog individually

The content sharing method you choose is dependent on your organization requirements, project structure, and the level of control required when releasing content to the catalog.

Policy Designfor Service Broker in vRealize Automation

You use policies to manage deployments requested from the Service Broker catalog with a defined set of rules.

Policies include definitions that are a set of rules or parameters for a specific use. The definition commonly includes the scope and enforcement type. The scope of a definition determines if the policy is applicable to all deployments in an organization or only to deployments in a selected project.

At the time of this publication, two policy types exist in Service Broker:

Lease Policy	Allows you to set a workload lease policy.
Day-2 Actions Policy	Allows you to select the actions that can be performed on a workload post deployment.

When a project member requests a blueprint, there might be more than one policy that applies. An enforcement type defines how multiple policies are evaluated, ranked, and, where applicable, merged, to produce an effective policy. An effective policy produces the intended results but is not always a specific named policy.

Table 2-57. Policy Enforcement Types

Enforcement Type	Description
Hard	Ranked higher than soft policies
Soft	Overridden by hard policies

During the evaluation phase, Service Broker first identifies and ranks policies.

- 1 Policy types are evaluated:
 - If there are hard and soft policies, then only the hard policies are considered and ranked.
 - If there are only soft policies, then the soft policies are ranked.
- 2 Policies with an organization scope are ranked higher than policies with a project scope.
- 3 Policies with older creation dates are ranked higher than policies with newer creation dates. Post ranking, policies are evaluated to identify the merge order.
- 4 The highest-ranking policy becomes the baseline.
- 5 The second-level ranking policy is applied next, and so on.
- 6 If a policy is incompatible with the preceding policies, then it is discarded and marked as ineffective.

Table 2-58. Example Lease Policy Configuration

Goal	Configuration	Effective Policy Behavior
A default organization-level policy that allows the project-level policy values to influence the applied values.	Organization Policy = Soft ■ Grace period: 10 ■ Lease: 100 ■ Total Lease: 100 Project A Policy 1 = Soft ■ Lease: 20 ■ Total Lease: 50 Project B Policy 1 = Soft ■ Lease: 10 ■ Total Lease: 30	A member of Project A requests a catalog item. Project B is not considered because the it is not applicable to Project A deployments. Merged Effective Policy: ■ Grace period: 10 ■ Lease: 20 ■ Total Lease: 100

Table 2-59. Design Decisions on Policies for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-SB-006	Identify and apply goals for your organization and each project based on the applicability of available policy types.	By understanding how the policies are processed, you can meet organization goals without creating an excessive and unmanageable number of policies.	For each policy type, you must determine the applicability and your organization goals to design policy enforcement and scope that result in the necessary effective policy.

Notifications Design for Service Broker in vRealize Automation

You configure notifications in Service Broker to send event messages to the user.

You configure a mail server in Service Broker to send outbound SMTP messages to users about system events.

Table 2-60. Design Decisions on Notifications for vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-CAS-SB-007	Configure Service Broker to use an outbound SMTP mail server to route notifications for system events.	Integrates vRealize Automation system events notifications to users by email to provide an enhanced user experience.	You must configure an SMTP server to relay messages from vRealize Automation.

Monitoring and Alerting Design for vRealize Automation

You integrate vRealize Operations with vRealize Automation to provide operational visibility.

The native integration to vRealize Automation from vRealize Operations provides the ability to monitor the health, efficiency, and capacity risks associated with cloud accounts. You can use the integration to perform the following:

- View the performance and health of cloud zones.
- Integrate and troubleshoot vCenter Server adapters problems associated with vRealize Automation.

For more information, see the vRealize Operations section of the design.

Data Protection and Backup Design for vRealize Automation

To preserve the cloud automation services functionality when data or system loss occurs, the design supports the use of data protection.

vRealize Automation supports data protection through the creation of consistent image-level backups, using backup software that is based on the vSphere Storage APIs - Data Protection (VADP).

Disaster Recovery Design for vRealize Automation

To preserve the cloud automation services functionality when a disaster occurs, the design supports the failover of vRealize Automation between regions.

You place vRealize Automation and its Workspace ONE Access components on the cross-region virtual network in the management domain. As a result, after recovery, you continue using the same IP addresses, DNS records, and routing configuration. Workspace ONE Access also use this network for their cross-region failover capabilities.

If a planned migration or disaster occurs, you use Site Recovery Manager and vSphere Replication for an orchestrated recovery of the vRealize Automation cluster and the supporting Workspace ONE Access cluster. After the recovery, vRealize Automation and Workspace ONE Access to provide cloud automation services.

Cloud Operations Design

The operations management design includes the software components that make up the operations management layer. The design provides guidance on the main elements of a product design such as deployment, sizing, networking, diagnostics, security, and integration with management solutions.

- Features of vRealize Suite Lifecycle Manager support initial installation and configuration of vRealize Suite products. Additional features support the life cycle management capabilities and configuration drift analysis for the vRealize Suite products.
- Monitoring operations support in vRealize Operations Manager and vRealize Log Insight provides performance, capacity management, and real-time logging of related physical and virtual infrastructure and cloud management components.

vRealize Suite Lifecycle Manager Design

vRealize Suite Lifecycle Manager provides life cycle management capabilities for vRealize components including automated deployment, configuration, patching, and upgrade as well as content management across vRealize products. You deploy vRealize Suite Lifecycle Manager as a single virtual appliance. In a multi-region SDDC, you can fail over the vRealize Suite Lifecycle Manager appliance across regions.

In this design, vRealize Suite Lifecycle Manager supports the following products:

- Workspace ONE Access (cross-region cluster)
- vRealize Log Insight
- vRealize Operations Manager

■ vRealize Automation

Logical Design for vRealize Suite Lifecycle Manager

vRealize Suite Lifecycle Manager provides life cycle management capabilities for vRealize components including automated deployment, configuration, patching, and upgrade, as well as content management across vRealize products. You deploy vRealize Suite Lifecycle Manager as a single virtual appliance. In a multi-region SDDC, you can fail over the vRealize Suite Lifecycle Manager appliance across regions.

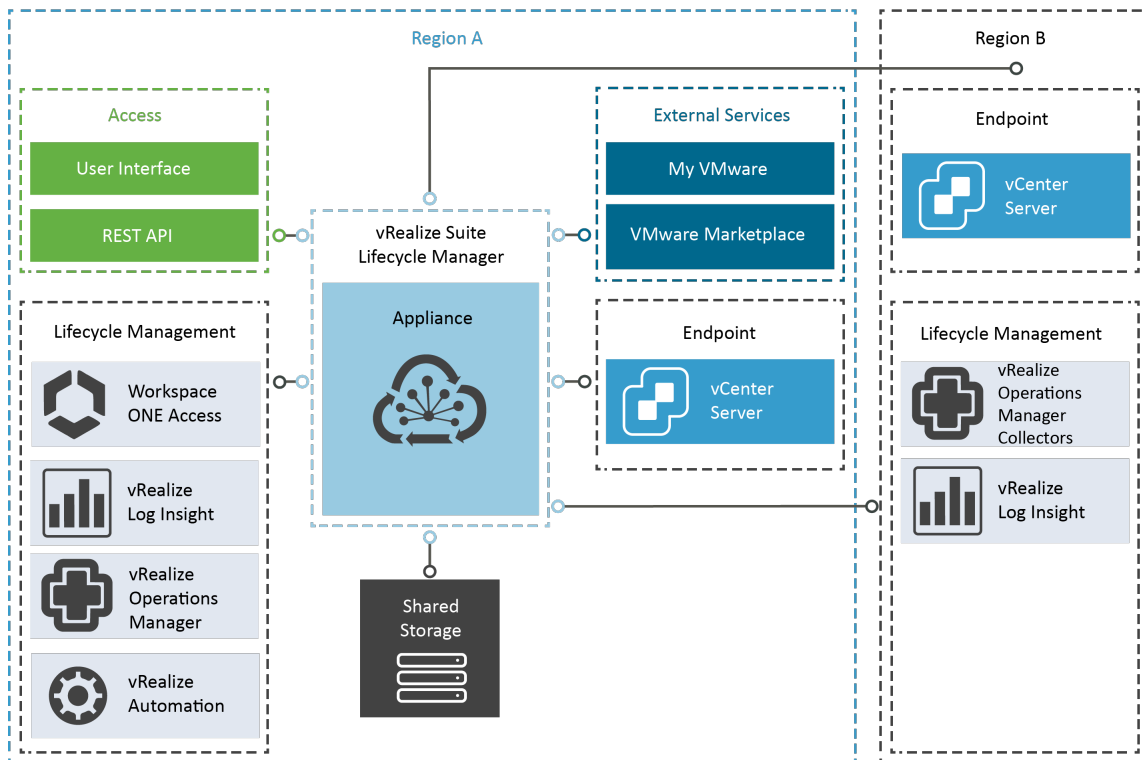
Logical Design

To orchestrate the deployment, patching, and upgrade of the vRealize products in the SDDC, vRealize Suite Lifecycle Manager communicates with each Management Domain vCenter Server instance in the SDDC.

vRealize Suite Lifecycle Manager consists of a single virtual appliance that deploys and upgrades the vRealize products across the virtual infrastructure that is controlled by one or more vCenter Server instances.

vRealize Suite Lifecycle Manager manages the lifecycle operations of cross-region and regional product instances and components.

Figure 2-13. Logical Design of vRealize Suite Lifecycle Manager in a Multi-Region Deployment



vRealize Suite Lifecycle Manager operates with the following elements and components:

Table 2-61. Comment Elements and Components in vRealize Suite Lifecycle Manager

Element	Components
Locker	<ul style="list-style-type: none"> ■ Passwords ■ Certificates ■ Licenses
Product Support	<ul style="list-style-type: none"> ■ Product binaries for install and upgrade (.ova, .pak, .iso) ■ Patch binaries ■ Product supports packs (.pspak)
Data Center	<ul style="list-style-type: none"> ■ Geographic location (optional) ■ vCenter Server instances
Environment	<ul style="list-style-type: none"> ■ Product deployments ■ Product import ■ Product operations, for example scale out, add a license, etc. ■ Product health
My VMware	<ul style="list-style-type: none"> ■ Product entitlement ■ Product downloads ■ Product licensing
Marketplace	<ul style="list-style-type: none"> ■ My VMware account ■ Marketplace content download and compatibility ■ vRealize Log Insight content packs ■ vRealize Operations Manager management packs ■ vRealize Automation blueprints ■ vRealize Orchestrator workflow packages

Configuration Design for vRealize Suite Lifecycle Manager

To deploy the vRealize products by using vRealize Suite Lifecycle Manager, you configure data centers, product support, environment structures, and product configuration drift.

Deployment Model for vRealize Suite Lifecycle Manager

vRealize Suite Lifecycle Manager is distributed as a virtual appliance in OVA format.

To accomplish this design objective, you deploy or reuse the following components to deploy this cloud operations solution for the SDDC.

- SDDC Manager
- vRealize Suite Lifecycle Manager
- NSX for vSphere Load Balancer for Workspace ONE Access
- Workspace ONE Access cluster
- Supporting infrastructure services, such as Active Directory, DNS, and NTP.

vRealize Suite Lifecycle Manager is distributed as a virtual appliance in OVA format. The vRealize Suite Lifecycle Manager appliance includes the cloud operations services for the life cycle management of vRealize products.

You place the vRealize Suite Lifecycle Manager appliance on a specific application virtual network for isolation and failover.

In the design, you deploy a single vRealize Suite Lifecycle Manager appliance instance on the first vSphere cluster in the management domain of Region A. With this configuration, you can centrally manage the life cycle of all vRealize products deployed across the entire SDDC. The SDDC can comprise multiple regions and multiple availability zones.

Sizing Compute and Storage Resources for vRealize Suite Lifecycle Manager

The vRealize Suite Lifecycle Manager appliance has the following resource requirements.

Table 2-62. vRealize Suite Lifecycle Manager CPU, Memory, and Storage Resources

Attribute	Appliance
CPU	2 vCPUs
Memory	6 GB
Storage	<ul style="list-style-type: none"> ■ 1.8 GB (thin provisioned) ■ 48 GB (thick provisioned)

When you deploy vRealize Suite Lifecycle Manager, consider the storage required for the following content:

- Product support for install, upgrade, patch, and support pack binaries
- Application and operating system logs
- Content management
- Marketplace content

Table 2-63. Design Decisions on the Deployment of vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-001	Deploy a single vRealize Suite Lifecycle Manager instance in the first vSphere cluster in the management domain in Region A to manage vRealize Suite products across all regions.	<ul style="list-style-type: none"> ■ vRealize Suite Lifecycle Manager can manage one or more regions and provides a single cloud operations service, regardless of region. ■ Because of the isolation of vRealize Suite Lifecycle Manager over virtual networking, it is independent of any physical site locations or hardware. 	None
SDDC-OPS-LCM-002	Deploy vRealize Suite Lifecycle Manager through SDDC Manager.	Allows SDDC Manager the ability to provide life cycle management of vRealize Suite components with its integration to vRealize Suite Lifecycle Manager.	None

Table 2-63. Design Decisions on the Deployment of vRealize Suite Lifecycle Manager (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-003	Increase the initial storage of the vRealize Suite Lifecycle Manager virtual appliance by 100 GB.	Support for product binaries, such as install, upgrade, and patch, and content management.	None
SDDC-OPS-LCM-004	When using two availability zones in Region A, add the vRealize Suite Lifecycle Manager virtual appliance to the primary availability zone VM group, for example, sfo01-m01-mgmt01-primary-az-vm-group.	Ensures the vRealize Suite Lifecycle Manager virtual appliance is powered on within the primary availability zone hosts group by default.	If vRealize Suite Lifecycle Manager is deployed after the creation of the stretched clusters for management domain availability zones, the VM group for the primary availability zone virtual machines must be updated to include the vRealize Suite Lifecycle Manager virtual appliance.
SDDC-OPS-LCM-005	Place the cross-region vRealize Suite Lifecycle Manager virtual appliance in a dedicated virtual machine folder in Region A, for example, xregion-sfo01-lax01-m01fd-vrslcm.	Provides the organization of the vRealize Suite Lifecycle Manager virtual appliance in the management domain inventory and preparation for Site Recovery Manager folder mappings for disaster recovery.	A corresponding virtual machine folder in Region B must be created in preparation for the Site Recovery Manager folder mapping, for example, xregion-lax01-sfo01-m01fd-vrslcm.

Logging Design for vRealize Suite Lifecycle Manager

You integrate vRealize Suite Lifecycle Manager with vRealize Log Insight to provide operational visibility.

The native integration to vRealize Log Insight from vRealize Suite Lifecycle Manager provides the ability to send logs from the service containers for aggregation and analysis, as needed.

Logging to a vRealize Log Insight instance through the ingestion API is established by updated in the appliance settings in the vRealize Suite Lifecycle Manager user interface or by updating the vRealize Log Insight `liagent.ini`.

For more information, see the vRealize Log Insight section of the design.

Table 2-64. Design Decisions on Logging for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS- LCM-006	Configure vRealize Suite Lifecycle Manager to send logs to the vRealize Log Insight cluster in Region A. Use the vRealize Log Insight Agent from the Region A region-specific vRealize Log Insight cluster for cross-region Workspace ONE Access cluster.	Allows logs from vRealize Suite Lifecycle Manager to be forwarded to a vRealize Log Insight cluster.	A vRealize Log Insight Content Pack for vRealize Suite Lifecycle Manager is not available.
SDDC-OPS- LCM-007	Communicate with the vRealize Log Insight using the default Ingestion API (cfapi) port 9000 and non-default No SSL.	Supports disaster recovery of vRealize Suite Lifecycle Manager in the SDDC.	Transmission traffic for logs is not secure.

For more information, refer to the vRealize Log Insight section of the design.

Life Cycle Management Design of vRealize Suite Lifecycle Manager

The life cycle management design details the design decisions covering the life cycle management of vRealize Suite Lifecycle Manager.

The life cycle management of vRealize Suite Lifecycle Manager involves the process of performing patch updates or upgrades to vRealize Suite Lifecycle Manager.

In this design, the life cycle management of vRealize Suite Lifecycle Manager is performed by using vRealize Suite Lifecycle Manager itself.

Table 2-65. Design Decisions on the Life Cycle Management of vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-008	Use vRealize Suite Lifecycle Manager to perform the life cycle management of vRealize Suite Lifecycle Manager.	vRealize Suite Lifecycle Manager upgrades vRealize Suite Lifecycle Manager through its own user interface.	None

Network Design for vRealize Suite Lifecycle Manager

For secure access to the UI and API and for failover of vRealize Suite Lifecycle Manager, you place the appliance in the shared cross-region application virtual network.

Application Network Segment

The vRealize Suite Lifecycle Manager virtual appliance is connected to the cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN, for secure access to the application UI and API, and for failover support.

This networking design has the following features:

- vRealize Suite Lifecycle Manager can be failed over between regions if there is a planned migration or disaster recovery without changing any IP address, DNS records, or routing configurations. Workspace ONE Access, vRealize Automation, and vRealize Operations also share this network for cross-region failover support.
- vRealize Suite Lifecycle Manager has a routed access to the VLAN-backed management network through the NSX Universal Distributed Logical Router.
- Routing to the VLAN-management network, application virtual networks, and external networks are dynamic and are based on the Border Gateway Protocol (BGP).

Figure 2-14. Networking Design of the vRealize Suite Lifecycle Manager Deployment

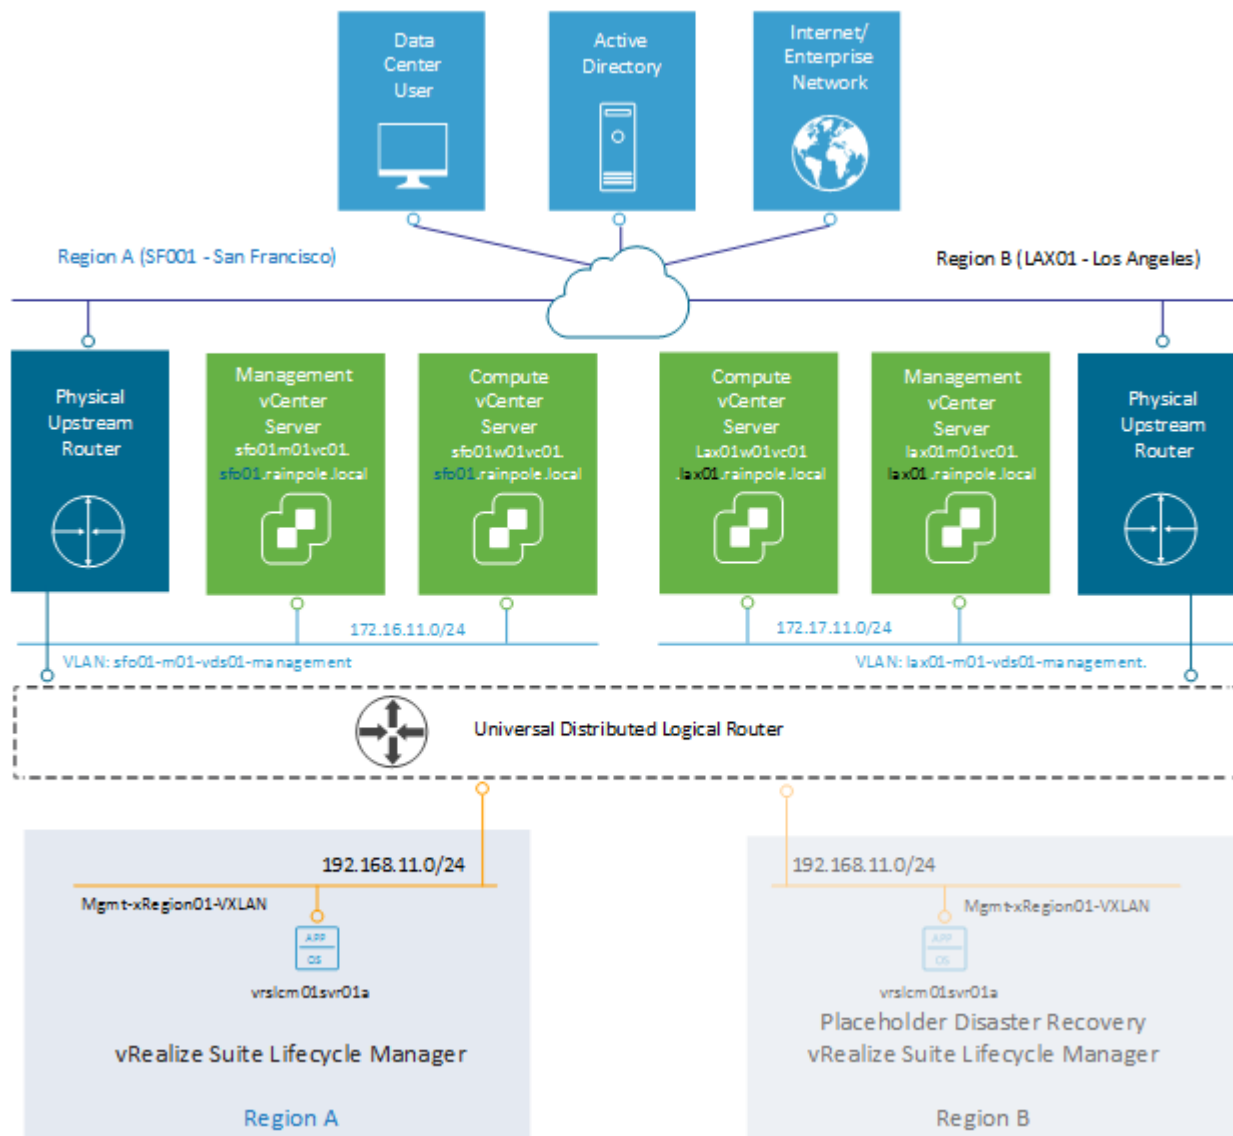


Table 2-66. Design Decisions on the Application Virtual Network for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-009	Place the vRealize Suite Lifecycle Manager appliance on the cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN.	Supports secure access from an external location and disaster recovery.	You must use an implementation in NSX for vSphere to support this networking configuration.

IP Addressing Scheme

You allocate a subnet for the cross-region network segment in the management domain and use it for the vRealize Suite Lifecycle Manager deployment.

Table 2-67. Example IP Subnet for vRealize Suite Lifecycle Manager

Solution	Example IP	Gateway	NSX Application Virtual Network
vRealize Suite Lifecycle Manager in Region A	192.168.11.0/24	192.168.11.1	Mgmt-xRegion01-VXLAN

Name Resolution

The host name of the vRealize Suite Lifecycle Manager appliance follows a specific domain name resolution:

- The IP address of the vRealize Suite Lifecycle Manager appliance is associated with a fully qualified name whose suffix is set to the root domain `rainpole.local`.

Table 2-68. Example FQDN and IP Address for vRealize Suite Lifecycle Manager

Fully Qualified Domain Name	IP Address	Description	Region	Failed Over to Region B
vrs lcm01svr01.rainpole.local	192.168.11.20	vRealize Suite Lifecycle Manager	Region A	<input checked="" type="checkbox"/>

Table 2-69. Design Decisions on DNS for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-010	Configure forward and reverse DNS records for the vRealize Suite Lifecycle Manager appliance.	vRealize Suite Lifecycle Manager is accessible by using a fully qualified domain name instead of by using the IP address only.	You must provide DNS records for the vRealize Suite Lifecycle Manager appliance.
SDDC-OPS-LCM-011	In a multi-region deployment, configure the DNS settings for the vRealize Suite Lifecycle Manager appliance to use DNS servers in each region.	vRealize Suite Lifecycle Manager can resolve DNS from regional DNS servers during a planned migration or disaster recovery between regions.	As you scale from a single region to multi-region deployment, the DNS settings the vRealize Suite Lifecycle Manager appliance must be updated.

Time Synchronization

vRealize Suite Lifecycle Manager is dependent on time synchronization.

Table 2-70. Design Decisions on NTP for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-012	Configure NTP on the vRealize Suite Lifecycle Manager appliance.	vRealize Suite Lifecycle Manager is dependent on time synchronization.	None
SDDC-OPS-LCM-013	In a multi-region deployment, configure the NTP settings for the vRealize Suite Lifecycle Manager appliance to use NTP servers in each region.	vRealize Suite Lifecycle Manager can query NTP from regional NTP servers to synchronize time during a planned migration or disaster recovery between regions.	As you scale from a single region to multi-region deployment, the NTP settings on the vRealize Suite Lifecycle Manager appliance must be updated.

Information Security and Access Control Design for vRealize Suite Lifecycle Manager

You protect the vRealize Suite Lifecycle Manager deployment by configuring the authentication and secure communication with the other components in the SDDC. You dedicate a service account to the communication between vRealize Suite Lifecycle Manager and vCenter Server.

You use a custom role in vSphere with permissions to perform life cycle operations on vRealize Suite components in the SDDC. A dedicated service account is assigned a custom role for communication between vRealize Suite Lifecycle Manager and the vCenter Server instances in the environment.

Authentication and Authorization Design for vRealize Suite Lifecycle Manager

Users can authenticate to vRealize Suite Lifecycle Manager by using local administrator accounts or by using Workspace ONE Access.

vRealize Suite Lifecycle Manager performs local authentication for the default administrator, that is, the **admin@local** account only. You enable authentication by using Workspace ONE Access to ensure accountability on user access. You can grant both users and groups access to vRealize Suite Lifecycle Manager to perform tasks, and initiate orchestrated operations, such as deployment and upgrade of vRealize Suite components and content.

Configure a service account for communication between vRealize Suite Lifecycle Manager and vCenter Server endpoint instances. You define a service account with only the minimum set of permissions to perform inventory data collection and life cycle management operations for the instances defined in the data center.

Table 2-71. Design Decision on Authentication and Authorization for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-014	Rotate the root password on or before 365 days post-deployment.	The password for the root user account expires 365 days after the initial deployment.	You must manage the password rotation schedule for the root user account in accordance with your organization policies and regulatory standards, as applicable.
SDDC-OPS-LCM-015	Enable the vRealize Suite Lifecycle Manager integration with your corporate identity source using Workspace ONE Access.	<ul style="list-style-type: none"> ■ Allows authentication, including multi-factor, to vRealize Suite Lifecycle Manager using your corporate identity source. ■ Allows authorization through the assignment of organization and cloud services roles to enterprise users and groups defined in your corporate identity source. 	You must deploy and configure the Workspace ONE Access to establish the integration between vRealize Suite Lifecycle Manager and your corporate identity sources.
SDDC-OPS-LCM-016	Define a custom vCenter Server role, for example, vRealize Suite Lifecycle Manager to vSphere Integration , for vRealize Suite Lifecycle Manager that has the minimum privileges required to support the deployment and upgrade of vRealize Suite products in the design.	vRealize Suite Lifecycle Manager accesses vSphere with the minimum set of permissions that are required to support the deployment and upgrade of vRealize Suite products in the design.	You must maintain the permissions required by the custom role.
SDDC-OPS-LCM-017	Configure a service account, for example, svc-vrslcm-vsphere@rainpole.local , in vCenter Server for application-to-application communication from vRealize Suite Lifecycle Manager to vSphere.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ vRealize Suite Lifecycle Manager accesses vSphere with the minimum set of required permissions. ■ If there is a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	You must maintain the life cycle and availability of the service account outside of the SDDC stack.

Table 2-71. Design Decision on Authentication and Authorization for vRealize Suite Lifecycle Manager (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS- LCM-018	Assign global permissions for the vRealize Suite Lifecycle Manager to vSphere service account, for example, svc-vrslcm-vsphere@rainpole.local , in vCenter Server using the custom role, for example, vRealize Suite Lifecycle Manager to vSphere Integration .	<ul style="list-style-type: none"> ■ vRealize Suite Lifecycle Manager accesses vSphere with the minimum set of permissions that are required to support the deployment and upgrade of vRealize Suite products in the design. ■ Allows for a content management integration with the vSphere Content Library, if necessary. 	<ul style="list-style-type: none"> ■ All vCenter Server instances must be in the same vSphere domain. ■ You must maintain the assignment of the service account and the custom role at a cluster level for each management cluster instead of using global permissions. ■ If you do not plan to use the content management feature of vRealize Suite Lifecycle Manager, you must set the role on each workload domain vCenter Server instance to No Access to ensure that the account cannot communicate with the workload domain.
SDDC-OPS- LCM-019	Create a security group in your organization directory services for the vRealize Suite Lifecycle Manager administrators role, for example rainpole.local\ug-vrslcm-admins , and synchronize the group in the Workspace ONE Access configuration for vRealize Suite Lifecycle Manager.	Allows you to streamline the management of vRealize Suite Lifecycle Manager roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the desired directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS- LCM-020	Assign the enterprise group for vRealize Suite Lifecycle Manager administrators, for example, rainpole.local\ug-vrslcm-admins , the LCM Admin role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to vRealize Suite Lifecycle Manager administration is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking organization owner access to vRealize Suite Lifecycle Manager. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-71. Design Decision on Authentication and Authorization for vRealize Suite Lifecycle Manager (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-021	Create a security group in your organization directory services for the vRealize Suite Lifecycle Manager content managers role, for example, rainpole.local\ug-vrslcm-content-managers , and synchronize the group in the Workspace ONE Access configuration for vRealize Suite Lifecycle Manager.	Allows you to streamline the management of vRealize Suite Lifecycle Manager roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS-LCM-022	<p>Assign the enterprise group for vRealize Suite Lifecycle Manager content managers, for example, rainpole.local\ug-vrslcm-content-managers, the Content Manager role.</p> <p>Note The content management feature is out of scope for this design. However, this design accounts for the Identity and Access Management controls for the feature.</p>	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to the vRealize Suite Lifecycle Manager content management is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking organization owner access to vRealize Suite Lifecycle Manager. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-OPS-LCM-023	Create a security group in your organization directory services for the vRealize Suite Lifecycle Manager content developers role, for example, rainpole.local\ug-vrslcm-content-developers , and synchronize the group in the Workspace ONE Access configuration for vRealize Suite Lifecycle Manager.	Allows you to streamline the management of vRealize Suite Lifecycle Manager roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS-LCM-024	<p>Assign the enterprise group for vRealize Suite Lifecycle Manager content developers, for example, rainpole.local\ug-vrslcm-content-developers, the Content Developer role.</p> <p>Note The content management feature is out of scope for this design. However, this design accounts for the Identity and Access Management controls for the feature.</p>	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to vRealize Suite Lifecycle Manager content development is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking organization owner access to vRealize Suite Lifecycle Manager. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Marketplace Integration Design for vRealize Suite Lifecycle Manager

You can use vRealize Suite Lifecycle Manager to add and manage content from VMware Marketplace. After you download Marketplace content, you can direct the content deployment to your SDDC directly from vRealize Suite Lifecycle Manager.

To use the integration with the VMware Marketplace, you must register the vRealize Suite Lifecycle Manager appliance with My VMware and the appliance must have Internet access.

You can download additional content packs from the Marketplace in vRealize Suite Lifecycle Manager for integration in the SDDC.

You can also use vRealize Suite Lifecycle Manager to download and install additional vRealize Operations management packs from the Marketplace.

For information about the content packs, management packs, and their versions in this design, see the Release Notes.

Table 2-72. Design Decisions on Marketplace for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-025	Use a dedicated My VMware account for vRealize Suite Lifecycle Manager instead of a named user account for the Marketplace integration.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Accessibility and privileges on the destination service remain restricted to an integration account. ■ Accountability in tracking interactions between the vRealize Suite Lifecycle Manager and My VMware. 	<ul style="list-style-type: none"> ■ The use of the Marketplace requires the vRealize Suite Lifecycle Manager instance to have outbound access to the Internet to pull content from the VMware Content Delivery Network. If vRealize Suite Lifecycle Manager is deployed on an isolated network that does not allow outbound traffic to the Internet, an HTTP proxy must be configured. ■ You must manage a dedicated My VMware account for use with vRealize Suite Lifecycle Manager. ■ If used for product downloads, the My VMware account must access to vRealize Suite product downloads and licenses.

Encryption Design for vRealize Suite Lifecycle Manager

Access to all vRealize Suite Lifecycle Manager endpoint interfaces requires an SSL connection. By default, vRealize Suite Lifecycle Manager uses a self-signed certificate for the appliance. To provide secure access to the vRealize Suite Lifecycle Manager and between SDDC endpoints, replace the default self-signed certificate with a CA-signed certificate.

Table 2-73. Design Decisions on Encryption for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-026	Replace the default self-signed certificate of the virtual appliance of vRealize Suite Lifecycle Manager with a CA-signed certificate.	Configuring a CA-signed certificate ensures that the communication to the externally facing Web UI and API for vRealize Suite Lifecycle Manager, and cross-product, is encrypted.	Replacing the default certificates with trusted CA-signed certificates from a certificate authority might increase the deployment preparation time as certificates requests are generated and delivered.

Locker Design for vRealize Suite Lifecycle Manager

Locker allows you to secure and manage passwords, certificates, and licenses for vRealize product solutions and integrations.

Locker Passwords Design in vRealize Suite Lifecycle Manager

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

Table 2-74. Lifecycle Operations Use of Locker Passwords in vRealize Suite Lifecycle Manager

Lifecycle Operations Element	Password Use
Data Centers	<ul style="list-style-type: none"> ■ vCenter Server credentials: vRealize Suite Lifecycle Manager to vSphere integration user, for example, svc-vrslcm-vsphere@rainpole.local.
Environments	<ul style="list-style-type: none"> ■ Global environment default configuration administrator, for example, configadmin ■ Environment password, for example, product default admin and root
Products	<ul style="list-style-type: none"> ■ Product administrator password, for example, admin for an individual product ■ Product appliance password, for example, root for and individual product
Integrations	<ul style="list-style-type: none"> ■ My VMware users, for example, svc-vrslcm-myvmware@rainpole.com, Licensing and Marketplace

Table 2-75. Design Decisions on Locker Passwords in vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-027	Replace the default Store passwords in the locker repository for use by the life cycle operations.	<p>Allows you to reference specific passwords for use across life cycle operations elements, such as:</p> <ul style="list-style-type: none"> ■ vCenter Server registration and updates (Management Domain vCenter Servers) ■ Environment creations ■ Product deployments and updates ■ My VMware registration and updates 	<ul style="list-style-type: none"> ■ Password items in the locker cannot be edited or deleted from the UI. However, they can be deleted using the API. You must register and use a new locker password when rotating a password. ■ When using the API, you must specify the locker ID for the password to be used in the JSON payload, for example, locker:password:39e54078-cefa-4979-ac5b-989e986b7aa5:20191023-svc-vrslcm-vsphere@rainpole.local.

Locker Certificates Design in vRealize Suite Lifecycle Manager

vRealize Suite Lifecycle Manager stores certificates in the locker repository which can be referenced during product life cycle operations. Externally provided certificates, such as Certificate Authority signed certificates, can be imported or certificates generated by the vRealize Suite Lifecycle Manager Appliance.

The certificate validity - such as the issued date, expiration date, time remaining - and certificate details - such as the issuer, subject, and subject alternative names - are available for reference along with the certificate health based on the expiration date. Also, you can review the certificate reference to see where the certificate is in use across environments and products. As certificates must be replaced, such as with expiration or a cluster scale-out, the locker provides the ability to replace certificates on referenced entities.

Table 2-76. Design Decision on Locker Certificates in vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-028	Import Certificate Authority signed certificates to the locker repository for product life cycle operations.	<ul style="list-style-type: none"> ■ Allows you to review the validity, details, and the environment and deployment use for the certificate across the vRealize products. ■ Allows you to reference and use Certificate Authority signed certificates during product life cycle operations, such as deployment and certificate replacement. 	<p>When using the API, you must specify the locker ID for the certificate to be used in the JSON payload, for example, locker:certificate:438bb1f7-fdad-4b41-ab22-9b41f0014b97:20191023-vra01svr01.rainpole.local.</p>

Locker Licenses Design in vRealize Suite Lifecycle Manager

vRealize Suite Lifecycle Manager stores licenses in the locker repository which can be referenced during product life cycle operations. Licenses can be validated and added to repository directory or imported through an integration with My VMware.

The license details - such as the issued date, expiration date, time remaining - and certificate details - such as the type, quantity, unit, and expiration - are available for reference. Also, you can review the license references to see where the license is in use across environments and products. As licenses must be replaced, such as with workload domain expansion, the locker provides the ability to replace the license on an individual or all referenced entities.

Table 2-77. Design Decision on Locker Licenses in vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-029	Import vRealize product licenses to the locker repository for product life cycle operations.	<ul style="list-style-type: none"> Allows you to see review the validity, details, and the environment and deployment use for the license across the vRealize products. Allows you to reference and use licenses during product life cycle operations, such as deployment and license replacement. 	When using the API, you must specify the locker ID for the license to be used in the JSON payload, for example, locker:license:09ca76a2-186a-44ac-a1b9-472a67c659a5:vRealize Suite 2019 Enterprise.

SDDC Life Cycle Operations Design

To deploy the vRealize products by using vRealize Suite Lifecycle Manager, you configure product support, data centers, environment structures, and product specifications.

Product Support Design in vRealize Suite Lifecycle Manager

vRealize Suite Lifecycle Manager provides two methods to obtain and store product binaries for the install, patch, and upgrade of the vRealize products.

Table 2-78. Methods for Obtaining and Storing Product Binaries

Method	Description
Product Upload	<ul style="list-style-type: none"> You can upload and discover product binaries to the vRealize Suite Lifecycle Manager appliance. You can upload the product and patch binaries directly to a folder, for example, /data/upload of the appliance or to an NFS share, after which you can discover and add the product binaries to the repository. If you remove the individual product or patch binaries that are discovered, vRealize Suite Lifecycle Manager removes the metadata from the repository but you must manage the related file on the file system.
My VMware	<ul style="list-style-type: none"> You can integrate vRealize Suite Lifecycle Manager with My VMware to access and download vRealize product entitlements. This method simplifies, automates, and organizes the repository. If you remove individual product or patch binaries that are downloaded from the My VMware integration, the vRealize Suite Lifecycle Manager removes the related files and metadata from the repository.

Table 2-79. Design Decisions on Product Support for vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-030	Upload and discover the vRealize product binaries for install, patch, and upgrade binaries.	<ul style="list-style-type: none"> ■ Allows product binaries for vRealize product install, patch, and upgrade to be provided through SDDC Manager. ■ It allows you to deploy and manage the design in an environment that does not allow access to the Internet or are dark sites. 	<ul style="list-style-type: none"> ■ Product binaries for install, patch, and upgrade are transferred from SDDC Manager to the vRealize Suite Lifecycle Manager repository for registration. ■ Because you use the local vRealize Suite Lifecycle Manager repository, product binaries are transferred across the WAN during product install, patch, and upgrade across regions.

Environments and Data Centers Design in vRealize Suite Lifecycle Manager

vRealize Suite Lifecycle Manager supports the deployment and upgrade of vRealize Suite products in a logical environment grouping.

These products included in this design include:

- Workspace ONE Access (cross-region cluster)
- vRealize Log Insight
- vRealize Operations Manager
- vRealize Automation

An environment is a logical object that is mapped to a data center object in vRealize Suite Lifecycle Manager. Each environment can contain only one instance of a vRealize product. For example, only one vRealize Log Insight cluster can exist in an environment. However, you can use vRealize Suite Lifecycle Manager to deploy and scale-out this vRealize Log Insight cluster in the environment to the required number of nodes. In this example, you can add an extra environment that can contain a second vRealize Log Insight cluster.

A data center is another logical object in vRealize Suite Lifecycle Manager to represent a geographical or logical location for an organization. Management Domain vCenter Server instances are added to specific data centers.

In this design, you create data centers and environments in vRealize Suite Lifecycle Manager to manage the life cycle operations on the vRealize products and to support the growth of the SDDC.

You create the following data center and environment objects:

Table 2-80. Logical Data Center to vCenter Server Mappings in vRealize Suite Lifecycle Manager

Logical Data Center	vCenter Server Type	vCenter Server Type	Description
Cross-Region	Management	<ul style="list-style-type: none"> ■ Management Domain vCenter Server in Region A ■ Management Domain vCenter Server in Region B 	Supports the deployment of cross-region components like Workspace ONE Access, vRealize Operations, and vRealize Automation, including any regional collector components.
Region A	Management	<ul style="list-style-type: none"> ■ Management Domain vCenter Server in Region A 	Supports the deployment of vRealize Log Insight in Region A. vRealize Log Insight has instances across the SDDC, each instance is designated to an SDDC region. You deploy each instance using a separate logical data center and environment.
Region B	Management	<ul style="list-style-type: none"> ■ Management Domain vCenter Server in Region B 	Supports the deployment of vRealize Log Insight in Region B.

Table 2-81. Environment Topologies

Environment Name	Logical Data Center	Product Components
Global Environment	Cross-Region	Workspace ONE Access cluster nodes
Cross-Region	Cross-Region	<ul style="list-style-type: none"> ■ vRealize Operations Manager analytics cluster nodes ■ vRealize Operations Manager remote collectors ■ vRealize Automation cluster
Region A	Region A	vRealize Log Insight cluster nodes
Region B	Region B	vRealize Log Insight cluster nodes

Table 2-82. Design Decision on Environment and Data Center Management in vRealize Suite Lifecycle Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS- LCM-031	<ul style="list-style-type: none"> ■ Create a data center object in vRealize Suite Lifecycle Manager for SDDC solutions that are managed across regions. ■ Assign the Management vCenter Server instance in each region to the data center. 	Allows you to deploy and manage the integrated vRealize Suite components across the SDDC as a group.	None
SDDC-OPS- LCM-032	<ul style="list-style-type: none"> ■ Create a data center object in vRealize Suite Lifecycle Manager for each region. ■ Assign each data center object the Management vCenter Server instance for the region. 	Supports deployment and management of vRealize products that are region-specific.	You must manage a separate data center object for the products that are specific to each region.
SDDC-OPS- LCM-033	Create the initial global environment, for example, globalenvironment, in vRealize Suite Lifecycle Manager for required Workspace ONE Access instance.	The global environment is required by vRealize Suite Lifecycle Manager.	None
SDDC-OPS- LCM-034	<p>Create an environment in vRealize Suite Lifecycle Manager for SDDC solutions that are cross-region:</p> <ul style="list-style-type: none"> ■ vRealize Operations Manager analytics cluster nodes ■ vRealize Operations remote collectors ■ Realize Automation cluster nodes 	<ul style="list-style-type: none"> ■ Supports deployment and management of the integrated vRealize products across the SDDC regions as a group. ■ Enables the deployment of region-specific components, such as vRealize Operations remote collectors. In vRealize Lifecycle Manager, you can deploy and manage vRealize Operations remote collector objects only in an environment that contains the associated cross-region instance. components. 	You can manage region-specific components, such as remote collectors, only in an environment that is cross-region.
SDDC-OPS- LCM-035	<p>Create an environment in vRealize Suite Lifecycle Manager for each region to deploy and manage the standalone vRealize products that are region-specific:</p> <ul style="list-style-type: none"> ■ vRealize Log Insight cluster nodes 	Supports the deployment of an instance of a management product in each region. Using vRealize Lifecycle Manager, you can deploy only one instance of a vRealize product per environment. You use a separate environment for each region where you deploy a product instance.	You must maintain an environment for each region to deploy and manage the standalone region-specific solutions.

Product Deployment Design in vRealize Suite Lifecycle Manager

Products are deployed by using the vRealize Suite Lifecycle Manager constructs named environments.

Each environment has the following high-level attributes:

- Environment name
- Administrator email
- Administrator password (from the Locker)
- Datacenter
- Customer Experience Improvement Program State

The vRealize Suite Lifecycle Manager UI provides the following installation methods for environment creation.

- Installation wizard (default)
- JSON configuration file

You can deploy new vRealize products to the SDDC environment or import existing product deployments.

When you add one or more products to an environment, the parameters differ for a new product deployment and for an existing product import.

For example, for a new deployment, you can provide the following parameters:

Table 2-83. Example Deployment Configuration Elements

Elements	Properties
Product Selection	<ul style="list-style-type: none"> ■ Selection ■ Version ■ Deployment type, for example, standard or cluster, if applicable ■ Node count if applicable
Certificate	Certificate Authority-signed certificate (from Locker)
Infrastructure	<ul style="list-style-type: none"> ■ vCenter Server ■ Cluster ■ Datastore ■ Disk Mode ■ Folder ■ Resource Pool ■ Affinity/Anti-Affinity Rules
Network Settings	<ul style="list-style-type: none"> ■ Network ■ Gateway ■ NTP ■ DNS ■ Domain Search Path

Table 2-83. Example Deployment Configuration Elements (continued)

Elements	Properties
Product Properties	<ul style="list-style-type: none"> ■ Product Properties <ul style="list-style-type: none"> ■ License (from Locker) ■ Cluster VIPs ■ Workspace ONE Integration ■ Virtual Machine Names ■ Virtual Machine IP Addresses ■ Fully Qualified Domain Names ■ Node Size ■ Password Overrides (from Locker) ■ Infrastructure Overrides ■ Network Overrides
Pre-checks	<ul style="list-style-type: none"> ■ Data Validation ■ Infrastructure Validation ■ Network Validation

You can deploy products by using a configuration file in JSON format. When using the API, the JSON configuration file is provided as the payload for the environment specification.

Example Cross-Region JSON Payload Specification

```
{
  "environmentId": "4980533e-000d-4a0d-a24c-7c8da162c9a8",
  "environmentName": "cross-region",
  "infrastructure": {
    "properties": {
      "cluster": "sfo01-m01-mgmt01",
      "dns": "172.16.11.6,172.17.11.5",
      "diskMode": "thin",
      "vCenterHost": "sfo01m01vc01.sfo01.rainpole.local",
      "storage": "sfo01-m01-vsan",
      "ntp": "ntp.sfo01.rainpole.local,ntp.lax01.rainpole.local",
      "vCenterName": "sfo01m01vc01.sfo01.rainpole.local",
      "network": "vxw-dvs-28-universalwire-3-sid-30002-Mgmt-xRegion01-VXLAN",
      "masterVidmEnabled": "false",
      "netmask": "255.255.255.0",
      "vcUsername": "svc-vrslcm-vsphere@rainpole.local",
      "domain": "rainpole.local",
      "vcPassword": "locker:password:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:svc-vrslcm-
vsphere@rainpole.local",
      "gateway": "192.168.11.1",
      "searchpath": "rainpole.local",
      "defaultPassword": "locker:password:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:admin-cross-region",
      "adminEmail": "no-reply@rainpole.local"
    }
  },
  "products": [
    {
      "id": "vrops",
      "version": "8.0.1",
```

```

    "patchHistory":null,
    "snapshotHistory":null,
    "logHistory":null,
    "clusterVIP":{
      "clusterVips":[

    ]
  },
  "nodes":[
    {
      "type":"master",
      "properties":{
        "hostName":"vrops01svr01a.rainpole.local",
        "vmName":"vrops01svr01a",
        "ip":"192.168.11.31",
        "extendedStorage":"1000",
        "network":"vxw-dvs-28-universalwire-3-sid-30002-Mgmt-xRegion01-VXLAN",
        "licenseRef":"locker:license:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:vRealize Suite
2019 Enterprise",
        "useMasterIdentityManager":"true",
        "folderName":"group-v58 (sfo01-m01fd-vrops)"
      }
    },
    {
      "type":"replica",
      "properties":{
        "hostName":"vrops01svr01b.rainpole.local",
        "vmName":"vrops01svr01b",
        "extendedStorage":"1000",
        "ip":"192.168.11.32"
      }
    },
    {
      "type":"data",
      "properties":{
        "hostName":"vrops01svr01c.rainpole.local",
        "vmName":"vrops01svr01c",
        "extendedStorage":"1000",
        "ip":"192.168.11.33"
      }
    },
    {
      "type":"remotecollector",
      "properties":{
        "hostName":"sfo01vropsc01a.sfo01.rainpole.local",
        "vmName":"sfo01vropsc01a",
        "ip":"192.168.31.31",
        "dns":"172.16.11.5,172.16.11.4",
        "network":"vxw-dvs-28-universalwire-1-sid-30000-Mgmt-RegionA01-VXLAN",
        "domain":"sfo01.rainpole.local",
        "folderName":"group-v52 (sfo01-m01fd-vropsrc)",
        "gateway":"192.168.31.1",
        "searchpath":"sfo01.rainpole.local",
        "deployOption":"smallrc"
      }
    }
  ]
}

```

```

    },
    {
      "type": "remotecollector",
      "properties": {
        "hostName": "sfo01vropsc01b.sfo01.rainpole.local",
        "vmName": "sfo01vropsc01b",
        "ip": "192.168.31.32",
        "dns": "172.16.11.5,172.16.11.4",
        "network": "vxw-dvs-28-universalwire-1-sid-30000-Mgmt-RegionA01-VXLAN",
        "domain": "sfo01.rainpole.local",
        "folderName": "group-v52 (sfo01-m01fd-vropsrc)",
        "gateway": "192.168.31.1",
        "searchpath": "sfo01.rainpole.local",
        "deployOption": "smallrc"
      }
    },
    {
      "type": "remotecollector",
      "properties": {
        "hostName": "lax01vropsc01b.lax01.rainpole.local",
        "cluster": "lax01-m01-mgmt01",
        "vmName": "lax01vropsc01b",
        "ip": "192.168.32.32",
        "dns": "172.17.11.5,172.17.11.4",
        "vCenterHost": "lax01m01vc01.lax01.rainpole.local",
        "storage": "lax01-m01-vsan",
        "network": "vxw-dvs-28-universalwire-4-sid-30003-Mgmt-RegionB01-VXLAN",
        "netmask": "255.255.255.0",
        "domain": "lax01.rainpole.local",
        "folderName": "group-v55 (lax01-m01fd-vropsrc)",
        "gateway": "192.168.32.1",
        "searchpath": "lax01.rainpole.local",
        "deployOption": "smallrc"
      }
    },
    {
      "type": "remotecollector",
      "properties": {
        "hostName": "lax01vropsc01a.lax01.rainpole.local",
        "cluster": "lax01-m01-mgmt01",
        "vmName": "lax01vropsc01a",
        "ip": "192.168.32.31",
        "dns": "172.17.11.5,172.17.11.4",
        "vCenterHost": "lax01m01vc01.lax01.rainpole.local",
        "storage": "lax01-m01-vsan",
        "network": "vxw-dvs-28-universalwire-4-sid-30003-Mgmt-RegionB01-VXLAN",
        "netmask": "255.255.255.0",
        "domain": "lax01.rainpole.local",
        "folderName": "group-v55 (lax01-m01fd-vropsrc)",
        "gateway": "192.168.32.1",
        "searchpath": "lax01.rainpole.local",
        "deployOption": "smallrc"
      }
    }
  ],

```

```

    "properties":{
      "disableTls":"TLSv1,TLSv1.1",
      "vropsAdminPassword":"locker:password:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:admin-
vrops01svr01",
      "certificateChain":"locker:certificate:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:vrops01svr01",
      "deployOption":"medium",
      "licenseRef":"locker:license:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:vRealize Suite 2019
Enterprise"
    }
  },
  {
    "id":"vra",
    "version":"8.0.1",
    "clusterVIP":{
      "clusterVips":[
        {
          "type":"vra-va",
          "properties":{
            "hostName":"vra01svr01.rainpole.local\n"
          }
        }
      ]
    },
    "nodes":[
      {
        "type":"vrava-primary",
        "properties":{
          "hostName":"vra01svr01a.rainpole.local",
          "vmName":"vra01svr01a",
          "ip":"192.168.11.51",
          "folderName":"group-v51 (sfo01-m01fd-vra)",
          "rootPassword":"locker:password:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:root-
vra01svr01"
        }
      },
      {
        "type":"vrava-secondary",
        "properties":{
          "hostName":"vra01svr01b.rainpole.local",
          "vmName":"vra01svr01b",
          "ip":"192.168.11.52",
          "folderName":"group-v51 (sfo01-m01fd-vra)",
          "rootPassword":"locker:password:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:root-
vra01svr01"
        }
      },
      {
        "type":"vrava-secondary",
        "properties":{
          "hostName":"vra01svr01c.rainpole.local",
          "vmName":"vra01svr01c",
          "ip":"192.168.11.53",
          "folderName":"group-v51 (sfo01-m01fd-vra)",
          "rootPassword":"locker:password:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:root-
vra01svr01"
        }
      }
    ]
  }
}

```

```

    }
  }
],
"properties":{
  "certificateChain":"locker:certificate:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:vra01svr01",
  "licenseRef":"locker:license:xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx:vRealize Suite 2019
Enterprise"
}
}
]
}

```

Data Protection and Backup Design for vRealize Suite Lifecycle Manager

To preserve the cloud operations services functionality when data or system loss occurs, the design supports the use of data protection.

vRealize Suite Lifecycle Manager supports data protection through the creation of consistent image-level backups, using backup software that is based on the vSphere Storage APIs - Data Protection (VADP).

Disaster Recovery Design for vRealize Suite Lifecycle Manager

To preserve the cloud operations services functionality when a disaster occurs, this design supports the failover of vRealize Suite Lifecycle Manager between regions.

You place vRealize Suite Lifecycle Manager on the cross-region application virtual network, Mgmt-xRegion01-VXLAN. As a result, after recovery, you continue using the same IP address, DNS records, and routing configuration. vRealize Automation and vRealize Operations also use this network for their cross-region failover capabilities.

If a planned migration or disaster occurs, you use Site Recovery Manager and vSphere Replication for an orchestrated recovery of the vRealize Suite Lifecycle Manager appliance. After the recovery, vRealize Suite Lifecycle Manager continues to provide cloud operations services functionality to manage the deployment of the available environments.

vRealize Operations Manager Design

The deployment of vRealize Operations Manager is a single instance of a three-node analytics cluster that is deployed in the primary region of the SDDC, and a two-node remote collector group in each region. The components run in the management domain in each region.

■ [Logical Design of vRealize Operations Manager](#)

vRealize Operations Manager communicates with management components in all regions of the SDDC to collect metrics which are presented through various dashboards and views.

■ [Configuration Design of vRealize Operations Manager](#)

Configuration design details the design decisions covering physical design and sizing for vRealize Operations Manager.

- **Life Cycle Management Design of vRealize Operations Manager**

The life cycle management design details the design decisions covering the life cycle management of vRealize Operations Manager.

- **Network Design of vRealize Operations Manager**

For secure access to the UI and API and for failover of vRealize Operations Manager, you place the appliance in the shared cross-region application virtual network. You provide isolation of the vRealize Operations Manager nodes by placing them in several network segments. This network design also supports public access to the analytics cluster nodes.

- **Information Security and Access Design for vRealize Operations Manager**

You protect the vRealize Operations Manager deployment by configuring authentication and secure communication with the other components in the SDDC. A dedicated service account is assigned a custom role for communication between vRealize Operations Manager and the management solutions in the data center.

- **Monitoring and Alerting Design in vRealize Operations Manager**

You use vRealize Operations Manager to monitor the state of the management components in the SDDC by using dashboards. You can use the self-monitoring capability of vRealize Operations Manager to receive alerts about issues that are related to its operational state.

- **Data Protection and Backup Design for vRealize Operations Manager**

To preserve the cloud operations services functionality when data or system loss occurs, the design supports the use of data protection.

- **Disaster Recovery Design for vRealize Operations Manager**

To preserve the cloud operations services functionality when a disaster occurs, this design supports the failover of vRealize Operations Manager between regions.

Logical Design of vRealize Operations Manager

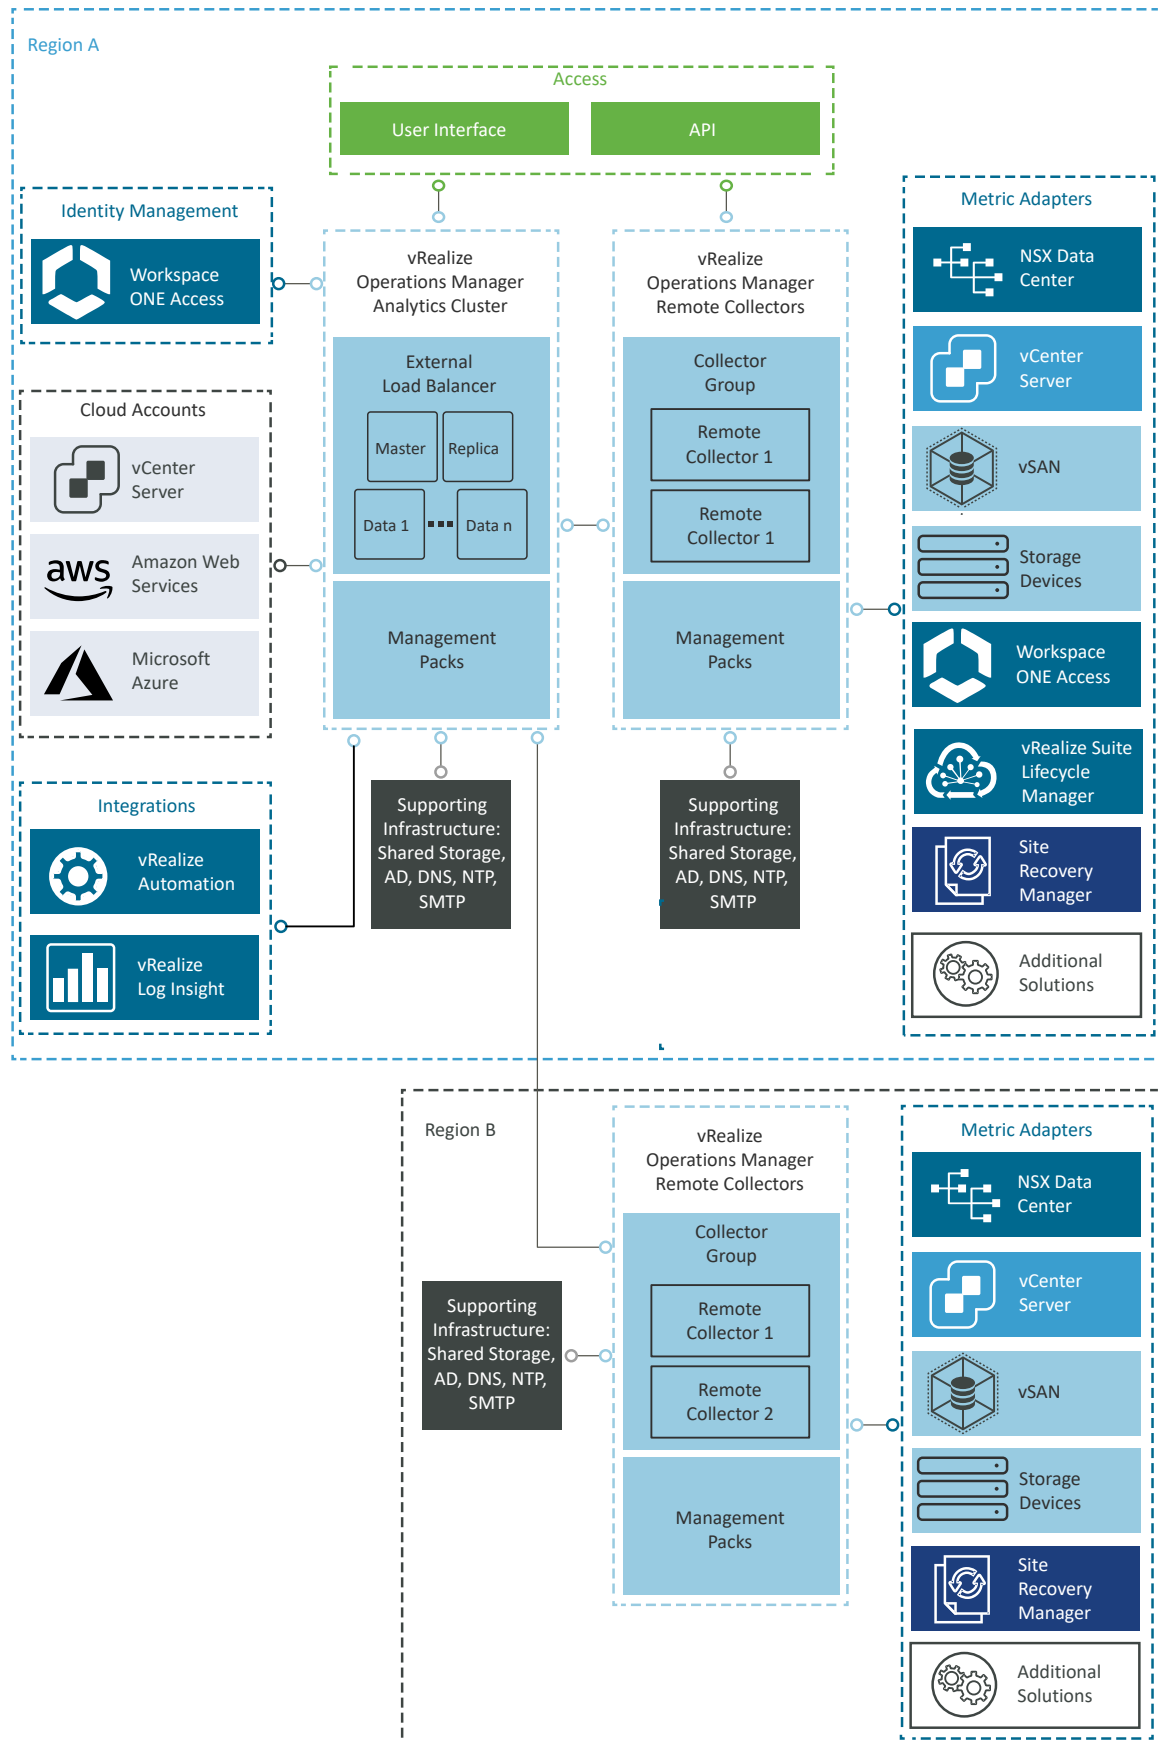
vRealize Operations Manager communicates with management components in all regions of the SDDC to collect metrics which are presented through various dashboards and views.

Logical Design

In a multi-region SDDC, you deploy a vRealize Operations Manager configuration that consists of the following entities.

- A three-node medium-size vRealize Operations Manager analytics cluster that is highly available (HA). This topology provides high availability, scale-out capacity up to 16 nodes, and failover.
- A group of two remote collector nodes in each region. The remote collectors communicate directly with the data nodes in the vRealize Operations Manager analytics cluster. Use two remote collectors in each region for load balancing and fault tolerance.
- Each region contains its own pair of remote collectors whose role is to ease scalability by performing the data collection from the applications that are not subject to failover, and periodically sending collected data to the analytics cluster.

Figure 2-15. Logical Design of a Multi-Region Deployment of vRealize Operations Manager



Configuration Design of vRealize Operations Manager

Configuration design details the design decisions covering physical design and sizing for vRealize Operations Manager.

Deployment Design of vRealize Operations Manager

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 in the design objectives of this design.

Deploy a three-node vRealize Operations Manager analytics cluster in the cross-region application virtual network. The analytics cluster consists of one master node, one master replica node, and one data node to enable scale out and high availability.

To accomplish this design objective, you deploy or reuse the following components to deploy this operations management solution for the SDDC.

- vRealize Suite Lifecycle Manager
- NSX for vSphere Load Balancer for vRealize Operations Manager
- Workspace ONE Access cluster
- Supporting infrastructure services, such as Active Directory, DNS, and NTP.

You place the vRealize Operations Manager on a specific application virtual network for isolation and failover.

vRealize Operations Manager is distributed as a virtual appliance in OVA format.

In the design, you deploy the vRealize Operations Manager analytics cluster nodes on the first vSphere cluster in the management domain of Region A. You deploy the vRealize Operations Manager remote collector nodes on the first vSphere cluster in the management domain of each region. With this configuration, you can centrally manage monitoring across the entire SDDC. The SDDC can comprise multiple regions and multiple availability zones.

Sizing Compute and Storage Resources for vRealize Operations Manager

You size resources for vRealize Operations Manager to provide enough resources to accommodate the analytics operations for monitoring the SDDC and the expected number of virtual machines in the SDDC.

Sizing Resources for the Analytics Cluster Nodes

Deploying three medium-size nodes satisfies the requirement for retention and for monitoring the expected number of objects and metrics for-dual-region environments based on the following assumptions.

Table 2-84. Dual-Region SDDC - Three Nodes

Number of Virtual Machines	Expected Number of Data Center Objects
10,000	12,500

As the environment expands, deploy additional data nodes to accommodate the higher expected number of objects and metrics. For detailed sizing and scaling guidance, you can use the VMware vRealize Operations Manager sizing guides at <http://vropssizer.vmware.com>.

Before deploying additional vRealize Operations Manager data nodes, you must ensure additional ESXi hosts are added to the first cluster in the management domain, so that you guarantee that the vSphere cluster has enough capacity to host additional data nodes, without violating the vSphere DRS anti-affinity rules.

You allocate storage capacity for analytics data collected from the management products and from the number of tenant virtual machines that is defined in the objectives of this SDDC design.

This design uses medium-size nodes for the analytics cluster and standard-size nodes for the remote collector group. To collect the required number of metrics, you must add a virtual disk with the size of 1 TB to each analytics cluster node.

Table 2-85. vRealize Operations Manager Analytics Cluster CPU, Memory, and Storage Resources

Attribute	Per Appliance	Cluster Deployment
Appliance size	Medium	-
vCPUs	8	24
Memory	32 GB	96 GB
Initial Storage	274 GB	822 GB
Additional Storage	1 TB	3 TB

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 region does not increase the capacity for monitored objects.

Table 2-86. vRealize Operations Manager Remote Collector CPU, Memory, and Storage Resources

Attribute	Appliance
Appliance size	Medium
CPU	8 vCPUs
Memory	16 GB
Additional Storage	N/A

Table 2-87. Design Decisions for vRealize Operations Manager Deployment

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-001	Deploy vRealize Operations Manager as a cluster of three nodes: one master, one master replica, and one data node in the first vSphere cluster in the management domain in Region A. Use this deployment to monitor all regions.	<ul style="list-style-type: none"> ■ Provides the scale capacity required for monitoring up to 10,000 virtual machines. ■ Supports scale-up with additional data nodes. 	You must identically size all nodes which increases the resource requirements in the SDDC.
SDDC-OPS-MON-002	Deploy two remote collector nodes in the first vSphere cluster in the management domain in each region.	Removes the load from the analytics cluster from collecting metrics from applications that do not fail over between regions.	You must assign a collector group when configuring the monitoring of a solution.
SDDC-OPS-MON-003	Deploy vRealize Operations Manager by using vRealize Suite Lifecycle Manager.	Allows vRealize Suite Lifecycle Manager the ability to provide life cycle management of vRealize Operations Manager.	None
SDDC-OPS-MON-004	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules to the vRealize Operations Manager analytics cluster.	Using vSphere DRS prevents the vRealize Operations Manager analytics cluster nodes from running on the same ESXi host and risking the high availability of the cluster.	<ul style="list-style-type: none"> ■ You must perform an additional configuration to set up an anti-affinity rule. ■ You must update the anti-affinity rule if additional data nodes are added. ■ You can put in maintenance mode only a single ESXi host at a time in a management cluster of four ESXi hosts.
SDDC-OPS-MON-005	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules to the vRealize Operations Manager remote collector group.	Using vSphere DRS prevents the vRealize Operations Manager remote collector nodes from running on the same ESXi host and risking the high availability of the cluster.	You must perform an additional configuration to set up an anti-affinity rule.

Table 2-87. Design Decisions for vRealize Operations Manager Deployment (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-006	Deploy each node in the analytics cluster as a medium-size appliance.	<ul style="list-style-type: none"> ■ If you use fewer large-size vRealize Operations Manager nodes, you must increase the minimum host memory size to handle the increased performance that is the result from stretching NUMA node boundaries. ■ Provides enough capacity for the metrics and objects generated by up to 12,500 objects while having high availability in the analytics cluster enabled. Metrics are collected from the following components: <ul style="list-style-type: none"> ■ vCenter Server instances ■ ESXi Hosts ■ NSX-T Data Center Components ■ vRealize Automation ■ vRealize Log Insight ■ Storage Array and data center infrastructure 	<p>ESXi hosts in the management cluster must have physical CPUs with a minimum of 8 cores per socket. In total, vRealize Operations Manager uses 24 vCPUs and 96 GB of memory in the management cluster.</p> <p>You must deploy additional nodes once you exceed 12,500 objects.</p>
SDDC-OPS-MON-007	Add more medium size nodes to the analytics cluster if the number of SDDC objects exceeds 12,500.	Ensures that the analytics cluster has enough capacity to meet the SDDC object and metric growth.	<ul style="list-style-type: none"> ■ The capacity of the physical ESXi hosts must be enough to accommodate virtual machines that require 32 GB RAM without bridging NUMA node boundaries. ■ The management cluster must have enough ESXi hosts so that vRealize Operations Manager can run according to vSphere DRS anti-affinity rules. ■ The number of nodes must not exceed number of ESXi hosts in the management cluster - 1. <p>For example, if the management cluster contains six ESXi hosts, you can deploy up to five vRealize Operations Manager nodes in the analytics cluster.</p>

Table 2-87. Design Decisions for vRealize Operations Manager Deployment (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-008	Increase the initial storage of each vRealize Operations Manager analytics cluster virtual appliance by 1 TB.	Supports the storage requirements for monitoring up to 12,500 objects.	None
SDDC-OPS-MON-009	Deploy the medium-size vRealize Operations Manager remote collector virtual appliances.	<ul style="list-style-type: none"> ■ Enables metric collection for the expected number of objects in the SDDC when at full capacity. ■ Remote collectors do not perform analytics operations or store data on disk, therefore no additional storage is required. 	You must provide 4 vCPUs and 8 GB of memory in the management cluster in each region.
SDDC-OPS-MON-010	When using two availability zones in Region A, add the vRealize Operations Manager virtual appliances to the primary availability zone VM group, for example, <code>sfo01-m01-mgmt01-primary-az-vm-group</code> .	Ensures the vRealize Operations Manager virtual appliance is powered on within the primary availability zone hosts group by default.	If vRealize Operations Manager is deployed after the creation of the stretched cluster for management domain availability zones, the VM Group for the primary availability zone virtual machines must be updated to include the vRealize Operations Manager virtual appliances.
SDDC-OPS-MON-011	Place the cross-region vRealize Operations Manager virtual appliances in a dedicated virtual machine folder in Region A, for example, <code>xregion-sfo01-lax01-m01fd-vrops</code> .	Provides an organization of the vRealize Operations Manager virtual appliance in the management domain inventory and a preparation for Site Recovery Manager folder mappings for disaster recovery.	A corresponding virtual machine folder in Region B must be created in preparation for Site Recovery Manager folder mapping, for example, <code>xregion-lax01-sfo01-m01fd-vrops</code> .

Logging Design for vRealize Operations Manager

You integrate vRealize Operations Manager with vRealize Log Insight to provide operational visibility.

The native integration to vRealize Log Insight from vRealize Operations provides the ability to send logs for aggregation and analysis, as necessary.

Logging to a vRealize Log Insight instance through the ingestion API is established by updated in the appliance settings in the vRealize Operations Manager user interface or by updating the vRealize Log Insight `liagent.ini`.

For more information, see the vRealize Log Insight section of the design.

Table 2-88. Design Decisions on Logging for vRealize Operations Manager

Design ID	Design Decision	Design Justification	Design Justification
SDDC-OPS-MON-012	Configure vRealize Operations Manager to send logs to the vRealize Log Insight cluster in Region A.	It allows logs from vRealize Operations Manager to be forwarded to a vRealize Log Insight cluster.	You must configure vRealize Operations Manager to send logs to the vRealize Log Insight cluster in Region A.
SDDC-OPS-MON-013	Communicate with the vRealize Log Insight using the default Ingestion API (cfapi) port 9000 and non-default No SSL.	Supports disaster recovery of vRealize Operations Manager in the SDDC.	Transmission traffic for logs is not secure.

For more information, see [Integration of vRealize Log Insight with vRealize Operations Manager and Workspace ONE Access](#).

Notifications Design for vRealize Operations Manager

You configure notifications in vRealize Operations Manager to send event messages to the user. You configure a Standard Email Plug-in to deliver outbound SMTP messages to users about system events.

Table 2-89. Design Decisions on Notifications for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-014	Configure vRealize Operations Manager to use an outbound SMTP mail server to route notifications for system events.	Integrates vRealize Operations Manager system events notifications to users by email to provide an enhanced user experience.	None

Costing Design for vRealize Operations Manager

To enable accurate costing in the correct currency, you set the global currency option in vRealize Operations Manager. vRealize Automation integration uses this currency setting when calculating costing.

Table 2-90. Design Decisions on Costing for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-015	Configure the correct currency in the vRealize Operations Manager global options.	Ensures accurate costing in the correct currency.	The currency cannot be changed after it is set.

Life Cycle Management Design of vRealize Operations Manager

The life cycle management design details the design decisions covering the life cycle management of vRealize Operations Manager.

The life cycle management of vRealize Operations Manager involves the process of performing patch updates or upgrades to the vRealize Operations Manager analytics cluster and remote collector nodes.

In this design, the life cycle management is performed using vRealize Suite Lifecycle Manager.

Table 2-91. Design Decisions on the Life Cycle Management of vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-016	Use vRealize Suite Lifecycle Manager to perform the life cycle management of vRealize Operations Manager.	vRealize Suite Lifecycle Manager automates the life cycle of vRealize Operations Manager.	You must deploy vRealize Suite Lifecycle Manager.

Network Design of vRealize Operations Manager

For secure access to the UI and API and for failover of vRealize Operations Manager, you place the appliance in the shared cross-region application virtual network. You provide isolation of the vRealize Operations Manager nodes by placing them in several network segments. This network design also supports public access to the analytics cluster nodes.

For secure access, load balancing, and portability, you deploy the vRealize Operations Manager analytics cluster in the shared cross-region application virtual network and the remote collector groups in the region-specific application virtual networks.

Application Virtual Network

The vRealize Operations Manager analytics cluster virtual appliances are connected to the cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN, for secure access to the application UI and API, and for failover support.

The vRealize Operations Manager remote collector virtual appliances are connected to the region-specific application virtual networks, for example, Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN, for collection of metrics locally per region.

This networking design has the following features:

- The vRealize Operations Manager analytics cluster can be failed over between regions if there is a planned migration or disaster recovery without changing any IP address, DNS records, or routing configurations. Workspace ONE Access, vRealize Automation, and vRealize Suite Lifecycle Manager also share this network for cross-region failover support.
- The vRealize Operations Manager remote collector groups are deployed together on the same network in each region. This configuration ensures collection of metrics locally per region in the event of a cross-region network outage. It also co-locates metric collection with the region-specific applications using the region-specific application virtual network.
- vRealize Operations Manager has routed access to the VLAN-backed management network through the NSX Universal Distributed Logical Router.
- Routing to the VLAN-management network, application virtual networks, and external networks are dynamic and are based on the Border Gateway Protocol (BGP).

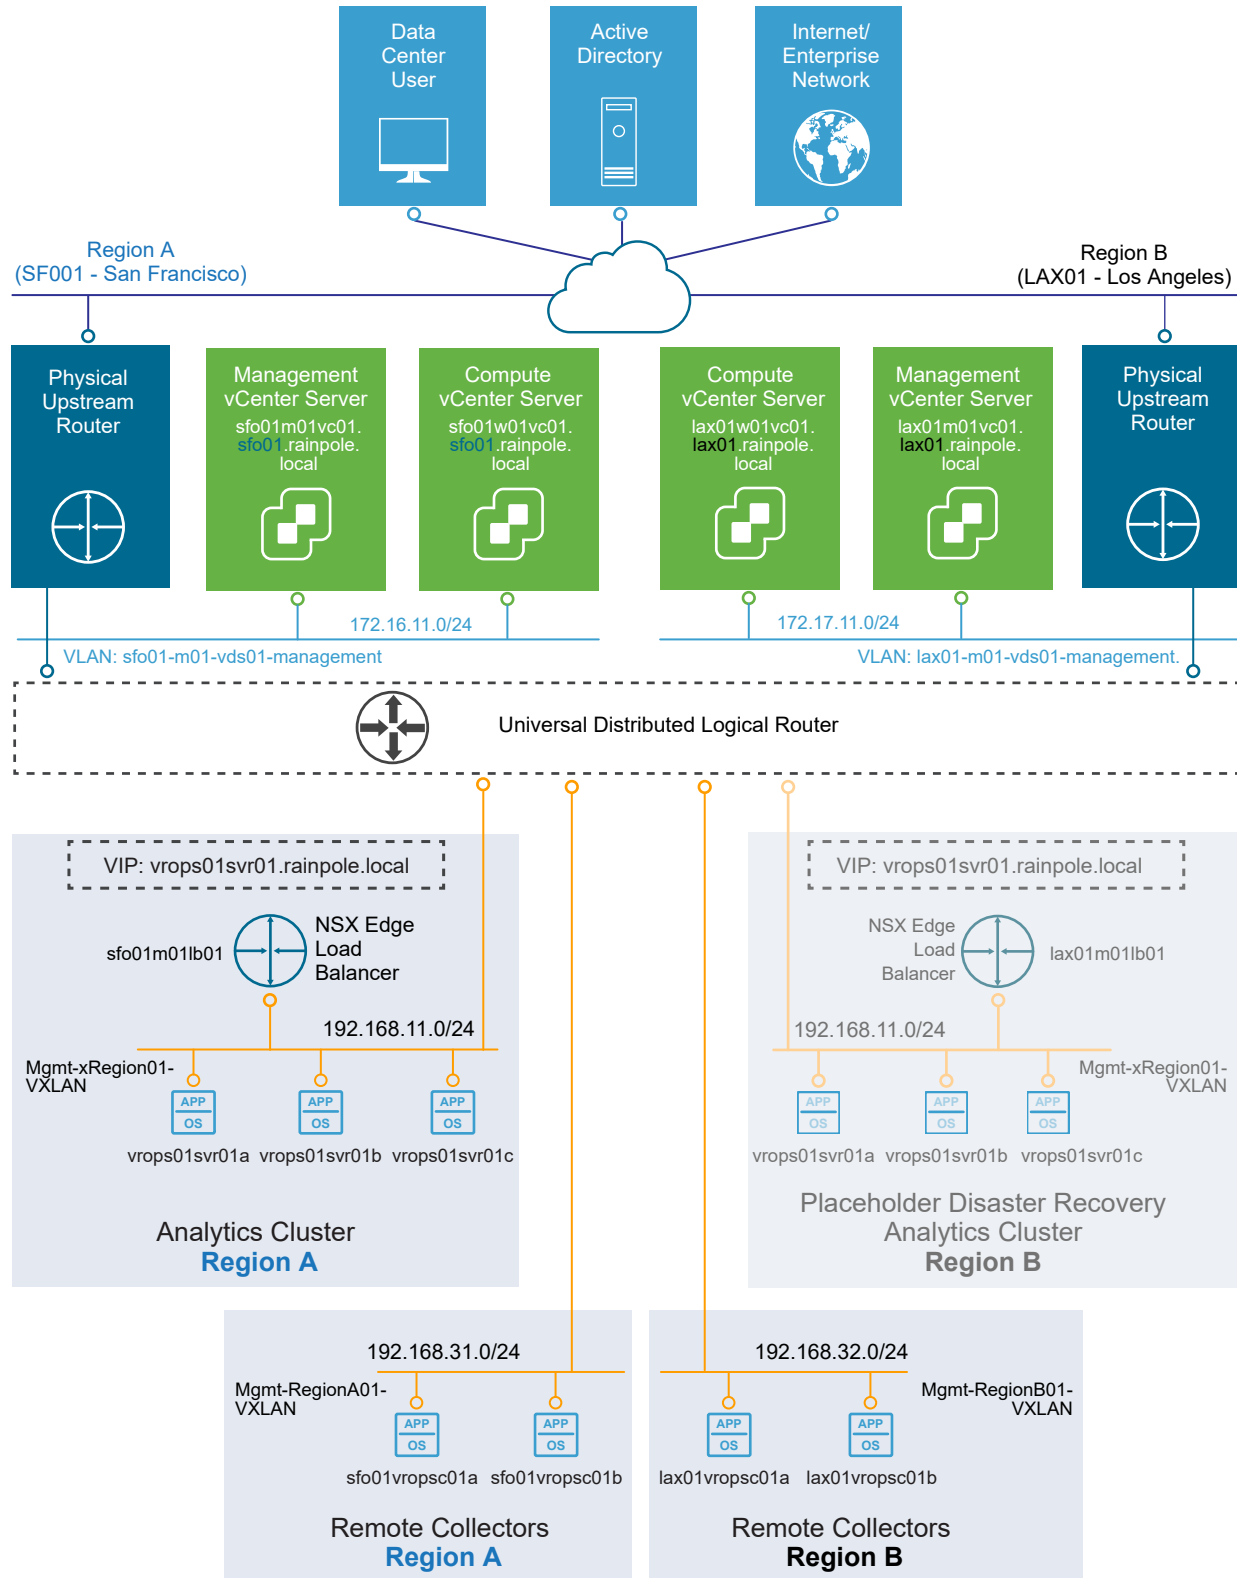
Figure 2-16. Networking Design of the vRealize Operations Manager Deployment

Table 2-92. Design Decisions on the Application Virtual Network for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-017	Place the vRealize Operations Manager analytics appliances on the cross-region application virtual network, for example, Mgmt-xRegion01-VXLAN.	Supports secure access from an external location and disaster recovery.	You must use an implementation in NSX for vSphere to support this networking configuration.
SDDC-OPS-MON-018	Place the vRealize Operations Manager remote collector appliances on the region-specific application virtual network, for example, Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN.	Supports collection of metrics locally per region.	You must use an implementation in NSX for vSphere to support this networking configuration.

IP Addressing Scheme

You allocate a subnet for the cross-region network segment, and the region-specific network segments in the management domain and use them for the vRealize Operations Manager deployment.

Table 2-93. Example IP Subnets for vRealize Operations Manager

Solution	IP Subnet	Gateway	NSX Application Virtual Network
Analytics Cluster in Region A	192.168.11.0/24	192.168.11.1	Mgmt-xRegion01-VXLAN
Remote collectors in Region A	192.168.31.0/24	192.168.31.1	Mgmt-RegionA01-VXLAN
Remote collectors in Region B	192.168.32.0/24	192.168.32.1	Mgmt-RegionB01-VXLAN

Table 2-94. Design Decisions on the IP Subnets for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-019	Allocate separate subnets for each application virtual network.	Placing the remote collectors on their own subnet enables them to communicate with the analytics cluster and not be a part of the failover group.	None

Name Resolution

The FQDNs of the vRealize Operations Manager nodes follow a certain domain name resolution:

- The IP addresses of the analytics cluster nodes and a load balancer virtual IP address (VIP) are associated with names whose suffix is set to the root domain `rainpole.local`.

From the public network, users access vRealize Operations Manager using the VIP address, the traffic to which is handled by an NSX for vSphere edge providing the load balancer services.

- Name resolution for the IP addresses of the remote collector group nodes uses a region-specific suffix, for example, `sfo01.rainpole.local` or `lax01.rainpole.local`.
- The IP addresses of the remote collector group nodes are associated with names whose suffix is set to the region-specific domain, for example, `sfo01.rainpole.local` or `lax01.rainpole.local`.

Table 2-95. Example FQDNs and IP Addresses for vRealize Operations Manager

Fully Qualified Domain Name	IP Address	Description	Region	Failed Over to Region B
vrops01svr01.rainpole.local	192.168.11.30	NSX load balancer Virtual IP of the analytics cluster	Region A	✓
vrops01svr01a.rainpole.local	192.168.11.31	Master node in the analytics cluster	Region A	✓
vrops01svr01b.rainpole.local	192.168.11.32	Master replica node in the analytics cluster	Region A	✓
vrops01svr01c.rainpole.local	192.168.11.33	First data node in the analytics cluster	Region A	✓
vrops01svr01x.rainpole.local	192.168.11. <i>n</i>	Additional data nodes in the analytics cluster	Region A	✓
sfo01vropsc01a.sfo01.rainpole.local	192.168.31.31	First remote collector node	Region A	x
sfo01vropsc01b.sfo01.rainpole.local	192.168.31.32	Second remote collector node	Region A	x
lax01vropsc01a.lax01.rainpole.local	192.168.32.31	First remote collector node	Region B	x
lax01vropsc01b.lax01.rainpole.local	192.168.32.32	Second remote collector node	Region B	x

Table 2-96. Design Decisions on Name Resolution for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-020	Configure forward and reverse DNS records for all vRealize Operations Manager nodes and the VIP address.	All nodes are accessible by using fully qualified domain names instead of by using IP addresses only.	You must provide DNS records for the vRealize Operations Manager appliances.
SDDC-OPS-MON-021	In a multi-region deployment, configure the DNS settings for the vRealize Operations Manager analytics cluster appliances to use DNS servers in each region.	vRealize Operations Manager can resolve DNS from regional DNS servers during a planned migration or disaster recovery between regions.	As you scale from a single-region to multi-region deployment, the DNS settings for the vRealize Operations Manager appliance must be updated.

Time Synchronization

Time synchronization provided by the Network Time Protocol (NTP) is important to ensure that all components within the SDDC are synchronized to the same time source. Configure the vRealize Operations Manager virtual appliances with time synchronization using an internal NTP time source.

Table 2-97. Design Decisions on Time Synchronization for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-022	Configure NTP on each vRealize Operations Manager appliance.	vRealize Operations Manager is dependent on time synchronization.	None
SDDC-OPS-MON-023	In a multi-region deployment, configure the NTP settings for the vRealize Operations Manager analytics cluster appliances to use NTP servers in each region.	vRealize Operations Manager can query NTP from regional NTP servers to synchronize time during a planned migration or disaster recovery between regions.	As you scale from a single-region to multi-region deployment, the NTP settings on the vRealize Operations Manager analytics cluster appliances must be updated.
SDDC-OPS-MON-024	Configure the timezone of vRealize Operations Manager to use UTC.	You must use UTC to enable the integration with vRealize Automation as vRealize Automation only supports UTC.	If you are in a timezone other than UTC, timestamps appear skewed.

Load Balancing

A vRealize Operations Manager cluster deployment requires a load balancer to manage connections to vRealize Operations Manager. The design uses load balancing services provided by NSX for vSphere in the management domain.

Table 2-98. Design Decisions on Load Balancing for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-025	Configure the load balancer that was created in NSX for vSphere in the management domain to load balance Workspace ONE Access to also load balance connections across vRealize Operations Manager analytics cluster members.	Required to deploy vRealize Operations Manager analytics cluster deployment type with distributed user interface access across members.	You must use an implementation in NSX for vSphere to support this network configuration.
SDDC-OPS-MON-026	<ul style="list-style-type: none"> ■ Add an NSX load balancer monitor, for example, vrops-https-monitor, for vRealize Operations Manager with an active HTTPS monitor on monitoring port 443. ■ Set the intervals and timeouts for the monitor: <ul style="list-style-type: none"> ■ Interval: 5 seconds ■ Timeout: 16 seconds ■ Max Retries: 3 seconds ■ Set the HTTP Request for the monitor: <ul style="list-style-type: none"> ■ HTTP Method: Get ■ Request URL: /suite-api/api/deployment/node/status?service=api&service=admin&service=ui ■ Set the HTTP Response for the monitor: <ul style="list-style-type: none"> ■ HTTP Response Code: 200, 204, 301 ■ HTTP Response Body: ONLINE 	<ul style="list-style-type: none"> ■ The vRealize Operations Manager health check is provided over HTTPS on port 443. ■ The Active Monitor uses HTTPS requests to monitor the application health reported by vRealize Operations Manager. ■ Ensures that connections to unhealthy vRealize Operations Manager analytics cluster members in the pool are disabled until a subsequent periodic health check finds the members to be healthy. 	None

Table 2-98. Design Decisions on Load Balancing for vRealize Operations Manager (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-027	<ul style="list-style-type: none"> ■ Add an NSX load balancer server pool, for example, <code>vrops-server-pool</code>, for vRealize Operations Manager to use the LEASTCONN algorithm. ■ Set the static members for the pool: <ul style="list-style-type: none"> ■ Name: <i>host name</i> ■ IP: <i>IP address</i> ■ Port: 443 ■ Weight: 1 ■ State: Enabled ■ Max Concurrent Connections: 10 	<ul style="list-style-type: none"> ■ Least Connection distributes requests to members based on the number of current connections. New connections are sent to the vRealize Operations Manager analytics cluster pool member with the fewest connections. ■ vRealize Operations Manager analytics cluster services respond on TCP 443. ■ Each vRealize Operations Manager analytics cluster node can accept up to a maximum of 10 concurrent user interface connections. 	When the level of connections exceeds the maximum number that the vRealize Operations Manager analytics cluster can accept, connections may be dropped.
SDDC-OPS-MON-028	<ul style="list-style-type: none"> ■ Add an NSX load balancer fast TCP application profile, for example <code>vrops-tcp-app-profile</code>, for vRealize Operations Manager. ■ Set the application profile type to SSL Passthrough. Set the persistence to None ■ Set the timeout to 1800 seconds (30 minutes). 	An application profile must set the required timeout for HTTPS requests to vRealize Operations Manager.	None
SDDC-OPS-MON-029	<ul style="list-style-type: none"> ■ Add an NSX load balancer Source IP persistence profile, for example <code>vrops-source-ip-persistence-profile</code>, for vRealize Operations Manager. ■ Set the timeout to 1800 seconds (30 minutes). 	A source IP persistence profile is required for the vRealize Operations Manager analytics cluster user interface access.	None

Table 2-98. Design Decisions on Load Balancing for vRealize Operations Manager (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-030	<ul style="list-style-type: none"> ■ Add an NSX load balancer virtual server, for example, <code>vrops-https</code>, for vRealize Operations Manager to use the L4 TCP type and port 443. ■ Set the acceleration to Enable. ■ Set the IP address for the virtual server. ■ Set the Source IP for the virtual server, for example, <code>vrops-source-ip-persistence-profile</code>. ■ Set the application profile, for example, <code>vrops-https-app-profile</code>. ■ Set the DefaultPool to use the vRealize Operations Manager server pool, for example, <code>vrops-server-pool</code>. 	The virtual server receives all the client connections and distributes them among the vRealize Operations analytics cluster pool members based on the state of those pool members.	None
SDDC-OPS-MON-031	<ul style="list-style-type: none"> ■ Add an NSX load balancer HTTP application profile, for example, <code>vrops-http-app-profile-redirect</code>, for vRealize Operations Manager to redirect HTTP to HTTPS. ■ Set the timeout to 1800 seconds (30 minutes). ■ Set redirection to HTTP to HTTPS Redirect. 	Ensures that connections to non-secure HTTP are automatically redirected to HTTPS for vRealize Operations Manager.	None

Table 2-98. Design Decisions on Load Balancing for vRealize Operations Manager (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-032	<ul style="list-style-type: none"> ■ Add another NSX load balancer virtual server, for example, vrops-http-redirect, for vRealize Operations Manager HTTP to HTTPS redirection to use the L7 HTTP type and port 80. ■ Set the acceleration to Disable. ■ Set the IP address for the Load Balancer to the same IP address used for the HTTPS virtual server, for example, vrops-https. ■ Set the application profile to the HTTP to HTTPS Redirect profile, for example, vrops-http-app-profile-redirect. 	Ensures that connections to non-secure HTTP are automatically redirected to HTTPS for vRealize Operations Manager.	None
SDDC-OPS-MON-033	Do not use a load balancer for the vRealize Operations Manager remote collector nodes.	<ul style="list-style-type: none"> ■ vRealize Operations Manager remote collector nodes must directly access the systems that they are monitoring. ■ vRealize Operations Manager remote collector nodes do not require access to and from the public network. 	None

Information Security and Access Design for vRealize Operations Manager

You protect the vRealize Operations Manager deployment by configuring authentication and secure communication with the other components in the SDDC. A dedicated service account is assigned a custom role for communication between vRealize Operations Manager and the management solutions in the data center.

This design incorporates the NIST 800-53 standard for password policies as a baseline as follows.

Table 2-99. Password Policy Settings

Setting	Value
Minimum Length	15
Maximum lifetime	60 days

Table 2-99. Password Policy Settings (continued)

Setting	Value
Complexity	At least one upper case, one lower case, one number, and one special char
Maximum failed login attempts	3

Authentication and Authorization Design for vRealize Operations Manager

Users can authenticate to vRealize Operations Manager by using the following account types:

Table 2-100. vRealize Operations Manager Account Types

Account Type	Description
Imported from an LDAP database	Users can use their LDAP credentials to log in to vRealize Operations Manager.
Integrated with Workspace ONE Access	Specified users and groups from upstream identity sources are synchronized to vRealize Operations Manager through Workspace ONE Access.
vCenter Server user accounts	<p>After a vCenter Server instance is registered with vRealize Operations Manager, the following vCenter Server users can log in to vRealize Operations Manager:</p> <ul style="list-style-type: none"> ■ Users that have administration access in vCenter Server. ■ Users that have one of the vRealize Operations Manager privileges, such as PowerUser, assigned to the account which appears at the root level in vCenter Server.
Local user accounts in vRealize Operations Manager	vRealize Operations Manager performs local authentication using the account information stored in its global database.

You enable authentication using Workspace ONE Access to ensure accountability on user access. You can grant both users and groups access to vRealize Operations Manager to perform tasks, such as creating and viewing dashboards.

Cloud Accounts Design for vRealize Operations Manager

You use cloud accounts to add cloud endpoints as adapter instances to enable vRealize Operations Manager to communicate with them. vRealize Operations Manager collects data from the following cloud accounts.

Table 2-101. Cloud Account Types

Cloud Account	Additional Option	Description
vCenter Server	-	Enables vRealize Operations Manager to communicate with vCenter Server.
	vSAN	Enables vRealize Operations Manager to gather vSAN metrics from vCenter Server.
	Application Discovery	Enables vRealize Operations Manager to discover applications running on virtual machines in vCenter Server

Table 2-101. Cloud Account Types (continued)

Cloud Account	Additional Option	Description
Azure	-	Enables vRealize Operations Manager to communicate with a Microsoft Azure Cloud Endpoint.
AWS	-	Enables vRealize Operations Manager to communicate with an Amazon AWS Cloud Endpoint.

Integrations Design for vRealize Operations Manager

vRealize Operations Manager includes direct integrations with vRealize Automation and vRealize Log Insight. These integrations provide the following functionality:

Table 2-102. vRealize Operations Manager Integrations

Integration	Description
vRealize Automation	<ul style="list-style-type: none"> ■ Ability to share common constructs such as cloud accounts, cloud zones, and projects across vRealize Operations Manager and vRealize Automation. ■ Ability to understand the deployment cost: <ul style="list-style-type: none"> ■ Evaluate upfront costs on vRealize Automation. ■ Monitor ongoing costs per virtual machine, deployment, or project.
vRealize Log Insight	<ul style="list-style-type: none"> ■ Enables Logs tab in vRealize Operations Manager ■ Enables Troubleshoot with Logs dashboard ■ Enables the vRealize Log Insight launch in context from vRealize Operations Manager

Management Packs Design for vRealize Operations Manager

You add and configure accounts associated with other solutions by installing and activating management packs. After you have configured the account, vRealize Operations Manager can collect data from or send data to the target system. You install and activate the following management packs.

Table 2-103. Management Packs

Management Pack	Description
Management Pack for NSX-T Data Center	Enables vRealize Operations Manager to communicate with an NSX-T Data Center Endpoint.
Management Pack for Storage Devices	Enables vRealize Operations Manager to communicate with storage devices and arrays.
Management Pack for Site Recovery Manager	Enables vRealize Operations Manager to communicate with VMware Site Recovery Manager Endpoints.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-034	Rotate the root password on or before 365 days post deployment.	The password for the root user account expires 365 days after the initial deployment.	You must manage the password rotation schedule for the root user account in accordance with your organization policies and regulatory standards, as applicable.
SDDC-OPS-MON-035	Enable the vRealize Operations Manager integration with your corporate identity source using Workspace ONE Access.	Allows authentication, including multi-factor, to vRealize Operations Manager using your corporate identity source. Allows authorization through the assignment of organization and cloud services roles to enterprise users and groups defined in your corporate identity source.	You must deploy and configure the Workspace ONE Access to establish the integration between vRealize Operations Manager and your corporate identity sources.
SDDC-OPS-MON-036	Create a security group in your organization directory services for the vRealize Operations Manager administrators role, for example, rainpole.local\ug-vrops-admins , and synchronize the group in the Workspace ONE Access configuration for vRealize Operations Manager.	Allows you to streamline the management of vRealize Operations Manager roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the desired directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS-MON-037	Assign the enterprise group for vRealize Operations Manager administrators, for example, rainpole.local\ug-vrops-admins , the Administrator role.	Provides the following access control features: <ul style="list-style-type: none"> ■ Access to vRealize Operations Manager administration is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking organization owner access to vRealize Operations Manager. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-038	Create a security group in your organization directory services for the vRealize Operations Manager content administrators role, for example, rainpole.local\ug-vrops-content-admins , and synchronize the group in the Workspace ONE Access configuration for vRealize Operations Manager.	Allows you to streamline the management of vRealize Operations Manager roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS-MON-039	Assign the enterprise group for vRealize Operations Manager content administrators, for example, rainpole.local\ug-vrops-content-admins , the ContentAdmin role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to the vRealize Operations Manager user interface is granted to a managed set of individuals that are members of the security group. ■ You can introduce an improved accountability and tracking organization owner access to vRealize Operations Manager. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-OPS-MON-040	Create a security group in your organization directory services for the vRealize Operations Manager read-only users role, for example, rainpole.local\ug-vrops-read-only , and synchronize the group in the Workspace ONE Access configuration for vRealize Operations Manager.	Allows you to streamline the management of vRealize Operations Manager roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-041	Assign the enterprise group for vRealize Operations Manager read-only users, for example, rainpole.local\ug-vrops-read-only , the ReadOnly role.	Provides the following access control features: <ul style="list-style-type: none"> ■ Access to vRealize Operations Manager user interface is granted to a managed set of individuals that are members of the security group. ■ You can introduce improved accountability and tracking organization owner access to vRealize Operations Manager. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-OPS-MON-042	Define a custom vCenter Server role for vRealize Operations Manager that has the minimum privileges required to support collecting metrics and performing actions against vSphere endpoints across the SDDC, for example, vRealize Operations to vSphere Integration – Actions .	vRealize Operations Manager accesses vSphere with the minimum set of permissions that are required to support performing actions against vSphere endpoints across the SDDC.	You must maintain the permissions required by the custom role.
SDDC-OPS-MON-043	Configure a service account in vCenter Server with global permissions, for application-to-application communication from vRealize Operations Manager to vSphere, for example, svc-vrops-vsphere@rainpole.local , and assign the Actions custom role, for example, vRealize Operations to vSphere Integration – Actions .	Provides the following access control features: <ul style="list-style-type: none"> ■ The adapters in vRealize Operations Manager access vSphere with the minimum set of permissions that are required to collect metrics and perform permitted actions. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	<ul style="list-style-type: none"> ■ You must maintain the life cycle and availability of the service account outside of the SDDC stack. ■ All vCenter Server instances must be in the same vSphere domain.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-044	Configure a vCenter Server cloud account for each vCenter Server instance in the SDDC using the vCenter Server service account, for example, svc-vrops-vsphere@rainpole.local .	Enables integration and data collection of all vCenter Server instances in the SDDC in vRealize Operations Manager.	You must manage the password life cycle of this Cloud Account.
SDDC-OPS-MON-045	Configure each vCenter Server cloud account to use the remote collector group for its region.	Components that are not failed over between regions are configured to use the remote collector group. This offloads data collection for local management components from the analytics cluster.	None
SDDC-OPS-MON-046	Define a custom vCenter Server role for vRealize Operations Manager that has the minimum privileges required to support collecting metrics from vSphere endpoints across the SDDC, for example, vRealize Operations to vSphere Integration – Metrics .	vRealize Operations Manager accesses vSphere with the minimum set of permissions that are required to support collecting metrics from vSphere endpoints across the SDDC.	You must maintain the permissions required by the custom role.
SDDC-OPS-MON-047	Configure a service account in vCenter Server with global permissions, for application-to-application communication from the vSAN adapters in vRealize Operations Manager to vSphere, for example, svc-vrops-vsan@rainpole.local , and assign the metrics custom role, for example, vRealize Operations to vSphere Integration – Metrics).	Provides the following access control features: <ul style="list-style-type: none"> ■ The adapters in vRealize Operations Manager access vSphere with the minimum set of permissions that are required to collect metrics about vSAN inventory objects. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	You must maintain the life cycle and availability of the service account outside of the SDDC stack.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-048	Configure the vCenter Server cloud account to enable vSAN integration using the vSAN service account, for example, svc-vrops-vsant@rainpole.local .	Enables integration and data collection of all vSAN in the SDDC in vRealize Operations Manager.	You must manage the password life cycle of this endpoint.
SDDC-OPS-MON-049	Configure a service account in vCenter Server for application-to-application communication from vRealize Operations Manager to NSX Data Center for vSphere.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ The adapters in vRealize Operations Manager access NSX Data Center for vSphere with the minimum set of permissions that are required for metric collection and topology mapping. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	You must maintain the life cycle of the service account outside of the SDDC stack to ensure its availability.
SDDC-OPS-MON-050	Configure a local service account in each NSX Data Center for vSphere instance for application-to-application communication from the NSX vSphere adapters in vRealize Operations Manager to NSX.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ The adapters in vRealize Operations Manager access NSX for vSphere with the minimum set of permissions that are required for metric collection and topology mapping. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	You must maintain the life cycle of the service account outside of the SDDC stack to ensure its availability.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-051	Install and configure the NSX-T management pack for vRealize Operations Manager. Configure the management pack endpoint to use the NSX-T admin account.	Enables integration and data collection of all workload domain NSX-T instances in the SDDC in vRealize Operations Manager.	You must manage the password life cycle of this endpoint.
SDDC-OPS-MON-052	Configure a service account in vCenter Server with global permissions, for application-to-application communication from the storage devices adapters in vRealize Operations Manager to vSphere, for example, svc-vrops-mpsd@rainpole.local , and assign the metrics custom role, for example, vRealize Operations to vSphere Integration – Metrics .	Provides the following access control features: <ul style="list-style-type: none"> ■ The adapters in vRealize Operations Manager access vSphere with the minimum set of permissions that are required to collect metrics about vSphere inventory objects. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	You must maintain the life cycle and availability of the service account outside of the SDDC stack.
SDDC-OPS-MON-053	Install and configure the Storage Devices management pack for vRealize Operations Manager. Configure the management pack endpoint to use the storage devices service account, for example, svc-vrops-mpsd@rainpole.local .	Enables integration and data collection of all storage devices in the SDDC in vRealize Operations Manager.	You must manage the password life cycle of this endpoint.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-054	Configure a service account in vCenter Server with global permissions, for application-to-application communication from the Site Recovery Manager adapters in vRealize Operations Manager to vSphere and Site Recovery Manager, for example, svc-vrops-srm@rainpole.local , and assign the Read Only role.	Provides the following access control features: <ul style="list-style-type: none"> ■ The adapters in vRealize Operations Manager access vSphere and Site Recovery Manager with the minimum set of permissions that are required to collect metrics. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	You must maintain the life cycle and availability of the service account outside of the SDDC stack.
SDDC-OPS-MON-055	Install and configure the Site Recovery Manager management pack for vRealize Operations Manager. Configure the management pack endpoints to use the Site Recovery Manager service account, for example, svc-vrops-srm@rainpole.local .	Enables integration and data collection of all Site Recovery Manager instances in the SDDC in vRealize Operations Manager.	You must manage the password life cycle of this endpoint.
SDDC-OPS-MON-056	Configure the following management pack adapter instances to use the remote collector group: <ul style="list-style-type: none"> ■ NSX for vSphere ■ NSX-T Data Center ■ Storage Devices ■ vSAN ■ Site Recovery Manager 	Components that are not failed over between regions are configured to use the remote collector group. This offloads data collection for local management components from the analytics cluster.	None

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-057	Configure a service account in vRealize Automation for application-to-application communication from the vRealize Automation integration in vRealize Operations Manager to vRealize Automation, for example, svc-vrops-vra@rainpole.local .	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ The integration in vRealize Operations Manager accesses vRealize Automation with the minimum set of permissions that are required for collecting metrics about provisioned virtual machines and capacity management. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the components of the SDDC. 	<ul style="list-style-type: none"> ■ You must maintain the life cycle of the service account outside of the SDDC stack to ensure its availability. ■ If you add more tenants to vRealize Automation, you must maintain the service account permissions to guarantee that metric uptake in vRealize Operations Manager is not compromised.
SDDC-OPS-MON-058	Configure the vRealize Automation integration in vRealize Operations Manager.	<ul style="list-style-type: none"> ■ Provides the ability to share common constructs, such as cloud accounts, cloud zones, and projects across vRealize Operations Manager and vRealize Automation. ■ Provides the ability to understand the deployment cost: <ul style="list-style-type: none"> ■ Evaluate upfront costs on vRealize Automation. ■ Monitor ongoing costs per virtual machine, deployment, or project. 	You must manage the password life cycle of this endpoint.
SDDC-OPS-MON-059	Configure the vRealize Automation integration to use the default collector group.	Components that are failed over between regions are configured to use the default collector group. This configuration provides monitoring of components during a failover.	The load on the analytics cluster, though minimal, increases.

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-060	Configure a service account in Workspace ONE Access for application-to-application communication from the Workspace ONE Access adapter in vRealize Operations Manager to Workspace ONE Access, for example, svc-vrops-wsa@rainpole.local .	The service account is used for application-to-application communication from vRealize Operations Manager to Workspace ONE Access.	■ You must maintain the life cycle of the service account outside of the SDDC stack to ensure its availability.
SDDC-OPS-MON-061	Configure the Workspace ONE Access integration in vRealize Operations Manager.	■ Enables integration and data collection of all Workspace ONE Access instances in the SDDC in vRealize Operations Manager.	You must manage the password life cycle of this endpoint.
SDDC-OPS-MON-062	Configure the cross-region Workspace ONE Access integration to use the default collector group.	Components that are failed over between regions are configured to use the default collector group. This configuration provides monitoring of components during a failover.	The load on the analytics cluster, though minimal, increases.
SDDC-OPS-MON-063	Configure the region-specific Workspace ONE Access integration to use the remote collector group for its region.	Components that are not failed over between regions are configured to use the remote collector group. This offloads data collection for local management components from the analytics cluster.	None

Table 2-104. Design Decision on Information and Security Access for vRealize Operations Manager (continued)

	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-064	Configure the vRealize Log Insight integration in vRealize Operations Manager.	<ul style="list-style-type: none"> ■ Enables the Logs tab in vRealize Operations Manager ■ Enables the Troubleshoot with Logs dashboard ■ Enables the vRealize Log Insight launch in context from vRealize Operations Manager 	You must manage the password life cycle of this endpoint.
SDDC-OPS-MON-065	Configure the vRealize Log Insight integration to use the remote collector group.	Components that are not failed over between regions are configured to use the remote collector group. This offloads data collection for local management components from the analytics cluster.	None

Encryption Design for vRealize Operations Manager

Access to all vRealize Operations Manager Web interfaces requires an SSL connection. By default, vRealize Operations Manager uses a self-signed certificate. To provide secure access to the vRealize Operations Manager user interface, replace the default self-signed certificate with a CA-signed certificate.

Table 2-105. Design Decisions on Encryption for vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-066	Replace the default self-signed certificate of the virtual appliance of vRealize Operations Manager with a CA-signed certificate containing the analytics and remote collector nodes in the SAN attributes.	Configuring a CA-signed certificate ensures that the communication to the externally facing Web UI and API for vRealize Operations Manager, and cross-product, is encrypted.	<ul style="list-style-type: none"> ■ Replacing the default certificates with trusted CA-signed certificates from a certificate authority might increase the deployment preparation time as certificates requests are generated and delivered. ■ Each time a node is added the certificate must be replaced to include the new node.

Monitoring and Alerting Design in vRealize Operations Manager

You use vRealize Operations Manager to monitor the state of the management components in the SDDC by using dashboards. You can use the self-monitoring capability of vRealize Operations Manager to receive alerts about issues that are related to its operational state.

Table 2-106. vRealize Operations Manager Administrative Alert Types

Alert Type	Description
System alert	There is a failed component of the vRealize Operations Manager application.
Environment alert	vRealize Operations Manager stopped receiving data from one or more resources. Such an alert might indicate a problem with system resources or network infrastructure.
Log Insight log event	The infrastructure on which vRealize Operations Manager is running has low-level issues. You can also use the log events for root cause analysis.
Custom dashboard	vRealize Operations Manager can show super metrics for data center monitoring, capacity trends, and single pane of glass overview.

Table 2-107. Design Decisions on Monitoring of vRealize Automation

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-067	Create a service account, for example, svc-vrops-vra@rainpole.local , in the directory services and ensure it is synchronized in Workspace ONE Access.	The service account is used for application-to-application communication from vRealize Operations Manager to vRealize Automation.	<ul style="list-style-type: none"> ■ You must maintain the life cycle and availability of the service account outside of the SDDC stack. ■ You must maintain the synchronization and availability of the service account in Workspace ONE Access.
SDDC-OPS-MON-068	Assign the service account, for example, svc-vrops-vra@rainpole.local , the Organization Owner organization role and Cloud Assembly User service role for the application-to-application communication from vRealize Operations Manager to vRealize Automation.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ vRealize Operations Manager accesses vRealize Automation with the minimum set of required permissions for the integration. ■ If there is a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce an improved accountability in tracking request-response interactions between the vRealize Operations Manager and vRealize Automation integration. 	

Table 2-107. Design Decisions on Monitoring of vRealize Automation (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-069	Configure the integration for vRealize Automation in vRealize Operations Manager using the service account.	It provides the ability to monitor the health, efficiency, and capacity risks associated with vRealize Automation and the resources that are managed.	<ul style="list-style-type: none"> ■ Must be assigned to the default collector group since vRealize Automation is portable between SDDC regions. ■ You must update the vRealize Automation integration in vRealize Operations Manager when the service account password changes during its lifecycle.
SDDC-OPS-MON-070	Configure the integration for vRealize Automation in vRealize Operations Manager using the default collector group.	Ensures that vRealize Automation monitoring continues during a failover - planned or disaster recovery.	The load on the vRealize Automation analytics cluster, though minimal, increases.

Table 2-108. Design Decisions on Monitoring of Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-071	Create a service account, for example, svc-vrops-wsa@rainpole.local , in the directory services and ensure it is synchronized in Workspace ONE Access.	The service account is used for application-to-application communication from vRealize Operations Manager to Workspace ONE Access.	You must maintain the life cycle and availability of the service account in Workspace ONE Access outside of the SDDC stack.
SDDC-OPS-MON-072	Assign the service account, for example, svc-vrops-wsa@rainpole.local , user the Super Admin role for the application-to-application communication from vRealize Operations Manager to Workspace ONE Access.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ vRealize Operations Manager accesses Workspace ONE Access with the minimum set of required permissions for the integration. ■ If there is a compromised account, the accessibility in the destination application remains restricted. ■ Improved accountability in tracking request-response interactions between the vRealize Operations Manager and Workspace ONE Access integration. 	None

Table 2-108. Design Decisions on Monitoring of Workspace ONE Access (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-MON-073	Configure a Workspace ONE Access management pack adapter instance for each Workspace ONE Access instance using the service account.	Provides the ability to monitor the health, efficiency, and capacity risks associated with Workspace ONE Access.	You must update the Workspace ONE Access integration in vRealize Operations Manager when the service account password changes during its life cycle. Must be assigned to the specific collector groups based on the region-specific or cross-region designation.
SDDC-OPS-MON-074	Configure the Workspace ONE Access management pack adapter instance for each region-specific Workspace ONE Access instance using the region-specific collector group.	The components that are not failed over between regions are configured to use the local, region-specific remote collector group. This offloads data collection for local, region-specific management components from the vRealize Operations Manager analytics cluster.	None
SDDC-OPS-MON-075	Configure the Workspace ONE Access management pack adapter instance for the cross-region Workspace ONE Access cluster using the default collector group.	It ensures that cross-region Workspace ONE Access cluster monitoring continues during a failover - planned or disaster recovery.	The load on the vRealize Operations Manager analytics cluster, though minimal, increases.

Data Protection and Backup Design for vRealize Operations Manager

To preserve the cloud operations services functionality when data or system loss occurs, the design supports the use of data protection.

vRealize Operations Manager supports data protection through the creation of consistent image-level backups, using backup software that is based on the vSphere Storage APIs - Data Protection (VADP).

Disaster Recovery Design for vRealize Operations Manager

To preserve the cloud operations services functionality when a disaster occurs, this design supports the failover of vRealize Operations Manager between regions.

You place vRealize Operations Manager on the cross-region application virtual network, Mgmt-xRegion01-VXLAN. As a result, after the recovery, you continue to use the same IP address, DNS records, and routing configuration. vRealize Automation and vRealize Operations Manager also use this network for their cross-region failover capabilities.

If a planned migration or disaster occurs, you use Site Recovery Manager and vSphere Replication for an orchestrated recovery of the vRealize Operations Manager appliance. After the recovery, vRealize Operations Manager continues to provide cloud operations services functionality to manage the deployment of the available environments.

vRealize Log Insight Design

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

- [Logical Design of vRealize Log Insight](#)

In a multi-region SDDC, deploy a vRealize Log Insight cluster that consists of three nodes in each region. This configuration provides continued availability and increased log ingestion rates.

- [Configuration Design of vRealize Log Insight](#)

Configuration design details the design decisions covering physical design and sizing for vRealize Log Insight.

- [Life Cycle Management Design of vRealize Log Insight](#)

Life cycle management design details the design decisions covering the life cycle management of vRealize Log Insight.

- [Network Design of vRealize Log Insight](#)

In each region, for isolation and co-location with logging sources, the vRealize Log Insight instances are connected to the region-specific application virtual networks Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN. The networking design also supports public access to the vRealize Log Insight cluster.

- [Information Security and Access Design for vRealize Log Insight](#)

You protect the vRealize Log Insight deployment by configuring the authentication and secure communication with the other components in the SDDC. A dedicated service account is assigned a custom role for communication between vRealize Log Insight and the management solutions in the data center.

- [Event Forwarding Between Regions with vRealize Log Insight](#)

vRealize Log Insight supports event forwarding to other clusters and standalone instances. Use log forwarding between SDDC regions to have access to all logs if a disaster occurs in a region.

- [Data Protection and Backup Design for vRealize Log Insight](#)

To preserve the cloud operations services functionality when data or system loss occurs, the design supports the use of data protection.

- [Disaster Recovery Design for vRealize Log Insight](#)

Each region is configured to forward log information to the vRealize Log Insight instance in the other region.

Logical Design of vRealize Log Insight

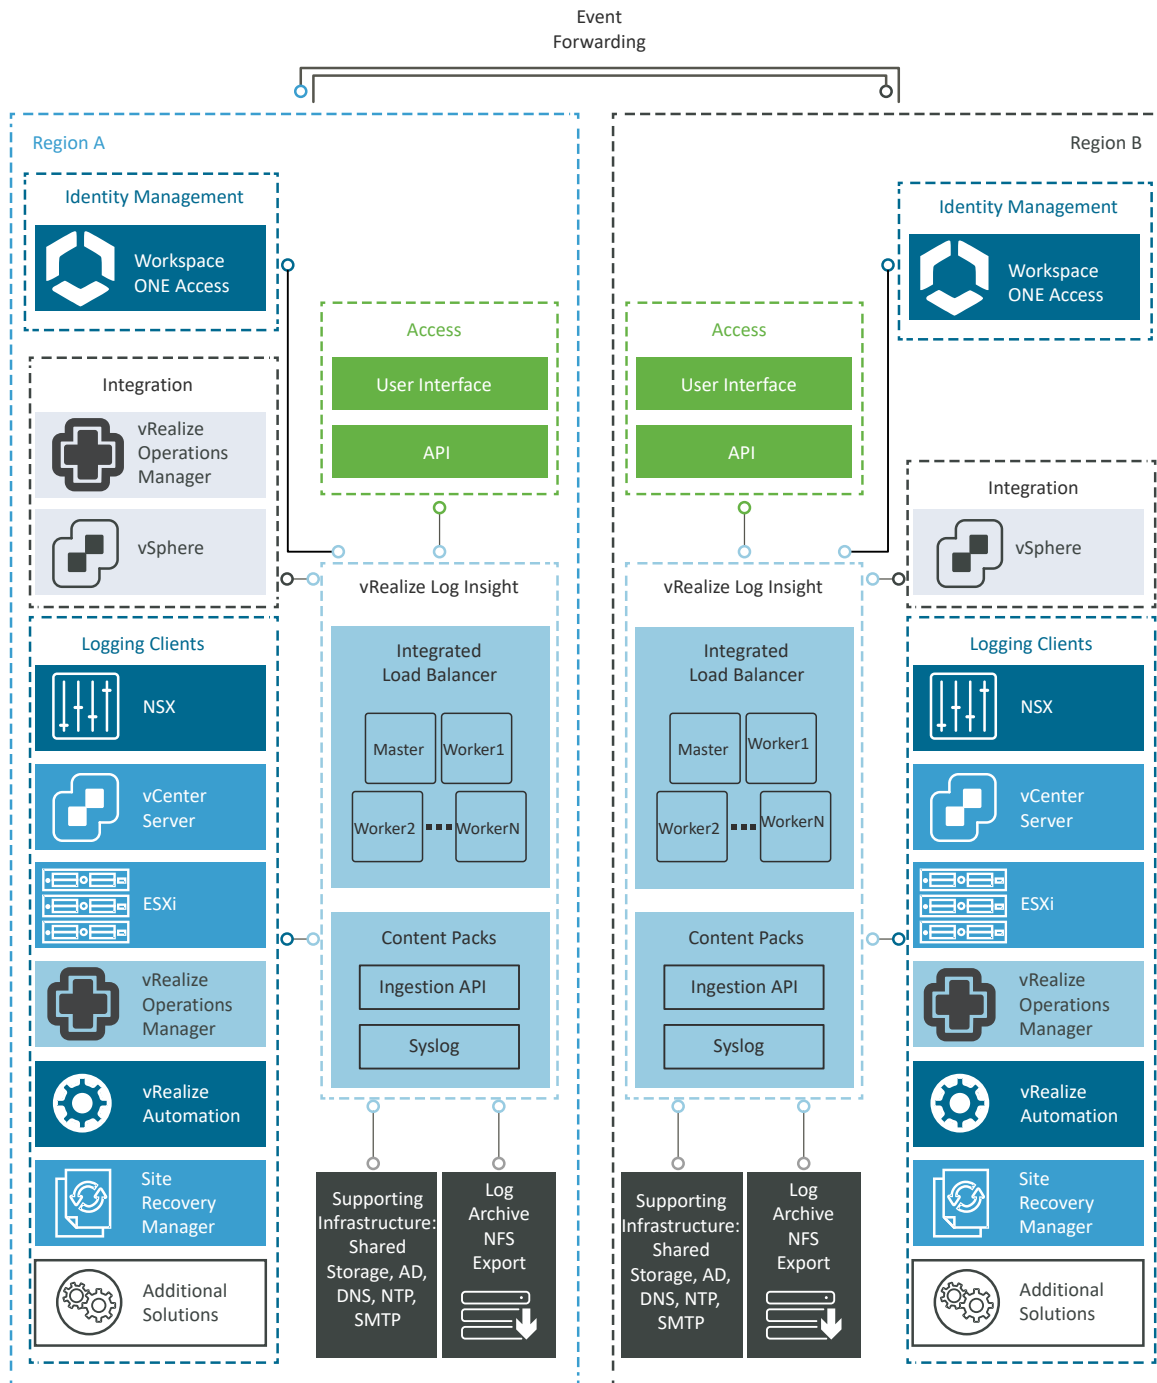
In a multi-region SDDC, deploy a vRealize Log Insight cluster that consists of three nodes in each region. This configuration provides continued availability and increased log ingestion rates.

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

Logical Design

In a multi-region SDDC, deploy a vRealize Log Insight cluster that consists of three nodes in each region. This configuration provides continued availability and increased log ingestion rates.

Figure 2-17. Logical Design of vRealize Log Insight



Configuration Design of vRealize Log Insight

Configuration design details the design decisions covering physical design and sizing for vRealize Log Insight.

Deployment Design of vRealize Log Insight

The vRealize Log Insight cluster consists of one master node and two worker nodes behind a load balancer.

You enable the integrated load balancer (ILB) on the three-node cluster so that all log sources can access the cluster by its ILB. By using the ILB, it is not necessary to reconfigure all log sources with a new destination address when there is a scale out. Using the ILB also guarantees that vRealize Log Insight accepts all incoming ingestion traffic.

vRealize Log Insight users, using both the Web user interface or API, and clients ingesting logs by using syslog or the Ingestion API, connect to vRealize Log Insight by using the ILB address.

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

To accomplish this design objective, you deploy or reuse the following components to deploy this operations management solution for the SDDC.

- vRealize Suite Lifecycle Manager
- Workspace ONE Access
- Supporting infrastructure services, such as Active Directory, DNS, and NTP.

You place the vRealize Log Insight on a specific application virtual network for isolation.

vRealize Log Insight is distributed as a virtual appliance in OVA format.

In the design, you deploy the vRealize Log Insight appliance instances on the first vSphere cluster in the management domain of each region. The SDDC can comprise multiple regions and multiple availability zones.

Sizing Compute and Storage Resources for vRealize Log Insight

You size resources for vRealize Log Insight to provide enough resources to accommodate the logging operations of the management components of the SDDC.

Compute Resources

To accommodate log data from the products in the SDDC, you must correctly size the compute resources and storage for the Log Insight nodes. For detailed sizing guidance, see the vRealize Log Insight sizing calculator at <https://kb.vmware.com/s/article/60355>.

By default, the vRealize Log Insight appliance uses the predefined values for medium configurations.

Sizing Nodes

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

Table 2-109. Compute Resources for a vRealize Log Insight Medium-Size Node

Attribute	Specification
Appliance size	Medium
CPU	8 vCPUs
Memory	16 GB
Disk capacity	530 GB
IOPS	1 000 IOPS
Amount of processed log data when using log ingestion.	75 GB/day of processing per node
Number of processed log messages.	5 000 event/second of processing per node
Environment	Up to 250 syslog connections per node

Storage Resources

Sizing is usually based on the requirements of the organization. This design provides calculations that are based on a single-region implementation and are implemented on a per-region basis. This sizing is calculated according to the following node configuration per region:

Table 2-110. Management Systems That Send Log Data to vRealize Log Insight

Category	Logging Sources	Quantity
Management Workload Domain	Platform Services Controller	2
	vCenter Server	1
	Site Recovery Manager	1
	vSphere Replication	1
	ESXi Hosts	4
	NSX Manager	1
	NSX Controller	3
	NSX Edge	5
VI Workload Domain	vCenter Server	1
	ESXi Hosts	64
	NSX Manager	3
	NSX Edge	5
Workspace ONE Access	Workspace ONE Access Appliances	4
SDDC Manager	SDDC Manager Appliance	1
vRealize Suite Lifecycle Manager	vRealize Suite Lifecycle Manager Appliance	1
vRealize Operations Manager	vRealize Operations Manager Appliances	3
vRealize Automation	vRealize Automation Appliances	3
Total Region A		103

In this design, to simplify the calculations, all calculations are done using the large 220-byte size which results in 190 MB of log data expected per day per source.

For 206 logging sources (two regions), at a base rate of approximately 190 MB of logs that are ingested per day per source over seven days, you need approximately 160 GB of storage per node. Based on this example, the storage space that is allocated per medium-size vRealize Log Insight virtual appliance is enough to monitor the SDDC.

Consider the following approaches when you must increase the Log Insight capacity:

- If you must maintain a log data retention for more than seven days in your SDDC, you can 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, extend the storage for all nodes.

- 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 do the following:
 - Increase the size of the vRealize Log Insight node to a large deployment size as defined in the vRealize Log Insight documentation.
 - Deploy more vRealize Log Insight nodes to scale out your environment. vRealize Log Insight can scale up to 12 nodes in an HA cluster.

Table 2-111. Design Decisions for vRealize Log Insight Deployment

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-001	In each region, deploy vRealize Log Insight in a cluster configuration of three nodes with an integrated load balancer: one master and two worker nodes, on the first vSphere cluster in the management domain.	<ul style="list-style-type: none"> ■ Provides high availability. ■ Using the integrated load balancer prevents a single point of failure. ■ Using the integrated load balancer simplifies the vRealize Log Insight deployment and subsequent integration. 	<ul style="list-style-type: none"> ■ You must deploy a minimum of three medium nodes. ■ You must size each node identically. ■ If the capacity of your vRealize Log Insight cluster must expand, identical capacity must be added to each node.
SDDC-OPS-LOG-002	Deploy vRealize Log Insight through vRealize Suite Lifecycle Manager.	Allows vRealize Suite Lifecycle Manager the ability to provide life cycle management of vRealize Log Insight.	You must deploy vRealize Suite Lifecycle Manager.
SDDC-OPS-LOG-003	Apply vSphere Distributed Resource Scheduler (DRS) anti-affinity rules to the vRealize Log Insight cluster nodes.	Using vSphere DRS prevents the vRealize Log Insight cluster nodes from running on the same ESXi host and risking the high availability of the cluster.	<ul style="list-style-type: none"> ■ You must perform an additional configuration to set up an anti-affinity rule. ■ You can put in maintenance mode only a single ESXi host at a time in a management cluster of four ESXi.

Table 2-111. Design Decisions for vRealize Log Insight Deployment (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-004	Deploy each node in the vRealize Log Insight cluster as a medium-size appliance.	<p>Accommodates the number of expected syslog and vRealize Log Insight Agent connections from the following sources:</p> <ul style="list-style-type: none"> ■ Management workload domain vCenter Server and Compute workload domain vCenter Server instances ■ Management ESXi hosts, and shared edge and compute ESXi hosts ■ Management Site Recovery Manager components ■ Management and compute components of NSX for vSphere. ■ Workload domain components of NSX-T Data Center ■ Workspace ONE Access ■ SDDC Manager ■ vRealize Suite Lifecycle Manager ■ vRealize Automation components ■ vRealize Operations Manager components ■ Skyline Collector ■ Cross- vRealize Log Insight cluster event forwarding. <p>These components generate approximately 200 syslog and vRealize Log Insight Agent sources.</p> <p>Using medium-size appliances ensures that the storage space for the vRealize Log Insight cluster is sufficient for seven days of data retention.</p>	You must increase the size of the nodes if you configure vRealize Log Insight to monitor additional syslog sources.
SDDC-OPS-LOG-005	When using two availability zones in Region A, add the vRealize Log Insight virtual appliances to the primary availability zone VM group, for example, sfo01-m01-	It ensures that the vRealize Log Insight virtual appliance is powered on within the primary availability zone hosts group by default.	If vRealize Log Insight is deployed after the creation of the stretched cluster for management domain availability zones, the VM

Table 2-111. Design Decisions for vRealize Log Insight Deployment (continued)

Decision ID	Design Decision	Design Justification	Design Implication
	mgmt01-primary-az-vm-group.		group for the primary availability zone virtual machines must be updated to include the vRealize Log Insight virtual appliances.

Notifications Design for vRealize Log Insight

Enable notifications for alerts in vRealize Log Insight.

Table 2-112. Design Decision on Alert Notifications for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-006	Enable alert notifications.	Enables administrators and operators to receive alerts by email and from vRealize Log Insight.	Requires access to an external SMTP server.

Retention and Archiving Design in vRealize Log Insight

Configure archive and retention parameters of vRealize Log Insight according to the company policy for compliance and governance.

Each vRealize Log Insight appliance has three default virtual disks and can use more virtual disks for storage.

Table 2-113. Virtual Disk Configuration in the vRealize Log Insight Appliance

Hard disk	Size	Usage
Hard disk 1	20 GB	Root file system
Hard disk 2	510 GB for medium-size deployment	Contains two partitions: <ul style="list-style-type: none"> ■ /storage/var for system logs ■ /storage/core for collected logs
Hard disk 3	512 MB	First boot only

Configure a retention period of seven days for the medium-size vRealize Log Insight appliance.

Table 2-114. Log Archiving Attributes for vRealize Log Insight

Attribute	Description
Archiving period	vRealize Log Insight archives log messages as soon as possible. At the same time, the logs are retained on the virtual appliance until the free local space is almost filled. Data exists on both the vRealize Log Insight appliance and the archive location for most of the retention period. The archiving period must be longer than the retention period.
Archive location	The archive location must be on an NFS version 3 shared storage. The archive location must be available and must have enough capacity to accommodate the archives.

You configure vRealize Log Insight to archive log data only if you must retain logs for an extended period for compliance, auditability, or a customer-specific reason.

Apply an archive policy of 90 days for the medium-size vRealize Log Insight appliance. The vRealize Log Insight cluster uses an estimated 400 GB of shared storage. To calculate required archiving storage, you can use the vRealize Log Insight sizing calculator. See <https://kb.vmware.com/s/article/60355>.

Table 2-115. Design Decision on Log Archive Policy for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS- LOG-007	Provide a minimum of 400 GB of NFS version 3 shared storage to the vRealize Log Insight cluster in each region.	Accommodates log archiving from 200 logging sources for 90 days.	<ul style="list-style-type: none"> ■ You must manually maintain the vRealize Log Insight archive blobs stored on the NFS store, selectively cleaning the datastore as more space is required. ■ You must increase the size of the NFS shared storage if you configure vRealize Log Insight to monitor more logging sources or more vRealize Log Insight workers are added. ■ You must enforce the archive policy directly on the shared storage. ■ If the NFS mount does not have enough free space or is unavailable for a period greater than the retention period of the virtual appliance, vRealize Log Insight stops ingesting new data until the NFS mount has enough free space, becomes available, or archiving is disabled. ■ When using two availability zones, ensure that the NFS share is available in both availability zones.

Monitoring and Alerting in vRealize Log Insight

vRealize Log Insight supports alerts that trigger notifications about its health and about the health of monitored solutions.

Table 2-116. vRealize Log Insight Administrative Alert Types

Alert Type	Description
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	<p>Administrators and users can define their own alerts based on data ingested by vRealize Log Insight.</p> <p>vRealize Log Insight handles alerts in two ways:</p> <ul style="list-style-type: none"> ■ Send an email over SMTP. ■ Send System Notifications to Third-Party Products using Webhooks. ■ Send to vRealize Operations Manager.

Integration of vRealize Log Insight with vRealize Operations Manager and Workspace ONE Access

vRealize Log Insight supports integration with vRealize Operations Manager to provide a central location for monitoring and diagnostics.

Table 2-117. vRealize Log Insight Integration Points

Integration Point	Description
Notification Events	Forward notification events from vRealize Log Insight to vRealize Operations Manager.
Launch in Context	Launch vRealize Log Insight from the vRealize Operation Manager user interface.
Embedded vRealize Log Insight	Access the integrated vRealize Log Insight user interface directly in the vRealize Operations Manager user interface.

Table 2-118. Design Decisions on Integration of vRealize Log Insight with vRealize Operations Manager

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS- LOG-008	In vRealize Operations Manager, add an application-to-application service account from Workspace ONE Access, for vRealize Log Insight Integration, for example, svc-vrli-vrops@rainpole.local . Assign this user the default Super Admin role.	Enables integration between vRealize Log Insight and vRealize Operations Manager.	You must maintain the life cycle and availability of the service account outside of the SDDC stack.
SDDC-OPS- LOG-009	Enable vRealize Operations Manager integration in vRealize Log Insight using the vRealize Operations Manager service account, for example, svc-vrli-vrops@rainpole.local .	Integrating vRealize Log Insight alerts with vRealize Operations Manager allows you to view all information about your environment in a single user interface.	<ul style="list-style-type: none"> ■ You must maintain the life cycle of this integration. ■ You must specify the user account in the <i>user@domain@source</i> format for the integration. <i>Source</i> is the name of the authentication source created in Workspace ONE Access, for example, svc-vrli-vrops@rainpole.local@Workspace ONE.
SDDC-OPS- LOG-010	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.	None
SDDC-OPS- LOG-011	Support launch in context with vRealize Operation Manager.	Provides access to vRealize Log Insight for context-based monitoring of an object in vRealize Operations Manager.	You can register only one vRealize Log Insight cluster with vRealize Operations Manager for launch in context at a time.
SDDC-OPS- LOG-012	Enable the embedded vRealize Log Insight user interface in vRealize Operations Manager.	Provides central access to the vRealize Log Insight user interface for improved context-based monitoring on an object in vRealize Operations Manager.	You can register only one vRealize Log Insight cluster with vRealize Operations Manager at a time.

Table 2-119. Design Decisions on Integration of vRealize Log Insight with Workspace ONE Access

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-013	Do not configure the Workspace ONE Access virtual appliances to use the Syslog protocol for logs.	Workspace ONE Access virtual appliances are configured to use the vRealize Log Insight ingestion API.	None
SDDC-OPS-LOG-014	<p>Install and configure the vRealize Log Insight Agent on each Workspace ONE Access virtual appliance to send logs to a vRealize Log Insight cluster.</p> <ul style="list-style-type: none"> ■ Use the vRealize Log Insight Agent from the region-specific vRealize Log Insight cluster for the corresponding region-specific Workspace ONE Access Instance. ■ Use the vRealize Log Insight Agent from the Region A region-specific vRealize Log Insight cluster for cross-region Workspace ONE Access cluster. 	<p>Provides a standardized configuration that is pushed to the vRealize Log Insight Agents for each Workspace ONE Access virtual appliance.</p> <p>Supports collection according to the context of the Workspace ONE Access using the vRealize Log Insight ingestion API and parses of the logs by the vRealize Log Insight agent such as specific log directories, log files, and logging formats.</p>	None
SDDC-OPS-LOG-015	Integrate with the vRealize Log Insight using the Ingestion API port=9000 (default) and ssl=no (non-default).	Supports disaster recovery of the cross-region Workspace ONE Access cluster in the SDDC.	Transmission traffic for logs is not secure.
SDDC-OPS-LOG-016	Configure an agent group in each vRealize Log Insight cluster for all Workspace ONE Access virtual appliances.	<p>Provides a standardized configuration that is pushed to the vRealize Log Insight Agents for each Workspace ONE Access virtual appliance.</p> <p>Supports collection according to the context of the Workspace ONE Access using the vRealize Log Insight ingestion API and parses of the logs by the vRealize Log Insight agent such as specific log directories, log files, and logging formats.</p>	Adds minimal load to the vRealize Log Insight cluster.

Content Packs in vRealize Log Insight

Use content packs to have the logs generated from the management components in the SDDC retrieved, extracted and parsed into a human-readable format. In this way, Log Insight saves log queries and alerts, and you can use dashboards for efficient monitoring.

Table 2-120. vRealize Log Insight Content Packs in This Design

Content Pack	Installed by Default
General	✓
VMware - vSphere	✓
VMware – vRealize Operations Manager	✓
VMware - NSX for vSphere	x
VMware – NSX-T Data Center	x
VMware - Linux	x
VMware - Site Recovery Manager	x
VMware – Workspace ONE Access	x

Table 2-121. Design Decisions on Content Packs for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS- LOG-017	Install the following content packs: <ul style="list-style-type: none"> ■ VMware - NSX for vSphere ■ VMware - Linux ■ VMware – NSX-T Date Center ■ VMware – Workspace ONE Access 	Provides additional granular monitoring on the virtual infrastructure. The following content packs are installed by default in vRealize Log Insight: <ul style="list-style-type: none"> ■ General VMware - ■ vSphere ■ VMware - vSAN ■ VMware - vRealize Operations Manager 	Requires installation and configuration of each non-default content pack.
SDDC-OPS- LOG-018	Configure the following agent groups that are related to content packs: <ul style="list-style-type: none"> ■ VMware Virtual Appliances ■ Linux 	<ul style="list-style-type: none"> ■ Provides a standardized configuration that is pushed to the all vRealize Log Insight Agents in each of the groups. ■ Supports collection according to the context of the applications and parsing of the logs generated from the SDDC components by the vRealize Log Insight agent such as specific log directories, log files, and logging formats. 	Adds minimal load to vRealize Log Insight.

Life Cycle Management Design of vRealize Log Insight

Life cycle management design details the design decisions covering the life cycle management of vRealize Log Insight.

Life cycle management of vRealize Log Insight involves the process of performing patch updates or upgrades to the vRealize Log Insight cluster.

In this design, life cycle management is performed using vRealize Suite Lifecycle Manager.

Table 2-122. Design Decisions on the Life Cycle Management of vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-019	Use vRealize Suite Lifecycle Manager to perform the life cycle management of vRealize Log Insight.	vRealize Suite Lifecycle Manager automates the life cycle of vRealize Log Insight.	You must deploy vRealize Suite Lifecycle Manager.

Network Design of vRealize Log Insight

In each region, for isolation and co-location with logging sources, the vRealize Log Insight instances are connected to the region-specific application virtual networks Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN. The networking design also supports public access to the vRealize Log Insight cluster.

Application Network Segment

The vRealize Log Insight virtual appliances are connected to the region-specific application virtual networks, for example, Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN, for collection of logs locally per region.

This networking design has the following features:

- The vRealize Log Insight cluster appliances are deployed together on the same network in each region. This configuration ensures collection of logs locally per region in the event of a cross-region network outage. It also co-locates log collection with the region-specific applications using the region-specific application virtual network.
- vRealize Log Insight has routed access to the VLAN-backed management network through the NSX Universal Distributed Logical Router.
- Routing to the VLAN-management network, application virtual networks, and external networks are dynamic and are based on the Border Gateway Protocol (BGP).

Figure 2-18. Networking Design of the vRealize Log Insight Deployment

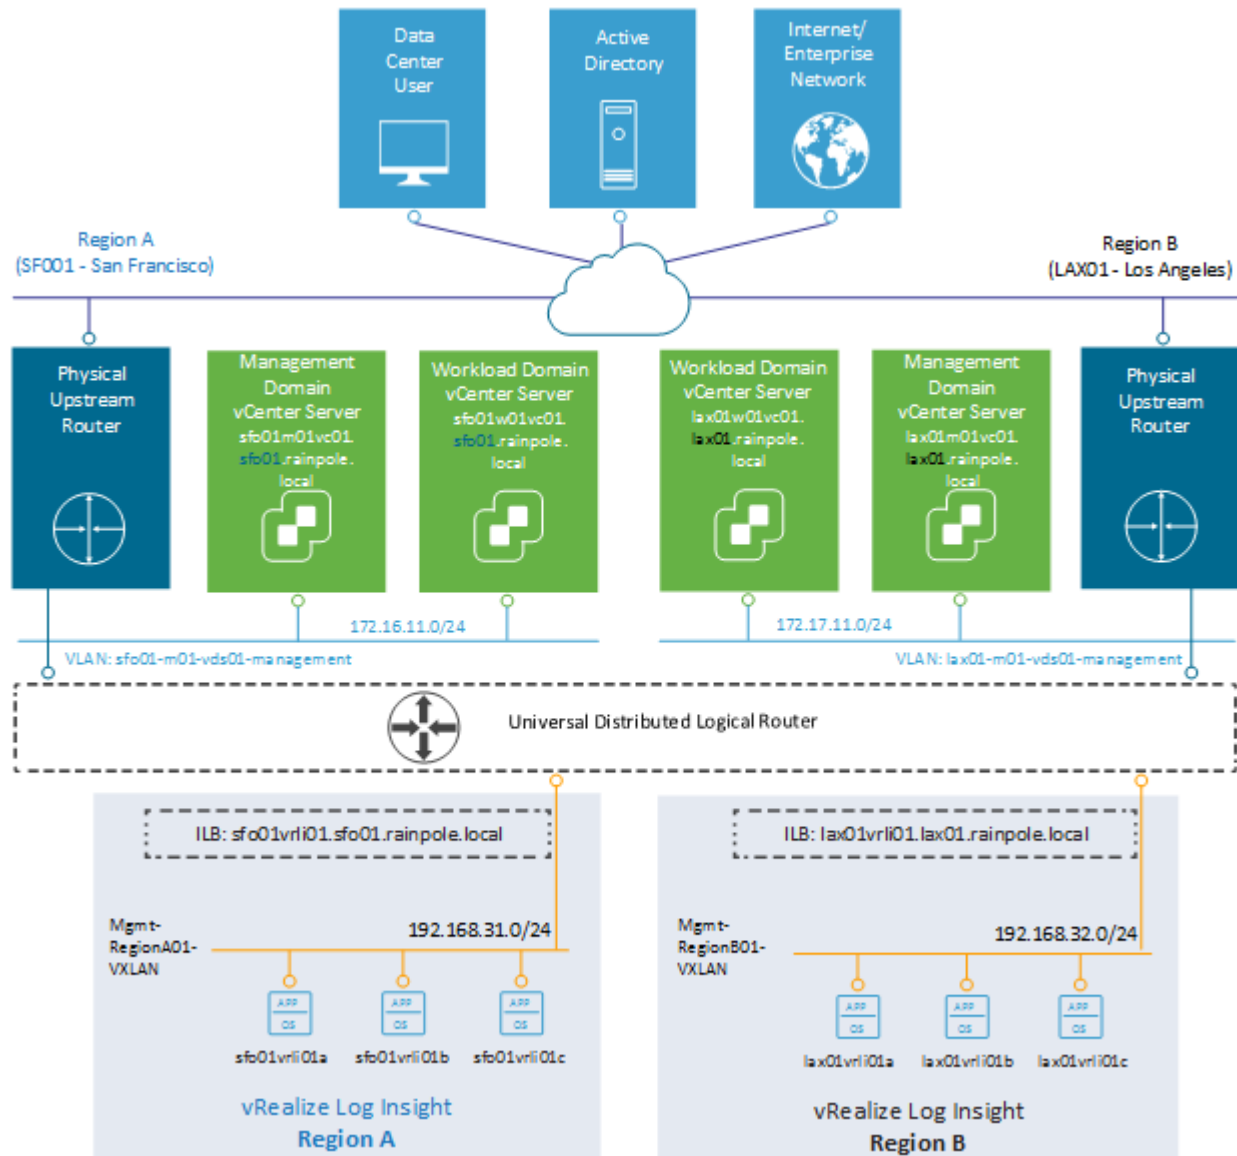


Table 2-123. Design Decisions on the Application Virtual Network for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-020	Place the vRealize Log Insight appliances on the region-specific application virtual network, for example, Mgmt-RegionA01-VXLAN and Mgmt-RegionB01-VXLAN.	<ul style="list-style-type: none"> Ensures centralized access to log data per region if a cross-region network outage occurs. Co-locates 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 cross-region network can impact event forwarding between the vRealize Log Insight clusters and cause gaps in log data. You must use an implementation in NSX for vSphere to support this networking configuration.

IP Addressing Scheme

You allocate a subnet for the cross-region network segment, and the region-specific network segments in the management domain and use them for the vRealize Log Insight deployment.

Table 2-124. Example IP Subnets for vRealize Log Insight

vRealize Log Insight Cluster	IP Subnet	Gateway	NSX Application Virtual Network
Region A	192.168.31.0/24	192.168.31.1	Mgmt-RegionA01-VXLAN
Region B	192.168.32.0/24	192.168.32.1	Mgmt-RegionB01-VXLAN

Table 2-125. Design Decisions on the IP Subnets for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-021	Allocate separate subnets for each region-specific application virtual network.	vRealize Log Insight does not fail over between regions. Allocating a region-specific application virtual network subnet enables communication with region-specific components for log collection.	None

Name Resolution

vRealize Log Insight node name resolution, including the integrated load balancer virtual IP addresses (VIPs), uses a region-specific suffix, such as sfo01.rainpole.local or lax01.rainpole.local. The Log Insight components in both regions have the following node names.

Table 2-126. Example FQDN and IP Address for vRealize Log Insight

Region	Fully Qualified Domain Name	IP Address	Description
Region A	sfo01vrli01.sfo01.rainpole.local	192.168.31.10	Integrated Load Balancer VIP
	sfo01vrli01a.sfo01.rainpole.local	192.168.31.11	Master node
	sfo01vrli01b.sfo01.rainpole.local	192.168.31.12	Worker node
	sfo01vrli01c.sfo01.rainpole.local	192.168.31.13	Worker node
	sfo01vrli01x.sfo01.rainpole.local	192.168.31.n	Additional worker nodes (not deployed)
Region B	lax01vrli01.lax01.rainpole.local	192.168.32.10	Log Insight ILB VIP
	lax01vrli01a.lax01.rainpole.local	192.168.32.11	Master node
	lax01vrli01b.lax01.rainpole.local	192.168.32.12	Worker node
	lax01vrli01c.lax01.rainpole.local	192.168.32.13	Worker node
	lax01vrli01x.lax01.rainpole.local	192.168.32.n	Additional worker nodes (not deployed)

Table 2-127. Design Decisions on DNS for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LCM-022	Configure forward and reverse DNS records for all vRealize Log Insight nodes and the Integrated Load Balancer VIP address.	All nodes are accessible by using fully qualified domain names instead of by using IP addresses only.	You must provide DNS records for the vRealize Log Insight appliances.
SDDC-OPS-LCM-023	For all applications that fail over between regions, such as vRealize Automation and vRealize Operations Manager, use the FQDN of the vRealize Log Insight integrated load balancer (ILB) in Region A when you configure logging.	Logging continues during a partial failover to Region B. For example, only one application is moved to Region B.	<ul style="list-style-type: none"> ■ If vRealize Automation and vRealize Operations Manager are failed over to Region B and the vRealize Log Insight cluster is no longer available in Region A, you must update the A record on the child DNS server to point to the vRealize Log Insight cluster in Region B. ■ You must set ssl=no for the vRLI agents.

Time Synchronization

Time synchronization provided by the Network Time Protocol (NTP) is important to ensure that all components within the SDDC are synchronized to the same time source. Configure time synchronization using an internal NTP time source across all vRealize Log Insight appliances.

Table 2-128. Design Decisions on Time Synchronization for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-024	Configure NTP on each vRealize Log Insight appliance.	vRealize Log Insight is dependent on time synchronization.	None

Information Security and Access Design for vRealize Log Insight

You protect the vRealize Log Insight deployment by configuring the authentication and secure communication with the other components in the SDDC. A dedicated service account is assigned a custom role for communication between vRealize Log Insight and the management solutions in the data center.

Authentication and Authorization Design for vRealize Log Insight

Users can authenticate to vRealize Log Insight in the following ways:

Table 2-129. vRealize Log Insight Account Types

Account Type	Description
Imported users or user groups from Microsoft Active Directory	Users can use their AD credentials to log in to vRealize Log Insight.
Integrated users and groups with Workspace ONE Access	Specified users and groups from upstream identity sources are synchronized to vRealize Log Insight through Workspace ONE Access.
Local user accounts created in vRealize Log Insight	vRealize Log Insight performs local authentication using the account information stored in its global database.

You enable authentication by using Workspace ONE Access to ensure accountability on user access. You can grant both users and groups access to vRealize Log Insight to perform tasks, such as analyzing logs and viewing dashboards.

Integrating vRealize Log Insight with the SDDC

vRealize Log Insight collects logs as to provide monitoring information about the SDDC from a central location. As a 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 by using the syslog TCP, syslog TCP over TLS/SSL, or syslog UDP protocols
- By using a vRealize Log Insight Agent
- By using vRealize Log Insight to query directly the vSphere Web Server APIs

- By using a vRealize Log Insight user interface.

vRealize Log Insight collects log events from the following virtual infrastructure and cloud management components:

Table 2-130. vRealize Log Insight Logging Sources

Logging Sources	Logging Type
vCenter Server	Syslog
ESXi Hosts	Syslog
NSX for vSphere Manager	Syslog
NSX for vSphere Controller instances	Syslog
NSX for vSphere edge instances	Syslog
NSX-T Manager	Syslog
NSX-T Edge	Syslog
Workspace ONE Access	Agent
SDDC Manager	Agent
vRealize Suite Lifecycle Manager	Agent
vRealize Operations Manager	Agent
Site Recovery Manager	Agent
vRealize Automation	Fluentd Plugin for vRealize Log Insight

Table 2-131. Design Decision on Authentication and Authorization for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-025	Rotate the root password on or before 365 days post deployment.	The password for the root user account expires 365 days after the initial deployment.	You must manage the password rotation schedule for the root user account in accordance with your organization policies and regulatory standards, as applicable.
SDDC-OPS-LOG-026	Enable vRealize Log Insight integration with your corporate identity source using the regional Workspace ONE Access instance.	Allows authentication, including multi-factor, to vRealize Log Insight using your corporate identity source. Allows authorization through the assignment of roles to enterprise users and groups defined in your corporate identity source.	You must deploy and configure the regional Workspace ONE Access instance to establish the integration between vRealize Log Insight and your corporate identity sources.

Table 2-131. Design Decision on Authentication and Authorization for vRealize Log Insight (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-027	Create a security group in your organization directory services for the vRealize Log Insight administrators role, for example, rainpole.locallug-vrli-admins , and synchronize the group in the Workspace ONE Access configuration for vRealize Log Insight.	Allows you to streamline the management of vRealize Log Insight roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS-LOG-028	Assign the enterprise group for vRealize Log Insight administrators, for example, rainpole.locallug-vrli-admins , the Super Admin role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to vRealize Log Insight administration is granted to a managed set of individuals that are members of the security group. ■ You can introduce improved accountability and tracking organization owner access to vRealize Log Insight. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-OPS-LOG-029	Create a security group in your organization directory services for the vRealize Log Insight content administrators role, for example, rainpole.locallug-vrli-content-admins , and synchronize the group in the Workspace ONE Access configuration for vRealize Log Insight.	Allows you to streamline the management of vRealize Log Insight roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS-LOG-030	Assign the enterprise group for vRealize Log Insight content administrators, for example, rainpole.locallug-vrli-content-admins , the ContentAdmin role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to vRealize Log Insight administration is granted to a managed set of individuals that are members of the security group. ■ You can introduce improved accountability and tracking organization owner access to vRealize Log Insight. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-131. Design Decision on Authentication and Authorization for vRealize Log Insight (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-031	Create a security group in your organization directory services for the vRealize Log Insight users role, for example, rainpole.local\ug-vrli-users , and synchronize the group in the Workspace ONE Access configuration for vRealize Log Insight.	Allows you to streamline the management of vRealize Log Insight roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the desired directory synchronization interval in Workspace ONE Access to ensure changes are available within a reasonable period.
SDDC-OPS-LOG-032	Assign the enterprise group for vRealize Log Insight users, for example, rainpole.local\ug-vrli-users , the User role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to vRealize Log Insight user interface is granted to a managed set of individuals that are members of the security group. ■ You can introduce improved accountability and tracking organization owner access to vRealize Log Insight. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.
SDDC-OPS-LOG-033	Create a security group in your organization directory services for the vRealize Log Insight viewers role, for example, rainpole.local\ug-vrli-viewers , and synchronize the group in the Workspace ONE Access configuration for vRealize Log Insight.	Allows you to streamline the management of vRealize Log Insight roles for users.	<ul style="list-style-type: none"> ■ You must create the security group outside of the SDDC stack. ■ You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.
SDDC-OPS-LOG-034	Assign the enterprise group for vRealize Log Insight viewers, for example, rainpole.local\ug-vrli-viewers , the Viewer role.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> ■ Access to the vRealize Log Insight user interface is granted to a managed set of individuals that are members of the security group. ■ You can introduce improved accountability and tracking organization owner access to vRealize Log Insight. 	You must maintain the life cycle and availability of the security group outside of the SDDC stack.

Table 2-131. Design Decision on Authentication and Authorization for vRealize Log Insight (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-035	Configure syslog sources and vRealize Log Insight Agents to send log data directly to the virtual IP (VIP) address of the vRealize Log Insight integrated load balancer (ILB).	<ul style="list-style-type: none"> ■ Allows for future scale-out without reconfiguring all log sources with a new destination address. ■ Simplifies the configuration of log sources in the SDDC 	<ul style="list-style-type: none"> ■ You must configure the integrated load balancer on the vRealize Log Insight cluster. ■ You must configure logging sources to forward data to the vRealize Log Insight VIP.
SDDC-OPS-LOG-036	Communicate with the vRealize Log Insight agents by using the default Ingestion API (cfapi using port 9000), default disk buffer of 200 MB and non-default No SSL.	<ul style="list-style-type: none"> ■ Supports multi-line message transmissions from logs. ■ Provides ability to add metadata to events generated from system. ■ Provides client-side compression, buffering, and throttling capabilities ensuring minimal to no message loss during intermittent connection problems. ■ Provides server-side administration, metric collection, configurations management of each deployed agent. ■ Supports disaster recovery of components in the SDDC. 	<ul style="list-style-type: none"> ■ Transmission traffic is not secure. ■ Agent presence increases the overall resources used on the system.
SDDC-OPS-LOG-037	Configure all vCenter Server instances as direct syslog sources to send log data directly to vRealize Log Insight.	Simplifies configuration 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 in vRealize Log Insight require the use of the vRealize Log Insight agent for proper ingestion. <p>Not all operating system level events are forwarded to vRealize Log Insight.</p>

Table 2-131. Design Decision on Authentication and Authorization for vRealize Log Insight (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-038	Define a custom vCenter Server role for vRealize Log Insight that has the minimum privileges required to support collecting logs from vSphere endpoints across the SDDC, for example, vRealize Log Insight to vSphere Integration .	vRealize Log Insight accesses vSphere with the minimum set of permissions that are required to support collecting logs from vSphere endpoints across the SDDC.	You must maintain the permissions required by the custom role.
SDDC-OPS-LOG-039	Configure a service account in vCenter Server with global permissions, for application-to-application communication from vRealize Log Insight to vSphere, for example, svc-vrli-vsphere@rainpole.local , and assign the custom role, for example, vRealize Log Insight to vSphere Integration .	Provides the following access control features: <ul style="list-style-type: none"> ■ vRealize Log Insight accesses vSphere with the minimum set of permissions that are required to collect logs. ■ In the event of a compromised account, the accessibility in the destination application remains restricted. ■ You can introduce improved accountability in tracking request-response interactions between the components of the SDDC. 	<ul style="list-style-type: none"> ■ You must maintain the life cycle and availability of the service account outside of the SDDC stack. ■ All vCenter Server instances must be in the same vSphere domain.
SDDC-OPS-LOG-040	Configure vRealize Log Insight to ingest events, tasks, and alarms from the Management vCenter Server and Compute vCenter Server instances by using the vRealize Log Insight service account, for example, svc-vrli-vsphere@rainpole.local .	Ensures that all tasks, events, and alarms generated across all vCenter Server instances in a specific region of the SDDC are captured and analyzed for the administrator.	You must manage the password life cycle of this service account.
SDDC-OPS-LOG-041	Configure the vRealize Log Insight agent on the SDDC Manager appliance.	Simplifies configuration of log sources in the SDDC that are pre-packaged with the vRealize Log Insight agent.	You must configure the vRealize Log Insight agent to forward logs to the vRealize Log Insight VIP.

Table 2-131. Design Decision on Authentication and Authorization for vRealize Log Insight (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-042	Define a custom vCenter Server role for vRealize Log Insight that has the minimum privileges required to support collecting metrics from vSphere endpoints across the SDDC, for example, vRealize Operations to vSphere Integration – Metrics .	vRealize Log Insight accesses vSphere with the minimum set of permissions that are required to support collecting metrics from vSphere endpoints across the SDDC.	You must maintain the permissions required by the custom role.
SDDC-OPS-LOG-043	Configure the vRealize Log Insight agent on the vRealize Suite Lifecycle Manager appliance.	Simplifies configuration of log sources in the SDDC that are pre-packaged with the vRealize Log Insight agent.	You must configure the vRealize Log Insight agent to forward logs to the vRealize Log Insight VIP.
SDDC-OPS-LOG-044	Configure the Fluentd vRealize Log Insight plug-in on the vRealize Automation appliance instances.	Enables the appliance and the containers to send logs to vRealize Log Insight.	You must configure the Fluentd vRealize Log Insight plug-in to forward logs to the vRealize Log Insight VIP.
SDDC-OPS-LOG-045	Configure the vRealize Log Insight agent for the vRealize Operations Manager appliances including: <ul style="list-style-type: none"> ■ Analytics nodes ■ Remote Collector instances 	Simplifies configuration of log sources in the SDDC that are pre-packaged with the vRealize Log Insight agent.	You must configure the vRealize Log Insight agent to forward logs to the vRealize Log Insight VIP.
SDDC-OPS-LOG-046	Configure the vRealize Log Insight agent on the Skyline Collector appliance.	Simplifies configuration of log sources in the SDDC that are pre-packaged with the vRealize Log Insight agent.	You must configure the vRealize Log Insight agent to forward logs to the vRealize Log Insight VIP.
SDDC-OPS-LOG-047	Configure the NSX for vSphere components as direct syslog sources for vRealize Log Insight including: <ul style="list-style-type: none"> ■ NSX Manager instances ■ NSX Controller instances ■ NSX Edge Services Gateway instances 	Simplifies configuration of log sources in the SDDC 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. ■ Not all operating system-level events are forwarded to vRealize Log Insight.
SDDC-OPS-LOG-048	Configure the NSX-T Data Center components as direct syslog sources for vRealize Log Insight including: <ul style="list-style-type: none"> ■ NSX-T Manager instances ■ NSX Edge Services Gateway instances 	Simplifies configuration of log sources in the SDDC 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. ■ Not all operating system-level events are forwarded to vRealize Log Insight.

Table 2-131. Design Decision on Authentication and Authorization for vRealize Log Insight (continued)

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-049	Communicate with the syslog clients, such as ESXi, vCenter Server, NSX for vSphere and NSX-T Data Center, using the TCP protocol.	Using the TCP syslog protocol ensures reliability and supports retry mechanisms. TCP syslog traffic is secure and more consistent with RFC 5424	TCP has a higher performance overhead compared to UDP.
SDDC-OPS-LOG-050	Include the syslog configuration for vRealize Log Insight in the host profile for the following clusters: <ul style="list-style-type: none"> ■ Management domain ■ Workload domains 	Simplifies the configuration of the hosts in the cluster and ensures that settings are uniform across the cluster.	Every time you make an authorized change to a host regarding the syslog configuration you must update the host profile to reflect the change or the status shows non-compliant.
SDDC-OPS-LOG-051	Configure the vRealize Log Insight agent on the Site Recovery Manager appliance.	Simplifies configuration of log sources in the SDDC that are pre- packaged with the vRealize Log Insight agent.	You must configure the vRealize Log Insight agent to forward logs to the vRealize Log Insight VIP.
SDDC-OPS-LOG-052	Install and configure the vRealize Log Insight agent on the vSphere Replication appliance.	Simplifies configuration of log sources in the SDDC that are pre- packaged with the vRealize Log Insight agent.	You must configure the vRealize Log Insight agent to forward logs to the vRealize Log Insight VIP.
SDDC-OPS-LOG-053	Do not configure vRealize Log Insight to automatically update all deployed agents.	Manually install updated versions of the Log Insight Agents for each of the specified components in the SDDC for precise maintenance.	You must maintain manually the vRealize Log Insight Agents on each of the SDDC components.

Encryption Design for vRealize Log Insight

Access to the vRealize Log Insight user interface and API require an SSL connection. By default, vRealize Log Insight uses a self-signed certificate. To provide secure access to the vRealize Log Insight user interface, replace the default self-signed certificate with a CA-signed certificate.

Table 2-132. Design Decisions on Encryption for vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS-LOG-054	Replace the default self-signed certificate of the virtual appliance of vRealize Log Insight with a CA-signed certificate.	Configuring a CA-signed certificate ensures that the communication to the externally facing UI and API for vRealize Log Insight, and cross-product, is encrypted.	Replacing the default certificates with trusted CA-signed certificates from a certificate authority might increase the deployment preparation time as certificates requests are generated and delivered.

Event Forwarding Between Regions with vRealize Log Insight

vRealize Log Insight supports event forwarding to other clusters and standalone instances. Use log forwarding between SDDC regions to have access to all logs if a disaster occurs in a region.

You forward syslog data in vRealize Log Insight by using the Ingestion API or a native syslog implementation. While forwarding events, the vRealize Log Insight instance still ingests, stores, and archives events locally.

The vRealize Log Insight Ingestion API uses TCP communication. In contrast to syslog, the forwarding module supports the following features for the Ingestion API:

- Forwarding to other vRealize Log Insight instances
- Support for both structured and unstructured data, that is, multi-line messages
- Metadata in the form of tags
- Client-side compression

Table 2-133. Design Decisions on Event Forwarding Across Regions in vRealize Log Insight

Decision ID	Design Decision	Design Justification	Design Implication
SDDC-OPS- LOG-055	Forward log events to the other region by using the Ingestion API.	<p>Supports the following operations:</p> <ul style="list-style-type: none"> ■ Structured and unstructured data for client-side compression ■ Event throttling from one vRealize Log Insight cluster to the other. <p>During a disaster recovery situation, the administrator has access to all logs from the two regions although one region is offline.</p>	<ul style="list-style-type: none"> ■ You must configure each region to forward log data to the other. The configuration introduces administrative overhead to prevent recursion of logging between regions using inclusion and exclusion tagging. ■ Log forwarding adds more load on each region. You must consider log forwarding in the sizing calculations for the vRealize Log Insight cluster in each region. ■ You must configure identical sizes on both source and destination clusters.
SDDC-OPS- LOG-056	Configure log forwarding to use SSL on port 9543.	Ensures that the log forward operations from one region to the other are secure.	<ul style="list-style-type: none"> ■ You must set up a custom CA- signed SSL certificate. ■ Event forwarding with SSL does not work with the self-signed certificate that is installed on the destination servers by default. ■ If you add more vRealize Log Insight nodes to a cluster in a region, the SSL certificate used by the vRealize Log Insight cluster in the other region must be installed in that the Java keystore of the nodes before SSL can be used.

Data Protection and Backup Design for vRealize Log Insight

To preserve the cloud operations services functionality when data or system loss occurs, the design supports the use of data protection.

vRealize Log Insight supports data protection through the creation of consistent image-level backups, using backup software that is based on the vSphere Storage APIs - Data Protection (VADP).

Disaster Recovery Design for vRealize Log Insight

Each region is configured to forward log information to the vRealize Log Insight instance in the other region.

Because of the forwarding configuration, an administrator of the SDDC can use either of the vRealize Log Insight clusters in the SDDC to query the available logs from one of the regions. As a result, you do not have to configure failover for the vRealize Log Insight clusters, and each cluster can remain associated with the region in which it was deployed.