

vCloud Availability for vCloud Director 2.0 Installation and Configuration Guide

vCloud Availability for vCloud Director 2.0



vmware®

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

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

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

docfeedback@vmware.com

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

Copyright © 2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

vCloud Availability for vCloud Director 2.0 Installation and Configuration Guide 4

Updated Information 5

1 Overview of vCloud Availability 7

- vCloud Availability Architecture 7
- vCloud Availability Components 8
- vCloud Availability Features 10
- vCloud Availability Portal Overview 11
- vCloud Availability Administration Portal Overview 12

2 Deployment of Components 14

- Production Deployment 15
- Firewall and Port Configuration 17
- Load Balancing 21
- Domain Name System Support 23

3 Service Provider Installation and Configuration 25

- Installation and Configuration Example 26
- Preparing Your Environment to Install vCloud Availability 26
- Automated Installation and Configuration 50
- Manual Installation and Configuration 62
- Post-Installation vCloud Director Configuration 88
- Unconfiguring vCloud Availability 90

4 Tenant Installation and Configuration 94

- Prepare Your Environment to Install vSphere Replication 94
- Deploy the vSphere Replication Virtual Appliance 95
- Register the vSphere Replication Appliance with vCenter Single Sign-On 97
- Update the vSphere Replication Appliances to Trust the vCloud Director Self-Signed Certificate in a Development Environment 100
- Configure Cloud Provider 101

vCloud Availability for vCloud Director 2.0 Installation and Configuration Guide

The *vCloud Availability for vCloud Director* Installation and Configuration Guide provides information on how to install and configure the vCloud Availability for vCloud Director 2.0 DRaaS solution on both the service provide and tenant sites.

Intended Audience

This information is intended for VMware Cloud Provider Program service providers and experienced system administrators who are familiar with virtual machine technology and data center operations including but not limited to the following areas:

- VMware vSphere[®]
- VMware vCloud Director[®]
- Virtual Infrastructure
- Secure Shell (SSH)
- Bash

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Updated Information

This *vCloud Availability for vCloud Director 2.0 Installation and Configuration Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *vCloud Availability for vCloud Director 2.0 Installation and Configuration Guide*.

Revision	Description
9 JUL 2018	Updated the information in topic Update the vSphere Replication Appliances to Trust the vCloud Director Self-Signed Certificate in a Development Environment .
20 JUN 2018	Updated the information in topic Create Cloud Proxy .

Revision	Description
22 FEB 2018	<ul style="list-style-type: none"> ■ Introduced <i>vCloud Availability</i> as the new short name for the <i>vCloud Availability for vCloud Director</i> solution. ■ The name of the <i>vCloud Availability for vCloud Director Portal</i> changes to <i>vCloud Availability Portal</i>. ■ The name of the <i>vCloud Availability for vCloud Director Service Manager Portal</i> changes to <i>vCloud Availability Administration Portal</i>. ■ Added the Deploy vCloud Availability Installer Appliance by Using the vSphere Web Client topic. ■ Moved topics Installing and Configuring Cassandra Servers and Installing and Configuring RabbitMQ Servers to chapter Preparing Your Environment to Install vCloud Availability. ■ Updated the information in the following topics: <ul style="list-style-type: none"> ■ Configuring vCloud Director ■ vCloud Availability Administration Portal Overview ■ Production Deployment ■ Firewall and Port Configuration ■ Download and Install RabbitMQ Servers ■ Configure a Primary RabbitMQ Server ■ Create Self-Signed Certificates for the Primary RabbitMQ Server ■ (Optional) Configure a RabbitMQ Cluster ■ (Optional) Registry Entry for a Docker Host ■ Registry Entry for a vCloud Availability Administration Portal Host ■ Defining Installation Variables ■ Update the vSphere Replication Appliances to Trust the vCloud Director Self-Signed Certificate in a Development Environment ■ (Optional) Update the vSphere Replication Appliances to Trust the Cloud Proxy Self-Signed Certificate ■ Replaced topics <i>Unregistering One or All vSphere Replication Solution Users</i> and <i>Unregistering the Cassandra Endpoints from the Lookup Service</i> with topic Unconfigure vCloud Availability. ■ Replaced topic <i>Install and Configure a Cassandra Server</i> with topics Installing and Configuring Cassandra Servers, Prepare Hosts for Cassandra Installation, Install Cassandra, Configure a Cassandra Cluster, Generate Certificates for Cassandra Nodes, and Enable Server and Client Communication with Cassandra over SSL.
19 OCT 2017	Initial release.

Overview of vCloud Availability

This section describes the core architecture of the vCloud Availability solution.

The architecture of the solution relies on the service provider environment that provides the replication target and the customer, or tenant, environment that employs vSphere Replication to move the data to the service provider. In the service provider environment, multiple components operate together to support replication, secure communication, and storage of the replicated data. Each service provider can support recovery for multiple customer environments that can scale to handle increasing loads for each tenant, and for multiple tenants.

On the tenant side, a single VM instance is deployed in the tenant vSphere environment. This deployment provides management service that is used to oversee the replication operation for each replicated VM. Standard vSphere Replication is used to exchange this information with the service provider infrastructure.

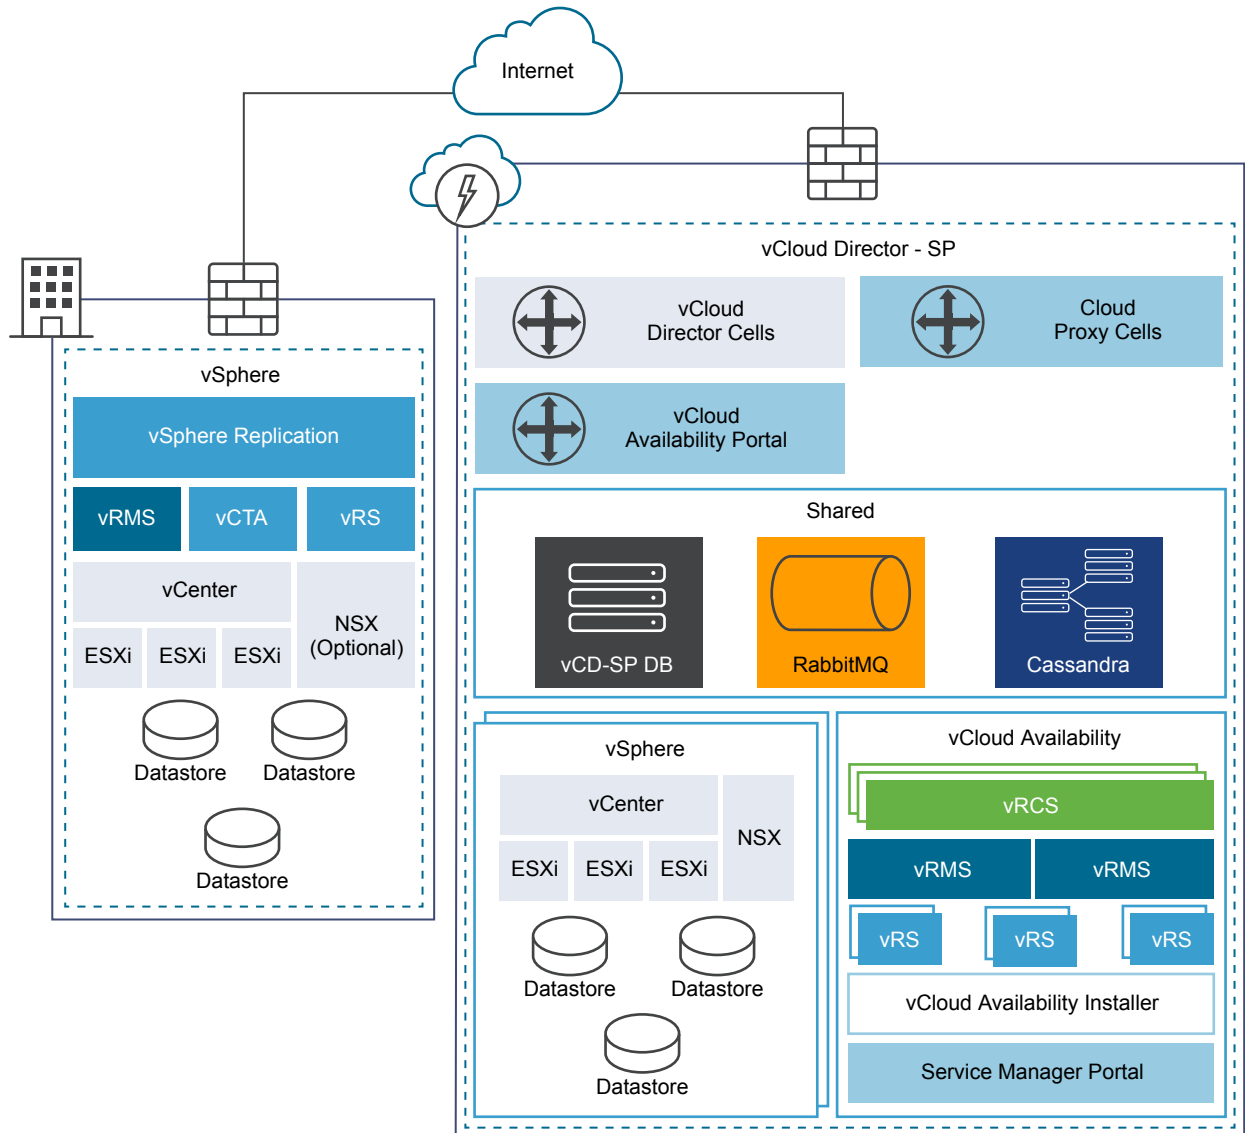
This chapter includes the following topics:

- [vCloud Availability Architecture](#)
- [vCloud Availability Components](#)
- [vCloud Availability Features](#)
- [vCloud Availability Portal Overview](#)
- [vCloud Availability Administration Portal Overview](#)

vCloud Availability Architecture

The vCloud Availability solution is composed of infrastructure and solution components on both service provider and tenant sides.

The gray cells in the following diagram represent existing components in the service provider and tenant environments. The remaining colored cells represent vCloud Availability components that you deploy during vCloud Availability installation and configuration procedures.



vCloud Availability Components

vCloud Availability consists of several different components which work together to provide the overall service.

Table 1-1. vCloud Availability Component Definitions.

Name	Abbreviation/Internal Name	Description
vSphere Replication Cloud Service	vRCS/HCS	A tenant-aware replication manager that provides the required API for managing the service and all the components. vSphere Replication Cloud Service registers as a VMware vCloud Director® extension enabling the functionality through the existing vCloud Director API.
vSphere Replication Manager	vRMS/HMS	The management server manages and monitors the replication process from tenant VMs to the service provider environment. A vSphere Replication Management Server runs for each vCenter Server and tracks changes to VMs and infrastructure related to replication.
vSphere Replication Server	vRS/HBR	The replication server receives and records delta information for each replicated VM. During to-cloud replication, delta information is sent by the tenant ESXi host and recorded by the provider vRS. During from-cloud replication, delta information is sent by the provider ESXi host and recorded by the on-premise vSphere Replication Server.
vCloud Tunneling Agent	vCTA	vCTA is a software component which supports tunneling functionality at the tenant data center. vCTA is responsible for orchestrating a secure tunnel creation for both to-the-cloud and from-the-cloud tunnels.
vCloud Director	vCD	With the vCloud Director solution service providers can build secure, multi-tenant private clouds by pooling infrastructure resources into virtual data centers and exposing them to users through Web-based portals and programmatic interfaces as fully automated, catalog-based services.
Cloud Proxy	n/a	Provides the vCloud Director endpoint for tunnels use to replicated data from tenant vCTA to and from vCloud Director.
Management vCenter Server	n/a	The Management vCenter Server environment is managed by the service provider and not accessible for tenants.
Resource vCenter Server	n/a	The Resource vCenter Server is a vCenter Server registered to vCloud Director and made available to tenants. Tenants do not have direct access to the Resource vCenter Server environment. Tenants can only locate workloads on the Resource vCenter Server instances using vCloud Director.
Tenant vCenter Server	n/a	The Tenant vCenter Server environment is used solely by the tenant users and is not connected to vCloud Director.
vCloud Availability Portal	n/a	The vCloud Availability Portal provides tenants with a graphic user interface to facilitate the management of the vCloud Availability solution. The vCloud Availability Portal also provides overall system and workload information.

Table 1-1. vCloud Availability Component Definitions. (Continued)

Name	Abbreviation/Internal Name	Description
vCloud Availability Administration Portal	n/a	The vCloud Availability Administration Portal provides a graphic user interface to facilitate the service providers to monitor and manage the vCloud Availability solution. The vCloud Availability Administration Portal also provides replications and IaaS consumption information.
VMware Platform Services Controller™	PSC	The Platform Services Controller provides common infrastructure services to the vSphere environment. Services include licensing, certificate management, and authentication with VMware vCenter® Single Sign-On.
Cassandra	C	Cassandra is a free and open-source distributed NoSQL database management system that stores metadata and supports storage of the metadata for replication services. Cassandra is used to store metadata about the replication, replicated VM instances, and infrastructure elements required to support the service. Cassandra is used as a fault-tolerant datastore.
RabbitMQ	n/a	An open source message broker that implements the Advanced Message Queuing Protocol (AMQP). When vSphere Replication Cloud Service registers as a vCloud Director extension, RabbitMQ is used to exchange information with vCloud Director.
Locator	n/a	The locator must be a valid path to be used with the VMware OVF Tool, as shown in the following examples. <ul style="list-style-type: none"> ■ <code>/datacenter-name/host/esx-name</code> ■ <code>/datacenter-name/host/cluster-name</code>
Datastore	n/a	The name of a vSphere datastore, accessible by the locator.

vCloud Availability Features

The vCloud Availability solution provides many features to support disaster recovery and migration services from an on-premises vCenter to a service provider cloud.

For the tenants, the vCloud Availability solution provides the following features:

- Self-service protection, failover, and failback workflows per virtual machine.
- Recovery point objective (RPO) from 15 minutes to 24 hours.
- Initial data seeding by shipping a disk.
- Full integration with the vSphere Web client.

For the service provider, the vCloud Availability solution:

- Integrates with existing vSphere environments.
- Provides multi-tenant support.

- Provides built-in encryption of replication traffic.
- Supports multiple vSphere versions.
- Supports multiple ESXi versions.
- Individual systems are isolated as virtual machine files.
- Provides simplified installation and configuration by using one single command.
- Provides automation through standard Web service APIs.

Failover from on-premises to Cloud

Replicates data from on-premises vSphere workloads to service provider cloud environments. After the virtual machines are replicated, failover support for running the workloads in the cloud. Recovery Point Objective (RPO) can be configured from 5 minutes to 24 hours.

Failback to on-premises

For failover workloads that have been migrated to the cloud, changes can be replicated back to the on-premise environment. You can then fail back workloads in the on-premise environment.

Multiple Points In Time (MPIT) Recovery

Up to 24 restore points can be created. Depending on the RPO configuration, restoration is available from any recovery point.

Orchestration

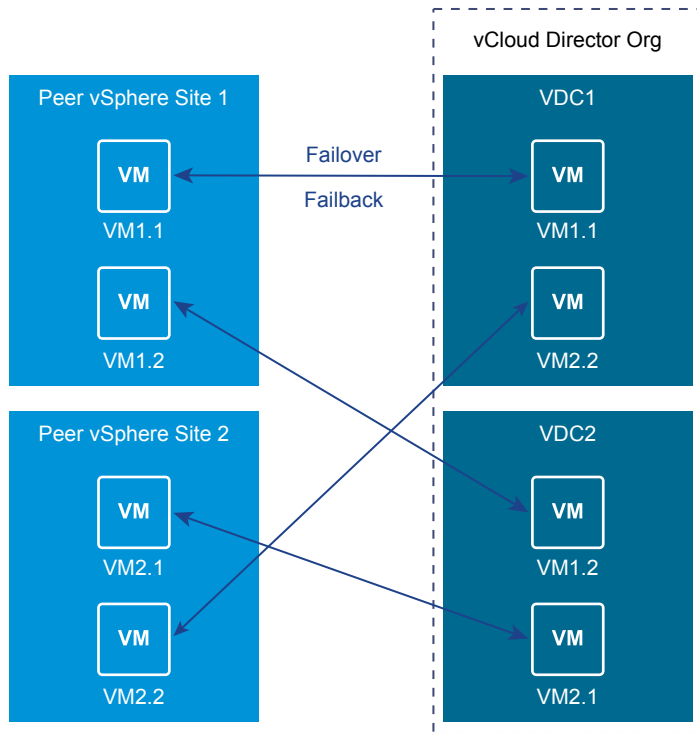
Using VMware vRealize[®] Orchestrator[™] Appliance and plug-in for vSphere Replication you can easily design and deploy scalable workflows that automate complex IT processes.

vCloud Availability Portal Overview


The vCloud Availability Portal provides tenants with a graphic user interface to facilitate the management of the vCloud Availability solution.

The vCloud Availability Portal uses vCloud Director APIs to perform common vCloud Availability tasks. The vCloud Availability Portal also provides overall system and workload information.

The first step in the disaster recovery workflow is protecting the VMs from your vSphere tenant environment into your vCloud Director cloud environment. You can replicate the protected VMs from multiple vSphere tenant environments to multiple vCloud Director VDC instances, and the reverse. The following diagram illustrates the replication topology for the use case.



You can fail over a virtual machine from one vSphere tenant environment to a vCloud Director cloud site and failback the same virtual machine to a different peer vSphere tenant environment. After a successful failover, you detach the workload by using the vCloud Availability Portal, and later from the second peer vSphere environment you can configure a reverse replication from cloud. The following diagram illustrates the replication topology for the use case.



vCloud Availability Administration Portal Overview

The vCloud Availability Administration Portal provides a graphic user interface to facilitate the service providers to monitor and manage the vCloud Availability solution.

The service providers can use the vCloud Availability Administration Portal as a dashboard providing information about replications, and as a tool to clean up stale replications and migrate replications from one datastore to another.

Every 12 hours, the vCloud Availability Administration Portal fetches data from vCloud Director. This data is stored into the local vCloud Availability Administration Portal database. The status information on the **Home** page and **Reports** page is extracted from the local database.

If you add a new tenant or an existing tenant creates a new replication, the vCloud Availability Administration Portal will display the new information after the next automatic synchronization between the local database and vCloud Director. You can force the synchronization and get real time data by clicking **GENERATE** on the **Home > Report** page.

If you need to move a replication to a different datastore or if you Scrub Replications, the vCloud Availability Administration Portal starts fetching real time data from vCloud Director to complete the operation.

By using the impersonation feature, the service providers can pose as tenants and perform all DR tasks in the tenant vCloud Availability Portal.

Deployment of Components

Deployment involves installing components within the service provider environment, and just one component within each tenant vCenter Server environment.

Deployment of vCloud Availability is based on configuring several different components. Some components, such as Cassandra, or NSX may already be deployed within your environment. All the components work together to provide the overall service. Some components are required only once, others are required two or three times to provide redundancy, some are required multiple times to support an increasing number of protected virtual machines.

A typical deployment includes the following components:

- **Tenant Service**
 - The tenant service, which consists of the vSphere Replication Appliance, is installed on-site in an existing vSphere environment and provides the necessary tools to replicate information to the service provider site.
- **Service Provider Service**
 - The service provider side of the system supports one or more tenants. A Cloud Proxy provides network connectivity, a vSphere Replication Manager handles the replication of data and controls the replication, and a cluster of appliances receives the disk updates and stores the information ready for the failover of operation from the tenant site to the service provider environment. The number of instances of each component required varies depending on the number of VMs on the tenant side that need to be protected.
- **Production Deployment**

Production deployments of vCloud Availability have specific sizing and component configurations.
- **Firewall and Port Configuration**

Several firewall network ports are required to be used between different components and systems.
- **Load Balancing**

Some components of vCloud Availability require external load balancer.
- **Domain Name System Support**

You can configure and work with the vCloud Availability solution using IP addresses or domain names.

Production Deployment

Production deployments of vCloud Availability have specific sizing and component configurations.

A production deployment uses multiple components to support many protected virtual machines, and to provide a fault-tolerance within the DRaaS environments.

Production Architecture

Production deployments must meet certain requirements.

- At each tenant site, there is one or more single-tenant environments to be protected.
- In the service provider disaster recovery site, one or more vCloud Director is configured with a specific number of components designed to handle the required number of VMs from each tenant.
- A single vCloud Director environment in a data center hosts up to 1000 individual tenants.

Using this information as a base, you must install and configure a new vCloud Availability service instance for every 100 tenants that use the service.

In a single vCloud Director deployment, there is a limit to the number of VMs that can be replicated as part of the DR solution. The exact combination depends on the number of VMs that must be supported combined with the system limits for each component.

Component Sizing

Individual components have a minimum installation count required for a base installation.

Table 2-1. Relative Component Sizing

Component	Related Component
vCloud Director	1 Management vCenter Server Instance 1 Resource vCenter Server Instance
vSphere Replication Manager	4 vSphere Replication Server

Component Limits

Individual components have limits for the maximum number of supported services, instances, or connections required.

Table 2-2. Component Counts and Limits per Instance

Component	Limit
Cloud Proxy	2000 Connections (vSphere Replication control connections and replicated VM payload)
vSphere Replication Server	250 Active Replications
Tenants	1000 per vCloud Director

Table 2-2. Component Counts and Limits per Instance (Continued)

Component	Limit
vCloud Director	15 vCenter Server Instances
vCloud Availability Portal	800 Concurrent Sessions

Sample Deployment Scaling

Using the information on sizing and configurations, for a single instance, supporting up to the maximum of 100 tenants.

Table 2-3. Component Counts for Production Deployments

Protected VMs	500	1000	2000	3000	5000	10000
vSphere Replication Server	2	4	8	12	20	40
vSphere Replication Cloud Service	2	2	3	3	3	3
Cloud Proxy	2	2	2	2	3	5

Minimum High Availability Component Counts

To sustain a highly available vCloud Availability solution, consider the following counts per component.

Table 2-4. Component Counts for Highly Available vCloud Availability Solution

Component	Number of Instances
vCloud Director	2
vCenter Server	1 Management vCenter Server Instance with 2 ESXi hosts. 1 Resource vCenter Server Instance with 2 ESXi hosts.
External Platform Services Controller (PSC)	1
vSphere Replication Manager	1
vSphere Replication Cloud Service host	2
vSphere Replication Server	4
Cloud Proxy	2
vCloud Availability Portal	2
vCloud Availability Administration Portal	1
Cassandra Host	3
RabbitMQ Host	2

Network Configuration Best Practice

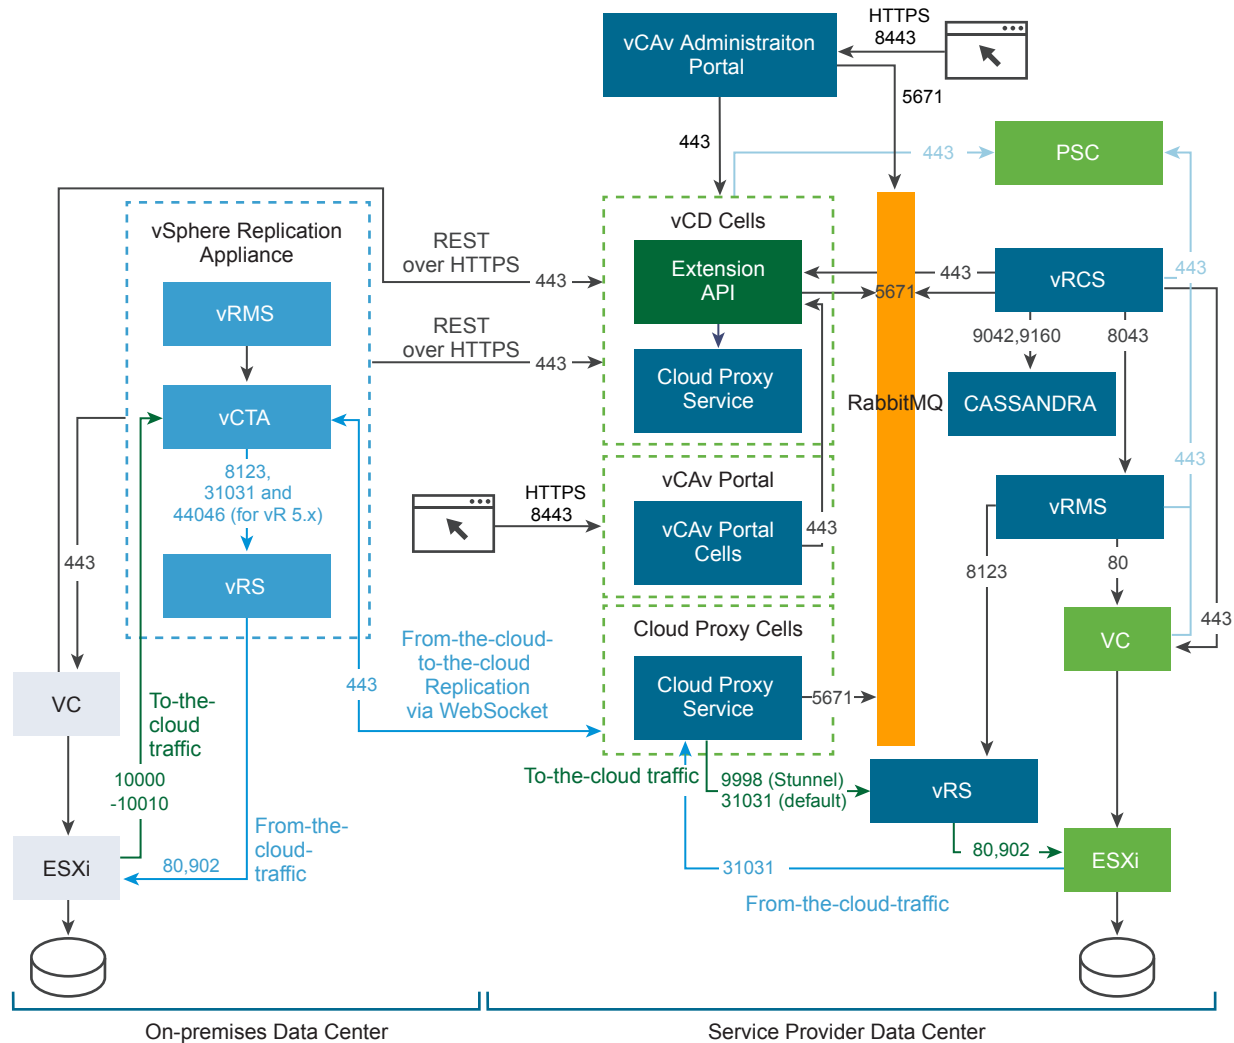
Within a production deployment, the underlying network architecture is important. The following describes a best practice network configuration.

- Each underlying physical ESXi installation is configured with a VMXNET3 high-speed network adapter, connected to two separate 10Gbe switches using NIC teaming.
- The two switches are connected to each other using two 40gbe QSFP cables.
- The switches are configured to present a VLAN that is configured as a port group on the ESXi hosts.
- The virtual machines that make up the environment are all configured in this flat broadcast domain.

Firewall and Port Configuration

Several firewall network ports are required to be used between different components and systems.

The following diagram shows the flow of network ports and data through a typical deployment on both the service provider and tenant side.



Note To avoid connectivity issues caused by firewall restrictions, make sure that vSphere Replication Server instances can connect to all ESXi hosts in a cluster that hosts replication workloads.

The following table provides a list of ports to be used between the different systems and components.

Table 2-5. Firewall Port Component Configurations within a Service Provider Deployment

Source	Destination	Port Number	Protocol or Description
Cloud Proxy	vCloud Director DB	1433 or 1521	TCP Port 1433 is the default Microsoft SQL Server database port. Port 1521 is the default Oracle database port.
Cloud Proxy	RabbitMQ	5671	AMQP

Table 2-5. Firewall Port Component Configurations within a Service Provider Deployment (Continued)

Source	Destination	Port Number	Protocol or Description
Cloud Proxy	vSphere Replication Server	31031	Initial and ongoing replication traffic
ESXi	Cloud Proxy	31031	Initial and ongoing replication traffic
External	Cloud Proxy	443	Initial and ongoing replication traffic
Web Browser	vSphere Replication Manager	5480	Virtual Appliance Management Interface (VAMI) Web UI. Administrator's Web browser.
Web Browser	vSphere Replication Server	5480	Virtual Appliance Management Interface (VAMI) Web UI. Administrator's Web browser.
vCloud Director	RabbitMQ	5671	Default RabbitMQ port. AMQP, API Extensibility, Notifications.
Cloud Proxy	vCloud Director	61616, 61611	JMS
vCloud Director	Cloud Proxy	61616, 61611	JMS
vCloud Director	PSC	7444, 443	SOAP
vSphere Replication Manager	PSC	7444, 443	SOAP
vSphere Replication Cloud Service	PSC	7444, 443	SOAP
vSphere Replication Cloud Service	vCenter Server	80	HTTP
vSphere Replication Cloud Service	vCloud Director	443	HTTP
vSphere Replication Cloud Service	vCenter Server	443	SOAP
vSphere Replication Cloud Service	RabbitMQ	5671	Default Rabbit MQ port. AMQP
vSphere Replication Cloud Service	vSphere Replication Manager	8043	SOAP
vSphere Replication Cloud Service	Cassandra	9042	Default Cassandra port. CQL Native Transport Port
vSphere Replication Cloud Service	Cassandra	9160	Default Cassandra port. Thrift.
vSphere Replication Manager	vCenter Server	80	SOAP
vSphere Replication Manager	vCenter Server	443	SOAP

Table 2-5. Firewall Port Component Configurations within a Service Provider Deployment (Continued)

Source	Destination	Port Number	Protocol or Description
vSphere Replication Manager	vSphere Replication Server	8123	SOAP
vSphere Replication Server	ESXi	80	SOAP
vSphere Replication Server	ESXi	902	NFC
vCloud Availability Installer Appliance	vCloud Director	443	HTTPS
vCloud Availability Installer Appliance	PSC	7444, 443	SOAP
vCloud Availability Installer Appliance	Resource vCenter Server	80, 7444, 443	SOAP
vCloud Availability Installer Appliance	Resource ESXi	443	SOAP
vCloud Availability Installer Appliance	Management vCenter Server	80, 7444, 443	SOAP
vCloud Availability Installer Appliance	Management ESXi	443	SOAP
vCloud Availability Installer Appliance	RabbitMQ	5671	AMQP
vCloud Availability Installer Appliance	Cassandra	22, 9042	CQL Native Transport Port SSH
vCloud Availability Installer Appliance	vSphere Replication Manager, vSphere Replication Cloud Service, vSphere Replication Server	22, 5480, 8043	Replication Management service. Administrator's Web browser.
vCloud Availability Installer Appliance	vCloud Availability Portal, vCloud Availability Administration Portal	22, 8443	SSH, HTTPS
vCloud Availability Portal	vCloud Director	443	HTTPS
vCloud Availability Administration Portal	vCloud Availability Portal	8443	HTTPS
vCloud Availability Administration Portal	vCloud Director	443	SOAP
vCloud Availability Administration Portal	RabbitMQ	5671	AMQP

For the deployment within a tenant, configure the following ports.

Table 2-6. Firewall Port Configurations within a Tenant Environment

Source	Destination	Port Number	Protocol or Description
vSphere Replication Appliance	vCenter Server	80	SOAP
vSphere Replication Server	ESXi	80	SOAP
vSphere Replication Server	ESXi	902 (TCP and UDP)	NFC
Web Browser	vSphere Replication Appliance	5480	Administrator's Web browser. VAMI.
vSphere Replication Server	vSphere Replication Appliance	8043	SOAP
Web Browser	vSphere Replication Server	5480	Administrator's Web browser. VAMI.
vSphere Replication Appliance	vSphere Replication Server	8123	SOAP
vSphere Replication Appliance	vCloud Availability Public Endpoints	443	HTTP
vCenter Server	vCloud Availability Public Endpoints	443	HTTP
ESXi	vSphere Replication Appliance at Service Provider	10000- 10010	Replication traffic, local vSphere Replication Appliance
Web Browser	vCloud Availability Portal	8443	Default vCloud Availability Portal port. HTTPS
vSphere Replication Appliance	External vSphere Replication Server	31031, 44046	TCP

Load Balancing

Some components of vCloud Availability require external load balancer.

External load balancer is not provided as part of the vCloud Availability solution. You can use virtual NSX Edge load balancer or third-party virtual or physical load balancers.

- [To-the-Cloud Traffic](#)

To support a large number of clients connecting to a public listening port, Service Providers running vCloud Availability can deploy multiple instances of Cloud Proxy and a load balancer to distribute requests to a public port across Cloud Proxy instances.

- [From-the-Cloud Traffic](#)

To establish a from-the-cloud connection, Cloud Proxy sends a request to the vSphere Replication Appliance on premises, using the vCloud Tunneling Agent control connection.

- [Using Layer 4 Load Balancing](#)

Layer 4 load balancers operate on the transport layer, using information defined at the networking transport layer as the basis for client request distribution decisions.

To-the-Cloud Traffic

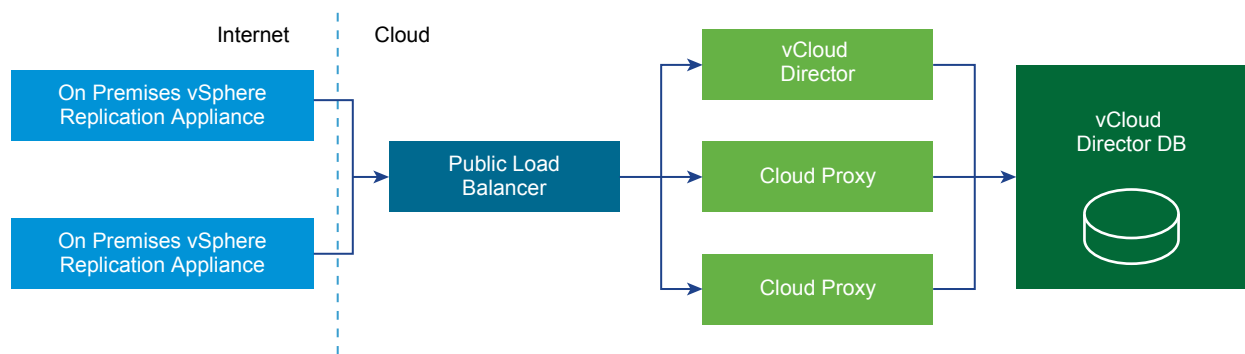
To support a large number of clients connecting to a public listening port, Service Providers running vCloud Availability can deploy multiple instances of Cloud Proxy and a load balancer to distribute requests to a public port across Cloud Proxy instances.

To-the-Cloud traffic setup has the following specifications .

- All Cloud Proxy instances must use the same persistence database as the rest of the vCloud Director instances
- All Cloud Proxy instances must use the same transfer share as the rest of the vCloud Director instances
- All Cloud Proxy instances must use the same NTP time source as the rest of the components (recommended to use an internal NTP source)
- The load balancer exposes a single public port
- Public ports on the Cloud Proxy instances are not exposed to internet
- Cloud Proxy instances do not require to be aware of the load balancer

The Cloud Proxy load balancer can be the same load balancer that is used to distribute REST API requests among vCloud Director instances. In case of high load, a best practice is to have dedicated load balancers for each replication direction, as the traffic can come from Internet and from the cloud .

Figure 2-1. Cloud Proxy Load Balancing Deployment Model



Public Cloud Proxy endpoint (URI) for to-the-cloud tunnel termination and internal IP address for from-the-cloud traffic (used by ESXi host-based replication) must be specifically configured in vCloud Director using the vCloud Director API call . For more information about Cloud Proxy configuration, see [Create Cloud Proxy](#).

From-the-Cloud Traffic

To establish a from-the-cloud connection, Cloud Proxy sends a request to the vSphere Replication Appliance on premises, using the vCloud Tunneling Agent control connection.

In response the vCloud Tunneling Agent initiates a from-the-cloud connection back to the Cloud Proxy. There is only one control connection per tenant (source site), so only one instance of Cloud Proxy can send requests to a particular vCloud Tunneling Agent. Since the load balancer can forward connections to any instance of the Cloud Proxy, the particular instance serving the client request may not have access to the control connection corresponding to the (source site). Furthermore, once from-the-cloud connection is established, there is no guarantee that this connection will be forwarded to the same instance of Cloud Proxy serving the tenant (source site).

The recommended solution consists of the following two elements.

- Using an internal message bus to request from-the-cloud connection
- Configure L4 or L7 load balancing mechanism (Service Providers are expected to configure the load balancer mechanism accordingly)

Using Layer 4 Load Balancing

Layer 4 load balancers operate on the transport layer, using information defined at the networking transport layer as the basis for client request distribution decisions.

A Layer 4 load balancer delivers messages with no regard to the content of the message. Such load balancers only forward network packets to and from an upstream host without inspecting the content of the packets.

If you use a Layer 4 load balancer, the solution requires a load balanced Virtual IP address for each Cloud Proxy instance, with a specific override of its public fully qualified domain name (FQDN) to a specific one .

To override, you must add the following property to the `/opt/vmware/vcloud-director/etc/global.properties` file. `cloudproxy.reverseconnection.fqdn = name.com:443`

Layer 4 load balancing advantages

- Better **performance** compared to Layer 7 load balancing. Layer 4 load balancers only use the network transport layer information, disregarding the content of the packages .
- Better **throughput**.

Layer 4 load balancing disadvantages

- You need multiple Virtual IP addresses.
- You need multiple IP addresses.
- You need multiple certificates.
- You need multiple one member IP pools.

Domain Name System Support

You can configure and work with the vCloud Availability solution using IP addresses or domain names.

To use Domain Name System (DNS), you must register the domain names with a DNS server. This DNS server must be accessible to the vCloud Availability appliances that are deployed in the service provider data center.

If you use static IP addresses, the vCenter Server IP pool must have a valid DNS server.

If you use Dynamic Host Configuration Protocol (DHCP) for IP address management, the DHCP server must provide a valid DNS server for each deployed appliance.

If you are utilizing the automated vCloud Availability deployment capability, you do not assign a hostname to VMs you deploy by default. To assign specific hostname to VMs you deploy with the registry file, add the `hostname` property and value in the respective sections of the registry file. For more information, see [Creating a Registry File for an Automated Installation](#).

If you are deploying vCloud Availability manually, you use the `vcav component-alias create` command to deploy all vCloud Availability components. By default, the value of the `--vm-name=` argument defines the hostname of the virtual machine you are creating. You can assign a specific hostname to all VMs you deploy by adding the `--hostname=desired-hostname` argument to the `vcav component-alias create` command.

Service Provider Installation and Configuration

3

vCloud Availability has multiple components that you must install and configure in a specific order. You can select between an automated and a manual installation and configuration.

In environments where the service provider and tenant workloads run on the same network, vCloud Availability components are installed, configured, and operate on a single vCloud Director instance.

Environments where the service provider and the tenant networks are separated are also supported. In such environments, the vSphere Replication appliances are deployed on the tenant, resource side. The remaining components are deployed in the service provider management environment.

1 [Installation and Configuration Example](#)

In the current example, the deployment environment consists of one management vSphere and two resource vSphere sites.

2 [Preparing Your Environment to Install vCloud Availability](#)

You must prepare your environment before installing the vCloud Availability solution.

3 [Automated Installation and Configuration](#)

The automated installation and configuration of the vCloud Availability solution components is repeatable, time saving, faster and less error prone compared to the manual deployment. You can easily adopt it to brownfield and mixed environments.

4 [Manual Installation and Configuration](#)

You can install and configure the vCloud Availability solution by running a set of commands in a step-by-step manner.

5 [Post-Installation vCloud Director Configuration](#)

After you install and configure the vCloud Availability solution, you must perform some configurations to vCloud Director.

6 [Unconfiguring vCloud Availability](#)

You unconfigure a vCloud Availability instance by using the vCloud Availability Installer Appliance scripts and the vSphere Web Client capabilities.

Installation and Configuration Example

In the current example, the deployment environment consists of one management vSphere and two resource vSphere sites.

The management vSphere details refer to the path of the environment managed by the service providers that is not available to tenant users. The resource vSphere details relate to the path of the environment that tenants use.

The resource vSphere sites are part of the same SSO domain to which the vCloud Director host is federated.

The management vSphere hosts the vSphere Replication Cloud Service and the vCloud Availability Portal.

A resource vSphere hosts the vSphere Replication Manager and the vSphere Replication Server.

For test and development environments, if you use docker to manage your Cassandra and RabbitMQ hosts, the commands in the current example place the docker host in the management vSphere environment.

If necessary, you can host all the components in the management vSphere.

Preparing Your Environment to Install vCloud Availability

You must prepare your environment before installing the vCloud Availability solution.

Procedure

1 [Creating vCloud Availability Installer Appliance](#)

To begin installing the components of vCloud Availability, you must first install and configure the vCloud Availability Installer Appliance.

2 [Configuring vCloud Director](#)

vCloud Director must be configured to support environments that can securely support multiple tenants.

3 [Create Cloud Proxy](#)

The Cloud Proxy is a standalone, optional component of vCloud Director that can act as a generic Transmission Control Protocol (TCP) connection proxy. It supports forwarding incoming TCP connections and listening incoming connections.

4 [Installing and Configuring Cassandra Servers](#)

Cassandra is a free and open-source distributed NoSQL database management system that stores metadata and supports storage of the metadata for replication services.

5 [Installing and Configuring RabbitMQ Servers](#)

RabbitMQ is an AMQP server that exchanges messages within a vCloud Director environment. You can install and configure up to three RabbitMQ servers on CentOS hosts. For high availability and scalability purposes, you can configure the RabbitMQ servers in a cluster.

Creating vCloud Availability Installer Appliance

To begin installing the components of vCloud Availability, you must first install and configure the vCloud Availability Installer Appliance.

The vCloud Availability Installer Appliance is deployed as an OVA file and includes the following components:

- vCloud Availability scripts for installation and maintenance operations
- SLES 12 SP3 image to provide a Docker container hosting

Installation and configuration procedures contain long commands with multiple arguments which you must run as single commands. For better visibility line breaks are marked with backslash (\) within a command. The beginning of a new command is marked with the number sign (#).

- [Deploy vCloud Availability Installer Appliance by Using the OVF Tool](#)

This procedure demonstrates how to deploy and configure a vCloud Availability Installer Appliance by using the VMware OVF Tool. Alternatively, you can use the vSphere Web Client to install the vCloud Availability Installer Appliance.

- [Deploy vCloud Availability Installer Appliance by Using the vSphere Web Client](#)

This procedure demonstrates how to deploy and configure a vCloud Availability Installer Appliance by using the vSphere Web Client.

- [Working with the vCloud Availability Installer Appliance](#)

You can use vCloud Availability Installer Appliance scripts to install, configure, and manage the vCloud Availability.

Deploy vCloud Availability Installer Appliance by Using the OVF Tool

This procedure demonstrates how to deploy and configure a vCloud Availability Installer Appliance by using the VMware OVF Tool. Alternatively, you can use the vSphere Web Client to install the vCloud Availability Installer Appliance.

Prerequisites

Download the vCloud Availability Installer Appliance.

- 1 In a Web browser, navigate to the product download page.
- 2 Download the `vcloud-availability-release_number-xxx-build_number_OVF10.ova` file.

Procedure

1 Define deployment variables.

The `VSPHERE_LOCATOR` value contains the target data center name, the tag `host`, the name of the target cluster, and the IP address or the fully qualified domain name (FQDN) of the target ESXi host. The `VSPHERE_LOCATOR` value depends on the topology of your vSphere environment. Following are examples for valid `VSPHERE_LOCATOR` values.

- `/data-center-name/host/cluster-1-name/fully-qualified-domain-name`
- `/data-center-name/host/cluster-2-name/host-IP-address`

If the target ESXi host is not part of a cluster, skip the `cluster-name` element, as shown in the following examples.

- `/data-center-name/host/fully-qualified-domain-name`
- `/data-center-name/host/host-IP-address`

The `VSPHERE_DATASTORE` value is the datastore name as it is displayed in the vSphere Web Client.

For more information about the `VSPHERE_LOCATOR` and `VSPHERE_DATASTORE` values, see *Specifying the Inventory Path for a Cluster, Host, or Resource Pool* in the [OVF Tool User's Guide](#).

```
# OVA_VM_NAME=vcav-installer-name
# VSPHERE_LOCATOR="vsphere-locator"
# VSPHERE_DATASTORE="vsphere-datastore"
# VSPHERE_ADDRESS=vsphere-ip-address
# VSPHERE_USER=vsphere-admin-user
# VSPHERE_NETWORK="VM-Network"
# OVA=local_client_path/vcloud-availability-installer-appliance-release_number-xxx-build_number.ova
# ROOT_PASSWORD=vcloud-availability-installer-appliance-root-password
```

2 Deploy vCloud Availability Installer Appliance OVA.

Note Password authentication is the default method for deploying the vCloud Availability Installer Appliance. You can deploy the appliance using SSH key authentication by adding the `"--prop:guestinfo.cis.appliance.root.sshkey=${SSH_KEY}"` argument in the installation command. You also must have a valid SSH public key to deploy vCloud Availability Installer Appliance using SSH key authentication method.

The following is a long, single command that should be run as one. There are breaks for better visibility marked with backslash (\).

```
# ovftool \
--acceptAllEulas \
--skipManifestCheck \
--X:injectOvfEnv \
--allowExtraConfig \
--X:enableHiddenProperties \
--sourceType=OVA \
--allowAllExtraConfig \
--powerOn \
--X:waitForIp \
"--net:VM Network=${VSPHERE_NETWORK}" \
--diskMode=thin \
--datastore=${VSPHERE_DATASTORE} \
--name=${OVA_VM_NAME} \
--prop:vami.hostname=${OVA_VM_NAME} \
--prop:guestinfo.cis.appliance.ssh.enabled=True \
"--prop:guestinfo.cis.appliance.root.password=${ROOT_PASSWORD}" \
${OVA} \
"vi://${VSPHERE_USER}@${VSPHERE_ADDRESS}${VSPHERE_LOCATOR}"
```

The system prints the IP address of the vCloud Availability Installer Appliance. Write down the IP address, because you are going to use it during the installation.

3 Create an SSH connection to the vCloud Availability Installer Appliance.

```
# ssh root@vCloud-Availability-Installer-Appliance-IP-Address
```

Deploy vCloud Availability Installer Appliance by Using the vSphere Web Client

This procedure demonstrates how to deploy and configure a vCloud Availability Installer Appliance by using the vSphere Web Client.

Prerequisites

- Download the vCloud Availability Installer Appliance.
 - a In a Web browser, navigate to the product download page.
 - b Download the `vcloud-availability-release_number-xxx-build_number_OVF10.ova` file.
- Download and install the Client Integration Plug-in through the **Deploy OVF Template** option in the vSphere Web Client if you deploy an OVF template for the first time.

Procedure

- 1 Log in to the vSphere Web Client.

- 2 Right-click the target location where the vCloud Availability Installer Appliance will be deployed (data center, folder, cluster, resource pool or host) and select the **Deploy OVF Template** option from the drop-down menu.

The **Deploy OVF Template** wizard is present.

- 3 Enter the source URL or browse for the vCloud Availability Installer Appliance location in the **Select source** page.

The syntax for the OVF filename is `vcloud-availability-release_number-xxx-build_number_OVF10.ova`.

- 4 Click **Next** and review the details.
- 5 Read and accept the license agreement, and click **Next**.
- 6 In the **Select name and folder** page, specify a name for the vCloud Availability Installer Appliance and select the data center or data center folder that contains the host or cluster on which you want to deploy the appliance, and click **Next**.
- 7 In the **Select a resource** page, select the target host or cluster where the deployed vCloud Availability Installer Appliance will run and click **Next**.
- 8 Select the virtual disk format and the storage policy for the virtual machine from the drop-down menu.
Thick Provision Lazy Zeroed and Datastore Default are selected by default.
- 9 Specify the target location for storing the virtual machine configuration files and virtual disks from the list of available datastores in the **Select Storage** page, and click **Next**.
- 10 Specify the settings for connecting the vCloud Availability Installer Appliance to the network in the **Setup network** page.
 - a Select the source network the appliance installer will use.
VM Network is used by default.
 - b Select the protocol version for the appliance IP address.
IPv4 is selected by default.
 - c Choose how to allocate the IP address of the vCloud Availability Installer Appliance.
Static – Manual is used by default.

Option	Description
Static - Manual	IP addresses are manually configured. No automatic allocation is performed.
Transient - IP Pool	IP addresses are automatically allocated using IP pools from a specified range when the vApp is powered on. The IP addresses are released when the appliance is powered off.
DHCP	A DHCP server is used to allocate the IP addresses. The addresses assigned by the DHCP server are visible in the OVF environments of virtual machines started in the vApp.
Static - IP Pool	IP addresses are automatically allocated from the managed IP network range of vCenter Server at power-on, and remain allocated at power-off.

11 Customize the deployment properties of the vCloud Availability Installer Appliance in the **Customize template** page and click **Next**.

- a Select the check box to enable SSH.
- b In the NTP Server section, enter the NTP server address the appliance installer uses.
For example, **ntime.vmware.com**.
- c Specify the port for the system logs in the Syslog Port section.
Port **516** is set by default.
- d In the System hostname section, type the hostname to use for the appliance installer.
- e In the root password section, enter and confirm the password of the root user for the appliance installer.
- f Specify the Networking Properties for the appliance installer if DHCP is not used.

12 Review all the settings configured for the vCloud Availability Installer Appliance and click **Finish** to start the OVF deployment process.

The wizard completes.

The **Recent Tasks** page shows the status for initializing the OVF deployment on the target host.

Working with the vCloud Availability Installer Appliance

You can use vCloud Availability Installer Appliance scripts to install, configure, and manage the vCloud Availability.

vCloud Availability Installer Appliance Basic Operations

You must create an SSH connection to the vCloud Availability Installer Appliance, to use the available scripts.

You run all installation and configuration commands from the vCloud Availability Installer Appliance, unless documentation instructs otherwise.

You must run vCloud Availability Installer Appliance commands from the **root** user's home directory of your vCloud Availability Installer Appliance.

The vCloud Availability Installer Appliance uses the following command-line syntax logic.

```
vcav [OPTIONS] COMMAND SUBCOMMAND [ARGUMENT]
```

All arguments that end with an equals sign (=) require a value.

Asterisk (*) marks required arguments for a command.

You can see the basic vCloud Availability Installer Appliance options and arguments in the following table.

Table 3-1. Basic vCloud Availability Installer Appliance Options and Arguments

Option	Argument	Description
--help (-h)	None	Displays a summary of options and arguments.
--session-dir=	File path	Define the location to store files and create a registry file.
--log-level=	Error, Warning, Info, and Debug	Define the level for logging in to the session directory.
--debug (-d)	None	Displays a DEBUG message in the vCloud Availability Installer Appliance console and creates an entry with the same information in the log file. By default, the value is set to <code>False</code> . You can append the option to every command that you run on the vCloud Availability Installer Appliance.
--info (-i)	None	Displays an INFO message to the vCloud Availability Installer Appliance console and creates an entry with the same information in the log file. By default, the value is set to <code>False</code> . You can append the option to every command that you run on the vCloud Availability Installer Appliance.
--ssh-cert=	File path	Use the option to provide the path to a private SSH certificate.
--registry=	File path	Use the option to provide the path to a registry file.
--dry-run	None	Use this command to validate a process without running the command. By default, the value is set to <code>False</code> . You can append the option to every command that you run on the vCloud Availability Installer Appliance.
-k	None	<p>Caution With this option, you can use SSL and SSH connections without a certificate validation.</p> <p>By default, the value is set to <code>False</code>. You can append the option to every command that you run on the vCloud Availability Installer Appliance.</p>
--non-interactive (-n)	None	Use the option to disable interactive prompts. By default, the value is set to <code>False</code> .

Configuring vCloud Director

vCloud Director must be configured to support environments that can securely support multiple tenants.

The vCloud Director environment must be fully configured to support workloads before you can continue with the vCloud Availability installation. You must create the Resource vSphere, Provider VDCs, Organizations, and Organization VDCs before installing the vCloud Availability solution. For more information, see the [vCloud Director Documentation](#).

vCloud Director is configured to use the following settings:

- Enabled Public URL and certificates. For more information, see the *vCloud Director Administrator's Guide*.
- Shared vCenter Single Sign On domain. vCloud Director must be federated with the same vCenter Single Sign On domain as the vCenter Server instances that are part of the vCloud Availability solution. For more information, see *Configure vCloud Director to use vCenter Single Sign On* topic in the *vCloud Director Administrator's Guide*.
- By default, vCloud Availability 2.0 supports the use of TLS 1.2 during the SSL handshake process. To build a pure TLS 1.2 environment for vCloud Availability operations, the vSphere Replication and vCenter Server instances that are deployed on-premise must also support TLS 1.2. For more information, select the *Transport Layer Security* category from the drop-down in the [Interoperability Pages](#) for vCloud Availability 2.0.

- Configure the timeout settings for the vCloud Director extensions.

Use a text editor to open `/opt/vmware/vcloud-director/etc/global.properties` file. Set the `extensibility.timeout` value to 60.

- The following files are required to create and configure a Cloud Proxy:
 - Copy of the `/opt/vmware/vcloud-director/etc/responses.properties` file from an existing vCloud Director instance.
 - A `certificates.ks` file that contains a wild card certificate(`*.provider.com`) with `http` and `consoleproxy` aliases.

Create Cloud Proxy

The Cloud Proxy is a standalone, optional component of vCloud Director that can act as a generic Transmission Control Protocol (TCP) connection proxy. It supports forwarding incoming TCP connections and listening incoming connections.

By default, the Cloud Proxy can create virtual connections for data to travel from the tenant (on-premise) site to the service provider (cloud) site and reverse. A vCloud Director instance runs Cloud Proxy services on the same Java Virtual Machine (JVM). For scalability purposes, vCloud Director appliances can be configured to act as a different type of cell, each one with its own JVM. For example, you can configure a vCloud Director appliance to act as an Application cell, as a Cloud Proxy cell, or as a combination of both types on the same cell.

Important When creating a Cloud Proxy, if you clone an existing vCloud Director instance, or you start the vCloud Director services before configuring the cell as a dedicated Cloud Proxy, you negatively impact existing vCloud Director instances. By cloning an existing cell, or starting vCloud Director services before configuring the cell as a dedicated Cloud Proxy, the cell registers in the vCloud Director data base as one that can run vCenter proxy listener. Then, when you configure the same cell as a dedicated Cloud Proxy, it is no longer able to run the vCenter proxy listener and you receive a None of the cells have a vCenter proxy service running error in the **System > Manage & Monitor > Cloud Cells** menu of the vCloud Director user interface. For more information, see <https://kb.vmware.com/kb/53172>.

If you are installing vCloud Availability on top of an existing vCloud Director infrastructure, you can configure existing vCloud Director appliances to serve as Cloud Proxy instances, by disabling most of the vCloud Director services. Cloud Proxy hosts must have access to the vCloud Director data base and the transfer share.

You can load balance Cloud Proxy instances with different public Virtual IP addresses (VIPs). You can also use SSL certificates different from the other vCloud Director instances.

Cloud Proxy scales out horizontally, depending on the number of concurrent connections.

Cloud Proxy provides the endpoints used for replicating data for the vCloud Availability solution. Cloud Proxy installation and configuration for vCloud Availability requires configuration of a vCloud Director instance and network interface.

For testing and developing deployments, you can use the primary vCloud Director host as a Cloud Proxy. To expand capacity, deploy additional Cloud Proxy hosts and register them with vCloud Director.

Prerequisites

- Create a virtual machine to run the Cloud Proxy. The Cloud Proxy uses the same OS and configuration as the vCloud Director hosts. For more information about supported operating systems, see the *vCloud Director for Service Providers Release Notes*.
- Verify that all vCloud Director and Cloud Proxy instances have FQDN configured.
- Verify that NTP is configured.
- Verify that the OpenSSL version used in the Guest OS of vCloud Director instance is 1.0.1e-30 or later.
- Verify that the Cloud Proxy hosts use a wildcard certificate and cover all Cloud Proxy host names. If the Cloud Proxy certificate differs from the one used on your vCloud Director instances, you must update the SSL certificates on the Cloud Proxy hosts. For more information about creating and importing SSL certificates, see the *vCloud Director Installation and Upgrade Guide*.

Procedure

1 Pre-installation

- a Copy the `vmware-vcld-director-X.X.X-YYYY.bin` file to the `/tmp` folder of the new Cloud Proxy virtual machine by running the following command.

```
# scp root@vcd-address:/file-path/vmware-vcld-director-X.X.X-YYYY.bin /tmp
# chmod 755 /tmp/vmware-vcld-director-X.X.X-YYYY.bin
```

The `certificates.ks` file is located in the same location as on the primary vCloud Director host. You can find the exact path at `user.keystore.path` in the `responses.properties` file. Update the `user.keystore.path` value to reflect the new path to the certificates file.

- b Copy the configuration file to the `/tmp` folder of the new Cloud Proxy virtual machine by running the following command.

```
# scp root@vcd-address:/opt/vmware/vcloud-director/etc/responses.properties /tmp
# chmod 644 /tmp/responses.properties
```

- c Copy the certificates file to the `/tmp` folder of the new Cloud Proxy virtual machine by running the following command.

```
# scp root@vcd-address:/root/certificates.ks /tmp
# chmod 644 /tmp/certificates.ks
```

- d Update the `database.jdbcUrl` value in the `responses.properties` file to use FQDN for a database host.

- e Mount shared NFS storage.

Verify that you have mounted the shared NFS storage to your Cloud Proxy `/opt/vmware/vcloud-director/data/transfer`.

- f Cloud Proxy Second Network Interface

The vCloud Director installation requires a second NIC to be present, but the Cloud Proxy does not use the second NIC. If you have already provisioned your virtual machine with a second NIC you can set the IP address to a single CIDR address, for example `192.168.254.254/32`. In this case, you do not need to configure the alias NIC.

- g If necessary, set up an alias NIC:

```
# ifconfig eth0:5 192.168.254.254 up
```

2 Install

Run the vCloud Director install script: `vmware-vcld-director-X.X.X-YYYY.bin`

- Do not run the configuration.
- Do not start the `vmware-vcd` service.

3 Configure

Use the responses.properties file to configure the vCloud Director host. Make sure that you do not start the vmware-vcd service.

```
# /opt/vmware/vcloud-director/bin/configure -r /tmp/responses.properties
```

This operation takes a few minutes to finish. The system does not display any output during this time.

4 Specialize a vCloud Director cell to become a dedicated Cloud Proxy cell.

Edit /opt/vmware/vcloud-director/etc/global.properties:

Add the following property:

```
com.vmware.cell.runtime.application=com.vmware.vcloud.cloud-proxy-server.cloudProxyApplication
```

5 Second NIC

The second NIC or alias that you used for the install is no longer required. You can safely turn off the interface.

```
# ifconfig eth0:5 192.168.254.254 down
```

6 Start the vCloud Director service.

```
service vmware-vcd start
```

7 Modify Cloud Proxy address.

If you are running separate Cloud Proxy instances, you must change the address for the Cloud Proxy server.

- a Create a protected password files on your vCloud Availability Installer Appliance in the `~/ .ssh` directory.

```
# mkdir ~/.ssh
# chmod 0700 ~/.ssh
# echo 'vcd-password' > ~/.ssh/.vcd
# find ~/.ssh -type f -name '.*' -print0 | xargs -0 chmod 0600
```

- b To see the currently configured Cloud Proxy address, run the following command on the vCloud Availability Installer Appliance.

```
# vcav vcd get-cloud-proxy \
--type=to-the-cloud \
--vcd-address=vcd-address \
--vcd-user=vcd-user \
--vcd-password-file=~/.ssh/.vcd
```

The vCloud Availability Installer Appliance returns the following message.

```
wss://cloud-proxy-IP-address:to-the-cloud-port/socket/cloudProxy
```

- c Modify the Cloud Proxy by using the following command.

You can modify `--to-the-cloud-address`, `--to-the-cloud-port`, and `--from-the-cloud-address`. For this example, `--to-the-cloud-address` is modified.

```
# vcav vcd set-cloud-proxy \
--to-the-cloud-address=cloud-proxy-FQDN \
--vcd-address=vcd-address \
--vcd-user=vcd-user \
--vcd-password-file=~/.ssh/.vcd
```

The vCloud Availability Installer Appliance returns an OK message.

Installing and Configuring Cassandra Servers

Cassandra is a free and open-source distributed NoSQL database management system that stores metadata and supports storage of the metadata for replication services.

For test and development environments, you can skip this procedure and follow the instructions in [Deploy Cassandra and RabbitMQ as Containers for Test and Development Environments](#) to have Cassandra run in a Docker container instead.

The following is an example of the installation and configuration of a Cassandra server on a CentOS 7.x host.

Prepare Hosts for Cassandra Installation

Before you install and configure a Cassandra host, you must update the Python interpreter to version 2.7, install and configure Java 1.8.x, and configure NTP servers.

Prerequisites

- Verify that you have a CentOS 7 VM to host a Cassandra instance.
- Verify that routing, NTP, forward, and reverse DNS resolutions are working correctly.
- Make sure that SELinux and your firewall are not impacting Cassandra functionalities.
- For best Cassandra performance, it is a best practice to disable Linux swap.

Procedure

- 1 In the CentOS host console, to verify that Python 2.7.5 is installed, run the `# python --version` command.

If the system returns an earlier Python version than 2.7.5, or a `-bash: python: command not found` error, install Python 2.7.5 by running the `# yum install -y python2` command.

- 2 Disable the Linux firewall by running the following commands:

```
# systemctl disable firewalld
# systemctl stop firewalld
```

- 3 Download Oracle Java Development Kit (JDK) 8u152.

You can download Oracle JDK 8u152 from the Oracle archive downloads at <http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html>. Log in with a free Oracle account, accept the Oracle License Agreement, and download the `jdk-8u152-linux-x64.rpm` file locally and copy it to the `/tmp` directory of your CentOS host.

- 4 To install Oracle JDK, in the CentOS host console, navigate to the `/tmp` directory and run the following command:

```
# rpm -ivh jdk-8u152-linux-x64.rpm
```

5 Install the Oracle Java Cryptography Extension (JCE) to increase the java encryption level support.

- a Download the JCE by running the following command:

```
# wget -c --header "Cookie: oraclelicense=accept-securebackup-cookie"
http://download.oracle.com/otn-pub/java/jce/8/jce_policy-8.zip
```

- b Unzip the file by running the following command:

```
# unzip jce_policy-8.zip
```

If the unzip utility is not installed, run the `# yum install zip unzip -y` command to install it.

- c Copy the unzipped JCE file to the respective directory by running the following command:

```
# cp UnlimitedJCEPolicyJDK8/*.jar /usr/java/jdk1.8.0_152/jre/lib/security/
```

6 Verify that the default Java version is correct by running the following command:

```
# java -version
```

The system returns the following:

```
java version "1.8.0_152"
Java(TM) SE Runtime Environment (build 1.8.0_152-b16)
Java HotSpot(TM) 64-Bit Server VM (build 25.152-b16, mixed mode)
```

7 If you are configuring a Cassandra cluster, after the first Cassandra host is deployed, deploy two additional hosts.

It is a best practice to configure a DRS anti-affinity rule to prevent the Cassandra systems from residing on the same host. For more information, see *Create Intra-VM Anti-Affinity Rules* in the *vSphere Resource Management* documentation.

Install Cassandra

Register the Apache Cassandra repository, install Cassandra services, and verify that the installation is successful.

Prerequisites

Verify that you have Python 2.7.5 or later, that all Cassandra hosts are NTP configured, and that Java 1.8.x is installed and configured. For more information, see [Prepare Hosts for Cassandra Installation](#).

Procedure

- 1 To register the Apache Cassandra repository, create the `/etc/yum.repos.d/cassandra.repo` file with the following content:

```
[cassandra]
name=Apache Cassandra
baseurl=https://www.apache.org/dist/cassandra/redhat/311x/
gpgcheck=1
repo_gpgcheck=1
gpgkey=https://www.apache.org/dist/cassandra/KEYS
```

- 2 Install Cassandra by running the following command.

```
# yum install cassandra
```

- 3 If a Cassandra PID directory does not exist at `/var/run/cassandra`, run the following command to create it.

```
# mkdir -p /var/run/cassandra; chown -R cassandra:cassandra /var/run/cassandra
```

- 4 If you are configuring a Cassandra cluster, configure all nodes to use `GossipingPropertyFileSnitch`.

It is a best practice to use `GossipingPropertyFileSnitch` for production environments. The `GossipingPropertyFileSnitch` uses rack and data center information for the local node defined in the `/etc/cassandra/default.conf/cassandra-rackdc.properties` file and propagates this information to other nodes by using the gossip protocol.

To configure a node to use `GossipingPropertyFileSnitch`, edit the `cassandra-rackdc.properties` and the `/etc/cassandra/conf/cassandra.yaml` files.

- a Define data center and rack that include this node. By default, the following values are set:

```
dc=DC1
rack=RAC1
```

Note Data center and rack names are case-sensitive.

- b To save bandwidth, add the `prefer_local=true` option. This option tells the Apache Cassandra platform to use the local IP address when the communication is not across different data centers.
 - c In the `/etc/cassandra/conf/cassandra.yaml` file, locate the `endpoint_snitch` line and change the value to `GossipingPropertyFileSnitch`.
- 5 Enable the Cassandra service.

```
# systemctl enable cassandra
```

- 6 If you are configuring a Cassandra cluster, set the cluster name.

You set a Cassandra cluster name by editing the `cluster_name` option in the `/etc/cassandra/conf/cassandra.yaml` file. By default, the `cluster_name` is **Test Cluster**.

- 7 Start the Cassandra service.

```
# systemctl start cassandra
```

- 8 Check the Cassandra service status.

```
# systemctl status cassandra
```

Configure a Cassandra Cluster

You configure a Cassandra cluster by editing the `/etc/cassandra/conf/cassandra.yaml` configuration file available on each Cassandra node.

Repeat the following steps on every Cassandra node that you want to join to a cluster.

Prerequisites

Verify that you have enabled client communication with Cassandra over SSL. For more information, see [Enable Server and Client Communication with Cassandra over SSL](#).

Procedure

- 1 Set the `seeds_provider` value to list the IP addresses of the Cassandra cluster nodes. Use commas to separate the IP addresses.

```
seed_provider:
  # Addresses of hosts that are deemed contact points.
  # Cassandra Nodes use this list of hosts to find each other and learn
  # The Rerminology of the ring. You must change this if you are running
  # Multiple nodes!
  -class_name: org.apache.cassandra.locator.simpleSeedProvider
    parameters:
      #seeds is actually a comma-delimited list of addresses
      #Ex: "<ip1>,<ip2>,<ip3>"
      - seeds: "ip-Cassandra-node-1,ip-Cassandra-node-2,ip-Cassandra-node-3"
```

- 2 Save the changes and exit the file.
- 3 Restart the Cassandra node by running the following command.

```
# systemctl restart cassandra
```

- 4 Verify that the node is operational by running the following command.

Repeat this step for each Cassandra node before you proceed.

```
# nodetool status
```

If the Cassandra node is operational, the system returns the following information:

```
Datcenter:Datcenter-Name
=====
Status=Up/Down
|/State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens      Owns (effective)  Host ID      Rack
UN Node-IP-Address  190.68 KB  256        100.0%           9c428c6f-e792-4125-8f4c-1c88dae82ebb  rack-
name
```

- 5 After the `seeds_provider` value is properly configured on all nodes, verify that all nodes in the Cassandra cluster are communicating with each other as expected.

```
# nodetool status
```

A successful cluster configuration returns the following status:

```
Datacenter:Datacenter-Name
=====
Status=Up/Down
|/State=Normal/Leaving/Joining/Moving
-- Address      Load       Tokens     Owns (effective)  Host ID           Rack
UN Node-IP-Address1  213.07 KB   256        66.5%             9c428c6f-e792-4125-8f4c-1c88dae82ebb  rack-
name
UN Node-IP-Address2  203.85 KB   256        67.5%             cae13946-04a5-4f3a-8225-842850dbe607  rack-
name
UN Node-IP-Address3  156.22 KB   256        66.0%             f5ee3d7b-e83f-4fa7-b609-dea989b2b415  rack-
name
```

Generate Certificates for Cassandra Nodes

Cassandra requires an SSL connection between the client and the node to enable vSphere Replication Cloud Service to communicate with Cassandra.

Repeat the steps on every Cassandra node.

Procedure

- 1 Generate an SSL certificate by running the following command.

```
# /opt/jdk_version/bin/keytool -keystore /etc/cassandra/conf/.keystore \
-storepass vmware -validity 365 -storetype JKS -genkey -keyalg RSA \
-alias cass-node-ip-address -dname 'cn=cass-node-ip-address, ou=DR2C, o=VMware, c=US' \
-keypass vmware
```

- 2 Export the Cassandra certificate to a PEM-formatted file.

```
# /opt/jdk_version/bin/keytool -export -rfc \
-keystore /etc/cassandra/conf/.keystore -storepass vmware \
-file /root/cloud-cass-node-ip-address-node_number.pem -alias cass-node-ip-address
```

- 3 Copy the `.pem` certificate file of each Cassandra node to the `/root` directory of the rest of the Cassandra nodes.
- 4 Import each certificate to the truststore of every Cassandra host.

```
# /opt/jdk_version/bin/keytool -noprompt -import -trustcacerts \
-alias cass-node-ip-address -file /root/cloud-cass-node-ip-address-node_number.pem \
-keystore /etc/cassandra/conf/.truststore -storepass vmware
```

Every Cassandra truststore contains a copy of the `.pem` certificate of all the remaining nodes.

Enable Server and Client Communication with Cassandra over SSL

You enable the server and the client communication with Cassandra over SSL by editing the `/etc/cassandra/conf/cassandra.yaml` configuration file available on each Cassandra node.

Repeat the following steps on every Cassandra node that you want to join to a cluster.

Procedure

- 1 Set the `listen_address` and `rpc_address` values to the Cassandra node IP address.

```
listen_address: Cass-Node-IP
rpc_address: Cass-Node-IP
```

- 2 Update the values of the `server_encryption_options` properties.

```
server_encryption_options:
  internode_encryption: all
  keystore: /etc/cassandra/conf/.keystore
  keystore_password: vmware
  truststore: /etc/cassandra/conf/.truststore
  truststore_password: vmware
  # More advanced defaults below:
  # protocol: TLS
  # algorithm: SunX509
  store_type: JKS
  require_client_auth: true
  # require_endpoint_verification: false
```

- 3 Update the values of the `client_encryption_options` properties.

The keystore and truststore passwords are the same passwords that you used to create the keystore and the truststore.

```
Client_encryption_options:
  enabled: true
  # If enabled and optional is set to true encrypted and unencrypted connections are handled.
  optional: true
  keystore: /etc/cassandra/conf/.keystore
  keystore_password: vmware
  require_client_auth: true
  # Set truststore and truststore_password if require_client_auth is true
  truststore: /etc/cassandra/conf/.truststore
  truststore_password: vmware
  # More advanced defaults below:
  # protocol: TLS
  # algorithm: SunX509
  store_type: JKS
```

- 4 Save the changes and exit the file.

- Restart the Cassandra node by running the following command.

```
# systemctl restart cassandra
```

Installing and Configuring RabbitMQ Servers

RabbitMQ is an AMQP server that exchanges messages within a vCloud Director environment. You can install and configure up to three RabbitMQ servers on CentOS hosts. For high availability and scalability purposes, you can configure the RabbitMQ servers in a cluster.

Note If you have RabbitMQ already installed, make sure that the host is configured to support SSL connections.

For test and development environments, you can have RabbitMQ run in a Docker container instead. For more information, see [Deploy Cassandra and RabbitMQ as Containers for Test and Development Environments](#) in the *vCloud Availability for vCloud Director Installation and Configuration Guide*.

Procedure

- [Download and Install RabbitMQ Servers](#)

You can download and install RabbitMQ server on up to three CentOS hosts.

- [Configure a Primary RabbitMQ Server](#)

After you download and install RabbitMQ on all hosts in your environment, you must configure the primary RabbitMQ server.

- [Create Self-Signed Certificates for the Primary RabbitMQ Server](#)

After you configure the primary RabbitMQ server, you must create self-sign SSL certificates for it.

- [\(Optional\) \(Optional\) Configure a RabbitMQ Cluster](#)

For high availability and scalability purposes, you can set up a RabbitMQ cluster.

Download and Install RabbitMQ Servers

You can download and install RabbitMQ server on up to three CentOS hosts.

You must complete the install procedure on all hosts within your environment:

Procedure

- Login as **root** on the CentOS host.
- Download and install a zero dependency Erlang RPM package that is required to run RabbitMQ.

```
# wget http://www.rabbitmq.com/releases/erlang/erlang-18.3-1.e16.x86_64.rpm
# rpm -i erlang-18.3-1.e16.x86_64.rpm
```

3 Download and import the public signing key.

```
# wget https://www.rabbitmq.com/rabbitmq-signing-key-public.asc
# rpm --import rabbitmq-signing-key-public.asc
```

4 Download and install the RabbitMQ application.

```
# wget http://www.rabbitmq.com/releases/rabbitmq-server/v3.6.1/rabbitmq-server-3.6.1-1.noarch.rpm
# rpm -i rabbitmq-server-3.6.1-1.noarch.rpm
```

5 Configure an NTP server.

Configure an NTP server for every RabbitMQ host in your environment.

- a Run the following command to install the ntp service.

```
# yum install ntp
```

- b Add the following line for each time server to the end of the `/etc/ntp.conf` file.

```
server your.time.server iburst
```

- c Restart the NTP service.

```
# service ntp restart
```

Configure a Primary RabbitMQ Server

After you download and install RabbitMQ on all hosts in your environment, you must configure the primary RabbitMQ server.

Procedure

1 Create directories.

```
cd /etc/rabbitmq;
mkdir testca server client
```

2 Create the `/etc/rabbitmq/rabbitmq.config` file with the following content.

```
{[ssl, [{versions, [tlsv1, 'tlsv1.2', 'tlsv1.1']},
        {ciphers, ["ECDHE-ECDSA-AES256-GCM-SHA384","ECDHE-RSA-AES256-GCM-SHA384",
"ECDHE-ECDSA-AES256-SHA384","ECDHE-RSA-AES256-SHA384",
"ECDH-ECDSA-AES256-GCM-SHA384","ECDH-RSA-AES256-GCM-SHA384",
"ECDH-ECDSA-AES256-SHA384","ECDH-RSA-AES256-SHA384",
"DHE-RSA-AES256-GCM-SHA384","DHE-DSS-AES256-GCM-SHA384",
"DHE-RSA-AES256-SHA256","DHE-DSS-AES256-SHA256","AES256-GCM-SHA384",
"AES256-SHA256","ECDHE-ECDSA-AES128-GCM-SHA256",
"ECDHE-RSA-AES128-GCM-SHA256","ECDHE-ECDSA-AES128-SHA256",
"ECDHE-RSA-AES128-SHA256","ECDH-ECDSA-AES128-GCM-SHA256",
"ECDH-RSA-AES128-GCM-SHA256","ECDH-ECDSA-AES128-SHA256",
"ECDH-RSA-AES128-SHA256","DHE-RSA-AES128-GCM-SHA256",
"DHE-DSS-AES128-GCM-SHA256","DHE-RSA-AES128-SHA256","DHE-DSS-AES128-SHA256",
"AES128-GCM-SHA256","AES128-SHA256","ECDHE-ECDSA-AES256-SHA",
```

```

"ECDHE-RSA-AES256-SHA", "DHE-RSA-AES256-SHA", "DHE-DSS-AES256-SHA",
"ECDH-ECDSA-AES256-SHA", "ECDH-RSA-AES256-SHA", "AES256-SHA",
"ECDHE-ECDSA-DES-CBC3-SHA", "ECDHE-RSA-DES-CBC3-SHA", "EDH-RSA-DES-CBC3-SHA",
"EDH-DSS-DES-CBC3-SHA", "ECDH-ECDSA-DES-CBC3-SHA", "ECDH-RSA-DES-CBC3-SHA",
"DES-CBC3-SHA", "ECDHE-ECDSA-AES128-SHA", "ECDHE-RSA-AES128-SHA",
"DHE-RSA-AES128-SHA", "DHE-DSS-AES128-SHA", "ECDH-ECDSA-AES128-SHA",
"ECDH-RSA-AES128-SHA", "AES128-SHA", "EDH-RSA-DES-CBC-SHA", "DES-CBC-SHA"]]],
{rabbit, [
  {ssl_listeners, [5671]},
  {ssl_options, [{cacertfile, "/etc/rabbitmq/testca/cacert.pem"},
    {certfile, "/etc/rabbitmq/server/cert.pem"},
    {keyfile, "/etc/rabbitmq/server/key.pem"},
    {versions, [tlsv1, 'tlsv1.2', 'tlsv1.1']},
    {verify, verify_peer},
    {padding_check, true},
    {ciphers, ["ECDHE-ECDSA-AES256-GCM-SHA384", "ECDHE-RSA-AES256-GCM-SHA384",
"ECDHE-ECDSA-AES256-SHA384", "ECDHE-RSA-AES256-SHA384",
"ECDH-ECDSA-AES256-GCM-SHA384", "ECDH-RSA-AES256-GCM-SHA384",
"ECDH-ECDSA-AES256-SHA384", "ECDH-RSA-AES256-SHA384",
"DHE-RSA-AES256-GCM-SHA384", "DHE-DSS-AES256-GCM-SHA384",
"DHE-RSA-AES256-SHA256", "DHE-DSS-AES256-SHA256", "AES256-GCM-SHA384",
"AES256-SHA256", "ECDHE-ECDSA-AES128-GCM-SHA256",
"ECDHE-RSA-AES128-GCM-SHA256", "ECDHE-ECDSA-AES128-SHA256",
"ECDHE-RSA-AES128-SHA256", "ECDH-ECDSA-AES128-GCM-SHA256",
"ECDH-RSA-AES128-GCM-SHA256", "ECDH-ECDSA-AES128-SHA256",
"ECDH-RSA-AES128-SHA256", "DHE-RSA-AES128-GCM-SHA256",
"DHE-DSS-AES128-GCM-SHA256", "DHE-RSA-AES128-SHA256", "DHE-DSS-AES128-SHA256",
"AES128-GCM-SHA256", "AES128-SHA256", "ECDHE-ECDSA-AES256-SHA",
"ECDHE-RSA-AES256-SHA", "DHE-RSA-AES256-SHA", "DHE-DSS-AES256-SHA",
"ECDH-ECDSA-AES256-SHA", "ECDH-RSA-AES256-SHA", "AES256-SHA",
"ECDHE-ECDSA-DES-CBC3-SHA", "ECDHE-RSA-DES-CBC3-SHA", "EDH-RSA-DES-CBC3-SHA",
"EDH-DSS-DES-CBC3-SHA", "ECDH-ECDSA-DES-CBC3-SHA", "ECDH-RSA-DES-CBC3-SHA",
"DES-CBC3-SHA", "ECDHE-ECDSA-AES128-SHA", "ECDHE-RSA-AES128-SHA",
"DHE-RSA-AES128-SHA", "DHE-DSS-AES128-SHA", "ECDH-ECDSA-AES128-SHA",
"ECDH-RSA-AES128-SHA", "AES128-SHA", "EDH-RSA-DES-CBC-SHA", "DES-CBC-SHA"]]},
    {fail_if_no_peer_cert, false}]]
  ]}
].

```

To configure RabbitMQ to listen on both SSL and non-SSL ports, add the `{tcp_listeners, [5672]}` line in the `/etc/rabbitmq/rabbitmq.config` file:

```

...
{rabbit, [
  {tcp_listeners, [5672]},
  {ssl_listeners, [5671]},
  ...]}.

```

3 Start the RabbitMQ service.

```
# service rabbitmq-server start
```

4 Enable RabbitMQ UI on **http://server-name:15672/**.

```
# rabbitmq-plugins enable rabbitmq_management
```

5 Create **admin** user with password **vmware** and set the user permissions.

```
# rabbitmqctl add_user admin vmware
# rabbitmqctl set_permissions -p / admin ".*" ".*" ".*"
# rabbitmqctl set_user_tags admin administrator
```

Create Self-Signed Certificates for the Primary RabbitMQ Server

After you configure the primary RabbitMQ server, you must create self-sign SSL certificates for it.

In the following example, there are two CentOS hosts and one Load Balancer server.

Prerequisites

To use `keytool` command, you must have Java installed. You have two options:

- Install Java on your system.
- Run the `keytool` command on another system that has it installed and copy the certificates to the RabbitMQ servers after their creation.

Procedure

- 1 Create a public and a private key.

Important The SAN attribute contains DNS names and IP addresses of all of the RabbitMQ hosts and the load balancer. You must update the values in the command to match your environment. The CN attribute must contain a wildcard for the domain. Because this is a self-signed certificate, the root certificate and the server certificate are the same.

```
keytool -genkeypair \
-keystore rootca.jks \
-storepass vmware \
-keyalg RSA \
-validity 365 \
-keypass vmware \
-alias rabbitmq \
-dname "CN=*.corp-ext.local,OU=Test, O=Corp, L=Palo Alto S=CA C=US" \
-ext san=\
dns:test2.corp-ext.local,dns:test3.corp-ext.local,dns:testrabbitmq1b.corp-ext.local,\
ip:172.31.3.39,ip:172.31.3.40,ip:172.31.3.41
```

Note You can change the validity period of the certificate by adjusting the *validity* value in the command. In the example, the created certificate is valid for 365 days.

2 Import the RabbitMQ key pair to the PKCS12 trust store.

```
keytool -importkeystore -srckeystore rootca.jks \
-destkeystore foo.p12 -deststoretype pkcs12 \
-srcstorepass vmware -deststorepass vmware \
-alias rabbitmq
```

3 Convert the key pair file to PEM format.

```
openssl pkcs12 -in foo.p12 \
-out foo.pem -passin pass:vmware \
-passout pass:vmware
```

4 Extract the encrypted private key.

```
sed -n '/-----BEGIN ENCRYPTED PRIVATE KEY-----/,/-----END ENCRYPTED PRIVATE KEY-----/p' \
foo.pem > enc.pem
```

5 Decrypt the private key.

```
openssl rsa -in enc.pem \
-out unenc.pem -passin pass:vmware
```

6 Extract the certificate.

```
sed -n '/-----BEGIN CERTIFICATE-----/,/-----END CERTIFICATE-----/p' \
foo.pem > cert.pem
```

7 Install the self-signed certificates by copying them to the newly created directories.

```
cp cert.pem /etc/rabbitmq/testca/cacert.pem
cp cert.pem /etc/rabbitmq/server/cert.pem
cp unenc.pem /etc/rabbitmq/server/key.pem
cp cert.pem /etc/rabbitmq/client/cert.pem
cp unenc.pem /etc/rabbitmq/client/key.pem
```

8 Change the ownership of the newly created directories.

```
chown -R rabbitmq: /etc/rabbitmq/testca
chown -R rabbitmq: /etc/rabbitmq/server
chown -R rabbitmq: /etc/rabbitmq/client
```

9 Start and enable the rabbitmq-server service.

```
service rabbitmq-server start
chkconfig rabbitmq-server on
```


(Optional) (Optional) Configure a RabbitMQ Cluster

For high availability and scalability purposes, you can set up a RabbitMQ cluster.

Run the following procedure on all hosts within your environment:

Procedure

- 1 Log in as **root** on the current CentOS host.
- 2 Copy the entire `/etc/rabbitmq` directory from the primary RabbitMQ node to the current node.
- 3 Enable starting the RabbitMQ service on boot.
`chkconfig rabbitmq-server on`
- 4 Start the RabbitMQ service.
`service rabbitmq-server start`
- 5 Stop the RabbitMQ application.
`rabbitmqctl stop_app`
- 6 Reset the RabbitMQ server.
`rabbitmqctl reset`
- 7 Stop the RabbitMQ service.
`service rabbitmq-server stop`
- 8 Navigate to the `/var/lib/rabbitmq` directory.
`cd /var/lib/rabbitmq/`
- 9 To back up the `.erlang.cookie` file, run the following command.
`mv .erlang.cookie .erlang.cookie.OLD`
- 10 Run the following command.
`scp first_node:/var/lib/rabbitmq/ .erlang.cookie`
- 11 Change the ownership of the `.erlang.cookie` file.
`chown rabbitmq:rabbitmq .erlang.cookie`
- 12 Start the RabbitMQ service.
`service rabbitmq-server start`
- 13 Stop the RabbitMQ application.
`rabbitmqctl stop_app`
- 14 Join the current RabbitMQ node to the primary RabbitMQ node `rabbit@first_node`, as a RAM node.
`rabbitmqctl join_cluster --ram rabbit@first_node`

15 Start the RabbitMQ application.

```
rabbitmqctl start_app
```

16 Configure a high availability policy for queue mirroring.

```
rabbitmqctl set_policy ha-all "" '{"ha-mode":"all","ha-sync-mode":"automatic"}'
```

17 Verify the status of the cluster configuration.

```
rabbitmqctl cluster_status
```

What to do next

You must configure the load balancer to route the traffic between the RabbitMQ nodes on SSL port 5671 and non-SSL port 5672.

You must configure TCP (not HTTP) to route on both ports on the load balancer.

Automated Installation and Configuration

The automated installation and configuration of the vCloud Availability solution components is repeatable, time saving, faster and less error prone compared to the manual deployment. You can easily adopt it to brownfield and mixed environments.

Creating a Registry File for an Automated Installation

Before you start the automated installation of vCloud Availability, you must create a registry file containing information about the infrastructure and solution components you are about to deploy. Some of the components might exist or you can deploy them through the vCloud Availability Installer Appliance.

Use a text editor to create a file named `registry` and upload it to the `/root/.vcav/` directory on the vCloud Availability Installer Appliance.

Important You can find a template to create your registry file at `/root/.vcav/.registry.tpl`

General Properties for the Registry File

Some of the properties in the registry file are applicable to all virtual machines you deploy. Change the properties values as appropriate.

```
# DISABLE_CERTIFICATE_VALIDATION = 1    # Disable all SSL and SSH certificate validation
# NTP = time.example.com                # Default NTP server for the deployed VMs
# SSH_PASSWORD =                       # Default SSH password for the deployed VMs
```

Structure of the Registry Entries

Each entry in the registry file requires a specific format. The order of the components is irrelevant.

```
{entry-type} {entry-name}
  property-name1 property-value1
  property-name2 property-value2
  .....
  property-nameN property-valueN
```

The `{entry_type}` field supports entries for the *vsphere*, *vcd*, *docker*, *amqp*, *cassandra*, *hms*, *hcs*, *hbr*, *ui*, and *smp* components.

Note By default, you do not assign a hostname to VMs you deploy. To assign a hostname to the VMs you deploy with the registry file, add the `hostname` property and value in the respective sections of the registry file.

Use of the Registry File

You can specify the registry entries as arguments to some individual commands of the vCloud Availability Installer Appliance and use them for cleanup and unconfigure operations.

Registry Entry for a vSphere Host

You must create a registry entry for each management or resource vSphere host in your environment and update the values according to your environment.

The following list contains all vSphere properties that the registry file supports and the required structure:

```
vsphere vsphere-name
  address vsphere-address
  api-port 443
  api-user admin-user
  api-password admin-user-password
  api-thumbprint api-thumbprint-value
  placement-locator mgmt-locator
  placement-datastore mgmt-datastore
  placement-network mgmt-network
  placement-folder mgmt-folder
  network-profile-name network-name
```

```

network-gateway network-gateway-IP-address
network-netmask network-netmask
network-dns DNS-server1-IP, DNS-server1-IP
network-subnet network-subnet-IP-address

```

Note In the current example, you create three registry entries: one for the management vSphere host and two for the resource vSphere hosts. The following example contains only the mandatory vSphere properties. If necessary, you can add more.

```

vsphere mgmt-vsphere-name
  address mgmt-vsphere-address
  api-port 443
  api-user admin-user
  api-password admin-user-password
  placement-locator mgmt-locator
  placement-datastore mgmt-datastore
  placement-network mgmt-network

vsphere vsphere-01-name
  address vsphere-01-address
  api-port 443
  api-user vsphere-01-admin-user
  api-password vsphere-01-admin-password
  placement-locator vsphere-01-locator
  placement-datastore vsphere-01-datastore
  placement-network vsphere-01-network

vsphere vsphere-02-name
  address vsphere-02-address
  api-port 443
  api-user vsphere-02-admin-user
  api-password admin-user-password
  placement-locator vsphere-02-locator
  placement-datastore vsphere-02-datastore
  placement-network vsphere-02-network

```

To enable a static IP addresses deployment, make sure that you have a vSphere network pool created. You can create it in the following ways:

- Set the network profile properties in the corresponding registry entry and the vCloud Availability Installer Appliance creates and configures the IP pools.
Set the *network-profile-name*, *network-gateway*, *network-netmask*, *network-dns*, and optionally *network-subnet* property values.
- Create an IP pool in the vSphere Web Client. For more information about creating an IP pool, see the *VMware NSX for vSphere* documentation.
- Run the `vcav ip-pool` command. For more information, see [Enable Static IP Addresses Deployment](#).

Registry Entry for a vCloud Director Host

You must create a registry entry for the vCloud Director host and update the values according to your environment.

```
vcd vcd-01-name
  address vcd-01-address
  api-port 443
  api-user root@system
  api-password vcd-root-password
  sso-user administrator@vsphere.local
  sso_password sso-password
  amqp amqp-registry-name
```

You must provide the corresponding registry name of the RabbitMQ endpoint to which the vCloud Director host connects.

(Optional) Registry Entry for a Docker Host

If the Cassandra database or the RabbitMQ event broker runs as Docker container, you must create a registry entry for the Docker instance and update the values according to your environment.

```
docker docker-name
  placement-vsphere vsphere-address-registry-entry
  placement-address static-IP-address
  ovf-url URL
  ssh-password ssh-password
  ntp ntp-address
```

Note If you already set either the `ssh-password` or the `ntp-address` as global properties at the beginning of the registry file, you do not need to specify them in this registry entry.

The `placement-address` property name is optional. Provide it if you use static network addresses.

The `ovf-url` property name is optional. Provide it if you want to specify a custom location for the OVA deployment.

Prepare the vCloud Availability Installer Appliance for Docker Hosts

Before you start the automated vCloud Availability deployment, you must start the Docker service on your vCloud Availability Installer Appliance and download Cassandra and RabbitMQ images.

To start the Docker service on the vCloud Availability Installer Appliance, run the `systemctl start docker` command.

Download the Cassandra image by running the `docker pull cassandra:3.9` command.

Download the RabbitMQ image by running the `docker pull rabbitmq:3.4` command.

Registry Entry for a RabbitMQ Host

You can configure an already deployed RabbitMQ host as part of the automated installation and configuration process.

For test or development environments, you can create a RabbitMQ container as part of the automated installation process.

Important If you use a RabbitMQ cluster, you should install and configure the cluster before installing and configuring the vCloud Availability solution. For more information about installing RabbitMQ servers and configuring a RabbitMQ cluster, see [Installing and Configuring RabbitMQ Servers](#). If you are using a RabbitMQ cluster as part of the vCloud Availability solution, enter the load balancer virtual IP address in the registry file entry.

Update the property values according to your environment.

Registry Entry for an Existing RabbitMQ Server

```
amqp amqp.0
  address amqp-address
  port 5671
  user admin
  password vmware
  vhost /
  exchange systemExchange
```

Registry Entry for the Deployment of a RabbitMQ Container

```
amqp amqp.1
  placement-docker docker-name
  port 5671
  user admin
  password vmware
  vhost /
  exchange systemExchange
```

Registry Entry for a Cassandra Host

You can have an already deployed Cassandra server in your environment or you can deploy a Cassandra container as part of the automated installation process.

Important Update the property values according to your environment.

You provide a comma-separated list of registry names of the vSphere Replication Cloud Service appliances and add a record for the Cassandra server as a property entry in the *hcs-list*.

Registry Entry for an Existing Cassandra Server

```
cassandra cass.0
  address cassandra-address
  port 9042
  ssh-password ssh-password
  hcs-list hcs.0,hcs.1,...,hcs.N
```

Registry Entry for Deployment of a Cassandra Container

```
cassandra cass.1
  placement-docker docker-name
  port 9042
  hcs-list hcs.0,hcs.1,...,hcs.N
```

Registry Entry for a vSphere Replication Manager Appliance

You can have an already deployed vSphere Replication Manager appliance or you can deploy it as part of the automated installation process.

Important Update the property values according to your environment.

You provide the registry name of vCloud Director for which the vSphere Replication Manager appliance is configured as the *vcd* property name.

You provide the registry name of vSphere for which the vSphere Replication Manager appliance is configured as the *vsphere* property name.

If you already set either the *ssh-password* or the *ntp-address* as global properties at the beginning of the registry file, you do not need to specify them in the registry entry.

Registry Entry for an Existing vSphere Replication Manager Appliance

```
hms hms.0
  address hms-address
  ssh-password ssh-password
  ssh-thumbprint ssh-thumbprint
  ntp ntp-address
  vcd vcd-01-name
  vsphere vsphere-registry-name
```

Registry Entry for the Deployment of a vSphere Replication Manager Appliance

```
hms hms.1
  placement-vsphere vsphere-registry-name
  placement-address static-IP-address
  ovf-url URL
```

```
ssh-password ssh-password
ntp ntp-address
vcd vcd-registry-name
vsphere vsphere-registry-name
```

The *placement-address* property name is optional. Provide it if you use static network addresses.

The *ovf-url* property name is optional. Provide it if you want to specify a custom location for the OVA deployment.

Registry Entry for a vSphere Replication Cloud Service Appliance

You can have an already deployed vSphere Replication Cloud Service appliance or you can deploy it as part of the automated installation process.

Important Update the property values according to your environment.

You provide the registry name of vCloud Director for which the vSphere Replication Cloud Service appliance is configured as the *vcd* property name.

You specify the number of Cassandra servers to be used as the *cassandra-replication-factor* property name.

You provide the corresponding registry name of the RabbitMQ endpoint to which the vCloud Director host connects as the *amqp* property name.

If you already set either the *ssh-password* or the *ntp-address* as global properties at the beginning of the registry file, you do not need to specify them in the registry entry.

Registry Entry for an Existing vSphere Replication Cloud Service Appliance

```
hcs hcs.0
  address hcs-address
  ssh-password ssh-password
  ssh-thumbprint ssh-thumbprint
  ntp ntp-address
  vcd vcd-registry-name
  cassandra-replication-factor number-of-servers
  amqp amqp-registry-name
```

Registry Entry for the Deployment of a vSphere Replication Cloud Service Appliance

```
hcs hcs.1
  placement-vsphere vsphere-registry-name
  placement-address static-IP-address
  ovf-url URL
  ssh-password ssh-password
  ntp ntp-address
  vcd vcd-registry-name
  cassandra-replication-factor number-of-servers
  amqp amqp-registry-name
```


The *placement-address* property name is optional. Provide it if you use static network addresses.

The *ovf-url* property name is optional. Provide it if you want to specify a custom location for the OVA deployment.

Registry Entry for a vSphere Replication Server Appliance

You can have an already deployed vSphere Replication Server appliance or you can deploy it as part of the automated installation process.

Important Update the property values according to your environment.

You provide the registry name of vCloud Director for which the vSphere Replication Server appliance is configured as the *vcd* property name.

You provide the registry name of vSphere for which the vSphere Replication Server appliance is configured as the *vsphere* property name.

If you already set either the *ssh-password* or the *ntp-address* as global properties at the beginning of the registry file, you do not need to specify them in the registry entry.

Registry Entry for an Existing vSphere Replication Server Appliance

```
hbr hbr.0
  address static-IP-address
  ssh-password ssh-password
  ssh-thumbprint ssh-thumbprint
  ntp ntp-address
  vcd vcd-registry-name
  vsphere vsphere-registry-name
```

Registry Entry for the Deployment of a vSphere Replication Server Appliance

```
hbr hbr.1
  placement-vsphere vsphere-registry-name
  placement-address static-IP-address
  ovf-url          URL
  ssh-password ssh-password
  ntp ntp-address
  vcd vcd-registry-name
  vsphere vsphere-registry-name
```

The *placement-address* property name is optional. Provide it if you use static network addresses.

The *ovf-url* property name is optional. Provide it if you want to specify a custom location for the OVA deployment.

Registry Entry for a vCloud Availability Portal Host

You can have an already deployed vCloud Availability Portal host or you can deploy it as part of the automated installation process.

Important Update the property values according to your environment.

You provide the registry name of vCloud Director for which the vCloud Availability Portal host is configured as the `vcd` property name.

The `https-certificate` property value is optional. If you provide it, you must also set the set `https-certificate` property value. If you do not provide it, a self-signed certificate is created.

If you already set either the `ssh-password` or the `ntp-address` as global properties at the beginning of the registry file, you do not need to specify them in the registry entry.

Registry Entry for an Existing vCloud Availability Portal Host

```
ui ui.0
  address portal-address
  ssh-password ssh-password
  ssh-thumbprint ssh-thumbprint
  ntp ntp-address
  vcd vcd-registry-name
  truststore-password truststore-password
  https-certificate path-to-certificate-file
  https-key path-to-key-file
```

Registry Entry for the Deployment of a vCloud Availability Portal Host

```
ui ui.1
  placement-vmware vmware-registry-name
  placement-address static-IP-address
  ovf-url URL
  ssh-password ssh-password
  ntp ntp-address
  vcd vcd-registry-name
  truststore-password truststore-password
  https-certificate path-to-certificate-file
  https-key path-to-key-file
  deployment-type small
```

The `deployment-type` property value is optional. You can set it to one of `small`, `medium`, and `large` value. Default value is `small`. For more information, see [Create vCloud Availability Portal Host](#).

The `placement-address` property name is optional. Provide it if you use static network addresses.

The `ovf-url` property name is optional. Provide it if you want to specify a custom location for the OVA deployment.

Registry Entry for a vCloud Availability Administration Portal Host

You can have an already deployed vCloud Availability Administration Portal host or you can deploy it as part of the automated installation process..

Important Update the property values according to your environment.

You provide the registry name of vCloud Director for which the vCloud Availability Administration Portal host is configured as the *vcd* property name.

The *https-certificate* property value is optional. If you provide it, you must also set the set *https-certificate* property value. If you do not provide it, a self-signed certificate is created.

The *mongodb-database* property value is optional. Default value is *vcav-smp*.

The *mongodb-user* property value is optional. Default value is *vcav-smp*.

The *max-jvm-size* property value is optional. Default value is 1024.

If you already set either the *ssh-password* or the *ntp-address* as global properties at the beginning of the registry file, you do not need to specify them in the registry entry.

Registry Entry for an Existing vCloud Availability Administration Portal Host

```
smp smp.0
  address smp-portal-address
  ssh-password ssh-password
  ssh-thumbprint ssh-thumbprint
  ntp ntp-address
  vcd vcd-registry-name
  truststore-password truststore-password
  https-certificate path-to-certificate-file
  https-key path-to-key-file
  mongodb-database mongodb-database-name
  mongodb-user mongodb-user
  mongodb-password mongodb-password
  max-jvm-size 1024
```

Registry Entry for the Deployment of a vCloud Availability Administration Portal Host

```
smp smp.1
  placement-vsphere vsphere-registry-name
  placement-address static-IP-address
  ovf-url URL
  ssh-password ssh-password
  ntp ntp-address
  vcd vcd-registry-name
  truststore-password truststore-password
  https-certificate path-to-certificate-file
  https-key path-to-key-file
  mongodb-database mongodb-database-name
  mongodb-user mongodb-user
```

```

mongodb-password mongodb-password
max-jvm-size 1024
amqp amqp-registry-name
tenant-ui tenant-UI-registry-name
tenant-ui-url tenant-UI-URL

```

The *placement-address* property name is optional. Provide it if you use static network addresses.

The *ovf-url* property name is optional. Provide it if you want to specify a custom location for the OVA deployment.

You provide the corresponding registry name of the RabbitMQ endpoint to which the vCloud Director host connects as the *amqp* property name.

Starting a vCloud Availability Installation

After you create a registry file containing the deployment details for all vCloud Availability components, you can run a single command to initiate the installation.

The vCloud Availability installation consists of two consecutive steps, pre-validation and actual deployment. The `prevalidate` script performs basic environment checks and reports potential installation and configuration mismatches and issues. After you successfully complete the pre-validation step, a list of tasks is being created. A task can be a VM creation, a container creation, a VM, or a container configuration. During the installation process, the tasks from the list run one-by-one.

You can verify the information collected from the registry file for the vCloud Availability endpoints creation by running the `vcav registry list` command. The vCloud Availability Installer Appliance collects all the information for the endpoints and pre-validates all tasks needed for deployment. The vCloud Availability Installer Appliance prompts you for missing information.

To begin the installation process, you can do one of the following:

- Run the `vcav prevalidate` command and then run the `vcav resume` command.
- Run the `vcav start` command and interactively participate in the installation process.

Both options bring the same result. With the `vcav start` command, the necessary pre-installation validations and checks are implicitly performed.

Note If you run the `vcav prevalidate` command and try to run `vcav start` command, you get an error. You must run `vcav resume` command instead.

If you interrupt the installation process, or you want to modify an existing component, you can run one of the installation commands with an extra argument. You can add the following arguments to the `vcav prevalidate`, the `vcav resume`, and the `vcav start` commands depending on your case.

Argument	Description	Default Value
<code>--overwrite</code>	Overwrite existing VMs or containers	False
<code>--no-overwrite</code>	Do not overwrite existing VMs or containers	True

Argument	Description	Default Value
<code>--reconfigure</code>	Reconfigure VMs or containers	False
<code>--no-reconfigure</code>	Do not reconfigure VMs or containers	True
<code>--unregister-hms-extension</code>	Unregister existing HMS extensions	False
<code>--no-unregister-hms-extension</code>	Do not unregister existing HMS extensions	True

During the execution of the tasks, if a single vCloud Availability component exists, the vCloud Availability Installer Appliance first prompts you to overwrite it and then to reconfigure it.

Managing the vCloud Availability Installation

You can install vCloud Availability by running a single command and you are also given an option to participate in the installation process if needed. You can view the status of the installation, update the registry file, restart or resume the installation process, or run or skip a given task at any time.

Viewing the Installation Tasks and the Deployment State

After a successful pre-validation, the vCloud Availability Installer Appliance creates a list of installation and configuration tasks. The vCloud Availability Installer Appliance runs these tasks one at a time. Each task creates or configures vCloud Availability VMs and containers.

To view the list of all numbered tasks required for a complete installation and configuration, you run the `vcav status` command at any time of the installation and configuration process. You can see the status of each task which can be *Started* or *Not Started*. For a started task, you can see when it was started. To view the endpoints related to each task, run the `vcav status` command with an extra `--verbose` argument.

Running a Single Task

You can run a single task by running the `vcav next` command. The vCloud Availability Installer Appliance detects the first task that is not completed and runs it.

You can indicate which task you want to run by adding the `--task=Task-Number` argument.

To rerun a task, you add the `--rerun` argument to the command.

You can add the following arguments to the `vcav next` command.

Argument	Description	Default Value
<code>--overwrite</code>	Overwrite existing VMs or containers	False
<code>--no-overwrite</code>	Do not overwrite existing VMs or containers	True
<code>--reconfigure</code>	Reconfigure VMs or containers	False
<code>--no-reconfigure</code>	Do not reconfigure VMs or containers	True

Argument	Description	Default Value
<code>--unregister-hms-extension</code>	Unregister existing HMS extensions	False
<code>--no-unregister-hms-extension</code>	Do not unregister existing HMS extensions	True

Skipping a Single Task

To skip a task, run the `vcav skip --task =Task-Number` command.

Resuming the Installation Process

If the installation and configuration process is interrupted, you can resume it to the last completed task by running the `vcav resume` command. If a command fails, you can run the `vcav resume` command and rerun the failed command.

Restarting the Installation Process

To restart the installation and configuration process, run the `vcav cancel` command.

By running this command, you force the list of tasks to be cleaned. Then you must start a new deployment.

Updating the Registry File During the Installation Process

If you need to add or remove entries from the registry file after the installation start, without removing any existing VMs and configuration, you must update the task list. In this case, you run the `vcav cancel` command and do the necessary changes in the registry file. To resume the installation and configuration process, run the `vcav prevalidate` command and add the `--no-overwrite` `--no-reconfigure` arguments.

Handling Installation Errors

If the deployment fails during the execution of a single task, resolve the issue causing the error and resume the installation process by running the `vcav resume` command.

In there is an error, you can see the log file `.vcav/vcav.log`. You can check the log of the affected appliance. For more information, see [Service Provider Diagnostics](#).

Getting the IP Addresses of the Components

To get the IP addresses of all vCloud Availability components, run the `vcav registry list-endpoints` command.

Manual Installation and Configuration

You can install and configure the vCloud Availability solution by running a set of commands in a step-by-step manner.

Prerequisites for Manual Installation and Configuration

You must perform a list of tasks before you begin the installation of the vCloud Availability solution.

Procedure

1 [Create Password Files on Your vCloud Availability Installer Appliance](#)

You must create a protected password files on your vCloud Availability Installer Appliance for each of the vCloud Availability components.

2 [Defining Installation Variables](#)

You can simplify the deployment of individual components by defining installation variables or by creating a registry file on your vCloud Availability Installer Appliance.

3 [Add Trusted Thumbprints to the vCloud Availability Installer Appliance](#)

The vCloud Availability Installer Appliance must be able to verify the thumbprint of the vCenter Server and vCloud Director hosts that it works with.

4 [Enable Static IP Addresses Deployment](#)

By default, the vCloud Availability Installer Appliance creates VMs with DHCP. You can apply static IP addresses by adding the `--vm-address` option to any command that deploys an OVF.

5 [Configure vCloud Director](#)

You must perform an additional vCloud Director configuration in case of manual deployment.

6 [Check vCloud Director Endpoints](#)

Verify that your environment is properly configured for vCloud Availability installation, by checking the vCloud Director endpoints for known problems.

Create Password Files on Your vCloud Availability Installer Appliance

You must create a protected password files on your vCloud Availability Installer Appliance for each of the vCloud Availability components.

Procedure

- 1 Connect to the vCloud Availability Installer Appliance over SSH.

2 Create protected password files on your vCloud Availability Installer Appliance.

OS credentials are stored in text files in `~/ .ssh` directory for all appliances. The files are only accessible to the system **root** user for security purposes. You provide the path to the respective password file during installation and configuration steps.

Note The *appliances-root-password* is the **root** password that is set for the vCloud Availability appliances that you create during installation procedures. The following example uses the same **root** password for all vCloud Availability appliances. You can set different passwords for all appliances, by creating a dedicated password file in the `~/ .ssh` directory. Provide the path to the correct password file in the respective installation and configuration step.

```
# mkdir ~/.ssh
# chmod 0700 ~/.ssh
# echo 'appliances-root-password' > ~/.ssh/.root
# echo 'vcd-password' > ~/.ssh/.vcd
# echo 'sso-password' > ~/.ssh/.sso
# echo 'management-vmware-password' > ~/.ssh/.vsphere.mgmt
# for RabbitMQ and...others
# find ~/.ssh -type f -name '.*' -print0 | xargs -0 chmod 0600
```

Defining Installation Variables

You can simplify the deployment of individual components by defining installation variables or by creating a registry file on your vCloud Availability Installer Appliance.

There are two ways to deploy and configure vCloud Availability

- You can use **Full Commands Installation**. The commands include addresses, user names, and the location of password files for all vCenter Server instances and vCloud Director hosts.
- With **Simple Command Installation** you are using a vCloud Availability Installer Appliance registry. This way all vCenter Server and vCloud Director details are contained in a registry file.

Both ways to deploy and configure vCloud Availability are displayed for your reference. The installation using variables is presented in the left column of the table in each step, containing standard installation and configuration commands. The installation with simple commands, using a vCloud Availability Installer Appliance registry file, is presented in the right column of the table in each step.

The *VSPHERE_PLACEMENT_LOCATOR* value contains the target data center name, the tag *host*, the name of the target cluster, and the IP address or the fully qualified domain name (FQDN) of the target ESXi host. The *VSPHERE_PLACEMENT_LOCATOR* value depends on the topology of your vSphere environment. Following are examples for valid *VSPHERE_PLACEMENT_LOCATOR* values.

- */data-center-name/host/cluster-1-name/fully-qualified-domain-name*
- */data-center-name/host/cluster-2-name/host-IP-address*

If the target ESXi host is not part of a cluster, skip the *cluster-name* element, as shown in the following examples.

- */data-center-name/host/fully-qualified-domain-name*
- */data-center-name/host/host-IP-address*

The *VSPHERE_PLACEMENT_DATASTORE* value is the datastore name as it is displayed in the vSphere Web Client.

For more information about the *VSPHERE_PLACEMENT_LOCATOR* and *VSPHERE_PLACEMENT_DATASTORE* values, see *Specifying the Inventory Path for a Cluster, Host, or Resource Pool* in the [OVF Tool User's Guide](#).

Important The *Variables* listed in the table are used as an example. Update values to match your environment.

Define Installation Variables	Create Registry File
<pre># export MGMT_VSPHERE_ADDRESS=mgmt-vsphere-address # export MGMT_VSPHERE_USER=mgmt-vsphere-admin-user # export MGMT_VSPHERE_LOCATOR='mgmt-locator' # export MGMT_VSPHERE_DATASTORE='mgmt-datastore' # export MGMT_VSPHERE_NETWORK='mgmt-network' # export VSPHERE01_ADDRESS=vsphere-01-address # export VSPHERE01_PLACEMENT_LOCATOR=vsphere-01-locator # export VSPHERE01_PLACEMENT_DATASTORE=vsphere-01-datastore # export VSPHERE01_PLACEMENT_NETWORK=vsphere-01-network # export VSPHERE02_ADDRESS=vsphere-02-address # export VSPHERE02_PLACEMENT_LOCATOR=vsphere-02-locator # export VSPHERE02_PLACEMENT_DATASTORE=vsphere-02-datastore # export VSPHERE02_PLACEMENT_NETWORK=vsphere-02-network # export VCD_ADDRESS=vcd-01-address # export VCD_USER=root@system # export SSO_USER=administrator@vsphere.local</pre>	<p>1 Create a <code>~/.vcav/registry</code> file to hold installation variables. Update the values to match your environment.</p> <pre>vsphere mgmt-vsphere-name address mgmt-vsphere-address api-port 443 api-user admin-user api-password admin-user-password placement-locator mgmt-locator placement-datastore mgmt-datastore placement-network mgmt-network vsphere vsphere-01-name address vsphere-01-address api-port 443 api-user vsphere-01-admin-user api-password vsphere-01-admin-password placement-locator vsphere-01-locator placement-datastore vsphere-01-datastore placement-network vsphere-01-network vsphere vsphere-02-name address vsphere-02-address api-port 443 api-user vsphere-02-admin-user api-password admin-user-password placement-locator vsphere-02-locator placement-datastore vsphere-02-datastore placement-network vsphere-02-network VCD vcd-01-name address vcd-01-address api-port 443 api-user administrator@system api-password vcd-root-password sso-user administrator@vsphere.local sso-password sso-password</pre> <p>2 Update the file permissions</p> <pre># chmod 0600 ~/.vcav/registry</pre>

If you are using a registry file to work with the vCloud Availability Installer Appliance, you can replace the `--vsphere-address`, `--vsphere-user`, and `--vsphere-password-file` options with the `--vsphere=vsphere-name` argument.

If you are using a registry file to work with the vCloud Availability Installer Appliance, you can replace the `--vcd-address`, `--vcd-user`, and `--vcd-password-file` options with `--vcd=vcd-name`.

Add Trusted Thumbprints to the vCloud Availability Installer Appliance

The vCloud Availability Installer Appliance must be able to verify the thumbprint of the vCenter Server and vCloud Director hosts that it works with.

To achieve this, you first import the SSL certificate thumbprint of these hosts into the vCloud Availability Installer Appliance, by running the `vcav trust add` command. The command displays the thumbprint that you are importing. For security purposes, you must verify that the displayed thumbprint matches the actual server certificate.

If the SSL certificate of one of the servers changes, rerun the `vcav trust add` command for that host.

Procedure

- 1 Create a trust between your vSphere instances and the vCloud Availability Installer Appliance.

Repeat this step for every vCenter Server.

a

Standard Command	Command Using Registry
<pre># vcav trust add --address=\$VSPHERE01_ADDRESS --port=443 --accept-all</pre>	<pre># vcav trust add --vsphere=vsphere-01-name --accept-all</pre>

b

Standard Command	Command Using Registry
<pre># vcav trust add --address=\$VSPHERE02_ADDRESS --port=443 --accept-all</pre>	<pre># vcav trust add --vsphere=vsphere-02-name --accept-all</pre>

c

Standard Command	Command Using Registry
<pre># vcav trust add --address=\$MGMT_VSPHERE_ADDRESS --port=443 --accept-all</pre>	<pre># vcav trust add --vsphere=mgmt-vsphere-name --accept-all</pre>

- 2 Create a trust with vCloud Director.

Standard Command	Command Using Registry
<pre># vcav trust add --address=\$VCD_ADDRESS --port=443 --accept-all</pre>	<pre># vcav trust add --vcd=vcd-01-name --accept-all</pre>

Enable Static IP Addresses Deployment

By default, the vCloud Availability Installer Appliance creates VMs with DHCP. You can apply static IP addresses by adding the `--vm-address` option to any command that deploys an OVF.

Depending on your environment topology, you may need to create several IP pools. For example, you might need to create one IP pool in the Resource vCenter Server, one in the Management vCenter Server, and another one in the Management vCenter Server for the vCloud Availability Portal in the DMZ with different network settings. In the current example, we create an IP pool in the Resource vCenter Server. If necessary, repeat the procedure for your Management vCenter Server and the DMZ networks.

You must add an IP Pool and IP Range in the Management or Resource vCenter Server for the network that you want to manage. The IP Pool objects assign all network parameters to VMs, except the IP address. The IP Pool object also ensures that the desired IP is supported on the requested network.

Procedure

- 1 List existing IP Pools, defined in your environment.

Standard Command	Command Using Registry
<pre># vcav ip-pool list \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso</pre>	<pre># vcav ip-pool list --vsphere=<i>vsphere-01-name</i></pre>

The system displays the following result if you have no IP pools.

```
BackingDC
  No IP Pools
VC4
  No IP Pools
```

- 2 Create an IP pool.

The values used in the following command are used as examples. Update the values in the command to match your environment.

Standard Command	Command Using Registry
<pre># vcav ip-pool create \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --datacenter=VC4 \ --name=WDC3-Routed \ --subnet=10.158.12.0 \ --gateway=10.158.15.253 \ --netmask=255.255.252.0 \ "--dns=10.158.12.104,10.158.12.105" \ "--networks=VM Network"</pre>	<pre># vcav ip-pool create \ --vsphere=<i>vsphere-01-name</i> \ --datacenter=VC4 \ --name=WDC3-Routed \ --subnet=10.158.12.0 \ --gateway=10.158.15.253 \ --netmask=255.255.252.0 \ "--dns=10.158.12.104,10.158.12.105" \ "--networks=VM Network" \</pre>

You created an IP pool. Running the `vcav ip-pool list` command now brings the following output.

```
BackingDC
  No IP Pools
VC4
  WDC3-Routed
    Networks:    VM Network
    IPv4 Subnet: 10.158.12.0
    IPv4 Gateway: 10.158.15.253
    IPv4 Netmask: 255.255.252.0
    IPv4 DNS:    10.158.12.104, 10.158.12.105
    IPv4 DHCP:   False
```

3 Associate the IP pool object with more networks.

Standard Command	Command Using Registry
<pre># vcav ip-pool update \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --datacenter=VC4 \ --name=WDC3-Routed \ "--networks=VM Network, Private Network"</pre>	<pre># vcav ip-pool update \ --vsphere=vsphere-01-name \ --datacenter=VC4 \ --name=WDC3-Routed \ "--networks=VM Network, Private Network"</pre>

The associated networks are updated. Running the `vcav ip-pool list` command now results in the following output:

```
BackingDC
  No IP Pools
VC4
  WDC3-Routed
    Networks:    VM Network, Private Network
    IPv4 Subnet: 10.158.12.0
    IPv4 Gateway: 10.158.15.253
    IPv4 Netmask: 255.255.252.0
    IPv4 DNS:    10.158.12.104, 10.158.12.105
    IPv4 DHCP:   False
```

You created an IP pool in your environment and can create VMs with static IP addresses by adding `--vm-address` in any command that deploys an OVF.

Configure vCloud Director

You must perform an additional vCloud Director configuration in case of manual deployment.

Port 5671 is used for AMQP messaging over SSL. SSL connections are recommended, but if there is a requirement to use non-SSL connections for vCloud Director, you can add the `--amqp-port=port-number` argument to the `vcav hcs configure` command. For more information, see [Configure vSphere Replication Cloud Service](#). You can configure RabbitMQ to listen on both SSL and non-SSL ports. For more information, see [Configure a Primary RabbitMQ Server](#). For more information about configuring an AMQP broker, see the *vCloud Director Administrator's Guide*.

To configure vCloud Director to use the RabbitMQ Servers, do the following:

- 1 Create a trusted connection between the RabbitMQ host and the vCloud Availability Installer Appliance.

```
# vcav trust add \  
--address=$AMQP_ADDRESS \  
--port=5671 \  
--accept-all
```

- 2 Register the RabbitMQ host with vCloud Director by running the following command on the vCloud Availability Installer Appliance:

Standard Command	Command Using Registry
<pre># vcav vcd configure-amqp \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --amqp-address=\$AMQP_ADDRESS \ --amqp-port=5671 \ --amqp-user=vcd \ --amqp-password-file=~/.ssh/.amqp \ --amqp-vhost=/ \ --amqp-exchange=systemExchange</pre>	<pre># vcav vcd configure-amqp \ --vcd=vcd-01-name \ --amqp-address=\$AMQP_ADDRESS \ --amqp-port=5671 \ --amqp-user=vcd \ --amqp-password-file=~/.ssh/.amqp \ --amqp-vhost=/ \ --amqp-exchange=systemExchange</pre>

- 3 Restart vCloud Director and Cloud Proxy hosts after configuring AMQP settings, by creating an SSH connection to the hosts and restarting the `vmware-vcd` service.

Check vCloud Director Endpoints

Verify that your environment is properly configured for vCloud Availability installation, by checking the vCloud Director endpoints for known problems.

This check verifies the connectivity between vCloud Director and all related vCenter Server instances.

Procedure

- ◆ To verify that your environment is properly configured, run the following command.

Standard Command	Command Using Registry
<pre># vcav vcd check \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav vcd check --vcd=vcd-01-name</pre>

The system returns an **OK** message upon successful validation.

Installing vCloud Availability

Before configuring the vCloud Availability solution, you must deploy virtual appliances to support vCloud Availability components.

You install all individual components of vCloud Availability by using the vCloud Availability Installer Appliance.

Both ways to deploy and configure vCloud Availability are demonstrated for your reference.

You run all installation commands from the vCloud Availability Installer Appliance, unless documentation instructs otherwise.

Note You use the `vcav component-alias create` command to deploy all vCloud Availability components. By default, the value of the `--vm-name=` argument defines the hostname of the virtual machine you are creating. You can assign a specific hostname to all VMs you deploy by adding the `--hostname=desired-hostname` argument to the `vcav component-alias create` command.

Create vSphere Replication Manager

The vSphere Replication Manager manages and monitors the replication process from tenant VMs to the service provider environment. A vSphere Replication management service runs for each vCenter Server and tracks changes to VMs and infrastructure related to replication.

The Resource vCenter Server is a vCenter Server registered to vCloud Director and made available to tenants.

Important Deploy one vSphere Replication Manager for each Resource vCenter Server. The total number of vSphere Replication Management servers depends on your environment and deployment requirements.

Procedure

- 1 Create an SSH connection to the vCloud Availability Installer Appliance using your **root** credentials. You run all installation and configuration commands from the vCloud Availability Installer Appliance.
- 2 Create a vSphere Replication Manager.

Standard Command	Command Using Registry
<pre># vcav hms create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ "--network=\$VSPHERE01_PLACEMENT_NETWORK" \ "--vsphere-locator= \$VSPHERE01_PLACEMENT_LOCATOR" \ --datastore=\$VSPHERE01_PLACEMENT_DATASTORE \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=hms01-name</pre>	<pre># vcav hms create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --vsphere=vsphere-01-name \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=hms01-name</pre>

The IP address of the new vSphere Replication Manager is displayed. Write it down because you need it during the configuration.

Repeat this step for every resource vCenter Server in your environment.

3 If necessary, unregister the vSphere Replication extension from the vSphere Web Client.

By default, the vSphere Replication Manager registers as an extension to the instance of vSphere it is deployed to. This model is called in inventory deployment. For in inventory deployments, the vSphere Replication Manager manages the replications to the vSphere instance it is deployed to. In such cases, you must skip this step.

You can deploy a vSphere Replication Manager to an infrastructure pool that tenants are not using and register the vSphere Replication Manager to a resource pool instance of vSphere. This model is called out of inventory deployment. For out of inventory deployments, the vSphere Replication Manager does not manage the replications to the vSphere instance it is deployed to. In such cases, you must unregister the vSphere Replication extension by running the following command.

Standard Command	Command Using Registry
<pre># vcav hms unregister-extension \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso</pre>	<pre># vcav hms unregister-extension \ --vsphere=vsphere-01-name</pre>

4 Set a variable to the address of the created virtual machine.

Standard Command	Command Using Registry
<pre># HMS01_ADDRESS=`vcav vsphere get-ip \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ "--network=\$VSPHERE01_PLACEMENT_NETWORK" \ --vm-name=hms01-name`</pre>	<pre># HMS01_ADDRESS=`vcav vsphere get-ip \ --vsphere=vsphere-01-name \ "--network=vsphere-01-network" \ --vm-name=hms01-name`</pre>

5 If you use Fully Qualified Domain Names (FQDN) to access and manage appliances, you must verify that the DNS record matches the vSphere Replication Manager IP address and trusts the SSH certificate for this FQDN.

- Check the DNS server to ensure that the entry matches the IP address of the vSphere Replication Manager.
- Run the following command to trust the certificate for the vSphere Replication Manager FQDN.

Standard Command	Command Using Registry
<pre># vcav vsphere trust-ssh \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --root-password-file=~/.ssh/.root \ --vm-name=hms01-name \ --vm-address=hms01-FQDN</pre>	<pre># vcav vsphere trust-ssh \ --vsphere=vsphere-01-name \ --root-password-file=~/.ssh/.root \ --vm-name=hms01-name \ --vm-address=hms01-FQDN</pre>

Create vSphere Replication Cloud Service Host

The vSphere Replication Cloud Service is a tenant-aware replication manager that provides the required API for managing the service and all the components. vSphere Replication Cloud Service registers as a vCloud Director extension and is accessible through the vCloud Director interface.

Procedure

- 1 Create a vSphere Replication Cloud Service host.

Standard Command	Command Using Registry
<pre># vcav hcs create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password-file=~/.ssh/.vsphere.mgmt \ "--network=\$MGMT_VSPHERE_NETWORK" \ "--vsphere-locator=\$MGMT_VSPHERE_LOCATOR" \ --datastore=\$MGMT_VSPHERE_DATASTORE \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=hcs01-name</pre>	<pre># vcav hcs create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --vsphere=mgmt-vsphere-name \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=hcs01-name</pre>

The IP address of the new vSphere Replication Cloud Service host is displayed. Write it down because you need it during the configuration.

- 2 Set a variable to the address of the created virtual machine.

Standard Command	Command Using Registry
<pre># HCS01_ADDRESS=`vcav vsphere get-ip \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password-file=~/.ssh/.vsphere.mgmt \ "--network=\$MGMT_VSPHERE_NETWORK" \ --vm-name=hcs01-name`</pre>	<pre># HCS01_ADDRESS=`vcav vsphere get-ip \ --vsphere=mgmt-vsphere-name \ "--network=mgmt-network" \ --vm-name=hcs01-name`</pre>

- 3 If you use Fully Qualified Domain Names (FQDN) to access and manage appliances, you must verify that the DNS record matches the vSphere Replication Cloud Service host IP address, and trusts the SSH certificate for this FQDN.
 - a Check the DNS server to ensure that the record matches the IP address of the vSphere Replication Cloud Service host.
 - b Run the following command to trust the certificate for the vSphere Replication Cloud Service host FQDN.

Standard Command	Command Using Registry
<pre># vcav vsphere trust-ssh \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password- file=~/.ssh/.vsphere.mgmt \ --root-password-file=~/.ssh/.root \ --vm-name=hcs01-name \ --vm-address=hcs01-FQDN</pre>	<pre># vcav vsphere trust-ssh \ --vsphere=mgmt-vsphere-name \ --root-password-file=~/.ssh/.root \ --vm-name=hcs01-name \ --vm-address=hcs01-FQDN</pre>

Create vSphere Replication Server

The vSphere Replication Server handles the replication process for each protected virtual machine.

Important Deploy at least one vSphere Replication Server for each vSphere Replication Manager.

Procedure

- 1 Create vSphere Replication Server.

Standard Command	Command Using Registry
<pre># vcav hbr create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ "--network=\$VSPHERE01_PLACEMENT_NETWORK" \ "--vsphere-locator= \$VSPHERE01_PLACEMENT_LOCATOR" \ --datastore=\$VSPHERE01_PLACEMENT_DATASTORE \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=hbr01-name</pre>	<pre># vcav hbr create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --vsphere=vsphere-01-name \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=hbr01-name</pre>

The IP address of the new vSphere Replication Server is displayed. Write it down because you need it during the configuration.

Important Repeat this step for every vSphere Replication Manager in your environment.

2 Set a variable to the address of the created virtual machine.

Standard Command	Command Using Registry
<pre># HBR_ADDRESS=`vcav vsphere get-ip \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ "--network=\$VSPHERE01_PLACEMENT_NETWORK" \ --vm-name=hbr01-name`</pre>	<pre># HBR_ADDRESS=`vcav vsphere get-ip \ --vsphere=vsphere-01-name \ "--network=vsphere-01-network" \ --vm-name=hbr01-name`</pre>

3 If you use Fully Qualified Domain Names (FQDN) to access and manage appliances, you must verify that the DNS record matches the vSphere Replication Server IP address, and trust the SSH certificate for this FQDN.

- a Check the DNS server to ensure that the entry matches the IP address of the vSphere Replication Server.
- b Run the following command to trust the certificate for the vSphere Replication Server FQDN.

Standard Command	Command Using Registry
<pre># vcav vsphere trust-ssh \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --root-password-file=~/.ssh/.root \ --vm-name=hbr01-name \ --vm-address=hbr01-FQDN</pre>	<pre># vcav vsphere trust-ssh \ --vsphere=vsphere-01-name \ --root-password-file=~/.ssh/.root \ --vm-name=hbr01-name \ --vm-address=hbr01-FQDN</pre>

Create vCloud Availability Portal Host

The vCloud Availability Portal provides a graphic user interface to facilitate the management of vCloud Availability operations.

The vCloud Availability Portal back end (PBE) scales horizontally. You can deploy a new vCloud Availability Portal instance on demand connected to the same load balancer that all the vCloud Availability Portal instances are under. The load balancer must support sticky sessions, so that the same PBE instance processes user requests within a session. This setting ensures that all the information displayed in the vCloud Availability Portal is consistent.

Depending on the number of concurrent sessions that the vCloud Availability Portal is expected to host, you can deploy *small*, *medium*, or *large* vCloud Availability Portal host. The vCloud Availability Portal sends requests to a vCloud Director instance and receives data from the same vCloud Director instance. To host the maximum number of concurrent sessions, ensure that the vCloud Director database can use similar compute resources that you allocate to the vCloud Availability Portal host. You can find details about the vCloud Availability Portal deployment types in the following table.

Table 3-2. vCloud Availability Portal Host Deployment Types

Deployment Type	Description
Small	Deploys an appliance with 2 CPUs, 2 GB of memory, 10 GB of disk space, and 512 MB of Java Virtual Memory. Suitable for hosting up to 150 concurrent sessions.
Medium	Deploys an appliance with 2 CPUs, 4 GB of memory, 10 GB of disk space, and 1.5 GB of Java Virtual Memory. Suitable for hosting up to 400 concurrent sessions.
Large	Deploys an appliance with 4 CPUs, 6 GB of memory, 10 GB of disk space, and 3 GB of Java Virtual Memory. Suitable for hosting up to 800 concurrent sessions.

Procedure

- 1 Create a vCloud Availability Portal host by running the following command.

Important The `--deployment-type` argument in the following command defines the compute resources that you allocate to the vCloud Availability Portal host. By default, the value is `small`. You can change the value depending on your requirements.

Standard Command	Command Using Registry
<pre># vcav vcd-ui create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --deployment-type=small \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password-file=~/.ssh/.vsphere.mgmt \ "--network=\$MGMT_VSPHERE_NETWORK" \ "--vsphere-locator=\$MGMT_VSPHERE_LOCATOR" \ --datastore=\$MGMT_VSPHERE_DATASTORE \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=ui01-name</pre>	<pre># vcav vcd-ui create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --deployment-type=small \ --vsphere=mgmt-vsphere-name \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=ui01-name</pre>

The IP address of the new vCloud Availability Portal virtual machine is displayed. Write it down because you need it during the configuration.

- 2 Set a variable to the address of the created virtual machine.

Standard Command	Command Using Registry
<pre># UI01_ADDRESS=`vcav vsphere get-ip \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password-file=~/.ssh/.vsphere.mgmt \ "--network=\$MGMT_VSPHERE_NETWORK" \ --vm-name=ui01-name`</pre>	<pre># UI01_ADDRESS=`vcav vsphere get-ip \ --vsphere=mgmt-vsphere-name \ "--network=mgmt-network" \ --vm-name=ui01-name`</pre>

3 Update the truststore file with the vCloud Availability Portal virtual machine credentials.

```
# echo 'Portal-VM-Password' > ~/.ssh/.truststore

# chmod 0600 ~/.ssh/.truststore
```

4 If you use Fully Qualified Domain Names (FQDN) to access and manage appliances, you must verify that the DNS record matches the vCloud Availability Portal IP address, and trusts the SSH certificate for this FQDN.

- a Check the DNS server to ensure that the entry matches the IP address of the vCloud Availability Portal host.
- b Run the trust-ssh command to trust the certificate for the vCloud Availability Portal FQDN.

Standard Command	Command Using Registry
<pre># vcav vsphere trust-ssh \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password- file=~/.ssh/.vsphere.mgmt \ --root-password-file=~/.ssh/.root \ --vm-name=ui01-name \ --vm-address=ui01-FQDN</pre>	<pre># vcav vsphere trust-ssh \ --vsphere=mgmt-vsphere-name \ --root-password-file=~/.ssh/.root \ --vm-name=ui01-name \ --vm-address=ui01-FQDN</pre>

Create vCloud Availability Administration Portal Host

The vCloud Availability Administration Portal provides a graphic user interface to facilitate the service providers to monitor and manage their DR environments.

Important You must deploy vCloud Availability Administration Portal separately from vCloud Availability Portal. The best practice is to deploy it behind a VPN to limit access.

Depending on the number of concurrent sessions that the vCloud Availability Administration Portal is expected to host, you can deploy `small`, `medium`, or `large` vCloud Availability Administration Portal host. The vCloud Availability Administration Portal sends requests to a vCloud Director instance and receives data from the same vCloud Director instance. To host the maximum number of concurrent sessions, ensure that the vCloud Director database can use similar compute resources that you allocate to the vCloud Availability Administration Portal host. You can find details about the vCloud Availability Administration Portal deployment types in the following table.

Table 3-3. vCloud Availability Administration Portal Host Deployment Types

Deployment Type	Description
Small	Deploys an appliance with 2 CPUs, 2 GB of memory, 10 GB of disk space, and 512 MB of Java Virtual Memory. Suitable for hosting up to 150 concurrent sessions.
Medium	Deploys an appliance with 2 CPUs, 4 GB of memory, 10 GB of disk space, and 1.5 GB of Java Virtual Memory. Suitable for hosting up to 400 concurrent sessions.
Large	Deploys an appliance with 4 CPUs, 6 GB of memory, 10 GB of disk space, and 3 GB of Java Virtual Memory. Suitable for hosting up to 800 concurrent sessions.

Procedure

- 1 Create a vCloud Availability Administration Portal host by running the following command.

Important The `--deployment-type` argument in the following command defines the compute resources that you allocate to the vCloud Availability Administration Portal host. By default, the value is `small`. You can change the value depending on your requirements.

Standard Command	Command Using Registry
<pre># vcav vcd-ui create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --deployment-type=small \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password-file=~/.ssh/.vsphere.mgmt \ "--network=\$MGMT_VSPHERE_NETWORK" \ "--vsphere-locator=\$MGMT_VSPHERE_LOCATOR" \ --datastore=\$MGMT_VSPHERE_DATASTORE \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=ui02-name</pre>	<pre># vcav vcd-ui create \ --ovf-url=vcloud-availability-release_number- xxx-build_number_OVF10.ova \ --deployment-type=small \ --vsphere=mgmt-vsphere-name \ --ntp=pool.ntp.org \ --root-password-file=~/.ssh/.root \ --vm-name=ui02-name</pre>

The IP address of the new vCloud Availability Administration Portal virtual machine is displayed. Write it down because you need it during the configuration.

- 2 Set a variable to the address of the created virtual machine.

Standard Command	Command Using Registry
<pre># UI02_ADDRESS=`vcav vsphere get-ip \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password-file=~/.ssh/.vsphere.mgmt \ "--network=\$MGMT_VSPHERE_NETWORK" \ --vm-name=ui02-name`</pre>	<pre># UI02_ADDRESS=`vcav vsphere get-ip \ --vsphere=mgmt-vsphere-name \ "--network=mgmt-network" \ --vm-name=ui02-name`</pre>

- 3 Update the truststore file with the vCloud Availability Administration Portal virtual machine credentials.

```
# echo 'SMP-Portal-VM-Password' > ~/.ssh/.truststore

# chmod 0600 ~/.ssh/.truststore
```

- 4 If you use Fully Qualified Domain Names (FQDN) to access and manage appliances, you must verify that the DNS record matches the vCloud Availability Administration Portal IP address, and trusts the SSH certificate for this FQDN.
- Check the DNS server to ensure that the entry matches the IP address of the vCloud Availability Administration Portal host.
 - Run the `trust-ssh` command to trust the certificate for the vCloud Availability Administration Portal FQDN.

Standard Command	Command Using Registry
<pre># vcav vsphere trust-ssh \ --vsphere-address=\$MGMT_VSPHERE_ADDRESS \ --vsphere-user=\$MGMT_VSPHERE_USER \ --vsphere-password- file=~/.ssh/.vsphere.mgmt \ --root-password-file=~/.ssh/.root \ --vm-name=ui02-name \ --vm-address=ui02-FQDN</pre>	<pre># vcav vsphere trust-ssh \ --vsphere=mgmt-vsphere-name \ --root-password-file=~/.ssh/.root \ --vm-name=ui02-name \ --vm-address=ui02-FQDN</pre>

Validate Deployment

Before you configure vCloud Availability, you must confirm that all appliances are ready to be configured.

Procedure

- Verify that all components are ready for configuration.

Standard Command	Command Using Registry
<pre># vcav vcd wait-for-api \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --timeout=300 # vcav vcd is-federation-enabled \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd</pre>	<pre># vcav vcd wait-for-api \ --vcd=vcd-01-name \ --timeout=300 # vcav vcd is-federation-enabled \ --vcd=vcd-01-name</pre>

Configuring vCloud Availability for vCloud Director

After you deploy all individual components of vCloud Availability, you must configure them to support DRaaS.

Procedure

1 Configure vSphere Replication Manager

Each vSphere Replication Manager must be registered to a single vCenter Server.

2 Configure Cassandra Servers

Update each Cassandra server to trust every vSphere Replication Cloud Service appliance and register each Cassandra server with the lookup service used by vCloud Director.

3 Configure vSphere Replication Cloud Service

To configure the vSphere Replication Cloud Service host, you must register each vSphere Replication Cloud Service appliance to your vCloud Director appliance, resource vCenter Server, and RabbitMQ.

4 Configure vSphere Replication Server

Attach each vSphere Replication Server to your vSphere Replication Manager and vCenter Server.

5 Configure vCloud Availability Portal Host

You must configure the vCloud Availability Portal host.

6 Configure vCloud Availability Administration Portal Host

You must configure the vCloud Availability Administration Portal host with both the vCloud Director server and the embedded MongoDB server, and start the system services.

Configure vSphere Replication Manager

Each vSphere Replication Manager must be registered to a single vCenter Server.

Procedure

1 Configure the vSphere Replication Manager.

Standard Command	Command Using Registry
<pre># vcav hms configure \ --hms-address=\$HMS01_ADDRESS \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav hms configure \ --hms-address=\$HMS01_ADDRESS \ --vsphere=vsphere-01-name \ --vcd=vcd-01-name</pre>

The system returns an OK message, after the process finishes.

- 2 Run the following command to verify that the hms service starts successfully.

Standard Command	Command Using Registry
<pre># vcav hms wait-for-extension \ --hms-address=\$HMS01_ADDRESS \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav hms wait-for-extension \ --hms-address=\$HMS01_ADDRESS \ --vsphere=vsphere-01-name \ --vcd=vcd-01-name</pre>

If the hms service starts successfully, the system displays an OK message.

If the system returns an error, or there is no output in 5 minutes, check the `/opt/vmware/logs/hms/hms.log` file for errors.

Important Repeat these steps for every vSphere Replication Manager that you deployed.

Configure Cassandra Servers

Update each Cassandra server to trust every vSphere Replication Cloud Service appliance and register each Cassandra server with the lookup service used by vCloud Director.

Note The current procedure pertains to configuring Cassandra servers for production deployments. If you use a Docker container to manage your Cassandra servers in test and development environments, perform the steps documented in [Deploy Cassandra and RabbitMQ as Containers for Test and Development Environments](#).

Procedure

- 1 Create a password file for the Cassandra host `root` user in `/.ssh/cassandra.root.password`.
- 2 Create a trusted connection between the vCloud Availability Installer Appliance and your Cassandra hosts. This connection allows the vCloud Availability Installer Appliance to trust the Cassandra certificate and is required before you can add the Cassandra hosts to the lookup service used by vCloud Director.

Repeat this step for every Cassandra host in your environment.

```
# vcav trust add-ssh --address=$CASSANDRA_ADDRESS \
--root-password-file=/.ssh/cassandra.root.password \
--accept-all
```

- 3 Add the vSphere Replication Cloud Service certificate to the Cassandra truststore, so that the Cassandra host accepts SSL connections from the vSphere Replication Cloud Service.

The Cassandra truststore stores the certificates that are accepted for connection. The Cassandra keystore only stores the certificate that the Cassandra server publishes.

Run the following command on every Cassandra server before you finish the vCloud Availability configuration. Run the commands for each vSphere Replication Cloud Service host.

```
# vcav cassandra import-hcs-certificate \  
--cassandra-address=$CASSANDRA_ADDRESS \  
--hcs-address=$HCS01_ADDRESS
```

If the command cannot find the Cassandra configuration file, you can specify the path to the file by adding the `--cassandra-config-file=path-to-Cassandra-config-file`.

- 4 Register the Cassandra hosts with the lookup service by running the following command.

Repeat this step for every Cassandra host in your environment.

Standard Command	Command Using Registry
<pre># vcav cassandra register \ --hcs-address=\$HCS01_ADDRESS \ --cassandra-address=\$CASSANDRA_ADDRESS \ --cassandra-port=9042 \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav cassandra register \ --hcs-address=\$HCS01_ADDRESS \ --cassandra-address=\$CASSANDRA_ADDRESS \ --cassandra-port=9042 \ --vcd=vcd-01-name</pre>

The system displays an OK message upon a successful registration.

Configure vSphere Replication Cloud Service

To configure the vSphere Replication Cloud Service host, you must register each vSphere Replication Cloud Service appliance to your vCloud Director appliance, resource vCenter Server, and RabbitMQ.

Important If you have more than one vCloud Director instance configured in your vCenter Server lookup service, the vSphere Replication Cloud Service VM registers to the first vCloud Director instance in the lookup service.

Procedure

- 1 Configure the vSphere Replication Cloud Service Appliance.

The `cassandra-replication-factor` argument in the following command defines the number of data replicas across the Cassandra cluster. A replication factor 4 means that there are four copies of each row, where each copy is on a different node. The replication factor must not exceed the number of nodes in the Cassandra cluster.

By default, the following command uses the AMQP settings from vCloud Director. If vCloud Director is not using an SSL port for AMQP, the `vcav hcs configure` operation returns an error. You can add the `--amqp-port=port-number` argument to override the vCloud Director port and point the AMQP service to an SSL port.

Standard Command	Command Using Registry
<pre># vcav hcs configure \ --hcs-address=\$HCS01_ADDRESS \ --amqp-password-file=~/.ssh/.amqp \ --cassandra-replication-factor=number-of- Cassandra-nodes \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav hcs configure \ --hcs-address=\$HCS01_ADDRESS \ --amqp-password-file=~/.ssh/.amqp \ --cassandra-replication-factor=number-of- Cassandra-nodes \ --vcd=vcd-01-name</pre>

The system returns an OK message, after the process finishes.

- 2 Run the following command to verify that the hcs service starts successfully.

Standard Command	Command Using Registry
<pre># vcav hcs wait-for-extension \ --hcs-address=\$HCS01_ADDRESS \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav hcs wait-for-extension \ --hcs-address=\$HCS01_ADDRESS \ --vcd=vcd-01-name</pre>

If the hcs service starts successfully, the system displays an OK message.

If the system returns an error, or there is no output in 5 minutes, check the /opt/VMware/logs/hms/hcs.log file for errors.

Configure vSphere Replication Server

Attach each vSphere Replication Server to your vSphere Replication Manager and vCenter Server.

Procedure

- 1 Configure your vSphere Replication Server.

Standard Command	Command Using Registry
<pre># vcav hbr configure \ --hbr-address=\$HBR_ADDRESS \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd</pre>	<pre># vcav hbr configure \ --hbr-address=\$HBR_ADDRESS \ --vsphere=vsphere-01-name \ --vcd=vcd-01-name</pre>

The system returns an OK message, after the process finishes.

Important Repeat this step for every vSphere Replication Server in your environment.

2 Verify that the hbr service starts successfully.

Standard Command	Command Using Registry
<pre># vcd hbr wait-for-extension \ --hbr-address=\$HBR_ADDRESS \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd</pre>	<pre># vcd hbr wait-for-extension \ --hbr-address=\$HBR_ADDRESS \ --vsphere=vsphere-01-name \ --vcd=vcd-01-name</pre>

If the hbr service starts successfully, the system displays an OK message.

If the system returns an error, or there is no output in 5 minutes, check the `/var/log/vmware/hbrsrv.log` file for errors.

Configure vCloud Availability Portal Host

You must configure the vCloud Availability Portal host.

Procedure

- 1 Configure the vCloud Availability Portal host by running the following command.

Standard Command	Command Using Registry
<pre># vcav vcd-ui configure \ --ui-address=\$UI01_ADDRESS \ --keep-self-signed-certificate \ --truststore-password- file=~/.ssh/.truststore \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=administrator@vsphere.local \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav vcd-ui configure \ --ui-address=\$UI01_ADDRESS \ --keep-self-signed-certificate \ --truststore-password- file=~/.ssh/.truststore \ --vcd=vcd-01-name</pre>

The system returns an OK message, after the process finishes.

In the example, the vCloud Availability Portal is configured to operate with a new generated self-signed SSL certificate. You can set up the vCloud Availability Portal to use an externally signed SSL certificate, by replacing the `--keep-self-signed-certificate` argument with `--https-certificate=/file-path-to-certificate-file` and `--https-key=/file-path-to-certificate-public-key`. The vCloud Availability Portal appliance provides the certificate and key files to an nginx process.

- 2 You allocate small, medium, and large sizes of Java Virtual Memory (JVM) to the vCloud Availability Portal service process during the host deployment. You can optionally update the allocated JVM by completing the following steps. For more information about the vCloud Availability Portal deployment types and related JVM configuration, see [Create vCloud Availability Portal Host](#).

- a Use SSH to connect to the vCloud Availability Portal host.
- b Use a text editor to open the `/opt/vmware/conf/vcav-ui/nginx/nginx.conf` file.
- c The initial size of the memory allocation pool is defined in the following line. Change the numeric value to designate more JVM to the vCloud Availability Portal host.

```
jvm_options "-Xms1024m";
```

- d The following line defines the maximum size of memory allocation pool for the `nginx` process. The numeric value must be equal to or greater than the numeric value you defined in the previous step.

```
jvm_options "-Xmx1024m";
```

- e You can optionally uncomment the following lines to enable JVM heap dump and define the heap dump file path.

```
jvm_options "-XX:+HeapDumpOnOutOfMemoryError";
jvm_options "-XX:HeapDumpPath=/opt/vmware/logs/vcav-ui/jvm.hprof";
```

- f By default, the maximum number of concurrent client sessions is set to 1024. To increase this number, use a text editor to open the `/usr/lib/systemd/system/vcav-ui.service` and add the following line after the `[Service]` line.

```
LimitNOFILE=8192
```

- g Restart the vCloud Availability Portal service to complete this configuration, by running the following command.

```
systemctl restart vcav-ui
```

- 3 Configure the timeout settings for the vCloud Availability Portal host.

- a Use a text editor to open the `opt/vmware/conf/vcav-ui/config.yml` file.
- b Set the `connectTimeout` value to 60000 and the `socketTimeout` value to 60000.

4 Configure the nginx process to run for a non-root user.

The vCloud Availability Portal host nginx process runs under the system root user by default. You can change the user that the nginx process uses by modifying the vCloud Availability Portal service script. If you do not want to edit the nginx process user, you can skip this step.

- a Use SSH to connect to the vCloud Availability Portal host as root.
- b Stop the vCloud Availability Portal service by running the following command.

```
# systemctl stop vcav-ui
```

- c Use a text editor to modify the `/usr/lib/systemd/system/vcav-ui.service` file, by adding `User=new-user-name` line after the `[Service]` line.
- d Change the line that provides the PID file location to read `PIDFile=/opt/vmware/logs/vcav-ui/vcav-ui.pid`.
- e Using a text editor open the `/opt/vmware/conf/vcav-ui/nginx/nginx.conf` and change the line that provides the PID file location to read `pid /opt/vmware/logs/vcav-ui/vcav-ui.pid`.
- f Change the ownership of the log files that the service uses by running the following commands.

```
# chown -R new-user-name /opt/vmware/logs/vcav-ui
# chown -R new-user-name /opt/vmware/vcav-ui/logs
```

- g Start the vCloud Availability Portal service by running the following command.

```
# systemctl start vcav-ui
```

5 Assign a domain name to your vCloud Availability Portal host.

It is a best practice to assign a domain name to your vCloud Availability Portal VM for production deployments.

- 6 Verify that the vCloud Availability Portal is configured correctly, by running the following command.

```
'curl -k https://$UI01_ADDRESS:8443/
```

Configure vCloud Availability Administration Portal Host

You must configure the vCloud Availability Administration Portal host with both the vCloud Director server and the embedded MongoDB server, and start the system services.

Procedure

- 1 To configure vCloud Availability Administration Portal, run the following command.

Standard Command	Command Using Registry
<pre># vcav vcd-ui configure-smp \ --ui-address=\$UI02_ADDRESS \ --keep-self-signed-certificate \ --truststore-password- file=~/.ssh/.truststore \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=administrator@vsphere.local \ --sso-password-file=~/.ssh/.sso \ --mongodb-password-file=~/.ssh/.root \ --amqp-user= \$AMQP_USER \ --amqp-password-file= ~/.ssh/.amqp \ --tenant-ui-url=https://tenant-ui-FQDN:8443 \ --max-jvm-memory=1024</pre>	<pre># vcav vcd-ui configure-smp \ --ui-address=\$UI02_ADDRESS \ --keep-self-signed-certificate \ --truststore-password- file=~/.ssh/.truststore \ --vcd=vcd-01-name \ --mongodb-password-file=~/.ssh/.root \ --amqp-user= \$AMQP_USER \ --amqp-password-file= ~/.ssh/.amqp \ --tenant-ui-url=https://tenant-ui-FQDN:8443 \ --max-jvm-memory=1024</pre>

The system returns an OK message, after the process finishes.

Important To enable tenant impersonation, you must set `--tenant-ui-url` argument value to the base URL of the tenant vCloud Availability Portal.

The `--amqp-user` and `--amqp-password-file` argument values are mandatory.

In the example, the vCloud Availability Administration Portal is configured to operate with a new generated self-signed SSL certificate. You can set up the vCloud Availability Administration Portal to use an externally signed SSL certificate, by replacing the `--keep-self-signed-certificate` argument with `--https-certificate=/file-path-to-certificate-file` and `--https-key=/file-path-to-certificate-public-key`. The vCloud Availability Administration Portal appliance provides the certificate and key files to a `java` process.

- 2 Assign a domain name to your vCloud Availability Administration Portal host. It is a best practice to assign a domain name to your vCloud Availability Administration Portal host for production deployments.
- 3 Verify that the vCloud Availability Administration Portal is configured correctly, by running the following command.

```
curl -k https://$UI02_ADDRESS:8443/
```

Post-Installation vCloud Director Configuration

After you install and configure the vCloud Availability solution, you must perform some configurations to vCloud Director.

Note If RabbitMQ was not configured in vCloud Director before you installed vCloud Availability, you must restart the vCloud Director service.

- 1 Create an SSH connection to the vCloud Director host.
- 2 Restart the `vmware-vcd` service by running the `service vmware-vcd restart` command.

Procedure

1 [Assign vSphere Replication Cloud Service Rights to the vCloud Director Organization Administrator Role](#)

You must assign vSphere Replication Cloud Service rights to the vCloud Director organization administrator role before you enable vCloud Director organization VDC for replication.

2 [Enable a vCloud Director Organization VDC for Replication](#)

You must enable a vCloud Director organization VDC for replication before you start using the vCloud Availability solution.

Assign vSphere Replication Cloud Service Rights to the vCloud Director Organization Administrator Role

You must assign vSphere Replication Cloud Service rights to the vCloud Director organization administrator role before you enable vCloud Director organization VDC for replication.

Procedure

- 1 Create an SSH connection to the vCloud Availability Installer Appliance.
- 2 Assign vSphere Replication Cloud Service rights to the vCloud Director *organization administrator* role.
 - For vCloud Director 8.10 and earlier, you assign vSphere Replication Cloud Service rights to the *organization administrator* role and it applies to all organizations.

Standard Command	Command Using Registry
<pre># vcav hcs add-rights-to-role \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ "--role=Organization Administrator"</pre>	<pre># vcav hcs add-rights-to-role \ --vcd=vcd-01-name \ "--role=Organization Administrator"</pre>

- For vCloud Director 8.20 and above, you assign vSphere Replication Cloud Service rights to the *organization administrator* role for each organization or for all organizations.

	Standard Command	Command Using Registry
For each organization	<pre># vcav hcs add-rights-to- role \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password- file=~/.ssh/.vcd \ "--role=Organization Administrator" \ --org=org-name</pre>	<pre># vcav hcs add-rights-to- role \ --vcd=vcd-01-name \ "--role=Organization Administrator" \ --org=org-name</pre>
For all organizations	<pre># vcav hcs add-rights-to- role \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password- file=~/.ssh/.vcd \ "--role=Organization Administrator" \ --org=*</pre>	<pre># vcav hcs add-rights-to- role \ --vcd=vcd-01-name \ "--role=Organization Administrator" \ --org=*</pre>

Note You do not need to restart any component for the changes to take effect.

Enable a vCloud Director Organization VDC for Replication

You must enable a vCloud Director organization VDC for replication before you start using the vCloud Availability solution.

Prerequisites

Verify that you have assigned vSphere Replication Cloud Service rights to the vCloud Director *organization administrator* role for the organization. For more information, see the `vcav hcs add-rights-to-role` command in [Assign vSphere Replication Cloud Service Rights to the vCloud Director Organization Administrator Role](#).

Procedure

- 1 Create an SSH connection to the vCloud Availability Installer Appliance.

- 2 To enable a vCloud Director organization VDC for replication, run the following commands.

Standard Command	Command Using Registry
<pre># vcav org list \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd # vcav org-vdc list \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --org=org1 # vcav org-vdc enable-replication \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --org=org1 \ --vdc=vdc_org1</pre>	<pre># vcav org list \ --vcd=vcd-01-name # vcav org-vdc list \ --vcd=vcd-01-name \ --org=org1 # vcav org-vdc enable-replication \ --vcd=vcd-01-name \ --org=org1 \ --vdc=vdc_org1</pre>

After the process finishes, you get an OK message.

Unconfiguring vCloud Availability

You unconfigure a vCloud Availability instance by using the vCloud Availability Installer Appliance scripts and the vSphere Web Client capabilities.

Before you unconfigure vCloud Availability, all replications must be stopped and all cloud sites must be disconnected from the tenant on-premises environments.

Unconfigure vCloud Availability

You unconfigure the vCloud Availability solution by cleaning up the components configuration in a specific order.

Procedure

1 Unconfigure the vSphere Replication Server.

Standard Command	Command Using Registry
<pre># vcav hbr unconfigure \ --hbr-address=\$HBR_ADDRESS \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd</pre>	<pre># vcav hbr unconfigure \ --hbr-address=\$HBR_ADDRESS \ --vsphere=vsphere-01-name \ --vcd=vcd-01-name</pre>

Important Repeat this step for every vSphere Replication Server in your vCloud Availability environment. The vCloud Availability Installer Appliance returns an error when you try to unconfigure the last vSphere Replication Server in your vCloud Availability environment. You can ignore the error and proceed to the next step.

2 Unconfigure the vSphere Replication Cloud Service Appliance.

Standard Command	Command Using Registry
<pre># vcav hcs unconfigure \ --hcs-address=\$HCS01_ADDRESS \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav hcs unconfigure \ --hcs-address=\$HCS01_ADDRESS \ --vcd=vcd-01-name</pre>

Important Repeat this step for every vSphere Replication Cloud Service host in your vCloud Availability environment.

3 Remove all **com.vmware.vr** rights from vCloud Director roles that are using them.

4 Unregister the vSphere Replication Cloud Service extension from vCloud Director.

Standard Command	Command Using Registry
<pre># vcav hcs unregister-extension \ --hcs-address=\$HCS01_ADDRESS \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav hcs unregister-extension \ --hcs-address=\$HCS01_ADDRESS \ --vcd=vcd-01-name</pre>

You run this command once per vCloud Availability instance.

5 Unconfigure the vSphere Replication Manager.

Standard Command	Command Using Registry
<pre># vcav hms unconfigure \ --hms-address=\$HMS01_ADDRESS \ --vsphere-address=\$VSPHERE01_ADDRESS \ --vsphere-user=\$SSO_USER \ --vsphere-password-file=~/.ssh/.sso \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --sso-user=\$SSO_USER \ --sso-password-file=~/.ssh/.sso</pre>	<pre># vcav hms unconfigure \ --hms-address=\$HMS01_ADDRESS \ --vsphere=vsphere-01-name \ --vcd=vcd-01-name</pre>

Important Repeat this step for every vSphere Replication Manager in your vCloud Availability environment.

6 Clean up Cassandra endpoints from the lookup service.

- a List all Cassandra hosts that are configured with a vSphere Replication Cloud Service appliance and are registered to a lookup service used by vCloud Director.

Standard Command	Command Using Registry
<pre>vcav hcs list-cassandra \ --hcs-address=hcs-IP-address \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd</pre>	<pre>vcav hcs list-cassandra \ --hcs-address=hcs-IP-address \ --vcd=vcd-01-name</pre>

- b Remove all Cassandra entries from the vCloud Director lookup service.

Standard Command	Command Using Registry
<pre>vcav hcs unregister-cassandra \ --hcs-address=hcs-IP-address \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --all</pre>	<pre>vcav hcs unregister-cassandra \ --hcs-address=hcs-IP-address \ --vcd=vcd-01-name \ --all</pre>

7 Delete the Cassandra key space.

- a Connect to the Cassandra host over SSH.
- b To delete the Cassandra key space, run the following command.

```
drop keyspace keyspace_name
```

By default, the Cassandra key space name is `vr2c` and is defined in the `/opt/vmware/hms/conf/hcs-config.properties` file of the associated vSphere Replication Cloud Service host.

Note If you configured client encryption for the Cassandra host, the `cqlsh` utility requires additional configuration. For more information about configuring Cassandra, see [Enable Server and Client Communication with Cassandra over SSL](#). For more information about using the `cqlsh` utility in a Cassandra host with client encryption, see the *Using cqlsh with SSL encryption* topic in the DataStax documentation.

8 Remove the vSphere Replication solution user from vCloud Director.

- a List all vSphere Replication solution user in vCloud Director

Standard Command	Command Using Registry
<pre>vcav vcd list-vr-users \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd</pre>	<pre>vcav vcd list-vr-users \ --vcd=vcd-01-name</pre>

- b Remove the vSphere Replication solution user from vCloud Director

Standard Command	Command Using Registry
<pre>vcav vcd remove-vr-user \ --vcd-address=\$VCD_ADDRESS \ --vcd-user=\$VCD_USER \ --vcd-password-file=~/.ssh/.vcd \ --user=user_id@vsphere.local</pre>	<pre>vcav vcd remove-vr-user \ --vcd=vcd-01-name \ --user=user_id@vsphere.local</pre>

To remove all vSphere Replication solution users from vCloud Director, replace the `--user=user_id@vsphere.local` argument with the `--all` argument.

9 Delete all vCloud Availability VMs by using the vSphere Web Client.

Tenant Installation and Configuration

4

Client configuration relies on the configuration of vSphere Replication within the tenant environment

1 [Prepare Your Environment to Install vSphere Replication](#)

Before you deploy the vSphere Replication appliance, you must prepare the environment.

2 [Deploy the vSphere Replication Virtual Appliance](#)

vSphere Replication is distributed as an OVF virtual appliance.

3 [Register the vSphere Replication Appliance with vCenter Single Sign-On](#)

You must register the vSphere Replication Management Server with vCenter Single Sign-On on both the source and the target sites.

4 [Update the vSphere Replication Appliances to Trust the vCloud Director Self-Signed Certificate in a Development Environment](#)

By using a certificate in the tenant configuration, you ensure security and encryption for tenant deployments. If the service provider vCloud Director instances use a self-signed certificate, you must update the on-premise vSphere Replication appliances to trust the self-signed certificate.

5 [Configure Cloud Provider](#)

You configure the Cloud Provider to assign the correct service provider destination for replication.

Prepare Your Environment to Install vSphere Replication

Before you deploy the vSphere Replication appliance, you must prepare the environment.

Procedure

- 1 Verify that you have vSphere and vSphere Web Client installations for the source and target sites.

- 2 In the vSphere Web Client, select the vCenter Server instance on which you are deploying vSphere Replication, click **Manage > Settings > Advanced Settings**, and verify that the `VirtualCenter.FQDN` value is set to a fully-qualified domain name or a literal address.

Note vSphere Replication can be deployed with either IPv4 or IPv6 address. Mixing IP addresses, for example having a single appliance with an IPv4 and an IPv6 address, is not supported. To register as an extension, vSphere Replication relies on the `VirtualCenter.FQDN` property of the vCenter Server. When an IPv6 address is used for vSphere Replication, the `VirtualCenter.FQDN` property must be set to a fully qualified domain name that can be resolved to an IPv6 address or to a literal address. When operating with an IPv6 address, vSphere Replication requires that all components in the environment, such as vCenter Server and ESXi hosts are accessible using the IPv6 address.

What to do next

You can deploy the vSphere Replication appliance.

Deploy the vSphere Replication Virtual Appliance

vSphere Replication is distributed as an OVF virtual appliance.

You deploy the vSphere Replication appliance by using the standard vSphere OVF deployment wizard.

Note vSphere Replication can be deployed with either IPv4 or IPv6 address. Mixing IP addresses, for example having a single appliance with an IPv4 and an IPv6 address, is not supported. To register as an extension, vSphere Replication relies on the `VirtualCenter.FQDN` property of the vCenter Server. When an IPv6 address is used for vSphere Replication, the `VirtualCenter.FQDN` property must be set to a fully qualified domain name that can be resolved to an IPv6 address or to a literal address. When operating with an IPv6 address, vSphere Replication requires that all components in the environment, such as vCenter Server and ESXi hosts are accessible using the IPv6 address.

Prerequisites

Download the vSphere Replication ISO image and mount it on a system in your environment.

Procedure

- 1 Log in to the vSphere Web Client on the source site.
- 2 Select **vCenter > Hosts and Clusters**.
- 3 Right-click a host and select **Deploy OVF template**.

- 4 Provide the location of the OVF file from which to deploy the vSphere Replication appliance, and click **Next**.

- Select **URL** and provide the URL to deploy the appliance from an online URL.
- If you downloaded and mounted the vSphere Replication ISO image on a system in your environment, select **Local file > Browse** and navigate to the `\bin` directory in the ISO image, and select the `vSphere_Replication_OVF10.ovf`, `vSphere_Replication-system.vmdk`, and `vSphere_Replication-support.vmdk` files.

- 5 Accept the name, select or search for a destination folder or datacenter for the virtual appliance, and click **Next**.

You can enter a new name for the virtual appliance. The name must be unique within each vCenter Server virtual machine folder.

- 6 Select a cluster, host, or resource pool where you want to run the deployed template, and click **Next**.
- 7 Review the virtual appliance details and click **Next**.
- 8 Accept the end user license agreements (EULA) and click **Next**.
- 9 Select the number of vCPUs for the virtual appliance and click **Next**.

Note Selecting higher number of vCPUs ensures better performance of the vSphere Replication Management Server, but might slow down the replications that run on ESXi host systems that have 4 or less cores per NUMA node. If you are unsure what the hosts in your environment are, select 2 vCPUs.

- 10 Select a destination datastore and disk format for the virtual appliance and click **Next**.
- 11 Select a network from the list of available networks, set the IP protocol and IP allocation, and click **Next**.

vSphere Replication supports both DHCP and static IP addresses. You can also change network settings by using the virtual appliance management interface (VAMI) after installation.

- 12 Set the password for the root account for the customized template, and click **Next**.

The password must be at least eight characters long.

- 13 Review the binding to the vCenter Extension vService and click **Next**.

- 14 Review the settings and click **Finish**.

The vSphere Replication appliance is deployed.

- 15 Power on the vSphere Replication appliance. Take a note of the IP address of the appliance and log out of the vSphere Web Client.

- 16 Repeat the procedure to deploy vSphere Replication on the target site.

What to do next

Register the vSphere Replication appliance with the SSO service.

Register the vSphere Replication Appliance with vCenter Single Sign-On

You must register the vSphere Replication Management Server with vCenter Single Sign-On on both the source and the target sites.

After you deploy the vSphere Replication appliance, you use the Virtual Appliance Management Interface (VAMI) to register the endpoint and the certificate of the vSphere Replication Management Server with the vCenter Lookup Service, and to register the vSphere Replication solution user with the vCenter Single Sign-On administration server.

If you do not register vSphere Replication with vCenter Single Sign-On on the target site, vSphere Replication cannot operate as expected. In addition, storage DRS does not detect the replicated data that vSphere Replication stores on the target site and might destroy it.

Prerequisites

- Verify that the vSphere Replication appliance is powered on.
- Verify that you have administrator privileges to configure the vSphere Replication appliance.
- Verify that the vSphere Replication management server is synchronized with the time of the Single Sign-On server.

Procedure

- 1 Use a supported browser to log in to the vSphere Replication VAMI.
The URL for the VAMI is `https://vr-appliance-address:5480`.
- 2 Type the root user name and password for the appliance.
You configured the root password during the OVF deployment of the vSphere Replication appliance.
- 3 Click the **Configuration** tab.
- 4 In the **LookupService Address** text box, enter the IP address or domain name of the server where the lookup service runs.
- 5 Enter the credentials of a user with administrator privileges to vCenter Single Sign-On.
- 6 Click **Save and Restart Service**.
- 7 Repeat the procedure to register vSphere Replication on the target site.

vSphere Replication appears on the **Home** tab in the vSphere Web Client.

What to do next

Note If you registered the vSphere Replication appliance with SSO as part of the upgrade procedure, all existing connections turn into `Connection` issue status. See [Reconnect to a Remote Site](#).

If you completed this procedure as part of the installation process, you can configure connections between the source and target sites.

Perform optional reconfiguration of the vSphere Replication appliance by using the VAMI. You can install a certificate, change the appliance root password, change the trust policy, or configure vSphere Replication to use an external database.



States of vSphere Replication Displayed in the vSphere Web Client

Before you can start using vSphere Replication, you must register the vSphere Replication appliance with the vCenter Lookup Service and the Single Sign-On administration server in the environment.

In the vSphere Web Client, on the vSphere Replication **Home** tab, you can check the list of vCenter Server instances in the Single-Sign On domain, and the status of vSphere Replication on each vCenter Server instance.

The following table lists the vSphere Replication states that you can observe, their meanings, and what you can do to change a state back to normal.

Table 4-1. vSphere Replication States on vCenter Server Instances

Status	Description	Remediation
Not installed	<p>The vSphere Replication extension is not registered in the vCenter Server Extension Manager.</p> <p>The vSphere Replication appliance is either not deployed or the vSphere Replication extension has been deleted from the vCenter Server Extension Manager.</p>	<p>If a vSphere Replication appliance is deployed on this vCenter Server, restart the appliance or the vSphere Replication Management service on the appliance.</p> <ol style="list-style-type: none"> 1 Use a supported browser to log in to the vSphere Replication VAMI as the root user. <p>The URL for the VAMI is <code>https://vr-appliance-address:5480</code>.</p> <ol style="list-style-type: none"> 2 On the Configuration tab, click Save and Restart Service.
Enabled (Configuration issue)	<p>A configuration error occurred.</p> <p>The vSphere Replication Management Server is either not registered with the vCenter SSO components, or the configuration is incorrect and must be updated.</p> <p>You cannot manage existing replications, or configure new replications to this server .</p>	<p>Configure the vSphere Replication appliance.</p> <ol style="list-style-type: none"> 1 Select the row that indicates the Enabled (Configuration issue) status. 2 Point to the Enabled (Configuration issue) status. <p>The detailed error message appears in a tooltip.</p> <ol style="list-style-type: none"> 3 Click the Configure icon () above the list of vCenter Server instances. <p>The vSphere Replication VAMI opens.</p> <ol style="list-style-type: none"> 4 On the Configuration tab, enter the parameters that were indicated in the error message and click Save and Restart Service .
Enabled (Not accessible)	<p>The vSphere Replication Management Server is not accessible.</p> <p>The vSphere Replication extension is registered in the vCenter Server Extension Manager, but the vSphere Replication appliance is missing or powered off, or the vSphere Replication Management service is not running.</p> <p>You cannot manage existing replications, or configure new replications to this server .</p>	<ul style="list-style-type: none"> ■ Verify that the vSphere Replication appliance exists on the vCenter Server. ■ Verify that the vSphere Replication appliance is powered on. ■ Restart the VRM service. <ol style="list-style-type: none"> a On the vSphere Replication Home tab, select the row that indicates the Enabled (Not accessible) status and click the Configure icon () above the list of replication servers. b On the Configuration tab, restart the VRM service.
Enabled (OK)	The vSphere Replication appliance is installed, configured, and functioning properly.	Not needed.

Update the vSphere Replication Appliances to Trust the vCloud Director Self-Signed Certificate in a Development Environment

By using a certificate in the tenant configuration, you ensure security and encryption for tenant deployments. If the service provider vCloud Director instances use a self-signed certificate, you must update the on-premise vSphere Replication appliances to trust the self-signed certificate.

Note The following procedure contains long, single commands that should be run as one. There are breaks in the command for better visibility marked with "\". "#" marks the beginning of a new command.

Prerequisites

Make sure that SSH is enabled on your vSphere Replication Appliance. For more information, see <https://kb.vmware.com/s/article/2112307>.

Procedure

- 1 Export the vCloud Director self-signed certificate and import it into the vSphere Replication Appliance keystore.
 - a Log in to vSphere Replication Appliance.
 - b Back up the appliance keystore by running the following command.

```
# cp /usr/java/default/lib/security/cacerts /usr/java/default/lib/security/cacerts.bak
```

- c Export the vCloud Director self-signed certificate by running the following command.

```
# openssl s_client -connect vCD-IP:443 </dev/null 2>/dev/null \  
| openssl x509 > /tmp/vcloud.pem
```

- d Import the vCloud Director self-signed certificate into the vSphere Replication Appliance keystore by running the following command.

```
# /usr/java/default/bin/keytool -noprompt \  
-import -trustcacerts -alias cloudproxy -file /tmp/vcloud.pem \  
-keystore /usr/java/default/lib/security/cacerts -storepass changeit
```

Note The keytool can be located in a different folder depending on the vSphere Replication release.

- 2 Restart the services that use the keystore file by running the following commands.

```
# service hms restart  
  
# service vmware-vcd restart
```

(Optional) Update the vSphere Replication Appliances to Trust the Cloud Proxy Self-Signed Certificate

If you have Cloud Proxy instances in the service provider environment that use different self-signed certificates than vCloud Director, you must update vSphere Replication appliances to trust them.

Repeat the following procedure for each Cloud Proxy instance in the service provider environment.

Procedure

- 1 Copy the self-signed certificate to the client vSphere Replication Appliance and load it into the keystore.
 - a Log in to vSphere Replication Appliance.
 - b Export the Cloud Proxy certificate and import it into the Java keystore:

```
# openssl s_client -connect $CLOUD_PROXY_IP:443 </dev/null 2>/dev/null \
| openssl x509 > /tmp/vcloud.pem

# /usr/java/default/bin/keytool -noprompt \
-import -trustcacerts -alias cloudproxy -file /tmp/vcloud.pem \
-keystore /usr/java/default/lib/security/cacerts -storepass changeit
```

Note Keytools can be on a different folder depending on the vSphere Replication release.

- 2 Restart the services that use the keystore file by running the following commands.

```
# service hms restart

# service vmware-vcd restart
```

Configure Cloud Provider

You configure the Cloud Provider to assign the correct service provider destination for replication.

Prerequisites

Open the vCenter Server administration interface.

Procedure

- 1 Open the vCenter Server by using the Web Client.
- 2 Navigate to the **Connect to a Cloud Provider** tab.
- 3 In the **Manage** tab, click **vSphere Replication**.
- 4 Click **Target Sites** and click the **Connect to Cloud Provider** icon.

5 Connect to Cloud Provider.

- a Enter the Cloud Provider Address: `vcd.provider.com` without the `/cloud` suffix
- b Enter the Organization Name.
- c Enter user name and password. The user profile must be assigned with Replication Rights.

6 After you configure the Cloud Provider, right-click and select **Configure Target Networks**.