

vRealize Operations Manager Installation and Configuration Guide for Linux

vRealize Operations Manager 6.5



vmware®

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

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

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

docfeedback@vmware.com

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

Copyright © 2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About Installation and Configuration for Linux 5

1 About Installing 6

Installation Overview 6

Workflow of vRealize Operations Manager Installation 6

Sizing the Cluster 7

 Add Data Disk Space to a Linux Node 8

Complexity of Your Environment 9

Cluster Nodes 11

About Remote Collector Nodes 12

About High Availability 12

2 Preparing for Installation 15

Platform requirements for vRealize Operations Manager 15

 Platform Requirements for Linux 15

Requirements 17

 Requirements for IPv6 17

 Cluster Requirements 18

3 Installing vRealize Operations Manager 24

Deployment of vRealize Operations Manager 24

 Create a vRealize Operations Node using Installers 24

Installation Types 27

 Installing vRealize Operations Manager for a New User 27

 Installing vRealize Operations Manager as an Administrator 30

 Expand an Existing Installation of vRealize Operations Manager 31

4 Resize your Cluster by Adding Nodes 34

Gathering More Data by Adding a Remote Collector Node 35

 Run the Setup Wizard to Create a Remote Collector Node 35

Adding High Availability 36

 Run the Setup Wizard to Add a Master Replica Node 36

Cluster and Node Maintenance 38

 Cluster Management 39

5 Post-Installation Considerations 41

About Logging In 41

About New Installations 42

[Log In and Continue with a New Installation](#) 42

6 [Updating, Migrating and Restoring](#) 44

[Obtain the Software Update PAK File](#) 44

[Create a Snapshot as Part of an Update](#) 45

[How To Preserve Customized Content](#) 45

[Backup and Restore](#) 46

[Software Updates](#) 47

[Install a Software Update](#) 47

[Install a vRealize Operations Manager Software Update from the Administration Interface](#) 49

[Migrate a vCenter Operations Manager Deployment into this Version](#) 49

7 [Uninstalling](#) 51

[Uninstallation from Linux](#) 51

About Installation and Configuration for Linux

The *vRealize Operations Manager Installation and Configuration Guide for Linux* provides information about installing VMware® vRealize Operations Manager on the Linux operating system, including how to create and configure the vRealize Operations Manager cluster.

The vRealize Operations Manager installation process consists of running the vRealize Operations Manager Enterprise installer on 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 on Linux machines. The information is written for experienced Linux system administrators who are familiar with enterprise management applications and datacenter operations.

VMware Technical Publications Glossary

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

About Installing

When you install vRealize Operations Manager, you can install the product in an environment that has never been monitored by vRealize Operations Manager. You can also migrate, which captures an environment monitored by a previous version of vRealize Operations Manager so that the new copy of vRealize Operations Manager can monitor that environment.

You can migrate at installation time, or you can postpone a migration until after your copy of vRealize Operations Manager is in production use. In other words, you can run vRealize Operations Manager to monitor a fresh environment, and at any time, decide to add an environment that was being monitored by a previous vRealize Operations Manager.

This chapter includes the following topics:

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

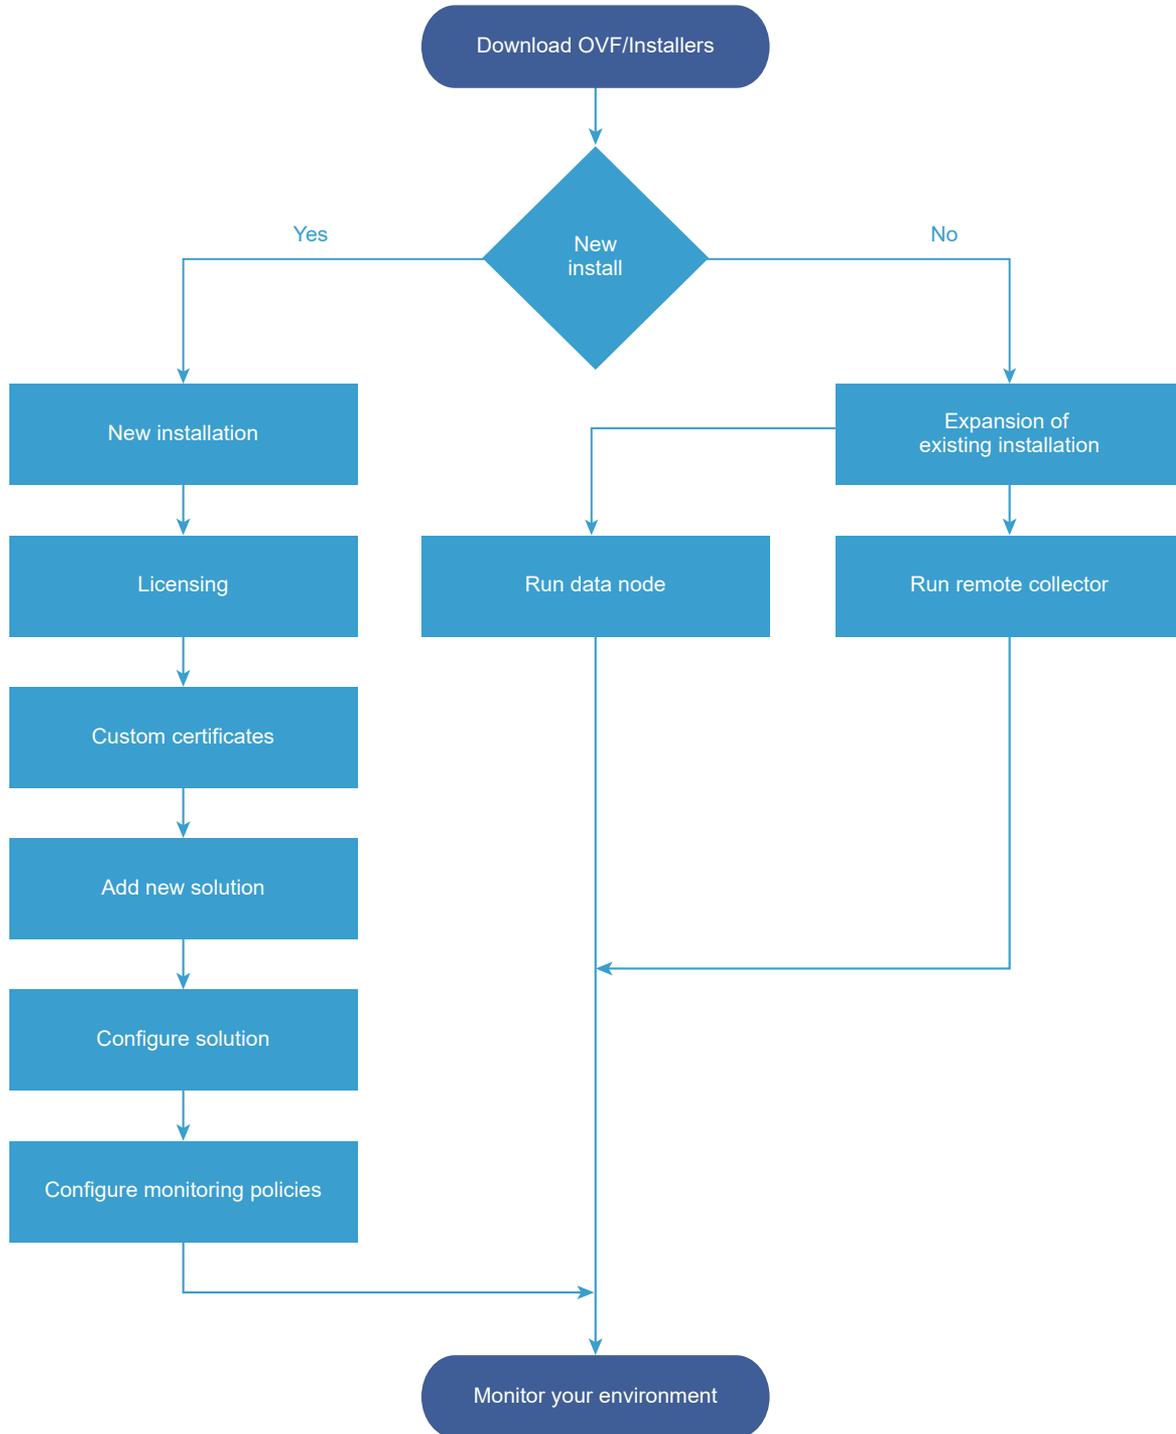
Installation Overview

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

Workflow of vRealize Operations Manager Installation

The vRealize Operations Manager virtual appliance installation process consists of deploying the vRealize Operations Manager OVF or installer 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



Sizing the vRealize Operations Manager Cluster

The resources needed for vRealize Operations Manager 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 Manager 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 Manager.

[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 Manager itself.

- 1 Review this guide to understand how to deploy and configure a vRealize Operations Manager node.
- 2 Deploy a temporary vRealize Operations Manager 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 Manager 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.

Add Data Disk Space to a vRealize Operations Manager Linux Node

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

The following example is for a Linux system.

Prerequisites

Note the disk size of the analytics cluster nodes. When adding disk, you must maintain uniform size across analytics cluster nodes.

Procedure

- 1 Add a new disk to the system, and partition and format the disk as needed.
- 2 Use the vRealize Operations Manager administration interface to take the cluster offline.
- 3 Stop the `vmware-casa` service.
- 4 Move the contents of `/storage/db` into a directory on the new disk.
- 5 Create a symbolic link from the new directory back to `/storage/db`, so that `/storage/db` now references the new disk.
- 6 Start the `vmware-casa` service.
- 7 Bring the cluster online.

Complexity of Your Environment

When you deploy vRealize Operations Manager, 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 Manager and in partnership with Professional Services, must be taken into account when deploying vRealize Operations Manager.

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 Manager deployment. | Lone instances are usually easy to create in vRealize Operations Manager. |
| Green | Your deployment includes a management pack that is listed as Green according to the compatibility guide on the VMware Solutions Exchange Web site. | The compatibility guide indicates whether the supported management pack for vRealize Operations Manager 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. Note that the terms <i>solution</i> , <i>management pack</i> , <i>adapter</i> , and <i>plug-in</i> are used somewhat interchangeably. |
| Yellow | You run multiple instances of vRealize Operations Manager. | 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 Manager 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 Manager 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 Manager cluster. | Multiple nodes are typically used for scaling out the monitoring capability of vRealize Operations Manager. |
| Yellow | Your new vRealize Operations Manager 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 Manager 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. |
| Yellow | You want help in understanding the new or changed features in vRealize Operations Manager and how to use them in your environment. | vRealize Operations Manager is different than vCenter Operations Manager in areas such as policies, alerts, compliance, custom reporting, or badges. In addition, vRealize Operations Manager uses one consolidated interface. |
| Red | You run multiple instances of vRealize Operations Manager, 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. |

| Complexity Level | Current or New Deployment Condition | Additional Notes |
|------------------|--|---|
| 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 Manager 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 Manager clusters. | Multiple clusters are typically used to isolate business operations or functions. |
| Red | Your current vRealize Operations Manager 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 Manager 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 Manager Cluster Nodes

All vRealize Operations Manager clusters consist of a master node, an optional replica node for high availability, optional data nodes, and optional remote collector nodes.

When you install vRealize Operations Manager, you use a vRealize Operations Manager vApp deployment or Linux installer 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 Manager to monitor an environment as the environment grows larger.

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

Master Node

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

In a single-node installation, the master 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 master and replica node resources can be dedicated to cluster management.

Replica Node

To use vRealize Operations Manager high availability (HA), the cluster requires that you convert a data node into a replica of the master node.

The following node type is a member of the vRealize Operations Manager 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 bandwidth across data centers, or reduce the load on the vRealize Operations Manager analytics cluster. Remote collectors only gather objects for the inventory, without storing data or performing analysis. In addition, remote collector nodes may be installed on a different operating system than the rest of the cluster.

About vRealize Operations Manager Remote Collector Nodes

A remote collector node is an additional cluster node that allows vRealize Operations Manager to gather more objects into its inventory for monitoring. Unlike data nodes, remote collector nodes only include the collector role of vRealize Operations Manager, without storing data or processing any analytics functions.

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 Manager analytics cluster.

Remote collectors do not buffer data while the network is experiencing a problem. If the connection between remote collector and 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 Manager does not retroactively incorporate associated events from that time into any monitoring or analysis.

You must have at least a master node before adding remote collector nodes.

About vRealize Operations Manager High Availability

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

With HA, data stored on the master 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 master 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 master provides, were the master to fail for any reason. If the master fails, failover to the replica is automatic and requires only two to three minutes of vRealize Operations Manager downtime to resume operations and restart data collection.

When a master node problem causes failover, the replica node becomes the master 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 master node. When a master node exits an HA-enabled cluster, master 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.
- Return to HA mode by converting a data node into a new replica node and then removing the old, failed master node. Removed master nodes cannot be repaired and re-added to vRealize Operations Manager.
- Change to non-HA operation by disabling HA and then removing the old, failed master node. Removed master nodes cannot be repaired and re-added to vRealize Operations Manager.
- In the administration interface, after an HA replica node takes over and becomes the new master node, you cannot remove the previous, offline master node from the cluster. In addition, the previous node continues to be listed as a master 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 master/replica nodes, or two or more data nodes, is not supported. Instead, vRealize Operations Manager HA provides additional application level data protection to ensure application level availability.
- When HA is enabled, it lowers vRealize Operations Manager capacity and processing by half, because HA creates a redundant copy of data throughout the cluster, as well as the replica backup of the master node. Consider your potential use of HA when planning the number and size of your vRealize Operations Manager cluster nodes. See [Sizing the vRealize Operations Manager 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 would cause the loss of more than one node, which is not supported, and all of vRealize Operations Manager would become unavailable.

The opposite is also true. Without HA, you could keep nodes on the same host, and it would not make a difference. Without HA, the loss of even one node would make all of vRealize Operations Manager 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 Explorer 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 explorer displays a different quantity of objects from that shown in the Administration UI.



Creating a Replica Node for High Availability

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vrops_create_replica_node_ha)

Preparing for Installation

You preparing for your installation, consider some of these best practises, platform, and cluster requirements.

This chapter includes the following topics:

- [Platform requirements for vRealize Operations Manager](#)
- [Requirements](#)

Platform requirements for vRealize Operations Manager

vRealize Operations Manager requires the following hardware and software when you install on any platform.

vRealize Operations Manager Platform Requirements for Linux

vRealize Operations Manager requires the following hardware and software when you install on Linux.

CPU and Memory Requirements

vRealize Operations Manager is supported for installation with the following CPU and memory.

Table 2-1. vRealize Operations Manager Linux Virtual CPU and Memory Requirements

| Node Size | Virtual CPU and Memory |
|---------------------------|------------------------|
| Small | 4 vCPU |
| | 16 GB vRAM |
| Medium | 8 vCPU |
| | 32 GB vRAM |
| Large | 16 vCPU |
| | 48 GB vRAM |
| Standard Remote Collector | 2 vCPU |
| | 4 GB vRAM |
| Large Remote Collector | 4 vCPU |
| | 16 GB vRAM |

Disk Requirements

Disk space for vRealize Operations Manager is not driven solely by how much space the application needs in order to successfully install. In addition, you must consider data collection and retention requirements, which might vary from site to site.

See [Sizing the vRealize Operations Manager Cluster](#).

The default disk requirement for a new, single-node cluster is 250 GB. Thereafter, one approach to prevent disk capacity shortages is by using vRealize Operations Manager for self monitoring and by adding disk or data nodes as needed.

Software Version Requirements

vRealize Operations Manager is supported for installation on the following Linux versions.

- Red Hat Enterprise Linux (RHEL) 6, starting with version 6.5.

Required Linux Packages for vRealize Operations Manager

vRealize Operations Manager requires that certain Linux packages be installed before running the product installer. Also, vRealize Operations Manager installs additional packages.

Prerequisite Linux Packages

The following packages must be present before running the vRealize Operations Manager installer. Furthermore, if a package is a Linux default, it must not be removed after installation.

- bash
- chkconfig
- coreutils
- db4
- expat
- glibc
- initscripts
- libaio
- libselinux
- libstdc++
- libuuid
- mailcap
- openldap
- pcre
- python
- sudo

- redhat-logos
- rpm-libs
- shadow-utils
- zlib

Packages that vRealize Operations Manager Installs

vRealize Operations Manager installs its own copies of the following packages.

- apr
- apr-util
- apr-util-ldap
- httpd
- httpd-tools
- mod_ssl
- openssl
- python
- VMware-Postgres-libs
- VMware-Postgres-osslibs
- VMware-Postgres-osslibs-server
- VMware-Postgres-server

Requirements

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

Using IPv6 with vRealize Operations Manager

vRealize Operations Manager supports Internet Protocol version 6 (IPv6), the network addressing convention that will eventually replace IPv4. Use of IPv6 with vRealize Operations Manager requires that certain limitations be observed.

Using IPv6

- All vRealize Operations Manager cluster nodes, including remote collectors, must have IPv6 addresses. Do not mix IPv6 and IPv4.
- All vRealize Operations Manager cluster nodes, including remote collectors, must be vApp or Linux based.
- Use global IPv6 addresses only. Link-local addresses are not supported.
- If any nodes use DHCP, your DHCP server must be configured to support IPv6.

- DHCP is only supported on data nodes and remote collectors. Master nodes and replica nodes still require fixed addresses, which is true for IPv4 as well.
- Your DNS server must be configured to support IPv6.
- When adding nodes to the cluster, remember to enter the IPv6 address of the master node.
- When registering a VMware vCenter® instance within vRealize Operations Manager, 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 that, even when vRealize Operations Manager is using IPv6, vCenter Server may still have an IPv4 address. In that case, vRealize Operations Manager does not need the square brackets.

- You cannot register an Endpoint Operations Management agent in an environment that supports both IPv4 and IPv6. In the event that you attempt to do so, the following error appears:

```
Connection failed. Server may be down (or wrong IP/port were used). Waiting for 10 seconds before
retrying.
```

Cluster Requirements

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

General vRealize Operations Manager Cluster Node Requirements

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

General Requirements

- vRealize Operations Manager version. All nodes must run the same vRealize Operations Manager version.

For example, do not add a version 6.1 data node to a cluster of vRealize Operations Manager 6.2 nodes.

- Analytics Cluster Deployment Type. In the analytics cluster, all nodes must be the same kind of deployment: vApp or Linux.

Do not mix vApp, Linux nodes in the same analytics cluster.

- 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
- Linux analytics cluster
- Analytics Cluster Node Sizing. In the analytics cluster, CPU, memory, and disk size must be identical for all nodes.

Master, 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.
- **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.

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

For example, going outside of vRealize Operations Manager 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 Manager High 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 Manager itself. For example, vRealize Operations Manager 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 somewhat interchangeably.

How vRealize Operations Manager Uses Network Ports

vRealize Operations Manager uses network ports to communicate with a VMware vCenter Server system and vRealize Operations Manager components.

In Linux deployments, you must manually verify or configure ports.

Important vRealize Operations Manager does not support the customization of server ports.

Network Ports

Configure firewalls so that the following ports are open for bidirectional traffic.

Table 2-2. Network Port Access Requirements for vRealize Operations Manager

| Port Number | Description |
|-------------|---|
| 22 (TCP) | Used for SSH access to the vRealize Operations Manager cluster. |
| 80 (TCP) | Redirects to port 443. |
| 123 (UDP) | Used by vRealize Operations Manager for Network Time Protocol (NTP) synchronization to the master node. |

| Port Number | Description |
|---------------------------|--|
| 443 (TCP) | Used to access the vRealize Operations Manager product user interface and the vRealize Operations Manager administrator interface. |
| 10443 (TCP) | Used by vRealize Operations Manager to communicate with the vCenter Server Inventory service. |
| 3091–3094 (TCP) | When Horizon View (V4V) is installed, used to access data for vRealize Operations Manager from V4V. |
| 5433 (TCP) | When high availability is enabled, used by the master and replica nodes to replicate the global database. |
| 6061 (TCP) | Used by clients to connect to the GemFire Locator to get connection information to servers in the distributed system. Also monitors server load to send clients to the least-loaded servers. |
| 7001 (TCP) | Used by Cassandra for secure inter-node cluster communication. |
| 9042 (TCP) | Used by Cassandra for secure client related communication amongst nodes. |
| 10000–10010 (TCP and UDP) | GemFire Server ephemeral port range used for unicast UDP messaging and for TCP failure detection in the peer-to-peer distributed system. |
| 20000–20010 (TCP and UDP) | GemFire Locator ephemeral port range used for unicast UDP messaging and for TCP failure detection in the peer-to-peer distributed system. |

Localhost Ports

Verify that your port configuration allows localhost access to the following ports. You may restrict off-host access to these ports if site policies are a concern.

Table 2-3. Localhost Port Access Requirements for vRealize Operations Manager

| Port Number | Description |
|-------------|--|
| 1099 | GemFire Locator Java Management Extensions (JMX) Manager |
| 9004 | Analytics JMX Manager |
| 9008 | Cassandra database JMX Manager |
| 9160 | Cassandra Thrift client port |

vRealize Operations Manager Cluster Node Networking Requirements

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

Networking Requirements

Important vRealize Operations Manager 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 master and replica nodes must use 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.
- One-way latency 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.

vRealize Operations Manager Cluster Node Best Practices

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

Best Practices

- Deploy vRealize Operations Manager analytics cluster nodes in the same vSphere cluster in a single datacenter 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 Manager analytics cluster nodes need guaranteed resources when running at scale. The vRealize Operations Manager 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 Master, 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 master node role.

- The NUMA placement is removed in the vRealize Operations Manager 6.3 and later. Procedures related to NUMA settings from the OVA file follow:

Table 2-4. NUMA Setting

| Action | Description |
|---|---|
| Set the vRealize Operations Manager cluster status to offline | <ol style="list-style-type: none"> 1 Shut down the vRealize Operations Manager 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.

Installing vRealize Operations Manager

3

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

This chapter includes the following topics:

- [Deployment of vRealize Operations Manager](#)
- [Installation Types](#)

Deployment of vRealize Operations Manager

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

In general, there are two ways to install vRealize Operations Manager product.

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

Installers vRealize Operations Manager consists of one or more nodes, in a cluster. To create nodes, you download and run the vRealize Operations Manager Enterprise installer for Linux.

Create a vRealize Operations Node using Installers

You can create one or more nodes to form a cluster by installing the vRealize Operations Manager installers depending on the type of operating environment.

Create a Node by Running the vRealize Operations Manager Linux Installer

vRealize Operations Manager consists of one or more nodes, in a cluster. To create nodes, you download and run the vRealize Operations Manager Enterprise installer for Linux.

Prerequisites

- Plan to use the system only as a vRealize Operations Manager node. Do not host other applications on the same machine.
- Verify that vRealize Operations Manager ports are open at the firewall. See [How vRealize Operations Manager Uses Network Ports](#).

- Verify that prerequisite packages are installed. See [Required Linux Packages for vRealize Operations Manager](#) .
- If this node is to be the master node, reserve a static IP address for the virtual machine, and know the associated domain name server, 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 replica node, reserve a static IP address for the virtual machine, and know the associated domain name server, default gateway, and network mask values.

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

In addition, familiarize yourself with HA node placement as described in [About vRealize Operations Manager High Availability](#) .
- Preplan your domain and machine naming so that the Linux machine name will begin and end with 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.
- Preplan node placement and networking to meet the requirements described in [General vRealize Operations Manager Cluster Node Requirements](#) and [vRealize Operations Manager Cluster Node Networking Requirements](#).
- If you expect the vRealize Operations Manager cluster to use IPv6 addresses, review the IPv6 limitations described in [Using IPv6 with vRealize Operations Manager](#).
- Be aware that vRealize Operations Manager uninstalls `httpd` if it is installed, because vRealize Operations Manager installs its version of Apache.

If vRealize Operations Manager uninstalls `httpd`, it backs up the `/etc/httpd` configuration directory.
- Uninstall any existing copies of PostgreSQL, and remove PostgreSQL directories and data.

vRealize Operations Manager must install its own copy of PostgreSQL.
- Verify that all machines in the file `ntp.conf` are resolvable. If you are unsure about the contents of `ntp.conf`, make a backup copy of the file, and overwrite the original with the default version from a new machine installation.
- Locate your copy of the vRealize Operations Manager Enterprise bin installer for Linux.

Procedure

- 1 Log in with an account that has root privileges.

2 Turn off the firewall.

If using IPv4:

```
# su -
# service iptables save
iptables: Saving firewall rules to /etc/sysconfig/iptables: [ OK ]
# service iptables stop
iptables: Flushing firewall rules: [ OK ]
iptables: Setting chains to policy ACCEPT: filter [ OK ]
iptables: Unloading modules: [ OK ]
# chkconfig iptables off
# service iptables status
iptables: Firewall is not running.
```

If using IPv6:

```
# su -
# service ip6tables save
ip6tables: Saving firewall rules to /etc/sysconfig/ip6tables: [ OK ]
# service ip6tables stop
ip6tables: Flushing firewall rules: [ OK ]
ip6tables: Setting chains to policy ACCEPT: filter [ OK ]
ip6tables: Unloading modules: [ OK ]
# chkconfig ip6tables off
# service ip6tables status
ip6tables: Firewall is not running.
```

3 Ensure that the open file limit is appropriate by configuring the required minimum.

```
echo "* - nofile 64000" >> /etc/security/limits.conf
```

4 Set SELinux to Permissive.

```
setenforce 0
sed -i "s/SELINUX=[^ ]*/SELINUX=permissive/g" /etc/selinux/config
```

5 Ensure that node hostname is resolvable.**6** Run the vRealize Operations Manager Enterprise bin installer, and follow the prompts.

Add `-i console`, `-i silent`, or `-i gui` to set the installation mode. The default mode conforms to your session type, for example, console for terminal connections or gui for X-Windows.

```
cd /tmp
sh ./vRealize_Operations_Manager_Enterprise.bin -i gui
```

7 If you are creating a multiple node vRealize Operations Manager cluster, repeat all the steps on each Linux machine that will serve as a node in your vRealize Operations Manager cluster.

What to do next

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

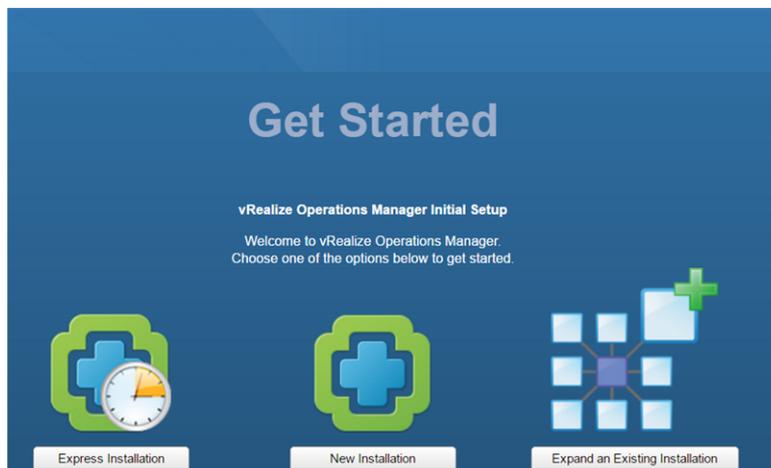
Caution For security, do not access vRealize Operations Manager from untrusted or unpatched clients, or from clients using browser extensions.

Installation Types

After you have installed vRealize Operations Manager 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

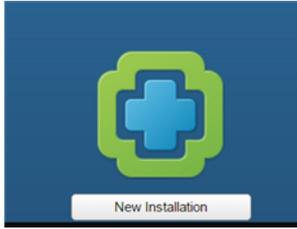


Installing vRealize Operations Manager for a New User

After you install vRealize Operations Manager 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 screen

Perform a New Installation on the vRealize Operations Manager product UI

You can create a single node and configure this as a master node or create a master node in a cluster to handle additional data. All vRealize Operations Manager installations require a master node. With a single node cluster, administration and data functions are on the same master node. A multiple-node vRealize Operations Manager cluster contains one master node and one or more nodes for handling additional data.

Prerequisites

- Create a node by running the vRealize Operations Manager Enterprise installer for Linux.
- 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 Manager.

Procedure

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

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

- 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 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 Select whether to use the certificate included with vRealize Operations Manager 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 Manager.
- 6 Click **Next**.

7 Enter a name for the master node.

For example: **Ops-Master**

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

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

9 Click **Add**.

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

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

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

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

What to do next

After creating the master 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 Manager Master Node

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

The master node performs administration for the cluster and must be online before you configure any new nodes. In addition, the master node must be online before other nodes are brought online. If the master node and replica node go offline together, bring them back online separately. Bring the master node completely online first, and then bring the replica node online. For example, if the entire cluster were offline for any reason, you would bring the master node online first.



Creating the Master Node

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vrops_create_master_node)

Advantages of a New installation

You can use the new installation to create a new master node during the first installation of vRealize Operations Manager. With the master 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 are on the same master node. A multiple-node cluster includes one master 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 more information on creating a master node, see [About the vRealize Operations Manager Master Node](#).

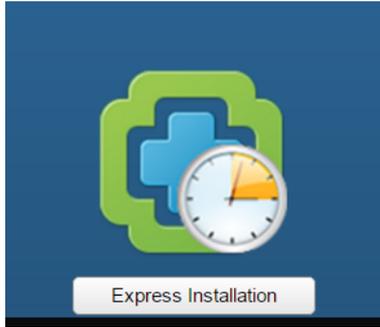
Installing vRealize Operations Manager as an Administrator

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

Introduction to Express Installation

Express installation is one possible way to create master 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 new installation. It is recommended 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 Manager product UI

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

Prerequisites

Verify that you have a static IP address created from either an OVF file or a Linux installer.

Procedure

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

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

- 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**.

You have created a master 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 master node. The express installation uses the default certificates, which differs from one organization to another. This feature is mainly used by the developers or the administrators.

Expand an Existing Installation of vRealize Operations Manager

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

Introduction to expand an existing installation

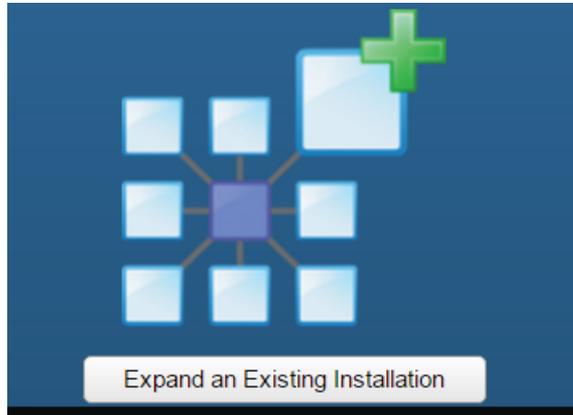
You can deploy and configure additional nodes so that vRealize Operations Manager can support larger environments. A master node always requires an additional node for a cluster to monitor your environment. With expanding your installation, you can add more than one node to your cluster.

Adding Data Nodes

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

You can dynamically scale out vRealize Operations Manager by adding data nodes without stopping the vRealize Operations Manager cluster. When you scale out the cluster by 25% or more, you should restart the cluster to allow vRealize Operations Manager 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 Manager 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 Manager 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 Manager interface. Shut down a node externally only after taking it offline in the vRealize Operations Manager interface.



Creating a Data Node

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vrops_create_data_node)

Expand an existing installation to add a data node

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

Prerequisites

- Create nodes by running the vRealize Operations Manager Enterprise installer for Linux.
- Create and configure the master node.
- Note the fully qualified domain name (FQDN) or IP address of the master 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 Manager.
- 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 master node and verify the thumbprint.

8 Verify the vRealize Operations Manager administrator username of admin.

9 Enter the vRealize Operations Manager administrator password.

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

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

The administration interface appears, and it takes a moment for vRealize Operations Manager 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 master 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 more data nodes.
- Create and add remote collector nodes.
- Create a high availability master replica node, which requires a cluster restart.

Advantages of an Expanding an Installation

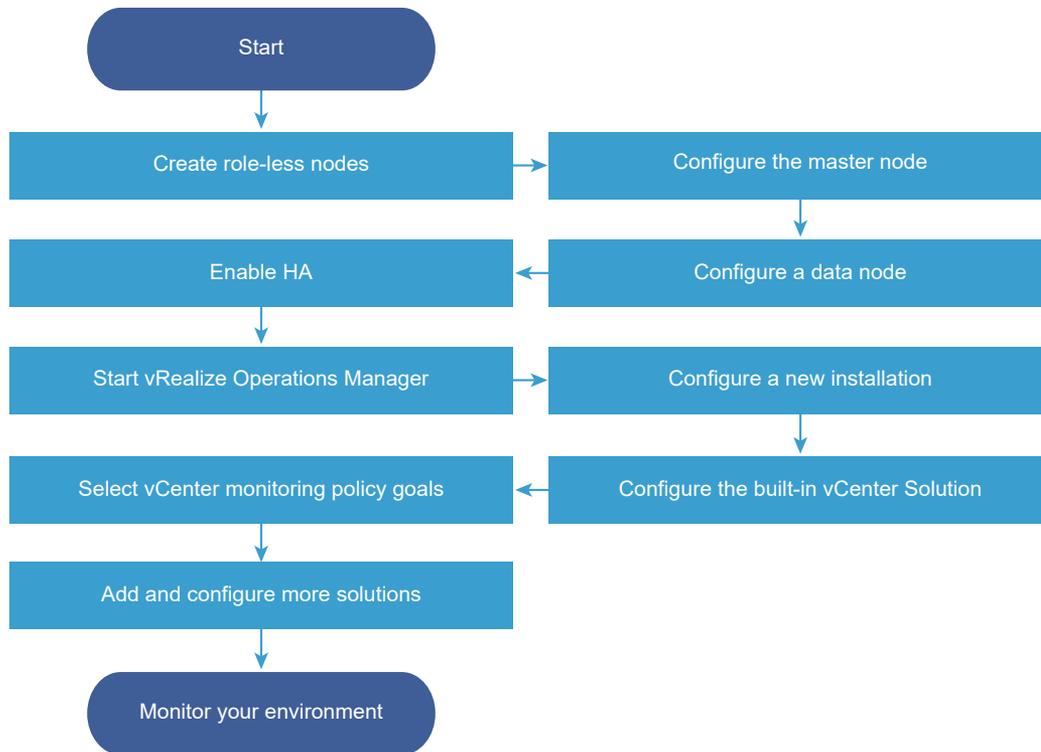
A data node shares the load of performing vRealize Operations Manager analysis and it can also have an adapter installed to perform collection and data storage from the environment. You must have a master node before you add data nodes to form a cluster.

Resize your Cluster by Adding Nodes

4

You can deploy and configure additional nodes so that vRealize Operations Manager 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 Manager Remote Collector Node](#)
- [Adding High Availability to vRealize Operations Manager](#)
- [vRealize Operations Manager Cluster and Node Maintenance](#)

Gathering More Data by Adding a vRealize Operations Manager Remote Collector Node

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

Run the Setup Wizard to Create a Remote Collector Node

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

Prerequisites

- Create nodes by running the vRealize Operations Manager Enterprise installer for Linux.
- 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 master node.
- Note the fully qualified domain name (FQDN) or an IP address of the master 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 Manager.

- 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 master node and verify the thumbprint.

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

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

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

The administration interface appears, and it takes several minutes for vRealize Operations Manager 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 master 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 master replica node, which requires a cluster restart.

Adding High Availability to vRealize Operations Manager

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

Run the Setup Wizard to Add a Master Replica Node

You can convert a vRealize Operations Manager data node to a replica of the master node, which adds high availability (HA) for vRealize Operations Manager.

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

If you convert a data node that is already in use for data collection and analysis, adapters and data connections that were provided through that data node fail over to other data nodes.

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

Prerequisites

- Create nodes by running the vRealize Operations Manager Enterprise installer for Linux.
- Create and configure the master node.

- Create and configure a data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the master 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 Manager administrator username of **admin**.
- 3 Enter the vRealize Operations Manager administrator password and click **Log In**.
- 4 Under High Availability, click **Enable**.
- 5 Select a data node to serve as the replica for the master 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 Manager configures, synchronizes, and rebalances the cluster for HA.

- 7 If the master node and replica node go offline, and the master remains offline for any reason while the replica goes online, the replica node does not take over the master 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 master node and bring the remaining cluster nodes online.

What to do next

After creating a master 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.

vRealize Operations Manager Cluster and Node Maintenance

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

You perform most vRealize Operations Manager 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 master or replica offline causes vRealize Operations Manager to run from the remaining node and for HA status to be degraded.</p> <p>Any manual or system action that restarts the cluster brings all vRealize Operations Manager nodes online, including any nodes that you had taken offline.</p> <p>If you take a data node that is part of a multi-node cluster offline and then bring it back online, the Endpoint Operations Management adapter does not automatically come back online. To bring the Endpoint Operations Management adapter online, select the Endpoint Operations Management adapter in the Inventory Explorer and click the Start Collector icon .</p> |
| Enable or Disable High Availability | Administration | <p>Enabling or disabling 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>Disabling high availability removes the replica node and restarts the vRealize Operations Manager cluster. After you disable high availability, the replica node vRealize Operations Manager converts back to a data node and restarts the 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. The passphrase is only valid for a single use.</p> |

| 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 Manager 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 Manager cluster synchronize with each other by standardizing on the master 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 Manager cluster nodes to increase the efficiency of your environment. |

Cluster Management

vRealize Operations Manager includes a central page where you can monitor and manage the nodes in your vRealize Operations Manager cluster as well as 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 Manager 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

In the left pane, select **Administration > Cluster Management**.

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 Manager cluster. |
| High Availability | Indicates whether HA is enabled, disabled, or degraded. |

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

Table 4-3. Nodes in the vRealize Operations Manager 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. Master and replica nodes require static IP addresses. Data nodes may use DHCP or static IP. |
| Cluster Role | Type of vRealize Operations Manager node: master, data, replica, or remote collector. |
| 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. |
| Metrics in Process | Total metrics that the node has collected since being added to the cluster. |
| Build | vRealize Operations Manager software build number installed on the node. |
| Version | vRealize Operations Manager software version installed on the node. |
| Deployment Type | Type of machine on which the node is running: vApp or Linux |

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. |

vRealize Operations Manager Post-Installation Considerations

5

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

This chapter includes the following topics:

- [About Logging In to vRealize Operations Manager](#)
- [About New vRealize Operations Manager Installations](#)

About Logging In to vRealize Operations Manager

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

When you log in to vRealize Operations Manager, 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 Manager contains further detail about simultaneous login support. See [Knowledge Base article 2093783](#).

- You cannot log in to a vRealize Operations Manager interface with user accounts that are internal to vRealize Operations Manager, 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 Manager Release Notes for your version.

About New vRealize Operations Manager Installations

A new vRealize Operations Manager 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.



Logging In for the First Time

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vrops_first_time_login)

Log In and Continue with a New Installation

To finish a new vRealize Operations Manager 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 Manager nodes.
- Verify that the cluster has enough capacity to monitor your environment. See [Sizing the vRealize Operations Manager Cluster](#).

Procedure

- 1 In a Web browser, navigate to the IP address or fully qualified domain name of the master node.
- 2 Enter the username **admin** and the password that you defined when you configured the master 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 Manager 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 Manager 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 Manager interface appears.

What to do next

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

Updating, Migrating and Restoring

6

You can update your existing vRealize Operations Manager 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 Manager 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.

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](#)
- [Backup and Restore](#)
- [vRealize Operations Manager Software Updates](#)

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 Manager environment, you need to download the right PAK file for the clusters you wish to upgrade. Notice that only the Virtual Appliance clusters use an OS Update PAK file. Host name entries in the `/etc/hosts` of each node might be reset when applying the OS update PAK file for an update from vRealize Operations 6.0.x to version 6.1. You can manually update the hosts file after completing the software update.

Table 6-1. Specific PAK Files for Different Cluster Types

| Cluster Type | OS Update | Product Update |
|--|---|--|
| Virtual Appliance clusters. Use both the OS and the product update PAK files. | vRealize_Operations_Manager-VA-OS- xxx.pak | vRealize_Operations_Manager-VA- xxx.pak |
| RHEL standalone clusters. | | vRealize_Operations_Manager-RHEL- xxx.pak |

Create a Snapshot as Part of an Update

It's a good practice to create a snapshot of each node in a cluster before you update a vRealize Operations Manager 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 Manager 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 Manager 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 Manager, 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 Manager environment.

- If you customized any of the alert definitions, symptom definitions, recommendations, or views that were provided with previous versions of vRealize Operations Manager, 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 Manager, 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 Manager, back up the changes to your alert definitions, symptom definitions, recommendations, and views by cloning them.
- 2 Start the upgrade of vRealize Operations Manager.
- 3 During the upgrade, select the check box named **Reset out-of-the-box content**.

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.

Backup and Restore

Backup and restore your vRealize Operations Manager system on a regular basis 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 backup and restore vRealize Operations Manager 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 backup and restore vRealize Suite components by using vSphere Data Protection and NetBackup, see the Backup and Restore section in the [vRealize Suite Information Center](#).

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

vRealize Operations Manager Software Updates

vRealize Operations Manager 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 Manager product itself.

Where You Find Software Updates

Log in to the vRealize Operations Manager 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 Manager cluster nodes on which they are installed.

Table 6-2. 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 Manager 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. Master 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 | Success, failure, in-progress, or unknown condition of the software update |

Install a Software Update

If you have already installed vRealize Operations Manager, 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.

Prerequisites

- Create a snapshot of each node in your cluster. For information about how to perform this task, see the vRealize Operations Manager Information Center.
- Obtain the PAK file for your cluster. For information about which file to use, see the vRealize Operations Manager Information Center.

- Before you install the PAK file, or upgrade your vRealize Operations Manager 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**.
- The version 6.2.1 vRealize Operations Manager 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 Manager Administrator interface of your cluster at `https://master-node-FQDN-or-IP-address/admin`.
- 2 Click **Software Update** in the left panel.
- 3 Click **Install a Software Update** in the main panel.
- 4 Follow the steps in the wizard to locate and install your PAK file.
 - a If you are updating a Virtual Appliance deployment, perform the OS update.
This updates the OS on the virtual appliance and restarts each virtual machine.
 - b Install the product update PAK file.
Wait for the software update to complete. When it does, the Administrator interface logs you out.
- 5 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.
- 6 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.

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

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

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.

Install a vRealize Operations Manager Software Update from the Administration Interface

You activate the vRealize Operations Manager product or its additional solutions by registering licenses.

Prerequisites

- Know the name and location of the software update PAK file.
- Before you install the PAK file, or upgrade your vRealize Operations Manager 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**.

Procedure

- 1 In a Web browser, navigate to the vRealize Operations Manager administration interface at `https://master-node-name-or-ip-address/admin`.
- 2 Log in with the admin username and password for the master node.
- 3 On the left, click **Software Update**.
- 4 Click **Install a Software Update**.
- 5 Follow the wizard to locate and install your copy of *update-filename.pak*.
Installation completes in a couple of minutes, and the administrator interface logs you out. If you are not logged out automatically after 5 minutes, refresh the page in your browser.
- 6 Log back in to the master node administrator interface, and click **Software Update** again.
- 7 Verify that update name appears on the right. If the update does not appear, wait a few minutes, and refresh the page in your browser.

Migrate a vCenter Operations Manager Deployment into this Version

By importing data, an established or production version of vRealize Operations Manager can assume the monitoring of a vCenter Operations Manager deployment.

You cannot migrate vCenter Operations Manager directly to this version of vRealize Operations Manager. Instead, you follow a two-step process:

- 1 Migrate and import vCenter Operations Manager 5.8.x into vRealize Operations Manager 6.0.x as described in the version 6.0.x documentation.
- 2 Use the vRealize Operations Manager **Software Update** option to update vRealize Operations Manager 6.0.x to this version.

Note Make sure your vCenter Operations Manager 5.8.x and vRealize Operations Manager 6.0.x instances are on the same physical network. Otherwise the data import may not work. Data import process fails when source (vCenter Operations Manager 5.x) is separated from the destination vRealize Operations Manager 6.x environment by a slow network connection (WAN). Data import over a connection that is slower than LAN speed is not supported. For more information, see the Knowledge Base article [2141964](#).

Uninstalling

You can uninstall vRealize Operations Manager instances from your Linux environment.

This chapter includes the following topics:

- [Uninstallation from Linux](#)

Uninstallation from Linux

This release of vRealize Operations Manager for Linux does not include a clean uninstall option. To remove the product, you run the uninstall command and manually remove the remaining artifacts that vRealize Operations Manager installs.

Prerequisites

Log in to the console as root, in vCenter Server or by direct access. In vCenter Server, use Alt+F1 to access the login prompt.

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

Procedure

- 1 Uninstall the product by running the following command:

```
/usr/bin/sh /usr/lib/vmware-vcopsuite-installsupport/_vRealize\ Operations\ Manager\ Enterprise_installation/Uninstall\ vRealize\ Operations\ Manager\ Enterprise -i silent
```

Alternatively, if you are removing the Beta version, run the following command:

```
/usr/bin/sh /usr/lib/vmware-vcopsuite-installsupport/_vCenter\ Operations\ Manager\ Enterprise_installation/Uninstall\ vCenter\ Operations\ Manager\ Enterprise -i silent
```

- 2 Stop the HTTPD service by running the following command:

```
/sbin/service httpd stop
```

- 3 Remove the RPMs by running the following commands:

```
/bin/rpm -e --nodeps httpd  
/bin/rpm -e --nodeps httpd-tools  
/bin/rpm -e --nodeps VMware-Postgres  
/bin/rpm -e --nodeps VMware-Postgres-libs  
/bin/rpm -e --nodeps VMware-Postgres-osslibs  
/bin/rpm -e --nodeps VMware-Postgres-osslibs-server
```

4 Remove the extra users and groups by running the following commands:

```
/usr/sbin/userdel -fr admin
/usr/sbin/userdel -fr postgres
/usr/sbin/groupdel admin
```

5 Remove the extra files and directories by running the following commands:

```
/bin/rm -rf /usr/lib/openssl/lib/libcrypto.so.10
/bin/rm -rf /usr/lib/openssl/lib/libssl.so.10
/bin/rm -rf /usr/lib/openssl/lib/
/bin/rm -rf /usr/lib/openssl/
/bin/rm -rf /usr/lib/vmware-vcopssuite-installsupport/.buildInfo.<build_number>
/bin/rm -rf /usr/lib/vmware-vcopssuite-installsupport/
/bin/rm -rf /etc/rc.d/*/*vmware-vcops-watchdog
/bin/rm -rf /etc/rc.d/*/*vmware-casa
/bin/rm -rf /etc/rc.d/*/*vmware-vcops
/bin/rm -rf /etc/rc.d/*/*vmware-vcops-web
/bin/rm -rf /etc/rc.d/*/*vmware-vcops-reboot-config
/bin/rm -rf /var/log/firstboot
/bin/rm -rf /var/log/preb2b
/bin/rm -rf /var/log/postb2b
/bin/rm -rf /var/log/firstboot
/bin/rm -rf /var/log/casa_logs
/bin/rm -rf /var/log/tomcat_logs
/bin/rm -rf /var/log/vcops_logs
/bin/rm -rf /var/.com.zerog.registry.xml
/bin/rm -rf /var/log/log
```

6 Remove the sudoers entries by running the following commands. If you ran the installer multiple times, you might need to run the following commands multiple times.

```
/bin/sed -i '/# ----- vCenter Operations Manager Settings for VCOPS_USER/,/# ----- End of
vCenter Operations Manager Settings for VCOPS_USER/d' /etc/sudoers
/bin/sed -i '/# ----- vCenter Operations Manager Settings for CaSA/,/# ----- End of vCenter
Operations Manager Settings for CaSA/d' /etc/sudoers
/bin/sed -i '/# ----- vCenter Operations Manager Settings for vsutilities/,/# ----- End of
vCenter Operations Manager Settings for vsutilities/d' /etc/sudoers
```

7 Review the sudoers file /etc/sudoers to ensure that there are no vRealize Operations Manager entries.