

Installing

12 JAN 2023

vRealize Operations 8.6

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

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

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

Copyright © 2023 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About vApp Deployment and Configuration 5

1 About Installing 6

- Workflow of vRealize Operations Installation 6
- Sizing the Cluster 8
 - Add Data Disk Space to a vApp Node 9
- Complexity of Your Environment 9
- Cluster Nodes 11
- About Remote Collector Nodes 12
- About High Availability 13
- About vRealize Operations Continuous Availability 15

2 Preparing for Installation 17

- Requirements 17
 - Requirements for IPv6 17
 - Cluster Requirements 18
 - Sizing and Scaling Requirements 21

3 Installing vRealize Operations 22

- Deployment of vRealize Operations 22
 - Create a Node by Deploying an OVF 22
- Installation Types 25
 - Installing vRealize Operations for a New User 25
 - Installing vRealize Operations as an Administrator 28
 - Expand an Existing Installation of vRealize Operations 29
- Installing vRealize Operations on VMware Cloud on AWS 31
 - Using vRealize Operations on-premises on VMware Cloud on AWS 32
 - Deploying vRealize Operations on VMware Cloud on AWS 34
- Installing vRealize Operations for Azure VMware Solution 36
 - Using vRealize Operations on-premises for Azure VMware Solution 37
 - Deploying vRealize Operations on Azure VMware Solution 40
- Installing vRealize Operations for Oracle Cloud VMware Solution 40
 - Using vRealize Operations on-premises for Oracle Cloud VMware Solution 41
 - Deploying vRealize Operations on Oracle Cloud VMware Solution 44
- Installing vRealize Operations for Google Cloud VMware Engine 44
 - Using vRealize Operations on-premises for Google Cloud VMware Engine 44
 - Deploying vRealize Operations on Google Cloud VMware Engine 47
- Installing vRealize Operations for VMware Cloud on Dell EMC 47

Using vRealize Operations on-premises for VMware Cloud on Dell EMC	48
Deploying vRealize Operations on VMware Cloud on Dell EMC	49
4 Resize your Cluster by Adding Nodes	52
Gathering More Data by Adding a Remote Collector Node	53
Run the Setup Wizard to Create a Remote Collector Node	53
Adding High Availability	54
Run the Setup Wizard to Add a Primary Replica Node	54
Adding Continuous Availability	56
Enable Continuous Availability in vRealize Operations	56
Cluster and Node Maintenance	57
Cluster Management	60
Troubleshooting	62
Troubleshooting Cluster Problems	62
5 Installing Cloud Proxy	64
Configuring Cloud Proxies in vRealize Operations	64
Managing Cloud Proxies in vRealize Operations	67
Adding Cloud Proxies To a Collector Group	68
Monitoring the Health of Cloud Proxies	68
Cloud Proxy FAQ	71
Cloud Proxy Troubleshooting	74
6 Post-Installation Considerations	78
About Logging In	78
After You Log In	79
Secure the Console	81
Log in to a Remote Console Session	82
About New Installations	82
Log In and Continue with a New Installation	82
7 Upgrade, Backup and Restore	85
Obtain the Software Update PAK File	85
Create a Snapshot as Part of an Update	86
How To Preserve Customized Content	87
Back Up and Restore	88
Software Updates	88
Install a Software Update	90
Before Upgrading to vRealize Operations 8.6	92
Running the vRealize Operations 8.6 Pre-Upgrade Readiness Assessment Tool	92

About vApp Deployment and Configuration

The *vRealize Operations Manager vApp Deployment and Configuration Guide* provides information about deploying the VMware[®] vRealize Operations Manager virtual appliance, including how to create and configure the vRealize Operations Manager cluster.

The vRealize Operations Manager installation process consists of deploying the vRealize Operations Manager virtual appliance once for each cluster node, and accessing the product to finish setting up the application.

Intended Audience

This information is intended for anyone who wants to install and configure vRealize Operations Manager by using a virtual appliance deployment. The information is written for experienced virtual machine administrators who are familiar with enterprise management applications and data center operations.

For administrators who want to deploy the vRealize Operations Manager virtual appliance programmatically, the VMware vRealize Operations Manager CaSA API documentation is available in HTML format and is installed with your vRealize Operations Manager instance. For example, if the URL of your instance is `https://vrealize.example.com`, the API reference is available from `https://vrealize.example.com/casa/api-guide.html`.

About Installing

1

You prepare for vRealize Operations installation by evaluating your environment and deploying enough vRealize Operations cluster nodes to support how you want to use the product.

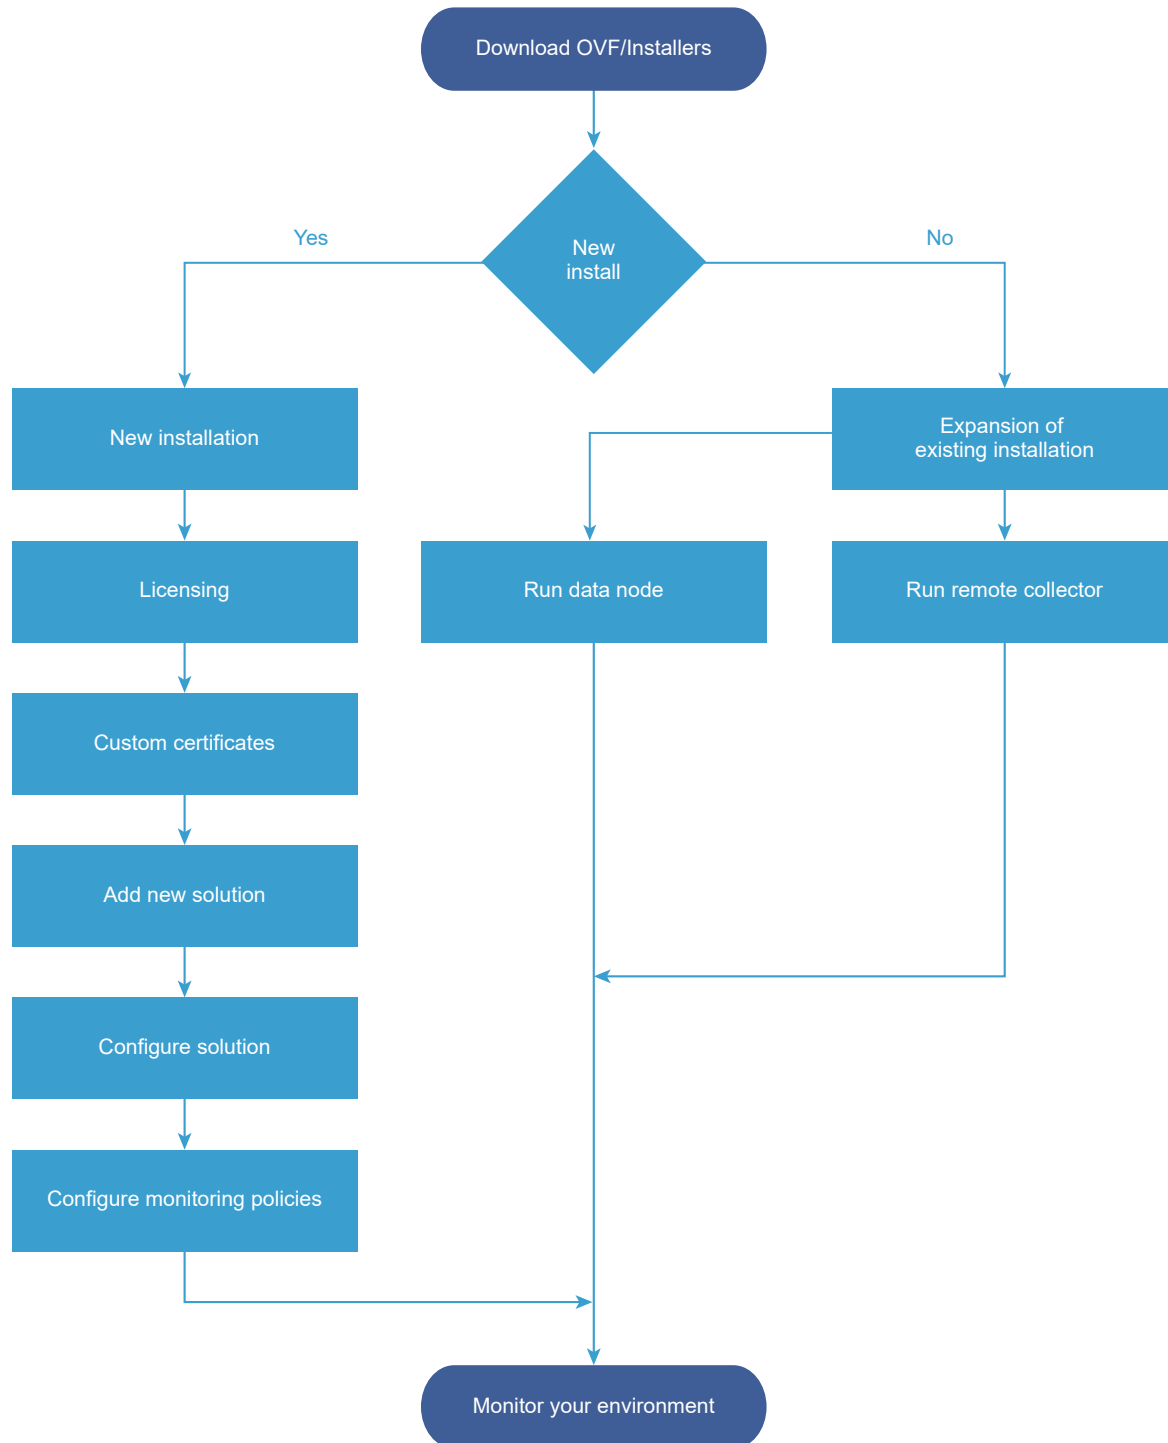
This chapter includes the following topics:

- [Workflow of vRealize Operations Installation](#)
- [Sizing the vRealize Operations Cluster](#)
- [Complexity of Your Environment](#)
- [About vRealize Operations Cluster Nodes](#)
- [About vRealize Operations Remote Collector Nodes](#)
- [About vRealize Operations High Availability](#)
- [About vRealize Operations Continuous Availability](#)

Workflow of vRealize Operations Installation

The vRealize Operations virtual appliance installation process consists of deploying the vRealize Operations OVA, once for each cluster node, accessing the product to set up cluster nodes according to their role, and logging in to configure the installation.

Figure 1-1. vRealize Operations Manager Installation Architecture



To automate installation, configuration, upgrade, patch, configuration management, drift remediation and health from within a single pane of glass, you can use vRealize Suite Lifecycle Manager. If you are a new user, click [here](#) to install **vRealize Suite Lifecycle Manager**. This provides the IT Managers of Cloud admin resources to focus on business-critical initiatives, while improving time to value (TTV), reliability, and consistency.

You can also install upgrade vRealize Operations by using vRealize Suite Lifecycle Manager. For more information, see the [Creating an Environment from Configure vRealize Products](#).

Sizing the vRealize Operations Cluster

The resources needed for vRealize Operations depend on how large of an environment you expect to monitor and analyze, how many metrics you plan to collect, and how long you need to store the data.

It is difficult to broadly predict the CPU, memory, and disk requirements that will meet the needs of a particular environment. There are many variables, such as the number and type of objects collected, which includes the number and type of adapters installed, the presence of HA, the duration of data retention, and the quantity of specific data points of interest, such as symptoms, changes, and so on.

VMware expects vRealize Operations sizing information to evolve, and maintains Knowledge Base articles so that sizing calculations can be adjusted to adapt to usage data and changes in versions of vRealize Operations.

[Knowledge Base article 2093783](#)

The Knowledge Base articles include overall maximums, plus spreadsheet calculators in which you enter the number of objects and metrics that you expect to monitor. To obtain the numbers, some users take the following high-level approach, which uses vRealize Operations itself.

- 1 Review this guide to understand how to deploy and configure a vRealize Operations node.
- 2 Deploy a temporary vRealize Operations node.
- 3 Configure one or more adapters, and allow the temporary node to collect overnight.
- 4 Access the Cluster Management page on the temporary node.
- 5 Using the Adapter Instances list in the lower portion of the display as a reference, enter object and metric totals of the different adapter types into the appropriate sizing spreadsheet from [Knowledge Base article 2093783](#).
- 6 Deploy the vRealize Operations cluster based on the spreadsheet sizing recommendation. You can build the cluster by adding resources and data nodes to the temporary node or by starting over.

If you have a large number of adapters, you might need to reset and repeat the process on the temporary node until you have all the totals you need. The temporary node will not have enough capacity to simultaneously run every connection from a large enterprise.

Another approach to sizing is through self monitoring. Deploy the cluster based on your best estimate, but create an alert for when capacity falls below a threshold, one that allows enough time to add nodes or disk to the cluster. You also have the option to create an email notification when thresholds are passed.

During internal testing, a single-node vApp deployment of vRealize Operations that monitored 8,000 virtual machines ran out of disk storage within one week.

Add Data Disk Space to a vRealize Operations vApp Node

You add to the data disk of vRealize Operations vApp nodes when space for storing the collected data runs low.

Prerequisites

- Note the disk size of the analytics cluster nodes. When adding disk, you must maintain uniform size across analytics cluster nodes.
- Use the vRealize Operations administration interface to take the node offline.
- Verify that you are connected to a vCenter Server system with a vSphere Client, and log in to the vSphere Client.

Procedure

- 1 Shut down the virtual machine for the node.
- 2 Edit the hardware settings of the virtual machine, and add another disk.

Note Do not expand disks. vRealize Operations does not support expanding disks.

- 3 Power on the virtual machine for the node.

Results

During the power-on process, the virtual machine expands the vRealize Operations data partition.

Complexity of Your Environment

When you deploy vRealize Operations, the number and nature of the objects that you want to monitor might be complex enough to recommend a Professional Services engagement.

Complexity Levels

Every enterprise is different in terms of the systems that are present and the level of experience of deployment personnel. The following table presents a color-coded guide to help you determine where you are on the complexity scale.

- Green
Your installation only includes conditions that most users can understand and work with, without assistance. Continue your deployment.
- Yellow
Your installation includes conditions that might justify help with your deployment, depending on your level of experience. Consult your account representative before proceeding, and discuss using Professional Services.
- Red

Your installation includes conditions that strongly recommend a Professional Services engagement. Consult your account representative before proceeding, and discuss using Professional Services.

Note that these color-coded levels are not firm rules. Your product experience, which increases as you work with vRealize Operations and in partnership with Professional Services, must be taken into account when deploying vRealize Operations.

Table 1-1. Effect of Deployment Conditions on Complexity

Complexity Level	Current or New Deployment Condition	Additional Notes
Green	You run only one vRealize Operations deployment.	Lone instances are usually easy to create in vRealize Operations.
Green	Your deployment includes a management pack that is listed as Green according to the compatibility guide on the VMware Solutions Exchange Web site.	<p>The compatibility guide indicates whether the supported management pack for vRealize Operations is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected.</p> <p>Note that the terms <i>solution</i>, <i>management pack</i>, <i>adapter</i>, and <i>plug-in</i> are used somewhat interchangeably.</p>
Yellow	You run multiple instances of vRealize Operations.	Multiple instances are typically used to address scaling or operator use patterns.
Yellow	Your deployment includes a management pack that is listed as Yellow according to the compatibility guide on the VMware Solutions Exchange Web site.	The compatibility guide indicates whether the supported management pack for vRealize Operations is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected.
Yellow	You are deploying vRealize Operations remote collector nodes.	Remote collector nodes gather data but leave the storage and processing of the data to the analytics cluster.
Yellow	You are deploying a multiple-node vRealize Operations cluster.	Multiple nodes are typically used for scaling out the monitoring capability of vRealize Operations.
Yellow	Your new vRealize Operations instance will include a Linux based deployment.	Linux deployments are not as common as vApp deployments and often need special consideration.
Yellow	Your vRealize Operations instance will use high availability (HA).	High availability and its node failover capability is a unique multiple-node feature that you might want additional help in understanding.

Table 1-1. Effect of Deployment Conditions on Complexity (continued)

Complexity Level	Current or New Deployment Condition	Additional Notes
Yellow	You want help in understanding the new or changed features in vRealize Operations and how to use them in your environment.	vRealize Operations is different than vCenter Operations Manager in areas such as policies, alerts, compliance, custom reporting, or badges. In addition, vRealize Operations uses one consolidated interface.
Red	You run multiple instances of vRealize Operations, where at least one includes virtual desktop infrastructure (VDI).	Multiple instances are typically used to address scaling, operator use patterns, or because separate VDI (V4V monitoring) and non-VDI instances are needed.
Red	Your deployment includes a management pack that is listed as Red according to the compatibility guide on the VMware Solutions Exchange Web site.	The compatibility guide indicates whether the supported management pack for vRealize Operations is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected.
Red	You are deploying multiple vRealize Operations clusters.	Multiple clusters are typically used to isolate business operations or functions.
Red	Your current vRealize Operations deployment required a Professional Services engagement to install it.	If your environment was complex enough to justify a Professional Services engagement in the previous version, it is possible that the same conditions still apply and might warrant a similar engagement for this version.
Red	Professional Services customized your vRealize Operations deployment. Examples of customization include special integrations, scripting, nonstandard configurations, multiple level alerting, or custom reporting.	If your environment was complex enough to justify a Professional Services engagement in the previous version, it is possible that the same conditions still apply and might warrant a similar engagement for this version.

About vRealize Operations Cluster Nodes

All vRealize Operations clusters consist of a master node (primary node), an optional replica node for high availability or continuously availability, optional data nodes, and optional remote collector nodes.

When you install vRealize Operations, you use a vRealize Operations vApp deployment to create role-less nodes. After the nodes are created and have their names and IP addresses, you use an administration interface to configure them according to their role.

You can create role-less nodes all at once or as needed. A common as-needed practice might be to add nodes to scale out vRealize Operations to monitor an environment as the environment extends larger.

The following node types make up the vRealize Operations analytics cluster:

Master Node

The master node is the primary node and the initial, required node in vRealize Operations. All other nodes are managed by the primary node.

In a single-node installation, the primary node manages itself, has adapters installed on it, and performs all data collection and analysis.

Data Node

In larger deployments, additional data nodes have adapters installed and perform collection and analysis.

Larger deployments usually include adapters only on the data nodes so that primary and replica node resources can be dedicated to cluster management.

Replica Node

To use vRealize Operations high availability (HA) and continuous availability (CA) the cluster requires that you convert a data node into a replica of the primary node.

The following node types are a member of the vRealize Operations cluster but not part of the analytics cluster:

Remote Collector Node

Distributed deployments might require a remote collector node that can navigate firewalls, interface with a remote data source, reduce the bandwidth across data centers, or reduce the load on the vRealize Operations analytics cluster. Remote collectors only gather objects for the inventory, without storing data or performing analysis. In addition, remote collector nodes might be installed on a different operating system than the rest of the cluster.

Witness Node

To use vRealize Operations continuous availability (CA), the cluster requires that you have a witness node. Each vRealize Operations cluster can have only one witness node. If the network connection between the two fault domains is lost, the witness node acts as a decision maker regarding the availability of vRealize Operations.

About vRealize Operations Remote Collector Nodes

A remote collector node is an additional cluster node that allows vRealize Operations to gather more objects into its inventory for monitoring purposes. Unlike the data nodes, the remote collector nodes only perform the collector role of vRealize Operations. These remote collectors do not store data or process any analytics functions. Remote collectors collect data from integrated objects and then forward the data back to the primary node. The primary node then processes the data which you then view as reports and analytics.

Remote collectors are very useful when you have multiple locations. You can deploy remote collectors on remote location sites and only deploy the primary node at the primary location.

You must have at least one primary node before adding remote collector nodes.

A remote collector node is usually deployed to navigate firewalls, reduce bandwidth across data centers, connect to remote data sources, or reduce the load on the vRealize Operations analytics cluster. To deploy a remote collector node, see [Run the Setup Wizard to Create a Remote Collector Node](#).

Remote collectors do not buffer data while the network is experiencing a problem. If the connection between the remote collector and the analytics cluster is lost, the remote collector does not store data points that occur during that time. In turn, and after the connection is restored, vRealize Operations does not retroactively incorporate associated events from that time into any monitoring or analysis.

Ports information for vRealize Operations is available on [Ports and Protocol](#).

About vRealize Operations High Availability

vRealize Operations supports high availability (HA). HA creates a replica for the vRealize Operations primary node and protects the analytics cluster against the loss of a node.

With HA, data stored in the primary node is always 100% backed up on the replica node. To enable HA, you must have at least one data node deployed, in addition to the primary node. If you have more than one data node, the data stored in the primary node can be stored and replicated in any of the other nodes. But in case the primary node fails, only the replica node can function as the replacement of the primary node.

- HA is not a disaster recovery mechanism. HA protects the analytics cluster against the loss of only one node, and because only one loss is supported, you cannot stretch nodes across vSphere clusters in an attempt to isolate nodes or build failure zones.
- When HA is enabled, the replica can take over all functions that the primary provides, were the primary to fail for any reason. If the primary fails, failover to the replica is automatic and requires only two to three minutes of vRealize Operations downtime to resume operations and restart data collection.

When a primary node problem causes failover, the replica node becomes the primary node, and the cluster runs in degraded mode. To get out of degraded mode, take one of the following steps.

- Return to HA mode by correcting the problem with the primary node. When a primary node exits an HA-enabled cluster, primary node does not rejoin with the cluster without manual intervention. Therefore, restart the vRealize Operations Analytics process on the downed node to change its role to replica and rejoin the cluster.
- Remove the failed primary node then re-enable HA by converting a data node into replica. Removed primary nodes cannot be repaired and readadded to vRealize Operations.

- Remove the old, failed primary node and then change to non-HA operation by disabling HA. Removed primary nodes cannot be repaired and readded to vRealize Operations.
- In the administration interface, after an HA replica node takes over and becomes the new primary node, you cannot remove the previous, offline primary node from the cluster. In addition, the previous node remains listed as a primary node. To refresh the display and enable removal of the node, refresh the browser.
- When HA is enabled, the cluster can survive the loss of one data node without losing any data. However, HA protects against the loss of only one node at a time, of any kind, so simultaneously losing data and primary/replica nodes, or two or more data nodes, is not supported. Instead, vRealize Operations HA provides additional application level data protection to ensure application level availability.
- When HA is enabled, it lowers vRealize Operations capacity and processing by half, because HA creates a redundant copy of data throughout the cluster, and the replica backup of the primary node. Consider your potential use of HA when planning the number and size of your vRealize Operations cluster nodes. See [Sizing the vRealize Operations Cluster](#).
- When HA is enabled, deploy analytics cluster nodes on separate hosts for redundancy and isolation. One option is to use anti-affinity rules that keep nodes on specific hosts in the vSphere cluster.

If you cannot keep the nodes separate, you should not enable HA. A host fault might cause the loss of more than one node, which is not supported, and all of vRealize Operations can become unavailable.

The opposite is also true. Without HA, you can keep nodes on the same host, and it will not make a difference. Without HA, the loss of even one node can make all of vRealize Operations unavailable.

- When you power off the data node and change the network settings of the VM, this affects the IP address of the data node. After this point, the HA cluster is no longer accessible and all the nodes have a status of "Waiting for analytics". Verify that you have used a static IP address.
- When you remove a node that has one or more vCenter adapters configured to collect data from a HA-enabled cluster, one or more vCenter adapters associated with that node stops collecting. You change the adapter configuration to pin them to another node before removing the node.
- Administration UI shows the resource cache count, which is created for active objects only, but the Inventory displays all objects. Therefore, when you remove a node from a HA-enabled cluster allowing the vCenter adapters collect data and rebalance each node, the Inventory displays a different quantity of objects from that shown in the Administration UI.

About vRealize Operations Continuous Availability

vRealize Operations supports continuous availability (CA). CA separates the vRealize Operations cluster into two fault domains, stretching across vSphere clusters, and protects the analytics cluster against the loss of an entire fault domain.

You can configure the analytics cluster with Continuous Availability. This allows the cluster nodes to be stretch across two fault-domains. A fault domain consists of one or more analytics nodes grouped according to their physical location in the data center. With CA, the two fault domains permit vRealize Operations to tolerate failures of an entire physical location and failures from resources dedicated to a single fault domain.

To enable continuous availability within vRealize Operations, the witness node must be deployed in the cluster. The vRealize Operations cluster can have only one witness node. The witness node does not collect nor store data. In a situation where network connectivity the two fault-domains is lost, the cluster would go into a split-brain situation. This situation is detected by the Witness Node and one of the fault domains will go offline to avoid data inconsistency issues. You will see a **Bring Online** button on the admin UI of the nodes which are made offline by the witness node. Before using this option to bring the fault domain online, ensure that the network connectivity between the nodes across the two fault domains is restored and stable. Once confirmed you can bring the fault domain online.

With CA, the data stored in the primary node and data nodes grouped in fault domain 1 is always 100% synced to the replica node and data nodes paired in fault domain 2. To enable CA, you must have at least one data node deployed, in addition to the primary node. If you have more than one data node, there must be an even number of data nodes including the primary node. For example, the cluster must have 2, 4, 6, 8, 10, 12, 14 or 16 nodes based on the appropriate sizing requirements. The data stored in the primary node in fault domain 1 is stored and replicated in the replica node in fault domain 2. The data stored in the data nodes in fault domain 1 is stored and replicated in the paired data nodes in fault domain 2. But in case the primary node fails, only the replica node can function as the replacement of the primary node.

- CA protects the analytics cluster against the loss of half the analytics nodes specific to one fault domain. You can stretch nodes across vSphere clusters in an attempt to isolate nodes or build failure zones.
- When CA is enabled, the replica node can take over all functions that the primary node provides, in case of a primary node failure. The failover to the replica is automatic and requires only two to three minutes of vRealize Operations downtime to resume operations and restart data collection.

Note In case of a primary node failure, the replica node becomes the primary node, and the cluster runs in degraded mode. To fix this, perform any one of the following actions.

- Correct the primary node failure manually.
 - Return to CA mode by replacing the primary node. Replacement nodes do not repair the node failure, instead a new node assumes the primary node role.
-

- In the administration interface, after a CA replica node takes over and becomes the new primary node, you cannot remove the previous, offline primary node from the cluster. In addition, the previous node remains listed as a primary node. To refresh the display and enable the removal of the node, refresh the browser.
- When CA is enabled, the cluster can survive the loss of half the data nodes, all in one fault domain, without losing any data. CA protects against the loss of only one fault domain at a time. Simultaneously losing data and primary/replica nodes, or two or more data nodes in both fault domains, is not supported.
- A CA enabled cluster will be non-functional if you power off the primary node or the primary node replica while one of the fault domains is down.
- When CA is enabled, it lowers the vRealize Operations capacity and processing by half, because CA creates a redundant copy of data throughout the cluster, and the replica backup of the primary node. Consider your potential use of CA when planning the number and size of your vRealize Operations cluster nodes. See [Sizing the vRealize Operations Cluster](#).
- When CA is enabled, deploy analytics cluster nodes, in each fault domain, on separate hosts for redundancy and isolation. You can also use anti-affinity rules that keep nodes on specific hosts in the vSphere clusters.
- If you cannot keep the nodes separate in each fault domain, you can still enable CA. A host fault might cause the loss of the data nodes in the fault domain, and vRealize Operations can still be available in the other fault domain.
- If you cannot split the data nodes into different vSphere clusters, do not enable CA. A cluster failure can cause the loss of more than half of the data nodes, which is not supported, and all of vSphere might become unavailable.
- Without CA, you can keep nodes on the same host in the same vSphere. Without CA, the loss of even one node might make all of vRealize Operations unavailable.
- When you power off data nodes in both fault domains and change the network settings of the VMs, it affects the IP address of the data nodes. After this point, the CA cluster is no longer accessible and all the nodes status change to "Waiting for analytics". Verify that you have used a static IP address.
- When you remove a node that has one or more vCenter adapters configured to collect data from a CA-enabled cluster, one or more vCenter adapters associated with that node stops collecting. You must change the adapter configuration to pin them to another node before removing the node.
- The administration interface displays the resource cache count, which is created for active objects only, but the inventory displays all objects. When you remove a node from a CA-enabled cluster allowing the vCenter adapters to collect data and rebalance each node, the inventory displays a different quantity of objects from that shown in the administration interface.

Preparing for Installation

2

When you prepare for your installation, consider some of these best practices, cluster, sizing and scaling requirements.

This chapter includes the following topics:

- [Requirements](#)

Requirements

You have to consider important requirements while creating nodes in a vRealize Operations.

Using IPv6 with vRealize Operations

vRealize Operations supports both, Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6). All nodes in the cluster must follow the same protocol. For endpoint communications, you can use IPv4 or IPv6. If the environment only supports the IPv6 protocol, the **Prefer IPv6** flag must be enabled during the OVF deployment for each node. If you set the **Prefer IPv6** flag, then vRealize Operations uses IPv6 for all communications.

Considerations While Using IPv6

- If any nodes use DHCP, your DHCP server must be configured to support IPv6.
- IPv6 DHCP or Static configuration must have Global Scope.
- DHCP is only supported on data nodes and remote collectors. Primary nodes and replica nodes require static addresses.
- Your DNS server must be configured to support IPv6.
- When adding nodes to the cluster, enter the IPv6 address of the primary node.
- When registering a VMware vCenter® instance within vRealize Operations, place square brackets around the IPv6 address of your VMware vCenter Server® system if vCenter is also using IPv6.

For example: [2015:0db8:85a3:0042:1000:8a2e:0360:7334]

Note When vRealize Operations is using IPv6, vCenter Server might still have an IPv4 address. In that case, vRealize Operations does not need the square brackets.

Cluster Requirements

When you create the cluster nodes that make up vRealize Operations, you have general requirements that you must meet.

General vRealize Operations Cluster Node Requirements

You have to follow some general requirements to create a node on your environment.

General Requirements

- vRealize Operations version. All nodes must run the same vRealize Operations version.
For example, do not add a version 6.1 data node to a cluster of vRealize Operations 6.2 nodes.
- Analytics Cluster Deployment Type. In the analytics cluster, all nodes must be the same kind of deployment: vApp.
- Remote Collector Deployment Type. A remote collector node does not need to be the same deployment type as the analytics cluster nodes.

When you add a remote collector of a different deployment type, the following clusters are supported:

- vApp analytics cluster
- Witness Node Deployment Type. The witness node must be the same vApp deployment.
- Analytics Cluster Node Sizing. In the analytics cluster, CPU, memory, and disk size must be identical for all nodes.
Primary, replica, and data nodes must be uniform in sizing.
- Remote Collector Node Sizing. Remote collector nodes may be of different sizes from each other or from the uniform analytics cluster node size.
- Witness Node Sizing. The witness node has only one size and may be of different sizes from remote collectors or from the uniform analytics cluster node size
- Geographical Proximity. You may place analytics cluster nodes in different vSphere clusters, but the nodes must reside in the same geographical location.

Different geographical locations are not supported.

- Witness Node Placement. You may place the witness node in a different vSphere cluster separate from the analytics nodes.

Note A vRealize Operations cluster can have only one witness node.

- Virtual Machine Maintenance. When any node is a virtual machine, you may only update the virtual machine software by directly updating the vRealize Operations software.

For example, going outside of vRealize Operations to access vSphere to update VMware Tools is not supported.

- Redundancy and Isolation. If you expect to enable HA, place analytics cluster nodes on separate hosts. See [About vRealize Operations High Availability](#) .
- If you expect to enable CA, place analytics cluster nodes on separate hosts in fault domains, stretched across vSphere clusters. See [About vRealize Operations Continuous Availability](#).
- You can deploy remote collectors behind a firewall. You cannot use NAT between remote collectors and analytics nodes.

Requirements for Solutions

Be aware that solutions might have requirements beyond those for vRealize Operations itself. For example, vRealize Operations for Horizon View has specific sizing guidelines for its remote collectors.

See your solution documentation, and verify any additional requirements before installing solutions. Note that the terms *solution*, *management pack*, *adapter*, and *plug-in* are used interchangeably.

vRealize Operations Cluster Node Networking Requirements

When you create the cluster nodes that make up vRealize Operations, the associated setup within your network environment is critical to the inter-node communication and proper operation.

Networking Requirements

Important vRealize Operations analytics cluster nodes need frequent communication with one another. In general, your underlying vSphere architecture might create conditions where some vSphere actions affect that communication. Examples include, but are not limited to, vMotions, storage vMotions, HA events, and DRS events.

- The primary and replica nodes must use a static IP address, or fully qualified domain name (FQDN) with a static IP address.
Data and remote collector nodes can use dynamic host control protocol (DHCP).
- You can successfully reverse-DNS all nodes, including remote collectors, to their FQDN, currently the node hostname.
Nodes deployed by OVF have their hostnames set to the retrieved FQDN by default.
- All nodes, including remote collectors, must be bidirectionally routable by IP address or FQDN.
- Do not separate analytics cluster nodes with network address translation (NAT), load balancer, firewall, or a proxy that inhibits bidirectional communication by IP address or FQDN.
- Analytics cluster nodes must not have the same hostname.
- Place analytics cluster nodes within the same data center and connect them to the same local area network (LAN).
- Place analytics cluster nodes on same Layer 2 network and IP subnet.

A stretched Layer 2 or routed Layer 3 network is not supported.

- Do not span the Layer 2 network across sites, which might create network partitions or network performance issues.
- With Continuous Availability enabled, separate analytics cluster nodes into fault domains, stretched across vSphere clusters
- Packet Round Trip Time between the analytics cluster nodes must be 5 ms or lower.
- Network bandwidth between the analytics cluster nodes must be one gbps or higher.
- Do not distribute analytics cluster nodes over a wide area network (WAN).

To collect data from a WAN, a remote or separate data center, or a different geographic location, use remote collectors.

- Remote collectors are supported through a routed network but not through NAT.
- Do not include an underscore in the hostname of any cluster node.
- Cloud proxies must have a proper DNS resolution to the vRealize Operations nodes when using short/long FQDN names. This is applicable to on-prem cloud proxy.

vRealize Operations Cluster Node Best Practices

When you create the cluster nodes that make up vRealize Operations, additional best practices improve performance and reliability in vRealize Operations.

Best Practices

- Deploy vRealize Operations analytics cluster nodes in the same vSphere cluster in a single data center and add only one node at a time to a cluster allowing it to complete before adding another node.
- If you deploy analytics cluster nodes in a highly consolidated vSphere cluster, you might need resource reservations for optimal performance.

Determine whether the virtual to physical CPU ratio is affecting performance by reviewing CPU ready time and co-stop.

- Deploy analytics cluster nodes on the same type of storage tier.
- To continue to meet analytics cluster node size and performance requirements, apply storage DRS anti-affinity rules so that nodes are on separate datastores.
- To prevent unintentional migration of nodes, set storage DRS to manual.
- To ensure balanced performance from analytics cluster nodes, use ESXi hosts with the same processor frequencies. Mixed frequencies and physical core counts might affect analytics cluster performance.

- To avoid a performance decrease, vRealize Operations analytics cluster nodes need guaranteed resources when running at scale. The vRealize Operations Knowledge Base includes sizing spreadsheets that calculate resources based on the number of objects and metrics that you expect to monitor, use of HA, and so on. When sizing, it is better to over-allocate than under-allocate resources.

See [Knowledge Base article 2093783](#).

- Because nodes might change roles, avoid machine names such as Primary, Data, Replica, and so on. Examples of changed roles might include making a data node into a replica for HA, or having a replica take over the primary node role.
- The NUMA placement is removed in the vRealize Operations 6.3 and later. Procedures related to NUMA settings from the OVA file follow:

Table 2-1. NUMA Setting

Action	Description
Set the vRealize Operations cluster status to offline	<ol style="list-style-type: none"> 1 Shut down the vRealize Operations cluster. 2 Right-click the cluster and click Edit Settings > Options > Advanced General. 3 Click Configuration Parameters. In the vSphere Client, repeat these steps for each VM.
Remove the NUMA setting	<ol style="list-style-type: none"> 1 From the Configuration Parameters, remove the setting <code>numa.vcpu.preferHT</code> and click OK. 2 Click OK. 3 Repeat these steps for all the VMs in the vRealize Operations cluster. 4 Power on the cluster.

Note To ensure the availability of adequate resources and continued product performance, monitor vRealize Operations performance by checking its CPU usage, CPU ready and CPU contention time.

Sizing and Scaling Requirements

The CPU, memory, and disk requirements that meet the needs of a particular environment depend on the number and type of objects in your environment and the data collected. This includes the number and type of adapters installed, the use of HA (High Availability) or CA (Continuous Availability), the duration of data retention, and the quantity of specific data points of interest.

VMware updates [Knowledge Base article 2093783](#) with the most current information about sizing and scaling. The Knowledge Base article includes overall maximums and spreadsheet calculations that provide a recommendation based on the number of objects and metrics you expect to monitor.

Installing vRealize Operations

3

vRealize Operations nodes are virtual appliance (vApp) based systems.

This chapter includes the following topics:

- [Deployment of vRealize Operations](#)
- [Installation Types](#)
- [Installing vRealize Operations on VMware Cloud on AWS](#)
- [Installing vRealize Operations for Azure VMware Solution](#)
- [Installing vRealize Operations for Oracle Cloud VMware Solution](#)
- [Installing vRealize Operations for Google Cloud VMware Engine](#)
- [Installing vRealize Operations for VMware Cloud on Dell EMC](#)

Deployment of vRealize Operations

vRealize Operations consists of one or more nodes in a cluster. To create these nodes, you have to download and install the vRealize Operations suitable to your environment.

Create a Node by Deploying an OVF

vRealize Operations consists of one or more nodes, in a cluster. To create nodes, you use the vSphere client to download and deploy the vRealize Operations virtual machine, once for each cluster node.

Prerequisites

- Verify that you have permissions to deploy OVF templates to the inventory.
- If the ESXi host is part of a cluster, enable DRS in the cluster. If an ESXi host belongs to a non-DRS cluster, all resource pool functions are disabled.
- If this node is to be the primary node, reserve a static IP address for the virtual machine, and know the associated domain name, domain search path, domain name servers, default gateway, and network mask values.

Plan to keep the IP address because it is difficult to change the address after installation.

- If this node is to be a data node that will become the HA/CA replica node, reserve a static IP address for the virtual machine, and store the associated domain name, domain search path, domain name servers, default gateway, and network mask values for later use.

In addition, familiarize yourself with HA node placement as described in [About vRealize Operations High Availability](#) and CA node allocation as described in [About vRealize Operations Continuous Availability](#).

- Plan your domain and machine naming so that the deployed virtual machine name begins and ends with an alphabet (a–z) or digit (0–9) characters, and will only contain alphabet, digit, or hyphen (-) characters. The underscore character (_) must not appear in the host name or anywhere in the fully qualified domain name (FQDN).

Plan to keep the name because it is difficult to change the name after installation.

For more information, review the host name specifications from the Internet Engineering Task Force. See www.ietf.org.

- Plan node placement and networking to meet the requirements described in [General vRealize Operations Cluster Node Requirements](#) and [vRealize Operations Cluster Node Networking Requirements](#).
- If you expect the vRealize Operations cluster to use IPv6 addresses, review the IPv6 limitations described in [Using IPv6 with vRealize Operations](#).
- Download the vRealize Operations .ova file to a location that is accessible to the vSphere client.
- If you download the virtual machine and the file extension is .tar, change the file extension to .ova.
- Verify that you are connected to a vCenter Server system with a vSphere client, and log in to the vSphere client.

Do not deploy vRealize Operations from an ESXi host. Deploy only from vCenter Server.

Procedure

- 1 Select the vSphere **Deploy OVF Template** option.
- 2 Enter the path to the vRealize Operations .ova file.
- 3 Follow the prompts until you are asked to enter a name for the node.
- 4 Enter a node name. Examples might include **Ops1**, **Ops2**, **Ops-A**, **Ops-B**.

Do not include nonstandard characters such as underscores (_) in node names.

Use a different name for each vRealize Operations node.

- 5 Follow the prompts until you are asked to select a configuration size.

- 6 Select the size configuration that you need. Your selection does not affect the disk size.

Default disk space is allocated regardless of which size you select. If you need additional space to accommodate the expected data, add more disk after deploying the vApp, see [Add Data Disk Space to a vRealize Operations vApp Node](#).

- 7 Follow the prompts until you are asked to select the disk format.

Option	Description
Thick Provision Lazy Zeroed	Creates a virtual disk in a default thick format.
Thick Provision Eager Zeroed	Creates a type of thick virtual disk that supports clustering features such as Fault Tolerance. Thick provisioned eager-zeroed format can improve performance depending on the underlying storage subsystem. Select the thick provisioned eager-zero option when possible.
Thin Provision	Creates a disk in thin format. Use this format to save storage space.

Snapshots can negatively affect the performance of a virtual machine and typically result in a 25–30 percent degradation for the vRealize Operations workload. Do not use snapshots.

- 8 Click **Next**.
- 9 From the drop-down menu, select a Destination Network, for example, **Network 1 = TEST**, and click **Next**.
- 10 Under Networking Properties, in case of a static IP, specify the associated **Default Gateway**, **Domain Name**, **Domain Search Path**, **Domain Name Servers**, **Network 1 IP Address**, and **Network 1 Netmask** values. In case of DHCP, leave all the fields blank. The primary node and replica node require a static IP. A data node or remote collector node can use DHCP or a static IP.

Note The hostname is configured using DHCP and DNS. If a static IP is used the hostname is configured according to the node name specified during node configuration, after deployment.

- 11 In the Timezone Setting, leave the default of UTC or select a time zone.

The preferred approach is to standardize on UTC. Alternatively, configure all nodes to the same time zone.

Note You cannot configure nodes to different time zones.

- 12 (Optional) In Properties, under Application, select the option for IPv6 .
- 13 (Optional) If you want to deploy a FIPS enabled vRealize Operations setup, in the FIPS setting, select the **Enable FIPS Mode** check box.
- 14 Click **Next**.
- 15 Review the settings and click **Finish**.
- 16 If you are creating a multiple-node vRealize Operations cluster, repeat through all the steps to deploy each node.

What to do next

Use a Web browser client to configure a newly added node as the vRealize Operations primary node, a data node, a high availability primary replica node, or a remote collector node. The primary node is required first.


Caution For security, do not access vRealize Operations from untrusted or unpatched clients, or from clients using browser extensions.

Installation Types

After you have installed vRealize Operations product, you can either perform a new installation, an express installation, or expand an existing installation.

- Express Installation
- New installation
- Expand Installation

Figure 3-1. Getting Started Setup



Installing vRealize Operations for a New User

After you install vRealize Operations using an OVF or an installer, you are notified to the main product UI page. You can create a single node or multiple nodes depending on your environment.

Introduction to a New Installation

You can perform a new installation as a first-time user and create a single node to handle both administration and data handling.

Figure 3-2. New Installation from the Setup page



Perform a New Installation on the vRealize Operations Product UI

You can create a single node and configure it as a primary node or create a data node in a cluster to handle additional data. All vRealize Operations installations require a primary node.

With a single node cluster, administration and data functions are on the same primary node. A multiple-node vRealize Operations cluster contains one primary node and one or more nodes for handling additional data.

Prerequisites

- Create a node by deploying the vRealize Operations vApp.
- After it is deployed, note the fully qualified domain name (FQDN) or IP address of the node.
- If you plan to use a custom authentication certificate, verify that your certificate file meets the requirements for vRealize Operations.

Procedure

- 1 Navigate to the name or IP address of the node that will be the primary node of vRealize Operations.

The setup wizard appears, and you do not need to log in to vRealize Operations.

- 2 Click **New Installation**.

- 3 Click **Next**.

- 4 Enter and confirm a password for the admin user account, and click **Next**.

Passwords require a minimum of eight characters, one uppercase letter, one lowercase letter, one digit, and one special character.

The user account name is admin by default and cannot be changed.

- 5 Select whether to use the certificate included with vRealize Operations or to install one of your own.
 - a To use your own certificate, click **Browse**, locate the certificate file, and click **Open** to load the file in the Certificate Information text box.
 - b Review the information detected from your certificate to verify that it meets the requirements for vRealize Operations.

- 6 Click **Next**.

- 7 Enter a name for the primary node.

For example: **Ops-Master**

- 8 Enter the URL or IP address for the Network Time Protocol (NTP) server with which the cluster synchronizes.

For example: **nist.time.gov**

- 9 Click **Add**.

Leave the NTP blank to have vRealize Operations manage its own synchronization by having all nodes synchronize with the primary node and replica node.

- 10 Click **Next**.

- 11 Configure the vRealize Operations availability. To install vRealize Operations with availability, enable the **Availability Mode** and select High Availability or Continuous Availability. To continue your installation on full capacity, click **Next**.

Note You can enable High Availability or Continuous Availability after installation from the administrator interface.

- 12 Click the Add icon to add a node.
 - a Enter the **Node Name** and **Node Address**.
 - b Select the **Current Cluster Role**.

Note This step is optional if you use the default configuration. If you select High Availability for this cluster option, you can select a node from the added list of nodes to be the replica node. Although, only one node from the list can be selected as a replica node. For more information on High Availability, see [Adding High Availability to vRealize Operations](#). If you select Continuous Availability for this cluster, add at least one witness node and an even number of data nodes including the primary node and divide them across two fault domains. For more information, see [Adding Continuous Availability](#).

- 13 Click **Next**, and click **Finish**.

The administration interface appears, and it takes a moment for vRealize Operations to finish adding the primary node.

Results

You have created a primary node to which you can add more nodes.

What to do next

After creating the primary node, you have the following options.

- Create and add data nodes to the unstarted cluster.
- Create and add remote collector nodes to the unstarted cluster.
- Click **Start vRealize Operations Manager** to start the single-node cluster, and log in to finish configuring the product.

The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

About the vRealize Operations Master Node

The master node is the primary node that is the required, initial node in your vRealize Operations cluster.

The primary node performs administration for the cluster and must be online before you configure any new nodes. In addition, the primary node must be online before other nodes are brought online. If the primary node and replica node go offline together, bring them back online separately. Bring the primary node online first, and then bring the replica node online.

Advantages of a New Installation

You can use the new installation to create a primary node during the first installation of vRealize Operations. With the primary node in place, you can then start adding more nodes to form a cluster and then define an environment for your organization.

In a single-node clusters, administration and data is on the same primary node. A multiple-node cluster includes one primary node and one or more data nodes. In addition, there might be remote collector nodes, and there might be one replica node used for high availability. For continuous availability, you need a witness node and an even number of data nodes including the primary node. For more information on creating a primary node, see [About the vRealize Operations Master Node](#).


Installing vRealize Operations as an Administrator

As an administrator, you can install several instances of vRealize Operations build in your VM environment.

Introduction to Express Installation

Express installation is one possible way to create primary nodes, add data nodes, form clusters, and test your connection status. You can use express installation to save time and speed up the process of installation when compared to a new installation. Do not to use this feature unless the user is an administrator.

Figure 3-3. Express Installation from the Setup screen



Perform an Express Installation on the vRealize Operations product UI

Use express installation on the vRealize Operations cluster to create a primary node. Select express installation option when installing for the first time.

Prerequisites

Verify that you have a static IP address created from an OVF file.

Procedure

- 1 Navigate to the name or IP address of the node that will be the primary node of vRealize Operations.

The setup wizard appears, and you do not need to log in to vRealize Operations.

- 2 Click **Express Installation**.

- 3 Click **Next**.

- 4 Enter and confirm a password for the admin user account, and click **Next**.

Passwords require a minimum of 8 characters, one uppercase letter, one lowercase letter, one digit, and one special character.

The user account name is admin by default and cannot be changed.

- 5 Click **Next**.

- 6 Click **Finish**.

Results

You have created a primary node to which you can add more nodes.

Advantages of an Express Installation

Express installation saves time when compared to a new installation to create a new primary node. The express installation uses the default certificates, which differ from one organization to another. This feature is mainly used by the developers or the administrators.

Expand an Existing Installation of vRealize Operations

Use this option to add a node to an existing vRealize Operations cluster. You can use this option if you have already configured a primary node and you want to increase the capacity by adding more nodes to your cluster.

Introduction to Expand an Existing Installation

You can deploy and configure additional nodes so that vRealize Operations can support larger environments. A primary node always requires an additional node for a cluster to monitor your environment. With expanding your installation, you can add more than one node to your cluster.

Adding Data Nodes

Data nodes are the additional cluster nodes that allow you to scale out vRealize Operations to monitor larger environments.

You can dynamically scale out vRealize Operations by adding data nodes without stopping the vRealize Operations cluster. When you scale out the cluster by 25% or more, you should restart the cluster to allow vRealize Operations to update its storage size, and you might notice a decrease in performance until you restart. A maintenance interval provides a good opportunity to restart the vRealize Operations cluster.

In addition, the product administration options include an option to re-balance the cluster, which can be done without restarting. Rebalancing adjusts the vRealize Operations workload across the cluster nodes.

Figure 3-4. Expand an existing installation from the Setup screen



Note Do not shut down online cluster nodes externally or by using any means other than the vRealize Operations interface. Shut down a node externally only after taking it offline in the vRealize Operations interface.

Expand an Existing Installation to Add a Data Node

Larger environments with multiple-node vRealize Operations clusters contain one primary node and one or more data nodes for additional data collection, storage, processing, and analysis.

Prerequisites

- Create nodes by deploying the vRealize Operations vApp.
- Create and configure the primary node.
- Note the fully qualified domain name (FQDN) or IP address of the primary node.

Procedure

- 1 In a Web browser, navigate to the name or IP address of the node that will become the data node.

The setup wizard appears, and you do not need to log in to vRealize Operations.

- 2 Click **Expand an Existing Installation**.
- 3 Click **Next**.
- 4 Enter a name for the node (for example, **Data-1**).
- 5 From the Node Type drop-down, select **Data**.
- 6 Enter the FQDN or IP address of the master node and click **Validate**.
- 7 Select **Accept this certificate** and click **Next**.

If necessary, locate the certificate on the primary node and verify the thumbprint.

- 8 Verify the vRealize Operations administrator username of admin.
- 9 Enter the vRealize Operations administrator password.

Alternatively, instead of a password, type a pass-phrase that you were given by your vRealize Operations administrator.

10 Click **Next**, and click **Finish**.

The administration interface appears, and it takes a moment for vRealize Operations to finish adding the data node.

What to do next

After creating a data node, you have the following options.

- New, unstarted clusters:
 - Create and add more data nodes.
 - Create and add remote collector nodes.
 - Create a high availability primary replica node.
 - In a Web browser, navigate to the master node administration interface at **`https://master-node-name-or-ip-address/admin`**. Verify that all the nodes are listed under the **Nodes in the vRealize Operations Manager Cluster**. Then, click **Start vRealize Operations Manager** to start the cluster and to finish configuring the product.

The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

- Established, running clusters:
 - Create and add more data nodes.
 - Create and add remote collector nodes.
 - Create a high availability primary replica node, which requires a cluster restart.

Advantages of an Expanding an Installation

A data node shares the load of performing vRealize Operations analysis and it can also have an adapter installed to perform collection and data storage from the environment. You must have a primary node before you add data nodes to form a cluster.

Installing vRealize Operations on VMware Cloud on AWS

You can use your on-premises vRealize Operations to manage and monitor your cloud infrastructure on VMware Cloud by simply adding your VMware Cloud based vCenter Server into vRealize Operations. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations on-premises and VMware Cloud.

- Scale the existing vRealize Operations cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Manager Online Sizer](#).

Known Limitations

- Only migration planning and add/remove workload scenarios with VMware Cloud are supported.
- The compliance workflows in vRealize Operations work for virtual machines running on a vCenter Server in VMware Cloud on AWS. The compliance checks for VMware management objects such as Hosts, vCenter, and so on, are not available.
- Workload optimization including pDRS and host-based business intent does not work because VMware managers cluster configurations.
- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with vRealize Operations. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter Server interface.
- VMware Cloud does not support vRealize Operations plugin.
- You cannot log in to vRealize Operations using your VMware Cloud vCenter Server credentials.

Using vRealize Operations on-premises on VMware Cloud on AWS

Extend the monitoring capabilities of your vRealize Operations to monitor the VMware Cloud on AWS by creating a cloud account. Ensure that you have a cloud proxy or remote collector deployed on the VMware Cloud SDDC.

Procedure

- 1 Deploy the vRealize Operations remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Data Center level and select the **Compute Resource Pools** and validate your deployment. You can only select the workload datastore for storage when deploying the OVF in VMware Cloud.

Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations cluster in VMware Cloud. For more information, see the *Configuring a VMware Cloud on AWS Instance in vRealize Operations* topic in the *vRealize Operations Configuration Guide*.

Note In case of a vCenter adapter instance, set the **Cloud Type** to **VMware Cloud on AWS**.

Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-5. vRealize Operations On-Premises collecting data from VMware Cloud and AWS without remote data collectors

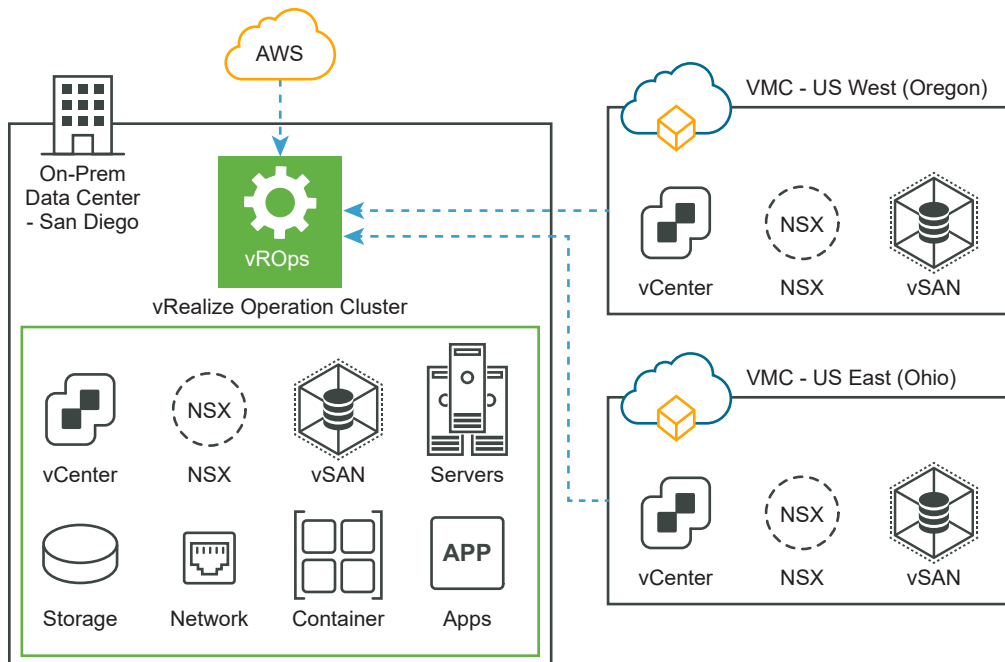
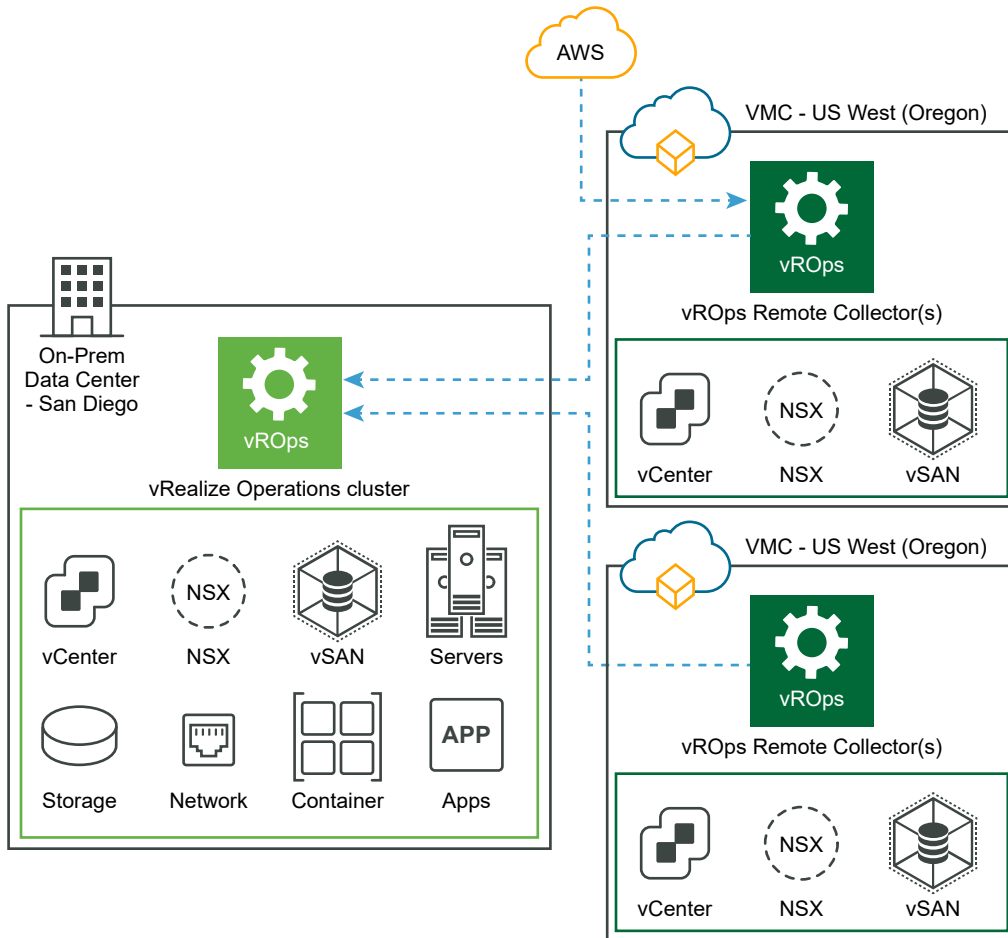


Figure 3-6. vRealize Operations On-Premises collecting data from VMware Cloud and AWS with remote data collectors



Deploying vRealize Operations on VMware Cloud on AWS

If you have moved a large part of your environment into VMware Cloud, you can deploy or migrate your vRealize Operations instance into VMware Cloud directly. After the vRealize Operations cluster is deployed on VMware Cloud, you can collect data from other VMware Cloud SDDCs and the SDDC located on-prem using remote collectors. You can deploy remote collectors to send over data into the centralized analytics cluster deployed in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations cluster in VMware Cloud, see [Deployment of vRealize Operations Manager](#).

Note Deploy the OVF template in the VMware Cloud on the data center level. VMware Cloud has two resource pools, the regular workload and the administrative workload. You can only deploy the new OVF template in the workload resource pool.

- 2 Deploy the remote collectors in vRealize Operations , see [Create a Remote Collector](#).

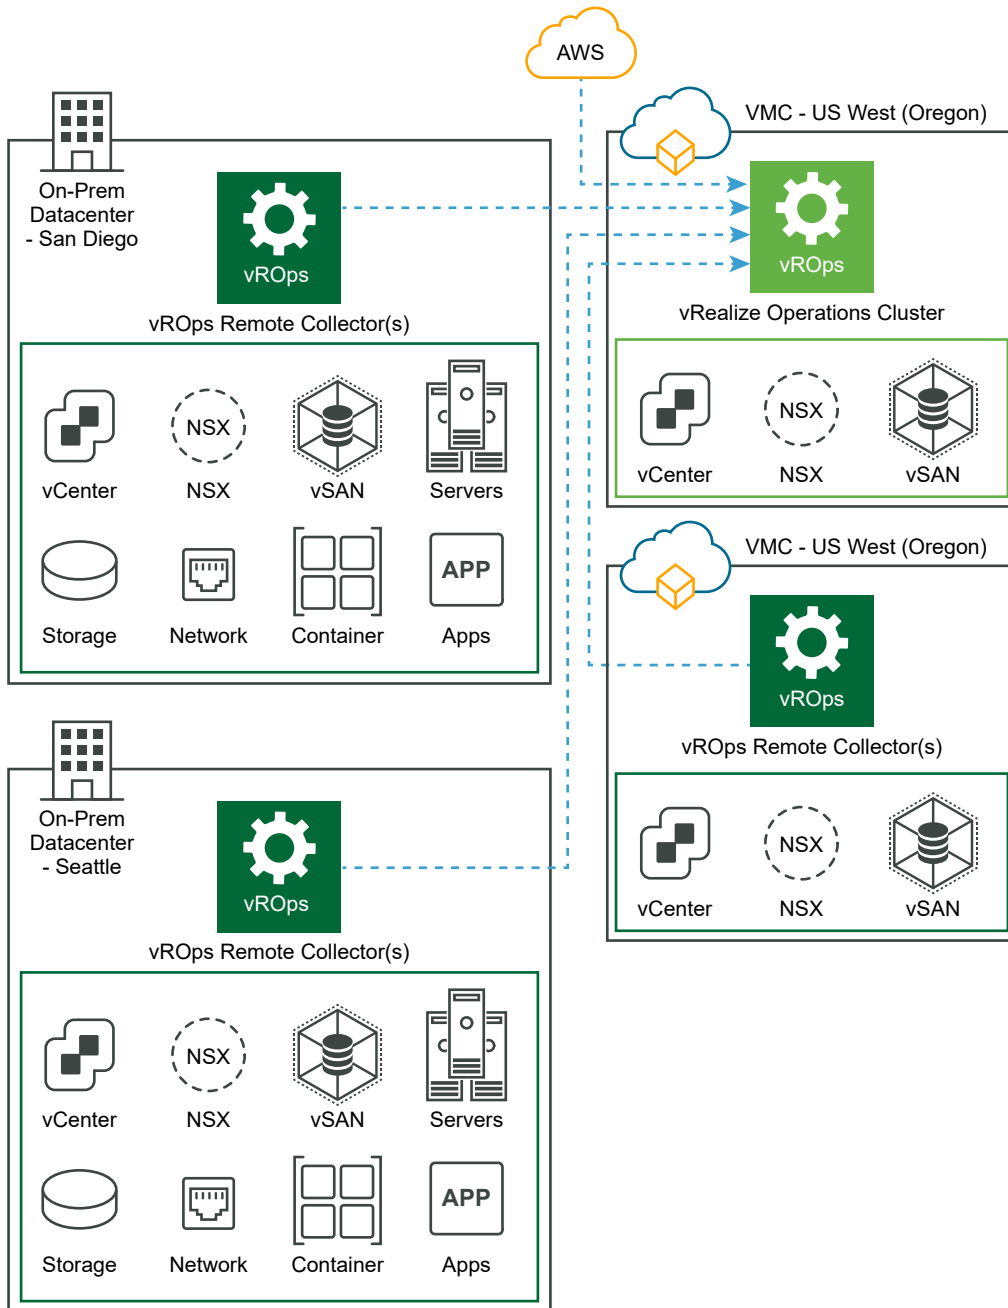
Note VMware Cloud is set in an isolated network and so, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations primary node and the remote collector you have created. To do so, you can use a VPN or a direct connection with no NAT.

- 3 Add and configure an adapter instance in the vRealize Operations cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Note If the Remote collectors are deployed on-premises, set **Cloud Type** to **Private Cloud**. However, if you deploy remote collectors in another VMware Cloud, set the **Cloud Type** to **VMware Cloud on AWS**.

Ensure that the remote collector is assigned to the adapter instance and the data collection of the adapter instance happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-7. vRealize Operations in VMware Cloud collecting data from other VMware Cloud SDDC, AWS, and On-Premise with remote data collectors



Installing vRealize Operations for Azure VMware Solution

You can use your on-premises vRealize Operations to manage and monitor your cloud infrastructure on VMware Cloud by adding the Azure VMware Solution cloud account. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations on-premises and VMware Cloud.
- Scale the existing vRealize Operations cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Manager Online Sizer](#).

Known Limitations

- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters might appear lower than expected and capacity remaining may appear higher than expected.
- Cost calculation based on reference database is supported on Azure VMware Solution.
- The end-user on the vCenter Server on Azure VMware Solution has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to vRealize Operations using the credentials of the vCenter Server on Azure VMware Solution.
- The vCenter Server on Azure VMware Solution does not support the vRealize Operations plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.

Using vRealize Operations on-premises for Azure VMware Solution

Extend the monitoring capabilities of your on-premises vRealize Operations to monitor the VMware Cloud vCenter Server by adding the Azure VMware Solution cloud account.

For more details, see [Configuring an Azure VMware Solution Instance in vRealize Operations](#).

Note If the network latency between vRealize Operations primary node and VMware Cloud is greater than 5 milliseconds, you should deploy remote collectors in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Data Center level and select the **Compute Resource Pools** and validate your deployment. You can only select the workload datastore for storage when deploying the OVF in VMware Cloud.

Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Note Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-8. (Recommended) vRealize Operations On-Premises collecting data from Azure VMware Solution with remote data collectors

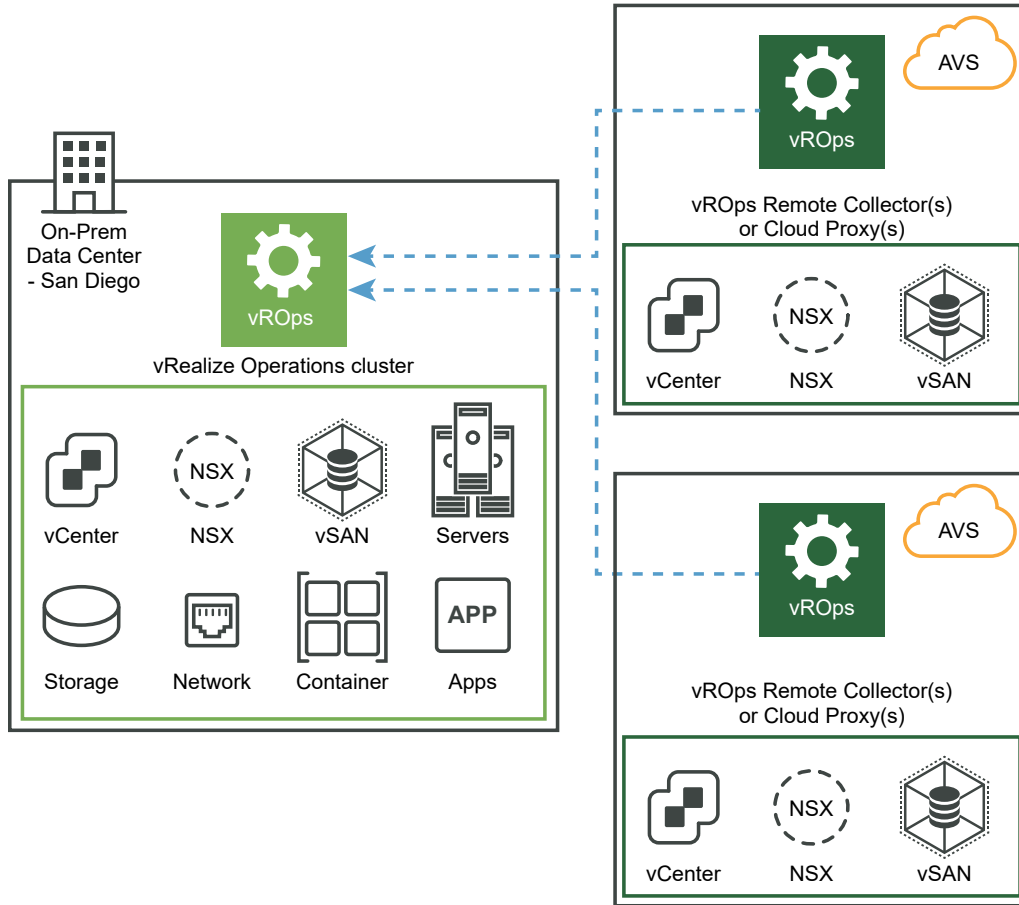



Figure 3-9. vRealize Operations On-Premises collecting data from Azure VMware Solution without remote data collectors



Deploying vRealize Operations on Azure VMware Solution

Deployment of vRealize Operations on Azure VMware Solution is not supported.

The options that are supported to monitor the Azure VMware Solution via vRealize Operations are as follows:

- Either through vRealize Operations deployed on-prem, or
- Through vRealize Operations Cloud

Installing vRealize Operations for Oracle Cloud VMware Solution

You can use your on-premises vRealize Operations to manage and monitor your cloud infrastructure on VMware Cloud by simply adding your VMware Cloud based vCenter Server into vRealize Operations. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations on-premises and VMware Cloud.
- Scale the existing vRealize Operations cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Manager Online Sizer](#).

Known Limitations

- Migration scenario is not supported in the What-if Analysis.
- Cost calculation is not supported on Oracle Cloud VMware Solution. Ignore all the cost metrics.

For configuration details, see [Configuring an Oracle Cloud VMware Solution Instance in vRealize Operations](#).

Using vRealize Operations on-premises for Oracle Cloud VMware Solution

Extend the monitoring capabilities of your on-premises vRealize Operations to monitor the VMware Cloud vCenter Server by connecting the VMware Cloud vCenter Server as an end point inside vRealize Operations.

Create an adapter instance both for vCenter Server and VMware vSAN to collect data from VMware Cloud and bring that into vRealize Operations. For more details, see [Configuring an Oracle Cloud VMware Solution Instance in vRealize Operations](#). You can either connect directly to the vCenter Server or use a remote collector which can be deployed inside a VMware Cloud SDDC to ensure that the data can be compressed and encrypted.

Note If the network latency between vRealize Operations primary node and VMware Cloud is greater than 5 milliseconds, you should deploy remote collectors in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Data Center level and select the **Workload resource pool** and validate your deployment.

Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Note Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-10. (Recommended) vRealize Operations On-Premises collecting data from Oracle Cloud VMware Solution with remote data collectors

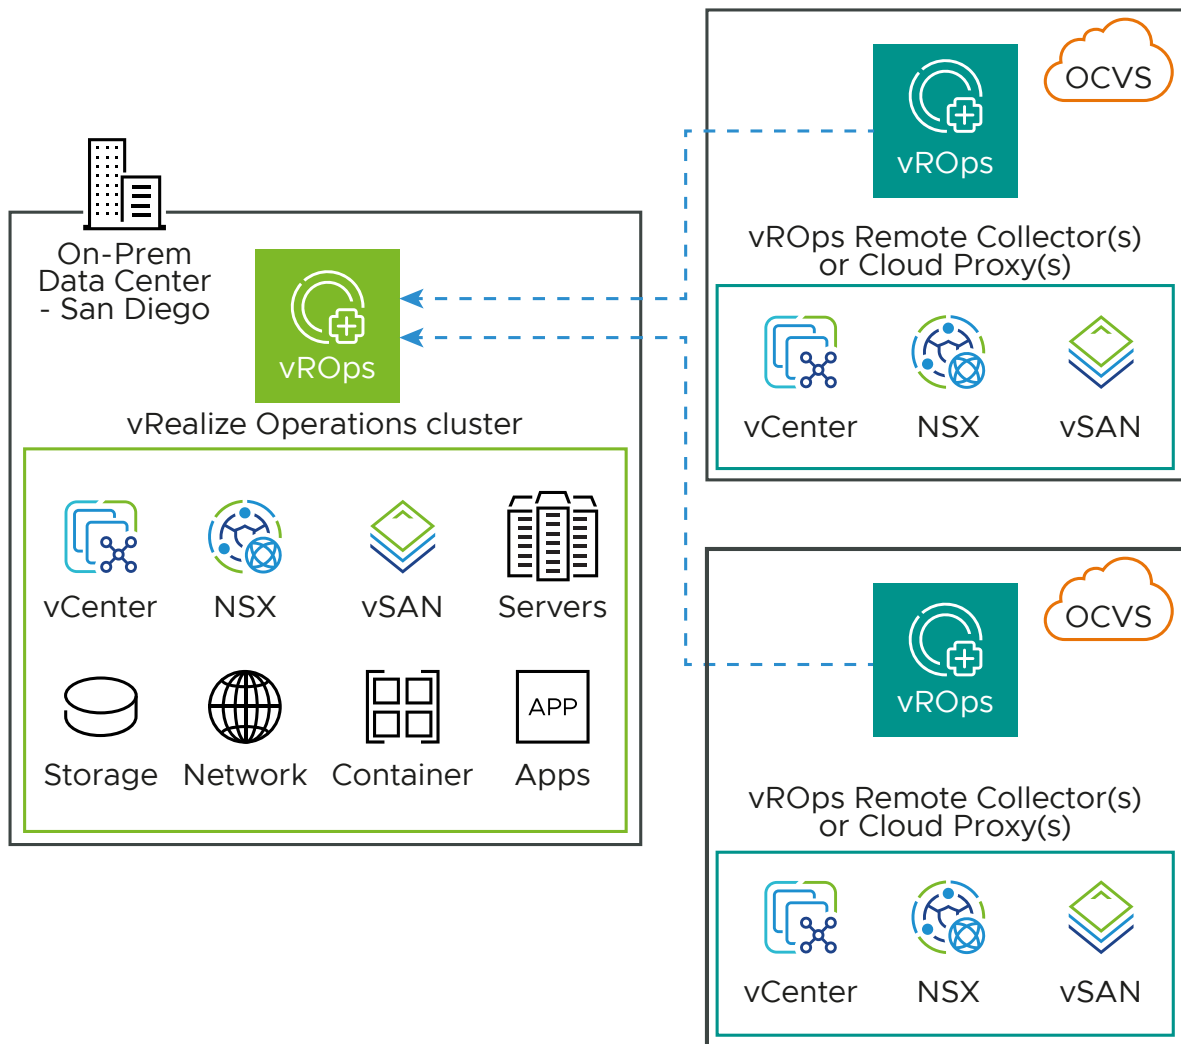
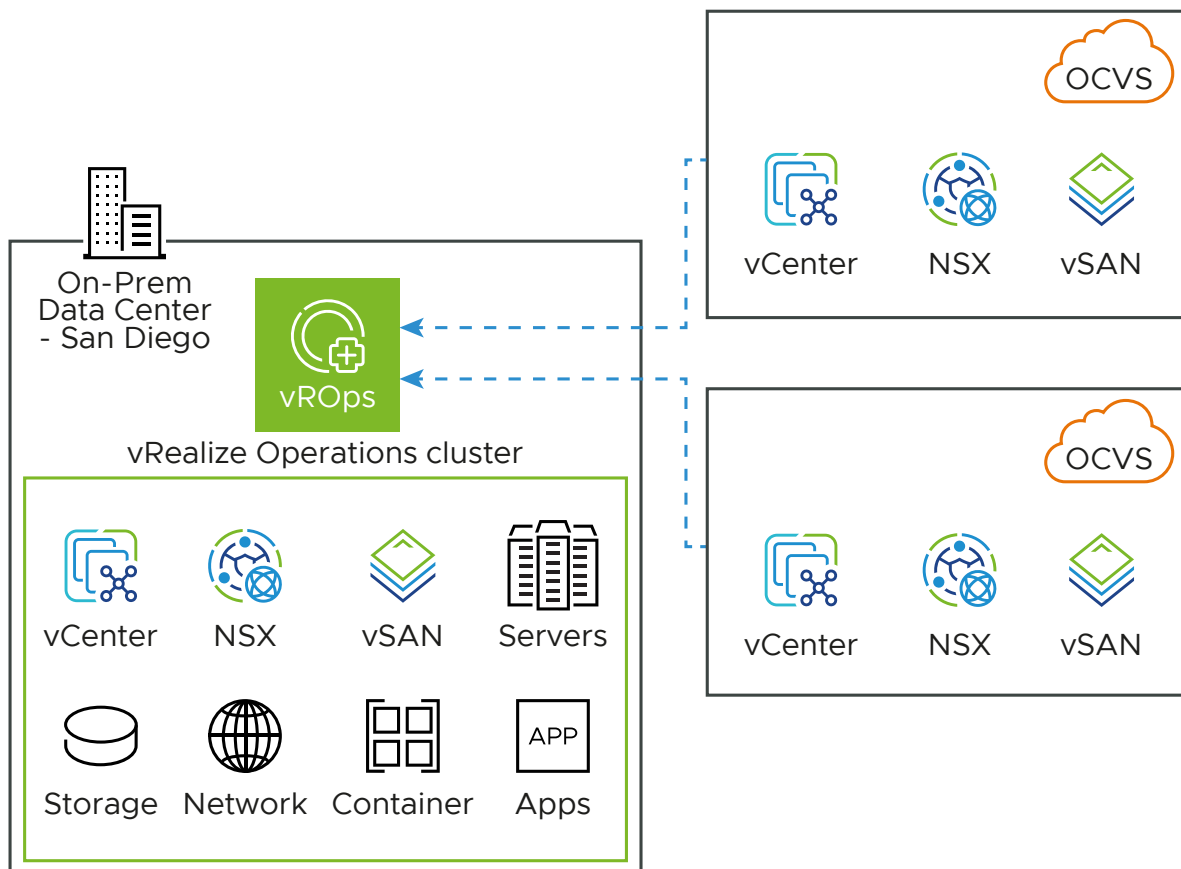


Figure 3-11. vRealize Operations On-Premises collecting data from Oracle Cloud VMware Solution without remote data collectors



Deploying vRealize Operations on Oracle Cloud VMware Solution

Deployment of vRealize Operations on Oracle Cloud VMware Solution is not supported.

Installing vRealize Operations for Google Cloud VMware Engine

You can use your on-premises vRealize Operations to manage and monitor your cloud infrastructure on VMware Cloud by simply adding your VMware Cloud based vCenter Server into vRealize Operations. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations on-premises and VMware Cloud.
- Scale the existing vRealize Operations cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Manager Online Sizer](#).

Known Limitations

- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters may appear lower than expected and capacity remaining may appear higher than expected.
- Cost calculation based on reference database is supported for Google Cloud VMware Engine.
- The end-user on the vCenter Server on Google Cloud VMware Engine has limited privileges. In-guest memory collection using VMware Tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to vRealize Operations using the credentials of the vCenter Server on Google Cloud VMware Engine.
- The vCenter Server on Google Cloud VMware Engine does not support the vRealize Operations plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.

Using vRealize Operations on-premises for Google Cloud VMware Engine

Extend the monitoring capabilities of your on-premises vRealize Operations to monitor the VMware Cloud vCenter Server by connecting the VMware Cloud vCenter Server as an end point inside vRealize Operations. Create an adapter instance both for vCenter Server and VMware

vSAN to collect data from VMware Cloud and bring that into vRealize Operations. You can either connect directly to the vCenter Server or use a remote collector which can be deployed inside a VMware Cloud SDDC to ensure that the data can be compressed and encrypted.

Note If the network latency between vRealize Operations primary node and VMware Cloud is greater than 5 milliseconds, you should deploy remote collectors in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Data Center level and select the **Compute Resource Pools** and validate your deployment. You can only select the workload datastore for storage when deploying the OVF in VMware Cloud.

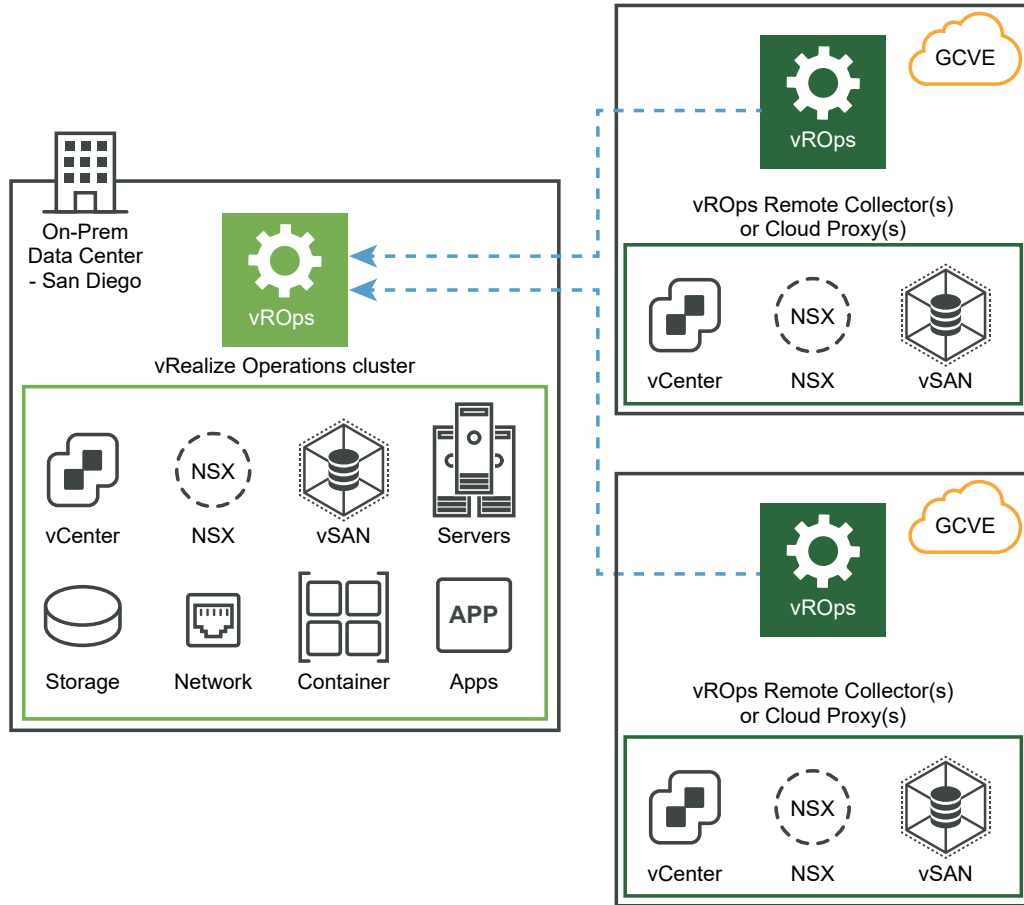
Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Important: When you configure the vCenter Server adapter, set the **Cloud Type** property to Google Cloud VMware Engine in the Advanced Settings.

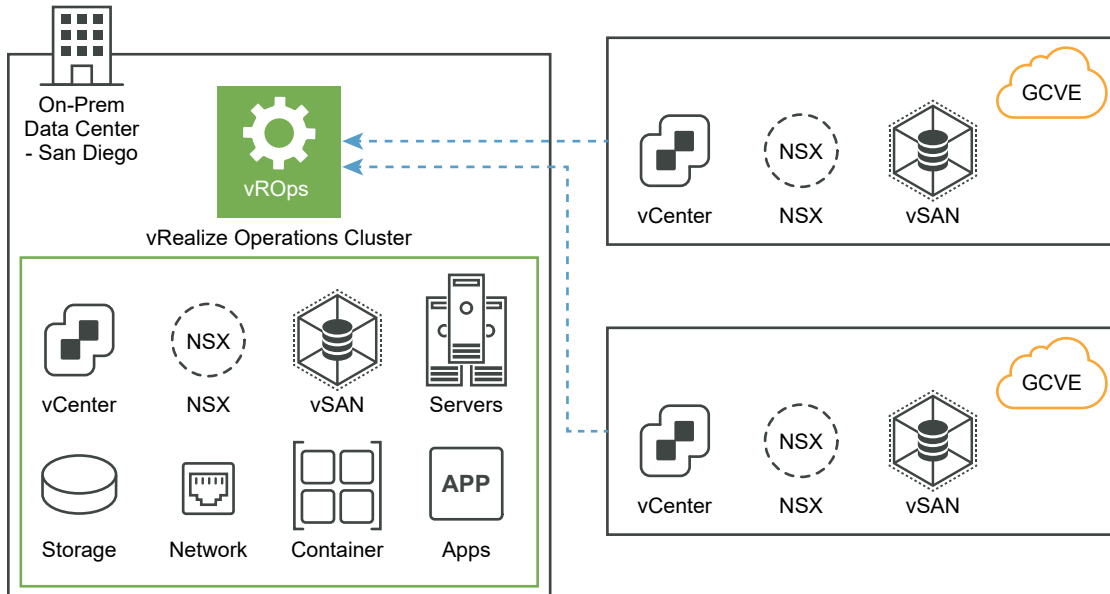
Note Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-12. (Recommended) vRealize Operations On-Premises collecting data from Google Cloud VMware Engine with remote data



collectors

Figure 3-13. vRealize Operations On-Premises collecting data from Google Cloud VMware Engine without remote data collectors



Deploying vRealize Operations on Google Cloud VMware Engine

Deployment of vRealize Operations on Google Cloud VMware Engine is not supported.

Installing vRealize Operations for VMware Cloud on Dell EMC

You can use your on-premises vRealize Operations to manage and monitor your cloud infrastructure on VMware Cloud by simply adding your VMware Cloud based vCenter Server into vRealize Operations. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations on-premises and VMware Cloud.
- Scale the existing vRealize Operations cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Online Sizer](#).

Known Limitations

- Cost calculation is not available for VMware Cloud on Dell EMC.
- The end-user on the vCenter Server on VMware Cloud on Dell EMC has limited privileges. In-guest memory collection using VMware Tools is not supported with virtual machines. Active and consumed memory utilization continue to work in this case.

- You cannot log in to vRealize Operations using the credentials of the vCenter Server on VMware Cloud on Dell EMC.
- The vCenter Server on VMware Cloud on Dell EMC does not support the vRealize Operations plugin.
- Workload optimization is not supported on VMware Cloud on Dell EMC because some management VMs could be moved improperly.
- Service Discovery on VMware Cloud on Dell EMC is supported in vRealize Operations FIPS disabled mode.
- Credential-less service discovery is not supported for VMware Cloud on Dell EMC.

Using vRealize Operations on-premises for VMware Cloud on Dell EMC

Extend the monitoring capabilities of your on-premises vRealize Operations to monitor the VMware Cloud vCenter Server by connecting the VMware Cloud vCenter Server as an end point inside vRealize Operations. Create an adapter instance both for vCenter Server and VMware vSAN to collect data from VMware Cloud and bring that into vRealize Operations. You can either connect directly to the vCenter Server or use a remote collector which can be deployed inside a VMware Cloud SDDC to ensure that the data can be compressed and encrypted.

Note If the network latency between vRealize Operations primary node and VMware Cloud is greater than 5 milliseconds, you should deploy remote collectors in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Data Center level and select the **Compute Resource Pools** and validate your deployment. You can only select the workload datastore for storage when deploying the OVF in VMware Cloud.

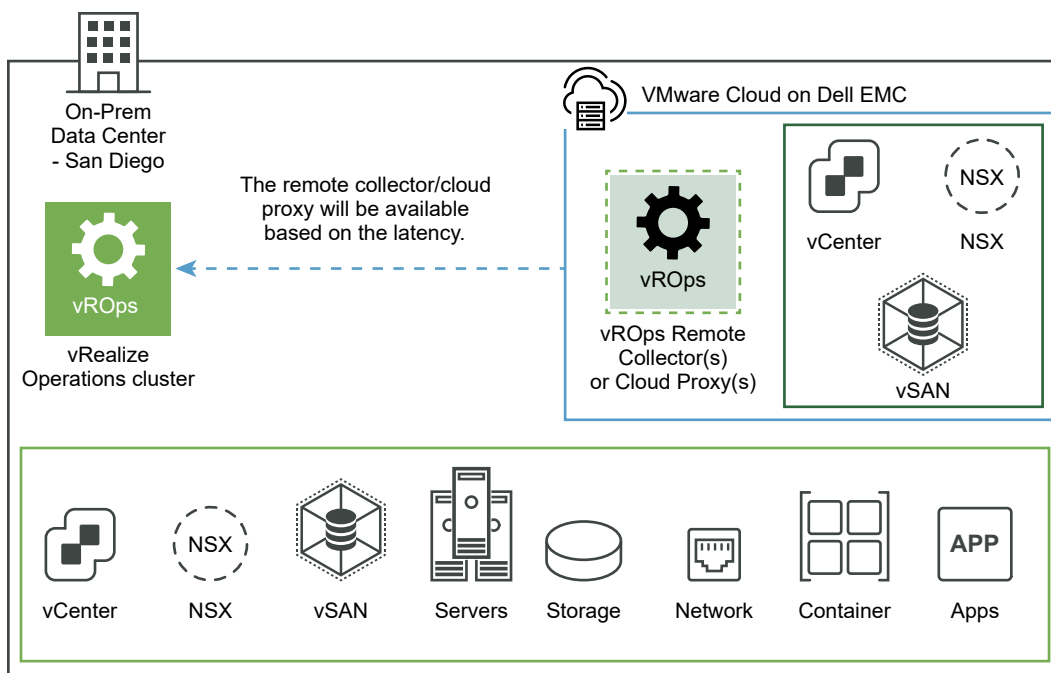
Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Server Cloud Account in vRealize Operations](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Important: When you configure the vCenter Server adapter, set the **Cloud Type** property to VMware Cloud on Dell EMC in the Advanced Settings.

Note Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-14. vRealize Operations On-Premises collecting data from VMware Cloud on Dell EMC



Deploying vRealize Operations on VMware Cloud on Dell EMC

If you have moved a large part of your environment into VMware Cloud, you can deploy or migrate your vRealize Operations instance into VMware Cloud directly. After the vRealize Operations cluster is deployed on VMware Cloud, you can collect data from other VMware Cloud SDDCs and the SDDC located on-prem using remote collectors. You can deploy remote collectors to send over data into the centralized analytics cluster deployed in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations cluster in VMware Cloud, see [Deployment of vRealize Operations](#).

Note Deploy the OVF template in the VMware Cloud on the data center level. VMware Cloud has two resource pools, the regular workload and the administrative workload. You can only deploy the new OVF template in the workload resource pool.

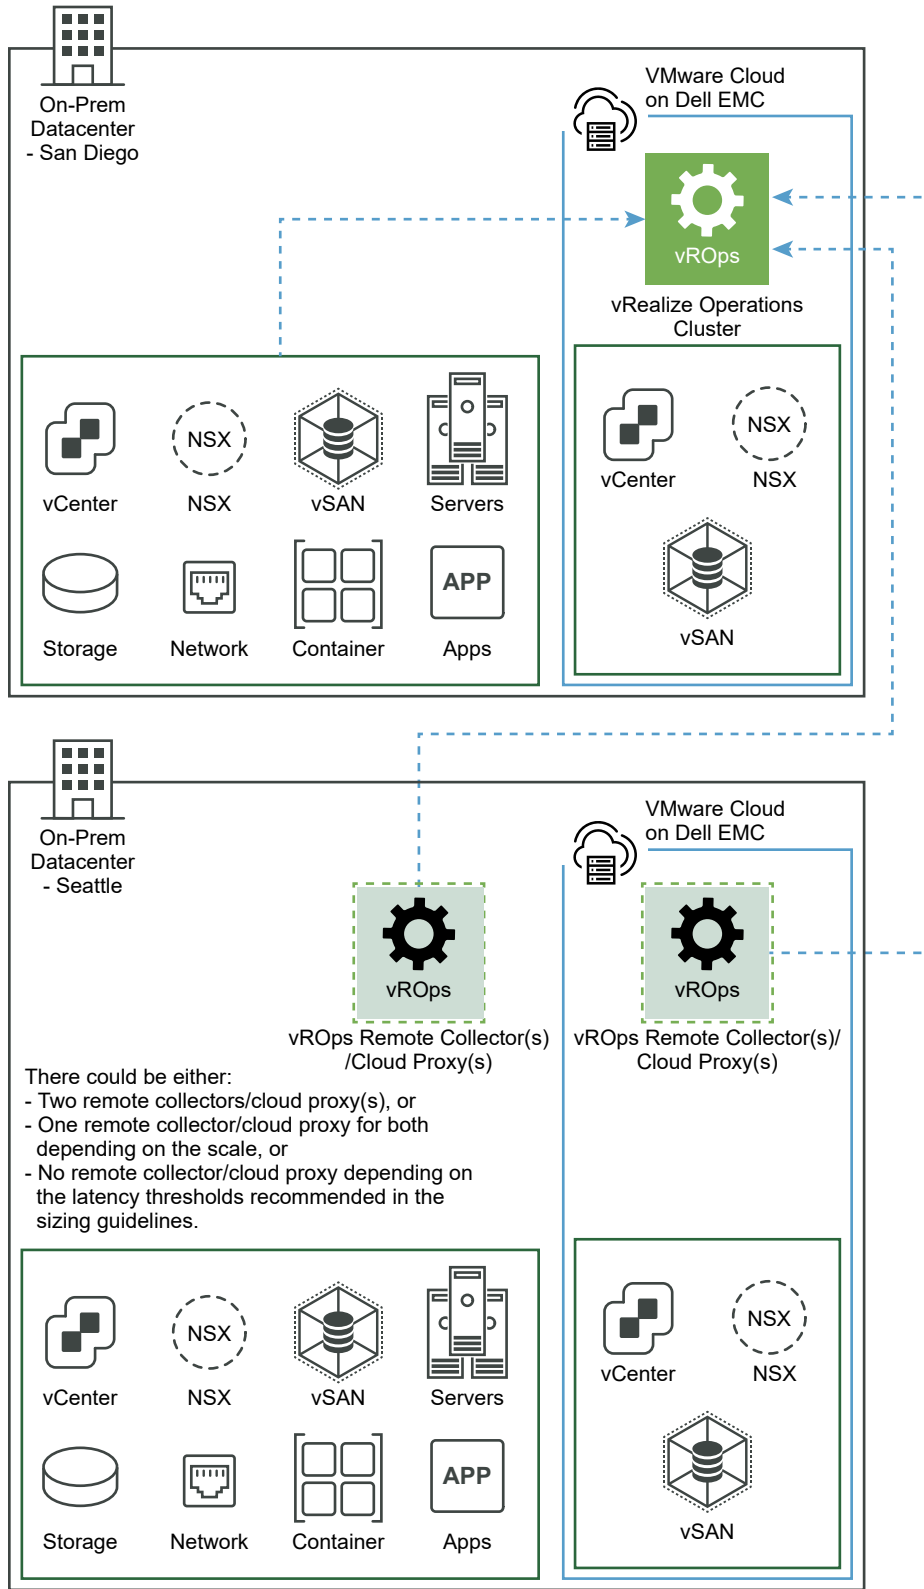
- 2 Deploy the remote collectors in vRealize Operations , see [Create a Remote Collector](#).

Note VMware Cloud is set in an isolated network and so, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations primary node and the remote collector you have created. To do so, you can use a VPN or a direct connection with no NAT.

- 3 Add and configure an adapter instance in the vRealize Operations cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Server Cloud Account in vRealize Operations](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Ensure that the remote collector is assigned to the adapter instance and the data collection of the adapter instance happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-15. vRealize Operations in VMware Cloud collecting data from VMware Cloud on Dell EMC and On-Premise with or without remote data collectors




Resize your Cluster by Adding Nodes

4

You can deploy and configure additional nodes so that vRealize Operations can support larger environments.

Figure 4-1. Workflow - Resize your cluster



This chapter includes the following topics:

- Gathering More Data by Adding a vRealize Operations Remote Collector Node
- Adding High Availability to vRealize Operations
- Adding Continuous Availability
- vRealize Operations Cluster and Node Maintenance
- Troubleshooting

Gathering More Data by Adding a vRealize Operations Remote Collector Node

You deploy and configure remote collector nodes so that vRealize Operations can add to its inventory of objects to monitor without increasing the processing load on vRealize Operations analytics.

Run the Setup Wizard to Create a Remote Collector Node

In distributed vRealize Operations environments, remote collector nodes increase the inventory of objects that you can monitor without increasing the load on vRealize Operations in terms of data storage, processing, or analysis.

Prerequisites

- Create nodes by deploying the vRealize Operations vApp.
During vApp deployment, select a remote collector size option.
- Ensure any remote adapter instance is running on the correct remote collector. If you have only one adapter instance, select Default collector group.
- Create and configure the primary node.
- Note the fully qualified domain name (FQDN) or an IP address of the primary node.
- Verify that there is one remote collector already added before you add another remote collector.

Note Remote collectors when added in parallel cause a cluster to crash.

Procedure

- 1 In a Web browser, navigate to the name or IP address of the deployed OVF that will become the remote collector node.

The setup wizard appears, and you do not need to log in to vRealize Operations.

- 2 Click **Expand an Existing Installation**.
- 3 Click **Next**.
- 4 Enter a name for the node, for example, **Remote-1**.
- 5 From the **Node Type** drop-down menu, select **Remote Collector**.
- 6 Enter the FQDN or IP address of the master node and click **Validate**.
- 7 Select **Accept this certificate** and click **Next**.

If necessary, locate the certificate on the primary node and verify the thumbprint.

- 8 Verify the vRealize Operations administrator username of **admin**.

- 9 Enter the vRealize Operations administrator password.

Alternatively, instead of a password, type a passphrase that you were given by the vRealize Operations administrator.

- 10 Click **Next**, and click **Finish**.

The administration interface appears, and it takes several minutes for vRealize Operations to finish adding the remote collector node.

What to do next

After creating a remote collector node, you have the following options.

- New, unstarted clusters:
 - Create and add data nodes.
 - Create and add more remote collector nodes.
 - Create a high availability primary replica node.
 - Click **Start vRealize Operations Manager** to start the cluster, and log in to finish configuring the product.

The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

- Established, running clusters:
 - Create and add data nodes.
 - Create and add more remote collector nodes.
 - Create a high availability primary replica node, which requires a cluster restart.

Adding High Availability to vRealize Operations

You can dedicate one vRealize Operations cluster node to serve as a replica node for the vRealize Operations primary node.

Run the Setup Wizard to Add a Primary Replica Node

To enable high availability (HA) for a vRealize Operations cluster, specify one of the data nodes to become a replica of the primary node.

Note If the cluster is running, enabling HA restarts the cluster.

You can add HA to the vRealize Operations cluster at installation time or after vRealize Operations is up and running. Adding HA at installation is less intrusive because the cluster has not yet started.

Prerequisites

- Create nodes by deploying the vRealize Operations vApp.
- Create and configure the primary node.
- Create and configure a data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the primary node.

Procedure

- 1 In a Web browser, navigate to the master node administration interface.

`https://master-node-name-or-ip-address/admin`

- 2 Enter the vRealize Operations administrator user name of **admin**.
- 3 Enter the vRealize Operations administrator password and click **Log In**.
- 4 Under High Availability, click **Enable**.
- 5 Select a data node to serve as the replica for the primary node.
- 6 Select the **Enable High Availability for this cluster** option, and click **OK**.

If the cluster was online, the administration interface displays progress as vRealize Operations configures, synchronizes, and rebalances the cluster for HA.

- 7 If the primary node and replica node go offline, and the primary remains offline for any reason while the replica goes online, the replica node does not take over the primary role, take the entire cluster offline, including data nodes and log in to the replica node command-line console as a root.
- 8 Open `$ALIVE_BASE/persistence/persistence.properties` in a text editor.
- 9 Locate and set the following properties:

```
db.role=MASTER
db.driver=/data/vcops/xdm/vcops.bootstrap
```

- 10 Save and close *persistence.properties*.
- 11 In the administration interface, bring the replica node online, and verify that it becomes the primary node and bring the remaining cluster nodes online.

What to do next

After creating a primary replica node, you have the following options.

- New, unstarted clusters:
 - Create and add data nodes.
 - Create and add remote collector nodes.
 - Click **Start vRealize Operations Manager** to start the cluster, and log in to finish configuring the product.

The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

- Established, running clusters:
 - Create and add data nodes.
 - Create and add remote collector nodes.

Adding Continuous Availability

Continuous availability prevents data loss in the event of one or more node failures. This mode requires one witness node, one primary node, and one data node divided across two fault domains. The witness node lies outside the fault domains. By default, the primary node is assigned to **Fault Domain 1**. The data node becomes the replica node and is assigned to **Fault Domain 2**. The primary node and the replica node create a pair. The number of data nodes including the primary node should always be an even number not exceeding 16. Each data node added to **Fault Domain 1** must have a pair in **Fault Domain 2** to preserve and replicate data that is added to its peer.

Enable Continuous Availability in vRealize Operations

You can enable continuous availability (CA) for vRealize Operations to protect your data if there is one or more node failures.

Note If the cluster is running, enabling CA restarts the cluster.

You can enable CA in the vRealize Operations cluster at the installation time or after vRealize Operations is up and running. Adding CA at installation is less intrusive because the cluster has not yet started.

Prerequisites

- Create nodes by deploying the vRealize Operations vApp.
- Create and configure the primary node.
- Create and configure the witness node.

Note vRealize Operations can have only one witness node in its cluster. While deploying an OVA file, you can select the recommended CPU/RAM configuration for the witness node.

- Create and configure one data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the primary node.

Procedure

- 1 In a Web browser, navigate to the master node administration interface.

`https://master-node-name-or-ip-address/admin`

- 2 Enter the vRealize Operations administrator user name of **admin**.
- 3 Enter the vRealize Operations administrator password and click **Log In**.
- 4 Under Continuous Availability, click **Enable CA**.

The Continuous Availability wizard opens. The Witness node exists outside the fault domains. The primary node is already assigned to **Fault Domain 1**.

Note You can enter names for each Fault Domain during installation. You can also edit the fault domain names after enabling continuous availability.

- 5 To create a pair with the primary node, drag the data nodes to **Fault Domain 2**.

Note You can add a maximum of 16 data nodes including the primary node and divide them between the fault domains to create eight pairs. You can also add remote collector nodes outside the fault domains as required.

- 6 Click **Ok**.

vRealize Operations Cluster and Node Maintenance

You perform cluster and node maintenance procedures to help your vRealize Operations perform more efficiently. Cluster and node maintenance involves activities such as changing the online or offline state of the cluster, fault domains, or individual nodes, enabling or disabling high availability (HA) or continuous availability (CA), reviewing statistics related to the installed adapters, and rebalancing the workload for a better performance.

You perform most vRealize Operations cluster and node maintenance using the Cluster Management page in the product interface, or the Cluster Status and Troubleshooting page in the administration interface. The administration interface provides more options than the product interface.

Table 4-1. Cluster and Node Maintenance Procedures

Procedure	Interface	Description
Change Cluster Status	Administration/Product	<p>You can change the status of a node to online or offline.</p> <p>In a high availability (HA) cluster, taking the primary or replica offline causes vRealize Operations to run from the remaining node and for HA status to be degraded.</p> <p>In continuous availability (CA) cluster, taking the primary or replica offline causes vRealize Operations to run in a degraded status.</p> <hr/> <p>Note You cannot convert a High Availability (HA) enabled cluster to a Continuous Availability cluster and vice versa. You must first disable the cluster availability, so that the cluster becomes a standard cluster and then enable HA or CA as required.</p> <hr/> <p>Any manual or system action that restarts the cluster brings all vRealize Operations nodes online, including any nodes that you had taken offline.</p>
Enable or Disable High Availability	Administration	<p>Enabling high availability requires the cluster to have at least one data node, with all nodes online or all offline. You cannot use Remote Collector nodes.</p> <p>To enable high availability, see Adding High Availability to vRealize Operations.</p> <p>Disabling high availability restarts the vRealize Operations cluster.</p> <p>After you disable high availability, the replica node in vRealize Operations converts back to a data node and restarts the cluster.</p>

Table 4-1. Cluster and Node Maintenance Procedures (continued)

Procedure	Interface	Description
Enable or Disable Continuous Availability	Administration	<p>Enabling continuous availability requires the cluster to have at least one witness node, and at least two data node, with all nodes online or all offline. You cannot use Remote Collector nodes. To enable continuous availability, see Adding Continuous Availability.</p> <p>Disabling continuous availability restarts the vRealize Operations cluster.</p> <p>When you disable continuous availability, you can choose to keep all your nodes or cut out one of the fault domains.</p> <ul style="list-style-type: none"> ■ Click Simply Disable with keeping all nodes to keep all your nodes when you disable continuous availability. <hr/> <p>Note You cannot disable continuous availability if one of your nodes is faulty. If you want to keep all your nodes, you must fix or replace the faulty node before you proceed.</p> <hr/> <ul style="list-style-type: none"> ■ Click Cut-Out one Fault Domain and then select the fault domain you want to keep. The other fault domain and the witness node are deleted. <p>After you disable continuous availability, the replica node in vRealize Operations converts back to a data node and restarts the cluster.</p>
Add Nodes	Administration	<p>You can add one or more nodes for your cluster. In a FIPS enabled environment, new nodes must be FIPS compliant. In a FIPS disabled environment, new nodes must be FIPS disabled. Enabling continuous availability requires one witness node, and an even number of data nodes including the primary node. For example, the cluster must have 2, 4, 6, 8, 10, 12, 14 or 16 nodes.</p>
Replace Nodes	Administration	<p>You can add nodes and replace them with a downed or non-functional node in a cluster.</p>
Generate Passphrase	Administration	<p>You can generate a passphrase to use instead of the administrator credentials to add a node to this cluster.</p> <p>The passphrase is only valid for a single use.</p>

Table 4-1. Cluster and Node Maintenance Procedures (continued)

Procedure	Interface	Description
Remove a Node	Administration	<p>When you remove a node, you lose data that the node had collected unless you are running in high availability (HA) mode. HA protects against the removal or loss of one node.</p> <p>You must not re-add nodes to vRealize Operations that you already removed. If your environment requires more nodes, add new nodes instead.</p> <p>When you perform maintenance and migration procedures, you should take the node offline, not remove the node.</p>
Configure NTP	Product	The nodes in vRealize Operations cluster synchronize with each other by standardizing on the primary node time or by synchronizing with an external Network Time Protocol (NTP) source.
Rebalance the Cluster	Product	You can rebalance adapter, disk, memory, or network load across vRealize Operations cluster nodes to increase the efficiency of your environment.

Cluster Management

vRealize Operations includes a central page where you can monitor and manage the nodes in your vRealize Operations cluster and the adapters that are installed on the nodes.

How Cluster Management Works

Cluster management lets you view and change the online or offline state of the overall vRealize Operations cluster or the individual nodes. In addition, you can enable or disable high availability (HA) and view statistics related to the adapters that are installed on the nodes.

Where You Find Cluster Management

From the left menu, click **Administration**, and then click the **Cluster Management** tile.

Cluster Management Options

The options include cluster-level monitoring and management features.

Table 4-2. Initial Setup Status Details

Option	Description
Cluster Status	Displays the online, offline, or unknown state of the vRealize Operations cluster. Once CA is enabled, it displays the status of the two fault domains.
High Availability	Indicates whether HA is enabled, disabled, or degraded.
Continuous Availability	Indicates whether CA is enabled, disabled, or degraded.

vRealize Operations provides node-level information and a toolbar for taking nodes online or offline.

Table 4-3. Nodes in the vRealize Operations Cluster

Option	Description
Node Name	Machine name of the node. The node that you are logged into displays a dot next to the name.
Node Address	Internet protocol (IP) address of the node. Primary and replica nodes require static IP addresses. Data nodes can use DHCP or static IP.
Cluster Role	Type of vRealize Operations node: primary, data, replica, or remote collector.
Fault Domain	Displays the fault domain a node is associated to in a CA enabled cluster. Note This column appears only if CA is enabled.
Node Pair	Displays which pair the node belongs to. For example, in CA, nodes are added in pairs. If there are four nodes, the column displays whether the node is part of pair one or two. Note This column appears only if CA is enabled.
State	Running, Not Running, Going Online, Going Offline, Inaccessible, Failure, Error
Status	Online, offline, unknown, or other condition of the node.
Objects in Process	Total environment objects that the node currently monitors.
Objects Being Collected	Total environment objects that the node collected.
Metrics in Process	Total metrics that the node has discovered since being added to the cluster.

Table 4-3. Nodes in the vRealize Operations Cluster (continued)

Option	Description
Metrics Being Collected	Total metrics the node has collected since being added to the cluster.
Version	Displays the vRealize Operations software version and the build number installed on the node.

In addition, there are adapter statistics for the selected node.

Table 4-4. Adapters on Server

Option	Description
Name	Name that the installing user gave to the adapter.
Status	Indication of whether the adapter is collecting data or not.
Objects Being Collected	Total environment objects that the adapter currently monitors.
Metrics Being Collected	Total metrics that the adapter has collected since being installed on the node.
Last Collection Time	Date and time of the most recent data collection by the adapter.
Added On	Date and time when the adapter was installed on the node.

Troubleshooting

Troubleshooting Cluster Problems

A multi-node vRealize Operations cluster does not behave as expected.

Problem

A multi-node vRealize Operations cluster does not behave as expected because of general problems within the cluster or because of suspected firewall concerns.

The problems might occur because of multiple reasons:

- You may be unable to install or uninstall management packs.
- The node shows as offline in the user interface even though it is online.
- You might face problems with new nodes joining the cluster.

Solution

Login to each vRealize Operations node in the cluster and run the following script:

```
$VMWARE_PYTHON_3_BIN /usr/lib/vmware-casa/bin/Netcheck.py
```

On each node, you are presented with a list of attempted connections. If a node cannot connect to the required port, it is reported in the list. Ports that do not connect must be investigated.

Note Only one port is required within the range of 10002-10010 and 20002-20010.

For more information see KB article [82421](#).

Installing Cloud Proxy

5

Install cloud proxy on your on-premise vRealize Operations to collect data across different geo locations.

Note FIPS mode is supported in cloud proxy. To leverage this functionality, make sure your cluster is in FIPS mode.

This chapter includes the following topics:

- [Configuring Cloud Proxies in vRealize Operations](#)
- [Managing Cloud Proxies in vRealize Operations](#)
- [Cloud Proxy FAQ](#)
- [Cloud Proxy Troubleshooting](#)

Configuring Cloud Proxies in vRealize Operations

Using cloud proxies in vRealize Operations, you can collect and monitor data from your remote data centers. Typically, you need only one cloud proxy per physical data center. You can deploy one or more cloud proxies in vRealize Operations to create a one-way communication between your remote environment and vRealize Operations. The cloud proxies work as one-way remote collectors and upload data from the remote environment to vRealize Operations. Cloud proxies can support multiple vCenter Server accounts.

Prerequisites

- Verify that you have an IP address, a DNS entry, and permissions to deploy OVF templates in vSphere.
- Log in to vSphere and verify that you are connected to a vCenter Server system.
- Allow outgoing HTTPS traffic for cloud proxy over port 443.
- Allow outgoing traffic from the endpoints to cloud proxy over 443, 4505, and 4506.
- Add a vCenter cloud account and provide an account with the following read and write privileges:
 - vCenter IP address or FQDN
 - Permissions required to install a cloud proxy on the vCenter Server.

For more information on privileges, see the topic called "Privileges Required for Configuring a vCenter Adapter Instance" in the *vRealize Operations Configuration Guide*.

- Cloud proxies must have a proper DNS resolution to the vRealize Operations nodes when using short/long FQDN names. This is applicable to on-prem cloud proxy.

Procedure

- 1 Log in to vRealize Operations.
- 2 From the left menu, click **Data Sources > Cloud Proxy**, and then click **New**.
- 3 Save the OVA path. Optionally, click **Download Cloud Proxy OVA** to download and save the OVA file locally.
 - To copy the link for the VMware vRealize® Operations Cloud Appliance™, click the **Copy Path** icon for the Cloud Proxy OVA.
 - To download and save the OVA file locally, click **Download Cloud Proxy OVA**.
- 4 Navigate to your vSphere, select the name of your vCenter Server cluster, and select **Deploy OVF Template** from the **Actions** menu.
- 5 Insert the ova link and then click **Next**.
 - Paste the cloud proxy ova link in the **URL** field.
 - Click the **Local File** option, browse, and select the downloaded OVA file.
- 6 Follow the prompts to install the OVA on your vCenter Server.

For the most current information about sizing and scaling, see [Knowledge Base article 78491](#).
- 7 When prompted to enter the One Time Key (OTK) in the **Customize template** screen, return to the Install Cloud Proxy page in vRealize Operations, and click the **Copy Key** icon.

The One Time Key expires 24 hours after generation. To avoid using an expired key, click **Regenerate Key** before proceeding. The one time key is used by the cloud proxy to authenticate to vRealize Operations.
- 8 Return to vSphere and paste the key in the **One Time Key** text box to install the vRealize Operations Cloud Appliance.
- 9 Select **Use IPv6** to use IPv6 for internal communications. For more information, see [Using IPv6 with vRealize Operations](#).
- 10 (Optional) Set up a proxy server in the **Customize template** screen.
 - a Enter details in the **Network Proxy IP Address** and **Network Proxy Password** properties.
 - b To enable SSL, select the **Use SSL connection to proxy** check box.
 - c If you are using SSL, you can verify the certificate of the proxy server. Public certificate authorities are used to verify the proxy server certificate. To enable this, select the **Verify proxy's SSL cert** check box in the **Verify SSL cert** property.

- d You can specify the IP /FQDN URL that is used to access the system when a load balancer is used.
- e If you have a custom certificate authority, paste the root certificate authority in the **Custom CA** property to verify the certificate of the proxy server. The root certificate authority is passed on to the cloud proxy. Do not include the following lines from the certificate authority:

```
"-----BEGIN CERTIFICATE-----"
```

```
"-----END CERTIFICATE-----"
```

For more information on adding CA certificates while deploying a cloud proxy in vRealize Operations, see the VMware KB article [83698](#).

You can use the Load Balancer Custom CA for the vRealize Operations environment

- 11 Click **Finish**.

The deployment takes a few minutes to finish.

- 12 Locate the cloud proxy you just installed, select the vRealize Operations Cloud Appliance, and click **Power on**.

Note You must power on the vRealize Operations Cloud Appliance within 24 hours of registering it. After 24 hours, the One Time Key expires, and you must delete the vRealize Operations Cloud Appliance and deploy another cloud proxy.

- 13 Return to the Cloud Proxy page in vRealize Operations to view the status of the cloud proxy you just installed.

Option	Description
Name	The name of the cloud proxy.
IP	The IP address of the cloud proxy.
Status	Status of the cloud proxy. For example, the Getting Online status is displayed for a few minutes when you add a new cloud proxy. Once the cloud proxy is connected to vRealize Operations, the status changes to Online. If the vRealize Operations is not connected, the Offline status is displayed.
Version	The version used to install the cloud proxy.
Accounts	The number of accounts that are created and associated with the cloud proxy.
Network Proxy Address	The network proxy address of the cloud proxy.
Network Proxy Port	The network proxy port number of the cloud proxy.

- 14** To view the accounts that are using this connection, click the Cloud Proxy.

The communication from the cloud proxy to cloud is one way. The cloud proxy initiates this connection and if necessary, it also pulls data from cloud (like the adapters configuration or upgrade pak). The cloud proxy requires a regular Internet access over the https protocol but it does not need any special firewall configuration. The cloud proxy verifies the certificate of the cloud service it connects to and if there are transparent proxy servers which do stop SSL, it might cause connectivity problems for the cloud proxy.

The cloud proxy also supports connection through the corporate proxy server. The proxy settings are given during OVF deployment.

- 15** (Optional) To remove a cloud proxy, click **Remove**.

What to do next

Upgrade your cloud proxy. For more information, see the topic called Upgrading Cloud Proxy in the *VMware vRealize Operations vApp Deployment Guide*.

The VMware vSphere solution connects vRealize Operations to one or more vCenter Server instances. For more information see the topic called Configure a vCenter Server Cloud Account in vRealize Operations in the Connecting to Data Sources section in the *VMware vRealize Operations Configuration Guide*.

Managing Cloud Proxies in vRealize Operations

You can use cloud proxies in vRealize Operations to collect and monitor data from your on-premises data centers.

Cloud proxies provide high availability within your cloud environment, you can group two or more cloud proxies to form a collector group. The cloud proxy collector group ensures that there is no single point of failure in your cloud environment. If one of the cloud proxies experiences a network interruption or becomes unavailable, the other cloud proxy from the collector group takes charge and ensures that there is no downtime. All other user-initiated manual operations on the collector, such as to stop or restart the collector manually, do not result in automated rebalancing.

Note When cloud proxies provide high availability within your cloud environment, the cluster can survive the loss of one data node without losing any data. However, the cloud proxy does not guarantee that the adapter instance will collect all the data during adapter instance failover (or when reassigning the adapter instance). vRealize Operations cloud proxy only provides additional application level data protection to ensure application level availability

You can also use cloud proxies to rebalance the resources across the collectors in your collector group. The Rebalance option is available as part of the Edit menu in the Collector Groups page.

Note You can use the rebalance option before the vCenter Adapter initiates data collection. Once the data collection starts, the rebalance option is disabled.

Adding Cloud Proxies To a Collector Group

You can create a cloud proxy collector group from the available collectors in your cloud environment. You can add two or more cloud proxies to a collector group.

Where You Add New Cloud Proxies

From the left menu, click **Data Sources > Cloud Proxy**. Click **New**.

Add New Cloud Proxy Workspace

Option	Description
Name	Name of the cloud proxy.
IP	IP address of the cloud proxy VM in the vCenter Server.
Status	Status of the cloud proxy when connected. For example, the Online status is displayed when the VM is connected.
Cloud Accounts	Number of cloud accounts that are created and associated with the cloud proxy.
Monitoring Accounts	Number of cloud accounts that are monitored using the cloud proxy.
IP/FQDN	The IP or FQDN address of the vCenter Server instance to which the cloud proxy is connected.
Port	The network port that vRealize Operations uses to communicate with a vCenter Server system and vRealize Operations components.

Linking Cloud Proxy with a Collector Group

When you create collector groups in your cloud environment, you have the option to include one or more cloud proxies in the Collector Group.

Note It is recommended that you do not add cloud proxy to a collector group from remote collectors. For cloud proxy, a separate cloud proxies group can be created which contains only cloud proxies.

From the **Add New Collector Group** page, select one or more cloud proxy accounts you want to link with the collector group and click **Save**. The selected cloud proxy accounts are now part of the collector group.

Monitoring the Health of Cloud Proxies

You can view the status and health of your cloud proxy after you add it in vRealize Operations. You can then monitor the health and view alerts and metrics of your cloud proxy using the vRealize Operations Cloud Proxy object.

Procedure

- 1 Log in to vRealize Operations.
- 2 From the left menu, click **Data Sources > Cloud Proxy**.

The list of cloud proxies is displayed.

- 3 Click a **Cloud Proxy**.

The **Cloud Proxy Details** page opens.

Each cloud proxy might have one or more adapters. You can also view the health and status of these adapters from this page.

Table 5-1. Cloud Proxy Page Options

Option	Description
Proxy ID	ID of the cloud proxy.
IP Address	IP address of the cloud proxy.
OVA Version	The OVA file version used to install the cloud proxy.
Creation Date	Date of creation of the cloud proxy.
Status	Status of the cloud proxy. For example, the Getting Online status is displayed for a few minutes when you add a cloud proxy. Once the cloud proxy is connected to vRealize Operations, the status changes to Online. If the vRealize Operations is not connected, the Offline status is displayed.
Last Heartbeat	Last time stamp when vRealize Operations ran a Health Check for this cloud proxy. When you click a cloud proxy to view its details, vRealize Operations sends a heartbeat to check if the cloud proxy is still reachable.
CPU	CPU usage.
Memory	Memory usage.

- 4 If your cloud proxy is not collecting data, you can view the health of the cloud proxy. From the left menu, click **Environment > Inventory**, select the **vRealize Operations Cloud Proxy Object** from the list, and then click **Show Detail**.

For more details, see [Inventory Tab](#) and [Inventory: List of Objects](#).

- 5 After you locate the vRealize Operations Cloud Proxy object, you can view the object details using the Summary tab. For more information, see [Summary Tab](#).

- 6 Use the [Alerts](#) tab to monitor the health of the cloud proxy. If there are any issues, troubleshoot them using the [Metrics](#) tab.

If your cloud proxy is not working properly, an alert is displayed.

```
One or more vRealize Operations services on a cloud proxy are down
```

To clear this alert, perform the following steps:

- Check the network connectivity and configuration for the cloud proxy.
- Take the cloud proxy offline and then bring it online.

If the problem still persists contact VMware support.

Note It is recommended that you create a notification rule for this alert so that, quick remediation steps can be taken, if necessary.

- 7 (Optional) You can use the cloud proxy command line interface for other cloud proxy related actions. For more details, see [Using the Cloud Proxy Command-Line Interface](#).

Upgrading Cloud Proxy

Cloud Proxies are upgraded to a compatible cluster version automatically after the cluster upgrade. Expect a downtime of one or two cycles, as the cloud proxy does not collect any data during this period. Data collection resumes after the upgrade is complete. In case the automatic upgrade fails, you can upgrade your cloud proxy manually using the CLI.

For more information on what data gets collected, see the topic called "VMware vSphere Solution in vRealize Operations" in the *VMware vRealize Operations Configuration Guide*.

You can manually upgrade your cloud proxy [Using the Cloud Proxy Command-Line Interface](#).

Using the Cloud Proxy Command-Line Interface

You can use SSH to access the cloud proxy instance and use its Command-Line Interface to run the following actions:

- Manually upgrade your cloud proxy in case the automatic download of the latest binary fails. When automatic download fails, you see a notification on the vRealize Operations user interface. To manually upgrade your cloud proxy instance to latest version, see the following KB article [80590](#).
- Generate support bundle.
- Gather the status of the cloud proxy's health and connectivity details.

Command Line	Description
<code>cprc-cli -h, --help</code>	Displays the help message and use of command-line interface.
<code>cprc-cli -s, --status</code>	Prints the cloud proxy life-cycle status, configuration details, upgrade related information and more. It is useful to catch necessary information related to support and troubleshooting, or to check the connection to vRealize Operations Cloud, or to check the product version number, and so on.
<code>cprc-cli -u PRODUCT_PAK, --upgrade PRODUCT_PAK</code>	The cloud proxy instance is enabled for an automated upgrade by default. But if the automated upgrade fails due to any exceptional issue, use this command line to upgrade your cloud proxy instance to the desired version.
<ul style="list-style-type: none"> ■ 8.3 Release <code>cprc-cli -sb, --generate-support-bundle</code> ■ 8.4 Release <code>cprc-cli -sb, --generate-support-bundle</code> ■ 8.5 Release <code>cprc-cli IS_HEAVY -sb, --generate-support-bundle IS_HEAVY</code> <p>The <code>IS_HEAVY</code> option should be specified as true or false. For example:</p> <pre>cprc-cli -sb true</pre> <pre>cprc-cli -sb false</pre> <p>With the true option, the support bundle is generated with journalctl logs. With the false option, the support bundle is generated without journalctl logs</p>	Generates the cloud proxy support bundle which is a package of logs, configurations, and status files. The support bundles are necessary for product support and troubleshooting. Generated support bundles can be found at the <code>/storage/db/vmware-vrops-cprc/support/</code> location.
<code>cprc-cli -rsb SUPPORT_BUNDLE, --remove-support-bundle SUPPORT_BUNDLE</code>	Removes any specified support bundle. Although generated support bundle packages can be removed using system embedded commands, it is recommended to use this command for that action.
<code>cprc-cli -fm, --enable-fips-mode</code>	Enables FIPS mode for cloud proxy.

Cloud Proxy FAQ

This topic covers some frequently asked questions about vRealize Operations Proxy.

Configuration

1 What are the prerequisites for setting up a cloud proxy account?

Prerequisites are given in the topic, [Configuring Cloud Proxies in vRealize Operations](#).

2 What does one-way connection mean?

Only outbound connections are initiated from cloud proxy to vRealize Operations, over `tcp/443`. Inbound ports to cloud proxy are not required. This ensures higher security as firewall ports need not be open to allow incoming connections. Also, cloud proxy can facilitate vCenter actions.

3 How do I edit environment settings for cloud proxy?

You can edit vApp options. For more information, see [Edit OVF Details for a Virtual Machine](#).

4 How are certificates managed?

Certificates are managed by cloud proxies. But for any additional proxy servers with SSL communication, you need to provide certificate(s).

5 What credential is used to login to cloud proxy?

You can login as the “root” user. You are expected to set a new password on the first login to cloud proxy VM.

SSH access is disabled by default, so the first login must be done via the vCenter console. You can run the following command to start SSH service:

```
systemctl start ssh
systemctl enable sshd
```

To reset password, see the VMWare KB Article, [2001476](#).

6 Where can I configure the local HTTP proxy for VMC on AWS?

Perform the following steps:

- a Login to vRealize Operations and go to the Administration page.
- b Go to Cloud Accounts.
- c Select VMC on AWS.
- d Click + next to credentials to add a credential.
- e In proxy details, add details for the local HTTP proxy. (Do not add details for cloud proxy here).

For more details, see the Configuring VMware Cloud on AWS in vRealize Operations Cloud topic in the *vRealize Operations Configuration Guide*.

7 Will I be notified if the connection between cloud proxy and vRealize Operations breaks down?

You can configure alerts/notifications on the *vRealize Operations cloud proxy* object. For more information, see [Monitoring the Health of Cloud Proxies](#).

vRealize Operations automatically generates notifications for the following scenarios:

- Cloud proxy is not reachable.
- Cloud proxy is nearing sizing limits.

8 How do I change account for cloud proxy?

You can edit vApp options. For more information, see [Edit OVF Details for a Virtual Machine](#).

9 How can I check the status of cloud proxy?

For more information, see [Monitoring the Health of Cloud Proxies](#).

Sizing

1 How should I size the cloud proxy?

For information on sizing, see the VMWare KB article [85832](#)

2 How would I know if cloud proxy is nearing sizing limit?

vRealize Operations customers will receive an email when cloud proxy is nearing sizing limit.

Upgrade

1 How do I upgrade cloud proxy?

Cloud proxy is upgraded automatically. In case the upgrade fails, see the VMWare KB article [80590](#).

Migration

1 What is the difference between Remote Collector, Application Remote Collector and cloud proxy?

The Remote Collector performs the data collection role from remote location sites and uploads data to the analytics nodes. Bidirectional connectivity is required between Remote Collector and analytics nodes. The Application Remote Collector discovers and collects data for applications running in Guest operating systems at a scale .

Cloud proxy takes the role for both Remote Collector and Application Remote Collector, in addition, it needs only one-way connectivity to analytics nodes and does not require connectivity from analytics nodes to itself.

The best practice for the on-prem vRealize Operations users is to leverage cloud proxy, for the vRealize Operations Cloud users this is the only supported option.

2 Should I use Remote Collector or cloud proxy for monitoring?

VMWare recommends that you use cloud proxy to take advantage of the latest enhancements. Also, application monitoring is only supported through cloud proxy.

High Availability

1 Is high availability supported?

Cloud proxy supports high availability. You can add multiple cloud proxies to a collector group. If the collecting cloud proxy fails or gets disconnected, collection can be picked up by another proxy in the group.

Note Since the failover is initiated after a period of 10 minutes, few collection cycles are lost.

To troubleshoot cloud proxy issues, see [Cloud Proxy Troubleshooting](#).

Cloud Proxy Troubleshooting

Cloud proxy troubleshooting steps are provided to help you easily resolve issues that you may come across in vRealize Operations.

Before you proceed with troubleshooting, see the [Cloud Proxy FAQ](#).

Installation and/or First Boot Failure

To verify the issue, check if `/var/log/firstboot` contains a file named "Succeeded".

If not, the following problems could result in vRealize Operations installation and/or first boot failure:

- 1 OTK used while deploying Cloud Proxy is invalid. To verify, check the cloud proxy console.

Solution: Redeploy cloud proxy.

Cloud Proxy VM is running, but the status is Offline in vRealize Operations.

Cloud Proxies ?						
<div>NEW</div> <div>ALL FILTERS ▼ Quick filter (Name)</div>						
Name	IP	Status	Version	Accounts	Network Proxy Address	Network Proxy Port
CP_TG	10.192.198.5	Offline	8.6.0.51997631	2 accounts	-	-

To verify the connection, use the following commands: (For the complete list of commands, please see [Using the Cloud Proxy Command-Line Interface](#).)

```
# Overall status of cloud proxy:cprc-cli -s

# Ping itself:
ip addr
ping <address>

# Ping gateway:
ip route
ping <gateway>

# Verify the connection outside the cloud proxy,
ping 8.8.8.8

Note: If you are using a network proxy,
use the /opt/vmware/share/vami/vami_config_net option#5 command
to ensure you have the correct configuration for the testings.
```

The following problems could result in vRealize Operations displaying the status of cloud proxy as offline.

1 Incorrect network proxy information in cloud proxy configuration.

To verify the connection via a network proxy, use the following:

```
curl -vvv --proxy http(s)://proxy_user:proxy_pass@proxy_ip:proxy_port -H 'Accept: application/json' -H 'Content-Type: application/json' -X GET https://<gateway url>/casa/security/ping (gateway url example - 10238.gw.dev.vrops-ops.com)
```

To ignore SSL validation for a proxy server, use `curl --proxy-insecure`. With SSL validation the customer can provide Proxy Server certificate during cloud proxy deployment or re-configuration so that provided certificate from customer can be used to check the connection with curl with SSL certificate validation.

Solution:

- a SSH to the Cloud Proxy VM and set the `connectretry` to 0 in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.
 - b Shutdown the Cloud Proxy VM.
 - c Update the network proxy configurations from the vCenter Server VM options using the vApp options [Edit OVF Details for a Virtual Machine](#).
 - d Boot the Cloud Proxy VM.
- 2 Required ports are not open.

To verify:

```
openssl s_client -showcerts -connect {address}:443

curl -v telnet://{address}:443

# Or, change the address to the machine you want to check:
python -c "import socket; print(socket.socket(socket.AF_INET, socket.SOCK_STREAM).connect_ex(('127.0.0.1', 443)))"

# If you get a !=0 response, the server is not listening to the port.
```

Solution:

- a SSH to the Cloud Proxy VM and set `connectretry` to 0 in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.
- b Provide port access as mentioned in the prerequisite section of [Configuring Cloud Proxies in vRealize Operations](#)
- c Boot the Cloud Proxy VM.

3 Invalid certificate.

To verify:

```
openssl s_client -showcerts -connect {address}:443
```

Solution:

a SSH to the Cloud Proxy VM and set `connectretry` to 0 in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.

b Follow the steps mentioned in VMware KB Article, [83698](#).

4 The logs folder `/storage/log` is running out of partition space.

Solution: Remove log files to ensure that enough space is available. Note that this is an exceptional case. In normal conditions, log files are auto archived.

5 One or more of the following services are down: `httpd-north.service`, `haproxy.service` and `collector.service`.

Solution:

- Check service status by running the following command: `systemctl status <service name>`.

- To start service, use the following command: `systemctl start <service name>`.

6 OTK expired.

Solution: Redeploy Cloud Proxy with fresh OTK.

Cloud proxy is online, and state of Cloud Account is Collecting, but status is Object Down.


vCenter 2 Accounts				
	Name	Status	Description	Colle
<input type="checkbox"/>	CA_TG	Warning		CP
<input type="checkbox"/>	API ADAPTER SAMPLE	State: Collecting Status: Object down Message: Unable to connect to VC	r Adapter Instance	Clo
VMware vRealize Application Management Pack 2 Accounts				

The following problem could result in vRealize Operations displaying the state of Cloud Account as `Collecting`, while the status is, `Object Down`.

1 Incorrect account credentials.

Solution: Check and update the credentials used while setting up the cloud account.

Cloud proxy status is stuck in Going Online.

Cloud Proxies ?						
<div>NEW</div> <div>ALL FILTERS ▾ Quick filter (Name)</div>						
Name	IP	Status	Version	Accounts	Network Proxy Address	Network Proxy Port
CP_TG	10.192.198.5	 Going Online	8.6.0.51997631	2 accounts	-	-

It can take up to 20 mins on first reboot, for the cloud proxy to be registered and come online. Wait for the specified time to see if cloud proxy comes online. If it still does not come online, one or more of the following services are down: `httpd-north.service`, `haproxy.service`, and `collector.service`.

Solution:

- 1 Check service status by running the following command: `systemctl status <service name>`
- 2 To start service, use the following command: `systemctl start <service name>`.

Cloud proxy does not upgrade automatically, after the upgrade of vRealize Operations

There could be a few possible reasons why cloud proxy does not upgrade automatically after an upgrade of vRealize Operations.

- 1 High network latency leading to PAK download failure. Latency of >500ms is not supported.

Solution: See the VMWare KB article [80590](#) on how to manually upgrade cloud proxy via CLI.

- 2 Upgrade status is stuck at `Running` since the previous upgrade had failed.

Solution: Follow the steps given below to change the upgrade status.

- a Stop the casa service: `systemctl stop vmare-casa.service`.
- b Change the upgrade status from `RUNNING` to `NONE` in the following files:

```
./storage/db/vmware-vrops-cprc/status/cprc.upgrade.status
./storage/db/vmware-vrops-cprc/status/cprc.pak.status
```

- c See the VMware KB article [80590](#) and run the manual upgrade.

Cloud proxy gets disconnected at regular intervals

There could be a few possible reasons why cloud proxy gets disconnected at regular intervals.

- 1 Check the network connectivity and latency.
- 2 Check if the cloud proxy VM can reach the DNS and use the `NSlookup` to validate the DNS connectivity.

vRealize Operations Post-Installation Considerations

6

After you install vRealize Operations, there are post-installation tasks that might need your attention.

This chapter includes the following topics:

- [About Logging In to vRealize Operations](#)
- [After You Log In](#)
- [Secure the vRealize Operations Console](#)
- [Log in to a Remote vRealize Operations Console Session](#)
- [About New vRealize Operations Installations](#)

About Logging In to vRealize Operations

Logging in to vRealize Operations requires that you point a Web browser to the fully qualified domain name (FQDN) or IP address of a node in the vRealize Operations cluster.

When you log in to vRealize Operations, there are a few things to keep in mind.

- After initial configuration, the product interface URL is:
`https://node-FQDN-or-IP-address`
- Before initial configuration, the product URL opens the administration interface instead.
- After initial configuration, the administration interface URL is:
`https://node-FQDN-or-IP-address/admin`
- The administrator account name is admin. The account name cannot be changed.
- The admin account is different from the root account used to log in to the console, and does not need to have the same password.
- When logged in to the administration interface, avoid taking the node that you are logged into offline and shutting it down. Otherwise, the interface closes.

- The number of simultaneous login sessions before a performance decrease depends on factors such as the number of nodes in the analytics cluster, the size of those nodes, and the load that each user session expects to put on the system. Heavy users might engage in significant administrative activity, multiple simultaneous dashboards, cluster management tasks, and so on. Light users are more common and often require only one or two dashboards.

The sizing spreadsheet for your version of vRealize Operations contains further detail about simultaneous login support. See [Knowledge Base article 2093783](#).

- You cannot log in to a vRealize Operations interface with user accounts that are internal to vRealize Operations, such as the maintenance Admin account.
- You cannot open the product interface from a remote collector node, but you can open the administration interface.
- For supported Web browsers, see the vRealize Operations Release Notes for your version.

After You Log In

After you log in to vRealize Operations from a web browser, you see the Quick Start page. You can set any dashboard to be the landing page instead of the Quick Start page. Click the **Actions** menu on a dashboard that you want to set as the landing page and select **Set as Home landing page**. To remove the dashboard as the home landing page, click the **Actions** menu on the relevant dashboard and select **Reset from Home landing page**.

The Quick Start page provides an overview of key areas of vRealize Operations.

Quick Start Page Before Cloud Accounts Are Configured

When you log in to vRealize Operations and no cloud accounts are configured, the Quick Start page displays guided tours in the Optimize Performance, Optimize Capacity, Troubleshoot, and Manage Configuration sections. Watch these guided tours to understand how the product functions. If your user account does not have administrative rights, then the Quick Start page prompts you to contact the administrator for configuration of cloud accounts.

If you have logged in using an administrative account, you must set the currency in the **Global Settings** page. From the left menu, click **Administration**, and then click the **Global Settings** tile. You can do so from the message that you see in the Quick Start page when you log in for the first time. Optionally, you can close the message. Once you set a currency, you cannot change it. As an administrator, you must also first set up a cloud account or configure an adapter before you can start using vRealize Operations. Until you do so, you see links to guided tours about vRealize Operations.

A new license key is required for vRealize Operations 7.0 and later versions. All license keys except vSOM Enterprise Plus and its add-ons are invalidated. The product works in evaluation mode until a new valid license key, which can be obtained from the [MyVMware](#) portal, is installed. After login, if you see the "You are using an evaluation license. Please consider applying a new license by the end of the evaluation period." message in the Quick Start page, you must add a new license before the end of the 60-day evaluation period in the Licensing page. To add a new license, from the message, click **Actions > Go to Licensing**.

Note If you added new licenses when you upgraded to vRealize Operations 7.0, you can skip this step.

After logging in, if you see a message like, "vRealize Operations Manager internal certificates will expire on dd/mm/yyyy. Please install a new certificate before the expiry date. For details, see KB 71018" in the Quick Start page, you must upgrade your internal certificates for vRealize Operations using the certificate renewal PAK file from the vRealize Operations Administrator interface. For more information, see the following KB article [71018](#).

Quick Start Page After Cloud Accounts Are Configured

When you log in to vRealize Operations after the cloud accounts or adapter instances are configured, and the initial setup is complete, the Quick Start displays the following sections.

Optimize Performance

Displays links to workload optimization, right sizing, recommendations, and optimization history.

Optimize Capacity

Displays links to assess capacity, reclaim resources, plan scenarios, assess costs and optimize cost.

Troubleshoot

Displays links to the troubleshooting workbench, alerts, logs, dashboards and applications.

Manage Configuration

Displays links to the compliance, configuration of virtual machines, hosts, clusters and distributed switches, and the sustainability dashboards.

The other tiles you can see are:

Extend Monitoring

Displays links to the following VMware website:

- True Visibility Suite
- SDDC Management Health

- vRealize Operations Aggregator
- Explore vRealize Operations REST APIs

Learn and Evaluate

Displays links to the following sites:

- Introduction to vRealize Operations
- Evaluate vRealize Suite
- vRealize Operations Guided Tour
- Additional Learning
- Evaluate Sample Dashboards
- Browse and download code samples from VMware

Run Assessments

Displays shortcut links to the VMware vRealize Cloud Management Assessment and vSphere Optimization Assessment (Deprecated) sites.

Secure the vRealize Operations Console

After you install vRealize Operations, you secure the console of each node in the cluster by logging in for the first time.

Procedure

- 1 Locate the node console in vCenter or by direct access. In vCenter, use Alt+F1 to access the login prompt.

For security, vRealize Operations remote terminal sessions are disabled by default.

- 2 Log in as **root**.

vRealize Operations prevents you from accessing the command prompt until you create a root password.

- 3 When prompted for a password, press Enter.
- 4 When prompted for the old password, press Enter.
- 5 When prompted for the new password, enter the root password that you want, and note it for future reference.
- 6 Re-enter the root password.
- 7 Log out of the console.

Log in to a Remote vRealize Operations Console Session

As part of managing or maintaining the nodes in your vRealize Operations cluster, you might need to log in to a vRealize Operations node through a remote console.

For security, remote login is disabled in vRealize Operations by default. To enable remote login, perform the following steps.

Procedure

- 1 Log in to a vCenter Server system using a vSphere Web Client and select a vCenter Server instance in the vSphere Web Client navigator.

- a Find the **Virtual Machine** in the hierarchy and click **Launch Console**.

Note You can also use the vSphere Client to launch the node console by direct access after enabling the SSHD service.

The virtual machine console opens in a new tab of the Web browser.

- 2 Locate the node console and click **Launch Console**.
- 3 In vCenter, use Alt+F1 to access the login prompt and log in as **root**. If this is the first time logging in, you must set a root password.
 - a When prompted for a password, press Enter.
 - b When prompted for the old password, press Enter.
 - c When prompted for the new password, enter the root password that you want, and note it for future reference.
 - d Re-enter the root password.
- 4 To enable remote login, enter the following command:

```
service sshd start
```

About New vRealize Operations Installations

A new vRealize Operations installation requires that you deploy and configure nodes. Then, you add solutions for the kinds of objects to monitor and manage.

After you add solutions, you configure them in the product and add monitoring policies that gather the kind of data that you want.

Log In and Continue with a New Installation

To finish a new vRealize Operations installation, you log in and complete a one-time process to license the product and configure solutions for the kinds of objects that you want to monitor.

Prerequisites

- Create the new cluster of vRealize Operations nodes.

- Verify that the cluster has enough capacity to monitor your environment. See [Sizing the vRealize Operations Cluster](#).

Procedure

- 1 In a Web browser, navigate to the IP address or fully qualified domain name of the primary node.

- 2 Enter the username **admin** and the password that you defined when you configured the primary node, and click **Login**.

Because this is the first time you are logging in, the administration interface appears.

- 3 To start the cluster, click **Start vRealize Operations Manager**.

- 4 Click **Yes**.

The cluster might take from 10 to 30 minutes to start, depending on your environment. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

- 5 When the cluster finishes starting and the product login page appears, enter the admin username and password again, and click **Login**.

A one-time licensing wizard appears.

- 6 Click **Next**.

- 7 Read and accept the End User License Agreement, and click **Next**.

- 8 Enter your product key, or select the option to run vRealize Operations in evaluation mode.

Your level of product license determines what solutions you may install to monitor and manage objects.

- Standard. vCenter only
- Advanced. vCenter plus other infrastructure solutions
- Enterprise. All solutions

vRealize Operations does not license managed objects in the same way that vSphere does, so there is no object count when you license the product.

Note When you transition to the Standard edition, you no longer have the Advanced and Enterprise features. After the transition, delete any content that you created in the other versions to ensure that you comply with EULA and verify the license key which supports the Advanced and Enterprise features.

- 9 If you entered a product key, click **Validate License Key**.

- 10 Click **Next**.

- 11 Select whether or not to return usage statistics to VMware, and click **Next**.

- 12 Click **Finish**.

The one-time wizard finishes, and the vRealize Operations interface appears.

What to do next

- Use the vRealize Operations interface to configure the solutions that are included with the product.
- Use the vRealize Operations interface to add more solutions.
- Use the vRealize Operations interface to add monitoring policies.

Upgrade, Backup and Restore

7

You can update your existing vRealize Operations deployments to a newly released version.

When you perform a software update, you need to make sure you use the correct PAK file for your cluster. A good practice is to take a snapshot of the cluster before you update the software, but you must remember to delete the snapshot once the update is complete.

If you have customized the content that vRealize Operations provides such as alerts, symptoms, recommendations, and policies, and you want to install content updates, clone the content before performing the update. In this way, you can select the option to reset out-of-the-box content when you install the software update, and the update can provide new content without overwriting customized content.

Starting with version 8.6 of vRealize Operations, internal certificates are renewed when you upgrade a cluster, except when the cloud proxy version 8.4, 8.5, or earlier is present. Automatic root-CA certificate renewal will be available when cloud proxy is version 8.6 and is upgraded to higher versions. After each product upgrade, the cluster will have a new root-CA certificate with a 5-year validity period.

Note Automatic certificate renewal does not affect custom certificates.

This chapter includes the following topics:

- [Obtain the Software Update PAK File](#)
- [Create a Snapshot as Part of an Update](#)
- [How To Preserve Customized Content](#)
- [Back Up and Restore](#)
- [vRealize Operations Software Updates](#)
- [Before Upgrading to vRealize Operations 8.6](#)

Obtain the Software Update PAK File

Each type of cluster update requires a specific PAK file. Make sure you are using the correct one.

Download the Correct PAK files

To update your vRealize Operations environment, you need to download the right PAK file for the clusters you wish to upgrade. In case modifications are required, you can manually update the hosts file after completing the software update.

To download the PAK file for vRealize Operations, go to [Download VMware vRealize Operations](#) page and select the correct version from the drop-down list.

If you are using cloud proxy, download the *vRealize Operations Manager - Virtual Appliance upgrade .pak file with Cloud Proxy* file from the Product Downloads tab, to update the vRealize Operations environment and your cloud proxy together.

Create a Snapshot as Part of an Update

It is mandatory to create a snapshot of each node in a cluster before you update a vRealize Operations cluster. Once the update is complete, you must delete the snapshot to avoid performance degradation.

For more information about snapshots, see the vSphere Virtual Machine Administration documentation.

Procedure

- 1 Log into the vRealize Operations Administrator interface at `https://<master-node-FQDN-or-IP-address>/admin`.
- 2 Click **Take Offline** under the cluster status.
- 3 When all nodes are offline, open the vSphere client.
- 4 Right-click a vRealize Operations virtual machine.
- 5 Click **Snapshot** and then click **Take Snapshot**.
 - a Name the snapshot. Use a meaningful name such as "Pre-Update."
 - b Uncheck the **Snapshot the Virtual Machine Memory** check box.
 - c Uncheck the **Ensure Quiesce Guest File System (Needs VMware Tools installed)** check box.
 - d Click **OK**.
- 6 Repeat these steps for each node in the cluster.

What to do next

Start the update process as described in [Install a Software Update](#).

How To Preserve Customized Content

When you upgrade vRealize Operations, it is important that you upgrade the current versions of content types that allow you to alert on and monitor the objects in your environment. With upgraded alert definitions, symptom definitions, and recommendations, you can alert on the various states of objects in your environment and identify a wider range of problem types. With upgraded views, you can create dashboards and reports to easily identify and report on problems in your environment.

You might need to perform certain steps before you upgrade the alert definitions, symptom definitions, recommendations, and views in your vRealize Operations environment.

- If you customized any of the alert definitions, symptom definitions, recommendations, or views that were provided with previous versions of vRealize Operations, and you want to retain those customized versions, perform the steps in this procedure.
- If you did not customize any of the alert definitions, symptom definitions, recommendations, or views that were provided with previous versions of vRealize Operations, you do not need to back them up first. Instead, you can start the upgrade, and during the upgrade select the check box named **Reset out-of-the-box content**.

Prerequisites

You previously customized versions of your alert definitions, symptom definitions, recommendations, or views.

Procedure

- 1 Before you begin the upgrade to vRealize Operations, back up the changes to your alert definitions, symptom definitions, recommendations, and views by cloning them.
- 2 Start the upgrade of vRealize Operations.
- 3 During the upgrade, select the check box named **Reset out-of-the-box content**.

Results

After the upgrade completes, you have preserved your customized versions of alert definitions, symptom definitions, recommendations, and views, and you have the current versions that were installed during the upgrade.

What to do next

Review the changes in the upgraded alert definitions, symptom definitions, recommendations, and views. Then, determine whether to keep your previously modified versions, or to use the upgraded versions. For more information, see *Creating a Backup and Importing Content* in the *Managing Content* chapter of the *Configuration Guide*.

Back Up and Restore

Back up and restore your vRealize Operations system regularly to avoid downtime and data loss in case of a system failure. If your system does fail, you can restore the system to the last full or incremental backup.

You can back up and restore vRealize Operations single or multi-node clusters by using vSphere Data Protection or other backup tools. You can perform full, differential, and incremental backups and restores of virtual machines.

To back up and restore vRealize Suite components by using vSphere Data Protection and NetBackup, see [Create a Backup Policy for vRealize Suite with NetBackup](#).

To back up and restore vRealize Suite single or multi-node clusters using EMC Avamar and to perform on-demand group backup, see [vRealize Suite Backup and Restore by Using EMC Avamar](#).

To back up and restore vRealize Operations single or multi-node clusters using the Veeam Backup & Replication tool, see [About Veeam Backup & Replication](#).

It is highly recommended to take a backup during quiet periods. Since a snapshot based backup happens at the block level, it is important that there are limited or no changes being performed by a user on the cluster configuration. This will ensure that you have a healthy backup.

It is best to take the cluster offline before you back up the vRealize Operations nodes. This will ensure the data consistency across the nodes and internally in the node. You can either shut down the VM before the backup or enable quiescing.

If the cluster remains online, backup your vRealize Operations multi-node cluster by using vSphere Data Protection or other backup tools, disable quiescing of the file system.

Note All nodes are backed up and restored at the same time. You cannot back up and restore individual nodes.

You can use the Site Recovery Manager to protect your vRealize Suite components. The vRealize Suite Disaster Recovery by Using Site Recovery Manager is a disaster recovery automation software that provides policy-based management, non-disruptive testing, and automated orchestration. For more information, see [vRealize Suite Disaster Recovery by Using Site Recovery Manager](#).

vRealize Operations Software Updates

vRealize Operations includes a central page where you can manage updates to the product software.

How Software Updates Work

The Software Update option lets you install updates to the vRealize Operations product itself.

Where You Find Software Updates

Log in to the vRealize Operations administration interface at <https://master-node-name-or-ip-address/admin>. On the left, click **Software Update**.

Software Update Options

The options include a wizard for locating the update PAK file and starting the installation, plus a list of updates and the vRealize Operations cluster nodes on which they are installed.

Table 7-1. Software Update Options

Option	Description
Install a Software Update	Launch a wizard that allows you to locate, accept the license, and start the installation of a vRealize Operations software update.
Node Name	Machine name of the node where the update is installed
Node IP Address	Internet protocol (IP) address of the node where the update is installed. Primary and replica nodes require static IP addresses. Data nodes may use DHCP or static IP.
Update Step	Software update progress in step x of y format
Status	<p>Success, failure, in-progress, or unknown condition of the software update.</p> <p>For cloud proxy upgrade, every stage of the upgrade process is displayed. Hover the mouse near the status message to see more details in the pop-up window. The Cloud Proxy upgrade stages are as follows:</p> <ul style="list-style-type: none">■ Stage 1 - Downloading■ Stage 2: Extracting■ Stage 3: Upgrading■ Stage 4: Rebooting■ Stage 5: Success

Install a Software Update

If you have already installed vRealize Operations, you can update your software when a newer version becomes available.

Note Installation might take several minutes or even a couple hours depending on the size and type of your clusters and nodes.

Note vRealize Application Remote Collector virtual appliance is deprecated and is no longer available for download from the vRealize Operations user interface when you upgrade to vRealize Operations 8.6. VMware recommends that you use cloud proxy to monitor your application services. You can migrate on-prem standalone vRealize Application Remote Collector to on-prem cloud proxy. For information about migrating from vRealize Application Remote Collector to cloud proxy, see [KB 83059](#).

Prerequisites

- Create a snapshot of each node in your cluster. For information about how to perform this task, see the vRealize Operations Information Center.
- Obtain the PAK file for your cluster. For information about which file to use, see the vRealize Operations Information Center.
- Before you install the PAK file, or upgrade your vRealize Operations instance, clone any customized content to preserve it. Customized content can include alert definitions, symptom definitions, recommendations, and views. Then, during the software update, you select the options named **Install the PAK file even if it is already installed** and **Reset out-of-the-box content**.
- Since version 6.2.1, vRealize Operations update operation has a validation process that identifies issues before you start to update your software. Although it is good practice to run the pre-update check and resolve any issues found, users who have environmental constraints can disable this validation check.

To disable the pre-update validation check, perform the following steps:

- Edit the update file to `/storage/db/pakRepoLocal/bypass_prechecks_vRealizeOperationsManagerEnterprise-buildnumberofupdate.json`.
- Change the value to TRUE and run the update.

Note If you disable the validation, you might encounter blocking failures during the update itself.

Procedure

- 1 Log into the master node vRealize Operations administrator interface of your cluster at `https://master-node-FQDN-or-IP-address/admin`.

- 2 Click **Software Update** in the left pane.
- 3 Click **Install a Software Update** in the main pane.
- 4 Follow the steps in the wizard to locate and install your PAK file.

This updates the OS on the virtual appliance and restarts each virtual machine.
- 5 Read the **End User License Agreement** and **Update Information**, and click **Next**.
- 6 Click **Install** to complete the installation of software update.

Note After you click **Install**, the installer will restart the vRealize Operations administrator interface, and you will be logged out. Log in once again to the vRealize Operations administrator interface when it is ready, and follow the update status in the software update page.

- 7 Log back into the master node administrator interface.

The main Cluster Status page appears and cluster goes online automatically. The status page also displays the Bring Online button, but do not click it.
- 8 Clear the browser caches and if the browser page does not refresh automatically, refresh the page.

The cluster status changes to Going Online. When the cluster status changes to Online, the upgrade is complete.

Note If a cluster fails and the status changes to offline during the installation process of a PAK file update, then some nodes become unavailable. To fix this, you can access the administrator interface and manually take the cluster offline and click **Finish Installation** to continue the installation process.

- 9 Click **Software Update** to check that the update is done.

A message indicating that the update completed successfully appears in the main pane.

Note When you update vRealize Operations to a latest version, all nodes get upgraded by default.

If you are using cloud proxies, the cloud proxy upgrades start after the vRealize Operations upgrade is complete successfully. For more information, see the Monitoring the Health of Cloud Proxies from the Admin UI topic in the *vRealize Operations Configuration Guide*.

What to do next

Delete the snapshots you made before the software update.

Note Multiple snapshots can degrade performance, so delete your pre-update snapshots after the software update completes.

Before Upgrading to vRealize Operations 8.6

With every vRealize Operations release, many metrics are either discontinued or disabled. These changes update the capacity analytics and improve the product scale. VMware has made many of these changes transparent or nearly so. Still, multiple changes can impact management packs that you might be using, along with the dashboards and reports that you have created. Therefore, before upgrading, run the vRealize Operations Pre-upgrade Readiness Assessment Tool (Assessment Tool) that helps you understand the precise impact on your environment through a detailed report.

Why Run the Assessment Tool

Various changes in vRealize Operations can impact the user experience. When you run the Assessment Tool, you get an HTML-formatted report identifying all the points in your system affected by the changes. Further, the Assessment Tool gives recommendations for the correct changes to be made in your content for when you upgrade from a previous release.

Note You must run the Assessment Tool on the instance of the vRealize Operations installation that you want to assess - typically your production system. The Assessment Tool does not alter anything in your system, and deletes itself when it has completed its run. It leaves behind only the assessment result - a support bundle that you download from the Support Bundles section of the vRealize Operations Administration user interface.

The Assessment tool validates your environment to ensure it is ready for the upgrade. For example, if the ESXi version does not match the product requirements, the assessment tool will identify the issue and provide you with a recommendation in the Systems Validation tab.

For detailed instructions on running the Assessment Tool, see [Running the vRealize Operations 8.6 Pre-Upgrade Readiness Assessment Tool](#).

To view the upgrade path from an earlier version of vRealize Operations to 8.6, see [vRealize Operations Upgrade Path](#).

Running the vRealize Operations 8.6 Pre-Upgrade Readiness Assessment Tool

Before upgrading, you can gauge the impact on your system by running the vRealize Operations Pre-Upgrade Readiness Assessment Tool (Assessment Tool). The tool generates a report detailing the precise impact on your environment and gives suggestions for replacement metrics.

Using the Assessment Tool consists of four distinct steps:

- 1 Download the PAK file from <https://my.vmware.com/group/vmware/get-download?downloadGroup=VROPS-860>.
- 2 Run the vRealize Operations Pre-Upgrade Readiness Assessment Tool.
- 3 Extract the report from the generated ZIP file.

- 4 Click the various items in the report to link to the solutions grid.

Note You must run the Assessment Tool on the instance of the vRealize Operations installation that you want to assess - typically your production system. The Assessment Tool does not alter anything in your system, and deletes itself when it has completed its run. It leaves behind only the assessment result - a support bundle that you download from the Support Bundles section of the vRealize Operations Administration user interface.

Prerequisites

You must have administrator privileges in your current installation of vRealize Operations to download and run the Assessment Tool. For more information on using the upgrade assessment tool, see the following KB article [67311](#).

Procedure

- 1 Download the Assessment Tool PAK from <https://my.vmware.com/group/vmware/get-download?downloadGroup=VROPS-860> to your local machine. Search for APUAT or vRealize Operations - Upgrade Assessment Tool.
- 2 Open a browser and navigate to the vRealize Operations administrator console: `https://<master_node_IP>/admin`.

Then log into the administrator user interface with the user ID **admin** and the associated password.
- 3 In the left pane of the administration home page, click **Software Update**.

The Software Update screen appears.
- 4 Click **Install a Software Update** at the top of the screen.

The Add Software Update workspace appears.
- 5 Click the **Browse** link and navigate to the PAK file you downloaded in Step 1.

A check mark appears next to the statement: **The selected file is ready to upload and install. Click UPLOAD to continue.**
- 6 Ensure that a check mark appears next to the statement: **Install the PAK file even if it is already installed.**

Leave blank the check box next to Reset Default Content...
- 7 Click the **UPLOAD** link.

The PAK file is uploaded from your local machine to vRealize Operations. Uploading may take a few minutes.
- 8 Once the PAK file is uploaded, click **NEXT**.

The End User License Agreement appears.

- 9 Click the check box next to the statement: **I accept the terms of this agreement.**

Click **NEXT**. The Important Update and Release Information screen appears.

- 10 Review the release information and click **NEXT**. At the Install Software Update screen, click **INSTALL**.

The Software Update screen appears again, this time with a rotating icon and an **installation in progress...** bar marking the progress of the PAK file and assessment as they run on your environment. The process can take from five to 20 minutes, depending on the size of your system.

- 11 When the process is complete, click **Support** in the left pane.

The Support screen appears.

- 12 Select the **Support Bundles** option above the toolbar.

The available support bundles are listed.

- 13 Locate the support bundle most recently created. Click the chevron next to the bundle name to open the file and select it, then click the download link on the toolbar to save the support bundle ZIP file to your local files.

- 14 To review the report, extract the files from the ZIP file and open the HTML file. (Do not open the CSV file, it is for VMware use only.)

The report is a graphical depiction of your vRealize Operations UI components - dashboards, reports, management packs, alerts, heat maps, and so on - and includes the number of deprecated metrics impacting each component. For example, you might find that 10 of your 25 dashboards contain a total of 15 deprecated metrics.

- 15 Click a component.

The report details for that component are listed following the graphics, under Impacted Component Details. Taking dashboards as an example, the list provides - for each dashboard - the dashboard name, owner, widgets removed, metric-impacted views, and metric-impacted widgets. The deprecated metrics are live links.

- 16 Click a live metric link.

A browser window opens at URL <http://partnerweb.vmware.com/programs/vrops/DeprecatedContent.html> with the selected metric highlighted in a table of like metrics. If a replacement metric is available for the deprecated metric, it is listed in the same row by name and metric key. You might choose to install the new metric in place of the deprecated metric.

- 17 Repeat Steps 15 and 16 for all your components.

If you replace the deprecated metrics with new metrics, or update each component to provide needed information without the deprecated metrics, your system is ready for the upgrade.

- 18 Rerun the entire assessment process from Step 1 to confirm that your system is no longer impacted or at least mostly not impacted by the metrics changes.

- 19 Once you have upgraded to vRealize Operations 8.6, fix the remaining issues with replacement metrics available in the new release.

Results

Your vRealize Operations components are updated to work correctly in the 8.6 release.

What to do next

Once you have installed vRealize Operations 8.6, conduct, at a minimum, random testing to determine if system metrics are operating as you expect. Monitor the platform on an ongoing basis to confirm that you are receiving the correct data.