

# Installing and Configuring VMware vRealize Orchestrator

vRealize Orchestrator 7.6

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](http://www.vmware.com)

Copyright © 2008-2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

# Contents

Installing and Configuring VMware vRealize Orchestrator	6
<b>1 Introduction to VMware vRealize Orchestrator</b>	<b>7</b>
Key Features of the Orchestrator Platform	7
Orchestrator User Types and Related Responsibilities	9
Orchestrator Architecture	10
Orchestrator Plug-Ins	11
<b>2 Orchestrator System Requirements</b>	<b>12</b>
Hardware Requirements for the Orchestrator Appliance	12
Browsers Supported by vRealize Orchestrator	12
Orchestrator Database Requirements	13
Software Included in the Orchestrator Appliance	13
Level of Internationalization Support	13
vRealize Orchestrator Network Ports	14
<b>3 Setting Up vRealize Orchestrator Components</b>	<b>16</b>
vCenter Server Setup	16
Authentication Methods	16
<b>4 Installing vRealize Orchestrator</b>	<b>18</b>
Download and Deploy the vRealize Orchestrator Appliance	18
Power on the vRealize Orchestrator Appliance and Open the Home Page	19
Change the Root Password	20
Enable or Disable SSH Administrator Login on the vRealize Orchestrator Appliance	20
Configure Network Settings for the vRealize Orchestrator Appliance	21
<b>5 Initial Configuration</b>	<b>23</b>
Configuring a Standalone Orchestrator Server	23
Configure a Standalone Orchestrator Server with vRealize Automation Authentication	23
Configure a Standalone Orchestrator Server with vSphere Authentication	25
vRealize Orchestrator Network Ports	27
Orchestrator Database Connection	28
Manage Certificates	28
Manage Orchestrator Certificates	29
Configure the vRealize Orchestrator Plug-Ins	31
Manage vRealize Orchestrator Plug-Ins	31
Install or Update a vRealize Orchestrator Plug-In	32

Uninstall a Plug-In	33
Orchestrator Availability and Scalability	34
Configure a Cluster of vRealize Orchestrator Instances in VAMI	34
Monitoring an Orchestrator Cluster	36
Enable Sync Mode for the Orchestrator Cluster	36
Promote an Orchestrator Replica Node to Primary Node	37
Delete an Orchestrator Cluster Node	37
Configuring the Customer Experience Improvement Program	38
Categories of Information that VMware Receives	38
Join the Customer Experience Improvement Program	38
<b>6 Using the API services</b>	<b>39</b>
Managing SSL Certificates Through the REST API	39
Delete an SSL Certificate by Using the REST API	40
Import SSL Certificates by Using the REST API	40
Create a Keystore by Using the REST API	41
Delete a Keystore by Using the REST API	42
Add a Key by Using the REST API	42
Automating the Orchestrator Configuration by Using the Control Center REST API	43
<b>7 Additional Configuration Options</b>	<b>44</b>
Reconfiguring Authentication	44
Change the Authentication Provider	44
Change the Authentication Parameters	45
Export the Orchestrator Configuration	46
Import the Orchestrator Configuration	46
Configuring the Workflow Run Properties	47
Orchestrator Log Files	48
Logging Persistence	49
Orchestrator Logs Configuration	49
Filter the Orchestrator Logs	50
Configuring Logging Integration with Remote Server	50
Add Network Interface Controllers	51
Configure Static Routes	52
Enabling the Opentracing and Wavefront Extensions	52
Configure the Opentracing Extension	53
Configure the Wavefront Extension	54
<b>8 Configuration Use Cases and Troubleshooting</b>	<b>56</b>
Configure vRealize Orchestrator Plugin for vSphere Web Client	56
Unregister Orchestrator Authentication	57

Changing SSL Certificates	57
Adding a Certificate to the Local Store	58
Change the Certificate of the Orchestrator Appliance Management Site	58
Cancel Running Workflows	59
Enable Orchestrator Server Debugging	60
Back Up the Orchestrator Configuration and Elements	60
Backing Up and Restoring vRealize Orchestrator	63
Back Up vRealize Orchestrator	63
Restore a vRealize Orchestrator Instance	64
Disaster Recovery of Orchestrator by Using Site Recovery Manager	65
Configure Virtual Machines for vSphere Replication	66
Create Protection Groups	66
Create a Recovery Plan	67
Organize Recovery Plans in Folders	68
Edit a Recovery Plan	69
<b>9 Setting System Properties</b>	<b>70</b>
Disable Access to the Orchestrator Client By Nonadministrators	70
Setting Server File System Access for Workflows and Actions	71
Rules in the js-io-rights.conf File Permitting Write Access to the Orchestrator System	71
Set Server File System Access for Workflows and Actions	72
Set Access to Operating System Commands for Workflows and Actions	73
Set JavaScript Access to Java Classes	74
Set Custom Timeout Property	75
<b>10 Where to Go From Here</b>	<b>76</b>
Log In to the Orchestrator Legacy Client from the Orchestrator Appliance Web Console	76

# Installing and Configuring VMware vRealize Orchestrator

*Installing and Configuring VMware vRealize Orchestrator* provides information and instructions about installing, upgrading and configuring VMware® vRealize Orchestrator.

## Intended Audience

This information is intended for advanced vSphere administrators and experienced system administrators who are familiar with virtual machine technology and datacenter operations.

# Introduction to VMware vRealize Orchestrator

# 1

VMware vRealize Orchestrator is a development- and process-automation platform that provides a library of extensible workflows to allow you to create and run automated, configurable processes to manage VMware products as well as other third-party technologies.

vRealize Orchestrator automates management and operational tasks of both VMware and third-party applications such as service desks, change management systems, and IT asset management systems.

This chapter includes the following topics:

- [Key Features of the Orchestrator Platform](#)
- [Orchestrator User Types and Related Responsibilities](#)
- [Orchestrator Architecture](#)
- [Orchestrator Plug-Ins](#)

## Key Features of the Orchestrator Platform

vRealize Orchestrator is composed of three distinct layers: an orchestration platform that provides the common features required for an orchestration tool, a plug-in architecture to integrate control of subsystems, and a library of workflows. Orchestrator is an open platform that can be extended with new plug-ins and libraries, and can be integrated into larger architectures through a REST API.

Orchestrator includes several key features that help with running and managing workflows.

### **Persistence**

Production-grade databases are used to store relevant information, such as processes, workflow states, and the Orchestrator configuration.

### **Central management**

Orchestrator provides a central way to manage your processes. The application server-based platform, with full version history, can store scripts and process-related primitives in the same storage location. This way, you can avoid scripts without versioning and proper change control on your servers.

## Check-pointing

Every step of a workflow is saved in the database, which prevents data-loss if you must restart the server. This feature is especially useful for long-running processes.

## Control Center

Control Center is a Web-based portal that increases the administrative efficiency of vRealize Orchestrator instances by providing a centralized administrative interface for runtime operations, workflow monitoring, unified log access and configurations, and correlation between the workflow runs and system resources. The Orchestrator logging mechanism is optimized with an extra log file that gathers various performance metrics for the Orchestrator engine throughput.

## Versioning

All Orchestrator Platform objects have an associated version history. Version history is useful for basic change management when distributing processes to project stages or locations.

## Scripting engine

The Mozilla Rhino JavaScript engine provides a way to create building blocks for the Orchestrator platform. The scripting engine is enhanced with basic version control, variable type checking, name space management, and exception handling. The engine can be used in the following building blocks:

- Actions
- Workflows
- Policies

## Workflow engine

The workflow engine allows you to automate business processes. It uses the following objects to create a step-by-step process automation in workflows:

- Workflows and actions that Orchestrator provides
- Custom building blocks created by the customer
- Objects that plug-ins add to Orchestrator

Users, other workflows, schedules, or policies can start workflows.

## Policy engine

You can use the policy engine to monitor and generate events to react to changing conditions in the Orchestrator server or a plugged-in technology. Policies can aggregate events from the platform or the plug-ins, which helps you to handle changing conditions on any of the integrated technologies.

## Orchestrator Client



Create, run, edit, and monitor workflows with the vRealize Orchestrator Client. You can also use the vRealize Orchestrator Client to manage action, configuration, policy, and resource elements. For more information, see *Using the vRealize Orchestrator Client*.

---

**Note** For information on the deprecated Java-based Orchestrator Legacy Client, see *Using the VMware vRealize Orchestrator Legacy Client*.

---

## Development and resources

The Orchestrator landing page provides quick access to resources to help you develop your own plug-ins, for use in vRealize Orchestrator. You will also find information about using the Orchestrator REST API to send requests to the Orchestrator server.

## Security

Orchestrator provides the following advanced security functions:

- Public Key Infrastructure (PKI) to sign and encrypt content imported and exported between servers.
- Digital Rights Management (DRM) to control how exported content can be viewed, edited, and redistributed.
- Secure Sockets Layer (SSL) to provide encrypted communications between the desktop client and the server and HTTPS access to the Web front end.
- Advanced access rights management to provide control over access to processes and the objects manipulated by these processes.

## Encryption

vRealize Orchestrator uses a FIPS-compliant Advanced Encryption Standard (AES) with a 256-bit cipher key for encryption of strings. The cipher key is randomly generated and is unique across appliances that are not part of a cluster. All nodes in a cluster share the same cipher key.

# Orchestrator User Types and Related Responsibilities

Orchestrator provides different tools and interfaces based on the specific responsibilities of the global user roles. In Orchestrator, you can have users with full rights, that are a part of the administrator group (Administrators) and users with limited rights, that are not part of the administrator group (End Users).

## Users with Full Rights

Orchestrator administrators and developers have equal administrative rights, but are divided in terms of responsibilities.

### Administrators

This role has full access to all of the Orchestrator platform capabilities. Basic administrative responsibilities include the following items:

- Installing and configuring Orchestrator
- Managing access rights for Orchestrator and applications
- Importing and exporting packages
- Running workflows and scheduling tasks
- Managing version control of imported elements
- Creating new workflows and plug-ins

### **Developers**

This user type has full access to all of the Orchestrator platform capabilities. Developers are granted access to the Orchestrator client interface and have the following responsibilities:

- Creating applications to extend the Orchestrator platform functionality
- Automating processes by customizing existing workflows and creating new workflows and plug-ins

## **Users with Limited Rights**

### **End Users**

End users can run and schedule workflows and policies that the administrators or developers make available in the Orchestrator client.


## **Orchestrator Architecture**

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

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

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

Figure 1-1. VMware vRealize Orchestrator Architecture



## Orchestrator Plug-Ins

Plug-ins allow you to use Orchestrator to access and control external technologies and applications. By exposing an external technology in an Orchestrator plug-in, you can incorporate objects and functions in workflows that access the objects and functions of that external technology.

The external technologies that you can access by using plug-ins include virtualization management tools, email systems, databases, directory services, and remote-control interfaces.

Orchestrator provides a set of standard plug-ins that you can use to incorporate into workflows such technologies as the VMware vCenter Server API and email capabilities. By using the plug-ins, you can automate the delivery of new IT services or adapt the capabilities of existing vRealize Automation infrastructure and application services. In addition, you can use the Orchestrator open plug-in architecture to develop plug-ins for accessing other applications.

The Orchestrator plug-ins that VMware develops are distributed as `.vmoapp` files. For more information about the Orchestrator plug-ins that VMware develops and distributes, see [vRealize Orchestrator External Plug-ins](#). For more information about third-party Orchestrator plug-ins, see [VMware Solution Exchange](#).

# Orchestrator System Requirements

# 2

Your system must meet the technical requirements that are necessary for Orchestrator to work properly.

For a list of the supported versions of vCenter Server, the vSphere Web Client, vRealize Automation, and other VMware solutions, as well as compatible database versions, see [VMware Product Interoperability Matrix](#).

This chapter includes the following topics:

- [Hardware Requirements for the Orchestrator Appliance](#)
- [Browsers Supported by vRealize Orchestrator](#)
- [Orchestrator Database Requirements](#)
- [Software Included in the Orchestrator Appliance](#)
- [Level of Internationalization Support](#)
- [vRealize Orchestrator Network Ports](#)

## Hardware Requirements for the Orchestrator Appliance

The Orchestrator Appliance is a preconfigured Linux-based virtual machine. Before you deploy the appliance, verify that your system meets the minimum hardware requirements.

The Orchestrator Appliance has the following hardware requirements:

- 2 CPUs
- 6 GB of memory
- 17 GB hard disk

Do not reduce the default memory size, because the Orchestrator server requires at least 2 GB of free memory.

## Browsers Supported by vRealize Orchestrator

Confirm that your browsers support vRealize Orchestrator.

To access the vRealize Orchestrator Client and Control Center, you must use one of the following browsers:

- Microsoft Edge
- Mozilla Firefox
- Google Chrome

## Orchestrator Database Requirements

The Orchestrator server includes a preconfigured PostgreSQL database that is production ready.

Starting with vRealize Orchestrator 7.5, external database integration is not supported. You can only use the preconfigured PostgreSQL database.

## Software Included in the Orchestrator Appliance

The Orchestrator Appliance is a preconfigured virtual machine optimized for running Orchestrator. The appliance is distributed with preinstalled software.

The Orchestrator Appliance package contains the following software:

- SUSE Linux Enterprise Server 11 Update 3 for VMware, 64-bit edition
- PostgreSQL
- Orchestrator

The default Orchestrator Appliance database configuration is production ready.

---

**Note** To use the Orchestrator Appliance in a production environment, you must configure the Orchestrator server to authenticate through vRealize Automation or vSphere. For more information about configuring an authentication provider, see [Configuring a Standalone Orchestrator Server](#).

---

## Level of Internationalization Support

The Orchestrator Control Center includes a Spanish, French, German, Traditional Chinese, Simplified Chinese, Korean, and Japanese locale. The Orchestrator client supports internationalization level 1.

## Non-ASCII Character Support in Orchestrator

Although Orchestrator the Orchestrator client is not localized, it can run on a non-English operating system and support non-ASCII text.

Table 2-1. Non-ASCII Character Support in Orchestrator GUI

Support for Non-ASCII Characters				
Orchestrator Item	Description Field	Name Field	Input and Output Parameters	Attributes
Action	Yes	No	No	No
Folder	Yes	Yes	-	-
Configuration element	Yes	Yes	-	No
Package	Yes	Yes	-	-
Policy	Yes	Yes	-	-
Policy template	Yes	Yes	-	-
Resource element	Yes	Yes	-	-
Workflow	Yes	Yes	No	No
Workflow presentation display group and input step	Yes	Yes	-	-

## vRealize Orchestrator Network Ports

vRealize Orchestrator uses specific ports to communicate with other systems. The ports are set with a default value that cannot be changed.

### Default Configuration Ports

To provide the vRealize Orchestrator service, you must set default ports and configure your firewall to allow incoming TCP connections.

**Note** Other ports might be required if you are using custom plug-ins.

Table 2-2. VMware vRealize Orchestrator Default Configuration Ports

Port	Number	Protocol	Source	Target	Description
Virtual Appliance Management Interface	5480	TCP			The access port to the appliance system settings interface.
vRealize Orchestrator appliance	5488 5489	TCP			The port internally used by the vRealize Orchestrator appliance for updates.
HTTP server port	80	TCP	End-user web browser	vRealize Orchestrator server	Optional. The requests sent to vRealize Orchestrator default HTTP web port 80 are redirected to the default HTTPS web port 8281.
HTTP server port	8280	TCP	End-user web browser	vRealize Orchestrator server	The requests sent to Orchestrator default HTTP web port 8280 are redirected to the default HTTPS web port 8281.

**Table 2-2. VMware vRealize Orchestrator Default Configuration Ports (continued)**

Port	Number	Protocol	Source	Target	Description
HTTPS server port	8281	TCP	End-user web browser	vRealize Orchestrator server	The access port for the vRealize Orchestrator home page.
Web configuration HTTPS access port	8283	TCP	End-user web browser	vRealize Orchestrator configuration	The SSL access port for the vRealize Orchestrator Client.

## External Communication Ports

You must configure your firewall to allow outgoing connections so that vRealize Orchestrator can communicate with external services.

**Table 2-3. VMware vRealize Orchestrator External Communication Ports**

Port number	Protocol	Source	Target	Description
123	UDP	vRealize Orchestrator server	NTP server	The default port for connecting directly to NTP instead of using host time.
25	TCP	vRealize Orchestrator server	SMTP server	The port used for email notifications.
443	TCP	vRealize Orchestrator server	vCenter Server API	The vCenter Server API communication port used by vRealize Orchestrator to obtain virtual infrastructure and virtual machine information from the orchestrated vCenter Server instances.
4000	UDP	vRealize Orchestrator server	SNMP server	The default port for listening to SNMP traps in the SNMP plug-in.
514	UDP	vRealize Orchestrator server	Syslog server	The port for sending syslog event messages.
5432	TCP	vRealize Orchestrator server	PostgreSQL server	The port used to communicate with the PostgreSQL server that is configured as the vRealize Orchestrator database.
5434	TCP	vRealize Orchestrator server	PostgreSQL server	The port used by the PostgreSQL manager service for database fault tolerance.

# Setting Up vRealize Orchestrator Components

# 3

When you download, and deploy the vRealize Orchestrator Appliance, the vRealize Orchestrator server is preconfigured. After deployment, the service starts automatically.

To enhance the availability and scalability of your vRealize Orchestrator setup, follow these guidelines:

- Install and configure an authentication provider and configure vRealize Orchestrator to work with the provider .
- For clustered vRealize Orchestrator environments, install and configure a load balancing server and configure it to distribute the workload between two or more vRealize Orchestrator servers.

This chapter includes the following topics:

- [vCenter Server Setup](#)
- [Authentication Methods](#)

## vCenter Server Setup

Increasing the number of vCenter Server instances in your Orchestrator setup causes Orchestrator to manage more sessions. Too many active sessions can cause Orchestrator to experience timeouts when more than 10 vCenter Server connections occur.

For a list of the supported versions of vCenter Server, see [VMware Product Interoperability Matrix](#).

---

**Note** You can run multiple vCenter Server instances on different virtual machines in your Orchestrator setup if your network has sufficient bandwidth and latency. If you are using LAN to improve the communication between Orchestrator and vCenter Server, a 100-Mb line is mandatory.

---

## Authentication Methods

To authenticate and manage user permissions, Orchestrator requires a connection to either vRealize Automation or a vSphere server instance.



When you download, and deploy the Orchestrator Appliance, you must set up a connection with a vRealize Automation or vSphere.

# Installing vRealize Orchestrator

# 4

vRealize Orchestrator consists of a server component and a client component.

To use vRealize Orchestrator, you must deploy the vRealize Orchestrator Appliance and configure the vRealize Orchestrator server.

You can change the default vRealize Orchestrator configuration settings by using the vRealize Orchestrator Control Center.

This chapter includes the following topics:

- [Download and Deploy the vRealize Orchestrator Appliance](#)

## Download and Deploy the vRealize Orchestrator Appliance

Download and install an vRealize Orchestrator Appliance by deploying it from a template.

### Prerequisites

- Verify that vCenter Server is installed and running.
- Verify that the host on which you are deploying the vRealize Orchestrator Appliance meets the minimum hardware requirements. For more information, see [Hardware Requirements for the Orchestrator Appliance](#).
- If your system is isolated and without Internet access, you must download the .ova file for the appliance from the VMware website.

### Procedure

- 1 Log in to the vSphere Web Client as an administrator.
- 2 In the vSphere Web Client, select an inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host.
- 3 Select **Actions > Deploy OVF Template**.
- 4 Enter the path or the URL to the .ova file and click **Next**.
- 5 Enter a name and location for the deployed vRealize Orchestrator Appliance, and click **Next**.

- 6 Select a host, cluster, resource pool, or vApp as a destination on which you want the appliance to run, and click **Next**.
- 7 Review the deployment details, and click **Next**.
- 8 Accept the terms in the license agreement and click **Next**.
- 9 Select the storage format you want to use for the deployed vRealize Orchestrator Appliance.

Format	Description
<b>Thick Provisioned Lazy Zeroed</b>	Creates a virtual disk in a default thick format. The space required for the virtual disk is allocated when the virtual disk is created. If any data remains on the physical device, it is not erased during creation, but is zeroed out on demand later on first write from the virtual machine.
<b>Thick Provisioned Eager Zeroed</b>	Supports clustering features such as Fault Tolerance. The space required for the virtual disk is allocated when the virtual disk is created. If any data remains on the physical device, it is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create disks in other formats.
<b>Thin Provisioned Format</b>	Saves hard disk space. For the thin disk, you provision as much datastore space as the disk requires based on the value that you select for the disk size. The thin disk starts small and, at first, uses only as much datastore space as the disk needs for its initial operations.

- 10 Click **Next**.
- 11 (Optional) Configure the network settings, and click **Next**.  
By default, the vRealize Orchestrator Appliance uses DHCP. You can change this setting and assign a fixed IP address from the appliance Web console.

- 12 Select the options that you want to enable and set the initial password for the root user account.

Your initial password must be at least eight characters long.

---

**Important** The password for the root account of the Orchestrator Appliance expires after 365 days. You can increase the expiry time for an account by logging in to the Orchestrator Appliance as root, and running `passwd -x number_of_days name_of_account`. If you want to increase the Orchestrator Appliance root password to infinity, run `passwd -x 99999 root`.

---

- 13 Review the **Ready to complete** page and click **Finish**.

## Results

The vRealize Orchestrator Appliance is successfully deployed.

## Power on the vRealize Orchestrator Appliance and Open the Home Page

To use the vRealize Orchestrator Appliance, you must first power it on and get an IP address for the virtual appliance.

### Procedure

- 1 Log in to the vSphere Web Client as an administrator.
- 2 Right-click the vRealize Orchestrator Appliance and select **Power > Power On**.
- 3 After the appliance powers on, select the **Summary** tab to view the vRealize Orchestrator Appliance IP address.
- 4 In a Web browser, navigate to the host address of your vRealize Orchestrator Appliance virtual machine.

`https://your_orchestrator_hostname/vco.`

## Change the Root Password

For security reasons, you can change the root password of the vRealize Orchestrator Appliance.

By default, the password for the root account of the vRealize Orchestrator Appliance expires after 365 days. You can increase the expiry for the root account by logging into the vRealize Orchestrator Appliance through an SSH client and running `passwd -x number_of_days name_of_account`. If you want to increase the vRealize Orchestrator Appliance root password to infinity, run `passwd -x 99999 root`.

### Prerequisites

- Download and deploy the vRealize Orchestrator Appliance.
- Verify that the vRealize Orchestrator Appliance is up and running.

### Procedure

- 1 Log in to the vRealize Orchestrator VAMI as **root**.  
.Access the VAMI at `https://your_orchestrator_hostname:5480`.
- 2 Select the **Admin** tab.
- 3 In the **Current administrator password** text box, enter the current root password.
- 4 Enter the new password in the **New administrator password** and **Retype new administrator password** text boxes.
- 5 Click **Save Settings**.

### Results

You successfully changed the password of the root Linux user of the vRealize Orchestrator Appliance.

## Enable or Disable SSH Administrator Login on the vRealize Orchestrator Appliance

You can enable or disable SSH access to the vRealize Orchestrator Appliance.

**Prerequisites**

- Download and deploy the vRealize Orchestrator Appliance.
- Verify that the vRealize Orchestrator Appliance is up and running.

**Procedure**

- 1 Log in to the vRealize Orchestrator VAMI as **root**.  
Access the VAMI at `https://your_orchestrator_hostname:5480`.
- 2 On the **Admin** tab, click **SSH service enabled** to enable or disable the vRealize Orchestrator SSH service.
- 3 (Optional) Click **Administrator SSH login enabled** to enable or disable root access to the vRealize Orchestrator Appliance using SSH.
- 4 Click **Save Settings**.

**Results**

When enabled, the **SSH Status** appears as *Running*. When disabled, the **SSH Status** appears as *Stopped*.

## Configure Network Settings for the vRealize Orchestrator Appliance

Configure network settings for the vRealize Orchestrator Appliance to assign a static IP address and define the proxy settings.

**Prerequisites**

- Download and deploy the vRealize Orchestrator Appliance.
- Verify that the vRealize Orchestrator Appliance is up and running.

**Procedure**

- 1 Log in to the vRealize Orchestrator VAMI as **root**.  
Access the VAMI at `https://your_orchestrator_hostname:5480`.
- 2 On the **Network** tab, click **Address**.
- 3 Select the method by which the vRealize Orchestrator Appliance obtains IP address settings.

Option	Description
<b>DHCP</b>	Obtains IP settings from a DHCP server. This is the default setting.
<b>Static</b>	Uses static IP settings. Selecting this option prompts you to enter an IP address, netmask (for IPv4), prefix (for IPv6), and gateway information.

Depending on your network settings, you might have to select IPv4 and IPv6 address types.

- 4 Click **Save Settings**.
- 5 (Optional) To configure a proxy server, select the **Proxy** tab.

- 6 (Optional) After configuring your proxy settings, click **Save Settings**.

# Initial Configuration

# 5

Before you begin automating tasks and managing systems and applications with vRealize Orchestrator, you must use the vRealize Orchestrator Control Center to configure an external authentication provider. You can also use the vRealize Orchestrator Control Center for additional configuration tasks like managing license and certificate information, installing plug-ins, or monitoring and configuring the vRealize Orchestrator logs.

This chapter includes the following topics:

- [Configuring a Standalone Orchestrator Server](#)
- [vRealize Orchestrator Network Ports](#)
- [Orchestrator Database Connection](#)
- [Manage Certificates](#)
- [Configure the vRealize Orchestrator Plug-Ins](#)
- [Orchestrator Availability and Scalability](#)
- [Configuring the Customer Experience Improvement Program](#)

## Configuring a Standalone Orchestrator Server

Although the Orchestrator Appliance is a preconfigured Linux-based virtual machine, you must follow the configuration wizard before you access the Orchestrator Control Center.

### Configure a Standalone Orchestrator Server with vRealize Automation Authentication

To prepare the Orchestrator Appliance for use, you must configure host settings and the authentication provider. You can configure Orchestrator to authenticate through the vRealize Automation component registry.

#### Prerequisites

- Download and deploy the latest version of the vRealize Orchestrator Appliance. See [Download and Deploy the vRealize Orchestrator Appliance](#).

- Install and configure vRealize Automation and verify that your vRealize Automation server is running. See the vRealize Automation documentation.

If you plan to create a cluster:

- Set up a load balancer to distribute traffic among multiple instances of vRealize Orchestrator. For more information, see the documentation for vRealize Orchestrator Load Balancing.

### Procedure

- 1 Access Control Center to start the configuration wizard.
  - a Navigate to `https://your_orchestrator_server_IP_or_DNS_name:8283/vco-controlcenter`.
  - b Log in as **root** with the password you entered during OVA deployment.
- 2 Click **CHANGE** to configure the host name on which Control Center will be accessible.

---

**Note** If you are about to configure an Orchestrator cluster, enter the host name of the load balancer virtual server.

---

- 3 Configure the authentication provider.
  - a On the **Configure Authentication Provider** page, select **vRealize Automation** from the **Authentication mode** drop-down menu.
  - b In the **Host address** text box, enter your vRealize Automation host address and click **CONNECT**.
  - c Click **Accept Certificate**.
  - d In the **User name** and **Password** text boxes, enter the credentials of the user account that is configured for SSO connection in vRealize Automation. Click **REGISTER**.  
  
By default, the SSO account is **administrator** and the name of the default tenant is **vsphere.local**.
  - e In the **Admin group** text box, enter the name of an administrators group and click **SEARCH**.  
  
For example, **vsphere.local\vcadmins**
  - f In the list of groups, click on the name of the group to select it.
  - g Click **SAVE CHANGES**.

A message indicates that your configuration is saved successfully.

### Results

You have successfully finished the Control Center configuration.

### What to do next

- Verify that **VRA** is the configured license provider at the **Licensing** page.



- Verify that the node is configured properly at the **Validate Configuration** page.

---

**Note** Following the configuration of the authentication provider, the Orchestrator server restarts automatically after 2 minutes. Verifying the configuration immediately after process completion can return an invalid configuration status.

---

## Configure a Standalone Orchestrator Server with vSphere Authentication

You register the Orchestrator server with a vCenter Single Sign-On server by using the vSphere Authentication mode. Use vCenter Single Sign-On authentication with vCenter Server 6.0 and later.

### Prerequisites

- Download and deploy the latest version of the vRealize Orchestrator Appliance. See [Download and Deploy the vRealize Orchestrator Appliance](#).
- Install and configure vCenter Server with vCenter Single Sign-On running. For information, see the vSphere documentation.

If you plan to create a cluster:

- Set up a load balancer to distribute traffic among multiple instances of vRealize Orchestrator. For more information, see the documentation for vRealize Orchestrator Load Balancing.

### Procedure

- 1 Access Control Center to start the configuration wizard.
  - a Navigate to `https://your_orchestrator_server_IP_or_DNS_name:8283/vco-controlcenter`.
  - b Log in as **root** with the password you entered during OVA deployment.
- 2 Click **CHANGE** to configure the host name on which Control Center will be accessible.

---

**Note** If you are about to configure an Orchestrator cluster, enter the host name of the load balancer virtual server.

---

### 3 Configure the authentication provider.

- a On the **Configure Authentication Provider** page, select **vSphere** from the **Authentication mode** drop-down menu.
- b In the **Host address** text box, enter the fully qualified domain name or IP address of the Platform Services Controller instance that contains the vCenter Single Sign-On and click **CONNECT**.

---

**Note** If you use an external Platform Services Controller or multiple Platform Services Controller instances behind a load balancer, you must manually import the certificates of all Platform Services Controllers that share the same vCenter Single Sign-On domain.

---

**Note** To integrate a different vSphere Client with your configured vRealize Orchestrator environment, you must configure vSphere to use the same Platform Services Controller registered to Orchestrator. For High Availability Orchestrator environments, you must replicate the PCS instances behind the Orchestrator load balancer server.

---

- c Click **Accept Certificate**.
- d In the **User name** and **Password** text boxes, enter the credentials of the local administrator account for the vCenter Single Sign-On domain. Click **REGISTER**.  
By default, this account is **administrator@vsphere.local** and the name of the default tenant is **vsphere.local**.
- e In the **Admin group** text box, enter the name of an administrators group and click **SEARCH**.  
For example, **vsphere.local\vcoadmins**
- f In the list of groups, click on the name of the group to select it.
- g Click **SAVE CHANGES**.  
A message indicates that your configuration is saved successfully.

#### Results

You have successfully completed the Control Center configuration.

#### What to do next

- Verify that **CIS** is the configured license provider at the **Licensing** page.
- Verify that the node is configured properly at the **Validate Configuration** page.

---

**Note** Following the configuration of the authentication provider, the Orchestrator server restarts automatically after 2 minutes. Verifying the configuration immediately after process completion can return an invalid configuration status.

---

## vRealize Orchestrator Network Ports

vRealize Orchestrator uses specific ports to communicate with other systems. The ports are set with a default value that cannot be changed.

### Default Configuration Ports

To provide the vRealize Orchestrator service, you must set default ports and configure your firewall to allow incoming TCP connections.

**Note** Other ports might be required if you are using custom plug-ins.

**Table 5-1. VMware vRealize Orchestrator Default Configuration Ports**

Port	Number	Protocol	Source	Target	Description
Virtual Appliance Management Interface	5480	TCP			The access port to the appliance system settings interface.
vRealize Orchestrator appliance	5488 5489	TCP			The port internally used by the vRealize Orchestrator appliance for updates.
HTTP server port	80	TCP	End-user web browser	vRealize Orchestrator server	Optional. The requests sent to vRealize Orchestrator default HTTP web port 80 are redirected to the default HTTPS web port 8281.
HTTP server port	8280	TCP	End-user web browser	vRealize Orchestrator server	The requests sent to Orchestrator default HTTP web port 8280 are redirected to the default HTTPS web port 8281.
HTTPS server port	8281	TCP	End-user web browser	vRealize Orchestrator server	The access port for the vRealize Orchestrator home page.
Web configuration HTTPS access port	8283	TCP	End-user web browser	vRealize Orchestrator configuration	The SSL access port for the vRealize Orchestrator Client.

### External Communication Ports

You must configure your firewall to allow outgoing connections so that vRealize Orchestrator can communicate with external services.

Table 5-2. VMware vRealize Orchestrator External Communication Ports

Port number	Protocol	Source	Target	Description
123	UDP	vRealize Orchestrator server	NTP server	The default port for connecting directly to NTP instead of using host time.
25	TCP	vRealize Orchestrator server	SMTP server	The port used for email notifications.
443	TCP	vRealize Orchestrator server	vCenter Server API	The vCenter Server API communication port used by vRealize Orchestrator to obtain virtual infrastructure and virtual machine information from the orchestrated vCenter Server instances.
4000	UDP	vRealize Orchestrator server	SNMP server	The default port for listening to SNMP traps in the SNMP plug-in.
514	UDP	vRealize Orchestrator server	Syslog server	The port for sending syslog event messages.
5432	TCP	vRealize Orchestrator server	PostgreSQL server	The port used to communicate with the PostgreSQL server that is configured as the vRealize Orchestrator database.
5434	TCP	vRealize Orchestrator server	PostgreSQL server	The port used by the PostgreSQL manager service for database fault tolerance.

## Orchestrator Database Connection

The Orchestrator server requires a database for storing data.

When you download, and deploy the Orchestrator Appliance, the Orchestrator server is configured to work with the PostgreSQL database preinstalled in the appliance.

The preconfigured Orchestrator PostgreSQL database is production ready. All transactions of the Orchestrator PostgreSQL are handled automatically through the VAMI interface.

**Note** Starting with vRealize Orchestrator 7.5, external databases like Oracle and Microsoft SQL are not supported.

## Manage Certificates

Issued for a particular server and containing information about the server public key, the certificate allows you to sign all elements created in vRealize Orchestrator and guarantee

authenticity. When the client receives an element from your server, typically a package, the client verifies your identity and decides whether to trust your signature.

- **Manage Orchestrator Certificates**

You can manage the Orchestrator certificates from the **Certificates** page in Control Center or through the Orchestrator client, by using the SSL Trust Manager workflows in the Configuration workflow category.

## Manage Orchestrator Certificates

You can manage the Orchestrator certificates from the **Certificates** page in Control Center or through the Orchestrator client, by using the SSL Trust Manager workflows in the Configuration workflow category.

### Import a Certificate to the Orchestrator Trust Store

Control Center uses a secure connection to communicate with vCenter Server, relational database management system (RDBMS), LDAP, Single Sign-On, and other servers. You can import the required SSL certificate from a URL or a PEM-encoded file. Each time you want to use an SSL connection to a server instance, you must import the corresponding certificate from the **Trusted Certificates** tab on the **Certificates** page and import the corresponding SSL certificate.

You can load the SSL certificate in Orchestrator from a URL address or a PEM-encoded file.

Option	Description
<b>Import from URL or proxy URL</b>	The URL of the remote server: <b><code>https://your_server_IP_address</code></b> or <b><code>your_server_IP_address:port</code></b>
<b>Import from file</b>	Path to the PEM-encoded certificate file. For more information on importing a PEM-encoded certificate file, see <a href="#">Import a Trusted Certificate Through Control Center</a> .

### Generate a Self-Signed Server Certificate

The Orchestrator Appliance includes a self-signed certificate that is generated automatically, based on the network settings of the appliance. If the network settings of the appliance change, you must generate a new self-signed certificate manually. You can create a self-signed certificate to guarantee encrypted communication and provide a signature for your packages. However, the recipient cannot be sure that the self-signed package is in fact a package issued by your server and not a third party claiming to be you. To prove the identity of your server, use a certificate signed by a Certificate Authority.

You can generate a self-signed certificate on the **Orchestrator Server SSL Certificate** tab from the **Certificates** page in Control Center.

Option	Description
<b>Signature Algorithm</b>	Encryption algorithm to generate a digital signature.
<b>Common Name</b>	Host name of the Orchestrator server.

Option	Description
<b>Organization</b>	Name of your organization. For example, <b>VMware</b> .
<b>Organizational Unit</b>	Name of your organizational unit. For example, <b>R&amp;D</b> .
<b>Country Code</b>	Country code abbreviation. For example, <b>US</b> .

Orchestrator generates a server certificate that is unique to your environment. The details about the public key of the certificate appear in the **Orchestrator Server SSL Certificate** tab. The private key is stored in the `vmo_keystore` table of the Orchestrator database.

## Import an Orchestrator Server SSL Certificate

vRealize Orchestrator uses an SSL certificate to identify itself to clients and remote servers during secure communication. By default, Orchestrator includes a self-signed SSL certificate that is generated automatically, based on the network settings of the appliance. You can import an SSL certificate signed by a Certificate Authority to avoid certificate trust errors.

You must import a certificate signed by a Certificate Authority as a PEM-encoded file that contains the public and the private key.

**Note** After generating or importing an SSL server certificate, refresh the **Orchestrator Server SSL Certificate** tab to view the details of the new certificate. You must also restart the Orchestrator configurator service.

```
service vco-configurator restart
```

## Package Signing Certificate

Packages exported from an Orchestrator server are digitally signed. Import, export, or generate a new certificate to be used for signing packages. Package signing certificates are a form of digital identification that is used to guarantee encrypted communication and a signature for your Orchestrator packages.

The Orchestrator Appliance includes a package signing certificate that is generated automatically, based on the network settings of the appliance. If the network settings of the appliance change, you must generate a new package signing certificate manually.

**Note** The Orchestrator Appliance includes a self-signed package signing certificate that is generated automatically during the initial Orchestrator configuration. You can change the package signing certificate, after which, all future exported packages are signed with the new certificate.

## Import a Trusted Certificate Through Control Center

To communicate with other servers securely, the Orchestrator server must be able to verify their identity. For this purpose, you might need to import the SSL certificate of the remote entity to the Orchestrator trust store. To trust a certificate, you can import it to the trust store either by establishing a connection to a specific URL, or directly as a PEM-encoded file.

## Prerequisites

Find the fully qualified domain name of the server to which you want Orchestrator to connect over SSL.

## Procedure

- 1 Log in to the Orchestrator Appliance over SSH as **root**.
- 2 Run a command to retrieve the certificate of the remote server.

```
openssl s_client -connect host_or_dns_name:secure_port
```

- a If you use a nonencrypted port, use `starttls` and the required protocol with the `openssl` command.

```
openssl s_client -connect host_or_dns_name:port -starttls smtp
```

- 3 Copy the text from the -----BEGIN CERTIFICATE----- to the -----END CERTIFICATE----- tag to a text editor and save it as a file.
- 4 Log in to Control Center as **root**.
- 5 Go to the **Certificates** page.
- 6 On the **Trusted Certificates** tab, click **Import** and select the **Import from a PEM-encoded file** option.
- 7 Browse to the certificate file and click **Import**.

## Results

You have successfully imported a remote server certificate to the Orchestrator trust store.

## Configure the vRealize Orchestrator Plug-Ins

The default vRealize Orchestrator plug-ins are configured through plug-in specific workflows run in the vRealize Orchestrator Client.

The vRealize Orchestrator Appliance provides access to a pre-installed library of default plug-ins. You can configure these default plug-ins, by running workflows specific to them from the vRealize Orchestrator Client.

For example, entering the tags *AMQP* and *Configuration* in the search text box of the workflow library, provides workflows that are used to manage AMQP brokers and subscriptions.

## Manage vRealize Orchestrator Plug-Ins

In the **Manage Plug-Ins** page of vRealize Orchestrator Control Center, you can view a list of all plug-ins that are installed in vRealize Orchestrator and perform basic management actions.

## Change Plug-Ins Logging Level

Instead of changing the logging level for vRealize Orchestrator, you can change it only for specific plug-ins.

## Install or Upgrade a New Plug-In

With the vRealize Orchestrator plug-ins, the vRealize Orchestrator server can integrate with other software products. The vRealize Orchestrator Appliance includes a set of pre-installed plug-ins. You can also further expand the capabilities of the vRealize Orchestrator platform by installing custom plug-ins.

You can install or upgrade plug-ins from the **Manage Plug-Ins** page of the vRealize Orchestrator. The file extensions that can be used are `.vmoapp` and `.dar`. A `.vmoapp` file can contain a collection of several `.dar` files and can be installed as an application. A `.dar` file contains all the resources associated with one plug-in.

---

**Note** The preferred file format for vRealize Orchestrator plug-ins is `.vmoapp`.

---

For more information on installing or upgrading vRealize Orchestrator plug-ins, see [Install or Update a vRealize Orchestrator Plug-In](#).

## Disable a Plug-In

You can disable a plug-in by deselecting the **Enable** check box next to the name of the plug-in.

This action does not remove the plug-in file. For more information on uninstalling a plug-in in Orchestrator, see [Uninstall a Plug-In](#).

## Install or Update a vRealize Orchestrator Plug-In

You can install or update third-party plug-ins with the vRealize Orchestrator Control Center.

### Prerequisites

Download the `.dar` or `.vmoapp` file of the plug-in.

---

**Note** The preferred file format for vRealize Orchestrator plug-ins is `.vmoapp`.

---

### Procedure

- 1 Log in the Control Center as **root**.
- 2 Select the **Manage Plug-ins** page.
- 3 Click **Browse** and select the `.dar` or `.vmoapp` file of the plug-in you want to install or update.
- 4 Click **Upload**.
- 5 Review the plug-in information, if applicable, accept the end-user license agreement, and click **Install**.

The plug-in is installed or updated and the vRealize Orchestrator server service is restarted.



## What to do next

Verify that the correct plug-in information is listed on the **Manage Plug-ins** page.

## Uninstall a Plug-In

You can use the vRealize Orchestrator Control Center to delete a third-party plug-in. After deleting the plug-in from the Control Center, you must delete the associated package from the vRealize Orchestrator Client.

### Procedure

- 1 Log in to Control Center as **root**.
- 2 Select **Manage Plug-ins**.
- 3 Select the plug-in you want to uninstall and click the right-side delete icon.
- 4 Confirm that you want to delete the plug-in, and click **Delete**.
- 5 Log in to the vRealize Orchestrator Client as an **administrator**
- 6 Select **Assets > Packages**.
- 7 Browse to the package associated with the deleted plug-in and click **Delete**.

---

**Note** You can find the relevant package by entering a plug-in specific tag in the search text box. For example, you can find the Site Recovery Manager plug-in package by entering the **SRM** tag.

---

- 8 Select **Delete the package, its content, but keep any shared items** and click **Delete**.
- 9 Log in to the Orchestrator Legacy Client as an administrator.
- 10 From the **Tools** menu in the upper-right corner, select **User preferences**.
- 11 On the **General** page, select the **Delete non empty folder permitted** check box.
- 12 Click **Save and Close**.

You can now delete an entire folder, including its subfolders and workflows, with a single click.

- 13 Click the **Workflow** tab.
- 14 Delete the folder of the plug-in that you want to uninstall.
- 15 Click the **Actions** tab.
- 16 Delete the action modules of the plug-in that you want to uninstall.
- 17 Restart the vRealize Orchestrator services.

### Results

You removed all custom workflows, actions, policies, configurations, settings, and resources related to the plug-in.

## Orchestrator Availability and Scalability

To increase the availability of the Orchestrator services, start multiple Orchestrator server instances in a cluster with a shared database. vRealize Orchestrator works as a single instance until it is configured to work as part of a cluster.

### Orchestrator Cluster

Multiple Orchestrator server instances with identical server and plug-ins configurations work together in a cluster and share one database.

All Orchestrator server instances communicate with each other by exchanging heartbeats. Each heartbeat is a timestamp that the node writes to the shared database of the cluster at a certain time interval. Network problems, an unresponsive database server, or overload might cause an Orchestrator cluster node to stop responding. If an active Orchestrator server instance fails to send heartbeats within the failover timeout period, it is considered non-responsive. The failover timeout is equal to the value of the heartbeat interval multiplied by the number of the failover heartbeats. It serves as a definition for an unreliable node and can be customized according to the available resources and the production load.

An Orchestrator node enters standby mode when it loses connection to the database, and remains in this mode until the database connection is restored. The other nodes in the cluster take control of the active work, by resuming all interrupted workflows from their last unfinished items, such as scriptable tasks or workflow invocations.

Orchestrator does not provide a built-in tool for monitoring the cluster status and sending failover notifications. You can monitor the cluster state by using an external component such as a load balancer. To check whether a node is running, you can use the health status REST API service at [https://your\\_orchestrator\\_server\\_IP\\_or\\_DNS\\_name:8281/vco/api/healthstatus](https://your_orchestrator_server_IP_or_DNS_name:8281/vco/api/healthstatus) and check the status of the node or at [https://your\\_orchestrator\\_server\\_IP\\_or\\_DNS\\_name:8283/vco-controlcenter/docs/](https://your_orchestrator_server_IP_or_DNS_name:8283/vco-controlcenter/docs/) to monitor the status of Control Center.

### Configure a Cluster of vRealize Orchestrator Instances in VAMI

Starting with vRealize Orchestrator 7.5, all clustering operations are done through the VAMI interface of the Orchestrator Appliance.

An Orchestrator cluster consists of at least two Orchestrator instances that share one database. You configure a new Orchestrator cluster or add new nodes to an existing cluster from the Orchestrator VAMI interface. There are three types of nodes in the Orchestrator cluster.

Node Type	Definition
Primary node	Each Orchestrator cluster has one primary node. All nodes in the cluster share the PostgreSQL database of the primary node. The primary database can run in both asynchronous and synchronous modes. The primary node must be in a healthy state for the cluster to function.
Replica node	Replica nodes are Orchestrator instances joined to the primary Orchestrator node.
Synced Replica node	When you enable sync mode, a replica node is promoted to the state of synced replica node. The synced replica enables the automatic failover of the primary node.

### Prerequisites

- Configure at least two standalone server nodes. For more information, see [Configuring a Standalone Orchestrator Server](#).
- Synchronize the clocks of the virtual machines that the Orchestrator instances are installed on.
- Set up a load balancer to distribute the traffic among multiple Orchestrator instances.

### Procedure

- 1 Log in to the VAMI interface as **root**.

Access the VAMI interface at `https://your_orchestrator_server_ip_or_DNS_name:5480`.

- 2 Select the **Cluster** tab and enter the credentials of the Orchestrator node that is going to be the primary node of the cluster.

For existing clustered Orchestrator environments, enter the credentials of the primary node of the Orchestrator cluster.

- 3 Click **Join Cluster**.
- 4 Review the certificate information of the node and click **Ok**.
- 5 The clustering operation synchronizes the content of the Orchestrator nodes and joins the replica node to the PostgreSQL database of the primary node.

### What to do next

Verify that the cluster is configured properly at the **Validate Configuration** page of the Orchestrator Control Center.

---

**Note** Following the configuration of the cluster node, the Orchestrator server restarts automatically after 2 minutes. Verifying the configuration immediately after the process is finished can return an invalid cluster status.

---

## Monitoring an Orchestrator Cluster

After you create a cluster, you can monitor the states of the cluster nodes.

You can monitor the configuration synchronization states of the Orchestrator instances that are joined in a cluster from the **Orchestrator Cluster Management** page in Control Center.

Configuration Synchronization State	Description
RUNNING	The Orchestrator service is available and can accept requests.
STANDBY	The Orchestrator service cannot process requests because: <ul style="list-style-type: none"> <li>■ The node is part of a High Availability (HA) cluster and remains in a standby mode until the primary node fails.</li> <li>■ The service cannot verify the configuration prerequisites, such as a valid connection to the database, authentication provider, and the Orchestrator instance license.</li> </ul>
Failed to retrieve the service's health status	The Orchestrator server service cannot be contacted because it is either stopped or a network issue is present.
Pending restart	Control Center detects a configuration change and the Orchestrator server restarts automatically.

## Enable Sync Mode for the Orchestrator Cluster

You can configure an Orchestrator database cluster to run in synchronous mode.

Sync Mode enables the automatic failover of the primary Orchestrator database. The process promotes one of the replica nodes to the state of **synced replica**. If the current primary node fails, the synced replica is automatically promoted to a primary node. The synced replica receives all finished transactions from the database of the primary node.

### Prerequisites

Configure an Orchestrator cluster consisting of at least three Orchestrator nodes.

### Procedure

- 1 Log in to the VAMI interface as **root**.

Access the VAMI interface at `https://your_orchestrator_server_ip_or_DNS_name:5480`.

- 2 Select the **Cluster** tab.
- 3 Click **Sync Mode**.
- 4 One of the nodes of the cluster is promoted to the state of a **synced replica**.

To confirm the success of the synchronizing operation, check that the replication mode status in the **Cluster** tab is **Database is in Synchronous mode**.

## Promote an Orchestrator Replica Node to Primary Node

You can reconfigure an Orchestrator cluster by promoting a replica node to a primary node.

Orchestrator nodes can be promoted in both async mode and sync mode.

---

**Note** Orchestrator clusters in sync mode have an automatic failover function, so if the current primary node fails, the synced replica node automatically becomes the new primary node.

---

### Prerequisites

Configure an Orchestrator cluster consisting of at least two Orchestrator instances.

### Procedure

- 1 Log in to the VAMI interface as **root**.

Access the VAMI interface at `https://your_orchestrator_server_ip_or_DNS_name:5480`.

- 2 Select the **Cluster** tab.
- 3 Click **Promote** next to the replica node that you want to promote to the state of new primary node.
- 4 The message **Successfully promoted to new master node** appears top left of the VAMI UI and the state of the node changes to **MASTER**.

## Delete an Orchestrator Cluster Node

You can delete an Orchestrator replica node from your Orchestrator cluster, so you can replace it or reduce its capacity.

You can only delete replica nodes from the cluster. To remove a primary node, you must first promote a replica node to replace it. For more information, see [Promote an Orchestrator Replica Node to Primary Node](#).

### Procedure

- 1 Log in to the VAMI interface as **root**.

Access the VAMI interface at `https://your_orchestrator_server_ip_or_DNS_name:5480`.

- 2 Select the **Cluster** tab.
- 3 Select the **Delete** command next to the replica node.
- 4 Confirm that you want to delete the replica node from the cluster and click **Ok**.

---

**Note** You must remove the hostname of the deleted replica node from the load balancer server.

---

- 5 The Orchestrator node is deleted from the cluster and the message **Node was successfully deleted** appears on the top left of the UI.

# Configuring the Customer Experience Improvement Program

If you choose to participate in the Customer Experience Improvement Program (CEIP), VMware receives anonymous information that helps to improve the quality, reliability, and functionality of VMware products and services.

## Categories of Information that VMware Receives

The Customer Experience Improvement Program (CEIP) provides VMware with information that enables VMware to improve our products and services and to fix problems.

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set in the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>. To join or leave the CEIP for this product, see [Join the Customer Experience Improvement Program](#).

## Join the Customer Experience Improvement Program

Join the Customer Experience Improvement Program from Control Center.

### Procedure

- 1 Log in to Control Center as **root** and open the **Customer Experience Improvement Program** page.
- 2 Select the **Join the Customer Experience Improvement Program** check box to enable CEIP or deselect the check box to disable the Program and then click **Save**.
- 3 (Optional) Deselect the **Automatic proxy discovery** check box if you want to add a proxy host manually.

# Using the API services

# 6

In addition to configuring Orchestrator by using Control Center, you can modify the Orchestrator server configuration settings by using the Orchestrator REST API, the Control Center REST API, or the command line utility, stored in the appliance.

The Configuration plug-in is included by default in the Orchestrator package. You can access the Configuration plug-in workflows from either the Orchestrator workflow library or the Orchestrator REST API. With these workflows, you can change the trusted certificate and keystore settings of the Orchestrator server. For information on all available Orchestrator REST API services calls, see the *Orchestrator REST API Reference* documentation, located at [https://orchestrator\\_server\\_IP\\_or\\_DNS\\_name:8281/vco/api/docs](https://orchestrator_server_IP_or_DNS_name:8281/vco/api/docs).

- [Managing SSL Certificates and Keystores by Using the REST API](#)

In addition to managing SSL certificates by using Control Center, you can also manage trusted certificates and keystores when you run workflows from the Configuration plug-in or by using the REST API.

- [Automating the Orchestrator Configuration by Using the Control Center REST API](#)

The Control Center REST API provides access to resources for configuring the Orchestrator server. You can use the Control Center REST API with third-party systems to automate the Orchestrator configuration.

## Managing SSL Certificates and Keystores by Using the REST API

In addition to managing SSL certificates by using Control Center, you can also manage trusted certificates and keystores when you run workflows from the Configuration plug-in or by using the REST API.

The Configuration plug-in contains workflows for importing and deleting SSL certificates and keystores. You can access these workflows by navigating to **Library > Configuration > SSL Trust Manager** and **Library > Configuration > Keystores** in the Workflows view of the Orchestrator client. You can also run these workflows by using the Orchestrator REST API.

## Delete an SSL Certificate by Using the REST API

You can delete an SSL certificate by running the Delete trusted certificate workflow of the Configuration plug-in or by using the REST API.

### Procedure

- 1 Make a GET request at the URL of the Workflow service of the Delete trusted certificate workflow.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows?conditions=name=Delete trusted certificate
```

- 2 Retrieve the definition of the Delete trusted certificate workflow by making a GET request at the URL of the definition.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows/8a70a326-ffd7-4fef-97e0-2002ac49f5bd
```

- 3 Make a POST request at the URL that holds the execution objects of the Delete trusted certificate workflow.

```
POST https://{orchestrator_host}:{port}/vco/api/workflows/8a70a326-ffd7-4fef-97e0-2002ac49f5bd/executions/
```

- 4 Provide the name of the certificate you want to delete as an input parameter of the Delete trusted certificate workflow in an execution-context element in the request body.

## Import SSL Certificates by Using the REST API

You can import SSL certificates by running a workflow from the Configuration plug-in or by using the REST API.

You can import a trusted certificate from a file or a URL. For information about importing certificates in Orchestrator by using Control Center, see [Manage Orchestrator Certificates](#).

### Procedure

- 1 Make a GET request at the URL of the Workflow service.

Option	Description
<b>Import trusted certificate from a file</b>	Imports a trusted certificate from a file.
<b>Import trusted certificate from URL</b>	Imports a trusted certificate from a URL address.
<b>Import trusted certificate from URL using proxy server</b>	Imports a trusted certificate from a URL address by using a proxy server.
<b>Import trusted certificate from URL with certificate alias</b>	Imports a trusted certificate with a certificate alias, from a URL address.



To import a trusted certificate from a file, make the following GET request:

```
GET https://{orchestrator_host}:{port}/vco/api/workflows?conditions=name=Import
trusted certificate from a file
```

- 2 Retrieve the definition of the workflow by making a GET request at the URL of the definition.

To retrieve the definition of the Import trusted certificate from a file workflow, make the following GET request:

```
GET https://{orchestrator_host}:{port}/vco/api/workflows/93a7bb21-0255-4750-9293-2437abe9d2e5
```

- 3 Make a POST request at the URL that holds the execution objects of the workflow.

For the Import trusted certificate from a file workflow, make the following POST request:

```
POST https://{orchestrator_host}:{port}/vco/api/workflows/93a7bb21-0255-4750-9293-2437abe9d2e5/
executions
```

- 4 Provide values for the input parameters of the workflow in an execution-context element of the request body.

Parameter	Description
<b>cer</b>	The CER file from which you want to import the SSL certificate. This parameter is applicable for the Import trusted certificate from a file workflow.
<b>url</b>	The URL from which you want to import the SSL certificate. For non-HTTPS services, the supported format is <i>IP_address_or_DNS_name:port</i> . This parameter is applicable for the Import trusted certificate from URL workflow.

## Create a Keystore by Using the REST API

You can create a keystore by running the Create a keystore workflow of the Configuration plug-in or by using the REST API.

### Procedure

- 1 Make a GET request at the URL of the Workflow service of the Create a keystore workflow.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows?conditions=name=Create a keystore
```

- 2 Retrieve the definition of the Create a keystore workflow by making a GET request at the URL of the definition.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows/6c301bff-e8fe-4ae0-ad08-5318178594b3/
```

- 3 Make a POST request at the URL that holds the execution objects of the Create a keystore workflow.

```
POST https://{orchestrator_host}:{port}/vco/api/workflows/6c301bff-e8fe-4ae0-ad08-5318178594b3/
executions/
```

- 4 Provide the name of the keystore you want to create as an input parameter of the Create a keystore workflow in an execution-context element in the request body.

## Delete a Keystore by Using the REST API

You can delete a keystore by running the Delete a keystore workflow of the Configuration plug-in or by using the REST API.

### Procedure

- 1 Make a GET request at the URL of the Workflow service of the Delete a keystore workflow.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows?conditions=name=Delete a keystore
```

- 2 Retrieve the definition of the Delete a keystore workflow by making a GET request at the URL of the definition.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows/7a3389eb-1fab-4d77-860b-81b66bb45b86/
```

- 3 Make a POST request at the URL that holds the execution objects of the Delete a keystore workflow.

```
POST https://{orchestrator_host}:{port}/vco/api/workflows/7a3389eb-1fab-4d77-860b-81b66bb45b86/
executions/
```

- 4 Provide the keystore you want to delete as an input parameter of the Delete a keystore workflow in an execution-context element in the request body.

## Add a Key by Using the REST API

You can add a key by running the Add key workflow of the Configuration plug-in or by using the REST API.

### Procedure

- 1 Make a GET request at the URL of the Workflow service of the Add key workflow.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows?conditions=name=Add key
```

- 2 Retrieve the definition of the Add key workflow by making a GET request at the URL of the definition.

```
GET https://{orchestrator_host}:{port}/vco/api/workflows/6c301bff-e8fe-4ae0-ad08-5318178594b3/
```

- 3 Make a POST request at the URL that holds the execution objects of the Add key workflow.

```
POST https://{orchestrator_host}:{port}/vco/api/workflows/6c301bff-e8fe-4ae0-ad08-5318178594b3/
executions/
```

- 4 Provide the keystore, key alias, PEM-encoded key, certificate chain and key password as input parameters of the Add key workflow in an execution-context element in the request body.

## Automating the Orchestrator Configuration by Using the Control Center REST API

The Control Center REST API provides access to resources for configuring the Orchestrator server. You can use the Control Center REST API with third-party systems to automate the Orchestrator configuration.

The root endpoint of the Control Center REST API is `https://{orchestrator_server_IP_or_DNS_name}:8283/vco-controlcenter/api`. For information on all available service calls that you can make to the Control Center REST API, see the *Control Center REST API Reference* documentation, at `https://{orchestrator_server_IP_or_DNS_name}:8283/vco-controlcenter/docs`.

### Command-Line Utility

You can use the Orchestrator command-line utility to automate the Orchestrator configuration.

Access the command-line utility by logging in to the Orchestrator Appliance as root over SSH. The utility is located in `/var/lib/vco/tools/configuration-cli/bin`. To see the available configuration options, run `./vro-configure.sh --help`.

# Additional Configuration Options

# 7

You can use Control Center to change the default Orchestrator behavior.

This chapter includes the following topics:

- [Reconfiguring Authentication](#)
- [Export the Orchestrator Configuration](#)
- [Import the Orchestrator Configuration](#)
- [Configuring the Workflow Run Properties](#)
- [Orchestrator Log Files](#)
- [Add Network Interface Controllers](#)
- [Configure Static Routes](#)
- [Enabling the Opentracing and Wavefront Extensions](#)
- [Configure the Opentracing Extension](#)
- [Configure the Wavefront Extension](#)

## Reconfiguring Authentication

After you set up the authentication method during the initial configuration of Control Center, you can change the authentication provider or the configured parameters at any time.

### Change the Authentication Provider

To change the authentication mode or the authentication provider connection settings, you must unregister the existing authentication provider first.

#### Prerequisites

#### Procedure

- 1 Log in to Control Center as **root**.

- 2 On the **Configure Authentication Provider** page, click the **UNREGISTER** button next to the host address text box to unregister the authentication provider that is in use.
- 3 In the **IDENTITY SERVICE** section, click **UNREGISTER** to delete the server credentials.

### Results

You have successfully unregistered the authentication provider.

### What to do next

Reconfigure the authentication in Control Center. For more information, see [Configure a Standalone Orchestrator Server with vRealize Automation Authentication](#) or [Configure a Standalone Orchestrator Server with vSphere Authentication](#).

## Change the Authentication Parameters

When you use vRealize Automation as an authentication provider in Control Center, you might want to change the default tenant of the Orchestrator administrators group. When you use vSphere authentication, you can change the administrators group.

### Prerequisites

- Log in to Control Center as **root**.
- Select the authentication mode and configure the connection settings of the authentication provider.

### Procedure

- 1 Change the default tenant.

---

**Note** You can change the default tenant only if you use the vRealize Automation authentication mode.

---

- a On the **Configure Authentication Provider** page in Control Center, click the **CHANGE** button next to the **Default tenant** text box.
- b In the text box, replace the name of the existing default tenant with the one you want to use.
- c Click the **CHANGE** button next to the **Admin group** text box.

---

**Note** If you do not reconfigure the administrators group, it remains empty and you are no longer able to access Control Center.

---

- d Enter the name of an administrators group and click **SEARCH**.
- e In the list of groups, double-click on the name of the group to select it.
- f Click **SAVE CHANGES**.

You are logged out of Control Center and redirected to the Single Sign-On login screen.

- 2 Change the administrators group.
  - a Click the **CHANGE** button next to the **Admin group** text box.
  - b Enter the name of an administrators group and click **SEARCH**.
  - c In the list of groups, double-click on the name of the group to select it.
  - d Click **SAVE CHANGES**.

You are logged out of Control Center and redirected to the Single Sign-On login screen.

## Export the Orchestrator Configuration

Control Center provides a mechanism to export the Orchestrator configuration settings to a local file. You can use the mechanism to take a snapshot of your system configuration at any moment and import this configuration into a new Orchestrator instance.

You should export and save your configuration settings regularly, especially when making modifications, performing maintenance tasks, or upgrading the system.

---

**Important** Keep the file with the exported configuration safe and secure, because it contains sensitive administrative information.

---

### Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **Export/Import Configuration**.
- 3 Select the type of files you want to export.

---

**Note** If you select **Export plug-in configurations** and the plug-in configurations contain encrypted properties, you must also select **Export server configuration** to successfully decrypt the data when importing.

---

- 4 (Optional) Enter a password to protect the configuration file.  
Use the same password when you import the configuration later.
- 5 Click **Export**.

### Results

Orchestrator creates an `orchestrator-config-export-hostname-dateReference.zip` file that is downloaded on your local machine. You can use this file to clone or to restore the system.

## Import the Orchestrator Configuration

You can restore a previously exported system configuration after you reinstall Orchestrator or if a system failure occurs.

If you use the import procedure to clone the Orchestrator configuration, the vCenter Server plug-in configuration becomes invalid and does not work, because a new vCenter Server plug-in ID is generated.

#### Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **Export/Import Configuration** and navigate to the **Import Configuration** tab.
- 3 Browse to and select the .zip file that you exported from your previous installation.

---

**Note** The default syntax for the exported configuration file is `orchestrator-config-export-hostname-dateofexport_timeofexport.zip`

---

- 4 (Optional) Enter the password that you used when exporting the configuration.  
This step is not necessary if you have not exported the configuration with a password.
- 5 Select the import type:

Option	Description
<b>Embedded</b>	Migrates to an Orchestrator instance that is embedded in vRealize Automation.
<b>External</b>	Migrates to an external Orchestrator.
<b>Replica</b>	Replicates the same Orchestrator instance.

- 6 Click **Import**.

#### Results

The new system replicates the old configuration, based on the selected import type. The Orchestrator server service restarts automatically.

#### What to do next

Verify that Orchestrator is configured properly at the **Validate Configuration** page in Control Center.

## Configuring the Workflow Run Properties

By default, you can run up to 300 workflows per node and up to 10,000 workflows can be queued if the number of actively running workflows is reached.

When the Orchestrator node has to run more than 300 concurrent workflows, the pending workflow runs are queued. When an active workflow run completes, the next workflow in the queue starts to run. If the maximum number of queued workflows is reached, the next workflow runs fail until one of the pending workflows starts to run.

On the **Advanced Options** page in Control Center, you can configure the workflow run properties.

Option	Description
<b>Enable safe mode</b>	If safe mode is enabled, all running workflows are canceled and are not resumed on the next Orchestrator node start.
<b>Number of concurrent running workflows</b>	The maximum number of concurrent Orchestrator node workflows that run simultaneously.
<b>Maximum amount of running workflows in the queue</b>	The number of workflow run requests that the Orchestrator node accepts before becoming unavailable.
<b>Maximum number of preserved runs per workflow</b>	The maximum number of finished workflow runs kept as history per workflow in a cluster. If the number is exceeded, the oldest workflow runs are deleted.
<b>Log events expiration days</b>	The number of days log events for the cluster are kept in the database before being purged.
<b>Profile all workflow runs</b>	Enable and disable automatic workflow profiling. When enabled, workflow profiling generates metric data on every workflow run.
<b>Interval for redistributing the Workflow Profiler statistics</b>	The interval during which Profiler statistics will be distributed to every Orchestrator instance in your environment.

## Orchestrator Log Files

VMware Technical Support routinely requests diagnostic information when you submit a support request. This diagnostic information contains product-specific logs and configuration files from the host on which the product runs.

You can download a zip bundle that includes the Orchestrator configuration files and log files from the **Export Logs** menu in Control Center.

**Table 7-1. Orchestrator Log Files list**

Filename	Location	Description
scripting.log	/var/log/vco/app-server	Provides scripting log messages of workflows and actions. Use the <code>scripting.log</code> file to isolate workflow runs and action runs from normal Orchestrator operations. This information is also included in the <code>server.log</code> file.
server.log	/var/log/vco/app-server	Provides information about all activities on the Orchestrator server. Analyze the <code>server.log</code> file when you debug Orchestrator or any application that runs on Orchestrator.
metrics.log	/var/log/vco/app-server	Contains runtime information about the server. The information is added to this log file once every 5 minutes.
localhost_access_log.txt	/var/log/vco/app-server	Contains the HTTP request log of the server.
localhost_access_log.date.txt	/var/log/vco/configuration	Contains the HTTP request log of the Control Center service.
controlcenter.log	/var/log/vco/configuration	The log file of the Control Center service.



## Logging Persistence

You can log information in any kind of Orchestrator script, for example workflow, policy, or action. This information has types and levels. The type can be either persistent or non-persistent. The level can be DEBUG, INFO, WARN, ERROR, TRACE, and FATAL.

**Table 7-2. Creating Persistent and Non-Persistent Logs**

Log Level	Persistent Type	Non-Persistent Type
DEBUG	<code>Server.debug("short text", "long text");</code>	<code>System.debug("text")</code>
INFO	<code>Server.log("short text", "long text");</code>	<code>System.log("text");</code>
WARN	<code>Server.warn("short text", "long text");</code>	<code>System.warn("text");</code>
ERROR	<code>Server.error("short text", "long text");</code>	<code>System.error("text");</code>

### Persistent Logs

Persistent logs (server logs) track past workflow run logs and are stored in the Orchestrator database. To view server logs, you must select a workflow, a completed workflow run, or a policy and click the **Events** tab in the Orchestrator client.

### Non-Persistent Logs

When you use a non-persistent log (system log) to create scripts, the Orchestrator server notifies all running Orchestrator applications about this log, but this information is not stored in the database. When the application is restarted, the log information is lost. Non-persistent logs are used for debugging purposes and for live information. To view system logs, you must select a completed workflow run in the Orchestrator client and click **Logs** on the **Schema** tab.

## Orchestrator Logs Configuration

On the **Configure Logs** page in Control Center, you can set the level of server log and the scripting log that you require. If either of the logs is generated multiple times a day, it becomes difficult to determine what causes problems.

The default log level of the server log and the scripting log is INFO. Changing the log level affects all new messages that the server enters in the logs and the number of active connections to the database. The logging verbosity decreases in descending order.

**Caution** Only set the log level to DEBUG or ALL to debug a problem. Do not use these settings in a production environment because it can seriously impair performance.

### Log Rotation Settings

To prevent the server log from becoming too large, you can set the maximum file size and count of the server log by modifying the values in the **Max file count** and **Max file size (MB)** text boxes.

## Orchestrator Log Files Export

From the **Export Logs** page in Control Center, you can generate a ZIP archive of troubleshooting information containing configuration, server, wrapper, and installation log files.

The log information is stored in a ZIP archive named `vco-logs-date_hour.zip`.

---

**Note** When you have more than one Orchestrator instance in a cluster, the ZIP archive includes the logs from all Orchestrator instances in the cluster.

---

## Filter the Orchestrator Logs

You can filter the Orchestrator server logs for a specific workflow run and collect diagnostic data about the workflow run.

The Orchestrator logs contain a lot of useful information which you can monitor in real time. When multiple instances of the same workflow are running at the same time, you can track the different workflow runs by filtering the diagnostic data about each run in the Orchestrator live log stream.

---

**Note** When you have more than one Orchestrator instance in a cluster, the live log stream shows only the logs of the local Orchestrator node.

---

### Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **Live Log Stream**.
- 3 In the search bar, enter your search parameters.  
For example, you can filter the logs by a user name, workflow name, workflow ID, or a token ID.
- 4 (Optional) Select **Case sensitive** and **Filter (grep)** to filter the search results further.  
By selecting **Filter (grep)** the live stream only shows the lines that match your search parameters.

### Results

The Orchestrator live log stream is filtered according to your search parameters.

### What to do next

You can use third-party log analyzing tools, if you want to filter old logs, that are not accessible through the **Live Log Stream** page in Control Center.

## Configuring Logging Integration with Remote Server

You can configure Orchestrator to send logs to remote logging systems like vRealize Log Insight or other Syslog servers.

**Procedure**

- 1 Log in to Control Center as **root**.
- 2 Navigate to the **Logging Integration** menu.
- 3 Turn on **Enable logging to a remote log server**.
- 4 Configure the logging integration options.
  - a Select logging system type.
  - b Enter the host name and port value of the remote logging server.
  - c Select the protocol used to send log events to the remote logging server.
- 5 To finish configuring logging integration to the remote server, click **Save**.

## Add Network Interface Controllers

vRealize Orchestrator supports multiple network interface controllers (NICs). After installation, you can add NICs to the Orchestrator appliance.

**Prerequisites**

Completely install vRealize Orchestrator to your vCenter Server environment.

**Procedure**

- 1 In vCenter Server, add NICs to each vRealize Orchestrator appliance.
  - a Right click the appliance and select **Edit Settings**.
  - b Add VMXNET3 NICs.
  - c If it is powered on, restart the appliance.
- 2 Log in to the vRealize Orchestrator appliance management interface as root.  
<https://orchestrator-appliance-IP:5480>
- 3 Select **Network**, and verify that multiple NICs are available.
- 4 Select **Address**, and configure the IP address for the NICs.

**Table 7-3. Example NIC Configuration**

Setting	Value
IPv4 Address Type	Static
IPv4 Address	172.22.0.2
Netmask	255.255.255.0

- 5 Click **Save Settings**.

## Configure Static Routes

When adding NICs to a vRealize Orchestrator installation, if you need static routes, open a command prompt session to configure them.

### Prerequisites

Add multiple NICs to the vRealize Orchestrator appliances.

### Procedure

- 1 Log in to the vRealize Orchestrator appliance command line as root.
- 2 Open the routes file in a text editor.  
`/etc/sysconfig/network/routes`
- 3 Locate the default line for the default gateway, but do not modify it.

---

**Note** In cases where the default gateway needs to change, use the vRealize Orchestrator management interface instead.

---

- 4 Below the default line, add new lines for static routes. For example:

```
default 10.10.10.1 - -
172.30.30.0 192.168.100.1 255.255.255.0 eth0
192.168.210.0 192.168.230.1 255.255.255.0 eth2
```

- 5 Save and close the routes file.
- 6 Restart the appliance.
- 7 In HA clusters, repeat the process for each appliance.

## Enabling the Opentracing and Wavefront Extensions

The Opentracing and Wavefront extensions for vRealize Orchestrator provide tools for gathering data about your vRealize Orchestrator environment . You can use this data for troubleshooting the vRealize Orchestrator system and workflows.

Before you can configure vRealize Orchestrator to use the Opentracing and Wavefront extensions, you must enable them in the vRealize Orchestrator Appliance.

### Prerequisites

Verify that the vRealize Orchestrator Appliance SSH service is enabled. You can verify the SSH service status on the **Admin** tab of the vRealize Orchestrator VAMI.

### Procedure

- 1 Log in to the vRealize Orchestrator Appliance with an SSH client as **root**.  
Access the vRealize Orchestrator Appliance at `https://your_orchestrator_hostname:5480`.

- 2 Navigate to the extensions directory.

```
cd /var/lib/vco/app-server/extensions
```

- 3 Run the `ls` command to list all available vRealize Orchestrator extensions.

- 4 Enable Wavefront and Opentracing extensions.

```
mv opentracing-7.6.0.jar.inactive opentracing-7.6.0.jar
```

```
mv wavefront-7.6.0.jar.inactive wavefront-7.6.0.jar
```

- 5 Run the `ls` command again and confirm that the extensions do not end with `.inactive`.

#### What to do next

Configure Opentracing and Wavefront integration with vRealize Orchestrator in the **Extension Properties** page of the Control Center. For more information, see [Configure the Opentracing Extension](#) and [Configure the Wavefront Extension](#).

## Configure the Opentracing Extension

The Opentracing extension sends data about workflow runs to a Jaeger server . Data includes the workflow status, input and output parameters, the user that initiated the workflow run, and the workflow ID data.

#### Prerequisites

- Verify sure that Opentracing is enabled in the Orchestrator Appliance. For more information, see [Enabling the Opentracing and Wavefront Extensions](#).
- Deploy a Jaeger server for use in the Orchestrator Opentracing Extension. For more information, see the [Getting Started with Jaeger documentation](#).

#### Procedure

- 1 Log in to the Orchestrator Control Center as **root**.
- 2 Navigate to the **Extension Properties** page.
- 3 Select the Opentracing extension.
- 4 Enter the Jaeger server host address and port.

---

**Note** Insert two forward slashes ("/") before entering the server address.

---

- 5 Click **Save**.

#### Results

You have configured the Orchestrator Opentracing Extension.

**What to do next**

- To access the Jaeger UI containing the data collected by the Opentracing extension, visit the host address entered during configuration.
- Under the **Service** option, select **Workflows**.
- To specify what data to view, use the **Tags** option. For example, to view data about failed workflows, enter **status=failed**.

## Configure the Wavefront Extension

Use the Wavefront extension to gather metric data about your Orchestrator system and workflows.

**Prerequisites**

- 1 Verify that Wavefront is enabled in the Orchestrator Appliance. For more information, see [Enabling the Opentracing and Wavefront Extensions](#).
- 2 Import the Wavefront Certificate:
  - a Log in to the Orchestrator Control Center as **root**.
  - b Navigate to the **Certificates** page.
  - c Click the **Import** drop-down menu and select **Import from URL**.
  - d Enter the Wavefront URL and click **Import**.
- 3 Configure a Wavefront Proxy. For more information, see [Installing and Managing Wavefront Proxies](#).

**Procedure**

- 1 Log in to the Orchestrator Control Center as **root**.
- 2 Navigate to the **Extension Properties** page.
- 3 Select the Wavefront extension.
- 4 Configure the Wavefront properties.

Option	Description
<b>Proxy</b>	The Wavefront proxy address.
<b>Host</b>	Optional. The Wavefront host address.
<b>Token</b>	Optional. The Wavefront API token. For more information on generating a Wavefront API token, see <a href="#">Generating an API Token</a> .
<b>Prefix</b>	Add prefix labels for each metric sent to Wavefront. Prefix labels are separated by a dot symbol.

- 5 (Optional) Select **Send default dashboard on next start**.

**6** Click **Save**.

**Results**

You have configured the Orchestrator Wavefront Extension.

**What to do next**

- To access the metrics collected by Wavefront, access the dashboard on the address entered during configuration.
- To get notifications about specific events in your Orchestrator environment, you can use Wavefront Alerts. For more information, see the [Wavefront Alerts documentation](#).

# Configuration Use Cases and Troubleshooting



You can configure the Orchestrator server to work with the vCenter Server appliance, you can also uninstall plug-ins from Orchestrator, or change the self-signed certificates.

The configuration use cases provide task flows that you can perform to meet specific configuration requirements of your Orchestrator server, as well as troubleshooting topics to understand and solve a problem, if a workaround exists.

This chapter includes the following topics:

- [Configure vRealize Orchestrator Plugin for vSphere Web Client](#)
- [Unregister Orchestrator Authentication](#)
- [Changing SSL Certificates](#)
- [Cancel Running Workflows](#)
- [Enable Orchestrator Server Debugging](#)
- [Back Up the Orchestrator Configuration and Elements](#)
- [Backing Up and Restoring vRealize Orchestrator](#)
- [Disaster Recovery of Orchestrator by Using Site Recovery Manager](#)

## Configure vRealize Orchestrator Plugin for vSphere Web Client

To use the vRealize Orchestrator plug-in for vSphere Web Client, you must register vRealize Orchestrator as an extension of vCenter Server.

After you register your vRealize Orchestrator server with vCenter Single Sign-On and configure it to work with vCenter Server, you must register vRealize Orchestrator as an extension of vCenter Server.

### Prerequisites

You must register vRealize Orchestrator with vSphere authentication to the same Platform Services Controller that your managed vCenter Server authenticates with.



### Procedure

- 1 Log in to the vRealize Orchestrator Client.
- 2 Navigate to **Library > Workflows**.
- 3 Search for the **Register vCenter Orchestrator as a vCenter Server extension** workflow and click **Run**.
- 4 Select the vCenter Server instance to register vRealize Orchestrator with.
- 5 (Optional) Enter `https://your_orchestrator_hostname:8281` or the service URL of the load balancer that redirects the requests to the vRealize Orchestrator server nodes.
- 6 Click **Run**.

## Unregister Orchestrator Authentication

Unregister Orchestrator as a Single Sign-On solution from the Configure Authentication Provider page in Control Center.

If you want to reconfigure the Orchestrator vCenter Single Sign-On or vRealize Automation authentication you must first unregister the Orchestrator authentication.

### Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **Configure Authentication Provider**.
- 3 Click **Unregister**.
- 4 (Optional) Enter your credentials if you want to delete registration data from the identity server.
- 5 Click **Unregister** from the **Identity service** section.

### Results

You have successfully unregistered your Orchestrator server instance.

## Changing SSL Certificates

By default, the Orchestrator server uses a self-signed SSL certificate to communicate remotely with the Orchestrator client. You can change the SSL certificates if, for example, your company security policy requires you to use its SSL certificates.

When you attempt to use Orchestrator over a trusted SSL Internet connection, and you open Control Center in a Web browser, you receive a warning that the connection is untrusted, if you use Mozilla Firefox, or that problems have been detected with the Web site's security certificate, if you use Internet Explorer.

After you click **Continue to this website (not recommended)**, even if you have imported the SSL certificate in the trusted store, you continue to see the Certificate Error red notification in the address bar of the Web browser. You can work with Orchestrator in the Web browser, but a third-party system might not work properly when attempting to access the API over HTTPS.

You might also receive a certificate warning when you start the Orchestrator client and attempt to connect to the Orchestrator server over an SSL connection.

You can resolve the problem by installing a certificate signed by a commercial certificate authority (CA). To stop receiving a certificate warning from the Orchestrator client, add your root CA certificate to the Orchestrator keystore on the machine on which the Orchestrator client is installed.

## Adding a Certificate to the Local Store

After you receive a certificate from a CA, you must add the certificate to your local storage to work with Control Center without receiving certificate warnings or error messages.

This workflow describes the process of adding the certificate to your local storage by using Internet Explorer.

- 1 Open Internet Explorer and go to `https://orchestrator_server_IP_or_DNS_name:8283/`.
- 2 When prompted, click **Continue to this website (not recommended)**.  
The certificate error appears on the right side of the address bar in Internet Explorer.
- 3 Click the Certificate Error and select **View Certificates**.
- 4 Click **Install Certificate**.
- 5 On the Welcome page of the **Certificate Import Wizard**, click **Next**.
- 6 In the **Certificate Store** window, select **Place all certificates in the following store**.
- 7 Browse and select **Trusted Root Certification Authorities**.
- 8 Complete the wizard and restart Internet Explorer.
- 9 Navigate to the Orchestrator server over your SSL connection.

You no longer receive warnings, and you do not receive a Certificate Error in the address bar.

Other applications and systems, such as VMware Service Manager, must have access to the Orchestrator REST APIs through an SSL connection.

## Change the Certificate of the Orchestrator Appliance Management Site

The Orchestrator Appliance uses Light HTTPd to run its own management site. You can change the SSL certificate of the Orchestrator Appliance management site if, for example, your company security policy requires you to use its SSL certificates.

## Prerequisites

By default the Orchestrator Appliance SSL certificate and private key are stored in a PEM file, which is located at: `/opt/vmware/etc/lighttpd/server.pem`. To install a new certificate, ensure that you export your new SSL certificate and private key from the Java keystore to a PEM file.

## Procedure

- 1 Log in to the Orchestrator Appliance Linux console as `root`.
- 2 Locate the `/opt/vmware/etc/lighttpd/lighttpd.conf` file and open it in an editor.
- 3 Find the following line:

```
#### SSL engine
ssl.engine = "enable"
ssl.pemfile = "/opt/vmware/etc/lighttpd/server.pem"
```

- 4 Change the `ssl.pemfile` attribute to point to the PEM file containing your new SSL certificate and private key.
- 5 Save the `lighttpd.conf` file.
- 6 Run the following command to restart the `light-httpd` server.

```
service vami-lighttpd restart
```

## Results

You successfully changed the certificate of the Orchestrator Appliance management site.

# Cancel Running Workflows

You can use Control Center to cancel workflows that do not finish properly.

## Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **Troubleshooting**.
- 3 Cancel running workflows.

Option	Description
<b>Cancel all workflow runs</b>	Enter a workflow ID, to cancel all tokens for that workflow.
<b>Cancel workflow runs by ID</b>	Enter all token IDs, you want to cancel. Separate IDs with a comma.
<b>Cancel all running workflows</b>	Cancel all running workflows on the server.

**Note** Operations where you cancel workflows by ID might not be successful, as there is no reliable way to cancel the run thread immediately.

## Results

On the next server start, the workflows are set in a canceled state.

## What to do next

Verify that the workflows are canceled from the **Inspect Workflows** page in Control Center.

# Enable Orchestrator Server Debugging

You can start the Orchestrator server in debug mode to debug issues when developing a plug-in.

## Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **Orchestrator Debugging**.
- 3 Click **Enable debugging**.
- 4 (Optional) Enter a port, different from the default one.
- 5 (Optional) Click **Suspend**.  
By selecting this option, you must attach a debugger before starting the Orchestrator server.
- 6 Click **Save**.
- 7 Open the Startup Options page in Control Center and click **Restart**.

## Results

The Orchestrator server is suspended upon start until you attach a remote Java debugger to the defined port.

# Back Up the Orchestrator Configuration and Elements

Back up your custom Orchestrator server configuration and workflow elements, to ensure their reusability by other Orchestrator instances.

If you edit any standard workflows, actions, policies, or configuration elements, and then import a package containing the same elements with a later Orchestrator version number, your changes to the elements are lost. You can prevent the loss of customized workflows and other elements by exporting them before you migrate your Orchestrator instance.

Each Orchestrator server instance has unique certificates, and each vCenter Server plug-in instance has a unique ID. The certificates and the unique ID define the identity of the Orchestrator server and the vCenter Server plug-in. If you do not back up the Orchestrator elements or export the Orchestrator configuration for backup purposes, make sure that you change these identifiers.

## Prerequisites

Deploy and configure a new Orchestrator server instance. See [Configuring a Standalone Orchestrator Server](#).

## Procedure

- 1 Export the Orchestrator configuration.
  - a Log in to Control Center as **root**.
  - b Click **Export/Import Configuration**.
  - c Select the types of files you want to export.
  - d (Optional) Protect the configuration file by entering a password.  
Use the same password when you import the configuration.
  - e Click **Export**.
- 2 Log in to the Orchestrator client application.
- 3 Create a package that contains all the Orchestrator elements that you created or edited.
  - a Under the **Administer** view, click the **Packages** tab.
  - b Click the menu button in the title bar of the Packages list and select **Add package**.
  - c Enter a name for the new package and click **OK**.  
The syntax for package names is *domain.your\_company.folder.package\_name*.  
For example, `com.vmware.myfolder.mypackage`.
  - d Right-click the package and select **Edit**.
  - e On the **General** tab, add a description for the package.
  - f On the **Workflows** tab, add workflows to the package.
  - g (Optional) Add policy templates, actions, configuration elements, resource elements, access rights, and plug-ins to the package.
  - h Click **Save and close**.
- 4 Export the package.
  - a Right-click the package you want to export and select **Export package**.
  - b Browse to and select a location where you want to save the package.
  - c (Optional) Use the corresponding certificate to sign the package.
  - d (Optional) Impose restrictions on the exported package.

- e (Optional) To apply restrictions for the contents of the exported package, deselect the options as required.

Option	Description
<b>Export version history</b>	The version history of the package is not exported.
<b>Export the values of the configuration settings</b>	The attribute values of the configuration elements in the package are not exported.
<b>Export global tags</b>	The global tags in the package are not exported.

---

**Note** The **Export the values of the configuration SecureString settings** option is deselected by default. Export of these configuration settings can cause a security problem. Use with caution.

---

- f Click **Save**.
- 5** Import the Orchestrator configuration that you exported earlier to the new Orchestrator server instance.
- a Log in to Control Center of the new Orchestrator instance as **root**.
  - b Click **Export/Import Configuration** and navigate to the **Import Configuration** tab.
  - c Browse to select the .zip file you exported from your previous installation.
  - d Type the password you used while exporting the configuration.  
This step is not necessary if you have not specified a password.
  - e Select the import type.
  - f Click **Import**.
- 6** Import the package that you exported to the new Orchestrator instance.
- a Log in to the Orchestrator client application of the new Orchestrator instance.
  - b From the drop-down menu in the Orchestrator client, select **Administer**.
  - c Click the **Packages** tab.
  - d Click the menu button in the title bar of the Packages list and select **Import package**.
  - e Browse to and select the package that you want to import and click **Open**.  
Certificate information about the exporter appears.
  - f Review the package import details and select **Import** or **Import and trust provider**.  
The Import package view appears. If the version of the imported package element is later than the version on the server, the system selects the element for import automatically.
  - g Select the elements that you want to import.

---

**Note** Deselect custom elements for which later versions exist.

---

- h (Optional) Deselect the **Import the values of the configuration settings** check box if you do not want to import the attribute values of the configuration elements from the package.
- i From the drop-down menu, select whether you want to import tags from the package.

Option	Description
<b>Import tags but preserve existing values</b>	Import tags from the package without overwriting existing tag values.
<b>Import tags and overwrite existing values</b>	Import tags from the package and overwrite their values.
<b>Do not import tags</b>	Do not import tags from the package.

- j Click **Import selected elements**.

## Results

You have successfully backed up the Orchestrator configuration and elements.

# Backing Up and Restoring vRealize Orchestrator

You can use vSphere Data Protection to back up and restore a virtual machine (VM) that contains a vRealize Orchestrator instance.

vSphere Data Protection is a VMware disk-based backup and recovery solution designed for vSphere environments. vSphere Data Protection is fully integrated with vCenter Server. With vSphere Data Protection, you can manage backup jobs and store backups in deduplicated destination storage locations. After you deploy and configure vSphere Data Protection, you can access vSphere Data Protection by using the vSphere Web Client interface to select, schedule, configure, and manage backups and recoveries of virtual machines. During a backup, vSphere Data Protection creates a quiesced snapshot of the virtual machine. Deduplication is automatically performed with every backup operation.

For information about how to deploy and configure vSphere Data Protection, see the *vSphere Data Protection Administration* documentation.

## Back Up vRealize Orchestrator

You can back up your vRealize Orchestrator instance as a virtual machine.

To ensure that all components of a VM in a single product are backed up together, store the VMs of your vRealize Orchestrator environment in a single vCenter Server folder and create a backup policy job for that folder.

### Prerequisites

- Verify that the vSphere Data Protection appliance is deployed and configured. For information about how to deploy and configure vSphere Data Protection, see the *vSphere Data Protection Administration* documentation.

- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.

#### Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Select your vSphere Data Protection appliance from the **VDP appliance** drop-down menu and click **Connect**.
- 3 On the **Getting Started** tab, click **Create Backup Job**.
- 4 Click **Guest Images** to back up your vRealize Orchestrator instance and click **Next**.
- 5 Select **Full Image** to back up the entire virtual machine and click **Next**.
- 6 Expand the **Virtual Machines** tree and select the check box of your vRealize Orchestrator VM.
- 7 Follow the prompts to set the backup schedule, retention policy, and name of the backup job.

For more information about how to back up and restore virtual machines, see the *vSphere Data Protection Administration* documentation.

Your backup job appears in the list of backup jobs on the **Backup** tab.

- 8 (Optional) Open the **Backup** tab, select your backup job and click **Backup now** to back up your vRealize Orchestrator.

---

**Note** Alternatively, you can wait for the backup to start automatically according to the schedule that you set.

---

The backup process appears on the **Recent Tasks** page.

#### Results

The image of your VM appears in the list of backups on the **Restore** tab.

#### What to do next

Open the **Restore** tab and verify that the image of your VM is in the list of backups.

## Restore a vRealize Orchestrator Instance

You can restore your vRealize Orchestrator instance on its original location or on a different location on the same vCenter Server.

#### Prerequisites

- Verify that the vSphere Data Protection appliance is deployed and configured. For information about how to deploy and configure vSphere Data Protection, see the *vSphere Data Protection Administration* documentation.
- Back up your vRealize Orchestrator instance. See [Back Up vRealize Orchestrator](#).



- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that you used during the vSphere Data Protection configuration.

#### Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Select your vSphere Data Protection appliance from the **VDP appliance** drop-down menu and click **Connect**.
- 3 Open the **Restore** tab.
- 4 From the list of backup jobs, select the vRealize Orchestrator backup that you want to restore.

---

**Note** If you have multiple VMs, you must restore them simultaneously so that they are synchronized.

---

- 5 To restore your vRealize Orchestrator instance on the same vCenter Server, click the **Restore** icon and follow the prompts to set the location on your vCenter Server where to restore your vRealize Orchestrator.

Do not select **Power On**, as the appliance must be the last component to be powered on. For information about how to back up and restore a virtual machine, see the *vSphere Data Protection Administration* documentation.

A message that states that the restore is successfully initiated appears.

- 6 (Optional) Power on your database hosts if they are external and restore your load balancer configuration.
- 7 Power on the vRealize Orchestrator Appliance.

#### Results

The restored vRealize Orchestrator VM appears in the vCenter Server inventory.

#### What to do next

Verify that vRealize Orchestrator is configured properly by opening the **Validate Configuration** page in Control Center.

## Disaster Recovery of Orchestrator by Using Site Recovery Manager

You must configure Site Recovery Manager to protect your vRealize Orchestrator. Secure this protection by completing the common configuration tasks for Site Recovery Manager.

## Prepare the Environment

You must ensure that you meet the following prerequisites before you start configuring Site Recovery Manager.

- Verify that vSphere 5.5 is installed on the protected and recovery sites.
- Verify that you are using Site Recovery Manager 5.8.
- Verify that vRealize Orchestrator is configured.

## Configure Virtual Machines for vSphere Replication

You must configure the virtual machines for vSphere Replication or array based replication in order to use Site Recovery Manager.

To enable vSphere Replication on the required virtual machines, perform the following steps.

### Procedure

- 1 In the vSphere Web Client, select a virtual machine on which vSphere Replication should be enabled and click **Actions > All vSphere Replication Actions > Configure Replication**.
- 2 In the **Replication type** window, select **Replicate to a vCenter Server** and click **Next**.
- 3 In the **Target site** window, select the vCenter for the recovery site and click **Next**.
- 4 In the **Replication server** window, select a vSphere Replication server and click **Next**.
- 5 In the **Target location** window, click **Edit** and select the target datastore, where the replicated files will be stored and click **Next**.
- 6 In the **Replication options** window, keep the default setting and click **Next**.
- 7 In the **Recovery settings** window, enter time for **Recovery Point Objective (RPO)** and **Point in time instances**, and click **Next**.
- 8 In the **Ready to complete** window, verify the settings and click **Finish**.
- 9 Repeat these steps for all virtual machines on which vSphere Replication must be enabled.

## Create Protection Groups

You create protection groups to enable Site Recovery Manager to protect virtual machines.

When you create protection groups, wait to ensure that the operations finish as expected. Make sure that Site Recovery Manager creates the protection group and that the protection of the virtual machines in the group is successful.

### Prerequisites

Verify that you performed one of the following tasks:

- Included virtual machines in datastores for which you configured array-based replication
- Configured vSphere Replication on virtual machines

- Performed a combination of some or all of the above

#### Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Protection Groups**.
- 2 On the **Objects** tab, click the icon to create a protection group.
- 3 On the Protection group type page, select the protected site, select the replication type, and click **Next**.

Option	Action
<b>Array-based replication groups</b>	Select <b>Array Based Replication (ABR)</b> and select an array pair.
<b>vSphere Replication protection group</b>	Select <b>vSphere Replication</b> .

- 4 Select datastore groups or virtual machines to add to the protection group.

Option	Action
<b>Array-based replication protection groups</b>	Select datastore groups and click <b>Next</b> .
<b>vSphere Replication protection groups</b>	Select virtual machines from the list, and click <b>Next</b> .

When you create vSphere Replication protection groups, only virtual machines that you configured for vSphere Replication and that are not already in a protection group appear in the list.

- 5 Review your settings and click **Finish**.

You can monitor the progress of the creation of the protection group on the **Objects** tab under **Protection Groups**.

#### Results

- If Site Recovery Manager successfully applied inventory mappings to the protected virtual machines, the protection status of the protection group is OK.
- If Site Recovery Manager successfully protected all of the virtual machines associated with the storage policy, the protection status of the protection group is OK.

## Create a Recovery Plan

You create a recovery plan to establish how Site Recovery Manager recovers virtual machines.

#### Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**.
- 2 On the **Objects** tab, click the icon to create a recovery plan.
- 3 Enter a name and description for the plan, select a folder, then click **Next**.

- 4 Select the recovery site and click **Next**.
- 5 Select the group type from the menu.

Option	Description
<b>VM protection groups</b>	Select this option to create a recovery plan that contains array-based replication and vSphere Replication protection groups.
<b>Storage policy protection groups</b>	Select this option to create a recovery plan that contains storage policy protection groups.

The default is **VM protection groups**.

**Note** If using stretched storage, select **Storage policy protection groups** for the group type.

- 6 Select one or more protection groups for the plan to recover, and click **Next**.
- 7 Click the **Test Network** value, select a network to use during test recovery, and click **Next**.  
The default option is to create an isolated network automatically.
- 8 Review the summary information and click **Finish** to create the recovery plan.

## Organize Recovery Plans in Folders

You can create folders in which to organize recovery plans.

Organizing recovery plans into folders is useful if you have many recovery plans. You can limit the access to recovery plans by placing them in folders and assigning different permissions to the folders for different users or groups.

### Procedure

- 1 In the Home view of the vSphere Web Client, click **Site Recovery**.
- 2 Expand **Inventory Trees** and click **Recovery Plans**.
- 3 Select the **Related Objects** tab and click **Folders**.
- 4 Click the **Create Folder** icon, enter a name for the folder to create, and click **OK**.
- 5 Add new or existing recovery plans to the folder.

Option	Description
<b>Create a new recovery plan</b>	Right-click the folder and select <b>Create Recovery Plan</b> .
<b>Add an existing recovery plan</b>	Drag and drop recovery plans from the inventory tree into the folder.

- 6 (Optional) To rename or delete a folder, right-click the folder and select **Rename Folder** or **Delete Folder**.

You can only delete a folder if it is empty.

## Edit a Recovery Plan

You can edit a recovery plan to change the properties that you specified when you created it. You can edit recovery plans from the protected site or from the recovery site.

### Procedure

**1** In the vSphere Web Client, select **Site Recovery > Recovery Plans**.

**2** Right-click a recovery plan, and select **Edit Plan**.

You can also edit a recovery plan by clicking the **Edit recovery plan** icon in the **Recovery Steps** view in the **Monitor** tab.

**3** (Optional) Change the name or description of the plan in the **Recovery Plan Name** text box, and click **Next**.

**4** On the Recovery site page, click **Next**.

You cannot change the recovery site.

**5** (Optional) Select or deselect one or more protection groups to add them to or remove them from the plan, and click **Next**.

**6** (Optional) Click the test network to select a different test network on the recovery site, and click **Next**.

**7** Review the summary information and click **Finish** to make the specified changes to the recovery plan.

You can monitor the update of the plan in the Recent Tasks view.

# Setting System Properties

# 9

You can set system properties to change the default Orchestrator behavior.

This chapter includes the following topics:

- [Disable Access to the Orchestrator Client By Nonadministrators](#)
- [Setting Server File System Access for Workflows and Actions](#)
- [Set Access to Operating System Commands for Workflows and Actions](#)
- [Set JavaScript Access to Java Classes](#)
- [Set Custom Timeout Property](#)

## Disable Access to the Orchestrator Client By Nonadministrators

You can configure the Orchestrator server to deny access to the Orchestrator client to all users who are not members of the Orchestrator administrator group.


By default, all users who are granted execute permissions can connect to the Orchestrator client. However, you can limit access to the Orchestrator client to Orchestrator administrators by setting an Orchestrator configuration system property.

---

**Important** If the property is not configured, or if the property is set to false, Orchestrator permits access to the Orchestrator client by all users.

---

### Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **System Properties**.
- 3 Click the **Add** icon (.
- 4 In the **Key** text box enter **com.vmware.o11n.smart-client-disabled**.
- 5 In the **Value** text box enter **true**.
- 6 (Optional) In the **Description** text box enter **Disable Orchestrator client connection**.

- 7 Click **Add**.
- 8 Click **Save changes** from the pop-up menu.  
A message indicates that you have saved successfully.
- 9 Restart the Orchestrator server.

## Results

You disabled access to the Orchestrator client to all users other than members of the Orchestrator administrator group.

## Setting Server File System Access for Workflows and Actions

In Orchestrator, the workflows and actions have limited access to specific file system directories. You can extend access to other parts of the server file system by modifying the `js-io-rights.conf` Orchestrator configuration file.

### Rules in the `js-io-rights.conf` File Permitting Write Access to the Orchestrator System

The `js-io-rights.conf` file contains rules that permit write access to defined directories in the server file system.

---

**Important** Before you modify the `js-io-rights.conf` file, you must stop the vRealize Orchestrator Control Center service. Otherwise, the `js-io-rights.conf` file reverts to its default configuration. See [Set Server File System Access for Workflows and Actions](#).

---

### Mandatory Content of the `js-io-rights.conf` File

Each line of the `js-io-rights.conf` file must contain the following information.

- A plus (+) or minus (-) sign to indicate whether rights are permitted or denied
- The read (r), write (w), and execute (x) levels of rights
- The path on which to apply the rights

### Default Content of the `js-io-rights.conf` File

The default content of the `js-io-rights.conf` configuration file in the Orchestrator Appliance is as follows:

```
-rwx /
+rxw /var/run/vco
+rx /etc/vco
-rwx /etc/vco/app-server/security/
+rx /var/log/vco/
```

The first two lines in the default `js-io-rights.conf` configuration file allow the following access rights:

```
-rwx /
```

All access to the file system is denied.

```
+rwx /var/run/vco
```

Read, write, and execute access is permitted in the `/var/run/vco` directory.

## Rules in the `js-io-rights.conf` File

Orchestrator resolves access rights in the order they appear in the `js-io-rights.conf` file. Each line can override the previous lines.

---

**Important** You can permit access to all parts of the file system by setting `+rwx /` in the `js-io-rights.conf` file. However, doing so represents a high security risk.

---

## Set Server File System Access for Workflows and Actions

To change which parts of the server file system that workflows and the vRealize Orchestrator API can access, modify the `js-io-rights.conf` configuration file. The `js-io-rights.conf` file is created when a workflow attempts to access the vRealize Orchestrator server file system.

### Procedure

- 1 Log in to the vRealize Orchestrator Appliance Linux console as **root**.
- 2 Stop the vRealize Orchestrator Control Center service.

```
service vco-configurator stop
```

- 3 Navigate to `/etc/vco/app-server`.
- 4 Open the `js-io-rights.conf` configuration file in a text editor.
- 5 Add the necessary lines to the `js-io-rights.conf` file.

For example, the following line denies the execution rights in the `/path_to_folder/noexec` directory:

```
-x /path_to_folder/noexec
```

`/path_to_folder/noexec` retains execution rights, but `/path_to_folder/noexec/bar` does not. Both directories remain readable and writable.

- 6 To apply the changes, run the following command.

```
/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh sync-local
```



- 7 Start the vRealize Orchestrator Control Center service.

```
service vco-configurator start
```

### Results


You modified the access rights to the file system for workflows and for the vRealize Orchestrator API.

## Set Access to Operating System Commands for Workflows and Actions

The Orchestrator API provides a scripting class, `Command`, that runs commands in the Orchestrator server host operating system. To prevent unauthorized access to the Orchestrator server host, by default, Orchestrator applications do not have permission to run the `Command` class. If Orchestrator applications require permission to run commands on the host operating system, you can activate the `Command` scripting class.

You grant permission to use the `Command` class by setting an Orchestrator configuration system property.

### Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **System Properties**.
- 3 Click the **Add** icon ()
- 4 In the **Key** text box, enter **com.vmware.js.allow-local-process**.
- 5 In the **Value** text box, enter **true**.
- 6 In the **Description** text box, enter a description for the system property.
- 7 Click **Add**.
- 8 Click **Save changes** from the pop-up menu.  
A message indicates that you have saved successfully.
- 9 Restart the Orchestrator server.

## Results

You granted permissions to Orchestrator applications to run local commands in the Orchestrator server host operating system.

---

**Note** By setting the `com.vmware.js.allow-local-process` system property to `true`, you allow the Command scripting class to write anywhere in the file system. This property overrides any file system access permissions that you set in the `js-io-rights.conf` file for the Command scripting class only. The file system access permissions that you set in the `js-io-rights.conf` file still apply to all scripting classes other than Command.

---

## Set JavaScript Access to Java Classes

By default, Orchestrator restricts JavaScript access to a limited set of Java classes. If you require JavaScript access to a wider range of Java classes, you must set an Orchestrator system property to allow this access.

Allowing the JavaScript engine full access to the Java virtual machine (JVM) presents potential security issues. Malformed or malicious scripts might have access to all of the system components to which the user who runs the Orchestrator server has access. Consequently, by default the Orchestrator JavaScript engine can access only the classes in the `java.util.*` package.


If you require JavaScript access to classes outside of the `java.util.*` package, you can list in a configuration file the Java packages to which to allow JavaScript access. You then set the `com.vmware.scripting.rhino-class-shutter-file` system property to point to this file.

### Procedure

- 1 Create a text configuration file to store the list of Java packages to which to allow JavaScript access.

For example, to allow JavaScript access to all the classes in the `java.net` package and to the `java.lang.Object` class, you add the following content to the file.

```
java.net.*
java.lang.Object
```

- 2 Save the configuration file with an appropriate name and in an appropriate place.
- 3 Log in to Control Center as **root**.
- 4 Click **System Properties**.
- 5 Click the **Add** icon (.
- 6 In the **Key** text box enter **`com.vmware.scripting.rhino-class-shutter-file`**.
- 7 In the **Value** text box enter the path to your configuration file.
- 8 In the **Description** text box enter a description for the system property.

- 9 Click **Add**.
- 10 Click **Save changes** from the pop-up menu.  
A message indicates that you have saved successfully.
- 11 Restart the Orchestrator server.

### Results

The JavaScript engine has access to the Java classes that you specified.


## Set Custom Timeout Property

When vCenter Server is overloaded, it takes more time to return the response to the Orchestrator server than the 20000 milliseconds set by default. To prevent this situation, you must modify the Orchestrator configuration file to increase the default timeout period.

If the default timeout period expires before the completion of certain operations, the Orchestrator server log contains errors.

```
Operation 'getPropertyContent' total time : '5742228' for 1823 calls, mean time :  
'3149.0', min time : '0', max time : '32313' Timeout, unable to get property 'info'  
com.vmware.vmo.plugin.vi4.model.TimeoutException
```

### Procedure

- 1 Log in to Control Center as **root**.
- 2 Click **System Properties**.
- 3 Click the **Add** icon (.
- 4 In the **Key** text box enter **com.vmware.vmo.plugin.vi4.waitUpdatesTimeout**.
- 5 In the **Value** text box enter the new timeout period in milliseconds.
- 6 (Optional) In the **Description** text box enter a description for the system property.
- 7 Click **Add**.
- 8 Click **Save changes** from the pop-up menu.  
A message indicates that you have saved successfully.
- 9 Restart the Orchestrator server.

### Results

The value you set overrides the default timeout setting of 20000 milliseconds.

# Where to Go From Here

# 10

When you have installed and configured vRealize Orchestrator, you can use Orchestrator to automate frequently repeated processes related to the management of the virtual environment.

- Log in to the Orchestrator client, run, and schedule workflows on the vCenter Server inventory objects or other objects that Orchestrator accesses through its plug-ins. See *Using the VMware vRealize Orchestrator Client*.
- Duplicate and modify the standard Orchestrator workflows and write your own actions and workflows to automate operations in vCenter Server.
- Develop plug-ins and Web services to extend the Orchestrator platform.
- Run workflows on your vSphere inventory objects by using the vSphere Web Client.

This chapter includes the following topics:

- [Log In to the Orchestrator Legacy Client from the Orchestrator Appliance Web Console](#)

## Log In to the Orchestrator Legacy Client from the Orchestrator Appliance Web Console

To perform general administration tasks or to edit and create workflows, you must log in to the Orchestrator Legacy Client interface.

The Orchestrator Legacy Client interface is designed for developers with administrative rights who want to develop workflows, actions, and other custom elements.

---

**Important** Ensure that the clocks of the Orchestrator Appliance and the Orchestrator Legacy Client machine are synchronized.

---

### Prerequisites

- Download and deploy the Orchestrator Appliance.
- Verify that the appliance is up and running.

- Install 64-bit Java on the workstation, on which you run the Orchestrator Legacy Client.

---

**Note** 32-bit Java is not supported

---

#### Procedure

- 1 In a Web browser, go to the IP address of your Orchestrator Appliance virtual machine.

`http://orchestrator_appliance_ip`

- 2 Click **Start Orchestrator Client**.

- 3 Enter the IP or the domain name of the Orchestrator Appliance in the **Host name** text box. The IP address of the Orchestrator Appliance is displayed by default.

- 4 Log in by using the Orchestrator Legacy Client user name and password.

Depending on whether you use vRealize Automation or vSphere as an authentication provider, enter the respective credentials to log in to the Orchestrator Legacy Client.

If multi tenancy is enabled on your Orchestrator environment, enter the respective system administrator or tenant administrator user name, password, and tenant ID.

- 5 In the **Security Warning** window, select an option to handle the certificate warning.

The Orchestrator Legacy Client communicates with the Orchestrator server by using an SSL certificate. A trusted CA does not sign the certificate during installation. You receive a certificate warning each time you connect to the Orchestrator server.

Option	Description
<b>Ignore</b>	Continue using the current SSL certificate. The warning message appears again when you reconnect to the same Orchestrator server, or when you try to synchronize a workflow with a remote Orchestrator server.
<b>Cancel</b>	Close the window and stop the login process.
<b>Install this certificate and do not display any security warnings for it anymore.</b>	Select this check box and click <b>Ignore</b> to install the certificate and stop receiving security warnings.

You can change the default SSL certificate with a certificate signed by a CA. For more information about changing SSL certificates, see *Installing and Configuring VMware vRealize Orchestrator*.

#### What to do next

You can import a package, start a workflow, or set root access rights on the system.